Fluctuations and mixing for planar random growth

Vittoria Silvestri

Queens' College

University of Cambridge

A thesis submitted for the degree of

Doctor of Philosophy

July 2016

Abstract

We study two models of random growth by aggregation on the plane, which turn out to share similar asymptotic features.

The first part of this thesis focuses on the Hastings–Levitov model HL(0), according to which clusters of particles are built on the complex plane by iterated composition of random conformal maps. Following the scaling limit result of Norris and Turner (2012), who proved that the limiting shape of large HL(0) clusters is a disc, we show that the fluctuations around this deterministic shape are described by a random holomorphic Gaussian field $\mathcal{F}$ on $\{|z| > 1\}$, of which we provide an explicit construction. We find that the boundary values of $\mathcal{F}$ perform an Ornstein–Uhlenbeck process on an infinite–dimensional Hilbert space, which can be characterised as the solution of a Stochastic Fractional Heat Equation. When the cluster is allowed to grow indefinitely, this boundary process converges to a log–correlated Gaussian Field, which coincides in law with the restriction of a Gaussian Free Field on the 2–dimensional torus to the unit circle $\{|z| = 1\}$.

The same scaling limit and boundary fluctuations are found by Jerison, Levine and Sheffield (2014) to arise in a different growth model, namely Internal Diffusion Limited Aggregation (IDLA). According to this discrete model, the aggregation process defines a Markov Chain on the infinite space of IDLA configurations, for which Jerison, Levine and Sheffield ask the following mixing question: how long does it take for IDLA dynamics to essentially forget where it started? We provide a partial answer to this question in the second part of this thesis, using coupling techniques to obtain an upper bound for this forget time. Finally, we specialise to IDLA on the cylinder graph $\mathbb{Z}_N \times \mathbb{Z}$, and show that our bound is polynomial in the size N of the base graph, as $N \to \infty$.

Alla mia famiglia

Ai miei insegnanti

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor James Norris, without whom this thesis would have never been written. Thank you for giving me the fantastic opportunity to be your student in Cambridge, and for introducing me to wonderful research projects. Thank you for always finding time to talk about my work with genuine enthusiasm and curiosity, and for guiding me through difficulties with patience. I have always enjoyed Mathematics, but from you I learnt the joy of research. Thank you for believing in me. I could not have hoped for a better guidance, and all I will accomplish in the future will reflect your example.

I am also in debt to Alessandra Faggionato, who led my first steps in the research world during my master studies, and with whom I had the pleasure to keep working during the last four years. Thank you for the repeated hospitality at "La Sapienza", and for the endless energy that you put into our work.

During the last four years I have been lucky enough to meet many wonderful people with whom I truly enjoyed talking about my research. In particular I would like to thank Lionel Levine for inviting me to visit Cornell University, and for many interesting discussions which led to our joint work presented in Chapter 4 of this thesis. I am also very thankful to Scott Sheffield for many insightful conversations, both in Cambridge and at MIT, and to his students Nina, Xin and Ewain for making me feel welcome during my time at the other Cambridge. To Amine Asselah I owe the fantastic opportunity to present my work in the form of a mini–course in Marseille, for which I am truly grateful. I am also indebted to Nathanaël Berestycki and Vincent Beffara for showing interest in my research, for many very nice discussions and for having kindly agreed to read my thesis and be my examiners. A big thank you also to Alan Sola and Amanda Turner for sharing with me their insights on Hastings–Levitov models, making me feel welcome in what was for me a new research area.

Thank you to Perla Sousi, for being both a great colleague and a wonderful friend. Thank you for always giving and never asking back, for always being there for me and for your contagious enthusiasm and love for research. I hope to be lucky enough to have you as a friend for many years to come.

Thank you to Jason Miller, Laure Dumaz, Romain Allez, Benoit Laslier, Gourab Ray, Antoine Dahlqvist and all past and current member of the Statslab for having created a nice and friendly environment, where it is a pleasure to work.

I am in debt to all the people who kindly allowed me to use their simulations in my thesis, whom I individually acknowledge throughout. I also gratefully recognise the support of the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/H023348/1.

Thank you to all the CCA students, with whom I shared joy and pain over the last four years. A special thanks to my mentor Gil Ramos and co-mentor Milana Gataric, who have looked after me all this time, while constantly feeding me with coffee. Thanks to Luca Calatroni, who is responsible for my decision to apply to Cambridge and who first made me feel welcome in what was for me a foreign city. Thanks to Lukas Vermach and Eoin Devane for the many nice chats over coffee. Thanks to Tom Holding for having read a draft of the introduction to this thesis, and for many insightful conversations: it is always a pleasure to discuss maths with you. Finally, a big thank you to all the 2012 CCA cohort for the time we spent together, it has been a great experience.

None of this would have been possible without the support of my friends outside the department. Thanks to Ece, Salema, Daniela, Freya, Emil and Veronica for always being there for me. Thanks to my current and past housemates, and to the Italian crew for the many dinners together. Thanks to Jessica and Marina, for having been my friends throughout these years despite the distance. Thanks to Giacomo, for having shared with me part of this experience, and much more.

Finally, I would like to thank my wonderful family for the support and unconditional love they have always given me. Thanks mum and dad for having accepted my decision to move abroad, and for having believed in me at all times. Thanks to my amazing sisters Aldina, Fioretta and Lorenza for having always encouraged me and for repeatedly helping me put things into perspective. Vi voglio bene ragazzi.

Vittoria Silvestri
Cambridge
18 July 2016

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as specified in the text. It is not substantially the same as any that I have submitted, or is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution. I further state that no substantial part of my dissertation has already been submitted, or is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University of similar institution.

Contents

Chapter 1

Introduction

The question of obtaining a rigorous description of spatial growth phenomena observed in nature has received much attention in the last few decades in the mathematical community. One possible approach to this problem, on which we focus here, is to model the growth as the result of subsequent aggregation of particles, which randomly move in the surrounding space. We refer to such dynamics as *random aggregation*.

1.1 Discrete aggregation models

A random aggregation model consists of an underlying infinite graph together with an aggregation rule. We focus here on deterministic, planar, undirected graphs, on whose vertex set the particles live. Say that a vertex v of the graph is *filled* if it is occupied by exactly one particle, and v is *empty* otherwise. Starting from one particle (i.e. filled vertex), we grow increasing clusters of particles by aggregating one vertex at the time, according to the prescribed aggregation rule. When a vertex is added to the cluster we declare it filled, so that the cluster coincides with the set of filled vertices.

1.1.1 The Eden model

Arguably, the simplest of such models is the so called *Eden model* [Ede61], introduced by Eden in 1961, according to which at each step one empty site adjacent to the cluster is chosen uniformly at random, and added to the cluster. More precisely, let the underlying graph be the d–dimensional lattice $\mathbb{Z}^d$, and take the initial cluster to be the singleton $E(0) = \{0\}$. Define a family of growing clusters $(E(n))_{n\geq 0}$ recursively as follows. Write $u \sim v$ if u is adjacent to v, that is $|u - v| = 1$, and let

$$\partial E(n-1) = \{v \notin E(n-1) : v \sim u \text{ for some } u \in E(n-1)\}$$

denote the outer boundary of the set $E(n-1)$. Then, if ω_n is a uniformly chosen element of $\partial E(n-1)$, set

$$E(n) = E(n-1) \cup \{\omega_n\}.$$

For such family of growing clusters it is natural to ask if there is a limiting shape as the number of particles diverges, i.e. as $n \to \infty$. This question was settled by Richardson in [Ric73], showing that there exists indeed a limiting shape, and it is compact and convex (see also [CD81, Kes93]). The problem of determining the limiting shape remains open.

1.1.2 Diffusion Limited Aggregation

A variant of the Eden model was introduced a few years later by the physicists Witten and Sander [WJS81, WJS83], seeking to explain the formation of arms, or dendrites, in metal–particles aggregation. This is the celebrated *Diffusion Limited Aggregation*, in short DLA, and it is defined similarly to the Eden model, with the crucial difference that the position of the next added particle is chosen according to the harmonic measure of the outer boundary of the cluster seen from infinity, rather than according to the uniform measure on the same set. More precisely, at each integer time $n \geq 1$ we start a simple symmetric random walk $\omega^{(n)} = (\omega^{(n)}(t))_{t\geq 0}$ on

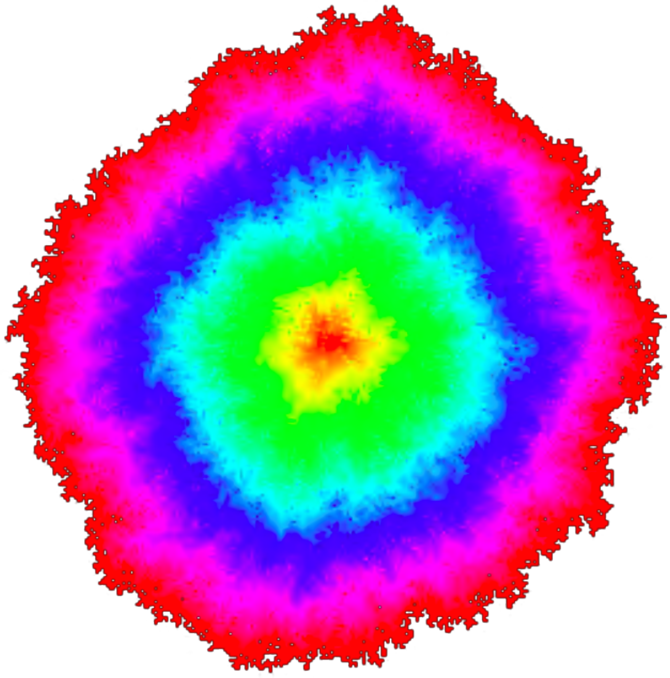

Figure 1.1: A large Eden cluster (simulation by Jason Miller).

$\mathbb{Z}^d$ from infinity, independent of everything else, and let $\tau_n := \inf\{t \geq 0 : \omega^{(n)}(t) \in \partial D(n-1)\}$ denote the first time the walk reaches the outer boundary of $D(n-1)$. Then set

$$D(n) = D(n-1) \cup \{\omega^{(n)}(\tau_n)\}.$$

Note that this aggregation rule promotes the formation of arms, or dendrites. Indeed, once an arm starts forming, it gains a large harmonic measure, and it therefore keeps on growing. This is also clear from simulations (see Figure 1.2).

The asymptotic behaviour of DLA clusters is very different from that of Eden clusters, and it is far less understood. In 1987 Kesten [Kes87, Kes90] obtained an upper bound for the length of arms in DLA, showing that in dimension d they can grow at most at rate $n^{2/(d+1)}$ as the number n of particles in the cluster diverges. After that, the only other result on DLA is due to Ebertz-Wagner [EW99], who showed that large 2–dimensional DLA clusters contain almost surely infinitely many holes. We refer the reader to [Hal00, MS13] for nice accounts on the topic, and to [SS16] for a multi–particle version of DLA growth.

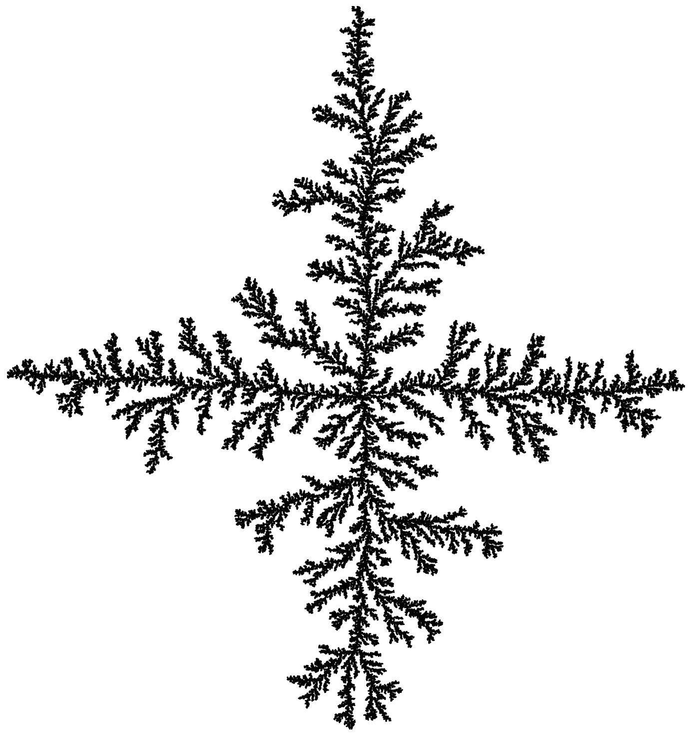

Figure 1.2: A large DLA cluster (simulation by Vincent Beffara).

1.1.3 ...and many more

Aside from DLA and the Eden model, many discrete models of random aggregation have been studied in the literature. In fact, to obtain a new model from the ones described above, one could either change the aggregation rule (i.e. change the measure with respect to which the location of the new added particle is chosen), or change the underlying graph on which the particles live, or change both. An example of aggregation rule change is given by the Internal DLA (IDLA) [DF91, LBG92, JLS14a, JLS14b, JLS13, JLS12, AG10, AG13a, AG13b, AG14], where the location of the next added particle is sampled according to the harmonic measure of the outer boundary of the cluster seen from the origin. Other important examples are the Dielectric Breakdown Models of parameter $\gamma > 0$, in short DBM$_\gamma$, according to which the location of the next particle is sampled from the γ^{th} power of the harmonic measure of the outer boundary of the cluster seen from infinity [NPW84]. Thus $\gamma = 0$ corresponds to the Eden model, $\gamma = 1$ to DLA, and $\gamma \in (0, 1)$ interpolates between the two. Finally, we would like to mention DLA on cylinder graphs [BY08] and on regular trees [BPP97], Directed DLA [Mar14], Stretched IDLA [BKP12], Diffusion Limited Deposition [ACSS16] and First Passage Percolation [HW65, Kes86, Kes03, CD81, Mar04] (see also the recent survey [AHD15] and references therein).

1.2 A continuum model: Hastings–Levitov growth

A common feature of some of the discrete aggregation models described above (e.g. Eden model and DLA) is that their asymptotic properties, such as the limiting shape and Hausdorff dimension, have been numerically found to depend on the underlying lattice structure (see for example [BBRT85, MBRS87], where simulations of such models are discussed). This led researchers to look, in the late 1980's, for off–lattice analogues [Roh11]. In 2000 Carleson and Makarov [CM01] introduced a deterministic growth model on the complex plane, defined in terms of Loewner flows driven by time-dependent measures on the unit circle. Following their ideas, Hastings and Levitov [HL98] proposed a one parameter family of random growth models, now called HL(α) for $\alpha \in [0, \infty)$. The first part of this thesis is concerned with the study of the $\alpha = 0$ case, that we now describe.

1.2.1 The HL(0) model

Let $\mathbb{D}$ denote the open unit disc in the complex plane, and let $K_0 = \overline{\mathbb{D}}$ be the initial cluster. We grow an increasing family $K_0 \subset K_1 \subset K_2 \ldots$ of compact subsets of the complex plane, that we call clusters, as follows. Fix $P \subset \mathbb{C} \setminus \mathbb{D}$ to be a (non–empty) connected compact set having 1 as a limit point, and such that the complement of $K = \overline{\mathbb{D}} \cup P$ in $\mathbb{C} \cup \{\infty\}$ is simply connected. We regard P as the basic particle. For concreteness, note that the slit $P = [1, 1+\delta]$ satisfies all the above assumptions (this is indeed the particle shape used in all HL(0) simulations included in this thesis).

At each step, a new particle P_n, consisting of a distorted copy of P, attaches to the cluster K_{n-1} according to the following growth mechanism. Let $D_0 =$

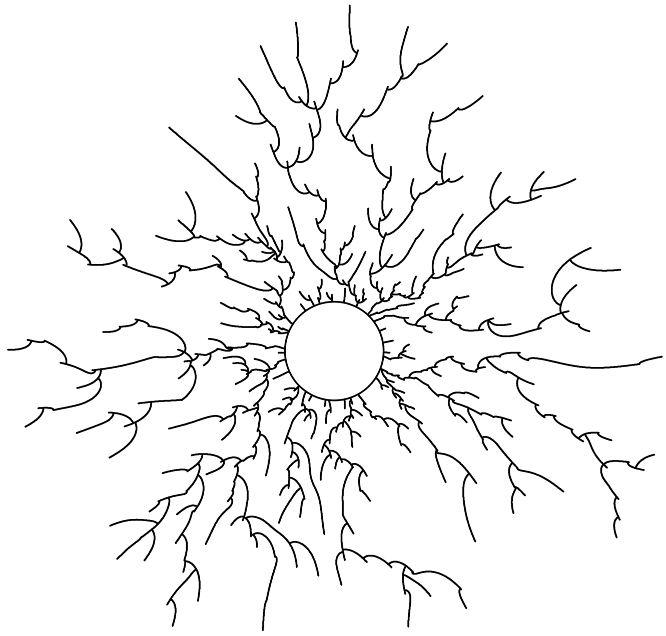

Figure 1.3: HL(0) cluster with 500 particles (simulation by Henry Jackson).

$(\mathbb{C} \cup \{\infty\}) \setminus K_0$ and $D = (\mathbb{C} \cup \{\infty\}) \setminus K$. Then there exist a unique conformal isomorphism $F : D_0 \to D$ and a unique constant $c \in \mathbb{R}_+$ such that $F(z) = e^c z + \mathcal{O}(1)$ as $|z| \to \infty$. We think of F as attaching the particle P to the closed unit disc $\overline{\mathbb{D}}$ at 1. The constant c is called *logarithmic capacity* of K, and can be interpreted as the expected value of $\log |B_T|$, for B planar Brownian Motion started at infinity, stopped at the first hitting time T of K (cf. Proposition 2.1).

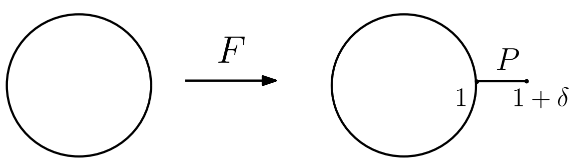

Figure 1.4: The conformal map F attaching the particle $P = [1, 1+\delta]$ at 1.

Set $G = F^{-1}$, and observe that $G(z) = e^{-c}z + \mathcal{O}(1)$ as $|z| \to \infty$. Let $(\Theta_n)_{n\geq 1}$ be a sequence of i.i.d. random variables with $\Theta_n \sim$Uniform$[-\pi, \pi)$, and set

$$F_n(z) := e^{i\Theta_n} F(e^{-i\Theta_n} z)\,, \qquad G_n(z) = F_n^{-1}(z)\,.$$

Then the map F_n attaches the particle P to the unit disc at the uniformly chosen point $e^{i\Theta_n}$. Define $\Phi_n(z) := F_1 \circ \cdots \circ F_n(z)$, $D_n = \Phi_n(D_0)$ and $K_n = (\mathbb{C} \cup \{\infty\}) \setminus D_n$. We say that the conformal map Φ_n grows an HL(0) cluster up to the n-th particle, while $\Gamma_n = \Phi_n^{-1}$ maps it out. Note that, by conformal invariance, choosing the attachment angles to be uniformly distributed corresponds to choosing the attachment point of the n-th particle according to the harmonic measure of the boundary of the cluster K_{n-1} seen from infinity. With this notation, then, one has

$$D_{n+1} = \Phi_n \circ F_{n+1} \circ \Gamma_n(D_n)\,.$$

This suggest the following interpretation for the attachment mechanism: given the cluster K_n, first map it out via Γ_n, then attach a new copy of the particle P to the unit circle at a uniformly chosen point $e^{i\Theta_n}$, and finally grow back the cluster K_n. It is then clear that, although we are attaching identical copies of the particle P at each step, the particle shape gets distorted each time by the application of the conformal map Φ_n, as shown in Figure 1.5 below.

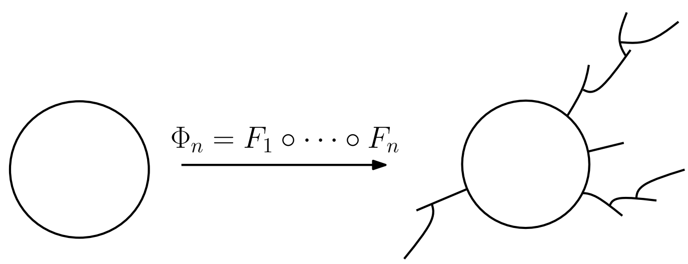

Figure 1.5: The conformal map Φ_n, which grows the cluster up to the n^{th} particle.

1.2.2 Overview of related work

As already mentioned, the HL(0) model described above belongs to a larger class of Hastings–Levitov models so called HL(α), indexed by a real parameter $\alpha \geq 0$. The case $\alpha = 0$ has already been described in detail. When $\alpha > 0$ the growth mechanism is the same, except that at each step the size of the new particle P_n is renormalised so that its logarithmic capacity is given by

$$c_n = \frac{c}{|\Phi'_{n-1}(e^{i\Theta_n})|^\alpha} \,. \tag{1.1}$$

If $\alpha = 2$ this results in growing clusters of particles roughly of the same size, and in general (1.1) has the effect of attenuating the natural distortion of the $\alpha = 0$ model (see (1.4) below). In [HL98] Hastings and Levitov argued by comparing local growth rates that the choice $\alpha = 1$ should correspond to the Eden model, $\alpha = 2$ to DLA and in general $\alpha \in (1, 2)$ to $\mathrm{DBM}_{\alpha-1}$.

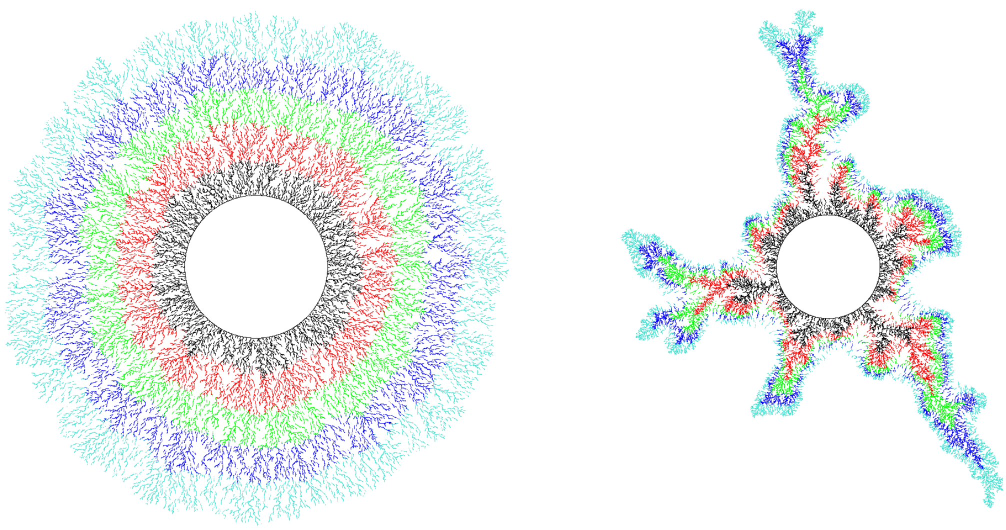

Figure 1.6: HL(1) cluster (left) and HL(2) cluster (right) with 25000 particles. Different colours correspond to different arrival times (simulations by Alan Sola).

Although $\alpha > 0$ is needed for these models to be realistic, the renormalization (1.1) creates long range dependences which make them very difficult to analyse. In fact, to the best of our knowledge there are no rigorous results on (the non–regularised version of) HL(α) models for $\alpha > 0$. A first regularised version of HL(α) appears in [RZ05], in which Rohde and Zinsmeister obtained bounds for the Hausdorff dimension of suitably regularised clusters for $\alpha \in [0, 2]$. More recently, a different type of regularization was considered by Sola, Turner and Viklund, in [JVST15], where they showed that the limiting shape of regularised clusters is given by a disc for any $\alpha > 0$, provided that the regularization is strong enough.

The case $\alpha = 0$ does not feature such long range dependences, and is much better understood. In [RZ05] Rohde and Zinsmeister obtained a scaling limit for HL(0) clusters as the particle size is kept fixed while $n \to \infty$. Moreover, they showed that the boundary of these limiting clusters is almost surely one-dimensional. More recently, in [NT15, NT12] Norris and Turner obtained a detailed description of HL(0) clusters in the small particle limit, that we now discuss.

1.2.3 Scaling limit

We are interested in studying asymptotic properties of HL(0) clusters, i.e. describing K_n as $n \to \infty$. While the number of particles diverges, to see a non-trivial shape of finite size we want to send the size of the particle to 0.

Assumption 1.1. Assume that P is such that the unique map $F : D_0 \to D$ extends continuously to $\overline{D_0} = \{z \in \mathbb{C} : |z| \geq 1\}$. Moreover, assume that there exists $\delta > 0$ such that

$$P \subseteq \{z \in \mathbb{C} : |z - 1| \leq \delta\}, \quad 1 + \delta \in P, \quad P = \{\bar{z} : z \in P\}. \tag{1.2}$$

Under the above assumption, which is in force throughout the thesis, Norris and Turner [NT12] prove the following.

Lemma 1.1 ([NT12],Corollary 4.2)**.** *Assume* (1.2) *and let c denote the logarithmic capacity of $\overline{\mathbb{D}} \cup P$. Then, as $\delta \to 0$, $\delta^2/6 \leq c \leq 3\delta^2/4$.*

Thus sending the size of P to 0 is equivalent to sending its logarithmic capacity $c \to 0$. Finally, fix $t > 0$ and assume that $nc \to t$ as $n \to \infty, c \to 0$. We call this *small particle regime.*

The existence of an asymptotic shape for HL(0) clusters in the small particle regime was proved in [NT12], where the authors show that large clusters converge to Euclidean balls. More precisely, they prove the following.

Theorem 1.1 ([NT12], Proposition 5.1)**.** *Let* $m = \lfloor \delta^{-6} \rfloor$ *and* $\varepsilon = \delta^{2/3}\log(1/\delta)$*, and let*

$$\mathcal{E}_\delta := \Big\{ \sup_{n \le m} \sup_{|z| \ge e^{5\varepsilon}} \big| \Phi_n(z) - e^{cn} z \big| < \varepsilon e^{6\varepsilon + cn} \Big\}.$$

Then, in the small particle regime, $\mathbb{P}(\mathcal{E}_\delta^c)$ *decays faster than any polynomial in* δ *as* $\delta \to 0$*. In particular, almost surely*[1] *the event* $\mathcal{E}_\delta$ *holds for* δ *small enough.*

Note that for any fixed n the function $z \mapsto \Phi_n(z) - e^{nc}z$ is holomorphic outside the unit disc $\mathbb{D}$, and it converges to 0 as $|z| \to \infty$. It therefore follows from the maximum principle that

$$\sup_{|z| \ge e^{5\varepsilon}} \big| \Phi_n(z) - e^{cn} z \big| = \sup_{|z| = e^{5\varepsilon}} \big| \Phi_n(z) - e^{cn} z \big| \to 0 \qquad \text{as } nc \to t \tag{1.3}$$

almost surely. Thus

$$\Phi_n(z) \approx e^t z \qquad \text{as} \quad nc \to t,$$

which shows that the map that grows the n–particles cluster is close, in the small particle limit, to the map that grows an Euclidean ball of radius e^t around the origin. In this sense we say that the scaling limit of HL(0) clusters is an Euclidean ball.

Remark. The convergence in (1.3) implies that $\Phi_n \to e^t z$ uniformly on all compact sets in $(\mathbb{C} \cup \{\infty\}) \setminus \overline{\mathbb{D}}$. Thus the associated sequence of random clusters K_n converges to the disc $e^t \overline{\mathbb{D}}$ in the Carathéodory sense.

Remark. Although one can let $|z| \to 1$ as $n \to \infty$ in Theorem 1.1 of Norris and Turner, the limiting shape result breaks down (and it should) if $|z| \to 1$ too fast with respect to δ. On the other hand, we will show in Chapter 3 that it does not break too much: it still holds in the sense of distributions, for a suitably small space of test functions (see Chapter 3, Section 3.3 for more details).

[1]Note that the only randomness is in the attachment angles. For each sequence $(\Theta_k)_{k \ge 1}$, and each $\delta > 0$, we build a sequence of conformal maps $(\Phi_n)_{n \ge 1}$ which grow the cluster with basic particle size δ. For different values of δ such sequences of conformal maps are coupled by using the same sequence of attachment angles, and by simply attaching a rescaled version of the basic particle P, according to the value of δ.

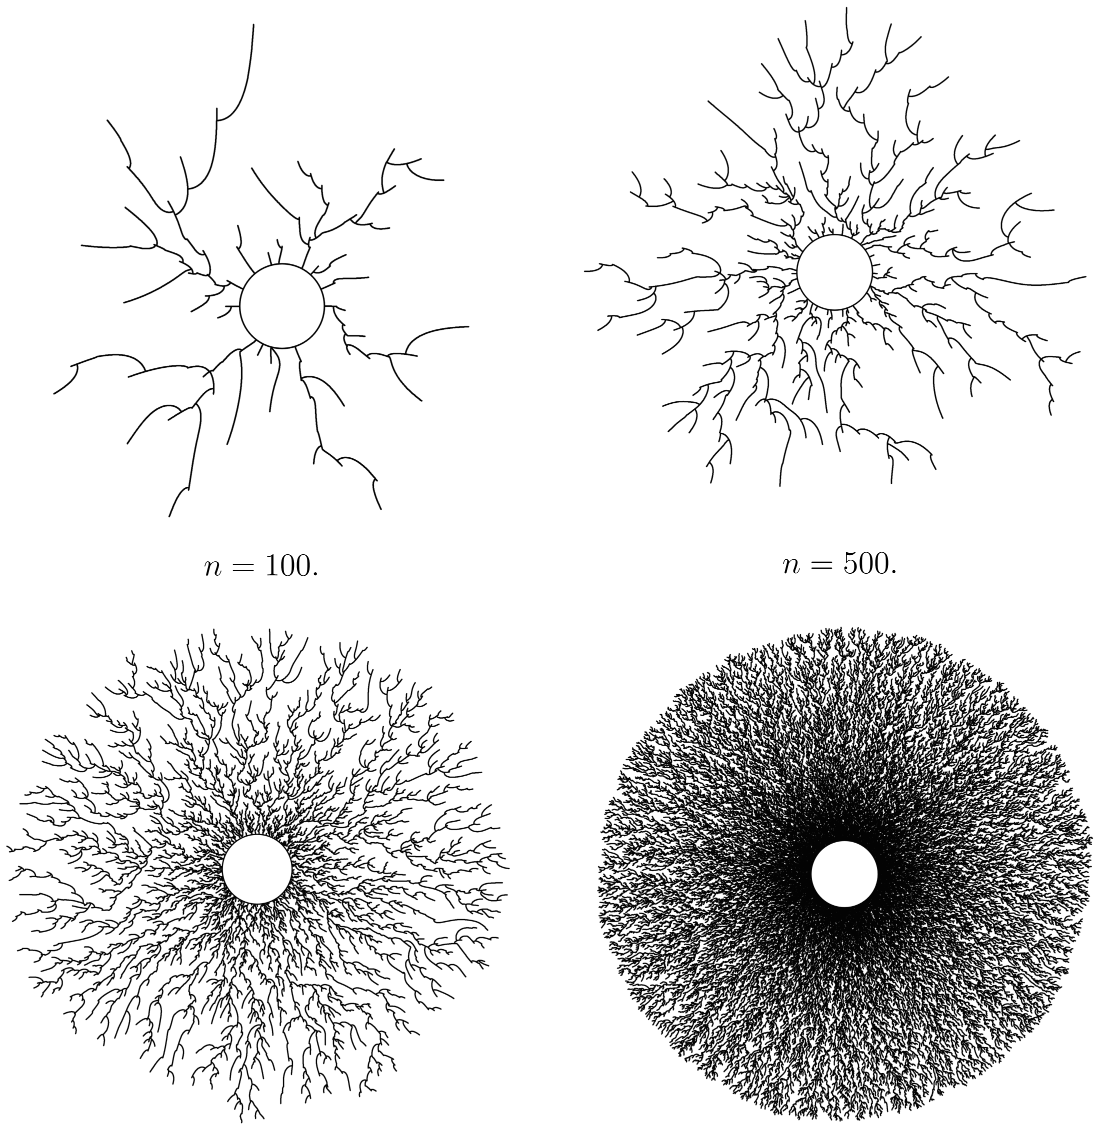

$n = 100$. $n = 500$.

$n = 5000$. $n = 100000$.

Figure 1.11: HL(0) clusters (simulations by Henry Jackson).

The analysis carried out by Norris and Turner in [NT12] is rather sophisticated, and it gives much more information than what stated in Theorem 1.1. In particular, it provides very accurate estimates on the distortion of the conformal map Φ_n, that are crucial for our analysis. On the other hand, the limiting shape of large clusters can be guessed by very simple heuristic arguments, that we now discuss.

Heuristics: a smoothing dynamics (distortion helps)

To fix ideas, take the particle P to be the slit $[1, 1+\delta]$ for some small $\delta > 0$, and let P_{n+1} denote the $(n+1)^{th}$ copy of P in the cluster, that is

$$P_{n+1} = \Phi_n\big(e^{i\Theta_{n+1}}P\big) = \big\{z : z = \Phi_n(w) \text{ for some } w \text{ such that } e^{-i\Theta_{n+1}}w \in P\big\}.$$

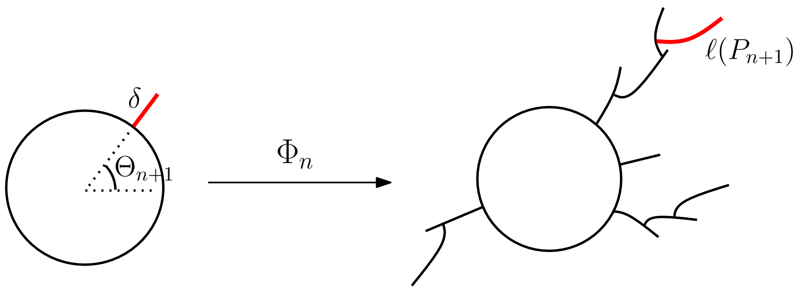

We measure the distortion of the conformal map Φ_n by comparing the size of P with the one of P_{n+1}, where by size of the particles we mean their arc length, denoted by $\ell(\cdot)$. By the mean value theorem, we have

$$\ell\big(P_{n+1}\big) = \int_1^{1+\delta} |\Phi_n'(re^{i\Theta_{n+1}})|\mathrm{d}r = \delta\,|\Phi_n'(r_0 e^{i\Theta_{n+1}})|, \tag{1.4}$$

for some $r_0 \in [1, 1+\delta]$. For small δ, $r_0 \approx 1$, from which we deduce that

$$\ell\big(P_{n+1}\big) \approx \delta\,|\Phi_n'(e^{i\Theta_{n+1}})|, \qquad \delta \ll 1.$$

On the other hand, it is easy to check that the density of harmonic measure at the attachment point $\Phi_n(e^{i\Theta_{n+1}})$, that is the probability density that the $(n+1)^{th}$ particle attaches to the cluster K_n at that point, is proportional to $|\Phi_n'(e^{i\Theta_{n+1}})|^{-1}$. Putting this together with (1.4), we conclude that:

- particles attaching to the cluster at unlikely locations are large, and

- particles attaching to the cluster at likely locations are small.

This shows that the distortion intrinsic in the model helps preventing the formation of arms, which, together with radial symmetry, suggests that the limiting shape could be an Euclidean ball.

Heuristics: a Loewner flow argument

This second heuristic argument is inspired by the discussion of anisotropic Hastings–Levitov models by Sola, Turner and Viklund in [JVST15], and it provides an exact guess for the limiting shape.

The main observation is that Hastings–Levitov clusters are Loewner hulls associated to the Loewner flow driven by a specific random measure, consisting of Dirac masses at the attachment points. More precisely, recall that $(\Theta_k)_{k\geq 1}$ denotes the sequence of attachment angles, and define

$$\xi^n(t) := \sum_{k=1}^{n} e^{i\Theta_k} 1_{[c(k-1),ck)}(t).$$

Let $(f_t)_{t\geq 0}$ be the Loewner flow driven by the random measure $\delta_{\xi^n(t)}$, where δ_x denotes the Dirac delta measure centred at x, i.e.

$$\partial_t f_t(z) = z f_t'(z) \int_{|z|=1} \frac{z+\zeta}{z-\zeta} \mathrm{d}\delta_{\xi^n(t)}(\zeta), \qquad f_0(z) = z. \tag{1.5}$$

Then $\Phi_n(z) = f_{nc}(z)$ for the slit particle $P = [1, 1+\delta]$. Using the Loewner equation (1.5) for f_t, we find

$$\begin{aligned}
\partial_t f_t(z) &= z f_t'(z) \int_{|z|=1} \frac{z+\zeta}{z-\zeta} \left(\sum_{k=1}^{n} \mathrm{d}\delta_{e^{i\Theta_k}}(\zeta) 1_{[c(k-1),ck)}(t) \right) \\
&= z f_t'(z) \sum_{k=1}^{n} \left(\int_{|z|=1} \frac{z+\zeta}{z-\zeta} \mathrm{d}\delta_{e^{i\Theta_k}}(\zeta) \right) 1_{[c(k-1),ck)}(t) \\
&= z f_t'(z) \sum_{k=1}^{n} \frac{z+e^{i\Theta_k}}{z-e^{i\Theta_k}} 1_{[c(k-1),ck)}(t).
\end{aligned}$$

Now, if we assume that there exists a deterministic scaling limit, then the random sum appearing in the last term above should converge to a deterministic limit itself.

To get an idea of how this limit should look like, we compute its expectation:

$$\mathbb{E}\left(\sum_{k=1}^{n}\frac{z+e^{i\Theta_k}}{z-e^{i\Theta_k}}1_{[c(k-1),ck)}(t)\right)=\sum_{k=1}^{n}\mathbb{E}\left(\frac{z+e^{i\Theta_k}}{z-e^{i\Theta_k}}\right)1_{[c(k-1),ck)}(t)$$
$$=\mathbb{E}\left(\frac{z+e^{i\Theta_1}}{z-e^{i\Theta_1}}\right)=\int_{|z|=1}\frac{z+\zeta}{z-\zeta}\frac{\mathrm{d}\zeta}{2\pi}=1.$$

Thus if there is a deterministic limit for the random map $\Phi_n = f_{nc}$, then it must coincide with the solution f_t to

$$\begin{cases}\partial f_t(z) = z f_t'(z)\\ f_0(z) = z\end{cases}$$

evaluated at time $t = nc$. The above is easy to solve, to get $f_t(z) = e^t z$. Thus we conclude that if there exists a random conformal map f_t such that $\Phi_n \to f_t$ as $n \to \infty, nc \to t$, then it must be $f_t(z) = e^t z$.

Remark. While the scaling limit result of Theorem 1.1 is concerned with the image, under the cluster growth, of points *outside the cluster*, Norris and Turner also look at the *internal structure* of large clusters. Under additional assumptions on the particle P, they identify a genealogical structure in the cluster, and prove in [NT15, NT12] that (in the small particle limit, and in logarithmic coordinates) this converges to the so called *Brownian Web*, which can be informally described as an uncountable family of coalescing 1–dimensional Brownian Motions starting from any point in time and space [Arr79, Arr81, TW98, FINR04]. This provides an interesting connection between these two seemingly unrelated models. We refer the reader to Section 3.4 of Chapter 3 for more details, and for some open questions in this direction.

1.2.4 Pointwise fluctuations

Once convergence to a deterministic limit for the sequence of conformal maps $(\Phi_n)_{n\geq 1}$ has been established, it is natural to ask what are the fluctuations around this limit. This is the question we answer in Part I of the thesis.

As it turns out, it is convenient to formulate fluctuation results in logarithmic coordinates, as this makes the underlying multiplicative structure of the model into an additive one. We thus look at the holomorphic function $\log \frac{\Phi_n(z)}{z}$ on $\{|z| > 1\}$, where the branch of the logarithm has been fixed so that $\log \frac{\Phi_n(z)}{z} \to nc$ as $|z| \to \infty$. Then the law of large numbers of Theorem 1.1 reads

$$\log \frac{\Phi_n(z)}{z} \approx nc \quad \text{as} \quad nc \to t. \tag{1.6}$$

Our results can be divided into local and global fluctuations.

Local fluctuations

Suppose we fix $z \in \mathbb{C} \setminus \overline{\mathbb{D}}$, say $z = e^{\sigma + ia}$ for some $\sigma > 0$, $a \in [-\pi, \pi)$, and look at the limiting fluctuations of $\log \frac{\Phi_n(z)}{z}$ around its mean as the HL cluster grows ($n \to \infty$) and the point z approaches the unit disc radially ($\sigma \to 0$). We prove that, provided $\sigma \to 0$ slowly enough, these limiting fluctuations are Gaussian. Moreover, approaching the unit disc radially from different angles results in asymptotically independent fluctuations. Our main result for this regime is the following.

Theorem (Local fluctuations)**.** *Pick any $t > 0$, and let $z = e^{ia+\sigma}$ for some $a \in [-\pi, \pi)$, $\sigma > 0$. Then as $n \to \infty$, $nc \to t$ and $\sigma \to 0$ slowly enough,*

$$\frac{\log \frac{\Phi_n(z)}{z} - nc}{\sqrt{c \log(\frac{1}{2\sigma})}} \longrightarrow \mathcal{N}(0,1)$$

in distribution, where $\mathcal{N}(0,1)$ denotes the law of a standard complex Gaussian random variable. Moreover, the correlation between fluctuations at two different points, say $z = e^{ia+\sigma}$ and $w = e^{ib+\sigma}$, vanishes in the limit, unless the angle $a - b$ converges to 0 fast enough with σ.

A precise statement of the above result is provided in Theorem A.1 of Appendix A.1. To prove it, we identify a sequence of (complex) backwards martingale difference arrays, defined in (2.10). We then invoke a Central Limit Theorem of McLeish [McL74], that we adapt to our setting in Appendix B, to show that such sequence has a Gaussian limit. The assumptions of such CLT are proved to hold by combining distortion estimates for the conformal maps F, G with a crucial good event of Norris and Turner [NT12]) (see also Theorem 2.3 below).

