

Internal DLA on cylinder graphs: fluctuations and mixing

Vittoria Silvestri*

Abstract

We use coupling ideas introduced in [13] to show that an IDLA process on a cylinder graph $G \times \mathbb{Z}$ forgets a typical initial profile in $\mathcal{O}(N\sqrt{\tau_N}(\log N)^2)$ steps for large N , where N is the size of the base graph G , and τ_N is the total variation mixing time of a simple random walk on G . The main new ingredient is a maximal fluctuations bound for IDLA on $G \times \mathbb{Z}$ which only relies on the mixing properties of the base graph G and the Abelian property.

Keywords: internal DLA; mixing; fluctuations.

MSC2020 subject classifications: 60J05.

Submitted to ECP on April 2, 2020, final version accepted on July 21, 2020.

1 Introduction

Internal Diffusion Limited Aggregation (IDLA) is a mathematical model for corrosion phenomena, introduced by Meakin and Deutch [14] and independently by Diaconis and Fulton [6]. It models the growth of a cluster of particles as being governed by the harmonic measure on its boundary seen from an internal point. More precisely, let $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ denote an infinite, locally finite graph with a marked vertex $Z_0 \in \mathbb{V}$, and assume that a simple random walk on \mathbb{G} starting from Z_0 exits every finite set containing Z_0 in finite time almost surely. We define an increasing family $(A(t))_{t \geq 0}$ of subsets of \mathbb{V} as follows. Set $A(0) = \{Z_0\}$, and recursively define

$$A(t) = A(t-1) \cup \{Z_t\},$$

where Z_t denotes the exit location from $A(t-1)$ of a simple random walk on \mathbb{G} starting from Z_0 , independent of everything else. Then $(A(t))_{t \geq 0}$ is a Markov chain on the space of connected subsets of \mathbb{G} , whose long-term behaviour is of interest. The case $\mathbb{G} = \mathbb{Z}^d$, $Z_0 = 0$ is by now well understood: as $t \rightarrow \infty$, $A(t)$ has been shown to converge to a Euclidean ball [10] with logarithmic fluctuations away from it in dimension $d \geq 2$ [2, 3, 8, 9].

Let $G = (V, E)$ be a finite, connected, non-oriented graph on N vertices. We take $\mathbb{G} = G \times \mathbb{Z}$, that is $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ with

$$\mathbb{V} = \{(v, y) : v \in G, y \in \mathbb{Z}\}$$

$$\mathbb{E} = \{((v, y), (v', y')) : v \sim v' \text{ and } y = y', \text{ or } v = v' \text{ and } |y - y'| = 1\},$$

where \sim denotes the adjacency relation in G . We refer to the first and second coordinate as the horizontal and vertical coordinate respectively. For $y \in \mathbb{Z}$ we call the set $G \times \{y\}$ the y^{th} level, while $R_y = G \times (-\infty, y]$ will denote the rectangle of height y .

*University of Rome La Sapienza, Italy. E-mail: silvestri@mat.uniroma1.it

Definition 1 (Simple random walk on \mathbb{G}). A simple random walk on \mathbb{G} is, for us, a discrete-time Markov chain which at each step moves either in the horizontal or in the vertical coordinate with probability $1/2$. A vertical move consists in stepping from (v, y) to $(v, y \pm 1)$ with equal probability, while a horizontal move is a lazy simple random walk step on G , i.e. $(v, y) \rightarrow (v, y)$ with probability $1/2$ and $(v, y) \rightarrow (w, y)$ with probability $1/(2 \deg v)$ if w is a neighbour of v in G .

Let π_N denote the stationary distribution of a simple random walk on G , that is

$$\pi_N(v) = \frac{\deg v}{2|E|}, \quad v \in V$$

where $\deg v$ is the degree of v in G . We make the following assumption throughout.

Assumption 1. There exist finite constants $\delta, \delta' > 0$, independent of N , such that

$$\frac{\delta}{N} \leq \pi_N(v) \leq \frac{\delta'}{N}, \quad \forall v \in V, \quad (1.1)$$

for all $N \geq 1$.

Graphs satisfying the above assumption are called *quasi-regular*. Example include the complete graph, expander graphs, the d -dimensional torus, and any graph with vertices of comparable degree. In particular, graphs in this class can have a wide range of mixing behaviour for large N (see the comments after the statement of Theorem 2).

We remark that Assumption 1 is not needed for our arguments to work: we choose to make it as it identifies the class of base graphs for which our approach gives best results.

Notation All constants in this note are allowed to depend on δ, δ' , and we suppress this dependence from the notation.

Definition 2 (IDLA on \mathbb{G}). Let $A(0)$ be any connected subset of \mathbb{V} containing the rectangle $R_0 = \{(v, y) : y \leq 0\}$ and finitely many sites above level 0. At each discrete time $t \geq 1$, a simple random walk is released from level 0 at a random location sampled from π_N , and its exit location Z_t from $A(t-1)$ is added to the cluster by setting $A(t) = A(t-1) \cup \{Z_t\}$.

Note that this is equivalent to adding a site to the cluster according to the harmonic measure on the cluster's external boundary seen from level $-\infty$. When $A(0) = R_0$ we say that the process starts from flat.

It is common to describe the IDLA dynamics as an interacting particle system. Particles are released one by one from level 0 according to π_N , and perform simple random walks until they hit an unexplored location, where they settle. The cluster is then the set of sites containing settled particles.

Let

$$\Omega = \{A \subset \mathbb{V} : A = R_0 \cup F \text{ for some finite } F\}$$

denote the state space of IDLA on \mathbb{G} , and note that the IDLA dynamics defines a transient Markov chain on Ω . The following shift procedure makes it recurrent.

Definition 3 (Shifted IDLA). Let $\mathcal{S} : \Omega \rightarrow \Omega$ be defined as follows. For a given cluster $A \in \Omega$, let

$$k_A := \max\{k \geq 0 : R_k \subseteq A\} \quad (1.2)$$

denote the height of the maximal filled rectangle contained in A . The *downshift* of A is the cluster

$$\mathcal{S}(A) := \{(v, y - k_A) : (v, y) \in A\}.$$

Note that $k_{\mathcal{S}(A)} = 0$, so $\mathcal{S}(\mathcal{S}(A)) = \mathcal{S}(A)$. If $(A(t))_{t \geq 0}$ is an IDLA process in Ω , we set $A^*(t) = \mathcal{S}(A(t))$ for all $t \geq 0$, and refer to $(A^*(t))_{t \geq 0}$ as the *shifted IDLA* process associated to $(A(t))_{t \geq 0}$.

