

Sharp asymptotics for Kawasaki dynamics on a finite box with open boundary

A. Bovier ^{*}
F. den Hollander [†]
F.R. Nardi [‡]

8th May 2005

Abstract

In this paper we study the metastable behavior of the lattice gas in two and three dimensions subject to Kawasaki dynamics in the limit of low temperature and low density. We consider the local version of the model, where particles live on a finite box and are created, respectively, annihilated at the boundary of the box in a way that reflects an infinite gas reservoir. We are interested in how the system nucleates, i.e., how it reaches a full box when it starts from an empty box. Our approach combines geometric and potential theoretic arguments.

In two dimensions, we identify the full geometry of the set of critical droplets for the nucleation, compute the average nucleation time up to a multiplicative factor that tends to one, show that the nucleation time divided by its average converges to an exponential random variable, express the proportionality constant for the average nucleation time in terms of certain capacities associated with simple random walk, and compute the asymptotic behavior of this constant as the system size tends to infinity. In three dimensions, we obtain similar results but with less control over the geometry and the constant.

A special feature of Kawasaki dynamics is that in the metastable regime particles move along the border of a droplet more rapidly than they arrive from the boundary of the box. The geometry of the critical droplet and the sharp asymptotics for the average nucleation time are highly sensitive to this motion.

AMS 2000 subject classifications. 60K35, 82B43, 82C43, 82C80.

Key words and phrases. Lattice gas, Kawasaki dynamics, metastability, critical droplet, discrete isoperimetric inequalities, potential theory, Dirichlet form, capacity.

Running title: Sharp asymptotics for Kawasaki dynamics.

Acknowledgment. The authors are grateful to Michiel van den Berg (Bristol, United Kingdom) and to Alex Gaudilli, Enzo Olivieri and Elisabetta Scoppola (Rome, Italy) for discussions during the preparation of this work. FRN was supported by a fellowship from NWO while being a postdoc at EURANDOM. AB and FdH were partially supported through the “Dutch-German Bilateral Research Group on Random Spatial Models from Physics and Biology”, which is financed by NWO and DFG. AB and FdH thank the Isaac Newton Institute in Cambridge, United Kingdom, for hospitality in the Summer of 2003.

^{*}Weierstrass-Institut für Angewandte Analysis und Stochastik, Mohrenstrasse 39, 10117 Berlin, Germany, and Institut für Mathematik, Technische Universität Berlin, Strasse des 17. Juni 136, 10623 Berlin, Germany

[†]EURANDOM, P.O. Box 513, 5600 MB Eindhoven, The Netherlands

[‡]Dipartimento di Matematica, Università di Roma “La Sapienza”, Piazzale Aldo Moro 2, 00100 Roma, Italy

1 Introduction and main results

In this paper we study the metastable behavior of the lattice gas in two and three dimensions subject to Kawasaki dynamics at low temperature and low density. Particles live on a finite box, hop between nearest-neighbor sites, have an attractive interaction when they sit next to each other, and are created, respectively, annihilated at the boundary of the box in a way that reflects an infinite gas reservoir. We are interested in how the system *nucleates*, i.e., how it reaches a full box when it starts from an empty box. Our goal is to improve on earlier work by combining a detailed analysis of the energy landscape for the dynamics with the potential theoretic approach to metastability that was developed in Bovier, Eckhoff, Gaynard, and Klein [5] and further exposed in Bovier [3].

Our main theorems sharpen those obtained by den Hollander, Olivieri, and Scoppola [9] in two dimensions and by den Hollander, Nardi, Olivieri, and Scoppola [8] in three dimensions. In particular, in two dimensions we identify the full geometry of the set of critical droplets, compute the average nucleation time up to a multiplicative factor that tends to one, show that the nucleation time divided by its average converges to an exponential random variable, express the proportionality constant for the average nucleation time in terms of certain capacities associated with simple random walk, and compute the asymptotic behavior of this constant as the system size tends to infinity. In three dimensions, we obtain similar results but with less control over the geometry and the constant.

Our results are comparable with those derived by Bovier and Manzo [6] for the Ising model on a finite box in two and three dimensions with periodic boundary conditions subject to Glauber dynamics at low temperature. This work sharpened earlier results by Neves and Schonmann [11] in two dimensions and by Ben Arous and Cerf [4] in three dimensions.

Kawasaki differs from Glauber in that it is a *conservative dynamics*: particles are conserved in the interior of the box. This creates a complication in controlling the growing and the shrinking of droplets, because particles have to travel between the droplet and the boundary of the box. Moreover, it turns out that in the metastable regime *particles move along the border of a droplet more rapidly than they arrive from the boundary of the box*. This leads to a shape of the critical droplet that is more complicated than the one for Ising spins under Glauber dynamics. This complexity needs to be handled in order to obtain the sharp asymptotics. For a critical comparison of Glauber and Kawasaki we refer to den Hollander [7].

The outline of the paper is as follows. In Section 1 we define the model, recall earlier results, and state our main theorems. In Section 2 we consider two dimensions, collect the key geometric facts that underlie our analysis, and prove our result identifying the full geometry of the set of critical droplets. In Section 3 we use this full geometry to prove our sharp asymptotics for the average nucleation time. In Section 4 we show to what extent these results can be extended to three dimensions.

1.1 Hamiltonian and Gibbs measure

Let $\Lambda \subseteq \mathbb{Z}^2$ be a large square box, centered at the origin. Let

$$\begin{aligned}\partial^- \Lambda &= \{x \in \Lambda : \exists y \notin \Lambda : |y - x| = 1\}, \\ \partial^+ \Lambda &= \{x \notin \Lambda : \exists y \in \Lambda : |y - x| = 1\},\end{aligned}\tag{1.1.1}$$

be the internal, respectively, external boundary of Λ , and put

$$\begin{aligned}\Lambda^- &= \Lambda \setminus \partial^- \Lambda, \\ \Lambda^+ &= \Lambda \cup \partial^+ \Lambda.\end{aligned}\tag{1.1.2}$$

With each site $x \in \Lambda$ we associate an occupation variable $\eta(x)$, assuming the values 0 or 1, indicating the absence or presence of a particle at x . A lattice configuration is denoted by $\eta \in \mathcal{X} = \{0, 1\}^\Lambda$. Each configuration $\eta \in \mathcal{X}$ has an energy given by the Hamiltonian

$$H(\eta) = -U \sum_{(x,y) \in \Lambda^{*, -}} \eta(x)\eta(y) + \Delta \sum_{x \in \Lambda} \eta(x), \quad (1.1.3)$$

where

$$\Lambda^{*, -} = \{(x, y) : x, y \in \Lambda^-, |x - y| = 1\} \quad (1.1.4)$$

is the set of non-oriented bonds in Λ^- . The interaction consists of a *binding energy* $-U < 0$ for each nearest-neighbor pair of particles in Λ^- . In addition, there is an *activation energy* $\Delta > 0$ for each particle in Λ .

The Gibbs measure associated with H is

$$\mu_\beta(\eta) = \frac{e^{-\beta H(\eta)}}{Z_\beta}, \quad \eta \in \mathcal{X}, \quad (1.1.5)$$

with inverse temperature $\beta > 0$ and partition sum

$$Z_\beta = \sum_{\eta \in \mathcal{X}} e^{-\beta H(\eta)}. \quad (1.1.6)$$

1.2 Kawasaki dynamics

We next define Kawasaki dynamics on Λ , with a *boundary condition that mimics the effect of an infinite gas reservoir* outside Λ with density

$$\rho_\beta = e^{-\Delta\beta}, \quad (1.2.1)$$

in accordance with the activation energy Δ appearing in (1.1.3).

Let $b = (x \rightarrow y)$ denote an oriented bond, i.e., an ordered pair of nearest-neighbor sites. Define

$$\begin{aligned} \Lambda^{*, orie} &= \{b = (x \rightarrow y) : x, y \in \Lambda\}, \\ \partial\Lambda^{*, in} &= \{b = (x \rightarrow y) : x \in \partial^+\Lambda, y \in \partial^-\Lambda\}, \\ \partial\Lambda^{*, out} &= \{b = (x \rightarrow y) : x \in \partial^-\Lambda, y \in \partial^+\Lambda\}, \end{aligned} \quad (1.2.2)$$

and put $\bar{\Lambda}^{*, orie} = \Lambda^{*, orie} \cup \partial\Lambda^{*, in} \cup \partial\Lambda^{*, out}$. Two configurations $\eta, \eta' \in \mathcal{X}$ with $\eta \neq \eta'$ are called *communicating configurations*, written $\eta \leftrightarrow \eta'$, if there exists a bond $b \in \bar{\Lambda}^{*, orie}$ such that $\eta' = T_b\eta$, where $T_b\eta$ is the configuration obtained from η as follows:

– $b = (x \rightarrow y) \in \Lambda^{*, orie}$:

$$(T_b\eta)(z) = \begin{cases} \eta(z) & \text{if } z \neq x, y, \\ \eta(x) & \text{if } z = y, \\ \eta(y) & \text{if } z = x. \end{cases} \quad (1.2.3)$$

– $b = (x \rightarrow y) \in \partial\Lambda^{*, in}$:

$$(T_b\eta)(z) = \begin{cases} \eta(z) & \text{if } z \neq y, \\ 1 & \text{if } z = y. \end{cases} \quad (1.2.4)$$

– $b = (x \rightarrow y) \in \partial\Lambda^{*, out}$:

$$(T_b\eta)(z) = \begin{cases} \eta(z) & \text{if } z \neq x, \\ 0 & \text{if } z = x. \end{cases} \quad (1.2.5)$$

These transitions correspond to particle motion in Λ , creation and annihilation in $\partial^-\Lambda$, respectively. Note that, for $b \in \Lambda^*$, *orie*, $T_b\eta$ is invariant under a change of orientation of b , while for $b \in \partial\Lambda^*$, *out* and $b \in \partial\Lambda^*$, *in* it is not.

The *Kawasaki dynamics* is defined to be the continuous-time Markov chain $(\eta_t)_{t \geq 0}$ on \mathcal{X} with transition rates

$$c_\beta(\eta, \eta') = 1_{\{\eta \leftrightarrow \eta'\}} e^{-\beta\{[H(\eta') - H(\eta)] \vee 0\}}, \quad \forall \eta, \eta' \in \mathcal{X}, \eta \neq \eta'. \quad (1.2.6)$$

This is a standard Metropolis dynamics with an open boundary: along each bond touching $\partial^-\Lambda$ from the outside, particles are created with rate ρ_β and are annihilated with rate 1, reflecting the activation energy, while inside Λ^- particles are conserved and jump at a rate that depends on the change in energy associated with the jump, reflecting the binding energy. Note that a move of particles inside $\partial^-\Lambda$ does not involve a change in energy because the interaction acts only inside Λ^- (see (1.1.3)).

The measure μ_β defined in (1.1.5) is the *reversible equilibrium* of the dynamics with transition rates c_β defined in (1.2.6):

$$\mu_\beta(\eta) c_\beta(\eta, \eta') = \mu_\beta(\eta') c_\beta(\eta', \eta) \quad \forall \eta, \eta' \in \mathcal{X}, \eta \neq \eta'. \quad (1.2.7)$$

1.3 Rough description of nucleation in two dimensions

1.3.1 Metastable regime and critical droplet size

In two dimensions, we will be interested in the *metastable regime*

$$\Delta \in (U, 2U), \quad \beta \rightarrow \infty. \quad (1.3.1)$$

In this regime, droplets tend to grow slowly: single particles attached to one side of a droplet typically detach before the arrival of a next particle (because $e^{U\beta} \ll e^{\Delta\beta}$), while bars of two or more particles typically do not detach (because $e^{\Delta\beta} \ll e^{2U\beta}$).

As was pointed out in den Hollander, Olivieri, and Scoppola [9], Section 1.2.3, the energy $E(\ell)$ of an $\ell \times \ell$ droplet in Λ^- equals (recall (1.1.3) and see Fig. 1)

$$E(\ell) = -U[2\ell(\ell - 1)] + \Delta\ell^2 = 2U\ell - (2U - \Delta)\ell^2, \quad (1.3.2)$$

which is maximal at $\ell = U/(2U - \Delta)$:




Fig. 1. $\ell \mapsto E(\ell)$.

The *critical droplet size* is therefore given by

$$\ell_c = \left\lceil \frac{U}{2U - \Delta} \right\rceil \quad (1.3.3)$$

($\lceil \cdot \rceil$ denotes the upper integer part), provided we assume that

$$\frac{U}{2U - \Delta} \notin \mathbb{N} \quad (1.3.4)$$

in order to avoid ties. Throughout the sequel we assume that (1.3.4) is in force. Thus, an $(\ell_c - 1) \times (\ell_c - 1)$ droplet is subcritical while an $\ell_c \times \ell_c$ droplet is supercritical.

1.3.2 Basic geometric definitions

Throughout the sequel, we identify a configuration $\eta \in \mathcal{X}$ with its support $\text{supp}(\eta) = \{x \in \Lambda : \eta(x) = 1\}$, and write $x \in \eta$ to indicate that η has a particle at x .

To state what is known about nucleation in two dimensions, we need some basic geometric definitions:

Definition 1.3.1 (a) A path ω is a sequence $\omega = (\omega_1, \dots, \omega_k)$, $k \in \mathbb{N}$, of communicating configurations, i.e., $\omega_i \in \mathcal{X}$ for $i = 1, \dots, k$ and $c_\beta(\omega_i, \omega_{i+1}) > 0$ for $i = 1, \dots, k - 1$. For $\eta, \eta' \in \mathcal{X}$, we write $\omega : \eta \rightarrow \eta'$ to denote a path from η to η' . For $\zeta \in \mathcal{X}$, we write $\zeta \in \omega$ when ω visits ζ . For $\mathcal{A} \subseteq \mathcal{X}$, we write $\omega \subseteq \mathcal{A}$ when ω stays inside \mathcal{A} .

(b) For $\eta, \eta' \in \mathcal{X}$, a path $\omega : \eta \rightarrow \eta'$ is called a U -path if

$$\begin{aligned} (i) \quad & H(\eta) = H(\eta'), \\ (ii) \quad & \max_i H(\omega_i) \leq H(\eta) + U, \\ (iii) \quad & |\omega_i| = |\eta| \text{ for all } i. \end{aligned} \quad (1.3.5)$$

(c) The configuration space \mathcal{X} can be partitioned as

$$\mathcal{X} = \bigcup_{n=0}^{|\Lambda|} \mathcal{V}_n, \quad (1.3.6)$$

where

$$\mathcal{V}_n = \{\eta \in \mathcal{X} : |\eta| = n\} \quad (1.3.7)$$

is the set of configurations with n particles, called the n -manifold.

(d) For $\mathcal{A} \subseteq \mathcal{X}$, the communication height between $\eta, \eta' \in \mathcal{A}$ inside \mathcal{A} is

$$\Phi_{\mathcal{A}}(\eta, \eta') = \min_{\substack{\omega : \eta \rightarrow \eta' \\ \omega \subseteq \mathcal{A}}} \max_{\zeta \in \omega} H(\zeta). \quad (1.3.8)$$

We write $\Phi(\eta, \eta') = \Phi_{\mathcal{X}}(\eta, \eta')$.

(e) For $\mathcal{A} \subseteq \mathcal{X}$, the communication level set between $\eta, \eta' \in \mathcal{A}$ inside \mathcal{A} is

$$\mathcal{S}_{\mathcal{A}}(\eta, \eta') = \left\{ \zeta \in \mathcal{A} : \exists \omega : \eta \rightarrow \eta', \omega \subseteq \mathcal{A}, \omega \ni \zeta : \max_{\xi \in \omega} H(\xi) = H(\zeta) = \Phi_{\mathcal{A}}(\eta, \eta') \right\}. \quad (1.3.9)$$

We write $\mathcal{S}(\eta, \eta') = \mathcal{S}_{\mathcal{X}}(\eta, \eta')$.

(f) For $\eta \in \mathcal{X}$, the law of $(\eta_t)_{t \geq 0}$ starting from $\eta_0 = \eta$ is denoted by \mathbb{P}_η . For $\mathcal{A} \subseteq \mathcal{X}$,

$$\tau_{\mathcal{A}} = \inf\{t > 0 : \eta_t \in \mathcal{A}, \exists 0 < s < t : \eta_s \neq \eta_0\} \quad (1.3.10)$$

is the first hitting/return time of \mathcal{A} .

Each configuration can be decomposed into maximally connected components, called *clusters*. The following sets of configurations will determine the geometry of the critical droplet, as will become clear later on.

Definition 1.3.2 (a) Let \mathcal{Q} denote the set of configurations having one cluster anywhere in Λ^- consisting of an $(\ell_c - 1) \times \ell_c$ quasi-square with a single protuberance attached to one of its sides.

(b) Let \mathcal{D} denote the set of configurations that can be reached from some configuration in \mathcal{Q} via a U -path, i.e.,

$$\mathcal{D} = \left\{ \eta' \in \mathcal{V}_{n_c} : \exists \eta \in \mathcal{Q} : H(\eta) = H(\eta'), \Phi_{\mathcal{V}_{n_c}}(\eta, \eta') \leq H(\eta) + U \right\}, \quad (1.3.11)$$

where $n_c = \ell_c(\ell_c - 1) + 1$ is the volume of the clusters in \mathcal{Q} .

(c) Let $\mathcal{C}^* = \mathcal{D}^{fp}$, where $(\cdot)^{fp}$ denotes addition of a free particle anywhere in Λ (see Fig. 2).

(d) Let

$$\begin{aligned} \Gamma^* &= H(\mathcal{C}^*) = H(\mathcal{D}^{fp}) = H(\mathcal{D}) + \Delta = H(\mathcal{Q}) + \Delta \\ &= -U[(\ell_c - 1)^2 + \ell_c(\ell_c - 2) + 1] + \Delta[\ell_c(\ell_c - 1) + 2] \\ &= 2U[\ell_c + 1] - (2U - \Delta)[\ell_c(\ell_c - 1) + 2] \end{aligned} \quad (1.3.12)$$

denote the energy of the configurations in \mathcal{C}^* .

As we will see shortly, \mathcal{Q} plays the role of the set of *canonical protocritical droplets* for the nucleation, $\mathcal{D} \supseteq \mathcal{Q}$ the set of *protocritical droplets*, and \mathcal{C}^* the set of *critical droplets*. Think of \mathcal{D} as the set of configurations the dynamics can reach *after* hitting \mathcal{Q} *before* the creation of the next free particle in $\partial^-\Lambda$ (which takes a time $e^{\Delta\beta} \gg e^{U\beta}$). This particle moves the configuration into \mathcal{C}^* and completes the formation of the critical droplet (= critical cluster + free particle) that triggers the nucleation. If subsequently the free particle moves to the critical cluster and attaches itself *properly* (i.e., in a corner), then the dynamics has “moved over the hill” and proceeds to fill Λ^- .




FIG. 2. A canonical critical droplet: an element of $\mathcal{Q}^{fp} \subseteq \mathcal{D}^{fp} = \mathcal{C}^*$.

1.3.3 Nucleation time and critical droplets

Let

$$\begin{aligned}\square &= \{\eta \in \mathcal{X} : \eta(x) = 0 \ \forall x \in \Lambda\}, \\ \blacksquare &= \{\eta \in \mathcal{X} : \eta(x) = 1 \ \forall x \in \Lambda^-, \eta(x) = 0 \ \forall x \in \partial^- \Lambda\},\end{aligned}\tag{1.3.13}$$

denote the configurations where Λ is empty, respectively, Λ^- is full and $\partial^- \Lambda$ is empty. We assume that Λ is so large that

$$H(\blacksquare) < H(\square) = 0.\tag{1.3.14}$$

In this case, \blacksquare is the global minimum of H . The main result known about nucleation in two dimensions reads as follows.

