

COLLAPSE AND DIFFUSION IN HARMONIC ACTIVATION AND TRANSPORT

JACOB CALVERT, SHIRSHENDU GANGULY, AND ALAN HAMMOND

ABSTRACT. For an n -element subset U of \mathbb{Z}^2 , select x from U according to harmonic measure from infinity, remove x from U , and start a random walk from x . If the walk leaves from y when it first enters the rest of U , add y to it. Iterating this procedure constitutes the process we call *Harmonic Activation and Transport* (HAT).

HAT exhibits a phenomenon we refer to as *collapse*: informally, the diameter shrinks to its logarithm over a number of steps which is comparable to this logarithm. Collapse implies the existence of the stationary distribution of HAT, where configurations are viewed up to translation, and the exponential tightness of diameter at stationarity. Additionally, collapse produces a renewal structure with which we establish that the center of mass process, properly rescaled, converges in distribution to two-dimensional Brownian motion.

To characterize the phenomenon of collapse, we address fundamental questions about the extremal behavior of harmonic measure and escape probabilities. Among n -element subsets of \mathbb{Z}^2 , what is the least positive value of harmonic measure? What is the probability of escape from the set to a distance of, say, d ? Concerning the former, examples abound for which the harmonic measure is exponentially small in n . We prove that it can be no smaller than exponential in $n \log n$. Regarding the latter, the escape probability is at most the reciprocal of $\log d$, up to a constant factor. We prove it is always at least this much, up to an n -dependent factor.

CONTENTS

1. Introduction	2
2. Conceptual overview	8
3. Harmonic measure estimates	11
4. Escape probability estimates	22
5. Clustering sets of relatively large diameter	25
6. Estimates of the time of collapse	26
7. Existence of the stationary distribution	48
8. Motion of the center of mass	59
Appendix A. Auxiliary lemmas	61
References	68

2020 *Mathematics Subject Classification.* 60J45, 82C24, and 82C41.

Key words and phrases. Markov chain, harmonic measure, random walk.

J.C. was partially supported by NSF grant DMS-1512908. S.G. was partially supported by NSF grant DMS-1855688, NSF CAREER Award DMS-1945172, and a Sloan Fellowship. A.H. was partially supported by NSF grants DMS-1512908 and DMS-1855550, and a Miller Professorship from the Miller Institute for Basic Research in Science.

1. INTRODUCTION

1.1. Harmonic activation and transport. Consider simple random walk $(S_j)_{j \geq 0}$ on \mathbb{Z}^2 and with $S_0 = x$, the distribution of which we denote by \mathbb{P}_x . For a finite, nonempty subset $A \subset \mathbb{Z}^2$, the hitting distribution of A from $x \in \mathbb{Z}^2$ is the function $\mathbb{H}_A(x, \cdot) : \mathbb{Z}^2 \rightarrow [0, 1]$ defined as $\mathbb{H}_A(x, y) = \mathbb{P}_x(S_{\tau_A} = y)$, where $\tau_A = \inf\{j \geq 1 : S_j \in A\}$ is the return time to A . (We use the notation τ_A instead of the more common τ_A^+ for brevity.) The recurrence of random walk on \mathbb{Z}^2 guarantees that τ_A is almost surely finite, and the existence of the limit $\mathbb{H}_A(y) = \lim_{|x| \rightarrow \infty} \mathbb{H}_A(x, y)$, called the harmonic measure of A , is well known [Law13]. Informally, the harmonic measure is the hitting distribution of a random walk “from infinity.”

In this paper, we introduce a Markov chain called *Harmonic Activation and Transport* (HAT), wherein the elements of a subset of \mathbb{Z}^2 (respectively styled as “particles” of a “configuration”) are iteratively selected according to harmonic measure and replaced according to the hitting distribution of a random walk started from the location of the selected element. We say that, with each step, a particle is “activated” and then “transported.”

Definition 1.1 (Harmonic activation and transport). *Given a finite subset U_0 of \mathbb{Z}^2 with at least two elements, HAT is the discrete-time Markov chain $(U_t)_{t \geq 0}$ on subsets of \mathbb{Z}^2 , the dynamics of which consists of the following steps (Figure 1).*

Activation. *At time t , remove a random element $X_t \sim \mathbb{H}_{U_t}$ from U_t , forming $V_t = U_t \setminus \{X_t\}$.*

Transport. *Then, add a random element $Y_t \sim \mathbb{P}_{X_t}(S_{\tau_{V_t}-1} \in \cdot \mid V_t)$ to V_t , forming $U_{t+1} = V_t \cup \{Y_t\}$.*

In other words, $(U_t)_{t \geq 0}$ has inhomogeneous transition probabilities given by

$$\mathbf{P}(U_{t+1} = (U_t \setminus \{x\}) \cup \{y\} \mid U_t) = \begin{cases} \mathbb{H}_{U_t}(x) \mathbb{P}_x(S_{\tau_{U_t \setminus \{x\}}-1} = y \mid U_t) & x \neq y, \\ \sum_{z \in U_t} \mathbb{H}_{U_t}(z) \mathbb{P}_z(S_{\tau_{U_t \setminus \{z\}}-1} = z \mid U_t) & x = y. \end{cases}$$

To guide the presentation of our results, we highlight four features of HAT. Two of these reference the diameter of a configuration, defined as $\text{diam}(U) = \sup_{x, y \in U} |x - y|$, where $|\cdot|$ is the Euclidean norm.

- *Conservation of mass.* HAT conserves the number of particles in the initial configuration.
- *Translation invariance.* For any configurations V, W and element $x \in \mathbb{Z}^2$,

$$\mathbf{P}(U_{t+1} = V \mid U_t = W) = \mathbf{P}(U_{t+1} = V + x \mid U_t = W + x).$$

In words, the HAT dynamics is invariant under translation by elements of \mathbb{Z}^2 . Accordingly, to each configuration U , we can associate an equivalence class

$$\widehat{U} = \{V \subseteq \mathbb{Z}^2 : \exists x \in \mathbb{Z}^2 : U = V + x\}.$$

- *Variable connectivity.* The HAT dynamics does not preserve connectivity. Indeed, a configuration which is initially connected will eventually be disconnected by the HAT dynamics, and the resulting components may “treadmill” away from one another, adopting configurations of arbitrarily large diameter.
- *Asymmetric behavior of diameter.* While the diameter of a configuration can increase by at most 1 with each step, it can decrease abruptly. For example, if the configuration is a pair of particles separated by d , then the diameter will decrease by $d - 1$ in one step.

We will shortly state the existence of the stationary distribution of HAT. By the translation invariance of the HAT dynamics, the stationary distribution will be supported on equivalence classes of configurations which, for brevity, we will simply refer to as configurations. In fact, the HAT dynamics cannot reach all such configurations. By an inductive argument, we will prove that the HAT dynamics is irreducible on

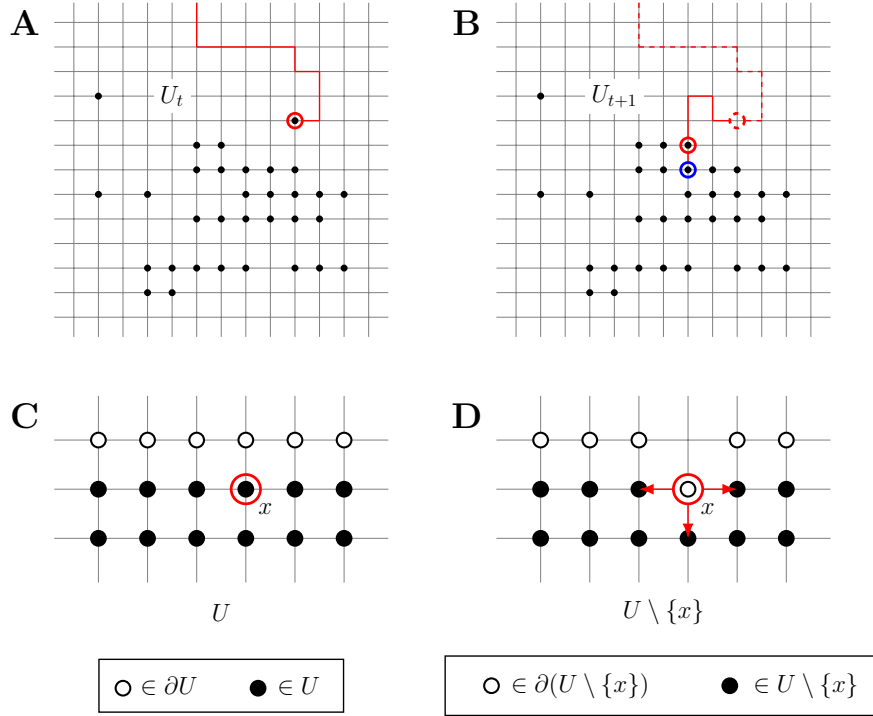


FIGURE 1. The harmonic activation and transport dynamics. **(A)** A particle (indicated by a solid, red circle) in the configuration U_t is activated according to harmonic measure. **(B)** The activated particle (following the solid, red path) hits another particle (indicated by a solid, blue circle); it is then fixed at the site visited during the previous step (indicated by a solid, red circle), giving U_{t+1} . **(C)** A particle of U (indicated by a red circle) is activated and **(D)** if it tries to move into $U \setminus \{x\}$, the particle will be placed at x . The notation ∂U refers to the exterior vertex boundary of U .




FIGURE 2. A configuration that HAT cannot reach.

the collection of configurations that have a non-isolated element with positive harmonic measure. Figure 2 depicts a configuration that HAT cannot reach because every element with positive harmonic measure has no neighbors in \mathbb{Z}^2 .

Definition 1.2. Denote by $\text{Iso}(n)$ the collection of n -element subsets U of \mathbb{Z}^2 such that every x in U with $\mathbb{H}_U(x) > 0$ belongs to a singleton connected component. In other words, all exposed elements of U are isolated: they lack nearest neighbors in U . We will denote the collection of all other n -element subsets of \mathbb{Z}^2 by $\text{NonIso}(n)$, and the corresponding equivalence class by

$$\widehat{\text{NonIso}}(n) = \{\widehat{U} : U \in \text{NonIso}(n)\}.$$

The variable connectivity of HAT configurations and concomitant opportunity for unchecked diameter growth seem to jeopardize the positive recurrence of the HAT dynamics on $\widehat{\text{NonIso}}(n)$. Indeed, if the diameter were to grow unabatedly, the HAT dynamics could not return to a configuration or equivalence class thereof, and would therefore be doomed to transience. However, due to the asymmetric behavior of diameter under the HAT dynamics and the recurrence of random walk in \mathbb{Z}^2 , this will not be the case. For an arbitrary initial configuration of $n \geq 2$ particles, we will prove—up to a factor depending on n —sharp bounds on the “collapse” time which, informally, is the first time the diameter is at most a certain function of n .

Definition 1.3. For a positive real number R , we define the level- R collapse time to be $\mathcal{T}(R) = \inf\{t \geq 0 : \text{diam}(U_t) \leq R\}$.

For a real number $r \geq 0$, we define $\theta_m = \theta_m(r)$ through

$$\theta_0 = r \quad \text{and} \quad \theta_m = \theta_{m-1} + e^{\theta_{m-1}} \quad \text{for } m \geq 1. \tag{1.1}$$

In particular, $\theta_n(r)$ is approximately the n^{th} iterated exponential of r .

Theorem 1. Let U be a finite subset of \mathbb{Z}^2 with $n \geq 2$ elements and a diameter of d . There exists a universal positive constant c such that, if d exceeds $\theta = \theta_{4n}(cn)$, then

$$\mathbf{P}_U \left(\mathcal{T}(\theta) \leq (\log d)^{1+o_n(1)} \right) \geq 1 - e^{-n}.$$

For the sake of concreteness, this is true with n^{-4} in the place of $o_n(1)$.

In words, for a given n , it typically takes $(\log d)^{1+o_n(1)}$ steps before the configuration of initial diameter d reaches a configuration with a diameter of no more than a large function of n . Here, $o_n(1)$ denotes a nonnegative function of n that is at most 1 and which tends to zero as n tends to ∞ . The d dependence in Theorem 1 is essentially the best possible, aside from the $o_n(1)$ term, because two pairs of particles separated by a distance of d typically exchange particles over $\log d$ steps. We elaborate this point in Section 2.

As a consequence of Theorem 1 and the preceding discussion, it will follow that the HAT dynamics constitutes an aperiodic, irreducible, and positive recurrent Markov chain on $\widehat{\text{NonIso}}(n)$. In particular, this means that, from any configuration of $\widehat{\text{NonIso}}(n)$, the time it takes for the HAT dynamics to return to that configuration is finite in expectation. Aperiodicity, irreducibility, and positive recurrence imply the existence and uniqueness of the stationary distribution π_n , to which HAT converges from any n -element configuration. Moreover—again, due to Theorem 1—the stationary distribution is exponentially tight.

Theorem 2. For every $n \geq 2$, from any n -element subset of \mathbb{Z}^2 , HAT converges to a unique probability measure π_n supported on $\widehat{\text{NonIso}}(n)$. Moreover, π_n satisfies the following tightness estimate. There exists a universal positive constant c such that, for any $r \geq 2\theta_{4n}(cn)$,

$$\pi_n(\text{diam}(\widehat{U}) \geq r) \leq \exp \left(-\frac{r}{(\log r)^{1+o_n(1)}} \right).$$

In particular, this is true with $6n^{-4}$ in the place of $o_n(1)$.

The r dependence in the tail bound of Theorem 2 is likely suboptimal, because its proof makes critical use of the fact that the diameter of a configuration increases by at most 1 with each step. A sufficiently high probability, sublinear bound on the growth rate of diameter would improve the rate of exponential decay. We note that an analogue of Kesten’s sublinear bound on the diameter growth rate of diffusion-limited aggregation [Kes87] would apply only to growth resulting from the exchange of particles between well separated “clusters” of particles, not to growth from intracluster transport.

As a further consequence of Theorem 1, we will find that the HAT dynamics exhibits a renewal structure which underlies the diffusive behavior of the corresponding center of mass process.

Definition 1.4. For a sequence of configurations $(U_t)_{t \geq 0}$ with n particles, define the corresponding center of mass process $(\mathcal{M}_t)_{t \geq 0}$ by $\mathcal{M}_t = n^{-1} \sum_{x \in U_t} x$.

For the following statement, denote by $\mathcal{C}([0, 1])$ the continuous functions $f : [0, 1] \rightarrow \mathbb{R}^2$ with $f(0)$ equal to the origin $o \in \mathbb{Z}^2$, equipped with the topology induced by the supremum norm $\sup_{0 \leq t \leq 1} |f(t)|$.

Theorem 3. If \mathcal{M}_t is linearly interpolated, then the law of the process $(t^{-1/2} \mathcal{M}_{st}, s \in [0, 1])$, viewed as a measure on $\mathcal{C}([0, 1])$, converges weakly as $t \rightarrow \infty$ to two-dimensional Brownian motion on $[0, 1]$ with coordinate diffusivity $\chi^2 = \chi^2(n)$. Moreover, for a universal positive constant c , χ^2 satisfies:

$$\theta_{6n}(cn)^{-1} \leq \chi^2 \leq \theta_{6n}(cn).$$

The same argument allows us to extend the convergence to $[0, \infty)$. The bounds on χ^2 are not tight and we have not attempted to optimize them; they primarily serve to show that χ^2 is positive and finite.

1.2. Extremal behavior of harmonic measure. As we elaborate in Section 2, the timescale of diameter collapse in Theorem 1 arises from novel estimates of harmonic measure and hitting probabilities, which control the activation and transport dynamics of HAT. Beyond their relevance to HAT, these results further the characterization of the extremal behavior of harmonic measure.

Estimates of harmonic measure often apply only to connected sets or depend on the diameter of the set. The discrete analogues of Beurling’s projection theorem [Kes87] and Makarov’s theorem [Law93] are notable examples. Furthermore, estimates of hitting probabilities often approximate sets by Euclidean balls which contain them (for example, the estimates in Chapter 2 of [Law13]). Such approximations work well for connected sets, but not for sets which are “sparse” in the sense that they have large diameters relative to their cardinality; we elaborate this in Section 2.2. For the purpose of controlling the HAT dynamics, which adopts such sparse configurations, existing estimates of harmonic and hitting measures are inapplicable.

To highlight the difference in the behavior of harmonic measure for general (i.e., potentially sparse) and connected sets, consider a finite subset A of \mathbb{Z}^2 with $n \geq 2$ elements. We ask: *What is the greatest value of $\mathbb{H}_A(x)$?* If we assume no more about A , then $\mathbb{H}_A(x)$ can be as large as $\frac{1}{2}$ (see Section 2.5 of [Law13] for an example). However, if A is connected, then the discrete analogue of Beurling’s projection theorem [Kes87] provides a finite constant c such that

$$\mathbb{H}_A(x) \leq cn^{-1/2}.$$

This upper bound is realized (up to a constant factor) when A is a line segment and x is one of its endpoints.

Our next result provides lower bounds of harmonic measure to complement the preceding upper bounds, addressing the question: *What is the least positive value of $\mathbb{H}_A(x)$?*

Theorem 4. There exists a universal positive constant c such that, if A is a subset of \mathbb{Z}^2 with $n \geq 1$ elements, then either $\mathbb{H}_A(x) = 0$ or

$$\mathbb{H}_A(x) \geq e^{-cn \log n}. \tag{1.2}$$

It is much easier to prove that, if A is connected, then (1.2) can be replaced by

$$\mathbb{H}_A(x) \geq e^{-cn}.$$

This lower bound is optimal in terms of its dependence on n , as we can choose A to be a narrow, rectangular “tunnel” with a depth of order n and an element just above the “bottom” of the tunnel, in which case the harmonic measure of this element is exponentially small in n ; we will shortly discuss a related example in greater detail. We expect that the bound in (1.2) can be improved to an exponential decay with a rate of order n instead of $n \log n$.

We believe that the best possible lower bound would be realized by the harmonic measure of the innermost element of a square spiral (Figure 3). The virtue of the square spiral is that, essentially, with each additional element, the shortest path to the innermost element lengthens by two steps. This heuristic suggests that the least positive value of harmonic measure should decay no faster than 4^{-2n} , as $n \rightarrow \infty$. Indeed, Example 1.6 suggests an asymptotic decay rate of $(2 + \sqrt{3})^{-2n}$. We formalize this observation as a conjecture. To state it, let \mathcal{H}_n be the collection of n -element subsets A of \mathbb{Z}^2 such that $\mathbb{H}_A(o) > 0$.