Shifted IDLA is itself a Markov chain¹ on the countable state space Ω . It is easy to see that this chain is irreducible, since all states communicate with the flat configuration R_0 . As in [13], it can be shown that shifted IDLA is positive recurrent, so it has a stationary distribution μ_N (cf. Remark 1 below). If a random cluster A is distributed according to μ_N , we say that A is a *stationary cluster*.

The aim of this note is to address the following questions:

- What do stationary clusters look like?
- How long does shifted IDLA take to forget a stationary initial state?

The first question concerns fluctuation bounds for IDLA clusters, while the second one looks as the mixing properties of shifted IDLA as a Markov chain. Both questions were answered in [13] for the cycle base $G = \mathbb{Z}_N$, which is by now well understood. Here we generalise the analysis to quasi-regular base graphs, for which far less is known. In particular, while [13] makes use of optimal fluctuation bounds which are specific to the case $G = \mathbb{Z}_N$, here we observe that near-optimal fluctuation bounds are not needed to obtain a near-optimal mixing result (see comments after the statement of Theorem 2).

In order to state our results, recall that for two probability measures μ, ν on a countable state space Ω , their total variation (TV) distance is defined as

$$\|\mu - \nu\|_{TV} = \sup_{A \subseteq \Omega} |\mu(A) - \nu(A)|.$$

Definition 4 (TV mixing time). Let $\omega = (\omega(k))_{k \geq 0}$, $\omega' = (\omega'(k))_{k \geq 0}$ denote two simple random walks on G starting from $\omega(0) = v$, $\omega'(0) = v'$. Then, if P_v^k and $P_{v'}^k$ denote the distribution of $\omega(k)$ and $\omega'(k)$ respectively, we define the ε -TV mixing time by

$$\tau_N(\varepsilon) = \inf\{k \geq 0 : \max_{v, v' \in V} \|P_v^k - P_{v'}^k\|_{TV} \leq \varepsilon\},$$

and write τ_N in place of $\tau_N(1/2)$ for brevity.

For $A \in \Omega$, let $|A|$ denote the number of sites in A above level 0, i.e. the cardinality of the finite set $A \cap \{y > 0\}$. We denote by

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

the height of A , so that $h(A) \geq 0$ for all $A \in \Omega$. Our first result bounds the height of stationary clusters in terms of τ_N , and it implies that the stationary distribution μ_N concentrates on a small subset of the state space Ω .

Theorem 1. Assume (1.1). Then for any $\gamma > 0$, there exists a constant c_γ , depending only on γ, δ, δ' , such that

$$\mu_N\left(\{A \in \Omega : h(A) > c_\gamma \sqrt{\tau_N} (\log N)^2\}\right) \leq N^{-\gamma}, \tag{1.3}$$

for N large enough.

¹To see this, it is convenient to modify the definition of the IDLA process $(A(t))_{t \geq 0}$ by releasing the t^{th} particle from level $k_{A(t-1)}$ rather than from level 0, with $k_{A(t-1)}$ defined as in (1.2). This clearly does not change the distribution of the process.

A similar result holds for any *lukewarm start* ν_N for shifted IDLA (cf. [13], Section 1), we omit the details. Note that the faster the mixing of the base graph, the smaller the height bound provided by Theorem 1. We remark that the proof of Theorem 1 only relies on the mixing properties of the base graph G and the Abelian property of IDLA (cf. Section 2), but does not require any Green's function estimate for the underlying random walks. We have not tried to optimise the exponent of the logarithmic factor.

We use the above control on stationary clusters to bound the time it takes for shifted IDLA to forget any such initial state.

Theorem 2. Assume (1.1). Then for any $\gamma > 0$ there exist a constant d_γ and a set $\Omega_\gamma \subseteq \Omega$, depending only on γ, δ, δ' , such that

$$\mu_N(\Omega_\gamma) \geq 1 - N^{-\gamma}$$

for N large enough, and such that the following holds. For any $A_0, A'_0 \in \Omega_\gamma$ with $|A_0| = |A'_0|$, writing $t_\gamma = d_\gamma N \sqrt{\tau_N} (\log N)^2$ for brevity, we have

$$\|P(t_\gamma) - P'(t_\gamma)\|_{TV} \leq N^{-\gamma} \quad (1.4)$$

where $(A(t))_{t \geq 0}$ and $(A'(t))_{t \geq 0}$ are IDLA processes starting from A_0 and A'_0 respectively, and $P(t)$ and $P'(t)$ denote the laws of $A(t)$ and $A'(t)$.

Thus IDLA forgets any stationary initial state in $\mathcal{O}(N \sqrt{\tau_N} (\log N)^2)$ steps, with high probability. The reader should compare this result with Theorem 1.3 of [13] for $G = \mathbb{Z}_N$, stating that (1.4) holds with $t_\gamma = d_\gamma N^2 \log N$. Although the stationary height bound (1.3) in Theorem 1 is far from optimal in this case, we remark that the conclusion of Theorem 2 only differs from the near-optimal result in [13] by a logarithmic factor. We conjecture that this is the case for all quasi-regular graphs.

Conjecture 1. Under Assumption 1, t_γ in Theorem 2 is optimal up to logarithmic factors. More explicitly, for any $\gamma, \varepsilon > 0$ there exist subsets $\Omega_\gamma, \Omega'_\gamma$ of Ω such that

$$\mu_N(\Omega_\gamma) \geq \frac{1}{2} - \gamma \quad \mu_N(\Omega'_\gamma) \geq \frac{1}{2} - \gamma,$$

and a constant $\alpha = \alpha(\gamma, \varepsilon, \delta, \delta') > 0$ such that if $A_0 \in \Omega_\gamma, A'_0 \in \Omega'_\gamma$ and $t < \alpha N \sqrt{\tau_N}$, then

$$\|P(t) - P'(t)\|_{TV} \geq 1 - \varepsilon$$

for N large enough.