Theorem 1.3.3 (den Hollander, Olivieri, and Scoppola [9], Theorem 1.53 and Proposition 4.24)

(i) $\Phi(\square, \blacksquare) = \Gamma^*$ and $\mathcal{S}(\square, \blacksquare) \supseteq \mathcal{C}^*$.

(ii)

$$\lim_{\beta \rightarrow \infty} \mathbb{P}_{\square} \left(e^{(\Gamma^* - \delta)\beta} < \tau_{\blacksquare} < e^{(\Gamma^* + \delta)\beta} \right) = 1 \quad \forall \delta > 0.\tag{1.3.15}$$

(iii)

$$\lim_{\beta \rightarrow \infty} \mathbb{P}_{\square} (\tau_{\mathcal{C}^*} < \tau_{\blacksquare} \mid \tau_{\blacksquare} < \tau_{\square}) = 1.\tag{1.3.16}$$

Theorem 1.3.3(i) identifies Γ^* as the communication height for the nucleation and \mathcal{C}^* as a subset of the communication level set for the nucleation. Theorem 1.3.3(ii) identifies the nucleation time to exponential order in β , with exponent Γ^* . Theorem 1.3.3(iii) states that \mathcal{C}^* is a *gate for the nucleation*.

1.4 Sharp description of nucleation in two dimensions

1.4.1 Goal and background

The goal of the present paper is to sharpen Theorem 1.3.3 in two ways:

- (I) Equation (1.3.11) defines \mathcal{D} as a certain neighborhood of \mathcal{Q} defined in terms of energies and communication heights. We will describe the configurations in \mathcal{D} *geometrically* and elaborate on the *gate structure* of $\mathcal{C}^* = \mathcal{D}^{fp}$.
- (II) We will sharpen (1.3.15) by computing the average nucleation time *up to a multiplicative factor that tends to one as $\beta \rightarrow \infty$* and by showing that the limit law is exponential. This will require the knowledge obtained in (I).

To achieve (II), we will apply the *potential theoretic approach to metastability* developed in Bovier, Eckhoff, Gaynard, and Klein [5] and further exposed in Bovier [3]. There it was shown that, for reversible Markov processes exhibiting metastability, the computation of average metastable exit times and of corresponding small eigenvalues of the generator reduces to the computation of certain *capacities*. The advantage of this reduction is that the variational representation of capacities given through the *Dirichlet form* allows for a sharp computation of these capacities up to multiplicative errors that tend to one as the small parameters in the theory tend to zero. Roughly speaking, the reason why this happens is that in metastable systems the full Dirichlet form effectively reduces to a Dirichlet form involving only a tiny fraction of the state space, namely, *the communication level set and its immediate vicinity*.

In Bovier and Manzo [6] it was shown that this situation arises for the Ising model with Glauber dynamics in finite volume in the limit of low temperature. For that model the situation turns out to be relatively simple, because the communication level set is completely disconnected, implying that the full Dirichlet form reduces to a sum of zero-dimensional Dirichlet forms. We will show that in our model the same approach can be followed, even though the structure of the communication level set is far more complicated. In particular, *in our model this set contains plateaus, wells embedded in these plateaus, and dead-ends*. Thus, the reduced Dirichlet forms remain multi-dimensional. However, we will be able to express them in terms of certain hitting probabilities of simple random walk. The latter will turn out to be sufficiently manageable so as to allow for a computation of the asymptotic behavior of the reduced Dirichlet forms as $\Lambda \rightarrow \mathbb{Z}^2$.

The idea behind the potential theoretic approach is explained in Section 3.3. Certain specific geometric information is needed for this approach to work, which is gathered in Section 2, but this information is relatively limited.

Throughout the paper we assume that $\ell_c \geq 3$. The case $\ell_c = 2$ is trivial: $\mathcal{Q} = \mathcal{D}$ is the set of configurations consisting of three particles forming a cluster anywhere in Λ^- , \mathcal{C}^* is the set of configurations obtained from these by adding a free particle anywhere in Λ , and $\Gamma^* = -2U + 4\Delta$.

1.4.2 Geometry of protocritical droplets

Our first theorem identifies the full geometry of the configurations in \mathcal{D} (see Fig. 3) and will be proved in Section 2.2.

Theorem 1.4.1 $\mathcal{D} = \bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$, where

- $\bar{\mathcal{D}}$ is the set of configurations having one cluster anywhere in Λ^- consisting of an $(\ell_c - 2) \times (\ell_c - 2)$ square with four bars of lengths \bar{k}_i , $i = 1, 2, 3, 4$, attached to its four sides satisfying

$$1 \leq \bar{k}_i \leq \ell_c - 1, \quad \sum_i \bar{k}_i = 3\ell_c - 3; \quad (1.4.1)$$

- $\tilde{\mathcal{D}}$ is the set of configurations having one cluster anywhere in Λ^- consisting of an $(\ell_c - 3) \times (\ell_c - 1)$ rectangle with four bars of lengths \tilde{k}_i , $i = 1, 2, 3, 4$, attached to its four sides satisfying

$$1 \leq \tilde{k}_i \leq \ell_c - 1, \quad \sum_i \tilde{k}_i = 3\ell_c - 2. \quad (1.4.2)$$




Fig. 3. Configurations in $\bar{\mathcal{Q}}$ and $\bar{\mathcal{D}}$ for $\ell_c = 14$. A similar picture applies for $\tilde{\mathcal{Q}}$ and $\tilde{\mathcal{D}}$ with a 11×13 rectangle in the center.

REMARK: In the first half of Theorem 1.4.1, the four bars may be placed anywhere in the ring around the square, i.e., anywhere in the union of the two rows, the two columns and the four corners forming the outer layer of the square. A total of $3\ell_c - 3$ particles must be accomodated in this ring in such a way that each side of the ring, i.e., each row or column with its two adjacent corners, contains precisely one bar. Thus, a bar may include a corner of the ring provided the neighboring bar also includes this corner. Similarly for the second half of Theorem 1.4.1.

Compare Definitions 1.3.2(a) and 1.4.1. Write $\mathcal{Q} = \bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}}$, where

- $\bar{\mathcal{Q}}$ are those configurations where the single particle is attached to one of the *longest* sides of the $(\ell_c - 1) \times \ell_c$ quasi-square.
- $\tilde{\mathcal{Q}}$ are those configurations where the single particle is attached to one of the *smallest* sides of the $(\ell_c - 1) \times \ell_c$ quasi-square.

Then $\bar{\mathcal{Q}}$ consists of precisely those configurations in $\bar{\mathcal{D}}$ where one \bar{k}_i equals 1 and the others are maximal. Similarly, $\tilde{\mathcal{Q}}$ consists of precisely those configurations in $\tilde{\mathcal{D}}$ where one \tilde{k}_i equals 1 and the others are maximal. We will see in Section 2.2 that the configurations in $\bar{\mathcal{D}}, \tilde{\mathcal{D}}$ arise from those in $\bar{\mathcal{Q}}, \tilde{\mathcal{Q}}$ via a *motion of particles along the border of the droplet*. This property is special for Kawasaki dynamics.

1.4.3 Minimal gates and entrance distribution

To formulate our sharpening of Theorem 1.3.3 we need some more definitions.

Definition 1.4.2 Fix $\eta, \eta' \in \mathcal{X}$.

- (a) The set of paths realizing the minimax in $\Phi(\eta, \eta')$ (recall (1.3.8)) is denoted by $(\eta \rightarrow \eta')_{opt}$.
- (b) A set $\mathcal{W} \subseteq \mathcal{X}$ is called a gate for $\eta \rightarrow \eta'$ if $\mathcal{W} \subseteq \mathcal{S}(\eta, \eta')$ and $\omega \cap \mathcal{W} \neq \emptyset$ for all $\omega \in (\eta \rightarrow \eta')_{opt}$.
- (c) A set $\mathcal{W} \subseteq \mathcal{X}$ is called a minimal gate for $\eta \rightarrow \eta'$ if it is a gate for $\eta \rightarrow \eta'$ and for any $\mathcal{W}' \subsetneq \mathcal{W}$ there exists an $\omega' \in (\eta \rightarrow \eta')_{opt}$ such that $\omega' \cap \mathcal{W}' = \emptyset$.
- (d) A priori there may be several (not necessarily disjoint) minimal gates. The union of all the minimal gates

$$\mathcal{G}(\eta, \eta') = \bigcup_{\mathcal{W} \text{ minimal gate for } \eta \rightarrow \eta'} \mathcal{W} \quad (1.4.3)$$

is called the essential gate for $\eta \rightarrow \eta'$.

- (e) The configurations in $\mathcal{S}(\eta, \eta') \setminus \mathcal{G}(\eta, \eta')$ are called dead-ends.

The notion of minimal gate for $\square \rightarrow \blacksquare$ is important: on its way from \square to \blacksquare the dynamics passes through each of the minimal gates for $\square \rightarrow \blacksquare$ with a probability tending to one as $\beta \rightarrow \infty$, i.e., (1.3.16) holds with \mathcal{C}^* replaced by any of the minimal gates, or any union of them. Thus, the essential gate $\mathcal{G}(\square, \blacksquare)$ plays the role of the minimal set of configurations in $\mathcal{S}(\square, \blacksquare)$ the dynamics can see on its way from \square to \blacksquare . For an elaborate dicussion of essential gates and their role for metastable transition times, we refer the reader to Manzo, Nardi, Olivieri and Scoppola [10].

Our second theorem extends Theorem 1.3.3(i–ii) and will be proved in Section 3.5.

Theorem 1.4.3 (i) $\mathcal{S}(\square, \blacksquare) \supsetneq \mathcal{G}(\square, \blacksquare)$, $\mathcal{S}(\square, \blacksquare) \supsetneq \mathcal{C}^*$.

(ii)

$$\lim_{\beta \rightarrow \infty} \mathbb{P}_{\square}(\tau_{\mathcal{Q}} < \tau_{\mathcal{C}^*} < \tau_{\blacksquare} \mid \tau_{\blacksquare} < \tau_{\square}) = 1. \quad (1.4.4)$$

(iii)

$$\lim_{\beta \rightarrow \infty} \mathbb{P}_{\square}(\eta_{\tau_{\mathcal{C}^*} -} = \eta \mid \tau_{\mathcal{C}^*} < \tau_{\square}) = \frac{1}{|\mathcal{D}|} \quad \forall \eta \in \mathcal{D} \quad (1.4.5)$$

with $\tau_{\mathcal{C}^*} -$ the time just prior to $\tau_{\mathcal{C}^*}$.

Theorem 1.4.3(i) shows that $\mathcal{S}(\square, \blacksquare)$ has dead-ends and is larger than the set of critical droplets \mathcal{C}^* . Theorem 1.4.3(ii) says that \mathcal{Q} is hit prior to \mathcal{C}^* . Theorem 1.4.3(iii) says that the *entrance* distribution of \mathcal{C}^* is uniform, i.e., the protocritical droplets in \mathcal{D} , seen just prior to the creation of the free particle in $\partial^- \Lambda$, occur with equal probability. (Incidentally, the *exit* distribution is not uniform and turns out to be hard to compute.)

In Section 3.5 we will see that neither $\mathcal{G}(\square, \blacksquare) \subseteq \mathcal{C}^*$ nor $\mathcal{C}^* \subseteq \mathcal{G}(\square, \blacksquare)$. We will identify a set $\mathcal{C}_{min}^* \subsetneq \mathcal{C}^*$ that is a union of minimal gates, each consisting of a protocritical droplet and a free particle in (part of) a ring around the protocritical droplet. Clearly, $\mathcal{C}_{min} \subseteq \mathcal{G}(\square, \blacksquare)$. The inequality will turn out to be strict. Since we have no full classification of the minimal gates, we have no full identification of $\mathcal{G}(\square, \blacksquare)$. This lack is due to motion of particles along the border of the droplet.

1.4.4 Sharp asymptotics

Our third and fourth theorem extend Theorem 1.3.3(iii) and will be proved in Sections 3.3–3.4.

Theorem 1.4.4 *There exists a constant $K = K(\Lambda, \ell_c)$ such that*

$$\mathbb{E}_{\square}(\tau_{\blacksquare}) = K e^{\Gamma^* \beta} [1 + o(1)] \quad \beta \rightarrow \infty. \quad (1.4.6)$$

Moreover,

$$\mathbb{P}_{\square}(\tau_{\blacksquare} > t \mathbb{E}_{\square}(\tau_{\blacksquare})) = [1 + o(1)] e^{-t[1+o(1)]}, \quad t \geq 0, \quad \beta \rightarrow \infty. \quad (1.4.7)$$

Theorem 1.4.4 provides the sharp asymptotics for the nucleation time:

$$\lim_{\beta \rightarrow \infty} \mathbb{P}_{\square}(\tau_{\blacksquare} > t K e^{\Gamma^* \beta}) = e^{-t}, \quad t \geq 0. \quad (1.4.8)$$

The exponential law is typical for “success only occurs after many unsuccessful attempts”.

In Section 3.3 we will derive a representation for the constant K in terms of certain capacities associated with two-dimensional simple random walk. *This representation will depend on the geometry of \mathcal{C}^* and its immediate vicinity*, i.e., those $\eta \in \mathcal{X} \setminus \mathcal{C}^*$ for which there is an $\eta' \in \mathcal{C}^*$ such that $\eta \leftrightarrow \eta'$. In Section 2.3 we will see that this immediate vicinity is actually *rather complex*, due to the fact that when the free particle attaches itself *improperly* to the protocritical droplet (i.e., not in a corner) it triggers a motion of particles along the border of the droplet. Consequently, no easily computable formula for K is available.

It turns out, however, that the behavior of K for large Λ can be computed explicitly.

Theorem 1.4.5 *As $\Lambda \rightarrow \mathbb{Z}^2$,*

$$K(\Lambda, \ell_c) \sim \frac{1}{4\pi N(\ell_c)} \frac{\log |\Lambda|}{|\Lambda|} \quad (1.4.9)$$

(\sim means that the ratio of the left and the right side tends to 1) with

$$N(\ell_c) = \frac{1}{3}(\ell_c - 1)\ell_c^2(\ell_c + 1) \quad (1.4.10)$$

the cardinality of $\mathcal{D} = \mathcal{D}(\Lambda, \ell_c)$ modulo shifts.

The intuition behind Theorem 1.4.5 is as follows. The average time it takes for the dynamics to enter \mathcal{C}^* when starting from \square is

$$\frac{1}{|\mathcal{D}|} \frac{1}{|\partial\Lambda^{*,in}|} e^{\Gamma^*\beta} [1 + o(1)] \quad \beta \rightarrow \infty, \quad (1.4.11)$$

where $|\mathcal{D}|$ counts the number of protocritical droplets and $|\partial\Lambda^{*,in}|$ counts the number of directed bonds from $\partial^+\Lambda$ to $\partial^-\Lambda$ along which the free particle can be created (recall (1.2.2)). Let $\pi(\Lambda, \ell_c)$ be the probability (averaged w.r.t. the uniform distribution for the protocritical droplet on \mathcal{D} and the uniform distribution for the free particle entering on $\partial\Lambda^{*,in}$) that the free particle moves from $\partial^-\Lambda$ to the protocritical droplet and attaches itself *properly* (i.e., in a corner). This is the probability that the dynamics after it enters \mathcal{C}^* moves onwards to \blacksquare rather than returns to \square . Then

$$\frac{1}{\pi(\Lambda, \ell_c)} [1 + o(1)] \quad \beta \rightarrow \infty \quad (1.4.12)$$

is the average number of times a free particle just created in $\partial^-\Lambda$ attempts to move to the protocritical droplet and attach itself properly before it finally manages to do so. The average nucleation time is the product of (1.4.11) and (1.4.12), and so we conclude that

$$K(\Lambda, \ell_c) = \frac{1}{|\mathcal{D}| |\partial\Lambda^{*,in}| \pi(\Lambda, \ell_c)}. \quad (1.4.13)$$

Now, we have

$$|\mathcal{D}| \sim |\Lambda| N(\ell_c) \quad \Lambda \rightarrow \mathbb{Z}^2. \quad (1.4.14)$$

Furthermore, we have

$$|\partial\Lambda^{*,in}| \pi(\Lambda, \ell_c) \sim \frac{4\pi}{\log|\Lambda|} \quad \Lambda \rightarrow \mathbb{Z}^2. \quad (1.4.15)$$

Indeed, as we will see in Section 3.4, the right-hand side of (1.4.15) is the probability for large Λ that a particle detaching itself from the protocritical droplet reaches $\partial^-\Lambda$ before re-attaching itself. Due to the recurrence of simple random walk in two dimensions, for large Λ this probability is independent of the shape and the location of the protocritical droplet, as long as it is far from $\partial^-\Lambda$. By reversibility, the reverse motion has the same probability, which explains (1.4.15). Combine (1.4.13–1.4.15) to get (1.4.9).

If the free particle attaches itself ‘improperly’ to the protocritical droplet, then either it may again detach itself or it may cause some motion of particles along the border of the droplet, after which another particle may detach itself, possibly leaving behind a different protocritical droplet. However, since for large Λ a free particle has a small probability to escape from the protocritical droplet and return to $\partial\Lambda$, it must eventually attach itself ‘properly’. We refer to Section 2.3.2 for details.

The asymptotics in (1.4.9) does not depend on the shape of Λ , e.g. it would be the same if Λ were a large circle rather than a large square.

1.5 Extension to three dimensions

The metastable regime, replacing (1.3.1), is

$$\Delta \in (U, 3U), \quad \beta \rightarrow \infty, \quad (1.5.1)$$

and we assume that

$$\frac{2U}{3U - \Delta} \notin \mathbb{N}. \quad (1.5.2)$$

The analogue of Definition 1.3.2 reads (see den Hollander, Nardi, Olivieri, and Scoppola [8] Eqs. (1.3.7), (1.3.11), (2.0.15), (2.0.17), (2.0.18) and (2.0.21)):

Definition 1.5.1 (a) Let \mathcal{Q} denote the set of configurations having one cluster anywhere in Λ^- consisting of an $(m_c - 1) \times (m_c - \delta_c) \times m_c$ quasi-cube with, attached to one of its faces, an $(\ell_c - 1) \times \ell_c$ quasi-square with, attached to one of its sides, a single particle. Here, $\delta_c \in \{0, 1\}$ depends on the arithmetic properties of U and Δ , while

$$\ell_c = \left\lceil \frac{U}{3U - \Delta} \right\rceil, \quad m_c = \left\lceil \frac{2U}{3U - \Delta} \right\rceil, \quad (1.5.3)$$

are the two-dimensional critical droplet size on a face, respectively, the three-dimensional critical droplet size, replacing (1.3.3). Note that $m_c \in \{2\ell_c - 1, 2\ell_c\}$.

(b) For $\Delta \in (2U, 3U)$, let \mathcal{D} denote the set of configurations that can be reached from some configuration in \mathcal{Q} via a $2U$ -path, i.e.,

$$\mathcal{D} = \left\{ \eta' \in \mathcal{V}_{n_c} : \exists \eta \in \mathcal{Q} : H(\eta) = H(\eta'), \Phi_{\mathcal{V}_{n_c}}(\eta, \eta') \leq H(\eta) + 2U \right\}, \quad (1.5.4)$$

where $n_c = m_c(m_c - \delta_c)(m_c - 1) + \ell_c(\ell_c - 1) + 1$ is the volume of the clusters in \mathcal{Q} . For $\Delta \in (U, 2U)$, use U instead of $2U$ in (1.5.4).

(c) Let $\mathcal{C}^* = \mathcal{D}^{fp}$ denote the set of configurations obtained from \mathcal{D} by adding a free particle anywhere in Λ (see Fig. 4).

(d) Let

$$\begin{aligned} \Gamma^* = H(\mathcal{C}^*) &= H(\mathcal{D}^{fp}) = H(\mathcal{D}) + \Delta = H(\mathcal{Q}) + \Delta \\ &= U[m_c(m_c - \delta_c) + m_c(m_c - 1) + (m_c - \delta_c)(m_c - 1) + 2\ell_c + 3] \\ &\quad - (3U - \Delta)[m_c(m_c - \delta_c)(m_c - 1) + \ell_c(\ell_c - 1) + 2] \end{aligned} \quad (1.5.5)$$

denote the energy of the configurations in \mathcal{C}^* .