Conjecture 1.5. *Asymptotically, the square spiral of Figure 3 realizes the least positive value of harmonic measure, in the sense that*

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log \inf_{A \in \mathcal{H}_n} \mathbb{H}_A(o) = 2 \log(2 + \sqrt{3}).$$

Example 1.6. Figure 3 depicts the construction of an increasing sequence of sets (A_1, A_2, \dots) such that, for all $n \geq 1$, A_n is an element of \mathcal{H}_n and the shortest path $\Gamma = (\Gamma_1, \Gamma_2, \dots, \Gamma_k)$ from the exterior boundary of $A_n \cup \partial A_n$ to $\Gamma_k = o$, which satisfies $\Gamma_i \notin A_n$ for $1 \leq i \leq k-1$, has a length of $k = 2(1 - o_n(1))n$. Since Γ_1 separates the origin from infinity in A_n^c , we have

$$\mathbb{H}_{A_n}(o) = \mathbb{H}_{A_n \cup \{\Gamma_1\}}(\Gamma_1) \cdot \mathbb{P}_{\Gamma_1}(S_{\tau_{A_n}} = o). \quad (1.3)$$

Concerning the first factor of (1.3), one can show that there exist positive constants $b, c < \infty$ such that, for all sufficiently large n ,

$$cn^{-b} \leq \mathbb{H}_{A_n \cup \{\Gamma_1\}}(\Gamma_1) \leq 1.$$

To address the second factor of (1.3), we sum over the last time $t < \tau_{A_n}$ that S_t visits Γ_1 :

$$\mathbb{P}_{\Gamma_1}(S_{\tau_{A_n}} = o) = \sum_{t=0}^{\infty} \mathbb{P}_{\Gamma_1}(S_t = \Gamma_1, t < \tau_{A_n}; \{S_{t+1}, \dots, S_{\tau_{A_n}}\} \subseteq \Gamma_{2:k}),$$

where $\Gamma_{2:k} = \{\Gamma_2, \dots, \Gamma_k\}$. The Markov property applied to t implies that

$$\mathbb{P}_{\Gamma_1}(\{S_{t+1}, \dots, S_{\tau_{A_n}}\} \subseteq \Gamma_{2:k} \mid S_t = \Gamma_1, t < \tau_{A_n}) = \mathbb{P}_{\Gamma_2}(\tau_o < \tau_{\mathbb{Z}^2 \setminus \Gamma_{2:k}}).$$

Therefore

$$\mathbb{P}_{\Gamma_1}(S_{\tau_{A_n}} = o) = \mathbb{P}_{\Gamma_2}(\tau_o < \tau_{\mathbb{Z}^2 \setminus \Gamma_{2:k}}) \sum_{t=0}^{\infty} \mathbb{P}_{\Gamma_1}(S_t = \Gamma_1, t < \tau_{A_n}). \quad (1.4)$$

Denote the first hitting time of a set $B \subseteq \mathbb{Z}^2$ by $\sigma_B = \inf\{j \geq 0 : S_j \in B\}$, or σ_x if $B = \{x\}$ for some $x \in \mathbb{Z}^2$. The first factor of (1.4) equals $\mathbb{P}_{\Gamma_2}(\sigma_o < \sigma_{\mathbb{Z}^2 \setminus \Gamma_{2:k}})$ because $\Gamma_2 \notin \{o\} \cup (\mathbb{Z}^2 \setminus \Gamma_{2:k})$. We calculate it as $f(2)$, where $f(i) = \mathbb{P}_{\Gamma_i}(\sigma_o < \sigma_{\mathbb{Z}^2 \setminus \Gamma_{2:k}})$ solves the system of difference equations

$$f(1) = 0, \quad f(k) = 1, \quad \text{and} \quad f(i) = \frac{1}{4}f(i+1) + \frac{1}{4}f(i-1), \quad 2 \leq i \leq k-1.$$

The solution of this system yields

$$\frac{1}{(2 + \sqrt{3})^{k-1}} \leq f(2) = \frac{2\sqrt{3}}{(2 + \sqrt{3})^{k-1} - (2 - \sqrt{3})^{k-1}} \leq \frac{1}{(2 + \sqrt{3})^{k-2}}. \quad (1.5)$$

Lastly, note that the second factor of (1.4) is the expected number of visits that S_t makes to Γ_1 before time τ_A , which equals the reciprocal of $\mathbb{P}_{\Gamma_1}(\tau_{\Gamma_1} < \tau_{A_n})$. Since Γ_1 is adjacent to τ_{A_n} , this probability is at least $\frac{1}{4}$, hence

$$1 \leq \sum_{t=0}^{\infty} \mathbb{P}_{\Gamma_1}(S_t = \Gamma_1, t < \tau_{A_n}) \leq 4.$$

Combining the preceding bounds, we conclude that, for all sufficiently large n ,

$$\frac{cn^{-b}}{(2 + \sqrt{3})^{k-1}} \leq \mathbb{H}_{A_n}(o) \leq \frac{4}{(2 + \sqrt{3})^{k-2}}.$$

Substituting $k = 2(1 - o_n(1))n$ and simplifying, we obtain

$$(2 + \sqrt{3})^{-2(1+o_n(1))n} \leq \mathbb{H}_{A_n}(o) \leq (2 + \sqrt{3})^{-2(1-o_n(1))n},$$

which implies

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log \mathbb{H}_{A_n}(o) = 2 \log(2 + \sqrt{3}).$$




FIGURE 3. A square spiral. The shortest path Γ (red) from Γ_1 to the origin, which first hits A_n (black and gray dots) at the origin, has a length of approximately $2n$. Some elements (gray dots) of A_n could be used to continue the spiral pattern (indicated by the black dots), but are presently placed to facilitate a calculation in Example 1.6.

We conclude the discussion of our main results by stating an estimate of hitting probabilities of the form $\mathbb{P}_x(\tau_{\partial A_d} < \tau_A)$, for $x \in A$ and where A_d is the set of all elements of \mathbb{Z}^2 within distance d of A ; we will call these *escape probabilities* from A . Among n -element subsets A of \mathbb{Z}^2 , when d is sufficiently large relative to the diameter of A , the greatest escape probability to a distance d from A is at most the reciprocal of $\log d$, up to a constant factor. We find that, in general, it is at least this much, up to an n -dependent factor.

Theorem 5. *There exists a universal positive constant c such that, if A is a finite subset of \mathbb{Z}^2 with $n \geq 2$ elements and if $d \geq 2 \text{diam}(A)$, then, for any $x \in A$,*

$$\mathbb{P}_x(\tau_{\partial A_d} < \tau_A) \geq \frac{c \mathbb{H}_A(x)}{n \log d}. \quad (1.6)$$

In particular,

$$\max_{x \in A} \mathbb{P}_x (\tau_{\partial A_d} < \tau_A) \geq \frac{c}{n^2 \log d}. \quad (1.7)$$

In the context of the HAT dynamics, we will use (1.7) to control the transport step, ultimately producing the $\log d$ timescale appearing in Theorem 1. In the setting of its application, A and d will respectively represent a subset of a HAT configuration and the separation of A from the rest of the configuration. Reflecting the potential sparsity of HAT configurations, d may be arbitrarily large relative to n .

Organization. HAT motivates the development of new estimates of harmonic measure and escape probabilities. We attend to these estimates in Section 3, after we provide a conceptual overview of the proofs of Theorems 1 and 2 in Section 2. To analyze configurations of large diameter, we will decompose them into well separated “clusters,” using a construction introduced in Section 5 and used throughout Section 6. The estimates of Section 3 control the activation and transport steps of the dynamics and serve as the critical inputs to Section 6, in which we analyze the “collapse” of HAT configurations. We then identify the class of configurations to which the HAT dynamics can return and prove the existence of a stationary distribution supported on this class; this is the primary focus of Section 7. The final section, Section 8, uses an exponential tail bound on the diameter of configurations under the stationary distribution—a result we obtain at the end of Section 7—to show that the center of mass process, properly rescaled, converges in distribution to two-dimensional Brownian motion.

Forthcoming Notation. We will denote expectation with respect to \mathbf{P}_U , the law of HAT from the configuration U , by \mathbf{E}_U ; the indicator of an event E by $\mathbf{1}_E$ or $\mathbf{1}(E)$; the Euclidean disk of radius r about x by $D_x(r) = \{y \in \mathbb{Z}^2 : |x-y| < r\}$, or $D(r)$ if $x = o$; its boundary by $C_x(r) = \partial D_x(r)$, or $C(r)$ if $x = o$; the radius of a finite set $A \subset \mathbb{Z}^2$ by $\text{rad}(A) = \sup\{|x| : x \in A\}$; the R -fattening of A by $A_R = \{x \in \mathbb{Z}^2 : \text{dist}(A, x) \leq R\}$; and the minimum of two random times τ_1 and τ_2 by $\tau_1 \wedge \tau_2$.

When we refer to a (universal) constant, we will always mean a positive real number. When we cite standard results from [Law13] and [Pop21], we will write $O(g)$ to denote a function f that uniformly satisfies $|f| \leq cg$ for an implicit constant c . However, in all other instances, f will be nonnegative and we will simply mean the estimate $f \leq cg$ by $O(g)$. We will include a subscript to indicate the dependence of the implicit constant on a parameter, e.g., $f = O_n(g)$. We will use $\Omega(g)$ and $\Omega_n(g)$ for the reverse estimate. We will use $o_n(1)$ to denote a nonnegative quantity that is at most 1 and which tends to 0 as $n \rightarrow \infty$, e.g., n^{-1} for $n \geq 1$.

Acknowledgements. J.C. thanks Joseph Slote for useful discussions concerning Conjecture 1.5 and Example 1.6. A.H. thanks Dmitry Belyaev for helpful discussions concerning the behavior of HAT configurations with well separated clusters and for simulating HAT dynamics.

2. CONCEPTUAL OVERVIEW

2.1. Estimating the collapse time and proving the existence of the stationary distribution. Before providing precise details, we discuss some of the key steps in the proofs of Theorems 1 and 2. Since the initial configuration U of n particles is arbitrary, it will be advantageous to decompose any such configuration into clusters such that the separation between any two clusters is at least exponentially large relative to their diameters. As we will show later, we can always find such a clustering when the diameter of U is large enough in terms of n . For the purpose of illustration, let us start by assuming that U consists of just two clusters with separation d and hence the individual diameters of the clusters are no greater than $\log d$ (Figure 4).

The first step in our analysis is to show that in time comparable to $\log d$, the diameter of U will shrink to $\log d$. This is the phenomenon we call *collapse*. Theorem 4 implies that every particle with positive harmonic measure has harmonic measure of at least $e^{-cn \log n}$. In particular, the particle in each cluster with the greatest escape probability from that cluster has at least this harmonic measure. Our choice of clustering will ensure that each cluster has positive harmonic measure. Accordingly, we will treat each cluster as the




FIGURE 4. Exponentially separated clusters.

entire configuration and Theorem 5 will imply that the greatest escape probability from each cluster will be at least $(\log d)^{-1}$, up to a factor depending upon n .

Together, these results will imply that, in $O_n(\log d)$ steps, with a probability depending only upon n , all the particles from one of the clusters in Figure 4 will move to the other cluster. Moreover, since the diameter of a cluster grows at most linearly in time, the final configuration will have diameter which is no greater than the diameter of the surviving cluster plus $O_n(\log d)$. Essentially, we will iterate this estimate—by clustering anew the surviving cluster of Figure 4—each time obtaining a cluster with a diameter which is the logarithm of the original diameter, until d becomes smaller than a deterministic function θ_{4n} , which is approximately the $4n^{\text{th}}$ iterated exponential of cn , for a constant c .

Let us denote the corresponding stopping time by \mathcal{T} (below θ_{4n}). In the setting of the application, there may be multiple clusters and we collapse them one by one, reasoning as above. If any such collapse step fails, we abandon the experiment and repeat it. Of course, with each failure, the set we attempt to collapse may have a diameter which is additively larger by $O_n(\log d)$. Ultimately, our estimates allow us to conclude that the attempt to collapse is successful within the first $(\log d)^{1+o_n(1)}$ tries with a high probability.

The preceding discussion roughly implies the following result, uniformly in the initial configuration U :

$$\mathbf{P}_U \left(\mathcal{T}(\text{below } \theta_{4n}) \leq (\log d)^{1+o_n(1)} \right) \geq 1 - e^{-n}.$$

At this stage, we prove that, given any configuration \widehat{U} and any configuration $\widehat{V} \in \widehat{\text{NonIso}}(n)$, if K is sufficiently large in terms of n and the diameters of \widehat{U} and \widehat{V} , then

$$\mathbf{P}_{\widehat{U}} \left(\mathcal{T}(\text{hits } \widehat{V}) \leq K^5 \right) \geq 1 - e^{-K},$$

where $\mathcal{T}(\text{hits } \widehat{V})$ is the first time the configuration is \widehat{V} . To prove this estimate, we split it into two, simpler estimates. Specifically, we show that the particles of \widehat{U} form a line segment of length n in K^4 steps with high probability, and we prove by induction on n that any other non-isolated configuration \widehat{V} is reachable from the line segment in K^5 steps, with high probability. In addition to implying irreducibility of the HAT dynamics on $\widehat{\text{NonIso}}(n)$, we use this result to obtain a finite upper bound on the expected return time to any non-isolated configuration (i.e., it proves the positive recurrence of HAT on $\widehat{\text{NonIso}}(n)$). Irreducibility and positive recurrence on $\widehat{\text{NonIso}}(n)$ imply the existence and uniqueness of the stationary distribution.

2.2. Improved estimates of hitting probabilities for sparse sets. HAT configurations may include subsets with large diameters relative to the number of elements they contain, and in this sense they are sparse. Two such cases are depicted in Figure 5. A key component of the proofs of Theorems 4 and 5 is a method which improves two standard estimates of hitting probabilities when applied to sparse sets, as summarized by Table 1.

For the scenario depicted in Figure 5 (left), we estimate the probability that a random walk from $x \in C(\frac{R}{3})$ hits the origin before any element of $A \setminus \{o\}$. Since $C(R)$ separates x from $A \setminus \{o\}$, this probability is at least $\mathbb{P}_x(\tau_o < \tau_{C(R)})$. We can calculate this lower bound by combining the fact that the potential kernel (defined in Section 3) is harmonic away from the origin with the optional stopping theorem (e.g., Proposition 1.6.7




FIGURE 5. Sparse sets like ones which appear in the proofs of Theorems 4 (left) and 5 (right). The elements of A are represented by dark green dots. On the left, $A \setminus \{o\}$ is a subset of $D(R)^c$. On the right, A is a subset of $D(r)$ and A_R , the R -fattening of A (shaded green), is a subset of $D(R+r)$. The figure is not to scale, as $R \geq e^n$ on the left, while $R \geq e^r$ on the right.

TABLE 1. Summary of improvements to standard estimates in sparse settings. The origin is denoted by o and A_R denotes the set of all points in \mathbb{Z}^d within a distance R of A .

Setting	Quantity	Standard estimate	New estimate
Fig. 5 (left), $R \geq e^n$	$\mathbb{P}_x(\tau_o < \tau_{A \cap D(R)^c})$	$\Omega\left(\frac{1}{\log R}\right)$	$\Omega\left(\frac{1}{n}\right)$
Fig. 5 (right), $R \geq e^r$	$\mathbb{P}_x(\tau_{\partial A_R} < \tau_A)$	$\Omega_n\left(\frac{1}{\log R}\right)$	$\Omega_n\left(\frac{\log r}{\log R}\right)$

of [Law13]):

$$\mathbb{P}_x(\tau_o < \tau_{C(R)}) = \frac{\log R - \log |x| + O(R^{-1})}{\log R + O(R^{-1})}.$$

This implies $\mathbb{P}_x(\tau_o < \tau_{A \cap D(R)^c}) = \Omega\left(\frac{1}{\log R}\right)$, since $x \in C\left(\frac{R}{3}\right)$.

We can improve the lower bound to $\Omega\left(\frac{1}{n}\right)$ by using the sparsity of A . We define the random variable $W = \sum_{y \in A \setminus \{o\}} \mathbf{1}(\tau_y < \tau_o)$ and write

$$\mathbb{P}_x(\tau_o < \tau_{A \setminus \{o\}}) = \mathbb{P}_x(W = 0) = 1 - \frac{\mathbb{E}_x W}{\mathbb{E}_x[W \mid W > 0]}.$$

We will show that $\mathbb{E}_x[W \mid W > 0] \geq \mathbb{E}_x W + \delta$ for some δ which is uniformly positive in A and n . We will be able to find such a δ because random walk from x hits a given element of $A \setminus \{o\}$ before o with a probability of at most $1/2$, so conditioning on $\{W > 0\}$ effectively increases W by $1/2$. Then

$$\mathbb{P}_x(\tau_o < \tau_{A \setminus \{o\}}) \geq 1 - \frac{\mathbb{E}_x W}{\mathbb{E}_x W + \delta} \geq 1 - \frac{n}{n + \delta} = \Omega\left(\frac{1}{n}\right).$$

The second inequality follows from the monotonicity of $\frac{\mathbb{E}_x W}{\mathbb{E}_x W + \delta}$ in $\mathbb{E}_x W$ and the fact that $|A| \leq n$, so $\mathbb{E}_x W \leq n$. This is a better lower bound than $\Omega\left(\frac{1}{\log R}\right)$ when R is at least e^n .

A variation of this method also improves a standard estimate for the scenario depicted in Figure 5 (right). In this case, we estimate the probability that a random walk from $x \in C(2r)$ hits ∂A_R before A , where A is contained in $D(r)$ and A_R consists of all elements of \mathbb{Z}^2 within a distance $R \geq e^r$ of A . We can bound below this probability by using the fact that

$$\mathbb{P}_x(\tau_{\partial A_R} < \tau_A) \geq \mathbb{P}_x(\tau_{C(R+r)} < \tau_{C(r)}).$$

A standard calculation with the potential kernel of random walk (e.g., Exercise 1.6.8 of [Law13]) shows that this lower bound is $\Omega_n(\frac{1}{\log R})$, since $R \geq e^r$ and $r = \Omega(n^{1/2})$.