In fact, as will become clear from the proof of Theorem 2, our methods give a bound on the time it takes for an IDLA process to forget *any* initial state. More explicitly, an IDLA process starting from an arbitrary profile $A_0 \in \Omega$ with height h_0 forgets its initial profile in $\mathcal{O}(Nh_0 + N \sqrt{\tau_N} (\log N)^2)$ steps, with high probability.

Structure of the paper

The remainder of the paper is organised as follows. In Section 2 we introduce the Abelian property of IDLA, and explain how it relates to couplings of IDLA processes. In Section 3 we collect two preliminary results on the mixing properties of a simple random walk on \mathbb{G} . In Section 4 we recall the coupling ideas introduced in [13], and adapt them to bound the maximal fluctuations for IDLA on \mathbb{G} for polynomial times (cf. Theorem 4). This is the core of the article. Finally, in Section 5 we briefly explain how to deduce Theorems 1 and 2 along the lines of [13], leaving the details to the reader.

2 On couplings of Abelian processes

In this section we introduce the Abelian property of Internal DLA, and comment on the role it plays in the construction of a valid coupling of two IDLA processes.

2.1 The Abelian property

There are several equivalent ways to state the Abelian property of IDLA: in this note we opt for the *stack of instructions* point of view [6, 16, 12]. Let $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ be a given graph (in our case $\mathbb{G} = G \times \mathbb{Z}$). To each vertex $x \in \mathbb{V}$ we associate an infinite stack of instructions

$$\xi_x = (\xi_x(k))_{k \geq 1},$$

with $\xi_x(k)$ neighbour of x in \mathbb{G} for all $k \geq 1$. For the moment we think of the stack of instructions $\xi = (\xi_x)_{x \in \mathbb{V}}$ as fixed, so there is no randomness in the model.

Given a configuration $\eta : \mathbb{V} \rightarrow \mathbb{N}$ of particles, with $\eta(x)$ counting the number of particles at x , we say that η is unstable at x if $\eta(x) \geq 2$. η is unstable if there exists $x \in \mathbb{V}$ such that η is unstable at x . We stabilise η as follows. At each discrete time step $t \geq 1$, we pick an unstable vertex $x \in \mathbb{V}$, and move one particle from x to an adjacent vertex according to the first unused instruction at x . More explicitly, if k_t denotes the number of visits to x up to and including time t , we move one particle from x to $\xi_x(k_t)$. We refer to this operation as *toppling* the vertex x . We keep toppling unstable vertices until η becomes stable, that is $\eta(x) \leq 1$ for all $x \in \mathbb{V}$.

A given sequence of topplings $\alpha = (x_1, x_2 \dots x_n)$ is said to be *legal* for a configuration η if for any $2 \leq k \leq n$ the vertex x_k is unstable for the configuration obtained by toppling η at $x_1, x_2 \dots x_{k-1}$. If a legal sequence α stabilises η , we write $\mathcal{S}_\alpha(\eta)$ for the stabilization of η according to α .

Theorem 3 (Abelian property). If $\alpha = (x_1, x_2 \dots x_n)$ and $\beta = (y_1, y_2 \dots y_n)$ are two finite, legal stabilising sequences for η , then they are permutations one of the other, and $\mathcal{S}_\alpha(\eta) = \mathcal{S}_\beta(\eta)$.

In particular, this tells us that the stabilization of a given configuration η according to a given stack of instructions $\xi = (\xi_x)_{x \in \mathbb{V}}$ does not depend on the sequence of topplings. This is particularly useful when randomness is introduced in the story.

2.2 Couplings of IDLA processes

With IDLA in mind, we randomize the stack of instructions $\xi = (\xi_x)_{x \in \mathbb{V}}$ as follows. Take $(\xi_x)_{x \in \mathbb{V}}$ to be i.i.d., each consisting of i.i.d. instructions. For example, for $\mathbb{G} = G \times \mathbb{Z}$ we could generate a random stack of instructions by letting a simple random walk on \mathbb{G} (as described in the introduction) run forever, and recording its jumps.

In this setting, the Abelian property tells us that if $\xi = (\xi_x)_{x \in \mathbb{V}}$ and $\xi' = (\xi'_x)_{x \in \mathbb{V}}$ are two identically distributed random stacks, then stabilizing a given configuration η with respect to ξ and ξ' will result in identically distributed stable configurations. This suggests a way of coupling two IDLA processes on \mathbb{G} : it suffices to couple the corresponding stacks of instructions.

While this is a natural approach, we often want to describe the coupling of IDLA processes in terms of couplings of the building random walks. This requires some care: allowing one random walker to look into the future of the other might break the Abelian property, since a walker might reveal instructions that the other one did not use. A simple way to overcome this problem is to force the random walks coupling to be co-adapted (see [4, 5] and references therein), which guarantees that different walkers in the same IDLA process use disjoint sets of instructions. On the other hand, requiring co-adaptedness could, a priori, substantially increase the coupling time [1]. To avoid this,

in Section 3 we build an explicit coupling that is not co-adapted, but with the property that different walkers in the same IDLA process use disjoint sets of instructions. Our construction follows Pitman [15] (see also [7] and references therein), and it achieves a coupling time that is within a constant factor from the optimal one.

3 Preliminaries

We collect here two preliminary results which were proved in [13] for $G = \mathbb{Z}_N$. To start with, recall that for any $\varepsilon > 0$ we have $\tau_N(\varepsilon) \leq \lceil \log_2(1/\varepsilon) \rceil \tau_N$ ([11], Chapter 4), which implies that

$$\tau_N(N^{-\gamma}) \leq 3\gamma\tau_N \log N \tag{3.1}$$

for any $\gamma > 0$. Let $\omega = (\omega(t))_{t \geq 0}$ be a simple random walk on \mathbb{G} , as described in the introduction, and let

$$\tau_n^{\mathbb{Z}} := \inf\{t \geq 0 : \omega(t) \in G \times \{n\}\}$$

denote the first time it reaches level n . For $(v, y) \in \mathbb{V}$ we denote by $\mathbb{P}_{(v,y)}$ the distribution of a simple random walk on \mathbb{G} starting from (v, y) .

Proposition 1. For any $\gamma > 0$, it holds

$$\max_{v \in \mathbb{V}} \mathbb{P}_{(v,0)}(\tau_n^{\mathbb{Z}} < 3\gamma\tau_N \log N) \leq N^{-\gamma}$$

for all $n \geq 10\gamma\sqrt{\tau_N} \log N$.