FIG. 4. An element of $\mathcal{Q}^{fp} \subseteq \mathcal{D}^{fp} = \mathcal{C}^*$ for $\ell_c = 10$, $m_c = 20$ and $\delta_c = 0$.

As is shown in den Hollander, Nardi, Olivieri, and Scoppola [8], Theorem 1.5.1, the results in Theorem 1.3.3 carry over from two to three dimensions. Unfortunately, we are not able to identify the full geometry of \mathcal{D} , and hence of $\mathcal{C}^* = \mathcal{D}^{fp}$, because *the motion of particles along the border of the droplet is much more complex in three than in two dimensions*, i.e., the analogue of Fig. 3 is not fully understood (see e.g. [8], Figure 7). Consequently, we have no result extending Theorems 1.4.1. Theorem 1.4.3 carries over. The following two theorems, proved in Sections 4.2–4.3, extend Theorems 1.4.4–1.4.5.

Theorem 1.5.2 *There exists a constant $K = K(\Lambda, \ell_c, m_c, \delta_c)$ such that*

$$\mathbb{E}_{\square}(\tau_{\blacksquare}) = K e^{\Gamma^* \beta} [1 + o(1)] \quad \beta \rightarrow \infty. \quad (1.5.6)$$

Moreover,

$$\mathbb{P}_{\square}(\tau_{\blacksquare} > t \mathbb{E}_{\square}(\tau_{\blacksquare})) = [1 + o(1)] e^{-t[1+o(1)]}, \quad t \geq 0, \quad \beta \rightarrow \infty. \quad (1.5.7)$$

We will derive a representation for the constant K in terms of certain capacities associated with three-dimensional simple random walk. As in two dimensions, this representation is so complex that no easily computable formula for K is available. We will deduce the following asymptotics, which is similar in spirit to the one obtained in two dimensions but less complete.

Theorem 1.5.3 *As $\Lambda \rightarrow \mathbb{Z}^3$,*

$$K(\Lambda, \ell_c, m_c, \delta_c) \sim \frac{1}{M(\ell_c, m_c, \delta_c) N(\ell_c, m_c, \delta_c)} \frac{1}{|\Lambda|}, \quad (1.5.8)$$

where $N(\ell_c, m_c, \delta_c)$ is the cardinality of $\mathcal{D} = \mathcal{D}(\Lambda, \ell_c, m_c, \delta_c)$ modulo shifts, and $M(\ell_c, m_c, \delta_c)$ satisfies the bounds

$$\kappa(m_c - \lceil \sqrt{m_c} \rceil) \leq M(\ell_c, m_c, \delta_c) \leq \kappa(m_c + 3) \quad (1.5.9)$$

with $\kappa(m)$ the capacity of the $m \times m \times m$ cube for simple random walk on \mathbb{Z}^3 .

The interpretation of the asymptotic formula for K is similar as in two dimensions. Instead of (1.4.13), we have

$$K = \frac{1}{|\mathcal{D}| |\partial \Lambda^{*,in}| \pi(\Lambda, \ell_c, m_c, \delta_c)} \quad (1.5.10)$$

with $\pi(\Lambda, \ell_c, m_c, \delta_c)$ the analogue of $\pi(\Lambda, \ell_c)$ in two dimensions (defined below (1.4.11)). By the transience of simple random walk in three dimensions, $|\partial \Lambda^{*,in}| \pi(\Lambda, \ell_c, m_c, \delta_c)$ converges to a limit $M(\ell_c, m_c, \delta_c)$ as $\Lambda \rightarrow \mathbb{Z}^3$.

The lower bound in (1.5.9) comes from the fact that all protocritical droplets contain a cube of side length $m_c - \sqrt{m_c}$. The upper bound comes from the fact that all protocritical droplets are contained in a cube of side length $m_c + 1$ and that as long as the free particle is at distance ≥ 2 from the protocritical droplet no border motion is possible (as shown in Section 4.1). Since

$$\kappa(m) \sim \kappa m \quad m \rightarrow \infty, \quad (1.5.11)$$

with κ the capacity of the unit cube for standard Brownian motion on \mathbb{R}^3 , which satisfies $\kappa \in (2\pi, 2\pi\sqrt{3})$, we have good control over $M(\ell_c, m_c, \delta_c)$ for m_c large, i.e., for Δ close to $3U$.

We have no formula for $N(\ell_c, m_c, \delta_c)$ analogous to (1.4.10). It would be nice to know its asymptotics for m_c large.

2 Geometry in two dimensions

In this section we collect the key geometric facts that underlie our analysis. In Section 2.1 we introduce some geometric definitions. In Section 2.2 we prove Theorem 1.4.1, which identifies the full geometry of the set of protocritical droplets. In Section 2.3 we obtain the structure of the communication level set for the nucleation. In Section 2.4 we prove two global geometric facts that will be needed in Section 3.

2.1 Some geometric definitions

Free particles and 1-protuberances are defined as follows:

- For $x \in \Lambda^-$, let $\text{NN}(x) = \{y \in \Lambda^- : |y - x| = 1\}$ be the set of nearest-neighbor sites of x in Λ^- .
- A *free particle* in $\eta \in \mathcal{X}$ is a site $x \in \eta \cap \partial^- \Lambda$ or a site $x \in \eta \cap \Lambda^-$ such that $\sum_{y \in \text{NN}(x)} \eta(y) = 0$, i.e., a particle not in interaction with any other particle (remember from (1.1.3) that particles in the interior boundary $\partial^- \Lambda$ have no interaction with particles in the interior Λ^-).
- A *1-protuberance* in $\eta \in \mathcal{X}$ is a site $x \in \eta \cap \Lambda^-$ such that $\sum_{y \in \text{NN}(x)} \eta(y) = 1$.
- A *corner* in $\eta \in \mathcal{X}$ is a site $x \in \Lambda^-$ such that $\sum_{y \in \text{NN}(x)} \eta(y) \geq 2$. A corner in η can either be occupied or vacant.

Given a configuration $\eta \in \mathcal{X}$, consider the set $C(\eta) \subseteq \mathbb{R}^2$ defined as the union of the closed unit squares centered at the sites inside Λ^- where η has a particle. The maximal connected components C_1, \dots, C_m , $m \in \mathbb{N}$, of $C(\eta)$ are called *clusters* of η (two unit squares touching only at the corners are not connected). There is a one-to-one correspondence between configurations $\eta \subseteq \Lambda^-$ and sets $C(\eta)$. A configuration $\eta \in \mathcal{X}$ is characterized by a set $C(\eta)$, depending only on $\eta \cap \Lambda^-$, plus possibly a set of particles in $\partial^- \Lambda$, namely, $\eta \cap \partial^- \Lambda$. Thus, we are actually identifying two different objects: a configuration $\eta \in \mathcal{X}$ and the pair $(C(\eta), \eta \cap \partial^- \Lambda)$.

For $\eta \in \mathcal{X}$, let $|\eta|$ be the number of particles in η , $\gamma(\eta)$ the Euclidean boundary of $C(\eta)$, called the *contour* of η , and $|\gamma(\eta)|$ the length of $\gamma(\eta)$, i.e., the number of broken bonds in η . Then the *energy* associated with η is given by

$$H(\eta) = \frac{U}{2} |\gamma(\eta)| - (2U - \Delta) |\eta \cap \Lambda^-| + \Delta |\eta \cap \partial^- \Lambda|. \quad (2.1.1)$$

For convenience we identify a configuration $\eta \in \mathcal{X}$ with its support $\text{supp}(\eta) = \{x \in \Lambda : \eta(x) = 1\}$ and write $x \in \eta$ to indicate that η has a particle at x .

Throughout the paper we assume that the square box $\Lambda \subseteq \mathbb{Z}^2$ is large enough to amply accommodate the critical droplet (say, it has side length $\geq 2\ell_c$).

- An $\ell_1 \times \ell_2$ *rectangle* is a union of closed unit squares centered at the sites inside Λ^- with side lengths $\ell_1, \ell_2 \geq 1$. We use the convention $\ell_1 \leq \ell_2$ and collect rectangles in *equivalence classes modulo translations and rotations*.
- A *bar* is a $1 \times k$ rectangle with $k \geq 1$. A bar is called a *row* or a *column* if it fills a side of a rectangle.
- A *corner* of a rectangle is an intersection of two bars attached to the rectangle.

- A *quasi-square* is an $\ell \times (\ell + \delta)$ rectangle with $\ell \geq 1$ and $\delta \in \{0, 1\}$. A square is a quasi-square with $\delta = 0$.
- If η is a configuration with a single contour, then we denote by $\text{CR}(\eta)$ the rectangle *circumscribing* η , i.e., the smallest rectangle containing η . We write

$$\begin{aligned}\partial^- \text{CR}(\eta) &= \{x \in \text{CR}(\eta) : \exists y \notin \text{CR}(\eta) : |y - x| = 1\}, \\ \partial^+ \text{CR}(\eta) &= \{x \notin \text{CR}(\eta) : \exists y \in \text{CR}(\eta) : |y - x| = 1\},\end{aligned}\tag{2.1.2}$$

to denote the interior, respectively, external boundary of $\text{CR}(\eta)$, and put

$$\begin{aligned}\text{CR}^-(\eta) &= \text{CR}(\eta) \setminus \partial^- \text{CR}(\eta), \\ \text{CR}^+(\eta) &= \text{CR}(\eta) \cup \partial^+ \text{CR}(\eta).\end{aligned}\tag{2.1.3}$$

Note that here we identify particles with unit squares.

- Given η such that $\eta \supseteq \text{CR}^-(\eta)$, we say that it is possible to move a particle from *row* $r_\alpha(\eta) \subseteq \partial^- \text{CR}(\eta)$ to *row* $r_{\alpha'}(\eta) \subseteq \partial^- \text{CR}(\eta)$ via *corner* $c_{\alpha, \alpha'}(\eta) \in \partial^- \text{CR}(\eta)$ if (see Fig. 6 below)

$$|c_{\alpha, \alpha'}(\eta) \cap \eta| = 0, \quad |r_\alpha(\eta) \cap \eta| \geq 1, \quad 1 \leq |r_{\alpha'}(\eta) \cap \eta| \leq |r_{\alpha'}(\eta)|,\tag{2.1.4}$$

where $\alpha, \alpha' \in \{ne, nw, se, sw\}$ with n = north, s = south, etc. By convention, corners are not part of rows. If equality holds in the last inequality, then we need to place the bar in the row opposite to $r_\alpha(\eta)$, say $r_{\alpha''}(\eta)$, a distance 1 away from $c_{\alpha, \alpha''}(\eta)$ in order to be able to accommodate the shift of a bar in $r_{\alpha'}(\eta)$ that is necessary to accomodate the particle that moves around the corner.

2.2 Protocritical droplets: Proof of Theorem 1.4.1

The proof of Theorem 1.4.1 will be given in two steps:

$$\begin{aligned}(i) \quad & \bar{\mathcal{D}} \cup \tilde{\mathcal{D}} \subseteq \mathcal{D}, \\ (ii) \quad & \bar{\mathcal{D}} \cup \tilde{\mathcal{D}} \supseteq \mathcal{D}.\end{aligned}\tag{2.2.1}$$

Proof of (i): Recall the definition of U -path in (1.3.5) and the definitions of $\bar{\mathcal{Q}}, \tilde{\mathcal{Q}}$ and $\bar{\mathcal{D}}, \tilde{\mathcal{D}}$ in Section 1.4. To prove (i) we must show that for all $\eta \in \bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$,

$$\begin{aligned}(i1) \quad & H(\eta) = H(\bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}}), \\ (i2) \quad & \exists \omega : \bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}} \rightarrow \eta : \max_i H(\omega_i) \leq H(\bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}}) + U, |\omega_i| = n_c \text{ for all } i.\end{aligned}\tag{2.2.2}$$

Proof of (i1): Any $\eta \in \bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$ has a single contour $\gamma(\eta)$ inside Λ^- of length $|\gamma(\eta)| = 4\ell_c$ and volume $|\eta \cap \Lambda^-| = \ell_c(\ell_c - 1) + 1 = n_c$, while $|\eta \cap \partial^- \Lambda| = 0$ (see Fig. 3). Thus, by (2.1.1), H is constant on $\bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$. Since $\bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}} \subseteq \bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$, this completes the proof of (i1).

Proof of (i2): Note that, because $\bar{\mathcal{Q}}$ and $\tilde{\mathcal{Q}}$ are connected via a U -path (disconnect the 1-protuberance and re-attach it to one of the neighboring sides of the $(\ell_c - 1) \times \ell_c$ quasi-square), we have

$$\mathcal{D} = \{\eta \in \mathcal{X} : \exists U\text{-path from } \bar{\mathcal{Q}} \text{ to } \eta\} = \{\eta \in \mathcal{X} : \exists U\text{-path from } \tilde{\mathcal{Q}} \text{ to } \eta\}.\tag{2.2.3}$$

First we prove that for any $\eta \in \bar{\mathcal{D}}$ there exists an $\omega: \bar{\mathcal{Q}} \rightarrow \eta$ such that $\max_i H(\omega_i) \leq H(\bar{\mathcal{Q}}) + U$ and $|\omega_i| = n_c$ for all i . We start the path from some $\zeta \in \bar{\mathcal{Q}}$. Then, recalling the labelling in Theorem 1.4.1, we have

- $\bar{k}_1(\zeta) = 1$ contained in $r_e(\zeta)$;
- $\bar{k}_2(\zeta) = \ell_c - 2$ contained in $r_n(\zeta)$;
- $\bar{k}_3(\zeta) = \bar{k}_4(\zeta) = \ell_c - 1$ contained in $r_w(\zeta) \cup c_{nw}(\zeta)$ and $r_s(\zeta) \cup c_{sw}(\zeta)$, respectively.

Here, without loss of generality, we assume that the 1-protuberance is attached to $r_e(\zeta)$ and proceed anti-clockwise. Using the mechanism described in Figs. 5 and 6, we move $\bar{k}_2(\zeta) - \bar{k}_2(\eta)$ particles from $r_n(\zeta)$ to $r_e(\zeta)$, one by one. After that we move $\bar{k}_3(\zeta) - \bar{k}_3(\eta) + \bar{k}_4(\zeta) - \bar{k}_4(\eta)$ particles from $r_s(\zeta) \cup c_{sw}(\zeta)$ to $r_e(\zeta)$. Finally, we move $\bar{k}_3(\zeta) - \bar{k}_3(\eta)$ particles from $r_w(\zeta) \cup c_{nw}(\zeta)$ to $r_s(\zeta) \cup c_{sw}(\zeta)$. The result is a configuration $\eta \in \bar{\mathcal{D}}$.




FIG. 5. Translation of a bar on a side of a rectangle at cost U .




FIG. 6. Motion of a particle around a corner of a rectangle at cost U .

Next we prove that for any $\eta \in \tilde{\mathcal{D}}$ there exists an $\omega: \tilde{\mathcal{Q}} \rightarrow \eta$ such that $\max_i H(\omega_i) \leq H(\tilde{\mathcal{Q}}) + U$ and $|\omega_i \cap \Lambda| = n_c$ for all i . We start the path from some $\zeta \in \tilde{\mathcal{Q}}$. We have

- $\tilde{k}_1(\zeta) = 1$ contained in $r_e(\zeta)$;
- $\tilde{k}_2(\zeta) = \tilde{k}_4(\zeta) = \ell_c - 1$ contained in $r_n(\zeta)$ and $r_s(\zeta)$;
- $\tilde{k}_3(\zeta) = \ell_c - 1$ contained in $r_w(\zeta) \cup c_{nw}(\zeta) \cup c_{sw}(\zeta)$.

We move $\tilde{k}_2(\zeta) - \tilde{k}_2(\eta)$ particles from $r_n(\zeta)$ to $r_e(\zeta)$. After that we move $\tilde{k}_3(\zeta) - \tilde{k}_3(\eta) + \tilde{k}_4(\zeta) - \tilde{k}_4(\eta)$ particles from $r_s(\zeta) \cup c_{sw}(\zeta)$ to $r_e(\zeta)$. Finally, we move $\tilde{k}_3(\zeta) - \tilde{k}_3(\eta)$ particles from $r_w(\zeta) \cup c_{nw}(\zeta)$ to $r_s(\zeta) \cup c_{sw}(\zeta)$. The result is a configuration $\eta \in \tilde{\mathcal{D}}$. This completes the proof of (i2).

Proof of (ii): By (2.2.2), all configurations in $\bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$ are connected via a U -path. Since $\bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}} \subseteq \mathcal{D} \cap (\bar{\mathcal{D}} \cup \tilde{\mathcal{D}})$, in order to prove (ii) it suffices to show that $\bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$ cannot be exited via a U -path (recall (2.2.3)).

Call a path *clustering* if all the configurations in the path consist of a single cluster and no free particles. Below we will prove that for any $\eta \in \bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$ and any η' connected to η by a clustering U -path,

$$\begin{aligned} (a) \quad & \text{CR}(\eta') = \text{CR}(\eta), \\ (b) \quad & \eta' \supseteq \text{CR}^-(\eta). \end{aligned} \tag{2.2.4}$$

What (2.2.4) says is that neither $\bar{\mathcal{D}}$ nor $\tilde{\mathcal{D}}$ can be exited via a clustering U -path. From this in turn we deduce that for any $\eta \in \bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$ and any η' connected to η by a U -path we must have that $\eta' \in \bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$, which is what we want to prove. The argument for the latter goes as follows. Detaching a particle costs $2U$ unless the particle is a 1-protuberance, in which case the cost is U . The only configurations in $\bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$ having a 1-protuberance are those in $\bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}}$ (recall the remarks made below Theorem 1.4.1). If we detach the 1-protuberance from a configuration in $\bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}}$, at cost U , then we obtain an $(\ell_c - 1) \times \ell_c$ quasi-square plus a free particle. Since now only moves at zero cost are allowed, only the free particle can move. Since in a U -path the particle number is conserved, the only way to regain U and complete the U -path is to re-attach the free particle to the quasi-square, in which case we return to $\bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}}$.

REMARK: Note that the motion of particles along the border a droplet may shift the droplet. Indeed, from any configuration in $\bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}}$ the 1-protuberance may detach itself and re-attach itself to a different side of the quasi-square or rectangle (recall Fig. 3). Thus, the U -path may shift the protocritical droplet to anywhere in Λ^- .

Proof of (a): Starting from any $\eta \in \mathcal{X}$, it is geometrically impossible to modify $\text{CR}(\eta)$ without detaching a particle.

Proof of (b): Fix $\eta \in \bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$. The proof is done in two steps.

1. Let us first consider clustering U -paths along which we do not move a particle from $\text{CR}^-(\eta)$. Along such paths we only encounter configurations in $\bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$ or configurations obtained from $\bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$ by breaking one of the bars in $\partial^-\text{CR}(\eta)$ into two pieces, at cost U (because there is no particle outside $\text{CR}(\eta)$ that can help to lower the cost). From the latter only moves at zero cost are possible, so no particle can be detached, and the only way to regain U and complete the U -path is to restore a bar.

2. Let us next consider clustering U -paths along which we move a particle from a corner of $\text{CR}^-(\eta)$. This move costs $2U$, which exceeds U . The overshoot U must be regained by letting the particle slide next to a bar that is attached to a side of $\text{CR}^-(\eta)$ (see Fig. 7). Since there are never two bars attached to the same side, we can at most gain U . This is why it is not possible to move a particle from $\text{CR}^-(\eta)$ other than from a corner.