We can improve the lower bound to $\Omega_n(\frac{\log r}{\log R})$ by using the sparsity of A . We define $W' = \sum_{y \in A} \mathbf{1}(\tau_y < \tau_{\partial A_R})$ and write

$$\mathbb{P}_x(\tau_{\partial A_R} < \tau_A) = 1 - \frac{\mathbb{E}_x W'}{\mathbb{E}_x[W' \mid W' > 0]} \geq 1 - \frac{n\alpha}{1 + (n-1)\beta},$$

where α bounds above $\mathbb{P}_x(\tau_y < \tau_{\partial A_R})$ and β bounds below $\mathbb{P}_z(\tau_y < \tau_{\partial A_R})$, uniformly for $x \in C(2r)$ and distinct $y, z \in A$. We will show that $\alpha \leq \beta$ and $\beta \leq 1 - \frac{\log(2r)}{\log R}$. The former is plausible because $|x - y|$ is at least as great as $|y - z|$; the latter because $\text{dist}(z, A) \geq R$ while $|y - z| \leq 2r$, and because of (3.9). We apply these facts to the preceding display to conclude

$$\mathbb{P}_x(\tau_{\partial A_R} < \tau_A) \geq n^{-1}(1 - \beta) = \Omega_n(\frac{\log r}{\log R}).$$

This is a better lower bound than $\Omega_n(\frac{1}{\log R})$ because r can be as large as $\log R$.

In summary, by analyzing certain conditional expectations, we can better estimate hitting probabilities for sparse sets than we can by applying standard results. This approach may be useful in obtaining other sparse analogues of hitting probability estimates.

3. HARMONIC MEASURE ESTIMATES

The purpose of this section is to prove Theorem 4. We will describe the proof strategy in Section 3.1, before stating several estimates in Section 3.2 that streamline the presentation of the proof in Section 3.3.

Consider a subset A of \mathbb{Z}^2 with $n \geq 2$ elements, which satisfies $\mathbb{H}_A(o) > 0$ (i.e., $A \in \mathcal{H}_n$). We frame the proof of (1.2) in terms of advancing a random walk from infinity to the origin in three or four stages, while avoiding all other elements of A . We continue to use the phrase ‘‘random walk from infinity’’ to refer to the limiting hitting probability of a fixed, finite subset of \mathbb{Z}^2 by a random walk from $x \in \mathbb{Z}^2$ as $|x| \rightarrow \infty$. We will write \mathbb{P}_∞ as a shorthand for this limit.

The stages of advancement are defined in terms of a sequence of annuli which partition \mathbb{Z}^2 . Denote by $\mathcal{A}(r, R) = D(R) \setminus D(r)$ the annulus with inner radius r and outer radius R . We will frequently need to reference the subset of A which lies within or beyond a disk. We denote $A_{<r} = A \cap D(r)$ and $A_{\geq r} = A \cap D(r)^c$. Define radii R_1, R_2, \dots and annuli $\mathcal{A}_1, \mathcal{A}_2, \dots$ through $R_1 = 10^5$, and $R_\ell = R_1^\ell$ and $\mathcal{A}_\ell = \mathcal{A}(R_\ell, R_{\ell+1})$ for $\ell \geq 1$. We fix $\delta = 10^{-2}$ for use in intermediate scales, like $C(\delta R_{\ell+1}) \subset \mathcal{A}_\ell$. Additionally, we denote by n_0, n_ℓ, m_ℓ , and $n_{>J}$ the number of elements of A in $D(R_1), \mathcal{A}_\ell, \mathcal{A}_\ell \cup \mathcal{A}_{\ell+1}$, and $D(R_{J+1})^c$, respectively.

We will split the proof of (1.2) into an easy case when $n_0 = n$ and a difficult case when $n_0 \neq n$. If $n_0 \neq n$, then $A_{\geq R_1}$ is nonempty and the following indices $I = I(A)$ and $J = J(A)$ are well defined:

$$\begin{aligned} I &= \min\{\ell \geq 1 : \mathcal{A}_\ell \text{ contains an element of } A \setminus \{o\}\}, \text{ and} \\ J &= \min\{\ell > I : \mathcal{A}_\ell \text{ contains no element of } A \setminus \{o\}\}. \end{aligned}$$

Figure 6 illustrates the definitions of I and J . We explain their roles in the following subsection.




FIGURE 6. The first annulus that intersects A (green dots) is \mathcal{A}_I ; the next empty annulus is \mathcal{A}_J .

3.1. Strategy for the proof of Theorem 4. This section outlines a proof of (1.2) by induction on n . The induction step is easy when $n_0 = n$ because it implies that A is contained in $D(10^5)$, hence $\mathbb{H}_A(o)$ is at least a universal positive constant. The following strategy concerns the difficult case when $n_0 \neq n$.

Stage 1: Advancing to $C(R_J)$. Assume $n_0 \neq n$ and $n \geq 3$. By the induction hypothesis, there is a universal constant c_1 such that the harmonic measure at the origin is at least $e^{-c_1 k \log k}$, for any set in \mathcal{H}_k , $1 \leq k < n$. Let $k = n_{>J} + 1$. Because a random walk from ∞ which hits the origin before $A_{\geq R_J}$ also hits $C(R_J)$ before A , the induction hypothesis applied to $A_{\geq R_J} \cup \{o\} \in \mathcal{H}_k$ implies that $\mathbb{P}_\infty(\tau_{C(R_J)} < \tau_A)$ is no smaller than exponential in $k \log k$. Note that $k < n$ because $A_{< R_{I+1}}$ has at least two elements by the definition of I .

The reason we advance the random walk to $C(R_J)$ instead of $C(R_{J-1})$ is that an adversarial choice of A could produce a “choke point” which likely dooms the walk to be intercepted by $A \setminus \{o\}$ in the second stage of advancement (Figure 7). To avoid a choke point when advancing to the boundary of a disk D , it suffices for the conditional hitting distribution of ∂D given $\{\tau_{\partial D} < \tau_A\}$ to be comparable to the uniform hitting distribution on ∂D . To prove this comparison, the annular region immediately beyond D and extending to a radius of, say, twice that of D , must be empty of A . This explains the need for exponentially growing radii and for \mathcal{A}_J to be empty of A .




FIGURE 7. An example of a choke point (left) and a strategy for avoiding it (right). The hitting distribution of a random walk conditioned to reach ∂D before A (green dots) may favor the avoidance of $A \cap D^c$ in a way which localizes the walk (e.g., as indicated by the dark red arc of ∂D) prohibitively close to $A \cap D$. The hitting distribution on $C(R_J)$ will be approximately uniform if the radii grow exponentially. The random walk can then avoid the choke point by “tunneling” through it (e.g., by passing through the tan-shaded region).

Stage 2: Advancing into \mathcal{A}_{I-1} . For notational convenience, assume $I \geq 2$ so that \mathcal{A}_{I-1} is defined; the argument is the same when $I = 1$. Each annulus \mathcal{A}_ℓ , $\ell \in \{I, \dots, J-1\}$, contains one or more elements of A , which the random walk must avoid on its journey to \mathcal{A}_{I-1} . Except in an easier sub-case, which we

address at the end of this subsection, we advance the walk into \mathcal{A}_{I-1} by building an overlapping sequence of rectangular and annular *tunnels*, through and between each annulus, which are empty of A and through which the walk can enter \mathcal{A}_{I-1} (Figure 8). Specifically, the walk reaches a particular subset Arc_{I-1} in \mathcal{A}_{I-1} at the conclusion of the tunneling process. We will define Arc_{I-1} in Lemma 3.3 as an arc of a circle in \mathcal{A}_{I-1} .

By the pigeonhole principle applied to the angular coordinate, for each $\ell \geq I + 1$, there is a sector of aspect ratio $m_\ell = n_\ell + n_{\ell-1}$, from the lower “ δ^{th} ” of \mathcal{A}_ℓ to that of $\mathcal{A}_{\ell-1}$, which contains no element of A (Figure 8). To reach the entrance of the analogous tunnel between $\mathcal{A}_{\ell-1}$ and $\mathcal{A}_{\ell-2}$, the random walk may need to circle the lower δ^{th} of $\mathcal{A}_{\ell-1}$. We apply the pigeonhole principle to the radial coordinate to conclude that there is an annular region contained in the lower δ^{th} of $\mathcal{A}_{\ell-1}$, with an aspect ratio of $n_{\ell-1}$, which contains no element of A .

The probability that the random walk reaches the annular tunnel before exiting the rectangular tunnel from \mathcal{A}_ℓ to $\mathcal{A}_{\ell-1}$ is no smaller than exponential in m_ℓ . Similarly, the random walk reaches the rectangular tunnel from $\mathcal{A}_{\ell-1}$ to $\mathcal{A}_{\ell-2}$ before exiting the annular tunnel in $\mathcal{A}_{\ell-1}$ with a probability no smaller than exponential in $n_{\ell-1}$. Overall, we conclude that the random walk reaches Arc_{I-1} without leaving the union of tunnels—and therefore without hitting an element of A —with a probability no smaller than exponential in $\sum_{\ell=I}^{J-1} n_\ell$.




FIGURE 8. Tunneling through nonempty annuli. We construct a contiguous series of sectors (tan) and annuli (blue) which contain no elements of A (green dots) and through which the random walk may advance from $C(R_{J-1})$ to $C(\delta R_{J-1})$ (dashed).

Stage 3: Advancing to $C(R_1)$. Figure 5 (left) essentially depicts the setting of the random walk upon reaching $x \in \text{Arc}_{I-1}$, except with $C(R_I)$ in the place of $C(R)$ and the circle containing Arc_{I-1} in the place of $C(\frac{R}{3})$, and except for the possibility that $D(R_1)$ contains other elements of A . Nevertheless, if the radius of Arc_{I-1} is at least e^n , then by pretending that $A_{<R_1} = \{o\}$, the method highlighted in Section 2.2 will show that $\mathbb{P}_x(\tau_{C(R_1)} < \tau_A) = \Omega(\frac{1}{n})$. A simple calculation will give the same lower bound (for a potentially smaller constant) in the case when the radius is less than e^n .

Stage 4: Advancing to the origin. Once the random walk reaches $C(R_1)$, we simply dictate a path for it to follow. There can be no more than $O(R_1^2)$ elements of $A_{<R_1}$, so there is a path of length $O(R_1^2)$ to the origin which avoids all other elements of A , and a corresponding probability of at least a constant that the random walk follows it.

Conclusion of Stages 1–4. The lower bounds from the four stages imply that there are universal constants c_1 through c_4 such that

$$\mathbb{H}_A(o) \geq e^{-c_1 k \log k - c_2 \sum_{\ell=I}^{J-1} n_\ell - \log(c_3 n) - \log c_4} \geq e^{-c_1 n \log n}.$$

It is easy to show that the second inequality holds if $c_1 \geq 8 \max\{1, c_2, \log c_3, \log c_4\}$, using the fact that $n - k = \sum_{\ell=I}^{J-1} n_\ell > 1$ and $\log n \geq 1$. We are free to adjust c_1 to satisfy this bound, because c_2 through c_4 do not depend on the induction hypothesis. This concludes the induction step.

A complication in Stage 2. If R_ℓ is not sufficiently large relative to m_ℓ , then we cannot tunnel the random walk through \mathcal{A}_ℓ into $\mathcal{A}_{\ell-1}$. We formalize this through the failure of the condition

$$\delta R_\ell > R_1(m_\ell + 1). \quad (3.1)$$

The problem is that, if (3.1) fails, then there are too many elements of A in \mathcal{A}_ℓ and $\mathcal{A}_{\ell-1}$, and we cannot guarantee that there is a tunnel between the annuli which avoids A .

Accordingly, we will stop *Stage 2* tunneling once the random walk reaches a particular subset Arc_{K-1} of a circle in \mathcal{A}_{K-1} , where \mathcal{A}_{K-1} is the outermost annulus which fails to satisfy (3.1). Specifically, we define K as:

$$K = \begin{cases} I, & \text{if (3.1) holds for } \ell \in \{I, \dots, J\}; \\ \min\{k \in \{I, \dots, J\} : (3.1) \text{ holds for } \ell \in \{k, \dots, J\}\}, & \text{otherwise.} \end{cases} \quad (3.2)$$

Informally, when $K = I$, the pigeonhole principle yields wide enough tunnels for the random walk to advance all the way to \mathcal{A}_{I-1} in *Stage 2*. When $K \neq I$, the tunnels become too narrow, so we must halt tunneling before the random walk reaches \mathcal{A}_{I-1} . However, this case is even simpler, as the failure of (3.1) for $\ell = K - 1$ implies that there is a path of length $O(\sum_{\ell=I}^{K-1} n_\ell)$ from Arc_{K-1} to the origin which otherwise avoids A . In this case, instead of proceeding to *Stages 3* and *4* as described above, the third and final stage consists of random walk from Arc_{K-1} following this path to the origin with a probability no smaller than exponential in $\sum_{\ell=I}^{K-1} n_\ell$.

Overall, if $K \neq I$, then *Stages 2* and *3* contribute a rate of $\sum_{\ell=I}^{J-1} n_\ell$. This rate is smaller than the one contributed by *Stages 2–4* when $K = I$, so the preceding conclusion holds.

3.2. Preparation for the proof of Theorem 4.

3.2.1. Input to Stage 1. Let $A \in \mathcal{H}_n$. Like in Section 3.1, we assume that $n_0 \neq n$ (i.e., $A_{\geq R_1} \neq \emptyset$) and defer the simpler complementary case to Section 3.3. Recall that the radii R_i must grow exponentially so that the conditional hitting distribution of $C(R_J)$ is comparable to the uniform distribution, thus avoiding potential choke points (Figure 7). The next two results accomplish this comparison. We state them in terms of the uniform distribution on $C(r)$, which we denote by μ_r .

Lemma 3.1. *Let $\varepsilon > 0$ and denote $\eta = \tau_{C(R)} \wedge \tau_{C(\varepsilon^2 R)}$. There is a constant c such that, if $\varepsilon \leq \frac{1}{100}$ and $R \geq 10\varepsilon^{-2}$, and if*

$$\min_{x \in C(\varepsilon R)} \mathbb{P}_x \left(\tau_{C(\varepsilon^2 R)} < \tau_{C(R)} \right) > \frac{1}{10}, \quad (3.3)$$

then, uniformly for $x \in C(\varepsilon R)$ and $y \in C(\varepsilon^2 R)$,

$$\mathbb{P}_x \left(S_\eta = y, \tau_{C(\varepsilon^2 R)} < \tau_{C(R)} \right) \geq c \mu_{\varepsilon^2 R}(y) \mathbb{P}_x \left(\tau_{C(\varepsilon^2 R)} < \tau_{C(R)} \right).$$

The proof, which is similar to that of Lemma 2.1 in [DPRZ06], approximates the hitting distribution of $C(\varepsilon^2 R)$ by the corresponding harmonic measure, which is comparable to the uniform distribution. The condition (3.3), and the assumptions on ε and R , are used to control the error of this approximation. We defer the proof to Section A.3, along with the proof of the following application of Lemma 3.1.

Lemma 3.2. *There is a constant c such that, for every $z \in C(R_J)$,*

$$\mathbb{P}_\infty(S_{\tau_{C(R_J)}} = z \mid \tau_{C(R_J)} < \tau_A) \geq c\mu_{R_J}(z). \quad (3.4)$$

Under the conditioning in (3.4), the random walk reaches $C(\delta R_{J+1})$ before hitting A . A short calculation shows that it typically proceeds to hit $C(R_J)$ before returning to $C(R_{J+1})$ (i.e., it satisfies (3.3) with R_{J+1} in the place of R and $\varepsilon^2 = 10^{-5}$). The inequality (3.4) then follows from Lemma 3.1.

3.2.2. Inputs to Stage 2. We continue to assume that $n_0 \neq n$, so that I, J , and K are well defined; the $n_0 = n$ case is easy and we address it in Section 3.3. In this subsection, we will prove an estimate of the probability that a random walk passes through annuli \mathcal{A}_{J-1} to \mathcal{A}_K without hitting A . First, in Lemma 3.3, we will identify a sequence of “tunnels” through the annuli, which are empty of A . Second, in Lemma 3.4 and Lemma 3.5, we will show that random walk traverses these tunnels through a series of rectangles, with a probability that is no smaller than exponential in the number of elements in $\mathcal{A}_K, \dots, \mathcal{A}_{J-1}$. We will combine these estimates in Lemma 3.6.

For each $\ell \in \mathbb{I} = \{K, \dots, J\}$, we define the annulus $\mathcal{B}_\ell = \mathcal{A}(R_{\ell-1}, \delta R_{\ell+1})$. The radial and angular tunnels from \mathcal{A}_ℓ into and around $\mathcal{A}_{\ell-1}$ will be subsets of \mathcal{B}_ℓ . The inner radius of \mathcal{B}_ℓ is at least R_1 because

$$\ell \in \mathbb{I} \implies R_\ell > \delta^{-1} R_1 (m_\ell + 1) \geq 10^7 \implies \ell \geq 2.$$

The first implication is due to (3.1) and (3.2); the second is due to the fact that $R_\ell = 10^{5\ell}$.




FIGURE 9. The regions identified in Lemma 3.3. The tan sectors and dark blue annuli are subsets of the overlapping annuli \mathcal{B}_ℓ and $\mathcal{B}_{\ell-1}$ that are empty of A .

The following lemma identifies subsets of \mathcal{B}_ℓ which are empty of A (Figure 9). Recall that $m_\ell = n_\ell + n_{\ell-1}$.