Note that the only Gaussian random field on $\mathbb{T} = \partial\mathbb{D}$ having the correlation structure defined above is the one given by an uncountable collection of i.i.d. $\mathcal{N}(0,1)$ random variables indexed by points on the unit circle. Although almost surely finite at every point, this random field is very wild, in the sense that, apart from not being continuous a.s., it is not even separable (i.e. it cannot be recovered by only looking at a countable collection of points).

Global fluctuations

At the price of keeping z away from the unit disc while the cluster grows, we see that the fluctuations of $\log \frac{\Phi_n(z)}{z}$ become rather well behaved. In order to emphasize the dependence on t in this regime, and to ultimately view the limit as a stochastic process, we now assume $nc \to 1$ as $n \to \infty$, and study the asymptotic behaviour of the map $\Phi_{\lfloor nt \rfloor}$ for $t \in [0, \infty)$.

Then, by the same techniques that allow us to prove the local fluctuations result, we can show (cf. Theorem 2.1) that, for any *fixed* $z \in \mathbb{C} \setminus \overline{\mathbb{D}}$, these fluctuations are again centred Gaussian, with variance now depending on $|z|$ and t. Moreover, the correlation structure is sufficiently well behaved to enable us to prove a functional central limit theorem for $\log \Phi_{\lfloor nt \rfloor}$ when restricted to any circle of the form $r\mathbb{T} := \{z : |z| = r\}$, $r > 1$ (cf. Theorem 3.1). Finally, we push our analysis forward to obtain limiting fluctuations of $\log \Phi_{\lfloor nt \rfloor}$ viewed as a càdlàg stochastic process taking values in the space of holomorphic functions on $\mathbb{C} \setminus \overline{\mathbb{D}}$. Our main result is the following.

Theorem (Global fluctuations)**.** *Let $\mathscr{H}$ denote the space of holomorphic functions on $\{|z| > 1\}$, and for $n \geq 1$ set $\mathcal{F}_n(t, z) = \frac{1}{\sqrt{c}}\big(\log \frac{\Phi_{\lfloor nt \rfloor}(z)}{z} - \lfloor nt \rfloor c\big)$. Then there exists a continuous stochastic process $\mathcal{F} = (\mathcal{F}(t, \cdot))_{t \geq 0}$ taking values in $\mathscr{H}$ such* *that $\mathcal{F}_n \to \mathcal{F}$ in distribution as $n \to \infty$, with respect to the Skorokhod topology on the space of càdlàg functions from $[0,\infty)$ to $\mathscr{H}$. Moreover, $\mathcal{F}$ can be obtained as the holomorphic extension of its boundary values on $\{|z| = 1\}$ to the outer unit disc $\{|z| > 1\}$. These boundary values are given by a distribution-valued stochastic process $\mathcal{W} = (\mathcal{W}(t,\cdot))_{t\geq 0}$, formally defined in Fourier space by*

$$\mathcal{W}(t,\vartheta) = \sum_{k\in\mathbb{Z}\setminus\{0\}} \Big(\frac{A_k(t) + iB_k(t)}{\sqrt{2}}\Big)\frac{e^{ik\vartheta}}{\sqrt{2\pi}}\,, \tag{1.7}$$

for $(A_k)_k$, $(B_k)_k$ independent collections of Ornstein-Uhlenbeck processes on $\mathbb{R}$, solutions to

$$\begin{cases} \mathrm{d}A_k(t) = -|k|A_k(t)\mathrm{d}t + \sqrt{2}\,\mathrm{d}\beta_k(t) \\ A_k(0) = 0, \end{cases} \qquad \begin{cases} \mathrm{d}B_k(t) = -|k|B_k(t)\mathrm{d}t + \sqrt{2}\,\mathrm{d}\beta'_k(t) \\ B_k(0) = 0\,, \end{cases} \tag{1.8}$$

where $(\beta_k)_k$, $(\beta'_k)_k$ are independent collections of i.i.d. Brownian Motions on $\mathbb{R}$.

A precise statement of this result is given in Theorem 3.3. This provides an explicit construction of the limiting Gaussian holomorphic field $\mathcal{F}$ and, perhaps more interestingly, of its boundary values $\mathcal{W}(t,\vartheta) = \mathcal{F}(t, e^{i\vartheta})$.

Note that, at least formally, equations (1.7)-(1.8) can be summarised by saying that $\mathcal{W}$ solves[2] the following Stochastic Fractional Heat Equation:

$$\begin{cases} \mathrm{d}\mathcal{W}(t,\cdot) = -\sqrt{-\Delta}\mathcal{W}(t,\cdot)\mathrm{d}t + \sqrt{2}\,\mathrm{d}\xi(t,\cdot), \\ \mathcal{W}(0,\cdot) = 0, \end{cases} \tag{1.9}$$

where Δ is the Laplace operator in the spatial variable, and ξ denotes complex space–time white noise on the unit circle $\mathbb{T} = \{|z| = 1\}$. We show in Chapter 3 that such $\mathcal{W}$ can be made sense of as an Ornstein–Uhlenbeck process in a suitable infinite–dimensional Hilbert space. This is a Gaussian Markov process, which converges to its stationary distribution as $t \to \infty$. Note that, since

[2] By this we mean that translating the above SPDE in Fourier space one recovers (1.7)-(1.8) for the Fourier coefficients of $\mathcal{W}$.

$A_k(t), B_k(t) \to \mathcal{N}\big(0, \frac{1}{|k|}\big)$ in law as $t \to \infty$, we have[3]

$$\mathcal{W}(t,\vartheta) \overset{t\to\infty}{\longrightarrow} \mathcal{W}_\infty(\vartheta) \overset{(d)}{=} \sum_{k\neq 0} \frac{1}{\sqrt{|k|}} \Big(\frac{A_k(\infty) + iB_k(\infty)}{\sqrt{2}} \Big) \frac{e^{ik\vartheta}}{\sqrt{2\pi}} \,, \tag{1.10}$$

for $(A_k(\infty))_k$, $(B_k(\infty))_k$ independent collections of i.i.d. standard Gaussian random variables on $\mathbb{R}$. This is (the complex version of) a well-known Fractional Gaussian Field on the unit circle $\mathbb{T}$, which be realised as

$$\mathcal{W}_\infty = (-\Delta)^{-1/4} W,$$

for W white noise on $\mathbb{T}$. Fractional Gaussian Fields can be studied in greater generality: we refer the interested reader to [DRSV14, LSSW14] and references therein. For $d \geq 1$, let Δ_d and W_d denote the Laplace operator and white noise on the d–dimensional torus $\mathbb{T}^d$ respectively, and let $s > 0$ be a real number. Then $(-\Delta_d)^{-s/2} W_d$ defines a Fractional Gaussian Field. This is a (random) continuous function if $s > d/2$, and otherwise it takes values in the space of distributions. Moreover, at the critical parameter $s = d/2$ it defines a log–correlated Gaussian Field (LGF), which coincides with the Gaussian Free Field (GFF) on $\mathbb{T}^d$ when $d = 2$. Restricting the GFF to $\mathbb{T}_1$ one obtains the LGF $(-\Delta_1)^{-1/4} W_1$, which by (1.10) has the same law as $\mathcal{W}_\infty$. In this sense we say that, as $t \to \infty$, the boundary fluctuations of HL(0) clusters are given by the restriction of a (complex) GFF on the 2–dimensional torus to the unit circle.

1.2.5 Heuristics: a universal fluctuations process

In a way, $\mathcal{W}$ defined in (1.9) is a very natural fluctuation process. Let us give some heuristics in support of this statement[4]. Write Φ_t in place of $\Phi_{\lfloor nt \rfloor}$ for simplicity, and for $s \leq t$ let $\Phi_{s,t} := \Phi_s^{-1} \circ \Phi_t$, i.e. $\Phi_{s,t} = F_{\lfloor ns \rfloor + 1} \circ \cdots \circ F_{\lfloor nt \rfloor}$. Then $\Phi_{0,t} = \Phi_t$,

[3]It is easy to see that the convergence in (1.10) holds in the sense of Finite Dimensional Distributions, i.e. for any finite collection of sufficiently smooth test functions $\phi_1, \ldots, \phi_k$ we have that the joint law of $(\mathcal{W}(t,\cdot), \phi_1), \ldots, (\mathcal{W}(t,\cdot), \phi_k)$ converges, as $t \to \infty$, to the joint law of $(\mathcal{W}_\infty, \phi_1), \ldots, (\mathcal{W}_\infty, \phi_k)$.

[4]We thank Vincent Beffara and James Norris for suggesting this heuristic argument.

and

$$\Phi_{0,t} = \Phi_{0,s} \circ \Phi_{s,t} \tag{1.11}$$

for all $s \leq t$. Let, informally, $\phi_{s,t}(z)$ denote the fluctuations of $\Phi_{s,t}(z)$ around its deterministic limit $e^{t-s}z$ in logarithmic coordinates, so that

$$\log \Phi_{s,t}(z) = \log z + (t-s) + \phi_{s,t}(z) \qquad \text{for } |z| > 1. \tag{1.12}$$

Thus

$$\begin{aligned} \log \Phi_{0,t}(z) &= \log z + t + \phi_{0,t}(z), \\ \log \big(\Phi_{0,s} \circ \Phi_{s,t}\big)(z) &= \log \Phi_{s,t}(z) + s + \phi_{0,s}\big(\Phi_{s,t}(z)\big) \\ &= \log z + (t-s) + \phi_{s,t}(z) + s + \phi_{0,s}\big(e^{t-s}z\big), \end{aligned}$$

where in the last equality we have used the first order approximation in the argument of the fluctuation term $\phi_{0,s}$. This together with (1.11) implies that

$$\phi_{0,t}(z) = \phi_{0,s}(e^{t-s}z) + \phi_{s,t}(z). \tag{1.13}$$

Note that $\phi_{s,t}$ is independent of $\phi_{0,s}$ for all $s \leq t$, since $\Phi_{s,t}$ is independent of $\Phi_{0,s}$. Moreover, all the functions appearing in (1.13) are holomorphic on $\{|z| > 1\}$ by (1.12). It follows that, assuming their boundary values are nice enough, they can be recovered as the holomorphic extension to $\{|z| > 1\}$ of their values on $\{|z| = 1\}$. Thus (1.13) also holds formally on the unit circle. For any $\vartheta \in [0, 2\pi)$, we take $z = e^{i\vartheta}$ in (1.13) and expand both sides in Fourier series, to find

$$\sum_k A_{0,t}(k) e^{ik\vartheta} = \sum_k A_{0,s}(k) e^{-(t-s)|k|} e^{ik\vartheta} + \sum_k A_{s,t}(k) e^{ik\vartheta},$$

where the exponential term in the second sum above comes from the holomorphic extension procedure. We therefore conclude that

$$A_{0,t}(k) = A_{0,s}(k) e^{-(t-s)|k|} + A_{s,t}(k)$$

for all k and all $s \leq t$. As s approaches t, it is natural to assume that $A_{s,t}(k)$ behaves like Gaussian noise, so that $A_{s,t}(k) \approx B_k(t)$, for B_k complex Gaussian process, with real and imaginary parts given by (possibly correlated) multiples of real Brownian Motions. Differentiating, still formally, the above equality at t then gives

$$\mathrm{d}A_{0,t}(k) = -|k|A_{0,t}(k) + \mathrm{d}B_k(t). \tag{1.14}$$

To conclude, we argue that the Brownian noises corresponding to different Fourier frequencies should be independent standard complex Brownian Motions. This is a simple consequence of rotation invariance. Indeed, writing ϕ_t in place of $\phi_{0,t}$ for brevity, for any $\vartheta, \varphi \in [0, 2\pi)$ we must have that $\phi_t(e^{i\vartheta})$ and $\phi_t(e^{i\varphi})$ have the same law, by rotation invariance of the model. Expand both terms in Fourier series, to get

$$\phi_t(e^{i\vartheta}) = \sum_k A_k(t)e^{ik\vartheta}, \qquad \phi_t(e^{i\vartheta}) = \sum_k A_k(t)e^{ik\varphi} = \sum_k \left[A_k(t)e^{ik(\varphi-\vartheta)}\right]e^{ik\vartheta}.$$

Assuming that the Fourier coefficients are jointly Gaussian, then, we deduce that the Gaussian vectors $\left(A_k(t)\right)_k$ and $\left(A_k(t)e^{i(\varphi-\vartheta)k}\right)_k$ must have the same law. This shows that the complex Fourier coefficients must be rotationally invariant in law, which forces the noises B_k to be (multiples of) standard complex Brownian Motions, since their covariance matrix must be a multiple of the identity matrix by rotation invariance. Moreover, all pairwise correlations vanish, since

$$\mathbb{E}[A_k(t)\overline{A_j(t)}] = \mathbb{E}[A_k(t)\overline{A_j(t)}]e^{i(k-j)(\varphi-\vartheta)}$$

for all φ, ϑ, which implies $\mathbb{E}[A_k(t)\overline{A_j(t)}] = 0$ for all $k \neq j$. Thus the Fourier coefficients $(A_k(t))_k$ must be pairwise independent. Since pairwise independence implies independence for jointly Gaussian random variables, this shows that the driving Brownian Motions $(B_k(t))_k$ appearing in the SDEs (1.14) should be independent.

The above argument explains, still at a heuristic level, the form of the SDEs (1.8) for the Fourier coefficients of the boundary process $\mathcal{W}$. Moreover, it suggests that the fluctuation process $\mathcal{W}$ should be in some sense universal, as we have made no use of the specific definition of the HL(0) model in the above derivation. Indeed, the same process is known to describe the fluctuations of another growth model, this time discrete, on which we focus in the second part of this thesis.

1.3 Back to discrete models: Internal DLA

Let us define another discrete aggregation model, so called Internal DLA (or IDLA), on the d–dimensional grid $\mathbb{Z}^d$. This model was introduced by Meakin and Deutch [MD86] in 1986 and, as the name suggests, is defined along the lines of Diffusion Limited Aggregation, with the crucial difference that the cluster grows from the inside rather than the outside. More precisely, assume that the initial cluster consists of the singleton $A(0) = \{0\}$. At each integer time $n \geq 1$, we start a simple symmetric random walk $\omega^{(n)}$ from the origin, independent of everything else. Let $\tau_n := \inf\{t \geq 0 : \omega^{(n)}(t) \notin A(n-1)\}$ denote the first time the walk exits the current cluster. Then we set

$$A(n) = A(n-1) \cup \{\omega^{(n)}(\tau_n)\}.$$

Although the IDLA dynamics looks similar to the one of external DLA, these two growth processes are very different in nature. Indeed, while external DLA promotes the formation of arms, IDLA tends to smooth them out, thus leading to much more regular clusters[5].

Contrary to external DLA, which has proven to be very difficult to analyse, much is known about asymptotic properties of IDLA clusters. In 1992 Lawler, Bramson and Griffeath [LBG92] show that large IDLA clusters in $\mathbb{Z}^d$ are Euclidean balls, apart from an error $o(\sqrt{r})$, where r is the radius of the ball, in all dimensions $d \geq 2$. Few years later Lawler [Law95] improves the error to $\mathcal{O}(r^{1/3})$, and asks

[5] Such smoothing effect is also present in HL(0) growth, as pointed out in Section 1.2.3, although of different nature. Indeed, while in IDLA we attach particles of the same size, and the smoothing comes from the attachment location, in HL(0) the attachment point is chosen exactly as in external DLA, and the smoothing is due to the fact that the size of the attached particle is random, depending on the distortion properties of the conformal map Φ_n.

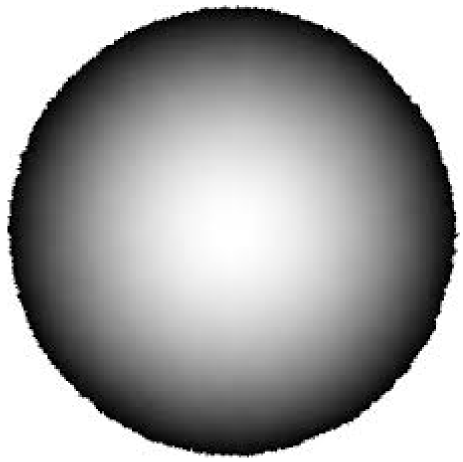

Figure 1.12: IDLA cluster with $n = 100000$ particles on $\mathbb{Z}^2$ (simulation by Wilfried Huss).

whether the same can be proved with a smaller power of the radius. The question remains open until 2010, when Asselah and Gaudillière prove, in a series of papers [AG10, AG13a, AG13b], that fluctuations from circularity are at most of size $\sqrt{\log r}$ in dimension $d \geq 3$, and at most $\log^2 r$ for $d = 2$. They also show in [AG14] that fluctuations of size $\sqrt{\log r}$ do occur all dimensions (matching the upper bound, and thus determining the right asymptotic order of maximal fluctuations, for $d = 3$). At the same time and independently, Jerison, Levine and Sheffield [JLS12, JLS13] prove that fluctuations from circularity are of order at most $\log r$ in dimension 2, and $\sqrt{\log r}$ in higher dimension. Their result for $d = 2$ is believed to be sharp, although this is still open. The same authors then look at average fluctuations from circular shape in [JLS14a], showing that they converge in law to the so called *augmented Gaussian Free Field* in all dimensions. Seeking to obtain exactly the Gaussian Free Field as the limit of average fluctuations, in the sequel paper [JLS14b] the same authors consider IDLA on the cylinder graph. More precisely, fix $N \geq 2$ integer, and let $\mathbb{Z}_N := \mathbb{Z}/N\mathbb{Z}$ denote the N–cycle. We consider an IDLA process starting from the half–filled configuration

$$A(0) = \{(x, y) \in \mathbb{Z}_N \times \mathbb{Z} : y \leq 0\}.$$

At each step, we choose one vertex of the form $(x, 0)$, $x \in \mathbb{Z}_N$ uniformly at random, and start a simple symmetric random walk on $\mathbb{Z}_N \times \mathbb{Z}$ from it. By recurrence of the vertical coordinate, such walk will exit the current cluster in a finite time almost surely. When this happens, we add its exit location to the cluster. Let $(A(n))_{n\geq 0}$ denote the resulting stochastic process on the space Ω of configurations,

$$\Omega = \{A \subset \mathbb{Z}_N \times \mathbb{Z},\ A \supseteq \{(x,y) : y \leq 0\}\}.$$

The first result of [JLS14b] states that the limiting shape of such clusters for large N is given by a cylinder, with fluctuations of size at most $\log N$.

Theorem 1.2 ([JLS14b], Theorem 2)**.** *For any finite $a, t > 0$, there exists a constant $C = C(a,t)$ such that, for $T = tN^2$ and N large enough,*

$$\left\{(x,y) : y \leq \frac{n}{N} - C\log N\right\} \subseteq A(n) \subseteq \left\{(x,y) : y \leq \frac{n}{N} + C\log N\right\}$$

for all $n \leq T$, with probability at least $1 - N^{-a}$.

In particular, rescaling the mesh size by $1/N$ we deduce that the limiting shape of such clusters above level 0 is the cylinder $\mathbb{T} \times [0,t]$, where $\mathbb{T}$ denotes the 1–dimensional torus.

Remark. Note that, as already pointed out in (1.6), the HL(0) scaling limit of Norris and Turner is equally given by the cylinder $\mathbb{T} \times [0,t]$ (in fact, this is the reason behind the choice of calling t the asymptotic height of the IDLA clusters).

Jerison, Levine and Sheffield also address the average fluctuations around this shape. As a measure of the deviations from the rectangular shape, they introduce the so called *discrepancy function*

$$D_{N,n} := N(1_{A_N(n)} - 1_{\{y \leq n/N^2\}}), \tag{1.15}$$

where $A_N(n)$ is obtained from the cluster $A(n)$ by filling all the squares of side–length 1 such that their top–right corner belongs to $A(n)$, and then rescaling the mesh–size by $1/N$ both in the horizontal and vertical direction:

$$A_N(n) = \Big\{ \Big[\frac{n_1-1}{N}, \frac{n_1}{N}\Big] \times \Big[\frac{n_2-1}{N}, \frac{n_2}{N}\Big] : (n_1, n_2) \in A(n) \Big\} \subset \mathbb{T} \times \mathbb{R}, \qquad (1.16)$$

For such discrepancy function they prove the following.

Theorem 1.3 ([JLS14b], Theorem 3)**.** *Let $\varphi \in C^\infty(\mathbb{T}\times\mathbb{R})$ be a test function of the form*

$$\varphi(x,y) = \sum_{|k|\le K} \alpha_k(y) e^{2\pi i k x},$$

for some finite integer K and some α_k's supported in the annulus $0 < c_1 \le |y| \le c_2$ and such that $\alpha_{-k} = \overline{\alpha_k}$, so that φ is real. Let $T = \lfloor tN^2 \rfloor$. Then, as $N \to \infty$,

$$D_{N,T}(\varphi) := \int_{\mathbb{T}\times\mathbb{R}} D_{N,T}(x,y)\varphi(x,y)\mathrm{d}x\mathrm{d}y \to \mathcal{N}\big(0, S_t^2(\varphi)\big)$$

in distribution, with

$$S_t^2(\varphi) = \sum_{0<|k|\le K} \Big(\frac{1-e^{-4\pi|k|t}}{4\pi|k|}\Big)|\alpha_k(t)|^2.$$

Informally, this tells us that $D_{N,T}(x,y) \to D_t(x)\delta_{y-t}$, where

$$D_t(x) = \sum_{k\in\mathbb{Z}\setminus\{0\}} \sqrt{\frac{1-e^{-4\pi|k|t}}{4\pi|k|}}(\alpha_k + i\beta_k)e^{2\pi i k x},$$

for $(\alpha_k)_{k\in\mathbb{Z}}$, $(\beta_k)_{k\in\mathbb{Z}}$ independent collections of (real) i.i.d. $\mathcal{N}(0,1)$ random variables. Thus, if $\mathcal{W}(t,x)$ denotes the boundary fluctuations of an HL(0) cluster at time t, as defined in (1.7), we find that[6]

$$D_t(x) \overset{(d)}{=} \mathcal{W}(t,x) \qquad (1.17)$$

in the sense of equality in law of random distributions. This shows that not only

[6]Here the equality is intended to hold apart from an absolute multiplicative constant, due to the different Fourier normalizations.

IDLA clusters on the cylinder graph and HL(0) clusters share the same limit, but also they have exactly the same fluctuation process. One could then ask whether they share other asymptotic properties: see Section 3.4 in Chapter 3 for some open questions in this direction.

1.3.1 Asymptotic fluctuations: a different scaling

In the IDLA fluctuations result of Jerison, Levine and Sheffield the authors show that at scaling $1/N$ the fluctuations of the discrepancy function in the vertical direction are killed, while the ones in the horizontal direction converge to the GFF. A natural question[7] then arises: is there a different scaling for the vertical coordinate, such that the associated discrepancy function $\tilde{D}_{N,T}$ converges, in that scaling, to a truly 2–dimensional limit? More precisely, recall the definition of the filled and rescaled cluster $A_N(n)$ from (1.16), and the associated discrepancy function $D_{N,n} = N(1_{A_N(n)} - 1_{\{y\leq n/N^2\}})$. For scaling coefficients a_N, b_N to be determined, define the rescaled cluster

$$\tilde{A}_N(n) = \Big\{ \Big[\frac{n_1-1}{N}, \frac{n_1}{N}\Big] \times \Big[\frac{n_2-1}{b_N}, \frac{n_2}{b_N}\Big] : (n_1,n_2) \in A(n) \Big\} \subset \mathbb{T}\times\mathbb{R},$$

and the rescaled discrepancy function $\tilde{D}_{N,n} = a_N(1_{\tilde{A}_N(n)} - 1_{\{y\leq n/(Nb_N)\}})$.

Question 1.1. Is there a choice of a_N, b_N such that $\tilde{D}_{N,n}$ converges to a non–trivial 2–dimensional limit

$$\tilde{D}_{N,n}(x,y) \to \tilde{D}_t(x,y),$$

as $N\to\infty$ and $n = \lfloor tN^2 \rfloor$, in the sense of distributions?

Although at present we are not able to answer this question, we believe that, due to Theorem 1.2, b_N should be of order at most $\log N$ for large N. To get some insight on how an IDLA cluster should look like in the strip of height $C\log N$ on which the discrepancy function is supported, one could ask the following, simpler question.

[7] We thank Tom Holding for this question.

Question 1.2. Fix $N \gg 1$, and let $(A(n))_{n\geq 0}$ be an IDLA process on the cylinder graph $\mathbb{Z}_N \times \mathbb{Z}$. How does $A(n)$ look like *above its lowest hole*, for $n \gg N$?

Answering this question is the aim of the second part of this thesis, which we devote to the study of the so called *shifted IDLA* process, first introduced in [JLS14b] (see Section 7 therein).

1.4 Shifted IDLA: a recurrent Markov chain

Let $(A(n))_{n\geq 0}$ denote, as in the previous section, an IDLA process on the cylinder graph $\mathbb{Z}_N \times \mathbb{Z}$ starting from $A(0) = \{(x, y) \in \mathbb{Z}_N \times \mathbb{Z} : y \leq 0\}$. We now define an associated Markov process $(A_S(n))_{n\geq 0}$, so called *shifted IDLA*.

Definition 1.1 (Shifted IDLA)**.** Let $\mathcal{S}$ be the map from the space of IDLA configurations Ω to itself defined as follows. For $A \in \Omega$, let

$$k_A := \max\{k \geq 0 : (x, y) \in A \text{ for all } x \in \mathbb{Z}_N, y \leq k\}.$$

Define the *shift operator* $\mathcal{S}$ by

$$\mathcal{S} : A \mapsto \mathcal{S}(A) := \{(x, y - k_A \mathbb{1}(k_A > 0)) : (x, y) \in A\}.$$

Then, if $(A(n))_{n\geq 0}$ denotes the time evolution of a IDLA cluster, define the *shifted chain* $(A_S(n))_{n\geq 0}$ by setting

$$A_S(n) := \mathcal{S}(A(n))$$

for all $n \geq 0$.

In words, $A_S(n)$ is obtained from $A(n)$ by shifting the cluster down in the vertical direction, so that the lowest empty site is at level 1 (i.e. it is of the form $(x, 1)$ for some $x \in \mathbb{Z}_N$). This construction defines a new Markov chain $(A(n))_{n\geq 0}$ on the infinite space Ω of configurations. Note that the chain consists of one communicating class, since from any configuration one can get to $\{(x,y) : y \leq 0\}$, and from there to any other configuration, in finitely many steps with positive probability. Thus either all configurations are recurrent or they are all transient. If it is the case that the chain is recurrent, one could ask if it is also positive recurrent, which would in turn imply the existence of a stationary distribution[8]. This turns out to be indeed the case.

Proposition 1.1. *Fix any integer $N > 0$, and let $(A_S(n))_{n\geq 0}$ denote a shifted IDLA Markov chain on the cylinder graph $\mathbb{Z}_N \times \mathbb{Z}$. Then the chain is positive recurrent, and hence it has a stationary distribution μ_N.*

The above result follows from Proposition 4.1 in Chapter 4, where we prove the same statement for all cylinder graphs of the form $G \times \mathbb{Z}$, with G any finite connected graph. By the ergodic theorem, then, we have that the law of the chain $(A_S(n))_{n\geq 0}$ converges to μ_N as $n \to \infty$. In other words, if we wait long enough then the configuration $A_S(n)$ looks like a configuration drawn from μ_N. But how long do we have to wait? This question, of independent interest, was posed in [JLS14b]: we provide a partial answer to it in Chapter 4. Before going into the details of our approach, let us give some heuristics on what this would tell us about the stationary measure μ_N.

Fix $N \gg 1$ large enough, and let $T_{mix} = T_{mix}(N)$ denote the first $n \geq 0$ such that the law of the cluster $A_S(n)$ is close, in some suitable sense, to the stationary measure μ_N. Suppose we can show that T_{mix} is at most polynomial in N, i.e. there exists a finite absolute constant $m > 0$ such that $T_{mix} \leq N^m$. Then $A_S(N^m)$ is close in law to μ_N. Suppose further that it is possible to generalise Theorem 1.2 to the case $T = N^m$. More precisely, suppose we can show that for any $a > 0$ and $m \geq 1$ there exists a finite constant $C = C(a,m)$ such that the maximal fluctuations of IDLA up to time $T = N^m$ have size at most $C \log N$ with high probability. Then at time T_{mix} the associated shifted IDLA cluster is both close in law to μ_N and

[8]Note that, as long as $N \geq 3$, such stationary distribution can never be a reversible distribution. Indeed, if the chain can go from configuration A to configuration B is one step, it cannot go from B to A in the next step. This makes the detailed balance equation fail, thus showing that the shifted IDLA chain is not reversible.

has height at most $C \log N$ with high probability. This motivates the following conjecture.

Conjecture 1.1. For $\ell > 0$, let $\mathcal{E}_\ell = \{A \in \Omega : h(A) \leq \ell\}$ denote the set of all IDLA configurations with height at most ℓ. Then for any $a > 0$ there exists a positive constant $C' = C'(a)$ such that, for N large enough,

$$\mu_N\big(\mathcal{E}_{C'(a)\log N}\big) \geq 1 - \frac{1}{N^a}.$$

The difficulty in proving this conjecture is twofold. First, one has to come up with a suitable notion of mixing, which can be used for Markov chains on an infinite state space, and which makes no use of the stationary distribution (since we have no information about it). Second, once the notion of mixing has been fixed, one needs to show that the associated mixing time is polynomial.

We attack this problem by introducing an auxiliary notion of mixing that adapts well to coupling techniques, that we call *forget time*.

1.4.1 The IDLA forget time

In [JLS14b] the authors ask the following question: how long does it take for an IDLA process on the cylinder graph to *mostly forget* its initial profile? In other words, let $(A(n))_{n\geq 0}$ and $(A'(n))_{n\geq 0}$ denote two IDLA processes on the cylinder graph $\mathbb{Z}_N \times \mathbb{Z}$, starting from different initial configurations $A(0), A'(0)$, both containing the half–space $\{(x,y) : y \leq 0\}$. Let

$$\mathscr{T} := \inf\{n \geq 0 : A(n) = A'(n)\} \tag{1.18}$$

denote the first time the two processes meet. How large can $\mathscr{T}$ be? Our main result in this direction, namely Theorem 1.4 below, shows that if we are allowed to couple the two processes in a clever way, then $\mathscr{T}$ is not too large.

Notation. Recall that $\Omega = \{E \subseteq \mathbb{Z}_N \times \mathbb{Z} : E \supseteq \{(x,y) : y \leq 0\}\}$. For any $E \in \Omega$, write $|E|$ for the cardinality of $E \cap \{(x,y) : y > 0\}$ and recall that $h(E) := \max\{y :$

$(x, y) \in E$ for some $x \in \mathbb{Z}_N\}$ denotes the height of E. Finally, for $E, F \in \Omega$ write $E\Delta F = (E \setminus F) \cup (F \setminus E)$ for their symmetric difference, and note that if $|E| = |F|$ then $|E\Delta F|$ is an even number.

Theorem 1.4. *Let $(A(n))_{n\geq 0}$, $(A'(n))_{n\geq 0}$ be two IDLA processes as above, and assume that $|A(0)| = |A'(0)| =: n_0$. Let $h_0 := \max\{h(A(0)), h(A'(0))\}$, and $2d := |A(0)\Delta A'(0)|$. Recall the definition of the forget time $\mathscr{T}$ from* (1.18)*. Then there exists a coupling of the two IDLA processes such that, if $\mathbb{P}$ is the probability measure associated to their joint law,*

$$\mathbb{P}(\mathscr{T} > t) \leq \bar{\mathscr{C}} e^{-\bar{\lambda} t} \tag{1.19}$$

for positive constants $\bar{\mathscr{C}}, \bar{\lambda} > 0$, with $\bar{\lambda} = \bar{\lambda}(d, N)$ and $\bar{\mathscr{C}} = \bar{\mathscr{C}}(\bar{\lambda}, d, N, h_0)$.

Note in particular that the decay rate $\bar{\lambda}$ does not depend on the initial height of the clusters.

Remark. The above theorem is true in greater generality. Indeed, we can prove that it holds for all cylinder graphs of the form $G \times \mathbb{Z}$, for any finite connected graph G (cf. Theorem 4.1 in Chapter 4). Moreover, we can allow the driving random walks to have a deterministic drift ∂ in the vertical direction. In this setting, the constants $\bar{\mathscr{C}}$ and $\bar{\lambda}$ would become $\bar{\lambda} = \bar{\lambda}(d, G, \partial)$ and $\bar{\mathscr{C}} = \bar{\mathscr{C}}(\bar{\lambda}, d, G, h_0, \partial)$.

To prove Theorem 1.4 we exhibit a coupling that achieves (1.19). This coupling is divided into several trials, each trial having a positive probability of success (by which we mean that the processes are coupled by the end of the trial). At the beginning of each trial, we have two IDLA configurations that differ by a controlled number of particles. We make them match as quickly as possible using the *freezing coupling*, by which we mean that we use the same driving random walks to grow both processes, and whenever a walk happens to exit from one cluster but not from the other we *freeze* it, to restart it at a later time. The fact that this is a valid coupling follows from the abelian property of IDLA. With this procedure the two configurations will match very quickly (apart from the frozen particles). Now we want to make sure that these matching configurations are nice enough (in some sense to be specified later) so that the frozen particles will couple with high probability upon release. We call this *extra releases phase*: as the name suggests, we achieve nice configurations simply by releasing enough extra particles while using the trivial coupling (so that the two configurations continue to match), until the processes hit a nice configuration. At this point we release the frozen particles. If they couple, we declare the trial successful, and otherwise we move to the next trial.

The main technical difficulty of this procedure is that all our estimates depend on the so called *discrepancy*[9] of the IDLA clusters, that is the difference between the height of the cluster and the height of the lowest empty site. Thus, before starting a new trial we must have a strong control on the discrepancy of the cluster.

1.5 Organisation of the thesis

The thesis is divided into two parts: in Part I (Chapters 2 and 3) we discuss Hastings–Levitov growth, while Part II (Chapter 4) is concerned with Internal DLA on cylinder graphs. The material presented in Part I consists of the paper *Fluctuation results for Hastings–Levitov planar growth*, published in Probability Theory and Related Fields [Sil15]. The material discussed in Part II is an extract of the paper *The reset time of Internal DLA on a cylinder* [LS], which is currently being written jointly with Lionel Levine.

We begin Chapter 2 by collecting some preliminary estimates for the basic conformal maps F (cf. Figure 1.4) and its inverse, to then prove the pointwise fluctuation result (cf. Theorem 2.1). We then explain in Chapter 3 how to use this result to obtain a global description of the fluctuations (cf. Theorem 3.3), and construct the boundary values process (cf. Section 3.3 therein). We conclude the chapter, and the first part of the thesis, by collecting some open questions. Chapter 4 is then devoted to the study of the forget time of IDLA on cylinder graphs. We start by considering general cylinder graphs of the form $G \times \mathbb{Z}$, and obtain an upper bound

[9] Not to be confused with the discrepancy function defined in (1.15).

for the forget time via coupling techniques (cf. Theorem 4.1). We then specialise to $G = \mathbb{Z}_N$, and show that this bound is at most polynomial in N, for N large enough.

Part I

Hastings–Levitov growth

Chapter 2

Pointwise fluctuations for HL(0)

In this chapter we present a detailed proof of the global fluctuations result. Theorem A.1 on local fluctuations is obtained by the same arguments, and we leave its discussion for Appendix A.1. The chapter is organised as follows. In Section 2.1 we collect some preliminary estimates for the basic conformal maps F, G. We then introduce in Section 2.2 our main tools, namely two sequences of backwards martingale difference arrays (2.10), and prove Theorem 2.1 on pointwise fluctuations.

2.1 Preliminary estimates

Fix a particle P as in the introduction, and, if $\mathbb{D}$ denotes the open unit disc, let $K_0 = \overline{\mathbb{D}}$, $K = \overline{\mathbb{D}} \cup P$. Moreover, set $D_0 = (\mathbb{C} \cup \{\infty\}) \setminus K_0$, $D = (\mathbb{C} \cup \{\infty\}) \setminus K$. It follows from the Riemann Mapping Theorem that there exists a unique conformal map[1] $F : D_0 \to D$ such that $F(z) = e^c z + \mathcal{O}(1)$ as $|z| \to \infty$, for some constant $c \in \mathbb{R}$. We assume that the particle P is regular enough so that F extends continuously to the boundary of D_0. Set $G = F^{-1}$.

Conformal maps are very rigid, and simply from the definition of F, G we can deduce the following properties:

(P1) $|F(z)| > |z|$ for all $z \in D_0$, $|G(z)| < |z|$ for all $z \in D$,

[1] Throughout this chapter and the next, by conformal map we mean a conformal isomorphism.

(P2) there exists a constant $C > 0$ such that $|F(z)| \le C|z|$ for all $z \in D_0$ and $|G(z)| \ge |z|/C$ for all $z \in D$.

Indeed, (P1) follows trivially from Schwarz lemma, while (P2) is a consequence of the prescribed behaviour at infinity for F, G.

Notation. Throughout the chapter, C denotes a finite, positive constant which can change from line to line, and which is either absolute or only depends on the maps F, G. Whenever this constant depends on other parameters, say $\alpha, \beta \ldots$, we make it explicit in the notation by using $C(\alpha, \beta \ldots)$.

In order to obtain finer distortion estimates for the maps F, G it is often useful to relate the logarithmic capacity c to geometric properties of the particle P. Assume the following:

Assumption 2.1. There exists $\delta > 0$ such that

$$P \subseteq \{z \in \mathbb{C} : |z-1| \le \delta\}, \quad 1+\delta \in P, \quad P = \{\bar{z} : z \in P\}. \tag{2.1}$$

We regard δ as measuring the diameter of the particle P.

Remark. Assumption 2.1 is in force throughout Chapters 2 and 3.

The following result appears in [NT12] (cf. Proposition 4.1 and Corollary 4.2 therein).

Proposition 2.1. *There exists an absolute constant $C > 0$ such that, for all $z \in D$: $|z-1| > 2\delta$, it holds:*

$$\left| \log\left(\frac{G(z)}{z}\right) + c \right| \le \frac{Cc}{|z-1|}, \qquad \left| \frac{\mathrm{d}}{\mathrm{d}z} \log\left(\frac{G(z)}{z}\right) \right| \le \frac{Cc}{|z-1|^2}. \tag{2.2}$$

Moreover,

$$\frac{\delta^2}{6} \le c \le \frac{3\delta^2}{4} \tag{2.3}$$

for δ small enough.

In light of (2.3) we use c and δ interchangeably, and all statements are intended to hold for c, δ small enough. Combining (2.2) and (2.3) we deduce the following improved bound.

Corollary 2.1. *There exists an absolute constant $C > 0$ such that, for all $z \in D$: $|z-1| > 2\delta$, it holds:*

$$\left| \log\left(\frac{G(z)}{z}\right) + c\,\frac{z+1}{z-1} \right| \le \frac{Cc^{3/2}|z|}{|z-1|^2}. \tag{2.4}$$

Proof. We follow the proof of [NT12], Proposition 4.1. Let u, v denote the real and imaginary part of $\log \frac{G(z)}{z}$ respectively, so that they are harmonic functions on D. Then by optional stopping $u(z) = -\mathbb{E}(\log|B_T|) < 0$, T being the first hitting time of K for a Brownian Motion B starting from z. Introduce the particle $P_1 \supset P$ defined by $P_1 = \left\{z \in D_0 : \left|\frac{z-1}{z+1}\right| \le r\right\}$, for $r = \delta/(2-\delta)$, and set $D^1 = (\mathbb{C} \cup \{\infty\}) \setminus (\overline{\mathbb{D}} \cup P_1)$. Then the unique conformal map $G_1 : D^1 \to D_0$ satisfying $G_1(\infty) = \infty$ and $G_1'(\infty) > 0$ is given by

$$G_1(z) = \frac{z(\gamma z - 1)}{z - \gamma} \qquad \text{for} \ \ \gamma = \frac{1-r^2}{1+r^2}.$$

Set $F_1 = G_1^{-1}$, and $A = \{z \in \partial P_1 : |z| > 1\}$. Then $G_1(A) = \{e^{i\vartheta} : |\vartheta| < \vartheta_0\}$ with $\vartheta_0 = \cos^{-1}\gamma$. Moreover, $u \circ F_1$ is harmonic and bounded on D_0 and, using that

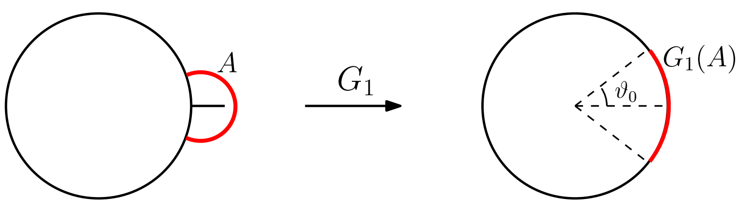

$\frac{1}{2\pi}\int_{|\vartheta|\le\vartheta_0} (u \circ F_1)(e^{i\vartheta})\mathrm{d}\vartheta = -c$, the optional stopping theorem yields

$$(u \circ F_1)(z) = -c + \frac{1}{2\pi}\int_{|\vartheta|\le\vartheta_0} (u \circ F_1)(e^{i\vartheta})\,\mathrm{Re}\left(\frac{2e^{i\vartheta}}{z - e^{i\vartheta}}\right)\mathrm{d}\vartheta\,.$$

It follows that

$$(u\circ F_1)(z)+c+\mathrm{Re}\Big(\frac{2c}{z-1}\Big)=\frac{1}{2\pi}\int_{|\vartheta|\leq\vartheta_0}(u\circ F_1)(e^{i\vartheta})\Big[\,\mathrm{Re}\Big(\frac{2e^{i\vartheta}}{z-e^{i\vartheta}}\Big)-\mathrm{Re}\Big(\frac{2}{z-1}\Big)\Big]\mathrm{d}\vartheta\,.$$

Using (2.3), then, we find $1-\cos\vartheta_0=1-\gamma\asymp\delta^2/2\leq C\sqrt{c}$, from which

$$\left|\,\mathrm{Re}\Big(\frac{2e^{i\vartheta}}{z-e^{i\vartheta}}\Big)-\mathrm{Re}\Big(\frac{2}{z-1}\Big)\right|\leq\frac{2|z||1-e^{i\vartheta_0}|}{|z-e^{i\vartheta}||z-1|}\leq\frac{C|z|\sqrt{c}}{|z-e^{i\vartheta}||z-1|}\leq\frac{C|z|\sqrt{c}}{\mathrm{dist}(z,G_1(A))^2}\,.$$

Since $(u\circ F_1)(e^{i\vartheta})<0$ for all ϑ in the integration range, we find

$$\left|(u\circ F_1)(z)+c+\mathrm{Re}\Big(\frac{2c}{z-1}\Big)\right|\leq\frac{C|z|\sqrt{c}}{\mathrm{dist}(z,G_1(A))^2}\left|\frac{1}{2\pi}\int_{|\vartheta|\leq\vartheta_0}(u\circ F_1)(e^{i\vartheta})\mathrm{d}\vartheta\right|=\frac{C|z|c^{3/2}}{\mathrm{dist}(z,G_1(A))^2}\,.$$

Now take $G_1(z)$ in place of z, to get

$$\begin{aligned}\left|u(z)+c+\mathrm{Re}\Big(\frac{2c}{z-1}\Big)\right|&\leq\left|u(z)+c+\mathrm{Re}\Big(\frac{2c}{G_1(z)-1}\Big)\right|+\left|\mathrm{Re}\Big(\frac{2c}{G_1(z)-1}\Big)-\mathrm{Re}\Big(\frac{2c}{z-1}\Big)\right|\\&\leq\frac{C|z|c^{3/2}}{\mathrm{dist}(G_1(z),G_1(A))^2}+\frac{2c|G_1(z)-z|}{|G_1(z)-1||z-1|}\,,\end{aligned}\tag{2.5}$$

where in the second inequality we have used that $|G(z)|<|z|$. Using the explicit expression for G_1, one shows that $|G_1(z)-1|>|z-1|/2$ and $|G_1(z)-z|\leq C\sqrt{c}|z|$ for δ small enough, from which

$$\frac{2c|G_1(z)-z|}{|G_1(z)-1||z-1|}<\frac{C|z|c^{3/2}}{|z-1|^2}\,.\tag{2.6}$$

Moreover, reasoning as in [NT12] one shows that there exists an absolute constant C_1 such that $\mathrm{dist}(G_1(z),G_1(A))\geq|z-1|/C_1$. Putting this together with (2.5) and (2.6) we finally obtain

$$\left|u(z)+c+\mathrm{Re}\Big(\frac{2c}{z-1}\Big)\right|\leq\frac{C|z|c^{3/2}}{|z-1|^2}$$

for all $z\in D$ such that $|z-1|>2\delta$.