From here only moves at zero cost are allowed. There are no 1-protuberances present anymore, because only the configurations in $\bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}}$ have a 1-protuberance. Thus, no particle outside $\text{CR}^-(\eta)$ can move, except the one that just detached itself from $\text{CR}^-(\eta)$. This particle can move back, in which case we return to the same configuration η . In fact, all possible moves at zero cost consist in moving the “hole” just created in $\text{CR}^-(\eta)$ along the side of $\text{CR}^-(\eta)$, until it reaches the height of the top of the bar attached to this side of $\text{CR}^-(\eta)$, after which it cannot advance anymore at zero cost (see Fig. 7). All these moves do not change the energy, except the one that returns the particle to its original position and regains U .




FIG. 7. Creation and motion of the hole at cost 0.

This proves our claim in (2.2.4), completes the proof of (ii) in (2.2.1), and hence of Theorem 1.4.1.

We saw above that U -paths cannot exit $\mathcal{D} = \bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$, but can make a crossover between $\bar{\mathcal{D}}$ and $\tilde{\mathcal{D}}$. This crossover can, however, only occur between $\bar{\mathcal{Q}}$ and $\tilde{\mathcal{Q}}$. A schematic picture of \mathcal{D} therefore is:




FIG. 8. Dumb-bell shape of $\mathcal{D} = \bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$ for U -paths.

2.3 Structure of the communication level set

2.3.1 Optimal paths

Most of this section is based on a line of argument in den Hollander, Nardi, Olivieri and Scoppola [8]. We repeat this argument here because it is vital for understanding the rest of the paper.

We begin by giving a precise description of $(\square \rightarrow \blacksquare)_{opt}$, the set of *optimal paths for the nucleation* (recall Definition 1.4.2(a)).

Proposition 2.3.1 (den Hollander, Olivieri, and Scoppola [9], Proposition 4.24)

(i) $\Phi(\square, \blacksquare) = \Gamma^*$.

(ii) $\mathcal{S}(\square, \blacksquare) \supseteq \mathcal{C}^*$.

Proof. The proof is different from that in [9].

(i) We prove that $\Phi(\square, \blacksquare) \leq \Gamma^*$ and $\Phi(\square, \blacksquare) \geq \Gamma^*$.

$\Phi(\square, \blacksquare) \leq \Gamma^*$: All we need to do is to construct a path that connects \square and \blacksquare without exceeding energy Γ^* . This is done in three steps.

1. We first show that the configurations in \mathcal{Q} are connected to \square by a path that stays below Γ^* .

Lemma 2.3.2 *For any $\eta^{1pr} \in \mathcal{Q}$ there exists an $\omega: \eta^{1pr} \rightarrow \square$ such that $\max_{\xi \in \omega} H(\xi) < \Gamma^*$.*

Proof. Fix $\eta^{1pr} \in \mathcal{Q}$. Note that, by (1.3.12), we have $H(\eta^{1pr}) = \Gamma^* - \Delta$. First, we detach the 1-protuberance from the $(\ell_c - 1) \times \ell_c$ quasi-square, which costs U and raises the energy to $\Gamma^* - \Delta + U (< \Gamma^*)$, move the particle to the boundary of the box, which costs nothing, and move it out of the box, which pays Δ . We are then left with a quasi-square of energy

$$\Gamma^* - 2\Delta + U. \quad (2.3.1)$$

Second, we detach a particle from a corner of the quasi-square, which costs $2U$, and move it out of the box, which pays Δ . Thus, the energy increases by $2U - \Delta$ when detaching and removing a particle from a corner of the quasi-square. We repeat this operation another $\ell_c - 3$ times, each time picking particles from the bar on the same shortest side. To guarantee that we never reach energy Γ^* , we have the condition that

$$(2U - \Delta)k + 2U < 2\Delta - U \quad \text{for } 0 \leq k \leq \ell_c - 3, \quad (2.3.2)$$

or

$$3 \leq \ell_c < \frac{U}{2U - \Delta} + 1. \quad (2.3.3)$$

The second inequality holds by the definition of ℓ_c in (1.3.3), the first inequality by our exclusion of $\ell_c = 2$ (recall the statement made just prior to Theorem 1.4.1). Third, detaching the last particle costs U instead of $2U$. To guarantee that we still do not reach energy Γ^* , we have the condition that

$$(2U - \Delta)(\ell_c - 2) + U < 2\Delta - U, \quad (2.3.4)$$

which is weaker than (2.3.2) because $2U - \Delta < U$. Removal of the last particle pays Δ , so that we arrive at energy

$$(\Gamma^* - 2\Delta + U) + (2U - \Delta)(\ell_c - 2) + (U - \Delta) = \Gamma^* - 2\Delta + (2U - \Delta)(\ell_c - 1), \quad (2.3.5)$$

which is strictly smaller than (2.3.1) by the second inequality in (2.3.3). Thus, removal of a row of length $\ell_c - 1$ from the $(\ell_c - 1) \times \ell_c$ quasi-square in $\eta^{1pr} \in \mathcal{Q}$ lowers the energy (see Fig. 9).




FIG. 9. Cost of adding or removing a row of length ℓ .

We now have a square of side length $\ell_c - 1$. It is obvious that we can remove further rows without encountering new conditions, until we reach \square . \heartsuit

2. For $\eta^{1pr} \in \mathcal{Q}$, let η^{2pr} be the configuration obtained from η^{1pr} by attaching an extra particle next to the 1-protuberance, thereby forming a 2-protuberance. We next show that η^{2pr} is connected to \blacksquare by a path that stays below Γ^* .

Lemma 2.3.3 *For any $\eta^{1pr} \in \mathcal{Q}$ there exists an $\omega: \eta^{2pr} \rightarrow \blacksquare$ such that $\max_{\xi \in \omega} H(\xi) < \Gamma^*$.*

Proof. Fix $\eta^{1pr} \in \mathcal{Q}$. Note that $H(\eta^{2pr}) = \Gamma^* - 2U$. First, we create a particle, which costs Δ and raises the energy to $\Gamma^* - 2U + \Delta (< \Gamma^*)$, move it to the droplet, which costs nothing, and attach it next to the 2-protuberance, which pays $2U$, thereby forming a bar of length 3. This operation pays $2U - \Delta$. We can repeat this operation another $\ell_c - 3$ times until the row is filled. By that time we have a square of side length ℓ_c and energy

$$\Gamma^* - 2U - (2U - \Delta)(\ell_c - 2). \quad (2.3.6)$$

Second, we create another particle and attach it anywhere to the square to form a new 1-protuberance. This operation costs $\Delta - U$. We must make sure that we can still create a particle without reaching energy Γ^* , which gives us the condition

$$(\Delta - U) + \Delta < 2U + (2U - \Delta)(\ell_c - 2), \quad (2.3.7)$$

or

$$\ell_c > \frac{U}{2U - \Delta}, \quad (2.3.8)$$

which holds by the definition of ℓ_c and the non-degeneracy hypothesis in (1.3.4). Third, we create another particle and attach it next to the new 1-protuberance. This brings us to energy

$$\Gamma^* - U - (2U - \Delta)\ell_c, \quad (2.3.9)$$

which is below the energy of η^{2pr} by (2.3.8). It is obvious that we can add further rows without encountering new conditions, until we reach \blacksquare . \heartsuit

3. We can now conclude the proof of $\Phi(\square, \blacksquare) \leq \Gamma^*$ by constructing a bridge between η^{1pr} and η^{2pr} that does not exceed Γ^* . Namely, create a particle at the boundary, which costs Δ and raises the energy to Γ^* , move it to the droplet, which costs nothing, and place it next to the 1-protuberance, which pays $2U$. The desired path $\omega: \square \rightarrow \blacksquare$ is realized by tracing the path in Lemma 2.3.2 in the reverse direction, back from \square to η^{1pr} , going over the bridge from η^{1pr} to η^{2pr} , and then following the path in Lemma 2.3.3 from η^{2pr} to \blacksquare . This ω will be called the *reference path through η for the nucleation*.

$\Phi(\square, \blacksquare) \geq \Gamma^*$: The proof comes in three steps.

1. The first crucial ingredient in the proof is the following observation:

Lemma 2.3.4 *Any $\omega \in (\square \rightarrow \blacksquare)_{\text{opt}}$ must pass through a configuration consisting of a single $(\ell_c - 1) \times \ell_c$ quasi-square somewhere in Λ^- .*

Proof. Any path $\omega: \square \rightarrow \blacksquare$ must cross the set $\mathcal{V}_{\ell_c(\ell_c-1)}$. As shown in Alonso and Cerf [1], Theorem 2.6, in $\mathcal{V}_{\ell_c(\ell_c-1)}$ the unique (modulo translations and rotations) configuration of minimal energy is the $(\ell_c - 1) \times \ell_c$ quasi-square, which we denote by η and which has energy

$$H(\eta) = \Gamma^* - 2\Delta + U. \quad (2.3.10)$$

All other configurations in $\mathcal{V}_{\ell_c(\ell_c-1)}$ have energy at least $\Gamma^* - 2\Delta + 2U$. To increase the particle number starting from any such configuration, we must create a particle at cost Δ . But the resulting configuration would have energy $\Gamma^* - \Delta + 2U (> \Gamma^*)$ and thus would lead to a path exceeding energy Γ^* . \heartsuit

2. The second crucial ingredient in the proof is the following observation:

Lemma 2.3.5 *Any $\omega \in (\square \rightarrow \blacksquare)_{opt}$ must pass through \mathcal{Q} .*

Proof. Follow the path until it hits the set $\mathcal{V}_{\ell_c(\ell_c-1)}$. According to Lemma 2.3.4, the configuration in this set must be an $(\ell_c - 1) \times \ell_c$ quasi-square. Since we need not consider any paths that return to the set $\mathcal{V}_{\ell_c(\ell_c-1)}$ afterwards, a first step beyond the quasi-square must be the creation of a new particle. This brings us to energy

$$\Gamma^* - \Delta + U. \quad (2.3.11)$$

Before any new particle is created, we must lower the energy by at least U . The obviously only possible way to do this is to move the particle to the quasi-square and attach it to one of its sides, which reduces the energy to

$$\Gamma^* - \Delta \quad (2.3.12)$$

and gives us a configuration in \mathcal{Q} . \heartsuit

3. It now suffices to show that to reach \blacksquare from \mathcal{Q} we must reach energy Γ^* . This goes as follows. Starting from \mathcal{Q} , it is impossible to reduce the energy without lowering the particle number. Indeed, this follows from Alonso and Cerf [1], Theorem 2.6, which asserts that the minimal energy in $\mathcal{V}_{\ell_c(\ell_c-1)+1}$ is realized (although not uniquely) by the configurations in \mathcal{Q} . Since any further move to increase the particle number involves the creation of a new particle, the energy must reach Γ^* .

This completes the proof of Proposition 2.3.1(i).

(ii) Our final observation is the following:

Lemma 2.3.6 *The set of configurations in $\mathcal{V}_{\ell_c(\ell_c-1)+1}$ that can be reached from \square by a path that stays below Γ^* and for which it is possible to add a particle without exceeding Γ^* coincides with the set \mathcal{D} defined in Definition 1.3.2(b).*

Proof. From step 2 above it is clear that the definition of \mathcal{D} precisely assures that the assertion holds true. Indeed, by Lemma 2.3.5, any $\omega \in (\square \rightarrow \blacksquare)_{opt}$ crosses $\mathcal{V}_{\ell_c(\ell_c-1)+1}$ in \mathcal{Q} . Once it is in \mathcal{Q} , before the arrival of the next particle, which costs Δ , it can reach all configurations that have the same energy, the same particle number, and can be reached at cost $\leq U < \Delta$. \heartsuit

By adding a particle to a configuration in \mathcal{D} we arrive in $\mathcal{C}^* = \mathcal{D}^{fp}$, the set defined in Definition 1.3.2(c). This completes the proof of Proposition 2.3.1(ii). \heartsuit

We conclude the following:

Proposition 2.3.7 Any $\omega \in (\square \rightarrow \blacksquare)_{\text{opt}}$ passes first through \mathcal{Q} , then possibly through $\mathcal{D} \setminus \mathcal{Q}$, and finally through \mathcal{C}^* .

Proof. Combine Lemmas 2.3.5–2.3.6 and Proposition 2.3.1(i). ♡

2.3.2 Motion on \mathcal{C}^*

The next proposition will be important later on.

Proposition 2.3.8 (i) Starting from $\mathcal{C}^* \setminus \mathcal{Q}^{fp}$, the only transitions that do not raise the energy are motions of the free particle, as long as the free particle is at lattice distance ≥ 3 from the protocritical droplet.

(ii) Starting from \mathcal{Q}^{fp} , the only transitions that do not raise the energy are motions of the free particle and motions of the 1-protuberance along the side of the quasi-square where it is attached, as long as the free particle is at lattice distance ≥ 3 from the protocritical droplet. When the lattice distance is 2, either the free particle can be attached to the protocritical droplet or the 1-protuberance can be detached from the protocritical droplet and attached to the free particle, to form a quasi-square plus a dimer. From the latter configuration the only transition that does not raise the energy is the reverse move.

(iii) Starting from \mathcal{C}^* , the only configurations that can be reached by a path that lowers the energy and does not decrease the particle number are those where the free particle is attached to the protocritical droplet.

Proof. Obvious. The restriction in (i) that the free particle must be at lattice distance ≥ 3 from the protocritical droplet is needed for the following reason: If the protocritical droplet is a configuration in $\mathcal{D} \setminus \mathcal{Q}$ and the free particle sits at lattice distance 2 from a corner of a bar, diagonally opposite the particle that sits in the corner of the bar, then at zero cost this particle may detach itself from the bar and slide inbetween the quasi-square and the free particle. For (iii) note the following: if we start from the configuration described above and slide the remaining particles in the bar one by one, all at zero cost except the last one, which pays U , then we reach a configuration where the free particle is attached to the protocritical droplet with the bar shifted. ♡

For $\eta \in \mathcal{C}^*$, we write $\eta = (\hat{\eta}, x)$ with $\hat{\eta} \in \mathcal{D}$ the protocritical droplet and $x \in \Lambda$ the position of the free particle. Let us denote the configurations that can be reached from $\eta = (\hat{\eta}, x)$ according to Proposition 2.3.8(iii) by

- $\mathcal{C}^G(\hat{\eta})$ if the particle is attached in $\partial^- \text{CR}(\hat{\eta})$.
- $\mathcal{C}^B(\hat{\eta})$ if the particle is attached in $\partial^+ \text{CR}(\hat{\eta})$,




Fig. 10. Good sites (G) and bad sites (B) for $\ell_c = 14$.

Let

$$C^G = \cup_{\hat{\eta} \in \mathcal{D}} C^G(\hat{\eta}), \quad C^B = \cup_{\hat{\eta} \in \mathcal{D}} C^B(\hat{\eta}). \quad (2.3.13)$$

The next proposition shows that when we reach \mathcal{C}^G we have made it “over the hill”, while when we reach \mathcal{C}^B we have not.

Proposition 2.3.9 (i) If $\eta \in C^G$, then there exists an $\omega: \eta \rightarrow \blacksquare$ such that $\max_{\xi \in \omega} H(\xi) < \Gamma^*$.
(ii) If $\eta \in C^B$, then there are no $\omega: \eta \rightarrow \square$ or $\omega: \eta \rightarrow \blacksquare$ such that $\max_{\xi \in \omega} H(\xi) < \Gamma^*$.

Proof. (i) If $\eta \in C^G$, then its energy is either $\Gamma^* - 2U$ or $\Gamma^* - U$, depending on whether the particle was attached in a corner or as a 1-protuberance. In the latter case we can move the particle at no cost into a corner and gain an extra $-U$. After that it is possible to create a new particle and re-attach it, which leads to energy $\Gamma^* - 2U - (2U - \Delta)$. We can continue in this way, filling up all rows in $\partial^- \text{CR}(\eta)$, until we reach either an $\ell_c \times \ell_c$ square or an $(\ell_c - 1) \times (\ell_c + 1)$ rectangle, depending on whether η arose from $\bar{\mathcal{D}}$ or $\bar{\mathcal{D}}$ (recall Fig. 3). In the first case we can proceed along the reference path for the nucleation constructed in the proof of Proposition 2.3.1. In the latter case, however, we can connect to this reference path as follows. The energy of the $(\ell_c - 1) \times (\ell_c + 1)$ rectangle is $\Gamma^* - 2U - (2U - \Delta)(\ell_c - 3)$. This is lower than $\Gamma^* - \Delta$, because $\ell_c \geq 3$. Create a particle, which costs Δ , and attach it to one of the longest sides of the rectangle, which pays U . Now slide particles along the corner of the rectangle, following the mechanism described in Figs. 6 and 7, until an $\ell_c \times \ell_c$ square is reached. This costs U and keeps the energy below Γ^* . From there again proceed along the reference path for the nucleation.

(ii) If $\eta \in C^B$, then $H(\eta) = \Gamma^* - U$, so as long as the energy stays below Γ^* it is impossible to create a new particle before further lowering the energy. But there are no moves available to lower the energy. The only moves available are those where the particle that was last attached is moving along the side or is detached again, which brings us back to \mathcal{C}^* , or those starting a motion of particles along the border of the droplet (as in Fig. 6), which may or may not bring us back to \mathcal{C}^* . In both cases the cost is U and the energy returns to Γ^* .




Fig. 11. An example of a path from \mathcal{C}^B to \blacksquare .

An example of a path from \mathcal{C}^B to \blacksquare that does not return to \mathcal{C}^* is obtained as follows (see Fig. 11). Suppose that $\hat{\eta} \in \mathcal{D}$ is such that one bar completes one side of $\partial^- \text{CR}(\hat{\eta})$, and suppose that the free particle attaches itself on top of that bar, forming a 1-protuberance. Then the energy is $\Gamma^* - U$. Slide this bar to the end of the side it is attached to (at cost and gain U) and slide the two bars on the neighboring sides to the end as well (at cost and gain U). Then the energy is again $\Gamma^* - U$. Now move the shorter bar on top of the longer bar via a motion as in Fig. 6. When the last particle of the bar is moved, it can be detached (at cost U) and re-attached (at gain $2U$). Then the energy is $\Gamma^* - 2U$. Now create a free particle (at cost Δ), move it to the droplet (at cost 0), and attach it in a corner of the droplet (at gain $2U$). Continue “downhill” in this way, adding on successive rows as in the reference path that was used above, until \blacksquare is reached. \heartsuit

Proposition 2.3.9(ii) shows that the configurations in \mathcal{C}^B are *wells*, i.e., their energy is $< \Gamma^*$, but to move to either \square or \blacksquare the energy must return to Γ^* . The configurations of the form “quasi-square plus dimer” described in (ii) in the proof of Proposition 2.3.8 are elements of $\mathcal{S}(\square, \blacksquare)$ but not of \mathcal{C}^* . Indeed, the only possible move at zero cost is the one where the free particle jumps back to the quasi-square. Thus, we see that

- \mathcal{C}^* is a union of *plateaus*, index by $\hat{\eta} \in \mathcal{D}$; each plateau consists of a protocritical droplet $\hat{\eta}$ and a collection of positions of the free particle, indexed by $\Lambda \setminus (\hat{\eta} \cup \partial^+ \hat{\eta})$; each plateau has *wells* and *dead-ends* when the free particle is close to the protocritical droplet.

This property is special for Kawasaki dynamics. We will not attempt to describe the wells and dead-ends in full detail. For our sharp asymptotics of the average nucleation time we will not need this detail.

2.3.3 Graph structure of the energy landscape

Let us summarize what we have shown so far:

Theorem 2.3.10 *View \mathcal{X} as a graph whose vertices are configurations and whose edges connect communicating configurations. Let*

- \mathcal{X}^* be the subgraph of \mathcal{X} obtained by removing all vertices η with $H(\eta) > \Gamma^*$ and all edges incident to these vertices;
- \mathcal{X}^{**} be the subgraph of \mathcal{X}^* obtained by removing all vertices η with $H(\eta) = \Gamma^*$ and all edges incident to these vertices;
- \mathcal{X}_{\square} and $\mathcal{X}_{\blacksquare}$ be the connected components of \mathcal{X}^{**} containing \square and \blacksquare , respectively.