This can be proved exactly as in [13], Lemma 3.1, replacing N^2 by τ_N . It tells us that, by the time a simple random walk on \mathbb{G} has travelled $\sqrt{\tau_N} \log N$ in the vertical coordinate, it has mixed well in the horizontal one.

Proposition 1 can be used to show that, as long as the walkers have time to mix in the horizontal coordinate before exiting the cluster, the IDLA process does not depend too much on their initial positions.

Proposition 2. Fix any $T > 0$ such that $T \leq N^m$ for some finite $m \in \mathbb{N}$. Let $\{(v_i, y_i)\}_{1 \leq i \leq T}$ and $\{(v'_i, y'_i)\}_{1 \leq i \leq T}$ denote two fixed collections of vertices of \mathbb{G} such that $y_i, y'_i \leq 0$ for all i . Let, moreover, $A = (A(t))_{t \leq T}$ and $A' = (A'(t))_{t \leq T}$ be two IDLA processes with starting configurations $A(0) = A'(0)$, such that the i^{th} walker in A (respectively A') starts from (v_i, y_i) (respectively from (v'_i, y'_i)). Then there exists a coupling of A, A' such that the following holds. For any $\gamma > 0$ there exists a finite constant $b_{\gamma,m}$, depending only on γ and m , such that if

$$A(0) = A'(0) \supseteq R_{b_{\gamma,m}\sqrt{\tau_N} \log N}, \tag{3.2}$$

then

$$\mathbb{P}(A(t) = A'(t) \text{ for all } t \leq T) \geq 1 - N^{-\gamma}. \tag{3.3}$$

This tells us that, as long as an IDLA cluster contains a tall enough rectangle, releasing T walkers from fixed initial locations below level 0 or from π_N -distributed locations at level 0 results in the same final cluster with high probability. We prove this via a coupling argument, showing that, as long as the rectangle is tall enough, the walkers have time to mix before leaving the cluster.

Proof. Let γ, m be given as in the statement, and set $\gamma' = \gamma + m + 1$. Recall that R_n denotes the rectangle of height n , and assume that $A(0) = A'(0) \supseteq R_n$ for $n = 10\gamma'\sqrt{\tau_N} \log N$.

For $i \leq T$, let $\omega_i = ((v_i(k), y_i(k)))_{k \geq 0}$ and $\omega'_i = ((v'_i(k), y'_i(k)))_{k \geq 0}$ denote the simple random walk trajectories of the i^{th} walkers starting from (v_i, y_i) and (v'_i, y'_i) respectively. These are coupled as follows. If, say, $y_i < y'_i$, then ω'_i stays in place until ω_i reaches level y'_i (and vice versa if $y_i > y'_i$). We can therefore assume that $y_i = y'_i$ without loss

of generality. We make the walks move together in the vertical coordinate, so that $y_i(0) = y'_i(0)$ implies that $y_i(k) = y'_i(k)$ for all $k \geq 0$.

We explain how to generate the trajectories of ω_1, ω'_1 in such a way that they coincide upon reaching level n , with high probability. Our construction is based on [15]. The same construction can be used for the remaining walks.

Under the assumption that both ω_1 and ω'_1 start from level y_1 , generate on the same probability space an auxiliary process $Y = (Y_t)_{t \geq 0}$ on \mathbb{Z} starting from $Y_0 = y_1$. The process sojourns at each site for a Geometric random time of mean 2, and then jumps to a uniform neighbour. This will be the projection of our walkers' trajectories onto the vertical coordinate. The process stops at

$$\tau_n^Y = \inf\{t \geq 0 : Y_t = n\},$$

that is upon reaching level n . Let

$$\tau_n^X = \sum_{t=1}^{\tau_n^Y} \mathbf{1}(Y_t = Y_{t-1})$$

denote the amount of time the walk stayed still. Then, if $s \geq 3\gamma' \tau_N \log N$ is an integer, we know from (3.1) that

$$\|P_{v_1}^s - P_{v'_1}^s\|_{TV} \leq N^{-\gamma'}$$

for N large enough. Hence, conditional on $\{\tau_n^X = s\}$, we can generate on the same probability space a pair of random variables $(X, X') \in G \times G$ such that $X \sim P_{v_1}^s$, $X' \sim P_{v'_1}^s$ and $\mathbb{P}(X \neq X') \leq N^{-\gamma'}$. Set

$$\omega(\tau_n^Y) = (X, n), \quad \omega'(\tau_n^Y) = (X', n).$$

Finally, given everything else, generate the horizontal steps of the two walkers as simple random walk bridges from v_1 to X (respectively from v'_1 to X') in time s , independent of each other. The trajectories of ω, ω' up to time τ_n^Y can then be built by following the Y process when it jumps, and following the random walk bridges at times when Y stays in place.

The fact that this is a valid coupling up to time τ_n^Y follows from [15]. Moreover, since τ_n^Y is a randomised stopping time for both ω_1 and ω'_1 by construction, the strong Markov property is in force [15, 5], and hence we can make the walkers stick together after time τ_n^Y by setting

$$\omega(k) = \omega'(k) \quad \text{for } k \geq \tau_n^Y.$$

It follows that, provided the coupling is successful, ω_1 and ω'_1 will exit $A(0)$ and $A'(0)$ at the same location, thus implying that $A(1) = A'(1)$. By Proposition 1 the probability of the coupling failing is bounded by

$$\mathbb{P}(\omega_1 \text{ and } \omega'_1 \text{ do not couple within } R_n) \leq \mathbb{P}(\tau_n^Y < 3\gamma' \tau_N \log N) + \mathbb{P}(X \neq X') \leq 2N^{-\gamma'}$$

for N large enough. This procedure can be iterated for $(\omega_i, \omega'_i)_{2 \leq i \leq T}$, and if all the couplings are successful then $A(t) = A'(t)$ for all $t \leq T$. The probability of failure is

$$\mathbb{P}(\exists t \leq T \text{ such that } A(t) \neq A'(t)) \leq 2TN^{-\gamma'} \leq N^{-\gamma}$$

for N large enough. Thus (3.3) holds with $b_{\gamma,m} = 10(\gamma + m + 1)$. □

4 IDLA maximal fluctuations

In this section we adapt the *water level coupling* introduced in [13] to prove the following result on the maximal fluctuations of IDLA on \mathbb{G} .