Now note that $u(z)+c+\mathrm{Re}\left(\frac{2c}{z-1}\right)$ and $v(z)+\mathrm{Im}\left(\frac{2c}{z-1}\right)$ are the real and imaginary part of the holomorphic function $z \mapsto \log \frac{G(z)}{z} + c + \frac{2c}{z-1}$ on $D \cap \{z : |z-1| > 2\delta\}$. It then follows by Cauchy's integral formula that

$$\left|\nabla\left(v(z) + \mathrm{Im}\left(\frac{2c}{z-1}\right)\right)\right| = \left|\nabla\left(u(z) + c + \mathrm{Re}\left(\frac{2c}{z-1}\right)\right)\right| \le \frac{C|z|c^{3/2}}{|z-1|^3}\,.$$

Finally, using that $\left|v(z) + \mathrm{Im}\left(\frac{2c}{z-1}\right)\right| \to 0$ as $|z| \to \infty$, we get

$$\left|v(z)+\mathrm{Im}\left(\frac{2c}{z-1}\right)\right| \le \int_0^\infty \left|\nabla\left(v(z+s(z-1))+\mathrm{Im}\left(\frac{2c}{z+s(z-1)-1}\right)\right)\right|\cdot|z-1|\mathrm{d}s \le \frac{C|z|c^{3/2}}{|z-1|^2}\,,$$

which concludes the proof. □

We combine Proposition 2.1 and Corollary 2.1 to obtain corresponding estimates for the function F, which are collected below.

Corollary 2.2. *There exists a constant $C > 0$ such that, for all $z \in D_0$ with $|F(z)-1| > 2\delta$, it holds:*

$$|F(z) - e^c z| \le \frac{Cc|z|}{|z|-1}\,, \qquad \left|\log\left(\frac{F(z)}{z}\right) - c\,\frac{z+1}{z-1}\right| \le \frac{Cc^{3/2}|z|^2}{(|z|-1)^3}.$$

Proof. For the first inequality, note that (2.2) readily implies that

$$|G(z) - e^{-c}z| \le \frac{Cc|z|}{|z-1|} \tag{2.7}$$

for all $z \in D : |z-1| > \delta$, and δ small enough. Therefore, by setting $w = F(z)$ and using (P1)-(P2), we obtain

$$|F(z) - e^c z| = e^c|G(w) - e^{-c}w| \le \frac{Cc|w|}{|w-1|} \le \frac{C'c|z|}{|z|-1}\,, \tag{2.8}$$

for all $z \in D_0$ such that $|F(z)-1| > 2\delta$ and δ small enough, as claimed.

For the second inequality, note that, since $|F(z)-1| > 2\delta$ by assumption, we can use (2.4) to get

$$\left|\log\frac{F(z)}{z}-c\frac{z+1}{z-1}\right|\le\left|\log\frac{G(w)}{w}+c\frac{w+1}{w-1}\right|+2c\left|\frac{1}{w-1}-\frac{1}{G(w)-1}\right|$$
$$\le\frac{Cc^{3/2}|w|}{|w-1|^2}+\frac{2c|G(w)-w|}{|w-1||G(w)-1|}\le\frac{Cc^{3/2}|z|}{(|z|-1)^2}+\frac{2c|G(w)-w|}{(|z|-1)^2}.$$

Moreover, it follows from (2.8) that

$$|G(w)-w|\le|F(z)-e^cz|+(1-e^{-c})|w|\le\frac{Cc|z|}{|z|-1}+Cc|z|\le\frac{Cc|z|^2}{|z|-1}$$

for c small enough. Putting all together, we end up with

$$\left|\log\frac{F(z)}{z}-c\frac{z+1}{z-1}\right|\le\frac{Cc^{3/2}|z|}{(|z|-1)^2}\left(1+\frac{\sqrt{c}|z|}{|z|-1}\right)\le\frac{2Cc^{3/2}|z|^2}{(|z|-1)^3}$$

for c small enough, as claimed. □

2.2 Pointwise fluctuations

In this section we prove that, fixed $t\ge 0$ and $z\in\mathbb{C}\backslash\overline{\mathbb{D}}$, the fluctuations of $\log\Phi_{\lfloor nt\rfloor}(z)$ around its mean are given by a complex Gaussian random variable, whose variance is independent of $\mathrm{Arg}(z)$.

Notation. Throughout the thesis $\mathcal{N}(\mu,\sigma^2)$ denotes the Gaussian distribution on $\mathbb{R}$ with mean μ and variance σ^2.

Our main result is the following.

Theorem 2.1. *Fix any $t>0$. Pick $a\in[-\pi,\pi)$, $\sigma>0$, and let $z=e^{ia+\sigma}$. Define $v_t^2(\sigma):=\log\frac{1-e^{-2(\sigma+t)}}{1-e^{-2\sigma}}$, and let $\mathcal{F}^\sigma(t,e^{ia})$ be a complex Gaussian random variable with i.i.d. real and imaginary part, distributed according to $\mathcal{N}(0,v_t^2(\sigma))$. Then it holds:*

$$\frac{1}{\sqrt{c}}\left(\log\frac{\Phi_{\lfloor nt\rfloor}(z)}{z}-\lfloor nt\rfloor c\right)\longrightarrow\mathcal{F}^\sigma(t,e^{ia})$$

in distribution as $n\to\infty$, $c\to 0$ and $nc\to 1$.

We prove Theorem 2.1 for $\sigma \leq 1$, which we assume without further notice. This entails no loss of generality (see discussion in Section 3.2), and it has the advantage of slightly simplifying the notation.

Remark. From this point onwards, with the exception of Appendix A.1, whenever we write $n \to \infty$ or $c \to 0$ we mean that $n \to \infty, c \to 0$ and $nc \to 1$.

Our main tool for the proof of Theorem 2.1 consists of two sequences of backwards martingale difference arrays, that we now define. Note that $\left|\frac{\Phi_{\lfloor nt \rfloor}(z)}{z}\right| > 1$ for all $z \in D_0$, so $\log \frac{\Phi_{\lfloor nt \rfloor}(z)}{z}$ defines a holomorphic function on D_0 (cf. [BF09], Theorem II.2.9). We fix the branch of the logarithm by requiring that $\log \frac{\Phi_{\lfloor nt \rfloor}(z)}{z} \to c$ as $z \to \infty$. For $\sigma > 0$ and $a \in [-\pi, \pi)$ as in Theorem 2.1, then, we find:

$$\log \frac{\Phi_{\lfloor nt \rfloor}(e^{ia+\sigma})}{e^{ia+\sigma}} = \sum_{k=1}^{\lfloor nt \rfloor} \log \frac{F_k \circ F_{k+1} \circ \cdots \circ F_{\lfloor nt \rfloor}(e^{ia+\sigma})}{F_{k+1} \circ \cdots \circ F_{\lfloor nt \rfloor}(e^{ia+\sigma})} = \sum_{k=1}^{\lfloor nt \rfloor} \log \frac{F(e^{-i\Theta_k} Z^\sigma_{k,\lfloor nt \rfloor}(a))}{e^{-i\Theta_k} Z^\sigma_{k,\lfloor nt \rfloor}(a)} ,$$

where $Z^\sigma_{k,\lfloor nt \rfloor}(a) := F_{k+1} \circ \cdots \circ F_{\lfloor nt \rfloor}(e^{ia+\sigma})$. Moreover, if $\mathscr{F}_{k,n} := \sigma(\Theta_k, \Theta_{k+1} \ldots \Theta_n)$, then

$$\begin{aligned} \mathbb{E}\left(\log \frac{F(e^{-i\Theta_k} Z^\sigma_{k,\lfloor nt \rfloor}(a))}{e^{-i\Theta_k} Z^\sigma_{k,\lfloor nt \rfloor}(a)} \middle| \mathscr{F}_{k+1,\lfloor nt \rfloor} \right) &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \log \frac{F(e^{-i\vartheta} Z^\sigma_{k,\lfloor nt \rfloor}(a))}{e^{-i\vartheta} Z^\sigma_{k,\lfloor nt \rfloor}(a)} \, \mathrm{d}\vartheta \\ &= \lim_{|z| \to \infty} \log \frac{F(z)}{z} = c \, . \end{aligned} \tag{2.9}$$

We therefore set

$$\begin{aligned} X^\sigma_{k,\lfloor nt \rfloor}(a) &= \frac{1}{\sqrt{c}} \left(\log \left| \frac{F(e^{-i\Theta_k} Z^\sigma_{k,\lfloor nt \rfloor}(a))}{e^{-i\Theta_k} Z^\sigma_{k,\lfloor nt \rfloor}(a)} \right| - c \right), \\ Y^\sigma_{k,\lfloor nt \rfloor}(a) &= \frac{1}{\sqrt{c}} \operatorname{Arg}\left(\frac{F(e^{-i\Theta_k} Z^\sigma_{k,\lfloor nt \rfloor}(a))}{e^{-i\Theta_k} Z^\sigma_{k,\lfloor nt \rfloor}(a)} \right) \end{aligned} \tag{2.10}$$

where $\operatorname{Arg}(z) \in [-\pi, \pi)$. The above computation then shows that $(X^\sigma_{k,\lfloor nt \rfloor}(a))_{k \leq \lfloor nt \rfloor}$ and $(Y^\sigma_{k,\lfloor nt \rfloor}(a))_{k \leq \lfloor nt \rfloor}$ form sequences of backwards martingale arrays with respect to the same filtration $(\mathscr{F}_{k,n})_{k \leq \lfloor nt \rfloor}$ (see [Bil95], Section 35 for the definition of martingale arrays). Moreover,

$$\sum_{k=1}^{\lfloor nt \rfloor} \big(X^{\sigma}_{k,\lfloor nt \rfloor}(a) + i\, Y^{\sigma}_{k,\lfloor nt \rfloor}(a) \big) = \frac{1}{\sqrt{c}} \left(\log \frac{\Phi_{\lfloor nt \rfloor}(e^{ia+\sigma})}{e^{ia+\sigma}} - \lfloor nt \rfloor c \right).$$

Remark. Throughout the chapter we identify $\mathbb{C}$ with $\mathbb{R}^2$, and often refer to complex random variables as random vectors and vice versa, depending on which point of view we seek to emphasize.

The following result appears in [McL74], Corollary 2.8, as a Central Limit Theorem for (forward) martingale difference arrays (see also [Bil95], Theorem 35.12). It is straightforward to adapt the proof to backwards martingale difference arrays (we do this is Appendix B for completeness), to obtain the following.

Theorem 2.2 (McLeish, [McL74])**.** *Let $(\mathcal{X}_{k,n})_{1 \le k \le n}$ be a backwards martingale difference array with respect to $\mathscr{F}_{k,n} = \sigma(\mathcal{X}_{k,n} \ldots \mathcal{X}_{n,n})$. Let $S_{k,n} = \sum_{j=k}^{n} \mathcal{X}_{j,n}$. Assume that:*

(I) for all $\eta > 0$, $\displaystyle\sum_{k=1}^{n} \mathcal{X}_{k,n}^2 \mathbb{1}(|\mathcal{X}_{k,n}| > \eta) \to 0$ in probability as $n \to \infty$,

(II) $\displaystyle\sum_{k=1}^{n} \mathcal{X}_{k,n}^2 \to s^2$ in probability as $n \to \infty$, for some $s^2 > 0$.

Then $S_{1,n}$ converges in distribution to $\mathcal{N}(0, s^2)$.

Note that the above theorem is concerned with scalar random variables, while we have 2–dimensional vectors. In order to reduce to the scalar case, recall that by the Cramér–Wold Theorem (cf. [Dur10], Theorem 3.9.5) it suffices to prove convergence in distribution of all linear combinations of the vector entries. To this end, pick any $\alpha, \beta \in \mathbb{R}$, and note that by linearity $\big(\alpha X^{\sigma}_{k,\lfloor nt \rfloor}(a) + \beta Y^{\sigma}_{k,\lfloor nt \rfloor}(a)\big)_{k \le \lfloor nt \rfloor}$ is again a backwards martingale difference array with respect to the filtration $(\mathscr{F}_{k,\lfloor nt \rfloor})_{k \le \lfloor nt \rfloor}$. We are going to apply Theorem 2.2 to this linear combination. To this end, we collect here some estimates for $(X^{\sigma}_{k,\lfloor nt \rfloor}(a))$ and $(Y^{\sigma}_{k,\lfloor nt \rfloor}(a))$. Since a and σ are fixed, we omit them from the notation throughout this section.

Lemma 2.1. *There exists a constant $C > 0$ such that, for c small enough, it holds $|X_{k,\lfloor nt\rfloor}| \leq C/\sqrt{c}$, $|Y_{k,\lfloor nt\rfloor}| \leq C/\sqrt{c}$ for all $n \geq 1$, $k \leq \lfloor nt\rfloor$.*

Proof. It follows from (P1) that $\log\left|\frac{F(e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor})}{e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor}}\right| > 0$, from which $X_{k,\lfloor nt\rfloor} > -\sqrt{c}$. Moreover, (P2) gives $|X_{k,\lfloor nt\rfloor}| \leq C/\sqrt{c} + \sqrt{c} \leq 2C/\sqrt{c}$ for some constant $C > 0$ and c small enough. Finally, since $\mathrm{Arg}\big(\frac{F(e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor})}{e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor}}\big) \in [-\pi, \pi)$, we have $|Y_{k,\lfloor nt\rfloor}| \leq \pi/\sqrt{c}$. □

Much better estimates can be obtained by restricting to a certain *good event*, which is shown in [NT12] to have high probability for large n. The following result identifies this event.

Theorem 2.3 ([NT12], Proposition 5.1)**.** *Fix a positive integer m and a constant $\varepsilon > 0$. For $n \leq m$ define the events*

$$\begin{aligned} E_n(\varepsilon) := &\{|e^{-cn}\Phi_n(z) - z| < \varepsilon e^{6\varepsilon} \text{ for all } z : |z| \geq e^{5\varepsilon}\} \\ &\cap \{|e^{cn}\Gamma_n(z) - z| < \varepsilon e^{5\varepsilon + cn} \text{ for all } z : |z| \geq e^{cn+4\varepsilon}\}, \end{aligned}$$

and set $E(m,\varepsilon) := \bigcap_{n=1}^m E_n(\varepsilon)$. Then it holds

$$\mathbb{P}(E(m,\varepsilon)^c) \leq C(m + \varepsilon^{-2})e^{-\varepsilon^3/(Cc)}$$

for some constant $C > 0$. In particular, by setting $m = \lfloor \delta^{-6} \rfloor$ and $\varepsilon = \delta^{2/3}\log(1/\delta)$, one obtains that $\delta^{-k}\mathbb{P}(E(m,\varepsilon)^c) \to 0$ as $\delta \to 0$, for any $k \geq 0$.

We refer to $E(m,\varepsilon)$ as the *good event.*

Remark. Without further notice, we take $m = \lfloor \delta^{-6} \rfloor$ and $\varepsilon = \delta^{2/3}\log(1/\delta)$ as in the last part of the above theorem, so that[2] $\varepsilon \gg \delta$ and $m \gg n$.

Lemma 2.2. *Assume that $\lim_{\delta\to 0}\frac{\delta}{\sigma} = 0$. Then there exists a constant $C(t) > 0$ such that, for n large enough, on the good event $E(m,\varepsilon)$ it holds*

$$\max_{k \leq \lfloor nt\rfloor}\Big\{|X_{k,\lfloor nt\rfloor}| \vee |Y_{k,\lfloor nt\rfloor}|\Big\} \leq C(t)\frac{\sqrt{c}}{\sigma}. \tag{2.11}$$

[2]If $a, b \to 0$ (respectively $a, b \to \infty$) write $a \gg b$ to mean $\lim_{a,b\to 0}\frac{b}{a} = 0$ (respectively $\lim_{a,b\to\infty}\frac{b}{a} = 0$).

Proof. For any $k \leq \lfloor nt \rfloor$ we have

$$\begin{aligned}\max\left\{|X_{k,\lfloor nt\rfloor}|;|Y_{k,\lfloor nt\rfloor}|\right\} &\leq \frac{1}{\sqrt{c}}\Big(\big|\log F(e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor}) - \log(e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor})\big| + c\Big)\\ &\leq \frac{1}{\sqrt{c}}\left[\Big(\sup_{|\xi|\geq e^{\sigma}}\frac{1}{|\xi|}\Big)\cdot\big|F(e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor}) - e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor}\big| + c\right],\end{aligned}$$

where the last inequality follows from the mean values theorem, and the fact that $|F(e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor})| > |e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor}| > e^{\sigma}$ almost surely by (P1). Now note that $|F(e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor}) - 1| > |Z_{k,\lfloor nt\rfloor}| - 1 \geq \sigma \gg 2\delta$, so by Corollary 2.2 we have

$$|F(z) - z| \leq |F(z) - e^c z| + (e^c - 1)|z| \leq \frac{Cc|z|}{|z|-1} + 2c|z| \leq \frac{2Cc|z|^2}{|z|-1} \tag{2.12}$$

for $z = e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor}$. Moreover, since we are on $E(m,\varepsilon)$, there exists a constant $C(t)$ depending only on t such that $e^{\sigma} < |Z_{k,\lfloor nt\rfloor}| \leq C(t)$. This, together with (2.12), yields

$$\max_{k\leq\lfloor nt\rfloor}\left\{|X_{k,\lfloor nt\rfloor}| \vee |Y_{k,\lfloor nt\rfloor}|\right\} \leq \frac{1}{\sqrt{c}}\left(\frac{2Cc|Z_{k,\lfloor nt\rfloor}|^2}{|Z_{k,\lfloor nt\rfloor}|-1} + c\right) \leq \frac{C(t)\sqrt{c}}{e^{\sigma}-1} + \sqrt{c} \leq 2C(t)\cdot\frac{\sqrt{c}}{\sigma}$$

as claimed. □

We now make use of the above bounds to prove that the backwards martingale difference array $(\mathcal{X}_{k,\lfloor nt\rfloor})_{k\leq\lfloor nt\rfloor}$, with $\mathcal{X}_{k,\lfloor nt\rfloor} := \alpha X_{k,\lfloor nt\rfloor} + \beta Y_{k,\lfloor nt\rfloor}$, satisfies Assumptions (I)-(II) of Theorem 2.2. In doing so, we provide an explicit formula for the limiting variance.

Lemma 2.3. *Assume that* $\lim_{\delta\to 0}\frac{\delta}{\sigma} = 0$. *Then for all* $\eta > 0$ *it holds*

$$\sum_{k=1}^{\lfloor nt\rfloor}\mathcal{X}^2_{k,\lfloor nt\rfloor}\mathbb{1}(|\mathcal{X}_{k,\lfloor nt\rfloor}| > \eta) \to 0$$

in probability as $n \to \infty$.

Proof. For any $\epsilon > 0$ we have:

$$\mathbb{P}\bigg(\sum_{k=1}^{\lfloor nt \rfloor} \mathcal{X}_{k,\lfloor nt \rfloor}^2 \mathbb{1}(|\mathcal{X}_{k,\lfloor nt \rfloor}| > \eta) > \epsilon\bigg) \le \mathbb{P}\bigg(\max_{1\le k\le \lfloor nt \rfloor} |\mathcal{X}_{k,\lfloor nt \rfloor}| > \eta\bigg) \le \frac{1}{\eta}\mathbb{E}\bigg(\max_{1\le k\le \lfloor nt \rfloor} |\mathcal{X}_{k,\lfloor nt \rfloor}|\bigg)$$
$$= \frac{1}{\eta}\mathbb{E}\bigg(\max_{1\le k\le \lfloor nt \rfloor} |\mathcal{X}_{k,\lfloor nt \rfloor}|\,; E(m,\varepsilon)^c\bigg) + \frac{1}{\eta}\mathbb{E}\bigg(\max_{1\le k\le \lfloor nt \rfloor} |\mathcal{X}_{k,\lfloor nt \rfloor}|\,; E(m,\varepsilon)\bigg).$$

The fact that the first term in the r.h.s. converges to zero as $n \to \infty$ follows from Lemma 2.1 and Theorem 2.3, while convergence to zero of the second term is a straightforward consequence of Lemma 2.2. □

We now concentrate on Assumption (II). The first step consists in replacing condition (II) with a more convenient one, involving conditional second moments. The following result shows that, provided σ is large enough with respect to c, this is allowed.

Lemma 2.4. *Assume* $\lim_{\delta\to 0} \frac{\sqrt{\delta}}{\sigma} = 0$, *and that* $\sum_{k=1}^{\lfloor nt \rfloor} \mathbb{E}(\mathcal{X}_{k,\lfloor nt \rfloor}^2 | \mathscr{F}_{k+1,\lfloor nt \rfloor}) \to s^2$ *in probability as* $n \to \infty$ *for some* $s^2 > 0$*. Then also* $\sum_{k=1}^{\lfloor nt \rfloor} \mathcal{X}_{k,\lfloor nt \rfloor}^2 \to s^2$ *in probability as* $n \to \infty$*.*

Proof. Let $M_{k,\lfloor nt \rfloor} := \mathcal{X}_{k,\lfloor nt \rfloor}^2 - \mathbb{E}(\mathcal{X}_{k,\lfloor nt \rfloor}^2 | \mathscr{F}_{k+1,\lfloor nt \rfloor})$. It is readily checked that $(M_{k,\lfloor nt \rfloor})_{k\le \lfloor nt \rfloor}$ is a backwards martingale difference array with respect to the filtration $(\mathscr{F}_{k,\lfloor nt \rfloor})_{k\le \lfloor nt \rfloor}$. We aim to show that for any $\eta > 0$ it holds

$$\mathbb{P}\bigg(\Big|\sum_{k=1}^{\lfloor nt \rfloor} M_{k,\lfloor nt \rfloor}\Big| > \eta\bigg) \to 0$$

as $n \to \infty$. Indeed,

$$\mathbb{P}\bigg(\Big|\sum_{k=1}^{\lfloor nt \rfloor} M_{k,\lfloor nt \rfloor}\Big| > \eta\bigg) \le \frac{1}{\eta^2}\mathbb{E}\bigg(\Big[\sum_{k=1}^{\lfloor nt \rfloor} M_{k,\lfloor nt \rfloor}\Big]^2\bigg) = \frac{1}{\eta^2}\sum_{k=1}^{\lfloor nt \rfloor} \mathbb{E}(M_{k,\lfloor nt \rfloor}^2) \overset{(*)}{\le} \frac{1}{\eta^2}\sum_{k=1}^{\lfloor nt \rfloor} \mathbb{E}(\mathcal{X}_{k,\lfloor nt \rfloor}^4)$$
$$= \frac{1}{\eta^2}\sum_{k=1}^{\lfloor nt \rfloor} \mathbb{E}(\mathcal{X}_{k,\lfloor nt \rfloor}^4; E(m,\varepsilon)^c) + \frac{1}{\eta^2}\sum_{k=1}^{\lfloor nt \rfloor} \mathbb{E}(\mathcal{X}_{k,\lfloor nt \rfloor}^4; E(m,\varepsilon))\,.$$

Above, $(*)$ follows from the general inequality $\mathbb{E}((X - \mathbb{E}(X))^2) \le \mathbb{E}(X^2)$, which we apply to each term with respect to $\mathbb{E}(\,\cdot\,|\mathscr{F}_{k,\lfloor nt\rfloor})$. The fact that both terms in the r.h.s. converge to zero as $n\to\infty$ is now a consequence of the bounds for $|X_{k,\lfloor nt\rfloor}|$ and $|Y_{k,\lfloor nt\rfloor}|$, and hence for $|\mathcal{X}_{k,\lfloor nt\rfloor}|$, obtained in Lemmas 2.1–2.2. $\square$

In light of the above result, it remains to compute the limit in probability of

$$\begin{aligned}\sum_{k=1}^{\lfloor nt\rfloor}\mathbb{E}(\mathcal{X}_{k,\lfloor nt\rfloor}^2|\mathscr{F}_{k+1,\lfloor nt\rfloor}) = \alpha^2\sum_{k=1}^{\lfloor nt\rfloor}\mathbb{E}(X_{k,\lfloor nt\rfloor}^2|\mathscr{F}_{k+1,\lfloor nt\rfloor}) + \beta^2\sum_{k=1}^{\lfloor nt\rfloor}\mathbb{E}(Y_{k,\lfloor nt\rfloor}^2|\mathscr{F}_{k+1,\lfloor nt\rfloor})\\ + 2\alpha\beta\sum_{k=1}^{\lfloor nt\rfloor}\mathbb{E}(X_{k,\lfloor nt\rfloor}Y_{k,\lfloor nt\rfloor}|\mathscr{F}_{k+1,\lfloor nt\rfloor})\,,\end{aligned}$$

and prove that it coincides with $(\alpha^2+\beta^2)v_t^2(\sigma)$, where $v_t^2(\sigma)$ is the limiting variance introduced in Theorem 2.1. The following result shows that, in fact, it suffices to compute the limit of the first term in the r.h.s. above.

Proposition 2.2. *Almost surely, it holds*

$$\begin{aligned}&\mathbb{E}(X_{k,\lfloor nt\rfloor}^2|\mathscr{F}_{k+1,\lfloor nt\rfloor}) = \mathbb{E}(Y_{k,\lfloor nt\rfloor}^2|\mathscr{F}_{k+1,\lfloor nt\rfloor})\\ &\mathbb{E}(X_{k,\lfloor nt\rfloor}Y_{k,\lfloor nt\rfloor}|\mathscr{F}_{k+1,\lfloor nt\rfloor}) = 0\end{aligned}$$

for all $k\le\lfloor nt\rfloor$.

Proof. All equalities in this proof are intended to hold almost surely. Introduce the holomorphic function $f(z):=\frac{1}{\sqrt{c}}\big(\log\frac{F(z)}{z}-c\big)$ defined for $|z|>1$, so that

$$\Big[f(e^{-i\Theta_k}Z_{k,\lfloor nt\rfloor})\Big]^2 = X_{k,\lfloor nt\rfloor}^2 - Y_{k,\lfloor nt\rfloor}^2 + 2iX_{k,\lfloor nt\rfloor}Y_{k,\lfloor nt\rfloor}\,.$$

Taking conditional expectations both sides, we find

$$\begin{aligned}\frac{1}{2\pi}\int_{-\pi}^{\pi}\Big[f(e^{-i\vartheta}Z_{k,\lfloor nt\rfloor})\Big]^2\mathrm{d}\vartheta =&\mathbb{E}(X_{k,\lfloor nt\rfloor}^2|\mathscr{F}_{k+1,\lfloor nt\rfloor}) - \mathbb{E}(Y_{k,\lfloor nt\rfloor}^2|\mathscr{F}_{k+1,\lfloor nt\rfloor})\\ &+ 2i\mathbb{E}(X_{k,\lfloor nt\rfloor}Y_{k,\lfloor nt\rfloor}|\mathscr{F}_{k+1,\lfloor nt\rfloor}).\end{aligned}\tag{2.13}$$

Set $z=e^{i\vartheta}$, and observe that the function $z\mapsto\Big[f\Big(\frac{Z_{k,\lfloor nt\rfloor}}{z}\Big)\Big]^2$ is holomorphic on $\mathbb{D}=\{|z|<1\}$, and extends continuously to $\overline{\mathbb{D}}$. Applying Cauchy's integral formula to this function, then, yields

$$\frac{1}{2\pi}\int_{-\pi}^{\pi}\Big[f(e^{-i\vartheta}Z_{k,\lfloor nt\rfloor})\Big]^2\mathrm{d}\vartheta = \frac{1}{2\pi i}\int_{|z|=1}\left[f\left(\frac{Z_{k,\lfloor nt\rfloor}}{z}\right)\right]^2\frac{\mathrm{d}z}{z} = \lim_{|z|\to 0}\left[f\left(\frac{Z_{k,\lfloor nt\rfloor}}{z}\right)\right]^2$$
$$= \lim_{|z|\to\infty} f^2(z) = 0.$$

Going back to (2.13), this implies that both real and imaginary part of the r.h.s. must vanish almost surely, which is what we wanted to show. □

Proposition 2.2 already shows that the limiting Gaussian vector $\mathcal{F}^\sigma(t, e^{ia})$ must have i.i.d. entries. It remains to compute the limiting variance, that is to show that for all $\eta > 0$ it holds

$$\mathbb{P}\left(\left|\sum_{k=1}^{\lfloor nt\rfloor}\mathbb{E}(X^2_{k,\lfloor nt\rfloor}|\mathscr{F}_{k+1,\lfloor nt\rfloor}) - v_t^2(\sigma)\right| > \eta\right) \to 0 \tag{2.14}$$

as $n\to\infty$. To this end it is clearly enough to work on $E(m,\varepsilon)$, the advantage being that on this event we have

$$\left|Z^\sigma_{k,\lfloor nt\rfloor}(a) - e^{ia+\sigma+(\lfloor nt\rfloor - k)c}\right| \le C(t)\varepsilon$$

for all $k \le \lfloor nt\rfloor$, as it follows directly from the definition of $E(m,\varepsilon)$ as long as $\sigma \gg \varepsilon$. Our strategy is then to replace each $Z_{k,\lfloor nt\rfloor}$ by its deterministic approximation, and show that, provided σ is large enough with respect to c, this does not affect the limiting variance.

Recall that the Poisson kernel for the unit disc $\mathbb{D}$ is given by $P_r(\vartheta) = \mathrm{Re}\Big(\dfrac{1+re^{i\vartheta}}{1-re^{i\vartheta}}\Big)$ for $r<1$, and that the function $re^{i\vartheta}\mapsto P_r(\vartheta)$ is harmonic in $\mathbb{D}$. Moreover, given any continuous function f on $\mathbb{T}=\partial\mathbb{D}$, its harmonic extension Hf inside $\mathbb{D}$ is given by Poisson's integral formula

$$(Hf)(re^{i\vartheta}) = \frac{1}{2\pi}\int_{-\pi}^{\pi}P_r(\vartheta - t)f(e^{it})\mathrm{d}t = (P_r * f)(\vartheta)\,. \tag{2.15}$$

We denote by $Q_r(\vartheta)$ the harmonic conjugate of $P_r(\vartheta)$ in $\mathbb{D}$ which is 0 at 0, i.e. $Q_r(\vartheta) = \operatorname{Im}\Big(\dfrac{1+re^{i\vartheta}}{1-re^{i\vartheta}}\Big)$.

Before proceeding with the proof of (2.14), let us give a short sketch of how the limiting variance is computed. Note that

$$\begin{aligned}\mathbb{E}(X_{k,n}^2|\mathscr{F}_{k+1,n}) &= \frac{1}{2\pi c}\int_{-\pi}^{\pi}\left[\operatorname{Re}\left(\log\frac{F(e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor})}{e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor}}\right)-c\right]^2 \mathrm{d}\vartheta \\ &= \frac{1}{2\pi c}\int_{-\pi}^{\pi}\left[\operatorname{Re}\left(\log\frac{F(e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor})}{e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor}}\right)\right]^2 \mathrm{d}\vartheta - c,\end{aligned}$$

where the second equality follows from (2.9). Now, using that on the event $E(m,c)$ we have[3]

$$Z^{\sigma}_{k,\lfloor nt\rfloor} \approx e^{\sigma+(\lfloor nt\rfloor-k)c},$$

and that by Corollary 2.2

$$\operatorname{Re}\left(\log\frac{F(e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor})}{e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor}}\right) \approx c\operatorname{Re}\left(\frac{e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor}+1}{e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor}-1}\right),$$

we obtain

$$\begin{aligned}\frac{1}{2\pi c}\int_{-\pi}^{\pi}\left[\operatorname{Re}\left(\log\frac{F(e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor})}{e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor}}\right)\right]^2 \mathrm{d}\vartheta &\approx \frac{c}{2\pi}\int_{-\pi}^{\pi}\left[\operatorname{Re}\left(\frac{e^{-i\vartheta+\sigma+(\lfloor nt\rfloor-k)c}+1}{e^{-i\vartheta+\sigma+(\lfloor nt\rfloor-k)c}-1}\right)\right]^2 \mathrm{d}\vartheta \\ &= \frac{c}{2\pi}\int_{-\pi}^{\pi}\Big[P_{e^{-\sigma-(\lfloor nt\rfloor-k)c}}(\vartheta)\Big]^2 \mathrm{d}\vartheta = cP_{e^{-2\sigma-2(\lfloor nt\rfloor-k)c}}(0),\end{aligned}$$

where the last integral has been computed using Poisson's integral formula (2.15). Finally, this yields

$$\begin{aligned}\sum_{k=1}^{\lfloor nt\rfloor}\mathbb{E}(X_{k,n}^2|\mathscr{F}_{k+1,n}) &\approx c\sum_{k=1}^{\lfloor nt\rfloor}P_{e^{-2\sigma-2(\lfloor nt\rfloor-k)c}}(0) - \lfloor nt\rfloor c \approx \int_{\sigma}^{\sigma+t}P_{e^{-2x}}(0)\mathrm{d}x - t \\ &= \log\frac{1-e^{-2(\sigma+t)}}{1-e^{-2\sigma}} = v_t^2(\sigma).\end{aligned}$$

[3] Write $a_n \approx b_n$ if the sequences $(a_n),(b_n)$ converge to the same limit as $n\to\infty$.

Let us now implement this strategy by showing that all the errors in the above approximations vanish in the limit.

Lemma 2.5. *Assume* $\lim_{\varepsilon\to 0}\frac{\varepsilon}{\sigma}=0$. *Then there exists a constant* $C(t)$, *depending only on* t, *such that on the event* $E(m,\varepsilon)$ *we have*

$$\left|\mathbb{E}(X_{k,\lfloor nt\rfloor}^2|\mathscr{F}_{k+1,\lfloor nt\rfloor})+c-\frac{c}{2\pi}\int_{-\pi}^{\pi}\left(P_{e^{-\sigma-(\lfloor nt\rfloor-k)c}}(\vartheta)\right)^2\mathrm{d}\vartheta\right|\le\frac{C(t)c\varepsilon}{(\sigma+(\lfloor nt\rfloor-k)c)^3}$$

for all $k\le\lfloor nt\rfloor$, *and* n *large enough.*

This result follows by a more general one, namely Lemma 3.2 stated in the next section, and the proof is therefore omitted. Assume now that $\sigma\gg\sqrt{\varepsilon}$. Then we deduce from Lemma 2.5 that on $E(m,\varepsilon)$ it holds:

$$\begin{aligned}\Bigg|\sum_{k=1}^{\lfloor nt\rfloor}\mathbb{E}(X_{k,\lfloor nt\rfloor}^2|\mathscr{F}_{k,\lfloor nt\rfloor})-&\sum_{k=1}^{\lfloor nt\rfloor}\frac{c}{2\pi}\int_{-\pi}^{\pi}\left(P_{e^{-\sigma-(\lfloor nt\rfloor-k)c}}(\vartheta)\right)^2\mathrm{d}\vartheta+\lfloor nt\rfloor c\Bigg|\\&\le\sum_{k=1}^{\lfloor nt\rfloor}\frac{C(t)c\varepsilon}{(\sigma+(\lfloor nt\rfloor-k)c)^3}\le C(t)\varepsilon\int_{\sigma}^{\sigma+\lfloor nt\rfloor c}\frac{\mathrm{d}x}{x^3}\\&=C(t)\varepsilon\left(\frac{1}{2\sigma^2}-\frac{1}{2(\sigma+\lfloor nt\rfloor c)^2}\right)\to 0\end{aligned}$$

as $n\to\infty$. Note that $\sqrt{\varepsilon}\asymp\delta^{1/3}$ (apart from logarithmic corrections), so the assumption $\sigma\gg\sqrt{\varepsilon}$ is stronger than the previous one $\sigma\gg\sqrt{\delta}$.

In conclusion, we have shown that, provided $\sigma\gg\sqrt{\varepsilon}$, the limiting variance is given by the deterministic expression

$$\begin{aligned}\lim_{n\to\infty}\left(\frac{c}{2\pi}\sum_{k=1}^{\lfloor nt\rfloor}\int_{-\pi}^{\pi}\left(P_{e^{-\sigma-(\lfloor nt\rfloor-k)c}}(\vartheta)\right)^2\mathrm{d}\vartheta-\lfloor nt\rfloor c\right)&=\int_{\sigma}^{\sigma+t}\frac{1}{2\pi}\int_{-\pi}^{\pi}\left(P_{e^{-x}}(\vartheta)\right)^2\mathrm{d}\vartheta\mathrm{d}x-t\\=\int_{\sigma}^{\sigma+t}(P_{e^{-x}}*P_{e^{-x}})(0)\mathrm{d}x-t=\int_{\sigma}^{\sigma+t}\frac{1+e^{-2x}}{1-e^{-2x}}\mathrm{d}x-t&=\log\frac{1-e^{-2(\sigma+t)}}{1-e^{-2\sigma}}=v_t^2(\sigma)\,.\end{aligned}$$

Finally, if $\sigma>0$ is kept fixed as $n\to\infty$ the assumption $\sigma\gg\sqrt{\varepsilon}$ is trivially satisfied, so this concludes the proof of Theorem 2.1.

Chapter 3

Global fluctuations for HL(0)

This chapter is concerned with the proof of the global fluctuations result, and the rigorous construction of the boundary process. We proceed in steps. In Section 3.1 we prove Theorem 3.1, describing the fluctuations of all points on the circle $\{|z| = e^{\sigma}\}$ simultaneously, for any $\sigma > 0$. This is then used in Section 3.2 to obtain a functional central limit theorem for the process $(t, z) \mapsto \log \Phi_{\lfloor nt \rfloor}(z)$, viewed as a stochastic process taking values in the space of holomorphic functions on $\{|z| > 1\}$ (cf. Theorem 3.3). Finally, in Section 3.3 we present an explicit construction of the boundary values of the limiting fluctuation process. These are shown to be given by a distribution–valued continuous process, which is rigorously defined as a Ornstein–Uhlenbeck dynamics in a suitable infinite–dimensional Hilbert space. We conclude the chapter by collecting some open questions in Section 3.4.

3.1 The fluctuation process on $C(\mathbb{T})$

Having a pointwise convergence result, it is natural to ask if this can be extended to obtain convergence of random fields. Recall that $\mathbb{T} = \{|z| = 1\}$ denotes the unit circle, and let $C(\mathbb{T})$ denote the space of continuous functions from $\mathbb{T}$ to $\mathbb{C}$, equipped with the supremum norm

$$\|x\|_{\infty} = \sup_{\vartheta \in [-\pi, \pi)} |x(e^{i\vartheta})| \,. \tag{3.1}$$

Moreover, let $D[0,\infty)$ denote the space of càdlàg functions $x:[0,\infty)\to C(\mathbb{T})$. The goal of this section is to prove the following result.

Theorem 3.1. *Fix any $\sigma>0$, and let $\mathcal{F}_n^\sigma$ denote the $C(\mathbb{T})$-valued càdlàg stochastic process defined by*

$$\mathcal{F}_n^\sigma(t,e^{ia})=\frac{1}{\sqrt{c}}\left(\log\frac{\Phi_{\lfloor nt\rfloor}(e^{ia+\sigma})}{e^{ia+\sigma}}-\lfloor nt\rfloor c\right)$$

for $e^{ia}\in\mathbb{T}$ and $t\geq 0$. Then there exists a continuous zero mean Gaussian process $\mathcal{F}^\sigma:[0,\infty)\to C(\mathbb{T})$ whose covariance structure is given by

$$\operatorname{Cov}\big(\mathcal{F}^\sigma(t,e^{ia})\big)=v_t^2(\sigma)\begin{pmatrix}1 & 0\\ \\ 0 & 1\end{pmatrix},$$

$$\operatorname{Cov}\big(\mathcal{F}^\sigma(t,e^{ia}),\mathcal{F}^\sigma(s,e^{ib})\big)=\begin{pmatrix}c_{s,t}(\sigma,a-b) & \hat{c}_{s,t}(\sigma,a-b)\\ \\ -\hat{c}_{s,t}(\sigma,a-b) & c_{s,t}(\sigma,a-b)\end{pmatrix},$$

where $v_t^2(\sigma)=c_{t,t}(\sigma,0)$, and for $s<t$

$$c_{s,t}(\sigma,\alpha):=\operatorname{Re}\left(\log\frac{1-e^{-2\sigma-(t+s)+i\alpha}}{1-e^{-2\sigma-(t-s)+i\alpha}}\right),\qquad \hat{c}_{s,t}(\sigma,\alpha)=\operatorname{Im}\left(\log\frac{1-e^{-2\sigma-(t+s)+i\alpha}}{1-e^{-2\sigma-(t-s)+i\alpha}}\right),$$

such that $\mathcal{F}_n^\sigma\to\mathcal{F}^\sigma$ in distribution as $n\to\infty$, in the sense of weak convergence of probability measures on the space $D[0,\infty)$ equipped with the Skorokhod topology.