Then

- (i) $\mathcal{X}_{\square} \neq \mathcal{X}_{\blacksquare}$, and so \mathcal{X}_{\square} and $\mathcal{X}_{\blacksquare}$ are disconnected in \mathcal{X}^{**} .
- (ii) $\mathcal{D} \subseteq \mathcal{X}_{\square}$, $\mathcal{C}^G \subseteq \mathcal{X}_{\blacksquare}$, $\mathcal{C}^B \subseteq \mathcal{X}^{**} \setminus (\mathcal{X}_{\square} \cup \mathcal{X}_{\blacksquare})$.

Propositions 2.3.7–2.3.9 and Theorem 2.3.10 will play a crucial role in Section 3.3, where we derive sharp estimates for the average nucleation time. We will see that they are in fact all that is needed for these estimates.

2.4 Two global geometric facts

In Sections 2.2–2.3 we have analysed the geometry of the configurations on and incident to \mathcal{C}^* that are relevant for the nucleation. This will be sufficient for the computation of the average nucleation time. To make full use of the results of Bovier, Eckhoff, Gaynard, and Klein [5], we must establish two further facts, both concerning the global geometry of the energy landscape.

Proposition 2.4.1 below shows that there are no valleys in the energy landscape whose depth equals or exceeds the communication height between \square , \blacksquare .

Proposition 2.4.1 *For all $\eta \in \mathcal{X} \setminus \{\square, \blacksquare\}$,*

$$\Phi(\eta, \{\square, \blacksquare\}) - H(\eta) < \Gamma^* = \Phi(\square, \blacksquare). \quad (2.4.1)$$

Proof. This is the analogue of Proposition 3.4.6 in den Hollander, Nardi, Olivieri, and Scoppola [8] for three dimensions. The proof can be carried over to two dimensions verbatim. \heartsuit

Proposition 2.4.2 below shows that \square is a proper metastable configuration because it lies at the bottom of its valley:

Proposition 2.4.2 *If $\eta \in \mathcal{X} \setminus \square$ is such that*

$$\Phi(\eta, \square) \leq \Phi(\eta, \blacksquare), \quad (2.4.2)$$

then $H(\eta) > 0$.

Proof. Recall that $n_c = \ell_c(\ell_c - 1) + 1$. Define

$$\mathcal{V}_{\leq n_c} = \bigcup_{0 \leq n \leq n_c} \mathcal{V}_n, \quad \mathcal{V}_{> n_c} = \mathcal{X} \setminus \mathcal{V}_{\leq n_c}. \quad (2.4.3)$$

First, we claim that if η satisfies (2.4.2) and $H(\eta) \leq 0$, then $\eta \in \mathcal{V}_{\leq n_c}$. Indeed, since $\Phi(\eta, \{\square, \blacksquare\}) = \Phi(\eta, \square) \wedge \Phi(\eta, \blacksquare)$, it follows from (2.4.1–2.4.2) that $\Phi(\eta, \square) < \Gamma^* + H(\eta)$. So, if $H(\eta) \leq 0$, then $\Phi(\eta, \square) < \Gamma^*$. But in the proof of Proposition 2.3.1(i) we have shown that $\Phi(\eta, \square) \geq \Gamma^*$ for all $\eta \in \mathcal{V}_{> n_c}$ (n_c is the volume of the clusters in \mathcal{D}).

Second, we claim that \square is the only configuration in $\mathcal{V}_{\leq n_c}$ with zero energy, while all other configurations have strictly positive energy. Indeed, inserting the isoperimetric inequality

$$|\eta \cap \Lambda^-| \leq \left(\frac{\gamma(\eta)}{4} \right)^2 \quad \forall \eta \neq \square \quad (2.4.4)$$

into (2.1.1), we get

$$\begin{aligned} H(\eta) &\geq \frac{U}{2} |\gamma(\eta)| - (2U - \Delta) |\eta \cap \Lambda^-| \\ &\geq \frac{U}{2} 4 \sqrt{|\eta \cap \Lambda^-|} - (2U - \Delta) |\eta \cap \Lambda^-| \\ &= (2U - \Delta) \sqrt{|\eta \cap \Lambda^-|} \left(2 \frac{U}{2U - \Delta} - \sqrt{|\eta \cap \Lambda^-|} \right) \\ &> (2U - \Delta) \sqrt{|\eta \cap \Lambda^-|} \left(2(\ell_c - 1) - \sqrt{\ell_c(\ell_c - 1) + 1} \right) \\ &> (2U - \Delta) \sqrt{|\eta \cap \Lambda^-|} (\ell_c - 1) > 0. \end{aligned} \quad (2.4.5)$$

♡

3 Average nucleation time in two dimensions

In this section we analyze the average nucleation time. Section 3.1 recalls the definition of Dirichlet form and capacity, and provides an a priori estimate for capacities between arbitrary sets. Section 3.2 shows that $\{\square, \blacksquare\}$ is a metastable pair in the proper sense, and provides the link between the average nucleation time and the capacity of the pair $\{\square, \blacksquare\}$. Section 3.3 contains the proof of Theorem 1.4.4 in two steps: (1) a priori estimates of the equilibrium potential associated with the capacity of the pair $\{\square, \blacksquare\}$; (2) reduction of the Dirichlet form for this capacity to one involving simple random walk. Section 3.4 gives the proof of Theorem 1.4.5, Section 3.5 of Theorem 1.4.3.

3.1 Dirichlet form and capacity

In the proof of Theorem 1.4.4, a key role is played by the *Dirichlet form*

$$\mathcal{E}_\beta(h) = \frac{1}{2} \sum_{\eta, \eta' \in \mathcal{X}} \mu_\beta(\eta) c_\beta(\eta, \eta') [h(\eta) - h(\eta')]^2, \quad h: \mathcal{X} \rightarrow [0, 1], \quad (3.1.1)$$

where μ_β is the Gibbs measure defined in (1.1.5) and c_β are the transition rates of the Kawasaki dynamics defined in (1.2.6). Given two non-empty disjoint sets $\mathcal{A}, \mathcal{B} \subseteq \mathcal{X}$, the *capacity* of the pair \mathcal{A}, \mathcal{B} is defined by

$$\text{CAP}_\beta(\mathcal{A}, \mathcal{B}) = \min_{\substack{h: \mathcal{X} \rightarrow [0, 1] \\ h|_{\mathcal{A}} \equiv 1, h|_{\mathcal{B}} \equiv 0}} \mathcal{E}_\beta(h), \quad (3.1.2)$$

where $h|_{\mathcal{A}} \equiv 1$ means that $h(\eta) = 1$ for all $\eta \in \mathcal{A}$ and $h|_{\mathcal{B}} \equiv 0$ means that $h(\eta) = 0$ for all $\eta \in \mathcal{B}$. The right-hand side of (3.1.2) has a unique minimizer $h_{\mathcal{A}, \mathcal{B}}^*$, called the *equilibrium potential* of the pair \mathcal{A}, \mathcal{B} , given by

$$h_{\mathcal{A}, \mathcal{B}}^*(\eta) = \mathbb{P}_\eta(\tau_{\mathcal{A}} < \tau_{\mathcal{B}}), \quad \eta \in \mathcal{X} \setminus (\mathcal{A} \cup \mathcal{B}) \quad (3.1.3)$$

(recall (1.3.10)). This is the solution of the equation

$$\begin{aligned} (c_\beta h)(\eta) &= 0, & \eta &\in \mathcal{X} \setminus (\mathcal{A} \cup \mathcal{B}), \\ h(\eta) &= 1, & \eta &\in \mathcal{A}, \\ h(\eta) &= 0, & \eta &\in \mathcal{B}, \end{aligned} \tag{3.1.4}$$

with $(c_\beta h)(\eta) = \sum_{\eta' \in \mathcal{X}} c_\beta(\eta, \eta') h(\eta')$. Moreover,

$$\text{CAP}_\beta(\mathcal{A}, \mathcal{B}) = \sum_{\eta \in \mathcal{A}} \mu_\beta(\eta) c_\beta(\eta, \mathcal{X} \setminus \eta) \mathbb{P}_\eta(\tau_{\mathcal{B}} < \tau_{\mathcal{A}}) \tag{3.1.5}$$

with $c_\beta(\eta, \mathcal{X} \setminus \eta) = \sum_{\eta' \in \mathcal{X} \setminus \eta} c_\beta(\eta, \eta')$ the rate of moving out of η . This rate enters because $\tau_{\mathcal{A}}$ is the first hitting time of \mathcal{A} after the initial configuration is left (recall (1.3.10)). Note from (1.2.7) and (3.1.1–3.1.2) that

$$\text{CAP}_\beta(\mathcal{A}, \mathcal{B}) = \text{CAP}_\beta(\mathcal{B}, \mathcal{A}). \tag{3.1.6}$$

The following elementary estimate will be important. Here $\Phi(\mathcal{A}, \mathcal{B}) = \min_{\eta \in \mathcal{A}, \eta' \in \mathcal{B}} \Phi(\eta, \eta')$ is the communication height between the pair \mathcal{A}, \mathcal{B} .

Lemma 3.1.1 *For every non-empty disjoint $\mathcal{A}, \mathcal{B} \subseteq \mathcal{X}$ there exist constants $0 < C_1 \leq C_2 < \infty$ (depending on \mathcal{A}, \mathcal{B}) such that for all β ,*

$$C_1 \leq e^{\beta \Phi(\mathcal{A}, \mathcal{B})} Z_\beta \text{CAP}_\beta(\mathcal{A}, \mathcal{B}) \leq C_2. \tag{3.1.7}$$

Proof. The proof uses basic properties of communication heights.

Upper bound: The upper bound is obtained from (3.1.2) by picking $h = 1_{K(\mathcal{A}, \mathcal{B})}$ with

$$K(\mathcal{A}, \mathcal{B}) = \{\eta \in \mathcal{X} : \Phi(\eta, \mathcal{A}) \leq \Phi(\eta, \mathcal{B})\}. \tag{3.1.8}$$

The key observation is that if $\eta \rightarrow \eta'$ is a transition from $K(\mathcal{A}, \mathcal{B})$ to $\mathcal{X} \setminus K(\mathcal{A}, \mathcal{B})$, then

$$\begin{aligned} (1) \quad & H(\eta') < H(\eta), \\ (2) \quad & H(\eta) \geq \Phi(\mathcal{A}, \mathcal{B}). \end{aligned} \tag{3.1.9}$$

To see (1), suppose that $H(\eta') \geq H(\eta)$. Clearly,

$$H(\eta') \geq H(\eta) \iff \Phi(\eta', \mathcal{F}) = \Phi(\eta, \mathcal{F}) \vee H(\eta') \quad \forall \mathcal{F} \subseteq \mathcal{X}. \tag{3.1.10}$$

But $\eta \in K(\mathcal{A}, \mathcal{B})$ tells us that $\Phi(\eta, \mathcal{A}) \leq \Phi(\eta, \mathcal{B})$, hence $\Phi(\eta', \mathcal{A}) \leq \Phi(\eta', \mathcal{B})$ by (3.1.10), and hence $\eta' \in K(\mathcal{A}, \mathcal{B})$, which is a contradiction.

To see (2), note that (1) implies the reverse of (3.1.10):

$$H(\eta) \geq H(\eta') \iff \Phi(\eta, \mathcal{F}) = \Phi(\eta', \mathcal{F}) \vee H(\eta) \quad \forall \mathcal{F} \subseteq \mathcal{X}. \tag{3.1.11}$$

Trivially, $\Phi(\eta, \mathcal{B}) \geq H(\eta)$. We claim that equality holds. Indeed, suppose that equality fails. Then we get

$$H(\eta) < \Phi(\eta, \mathcal{B}) = \Phi(\eta', \mathcal{B}) < \Phi(\eta', \mathcal{A}) = \Phi(\eta, \mathcal{A}), \tag{3.1.12}$$

where the two equalities come from (3.1.11), while the second inequality uses that $\eta' \in \mathcal{X} \setminus K(\mathcal{A}, \mathcal{B})$. Thus, $\Phi(\eta, \mathcal{A}) > \Phi(\eta, \mathcal{B})$, which contradicts $\eta \in K(\mathcal{A}, \mathcal{B})$. From $\Phi(\eta, \mathcal{B}) = H(\eta)$ we obtain $\Phi(\mathcal{A}, \mathcal{B}) \leq \Phi(\mathcal{A}, \eta) \vee \Phi(\eta, \mathcal{B}) = \Phi(\eta, \mathcal{B}) = H(\eta)$, which proves (2).

Combining (3.1.9) with (1.1.5), (1.2.6) and (1.2.7), we find that

$$\mu_\beta(\eta)c_\beta(\eta, \eta') \leq \frac{1}{Z_\beta} e^{-\beta\Phi(\mathcal{A}, \mathcal{B})} \quad \forall \eta \in K(\mathcal{A}, \mathcal{B}), \eta' \in \mathcal{X} \setminus K(\mathcal{A}, \mathcal{B}). \quad (3.1.13)$$

Hence

$$\text{CAP}_\beta(\mathcal{A}, \mathcal{B}) \leq \mathcal{E}_\beta(1_{K(\mathcal{A}, \mathcal{B})}) \leq C_2 \frac{1}{Z_\beta} e^{-\beta\Phi(\mathcal{A}, \mathcal{B})} \quad (3.1.14)$$

with $C_2 = |\{(\eta, \eta') : \eta \in K(\mathcal{A}, \mathcal{B}), \eta' \in \mathcal{X} \setminus K(\mathcal{A}, \mathcal{B})\}|$.

Lower bound: The lower bound is obtained by picking any path $\omega = (\omega_0, \omega_1, \dots, \omega_K)$ that realizes the minimax in $\Phi(\mathcal{A}, \mathcal{B})$ and ignore all the transitions that are not in this path, i.e.,

$$\text{CAP}_\beta(\mathcal{A}, \mathcal{B}) \geq \min_{\substack{h: \omega \rightarrow [0,1] \\ h(\omega_0)=1, h(\omega_K)=0}} \mathcal{E}_\beta^\omega(h), \quad (3.1.15)$$

where the Dirichlet form \mathcal{E}_β^ω is defined as \mathcal{E}_β in (3.1.1) but with \mathcal{X} replaced by ω . Due to the one-dimensional nature of the set ω , the variational problem in the right-hand side can be solved explicitly by elementary computations. One finds that the minimum equals

$$M = \left[\sum_{k=0}^{K-1} \frac{1}{\mu_\beta(\omega_k)c_\beta(\omega_k, \omega_{k+1})} \right]^{-1}, \quad (3.1.16)$$

and is uniquely attained at h given by

$$h(\omega_k) = M \sum_{l=0}^{k-1} \frac{1}{\mu_\beta(\omega_l)c_\beta(\omega_l, \omega_{l+1})}, \quad k = 0, 1, \dots, K. \quad (3.1.17)$$

We thus have

$$\begin{aligned} \text{CAP}_\beta(\mathcal{A}, \mathcal{B}) &\geq M \\ &\geq \frac{1}{K} \min_{k=0,1,\dots,K-1} \mu_\beta(\omega_k)c_\beta(\omega_k, \omega_{k+1}) \\ &= \frac{1}{K} \frac{1}{Z_\beta} \min_{k=0,1,\dots,K-1} e^{-\beta[H(\omega_k) \vee H(\omega_{k+1})]} \\ &= C_1 \frac{1}{Z_\beta} e^{-\beta\Phi(\mathcal{A}, \mathcal{B})} \end{aligned} \quad (3.1.18)$$

with $C_1 = 1/K$. ♡

Lemma 3.1.1 is a typical a priori bound for capacities. In particular, the use of one-dimensional subgraphs is a tool that with little effort produces rough estimates, which can be lifted to sharp estimates with some more effort, as we will see later on.

3.2 Metastable pair, link between average nucleation time and capacity

In Bovier, Eckhoff, Gaynard, and Klein [5] metastability is defined in terms of properties of capacities, namely:

Definition 3.2.1 *Consider a family of Markov chains, indexed by β , on a finite state space \mathcal{X} . A set $\mathcal{M} \subseteq \mathcal{X}$ is called metastable if*

$$\lim_{\beta \rightarrow \infty} \frac{\max_{\eta \notin \mathcal{M}} \mu_\beta(\eta) [\text{CAP}_\beta(\eta, \mathcal{M})]^{-1}}{\min_{\eta \in \mathcal{M}} \mu_\beta(\eta) [\text{CAP}_\beta(\eta, \mathcal{M} \setminus \eta)]^{-1}} = 0. \quad (3.2.1)$$

For our model we have:

Lemma 3.2.2 *The set $\{\square, \blacksquare\}$ is metastable in the sense of Definition 3.2.1.*

Proof. The numerator in (3.2.1) can be bounded above by $e^{(\Gamma^* - \delta)\beta}/C_1$, via Proposition 2.4.1 and Lemma 3.1.1. The denominator, on the other hand, can be bounded below by $e^{\Gamma^*\beta}/C_2$ (the minimum being attained at \square). Therefore the ratio is bounded above by $e^{-\delta\beta}(C_2/C_1)$. \heartsuit

Lemma 3.2.2 allows us to apply the theory in Bovier, Eckhoff, Gaynard, and Klein [5]. To obtain our sharp estimate of $\mathbb{E}_\square(\tau_\blacksquare)$, we will use the following key relation:

Proposition 3.2.3 $\mathbb{E}_\square(\tau_\blacksquare) = \frac{1}{Z_\beta \text{CAP}_\beta(\square, \blacksquare)} [1 + o(1)]$ as $\beta \rightarrow \infty$.

Proof. Bovier, Eckhoff, Gaynard, and Klein [5], Theorem 1.3(i), written in our notation, states that

$$\mathbb{E}_\square(\tau_\blacksquare) = \frac{\mu_\beta(\mathcal{R}_\square)}{\text{CAP}_\beta(\square, \blacksquare)} [1 + o(1)] \quad \beta \rightarrow \infty, \quad (3.2.2)$$

where

$$\mathcal{R}_\square = \{\eta \in \mathcal{X} : \mathbb{P}_\eta(\tau_\square < \tau_\blacksquare) \geq \mathbb{P}_\eta(\tau_\blacksquare < \tau_\square)\}. \quad (3.2.3)$$

It follows from the proof of Lemma 3.3.1 below that, for large enough β ,

$$\{\eta \in \mathcal{X} : \Phi(\eta, \square) < \Phi(\eta, \blacksquare)\} \subseteq \mathcal{R}_\square \subseteq \{\eta \in \mathcal{X} : \Phi(\eta, \square) \leq \Phi(\eta, \blacksquare)\} \quad (3.2.4)$$

and hence, via Proposition 2.4.2,

$$\min_{\eta \in \mathcal{R}_\square \setminus \square} H(\eta) > H(\square) = 0. \quad (3.2.5)$$

This in turn implies that $\mu_\beta(\mathcal{R}_\square)/\mu_\beta(\square) = 1 + o(1)$. Since $\mu_\beta(\square) = 1/Z_\beta$, we get the claim. \heartsuit

Proposition 3.2.3 shows that the computation of $\mathbb{E}_\square(\tau_\blacksquare)$ revolves around getting sharp bounds on $Z_\beta \text{CAP}_\beta(\square, \blacksquare)$. From Lemma 3.1.1 we know that $C_1 \leq e^{\beta\Gamma^*} Z_\beta \text{CAP}_\beta(\square, \blacksquare) \leq C_2$. In what follows we narrow down the constants.