Theorem 4 (Maximal fluctuations for polynomial times). Assume (1.1). Let $(A(t))_{t \geq 0}$ be an IDLA process on \mathbb{G} starting from the flat configuration $A(0) = R_0$. Suppose that $T = T(N) \leq N^m$ for some finite $m \in \mathbb{N}$. Then for any $\gamma > 0$ there exists a constant $C_{\gamma,m}$, depending only on $\gamma, \delta, \delta', m$, such that

$$\mathbb{P}\left(R_{\frac{T}{N} - C_{\gamma,m} \sqrt{\tau_N} (\log N)^2} \subseteq A(T) \subseteq R_{\frac{T}{N} + C_{\gamma,m} \sqrt{\tau_N} (\log N)^2}\right) \geq 1 - N^{-\gamma} \quad (4.1)$$

for N large enough.

The above theorem, of independent interest, tells us that the maximal fluctuations for IDLA on $G \times \mathbb{Z}$ from the rectangular shape are $\mathcal{O}(\sqrt{\tau_N} (\log N)^2)$ for polynomially many steps, when starting the process from flat. We split the proof into several lemmas, keeping Assumption 1 in force throughout.

Let us start with the following coupon-collector inequality stating that, starting from flat, after $\mathcal{O}(N \log N)$ releases level 1 is filled with high probability.

Lemma 1. Let $(A(t))_{t \geq 0}$ be an IDLA process on \mathbb{G} with $A(0) = R_0$. For arbitrary $\gamma > 0$, let $a_\gamma = (\gamma + 1)/\delta$. Then we have that

$$\mathbb{P}\left(A(a_\gamma N \log N) \supseteq R_1\right) \geq 1 - N^{-\gamma}$$

for N large enough.

Proof. By Assumption 1, each time a new random walk is released from level 0 each vertex at level 1 has probability at least δ/N of being filled, if empty. Thus

$$\mathbb{P}(A(a_\gamma N \log N) \not\supseteq R_1) \leq \sum_{v \in V} \mathbb{P}((v, 1) \notin A(a_\gamma N \log N)) \leq N e^{-(\gamma+1) \log N} = N^{-\gamma}. \quad \square$$

We use this simple observation to give a coupling proof of Theorem 4. Our proof generalises the one of Theorem 1.1 in [13], bypassing the lack of the log-fluctuations result via Lemma 1.

Proof of Theorem 4. The proof consists of two steps. In Step 1 we use Lemma 1 together with the Abelian property of IDLA to reduce the analysis to the case $T \leq \alpha N \sqrt{\tau_N} (\log N)^2$ for some finite constant α . In Step 2 we then bound the fluctuations of $A(T)$ for such small T by a standard argument due to Lawler et al [10].

Step 1. If $T \leq \alpha N \sqrt{\tau_N} (\log N)^2$ for some $\alpha < \infty$ then skip to Step 2. Otherwise, fix $\gamma > 0$ as in the statement. Let $a_{\gamma+2m}$ be defined as in Lemma 1, and $b_{\gamma+m+1,m}$ be as in Proposition 2. We write a, b in place of $a_{\gamma+2m}, b_{\gamma+m+1,m}$ for brevity. Define

$$\ell := \left\lfloor \frac{T}{aN \log N} \right\rfloor, \quad t_0 := (aN \log N)\ell,$$

and note that $b\sqrt{\tau_N} \log N + 1 \leq \ell \leq N^{m-1}$. At cost of increasing a and b , we may assume that t_0 and $b\sqrt{\tau_N} \log N$ are integers.

Recall that the IDLA process starts from the flat configuration $A(0) = R_0$. To start with, release $aN \log N$ walkers from level 0 according to π_N . We stop the walkers upon reaching level 1. A walker that reaches level 1 at an empty site settles there and never moves again, and the site is marked as occupied. Walkers that reach level 1 at occupied sites, on the other hand, are stopped and declared *frozen*. As the name suggests, the

motion of frozen walkers will be resumed at a later stage². Then, by Lemma 1, after $aN \log N$ releases level 1 is filled with probability at least $1 - N^{-(\gamma+2m)}$ and all walkers, whether settled or frozen, are at level 1. On this high probability event, we repeat this procedure by releasing $aN \log N$ additional walkers which stop, with the same rule as before, at level 2, and so on. Let W_1 denote the cluster of occupied sites after releasing t_0 walkers from level 0 in groups of $aN \log N$, stopping walkers of the k^{th} group at level k as described above. Then

$$\mathbb{P}(W_1 = R_\ell) \geq 1 - \ell N^{-(\gamma+2m)} \geq 1 - N^{-(\gamma+m+1)}$$

for N large enough. We stress that, while W_1 is not distributed as an IDLA cluster, the Abelian property implies that unfreezing all the frozen walkers and letting them move until settling at empty sites will result in a cluster with the same distribution as $A(t_0)$.

On the high probability event $\{W_1 = R_\ell\}$ all the frozen walkers between levels 1 and $\ell - b\sqrt{\tau_N} \log N$ are at distance at least $b\sqrt{\tau_N} \log N$ from the boundary of W_1 , so by Proposition 2 releasing them has, with high probability, the same effect as releasing the same number of walkers from level 0 according to π_N . To make this precise, note that on $\{W_1 = R_\ell\}$ there are exactly

$$t_1 := (aN \log N - N)(\ell - b\sqrt{\tau_N} \log N)$$

frozen walkers up to level $\ell - b\sqrt{\tau_N} \log N$. Let $W_1(t_1)$ denote the cluster obtained by releasing these t_1 frozen walkers. Then, if $(A'(t))_{t \geq 0}$ is an auxiliary IDLA process with $A'(0) = R_\ell$, by Proposition 2 we can couple $A'(t_1)$ and $W_1(t_1)$ so that

$$\mathbb{P}(W_1 = R_\ell, A'(t_1) = W_1(t_1)) \geq 1 - 2N^{-(\gamma+m+1)}$$

for N large enough. This suggests a way to couple the original IDLA process A with the auxiliary IDLA process A' . Indeed, having built $A'(t_1)$ as above, we can add $T_1 = T - N(\ell - b\sqrt{\tau_N} \log N)$ particles to $A'(t_1)$ using the locations of the frozen particles in $W_1(t_1)$ as starting locations. By the Abelian property, this gives