Note that the limiting process is rotationally invariant in the spatial coordinate, as one expects from the rotation invariance of the original model. The rest of this section is devoted to the proof of the above result.

Fix any $\sigma>0$. It is trivial to check that $\mathcal{F}_n^\sigma$ belongs to $D[0,\infty)$ for all $n\geq 0$. Since $D[0,\infty)$ equipped with the Skorokhod metric is a complete separable space, it follows by Prohorov's theorem that $\mathcal{F}_n^\sigma\to\mathcal{F}^\sigma$ weakly if and only if the finite dimensional distributions (FDDs) of $\mathcal{F}_n^\sigma$ converge to the ones of $\mathcal{F}^\sigma$, and $(\mathcal{F}_n^\sigma)_{n\geq 0}$ is tight (see [EK09], Lemma 4.3).

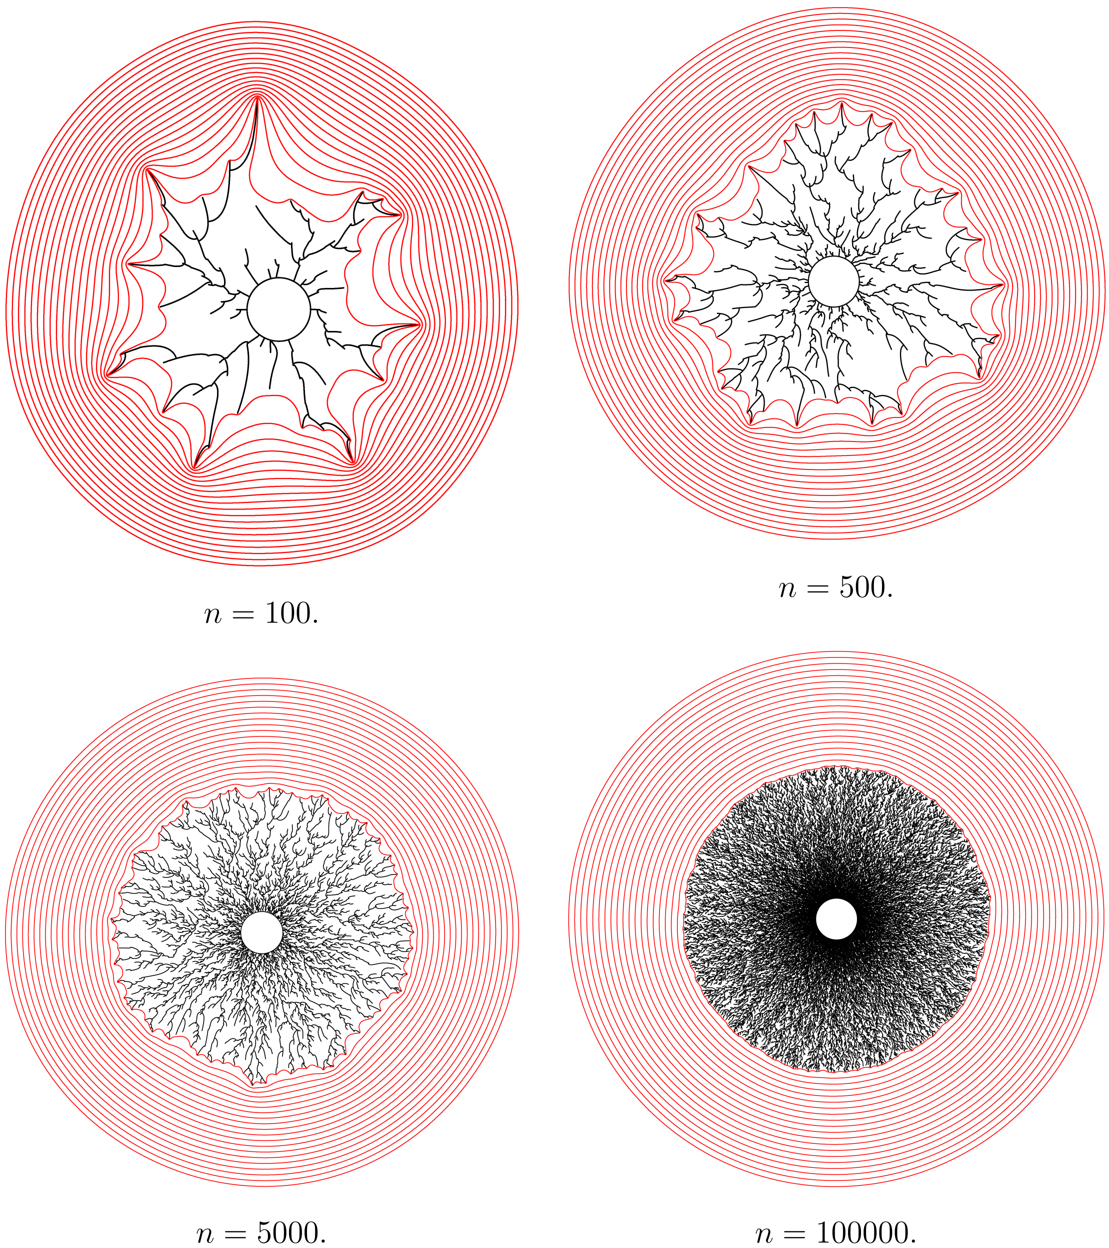

Figure 3.5: Image of 20 concentric circles of the form $e^{\sigma}\mathbb{T}$, $\sigma > 0$, under the HL(0) cluster growth (simulations by Henry Jackson).

3.1.1 Convergence of finite dimensional distributions

The following result is a direct consequence of the discussion in [Igl68, Whi70].

Lemma 3.1 ([Igl68, Whi70]). *Assume that:*

(i) for any $t > 0$ the family of probability measures of $(\mathcal{F}_n^\sigma(t,\cdot))_{n\geq 0}$ on $C(\mathbb{T})$ is tight, and

(ii) for any $0 \leq t_1 < \ldots < t_M$ and any $-\pi \leq a_1 < \ldots < a_M < \pi$, $M \in \mathbb{Z}_+$, it holds

$$\begin{pmatrix} \mathcal{F}_n^\sigma(t_1, e^{ia_1}) & \ldots & \mathcal{F}_n^\sigma(t_M, e^{ia_1}) \\ \vdots & & \vdots \\ \mathcal{F}_n^\sigma(t_1, e^{ia_M}) & \ldots & \mathcal{F}_n^\sigma(t_M, e^{ia_M}) \end{pmatrix} \longrightarrow \begin{pmatrix} \mathcal{F}^\sigma(t_1, e^{ia_1}) & \ldots & \mathcal{F}^\sigma(t_M, e^{ia_1}) \\ \vdots & & \vdots \\ \mathcal{F}^\sigma(t_1, e^{ia_M}) & \ldots & \mathcal{F}^\sigma(t_M, e^{ia_M}) \end{pmatrix}$$

in distribution as $n \to \infty$.

Then the FDDs of $\mathcal{F}_n^\sigma$ converge to the ones of $\mathcal{F}^\sigma$ as $n \to \infty$. In other words, for any $M \in \mathbb{N}_+$ and any $0 \leq t_1 < \ldots < t_M$, the random vector $(\mathcal{F}_n^\sigma(t_1,\cdot), \ldots, \mathcal{F}_n^\sigma(t_M,\cdot))$ in $C(\mathbb{T})^M$ converges in distribution to $(\mathcal{F}^\sigma(t_1,\cdot), \ldots, \mathcal{F}^\sigma(t_M,\cdot))$.

We start by showing that (ii) holds. To this end, let us again reduce to the scalar case by considering linear combinations. Recall the definition of the random variables $X_{k,n}^\sigma(\,\cdot\,)$, $Y_{k,n}^\sigma(\,\cdot\,)$ given in (2.10). For $(\alpha_{lj})_{1\leq l,j\leq M}$ in $\mathbb{R}^{M\times M}$, we look at the weak limit of

$$\sum_{l,j} \alpha_{lj} \mathcal{F}_n^\sigma(t_l, e^{ia_j}) = \sum_{k=1}^{\lfloor nt_M \rfloor} \sum_{l,j} \alpha_{lj} \Big[X_{k,\lfloor nt_l \rfloor}^\sigma(a_j) + i Y_{k,\lfloor nt_l \rfloor}^\sigma(a_j) \Big] \mathbf{1}(k \leq \lfloor nt_i \rfloor) = \sum_{k=1}^{\lfloor nt_M \rfloor} \mathscr{X}_{k,\lfloor nt_M \rfloor}^\sigma,$$

where we have set

$$\mathscr{X}_{k,\lfloor nt_M \rfloor}^\sigma = \sum_{l,j} \alpha_{lj} \Big[X_{k,\lfloor nt_l \rfloor}^\sigma(a_j) + i Y_{k,\lfloor nt_l \rfloor}^\sigma(a_j) \Big] \mathbf{1}(k \leq \lfloor nt_l \rfloor) \tag{3.2}$$

for $k \leq \lfloor nt_M \rfloor$. Then $\big(\mathscr{X}_{k,\lfloor nt_M \rfloor}^\sigma\big)_{k\leq \lfloor nt_M \rfloor}$ is a backwards martingale difference array with respect to the filtration $(\mathscr{F}_{k,\lfloor nt_M \rfloor})_{k\leq \lfloor nt_M \rfloor}$. We can therefore apply Theorem 2.2 to show weak convergence, provided that Assumptions (I) and (II) are satisfied. Note that all terms in the r.h.s. of (3.2) satisfy the estimates of Lemmas 2.1 and

2.2, from which one can easily show, reasoning as in Lemma 2.3, that Assumption (I) holds. Furthermore, Lemma 2.4 is still in force, from which we conclude that, provided $\sigma \gg \sqrt{\delta}$, the limiting variance is given by the limit in probability of $\sum_{k=1}^{\lfloor nt_M \rfloor} \mathbb{E}\big((\mathscr{X}^{\sigma}_{k,\lfloor nt_M \rfloor})^2 | \mathscr{F}_{k+1,\lfloor nt_M \rfloor}\big)$. We now focus on the computation of this limit. Expand the square and use linearity to see that the above sum equals

$$\sum_{l,j}\sum_{r,s} \alpha_{lj}\alpha_{rs} \sum_{k=1}^{\lfloor nt_l \rfloor \wedge \lfloor nt_r \rfloor} \mathbb{E}\Big[\Big(X^{\sigma}_{k,\lfloor nt_l \rfloor}(a_j) + iY^{\sigma}_{k,\lfloor nt_l \rfloor}(a_j)\Big)\Big(X^{\sigma}_{k,\lfloor nt_r \rfloor}(a_s) + iY^{\sigma}_{k,\lfloor nt_r \rfloor}(a_s)\Big)\Big| \mathscr{F}_{k,\lfloor nt_l \rfloor \vee \lfloor nt_r \rfloor}\Big].$$

It follows that it suffices to compute the limit in probability of

$$\sum_{k=1}^{\lfloor ns \rfloor} \mathbb{E}\Big[\Big(X^{\sigma}_{k,\lfloor nt \rfloor}(a) + iY^{\sigma}_{k,\lfloor nt \rfloor}(a)\Big)\Big(X^{\sigma}_{k,\lfloor ns \rfloor}(b) + iY^{\sigma}_{k,\lfloor ns \rfloor}(b)\Big)\Big| \mathscr{F}_{k,\lfloor nt \rfloor}\Big]$$

for arbitrary $a, b \in [-\pi, \pi)$ and $0 \le s \le t$. Moreover, by rotational invariance we can set $b = 0$ without loss of generality. The following result simplifies the computation.

Proposition 3.1. *Almost surely, it holds*

$$\mathbb{E}(X^{\sigma}_{k,\lfloor nt \rfloor}(a) X^{\sigma}_{k,\lfloor ns \rfloor}(0) | \mathscr{F}_{k+1,\lfloor nt \rfloor}) = \mathbb{E}(Y^{\sigma}_{k,\lfloor nt \rfloor}(a) Y^{\sigma}_{k,\lfloor ns \rfloor}(0) | \mathscr{F}_{k+1,\lfloor nt \rfloor})$$
$$\mathbb{E}(X^{\sigma}_{k,\lfloor nt \rfloor}(a) Y^{\sigma}_{k,\lfloor ns \rfloor}(0) | \mathscr{F}_{k+1,\lfloor nt \rfloor}) = -\mathbb{E}(X^{\sigma}_{k,\lfloor ns \rfloor}(0) Y^{\sigma}_{k,\lfloor nt \rfloor}(a) | \mathscr{F}_{k+1,\lfloor nt \rfloor})$$

for all $k \le \lfloor ns \rfloor$.

Proof. All equalities in this proof are intended to hold almost surely. Introduce the holomorphic function $f(z) := \frac{1}{\sqrt{c}}\big(\log \frac{F(z)}{z} - c\big)$ defined for $|z| > 1$, so that

$$\Big[f(e^{-i\Theta_k} Z^{\sigma}_{k,\lfloor nt \rfloor}(a)) + f(e^{-i\Theta_k} Z^{\sigma}_{k,\lfloor ns \rfloor}(0))\Big]^2 = \Big[X^{\sigma}_{k,\lfloor nt \rfloor}(a) + iY^{\sigma}_{k,\lfloor nt \rfloor}(a) + X^{\sigma}_{k,\lfloor ns \rfloor}(0) + iY^{\sigma}_{k,\lfloor ns \rfloor}(0)\Big]^2.$$

After taking conditional expectations with respect to $\mathcal{F}_{k+1,\lfloor nt \rfloor}$ both sides, and recalling Proposition 2.2, we find

$$\begin{aligned}\frac{1}{2\pi}\int_{-\pi}^{\pi}&\big[f(e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor}(a))+f(e^{-i\vartheta}Z^{\sigma}_{k,\lfloor ns\rfloor}(0))\big]^2\mathrm{d}\vartheta=\\&=2\Big[\mathbb{E}(X^{\sigma}_{k,\lfloor nt\rfloor}(a)X^{\sigma}_{k,\lfloor ns\rfloor}(0)|\mathscr{F}_{k+1,\lfloor nt\rfloor})-\mathbb{E}(Y^{\sigma}_{k,\lfloor nt\rfloor}(a)Y^{\sigma}_{k,\lfloor ns\rfloor}(0)|\mathscr{F}_{k+1,\lfloor nt\rfloor})\Big]\\&+2i\Big[\mathbb{E}(X^{\sigma}_{k,\lfloor nt\rfloor}(a)Y^{\sigma}_{k,\lfloor ns\rfloor}(0)|\mathscr{F}_{k+1,\lfloor nt\rfloor})+\mathbb{E}(X^{\sigma}_{k,\lfloor ns\rfloor}(0)Y^{\sigma}_{k,\lfloor nt\rfloor}(a)|\mathscr{F}_{k+1,\lfloor nt\rfloor})\Big].\end{aligned}\tag{3.3}$$

On the other hand, Cauchy's integral formula yields

$$\begin{aligned}\frac{1}{2\pi}\int_{-\pi}^{\pi}&\big[f(e^{-i\vartheta}Z^{\sigma}_{k,\lfloor nt\rfloor}(a))+f(e^{-i\vartheta}Z^{\sigma}_{k,\lfloor ns\rfloor}(0))\big]^2\mathrm{d}\vartheta=\\&=\lim_{|z|\to 0}\left[f\left(\frac{Z^{\sigma}_{k,\lfloor nt\rfloor}(a)}{z}\right)+f\left(\frac{Z^{\sigma}_{k,\lfloor ns\rfloor}(0)}{z}\right)\right]^2=0\,,\end{aligned}$$

and hence both real and imaginary part of the r.h.s. of (3.3) must vanish almost surely, which is what we wanted to show. □

It remains to compute the limit in probability of $\sum_k \mathbb{E}(X^{\sigma}_{k,\lfloor nt\rfloor}(a)X^{\sigma}_{k,\lfloor ns\rfloor}(0)|\mathscr{F}_{k+1,\lfloor nt\rfloor})$ and $\sum_k \mathbb{E}(X^{\sigma}_{k,\lfloor nt\rfloor}(a)Y^{\sigma}_{k,\lfloor ns\rfloor}(0)|\mathscr{F}_{k+1,\lfloor nt\rfloor})$. As in the previous section, we do this by approximating by a deterministic quantity.

Lemma 3.2. *Assume $\lim_{\varepsilon\to 0}\frac{\varepsilon}{\sigma}=0$. Then there exists a constant $C(t)$, depending only on t ($t>s$), such that on the event $E(m,\varepsilon)$ the following hold:*

$$\begin{aligned}\Big|\mathbb{E}(X^{\sigma}_{k,\lfloor nt\rfloor}(a)X^{\sigma}_{k,\lfloor ns\rfloor}(0)|\mathscr{F}_{k+1,\lfloor nt\rfloor})+c-\frac{c}{2\pi}\int_{-\pi}^{\pi}P_{e^{-\sigma-(\lfloor nt\rfloor-k)c}}(a-\vartheta)P_{e^{-\sigma-(\lfloor ns\rfloor-k)c}}(\vartheta)\mathrm{d}\vartheta\Big|\\\le\frac{C(t)c\varepsilon}{(\sigma+(\lfloor ns\rfloor-k)c)^3}\\\Big|\mathbb{E}(X^{\sigma}_{k,\lfloor nt\rfloor}(a)Y^{\sigma}_{k,\lfloor ns\rfloor}(0)|\mathscr{F}_{k+1,\lfloor nt\rfloor})+c-\frac{c}{2\pi}\int_{-\pi}^{\pi}P_{e^{-\sigma-(\lfloor nt\rfloor-k)c}}(a-\vartheta)Q_{e^{-\sigma-(\lfloor ns\rfloor-k)c}}(\vartheta)\mathrm{d}\vartheta\Big|\\\le\frac{C(t)c\varepsilon}{(\sigma+(\lfloor ns\rfloor-k)c)^3}\end{aligned}$$

for all $k\le\lfloor ns\rfloor$ and n large enough. Above $Q_r(\vartheta)=\mathrm{Im}\big(\frac{1+re^{i\vartheta}}{1-re^{i\vartheta}}\big)$, $r<1$, denotes the conjugate Poisson kernel.

We discuss the proof of Lemma 3.2, from which Lemma 2.5 also follows by setting $a=0$ and $s=t$, in Appendix A.2. As a consequence of the above result we obtain that, after further assuming $\sigma \gg \sqrt{\varepsilon}$, on the event $E(m,\varepsilon)$ it holds:

$$\Big|\sum_{k=1}^{\lfloor ns\rfloor}\mathbb{E}(X^\sigma_{k,\lfloor nt\rfloor}(a)X^\sigma_{k,\lfloor ns\rfloor}(0)|\mathscr{F}_{k+1,\lfloor nt\rfloor})+\lfloor ns\rfloor c-\frac{c}{2\pi}\sum_{k=1}^{\lfloor ns\rfloor}\int_{-\pi}^{\pi}P_{e^{-\sigma-(\lfloor nt\rfloor-k)c}}(a-\vartheta)P_{e^{-\sigma-(\lfloor ns\rfloor-k)c}}(\vartheta)\mathrm{d}\vartheta\Big|$$
$$\leq \sum_{k=1}^{\lfloor ns\rfloor}\frac{C(t)c\varepsilon}{(\sigma+(\lfloor ns\rfloor-k)c)^3}\leq C(t)\varepsilon\int_\sigma^{\sigma+\lfloor ns\rfloor c}\frac{\mathrm{d}x}{x^3}=C(t)\varepsilon\Big(\frac{1}{2\sigma^2}-\frac{1}{2(\sigma+\lfloor ns\rfloor c)^2}\Big)\to 0$$

as $n\to\infty$. This in turn implies that

$$\begin{aligned}\lim_{n\to\infty}\sum_{k=1}^{\lfloor ns\rfloor}\mathbb{E}(X^\sigma_{k,\lfloor nt\rfloor}(a)X^\sigma_{k,\lfloor ns\rfloor}(0)|\mathscr{F}_{k+1,\lfloor nt\rfloor}) &=\\ &=\lim_{n\to\infty}\Big(\frac{c}{2\pi}\sum_{k=1}^{\lfloor ns\rfloor}\int_{-\pi}^{\pi}P_{e^{-\sigma-(\lfloor nt\rfloor-k)c}}(a-\vartheta)P_{e^{-\sigma-(\lfloor ns\rfloor-k)c}}(\vartheta)\mathrm{d}\vartheta-\lfloor ns\rfloor c\Big)\\ &=\frac{1}{2\pi}\int_\sigma^{\sigma+s}\int_{-\pi}^{\pi}P_{e^{-(t-s)-x}}(a-\vartheta)P_{e^{-x}}(\vartheta)\mathrm{d}\vartheta-s=\int_{\sigma+\frac{t-s}{2}}^{\sigma+\frac{t+s}{2}}P_{e^{-2x}}(a)\mathrm{d}x-s\\ &=\log\left|\frac{1-e^{-2\sigma-(t+s)+ia}}{1-e^{-2\sigma-(t-s)+ia}}\right|,\end{aligned}$$

as claimed. Similarly,

$$\begin{aligned}\lim_{n\to\infty}\sum_{k=1}^{\lfloor ns\rfloor}\mathbb{E}(X^\sigma_{k,\lfloor nt\rfloor}(a)Y^\sigma_{k,\lfloor ns\rfloor}(0)|\mathscr{F}_{k+1,\lfloor nt\rfloor}) &=\frac{1}{2\pi}\int_\sigma^{\sigma+s}\int_{-\pi}^{\pi}P_{e^{-(t-s)-x}}(a-\vartheta)Q_{e^{-x}}(\vartheta)\mathrm{d}\vartheta\\ &=\int_{\sigma+\frac{t-s}{2}}^{\sigma+\frac{t+s}{2}}Q_{e^{-2x}}(a)\mathrm{d}x=\mathrm{Arg}\Big(\frac{1-e^{-2\sigma-(t+s)+ia}}{1-e^{-2\sigma-(t-s)+ia}}\Big).\end{aligned}$$

This concludes the proof of (ii) in Lemma 3.1. It remains to show that (i) holds.

Pick any $t>0$. Suppose we could show that:

(a) for all $\nu>0$, there exists $M_\nu>0$ and $n_0\in\mathbb{N}_+$ such that

$$\sup_{n\geq n_0}\mathbb{P}\big(\|\mathcal{F}^\sigma_n(0,\cdot)\|_\infty>M_\nu\big)\leq\nu,$$

(b) for all $\nu>0$ it holds

$$\lim_{\eta\to 0}\limsup_{n\to\infty}\mathbb{P}\Big(\sup_{|a-b|<\eta}|\mathcal{F}^\sigma_n(t,e^{ia})-\mathcal{F}^\sigma_n(t,e^{ib})|\geq\nu\Big)=0\,. \tag{3.4}$$

Then it would follow from Theorem 7.5 in [Bil99] that the sequence of probability measures on $C(\mathbb{T})$ associated to $(\mathcal{F}_n^\sigma(t,\cdot))_{n\geq 0}$ is tight.

Since $\mathcal{F}_n^\sigma(0,\cdot)\equiv 0$, (a) is trivially satisfied. The fact that also (b) holds follows by the same reasoning as in Section 3.1.2 below, and the proof is therefore omitted.

3.1.2 Tightness

To conclude the proof of Theorem 3.1 it remains to show that the family of probability measures associated to $(\mathcal{F}_n^\sigma)_{n\geq 0}$ on $D[0,\infty)$ is tight. Corollary 7.4 in [EK09] provides a characterization for tightness in this space. In particular, it tells us that if:

(a') for all $\nu>0$ and all $t\in[0,\infty)\cap\mathbb{Q}$, there exists a positive constant $M=M(\nu,t)$ such that

$$\limsup_{n\to\infty}\mathbb{P}(\|\mathcal{F}_n^\sigma(t,\cdot)\|_\infty > M) < \nu\,,$$

(b') for all $\nu, T>0$ there exists a positive constant $\eta=\eta(\nu,T)$ such that

$$\limsup_{n\to\infty}\mathbb{P}(\omega(\mathcal{F}_n^\sigma,\eta,T)\geq \nu)\leq \nu\,,$$

where ω denotes the modulus of continuity on $D[0,\infty)$, then the desired tightness follows. We start by showing (b'). It is clear that it suffices to restrict to the event $E(m,\varepsilon)$. We are set to show that for any $\nu,T>0$ it holds

$$\limsup_{n\to\infty}\mathbb{P}\left(\sup_{\substack{s,t\in[0,T],|t-s|<\eta\\ \alpha\in[-\pi,\pi)}}|\mathcal{F}_n^\sigma(t,e^{i\alpha})-\mathcal{F}_n^\sigma(s,e^{i\alpha})|\geq\nu\,;\,E(m,\varepsilon)\right)\to 0 \qquad \text{as } \eta\to 0\,. \tag{3.5}$$

It is convenient to switch to logarithmic coordinates. Following the notation introduced in [NT12], Section 5, we set $\tilde{D}_0=\{w\in\mathbb{C}:\operatorname{Re}(w)>0\}$, $\tilde{D}=\{w\in\mathbb{C}: e^w\in D\}$, and let $\tilde{F}$ be the unique conformal map from $\tilde{D}_0$ to $\tilde{D}$ such that $\tilde{F}(w)=w+c+o(1)$ as $\operatorname{Re}(w)\to\infty$. Moreover, let $\tilde{G}=\tilde{F}^{-1}$, so that $\tilde{G}(w)=w-c+o(1)$ as $\operatorname{Re}(w)\to\infty$. Finally, for all $k\leq n$ set $\tilde{F}_k(w)=\tilde{F}(w-i\Theta_k)+i\Theta_k$, $\tilde{G}_k=\tilde{F}_k^{-1}$ and

$\tilde{\Phi}_k = \tilde{F}_1 \circ \cdots \circ \tilde{F}_k$, $\tilde{\Gamma}_k = \tilde{\Phi}_k^{-1}$. Then $\mathcal{F}_n^\sigma(t,\vartheta) = \frac{1}{\sqrt{c}}\Big(\tilde{\Phi}_{\lfloor nt \rfloor}(i\vartheta + \sigma) - i\vartheta - \sigma - \lfloor nt \rfloor c\Big)$ and, assuming $t \geq s$ without loss of generality, we have:

$$\mathbb{P}\Bigg(\sup_{\substack{s,t\in[0,T],|t-s|<\eta \\ \alpha\in[-\pi,\pi)}} |\mathcal{F}_n^\sigma(t,e^{i\alpha}) - \mathcal{F}_n^\sigma(s,e^{i\alpha})| \geq \nu\,; E(m,\varepsilon)\Bigg) =$$
$$= \mathbb{P}\Bigg(\sup_{\substack{|t-s|<\eta \\ z:\mathrm{Re}(z)=\sigma}} \Big| \Big(\tilde{\Phi}_{\lfloor nt \rfloor}(z) - z - \lfloor nt \rfloor c\Big) - \Big(\tilde{\Phi}_{\lfloor ns \rfloor}(z) - z - \lfloor ns \rfloor c\Big) \Big| \geq \nu\sqrt{c}\,; E(m,\varepsilon)\Bigg). \tag{3.6}$$

Note that on $E(m,\varepsilon)$ we have $z = \tilde{\Gamma}_{\lfloor nt \rfloor}(w) = \tilde{\Gamma}_{\lfloor ns \rfloor}(w')$ for some $w \in \tilde{D}_{\lfloor nt \rfloor}$ and $w' \in \tilde{D}_{\lfloor ns \rfloor}$ such that $|\operatorname{Re}(w) - \sigma - \lfloor nt \rfloor c| < \varepsilon$, $|\operatorname{Re}(w') - \sigma - \lfloor ns \rfloor c| < \varepsilon$. Moreover, since $|t-s| < \eta$, it must be $|w - w'| < \lfloor nt \rfloor - \lfloor ns \rfloor + 4\varepsilon < 4\eta$ as long as n is large enough. For any $k \geq 0$, if $w \in \tilde{D}_k$ define $M_k(w) = \tilde{\Gamma}_k(w) - w + kc$. Finally, let

$$\mathcal{T} := \inf\big\{k \geq 0 : \xi \notin \tilde{D}_k \text{ for some } \xi \text{ such that } \operatorname{Re}(\xi) = \sigma + k - \varepsilon\big\},$$

and note that on $E(m,\varepsilon)$ we have $\mathcal{T} > \lfloor nt \rfloor$ for all $t \leq T$. It follows that the r.h.s. of (3.6) is bounded above by

$$\mathbb{P}\Bigg(\sup_{\substack{|t-s|<\eta \\ w:|\operatorname{Re}(w)-\sigma-\lfloor nt \rfloor c|<\varepsilon \\ w':|\operatorname{Re}(w')-\sigma-\lfloor ns \rfloor c|<\varepsilon \\ |w-w'|<4\eta}} \Big| \Big(\tilde{\Gamma}_{\lfloor nt \rfloor}(w) - w + \lfloor nt \rfloor c\Big) - \Big(\tilde{\Gamma}_{\lfloor ns \rfloor}(w') - w' + \lfloor ns \rfloor c\Big) \Big| \geq \nu\sqrt{c}\,; E(m,\varepsilon)\Bigg)$$
$$\leq \mathbb{P}\Bigg(\sup_{\substack{|t-s|<\eta \\ w:|\operatorname{Re}(w)-\sigma-\lfloor nt \rfloor c|<\varepsilon \\ w':|\operatorname{Re}(w')-\sigma-\lfloor ns \rfloor c|<\varepsilon \\ |w-w'|<4\eta}} \Big| M_{\lfloor nt \rfloor \wedge \mathcal{T}}(w) - M_{\lfloor ns \rfloor \wedge \mathcal{T}}(w') \Big| \geq \nu\sqrt{c}\Bigg).$$

We control the above probability by mean of a 3-dimensional version of Kolmogorov's continuity theorem (cf. [Dur96], Theorem 1.6). To this end, we show the following.

Lemma 3.3. *Fix any $s,t \in [0,T]$ with $s \leq t$, and any w, w' such that $|\operatorname{Re}(w) - \sigma - \lfloor nt \rfloor c| < \varepsilon$, $|\operatorname{Re}(w') - \sigma - \lfloor ns \rfloor c| < \varepsilon$ and $|w - w'| < 1$. Then there exists a constant*

$C = C(\sigma, T)$, depending only on σ and T, such that

$$\mathbb{E}\left(\left|M_{\lfloor nt\rfloor\wedge\mathcal{T}}(w) - M_{\lfloor ns\rfloor\wedge\mathcal{T}}(w')\right|^8\right) \leq C(\sigma,T)c^4|(t,w)-(s,w')|^4,$$

where $|(t,w)-(s,w')|$ denotes the Euclidean norm in $\mathbb{R}^3$.

Proof. Clearly the l.h.s. of the above inequality is upper bounded by

$$C\left[\mathbb{E}\left(\left|M_{\lfloor nt\rfloor\wedge\mathcal{T}}(w) - M_{\lfloor ns\rfloor\wedge\mathcal{T}}(w)\right|^8\right) + \mathbb{E}\left(\left|M_{\lfloor ns\rfloor\wedge\mathcal{T}}(w) - M_{\lfloor ns\rfloor\wedge\mathcal{T}}(w')\right|^8\right)\right] \tag{3.7}$$

for some absolute constant $C > 0$. We control the two terms separately.

Set $\tilde{G}_0(w) = \tilde{G}(w) - w$, so that $M_{k+1}(w) - M_k(w) = \tilde{G}_0(\tilde{\Gamma}_k(w) - i\Theta_{k+1}) + c$, and

$$\mathbb{E}(M_{k+1}(w) - M_k(w)|\sigma(\Theta_1,\dots,\Theta_k)) = \frac{1}{2\pi}\int_{-\pi}^{\pi}\tilde{G}_0(\tilde{\Gamma}_k(w) - i\vartheta)\mathrm{d}\vartheta + c = 0$$

for all $k \leq \lfloor nt\rfloor \wedge \mathcal{T}$. It follows that $\big(M_{k\wedge\mathcal{T}}(w)\big)_{k\leq\lfloor nt\rfloor}$ is a martingale. Moreover, we deduce from the estimates in (2.2), which now read

$$|\tilde{G}_0(w) + c| \leq \frac{Cc}{\mathrm{Re}(w) - \delta}\,, \qquad |\tilde{G}_0'(w)| \leq \frac{Cc}{(\mathrm{Re}(w)-\delta)^2}\,, \tag{3.8}$$

that $|M_{(k+1)\wedge\mathcal{T}}(w) - M_{k\wedge\mathcal{T}}(w)| = |\tilde{G}_0(\tilde{\Gamma}_k(w) - i\Theta_{k+1}) + c| \leq Cc/\sigma$ for all $k < \lfloor nt\rfloor$. Recall the following result, which appears in [DFJ68].

Theorem 3.2 ([DFJ68])**.** *Let $(M_n)_{n\geq 0}$ be a martingale with $M_0 = 0$. Then, for all $\nu \geq 2$ and $n \geq 1$, it holds*

$$\mathbb{E}(|M_n|^\nu) \leq C_\nu\, n^{\nu/2-1}\sum_{k=1}^{n}\mathbb{E}(|M_k - M_{k-1}|^\nu)\,,$$

where C_ν is an explicit constant, depending only on ν.

We use the above bound with $\nu = 8$, to get

$$
\begin{aligned}
\mathbb{E}\left(\left|M_{\lfloor nt\rfloor\wedge\mathcal{T}}(w) - M_{\lfloor ns\rfloor\wedge\mathcal{T}}(w)\right|^8\right) &= \mathbb{E}\left(\Big|\sum_{k=\lfloor ns\rfloor}^{\lfloor nt\rfloor-1}\big(M_{(k+1)\wedge\mathcal{T}}(w) - M_{k\wedge\mathcal{T}}(w)\big)\Big|^8\right) \\
&\le C(\lfloor nt\rfloor - \lfloor ns\rfloor)^3 \sum_{k=\lfloor ns\rfloor}^{\lfloor nt\rfloor-1}\mathbb{E}\left(\left|M_{(k+1)\wedge\mathcal{T}}(w) - M_{k\wedge\mathcal{T}}(w)\right|^8\right) \le \frac{Cc^4}{\sigma^8}|t-s|^4
\end{aligned}
\tag{3.9}
$$

for some absolute constant $C > 0$ and n large enough.

Let us now look at the second term in (3.7). For $k \le \lfloor ns\rfloor \wedge \mathcal{T}$ set $\tilde{M}_k = M_k(w) - M_k(w')$. Then $\big(\tilde{M}_{k\wedge\mathcal{T}}\big)_{k\le\lfloor nt\rfloor}$ is a martingale, and by (3.8)

$$
\left|\tilde{M}_{(k+1)\wedge\mathcal{T}} - \tilde{M}_{k\wedge\mathcal{T}}\right| = \left|\tilde{G}_0(\tilde{\Gamma}_k(w-i\Theta_{k+1}) - \tilde{G}_0(\tilde{\Gamma}_k(w'-i\Theta_{k+1})\right| \le \frac{Cc}{\sigma^2}(|w-w'| + |\tilde{M}_{k\wedge\mathcal{T}}|)\,.
$$

This, together with the bound in Theorem 3.2, yields

$$
\begin{aligned}
\mathbb{E}\left(\left|\tilde{M}_{\lfloor ns\rfloor\wedge\mathcal{T}}\right|^8\right) &= \mathbb{E}\left(\Big|\sum_{k=0}^{\lfloor ns\rfloor-1}\big(\tilde{M}_{(k+1)\wedge\mathcal{T}} - \tilde{M}_{k\wedge\mathcal{T}}\big)\Big|^8\right) \\
&\le \lfloor ns\rfloor^3 \sum_{k=0}^{\lfloor ns\rfloor-1}\mathbb{E}\left(\left|\tilde{M}_{(k+1)\wedge\mathcal{T}} - \tilde{M}_{k\wedge\mathcal{T}}\right|^8\right) \\
&\le \lfloor ns\rfloor^3 \frac{Cc^8}{\sigma^{16}} \sum_{k=0}^{\lfloor ns\rfloor-1}\Big(|w-w'|^8 + \mathbb{E}\big(|\tilde{M}_{k\wedge\mathcal{T}}|^8\big)\Big)\,.
\end{aligned}
$$

It then follows from Grönwall's inequality that

$$
\mathbb{E}\left(\left|\tilde{M}_{\lfloor ns\rfloor\wedge\mathcal{T}}\right|^8\right) \le \lfloor ns\rfloor^4 \frac{Cc^8}{\sigma^{16}}|w-w'|^8 e^{\frac{C\lfloor ns\rfloor^4 c^8}{\sigma^{16}}} \le C(\sigma,T)c^4|w-w'|^8 \tag{3.10}
$$

for n large enough and a constant $C(\sigma, T)$ depending only on σ and T. Putting (3.9) and (3.10) together, and using that $|w-w'|^8 \le |w-w'|^4$ since $|w-w'| < 1$

by assumption, we find

$$\mathbb{E}\left(\left|M_{\lfloor nt\rfloor\wedge\mathcal{T}}(w)-M_{\lfloor ns\rfloor\wedge\mathcal{T}}(w)\right|^8\right)+\mathbb{E}\left(\left|\tilde{M}_{\lfloor ns\rfloor\wedge\mathcal{T}}\right|^8\right)\leq C'(\sigma,T)c^4|(t,w)-(s,w')|^4\,,$$

for n large enough and $C'(\sigma,T)=\max\{C/\sigma^8;C(\sigma,T)\}$, as claimed. □

Kolmogorov's continuity theorem now yields the existence of a random variable $\mathscr{M}>0$ such that

$$\sup_{\substack{|t-s|<\eta\\ w:|\operatorname{Re}(w)-\sigma-\lfloor nt\rfloor c|<\varepsilon\\ w':|\operatorname{Re}(w')-\sigma-\lfloor ns\rfloor c|<\varepsilon\\ |w-w'|<4\eta}}\left|M_{\lfloor nt\rfloor\wedge\mathcal{T}}(w)-M_{\lfloor ns\rfloor\wedge\mathcal{T}}(w')\right|\leq\mathscr{M}|(t,w)-(s,w')|^{1/16}\,,$$

with $\mathbb{E}(\mathscr{M}^8)\leq C(\sigma,T)c^4$, for $C(\sigma,T)$ as in the statement of Lemma 3.3. Therefore we find

$$\begin{aligned}\mathbb{P}\Bigg(\sup_{\substack{|t-s|<\eta\\ w:|\operatorname{Re}(w)-\sigma-\lfloor nt\rfloor c|<\varepsilon\\ w':|\operatorname{Re}(w')-\sigma-\lfloor ns\rfloor c|<\varepsilon\\ |w-w'|<4\eta}}\left|M_{\lfloor nt\rfloor\wedge\mathcal{T}}(w)-M_{\lfloor ns\rfloor\wedge\mathcal{T}}(w')\right|\geq\nu\sqrt{c}\Bigg)&\leq\mathbb{P}\Big(\mathscr{M}^8\geq\frac{\nu^8c^4}{4\sqrt{\eta}}\Big)\\ &\leq 4C(\sigma,T)\frac{\sqrt{\eta}}{\nu^8}\,,\end{aligned}$$

and sending first $n\to\infty$ and then $\eta\to 0$ we conclude that (3.5) holds. This proves tightness, and hence it concludes the proof of Theorem 3.1.

3.2 The fluctuation process on $\mathscr{H}$

Notation. Recall that $\mathbb{T}=\{z\in\mathbb{C}:|z|=1\}$, and set $\alpha\mathbb{T}=\{z\in\mathbb{C}:|z|=\alpha\}$ for any $\alpha\in\mathbb{R}_+$. Let $\big(C(\alpha\mathbb{T}),\|\cdot\|_\infty\big)$ denote the space of continuous function on $\alpha\mathbb{T}$ equipped with the supremum norm. Moreover, for a subset D of the complex plane, denote by $\mathscr{H}(D)$ the space of holomorphic functions on D. Whenever $D=\{|z|>1\}$, denote $\mathscr{H}(D)$ simply by $\mathscr{H}$. Finally, let $C_{\mathscr{H}}[0,\infty)$ denote the space of continuous functions from $[0,\infty)$ to $\mathscr{H}$.

We have shown in the previous section that, for any fixed $\sigma > 0$, $\mathcal{F}_n^\sigma \to \mathcal{F}^\sigma$ as $n \to \infty$ in distribution with respect to the Skorokhod topology on the space $D[0,\infty)$ of càdlàg functions from $[0,\infty)$ to $C(\mathbb{T})$. Note that, for any $t \geq 0$, the continuous function $\mathcal{F}_n^\sigma(t,\cdot)$ coincides with the restriction of the holomorphic function

$$\mathcal{F}_n(t,z) = \frac{1}{\sqrt{c}}\Big(\log\frac{\Phi_{\lfloor nt\rfloor}(z)}{z} - \lfloor nt\rfloor c\Big)\,, \tag{3.11}$$

defined for all $|z| > 1$, to the circle $e^\sigma\mathbb{T}$. We show below that also the limit $\mathcal{F}^\sigma(t,\cdot)$ can be interpreted as the restriction of a holomorphic random function $\mathcal{F}(t,\cdot)$, defined for all $|z| > 1$, to $e^\sigma\mathbb{T}$. Moreover, we provide an explicit construction of $\mathcal{F}$, and prove that $\mathcal{F}_n \to \mathcal{F}$ in distribution as $n \to \infty$, with respect to the Skorokhod topology on the space of càdlàg functions from $[0,\infty)$ to $\mathscr{H}$.