Lemma 3.3. *Let $\ell \in \mathbb{I}$. Denote $\varepsilon_\ell = (m_\ell + 1)^{-1}$ and $\delta' = \delta/10$. For every $\ell \in \mathbb{I}$, there is an angle $\vartheta_\ell \in [0, 2\pi)$ and a radius $a_{\ell-1} \in [10R_{\ell-1}, \delta'R_\ell)$ such that the following regions contain no element of A :*

- the sector of \mathcal{B}_ℓ subtending the angular interval $[\vartheta_\ell, \vartheta_\ell + 2\pi\varepsilon_\ell)$, hence the “middle third” sub-sector

$$\text{Sect}_\ell = [R_\ell, \delta'R_{\ell+1}) \times \left[\vartheta_\ell + \frac{2\pi}{3}\varepsilon_\ell, \vartheta_\ell + \frac{4\pi}{3}\varepsilon_\ell \right); \text{ and}$$

- the sub-annulus $\text{Ann}_{\ell-1} = \mathcal{A}(a_{\ell-1}, b_{\ell-1})$ of \mathcal{B}_ℓ , where we define

$$b_{\ell-1} = a_{\ell-1} + \Delta_{\ell-1} \quad \text{for} \quad \Delta_{\ell-1} = \delta' \varepsilon_\ell R_\ell$$

and, in particular, the circle $\text{Circ}_{\ell-1} = C\left(\frac{a_{\ell-1} + b_{\ell-1}}{2}\right)$ and the “arc”

$$\text{Arc}_{\ell-1} = \text{Circ}_{\ell-1} \cap \{x \in \mathbb{Z}^2 : \arg x \in [\vartheta_\ell, \vartheta_\ell + 2\pi\varepsilon_\ell]\}.$$

We take a moment to explain the parameters and regions. Aside from \mathcal{B}_ℓ , which overlaps \mathcal{A}_ℓ and $\mathcal{A}_{\ell-1}$, the subscripts of the regions indicate which annulus contains them (e.g., $\text{Sec}_\ell \subset \mathcal{A}_\ell$ and $\text{Ann}_{\ell-1} \subset \mathcal{A}_{\ell-1}$). The proof uses the pigeonhole principle to identify regions which contain none of the m_ℓ elements of A in \mathcal{B}_ℓ and $\text{Ann}_{\ell-1}$; this motivates our choice of ε_ℓ . A key aspect of Sec_ℓ is that it is separated from $\partial\mathcal{B}_\ell$ by a distance of at least $R_{\ell-1}$, which will allow us to position one end of a rectangular tunnel of width $R_{\ell-1}$ in Sec_ℓ without the tunnel exiting \mathcal{B}_ℓ . We also need the inner radius of $\text{Ann}_{\ell-1}$ to be at least $R_{\ell-1}$ greater than that of \mathcal{B}_ℓ , hence the lower bound on $a_{\ell-1}$. The other key aspect of $\text{Ann}_{\ell-1}$ is its overlap with $\text{Sec}_{\ell-1}$. The specific constants (e.g., $\frac{2\pi}{3}$, 10, and δ') are otherwise unimportant.

Proof of Lemma 3.3. Fix $\ell \in \mathbb{I}$. For $j \in \{0, \dots, m_\ell\}$, form the intervals

$$2\pi\varepsilon_\ell [j, j+1) \quad \text{and} \quad 10R_{\ell-1} + \Delta_{\ell-1} [j, j+1).$$

\mathcal{B}_ℓ contains at most m_ℓ elements of A , so the pigeonhole principle implies that there are j_1 and j_2 in this range and such that, if $\vartheta_\ell = j_1 2\pi\varepsilon_\ell$ and if $a_{\ell-1} = 10R_{\ell-1} + j_2 \Delta_{\ell-1}$, then

$$\mathcal{B}_\ell \cap \{x \in \mathbb{Z}^2 : \arg x \in [\vartheta_\ell, \vartheta_\ell + 2\pi\varepsilon_\ell]\} \cap A = \emptyset, \quad \text{and} \quad \mathcal{A}(a_{\ell-1}, a_{\ell-1} + \Delta_{\ell-1}) \cap A = \emptyset.$$

Because $\mathcal{B}_\ell \supseteq \text{Sec}_\ell$ and $\text{Ann}_{\ell-1} \supseteq \text{Arc}_{\ell-1}$, for these choices of ϑ_ℓ and $a_{\ell-1}$, we also have $\text{Sec}_\ell \cap A = \emptyset$ and $\text{Arc}_{\ell-1} \cap A = \emptyset$. \square

The next result bounds below the probability that the random walk tunnels “down” from Sec_ℓ to $\text{Arc}_{\ell-1}$. We state it without proof, as it is a simple consequence of the known fact that random walk from the “bulk” of a rectangle exits its far, small side with a probability which is no smaller than exponential in the aspect ratio of the rectangle (Lemma A.4). In this case, the aspect ratio is $O(m_\ell)$.

Lemma 3.4. *There is a constant c such that, for any $\ell \in \mathbb{I}$ and every $y \in \text{Sec}_\ell$,*

$$\mathbb{P}_y(\tau_{\text{Arc}_{\ell-1}} < \tau_A) \geq c^{m_\ell}.$$

The following lemma bounds below the probability that the random walk tunnels “around” $\text{Ann}_{\ell-1}$, from $\text{Arc}_{\ell-1}$ to $\text{Sec}_{\ell-1}$. (This result applies to $\ell \in \mathbb{I} \setminus \{K\} = \{K+1, \dots, J\}$ because Lemma 3.3 defines Sec_ℓ for $\ell \in \mathbb{I}$.) Like Lemma 3.4, we state it without proof because it is a simple consequence of Lemma A.4. Indeed, random walk from $\text{Arc}_{\ell-1}$ can reach $\text{Sec}_{\ell-1}$ without exiting $\text{Ann}_{\ell-1}$ by appropriately exiting each rectangle in a sequence of $O(m_\ell)$ rectangles of aspect ratio $O(1)$. Applying Lemma A.4 then implies (3.5).

Lemma 3.5. *There is a constant c such that, for any $\ell \in \mathbb{I} \setminus \{K\}$ and every $z \in \text{Arc}_{\ell-1}$,*

$$\mathbb{P}_z(\tau_{\text{Sec}_{\ell-1}} < \tau_A) \geq c^{m_\ell}. \tag{3.5}$$

The next result combines Lemma 3.4 and Lemma 3.5 to tunnel from \mathcal{A}_J into \mathcal{A}_{K-1} . Because the random walk tunnels from \mathcal{A}_ℓ to $\mathcal{A}_{\ell-1}$ with a probability no smaller than exponential in $m_\ell = n_\ell + n_{\ell-1}$, the bound in (3.6) is no smaller than exponential in $\sum_{\ell=K-1}^{J-1} n_\ell$ (recall that $n_J = 0$).

Lemma 3.6. *There is a constant c such that*

$$\mathbb{P}_{\mu_{R,J}}(\tau_{\text{Arc}_{K-1}} < \tau_A) \geq c^{\sum_{\ell=K-1}^{J-1} n_\ell}. \tag{3.6}$$

Proof. Denote by G the event

$$\{\tau_{\text{Arc}_{J-1}} < \tau_{\text{Sec}_{J-1}} < \tau_{\text{Arc}_{J-2}} < \tau_{\text{Sec}_{J-2}} < \cdots < \tau_{\text{Sec}_K} < \tau_{\text{Arc}_{K-1}} < \tau_A\}.$$

Lemma 3.4 and Lemma 3.5 imply that there is a constant c_1 such that

$$\mathbb{P}_z(G) \geq c_1^{\sum_{\ell=K-1}^{J-1} n^\ell} \text{ for } z \in C(R_J) \cap \text{Sec}_J. \quad (3.7)$$

The intersection of Sec_J and $C(R_J)$ subtends an angle of at least n_{J-1}^{-1} , so there is a constant c_2 such that

$$\mu_{R_J}(\text{Sec}_J) \geq c_2 n_{J-1}^{-1}. \quad (3.8)$$

The inequality (3.6) follows from $G \subseteq \{\tau_{\text{Arc}_{K-1}} < \tau_A\}$, and (3.7) and (3.8):

$$\mathbb{P}_{\mu_{R_J}}(\tau_{\text{Arc}_{K-1}} < \tau_A) \geq \mathbb{P}_{\mu_{R_J}}(G) \geq c_2 n_{J-1}^{-1} \cdot c_1^{\sum_{\ell=K-1}^{J-1} n^\ell} \geq c_3^{\sum_{\ell=K-1}^{J-1} n^\ell}.$$

For the third inequality, we take $c_3 = (c_1 c_2)^2$. \square

3.2.3. Inputs to Stage 3 when $K = I$. We continue to assume that $n_0 \neq n$, as the alternative case is addressed in Section 3.3. Additionally, we assume $K = I$. We briefly recall some important context. When $K = I$, at the end of Stage 2, the random walk has reached $\text{Circ}_{I-1} \subseteq \mathcal{A}_{I-1}$, where Circ_{I-1} is a circle with a radius in $[R_{I-1}, \delta' R_I]$. (Note that $I > 1$ when $K = I$ because I must then satisfy (3.1), so the radius of Circ_{I-1} is at least R_1 .) Since \mathcal{A}_I is the innermost annulus which contains an element of A , the random walk from Arc_{I-1} must simply reach the origin before hitting $A_{>R_I}$. In this subsection, we estimate this probability.

We will use the potential kernel associated with random walk on \mathbb{Z}^2 . We denote it by \mathbf{a} . It equals zero at the origin, is harmonic on $\mathbb{Z}^2 \setminus \{o\}$, and satisfies

$$\left| \mathbf{a}(x) - \frac{2}{\pi} \log |x| - \kappa \right| \leq \lambda |x|^{-2}, \quad (3.9)$$

where $\kappa \in (1.02, 1.03)$ is an explicit constant and λ is less than 0.06882 [KS04]. In some instances, we will want to apply \mathbf{a} to an element which belongs to $C(r)$. It will be convenient to denote, for $r > 0$,

$$\mathbf{a}'(r) = \frac{2}{\pi} \log r + \kappa.$$

We will need the following standard hitting probability estimate (see, for example, Proposition 1.6.7 of [Law13]), which we state as a lemma because we will use it in other sections as well.

Lemma 3.7. *Let $y \in D_x(r)$ for $r \geq 2(|x| + 1)$ and assume $y \neq o$. Then*

$$\mathbb{P}_y(\tau_o < \tau_{C_x(r)}) = \frac{\mathbf{a}'(r) - \mathbf{a}(y) + O\left(\frac{|x|+1}{r}\right)}{\mathbf{a}'(r) + O\left(\frac{|x|+1}{r}\right)}. \quad (3.10)$$

The implicit constants in the error terms are less than one.

If $R_I < e^{8n}$, then no further machinery is needed to prove the Stage 3 estimate.

Lemma 3.8. *There exists a constant c such that, if $R_I < e^{8n}$, then*

$$\mathbb{P}_\infty(\tau_{C(R_I)} < \tau_A \mid \tau_{\text{Circ}_{I-1}} < \tau_A) \geq \frac{c}{n}.$$

The bound holds because the random walk must exit $D(R_I)$ to hit $A_{\geq R_I}$. By a standard hitting estimate, the probability that the random walk hits the origin first is inversely proportional to $\log R_I$ which is $O(n)$ when $R_I < e^{8n}$.

Proof of Lemma 3.8. Uniformly for $y \in \text{Circ}_{I-1}$, we have

$$\mathbb{P}_y(\tau_{C(R_I)} < \tau_A) \geq \mathbb{P}_y(\tau_o < \tau_{C(R_I)}) \geq \frac{\alpha'(R_I) - \alpha'(\delta R_{I-1}) - \frac{1}{R_I} - \frac{1}{\delta R_I}}{\alpha'(R_I) + \frac{1}{R_I}} \geq \frac{1}{\alpha'(R_I)}. \quad (3.11)$$

The first inequality follows from the observation that $C(R_I)$ and $C(R_I)$ separate y from o and A . The second inequality is due to Lemma 3.7, where we have replaced $\alpha(y)$ by $\alpha'(\delta R_I) + \frac{1}{\delta R_I}$ using Lemma A.2 and the fact that $|y| \leq \delta R_I$. The third inequality follows from $\delta R_I \geq 10^3$. To conclude, we substitute $\alpha'(R_I) = \frac{2}{\pi} \log R_I + \kappa$ into (3.11) and use assumption that $R_I < e^{8n}$. \square

We will use the rest of this subsection to prove the bound of Lemma 3.8, but under the complementary assumption $R_I \geq e^{8n}$. This is one of the two estimates we highlighted in Section 2.2.

Next is a standard result, which enables us to express certain hitting probabilities in terms of the potential kernel. We include a short proof for completeness.

Lemma 3.9. *For any pair of points $x, y \in \mathbb{Z}^2$, define*

$$M_{x,y}(z) = \frac{\alpha(x-z) - \alpha(y-z)}{2\alpha(x-y)} + \frac{1}{2}.$$

Then $M_{x,y}(z) = \mathbb{P}_z(\sigma_y < \sigma_x)$.

Proof. Fix $x, y \in \mathbb{Z}^2$. Theorem 1.4.8 of [Law13] states that for any proper subset B of \mathbb{Z}^2 (including infinite B) and bounded function $F : \partial B \rightarrow \mathbb{R}$, the unique bounded function $f : B \cup \partial B \rightarrow \mathbb{R}$ which is harmonic in B and equals F on ∂B is $f(z) = \mathbb{E}_z[F(S_{\sigma_{\partial B}})]$. Setting $B = \mathbb{Z}^2 \setminus \{x, y\}$ and $F(z) = \mathbf{1}(z = y)$, we have $f(z) = \mathbb{P}_z(\sigma_y < \sigma_x)$. Since $M_{x,y}$ is bounded, harmonic on B , and agrees with f on ∂B , the uniqueness of f implies $M_{x,y}(z) = f(z)$. \square

The next two results partly implement the first estimate that we discussed in Section 2.2.

Lemma 3.10. *For any $z, z' \in \text{Circ}_{I-1}$ and $y \in D(R_I)^c$,*

$$\mathbb{P}_z(\tau_y < \tau_o) \leq \frac{1}{2} \quad \text{and} \quad \left| \mathbb{P}_z(\tau_y < \tau_o) - \mathbb{P}_{z'}(\tau_y < \tau_o) \right| \leq \frac{2}{\log R_I}. \quad (3.12)$$

The first inequality in (3.12) holds because z is appreciably closer to the origin than it is to y . The second inequality holds because a Taylor expansion of the numerator of $M_{o,y}(z) - M_{o,y}(z')$ shows that it is $O(1)$, while the denominator of $2\alpha(y)$ is at least $\log R_I$.

Proof of Lemma 3.10. By Lemma 3.9,

$$\mathbb{P}_z(\tau_y < \tau_o) = \frac{1}{2} + \frac{\alpha(z) - \alpha(y-z)}{2\alpha(y)}.$$

The first inequality in (3.12) then follows from $\alpha(y-z) \geq \alpha(z)$, which is due to (1) of Lemma A.1. This fact applies because $|y-z| \geq 2|z| \geq 4$. The first of these bounds holds because $|z| \leq \delta R_I + 1$ and $|y-z| \geq (1-\delta)R_I - 1$, since $\text{Circ}_{I-1} \subseteq D(\delta R_I)$; the second holds because the radius of Circ_{I-1} is at least R_1 , since $I > 1$ when $K = I$.

Using Lemma 3.9, the difference in (3.12) can be written as

$$\left| M_{o,y}(z) - M_{o,y}(z') \right| \leq \frac{|\alpha(z) - \alpha(z')|}{2\alpha(y)} + \frac{|\alpha(y-z) - \alpha(y-z')|}{2\alpha(y)}. \quad (3.13)$$

By Lemma A.2, in terms of $r = \text{rad}(\text{Circ}_{I-1})$, $\alpha(z)$ and $\alpha(z')$ differ from $\alpha'(r)$ by no more than r^{-1} . Since $r \geq R_1$, this implies $|\alpha(z) - \alpha(z')| \leq 2R_1^{-1}$. Concerning the denominator, $\alpha(y)$ is at least $\frac{2}{\pi} \log |y| \geq$

$\frac{2}{\pi} \log R_I$ by (2) of Lemma A.1 because $|y|$ is at least 1. We apply (3) of Lemma A.1 with $R = R_I$ and $r = \text{rad}(\text{Circ}_{I-1}) \leq \delta R_I$ to bound the numerator by $\frac{4}{\pi}$. Substituting these bounds into (3.13) gives

$$\left| \mathbb{P}_z(\tau_y < \tau_o) - \mathbb{P}_{z'}(\tau_y < \tau_o) \right| \leq \frac{1}{\frac{2}{\pi} R_I \log R_I} + \frac{1}{\log R_I} \leq \frac{2}{\log R_I}.$$

□

Label the k elements in $A_{\geq R_1}$ by x_i for $1 \leq i \leq k$. Then let $Y_i = \mathbf{1}(\tau_{x_i} < \tau_o)$ and $W = \sum_{i=1}^k Y_i$. In words, W counts the number of elements of $A_{\geq R_1}$ which have been visited before the random walk returns to the origin.

Lemma 3.11. *If $R_I \geq e^{8n}$, then, for all $z \in \text{Circ}_{I-1}$,*

$$\mathbb{E}_z[W \mid W > 0] \geq \mathbb{E}_z W + \frac{1}{4}. \quad (3.14)$$

The constant $\frac{1}{4}$ in (3.14) is unimportant, aside from being positive, independently of n . The inequality holds because random walk from Circ_{I-1} hits a given element of $A_{\geq R_1}$ before the origin with a probability of at most $\frac{1}{2}$. Consequently, given that some such element is hit, the conditional expectation of W is essentially larger than its unconditional one by a constant.

Proof of Lemma 3.11. Fix $z \in \text{Circ}_{I-1}$. When $\{W > 0\}$ occurs, some labeled element, x_f , is hit first. After τ_{x_f} , the random walk may proceed to hit other x_i before returning to Circ_{I-1} at a time $\eta = \min \{t \geq \tau_{x_f} : S_t \in \text{Circ}_{I-1}\}$. Let \mathcal{V} be the collection of labeled elements that the walk visits before time η , $\{i : \tau_{x_i} < \eta\}$. In terms of \mathcal{V} and η , the conditional expectation of W is

$$\mathbb{E}_z[W \mid W > 0] = \mathbb{E}_z \left[|\mathcal{V}| + \mathbb{E}_{S_\eta} \sum_{i \notin \mathcal{V}} Y_i \mid W > 0 \right]. \quad (3.15)$$

Let V be a nonempty subset of the labeled elements and let $z' \in \text{Circ}_{I-1}$. We have

$$\left| \mathbb{E}_z \sum_{i \notin V} Y_i - \mathbb{E}_{z'} \sum_{i \notin V} Y_i \right| \leq \frac{2n}{\log R_I} \leq \frac{1}{4}.$$

The first inequality is due to Lemma 3.10 and the fact that there are at most n labeled elements outside of V . The second inequality follows from the assumption that $R_I \geq e^{8n}$.