$$\|A(T) - A'(T_1)\|_{TV} \leq \mathbb{P}(W_1 \neq R_\ell) + \mathbb{P}(W_1 = R_\ell, A'(t_1) \neq W_1(t_1)) \leq 2N^{-(\gamma+m+1)}$$

for N large enough. It thus suffices to prove that (4.1) holds with $A'(T_1)$ in place of $A(T)$. Moreover, since the process A' started from R_ℓ , the downshift of $A'(T_1)$ by ℓ levels has the same distribution as $A(T_1)$. This therefore reduces the problem to proving that (4.1) holds with T_1 in place of T . Now, if $T_1 \leq \alpha N \sqrt{\tau_N} (\log N)^2$ for some $\alpha < \infty$ stop and move to Step 2. Otherwise, iterate this procedure. Note that each iteration reduces T by at least $aN \log N - N$. Then, after $k < N^{m-1}$ iterations we are left with

$$T_k \leq (aN \log N)(2b\sqrt{\tau_N} (\log N)^2),$$

and

$$\begin{aligned} \mathbb{P}\left(R_{\frac{T}{N} - C_{\gamma,m}\sqrt{\tau_N}(\log N)^2} \subseteq A(T) \subseteq R_{\frac{T}{N} + C_{\gamma,m}\sqrt{\tau_N}(\log N)^2}\right) &\geq \\ &\geq \mathbb{P}\left(R_{\frac{T_k}{N} - C_{\gamma,m}\sqrt{\tau_N}(\log N)^2} \subseteq A(T_k) \subseteq R_{\frac{T_k}{N} + C_{\gamma,m}\sqrt{\tau_N}(\log N)^2}\right) - N^{-(\gamma+2)}, \end{aligned}$$

for N large enough. At this point we can move to Step 2.

Step 2. If $T \leq \alpha' N (\log N)^2$ for some finite constant α' then the conclusion is trivial. Otherwise, let $(A(t))_{t \geq 0}$ be an IDLA process starting from $A(0) = R_0$. We show that, if

²While freezing particles changes the distribution of the cluster, the Abelian property tells us that we can recover regular IDLA once we unfreeze all the frozen particles and let them evolve according to the usual IDLA rules.

$T \leq \alpha N \sqrt{\tau_N} (\log N)^2$ for some α (which may depend on γ, m), then there exists $\beta \in (0, 1)$, independent of N , such that

$$\mathbb{P}\left(A(T) \not\subseteq R_{\frac{3T}{N}}\right) \leq \beta^{T/N} \tag{4.2}$$

for N large enough. This would conclude the proof, as the inner bound is trivial since $T/N \leq \alpha \sqrt{\tau_N} (\log N)^2$. To see (4.2), let $Z_k(t) := |A(t) \cap \{y = k\}|$ denote the number of sites in $A(t)$ at level k , and set $\mu_k(t) := \mathbb{E}(Z_k(t))$. Then

$$\mathbb{P}\left(A(T) \not\subseteq R_{\frac{3T}{N}}\right) \leq \mathbb{P}\left(Z_{\frac{3T}{N}+1}(T) \geq 1\right) \leq \mu_{\frac{3T}{N}+1}(T).$$

We claim that

$$\mu_k(t) \leq N \left(\frac{1}{N}\right)^{k-1} \frac{t^k}{k!} \tag{4.3}$$

for all $k > 0$ and $t \geq 0$. Indeed, clearly $\mu_1(t) \leq N$ and $\mu_k(0) = 0$ for all $k > 0$. For other values of k, j we have

$$\mu_k(t+1) - \mu_k(t) = \mathbb{E}(Z_k(t+1) - Z_k(t)) \leq \frac{\delta'}{N} \mathbb{E}(|A(t) \cap \{y = k-1\}|) = \frac{\delta'}{N} \mu_{k-1}(t),$$

where in the above inequality we have used that the probability that the $(t+1)^{th}$ walker reaches level $k-1$ inside $A(t)$ is maximised when $A(t)$ is completely filled up to level $k-2$. Thus for $k > 0$

$$\mu_k(t) = \sum_{s=0}^{t-1} (\mu_k(s+1) - \mu_k(s)) \leq \frac{\delta'}{N} \sum_{s=0}^{t-1} \mu_{k-1}(s),$$

and (4.3) follows by a simple iteration. Now take $t = T$, $k = \frac{3T}{N} + 1$ and recall that $k! \geq k^k e^{-k}$ to get

$$\mu_{\frac{3T}{N}+1}(T) \leq N \left(\frac{1}{N}\right)^{\frac{3T}{N}} \frac{T^{\frac{3T}{N}+1}}{(\frac{3T}{N}+1)!} \leq N^2 \left[\left(\frac{e}{3}\right)^3\right]^{\frac{T}{N}} \leq \beta^{T/N}$$

for any $\beta \in \left(\left(\frac{e}{3}\right)^3, 1\right)$ and N large enough. This proves (4.2), hence concluding the proof of Theorem 4. □

5 Proof of Theorems 1 and 2

We sketch here the proofs of Theorems 1 and 2. We give precise statements of the several lemmas leading to the proof, but choose to leave the proof details to the reader, since these are straightforward generalisations of the arguments in [13]. This section also serves as a survey of the ideas introduced for the case $G = \mathbb{Z}_N$ in [13]. We split the argument into several steps, for ease of reading.

Step 1: Decay of excess height. For an IDLA cluster $A \in \Omega$, recall that $|A|$ counts the number of vertices in A above level 0. We define the *excess height* of A by

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

where $h(A)$ denotes the height of A . Thus $\mathcal{E}(R_k) = 0$ for all infinite rectangles R_k , while $\mathcal{E}(A)$ is large if A is tall and sparse. To start with, we show that if the excess height of a cluster is too large, it tends to decrease under IDLA dynamics.

Lemma 2. Let $(A(t))_{t \geq 0}$ be an IDLA process on $G \times \mathbb{Z}$ with $A(0) \supseteq R_0$ and filtration $(\mathcal{F}_t)_{t \geq 0}$. For $t \geq 0$ let $\mathcal{E}(t) = \mathcal{E}(A(t))$ denote the excess height of $A(t)$. Then there exists a constant $C_\mathcal{E}$, depending only on δ , such that if $\mathcal{E}(t) > C_\mathcal{E} N \sqrt{\tau_N} (\log N)^2$ then

$$\mathbb{E}(\mathcal{E}(t+1) - \mathcal{E}(t) | \mathcal{F}_t) < -\frac{1}{2N}$$

for N large enough.