Define, for all $N \geq 1$, a distance d_N on $\mathscr{H}$ by setting

$$d_N(\phi,\psi) = \sup_{|z|\geq e^{1/N}} |\phi(z)-\psi(z)| \wedge 1\,, \qquad \text{and let} \qquad d(\phi,\psi) = \sum_{N\geq 1}\frac{d_N(\phi,\psi)}{2^N}\,. \tag{3.12}$$

Since $(\mathscr{H}, d_N)$ is a complete separable metric space for all $N \geq 1$, this makes $(\mathscr{H}, d)$ into a complete separable metric space (i.e. Polish space). It follows that, if $D_{\mathscr{H}}[0,\infty)$ denotes the space of càdlàg functions from $[0,\infty)$ to $\mathscr{H}$ equipped with the Skorokhod metric $\mathbb{d}_{\mathscr{H}}$, then also $(D_{\mathscr{H}}[0,\infty), \mathbb{d}_{\mathscr{H}})$ is a complete separable metric space. It is clear that each $\mathcal{F}_n$ defined in (3.11) is a random variable in $D_{\mathscr{H}}[0,\infty)$.

Let us now describe the explicit construction of $\mathcal{F}$. Let $D = \mathbb{R}/2\pi\mathbb{Z}$, and recall that if we set $e_k(\vartheta) = e^{ik\vartheta}/\sqrt{2\pi}$ for $k \in \mathbb{Z}$, then $(e_k)_{k\in\mathbb{Z}}$ forms an orthonormal basis (in short ONB) for $L^2(D)$ with respect to the inner product $(f,g) = \int_{-\pi}^{\pi}\overline{f(\vartheta)}g(\vartheta)\mathrm{d}\vartheta$. On this basis, that we refer to as Fourier basis, the Poisson kernel $\mathrm{Re}\left(\frac{1+z}{1-z}\right)$ reads

$$P_{1/r}(\vartheta) = \mathrm{Re}\left(\frac{1+e^{i\vartheta}/r}{1-e^{i\vartheta}/r}\right) = \sqrt{2\pi}\Big(e_0 + \sum_{k\in\mathbb{Z}\setminus\{0\}} r^{-|k|}e_k(\vartheta)\Big)$$

for any $r > 1$. Take two independent collections $(\beta_k)_{k\in\mathbb{Z}}$ and $(\beta'_k)_{k\in\mathbb{Z}}$ of i.i.d. Brownian Motions on $\mathbb{R}$, and denote by $(A_k)_{k\in\mathbb{Z}}$ and $(B_k)_{k\in\mathbb{Z}}$ the solutions to

$$\begin{cases} \mathrm{d}A_k(t) = -|k|A_k(t)\mathrm{d}t + \sqrt{2}\,\mathrm{d}\beta_k(t)\,, \\ A_k(0) = 0 \end{cases} \qquad \begin{cases} \mathrm{d}B_k(t) = -|k|B_k(t)\mathrm{d}t + \sqrt{2}\,\mathrm{d}\beta'_k(t) \\ B_k(0) = 0\,. \end{cases}$$

Then A_k, B_k perform independent Ornstein-Uhlenbeck processes on $\mathbb{R}$ with invariant distribution $\mathcal{N}(0, 1/|k|)$. Define formally

$$\mathcal{W}(t,\vartheta) = \sum_{k\in\mathbb{Z}\setminus\{0\}} \Big(\frac{A_k(t) + iB_k(t)}{\sqrt{2}}\Big) e_k(\vartheta) \tag{3.13}$$

for $(t,\vartheta) \in [0,\infty) \times [-\pi,\pi)$. Finally, for $(t, re^{ia}) \in [0,\infty) \times \{|z| > 1\}$ set

$$\mathcal{F}(t, re^{ia}) = \frac{1}{\sqrt{2\pi}}\Big(P_{1/r} * \mathcal{W}(t,\cdot)\Big)(a) = \frac{1}{\sqrt{2\pi}}\int_{-\pi}^{\pi} \mathrm{Re}\,\Big(\frac{1+e^{i(a-\vartheta)}/r}{1-e^{i(a-\vartheta)}/r}\Big)\,\mathcal{W}(t,\vartheta)\mathrm{d}\vartheta\,.$$

We remark that, although $\mathcal{W}$ is only defined formally for now, its holomorphic extension $\mathcal{F}$ is a well defined function. This can be seen, for example, by expanding $\mathcal{F}$ in Fourier basis, and checking that $\mathrm{Var}\mathcal{F}(t,z) < \infty$ for all $(t,z) \in [0,\infty) \times \{|z| > 1\}$.

Theorem 3.3. *$\mathcal{F}$ is a random variable in $C_{\mathscr{H}}[0,\infty) \subset D_{\mathscr{H}}[0,\infty)$. Moreover, $\mathcal{F}_n \to \mathcal{F}$ as $n \to \infty$ in distribution with respect to the Skorokhod metric $\mathrm{d}_{\mathscr{H}}$ on $D_{\mathscr{H}}[0,\infty)$.*

The question of how to make sense of the boundary values $\mathcal{W}$ defined formally in (3.13) is addressed in the next section, where we show that $\mathcal{W}$ can be rigorously defined as an Ornstein-Uhlenbeck process on a suitable infinite-dimensional Hilbert space.

The rest of this section is devoted to the proof of Theorem 3.3.

Lemma 3.4. *$\mathcal{F}$ is a random variable in $C_{\mathscr{H}}[0,\infty)$. Moreover, $\mathcal{F}$ is Gaussian (in the sense that all finite–dimensional marginals of $\mathcal{F}(t,z)$ are complex Gaussian random vectors), and its restriction to $[0,\infty) \times e^{\sigma}\mathbb{T}$ agrees in distribution with $\mathcal{F}^{\sigma}$ defined in Theorem 3.1, for any $\sigma > 0$.*

Proof. The fact that $\mathcal{F}$ is Gaussian is true by construction. Fix any $t \geq 0$, and expand $\mathcal{F}(t,\cdot)$ in Fourier basis, to get

$$\begin{aligned}\mathcal{F}(t, re^{ia}) &= \Big(e_0 + \sum_{k\neq 0} r^{-|k|} e^{-ika} e_k, \sum_{k\neq 0}\Big(\frac{A_k(t)+iB_k(t)}{\sqrt{2}}\Big)e_k\Big) = \sum_{k\neq 0} r^{-|k|}\Big(\frac{A_k(t)+iB_k(t)}{\sqrt{2}}\Big)e^{ika} \\ &\overset{(d)}{=} \sum_{k\geq 1} r^{-k}\big[A_k(t)\cos ka - B_k(t)\sin ka\big] + i\sum_{k\geq 1} r^{-k}\big[B_k(t)\cos ka + A_k(t)\sin ka\big]\end{aligned} \tag{3.14}$$

for any $r > 1$ and $a \in [-\pi, \pi)$, where the last equality holds in law as stochastic processes, and it follows from the independence of the OU processes. This provides an almost surely convergent power series expansion for $\mathcal{F}(t,\cdot)$ at all z with $|z| > 1$, and hence it shows that $\mathcal{F}$ is a Gaussian stochastic processes taking values in $\mathscr{H}$. Recall from (3.12) the definition of the distances $(d_N), d$ on $\mathscr{H}$. To see that $\mathcal{F}$ is continuous, we have to show that for all compacts of the form $[0,T]$ it holds almost surely that, if $t_n \to t \in [0,T]$ as $n \to \infty$, then $d\big(\mathcal{F}(t_n,\cdot), \mathcal{F}(t,\cdot)\big) \to 0$, i.e. $d_N\big(\mathcal{F}(t_n,\cdot), \mathcal{F}(t,\cdot)\big) \to 0$ for all $N \geq 1$. We have:

$$\begin{aligned} d_N\big(\mathcal{F}(t_n,\cdot), \mathcal{F}(t,\cdot)\big) &= \sup_{|z|\geq e^{1/N}} \big|\mathcal{F}(t_n,z) - \mathcal{F}(t,z)\big| = \sup_{|z| = e^{1/N}} \big|\mathcal{F}(t_n,z) - \mathcal{F}(t,z)\big| \\ &= \sup_{a\in[-\pi,\pi)} \Big|\sum_{k\geq 1} e^{-k/N}\big[(A_k(t_n)+iB_k(t_n)) - (A_k(t)+iB_k(t))\big]e^{ika}\Big| \\ &\leq \sum_{k\geq 1} e^{-k/N}\Big(|A_k(t_n) - A_k(t)| + |B_k(t_n) - B_k(t)|\Big). \end{aligned}$$

To show that the last term converges to 0 as $n \to \infty$ almost surely for all $t_n, t \in [0,T]$, then, it suffices to prove that the sequence of continuous functions $g_M(t) = \sum_{k=1}^{M} e^{-k/N}|A_k(t)|$ converges uniformly on $[0,T]$ as $M \to \infty$. Indeed, we find

$$\sup_{t\in[0,T]} |g_M(t) - g_\infty(t)| = \sup_{t\in[0,T]} \sum_{k=M+1}^{\infty} e^{-k/N}|A_k(t)| \leq \sum_{k=M+1}^{\infty} e^{-k/N} \sup_{t\in[0,T]} |A_k(t)|\,.$$

On the other hand Doob's maximal inequality for the submartingale $(e^{-2kt}A_k^2(t))_{t\leq T}$

yields

$$\mathbb{P}\Big(\sup_{t\in[0,T]}|A_k(t)|\geq e^{k/2N}\Big)\leq\mathbb{P}\Big(\sup_{t\in[0,T]}\{e^{-2kt}A_k^2(t)\}\geq e^{k/N-2kT}\Big)\leq\frac{\mathbb{E}(A_k^2(T))}{e^{k/N}}\leq\frac{1}{ke^{k/N}}\,,$$

so by Borel-Cantelli $\sup_{t\in[0,T]}|A_k(t)|<e^{k/2N}$ for all but finitely many $k\geq 1$, almost surely. This proves uniform convergence on compacts, and hence almost sure continuity of $\mathcal{F}$.

To conclude the proof, we show that $\mathcal{F}$ has the same covariance structure of $\mathcal{F}^\sigma$ on the circle $e^\sigma\mathbb{T}$, for arbitrary $\sigma>0$. Indeed, it follows from (3.14) that real and imaginary parts of $\mathcal{F}(t,e^{\sigma+ia})$, $a\in[-\pi,\pi)$, are independent centred real Gaussian random variables, with

$$\mathbb{E}\big[(\mathrm{Re}\mathcal{F}(t,e^{\sigma+ia}))^2\big]=\mathbb{E}\big[(\mathrm{Im}\mathcal{F}(t,re^{\sigma+ia}))^2\big]=\sum_{k\geq1}\frac{e^{-2k\sigma}(1-e^{-2kt})}{k}=\log\frac{1-e^{-2(t+\sigma)}}{1-e^{-2\sigma}}\,.$$

Moreover, expanding $\mathcal{F}(s,e^\sigma)$ as in (3.14), one further checks that for $s<t$

$$\begin{aligned}\mathrm{Cov}\big(\mathrm{Re}\mathcal{F}(t,e^{\sigma+ia}),\mathrm{Re}\mathcal{F}(s,e^\sigma)\big)&=\mathrm{Cov}\big(\mathrm{Im}\mathcal{F}(t,e^{\sigma+ia}),\mathrm{Im}\mathcal{F}(s,e^\sigma)\big)\\&=\mathrm{Re}\bigg(\log\frac{1-e^{-(t+s)-2\sigma+ia}}{1-e^{-(t-s)-2\sigma+ia}}\bigg),\end{aligned}$$

and

$$\begin{aligned}\mathrm{Cov}\big(\mathrm{Re}\mathcal{F}(t,e^{\sigma+ia}),\mathrm{Im}\mathcal{F}(s,e^\sigma)\big)&=-\mathrm{Cov}\big(\mathrm{Im}\mathcal{F}(t,e^{\sigma+ia}),\mathrm{Re}\mathcal{F}(s,e^\sigma)\big)\\&=\mathrm{Im}\bigg(\log\frac{1-e^{-(t+s)-2\sigma+ia}}{1-e^{-(t-s)-2\sigma+ia}}\bigg).\end{aligned}$$

By rotational invariance in the spatial coordinate, this is enough to conclude that $\mathcal{F}$ and $\mathcal{F}^\sigma$ have the same covariance structure, and hence the same law, on every circle of the form $e^\sigma\mathbb{T}$, for arbitrary $\sigma>0$, as claimed. □

Proof of Theorem 3.3. For any $N\geq1$, denote by H_N the operator that maps a continuous function on $e^{1/N}\mathbb{T}$ to its holomorphic extension to the outer region $\{|z|\geq$

$e^{1/N}\}$. More precisely,

$$\begin{aligned} H_N : C\big(e^{1/N}\mathbb{T}\big) &\longrightarrow \mathscr{H}^0\big(\{|z| \geq e^{1/N}\}\big) \\ x &\mapsto \quad H_N(x) \end{aligned}$$

with $\mathscr{H}^0(\{|z| \geq e^{1/N}\}) := \big\{x \in \mathscr{H}(\{|z| \geq e^{1/N}\}) \text{ s.t. } \lim_{|z|\to\infty} x(z) = 0\big\}$, and

$$H_N(x)(re^{ia}) := \frac{1}{\sqrt{2\pi}} \int_{-\pi}^{\pi} \mathrm{Re}\left(\frac{1 + e^{1/N+i(a-\vartheta)}/r}{1 - e^{1/N+i(a-\vartheta)}/r}\right) x(e^{1/N+i\vartheta})\mathrm{d}\vartheta.$$

Moreover, let $D_N[0,\infty)$ and $D_{\mathscr{H}_N}[0,\infty)$ denote respectively the space of càdlàg functions from $[0,\infty)$ to $\big(C(e^{1/N}\mathbb{T}), \|\cdot\|_\infty\big)$, and the space of càdlàg functions from $[0,\infty)$ to $\big(\mathscr{H}(\{|z| \geq e^{1/N}\}), d_N\big)$, both equipped with the Skorokhod topology. Finally, let $f_N : D_N[0,\infty) \to D_{\mathscr{H}_N}[0,\infty)$ be defined for $x \in D_N[0,\infty)$ by

$$\begin{aligned} f_N(x) :\ [0,\infty) &\to \mathscr{H}(\{|z| \geq e^{1/N}\}) \\ t &\mapsto \quad H_N(x(t,\cdot))\,. \end{aligned}$$

We claim that the proof of Theorem 3.3 amounts to showing that the map f_N is continuous. Indeed, assume so. Then it follows from Theorem 3.1, together with the continuous mapping theorem (cf. [Bil99], Theorem 2.7), that $f_N(\mathcal{F}_n^{1/N}) \to f_N(\mathcal{F}^{1/N})$ in distribution with respect to the Skorokhod metric on $D_{\mathscr{H}_N}[0,\infty)$. Moreover, recall from (3.14) the definition of $\mathcal{F}$. Then, since by Lemma 3.4 the stochastic process $\mathcal{F}^{1/N}$ agrees in law with $\mathcal{F}$ on $[0,\infty) \times e^{1/N}\mathbb{T}$, we deduce that the corresponding holomorphic extensions $f_N(\mathcal{F}^{1/N})$ and $\mathcal{F}$ agree in law on $[0,\infty) \times \{|z| \geq e^{1/N}\}$. It follows that $f_N(\mathcal{F}_n^{1/N}) \to \mathcal{F}$ in distribution with respect to the Skorokhod metric on $D_{\mathscr{H}_N}[0,\infty)$. Since $f_N(\mathcal{F}_n^{1/N}) \equiv \mathcal{F}_n$ on $[0,\infty) \times \{|z| \geq e^{1/N}\}$, with $\mathcal{F}_n$ defined as in (3.11), this shows that $\mathcal{F}_n \to \mathcal{F}$ in distribution with respect to the Skorokhod metric on $D_{\mathscr{H}_N}[0,\infty)$. N being arbitrary, we conclude that $\mathcal{F}_n \to \mathcal{F}$ in distribution with respect to the Skorokhod metric $\mathbb{d}_{\mathscr{H}}$ on $D_{\mathscr{H}}[0,\infty)$, which is what we wanted to show.

It remains to prove that for all $N \geq 1$ the map f_N is continuous from $D_N[0,\infty)$

to $D_{\mathscr{H}_N}[0,\infty)$. Since f_N only acts on the spatial component, it suffices to show that the holomorphic extension map H_N is continuous from $\big(C(e^{1/N}\mathbb{T}), \|\cdot\|_\infty\big)$ to $\big(\mathscr{H}(\{|z| \geq e^{1/N}\}), d_N\big)$ for all $N \geq 1$. Indeed, take $(\phi_n)_n$, ϕ in $C(e^{1/N}\mathbb{T})$ and suppose that $\|\phi_n - \phi\|_\infty \to 0$ as $n \to \infty$. Then we have:

$$\begin{aligned} d_N(H_N(\phi_n), H_N(\phi)) &= \sup_{|z|\geq e^{1/N}} \big|(H_N\phi_n)(z) - (H_N\phi)(z)\big| \wedge 1 \\ &= \sup_{|z|= e^{1/N}} \big|(H_N\phi_n)(z) - (H_N\phi)(z)\big| \wedge 1 \\ &= \sup_{|z|= e^{1/N}} |\phi_n(z) - \phi(z)| \wedge 1 \leq \|\phi_n - \phi\|_\infty \to 0 \end{aligned}$$

as $n \to \infty$, where the second equality follows by applying the maximum principle, which is in force since $(H_N\phi_n)(z) \to 0$ as $|z| \to \infty$ by construction for all $N, n \geq 1$. This concludes the proof. □

3.3 The boundary process

We have proved weak convergence to a limiting Gaussian process $\mathcal{F}$ taking values in the space $\mathscr{H}$ of holomorphic functions on $\{|z| > 1\}$, of which we have provided an explicit construction. In this section we address the question of a rigorous definition of the boundary values $\mathcal{W}$ of $\mathcal{F}$, formally given by (3.13).

3.3.1 Abstract Wiener Space construction

It is clear that we have no hope to define $\mathcal{W}(t, \cdot)$ pointwise, since the formal series (3.13) diverges almost surely at each point. One could try, on the other hand, to make sense of it as a complex Gaussian process taking values in a suitable Hilbert space of functions on the unit circle.

For convenience of the reader, we review below the construction of Gaussian random variables, and then Gaussian stochastic processes, on infinite dimensional Hilbert spaces. This will then be applied to rigorously define the boundary process $\mathcal{W}$. Our presentation follows [DPZ14, She07].

Gaussian random variables on a Hilbert space

Definition 3.1. An Abstract Wiener Space is a triple (H, B, μ), where:

(i) $\big(H, (\cdot,\cdot)_H\big)$ is a Hilbert space,

(ii) $\big(B, \|\cdot\|_B\big)$ is the Banach space completion of H with respect to the measurable norm $\|\cdot\|_B$ on H, equipped with the Borel σ-algebra $\mathcal{B}$ induced by $\|\cdot\|_B$, and

(iii) μ is the unique Borel probability measure on $(B, \mathcal{B})$ such that, if B^* denotes the dual space of B, then $\mu \circ \phi^{-1} = \mathcal{N}(0, \|\tilde{\phi}\|_H^2)$ for all $\phi \in B^*$, where $\tilde{\phi}$ is the unique element of H such that $\phi(h) = (\tilde{\phi}, h)_H$ for all $h \in H$.

Note that (iii) also reads as follows: if X is a random variable on $(B, \mathcal{B})$ distributed according to μ, then $\phi(X) \sim \mathcal{N}(0, \|\tilde{\phi}\|_H)$ for all $\phi \in B^*$. In this case we say[1] that X is a standard Gaussian random variable on H. We refer the reader to [Gro67, She07] for the definition of measurable norm on a Hilbert space, and for existence and uniqueness of such a measure μ. Here we will only need the following two properties.

Facts. Let H be a Hilbert space with inner product $(\cdot,\cdot)_H$, and let $\|\cdot\|_H^2 = (\cdot,\cdot)_H$. Then:

(a) unless $\dim(H) < \infty$, any measurable norm on H is strictly weaker than $\|\cdot\|_H$, and

(b) if T is a Hilbert-Schmidt operator on H, i.e.

$$\sum_{i=1}^{\infty} \|Te_i\|_H^2 < \infty\,, \qquad (e_i)_{i=1}^{\infty} \text{ ONB of } \big(H, (\cdot,\cdot)_H\big)\,,$$

then $\|T\cdot\|_H$ is a measurable norm on H.

Note that if $\dim(H) = \infty$ then $\|\cdot\|_B \neq \|\cdot\|_H$ by (a). When $\|\cdot\|_B$ is itself induced by an inner product, which will turn out to be the case in our construction, B is itself a Hilbert space.

[1] Note that, unless $\dim H < \infty$, a standard Gaussian random variable X on H does not take values in H, but only in the larger Banach space B.

Brownian Motion on a Hilbert space

Definition 3.2. Let (H, B, μ) be an abstract Wiener space, and denote by μ_t the unique[2] probability measure on B, such that $\mu_t \circ \phi = \mathcal{N}(0, t\|\tilde{\phi}\|_H^2)$ for all $\phi \in B^*$. Finally, let $C_B[0, \infty)$ the space of continuous functions from $[0, \infty)$ to B, equipped with the σ-algebra generated by the coordinate functions $x \mapsto x(t)$. Then (cf. [DPZ14], pp.81-85) there exists a unique probability measure $\boldsymbol{\mu}$ on $C_B[0, \infty)$ such that, if W is a random variable in $C_B[0, \infty)$ distributed according to $\boldsymbol{\mu}$, then the following hold:

- $W(0) = 0$ $\boldsymbol{\mu}$-a.s.
- W has independent increments,
- for any $0 \le s < t$, $W(t) - W(s)$ is distributed according to μ_{t-s}.

If a random variable W on $C_B[0, \infty)$ is distributed according to $\boldsymbol{\mu}$ we call it (cylindrical) Brownian Motion on H.

Proposition 3.2 ([DPZ14], Proposition 4.3)**.** *Let (H, B, μ) be an Abstract Wiener Space, and let $(e_k)_k$ be an ONB of H with respect to $(\cdot, \cdot)_H$. If W is a Brownian Motion on $(H, (\cdot, \cdot)_H)$, then there exists a collection of i.i.d. real-valued Brownian Motions $(\beta_k)_k$ such that*

$$W(t) = \sum_k \beta_k(t) e_k \, ,$$

where the above series converges in $L^2(B)$ for all $t \ge 0$.

Note that by setting $t = 1$ in the above result we deduce the following.

Corollary 3.1. *Let (H, B, μ) be an Abstract Wiener Space. If X is a standard Gaussian random variable on H, i.e. $X \sim \mu$, then there exists a collection $(A_k)_k$ of i.i.d. $\mathcal{N}(0, 1)$ real random variables such that $X = \sum_k A_k e_k$.*

[2] Existence and uniqueness follow trivially by existence and uniqueness of μ.

OU process on a Hilbert space

Having constructed a Brownian Motion W on H, one could then go ahead and define stochastic integration with respect to it. For a detailed account on the theory of stochastic integrals with respect to a Brownian motion taking values in an infinite-dimensional Hilbert space we refer the reader to [DPZ14], Chapter 4. Here we are only interested in a very special case, namely the one of deterministic integrands which diagonalise on the ONB $(e_k)_k$ of H. Indeed, in this case the definition of stochastic integral with respect to W reduces to the one of a countable collection of stochastic integrals on $\mathbb{R}$.

Definition 3.3. Let (H, B, μ) be an Abstract Wiener Space, and assume that the norm $\|\cdot\|_B$ on B is induced by an inner product $(\cdot,\cdot)_B$, so that B is itself a Hilbert space. Let $(e_k)_k$ denote an ONB for H, and let $W(t) = \sum_k \beta_k(t) e_k$ be a Brownian Motion on H. If $\Psi(t) = \sum_k a_k(t) e_k$ for some collection of continuous real-valued functions $(a_k(t))_k$ such that $\sum_k \int_0^\infty |a_k(t)|^2 \mathrm{d}t < \infty$, then

$$\int_0^t \Psi(s)\mathrm{d}W(s) := \sum_k \int_0^t a_k(s)\mathrm{d}\beta_k(s) e_k\,, \qquad \forall t \geq 0\,.$$

3.3.2 The boundary process $\mathcal{W}$

Let us now see how the above definitions read in the case of our interest. Recall that $D = \mathbb{R}/2\pi\mathbb{Z}$, and for $f, g \in C^\infty(D)$ with Fourier expansion $f(\vartheta) = \sum_{k\in\mathbb{Z}} \hat{f}_k e_k(\vartheta)$, $g(\vartheta) = \sum_{k\in\mathbb{Z}} \hat{g}_k e_k(\vartheta)$ introduce the inner product

$$(f, g)_H := \sum_{k\neq 0} |k| \hat{f}_k\, \hat{g}_k\,.$$

Let $\sim$ be the equivalence relation on $C^\infty(D)$ which identifies two functions if they differ by a constant, and set $H_s(D) := C^\infty(D)/\sim$. We identify each equivalence class of $H_s(D)$ with its representative having zero average on D. Denote by H the Hilbert space completion of $H_s(D)$ with respect to $(\cdot,\cdot)_H$. We seek to define an abstract Wiener space (H, B, μ) for a suitable choice of measurable norm on H.

By property (b), in order to construct a measurable norm $\|\cdot\|_B$ on H it suffices to find a Hilbert-Schmidt operator T on H, and set $\|\cdot\|_B = \|T\cdot\|_H$. In fact, we are going to construct a one-parameter family of such operators, and hence of measurable norms on H.

For any $a \in \mathbb{R}$, let $(-\Delta)^a$ be the operator that acts on $L^2(D)$ functions by multiplying the Fourier coefficients by $|k|^{2a}$, that is

$$(-\Delta)^a\bigg(\sum_{k\in\mathbb{Z}} \hat{f}_k e_k\bigg)(\vartheta) := \sum_{k\neq 0} |k|^{2a}\hat{f}_k e_k(\vartheta).$$

Set

$$\mathcal{H}_a = \Big\{ f \in L^2(D) : (-\Delta)^a f \in L^2(D) \Big\} \big/ \sim ,$$

and equip $\mathcal{H}_a$ with the inner product $(f,g)_a := ((-\Delta)^a f, (-\Delta)^a g)$. Then $\big(\mathcal{H}_a, (\cdot,\cdot)_a\big)$ is a Hilbert space, and in fact it is the Hilbert space completion of $H_s(D)$ with respect to the inner product $(\cdot,\cdot)_a$. Denote by $\|\cdot\|_a$ the norm induced by $(\cdot,\cdot)_a$, and note that if $a > b$ then $\|\cdot\|_b \le \|\cdot\|_a$, so that $\mathcal{H}_a \subset \mathcal{H}_b$. It follows that, whenever $b < a$, $\|\cdot\|_b$ is well defined on $\mathcal{H}_a$, and moreover $\|f\|_b = \|(-\Delta)^{b-a} f\|_a$. In order for $\|\cdot\|_b$ to also be measurable on $\mathcal{H}_a$, then, it suffices to show that $(-\Delta)^{b-a}$ is a Hilbert-Schmidt operator on $\mathcal{H}_a$. To this end, note that if $f_k = (-\Delta)^{-a} e_k$ for $k \neq 0$, then $(f_k)_{k\neq 0}$ is an ONB for $\mathcal{H}_a$. Moreover,

$$\sum_{k\neq 0} \|(-\Delta)^{b-a} f_k\|_a^2 = \sum_{k\neq 0} \big((-\Delta)^b f_k, (-\Delta)^b f_k\big) = \sum_{k\neq 0} \big((-\Delta)^{b-a} e_k, (-\Delta)^{b-a} e_k\big) = 2\sum_{k\geq 1} k^{4(b-a)} ,$$

which converges if and only if $b < a - 1/4$. We have therefore proved the following.

Proposition 3.3. *Whenever $b < a - 1/4$ the norm $\|\cdot\|_b$ is measurable on $\mathcal{H}_a$.*

Note that $(f,g)_H = \big((-\Delta)^{1/4} f, (-\Delta)^{1/4} g\big)$ by definition, so $(-\Delta)^{-1/4}$ provides a Hilbert space isomorphism between the spaces H and $\mathcal{H}_{1/4}$, that when convenient we identify. It follows that in order to get a measurable norm on H it suffices to have a measurable norm on $\mathcal{H}_{1/4}$. By Proposition 3.3, any norm of the form $\|\cdot\|_{-\varepsilon}$

for $\varepsilon > 0$ will do.

Let $\mathcal{B}_a$ denote the Borel σ-algebra on $\mathcal{H}_a$ for the norm $\|\cdot\|_a$. Then for any $\varepsilon > 0$ there exists a unique Borel probability measure $\mu_{-\varepsilon}$ on $(\mathcal{H}_{-\varepsilon}, \mathcal{B}_{-\varepsilon})$ such that (iii) in Definition 3.1 holds, i.e. $\mu_{-\varepsilon} \circ \phi^{-1} = \mathcal{N}(0, \|\tilde{\phi}\|_H)$ for all ϕ continuous linear functionals on $\mathcal{H}_{-\varepsilon}$, where $\tilde{\phi}$ is the unique element of H such that $\phi(h) = (\tilde{\phi}, h)_H$ for all $h \in H$. It follows that $(H, \mathcal{H}_{-\varepsilon}, \mu_{-\varepsilon})$ is an Abstract Wiener Space for any $\varepsilon > 0$.

Definition 3.4. A random variable X is said to be a standard Gaussian on the Hilbert space $H \cong \mathcal{H}_{1/4}$ if $X \sim \mu_{-\varepsilon}$ as random variable on $\mathcal{H}_{-\varepsilon}$, for all $\varepsilon > 0$.

Remark. This definition is consistent, meaning that if $-\varepsilon_1 < -\varepsilon_2$, so that $\mathcal{H}_{-\varepsilon_2} \subset \mathcal{H}_{-\varepsilon_1}$, then by property (iii) of Definition 3.1 the restriction of $\mu_{-\varepsilon_1}$ to $\mathcal{H}_{-\varepsilon_2}$ coincides with $\mu_{-\varepsilon_2}$.

To conclude this abstract construction, we point out that X as above has a very simple expansion in Fourier basis. Indeed, if $f_k = (-\Delta)^{-1/4} e_k = e_k/\sqrt{|k|}$, then $(f_k)_k$ is an ONB for $\mathcal{H}_{1/4} \cong H$. It then follows from Corollary 3.1 that there exists a collection $(A_k)_{k \in \mathbb{Z}}$ i.i.d. real $\mathcal{N}(0,1)$ random variables such that

$$X(\vartheta) = \sum_{k \neq 0} A_k f_k(\vartheta) = \sum_{k \neq 0} \frac{A_k}{\sqrt{|k|}} e_k(\vartheta)\,.$$

Moreover, in light of Proposition 3.2 it is now easy to construct a Brownian motion W on H: simply take a collection of i.i.d. real-valued Brownian Motions $(\beta_k)_{k \in \mathbb{Z}}$, and set

$$W(t, \vartheta) = \sum_{k \neq 0} \beta_k(t) f_k(\vartheta) = \sum_{k \neq 0} \frac{\beta_k(t)}{\sqrt{|k|}} e_k(\vartheta)\,,$$

where for any $t \geq 0$ the above series converge in $L^2(\mathcal{H}_{-\varepsilon})$, for any $\varepsilon > 0$. From W we construct an Ornstein-Uhlenbeck process on H as follows. Let $\Psi : [0, \infty) \to B$ be defined by

$$\Psi(t) = \sum_{k \neq 0} \sqrt{2|k|} e^{-|k|t} f_k = \sum_{k \neq 0} e^{-|k|t} e_k\,,$$

and set $\tilde{\mathcal{W}}(t) := \int_0^t \Psi(t-s)\mathrm{d}W(s)$ for all $t \geq 0$. Then by Definition 3.3 we have

$$\tilde{\mathcal{W}}(t,\vartheta) = \sum_{k\neq 0} \left[\int_0^t e^{-|k|(t-s)} \mathrm{d}\beta_k(s) \right] e_k(\vartheta)\,,$$

so that the coefficients of $\tilde{\mathcal{W}}$ perform independent OU processes on $\mathbb{R}$. Take an independent copy $\tilde{\mathcal{W}}'$ of $\tilde{\mathcal{W}}$. Then we have

$$\frac{\tilde{\mathcal{W}}(t,\cdot) + i\tilde{\mathcal{W}}'(t,\cdot)}{\sqrt{2}} = \sum_{k\neq 0} \Big(\frac{A_k(t) + iB_k(t)}{\sqrt{2}} \Big) e_k(\cdot)\,,$$

for independent OU processes $(A_k)_k$, $(B_k)_k$ on $\mathbb{R}$, which equals the r.h.s. of (3.13). This provides a rigorous construction of the boundary process $\mathcal{W}$ of the limiting holomorphic field $\mathcal{F}$ as an OU process on $\mathcal{H}_{-\varepsilon}$, for any $\varepsilon > 0$.

Remark. Note that, by linearity,

$$\mathrm{d}\mathcal{W}(t,\cdot) = -\Big[\sum_{k\neq 0} |k| \Big(\frac{A_k(t) + iB_k(t)}{\sqrt{2}} \Big) e_k(\cdot) \Big] \mathrm{d}t + \sqrt{2} \Big[\sum_{k\neq 0} \frac{\mathrm{d}\beta_k(t) + i\mathrm{d}\beta_k'(t)}{\sqrt{2}} e_k(\cdot) \Big]\,,$$

where we have denoted by (β_k), (β_k') the i.i.d. real Brownian Motions driving the OU processes (A_k), (B_k). This shows that $\mathcal{W}$ solves the Fractional Stochastic Heat Equation

$$\mathrm{d}\mathcal{W}(t,\cdot) = -(-\Delta)^{1/2}\mathcal{W}(t,\cdot) + \sqrt{2}\,\mathrm{d}\xi(t,\cdot)$$

for ξ complex space-time white noise on the unit circle $\mathbb{T}$.

Remark. We point out that the convergence result $\mathcal{F}_n \to \mathcal{F}$ of Theorem 3.3 can be interpreted as a convergence result for the corresponding boundary values, seen as distributions acting on a suitable space of test functions. More precisely, let $\mathcal{W}_n$ denote the boundary values of $\mathcal{F}_n$, so that $\mathcal{W}_n(t,\vartheta) = \mathcal{F}_n(t, e^{i\vartheta})$ for $(t,\vartheta) \in [0,\infty) \times [-\pi,\pi)$. As space of test functions we take

$$\big\{ \varphi \in C^\infty(\mathbb{T}) : \varphi = P_{1/r} * \psi \text{ for some } \psi \in C^\infty(\mathbb{T}) \text{ and some } r > 1 \big\}\,.$$

For each such φ, we have:

$$(\mathcal{W}_n(t,\cdot),\varphi) = (\mathcal{W}_n(t,\cdot), P_{1/r} * \psi) = (P_{1/r} * \mathcal{W}_n(t,\cdot),\psi)$$
$$= (\mathcal{F}_n(t, re^{i\cdot}),\psi) \to (\mathcal{F}(t, re^{i\cdot}),\psi)$$

as $n \to \infty$ in distribution, as continuous functions on $\mathbb{T}$. It would be interesting to understand if such convergence holds for a larger class of test functions, and ultimately as stochastic processes taking values in $\mathcal{H}_{-\varepsilon}$ for $\varepsilon > 0$.

3.4 Further developments

As mentioned in the introduction, arguably the main open problem in this area is to obtain rigorous results on (the non-regularised version of) HL(α) models with $\alpha > 0$, with a particular interest for $\alpha \in [1,2]$. In this last section, on the other hand, we would like to collect some open questions for the $\alpha = 0$ case, which has proven to be already very interesting from a mathematical point of view, and it has the advantage of being more tractable than the case $\alpha > 0$ due to its intrinsic i.i.d. structure.

Maxima of the cluster boundary

Once the question of fluctuations of cluster boundary has been settled, one is led to wonder about the asymptotic behaviour of the maxima of the cluster. For $n \geq 0$, set $m_n = \sup\{|z| : z \in K_n\}$. It follows from the analysis carried out in [NT12] that $|m_n - e^{cn}| \to 0$ almost surely as $n \to \infty$, $c \to 0$ and $nc \to t$ for some $t > 0$. It would be interesting to have a description of the fluctuations around this deterministic behaviour. Do they depend on the particle shape P? A possible approach to this question could be to move away from the boundary of the cluster, and define $m_n(\sigma) = \sup\{|\Phi_{\lfloor nt \rfloor}(z)| : |z| = e^{\sigma}\}$. Our results control the fluctuations of $m_n(\sigma)$ around its mean for $\sigma > 0$ fixed as $n \to \infty$. How do these fluctuations look like if $\sigma \to 0$ quickly enough as $n \to \infty$?

Log-correlation and branching structure

Log–correlated Gaussian fields appear to arise in correspondence with underlying branching structures. To make this statement more precise, let us give some examples.

Branching Random Walk

For $N \in \mathbb{N}_+$ consider a binary tree $G_N = (V_N, E_N)$ of depth N, and assign to each edge $e \in E_N$ a standard Gaussian random variable X_e in an independent fashion. Let L_N denote the set of leaves of the binary tree (so that $|L_N| = 2^N$), and for each $v \in L_N$ denote by γ_v the unique path from v to the root. Then set

$$Y_v = \sum_{e \in \gamma_v} X_e \, .$$

It is easy to see that, when the set of leaves is seen as embedded in the unit interval $[0, 1]$ (i.e. the leaves are identified with 2^N equispaced points in $[0, 1]$), the correlation between two leaves $v, v' \in L_N$ decays logarithmically with their Euclidean distance $|v - v'|$, for all N. In other words, at each level $N \geq 1$ the random field $(Y_v)_{v \in L_N}$ is a log–correlated (Gaussian) field.

Branching in IDLA

A branching structure is also present in the Internal DLA model described in the Introduction of this thesis. Indeed, one could imagine to colour, each time a new particle is added to the cluster, the last edge along which the corresponding random walk has jumped in order to exit the cluster. More precisely, define the set of coloured edges recursively as follows. Given the cluster $A(n-1)$ with $n-1$ particle, start an independent random walk $\omega^{(n)} = (\omega^{(n)}_k)_{k \geq 0}$ from the origin and, if τ_n is the first exit time of $\omega^{(n)}$ from $A(n-1)$, colour the edge $(\omega^{(n)}_{\tau_n - 1}, \omega^{(n)}_{\tau_n})$. It is easy to see that the set of coloured edges forms a tree. Moreover, if one considers two points on the cluster boundary (which is asymptotic circular), then it follows from the work in [JLS14a] that the correlation between fluctuations from circularity at the two points decays logarithmically with their distance.

Branching in HL growth

The fact that a branching mechanism arises in the growth of HL(0) clusters was discovered by J. Norris and A. Turner in [NT12]. Think for simplicity of the particle P as being a slit, i.e. $P = [1, 1+\delta] \subset \mathbb{C}$. Then each time a new particle arrives it gets attached to exactly one particle in the cluster, at exactly one point, almost surely. Call this *parent particle*. This identifies an ancestral structure in the cluster. Given any finite set of points, it is shown in [NT12] that the corresponding ancestral lines converge to backwards coalescing Brownian Motions, i.e. Brownian Motions which evolve independently backwards in time until they coalesce, at which point they merge. This result can be rephrased by saying that the cluster boundary converges, in a finite dimensional sense, to the so called *Brownian Web*. Here the underlying (forward) branching structure is then really a backwards coalescing structure of the cluster boundary[3], but nonetheless one obtains log-correlated fluctuations.

In all the above examples log-correlated Gaussian fields arise as fluctuation fields of several different models, all featuring some kind of underlying branching structure. We believe it would be very interesting to understand to what extent this is a general phenomenon, and, if so, how robust it is with respect to variations of the branching mechanism (e.g. introduction of correlation between branches).

Genealogical structure of Internal DLA

As already observed, IDLA clusters on the cylinder $\mathbb{Z}_N \times \mathbb{Z}$ have the same scaling limit as (the logarithm of) HL(0) clusters, as $N \to \infty$. Moreover, they share the same boundary fluctuation process (see (1.17) in the Introduction). It is therefore natural to ask whether they have further properties in common. In particular, as

[3]Note that the dynamics of backwards coalescing Brownian Motions is different from the one of forward branching Brownian Motions, since for example in the latter paths can intersect after branching.

mentioned in the previous open question, we know that the genealogical structure of HL(0) clusters scales to the Brownian Web. Note that one could also identify a genealogical structure in IDLA clusters by saying that a vertex w is the parent of another vertex v if w and v are neighbours, and if v has been added to the cluster via a jump of the driving random walk along the edge connecting w to v. Then, given any finite collection of points on $\mathbb{Z}^2$ one could draw the associated random lines of ancestors connecting it to the origin. Do these random curves also converge to coalescing Brownian Motions? Do their scaling limit reflect the geometry of the underlying lattice structure?

Part II

Internal DLA on cylinder graphs

Chapter 4

IDLA on cylinder graphs

4.1 Introduction

Let $G = (V, E)$ be a finite, connected graph. We consider Internal Diffusion Limited Aggregation, in short IDLA, on the infinite cylinder $\mathcal{C}_G := G \times \mathbb{Z}$. This is a discrete–time Markov process $(A(t))_{t \geq 0}$ in the space of connected subsets of $G \times \mathbb{Z}$, which we call *clusters*, defined as follows.

Assume that the initial IDLA configuration is such that

$$A(0) \supseteq \{(x, y) \in V \times \mathbb{Z} : y \leq 0\} . \tag{4.1}$$

We say that $\omega = (\omega_k)_{k \geq 0}$ is a simple random walk on $\mathcal{C}_G$ if ω moves in the y coordinate with probability $1/2$, and otherwise it moves in the x coordinate by taking a simple random walk step on G. See Section 4.2.1 for a more precise definition. Let π denote the stationary measure associated to a simple random walk on G. At each discrete time $t \geq 1$, we start a nearest neighbour walk $\omega^{(t)} = (\omega_k^{(t)})_{k \geq 1}$ on $\mathcal{C}_G$ from a vertex $(x, 0)$, with x distributed according to π. Let $\tau^{(t)}$ denotes the first exit time of the walk from $A(t-1)$, i.e.

$$\tau^{(t)} := \inf\{k \geq 0 : \omega_k^{(t)} \notin A(t-1)\},$$

and define recursively

$$A(t) := A(t-1) \cup \{\omega^{(t)}_{\tau^{(t)}}\}$$

for all $t \geq 1$. This corresponds to choosing the location of the next added particle according to the harmonic measure of the outer boundary of the cluster seen from the zero level $G \times \{0\}$. We remark that the driving random walk is instantaneous, i.e. we do not take into account how long it takes for the walk to exit the cluster, but rather increase the time by 1 each time we add a new particle to it. Note further that each vertex of $\mathcal{C}_G$ contains at most one particle: if it does contain one, we say that it is occupied (or filled), and otherwise we say that it is empty.