We use this bound to replace S_η in (3.15) with z :

$$\mathbb{E}_z[W \mid W > 0] \geq \mathbb{E}_z \left[|\mathcal{V}| + \mathbb{E}_z \sum_{i \notin \mathcal{V}} Y_i \mid W > 0 \right] - \frac{1}{4}. \quad (3.16)$$

By Lemma 3.10, $\mathbb{P}_z(\tau_{x_i} < \tau_o) \leq \frac{1}{2}$. Accordingly, for a nonempty subset V of labeled elements,

$$\mathbb{E}_z \sum_{i \notin V} Y_i \geq \mathbb{E}_z W - \frac{1}{2}|V|.$$

Substituting this into the inner expectation of (3.16), we find

$$\begin{aligned} \mathbb{E}_z[W \mid W > 0] &\geq \mathbb{E}_z \left[|\mathcal{V}| + \mathbb{E}_z W - \frac{1}{2}|\mathcal{V}| \mid W > 0 \right] - \frac{1}{4} \\ &\geq \mathbb{E}_z W + \mathbb{E}_z \left[\frac{1}{2}|\mathcal{V}| \mid W > 0 \right] - \frac{1}{4}. \end{aligned}$$

Since $\{W > 0\} = \{|\mathcal{V}| \geq 1\}$, this lower bound is at least $\mathbb{E}_z W + \frac{1}{4}$. □

We use the preceding lemma to prove the analogue of Lemma 3.8 when $R_I \geq e^{8n}$. The proof uses the method highlighted in Section 2.2 and Figure 5 (left).

Lemma 3.12. *There exists a constant c such that, if $R_I \geq e^{8n}$, then*

$$\mathbb{P}_\infty(\tau_{C(R_1)} < \tau_A \mid \tau_{\text{Circ}_{I-1}} < \tau_A) \geq \frac{c}{n}. \quad (3.17)$$

Proof. Conditionally on $\{\tau_{\text{Circ}_{I-1}} < \tau_A\}$, let the random walk hit Circ_{I-1} at z . Denote the positions of the $k \leq n$ particles in $A_{\geq R_1}$ as x_i for $1 \leq i \leq k$. Let $Y_i = \mathbf{1}(\tau_{x_i} < \tau_o)$ and $W = \sum_{i=1}^k Y_i$, just as we did for Lemma 3.11. The claimed bound (3.17) follows from

$$\mathbb{P}_z(\tau_A < \tau_{C(R_1)}) \leq \mathbb{P}_z(W > 0) = \frac{\mathbb{E}_z W}{\mathbb{E}_z[W \mid W > 0]} \leq \frac{\mathbb{E}_z W}{\mathbb{E}_z W + 1/4} \leq \frac{n}{n + 1/4} \leq 1 - \frac{1}{5n}.$$

The first inequality follows from the fact that $C(R_1)$ separates z from the origin. The second inequality is due to Lemma 3.11, which applies because $R_I \geq e^{8n}$. Since the resulting expression increases with $\mathbb{E}_z W$, we obtain the third inequality by substituting n for $\mathbb{E}_z W$, as $\mathbb{E}_z W \leq n$. The fourth inequality follows from $n \geq 1$. \square

3.2.4. Inputs to Stage 4 when $K = I$ and Stage 3 when $K \neq I$. The results in this subsection address the last stage of advancement in the two sub-cases of the case $n_0 \neq n$: $K = I$ and $K \neq I$. In the former sub-case, the random walk has reached $C(R_1)$; in the latter sub-case, it has reached Circ_{K-1} . Both sub-cases will be addressed by corollaries of the following, known geometric fact, stated in a form convenient for our purposes.

Let \mathbb{Z}^{2*} be the graph with vertex set \mathbb{Z}^2 and with an edge between distinct x and y in \mathbb{Z}^2 when x and y differ by at most 1 in each coordinate. For $B \subseteq \mathbb{Z}^2$, we will define the $*$ -exterior boundary of B by:

$$\begin{aligned} \partial_{\text{ext}}^* B &= \{x \in \mathbb{Z}^2 : x \text{ is adjacent in } \mathbb{Z}^{2*} \text{ to some } y \in B, \\ &\text{and there is a path from } \infty \text{ to } x \text{ disjoint from } B\}. \end{aligned} \quad (3.18)$$

Lemma 3.13. *Let $A \in \mathcal{H}_n$ and $r > 0$. From any $x \in C(r) \setminus A$, there is a path Γ in $(A \setminus \{o\})^c$ from $\Gamma_1 = x$ to $\Gamma_{|\Gamma|} = o$ with a length of at most $10 \max\{r, n\}$. Moreover, if $A \subseteq D(r)$, then Γ lies in $D(r + 2)$.*

We choose the constant factor of 10 for convenience; it has no special significance. We use a radius of $r + 2$ in $D(r + 2)$ to contain the boundary of $D(r)$ in \mathbb{Z}^{2*} .

Proof of Lemma 3.13. Let $\{B_\ell\}_\ell$ be the collection of $*$ -connected components of $A \setminus \{o\}$. By Lemma 2.23 of [Kes86] (alternatively, Theorem 4 of [Tim13]), because B_ℓ is finite and $*$ -connected, $\partial_{\text{ext}}^* B_\ell$ is connected.

Fix $r > 0$ and $x \in C(r) \setminus A$. Let Γ be the shortest path from x to the origin. If Γ is disjoint from $A \setminus \{o\}$, then we are done, as $|\Gamma|$ is no greater than $2r$. Otherwise, let ℓ_1 be the label of the first $*$ -connected component intersected by Γ . Let i and j be the first and last indices such that Γ intersects $\partial_{\text{ext}}^* B_{\ell_1}$, respectively. Because $\partial_{\text{ext}}^* B_{\ell_1}$ is connected, there is a path Λ in $\partial_{\text{ext}}^* B_{\ell_1}$ from Γ_i to Γ_j . We then edit Γ to form Γ' as

$$\Gamma' = (\Gamma_1, \dots, \Gamma_{i-1}, \Lambda_1, \dots, \Lambda_{|\Lambda|}, \Gamma_{j+1}, \dots, \Gamma_{|\Gamma|}).$$

If Γ' is disjoint from $A \setminus \{o\}$, then we are done, as Γ' is contained in the union of Γ and $\bigcup_\ell \partial_{\text{ext}}^* B_\ell$. Since $\bigcup_\ell B_\ell$ has at most n elements, $\bigcup_\ell \partial_{\text{ext}}^* B_\ell$ has at most $8n$ elements. Accordingly, the length of Γ' is at most $2r + 8n \leq 10 \max\{r, n\}$. Otherwise, if Γ' intersects another $*$ -connected component of $A \setminus \{o\}$, we can simply relabel the preceding argument to continue inductively and obtain the same bound.

Lastly, if $A \subseteq D(r)$, then $\bigcup_\ell \partial_{\text{ext}}^* B_\ell$ is contained in $D(r + 2)$. Since Γ is also contained in $D(r + 2)$, this implies that Γ' is contained in $D(r + 2)$. \square

Lemma 3.13 implies two other results. The first addresses *Stage 4* when $K = I$. By the definition of K (3.2), I must satisfy (3.1) when $K = I$, which implies that $I > 1$, hence $A_{\geq R_1} \subseteq D(R_1 + 2)^c$. The next result follows from this observation and the fact that $|A_{< R_1}| = O(R_1^2)$.

Lemma 3.14. *There is a constant c such that*

$$\mathbb{P}_\infty(\tau_o \leq \tau_A \mid \tau_{C(R_1)} < \tau_A) \geq c. \quad (3.19)$$

Lemma 3.13 also addresses *Stage 3* when $K \neq I$.

Lemma 3.15. *Assume that $n_0 = 1$ and $K \neq I$. There is a constant c such that*

$$\mathbb{P}_\infty(\tau_o \leq \tau_A \mid \tau_{\text{Circ}_{K-1}} < \tau_A) \geq c^{\sum_{\ell=I}^{K-1} n_\ell}. \quad (3.20)$$

The bound (3.20) follows from Lemma 3.13 because $K \neq I$ implies that the radius r of Circ_{K-1} is at most a constant factor times $|A_{< r}|$. Lemma 3.13 then implies that there is a path Γ from Circ_{K-1} to the origin with a length of $O(|A_{< r}|)$, which remains in $D(r + 2)$ and otherwise avoids the elements of $A_{< r}$. In fact, because Circ_{K-1} is a subset of Ann_{K-1} , which contains no elements of A , by remaining in $D(r + 2)$, Γ avoids $A_{\geq r}$ as well. This implies (3.20).

3.3. Proof of Theorem 4. The proof is by induction on n . Since (1.2) clearly holds for $n = 1$ and $n = 2$, we assume $n \geq 3$. Let $A \in \mathcal{H}_n$. There are three cases: $n_0 = n$, $n_0 \neq n$ and $K = I$, and $n_0 \neq n$ and $K \neq I$. The first of these cases is easy: When $n_0 = n$, A is contained in $D(R_1)$, so Lemma 3.14 implies that $\mathbb{H}_A(o)$ is at least a universal constant. Accordingly, in what follows, we assume that $n_0 \neq n$ and address the two sub-cases $K = I$ and $K \neq I$.

First sub-case: $K = I$. If $K = I$, then we write

$$\mathbb{H}_A(o) = \mathbb{P}_\infty(\tau_o \leq \tau_A) \geq \mathbb{P}_\infty(\tau_{C(R_J)} < \tau_{\text{Circ}_{I-1}} < \tau_{C(R_1)} < \tau_o \leq \tau_A).$$

Because $C(R_J)$, Circ_{I-1} , and $C(R_1)$ respectively separate Circ_{I-1} , $C(R_1)$, and the origin from ∞ , we can express the lower bound as the following product:

$$\begin{aligned} \mathbb{H}_A(o) &\geq \mathbb{P}_\infty(\tau_{C(R_J)} < \tau_A) \times \mathbb{P}_\infty(\tau_{\text{Circ}_{I-1}} < \tau_A \mid \tau_{C(R_J)} < \tau_A) \\ &\quad \times \mathbb{P}_\infty(\tau_{C(R_1)} < \tau_A \mid \tau_{\text{Circ}_{I-1}} < \tau_A) \times \mathbb{P}_\infty(\tau_o \leq \tau_A \mid \tau_{C(R_1)} < \tau_A). \end{aligned} \quad (3.21)$$

We address the four factors of (3.21) in turn. First, by the induction hypothesis, there is a constant c_1 such that

$$\mathbb{P}_\infty(\tau_{C(R_J)} < \tau_A) \geq e^{-c_1 k \log k},$$

where $k = n_{> J} + 1$. Second, by the strong Markov property applied to $\tau_{C(R_J)}$ and Lemma 3.2, and then by Lemma 3.6, there are constants c_2 and c_3 such that

$$\mathbb{P}_\infty(\tau_{\text{Circ}_{I-1}} < \tau_A \mid \tau_{C(R_J)} < \tau_A) \geq c_2 \mathbb{P}_{\mu_{R_J}}(\tau_{\text{Arc}_{I-1}} < \tau_A) \geq e^{-c_3 \sum_{\ell=I}^{J-1} n_\ell}. \quad (3.22)$$

Third and fourth, by Lemma 3.8 and Lemma 3.12, and by Lemma 3.14, there are constants c_4 and c_5 such that

$$\mathbb{P}_\infty(\tau_{C(R_1)} \leq \tau_A \mid \tau_{\text{Circ}_{I-1}} < \tau_A) \geq (c_4 n)^{-1} \quad \text{and} \quad \mathbb{P}_\infty(\tau_o \leq \tau_A \mid \tau_{C(R_1)} \leq \tau_A) \geq c_5.$$

Substituting the preceding bounds into (3.21) completes the induction step for this sub-case:

$$\mathbb{H}_A(o) \geq e^{-c_1 k \log k - c_3 \sum_{\ell=I}^{J-1} n_\ell - \log(c_4 n) + \log c_5} \geq e^{-c_1 n \log n}.$$

The second inequality follows from $n - k = \sum_{\ell=I}^{J-1} n_\ell > 1$ and $\log n \geq 1$, and from potentially adjusting c_1 to satisfy $c_1 \geq 8 \max\{1, c_3, \log c_4, -\log c_5\}$. We are free to adjust c_1 in this way, since the other constants do not arise from the use of the induction hypothesis.

Second sub-case: $K \neq I$. If $K \neq I$, then we write $\mathbb{H}_A(o) \geq \mathbb{P}_\infty(\tau_{C(R_J)} < \tau_{\text{Circ}_{K-1}} < \tau_o \leq \tau_A)$. Because $C(R_J)$ and Circ_{K-1} separate Circ_{K-1} and the origin from ∞ , we can express the lower bound as:

$$\mathbb{H}_A(o) \geq \mathbb{P}_\infty(\tau_{C(R_J)} < \tau_A) \times \mathbb{P}_\infty(\tau_{\text{Circ}_{K-1}} < \tau_A \mid \tau_{C(R_J)} < \tau_A) \times \mathbb{P}_\infty(\tau_o \leq \tau_A \mid \tau_{\text{Circ}_{K-1}} < \tau_A). \quad (3.23)$$

As in the first sub-case, the first factor is addressed by the induction hypothesis and the lower bound (3.22) applies to the second factor of (3.23) with K in the place of I . Concerning the third factor, Lemma 3.15 implies that there is a constant c_6 such that

$$\mathbb{P}_\infty(\tau_o \leq \tau_A \mid \tau_{\text{Circ}_{K-1}} < \tau_A) \geq e^{-c_6 \sum_{\ell=I}^{K-1} n_\ell}.$$

Substituting the three bounds into (3.23) concludes the induction step in this sub-case:

$$\mathbb{H}_A(o) \geq e^{-c_1 k \log k - c_3 \sum_{\ell=K}^{J-1} n_\ell - c_6 \sum_{\ell=I}^{K-1} n_\ell} \geq e^{-c_1 n \log n}.$$

The second inequality follows from potentially adjusting c_1 to satisfy $c_1 \geq 8 \max\{1, c_3, c_6\}$. This completes the induction and establishes (1.2). \square

4. ESCAPE PROBABILITY ESTIMATES

The purpose of this section is to prove Theorem 5. It suffices to prove the escape probability lower bound (1.6), as (1.7) follows from (1.6) by the pigeonhole principle. Let A be an n -element subset of \mathbb{Z}^2 with at least two elements. We assume w.l.o.g. that $o \in A$. Denote $b = \text{diam}(A)$, and suppose $d \geq 2b$. We aim to show that there is a constant c such that, if $d \geq 2b$, then, for every $x \in A$,

$$\mathbb{P}_x(\tau_{\partial A_d} < \tau_A) \geq \frac{c \mathbb{H}_A(x)}{n \log d}.$$

In fact, by adjusting c , we can reduce to the case when $d \geq kb$ for $k = 200$ and when b is at least a large universal constant, b' . We proceed to prove (1.6) when $d \geq 200b$, for sufficiently large b . Since $C(kb)$ separates A from ∂A_d , we can write the escape probability as the product of two factors:

$$\mathbb{P}_x(\tau_{\partial A_d} < \tau_A) = \mathbb{P}_x(\tau_{C(kb)} < \tau_A) \mathbb{P}_x(\tau_{\partial A_d} < \tau_A \mid \tau_{C(kb)} < \tau_A). \quad (4.1)$$

Concerning the first factor of (4.1), we have the following lemma.

Lemma 4.1. *Let $x \in A$. Then*

$$\mathbb{P}_x(\tau_{C(kb)} < \tau_A) \geq \frac{\mathbb{H}_A(x)}{4 \log(kb)}. \quad (4.2)$$

The factor of $\log(kb)$ arises from evaluating the potential kernel at elements of $C(kb)$; the factor of 4 is unimportant. The proof is an application of the optional stopping theorem to the martingale $\mathbf{a}(S_{j \wedge \tau_o})$.

Proof of Lemma 4.1. Let $x \in A$. By conditioning on the first step, we have

$$\mathbb{P}_x(\tau_{C(kb)} < \tau_A) = \frac{1}{4} \sum_{y \notin A, y \sim x} \mathbb{P}_y(\tau_{C(kb)} < \tau_A), \quad (4.3)$$

where $y \sim x$ means $|x - y| = 1$. We apply the optional stopping theorem to the martingale $\mathbf{a}(S_{j \wedge \tau_o})$ with the stopping time $\tau_A \wedge \tau_{C(kb)}$ to find:

$$\frac{1}{4} \sum_{y \notin A, y \sim x} \mathbb{P}_y(\tau_{C(kb)} < \tau_A) = \frac{1}{4} \sum_{y \notin A, y \sim x} \frac{\mathbf{a}(y) - \mathbb{E}_y \mathbf{a}(S_{\tau_A})}{\mathbb{E}_y [\mathbf{a}(S_{\tau_{C(kb)}}) - \mathbf{a}(S_{\tau_A}) \mid \tau_{C(kb)} < \tau_A]}. \quad (4.4)$$

We need two facts. First, $\mathbb{H}_A(x)$ can be expressed as $\frac{1}{4} \sum_{y \notin A, y \sim x} (\mathbf{a}(y) - \mathbb{E}_y \mathbf{a}(S_{\tau_A}))$ [Pop21, Definition 3.15, Theorem 3.16]. Second, for any $z \in C(kb)$, $\mathbf{a}(z) \leq 4 \log(kb)$ by Lemma A.1. Applying these facts to (4.4), and the result to (4.3), we find

$$\mathbb{P}_x(\tau_{C(kb)} < \tau_A) \geq \frac{1}{4 \log(kb)} \cdot \frac{1}{4} \sum_{y \notin A, y \sim x} (\mathbf{a}(y) - \mathbb{E}_y \mathbf{a}(S_{\tau_A})) = \frac{\mathbb{H}_A(x)}{4 \log(kb)}.$$

□

Concerning the second factor of (4.1), given that $\{\tau_{C(kb)} < \tau_A\}$ occurs, we are essentially in the setting depicted on the right side of Figure 5, with $x = S_{\tau_{C(kb)}}$, $r = b$, kb in the place of $2r$, and $R = d$. The argument highlighted in Section 2.2 suggests that the second factor of (4.1) is at least proportional to $\frac{\log b}{n \log d}$. We will prove this lower bound and combine it with (4.1) and (4.2) to obtain (1.6) of Theorem 5.