To see this, note that if the excess height is large, then there are many empty sites below level $h(t)$. Call a level below $h(t)$ *bad* if it contains at least one empty site. If the excess height is high enough, then there are many bad levels at distance at least $20\sqrt{\tau_N} \log N$ one another. Thus a random walker has time to mix in the horizontal coordinate between bad levels, and each time it reaches a new bad level it has probability at least $\delta/(2N)$ to settle at an empty site. It follows that, if there are enough bad levels, the random walk will settle below height $h(t)$ with high probability, which will cause the excess height to decrease. We refer the reader to [13], Lemma 6.2, for details.

Lemma 2 shows that an IDLA process spends a small proportion of time in clusters with high excess height. From this it is standard (cf. [13], Proposition 6.1) to deduce that stationary clusters have low excess height, which gives the following result.

Lemma 3. For any $\gamma > 0$ there exists a constant $C_{\mathcal{E},\gamma}$, depending only on γ, δ , such that, with $\mathcal{E}_\gamma^* = C_{\mathcal{E},\gamma} \sqrt{\tau_N} N (\log N)^2$,

$$\mu_N(\{A \in \Omega : \mathcal{E}(A) > \mathcal{E}_\gamma^*\}) \leq N^{-\gamma}$$

for N large enough.

Step 2: From low excess height to low height. Let $(A(t))_{t \geq 0}$ be an IDLA process starting from a stationary configuration $A(0) \sim \mu_N$, and denote the associated shifted IDLA process by $(A^*(t))_{t \geq 0}$. While Lemma 3 shows that $A(0)$ has low excess height with high probability, this does not imply low height, since $|A(0)|$ is random. It gives, on the other hand, an upper bound on the number of empty sites, that is the number of sites in $R_{h(A(0))} \setminus A(0)$. This allows us to show that the height of the associated shifted IDLA process becomes $\mathcal{O}(N^4 \sqrt{\tau_N})$ within polynomial time, as stated below.

Lemma 4. Let $(A^*(t))_{t \geq 0}$ be a shifted IDLA process with $A(0) \sim \mu_N$, and let

$$T^* = \inf\{t \geq 0 : h(A^*(t)) \leq N^4 \sqrt{\tau_N}\}.$$

Then for any $\gamma > 0$ there exists a constant $C_{\mathcal{E},\gamma}^*$, depending only on γ, δ , such that

$$\mathbb{P}(T^* > N^6 \sqrt{\tau_N}) \leq N^{-\gamma}$$

for N large enough.

The above result is proved as follows. On the high probability event $\mathcal{E}(A(0)) \leq \mathcal{E}_{\gamma+2}^*$ there are at most $N\mathcal{E}_{\gamma+2}^*$ empty sites below level $h(A(0))$. Each time we add a new particle to the cluster, the lowest empty site has probability at least δ/N to be filled, independently of everything else. Hence it will take at most a geometric number of releases of parameter δ/N to fill it. We fill empty sites sequentially while the excess height does not exceed $2\mathcal{E}_{\gamma+2}^*$. When it does, we wait for the dynamics to bring the excess height below $\mathcal{E}_{\gamma+2}^*$, and try again to fill all the empty sites. It is easy to see that, with high probability, this procedure succeeds within $N^6 \sqrt{\tau_N}$ particle releases. See [13], Proposition 6.2, for details.

Remark 1. The same argument shows that shifted IDLA defines a positive recurrent Markov chain, since starting from the flat configuration R_0 , the chain comes back to a certain finite set containing R_0 in finite time.

Step 3: Forgetting polynomially high initial configurations. Lemma 4 tells us that stationary clusters have $\mathcal{O}(N^6\sqrt{\tau_N})$ height with high probability. We use the polynomial height bound to show that an IDLA process starting from a stationary configuration forgets that it did not start from flat within polynomially many steps.

Proposition 3. Let $A_0, A'_0 \in \Omega$ be such that $|A_0| = |A'_0|$, $A'_0 \supseteq R_{\lfloor |A_0|/N \rfloor}$ and

$$h(A_0) \leq \alpha N^6 \sqrt{\tau_N}, \quad h(A'_0) \leq \frac{|A_0|}{N} + 1$$

for some absolute constant $\alpha < \infty$. Let $(A(t))_{t \geq 0}$ and $(A'(t))_{t \geq 0}$ denote two IDLA processes starting from A_0 and A'_0 respectively, and denote the laws of $A(t)$, $A'(t)$ by $P(t)$, $P'(t)$. Then for any $\gamma > 0$ there exists a constant c_γ , depending only on γ, δ, δ' , such that, writing $s_\gamma = \alpha N^7 \sqrt{\tau_N} + c_\gamma N \sqrt{\tau_N} (\log N)^2$ for brevity, we have

$$\|P(s_\gamma) - P'(s_\gamma)\|_{TV} \leq N^{-\gamma}$$

for N large enough.

To prove the above result we use the coupling ideas introduced in Section 4. Indeed, let W_0 be a new IDLA cluster, independent of everything else, built by adding s_γ particles to the flat configuration R_0 . Let $\beta_{\gamma+1,9}$ be as in the statement of Proposition 2, and write β in place of $\beta_{\gamma+1,9}$ for brevity. By Theorem 4 we can choose c_γ large enough so that W_0 is completely filled up to height $\alpha N^6 \sqrt{\tau_N} + \beta \sqrt{\tau_N} (\log N)^2$ with high probability. In particular, W_0 contains both $A(0)$ and $A'(0)$, and each vertex of $A(0) \cup A'(0)$ is at distance at least $\beta \sqrt{\tau_N} (\log N)^2$ from the boundary of W_0 .