This procedure defines a Markov chain $(A(t))_{t\geq 0}$ on the infinite space of connected subsets of $G \times \mathbb{Z}$. The chain is clearly transient, since once a vertex has been filled it can never become empty again. In order to make it recurrent, we introduce the following *shift procedure*: each time the cluster is completely filled up to level $k > 0$, we shift the cluster down by k.

Definition 4.1 (Shifted chain)**.** Let $\mathcal{S}$ be the map from the space of IDLA configurations to itself defined as follows. For a given IDLA configuration A, let

$$k_A := \max\{k \geq 0 : (x, y) \in A \text{ for all } x \in G, y \leq k\}.$$

Define the *shift operator* $\mathcal{S}$ by

$$\mathcal{S} : A \mapsto \mathcal{S}(A) := \{(x, y - k_A \mathbb{1}(k_A > 0)) : (x, y) \in A\},$$

where $\mathbb{1}$ denotes the indicator function. Then, from the IDLA chain $(A(t))_{t\geq 0}$ we build the *shifted chain* $(A_S(t))_{t\geq 0}$ by setting

$$A_S(t) := \mathcal{S}(A(t))$$

for all $t \geq 0$.

This defines a new Markov Chain on the same configuration space Ω (in fact, the shifted chain only takes values in a strict subset of Ω, consisting of all configurations with at least one empty site at level 1). Note that recurrence of the shifted IDLA process is not clear (in principle, the process could grow arbitrarily long arms), nor is the existence of a stationary distribution.

The shifted chain was first introduced in [JLS14a], where the authors ask the following question: *how long does it take for the shifted IDLA process to forget where it started?* In this chapter we partially answer this question by providing an upper bound for this random time, that we call *forget time*. Our main result is the following.

Theorem 4.1. *Let $(A(t))_{t\geq 0}$, $(A'(t))_{t\geq 0}$ be two IDLA processes starting from $A(0), A'(0)$ respectively. We assume that $A(0), A'(0)$ satisfy* (4.1)*, and that they have the same number of vertices above level* 0*, i.e.* $|A(0)\cap\{y>0\}| = |A'(0)\cap\{y>0\}|$. *Define the* forget time $\mathscr{T}\in[0,+\infty]$ *to be the first time the two configurations match, i.e.*

$$\mathscr{T} := \inf\{t\geq 0 : A(t) = A'(t)\},$$

with the convention that $\inf\emptyset = +\infty$. *Finally, let $2d = |A(0)\Delta A'(0)|$ denote the cardinality of the symmetric difference of the initial configurations, and h_0 denote the maximal height among $A(0)$ and $A'(0)$. Then there exists a coupling of the two IDLA processes such that, if $\mathbb{P}$ denotes the joint law of $(A(t))_{t\geq 0}$ and $(A'(t))_{t\geq 0}$ under such coupling, it holds*

$$\mathbb{P}(\mathscr{T} > t) \leq \bar{\mathscr{C}} e^{-\bar{\lambda} t} \tag{4.2}$$

for positive constants $\bar{\mathscr{C}}, \bar{\lambda} > 0$, with

$$\bar{\lambda} = \bar{\lambda}(d, G), \qquad \bar{\mathscr{C}} = \bar{\mathscr{C}}(\bar{\lambda}, d, G, h_0).$$

We point out that the advantage of looking at the forget time (rather than a different notion of mixing time) is that its definition does not require any information about the stationary distribution, and in fact it makes sense also for transient chains.

On the way to proving the above theorem, we show that the flat configuration is positive recurrent for the shifted chain. Since the chain consists of only one communicating class, this settles both the question of recurrence and existence of the stationary distribution for $(A_S(n))_{n\geq 0}$.

Proposition 4.1. *The shifted chain $(A_S(t))_{t\geq 0}$ is positive recurrent, and hence it has a stationary distribution.*

4.1.1 Organisation of the chapter

This chapter is organised as follows. In Section 4.2 we fix the notation, review the Abelian property of IDLA and use it to define the *freezing coupling*. In Section 4.3 we focus on one IDLA process, obtaining an upper bound for the first time the discrepancy falls below a given threshold (see Proposition 4.3). On the way to controlling the discrepancy, we show in Remark 4.3.1 that the associated shifted IDLA process is positive recurrent, thus proving Proposition 4.1. In Section 4.4 we then obtain a lower bound for the probability that the frozen particles couple (and hence that the trial is successful) when released from good configurations. We put all together in Section 4.5, to conclude the proof of Theorem 4.1. Finally, we devote Section 4.6 to the case of the N–cycle base graph, and provide a lower bound for the decay rate $\bar{\lambda}$ of Theorem 4.1 in terms of the size of the base graph N, for N large enough.

4.2 The freezing coupling

4.2.1 Notation

Let us start by fixing the notation in use throughout the chapter.

The cylinder graph $\mathcal{C}_G$

Let $G = (V, E)$ be a finite, connected graph and write N for the cardinality of the vertex set V of G. We define the cylinder graph $\mathcal{C}_G$ by placing infinitely many copies of G one above the other. More precisely, let the vertex set of $\mathcal{C}_G$ be given by $\{(x,y) : x \in G, y \in \mathbb{Z}\}$. Moreover, two vertices (x,y) and (x',y') are said to be adjacent, i.e. joined by an edge, if and only if either $x = x'$ and $|y - y'| = 1$, or $y = y'$ and x is a neighbour of x' in G. See Figure 4.1 for an example.

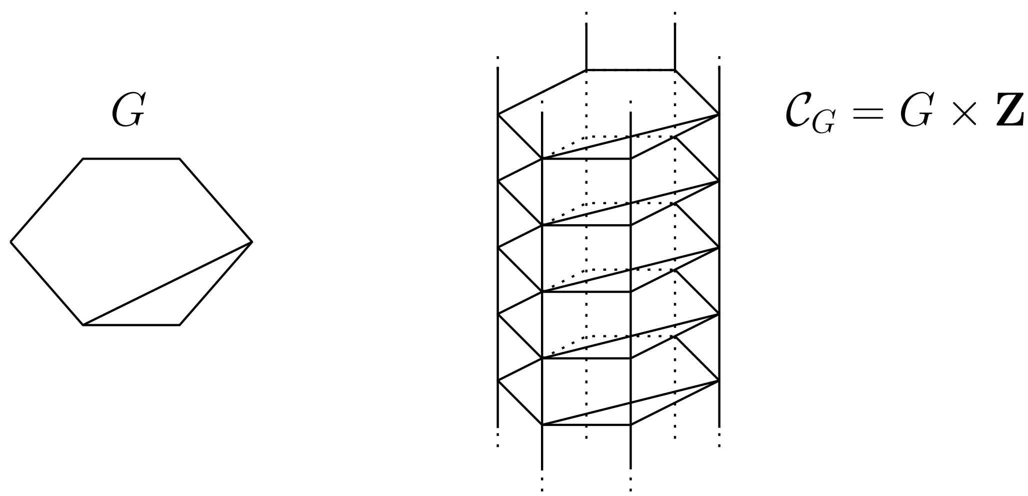

Figure 4.1: A graph G (left) and the associated cylinder graph $\mathcal{C}_G = G \times \mathbb{Z}$ (right).

We refer to x and y as the horizontal and vertical coordinate respectively, and call $G \times \{k\} = \{(x,k) : x \in G\}$ the k^{th} level of $\mathcal{C}_G$.

Simple random walk on G

Recall that $(x_k)_{k\geq 0}$ is a simple random walk on G if, at each step $k \geq 0$, given x_k the next location x_{k+1} is chosen uniformly at random among the nearest neighbours of x_k in G. An α–lazy simple random walk on G is a process that, at each integer time $k \geq 0$, stays in place with probability α and otherwise jumps to a uniformly chosen neighbour of the current location. We say that a simple random walk on G is lazy if it is $1/2$–lazy.

Simple random walk on $\mathcal{C}_G$

Fix $p \in (0,1)$. We say that $\omega = (x_k, y_k)_{k\geq 0}$ is a simple random walk on $\mathcal{C}_G$ if, at each step $k \geq 0$, ω moves in the y coordinate with probability $1/2$, stepping from (x_k, y_k) to $(x_k, y_k + 1)$ with probability p and to $(x_k, y_k - 1)$ with probability $1 - p$.

Moreover, if it does not move in the y coordinate then it performs a lazy simple random walk step on G, i.e. it either stays in place with probability $1/2$ or it jumps from (x_k, y_k) to (x'_k, y_k), for x'_k uniformly chosen neighbour of x_k in $G \times \{k\}$. Note that if a simple random walk on $\mathcal{C}_G$ starts from $(x_0, 0)$ with x_0 distributed according to the stationary measure π, then x_k is distributed according to π for all $k \geq 0$. For simplicity, we take $p = 1/2$ throughout the chapter, so that the vertical coordinate, when it moves, performs a simple symmetric random walk on $\mathbb{Z}$. All our results can be adapted to the transient case $p > 1/2$.

A *cluster* is simply a subset of the vertex set of the cylinder graph $\mathcal{C}_G$. Given an IDLA cluster A, we denote by $h(A)$ its height, that is

$$h(A) = \max\{y : (v, y) \in A \text{ for some } v \in G\}.$$

We assume (4.1) throughout the chapter, so that $h(A) \geq 0$. With a slight abuse of notation we write $|A|$ for $|A \cap \{y > 0\}|$, denoting the number of vertices in A above level 0.

4.2.2 The Abelian property

We now review the Abelian property of IDLA, first observed by Diaconis and Fulton in [DF91]. To this end it is useful to introduce the notion of *stack of cards*. For each vertex z of the cylinder graph $\mathcal{C}_G$ denote by E_z the set of edges of the graph incident to z. A stack of cards at z is an infinite sequence of edges in E_z. Let us assume that to each such vertex z we have associated an infinite stack of cards

$$z \quad \mapsto \quad (d_i^z)_{i \geq 1} \in E_z^{\mathbb{N}_+}.$$

Let σ denote an initial distribution of particles on $G \times \{0\}$, that is $\sigma(z)$ =number of particles at $(v, 0)$, $v \in G$, at time 0. A *legal move* consists of the following steps:

- Pick $z \in \mathcal{C}_G$ with at least 2 particles.

- Move the top particle to one of the nearest neighbours, following the edge of the top card at z.

- Discard the top card at z.

We write (z, e) for the legal move consisting of moving one particle from z to a neighbour of z following the edge $e \in E_z$, and then discarding the top card at z. Note that, since we have an infinite stack of cards at each point, we never run out of cards. Assume[1] that after finitely many steps the process stabilizes, i.e. it reaches a configuration with at most one particle at each site (such configuration is called *stable*). The Abelian property tells us that, as long as we start from the same initial configuration and use the same stack of instructions, any two sequences of legal moves will result in the same final configuration. We stress that for this to hold it is crucial for the particles to be indistinguishable.

Proposition 4.2 (see [DF91]). *Fix an initial configuration of particles σ, and assign to each vertex of $\mathcal{C}_G$ an infinite stack of instructions. Let $\alpha_1, \dots, \alpha_m$ and $\beta_1, \dots, \beta_n$ be any two finite sequences of legal moves resulting in stable configurations η_1 and η_2 respectively. Then $m = n$ and $\eta_1 = \eta_2$.*

Proof. We begin by observing that any two legal moves commute: this is trivially true from the definition of legal move. Now, let $\alpha_1 = (z_1, e_1)$ for some vertex z_1 and edge e_1 adjacent to z_1. If this is a legal move, then in the initial configuration there were at least 2 particles at z_1. It follows that there exists β_i in the second sequence of legal moves such that $\beta_i = (z_1, e_1)$. Indeed, there must be at least one instruction of the form $\beta_j = (z_1, e_j)$ in the sequence, otherwise there would still be at least 2 particles in the final configuration, which is not allowed. Moreover, if we take i to be the smallest among all indices j such that $\beta_j = (z_1, e_j)$ then it must be the case that $\beta_i = (z_1, e_1)$. Now, using that any two legal moves commute, we can exchange β_i and β_{i-1}, then β_i and β_{i-1} and so on until we obtain that the sequence $\beta_i, \beta_1 \dots \beta_{i-1}, \beta_{i+1} \dots \beta_n$ of legal moves gives the same final configuration

[1] We really have to assume this here, since we are on an infinite graph and the bottom half is filled.

as $\beta_1 \dots \beta_n$. It follows that it suffices to show that the two sequences of legal moves

$$\alpha_1, \dots, \alpha_m$$
$$\beta_i, \beta_1 \dots \beta_{i-1}, \beta_{i+1} \dots \beta_n$$

give the same final configuration. But they start with the same legal move $\alpha_1 = \beta_i$, so we are left to prove that

$$\alpha_2, \dots, \alpha_m$$
$$\beta_1 \dots \beta_{i-1}, \beta_{i+1} \dots \beta_n$$

give the same final configuration. Iterate to conclude that $m = n$ and the sequence $\alpha_1, \dots, \alpha_n$ is simply a permutation of $\beta_1, \dots, \beta_n$. Since any two legal moves commute, we gather that the two sequence must give the same final configuration. $\square$

To build an IDLA cluster with n particles above level 0, starting from the half filled cylinder graph, we choose the initial configuration by placing the n particles on vertices of $G \times \{0\}$ independently from each others, according to the stationary measure π of a simple random walk on G. We then assign to each vertex v of $\mathcal{C}_G$ an infinite stack of i.i.d. instructions, each consisting of a uniformly chosen edge among all possible edges adjacent to v. Thus the initial configuration and the stack of instructions are random, but the Abelian property tells us that there is no additional randomness in the model.

4.2.3 The freezing coupling

The Abelian property gives us some freedom on how to grow an IDLA cluster. For example, it implies that we can *freeze* some of the driving random walks before they exit the cluster, and restart them at a later time. Indeed, at each move we can either keep on evolving the current driving random walk, or start a new one (which would correspond to freezing the current one). This version of IDLA with freezing will generate stacks of instructions which are identical in law to the ones of standard IDLA. We can then construct an IDLA process using exactly these stacks of instructions, to recover, by the Abelian property, the same configuration of IDLA with freezing after all frozen particles have been released. This is the main idea behind the *freezing coupling*, that we now describe.

Let $A(0)$, $A'(0)$ be two initial IDLA configurations. We assume that they satisfy (4.1) and that they have the same number of vertices above level zero, i.e. $|A(0)| = |A'(0)| =: n_0$ (recall that by $|A|$ we mean the number of occupied vertices in A strictly above level zero). Finally, assume that at least one of the two configurations has an empty site at level 1 (if not, just shift everything down). To couple the IDLA processes $(A(t))_{t\geq 0}$, $(A'(t))_{t\geq 0}$, at each new release we use the same driving random walk on both clusters. If the walk exits both clusters at the same time, then we add the exit location to both clusters. If, on the other hand, the walk reaches, say, a vertex $z \in A(t) \setminus A'(t)$, then we add the vertex z to $A'(t)$, and one frozen particle at location z in $A(t)$. More formally, let $F(t), F'(t)$ denote the sets of frozen particles in $A(t), A'(t)$ respectively, and set $F(0) = F'(0) = \emptyset$. Then, if the driving random walk exits $A(t) \cap A'(t)$ at z, we proceed as follows.

- If $z \notin A(t) \cup A'(t)$ no freezing occurs, and we set

$$\begin{cases} A(t+1) = A(t) \cup \{z\} \\ F(t+1) = F(t) \end{cases} \qquad \begin{cases} A'(t+1) = A'(t) \cup \{z\} \\ F'(t+1) = F'(t). \end{cases}$$

- If $z \in A(t) \setminus A'(t)$, we set

$$\begin{cases} A(t+1) = A(t) \\ F(t+1) = F(t) \cup \{z\} \end{cases} \qquad \begin{cases} A'(t+1) = A'(t) \cup \{z\} \\ F'(t+1) = F'(t). \end{cases}$$

- If $z \in A'(t) \setminus A(t)$, we set

$$\begin{cases} A(t+1) = A(t) \cup \{z\} \\ F(t+1) = F(t) \end{cases} \qquad \begin{cases} A'(t+1) = A'(t) \\ F'(t+1) = F'(t) \cup \{z\}. \end{cases}$$

After an almost surely finite number of releases the two IDLA clusters will match. Let

$$T = \inf\{t \geq 0 : A(t) = A'(t)\}$$

denote this matching time. Then $A(T) = A'(T)$, but clearly $F(T) \neq F'(T)$. In fact, let[2] $2d = |A(0)\Delta A'(0)|$. Then $|F(T)| = |F'(T)| = d$ and $F(T) \cap F'(T) = \emptyset$, since

$$\begin{aligned} F(T) &= A(0) \setminus (A(0) \cap A'(0)), \\ F'(T) &= A'(0) \setminus (A(0) \cap A'(0)). \end{aligned}$$

We stress that the law of $A(T), A'(T)$ is not, in general, the law of an IDLA cluster with T particles, since we still have to finish off the evolution of the driving random walks which resulted in a freezing step. On the other hand, the Abelian property tells us that if we now release all the frozen particles, i.e. keep on evolving the corresponding random walks until they exit the cluster, then the final clusters will have the same law as IDLA clusters with $n_0 + T + d$ particles above level zero.

4.2.4 The freezing coupling phase

Let $A(0), A'(0)$ denote the two initial IDLA clusters, and $(A(t))_{t\geq 0}, (A'(t))_{t\geq 0}$ be the associated IDLA processes of which we would like to couple the dynamics. As a first attempt, we could think to evolve both processes according to the freezing coupling, and wait until they match. Once they do, we want to release (i.e. finish off the evolution of) the frozen particles. We could do this by releasing them in pairs, one frozen particle for each cluster, possibly coupled so as to maximise the probability for them to meet before they exit their clusters. If all pairs meet before exiting, then the final configurations have the law of IDLA clusters (since we have finished evolving all particles), and they match, so the forget time has occurred.

Clearly, releasing the frozen particles right after the matching time T is not a

[2] It is easy to see that if the two configurations have the same number of particles above level zero, and are completely filled below level zero, then the symmetric difference must consist of an even number of particles. We remind the reader that the symmetric difference of two sets is defined as $A\Delta B = (A \cup B) \setminus (A \cap B)$.

good strategy, as we have no information on the clusters $A(T) = A'(T)$, and hence on the probability that two frozen particles will couple upon release (for example, this probability will be small if the frozen particles are too close to the boundary of the clusters). Instead, we keep on adding new particles to both clusters while using the trivial coupling (so that they keep on matching), until the clusters become *good*. Only then we will release the frozen particles.

Definition 4.2 (Good clusters)**.** Let D, $\tilde{h}$ be two fixed integers to be specified later. We say that $A(t)$ is *good* if the following conditions hold:

(i) all levels of the cluster apart from the top D ones are completely filled, and

(ii) all the frozen particles in $F(t)$ are at least at distance $\tilde{h}$ from the boundary of the cluster. More precisely, if $h_F := \max\{y : (x,y) \in F(t) \text{ for some } x \in G\}$, with the convention that $h_F = 0$ if $F(t) = \emptyset$, then we ask that $A(t)$ is completely filled up to height $h_F + \tilde{h}$.

Condition (ii) is easy to ensure: it suffices to release enough extra particles. Indeed, if h_0 denotes the maximal initial height among the two initial configurations, then[3] all frozen particles will be at height at most h_0, i.e. $h_F \leq h_0$. It follows that (ii) holds as long as the cluster is completely full at least up to height $h_0 + \tilde{h}$. This is in turn guaranteed if the cluster contains at least $(h_0 + \tilde{h} + D)N$ particles above level zero, and (i) holds. We conclude that if $n_0 \leq (h_0 + \tilde{h} + D)N$ it suffices to release extra $(h_0 + \tilde{h} + D)N$ particles to ensure that as soon as (i) holds then also (ii) holds. Condition (i) requires more work, and we discuss it in detail in the next section.

We will show that, after releasing enough extra particles, the cluster becomes good. At this point we release the frozen particles in pairs. If their distance $\tilde{h}$ from the boundary of the clusters if large enough, then any pair of frozen particles has a positive probability to couple before one of them exits its cluster. Since we have at most d pairs of frozen particles, they will all couple with positive probability, in

[3] Note that frozen particles can only occupy locations in $A(0)\Delta A'(0)$ by definition of the freezing coupling. In particular, once the two clusters match no further freezing occurs.

which case the final IDLA clusters match, and the forget time has occurred. If, on the other hand, a pair fails to couple, then after releasing all the frozen particles we end up with two IDLA clusters of the same cardinality, differing by at most $2d$ particles. We can thus restart the coupling.

The method described above works well, but we can do better. To see this, note that the first matching time of two IDLA processes evolved with freezing coupling coincides with the first time both clusters contain the subset of vertices $A(0)\Delta A'(0)$. Indeed, freezing can only happen at locations in this set, and hence once the set is filled the clusters must match. Moreover, $A(0)\Delta A'(0)$ is itself contained in the rectangle $\{(v,y) : y \leq h_0\}$, so, as long as the two clusters contain such rectangle, they must match. But conditions (i)-(ii) force the clusters to contain the even larger rectangle $\{(v,y) : y \leq h_0+\tilde{h}\}$. It follows that, instead of waiting for the two clusters to match and *then* wait for them to become good, we can instead simply evolve both clusters according to the freezing coupling, and wait until the first time they become good: when they do, they must also match.

To implement this strategy, starting with initial clusters $A(0), A'(0)$ we release $(h_0+\tilde{h}+D)N$ particles ($\tilde{h}, D$ to be chosen later) while using the freezing coupling. This increases the initial height of the configurations by at most $(h_0+\tilde{h}+D)N$. Starting from the new height, we then need to bound the time it takes for the clusters to become good, i.e. for (i) to hold. This is the object of the discussion in the next section.

4.3 The discrepancy decay phase

In order to formalise condition (i) in the previous section, it is useful to introduce the notion of *dicrepancy*.

Definition 4.3 (Discrepancy)**.** Let A be a connected cluster in $\mathcal{C}_G$ satisfying (4.1). We define the *discrepancy* of A, denoted by $\delta(A)$, as the distance between the lowest empty site and the highest vertex in the cluster. More precisely,

$$\delta(A) := h(A) - \min\{y : (v, y) \notin A \text{ for some } v \in G\} + 1$$

(note that $\delta(A) = h(A_S)$, where A_S denotes the shifted version of A).

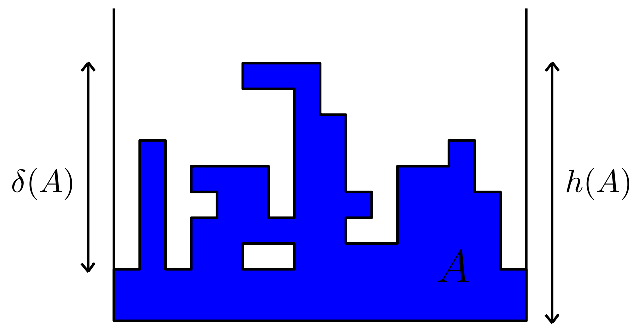

Figure 4.2: Summary of notation.

Then (i) could have been equivalently stated as

(i) the discrepancy of the cluster is at most D.

In this section we focus on one single IDLA process, that we keep on calling $(A(t))_{t\geq 0}$, and obtain bounds on the random time

$$\mathcal{T} := \inf\{t \geq 0 : \delta(A(t)) \leq D\}$$

as a function of the initial height. We use different methods to control the discrepancy, according to whether the density of the cluster is low or high, where the density of the cluster is defined as the proportion of occupied vertices in the smallest rectangle containing the cluster above level 0. When the cluster has low density, the discrepancy might decay quite slowly (for example, for a configuration consisting of a column of height δ, in order to fill one level we might have to release up to $\delta(N-1)+1$ particles). On the other hand, since there are many empty sites below the top level, it is very unlikely for the cluster to increase its height. This suggests that the density of the cluster increases on average, and so the cluster quickly becomes of high density. Once this happens, we can control the number of empty sites below the top level of the cluster, although we do not know their position. For the discrepancy to decay below D, then it suffices to fill all these empty sites, while making sure that the height of the configuration does not increase too much.

We distinguish between high and low density clusters by controlling the so called *excess height*.

Definition 4.4 (Excess height)**.** We call *excess height* of a cluster A the difference between the minimal height and the actual height of A, i.e.

$$\mathcal{E}(A) := h(A) - \frac{|A|}{N}.$$

We say that a given configuration has high density if its excess height is bounded by a constant C (to be chosen later), independent of the cardinality of the cluster.

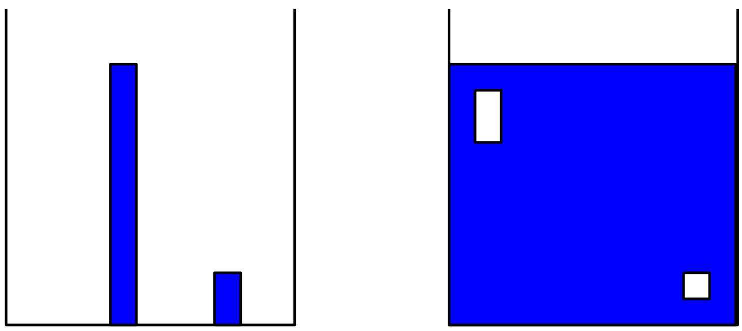

Figure 4.3: Clusters with high excess height (left) and low excess height (right).

As a first step, we show in Lemma 4.1 that, for a suitable choice of C, the excess height of a configuration drops below C quickly under the IDLA dynamics. Once the excess height is controlled, also the number of holes (i.e. empty vertices) below the top level is controlled. It only remains to fill those holes. To this end, we release a deterministic number of extra particles, enough (to have a good chance to fill all the holes) but not too many (not to increase the excess height too much). If we filled all the holes we are done: the discrepancy must now be controlled. If we failed, we again wait for the excess height to drop below a given constant, and iterate. We will succeed after a Geometric number of i.i.d. attempts.

4.3.1 Decay of excess height

The excess height is easier to control than the discrepancy. Indeed, as well as being independent of the positions of the holes in the configuration, the excess height only changes by a small amount at each step, thus allowing us to construct associated supermartingales with bounded increments (see (4.6) below).

The following result controls the time it takes for the excess height to fall below some given constant.

Lemma 4.1. *Let $(A(t))_{t\geq 0}$ denote an IDLA process on the infinite graph $\mathcal{C}_G$, and assume that $A(0)$ satisfies (4.1). For $t \geq 0$ let $\mathcal{E}(t) := \mathcal{E}(A(t))$ denote the excess height of $A(t)$. Finally, recall that π denotes the stationary measure for a simple random walk on the base graph G, and let*

$$m_\pi := \min_{v\in G} \pi(v). \tag{4.3}$$

Then for any $\eta \in (0,1)$ there exists a constant $C = C(G, \eta, m_\pi)$ such that, if

$$S := \inf\{t \geq 0 : \mathcal{E}(t) < C\}$$

denotes the first time that the excess height drops below C, then

$$\mathbb{P}(S > t) \leq e^{-\frac{\eta^2}{32N^2}t},$$

for t large enough depending on $C(G, \eta, m_\pi)$ and on the initial excess height $\mathcal{E}(0)$.

Remark. It is worth pointing out that knowing that the excess height is smaller that C does not guarantee that (i) holds, since there could still be an empty site at a very low level.

Lemma 4.1 is an easy consequence of the following result.

Lemma 4.2. *Let $\mathcal{F}_t := \sigma\{A(s) : s \leq t\}$, and write $h(t)$ in place of $h(A(t))$ for brevity. Then for all $\eta \in (0,1)$ there exists an integer constant $C(G, \eta, m_\pi)$ such* *that, while $\mathcal{E}(t) \geq C$, it holds*

$$\mathbb{E}(h(t+1) - h(t)|\mathcal{F}_t) < \frac{1-\eta}{N}\,.$$

Proof. Pick any $\eta \in (0,1)$, which is kept fixed throughout the proof. Assume that $\mathcal{E}(0) \geq C$, with C to be chosen later. Note that, while $\mathcal{E}(t) \geq C$, there are at least C particles above level $\lfloor |A(t)|/N \rfloor$.

We say that a level is *bad* if it contains at least one hole (i.e. empty site). It follows that, if $\mathcal{E}(t) \geq C$, then there are at least C bad levels below (and including) the top one $h(t)$. Indeed, since $h(t) \geq \frac{|A(t)|}{N} + C$ then there are at least $\frac{|A(t)|}{N} + C$ levels in the cluster, and at most $\left\lfloor \frac{|A(t)|}{N} \right\rfloor$ of them can be completely filled.

For each vertex x in G, let $\nu_x^{(k)}(x')$ denote the probability that a simple random walk on $\mathcal{C}_G$, starting from $(x,0)$, reaches level k for the first time at (x',k), for $x, x' \in G$. Let

$$l_G := \inf\left\{ n \geq 1 : \min_{x,x' \in G} \nu_x^{(n)}(x') \geq \frac{m_\pi}{2} \right\}. \tag{4.4}$$

In words, l_G is the lowest level such that a simple random walk starting at level zero has probability at least $m_\pi/2$ to reach that level for the first time at any given vertex, uniformly over the starting location.

Since there are at least C bad levels, we can find at least $\frac{C}{l_G}$ bad levels at distance at least l_G one from the other. We treat these bad levels as traps. More precisely, for the driving random walk to increase the height of the configuration, the walker has to survive all the bad levels without falling into a hole, until it reaches the top. Since we are taking the bad levels sufficiently far apart, the walk has time to mix in between, so it gets to each bad levels for the first time at an approximately π–distributed point. It follows that it has probability at least $m_\pi/2$ to fall into a hole. By taking enough bad levels, then, we can make the probability for a walker to survive all of them arbitrarily small.

Formally, since at each bad level the walker has probability at least $\frac{m_\pi}{2}$ to fall

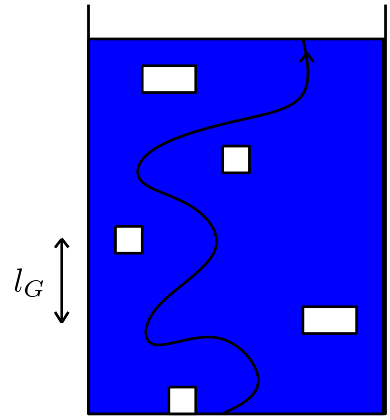

Figure 4.4: An example of the trajectory of a driving random walk surviving all the bad levels at distance at least l_G one from the other.

into a hole, we have

$$\mathbb{P}(h(t+1) - h(t) = 1 | \mathcal{F}_t) \leq \left(1 - \frac{m_\pi}{2}\right)^{C/l_G}. \tag{4.5}$$

Clearly, by choosing C large enough (depending on G, η, m_π) we can make sure that the right hand side is smaller than $\frac{1-\eta}{N}$, which concludes the proof. $\square$

Proof of Lemma 4.1. Let $n_0 := |A(0)|$. It follows from Lemma 4.2 that

$$M(t) := h(t) - \frac{n_0 + (1-\eta)t}{N}, \quad t \geq 0$$

is a supermartingale up to the stopping time S. Moreover

$$|M(t+1) - M(t)| \leq |h(t+1) - h(t)| + \frac{1-\eta}{N} \leq 2 \tag{4.6}$$

for all $t < S$. Finally, using that $|A(t)| = n_0 + t$ we see that

$$\{\mathcal{E}(t) > C\} = \left\{h(t) - \left\lceil \frac{n_0 + t}{N} \right\rceil > C\right\} \subseteq \left\{M(t) > C + \frac{\eta}{N} t\right\}.$$

It therefore follows from Azuma's inequality that

$$\mathbb{P}(S > t) \leq \mathbb{P}\big(\mathcal{E}(t) > C\big) = \mathbb{P}\Big(M(t) - M(0) > C + \frac{\eta}{N}t - M(0)\Big)$$
$$\leq \exp\Big\{ - \frac{(C + \eta t/N - M(0))^2}{8t}\Big\} \leq \exp\Big\{ - \frac{\eta^2}{32N^2}t\Big\},$$

where the last inequality holds for t large enough so that $\frac{\eta t}{2N} > M(0) - C = \mathcal{E}(0) - C$. □

It is useful to understand how the above time S depends on the initial excess height of the configuration.

Corollary 4.1. *Let $(A(t))_{t\geq 0}$ denote an IDLA process on the infinite graph $\mathcal{C}_G$, and assume that $A(0)$ satisfies (4.1). Then for λ small enough (depending only on η, N) it holds*

$$\mathbb{E}\big(e^{\lambda S}\big) \leq \hat{c}(\lambda) \exp\Big\{\frac{2N\lambda}{\eta}(\mathcal{E}(0) - C)\Big\},$$

with C as in Lemma 4.1, and $\hat{c}(\lambda) = \hat{c}(\lambda, \eta, N) := 1 + \frac{\lambda}{\frac{\eta^2}{32N^2} - \lambda}$.

Proof. In Lemma 4.1 we have shown that

$$\mathbb{P}(S > t) \leq e^{-\frac{\eta^2}{32N^2}t} \qquad \text{for all } t \geq \frac{2N}{\eta}(\mathcal{E}(0) - C).$$

It follows that, as long as $\lambda < \eta^2/32N^2$, integrating by parts we find

$$\mathbb{E}(e^{\lambda S}) = \int_0^\infty e^{\lambda t}\frac{\mathrm{d}}{\mathrm{d}t}\Big(- \mathbb{P}(S > t)\Big)\mathrm{d}t \leq 1 + \int_0^\infty \lambda e^{\lambda t}\mathbb{P}(S > t)\mathrm{d}t$$
$$\leq 1 + \int_0^{\frac{2N}{\eta}(\mathcal{E}(0)-C)} \lambda e^{\lambda t}\mathrm{d}t + \int_{\frac{2N}{\eta}(\mathcal{E}(0)-C)}^\infty \lambda e^{-\big(\frac{\eta^2}{32N^2} - \lambda\big)t}\mathrm{d}t$$
$$= e^{\lambda\frac{2N}{\eta}(\mathcal{E}(0)-C)}\left[1 + \frac{\lambda}{\frac{\eta^2}{32N^2} - \lambda}\right],$$

as claimed. □

Remark. Lemma 4.1 implies that the associated shifted Markov chain (4.1) is positive recurrent, thus proving Proposition 4.1. Indeed, since the chain consists of one communicating class, to see this it suffices to show, for example, that the flat configuration $A_0 = \{(x, y) \in \mathcal{C}_G : y \le 0\}$ is positive recurrent. This corresponds to the cluster being a perfect rectangle in the associated IDLA process. Let $(A(t))_{t \ge 0}$ be an IDLA process with $A(0) = A_0$, and define T_0 to be the first time t such that $A(t) = \{(x, y) \in \mathcal{C}_G : y \le k\}$ for some $k \ge 1$ (we call such configurations perfect rectangles). Here is a wasteful way to show that $\mathbb{E}_{A_0}(T_0) < \infty$. Let $C = C(G, \eta, m_\pi)$ be the integer constant in Lemma 4.1. Starting from A_0, release CN particles: with probability at least m_π^{-CN}, m_π defined as in (4.3), the final configuration will be a perfect rectangle. If not, then the excess height has increased by at most CN. We keep releasing particles until the excess height falls again below C: this takes a random time with exponential tails, and hence finite expectation, by Corollary 4.1. Once the excess height is at most C, there are at most CN holes below the top level in the configuration. We release as many particles as the number of holes below the top level: with probability at least m_π^{-CN} the final configuration will be a perfect rectangle. If not, the excess height is at most $C + CN$: we again wait for it to fall below C, and iterate. After at most a Geometric(m_π^{-CN}) number of attempt, the final configuration will be a perfect rectangle. Each attempt takes CN releases, plus the time it takes for the excess height to fall below C starting from $C + CN$, which has exponential tails. Thus, using for example Wald's identity, we conclude that the total time it takes to go from A_0 back to a flat configuration has finite expectation. This shows that the shifted chain is positive recurrent, therefore proving Proposition 4.1.

4.3.2 Filling the holes

At the stopping time S we have a cluster $A(S)$ with excess height at most C, for $C = C(G, \eta, m_\pi)$ as in the statement of Lemma 4.1. Note that, by definition, if the excess height decreases then it decreases by $1/N$. Since it must have decreased between time $S - 1$ and S by minimality of S, and it starts from $\mathcal{E}(0)$ multiple of $1/N$, we must have $\mathcal{E}(S) = C$. This implies that there are exactly CN holes below the top level $\{(x,y) : y = h(S)\}$. Indeed, if $|A(0)| = n_0$ then $|A(S)| = n_0 + S$, and

$$h(S) = \mathcal{E}(S) + \frac{n_0 + S}{N} = C + \frac{n_0 + S}{N},$$

thus there are $h(S)N - (n_0 + S) = CN$ holes up to (and including) the top level $h(S)$.

Observe that each time we add a new particle to the cluster the lowest hole has probability at least m_π to become filled (independently of everything else). Hence it will take at most a Geometric number of releases of parameter m_π to fill the lowest hole. In total, then, it will take at most the sum of CN i.i.d. Geometric(m_π) to fill all the holes up to the top level of $A(S)$. Some care is needed, though: we cannot let the excess height increase too much while releasing these extra particles. To ensure this we proceed as follows.

Let $(G_i)_{i=1}^{CN}$ be a collection of i.i.d. Geometric(m_π) random variables. For any $\varepsilon > 0$, then, we can find $D = D(C, N, m_\pi, \varepsilon)$ such that

$$\mathbb{P}\bigg(\sum_{i=1}^{CN} G_i > D \bigg) < \varepsilon \,. \tag{4.7}$$

It follows that if we release D particles then we have probability at least $1 - \varepsilon$ to fill all the holes below the top levels in $A(S)$. If we fail, then the excess height has increased by at most D. If $D > C$, then, we keep on releasing particles until the excess height falls again below C. When this happens we are again in the situation of having exactly CN holes below the top level. We can therefore iterate the previous procedure by releasing D particles: with probability at least $1 - \varepsilon$ we fill all the holes, and otherwise we wait for the excess height to fall again below C and restart. After a Geometric($1 - \varepsilon$) number of attempts we will have filled all the holes below the top level, which implies that the resulting cluster will have discrepancy at most D.

The next proposition is the main result of this section, and it tells us that the discrepancy falls below D quickly.

Proposition 4.3. *Let $(A(t))_{t\geq 0}$ denote an IDLA process on the infinite graph $\mathcal{C}_G$. Assume that $A(0)$ satisfies (4.1), and that $\mathcal{E}(0) = C - 1$, with C as in Lemma 4.1. Let $\delta(t) := \delta(A(t))$ denote the discrepancy of $A(t)$. Recall the definition of D from (4.7), and assume that $\delta(0) > D$. Recall that*

$$\mathcal{T} = \inf\{t \geq 0 : \delta(t) \leq D\}.$$

Then there exist $\lambda^* = \lambda^*(\varepsilon, C, D) > 0$ and an explicit constant $c = c(\varepsilon, \lambda^*, C, D)$ such that

$$\mathbb{P}(\mathcal{T} > t) \leq c\, e^{-\lambda^* t}. \tag{4.8}$$

Proof. Define the following random times:

$$\begin{aligned}
&\tau_0 := 0,\\
&\tau_k := \inf\{t \geq \tau_{k-1} + D : \mathcal{E}(t) < C\}, \quad k \geq 1,\\
&s_k := \tau_k - (\tau_{k-1} + D), \quad k \geq 1.
\end{aligned}$$

Then, by definition, at times τ_i there are at most CN holes below the top level. As already observed, the number of releases needed to fill all these holes is stochastically dominated by the sum of CN i.i.d. Geometric(m_π), so (4.7) implies that at times $\tau_i + D$ we have filled all the holes with probability at least $1 - \varepsilon$. If this happens, then $\delta(\tau_i + D) \leq D$, and we stop. Otherwise $\mathcal{E}(\tau_i + D) \leq C + D$, and we start again. Let

$$K := \min\{k \geq 0 : \delta(\tau_k + D) \leq D\}$$

denote the number of attempts needed to succeed. Then $K \preceq$ Geometric($1 - \varepsilon$), where $\preceq$ denotes stochastic domination. Note that the times s_i's are not in general identically distributed, since s_i denotes the time it takes for the excess height to fall below C, starting from $\mathcal{E}(\tau_{i-1} + D) \leq C + D$. Since we are only interested in an upper bound for $\mathcal{T}$, though, we can make them i.i.d. simply by assuming that they always start from excess height $C+D$. More precisely, let

$$\hat{s} := \inf\{t \geq 0 : \mathcal{E}(t) < C \text{ starting from } \mathcal{E}(0) = C + D\}, \tag{4.9}$$

and let $(\hat{s}_i)_{i\geq 1}$ be a sequence of i.i.d. random variables equal in law to $\hat{s}$. Then we have that

$$\mathcal{T} \preceq KD + \sum_{i=1}^{K} \hat{s}_i.$$

We use this to estimate the moment generating function of $\mathcal{T}$. For $\lambda \in \mathbb{R}$, set $M(\lambda) := \mathbb{E}(e^{\lambda \hat{s}})$. We have

$$\begin{aligned}
\mathbb{E}(e^{\lambda \mathcal{T}}) &\leq \mathbb{E}\Big[\exp\Big\{\lambda\Big(KD + \sum_{i=1}^{K} \hat{s}_i\Big)\Big\}\Big] \\
&= \sum_{k=1}^{\infty} \mathbb{E}\Big[\exp\Big\{\lambda\Big(kD + \sum_{i=1}^{k} s_i\Big)\Big\} \mathbb{1}(K = k)\Big] \\
&\leq \frac{1-\varepsilon}{\varepsilon} \sum_{k=1}^{\infty} \Big[e^{\lambda D} M(\lambda)\varepsilon\Big]^k,
\end{aligned} \tag{4.10}$$

where the first equality above follows from the Monotone Convergence Theorem. Note that the term in the square brackets is a continuous function of λ, and it converges to $\varepsilon < 1$ as $\lambda \to 0$. It follows that there exists $\lambda^* = \lambda^*(\varepsilon, C, D) > 0$ small enough so that the above series converges. This, together with Markov's inequality, yields

$$\mathbb{P}(\mathcal{T} > t) \leq e^{-\lambda^* t} \mathbb{E}(e^{\lambda^* \mathcal{T}}).$$

The inequality (4.8) then follows by setting $c = c(\varepsilon, \lambda^*, C, D)$ to be the last term in the chain of inequalities (4.10). $\square$

4.4 Releasing the frozen particles

Let $A(0) = A'(0)$ be two matching IDLA clusters satisfying (i)-(ii) of Definition 4.2, i.e. such that

$$\delta(A(0)) \vee \delta(A'(0)) \leq D, \qquad A(0) \cap A'(0) \supseteq \{(x, y) : x \in G, y \leq h_0 + \tilde{h}\},$$

and let $F(0) \subseteq A(0)$, $F'(0) \subseteq A'(0)$ denote the sets of frozen particles, with

$$|F(0)| = |F'(0)| = d, \quad F(0) \cap F'(0) = \emptyset, \quad F(0) \cup F'(0) \subseteq \{(x, y) : x \in G, y \leq h_0\}.$$

Note that the last condition above implies that all frozen particles are at distance at least $\tilde{h}$ from the boundary of the clusters. We now want to release the frozen particles in pairs, coupled so as to maximise the probability of them to exit the clusters at the same location.