Lemma 4.2. *Let $y \in C(kb)$. If $d \geq kb$ and if b is sufficiently large, then*

$$\mathbb{P}_y(\tau_{\partial A_d} < \tau_A) \geq \frac{\log b}{2n \log d}. \quad (4.5)$$

Proof. Let $y \in C(kb)$. We will follow the argument of Section 2.2. Label the points of A as x_1, x_2, \dots, x_n and define

$$Y_i = \mathbf{1}(\tau_{x_i} < \tau_{\partial A_d}) \quad \text{and} \quad W = \sum_{i=1}^n Y_i.$$

From the definition of W , we see that $\{W = 0\} = \{\tau_{\partial A_d} < \tau_A\}$. Thus to obtain the lower bound in (4.5), it suffices to get a complementary upper bound on

$$\mathbb{P}_y(W > 0) = \frac{\mathbb{E}_y W}{\mathbb{E}_y[W \mid W > 0]}. \quad (4.6)$$

We will find α and β such that, uniformly for $y \in C(kb)$ and $x_i, x_j \in A$,

$$\mathbb{P}_y(\tau_{x_i} < \tau_{\partial A_d}) \leq \alpha \quad \text{and} \quad \mathbb{P}_{x_i}(\tau_{x_j} < \tau_{\partial A_d}) \geq \beta. \quad (4.7)$$

Moreover, α and β will satisfy

$$\alpha \leq \beta \quad \text{and} \quad 1 - \beta \geq \frac{\log b}{2 \log d}. \quad (4.8)$$

The requirement that $\alpha \leq \beta$ prevents us from choosing $\beta = 0$. Essentially, we will be able to satisfy (4.7) and the first condition of (4.8) because $|x_i - x_j|$ is smaller than $|y - x_i|$. We will be able to satisfy the second condition because $\text{dist}(x_i, \partial A_d) \geq d$ while $|x_i - x_j| \leq b$, which implies that $\mathbb{P}_{x_i}(\tau_{x_j} < \tau_{\partial A_d})$ is roughly $1 - \frac{\log b}{\log d}$.

If α, β satisfy (4.7), then we can bound (4.6) as

$$\mathbb{P}_y(W > 0) \leq \frac{n\alpha}{1 + (n-1)\beta}. \quad (4.9)$$

Additionally, when α and β satisfy (4.8), (4.9) implies

$$\mathbb{P}_y(W = 0) \geq \frac{(1-\beta) + n(\beta-\alpha)}{(1-\beta) + n\beta} \geq \frac{1-\beta}{n} \geq \frac{\log b}{2n \log d},$$

which gives the claimed bound (4.5).

Identifying α . We now find the α promised in (4.7). Denote $F_i = C_{x_i}(d+b)$ (Figure 10). Since ∂A_d separates y from F_i , we have

$$\mathbb{P}_y(\tau_{x_i} < \tau_{\partial A_d}) \leq \mathbb{P}_y(\tau_{x_i} < \tau_{F_i}) = \mathbb{P}_{y-x_i}(\tau_o < \tau_{C(d+b)}). \quad (4.10)$$




FIGURE 10. Escape to ∂A_d , for $n = 3$. Each F_i is a circle centered on $x_i \in A$, separating A_d from infinity. Lemma 3.7 bounds above the probability that the walk hits x_i before F_i , uniformly for $y \in C(kb)$.

The hypotheses of Lemma 3.7 are met because $y - x_i \neq o$ and $y - x_i \in D(d + b)$. Hence (3.10) applies as

$$\mathbb{P}_{y-x_i}(\tau_o < \tau_{C(d+b)}) = \frac{\mathbf{a}'(d+b) - \mathbf{a}(y-x_i) + O(|y-x_i|^{-1})}{\mathbf{a}'(d+b) + O(|y-x_i|^{-1})}. \quad (4.11)$$

Ignoring the error terms, the expression in (4.11) is at most $\frac{\log(d+b) - \log(kb)}{\log(d+b)}$. A more careful calculation gives

$$\mathbb{P}_{y-x_i}(\tau_o < \tau_{C(d+b)}) = \frac{\log(d+b) - \log(kb)}{\log(d+b)} + \delta_1 \leq \frac{(1+\varepsilon)\log d - \log(kb)}{\log d} + \delta_1 =: \alpha,$$

where $\delta_1 = (\frac{\pi\kappa}{2} + O(b^{-1}))(\log d)^{-1}$ and $\varepsilon = \frac{b}{d \log d}$. (Recall that $\kappa \in (1.02, 1.03)$ is a constant associated with the potential kernel (3.9).) The inequality results from applying the inequality $\log(1+x) \leq x$, which holds for $x > -1$, to the $\log(d+b)$ term in the numerator, and reducing $\log(d+b)$ to $\log d$ in the denominator. By (4.10), α satisfies (4.7).

Identifying β . We now find a suitable β . Since $C_{x_i}(d)$ separates A from ∂A_d , we have

$$\mathbb{P}_{x_i}(\tau_{x_j} < \tau_{\partial A_d}) \geq \mathbb{P}_{x_i}(\tau_{x_j} < \tau_{C_{x_i}(d)}) = \mathbb{P}_{x_i-x_j}(\tau_o < \tau_{C(d)}). \quad (4.12)$$

The hypotheses of Lemma 3.7 are met because $x_i - x_j \neq o$ and $x_i - x_j \in D(d)$. Hence (3.10) applies as

$$\mathbb{P}_{x_i-x_j}(\tau_o < \tau_{C(d)}) = \frac{\mathbf{a}'(d) - \mathbf{a}(x_i-x_j) + O(|x_i-x_j|^{-1})}{\mathbf{a}'(d) + O(|x_i-x_j|^{-1})}. \quad (4.13)$$

Ignoring the error terms, (4.13) is at least $\frac{\log d - \log b}{\log d + \kappa}$. A more careful calculation gives

$$\mathbb{P}_{x_i-x_j}(\tau_o < \tau_{C(d)}) = \frac{\log d - \log b}{\log d} - \delta_2 =: \beta,$$

where $\delta_2 = (\frac{\pi\kappa}{2} + O(b^{-1}))(\log d)^{-1}$. By (4.12), β satisfies (4.7).

Verifying (4.8). To verify the first condition of (4.8), we calculate

$$(\beta - \alpha) \log d = \log k - \frac{b}{d} - \pi\kappa + O(b^{-1}) \geq 1 + O(b^{-1}).$$

The inequality is due to $k = 200$, $\frac{b}{d} \leq 0.5$, and $\pi\kappa < 3.5$. If b is sufficiently large, then $1 + O(b^{-1})$ is nonnegative, which verifies (4.8).

Concerning the second condition of (4.8), if b is sufficiently large, then

$$1 - \beta = \frac{\log b + 1}{\log d} \geq \frac{\log b}{2 \log d}.$$

We have identified α, β which satisfy (4.7) and (4.8) for sufficiently large b . By the preceding discussion, this proves (4.5). \square

Proof of Theorem 5. By (4.1), Lemma 4.1, and Lemma 4.2, we have

$$\mathbb{P}_x(\tau_{\partial A_d} < \tau_A) \geq \frac{\mathbb{H}_A(x)}{4 \log(kb)} \cdot \frac{\log b}{2n \log d} \geq \frac{\mathbb{H}_A(x)}{16n \log d}, \quad (4.14)$$

whenever $x \in A$ and $d \geq kb$, for sufficiently large b . The second inequality is due to the fact that $\log(kb) \leq 2 \log b$ for sufficiently large b .

By the reductions discussed at the beginning of this section, (4.14) implies that there is a constant c such that (1.6) holds for $x \in A$ if A has at least two elements and if $d \geq 2 \text{diam}(A)$. (1.7) follows from (1.6) because, by the pigeonhole principle, some element of A has harmonic measure of at least n^{-1} . \square

5. CLUSTERING SETS OF RELATIVELY LARGE DIAMETER

When a HAT configuration has a large diameter relative to the number of particles, we can decompose the configuration into clusters of particles, which are well separated in a sense. This is the content of Lemma 5.2, which will be a key input to the results in Section 6.

Definition 5.1 (Exponential clustering). *For a finite $A \subset \mathbb{Z}^2$ with $|A| = n$, an exponential clustering of A with parameter $r \geq 0$, denoted $A \mapsto_r \{A^i, x_i, \theta^{(i)}\}_{i=1}^k$, is a partition of A into clusters A^1, A^2, \dots, A^k with $1 \leq k \leq n$, such that each cluster arises as $A^i = A \cap D_{x_i}(\theta^{(i)})$ for $x_i \in \mathbb{Z}^2$, with $\theta^{(i)} \geq r$, and*

$$\text{dist}(A^i, A^j) > \exp(\max\{\theta^{(i)}, \theta^{(j)}\}) \text{ for } i \neq j. \quad (5.1)$$

We will call x_i the center of cluster i . In some instances, the values of r , x_i , or $\theta^{(i)}$ will be irrelevant and we will omit them from our notation. For example, $A \mapsto \{A^i\}_{i=1}^k$.

An exponential clustering of A with parameter r always exists because, if $A^1 = A$, $x_1 \in A$, and $\theta^{(1)} \geq \max\{r, \text{diam}(A)\}$, then $A \mapsto_r \{A^1, x_1, \theta^{(1)}\}$ is such a clustering. However, to ensure that there is an exponential clustering of A (with parameter r) with more than one cluster, we require that the diameter of A exceeds $2\theta_{n-1}(r)$. Recall that we defined $\theta_m(r)$ in (1.1) through $\theta_0(r) = r$ and $\theta_m(r) = \theta_{m-1}(r) + e^{\theta_{m-1}(r)}$ for $m \geq 1$.

Lemma 5.2. *Let $|A| = n$. If $\text{diam}(A) > 2\theta_{n-1}(r)$, then there exists an exponential clustering of A with parameter r into $k > 1$ clusters.*

To prove the lemma, we will identify disks with radii of at most $\theta_{n-1}(r)$, which cover A . Although it is not required of an exponential clustering, the disks will be centered at elements of A . These disks will give rise to at least two clusters, since $\text{diam}(A)$ exceeds $2\theta_{n-1}(r)$. The disks will be surrounded by large annuli which are empty of A , which will imply that the clusters are exponentially separated.

Proof of Lemma 5.2. For each $x \in A$ and $m \geq 1$, consider the annulus $\mathcal{A}_x(\theta_m) = D_x(\theta_m) \setminus D_x(\theta_{m-1})$. For each x , identify the smallest m such that $\mathcal{A}_x(\theta_m) \cap A$ is empty and call it m_x . Note that since $|A| = n$, m_x can

be no more than n and hence $\theta_{m_x} \leq \theta_n$. Call the corresponding annulus \mathcal{A}_x^* , and denote $D_x^* = D_x(\theta_{m_x-1})$. For convenience, we label the elements of A as x_1, x_2, \dots, x_n .

For $x_i \in A$, we collect those disks $D_{x_j}^*$ which contain it as

$$\mathcal{E}(x_i) = \{D_{x_j}^* : x_i \in D_{x_j}^*, 1 \leq j \leq n\}.$$

We observe that $\mathcal{E}(x_i)$ is always nonempty, as it contains $D_{x_i}^*$. Now observe that, for any two distinct $D_{x_j}^*, D_{x_\ell}^* \in \mathcal{E}(x_i)$, it must be that

$$D_{x_j}^* \cap A \subseteq D_{x_\ell}^* \cap A \quad \text{or} \quad D_{x_\ell}^* \cap A \subseteq D_{x_j}^* \cap A. \quad (5.2)$$

To see why, assume for the purpose of deriving a contradiction that each disk contains an element of A which the other does not. Without loss of generality, suppose $\theta_{m_{x_j}} \geq \theta_{m_{x_\ell}}$ and let $y_\ell \in (D_{x_\ell}^* \setminus D_{x_j}^*) \cap A$. Because each disk must contain x_i , we have $|y_\ell - x_i| \leq 2\theta_{m_{x_\ell}-1}$ and $|x_i - x_j| \leq \theta_{m_{x_j}-1}$. The triangle inequality implies

$$|y_\ell - x_j| \leq \theta_{m_{x_j}-1} + 2\theta_{m_{x_\ell}-1} \leq \theta_{m_{x_j}} \implies y_\ell \in D_{x_j}(\theta_{m_{x_j}}) \cap A.$$

By assumption, y_ℓ is not in $D_{x_j}(\theta_{m_{x_j}-1}) \cap A$, so y_ℓ must be an element of $\mathcal{A}_{x_j}(\theta_{m_{x_j}}) \cap A$, which contradicts the construction of m_{x_j} .

By (5.2), we may totally order the elements of $\mathcal{E}(x_i)$ by inclusion of intersection with A . For each x_i , we select the element of $\mathcal{E}(x_i)$ which is greatest in this ordering. If we have not already established it as a cluster, we do so. After we have identified a cluster for each x_i , we discard those $D_{x_j}^*$ which were not selected for any x_i . For the remainder of the proof, we only refer to those $D_{x_j}^*$ which were established as clusters, and we relabel the x_i so that the clusters can be expressed as the collection $\{D_{x_j}^*\}_{j=1}^k$, for some $1 \leq k \leq n$. We will show that k is strictly greater than one.

The collection of clusters contains all elements of A , and is associated to the collection of annuli $\{\mathcal{A}_{x_j}^*\}_{j=1}^k$, which contain no elements of A . We observe that, for some distinct x_j and x_ℓ , it may be that $\mathcal{A}_{x_j}^* \cap D_{x_\ell}^* \neq \emptyset$. However, because the annuli contain no elements of A , it must be that

$$\begin{aligned} \text{dist}(D_{x_j}^* \cap A, D_{x_\ell}^* \cap A) &> \max \left\{ \theta_{m_{x_j}} - \theta_{m_{x_j-1}}, \theta_{m_{x_\ell}} - \theta_{m_{x_\ell-1}} \right\} \\ &= e^{\max \left\{ \theta_{m_{x_j-1}}, \theta_{m_{x_\ell-1}} \right\}} \end{aligned}$$

As $D_{x_j}^* \cap A \subseteq D_{x_j}^*$ for any x_j in question, we conclude the desired separation of clusters by setting $A^i = D_{x_i}^* \cap A$ for each $1 \leq i \leq k$. Furthermore, since $m_{x_j} \leq n$ for all j , $\theta_{m_{x_j-1}} \leq \theta_{n-1}$ for all j . Since A is contained in the union of the clusters, if $\text{diam}(A) > 2\theta_{n-1}$, then there must be at least two clusters. Lastly, as $m_{x_j} \geq 1$ for all j , $\theta_{m_{x_j-1}} \geq r$ for all j . \square

6. ESTIMATES OF THE TIME OF COLLAPSE

We proceed to prove the main collapse result, Theorem 1. As the proof requires several steps, we begin by discussing the organization of the section and introducing some key definitions. We avoid discussing the proof strategy in detail before making necessary definitions; an in-depth proof strategy is covered in Section 6.2.

Briefly, to estimate the time until the diameter of the configuration falls below a given function of n , we will perform exponential clustering and consider the more manageable task of (i) estimating the time until some cluster loses all of its particles to the other clusters. By iterating this estimate, we can (ii) control the time it takes for the clusters to consolidate into a single cluster. We will find that the surviving cluster has a diameter which is approximately the logarithm of the original diameter. Then, by repeatedly applying this estimate, we can (iii) control the time it takes for the diameter of the configuration to collapse.

The purpose of Section 6.1 is to wield (ii) in the form of Proposition 6.3 and prove Theorem 1, thus completing (iii). The remaining subsections are dedicated to proving the proposition. An overview of our strategy will be detailed in Section 6.2. In particular, we describe how the key harmonic measure estimate of Theorem 4 and the key escape probability estimate of Theorem 5 contribute to addressing (i). We then develop basic properties of cluster separation and explore the geometric consequences of timely cluster collapse in Section 6.3. Lastly, in Section 6.4, we prove a series of propositions which collectively control the timing of individual cluster collapse, culminating in the proof of Proposition 6.3.

Implicit in this discussion is a notion of “cluster” which persists over multiple steps of the dynamics. Starting from an exponential clustering of U_0 , we define the clusters at later times by simply updating the original clusters to account for the movement of particles. Recall that an exponential clustering $U_0 \mapsto \{U_0^i, x_i, \theta^{(i)}\}_{i=1}^k$ of U_0 is a partition of U_0 that arises from intersecting U_0 with a collection of discs and which satisfies an exponential separation property (5.1).

Definition 6.1. *Let U_0 have an exponential clustering $U_0 \mapsto \{U_0^i, x_i, \theta^{(i)}\}_{i=1}^k$. We define a sequence $(\{U_t^i\}_{i=1}^k)_{t \geq 0}$, starting with $\{U_0^i\}_{i=1}^k$ and such that $\{U_t^i\}_{i=1}^k$ is a partition of U_t for each t , in the following way. Given U_t, X_t, Y_t , and the random walk S^t that accomplishes the transport step at time t , define*

$$U_{t+1}^i = \begin{cases} (U_t^i \setminus \{X_t\}) \cup \{Y_t\} & S_{\tau_{U_t \setminus \{X_t\}}}^t \in U_t^i \\ U_t^i \setminus \{X_t\} & S_{\tau_{U_t \setminus \{X_t\}}}^t \notin U_t^i. \end{cases} \quad (6.1)$$

We refer to the parts of $\{U_t^i\}_{i=1}^k$ as clusters.

Note that, since $\{U_t^i\}_{i=1}^k$ is a partition of U_t , the removal of X_t from each U_t^i in (6.1) only affects the cluster to which X_t belongs. For example, if X_t belongs to U_t^i and the random walk S^t first hits $U_t \setminus \{X_t\}$ at a different cluster U_t^j , then the i^{th} cluster loses X_t and the j^{th} cluster gains Y_t , while all other clusters remain the same.

Definition 6.2 (Cluster collapse times). *We define the ℓ -cluster collapse time inductively, according to $\mathcal{T}_0 \equiv 0$ and*

$$\mathcal{T}_\ell = \inf \left\{ t \geq \mathcal{T}_{\ell-1} : \exists 1 \leq j \leq k : U_{\mathcal{T}_{\ell-1}}^j \neq \emptyset \text{ and } U_t^j = \emptyset \right\}, \quad 1 \leq \ell \leq k.$$

6.1. Proving Theorem 1. We now state the proposition to which most of the effort in this section is devoted and, assuming it, prove Theorem 1. We will denote by

- n , the number of elements of U_0 ;
- $\Phi(r)$, the inverse function of $\theta_n(r)$ for all $r \geq 0$ ($\theta_n(r)$ is an increasing function of $r \geq 0$ for every n); and
- \mathcal{F}_t , the sigma algebra generated by the initial configuration U_0 , the first t activation sites X_0, X_1, \dots, X_{t-1} , and the first t random walks S^0, S^1, \dots, S^{t-1} , which accomplish the transport component of the dynamics.

We note that Φ is defined so that, if r equals $\Phi(\text{diam}(U_0))$, then the diameter of U_0 exceeds $2\theta_{n-1}(r)$. Lemma 5.2 states that, in this case, an exponential clustering of U_0 with parameter r has at least two clusters.