Let $W(0)$ (respectively $W'(0)$) denote the cluster obtained by adding to W_0 a frozen particle at each site of $A(0)$ (respectively $A'(0)$). Then $W(0), W'(0)$ contain the same number $T = |A_0|$ of frozen particles, that we release in pairs. To do this, fix an arbitrary enumeration of the locations of the frozen particles:

$$z_1, z_2 \dots z_T \in W(0), \quad z'_1, z'_2 \dots z'_T \in W'(0).$$

For each $1 \leq t \leq T$ we build $W(t), W'(t)$ from $W(t-1), W'(t-1)$ as follows. Start two simple random walks in $W(t-1), W'(t-1)$ from z_t, z'_t . The walks are coupled in such a way that their trajectories meet within W_0 with probability at least $N^{-(\gamma+10)}$, as in the proof of Proposition 2. They settle upon exiting $W(t-1), W'(t-1)$ respectively, and their settling locations are then added to $W(t-1), W'(t-1)$ to make $W(t), W'(t)$. Note that, importantly, if $W(t-1) = W'(t-1)$ and the t^{th} coupling is successful, then $W(t) = W'(t)$.

This defines two auxiliary processes $(W(t))_{t \leq T}$ and $(W'(t))_{t \leq T}$ on Ω . We stress that these are not IDLA processes, due to the presence of the frozen walkers. Nevertheless, the Abelian property implies that

$$A(s_\gamma) \stackrel{(d)}{=} W(T), \quad A'(s_\gamma) \stackrel{(d)}{=} W'(T),$$

since after T steps all the frozen walkers have settled. To finish the proof it then suffices to show that

$$\|W(T) - W'(T)\|_{TV} \leq N^{-(\gamma+1)}$$

for N large enough. But this follows from Proposition 2, since $T = |A_0|$ is at most polynomial in N , and all the frozen walkers' starting locations are at distance at least $\beta \sqrt{\tau_N} (\log N)^2$ from the cluster boundary, thus leaving time for the walkers to mix before exiting.

Step 4: Stationary height: proof of Theorem 1. We use Proposition 3 to argue that a stationary shifted IDLA process has, after polynomially many steps, both stationary law

and $\mathcal{O}(\sqrt{\tau_N}(\log N)^2)$ fluctuations, since it forgot that it did not start from flat with high probability. This tells us that stationary clusters have height $\mathcal{O}(\sqrt{\tau_N}(\log N)^2)$, and it thus proves Theorem 1.

To expand, let $(A(t))_{t \geq 0}$ denote a shifted IDLA process started at stationarity $A(0) \sim \mu_N$. Then $A(t) \sim \mu_N$ for all $t \geq 0$. In particular, if for arbitrary $\gamma > 0$ we define s_γ as in Proposition 3, we have that $A(s_\gamma) \sim \mu_N$. On the other hand, since $h(A_0) = \mathcal{O}(N^6 \sqrt{\tau_N})$ with high probability by Lemma 3, Proposition 3 tells us that $A(s_\gamma)$ is close in distribution to a shifted IDLA process started from height at most 1, and hence, using that $s_\gamma \leq N^9$ for N large enough,

$$h(A(s_\gamma)) \leq C_\gamma \sqrt{\tau_N} (\log N)^2$$

with high probability by Theorem 4.

Step 5: Proof of Theorem 2. For arbitrary $\gamma > 0$, let c_γ be defined as in Theorem 1, and set

$$\Omega_\gamma = \{A \in \Omega : h(A) \leq c_\gamma \sqrt{\tau_N} (\log N)^2\}.$$

Then $\mu_N(\Omega_\gamma^c) \leq N^{-\gamma}$, and Theorem 2 follows by exactly the same argument used in the proof of Proposition 3.

References

- [1] V.S. Anil Kumar and Hariharan Ramesh. Coupling vs. conductance for the jerrum–sinclair chain. *Random Structures & Algorithms*, 18(1):1–17, 2001. MR-1799799
- [2] Amine Asselah and Alexandre Gaudillière. From logarithmic to subdiffusive polynomial fluctuations for internal DLA and related growth models. *Ann. Probab.*, 41(3A):1115–1159, 2013. MR-3098673
- [3] Amine Asselah, Alexandre Gaudillière, et al. Sublogarithmic fluctuations for internal DLA. *The Annals of Probability*, 41(3A):1160–1179, 2013. MR-3098674
- [4] Krzysztof Burdzy and Wilfrid S. Kendall. Efficient markovian couplings: examples and counterexamples. *Annals of Applied Probability*, pages 362–409, 2000. MR-1768241
- [5] Stephen Connor and Saul Jacka. Optimal co-adapted coupling for the symmetric random walk on the hypercube. *Journal of Applied Probability*, 45(3):703–713, 2008. MR-2455179
- [6] 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. MR-1218674
- [7] David Griffeath. A maximal coupling for Markov chains. *Probability Theory and Related Fields*, 31(2):95–106, 1975. MR-0370771
- [8] David Jerison, Lionel Levine, and Scott Sheffield. Logarithmic fluctuations for internal DLA. *J. Amer. Math. Soc.*, 25(1):271–301, 2012. MR-2833484
- [9] David Jerison, Lionel Levine, and Scott Sheffield. Internal DLA and the Gaussian free field. *Duke Math. J.*, 163(2):267–308, 2014. MR-3161315
- [10] Gregory F. Lawler, Maury Bramson, and David Griffeath. Internal diffusion limited aggregation. *Ann. Probab.*, 20(4):2117–2140, 1992. MR-1188055
- [11] David A. Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017. MR-3726904
- [12] Lionel Levine and Vittoria Silvestri. How far do activated random walkers spread from a single source? *In preparation*.
- [13] Lionel Levine and Vittoria Silvestri. How long does it take for internal DLA to forget its initial profile? *Probability Theory and Related Fields*, pages 1–53, 2018. MR-3980315
- [14] Paul Meakin and John M. Deutch. The formation of surfaces by diffusion limited annihilation. *The Journal of Chemical Physics*, 85(4):2320–2325, 1986.

- [15] J.W. Pitman. On coupling of Markov chains. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 35(4):315–322, 1976. MR-0415775
- [16] Leonardo T. Rolla and Vladas Sidoravicius. Absorbing-state phase transition for driven-dissipative stochastic dynamics on \mathbb{Z} . *Inventiones Mathematicae*, 188(1):127–150, 2012. MR-2897694

Acknowledgments. I am thankful to Lionel Levine for several insightful discussions, and for a careful reading of an earlier version of this note. I am also indebted with Alexandre Stauffer for pointing me to Pitman’s construction of Markov chains optimal couplings, and for many stimulating discussions on this and related models.