Since we have very little information about the location of the particles, we choose an arbitrary pairing. More precisely, we label the elements of $F(0), F'(0)$ by

$$\begin{aligned} F(0) &= \{(v_1, y_1), \dots, (v_d, y_d)\} \\ F'(0) &= \{(v'_1, y'_1), \dots, (v'_d, y'_d)\}. \end{aligned}$$

At each step, we release two simple random walks ω_i and ω'_i from (v_i, y_i) and (v'_i, y'_i) in $A(i-1)$ and $A'(i-1)$ respectively. The walks are coupled as follows: if $y_i > y'_i$, then ω_i stays in place until ω'_i reaches level y_i. From the first time the two walks are at the same level y_i, they take the same steps in the y coordinate, while moving independently in the first coordinate. Since the first coordinate performs a lazy walk, there is a positive chance for ω_i and ω'_i to meet before exiting the rectangle of height $h_0 + \tilde{h}$. If they meet, they stick together until they exit the clusters.[4]

[4]If the paired walks fail to meet before one of them exits the cluster, we can use the freezing coupling to keep the clusters matching. If the clusters continue to match, then if two paired frozen particles couple before exiting the rectangle of height $h_0 + \tilde{h}$, they will exit their clusters in the same location, and the cardinality of the symmetric difference between the final IDLA clusters will decrease by 2. On the other hand, this is only useful if d is very large. We choose not to discuss this case here, and ask *all* pairs of frozen particles to couple.

Let p_i denote the probability that ω_i and ω'_i couple before either ω_i exits $A(i-1)$ or ω'_i exits $A'(i-1)$. Then

$$\min_{i\le d} p_i \ge \tilde{p}, \tag{4.11}$$

for some $\tilde{p} = \tilde{p}(\tilde{h}, G)$. We therefore conclude that

$$\mathbb{P}(\text{all pairs couple}) \ge \tilde{p}^d. \tag{4.12}$$

Remark. Clearly, the choice of evolving the paired driving random walks independently until they either meet or they exit the respective clusters is not optimal. Instead, one could use more clever couplings based on the geometry of the base graph G. In Section 4.6 we will discuss the case of the base graph being the N–cycle, i.e. $G = \mathbb{Z}_N$ with symmetric driving random walks.

At the end of this procedure we are left with two matching clusters with discrepancy at most $D+d$, together with two disjoint sets of frozen particles of cardinality at most d. Note that all frozen particles are at distance at most $D+d$ from the top level of their respective clusters.

4.5 Putting all together

We can finally conclude our discussion by proving Theorem 4.1. Let $A(0), A'(0)$ denote two IDLA clusters with $|A(0)| = |A'(0)| = n_0$ and $|A(0)\Delta A'(0)| = 2d$, and let $(A(t))_{t\ge 0}$, $(A'(t))_{t\ge 0}$ be two IDLA processes on the cylinder graph $\mathcal{C}_G$ starting from $A(0), A'(0)$ respectively. Recall from the statement of Theorem 4.1 that $\mathscr{T}$ denotes the first time the two clusters match, that is $A(\mathscr{T}) = A'(\mathscr{T})$ and there are no frozen particles. We prove the theorem by explaining how we can couple together the IDLA processes so that the stopping time $\mathscr{T}$ has exponential tails. Our coupling is divided into a random number of trials, each consisting of the extra releases phase, the discrepancy decay phase and the release of the frozen particles. If at the end of a trial the clusters do not match, we declare the trial unsuccessful and proceed with the next one.

Remark. It is important to note that all our estimates for the length of each trial depend on d, and on the initial height of the clusters. At the beginning of the first trial we have $h(A(0)) \vee h(A'(0)) =: h_0$, which is prescribed by the initial configurations. Note that, on the other hand, at the end of a trail we are left with two clusters with the same number of particles, both of discrepancy at most $D + d$. We can therefore shift both clusters down in the vertical direction by the same number of rows, so that at least one of the two has an empty site at level 1. This clearly does not change the law of the process, and it has the advantage of making the initial height for all subsequent trials at most $D + d$. We perform this shifting procedure at the start of each trial without further mention.

4.5.1 The coupling

The first phase consists of a deterministic number of extra releases while using the freezing coupling. Let

$$\begin{aligned} e_1 &= (h_0 + \tilde{h} + D)N, \\ e_i &= (\tilde{h} + 2D + d)N, \quad i \geq 2. \end{aligned}$$

At the beginning of each trial i we release enough extra particles to ensure that both clusters contain e_i particles above level 0 (note that $e_1 > h_0$ and $e_i > D{+}d$ for $i \geq 1$). This guarantees that when the discrepancy of the cluster falls below D all the frozen particles are at least at distance $\tilde{h}$ from the boundary of the cluster. Clearly, this requires at most $e_i + d$ releases for each trial $i \geq 1$ (the d extra releases come from the fact that up to d releases could result in freezing). Recall now the definition of the constant $C = C(G, \eta, m_\pi)$ from Lemma 4.1, and define $\bar{D} = \bar{D}(C + d, N, m_\pi, \varepsilon)$ as in (4.7) with $C + d$ in place of C, i.e. such that

$$\mathbb{P}\bigg(\sum_{i=1}^{(C+d)N} G_i > \bar{D} \bigg) < \varepsilon \,. \tag{4.13}$$

Consistently with the notation used in the previous sections, define the following random times.

- Let S_1 be the time it takes, under IDLA dynamics, for the excess height to fall below C, starting from excess height $e_1 + d$.
- For $i \geq 2$, let S_i be the time it takes, under IDLA dynamics, for the excess height to fall below C, starting from excess height $e_i + d$.
- For $i \geq 1$, let $\bar{\mathcal{T}}_i$ be the time it takes, under coupled IDLA dynamics, for the discrepancy of both clusters to fall below D, starting from excess height $C+d$.

Note that, for each trial $k \geq 1$, after the initial extra releases the two clusters contain at most $e_k + d$ particles above level 0, and so they trivially have excess height at most $e_k + d$. We then keep on releasing extra particles while using the freezing coupling, and wait for the excess height to fall below C. This takes S_k releases for each cluster. Since at most d releases will result in a freezing step, to make sure the excess height has fallen below C we in fact release $S_k + d$ particles. Thus the final excess height is at most $C + d$ in both clusters. Note that by Lemma 4.1 we know that S_k has some exponential moments for all $k \geq 1$, and in fact by Corollary 4.1 we have

$$\begin{aligned} &\mathbb{E}(e^{\lambda S_1}) \leq \hat{c}(\lambda)\, e^{\frac{2N\lambda}{\eta}(e_1+d-C)}, \\ &\mathbb{E}(e^{\lambda S_k}) \leq \hat{c}(\lambda)\, e^{\frac{2N\lambda}{\eta}(e_k+d-C)}, \quad k \geq 2 \end{aligned} \tag{4.14}$$

as long as $\lambda < \eta^2/32N^2$, with $\hat{c}(\lambda) = 1 + \lambda\big(\frac{\eta^2}{32N^2} - \lambda\big)^{-1}$ as in Corollary 4.1.

Once the excess height is controlled, we keep on releasing particles while using the freezing coupling until the discrepancy falls below D. As explained in the previous section, we do so by trying to fill all the holes below the top level in the cluster and, in case of failure, waiting again for the excess height to fall below C, to then iterate. This takes at most $\bar{\mathcal{T}}_k$ steps. We remark that, since we are now considering two coupled clusters, we have to take into account that freezing might occur. To this end, at each stage we release extra d particles. Exactly the same proof of Proposition 4.3 with $C + d$ and $\bar{D} + d$ in place of C, D then shows that, as in (4.10), for all $k \geq 1$ it holds

$$\mathbb{E}(e^{\lambda \bar{\mathcal{T}}_k}) \leq \frac{1-\varepsilon}{\varepsilon} \sum_{j=1}^{\infty} \Big[e^{\lambda(\bar{D}+d)} M(\lambda) \varepsilon \Big]^j = (1-\varepsilon) \frac{e^{\lambda(\bar{D}+d)} M(\lambda)}{1 - e^{\lambda(\bar{D}+d)} M(\lambda)\varepsilon} \tag{4.15}$$

for all $\lambda < \lambda^*$, where $\lambda^* = \lambda^*(\varepsilon, C, \bar{D})$ has been defined in Proposition 4.3, and $M(\lambda) = \mathbb{E}(e^{\lambda \hat{s}})$ is the moment generating function of the random time $\hat{s}$ defined in (4.9).

Finally, let $\tilde{p}$ be defined as in (4.12). We remind the reader that $\tilde{p}$ is a lower bound for the probability that any two frozen particles couple upon release, uniform over all pairs of starting points. Then each trial has success probability at least $\tilde{p}^d$. It follows that the random number of trials needed for the two IDLA processes to couple is stochastically dominated by a Geometric random variable K of parameter $\tilde{p}^d$, independent of everything else. Putting all together, then, we have

$$\mathscr{T} \preceq \sum_{i=1}^{K} \big[e_i + S_i + \bar{\mathcal{T}}_i + 2d \big].$$

This shows that, for λ small enough (so that all quantities below are finite),

$$\begin{aligned}
\mathbb{E}(e^{\lambda \mathscr{T}}) &\leq \sum_{k=1}^{\infty} e^{2\lambda k d}\, \mathbb{E}\Big[\exp\Big\{ \lambda \sum_{i=1}^{k} (e_i + S_i + \bar{\mathcal{T}}_i) \Big\} \Big] (1-\tilde{p}^d)^{k-1} \tilde{p}^d \\
&= e^{\lambda(2d+e_1)} \sum_{k=1}^{\infty} e^{\lambda(k-1)(2d+e_2)}\, \mathbb{E}(e^{\lambda S_1}) \mathbb{E}(e^{\lambda S_2})^{k-1} \mathbb{E}(e^{\lambda \bar{\mathcal{T}}_1})^k (1-\tilde{p}^d)^{k-1} \tilde{p}^d \\
&= e^{\lambda(2d+e_1)} \mathbb{E}(e^{\lambda S_1}) \mathbb{E}(e^{\lambda \bar{\mathcal{T}}_1}) \tilde{p}^d \sum_{k=1}^{\infty} \Big[e^{\lambda(2d+e_2)}\, \mathbb{E}(e^{\lambda S_2}) \mathbb{E}(e^{\lambda \bar{\mathcal{T}}_1}) (1-\tilde{p}^d) \Big]^{k-1} \\
&\leq \mathscr{C}(\lambda) \sum_{k=1}^{\infty} \Big[e^{\lambda(2d+e_2)} \hat{c}(\lambda) e^{\frac{2N\lambda}{\eta}(e_2+d-C)} (1-\varepsilon) \frac{e^{\lambda(\bar{D}+d)} M(\lambda)}{1 - e^{\lambda(\bar{D}+d)} M(\lambda)\varepsilon} (1-\tilde{p}^d) \Big]^{k-1}
\end{aligned} \tag{4.16}$$

with

$$\mathscr{C}(\lambda) := e^{\lambda(2d+e_1)} \hat{c}(\lambda) e^{\frac{2N\lambda}{\eta}(e_1+d-C)} \tilde{p}^d (1-\varepsilon) \frac{e^{\lambda(\bar{D}+d)} M(\lambda)}{1 - e^{\lambda(\bar{D}+d)} M(\lambda)\varepsilon}.$$

Since the term in the square brackets in (4.16) is continuous in λ, and it converges to $1 - \tilde{p}^d < 1$ as $\lambda \searrow 0$, there exists $\bar{\lambda} > 0$ such that the above series converges for all $\lambda \leq \bar{\lambda}$. It follows that Theorem 4.1 holds with $\bar{\mathscr{C}}$ equal to the right hand side of (4.16).

4.6 Example: the N–cycle

In this section we specialise to the case $G = \mathbb{Z}_N := \mathbb{Z}/(N\mathbb{Z})$, that is we take the base graph to be a cycle on N vertices. This is the model analysed in [JLS14b], where the authors let $N \to \infty$ and look at the fluctuations of the limiting cluster around its asymptotic deterministic shape.

The aim of this section is to understand how the decay rate of $\mathscr{T}$ depends on N, the size of the base graph, as N grows. The discussion below is intended to hold for N large enough.

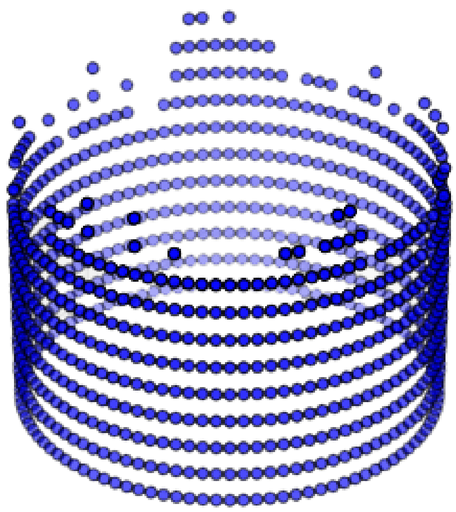

Figure 4.5: IDLA cluster on $\mathbb{Z}_N \times \mathbb{Z}$ (simulation by Tom Holding).

Notation

For sequences of non–negative real numbers $(a_N)_{N\geq 0}$ and $(b_N)_{N\geq 0}$, write $a_N = O(b_N)$ if $\lim_{N\to\infty} \frac{a_N}{b_N} < \infty$, and write $a_N \asymp b_N$ if both $a_N = O(b_N)$ and $b_N = O(a_N)$.

Let $A(0), A'(0)$ be two initial IDLA configurations. Recall that

$$h_0 := h(A(0)) \vee h(A'(0)),$$
$$|A(0)\Delta A'(0)| = 2d.$$

We stress that, while in the previous section h_0 and d where thought of as being constant, they are now allowed to depend on N, although we choose not to make this explicit in the notation.

Let us now compute the quantities that appear in (4.16), or better we are only interested in their asymptotic behaviour as $N \to \infty$.

We have $G = (V, E)$ with $V = \{1, 2, \ldots, N\}$, $E = \{(i, i+1), 1 \leq i \leq N-1\} \cup \{(1, N)\}$. Note that the graph might be periodic: to break periodicity, we assume the driving simple random walks on $\mathbb{Z}_N \times \mathbb{Z}$ to be lazy. It turns out to be convenient to make the driving walks only lazy in the horizontal coordinate: more precisely, we assume that they make a step in the horizontal direction with probability $1/4$, stay in place with probability $1/4$ and move in the vertical direction otherwise. This induces a $\left(\frac{1}{2}\right)$–lazy simple random walk on the base graph $G = \mathbb{Z}_N$ when looking at the walk only when it does not move in the vertical direction. The stationary measure of such projected walk on G is clearly the uniform one, $\pi(v) = 1/N$ for all $v \in V$, and so $m_\pi = 1/N$.

Recall from (4.4) that l_G denotes the lowest height such that when a simple random walk on $\mathbb{Z}_N \times \mathbb{Z}$ starting at level 0 reaches that height, it has probability at least $1/2N$ to be at any given vertex, i.e.

$$l_G := \inf \left\{ n \geq 1 : \min_{1 \leq i,j \leq N} \nu_i^{(n)}(j) \geq \frac{1}{2N} \right\}.$$

Lemma 4.3. *Let $G = \mathbb{Z}_N$. Then, for N large enough, $l_G \leq N \log N$.*

Proof. Let $\omega = (x(t), y(t))_{t \geq 0}$ be a simple random walk on $\mathcal{C}_G$ starting from $(0,0)$, and denote by τ_k the first hitting time of level k, i.e.

$$\tau_k := \inf\{t \geq 0 : y(t) = k\}.$$

We are interested in computing the law of $x(\tau_k)$ explicitly. To this end, recall from [LL10] that for $(x,y)\in\mathbb{Z}^2$ the potential kernel $a(x,y)$ is defined as

$$a(x,y) := \sum_{n=0}^{\infty}\Big[\mathbb{P}_{(0,0)}((x(n),y(n))=(0,0)) - \mathbb{P}_{(0,0)}((x(n),y(n))=(x,y))\Big].$$

Let $\Delta_{\mathbb{Z}^2}$ denote the discrete Laplacian on $\mathbb{Z}^2$, i.e. $\Delta_{\mathbb{Z}^2}f(x,y) := \frac{1}{4}\big[f(x+1,y)+f(x-1,y)+f(x,y+1)+f(x,y-1)\big]-f(x,y)$ for any function $f:\mathbb{Z}^2\to\mathbb{R}$. Then it is easy to check that

$$\Delta_{\mathbb{Z}^2}a(x,y) = \begin{cases} 1, & \text{if } (x,y)=(0,0), \\ 0, & \text{otherwise.}\end{cases}$$

We claim that for $i\in\{0,1,\dots,N-1\}$ and any $k\geq 1$, it holds

$$\mathbb{P}_{(0,0)}(x(\tau_k)=i) = \frac{1}{4}\sum_{j\in\mathbb{Z}}\Big[a(jN+i,k+1)-a(jN+i,k-1)\Big]. \tag{4.17}$$

Indeed, define $f(i,k) := \mathbb{P}_{(0,0)}(x(\tau_k)=i)$ for i,k as above, and extend it to the interior of the upper half plane by periodicity in the first coordinate. Further, set $f(i,0)=0$ for all $i\in\mathbb{Z}$. Then it is easily checked that $f(i,k)$ equals the r.h.s. of (4.17) for all $j\in\mathbb{Z}, k\geq 0$, as the two functions have the same discrete Laplacian and boundary values. Since $f(i,k)$ coincides with $\mathbb{P}_{(0,0)}(x(\tau_k)=i)$ for $k>1$ and $0\leq i\leq N-1$, the claim follows.

We now use the asymptotic expansion for the potential kernel

$$a(x,y) = \frac{1}{\pi}\log(x^2+y^2)+c+O\Big(\frac{1}{x^2+y^2}\Big),$$

for c absolute constant, to get, for $k,N\gg 1$,

$$a(jN+i,k+1)-a(jN+i,k-1) = \log\Big(1+\frac{4k}{(jN+i)^2+(k-1)^2}\Big)+O\Big(\frac{1}{(jN)^2+k^2}\Big).$$

Note that

$$\sum_{j\in\mathbb{Z}} O\Big(\frac{1}{(jN)^2+k^2}\Big) \le O\Big(\frac{1}{k^2}\Big) + O\Big(\frac{1}{N^2}\Big)$$

for $k, N \gg 1$. It follows that

$$\mathbb{P}_{(0,0)}(x(\tau_k)=i) = \sum_{j\in\mathbb{Z}} \log\Big(1+\frac{4k}{(jN+i)^2+(k-1)^2}\Big) + O\Big(\frac{1}{k^2}\Big) + O\Big(\frac{1}{N^2}\Big).$$

Splitting the sum in the right hand side into $jN \le k-1$ and $jN > k-1$ we then see that, as long as $k \gg N$,

$$\sum_{j\in\mathbb{Z}} \log\Big(1+\frac{4k}{(jN+i)^2+(k-1)^2}\Big) \asymp \frac{1}{k}\frac{k}{N} + \frac{2k}{N^2}\sum_{j>k/N}\frac{1}{j^2} \asymp \frac{1}{N},$$

where we have used the asymptotic estimate $\sum_{j\ge n}\frac{1}{j^2} \asymp \frac{1}{n}$. It follows that[5] for $k = N\log N$ and N large enough, it holds

$$\frac{1}{2N} \le \mathbb{P}_{(0,0)}(x(\tau_k)=i) \le \frac{3}{2N} \tag{4.18}$$

for all $0 \le i \le N-1$. Thus $l_G \le N\log N$ for N large enough. □

Fix $\eta = \varepsilon = 1/2$ (the precise value of η and ε does not play any role, as long as $0 < \eta, \varepsilon < 1$). To compute the constant $C = C(N)$ of Lemma 4.2, by (4.5) we impose

$$\Big(1-\frac{1}{2N}\Big)^{C/l_G} \le \frac{1-\eta}{N} = \frac{1}{2N},$$

which, for N large enough, is equivalent to $e^{-C/2Nl_G} \le \frac{1}{2N}$. Using that $l_G \le N\log N$, we see that it suffices to take

$$C = 4N^2\log^2 N.$$

Finally, to compute the constant $\bar{D} = \bar{D}(d,N)$ in (4.13), we take a collection

[5] In fact (4.18) holds for any $k \gg N$, but we do not try to optimise the logarithmic factor at this stage.

$(G_i)_{i=1}^{(C+d)N}$ of i.i.d. Geometric$(1/N)$ random variables, and ask that[6]

$$\mathbb{P}\Big(\sum_{i=1}^{(C+d)N} G_i > \bar{D}\Big) < \frac{1}{2}.$$

Since $\mathbb{P}\big(\sum_{i=1}^{(C+d)N} G_i > \bar{D}\big) \leq \frac{(C+d)N\mathbb{E}(G_1)}{\bar{D}} = \frac{(C+d)N^2}{\bar{D}}$, it suffices to ask that $\bar{D} > 2(C+d)N^2$, say

$$\bar{D} = 4(C+d)N^2. \tag{4.19}$$

If $d = O(N^2)$, for example, we could take $\bar{D} = 16N^4 \log^2 N$. It remains to compute $\tilde{h}$, the minimum distance of the frozen particles from the boundary of the cluster, and $\tilde{p}$, the minimum over all pairs of frozen particles of the probability that they couple before one of them exits the cluster, cf. (4.11).

Let $\omega = (x(t), y(t))_{t\geq 0}$ and $\omega' = (x'(t), y'(t))_{t\geq 0}$ denote two lazy simple random walks on $\mathbb{Z}_N \times \mathbb{Z}$, and assume that $y(0) \vee y'(0) \leq h_0$. We want to couple the walks so as to maximise the probability that they meet before one of them exits the rectangle of height $h_0 + \tilde{h}$ ($\tilde{h}$ to be determined). Assume that $y(0) > y'(0)$ without loss of generality. We couple the trajectories of the walks ω, ω' as follows. We let ω' evolve until it first reaches the level $y(0)$, while ω stays in place.

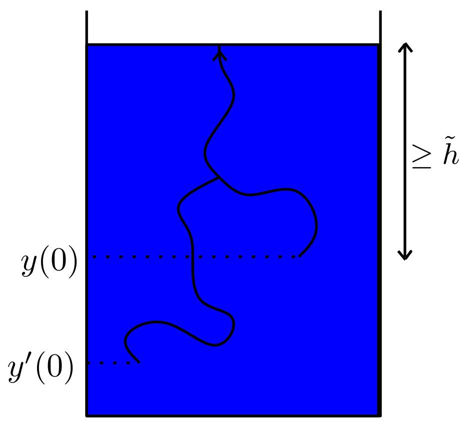

Figure 4.6: A schematic illustration of the coupling.

If

$$t_* := \inf\{t \geq 0 : y'(t) = y(0)\},$$

[6]Note that here there is room for improvement: we are only asking to each released particle to fill the lowest hole, but we could do much better.

then at time t_* the two walks have coupled in the y coordinate, so we use the trivial coupling in that coordinate to keep them matching:

$$y(t) = y'(t_* + t), \qquad t \geq 0.$$

Moreover, we couple the x coordinates by letting either x or x' move at each horizontal step, that is either x moves and x' stays in place or vice versa. Thus $x(t)-x'(t_*+t)$ is a (non–lazy) simple random walk on $\mathbb{Z}_N$. Let τ_x denote the first time at which the two walks ω, ω' match in the x coordinate, i.e.

$$\tau_x := \inf\{t \geq 0 : x(t) = x'(t_* + t)\}.$$

We denote by $\mathbb{P}_{i,j}$ the joint law of $\omega(\cdot)$ and $\omega'(t_* + \cdot)$ under the above coupling, and let $\tau^y_{\tilde{h}}$ be the first hitting time of level $y(0) + \tilde{h} \leq h_0 + \tilde{h}$, i.e.

$$\tau^y_{\tilde{h}} := \inf\{t \geq 0 : y(t) = y'(t_* + t) = y(0) + \tilde{h}\}.$$

Then, if $\hat{\omega} = (\hat{x}(t), \hat{y}(t))_{t \geq 0}$ is an auxiliary (non–lazy) simple random walk on the cylinder graph, independent of everything else, starting at $(j, 0)$ under $\hat{\mathbb{P}}_j$, and $\hat{\tau}_k$ denotes the first time $\hat{\omega}$ reaches level k, we have[7]

$$\begin{aligned}
1 - \tilde{p} &\leq \max_{1 \leq i \leq N} \hat{\mathbb{P}}_i\big(\hat{x}(t) \neq 0 \text{ for all } t \leq \hat{\tau}_{\tilde{h}}\big) \\
&\leq \max_{1 \leq i \leq N} \hat{\mathbb{P}}_i(\hat{x}(\hat{\tau}_2) \neq 0)\hat{\mathbb{P}}_{x(\hat{\tau}_2)}(\hat{x}(\hat{\tau}_4) \neq 0) \cdots \hat{\mathbb{P}}_{x(\hat{\tau}_{\tilde{h}-2})}(\hat{x}(\hat{\tau}_{\tilde{h}}) \neq 0) \\
&\leq \left(\max_{1 \leq i \leq N} \hat{\mathbb{P}}_i(\hat{x}(\hat{\tau}_2) \neq 0) \right)^{\tilde{h}/2} = \left(1 - \min_{1 \leq i \leq N} \hat{\mathbb{P}}_i(\hat{x}(\hat{\tau}_2) = 0) \right)^{\tilde{h}/2}.
\end{aligned} \tag{4.20}$$

[7] Assume that $\tilde{h}$ is even without loss of generality.

Using again the explicit formula (4.17) for the distribution of $\hat{x}(\hat{\tau}_2)$, we find[8]

$$\begin{aligned}\min_{1\le i\le N}\hat{\mathbb{P}}_i(\hat{x}(\hat{\tau}_2)=0) = \hat{\mathbb{P}}_{N/2}(\hat{x}(\hat{\tau}_2)=0) &= \sum_{j=0}^{\infty}\Big[a\Big(jN+\frac{N}{2},3\Big)-a\Big(jN+\frac{N}{2},1\Big)\Big]\\ &= \sum_{j\in\mathbb{Z}\setminus\{0\}}\log\left(1+\frac{8}{(j+\frac{1}{2})^2N^2+1}\right)+O\Big(\frac{1}{N^2}\Big)\\ &\asymp \sum_{j\in\mathbb{Z}\setminus\{0\}}\frac{1}{(j+\frac{1}{2})^2N^2+1}+O\Big(\frac{1}{N^2}\Big)\ge\frac{c}{N^2}\end{aligned}$$

for some absolute constant $c>0$ and N large enough. Finally, plugging this into (4.20) we obtain

$$1-\tilde{p}\le\left(1-\frac{c}{N^2}\right)^{\tilde{h}/2}$$

for N large enough. To have $\tilde{p}^d$ bounded away from 0 as $N\to\infty$, we ask $1-\tilde{p}\le 1/d$. To this end it suffices to choose $\tilde{h}$ such that

$$\left(1-\frac{c}{N^2}\right)^{\tilde{h}/2}\le\frac{1}{d}, \tag{4.21}$$

i.e. $\tilde{h}\ge\frac{2}{c}N^2\log d$. It follows that, for $c'=2/c$ absolute constant, we can choose

$$\tilde{h}=c'N^2\log d.$$

Note in particular that $\tilde{h}\ll\bar{D}$ as $N\to\infty$.

It only remains to plug all the above estimates in (4.16), to deduce the value of λ. Recall that (4.16) tells us that, with $\eta=\varepsilon=1/2$,

$$\mathbb{E}(e^{\lambda\mathscr{T}})\le\mathscr{C}(\lambda)\sum_{k=1}^{\infty}\left[e^{\lambda(2d+e_2)}\hat{c}(\lambda)e^{4N\lambda(e_2+d-C)}\frac{e^{\lambda(\bar{D}+d)}M(\lambda)}{2-e^{\lambda(\bar{D}+d)}M(\lambda)}(1-\tilde{p}^d)\right]^{k-1}, \tag{4.22}$$

with

$$\mathscr{C}(\lambda):=e^{\lambda(2d+e_1)}\hat{c}(\lambda)e^{4N\lambda(e_1+d-C)}\tilde{p}^d\frac{e^{\lambda(\bar{D}+d)}M(\lambda)}{2-e^{\lambda(\bar{D}+d)}M(\lambda)},$$

[8] Assume that N is even for simplicity of notation. This entails no loss of generality, since we are only concerned with the asymptotic behaviour as $N\to\infty$.

as long as λ is small enough so that $e^{\lambda(\bar{D}+d)}M(\lambda) < 2$. Above, $M(\lambda) = \mathbb{E}(e^{\lambda\hat{s}})$, for $\hat{s}$ defined as in (4.9) with $\bar{D}$ in place of D. In particular, by Corollary 4.1 with $\mathcal{E}(0) = \bar{D} + C$ we have

$$M(\lambda) \leq \hat{c}(\lambda)e^{4N\lambda\bar{D}},$$

as long as $\lambda < 1/128N^2$. Moreover, $\hat{c}(\lambda) \to 1$ as $\lambda \to 0$, as so $\hat{c}$ is bounded by an absolute constant for λ small enough. It follows that in order to have $e^{\lambda(\bar{D}+d)}M(\lambda) < 2$ it suffices to take $\lambda \ll 1/(\bar{D}N + d)$ as $N \to \infty$. This also ensures that

$$\hat{c}(\lambda)\frac{e^{\lambda(\bar{D}+d)}M(\lambda)}{2 - e^{\lambda(\bar{D}+d)}M(\lambda)} \to 0$$

as $N \to \infty$. Note further that $1 - \tilde{p}^d$ is bounded away from 1 as $N \to \infty$. Indeed, the inequality (4.21) guarantees that $\tilde{p} \geq 1 - 1/d$, from which

$$\tilde{p}^d \geq \Big(1 - \frac{1}{d}\Big)^d,$$

which converges to e^{-1} if $d \gg 1$ as $N \to \infty$, and is clearly strictly smaller than 1 if $d = O(1)$. It follows that for the series (4.22) to be convergent it suffices to take λ small enough so that

$$e^{\lambda(2d+e_2)+4N\lambda(e_2+d-C)} = O(1) \tag{4.23}$$

as $N \to \infty$. Recalling that $e_2 = (\tilde{h} + 2\bar{D} + d)N$, we see that (4.23) holds as long as

$$\lambda = O\Big(\frac{1}{N^2(\bar{D} + d)}\Big) \tag{4.24}$$

as $N \to \infty$. In particular, if $d = O(\bar{D})$ then by (4.19) we find

$$\lambda = O\Big(\frac{1}{N^2\bar{D}}\Big) = O\Big(\frac{1}{N^4(C + d)}\Big).$$

For example, if $d = O(N^2)$ then we find $\lambda = O\Big(\dfrac{1}{N^6\log^2 N}\Big)$. If, on the other hand, $d \gg \bar{D}$ as $N \to \infty$, we find

$$\lambda = O\Big(\frac{1}{N^2 d}\Big).$$

In particular, as long as d is polynomial in N, we can choose λ to be some finite inverse power of N. Finally, for λ satisfying (4.24) it is easily checked that

$$\mathscr{C}(\lambda) \leq e^{5\lambda[d+(h_0+\bar{D})N^2]} \tag{4.25}$$

for N large enough.

In conclusion, we have shown that there exists an absolute constant $c > 0$ such that, for $\bar{\lambda} = \dfrac{c}{N^2(\bar{D}+d)}$, we have

$$\mathbb{P}(\mathscr{T} > t) \leq \bar{\mathscr{C}}(\bar{\lambda}) e^{-\bar{\lambda} t},$$

where $\bar{\mathscr{C}}(\bar{\lambda})$ is the right hand side of (4.22) with $\lambda = \bar{\lambda}$. It follows that, for any $\epsilon > 0$ we can ensure that $\mathbb{P}(\mathscr{T} > t) < \epsilon$ as long as t is large enough so that $\bar{\mathscr{C}}(\bar{\lambda}) e^{-\bar{\lambda} t} < \epsilon$. This is in turn satisfied if

$$t > \frac{1}{\bar{\lambda}} \Big(\log \bar{\mathscr{C}}(\bar{\lambda}) + \log \frac{1}{\epsilon} \Big). \tag{4.26}$$

Note further that, with this choice of $\bar{\lambda}$, the term inside the square brackets in (4.22) converges to 0 as $N \to \infty$. By summing the series, then, we see that for N large enough the sum is bounded by an absolute constant $c' < \infty$. It follows that

$$\bar{\mathscr{C}}(\bar{\lambda}) \leq c' \mathscr{C}(\bar{\lambda}) \leq c' e^{5\bar{\lambda}[d+(h_0+\bar{D})N^2]}$$

by (4.25). This shows that (4.26) holds as long as

$$t > 10\big[d + (h_0 + \bar{D})N^2\big] + \frac{1}{\bar{\lambda}} \log \frac{1}{\epsilon}\,,$$

which is polynomial in N as long as both d and h_0 are, and ϵ is at most exponential in N.

4.7 Further developments

The techniques we used in this chapter to bound the forget time on IDLA on cylinder graphs are far from optimal. It would be very interesting to optimise the bound in Theorem 4.1, even just for the particular case $G = \mathbb{Z}_N$ treated in the above section. For this N-cycle graph we believe the optimal bound to be close to N^2, although our techniques cannot get down to this order without substantial modifications. A possible approach to improving our results could be to try to estimate the decay of the discrepancy directly, without going through the excess height.

Another interesting question is to identify which notions of mixing can be controlled by our methods. More precisely, note that the shifted IDLA chain lives on an infinite state space, and we can make the forget time arbitrarily large simply by picking suitably bad initial configurations (think about $A(0), A'(0)$ consisting of two disjoint, arbitrarily high columns). It follows that there is no hope of obtaining an upper bound for the forget time of the chain which is uniform in all starting configurations. On the other hand, some configurations are more likely then others, and it is natural to take this into account when choosing the starting configurations $A(0), A'(0)$. This is the idea behind the definition of the *reset time*, due to Lovász and Winkler [LW98], which is the minimum expected matching time for two coupled shifted IDLA processes starting from two independent samples of the stationary distribution of the chain. Estimating the reset time of the shifted IDLA process is work in progress [LS], and it requires to show that the stationary distribution gives high probability to low height configurations (or, at least, that starting from stationarity we get to a low height configuration quickly).

Finally, as mentioned in the introduction of the thesis, one could try to use our mixing result together with the maximal fluctuation result by Jerison, Levine and Sheffield [JLS14b] to show that the stationary distribution of shifted IDLA on the cylinder graph $\mathbb{Z}_N \times \mathbb{Z}$ gives high probability to configurations with height $O(\log N)$, for N large enough. This is also work in progress.

Appendix A

The local fluctuations result

A.1 Local fluctuations

In this appendix we prove the following local fluctuations theorem for HL(0) clusters.

Theorem A.1. *Pick any $t > 0$, and let $z = e^{ia+\sigma}$, $w = e^{\sigma}$ for some $a \in [-\pi, \pi)$, $\sigma > 0$. Define $\varepsilon = \delta^{2/3}\log(1/\delta)$ as in Theorem 2.3, and assume that $\sigma \to 0$ as $c \to 0$, with $\lim_{\varepsilon \to 0} \frac{\sqrt{\varepsilon}}{\sigma} = 0$. Then, as $n \to \infty$, $nc \to t$ and $\sigma \to 0$, it holds:*

$$\left(\frac{\log \frac{\Phi_n(z)}{z} - nc}{\sqrt{c \log(\frac{1}{2\sigma})}}, \frac{\log \frac{\Phi_n(w)}{w} - nc}{\sqrt{c \log(\frac{1}{2\sigma})}} \right) \longrightarrow (\mathcal{N}_1, \mathcal{N}_2)$$

in distribution, where $(\mathcal{N}_1, \mathcal{N}_2)$ is a random vector with centred complex Gaussian entries, and covariance structure given by

$$\mathbb{E}(\mathcal{N}_1 \mathcal{N}_2) = \begin{pmatrix} \frac{1}{1+\alpha^2} & -\frac{\alpha}{1+\alpha^2} \\ \\ \frac{\alpha}{1+\alpha^2} & \frac{1}{1+\alpha^2} \end{pmatrix},$$

for $\alpha = \lim_{\sigma \to 0} \frac{a}{2\sigma} \in [0, \infty]$, with the convention that $\frac{1}{1+\alpha^2} = \frac{\alpha}{1+\alpha^2} = 0$ when $\alpha = \infty$.

Theorem A.1 follows by the same arguments that lead to Theorem 3.1, considering now $\frac{\mathcal{X}^{\sigma}_{k,n}(\cdot)}{\sqrt{\log(\frac{1}{2\sigma})}}$ in place of $\mathcal{X}^{\sigma}_{k,n}(\cdot)$. Under the assumption $\sigma \gg \sqrt{\varepsilon}$ Lemmas

2.3-2.4 still apply, so that Theorem 2.2 holds. Together with Proposition 2.2, this shows that in the limit as $\sigma \to 0$, $n \to \infty$ and $nc \to t$, real and imaginary part of $\frac{\log \frac{\Phi_n(e^{ia+\sigma})}{e^{ia+\sigma}} - nc}{\sqrt{c \log(\frac{1}{2\sigma})}}$ are asymptotically i.i.d. centred Gaussians, with limiting variance given by

$$\lim_{\sigma\to 0} \frac{1}{\log\left(\frac{1}{2\sigma}\right)}\left[\frac{1}{2\pi}\int_{\sigma}^{\sigma+t}\int_{-\pi}^{\pi}\left[\mathrm{Re}\left(\frac{e^{-i\vartheta+x}+1}{e^{-i\vartheta+x}-1}\right)\right]^2 \mathrm{d}\vartheta \mathrm{d}x - t\right] = \lim_{\sigma\to 0}\frac{1}{\log\left(\frac{1}{2\sigma}\right)}\log\left|\frac{1-e^{-2(\sigma+t)}}{1-e^{-2\sigma}}\right| = 1\,.$$

For two point correlation we reason as in Section 3.1.1, to gather that the limiting covariance between $\mathrm{Re}\left(\frac{\log \frac{\Phi_n(e^{ia+\sigma})}{e^{ia+\sigma}} - nc}{\sqrt{c \log(\frac{1}{2\sigma})}}\right)$ and $\mathrm{Re}\left(\frac{\log \frac{\Phi_n(e^{\sigma})}{e^{\sigma}} - nc}{\sqrt{c \log(\frac{1}{2\sigma})}}\right)$ is given by

$$\begin{aligned}\lim_{\sigma\to 0} \frac{1}{\log\left(\frac{1}{2\sigma}\right)}&\left[\frac{1}{2\pi}\int_{\sigma}^{\sigma+t}\int_{-\pi}^{\pi}\mathrm{Re}\left(\frac{e^{-i(\vartheta-a)+x}+1}{e^{-i(\vartheta-a)+x}-1}\right)\mathrm{Re}\left(\frac{e^{-i\vartheta+x}+1}{e^{-i\vartheta+x}-1}\right)\mathrm{d}\vartheta \mathrm{d}x - t\right] = \\ &= \lim_{\sigma\to 0}\frac{1}{\log\left(\frac{1}{2\sigma}\right)}\log\left|\frac{1-e^{-2(\sigma+t)+ia}}{1-e^{-2\sigma+ia}}\right| = \lim_{\sigma\to 0}\frac{\log(1+e^{-4\sigma}-2e^{-2\sigma}\cos a)}{2\log 2\sigma} = \frac{1}{1+\alpha^2}\end{aligned}$$

whenever $a/2\sigma \to \alpha \in [0,\infty]$, where the last equality is obtained by Taylor expanding around $a=0, \sigma=0$. The same holds for imaginary parts correlation.

Finally, the asymptotic covariance between $\mathrm{Re}\left(\frac{\log \frac{\Phi_n(e^{ia+\sigma})}{e^{ia+\sigma}} - nc}{\sqrt{c \log(\frac{1}{2\sigma})}}\right)$ and $\mathrm{Im}\left(\frac{\log \frac{\Phi_n(e^{\sigma})}{e^{\sigma}} - nc}{\sqrt{c \log(\frac{1}{2\sigma})}}\right)$ is given by

$$\begin{aligned}\lim_{\sigma\to 0} \frac{1}{\log\left(\frac{1}{2\sigma}\right)}&\left[\frac{1}{2\pi}\int_{\sigma}^{\sigma+t}\int_{-\pi}^{\pi}\mathrm{Re}\left(\frac{e^{-i(\vartheta-a)+x}+1}{e^{-i(\vartheta-a)+x}-1}\right)\mathrm{Im}\left(\frac{e^{-i\vartheta+x}+1}{e^{-i\vartheta+x}-1}\right)\mathrm{d}\vartheta \mathrm{d}x - t\right] = \\ &= \lim_{\sigma\to 0}\frac{1}{\log\left(\frac{1}{2\sigma}\right)}\mathrm{Arg}\left(\frac{1-e^{-2(\sigma+t)+ia}}{1-e^{-2\sigma+ia}}\right) = \lim_{\sigma\to 0}\frac{1}{\log 2\sigma}\arctan\left(\frac{\sin a}{\cos a - e^{2\sigma}}\right) = \frac{\alpha}{1+\alpha^2}\end{aligned}$$

whenever $a/2\sigma \to \alpha \in [0,\infty]$, where the last equality is obtained by Taylor expanding around $a=0, \sigma=0$. This concludes the proof of Theorem A.1.