Proposition 6.3. *There is a constant c such that, if the diameter d of U_0 exceeds $\theta_{4n}(cn)$, then for any number of clusters k resulting from an exponential clustering of U_0 with parameter $r = \Phi(d)$, we have*

$$\mathbf{P}_{U_0} \left(\mathcal{T}_{k-1} \leq (\log d)^{1+7\delta} \right) \geq 1 - \exp \left(-2nr^\delta \right), \quad (6.2)$$

where $\delta = (2n)^{-4}$.

In words, if U_0 has a diameter of d , it takes no more than $(\log d)^{1+o_n(1)}$ steps to observe the collapse of all but one cluster, with high probability. Because no cluster begins with a diameter greater than $\log d$ (by exponential clustering) and, as the diameter of a cluster increases at most linearly in time, the remaining cluster at time \mathcal{T}_{k-1} has a diameter of no more than $(\log d)^{1+o_n(1)}$. We will obtain Theorem 1 by repeatedly applying Proposition 6.3. We prove the theorem here, assuming the proposition, and then prove the proposition in the following subsections.

Proof of Theorem 1. We organize time into “rounds” of collapse, which proceed in the following way. At the beginning of round $\ell \geq 1$, regardless of the diameter d_ℓ of the current configuration, we obtain an exponential clustering (with parameter $r_\ell = \Phi(d_\ell)$) of this configuration, and we define clusters and their collapse times in analogy with Definitions 6.1 and 6.2. The round lasts until all but one of the clusters have collapsed or until $(\log d_\ell)^{1+7\delta}$ steps have passed, at which time the next round begins.

The idea of the proof is to use Proposition 6.3 to bound below the probability that the diameter drops from d_ℓ to at most $(\log d_\ell)^2$, over each round ℓ until the first round η that begins with a diameter of at most $\theta_{4n}(cn)$ (c is the constant from Proposition 6.3). When this event occurs, the rounds shorten rapidly enough that the cumulative time it takes to reach round η is at most twice the maximum length $(\log d)^{1+7\delta}$ of the first round.

We define the rounds inductively. Set $\mathcal{R}_1 \equiv 0$. Round $\ell \geq 1$ begins at time \mathcal{R}_ℓ , in configuration $V_{\ell,0} = U_{\mathcal{R}_\ell}$. We further define $V_{\ell,t} = U_{\mathcal{R}_\ell+t}$ for $t \geq 0$. To define when the round ends, let $d_\ell = \text{diam}(V_{\ell,0})$ and $r_\ell = \Phi(d_\ell)$, and denote the exponential clustering of $V_{\ell,0}$ with parameter r_ℓ by $\{V_{\ell,0}^i\}_{i=1}^{k_\ell}$. We define the clusters $\{V_{\ell,t}^i\}_{i=1}^{k_\ell}$ in analogy with Definition 6.1, and the corresponding cluster collapse times in analogy with Definition 6.2. Specifically, we define the latter according to $\mathcal{T}_{\ell,0} = \mathcal{R}_\ell$ and

$$\mathcal{T}_{\ell,m} = \inf \left\{ t \geq \mathcal{T}_{\ell,m-1} : \exists 1 \leq j \leq k_\ell : V_{\ell,\mathcal{T}_{\ell,m-1}}^j \neq \emptyset \text{ and } V_{\ell,t}^j = \emptyset \right\}, \quad 1 \leq m \leq k_\ell.$$

Round ℓ ends, and round $\ell + 1$ begins, at time

$$\mathcal{R}_{\ell+1} = \mathcal{R}_\ell + \min \left\{ \mathcal{T}_{\ell,k_\ell-1} - \mathcal{R}_\ell, (\log d_\ell)^{1+7\delta} \right\}.$$

The key event is the event that the diameter drops from d_ℓ to at most $(\log d_\ell)^2$ over each round, until round

$$\eta = \inf \{ \ell \geq 1 : d_\ell \leq \theta_{4n}(cn) \}.$$

The round of the j^{th} diameter drop is

$$\xi_j = \inf \{ \ell > \xi_{j-1} : d_\ell \leq (\log d_{\xi_{j-1}})^2 \}, \quad j \geq 1,$$

where $\xi_0 \equiv 1$. In these terms, the key event is $\mathcal{E}_{\eta-1}$, where

$$\mathcal{E}_m = \cap_{j=1}^m E_j \quad \text{and} \quad E_j = \{ \xi_j - \xi_{j-1} = 1 \}.$$

If $\mathcal{E}_{\eta-1}$ occurs, then $\xi_{j-1} = j$ for every $j \leq \eta$, and some algebra shows that

$$\frac{\log d_j}{\log d_1} = \frac{\log d_{\xi_{j-1}}}{\log d_1} \leq 2^{-(j-1)}, \quad 1 \leq j \leq \eta. \quad (6.3)$$

We use (6.3) to bound the time \mathcal{R}_η that it takes to reach round η when $\mathcal{E}_{\eta-1}$ occurs:

$$\mathcal{R}_\eta \leq \sum_{j=1}^{\eta-1} (\log d_j)^{1+7\delta} \leq (\log d_1)^{1+7\delta} \sum_{j=1}^{\eta-1} 2^{-(j-1)} \leq (\log d_1)^{1+8\delta}.$$

These inequalities are respectively justified by the fact that the ℓ^{th} round lasts at most $(\log d_\ell)^{1+7\delta}$ steps, by (6.3), and by the fact that $(\log d_1)^\delta$ is at least 2.

Since \mathcal{R}_η is at least $\mathcal{T}(\theta_{4n}(cn))$, to complete the proof, it suffices to show that $\mathcal{E}_{\eta-1}^c$ occurs with a probability of at most e^{-n} . We express the probability of $\mathcal{E}_{\eta-1}^c$ in terms of $G_j = \mathcal{E}_{j-1} \cap \{\eta > \xi_{j-1}\}$, as

$$\mathbf{P}_{U_0}(\mathcal{E}_{\eta-1}^c) = \mathbf{E}_{U_0} \sum_{j=1}^{\eta-1} \mathbf{1}(E_j^c \cap \mathcal{E}_{j-1}) = \mathbf{E}_{U_0} \sum_{j=1}^{\infty} \mathbf{1}(E_j^c \cap G_j) = \sum_{j=1}^{\infty} \mathbf{P}_{U_0}(E_j^c \cap G_j). \quad (6.4)$$

The three equalities hold by the disjointness of $\{E_j^c \cap \mathcal{E}_{j-1}\}_{j=1}^{\infty}$; the fact that $\xi_{j-1} = j$ when \mathcal{E}_{j-1} occurs, hence $\{\eta > j\} = \{\eta > \xi_{j-1}\}$; and Tonelli's theorem, which applies because the summands are nonnegative.

For brevity, denote $\mathcal{R}_{\xi_{j-1}}$ by s_j . The summands on the right hand side of (6.4) satisfy

$$\mathbf{P}_{U_0}(E_j^c \cap G_j) = \mathbf{E}_{U_0} \left[\mathbf{P}_{U_0}(E_j^c \mid \mathcal{F}_{s_j}) \mathbf{1}_{G_j} \right] = \mathbf{E}_{U_0} \left[\mathbf{P}_{V_{j,0}}(E_1^c) \mathbf{1}_{G_j} \right] \leq \mathbf{E}_{U_0} \left[e^{-2nr_j^\delta} \mathbf{1}_{G_j} \right]. \quad (6.5)$$

The first equality holds by the tower rule and the fact that G_j is measurable with respect to \mathcal{F}_{s_j} ; the second equality follows from the strong Markov property applied to time s_j , which—when G_j occurs—is the time that the j^{th} round begins; the inequality is due to Proposition 6.3. Note that Proposition 6.3 applies because round j precedes round η when G_j occurs, hence the diameter of $V_{j,0}$ exceeds $\theta_{4n}(cn)$. We substitute (6.5) into (6.4) and use Tonelli's theorem a second time to conclude that

$$\mathbf{P}_{U_0}(\mathcal{E}_{\eta-1}^c) \leq \mathbf{E}_{U_0} \left[\sum_{j=1}^{\eta-1} e^{-2nr_j^\delta} \mathbf{1}_{G_j} \right].$$

To finish the proof, it therefore suffices to establish the first inequality of

$$\sum_{j=1}^{\eta-1} e^{-2nr_j^\delta} \mathbf{1}_{G_j} \leq \sum_{j=1}^{\infty} e^{-2nj} \leq e^{-n}. \quad (6.6)$$

Denote the number of drops before round η by $N = \sup\{j \geq 0 : \xi_j < \eta\}$. When G_j occurs, r_j^δ satisfies

$$r_j^\delta = r_{\xi_{j-1}}^\delta \geq r_{\xi_N}^\delta + N + 1 - j \geq N + 2 - j, \quad N \geq j - 1.$$

The first inequality follows from simple but cumbersome algebra, using the definitions of $r_\ell = \theta_n^{-1}(d_\ell)$ and ξ_{j-1} , and the fact that d_ℓ exceeds $\theta_{4n}(cn)$ when ℓ is at most ξ_N . In particular, this fact implies that $r_{\xi_N}^\delta$ is at least 1, which justifies the second inequality. The first sum in (6.6) satisfies

$$\sum_{j=1}^{\eta-1} e^{-2nr_j^\delta} \mathbf{1}_{G_j} \leq \sum_{j=1}^{\eta-1} e^{-2n(N+2-j)} \mathbf{1}_{G_j} \leq \sum_{j=1}^{N+1} e^{-2n(N+2-j)} \leq \sum_{j=1}^{\infty} e^{-2nj}.$$

The first bound is due to the preceding lower bound on r_j^δ ; the second holds because, by definition, N is at least $\eta - 2$; the third follows from reversing the order of the sum ($N + 2 - j \rightarrow j$) and then replacing N with ∞ in the resulting sum. \square

For applications in Section 7, we extend Theorem 1 to a more general tail bound of $\mathcal{T}(\theta_{4n})$.

Corollary 6.4 (Corollary of Theorem 1). *Let U be an n -element subset of \mathbb{Z}^2 with a diameter of d . There exists a universal positive constant c such that*

$$\mathbf{P}_U \left(\mathcal{T}(\theta_{4n}(cn)) > t(\log \max\{t, d\})^{1+o_n(1)} \right) \leq e^{-t} \quad (6.7)$$

for all $t \geq 1$. For the sake of concreteness, this is true with $2n^{-4}$ in the place of $o_n(1)$.

In the proof of the corollary, it will be convenient to have notation for the timescale of collapse after j failed collapses, starting from a diameter of d . Because diameter increases at most linearly in time, if the initial configuration has a diameter of d and collapse does not occur in the next $(\log d)^{1+o_n(1)}$ steps, then the diameter after this period of time is at most $d + (\log d)^{1+o_n(1)}$. In our next attempt to observe collapse, we would wait at most $(\log(d + (\log d)^{1+o_n(1)}))^{1+o_n(1)}$ steps. This discussion motivates the definition of the functions $g_j = g_j(d, \varepsilon)$ by

$$g_0 = (\log d)^{1+\varepsilon} \quad \text{and} \quad g_j = \left(\log \left(d + \sum_{i=0}^{j-1} g_i \right) \right)^{1+\varepsilon}, \quad j \geq 1.$$

We will use $t_j = t_j(d, \varepsilon)$ to denote the cumulative time $\sum_{i=0}^j g_i$.

Proof of Corollary 6.4. Let $\varepsilon = n^{-4}$ and use this as the ε parameter for the collapse timescales g_j and cumulative times t_j . Additionally, denote $\theta = \theta_{4n}(cn)$ for the constant c from Theorem 1 (this will also be the constant in the statement of the corollary). The bound (6.7) clearly holds when d is at most θ , so we assume $d \geq \theta$.

Because the diameter of U is d and as diameter grows at most linearly in time, conditionally on $F_j = \{\mathcal{T}(\theta) > t_j\}$, the diameter of U_{t_j} is at most $d + t_j$. Consequently, by the Markov property applied to time t_j , and by Theorem 1 (the diameter is at least θ) and the fact that $n \geq 1$, the conditional probability $\mathbf{P}_U(F_{j+1}|F_j)$ satisfies

$$\mathbf{P}_U(F_{j+1}|F_j) = \mathbf{E}_U \left[\mathbf{P}_{U_{t_j}}(\mathcal{T}(\theta) > g_{j+1}) \frac{\mathbf{1}_{F_j}}{\mathbf{P}_U(F_j)} \right] \leq e^{-1} \quad \text{for any } j \geq 0. \quad (6.8)$$

In fact, Theorem 1 implies that the inequality holds with e^{-n} in the place of e^{-1} , but this will make no difference to us.

Define J to be the greatest j such that $t_j \leq t$. There are at least J consecutive collapse attempts which must fail in order for $\mathcal{T}(\theta)$ to exceed t . By (6.8),

$$\mathbf{P}_U(\mathcal{T}(\theta) > t) \leq \prod_{i=0}^{J-1} \mathbf{P}_U(F_{i+1}|F_i) \leq e^{-J}. \quad (6.9)$$

To bound below J , note that g_j increases with j , hence $t_j \leq (j+1)g_j$ and $g_j \leq g_{j+1}$, which implies that

$$j \geq \frac{t_j - g_j}{g_j} \geq \frac{t_j - g_{j+1}}{g_{j+1}}.$$

By the definition of J , we have $t_J + g_{J+1} \geq t$ and $g_{J+1} \leq (\log(d+t))^{1+\varepsilon}$. By combining the preceding display with these two facts, we find that

$$J \geq \frac{t_J - g_{J+1}}{g_{J+1}} \geq \frac{t - 2g_{J+1}}{g_{J+1}} \geq \frac{t - 2(\log(d+t))^{1+\varepsilon}}{(\log(d+t))^{1+\varepsilon}}.$$

Returning to (6.9), we conclude that

$$\mathbf{P}_U(\mathcal{T}(\theta) > t) \leq e^{-\frac{t - 2(\log(d+t))^{1+\varepsilon}}{(\log(d+t))^{1+\varepsilon}}}.$$

We obtain (6.7) by replacing t with $t(\log \max\{t, d\})^{1+2\varepsilon}$ in the preceding inequality and noting that, because $d \geq \theta$ and $\varepsilon < 1$, the resulting bound is at most e^{-t} . \square

6.2. Proof strategy for Proposition 6.3. We turn our attention to the proof of Proposition 6.3, which finds a high-probability bound on the time it takes for all but one cluster to collapse. Heuristically, if there are only two clusters, separated by a distance ρ_1 , then one of the clusters will lose all its particles to the other cluster in $\log \rho_1$ steps (up to factors depending on n), due to the harmonic measure and escape probability lower bounds of Theorems 4 and 5. This heuristic suggests that, among k clusters, we should observe the collapse of *some* cluster on a timescale which depends on the smallest separation between any two of the k clusters. Similarly, at the time the ℓ^{th} cluster collapses, if the least separation among the remaining clusters is $\rho_{\ell+1}$, then we expect to wait $\log \rho_{\ell+1}$ steps for the $(\ell + 1)^{\text{st}}$ collapse.

If the timescale of collapse is small relative to the separation between clusters, then the pairwise separation and diameters of clusters cannot appreciably change while collapse occurs. In particular, the separation between any two clusters cannot significantly exceed the initial diameter d of the configuration, which suggests an overall bound of order $(\log d)^{1+o_n(1)}$ steps for all but one cluster to collapse, where the $o_n(1)$ factor accounts for various n -dependent factors. This is the upper bound that we establish.

We now highlight some key aspects of the proof.

6.2.1. Expiry time. As described above, over the timescale typical of collapse, the diameters and separation of clusters will not change appreciably. Because these quantities determine the probability with which the least separated cluster loses a particle, we will be able to obtain estimates of this probability which hold uniformly from the time $\mathcal{T}_{\ell-1}$ of the $(\ell - 1)^{\text{st}}$ cluster collapse and until the next time \mathcal{T}_ℓ that some cluster collapses, unless $\mathcal{T}_\ell - \mathcal{T}_{\ell-1}$ is atypically large. Indeed, if $\mathcal{T}_\ell - \mathcal{T}_{\ell-1}$ is as large as the separation ρ_ℓ of the least separated cluster at time $\mathcal{T}_{\ell-1}$, then two clusters may intersect or a cluster may no longer have an exposed element (e.g., due to a group of clusters surrounding it). We avoid this by defining a $\mathcal{F}_{\mathcal{T}_{\ell-1}}$ -measurable *expiry time* \mathfrak{t}_ℓ (which will effectively be $(\log \rho_\ell)^2$) and restricting our estimates to the interval from $\mathcal{T}_{\ell-1}$ to the minimum of $\mathcal{T}_{\ell-1} + \mathfrak{t}_\ell$ and \mathcal{T}_ℓ . An expiry time of $(\log \rho_\ell)^2$ is short enough that the relative separation of clusters will not change significantly before it, but long enough so that some cluster will collapse before it with overwhelming probability.

6.2.2. Midway point. From time $\mathcal{T}_{\ell-1}$ to time \mathcal{T}_ℓ or until expiry, we will track activated particles which reach a circle of radius $\frac{1}{2}\rho_\ell$ surrounding one of the least separated clusters, which we call the *watched* cluster. We will use this circle, called the *midway point*, to organize our argument with the following three estimates, which will hold uniformly over this interval of time (Figure 11).

- (1) Activated particles which reach the midway point deposit at the watched cluster with a probability of at most 0.51.
- (2) With a probability of at least $(\log \rho_\ell)^{-1-o_n(1)}$, the activated particle reaches the midway point.
- (3) Conditionally on the activated particle reaching the midway point, the probability that it originated at the watched cluster is at least $(\log \log \rho_\ell)^{-1}$.

To explain the third estimate, we make two observations. First, consider a cluster j separated from the watched cluster by a distance of ρ . In the relevant context, cluster j will essentially be exponentially separated, so its diameter will be at most $\log \rho$. Consequently, a particle activated at cluster j reaches the midway point with a probability of at most $\frac{\log \log \rho}{\log \rho}$. (This follows from the fact that random walk starting at a distance of at most $\log \rho$ from the origin escapes to a distance of ρ with a probability of at most $\frac{\log \log \rho}{\log \rho}$, roughly.) Because this probability is decreasing in ρ and because $\rho \geq \rho_\ell$, $\frac{\log \log \rho}{\log \rho}$ further bounds above it. Second, the probability that a particle activated at the watched cluster reaches the midway point is at least $(\log \rho_\ell)^{-1}$, up to a factor depending on n . Combining these two observations with Bayes's rule, a particle which reaches the midway point was activated at the watched cluster with a probability of at least $(\log \log \rho_\ell)^{-1}$, up to an n -dependent factor.