3.3 Average nucleation time: Proof of Theorem 1.4.4

In this section we will show how to turn the geometric information obtained in Theorem 2.3.10 into sharp bounds on $Z_\beta \text{CAP}_\beta(\square, \blacksquare)$. We follow the general strategy outlined in Bovier and Manzo [6] and Bovier [3]:

- Note that all terms in the Dirichlet form in (3.1.1) involving configurations η with $H(\eta) > \Gamma^*$, i.e., $\eta \in \mathcal{X} \setminus \mathcal{X}^*$, contribute at most $Ce^{-(\Gamma^* + \delta)\beta}$ for some $\delta > 0$ and can be neglected. Thus, effectively we can replace \mathcal{X} by \mathcal{X}^* .
- Show that $h_{\square, \blacksquare}^* = O(e^{-\delta\beta})$ on \mathcal{X}_\blacksquare and $h_{\square, \blacksquare}^* = 1 - O(e^{-\delta\beta})$ on \mathcal{X}_\square for some $\delta > 0$.
- Prove sharp upper and lower bounds for $h_{\square, \blacksquare}^*$ on $\mathcal{X}^* \setminus (\mathcal{X}_\square \cup \mathcal{X}_\blacksquare)$ in terms of a variational problem involving only the vertices and the bonds on and incident to $\mathcal{X}^* \setminus (\mathcal{X}_\square \cup \mathcal{X}_\blacksquare)$.

The last two steps are carried out in Sections 3.3.1–3.3.2. We identify the resulting variational problem with capacities associated with simple random walk. In Section 3.4 we analyse the asymptotics of these capacities for large Λ .

3.3.1 A priori estimates on the equilibrium potential

Note that

$$\begin{aligned}\mathcal{X}_\square &= \{\eta \in \mathcal{X}^*: \Phi(\eta, \square) < \Phi(\eta, \blacksquare)\}, \\ \mathcal{X}_\blacksquare &= \{\eta \in \mathcal{X}^*: \Phi(\eta, \blacksquare) < \Phi(\eta, \square)\}.\end{aligned}\tag{3.3.1}$$

The guiding idea behind the sharp estimate of $Z_\beta \text{CAP}_\beta(\square, \blacksquare)$ is that $h_{\square, \blacksquare}^*$ is exponentially close to 1 on \mathcal{X}_\square and exponentially close to 0 on \mathcal{X}_\blacksquare . This is the content of the following estimate, which will be needed later on.

Lemma 3.3.1 *There exist $C < \infty$ and $\delta > 0$ such that for all β ,*

$$\min_{\eta \in \mathcal{X}_\square} h_{\square, \blacksquare}^*(\eta) \geq 1 - Ce^{-\delta\beta}, \quad \max_{\eta \in \mathcal{X}_\blacksquare} h_{\square, \blacksquare}^*(\eta) \leq Ce^{-\delta\beta}.\tag{3.3.2}$$

Proof. A standard renewal argument gives the relations, valid for $\eta \notin \{\square, \blacksquare\}$,

$$\mathbb{P}_\eta(\tau_\blacksquare < \tau_\square) = \frac{\mathbb{P}_\eta(\tau_\blacksquare < \tau_{\square \cup \eta})}{1 - \mathbb{P}_\eta(\tau_{\square \cup \blacksquare} > \tau_\eta)}, \quad \mathbb{P}_\eta(\tau_\square < \tau_\blacksquare) = \frac{\mathbb{P}_\eta(\tau_\square < \tau_{\blacksquare \cup \eta})}{1 - \mathbb{P}_\eta(\tau_{\square \cup \blacksquare} > \tau_\eta)}.\tag{3.3.3}$$

For $\eta \in \mathcal{X}_\square \setminus \square$, we estimate

$$h_{\square, \blacksquare}^*(\eta) = 1 - \mathbb{P}_\eta(\tau_\blacksquare < \tau_\square) = 1 - \frac{\mathbb{P}_\eta(\tau_\blacksquare < \tau_{\square \cup \eta})}{\mathbb{P}_\eta(\tau_{\square \cup \blacksquare} < \tau_\eta)} \geq 1 - \frac{\mathbb{P}_\eta(\tau_\blacksquare < \tau_\eta)}{\mathbb{P}_\eta(\tau_\square < \tau_\eta)}\tag{3.3.4}$$

and, with the help of (3.1.5) and Lemma 3.1.1,

$$\frac{\mathbb{P}_\eta(\tau_\blacksquare < \tau_\eta)}{\mathbb{P}_\eta(\tau_\square < \tau_\eta)} = \frac{Z_\beta \text{CAP}_\beta(\eta, \blacksquare)}{Z_\beta \text{CAP}_\beta(\eta, \square)} \leq C(\eta) e^{-[\Phi(\eta, \blacksquare) - \Phi(\eta, \square)]\beta} \leq C(\eta) e^{-\delta\beta},\tag{3.3.5}$$

which proves the first claim with $C = \max_{\eta \in \mathcal{X}_\square \setminus \square} C(\eta)$. Note that $h_{\square, \blacksquare}^*(\square)$ is a convex combination of $h_{\square, \blacksquare}^*(\eta)$ with $\eta \in \mathcal{X}_\square \setminus \square$, namely, those η that communicate with \square . Hence the claim includes $\eta = \square$.

For $\eta \in \mathcal{X}_\blacksquare \setminus \blacksquare$, we estimate

$$h_{\square, \blacksquare}^*(\eta) = \mathbb{P}_\eta(\tau_\square < \tau_\blacksquare) = \frac{\mathbb{P}_\eta(\tau_\square < \tau_{\blacksquare \cup \eta})}{\mathbb{P}_\eta(\tau_{\square \cup \blacksquare} < \tau_\eta)} \leq \frac{\mathbb{P}_\eta(\tau_\square < \tau_\eta)}{\mathbb{P}_\eta(\tau_\blacksquare < \tau_\eta)}\tag{3.3.6}$$

and, with the help of (3.1.5) and Lemma 3.1.1,

$$\frac{\mathbb{P}_\eta(\tau_\square < \tau_\eta)}{\mathbb{P}_\eta(\tau_\blacksquare < \tau_\eta)} = \frac{Z_\beta \text{CAP}_\beta(\eta, \square)}{Z_\beta \text{CAP}_\beta(\eta, \blacksquare)} \leq C(\eta) e^{-[\Phi(\eta, \square) - \Phi(\eta, \blacksquare)]\beta} \leq C(\eta) e^{-\delta\beta},\tag{3.3.7}$$

which proves the second claim with $C = \max_{\eta \in \mathcal{X}_\blacksquare \setminus \blacksquare} C(\eta)$. \heartsuit

Knowing that $h_{\square, \blacksquare}^*$ is trivial on $\mathcal{X}_\square \cup \mathcal{X}_\blacksquare$, it remains to understand what $h_{\square, \blacksquare}^*$ looks like on the set

$$\mathcal{X}^* \setminus (\mathcal{X}_\square \cup \mathcal{X}_\blacksquare) = \{\eta \in \mathcal{X}^*: \Phi(\eta, \square) = \Phi(\eta, \blacksquare)\},\tag{3.3.8}$$

which separates \mathcal{X}_\square and \mathcal{X}_\blacksquare and contains $\mathcal{S}(\square, \blacksquare)$ (recall (1.3.9)). This will be carried out in Section 3.3.2.

Before doing so, we first show that $h_{\square, \blacksquare}^*$ is also trivial on $\mathcal{X}^{**} \setminus (\mathcal{X}_\square \cup \mathcal{X}_\blacksquare)$. This set can be partitioned into maximally connected components,

$$\mathcal{X}^{**} \setminus (\mathcal{X}_\square \cup \mathcal{X}_\blacksquare) = \bigcup_{i=1}^I \mathcal{X}_i,\tag{3.3.9}$$

where each \mathcal{X}_i is a *well* in $\mathcal{S}(\square, \blacksquare)$, i.e., a set of communicating configurations with energy $< \Gamma^*$ but with communication height Γ^* towards both \square and \blacksquare .

Lemma 3.3.2 *There exist $C < \infty$ and $\delta > 0$ such that for all $i = 1, \dots, I$ and all β ,*

$$\max_{\eta, \eta' \in \mathcal{X}_i} |h_{\square, \blacksquare}^*(\eta) - h_{\square, \blacksquare}^*(\eta')| \leq C e^{-\delta\beta}. \quad (3.3.10)$$

Proof. Fix $i = 1, \dots, I$ and $\eta, \eta' \in \mathcal{X}_i$. Estimate

$$h_{\square, \blacksquare}^*(\eta) = \mathbb{P}_\eta(\tau_\square < \tau_\blacksquare) \leq \mathbb{P}_\eta(\tau_\square < \tau_{\eta'}) + \mathbb{P}_\eta(\tau_{\eta'} < \tau_\square < \tau_\blacksquare). \quad (3.3.11)$$

First, as in the proof of Lemma 3.3.1, we have

$$\begin{aligned} \mathbb{P}_\eta(\tau_\square < \tau_{\eta'}) &= \frac{\mathbb{P}_\eta(\tau_\square < \tau_{\eta \cup \eta'})}{1 - \mathbb{P}_\eta(\tau_{\square \cup \eta'} > \tau_\eta)} \leq \frac{\mathbb{P}_\eta(\tau_\square < \tau_\eta)}{\mathbb{P}_\eta(\tau_{\eta'} < \tau_\eta)} \\ &= \frac{Z_{\beta \text{CAP}_\beta}(\eta, \square)}{Z_{\beta \text{CAP}_\beta}(\eta, \eta')} \leq C(\eta, \eta') e^{-[\Phi(\eta, \square) - \Phi(\eta, \eta')]\beta} \leq C(\eta, \eta') e^{-\delta\beta}, \end{aligned} \quad (3.3.12)$$

where we use that $\Phi(\eta, \square) = \Gamma^*$ and $\Phi(\eta, \eta') < \Gamma^*$. Second,

$$\mathbb{P}_\eta(\tau_{\eta'} < \tau_\square < \tau_\blacksquare) = \mathbb{P}_\eta(\tau_{\eta'} < \tau_{\square \cup \blacksquare}) \mathbb{P}_{\eta'}(\tau_\square < \tau_\blacksquare) \leq \mathbb{P}_{\eta'}(\tau_\square < \tau_\blacksquare) = h_{\square, \blacksquare}^*(\eta'). \quad (3.3.13)$$

Combining (3.3.11–3.3.13), we get

$$h_{\square, \blacksquare}^*(\eta) \leq C(\eta, \eta') e^{-\delta\beta} + h_{\square, \blacksquare}^*(\eta'). \quad (3.3.14)$$

Interchange η and η' to get the claim with $C = \max_i \max_{\eta, \eta' \in \mathcal{X}_i} C(\eta, \eta')$. \heartsuit

REMARK: We saw in Proposition 2.3.9(i) that for each $\hat{\eta} \in \mathcal{D}$ the four bars of bad sites in $\partial^+ \text{CR}(\hat{\eta})$ (see Fig. 10) each form a well. Lemma 3.3.2 shows that $h_{\square, \blacksquare}^*$ is close to a constant on each of these wells. These are not the only wells, but Lemma 3.3.2 shows that we not need care too much about wells anyway: *only the transitions in and out of the wells contribute to the Dirichlet form* at the order we are after, not those inside the wells. Later we shall see that we can even ignore the wells altogether, provided we are content with obtaining bounds. Indeed, in Proposition 2.3.8 we saw that the wells only occur when the free particle is at distance 2 from the protocritical droplet.

3.3.2 Reduction of the Dirichlet form

The reduction is done in two steps. First we reduce the full Dirichlet form to a Dirichlet form involving only the immediate vicinity of the communication level set.

Proposition 3.3.3 *There exists $\delta > 0$ such that for $\beta \rightarrow \infty$,*

$$Z_{\beta \text{CAP}_\beta}(\square, \blacksquare) = [1 + O(e^{-\delta\beta})] \Theta e^{-\Gamma^* \beta}, \quad (3.3.15)$$

where

$$\Theta = \min_{C_1, \dots, C_I} \min_{h|_{\mathcal{X}_\square} \equiv 1, h|_{\mathcal{X}_\blacksquare} \equiv 0, h|_{\mathcal{X}_i} \equiv C_i \forall i=1, \dots, I} \min_{h: \mathcal{X}^* \rightarrow [0,1]} \frac{1}{2} \sum_{\eta, \eta' \in \mathcal{X}^*} 1_{\{\eta \leftrightarrow \eta'\}} [h(\eta) - h(\eta')]^2. \quad (3.3.16)$$

Proof. First, recalling (1.1.5–1.1.6), (1.2.6) and (3.1.1–3.1.2), we have

$$\begin{aligned} Z_\beta \text{CAP}_\beta(\square, \blacksquare) &= Z_\beta \min_{\substack{h: \mathcal{X} \rightarrow [0,1] \\ h(\square)=1, h(\blacksquare)=0}} \frac{1}{2} \sum_{\eta, \eta' \in \mathcal{X}} \mu_\beta(\eta) c_\beta(\eta, \eta') [h(\eta) - h(\eta')]^2 \\ &= O\left(e^{-(\Gamma^* + \delta)\beta}\right) + Z_\beta \min_{\substack{h: \mathcal{X}^* \rightarrow [0,1] \\ h(\square)=1, h(\blacksquare)=0}} \frac{1}{2} \sum_{\eta, \eta' \in \mathcal{X}^*} \mu_\beta(\eta) c_\beta(\eta, \eta') [h(\eta) - h(\eta')]^2. \end{aligned} \quad (3.3.17)$$

Next, with the help of Lemmas 3.3.1–3.3.2, we get

$$\begin{aligned} &\min_{\substack{h: \mathcal{X}^* \rightarrow [0,1] \\ h(\square)=1, h(\blacksquare)=0}} \frac{1}{2} \sum_{\eta, \eta' \in \mathcal{X}^*} \mu_\beta(\eta) c_\beta(\eta, \eta') [h(\eta) - h(\eta')]^2 \\ &= \min_{\substack{h: \mathcal{X}^* \rightarrow [0,1] \\ h=h_{\square, \blacksquare}^* \text{ on } \mathcal{X}_{\square} \cup \mathcal{X}_{\blacksquare} \cup (\mathcal{X}_1, \dots, \mathcal{X}_I)}} \frac{1}{2} \sum_{\eta, \eta' \in \mathcal{X}^*} \mu_\beta(\eta) c_\beta(\eta, \eta') [h(\eta) - h(\eta')]^2 \\ &= [1 + O(e^{-\delta\beta})] \min_{C_1, \dots, C_I} \min_{\substack{h: \mathcal{X}^* \rightarrow [0,1] \\ h|_{\mathcal{X}_{\square}} \equiv 1, h|_{\mathcal{X}_{\blacksquare}} \equiv 0, h|_{\mathcal{X}_i} \equiv C_i \forall i=1, \dots, I}} \frac{1}{2} \sum_{\eta, \eta' \in \mathcal{X}^*} \mu_\beta(\eta) c_\beta(\eta, \eta') [h(\eta) - h(\eta')]^2, \end{aligned} \quad (3.3.18)$$

where the error term $O(e^{-\delta\beta})$ arises after we replace the approximate boundary conditions

$$h = \begin{cases} 1 - O(e^{-\delta\beta}) & \text{on } \mathcal{X}_{\square}, \\ O(e^{-\delta\beta}) & \text{on } \mathcal{X}_{\blacksquare}, \\ C_i + O(e^{-\delta\beta}) & \text{on } \mathcal{X}_i, i = 1, \dots, I, \end{cases} \quad (3.3.19)$$

by the sharp boundary conditions

$$h = \begin{cases} 1 & \text{on } \mathcal{X}_{\square}, \\ 0 & \text{on } \mathcal{X}_{\blacksquare}, \\ C_i & \text{on } \mathcal{X}_i, i = 1, \dots, I. \end{cases} \quad (3.3.20)$$

Finally, by (1.1.5–1.1.6) and (1.2.6–1.2.7) we have

$$Z_\beta \mu_\beta(\eta) c_\beta(\eta, \eta') = 1_{\{\eta \leftrightarrow \eta'\}} e^{-\Gamma^* \beta} \text{ for all } \eta, \eta' \in \mathcal{X}^* \text{ that are not both in } \mathcal{X}_{\square} \text{ or both in } \mathcal{X}_{\blacksquare} \text{ or both in } \mathcal{X}_i \text{ for some } i = 1, \dots, I. \quad (3.3.21)$$

Indeed, in each of these cases either $H(\eta) = \Gamma^* > H(\eta')$ or $H(\eta) < \Gamma^* = H(\eta')$, because there are no direct transitions between \mathcal{X}_{\square} , $\mathcal{X}_{\blacksquare}$ and \mathcal{X}_i , $i = 1, \dots, I$ (use Proposition 2.3.10(i) and recall the decomposition in (3.3.9)). Combining (3.3.17–3.3.18) and (3.3.21), we arrive at the claim. \heartsuit

Next we estimate Θ in terms of capacities associated with simple random walk.

Proposition 3.3.4 $\Theta \in [\Theta_1, \Theta_2]$ with

$$\begin{aligned} \Theta_1 &= \sum_{\hat{\eta} \in \mathcal{D}} \text{CAP}^{\Lambda^+}(\partial^+ \Lambda, \text{CR}(\hat{\eta})), \\ \Theta_2 &= \sum_{\hat{\eta} \in \mathcal{D}} \text{CAP}^{\Lambda^+}(\partial^+ \Lambda, \text{CR}^{++}(\hat{\eta})), \end{aligned} \quad (3.3.22)$$

where $\text{CR}^{++} = (\text{CR}^+)^+$ and

$$\text{CAP}^{\Lambda^+}(\partial^+ \Lambda, F) = \min_{\substack{g: \Lambda^+ \rightarrow [0,1] \\ g|_{\partial^+ \Lambda} \equiv 1, g|_F \equiv 0}} \frac{1}{2} \sum_{\substack{x, x' \in \Lambda^+ \\ x \sim x'}} [g(x) - g(x')]^2, \quad F \subseteq \Lambda^+, \quad (3.3.23)$$

and $x \sim x'$ means that x and x' are nearest-neighbor sites.

Proof. The variational problem in (3.3.16) decomposes into disjoint variational problems for the maximally connected components of \mathcal{X}^* . Only those components that contain \mathcal{X}_\square or \mathcal{X}_\blacksquare contribute, since for the other components the minimum is achieved by picking h constant.

$\Theta \geq \Theta_1$: The lower bound is obtained from (3.3.16) by removing all transitions that do not involve a protocritical droplet and a free particle that is moving. This removal gives

$$\Theta \geq \sum_{\hat{\eta} \in \mathcal{D}} \min_{C_j(\hat{\eta}), j=1,2,3,4} \min_{\substack{g: \Lambda^+ \rightarrow [0,1] \\ g|_{\partial^G \hat{\eta}} \equiv 0, g|_{\partial^B \hat{\eta}} \equiv C_j(\hat{\eta}), j=1,2,3,4, g|_{\partial^+ \Lambda} \equiv 1}} \frac{1}{2} \sum_{\substack{x, x' \in \Lambda^+ \\ x \sim x'}} [g(x) - g(x')]^2, \quad (3.3.24)$$

where $\partial^G \hat{\eta}$ denotes the set of good sites in $\partial^- \text{CR}(\hat{\eta})$ and $\partial_j^B \hat{\eta}$, $j = 1, 2, 3, 4$, denote the four bars of bad sites in $\partial^+ \text{CR}(\hat{\eta})$ (see Fig. 10). To see how this bound arises from (3.3.16), pick

$$h(\eta) = h(\hat{\eta}, x) = g(x), \quad \hat{\eta} \in \mathcal{D}, x \in \Lambda^+ \setminus \hat{\eta}, \quad (3.3.25)$$

and use Proposition 2.3.10(ii) to match the boundary conditions in (3.3.16) (recall the decomposition in (3.3.9)). Note that $x \in \partial^+ \Lambda$ in $\eta = (\hat{\eta}, x)$ corresponds to $\eta \in \mathcal{D}$ (i.e., the free particle at x is outside Λ), while $x \in \partial^+ \hat{\eta}$ corresponds to $\eta \in \mathcal{C}^G(\hat{\eta}) \cup \mathcal{C}^B(\hat{\eta})$. The right-hand side of (3.3.24) may be further bounded below by Θ_1 , because the latter has less stringent boundary conditions.