A.2 Proof of Lemma 3.2

Fix any $k \le \lfloor ns \rfloor$ and set for simplicity

$$Z_1 = Z_{k,\lfloor nt\rfloor}(a), \quad Z_2 = Z_{k,\lfloor ns\rfloor}(0), \quad W_1 = e^{ia+\sigma+(\lfloor nt\rfloor-k)c}, \quad W_2 = e^{\sigma+(\lfloor ns\rfloor-k)c}.$$

Moreover, introduce the functions $g_\vartheta(z) = \mathrm{Re}\big(\log \frac{F(e^{-i\vartheta}z)}{e^{-i\vartheta}z}\big)$, $h_\vartheta(z) = c\,\mathrm{Re}\big(\frac{e^{-i\vartheta}z+1}{e^{-i\vartheta}z-1}\big)$. Then

$$\mathbb{E}(X_{k,\lfloor nt\rfloor}(a)X_{k,\lfloor ns\rfloor}(0)|\mathscr{F}_{k+1,\lfloor nt\rfloor}) = \frac{1}{2\pi c}\int_{-\pi}^{\pi} g_\vartheta(Z_1)g_\vartheta(Z_2)\mathrm{d}\vartheta - c\,,$$

and we have to show that

$$\left|\frac{1}{2\pi c}\int_{-\pi}^{\pi} g_\vartheta(Z_1)g_\vartheta(Z_2)\mathrm{d}\vartheta - \frac{1}{2\pi c}\int_{-\pi}^{\pi} h_\vartheta(W_1)h_\vartheta(W_2)\mathrm{d}\vartheta\right| \leq \frac{C(t)c\varepsilon}{(\sigma+(\lfloor ns\rfloor-k)c)^3}\,. \tag{A.1}$$

Trivially, the l.h.s. is bounded above by

$$\frac{1}{2\pi c}\int_{-\pi}^{\pi}\Big(|g_\vartheta(Z_2)|\cdot|g_\vartheta(Z_1)-h_\vartheta(W_1)| + |h_\vartheta(W_1)|\cdot|g_\vartheta(Z_2)-h_\vartheta(W_2)|\Big)\mathrm{d}\vartheta\,.$$

We bound each term separately. Recall the definition of $E(m,\varepsilon)$, from which it follows that

$$\begin{aligned}\max_{\vartheta\in[-\pi,\pi)}\Big\{|Z_{k,\lfloor nt\rfloor}(\vartheta)|\vee|Z_{k,\lfloor ns\rfloor}(\vartheta)|\Big\} &\leq e^{\sigma+(\lfloor nt\rfloor-k)c}(1+2\varepsilon)\leq 2e^{t+2}=C(t)\,,\\ \min_{\vartheta\in[-\pi,\pi)}\Big\{|Z_{k,\lfloor nt\rfloor}(\vartheta)|\wedge|Z_{k,\lfloor ns\rfloor}(\vartheta)|-1\Big\} &\geq e^{\sigma+(\lfloor ns\rfloor-k)c}(1-2\varepsilon)\geq\frac{\sigma+(\lfloor ns\rfloor-k)c}{2}\end{aligned}$$

as long as n is large enough and $\sigma \gg \varepsilon$. We combine the above estimates with the bounds in Corollary 2.2, to get

$$\begin{aligned}|g_\vartheta(Z_2)| &\leq \left|\log\frac{F(e^{-i\vartheta}Z_2)}{e^{-i\vartheta}Z_2} - c\frac{e^{-i\vartheta}Z_2+1}{e^{-i\vartheta}Z_2-1}\right| + c\left|\frac{e^{-i\vartheta}Z_2+1}{e^{-i\vartheta}Z_2-1}\right| \leq \frac{Cc^{3/2}|Z_2|^2}{(|Z_2|-1)^3}+\frac{2c|Z_2|}{|Z_2|-1}\\ &\leq \frac{C(t)c^{3/2}}{(\sigma+(\lfloor ns\rfloor-k)c)^3}+\frac{C(t)c}{\sigma+(\lfloor ns\rfloor-k)c}\leq\frac{2C(t)c}{\sigma+(\lfloor ns\rfloor-k)c}\end{aligned} \tag{A.2}$$

for n large enough. Similarly, we find

$$|h_\vartheta(W_1)| \leq c\bigg(1+\frac{2}{|W_1|-1}\bigg) \leq c\bigg(1+\frac{2}{\sigma+(\lfloor ns\rfloor-k)c}\bigg) \leq \frac{C(t)c}{\sigma+(\lfloor ns\rfloor-k)c} \tag{A.3}$$

on the event $E(m,\varepsilon)$ for, say, $C(t)=t+3$ and n large enough.

In order to bound the remaining terms, observe that, by definition, on $E(m,\varepsilon)$ we have

$$\max_{\vartheta\in[-\pi,\pi)}\left\{\left|Z_{k,\lfloor ns\rfloor}(\vartheta)-e^{i\vartheta+\sigma+(\lfloor ns\rfloor-k)c}\right|\vee\left|Z_{k,\lfloor nt\rfloor}(\vartheta)-e^{i\vartheta+\sigma+(\lfloor nt\rfloor-k)c}\right|\right\}\le C(t)\,\varepsilon\,,$$

from which we get $|Z_1-W_1|\le C(t)\varepsilon$ and $|Z_2-W_2|\le C(t)\varepsilon$. Combining this with Corollary 2.2 we finally obtain

$$\begin{aligned}|g(Z_1)-h(W_1)|&\le|g(Z_1)-h(Z_1)|+|h(Z_1)-h(W_1)|\\&\le\frac{Cc^{3/2}|Z_1|^2}{(\sigma+(\lfloor ns\rfloor-k)c)^3}+\frac{2c|Z_1-W_1|}{(|Z_1|-1)(|W_1|-1)}\\&\le\frac{C(t)c^{3/2}}{(\sigma+(\lfloor ns\rfloor-k)c)^3}+\frac{C(t)c\varepsilon}{(\sigma+(\lfloor ns\rfloor-k)c)^2}\\&\le\frac{2C(t)c\varepsilon}{(\sigma+(\lfloor ns\rfloor-k)c)^2}\end{aligned}$$

for n large enough. Similarly one shows that the same bound holds for $|g(Z_2)-h(W_2)|$. Putting now this together with (A.2) and (A.3) gives (A.1).

The second statement of the lemma follows by the same arguments, and the proof is omitted.

Appendix B

McLeish Central Limit Theorem

In this section we prove of Theorem 2.2. For convenience of the reader, and to fix an easier notation, we recall the statement here.

Theorem B.1. *Let $X_{k,n}$, $k = 1 \dots n$ be a backwards martingale difference array with respect to $\mathcal{F}_{k,n} := \sigma(X_{k,n} \dots X_{n,n})$. Let $S_{k,n} := \sum_{j=k}^{n} X_{j,n}$. Assume that:*

(i) for all $\varepsilon > 0$, $\displaystyle\sum_{k=1}^{n} X_{k,n}^2 \mathbb{1}(|X_{k,n}| > \varepsilon) \to 0$ in probability as $n \to \infty$,

(ii) $\displaystyle\sum_{j=1}^{n} X_{j,n}^2 \to \sigma^2$ in probability as $n \to \infty$, for some $\sigma > 0$.

Then $S_{n,n} \to \mathcal{N}(0, \sigma^2)$ in distribution as $n \to \infty$.

The proof of the above result follows the one of [McL74] with minor modifications due to the fact that we are dealing with backwards martingales instead of forward ones. We start with a simple preliminary result.

Lemma B.1 ([McL74])**.** *Consider two collections of integrable random variables $(T_n)_{n\in\mathbb{N}}$ and $(U_n)_{n\in\mathbb{N}}$ satisfying the following:*

(I) $U_n \to a$ in probability as $n \to \infty$,

(II) (T_n), $(T_n U_n)$ are Uniformly Integrable (UI) sequences,

(III) $\mathbb{E}(T_n) \to b$ as $n \to \infty$,

for some real constants a, b. Then $\mathbb{E}(T_n U_n) \to ab$ as $n \to \infty$.

Proof. It holds $\mathbb{E}(T_n U_n) = \mathbb{E}(T_n(U_n - a)) + a\mathbb{E}(T_n)$. The last term converges to ab by the third assumption, so it is enough to show that $\mathbb{E}(T_n(U_n - a)) \to 0$ as $n \to \infty$. In fact, we can show that $T_n(U_n - a) \to 0$ in probability, which (by uniform integrability of the sequence) implies convergence in L^1, which in turn implies convergence in expectation. To this end, note that by uniform integrability for any $\eta > 0$ we can find a constant $K = K(\eta) > 0$ such that $\mathbb{P}(|T_n(U_n - a)| > K) < \eta$ for n large enough. Hence for any $\varepsilon > 0$ we have

$$\begin{aligned}\mathbb{P}(|T_n(U_n - a)| > \varepsilon) &= \mathbb{P}(|T_n(U_n - a)| > \varepsilon, |T_n| > K) + \mathbb{P}(|T_n(U_n - a)| > \varepsilon, |T_n| \leq K)\\ &\leq \mathbb{P}(|T_n| > K) + \mathbb{P}(|U_n - a| > \varepsilon/K) < 2\eta\,,\end{aligned}$$

where the last inequality follows from the uniform integrability of (T_n), and by choosing n large enough thanks to the convergence in probability $U_n \to a$. □

We can now prove the above theorem.

Proof of Theorem B.1. The proof consists in showing that the characteristic function of $S_{n,n}$, that with a slight abuse of notation we call S_n here, converges to the one of a $\mathcal{N}(0, \sigma^2)$ random variable. Note that, as a consequence of the dependence of the summands $(X_{k,n})_{k \leq n}$, the characteristic function of S_n does not split in the product of n characteristic functions. On the other hand, we can make use of the Martingale property at the price of Taylor expanding the complex exponential. Recall the identity[1]

$$e^{ix} = (1 + ix)\exp\big(-\frac{x^2}{2} + r(x)\big), \quad \text{with } |r(x)| \leq |x|^3 \quad \forall x : |x| < 1\,. \tag{B.1}$$

For t real we have:

$$e^{itS_n} = \prod_{k=1}^{n} e^{itX_{k,n}} = \prod_{k=1}^{n}(1 + itX_{k,n})\exp\Bigg(-\frac{t^2}{2}\sum_{k=1}^{n} X_{k,n}^2 + \sum_{k=1}^{n} r(tX_{k,n})\Bigg)\,.$$

[1]This identity can be seen by Taylor expanding $\log(1 + z)$ around 0.

To prove convergence of the expectation of the above quantity we would like to apply Lemma B.1, but we need the integrability assumptions to be verified. To this end we introduce a cut off. Let

$$Z_{k,n} := X_{k,n} \cdot \mathbb{1}\bigg(\sum_{j=k+1}^{n} X_{j,n}^2 \leq \sigma^2 + 1 \bigg).$$

Define $J_n := \max\big\{1 \leq k \leq n : \sum_{j=k}^{n} X_{j,n}^2 > \sigma^2 + 1\big\}$, with the convention that the $\max \emptyset = 0$. Note that $Z_{k,n} = 0$ if $k < J_n$. Let $\tilde{S}_n := \sum_{k=1}^{n} Z_{k,n}$.

Claim B.1. It holds $S_n - \tilde{S}_n \to 0$ in probability as $n \to \infty$.

Indeed, for any $\varepsilon > 0$ we have

$$\mathbb{P}(|S_n - \tilde{S}_n| > \varepsilon) = \mathbb{P}(J_n \geq 1) = \mathbb{P}\bigg(\sum_{j=1}^{n} X_{j,n}^2 > \sigma^2 + 1 \bigg) \to 0$$

as $n \to \infty$.

In light of this claim and of the continuity of the function $x \mapsto e^{ix}$ it is enough to prove the theorem for $\tilde{S}_n$ in place of S_n, that is to show that $\mathbb{E}(e^{it\tilde{S}_n}) \to e^{-t^2\sigma^2/2}$ as $n \to \infty$. We have:

$$\mathbb{E}\big(e^{it\tilde{S}_n}\big) = \mathbb{E}\Bigg[\bigg(\prod_{k=1}^{n}(1 + itZ_{k,n})\bigg) \cdot \exp\bigg(-\frac{t^2}{2} \sum_{k=1}^{n} Z_{k,n}^2 + \sum_{k=1}^{n} r(tZ_{k,n}) \bigg)\Bigg] = \mathbb{E}(T_n U_n) \tag{B.2}$$

with

$$T_n = \bigg(\prod_{k=1}^{n}(1 + itZ_{k,n}) \bigg), \qquad U_n = \exp\bigg(-\frac{t^2}{2} \sum_{k=1}^{n} Z_{k,n}^2 + \sum_{k=1}^{n} r(tZ_{k,n}) \bigg).$$

It only remains to verify Assumptions (I)-(II)-(III) of Lemma B.1 to conclude the proof.

We first prove (I), i.e. that $U_n \to \sigma^2$ in probability as $n \to \infty$, for σ^2 defined as in the statement of the theorem. Pick any $\eta > 0$, then:

$$\begin{aligned}\mathbb{P}\bigg(\Big|\sum_{k=1}^{n} X_{k,n}^2 - \sum_{k=1}^{n} Z_{k,n}^2\Big| > \eta\bigg) &= \mathbb{P}\bigg(\sum_{k=1}^{n} X_{k,n}^2 \mathbb{1}\bigg(\sum_{j=k+1}^{n} X_{j,n}^2 > \sigma^2 + 1\bigg) > \eta\bigg) \\ &\leq \mathbb{P}\bigg(\sum_{j=1}^{n} X_{j,n}^2 > \sigma^2 + 1\bigg) \to 0 \qquad \text{as } n \to \infty\,.\end{aligned}$$

We have shown that $\sum_{k=1}^{n} Z_{k,n}^2 \to \sigma^2$ in probability. Hence

$$\exp\bigg(-\frac{t^2}{2}\sum_{k=1}^{n} Z_{k,n}^2\bigg) \to \exp\bigg(-\frac{t^2\sigma^2}{2}\bigg) \tag{B.3}$$

in probability, by the continuity of the exponential. Now recall that

$$U_n = \exp\bigg(-\frac{t^2}{2}\sum_{k=1}^{n} Z_{k,n}^2 + \sum_{k=1}^{n} r(tZ_{k,n})\bigg)\,.$$

In light of (B.3), it is enough to show that $\exp\big(\sum_{k=1}^{n} r(tZ_{k,n})\big) \to 1$ in probability. We have:

$$\bigg|\sum_{k=1}^{n} r(tZ_{k,n})\bigg| \leq \sum_{k=1}^{n} |t|^3 \cdot \big|Z_{k,n}\big|^3 \leq |t|^3 \sum_{k=1}^{n} \big|X_{k,n}\big|^3 \leq |t|^3 \bigg(\max_{j=1\ldots n} |X_{j,n}|\bigg) \sum_{k=1}^{n} X_{k,n}^2\,,$$

where the first inequality follows from (B.1). Note that for any $\eta > 0$

$$\Big\{\max_{j=1\ldots n} |X_{j,n}| > \eta\Big\} = \Big\{\sum_{k=1}^{n} X_{k,n}^2 \mathbb{1}(|X_{k,n}| > \eta) > \eta^2\Big\}\,,$$

and so assumption (i) implies that $\max_k |X_{k,n}| \to 0$ in probability as $n \to \infty$. Using also assumption (ii), then, we conclude that the far r.h.s. of the above chain of inequalities converges to 0 in probability as $n \to \infty$, which is what we aimed to show.

We now prove (II). To see that (T_nU_n) is a UI sequence it suffices to observe that the sequence is uniformly bounded: $|T_n U_n| = |e^{it\tilde{S}_n}| = 1$ for any t real. Moreover

$$\begin{aligned}|T_n| &= \prod_{k=1}^{n} |1 + itZ_{k,n}| = \prod_{k=J_n}^{n} |1 + itX_{k,n}| \leq \prod_{k=J_n}^{n} \exp\left(\frac{t^2}{2} X_{k,n}^2\right) \\ &= \exp\left(\frac{t^2}{2} \sum_{k=J_n}^{n} X_{k,n}^2\right) \leq \exp\left(\frac{t^2}{2}(\sigma^2 + 1)\right).\end{aligned}$$

It follows that also (T_n) is a uniformly bounded sequence, which is therefore UI.

It remains to prove that (III) holds. This follows from the fact that $Z_{k,n}$ is itself a sequence of backwards martingale differences.

$$\begin{aligned}\mathbb{E}(T_n) &= \mathbb{E}\bigg(\prod_{k=1}^{n}(1 + itZ_{k,n})\bigg) = \mathbb{E}\bigg((1 + itZ_{1,n})\prod_{k=2}^{n}(1 + itZ_{k,n})\bigg) \\ &= \mathbb{E}\bigg[\bigg(\prod_{k=2}^{n}(1 + itZ_{k,n})\bigg)\mathbb{E}\big(1 + itZ_{1,n}|\mathcal{F}_{2,n}\big)\bigg] = \mathbb{E}\bigg(\prod_{k=2}^{n}(1 + itZ_{k,n})\bigg)\end{aligned}$$

which, iterating, gives $\mathbb{E}(T_n) = 1$ for all n. It follows trivially that $\mathbb{E}(T_n) \to 1$ as $n \to \infty$.

We have shown that the sequences (T_n), (U_n) defined in (B.2) satisfy the assumptions of Lemma B.1 with $a = \sigma^2$ and $b = 1$. Therefore we conclude that

$$\mathbb{E}(e^{it\tilde{S}_n}) = \mathbb{E}(T_n U_n) \to \sigma^2$$

as $n \to \infty$, which is the thesis. $\square$

Appendix C

Weak convergence of Banach–space valued stochastic processes

We collect here some results concerning weak convergence of càdlàg stochastic processes taking values in a Banach space. Our discussion is based on [EK09, Igl68], to which we refer the reader for the details.

Let $(E, \|\cdot\|)$ be a complete separable metric space (i.e. Polish space), and denote by $D_E[0,\infty)$ the space of càdlàg functions $x : [0,\infty) \to E$. Let Λ' denote the set of increasing continuous functions $\lambda : [0,\infty) \to [0,\infty)$ such that $\lambda(0) = 0$, $\lambda(t) \to \infty$ as $t \to \infty$. Finally, write $\Lambda \subset \Lambda'$ for the subset of Lipschitz $\lambda \in \Lambda'$ such that $\gamma(\lambda) := \sup_{0 \le s < t} \left| \log \frac{\lambda(t) - \lambda(s)}{t - s} \right| < \infty$.

Definition C.1 (Skorohod metric on $D_E[0,\infty)$)**.** For $x, y \in D_E[0,\infty)$, set

$$\mathbb{d}(x,y) := \inf_{\lambda \in \Lambda} \left[\gamma(\lambda) \vee \int_0^\infty e^{-u} d(x,y,\lambda,u) \mathrm{d}u \right],$$

where $d(x,y,\lambda,u) := \sup_{t \ge 0} \|x(t \wedge u) - y(\lambda(t) \wedge u)\| \wedge 1$.

The following theorem justifies the introduction of the above metric.

Theorem C.1 ([EK09], Theorem 5.6, p.121)**.** *If $(E, \|\cdot\|)$ is a complete and separable* *metric space, then so is* $(D_E[0,\infty), \mathbb{d})$*.*

Let $\mathcal{E}$ denote the Borel σ-algebra associated to the metric $\mathbb{d}$ on $D_E[0,\infty)$. We seek for sufficient conditions for weak convergence of Borel probability measures on $(D_E[0,\infty), \mathcal{E})$. To this end, recall the definition of finite dimensional distributions (FDDs).

Definition C.2 (FDDs)**.** For any $k \in \mathbb{N}_+$ pick $0 \le t_1 < \cdots < t_k$, and let $\pi_{t_1 \dots t_k} : D_E[0,\infty) \to E^k$ denote the projection map

$$\pi_{t_1 \dots t_k}(x) = \big(x(t_1), \dots, x(t_k)\big)\,.$$

Then each $\pi_{t_1 \dots t_k}$ induces a probability measure on E^k equipped with the product σ-algebra $\mathcal{E}^k$. For every $x \in D_E[0,\infty)$, the probability measures $\{\pi_{t_1 \dots t_k}(x) : k \in \mathbb{N}_+,\, 0 \le t_1 < \cdots < t_k\}$ are called finite dimensional distributions (FDDs) of x.

Fact C.1 ([EK09], Proposition 7.1, p.127)**.** A probability measure on $D_E[0,\infty)$ is uniquely determined by the totality of its FDDs.

Proposition C.1 ([EK09], Lemma 4.3, p.112)**.** *Let* $(x_n)_{n\ge 0}$ *be a sequence of random variables on* $(D_E[0,\infty), \mathcal{E})$*, and denote by* $(\mu_n)_{n\ge 0}$ *the associated probability measures. Assume that:*

(i) the family $(\mu_n)_{n\ge 0}$ *is relatively compact, and*

(ii) the FDDs of x_n *converge weakly to the ones of* x *as* $n \to \infty$*.*

Then $x_n \to x$ *weakly as* $n \to \infty$*.*

In the rest of this section we discuss sufficient conditions for (i) and (ii) to hold.

C.1 Tightness of probability measures on $D_E[0,\infty)$

Pre-compactness of family of probability measures is often difficult to check. The concept of tightness reduces it to checking compactness of certain sets, the advantage being that, on a complete separable metric space, these are completely characterized by the Ascoli-Arzelà theorem. Recall that a family $(\mu_n)_{n\geq 0}$ of probability measures on $D_E[0,\infty)$ is said to be tight if and only if for any $\varepsilon > 0$ there exists a compact set $K_\varepsilon \subset D_E[0,\infty)$ such that $\inf_{n\geq 0}\mu_n(K_\varepsilon) > 1-\varepsilon$.

Theorem C.2 (Prohorov's theorem, [EK09], Theorem 2.2, p.104)**.** *A family of Borel probability measures on $(D_E[0,\infty),\mathcal{E})$ is pre-compact if and only if it is tight.*

We conclude that if the FDDs of μ_n converge to the ones of μ, and $(\mu_n)_{n\geq 0}$ is tight, then $\mu_n \to \mu$ weakly on $D_E[0,\infty)$. In order to state a criterion for tightness we need the following definition.

Definition C.3 (Modified modulus of continuity)**.** For any $\eta, T > 0$ denote by $\Delta_{\eta,T}$ the set of all partitions (t_i) of the form $0 = t_0 < t_1 < \cdots < t_{k-1} < T \leq t_k$ for some $k \in \mathbb{N}_+$, such that $\min_{i\leq k}(t_i - t_{i-1}) > \eta$. Then, given $x \in D_E[0,\infty)$, the modified modulus of continuity of x is defined as

$$\omega'(x,\eta,T) := \sup_{(t_i)\in\Delta_{\eta,T}} \max_i \sup_{s,t\in[t_{i-1},t_i)} \|x(t)-x(s)\|\,.$$

Remark. Note that, if $\omega(x,\eta,T) := \sup_{\substack{s,t\in[0,T]\\ |t-s|<\eta}} \|x(t)-x(s)\|$ denotes the usual modulus of continuity, then it holds

$$\omega'(x,\eta,T) \leq \omega(x,2\eta,T) \tag{C.1}$$

for all $x \in D_E[0,\infty)$, $\eta, T > 0$.

The following result follows from the Ascoli-Arzelà characterization of compact sets in $D_E[0,\infty)$, and provides a useful criterion for tightness.

Theorem C.3 ([EK09], Corollary 7.4, p.129)**.** *Let $(x_n)_{n\geq 0}$ be a sequence of random variables taking values in $D_E[0,\infty)$, and denote by $(\mu_n)_{n\geq 0}$ the associated Borel probability measures. Then $(\mu_n)_{n\geq 0}$ is tight if and only if the following hold:*

(i) for all $\nu > 0, t \in [0,\infty) \cap \mathbb{Q}$*, there exists* $\Gamma_{\nu,t} \subset E$ *compact such that*

$$\liminf_{n\to\infty} \mathbb{P}(x_n(t) \in \Gamma_{\nu,t}) \geq 1 - \nu\,,$$

(ii) for all $\nu, T > 0$ *there exists* $\eta > 0$ *such that*

$$\limsup_{n\to\infty} \mathbb{P}(\omega'(x_n, \eta, T) \geq \nu) \leq \nu\,.$$

C.2 Convergence of probability measures on product spaces

Once tightness has been checked, in light of Proposition C.1 in order to show weak convergence one has to show convergence of FDDs. These are probability measures on E^k equipped with the product σ-algebra $\mathcal{E}^k$. In this section we discuss some criteria for weak convergence of such measures in the case $E = C(\mathbb{T})$ and $\|\cdot\| = \|\cdot\|_\infty$ (the supremum norm on $\mathbb{T}$). Our presentation is based on [Igl68].

Fix any $k \in \mathbb{N}_+$. Since $(C(\mathbb{T}), \|\cdot\|_\infty)$ is a complete separable metric space, it is well known that so is $C(\mathbb{T})^k$ with the metric

$$\rho(\underline{x}, \underline{y}) = \max_{i=1\ldots k} \|x_i - y_i\|_\infty\,,$$

for any $\underline{x} = (x_1 \ldots x_k)$, $\underline{y} = (y_1 \ldots y_k) \in C(\mathbb{T})^k$. Moreover, the Borel σ-algebra on $(C(\mathbb{T})^k, \rho)$ coincides with the product σ-algebra. It then follows by Prohorov's Theorem that a family of probability measures $(\mu_n)_{n\geq 0}$ on this space is relatively compact if and only if it is tight. Moreover, if $(\mu_n)_{n\geq 0}$ is tight, and the FDDs converge to the ones of a probability measure μ on the same space, then $\mu_n \to \mu$ weakly.

Proposition C.2 ([Igl68], Lemma 1)**.** *A family of probability measures* $(\mu_n)_{n\geq 0}$ *on* $C(\mathbb{T})^k$ *is tight if and only if each family of marginals* $(\mu_n^i)_{n\geq 0}$ *is tight on* $C(\mathbb{T})$*.*

Definition C.4 (FDDs of a product measure)**.** For any $l \in \mathbb{N}_+$ pick $0 \leq \vartheta_1 < \cdots < \vartheta_l$, and let $\pi_{\vartheta_1 \ldots \vartheta_k} : C(\mathbb{T})^k \to \mathbb{C}^{k \times l}$ denote the projection map

$$\pi_{\vartheta_1 \ldots \vartheta_l}(\underline{x}) = \begin{pmatrix} x_1(e^{i\vartheta_1}) & \ldots & x_1(e^{i\vartheta_l}) \\ \vdots & & \vdots \\ x_k(e^{i\vartheta_1}) & \ldots & x_k(e^{i\vartheta_l}) \end{pmatrix}.$$

Then each $\pi_{\vartheta_1 \ldots \vartheta_l}$ induces a probability measure on $\mathbb{C}^{k \times l}$ equipped with the usual Borel σ-algebra. For every $\underline{x} \in C(\mathbb{T})^k$, the probability measures $\{\pi_{\vartheta_1 \ldots \vartheta_l}(\underline{x}) : l \in \mathbb{N}_+, 0 \leq \vartheta_1 < \cdots < \vartheta_k < 2\pi\}$ are called finite dimensional distributions (FDDs)[1] of $\underline{x}$.

Proposition C.3 ([Igl68], Lemma 3)**.** *Let μ, μ' be probability measures on $C(\mathbb{T})^k$, and assume that the FDDs of μ and μ' coincide. Then $\mu = \mu'$.*

Combining the above result with Proposition C.2 one concludes that the following holds.

Theorem C.4 ([Igl68], Theorem 4)**.** *Let $(\mu_n)_{n \geq 0}$, μ be Borel probability measures on $C(\mathbb{T})^k$. Then $\mu_n \to \mu$ weakly if and only if:*

(i) the FDDs of μ_n converge to the ones of μ, and

(ii) for all $i = 1 \ldots k$, the family of marginal measures $(\mu_n^i)_{n \geq 0}$ on $C(\mathbb{T})$ is tight.

[1]Note the slight abuse of notation: this definition of FDDs is different from Definition C.2, since it concerns probability measures on a different space. Whenever we will use the term FDDs, it will be clear from the context whether we refer to Definition C.2 or C.4.

Bibliography

[ACSS16] Amine Asselah, Emilio N.M. Cirillo, Benedetto Scoppola, and Elisabetta Scoppola. On diffusion limited deposition. *Electronic Journal of Probability*, 21, 2016. (Cited on page 7).

[AG10] Amine Asselah and Alexandre Gaudillière. A note on fluctuations for internal diffusion limited aggregation. *arXiv preprint arXiv:1004.4665*, 2010. (Cited on pages 6 and 24).

[AG13a] Amine Asselah and Alexandre Gaudillière. From logarithmic to subdiffusive polynomial fluctuations for internal DLA and related growth models. *The Annals of Probability*, 41(3A):1115–1159, 2013. (Cited on pages 6 and 24).

[AG13b] Amine Asselah and Alexandre Gaudillière. Sublogarithmic fluctuations for internal DLA. *The Annals of Probability*, 41(3A):1160–1179, 2013. (Cited on pages 6 and 24).

[AG14] Amine Asselah and Alexandre Gaudillière. Lower bounds on fluctuations for internal DLA. *Probability Theory and Related Fields*, 158(1-2):39–53, 2014. (Cited on pages 6 and 24).

[AHD15] Antonio Auffinger, Jack Hanson, and Michael Damron. 50 years of first passage percolation. *arXiv preprint arXiv:1511.03262*, 2015. (Cited on page 7).

[Arr79] Richard Arratia. *Coalescing Brownian motions on the line.* University of Wisconsin–Madison, 1979. (Cited on page 16).

[Arr81] Richard Arratia. Coalescing Brownian motions and the voter model on Z. *Unpublished partial manuscript (circa 1981), available from rarratia@math.usc.edu*, 1981. (Cited on page 16).

[BBRT85] Robin C. Ball, Robert M. Brady, Giuseppe Rossi, and Bernard R. Thompson. Anisotropy and Cluster Growth by Diffusion–Limited aggregation. *Physical review letters*, 55(13):1406, 1985. (Cited on page 7).

[BF09] Rolf Busam and Eberhard Freitag. *Complex Analysis*. Springer Berlin Heidelberg, 2009. (Cited on page 43).

[Bil95] Patrick Billingsley. *Probability and measure*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, Inc., New York, third edition, 1995. A Wiley-Interscience Publication. (Cited on pages 43 and 44).

[Bil99] Patrick Billingsley. *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons, Inc., New York, second edition, 1999. A Wiley-Interscience Publication. (Cited on pages 60 and 69).

[BKP12] Noam Berger, Jacob J. Kagan, and Eviatar B. Procaccia. Stretched IDLA. *arXiv preprint arXiv:1209.3112*, 2012. (Cited on page 7).

[BPP97] Martin T. Barlow, Robin Pemantle, and Edwin A. Perkins. Diffusion–limited aggregation on a tree. *Probability Theory and Related Fields*, 107(1):1–60, 1997. (Cited on page 7).

[BY08] Itai Benjamini and Ariel Yadin. Diffusion Limited Aggregation on a Cylinder. *Communications in Mathematical Physics*, 279(1):187–223, 2008. (Cited on page 7).

[CD81] J. Theodore Cox and Richard Durrett. Some limit theorems for percolation processes with necessary and sufficient conditions. *The Annals of Probability*, pages 583–603, 1981. (Cited on pages 4 and 7).

[CM01] Lennart Carleson and Nikolai Makarov. Aggregation in the plane and Loewner's equation. *Communications in Mathematical Physics*, 216(3):583–607, 2001. (Cited on page 7).

[DF91] Persi Diaconis and William Fulton. A growth model, a game, an algebra, Lagrange inversion, and characteristic classes. *Rend. Sem. Mat. Univ. Pol. Torino*, 49(1):95–119, 1991. (Cited on pages 6, 88, and 89).

[DFJ68] S.W. Dharmadhikari, V. Fabian, and K. Jogdeo. Bounds on the moments of martingales. *The Annals of Mathematical Statistics*, pages 1719–1723, 1968. (Cited on page 62).

[DPZ14] Giuseppe Da Prato and Jerzy Zabczyk. *Stochastic equations in infinite dimensions*, volume 152 of *Encyclopedia of Mathematics and its Applications.* Cambridge University Press, Cambridge, second edition, 2014. (Cited on pages 70, 72, and 73).

[DRSV14] Bertrand Duplantier, Rémi Rhodes, Scott Sheffield, and Vincent Vargas. Log-correlated Gaussian fields: an overview. *arXiv:1407.5605*, 2014. (Cited on page 20).

[Dur96] Rick Durrett. *Stochastic calculus: a practical introduction.* Probability and Stochastics Series. CRC Press, Boca Raton, FL, 1996. (Cited on page 61).

[Dur10] Rick Durrett. *Probability: theory and examples.* Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, fourth edition, 2010. (Cited on page 44).

[Ede61] Murray Eden. A two-dimensional growth process. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. IV*, pages 223–239. Univ. California Press, Berkeley, Calif., 1961. (Cited on page 4).

[EK09] Stewart N. Ethier and Thomas G. Kurtz. *Markov processes: characterization and convergence*, volume 282. John Wiley & Sons, 2009. (Cited on pages 55, 60, 131, 132, and 133).

[EW99] Dorothea M. Eberz-Wagner. Discrete growth models. *arXiv preprint math/9908030*, 1999. (Cited on page 5).

[FINR04] L. R. G. Fontes, M. Isopi, C. M. Newman, and K. Ravishankar. The Brownian web: characterization and convergence. *The Annals of Probability*, 32(4):2857–2883, 2004. (Cited on page 16).

[Gro67] Leonard Gross. Abstract Wiener spaces. In *Proc. Fifth Berkeley Sympos. Math. Statist. and Probability (Berkeley, Calif., 1965/66), Vol. II: Contri-*

butions to Probability Theory, Part 1, pages 31–42. Univ. California Press, Berkeley, Calif., 1967. (Cited on page 71).

[Hal00] Thomas C. Halsey. Diffusion-limited aggregation: a model for pattern formation. *Physics Today*, 53(11):36–41, 2000. (Cited on page 5).

[HL98] Matthew B. Hastings and Leonid S. Levitov. Laplacian growth as one-dimensional turbulence. *Physica D: Nonlinear Phenomena*, 116(1):244–252, 1998. (Cited on pages 7 and 10).

[HW65] John M. Hammersley and D.J.A. Welsh. First-Passage Percolation, Subadditive Processes, Stochastic Networks, and Generalized Renewal Theory. In *Bernoulli 1713 Bayes 1763 Laplace 1813*, pages 61–110. Springer, 1965. (Cited on page 7).

[Igl68] Donald L. Iglehart. Weak convergence of probability measures on product spaces with applications to sums of random vectors. 1968. (Cited on pages 55, 131, 134, and 135).

[JLS12] David Jerison, Lionel Levine, and Scott Sheffield. Logarithmic fluctuations for internal DLA. *J. Amer. Math. Soc.*, 25(1):271–301, 2012. (Cited on pages 6 and 24).

[JLS13] David Jerison, Lionel Levine, and Scott Sheffield. Internal DLA in higher dimensions. *Electron. J. Probab*, 18(98):1–14, 2013. (Cited on pages 6 and 24).

[JLS14a] David Jerison, Lionel Levine, and Scott Sheffield. Internal DLA and the Gaussian free field. *Duke Math. J.*, 163(2):267–308, 2014. (Cited on pages 6, 24, 79, and 85).

[JLS14b] David Jerison, Lionel Levine, and Scott Sheffield. Internal DLA for cylinders. *Advances in Analysis: The Legacy of Elias M. Stein*, page 189, 2014. (Cited on pages 6, 24, 25, 26, 28, 29, 30, 110, and 119).

[JVST15] Fredrik Johansson Viklund, Alan Sola, and Amanda Turner. Small-particle limits in a regularized Laplacian random growth model. *Comm. Math. Phys.*, 334(1):331–366, 2015. (Cited on pages 10 and 15).

[Kes86] Harry Kesten. Aspects of first passage percolation. In *École d'Été de Probabilités de Saint Flour XIV-1984*, pages 125–264. Springer, 1986. (Cited on page 7).

[Kes87] Harry Kesten. How long are the arms in DLA? *Journal of Physics A: Mathematical and General*, 20(1):L29, 1987. (Cited on page 5).

[Kes90] Harry Kesten. Upper bounds for the growth rate of DLA. *Physica A: Statistical Mechanics and its Applications*, 168(1):529–535, 1990. (Cited on page 5).

[Kes93] Harry Kesten. On the speed of convergence in first-passage percolation. *The Annals of Applied Probability*, pages 296–338, 1993. (Cited on page 4).

[Kes03] Harry Kesten. First-passage percolation. In *From classical to modern probability*, pages 93–143. Springer, 2003. (Cited on page 7).

[Law95] Gregory F. Lawler. Subdiffusive fluctuations for internal diffusion limited aggregation. *Ann. Probab.*, 23(1):71–86, 1995. (Cited on page 23).

[LBG92] Gregory F. Lawler, Maury Bramson, and David Griffeath. Internal diffusion limited aggregation. *Ann. Probab.*, 20(4):2117–2140, 1992. (Cited on pages 6 and 23).

[LL10] Gregory F. Lawler and Vlada Limic. *Random walk: a modern introduction*, volume 123. Cambridge University Press, 2010. (Cited on page 112).

[LS] Lionel Levine and Vittoria Silvestri. The reset time of Internal DLA on a cylinder. *In preparation.* (Cited on pages 32 and 119).

[LSSW14] Asad Lodhia, Scott Sheffield, Xin Sun, and Samuel S Watson. Fractional Gaussian fields: a survey. *arXiv:1407.5598*, 2014. (Cited on page 20).

[LW98] László Lovász and Peter Winkler. Reversal of markov chains and the forget time. *Combinatorics, Probability and Computing*, 7(02):189–204, 1998. (Cited on page 119).

[Mar04] James B. Martin. Limiting shape for directed percolation models. *Annals of probability*, pages 2908–2937, 2004. (Cited on page 7).

[Mar14] Sébastien Martineau. Directed Diffusion-Limited Aggregation. *arXiv preprint arXiv:1411.3667*, 2014. (Cited on page 7).

[MBRS87] Paul Meakin, Robin C. Ball, P. Ramanlal, and Leonard Michael Sander. Structure of large two-dimensional square-lattice diffusion-limited aggregates: Approach to asymptotic behavior. *Physical Review A*, 35(12):5233, 1987. (Cited on page 7).

[McL74] D. L. McLeish. Dependent central limit theorems and invariance principles. *Ann. Probability*, 2:620–628, 1974. (Cited on pages 18, 44, and 125).

[MD86] Paul Meakin and John M. Deutch. The formation of surfaces by diffusion limited annihilation. *The Journal of chemical physics*, 85(4):2320–2325, 1986. (Cited on page 23).

[MS13] Jason Miller and Scott Sheffield. Quantum Loewner Evolution. *arXiv preprint arXiv:1312.5745*, 2013. (Cited on page 5).

[NPW84] L. Niemeyer, L. Pietronero, and H. J. Wiesmann. Fractal dimension of dielectric breakdown. *Phys. Rev. Lett.*, 52(12):1033–1036, 1984. (Cited on page 6).

[NT12] James Norris and Amanda Turner. Hastings-Levitov aggregation in the small-particle limit. *Comm. Math. Phys.*, 316(3):809–841, 2012. (Cited on pages 11, 12, 13, 16, 18, 38, 39, 40, 45, 60, 77, and 79).

[NT15] James Norris and Amanda Turner. Weak convergence of the localized disturbance flow to the coalescing Brownian flow. *The Annals of Probability*, 43(3):935–970, 2015. (Cited on pages 11 and 16).

[Ric73] Daniel Richardson. Random growth in a tessellation. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 74, pages 515–528. Cambridge Univ Press, 1973. (Cited on page 4).

[Roh11] Steffen Rohde. Oded Schramm: from circle packing to SLE. In *Selected Works of Oded Schramm*, pages 3–45. Springer, 2011. (Cited on page 7).

[RZ05] Steffen Rohde and Michel Zinsmeister. Some remarks on Laplacian growth. *Topology Appl.*, 152(1-2):26–43, 2005. (Cited on pages 10 and 11).

[She07] Scott Sheffield. Gaussian free fields for mathematicians. *Probab. Theory Related Fields*, 139(3-4):521–541, 2007. (Cited on pages 70 and 71).

[Sil15] Vittoria Silvestri. Fluctuation results for Hastings–Levitov planar growth. *Probability Theory and Related Fields*, pages 1–44, 2015. (Cited on page 32).

[SS16] Vladas Sidoravicius and Alexandre Stauffer. Multi-Particle Diffusion Limited Aggregation. *arXiv preprint arXiv:1603.03218*, 2016. (Cited on page 5).

[TW98] Bálint Tóth and Wendelin Werner. The true self-repelling motion. *Probability Theory and Related Fields*, 111(3):375–452, 1998. (Cited on page 16).

[Whi70] Ward Whitt. Weak convergence of probability measures on the function space C[0,∞). *The Annals of Mathematical Statistics*, pages 939–944, 1970. (Cited on page 55).

[WJS81] T.A. Witten Jr and Leonard M. Sander. Diffusion-Limited Aggregation, a Kinetic Critical Phenomenon. *Physical review letters*, 47(19):1400, 1981. (Cited on page 4).

[WJS83] T.A. Witten Jr and Leonard M. Sander. Diffusion-Limited Aggregation. *Physical Review B*, 27(9):5686, 1983. (Cited on page 4).