$\Theta \leq \Theta_2$: The upper bound is obtained from (3.3.16) by picking $C_i = 0$, $i = 1, \dots, I$, and

$$h(\eta) = \begin{cases} 1 & \text{for } \eta \in \mathcal{X}_\square, \\ g(x) & \text{for } \eta \in \mathcal{C}^{++}, \\ 0 & \text{for } \eta \in \mathcal{X}^* \setminus [\mathcal{X}_\square \cup \mathcal{C}^{++}], \end{cases} \quad (3.3.26)$$

where

$$\mathcal{C}^{++} = \{\eta = (\hat{\eta}, x): \hat{\eta} \in \mathcal{D}, x \in \Lambda \setminus \text{CR}^{++}(\hat{\eta})\}. \quad (3.3.27)$$

This choice satisfies the boundary conditions in (3.3.16), because

$$\mathcal{C}^{++} \subseteq \mathcal{C}^* \quad \text{and} \quad \mathcal{C}^* \cap [\mathcal{X}_\blacksquare \cup (\cup_{i=1}^I \mathcal{X}_i)] = \emptyset. \quad (3.3.28)$$

By Proposition 2.3.10(ii), $\mathcal{D} \subseteq \mathcal{X}_\square$, so that $h(\eta) = 1$ for $\eta = (\hat{\eta}, x)$ with $\hat{\eta} \in \mathcal{D}$ and $x \in \partial^+ \Lambda$, which is consistent with the boundary condition $g|_{\partial^+ \Lambda} \equiv 1$ in (3.3.23). Moreover, $h(\eta) = 0$ for $\eta = (\hat{\eta}, x)$ with $\hat{\eta} \in \mathcal{D}$ and $x \in \text{CR}^{++}(\hat{\eta})$, which is consistent with the boundary condition $g|_{\text{CR}^{++}(\hat{\eta})} \equiv 0$ in (3.3.23) with $F = \text{CR}^{++}(\hat{\eta})$. Note that, by Proposition 2.3.7, the only transitions in \mathcal{X}^* between \mathcal{X}_\square and \mathcal{C}^* are those where a free particle is entering at $\partial^- \Lambda$. Hence, there are no transitions between \mathcal{X}_\square and $\mathcal{X}^* \setminus [\mathcal{X}_\square \cup \mathcal{C}^{++}]$. Also note that, by Proposition 2.3.8(i–ii), the only transitions in \mathcal{X}^* between \mathcal{C}^{++} and $\mathcal{X}^* \setminus [\mathcal{X}_\square \cup \mathcal{C}^{++}]$ are those where the free particle moves to distance 1 from the protocritical droplet. Thus, (3.3.23) includes all the relevant transitions. \heartsuit

Propositions 3.3.3–3.3.4 complete the proof of the first half of Theorem 1.4.4, with K identified as $K = 1/\Theta$ with Θ defined in (3.3.16) and bounded in (3.3.22). The second half, i.e., the exponential limit law in (1.4.7), follows from Bovier, Eckhoff, Gaynard, and Klein [5], Theorem 1.3(iv).

The capacity defined in (3.3.23) is the capacity of the pair $\{\partial^+ \Lambda, D\}$ for continuous-time simple random walk on Λ^+ where transitions between sites occur at rate 1. In Section 3.4 we will show that Θ_1 and Θ_2 have the same asymptotics for $\Lambda \rightarrow \mathbb{Z}^2$.

3.4 Capacity asymptotics: Proof of Theorem 1.4.5

As $\Lambda \rightarrow \mathbb{Z}^2$, the capacities $\text{CAP}^{\Lambda^+}(\partial^+ \Lambda, \text{CR}(\hat{\eta}))$ and $\text{CAP}^{\Lambda^+}(\partial^+ \Lambda, \text{CR}^{++}(\hat{\eta}))$ tend to zero in a way that depends neither on the shape of the protocritical droplet $\hat{\eta}$ nor on its location in Λ , provided it is far from $\partial^+ \Lambda$:

Lemma 3.4.1 *Write $\Lambda = B_M = [-M, +M]^2 \cap \mathbb{Z}^2$. For any $\varepsilon > 0$,*

$$\lim_{M \rightarrow \infty} \max_{\substack{\hat{\eta} \in \mathcal{D} \\ d(\partial^+ B_M, \hat{\eta}) \geq \varepsilon M}} \left| \frac{\log M}{2\pi} \text{CAP}^{B_M^+}(\partial^+ B_M, \text{CR}(\hat{\eta})) - 1 \right| = 0 \quad (3.4.1)$$

and

$$\lim_{M \rightarrow \infty} \max_{\substack{\hat{\eta} \in \mathcal{D} \\ d(\partial^+ B_M, \hat{\eta}) \geq \varepsilon M}} \left| \frac{\log M}{2\pi} \text{CAP}^{B_M^+}(\partial^+ B_M, \text{CR}^{++}(\hat{\eta})) - 1 \right| = 0, \quad (3.4.2)$$

where $d(\partial^+ B_M, \hat{\eta}) = \min\{|x - y| : x \in \partial^+ B_M, y \in \hat{\eta}\}$.

Proof. Let us first prove (3.4.1). For $\hat{\eta} \in \mathcal{D}$, let $y \in \text{CR}(\hat{\eta}) \subseteq B_M$ denote the site closest to the center of $\text{CR}(\hat{\eta})$. The capacity decreases when we enlarge the set over which the Dirichlet form is minimized. Therefore we have

$$\begin{aligned} \text{CAP}^{B_M^+}(\partial^+ B_M, \text{CR}(\hat{\eta})) &\geq \text{CAP}^{B_M^+}(\partial^+ B_M, y) \\ &= \text{CAP}^{B_M^+ - y}(\partial^+(B_M - y), 0) \geq \text{CAP}^{B_{2M}^+}(\partial^+ B_{2M}, 0). \end{aligned} \quad (3.4.3)$$

By the analogue of (3.1.5–3.1.6), we have (compare (3.3.23) with (3.1.1–3.1.2))

$$\begin{aligned} \text{CAP}^{B_{2M}^+}(\partial^+ B_{2M}, 0) &= \text{CAP}^{B_{2M}^+}(0, \partial^+ B_{2M}) \\ &= \left(\sum_{x \in B_{2M}^+} 1_{\{x \sim 0\}} \right) \mathbb{P}_0(\tau_{\partial^+ B_{2M}} < \tau_0) = 4 \mathbb{P}_0(\tau_{\partial^+ B_{2M}} < \tau_0), \end{aligned} \quad (3.4.4)$$

where \mathbb{P}_0 is the law on path space of the *discrete-time* simple random walk on \mathbb{Z}^2 starting at 0. According to Révész [12], Lemma 22.1, we have

$$\mathbb{P}_0(\tau_{\partial^+ B_{2M}} < \tau_0) \sim \frac{\pi}{2 \log(2M)} \quad M \rightarrow \infty. \quad (3.4.5)$$

This proves the desired lower bound.

Similarly, by monotonicity we have

$$\text{CAP}^{B_M^+}(\partial^+ B_M, \text{CR}(\hat{\eta})) \leq \text{CAP}^{B_M^+}(\partial^+ B_M, S_{\ell_c}(y)) \leq \text{CAP}^{B_{\varepsilon M}^+}(\partial^+ B_{\varepsilon M}, S_{\ell_c}(0)), \quad (3.4.6)$$

where $S_{\ell_c}(y)$ is the $\ell_c \times \ell_c$ square or $(\ell_c - 1) \times (\ell_c + 1)$ rectangle centered at y containing $\text{CR}(\hat{\eta})$, and the last inequality uses that $d(\partial^+ B_M, \hat{\eta}) \geq \varepsilon M$. By the recurrence of simple random walk, we have

$$\text{CAP}^{B_{\varepsilon M}^+}(\partial^+ B_{\varepsilon M}, S_{\ell_c}(0)) \sim \text{CAP}^{B_{\varepsilon M}^+}(\partial^+ B_{\varepsilon M}, 0) \quad M \rightarrow \infty. \quad (3.4.7)$$

Therefore the desired upper bound follows from (3.4.5).

The proof of (3.4.2) is similar. ♡

Combining (3.3.22) and Lemma 3.4.1, we find

$$\begin{aligned}
\Theta_1 &= O(\varepsilon M) + \sum_{\substack{\hat{\eta} \in \mathcal{D} \\ d(\partial^+ B_M, \hat{\eta}) \geq \varepsilon M}} \text{CAP}^{B_M^+}(\partial^+ B_M, \text{CR}(\hat{\eta})) \\
&= O(\varepsilon M) + \frac{2\pi}{\log M} |\{\hat{\eta} \in \mathcal{D} : d(\partial^+ B_M, \hat{\eta}) \geq \varepsilon M\}| [1 + o(1)] \\
&= O(\varepsilon M) + \frac{2\pi}{\log M} N(\ell_c) [(1 - \varepsilon)M]^2 [1 + o(1)]
\end{aligned} \tag{3.4.8}$$

and the same expression for Θ_2 (recall that $N(\ell_c)$ is defined to be the cardinality of \mathcal{D} modulo shifts). Let $M \rightarrow \infty$ followed by $\varepsilon \downarrow 0$, to conclude that $\Theta \sim (2\pi/\log M)N(\ell_c)M^2$. Since $|\Lambda| = M^2$ and $K = 1/\Theta$, this proves the claim in Theorem 1.4.5 after we prove the formula for $N(\ell_c)$ stated in (1.4.10). This is done in Lemmas 3.4.2–3.4.3 below.

REMARK: The asymptotics in Lemma 3.4.1 shows that $\Theta \sim 4 \sum_{x \in \Lambda} \mathbb{P}_x(\tau_{\partial^+ \Lambda} < \tau_x)$ as $\Lambda \rightarrow \mathbb{Z}^2$ (recall (3.4.5)). In van den Berg [2] this sum is studied in more detail and for more general domains than the square box Λ .

Lemma 3.4.2 $|\bar{\mathcal{D}}| = \frac{1}{6}(\ell_c - 1)\ell_c(\ell_c + 1)(\ell_c + 2)$.

Proof. We have to count how many different shapes the clusters in $\bar{\mathcal{D}}$ can take on (recall Fig. 3). Return to Theorem 1.4.1. We will do the counting by starting from an $\ell_c \times \ell_c$ square and counting in how many ways $\ell_c - 1$ particles can be removed from the four bars. We will split the counting according to the number $k = 1, 2, 3, 4$ of corner particles that are removed.

$k = 1$: There are 4 choices for the one corner. Let m_{1+}, m_{1-} denote the number of particles that are removed in the two directions away from the corner. Then $m_{1+}, m_{1-} \geq 0$ and $m_{1+} + m_{1-} = \ell_c - 2$. There are $\ell_c - 1$ ways to choose these. Therefore the contribution to $|\bar{\mathcal{D}}|$ is $4(\ell_c - 1)$.

$k = 2$: There are 6 choices for the two corners. Let m_{1+}, m_{1-} and m_{2+}, m_{2-} denote the number of particles that are removed in the two directions away from the two corners. Then $m_{1+}, m_{1-}, m_{2+}, m_{2-} \geq 0$ and $m_{1+} + m_{1-} + m_{2+} + m_{2-} = \ell_c - 3$. There are $(\ell_c - 1)(\ell_c - 2)$ ways to choose these. Therefore the contribution to $|\bar{\mathcal{D}}|$ is $6(\ell_c - 1)(\ell_c - 2)$.

$k = 3$: There are 4 choices for the three corners. A similar argument as above shows that there are $\frac{1}{2}(\ell_c - 1)(\ell_c - 2)(\ell_c - 3)$ ways to remove $\ell_c - 4$ particles in the two directions away from the three corners. Therefore the contribution to $|\bar{\mathcal{D}}|$ is $2(\ell_c - 1)(\ell_c - 2)(\ell_c - 3)$.

$k = 4$: There is 1 choice for the four corners. There are $\frac{1}{6}(\ell_c - 1)(\ell_c - 2)(\ell_c - 3)(\ell_c - 4)$ ways to remove $\ell_c - 5$ particles in the two directions away from the four corners. Therefore the contribution to $|\bar{\mathcal{D}}|$ is $\frac{1}{6}(\ell_c - 1)(\ell_c - 2)(\ell_c - 3)(\ell_c - 4)$.

Sum the contributions to get the claim. ♡

Lemma 3.4.3 $|\tilde{\mathcal{D}}| = \frac{1}{6}(\ell_c - 2)(\ell_c - 1)\ell_c(\ell_c + 1)$.

Proof. Similar. Start from an $(\ell_c - 1) \times (\ell_c + 1)$ rectangle and count in how many ways $\ell_c - 2$ particles can be removed from the four bars. The answer is the same as in Lemma 3.4.2 with $\ell_c - 1$ replaced by $\ell_c - 2$. ♡

It follows from Lemmas 3.4.2–3.4.3 that $N(\ell_c) = |\mathcal{D}| = |\bar{\mathcal{D}}| + |\tilde{\mathcal{D}}| = \frac{1}{3}(\ell_c - 1)\ell_c^2(\ell_c + 1)$, as claimed in (1.4.10).

3.5 Gate for the nucleation: Proof of Theorem 1.4.3

(i) We saw in Proposition 2.3.8(ii) that the configuration consisting of an $(\ell_c - 1) \times \ell_c$ quasi-square plus a dimer at distance 1 is a dead-end in $\mathcal{S}(\square, \blacksquare)$. Therefore $\mathcal{S}(\square, \blacksquare) \supsetneq \mathcal{G}(\square, \blacksquare)$, which is the first part of Theorem 1.4.3(i).

To prove the second part of Theorem 1.4.3(i), we must exhibit a configuration in $\mathcal{S}(\square, \blacksquare)$ that is not in \mathcal{C}^* . For that we return to the proof of Proposition 2.3.9(ii), where we exhibited a path from \mathcal{C}^B to \blacksquare that does not exceed energy Γ^* and avoids \mathcal{C}^* . The configurations with energy Γ^* visited by this path are elements of $\mathcal{S}(\square, \blacksquare)$.

For $\eta \in \mathcal{C}^*$, let $\hat{\eta} \in \mathcal{D}$ be the configuration obtained from η by removing the free particle. For $A \subseteq \Lambda$ and $x \in \Lambda$, let $d(x, A)$ denote the lattice distance between x and A . Let Λ_4 be Λ without its four corners. Define, recursively,

$$B_1(\hat{\eta}) = \{x \in \Lambda_4 : x \notin \hat{\eta}, d(x, \hat{\eta}) = 1\} \quad (3.5.1)$$

and

$$\begin{aligned} B_2(\hat{\eta}) &= \{x \in \Lambda_4 : x \notin \hat{\eta}, d(x, B_1(\hat{\eta})) = 1\}, \\ \bar{B}_2(\hat{\eta}) &= B_2(\hat{\eta}), \end{aligned} \quad (3.5.2)$$

and

$$\begin{aligned} B_3(\hat{\eta}) &= \{x \in \Lambda_4 : x \notin B_1(\hat{\eta}), d(x, B_2(\hat{\eta})) = 1\}, \\ \bar{B}_3(\hat{\eta}) &= B_3(\hat{\eta}) \cup [\bar{B}_2(\hat{\eta}) \cap \partial^- \Lambda_4], \end{aligned} \quad (3.5.3)$$

and, for $i = 4, 5, \dots, L - \ell_c$ (with L the side length of Λ),

$$\begin{aligned} B_i(\hat{\eta}) &= \{x \in \Lambda_4 : x \notin B_{i-2}(\hat{\eta}), d(x, B_{i-1}(\hat{\eta})) = 1\}, \\ \bar{B}_i(\hat{\eta}) &= B_i(\hat{\eta}) \cup [\bar{B}_{i-1}(\hat{\eta}) \cap \partial^- \Lambda_4]. \end{aligned} \quad (3.5.4)$$

In words, $B_1(\hat{\eta})$ is the ring of sites in Λ_4 at distance 1 from $\hat{\eta}$, while $\bar{B}_i(\hat{\eta})$ is the ring of sites in Λ_4 at distance i from $\hat{\eta}$ plus all the sites in $\partial^- \Lambda_4$ at distance $1 < j < i$ from $\hat{\eta}$ ($i = 2, 3, \dots, L - \ell_c$). Note that, depending on the location of $\hat{\eta}$ in Λ , the $\bar{B}_i(\hat{\eta})$ coincide for large enough i . The maximal number of rings is $L - \ell_c$.

The following sets are minimal gates:

$$\mathcal{C}_i^* = \{(\hat{\eta}, x) : \hat{\eta} \in \mathcal{D}, x \in \bar{B}_i(\hat{\eta})\}, \quad i = 2, 3, \dots, L - \ell_c. \quad (3.5.5)$$

(Note that the sites in $B_1(\hat{\eta}) \cap \partial^- \Lambda_4$ are not in any minimal gate.) This may be seen as follows. First, \mathcal{C}_i^* is a gate. Indeed, any $\omega \in (\square \rightarrow \blacksquare)_{opt}$ enters \mathcal{C}^* through a configuration of the form $(\hat{\eta}, x)$, with $\hat{\eta} \in \mathcal{D}$ a protocritical droplet and x a free particle such that: $x \in B_i(\hat{\eta})$ when $d(\partial^- \Lambda_4, \hat{\eta}) > i$ and $x \in \bar{B}_i(\hat{\eta})$ when $d(\partial^- \Lambda_4, \hat{\eta}) \leq i$. Second, \mathcal{C}_i^* is a minimal gate. Indeed, for any $\eta \in \mathcal{C}_i^*$ there exists an $\omega \in (\square \rightarrow \blacksquare)_{opt}$ that avoids $\mathcal{C}_i^* \setminus \eta$, namely, any ω that enters \mathcal{C}_i^* at $\eta = (\hat{\eta}, x)$ and proceeds by moving the free particle at x inside $B_{i-1}(\hat{\eta})$ and from there towards $\hat{\eta}$, where it attaches itself ‘properly’ (i.e., in a corner). Note that the \mathcal{C}_i^* are not disjoint.

The union of the \mathcal{C}_i^* is the set \mathcal{C}_{min}^* that was announced at the end of Section 1.4.3. Clearly, $\mathcal{C}_{min}^* \subseteq \mathcal{G}(\square, \blacksquare)$. The configuration exhibited in the proof of the second part of Theorem 1.4.3(i) is an element of $\mathcal{G}(\square, \blacksquare)$ but not of \mathcal{C}_{min}^* .

(ii) We will show that there exist $\delta > 0$ and $C < \infty$ such that for all β ,

$$\mathbb{P}_{\square}(\tau_{\square} < \tau_{\mathcal{C}^*} < \tau_{\blacksquare} | \tau_{\blacksquare} < \tau_{\square}) \geq 1 - Ce^{-\delta\beta}, \quad (3.5.6)$$

which implies (1.4.4). The proof goes as follows. Recall (1.3.10).

By (3.1.5), $\text{CAP}_\beta(\square, \blacksquare) = \mu_\beta(\square) c_\beta(\square, \mathcal{X} \setminus \square) \mathbb{P}_\square(\tau_\blacksquare < \tau_\square)$ with $\mu_\beta(\square) = 1/Z_\beta$. From the lower bound in Lemma 3.1.1 it therefore follows that

$$\mathbb{P}_\square(\tau_\blacksquare < \tau_\square) \geq C_1 e^{-\Gamma^* \beta} \frac{1}{c_\beta(\square, \mathcal{X} \setminus \square)}. \quad (3.5.7)$$

We will show that

$$\mathbb{P}_\square(\{\tau_\mathcal{Q} < \tau_{\mathcal{C}^*} < \tau_\blacksquare\}^c, \tau_\blacksquare < \tau_\square) \leq C_2 e^{-(\Gamma^* + \delta)\beta} \frac{1}{c_\beta(\square, \mathcal{X} \setminus \square)}. \quad (3.5.8)$$

Combining (3.5.7–3.5.8), we get (3.5.6) with $C = C_2/C_1$.

In Proposition 2.3.7 we saw that any path from \square to \blacksquare that does not pass first through \mathcal{Q} and then through \mathcal{C}^* must pass the set $\mathcal{V}_{\ell_c(\ell_c-1)+2} \supseteq \mathcal{S}(\square, \blacksquare)$ in a configuration η with $H(\eta) > \Gamma^*$. Therefore there exists a set \mathcal{S} , with $H(\eta) \geq \Gamma^* + \delta$ for all $\eta \in \mathcal{S}$ and some $\delta > 0$, such that

$$\mathbb{P}_\square(\{\tau_\mathcal{Q} < \tau_{\mathcal{C}^*} < \tau_\blacksquare\}^c, \tau_\blacksquare < \tau_\square) \leq \mathbb{P}_\square(\tau_\mathcal{S} < \tau_\square). \quad (3.5.9)$$

Now estimate, with the help of reversibility (use (1.2.7) or (3.1.5–3.1.6)),

$$\begin{aligned} \mathbb{P}_\square(\tau_\mathcal{S} < \tau_\square) &\leq \sum_{\eta \in \mathcal{S}} \mathbb{P}_\square(\tau_\eta < \tau_\square) = \sum_{\eta \in \mathcal{S}} \frac{\mu_\beta(\eta) c_\beta(\eta, \mathcal{X} \setminus \eta)}{\mu_\beta(\square) c_\beta(\square, \mathcal{X} \setminus \square)} \mathbb{P}_\eta(\tau_\square < \tau_\eta) \\ &\leq \frac{1}{c_\beta(\square, \mathcal{X} \setminus \square)} \sum_{\eta \in \mathcal{S}} |\{\eta' \in \mathcal{X} \setminus \eta: \eta \leftrightarrow \eta'\}| e^{-\beta H(\eta)} \\ &\leq \frac{1}{c_\beta(\square, \mathcal{X} \setminus \square)} C_2 e^{-(\Gamma^* + \delta)\beta} \end{aligned} \quad (3.5.10)$$

with $C_2 = |\{(\eta, \eta') \in \mathcal{S} \times \mathcal{X} \setminus \eta: \eta \leftrightarrow \eta'\}|$, where we use that $c_\beta(\eta, \eta') \leq 1$. Combine (3.5.9–3.5.10) to get the claim in (3.5.8).

(iii) Let $\partial^- \mathcal{C}^*$ be those configurations in \mathcal{C}^* where the free particle is in $\partial^- \Lambda$. Write

$$\mathbb{P}_\square(\eta_{\tau_{\partial^- \mathcal{C}^*}} = \eta | \tau_{\partial^- \mathcal{C}^*} < \tau_\square) = \frac{\mathbb{P}_\square(\eta_{\tau_{\partial^- \mathcal{C}^*}} = \eta, \tau_{\partial^- \mathcal{C}^*} < \tau_\square)}{\mathbb{P}_\square(\tau_{\partial^- \mathcal{C}^*} < \tau_\square)}, \quad \eta \in \partial^- \mathcal{C}^*. \quad (3.5.11)$$

Again by reversibility (use (1.2.7)),

$$\begin{aligned} \mathbb{P}_\square(\eta_{\tau_{\partial^- \mathcal{C}^*}} = \eta, \tau_{\partial^- \mathcal{C}^*} < \tau_\square) &= \frac{\mu_\beta(\eta) c_\beta(\eta, \mathcal{X} \setminus \eta)}{\mu_\beta(\square) c_\beta(\square, \mathcal{X} \setminus \square)} \mathbb{P}_\eta(\tau_\square < \tau_{\partial^- \mathcal{C}^*}) \\ &= e^{-\Gamma^* \beta} \frac{c_\beta(\eta, \mathcal{X} \setminus \eta)}{c_\beta(\square, \mathcal{X} \setminus \square)} \mathbb{P}_\eta(\tau_\square < \tau_{\partial^- \mathcal{C}^*}), \quad \eta \in \partial^- \mathcal{C}^*. \end{aligned} \quad (3.5.12)$$

Moreover,

$$\mathbb{P}_\eta(\tau_\square < \tau_{\partial^- \mathcal{C}^*}) = \sum_{\substack{\eta' \in \mathcal{X} \setminus \partial^- \mathcal{C}^* \\ \eta \leftrightarrow \eta'}} \frac{c_\beta(\eta, \eta')}{c_\beta(\eta, \mathcal{X} \setminus \eta)} h_{\square, \partial^- \mathcal{C}^*}^*(\eta'), \quad \eta \in \partial^- \mathcal{C}^*, \quad (3.5.13)$$

where

$$h_{\square, \partial^- \mathcal{C}^*}^*(\eta') = \begin{cases} 0 & \text{if } \eta' \in \partial^- \mathcal{C}^*, \\ \mathbb{P}_{\eta'}(\tau_\square < \tau_{\partial^- \mathcal{C}^*}) & \text{otherwise.} \end{cases} \quad (3.5.14)$$

Because $\mathcal{D} \subseteq \mathcal{X}_\square$ by Theorem 2.3.10(ii), it follows from Lemma 3.3.1 and $\mathcal{C}^* \subseteq \mathcal{S}(\square, \blacksquare)$ that

$$\min_{\eta' \in \mathcal{D}} h_{\square, \partial^- \mathcal{C}^*}^*(\eta') \geq 1 - Ce^{-\delta\beta}, \quad (3.5.15)$$

Moreover, letting $\partial^{--}\mathcal{C}^*$ be the set of configurations obtained from $\partial^- \mathcal{C}^*$ by moving the free particle from $\partial^- \Lambda$ to $\partial^{--} \Lambda = \partial^-(\Lambda^-)$, we have

$$\max_{\eta' \in \partial^{--}\mathcal{C}^*} h_{\square, \partial^- \mathcal{C}^*}^*(\eta') \leq Ce^{-\delta\beta} \quad (3.5.16)$$

This is because moving a particle from the protocritical droplet costs at least U , while attaching the free particle to the protocritical droplet leads to a configuration that, by Theorem 2.3.10(ii), is either in $\mathcal{X}_{\blacksquare}$ or in $\mathcal{X}^{**} \setminus (\mathcal{X}_\square \cup \mathcal{X}_{\blacksquare})$. By restricting the sum in (3.5.13) to $\eta' \in \mathcal{D}$ and inserting (3.5.15), we get

$$\mathbb{P}_\eta(\tau_\square < \tau_{\partial^- \mathcal{C}^*}) \geq (1 - Ce^{-\delta\beta}) \frac{c_\beta(\eta, \mathcal{D} \setminus \eta)}{c_\beta(\eta, \mathcal{X} \setminus \eta)}, \quad \eta \in \partial^- \mathcal{C}^*. \quad (3.5.17)$$

On the other hand, by inserting (3.5.16), we get

$$\mathbb{P}_\eta(\tau_\square < \tau_{\partial^- \mathcal{C}^*}) \leq \frac{c_\beta(\eta, \mathcal{D} \setminus \eta)}{c_\beta(\eta, \mathcal{X} \setminus \eta)} + Ce^{-\delta\beta} |\partial^{--}\mathcal{C}^*|, \quad \eta \in \partial^- \mathcal{C}^*. \quad (3.5.18)$$

Next, we note that for all $\eta \in \partial^- \mathcal{C}^*$,

$$\frac{c_\beta(\eta, \mathcal{D} \setminus \eta)}{c_\beta(\eta, \mathcal{X} \setminus \eta)} = O(e^{-U\beta}) + \begin{cases} \frac{1}{2} & \text{if the free particle in } \partial^- \Lambda \text{ sits in a corner,} \\ \frac{1}{4} & \text{if the free particle in } \partial^- \Lambda \text{ sits not in a corner,} \end{cases} \quad (3.5.19)$$

because moves of the free particle from $\partial^- \Lambda$ do not raise the energy (whether it stays in Λ or exits Λ), while all other moves raise the energy by at least U . Combining (3.5.18–3.5.19), we obtain

$$\mathbb{P}_\eta(\tau_\square < \tau_{\partial^- \mathcal{C}^*}) \leq (1 + Ce^{-\delta\beta}) \frac{c_\beta(\eta, \mathcal{D} \setminus \eta)}{c_\beta(\eta, \mathcal{X} \setminus \eta)}, \quad \eta \in \partial^- \mathcal{C}^*. \quad (3.5.20)$$

Inserting (3.5.17) and (3.5.20) into (3.5.12), we deduce from (3.5.11) that

$$\begin{aligned} \mathbb{P}_\square(\eta_{\tau_{\partial^- \mathcal{C}^*}} = \eta \mid \tau_{\partial^- \mathcal{C}^*} < \tau_\square) &= \frac{c_\beta(\eta, \mathcal{X} \setminus \eta) \mathbb{P}_\eta(\tau_\square < \tau_{\partial^- \mathcal{C}^*})}{\sum_{\eta' \in \partial^- \mathcal{C}^*} c_\beta(\eta', \mathcal{X} \setminus \eta') \mathbb{P}_{\eta'}(\tau_\square < \tau_{\partial^- \mathcal{C}^*})} \\ &= [1 + O(e^{-\delta\beta})] \frac{c_\beta(\eta, \mathcal{D} \setminus \eta)}{\sum_{\eta' \in \partial^- \mathcal{C}^*} c_\beta(\eta', \mathcal{D} \setminus \eta')}, \quad \eta \in \partial^- \mathcal{C}^*. \end{aligned} \quad (3.5.21)$$

Via (3.5.19) this proves the assertion in (1.4.5), because the free particle is created in $\partial^- \Lambda$ twice as fast in a corner as not in a corner.

4 Extension to three dimensions

The extension of our results to three dimensions is in principle straightforward and involves no new ideas. However, the geometry of the communication level set is more difficult and we are unable to fully identify the set \mathcal{D} . In Section 4.1 we look at the structure of $\mathcal{S}(\square, \blacksquare)$. Section 4.2 gives the proof of Theorem 1.5.2, Section 4.3 of Theorem 1.5.3.

4.1 Structure of the communication level set

We use the notation of Section 1.5.

Proposition 4.1.1 (den Hollander, Nardi, Olivieri, and Scoppola [8], Eq. (2.0.23) and Proposition 3.3.1) $\Phi(\square, \blacksquare) = \Gamma^*$ and $\mathcal{S}(\square, \blacksquare) \supseteq \mathcal{C}^*$, with Γ^* and \mathcal{C}^* given by Definition 1.5.1(c-d).

Proof. The argument is similar as for $d = 2$. A key ingredient is the following fact, shown in Alonso and Cerf [1], Theorem 3.5: the configurations consisting of a single $(m_c - 1) \times (m_c - \delta_c) \times m_c$ quasi-cube anywhere in Λ^- with, attached anywhere to one of its faces, an $(\ell_c - 1) \times \ell_c$ quasi-square are the unique (modulo translations and rotations) minimizers of H in \mathcal{V}_{n_c-1} . The energy of these configurations is $\Gamma^* - 2\Delta + 2U$, while all other configurations in \mathcal{V}_{n_c-1} have energy at least $\Gamma^* - 2\Delta + 3U > \Gamma^* - \Delta$ and therefore do not permit the creation of a particle without exceeding energy Γ^* . Thus, all optimal nucleation paths must visit this set, i.e., the analogue of Lemma 2.3.4 holds. Similarly, Lemmas 2.3.5–2.3.6 and Proposition 2.3.7 carry over. \heartsuit

Thus, the only difficult part in identifying the reduced graph \mathcal{X}^* , analogous to the one in Theorem 2.3.10, is the explicit construction of the set \mathcal{D} and the analogues of the sets \mathcal{C}^B and \mathcal{C}^G , which remains open. Nonetheless, a few facts about \mathcal{D} are easy to establish:

Proposition 4.1.2 For all $\hat{\eta} \in \mathcal{D}$,

- (i) $\text{CR}(\hat{\eta})$ is contained in a cube of side length $m_c + 1$.
- (ii) $\text{CR}(\hat{\eta})$ contains a cube of side length $m_c - \lceil \sqrt{m_c} \rceil$.

Proof. (i) In den Hollander, Nardi, Olivieri, and Scoppola [8], Proposition 5.2.1, it is shown that

$$\text{CR}(\hat{\eta}) = \text{CR}(\hat{\eta}') \quad \text{for all } \hat{\eta}, \hat{\eta}' \in \mathcal{D}. \quad (4.1.1)$$

Clearly, this is stronger than (i). For reasons of completeness we give the proof of (i).

Note that any configuration in \mathcal{D} can, on either of its faces, have a protruding rectangle with a 1-protuberance attached to it. Indeed, if we fix the number of particles sitting on top of each of the faces of $\text{CR}^-(\hat{\eta})$, then it is clear that these “two-dimensional configurations on a face” must minimize their energy. Obviously, none of them can have two 1-protuberances, since detaching one 1-protuberance (which costs $2U$) and moving it next to the other 1-protuberance (which pays $3U$) would lead to a lowering of the energy. Moreover, if any of the six clusters attached to the faces is not a rectangle, then none of the other clusters can have a 1-protuberance, since detaching this 1-protuberance (which costs $2U$) and moving it into a corner of the cluster that is not a rectangle (which pays $3U$) would lead to a lowering of the energy.

From any configuration of the above form, if we detach the 1-protuberance and place it on top of one of the rectangles in $\partial^-\text{CR}(\hat{\eta})$, then we raise the energy to $\Gamma^* - \Delta + U$. From there, moving any particle except this 1-protuberance costs energy $2U$ and leads to an energy exceeding Γ^* . Therefore all we can do is move the 1-protuberance around on top of the rectangle, until finally we have to detach it again and re-attach it to $\text{CR}^-(\hat{\eta})$.

(ii) All configurations in \mathcal{D} have volume n_c and are “minimal polyominoes”, i.e., among the configurations with volume n_c their surface is minimal. Pick $\hat{\eta} \in \mathcal{D}$. Let j_1, j_2, j_3 be the smallest integers such that $\hat{\eta}$ is contained in the $j_1 \times j_2 \times j_3$ parallelepiped. Then $\hat{\eta}$ can be obtained from this parallelepiped by removing $j_1 j_2 j_3 - n_c$ unit cubes. By (4.1.1) and Definition 1.5.1(a), we have $j_1 j_2 j_3 - n_c \leq m_c^2 - (\ell_c - 1)\ell_c - 1$ (the bound corresponding to the case where

the $(\ell_c - 1) \times \ell_c$ quasi-square is attached to an $m_c \times m_c$ face). Since $m_c \in \{2\ell_c - 1, 2\ell_c\}$, it follows that $j_1 j_2 j_3 - n_c \leq 3m_c^2/4$. Thus, no more than $3m_c^2/4$ unit cubes need to be removed from the parallelepiped to obtain $\hat{\eta}$.

Next, according to Alonso and Cerf [1], Corollary 3.26, all minimal polyominoes can be obtained from their circumscribing parallelepiped by removing a succession of bars, as many as possible, and then removing a succession of corner cubes. In our case, by (4.1.1), each bar has length either $m_c - 1$ or m_c , so no more than m_c bars and m_c corner cubes can be removed. But any such removal can only involve bars and corner cubes that lie in a layer of thickness at most $\lceil \sqrt{m_c} \rceil$ of $\text{CR}(\hat{\eta})$ (the bound corresponding to the case where the bars form a parallelepiped with an $\lceil \sqrt{m_c} \rceil \times \lceil \sqrt{m_c} \rceil$ face). \heartsuit

REMARK: Recall from the remark made below (2.2.4) that in two dimensions a U -path can shift the protocritical droplet. In contrast, (4.1.1) shows that in three dimensions a $2U$ -path cannot (see [8], Section 5).

The two global geometric facts proved in Section 2.4 continue to hold in three dimensions as well.

4.2 Average nucleation time: Proof of Theorem 1.5.2

Based on the information obtained so far, we can proceed to estimate $Z_\beta \text{CAP}_\beta(\square, \blacksquare)$ in exactly the same way as was done in Section 3.3 for two dimensions. Lemmas 3.3.1–3.3.2 and Propositions 3.3.3–3.3.4 carry over verbatim. The resulting reduction of the Dirichlet form, together with Proposition 3.2.3, proves the first half of Theorem 1.5.2. As before, the second half follows from Bovier, Eckhoff, Gaynard, and Klein [5], Theorem 1.3(iv).

4.3 Capacity asymptotics: Proof of Theorem 1.5.3

By the transience of simple random walk in three dimensions,

$$\lim_{\Lambda \rightarrow \mathbb{Z}^3} \text{CAP}^{\Lambda^+}(\partial^+ \Lambda, F) = \text{CAP}^{\mathbb{Z}^3}(F) \quad (4.3.1)$$

exists for any finite nonempty $F \subseteq \mathbb{Z}^3$. The limit, which is positive and finite, is the capacity of F . This proves Theorem 1.5.3. The bounds in (1.5.9) come from Proposition 4.1.2 in combination with Proposition 3.3.4.

If F_m is a cube of side length m , then

$$\lim_{m \rightarrow \infty} \frac{\text{CAP}^{\mathbb{Z}^3}(F_m)}{m} = \kappa \quad (4.3.2)$$

with κ the capacity of the unit cube for standard Brownian motion on \mathbb{R}^3 . This explains (1.5.11). Since $2\pi R$ is the capacity of the ball with radius R for standard Brownian motion on \mathbb{R}^3 , we have that $\kappa \in (2\pi, 2\pi\sqrt{3})$ as claimed below (3.4.1).

References

- [1] L. Alonso and R. Cerf, The three-dimensional polyominoes of minimal area, *Electron. J. Combin.* 3 (1996) Research Paper 27.
- [2] M. van den Berg, Exit and return of a simple random walk, to appear in *Potential Analysis*.

- [3] A. Bovier, Metastability and ageing in stochastic dynamics, in: *Dynamics and Randomness II, Santiago de Chile, 2002* (eds. A. Maass, S. Martinez and J. San Martin), *Nonlinear Phenomena and Complex Systems*, Vol. 10, Kluwer Ac. Publ., Dordrecht, 2004, pp. 17–81.
- [4] G. Ben Arous and R. Cerf, Metastability of the three-dimensional Ising model on a torus at very low temperature, *Electron. J. Probab.* 1 (1996) Research Paper 10.
- [5] A. Bovier, M. Eckhoff, V. Gayrard, and M. Klein, Metastability and low lying spectra in reversible Markov chains, *Commun. Math. Phys.* 228 (2002) 219–255.
- [6] A. Bovier and F. Manzo, Metastability in Glauber dynamics in the low-temperature limit: beyond exponential asymptotics, *J. Stat. Phys.* 107 (2002) 757–779.
- [7] F. den Hollander, Metastability under stochastic dynamics, *Stoch. Proc. Appl.* 114 (2004) 1–26.
- [8] F. den Hollander, F.R. Nardi, E. Olivieri, and E. Scoppola, Droplet growth for three-dimensional Kawasaki dynamics, *Probab. Theory Relat. Fields* 125 (2003) 153–194.
- [9] F. den Hollander, E. Olivieri, and E. Scoppola, Metastability and nucleation for conservative dynamics, *J. Math. Phys.* 41 (2000) 1424–1498.
- [10] F. Manzo, F.R. Nardi, E. Olivieri, and E. Scoppola, On the essential features of metastability: tunnelling time and critical configurations, *J. Stat. Phys.* 115 (2004) 591–642.
- [11] E.J. Neves and R.H. Schonmann, Critical droplets and metastability for a Glauber dynamics at very low temperature, *Commun. Math. Phys.* 137 (1991) 209–230.
- [12] P. Révész, *Random Walk in Random and Non-Random Environments*, World Scientific, Singapore, 1990.