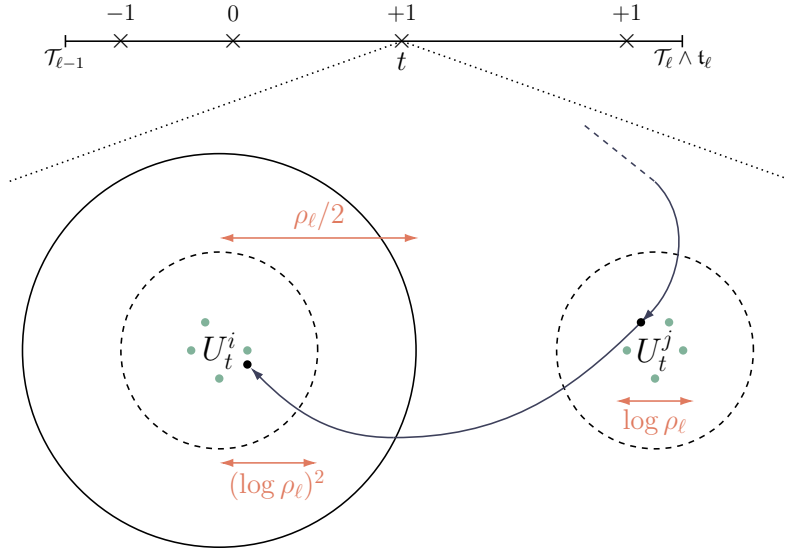


FIGURE 11. Setting of the proof of Proposition 6.3. Least separated clusters i and j (cluster i is the watched cluster), each with a diameter of approximately $\log \rho_\ell$, are separated by a distance ρ_ℓ at time $\mathcal{T}_{\ell-1}$. The diameters of the clusters grow at most linearly in time, so over approximately $(\log \rho_\ell)^2$ steps, the clusters remain within the dotted circles. Crosses on the timeline indicate times before collapse and expiry at which an activated particle reaches the midway point (solid circle). At these times, the number of particles in the watched cluster may remain the same or increase or decrease by one (indicated by $0, \pm 1$ above the crosses). At time t , the watched cluster gains a particle from cluster j .

6.2.3. *Coupling with random walk.* Each time an activated particle reaches the midway point while the watched cluster persists, there is a chance of at least $(\log \log \rho_\ell)^{-1}$ up to an n -dependent factor that the particle originated at the watched cluster and will ultimately deposit at another cluster. When this occurs, the watched cluster loses a particle. Alternatively, the activated particle may return to its cluster of origin—in which case the watched cluster retains its particles—or it deposits at the watched cluster, having originated at a different one—in which case the watched cluster gains a particle (Figure 11).

We will couple the number of elements in the watched cluster with a lazy, one-dimensional random walk, which will never exceed n and never hit zero before the size of the watched cluster does. It will take no more than $(\log \log \rho_\ell)^n$ instances of the activated particle reaching the midway point, for the random walk to make n consecutive down-steps. This is a coarse estimate; with more effort, we could improve the n -dependence of this term, but it would not qualitatively change the result. On a high probability event, ρ_ℓ will be sufficiently large to ensure that $(\log \log \rho_\ell)^n \leq (\log \rho_\ell)^{o_n(1)}$. Then, because it will typically take no more than $(\log \rho_\ell)^{1+o_n(1)}$ steps to observe a visit to the midway point, we will wait a number of steps on the same order to observe the collapse of a cluster.

6.3. **Basic properties of clusters and collapse times.** We will work in the following setting.

- For brevity, if we write θ_m with no parenthetical argument, we will mean $\theta_m(\gamma n)$ for the constant γ given by

$$\gamma = 18 \max\{c_1, c_2^{-1}\} + 36, \tag{6.10}$$

where c_1 and c_2 are the constants in Theorems 4 and 5. Any constant larger than γ would also work in its place.

- U_0 has $n \geq 2$ elements and $\text{diam}(U_0)$ is at least θ_{4n} .

- The clustering parameter r equals $\Phi(\text{diam}(U_0))$, where we continue to denote by $\Phi(\cdot)$ the inverse function of $\theta_n(\cdot)$. In particular, r satisfies

$$r \geq \Phi(\theta_{4n}) = \theta_{3n} \geq e^n. \quad (6.11)$$

- We will assume that the initial configuration is exponentially clustered with parameter r as $U_0 \mapsto_r \{U_0^i, x_i, \theta^{(i)}\}_{i=1}^k$. In particular, we assume that clustering produces k clusters. We note that the choice of r guarantees $\text{diam}(U_0) > 2\theta_{n-1}(r)$ which, by Lemma 5.2, guarantees that $k > 1$.
- We denote a generic element of $\{1, 2, \dots, k-1\}$ by ℓ .

6.3.1. *Properties of cluster separation and diameter.* We will use the following terms to describe the separation of clusters.

Definition 6.5. *We define pairwise cluster separation and the least separation by*

$$\text{sep}(U_t^i) = \min_{j \neq i} \text{dist}(U_t^i, U_t^j) \quad \text{and} \quad \text{sep}(U_t) = \min_i \text{sep}(U_t^i).$$

(By convention, the distance to an empty set is ∞ , so the separation of a cluster is ∞ at all times following its collapse.) If U_t^i satisfies $\text{sep}(U_t^i) = \text{sep}(U_t)$, then we say that U_t^i is least separated. Whenever there are at least two clusters, at least two clusters will be least separated. The least separation at a cluster collapse time will be an important quantity; we will denote it by

$$\rho_\ell = \text{sep}(U_{\mathcal{T}_{\ell-1}}).$$

Next, we introduce the *expiry time* \mathfrak{t}_ℓ and the *truncated collapse time* \mathcal{T}_ℓ^- . As discussed in Section 6.2, if the least separation at time $\mathcal{T}_{\ell-1}$ is ρ_ℓ , then we will obtain a lower bound on the probability that a least separated cluster loses a particle, which holds uniformly from time $\mathcal{T}_{\ell-1}$ to the first of $\mathcal{T}_{\ell-1} + \mathfrak{t}_\ell$ and $\mathcal{T}_\ell - 1$ (i.e., the time immediately preceding the ℓ^{th} collapse), which we call the truncated collapse time, \mathcal{T}_ℓ^- . Here, \mathfrak{t}_ℓ is an $\mathcal{F}_{\mathcal{T}_{\ell-1}}$ -measurable random variable which will effectively be $(\log \rho_\ell)^2$. It will be rare for \mathcal{T}_ℓ to exceed $\mathcal{T}_{\ell-1} + \mathfrak{t}_\ell$, so \mathcal{T}_ℓ^- can be thought of as $\mathcal{T}_\ell - 1$.

Definition 6.6. *Given the $\mathcal{F}_{\mathcal{T}_{\ell-1}}$ data (in particular ρ_ℓ and $\mathcal{T}_{\ell-1}$), we define the expiry time \mathfrak{t}_ℓ to be the integer part of*

$$(\log \rho_\ell)^2 - 4 \log(\rho_\ell + \mathcal{T}_{\ell-1}) - \mathcal{T}_{\ell-1}.$$

We emphasize that \mathfrak{t}_ℓ should be thought of as $(\log \rho_\ell)^2$; the other terms will be much smaller and are included to simplify calculations which follow. Additionally, we define the ℓ^{th} truncated cluster collapse time to be

$$\mathcal{T}_\ell^- = (\mathcal{T}_{\ell-1} + \mathfrak{t}_\ell) \wedge (\mathcal{T}_\ell - 1).$$

Note that \mathcal{T}_ℓ^- is not a stopping time, but it will be useful to us anyway.

Cluster diameter and separation have complementary behavior in the sense that diameter increases at most linearly in time but may decrease abruptly, while separation decreases at most linearly in time but may increase abruptly. We express these properties in the following lemma.

Lemma 6.7. *Cluster diameter and separation obey the following properties.*

- (1) *Cluster diameter increases by at most 1 each step:*

$$\text{diam}(U_t^i) \leq \text{diam}(U_{t-1}^i) + 1. \quad (6.12)$$

- (2) *Cluster separation decreases by at most 1 each step:*

$$\text{dist}(U_t^i, U_t^j) \geq \text{dist}(U_{t-1}^i, U_{t-1}^j) - 1 \quad \text{and} \quad \text{sep}(U_t^i) \geq \text{sep}(U_{t-1}^i) - 1. \quad (6.13)$$

(3) For any two times s and t satisfying $\mathcal{T}_{\ell-1} \leq s < t < \mathcal{T}_\ell$ and any two clusters i and j :

$$\text{dist}(U_t^i, U_t^j) \leq \text{dist}(U_s^i, U_s^j) + \text{diam}(U_s^i) + \text{diam}(U_s^j) + (t - s).$$

Proof. The first two properties are obvious; we prove the third. Let i, j label two clusters which are nonempty at time $\mathcal{T}_{\ell-1}$ and let s, t satisfy the hypotheses. If there are m_i activations at the i^{th} cluster from time s to time t , then for any x' in U_t^i , there is an x in U_s^i such that $|x - x'| \leq m_i$. The same is true of any y' in the j^{th} cluster with m_j in the place of m_i . Since the sum of m_i and m_j is at most $t - s$, two uses of the triangle inequality give

$$\text{dist}(U_t^i, U_t^j) \leq \max_{x' \in U_t^i, y' \in U_t^j} |x' - y'| \leq \max_{x \in U_s^i, y \in U_s^j} |x - y| + t - s.$$

This implies property (3) because, by two more uses of the triangle inequality,

$$\max_{x \in U_s^i, y \in U_s^j} |x - y| \leq \text{dist}(U_s^i, U_s^j) + \text{diam}(U_s^i) + \text{diam}(U_s^j).$$

□

6.3.2. *Consequences of timely collapse.* If clusters collapse before their expiry times—i.e., if the event

$$\text{Timely}(\ell) = \cap_{m=1}^{\ell} \{\mathcal{T}_m - \mathcal{T}_{m-1} \leq t_m\}$$

occurs—then we will be able to control the separation (Lemma 6.8) and diameters (Lemma 6.10) of the clusters by combining the initial exponential separation of the clusters with the properties of Lemma 6.7. Moreover, we will be able to guarantee that every cluster that has not yet collapsed has an element with positive harmonic measure.

The next lemma states that, when cluster collapses are timely, cluster separation decreases little. To state it, we recall that $\text{sep}(U_t^i)$ is the distance between U_t^i and the nearest other cluster, and that ρ_ℓ is the least of these distances among all pairs of distinct clusters at time $\mathcal{T}_{\ell-1}$. In particular, $\text{sep}(U_{\mathcal{T}_{\ell-1}}^i) \geq \rho_\ell$ for each i .

Lemma 6.8. *For any cluster i , when $\text{Timely}_{\ell-1}$ occurs and when t is at most \mathcal{T}_ℓ^- ,*

$$\text{sep}(U_t^i) \geq (1 - e^{-n}) \text{sep}(U_{\mathcal{T}_{\ell-1}}^i). \quad (6.14)$$

Additionally, when $\text{Timely}_{\ell-1}$ occurs,

$$\rho_\ell \geq \frac{1}{2} \rho_1 \geq e^{\theta_{2n}}. \quad (6.15)$$

The factor of $1 - e^{-n}$ in (6.14) does not have special significance; other factors of $1 - o_n(1)$ would work, too. (6.14) and the first inequality in (6.15) are consequences of the fact (6.13) that separation decreases at most linearly in time and, when $\text{Timely}_{\ell-1}$ occurs, $\mathcal{T}_{\ell-1}$ is small relative to the separation of the remaining clusters. The second inequality in (6.15) follows from our choice of r in (6.11).

Proof Lemma 6.8. We will prove (6.14) by induction, using the fact that separation decreases at most linearly in time (6.13) and that (by the definition of \mathcal{T}_ℓ^-) at most t_ℓ steps elapse between $\mathcal{T}_{\ell-1}$ and \mathcal{T}_ℓ^- .

For the base case, take $\ell = 1$. Suppose cluster i is nonempty at time $\mathcal{T}_{\ell-1}$. We must show that, when $t \leq \mathcal{T}_1^-$,

$$\text{sep}(U_t^i) \geq (1 - e^{-n}) \text{sep}(U_0^i).$$

Because separation decreases at most linearly in time (6.13) and because $t \leq \mathcal{T}_1^-$,

$$\text{sep}(U_t^i) \geq \text{sep}(U_0^i) - t \geq \text{sep}(U_0^i) - \mathcal{T}_1^-.$$

This implies (6.14) for $\ell = 1$ because

$$\text{sep}(U_0^i) - \mathcal{T}_1^- \geq \left(1 - \frac{(\log \rho_1)^2}{\rho_1}\right) \text{sep}(U_0^i) \geq (1 - e^{-n}) \text{sep}(U_0^i).$$

The first inequality is a consequence of the definitions of \mathcal{T}_1^- , \mathbf{t}_1 , and ρ_1 , which imply $\mathcal{T}_1^- \leq \mathbf{t}_1 \leq (\log \rho_1)^2$ and $\text{sep}(U_0^i) \geq \rho_1$. Since the ratio of $(\log \rho_1)^2$ to ρ_1 decreases as ρ_1 increases, the second inequality follows from the bound $\rho_1 \geq e^{e^n}$, which is implied by the fact that U_0 satisfies the exponential separation property (5.1) with parameter $r \geq e^n$ (6.11).

The argument for $\ell > 1$ is similar. Assume (6.14) holds for $\ell - 1$. We have

$$\text{sep}(U_t^i) \geq \text{sep}(U_{\mathcal{T}_{\ell-1}^-}^i) - (t - \mathcal{T}_{\ell-1}^-) \geq \text{sep}(U_{\mathcal{T}_{\ell-1}^-}^i) - \mathbf{t}_\ell \geq \left(1 - \frac{(\log \rho_\ell)^2}{\rho_\ell}\right) \text{sep}(U_{\mathcal{T}_{\ell-1}^-}^i). \quad (6.16)$$

The first inequality is implied by (6.13). The second inequality follows from the definitions of \mathcal{T}_ℓ^- and \mathbf{t}_ℓ , which imply $\mathcal{T}_\ell^- - \mathcal{T}_{\ell-1}^- \leq \mathbf{t}_\ell \leq (\log \rho_\ell)^2$, and $t \leq \mathcal{T}_\ell^-$. The third inequality is due to the same upper bound on \mathbf{t}_ℓ and the fact that $\text{sep}(U_{\mathcal{T}_{\ell-1}^-}^i) \geq \rho_\ell$ by definition.

We will bound below ρ_ℓ to complete the induction step with (6.16), because the ratio of $(\log \rho_\ell)^2$ to ρ_ℓ decreases as ρ_ℓ increases. Specifically, we will prove (6.15). By definition, when $\text{Timely}_{\ell-1}$ occurs, so too does $\text{Timely}(\ell - 2)$. Accordingly, the induction hypothesis applies and we apply it $\ell - 1$ times:

$$\rho_{\ell-1} = \min_i \text{sep}(U_{\mathcal{T}_{\ell-2}^-}^i) \geq (1 - e^{-n})^{\ell-1} \min_i \text{sep}(U_0^i) = (1 - e^{-n})^{\ell-1} \rho_1.$$

The equalities follow from the definitions of $\rho_{\ell-1}$ and ρ_1 . We also have

$$\rho_\ell \geq \rho_{\ell-1} - \mathbf{t}_{\ell-1} \geq \left(1 - \frac{(\log \rho_{\ell-1})^2}{\rho_{\ell-1}}\right) \rho_{\ell-1} \geq (1 - e^{-n}) \rho_{\ell-1}.$$

The first inequality is due to (6.13) and the fact that at most $\mathbf{t}_{\ell-1}$ steps elapse between $\mathcal{T}_{\ell-2}^-$ and $\mathcal{T}_{\ell-1}^-$ when $\text{Timely}_{\ell-1}$ occurs. The second inequality is due to $\mathbf{t}_{\ell-1} \leq (\log \rho_{\ell-1})^2$ and the third is due to the fact that the ratio of $(\log \rho_{\ell-1})^2$ to $\rho_{\ell-1}$ decreases as $\rho_{\ell-1}$ increases.

Combining the two preceding displays and then using the fact that $\ell \leq n$ and $\rho_1 \geq e^{e^n}$, and the inequality $(1+x)^r \geq 1+rx$, which holds for $x > -1$ and $r > 1$, we find

$$\rho_\ell \geq (1 - e^{-n})^\ell \rho_1 \geq (1 - ne^{-n}) \rho_1.$$

Because $ne^{-n} \leq \frac{1}{2}$ when $n \geq 2$, this proves $\rho_\ell \geq \frac{1}{2} \rho_1$, which is the first inequality of (6.15). To prove the second inequality in (6.15), we note that ρ_1 is at least θ_{3n} by (6.11).

We now apply $\rho_\ell \geq \frac{1}{2} \rho_1$ to the ratio in (6.16):

$$\frac{(\log \rho_\ell)^2}{\rho_\ell} \leq \frac{2(\log \rho_1)^2}{\rho_1} \leq e^{-n}.$$

The second inequality uses $\rho_1 \geq e^{e^n}$. We complete the induction step, proving (6.14), by substituting this bound into (6.16). \square

When cluster collapses are timely, \mathcal{T}_ℓ^- is at most $(\log \rho_\ell)^2$, up to a factor depending on n .

Lemma 6.9. *When $\text{Timely}_{\ell-1}$ occurs,*

$$\mathcal{T}_\ell^- \leq 2n(\log \rho_\ell)^2. \quad (6.17)$$

The factor of 2 is for brevity; it could be replaced by $1 + o_n(1)$. The lower bound on the least separation ρ_ℓ at time $\mathcal{T}_{\ell-1}^-$ in (6.15) indicates that, while ρ_ℓ may be much larger than ρ_1 , it is at least half of ρ_1 . Since the expiry time \mathbf{t}_ℓ is approximately $(\log \rho_\ell)^2$, the truncated collapse time \mathcal{T}_ℓ^- —which is at most the sum of the first ℓ expiry times—should be of the same order, up to a factor depending on ℓ (which we will replace with n since $\ell \leq n$).

Proof of Lemma 6.9. We write

$$\mathcal{T}_\ell^- = \mathcal{T}_\ell^- - \mathcal{T}_{\ell-1} + \sum_{m=1}^{\ell-1} (\mathcal{T}_m - \mathcal{T}_{m-1}) \leq \sum_{m=1}^{\ell} t_m \leq \sum_{m=1}^{\ell} (\log \rho_m)^2.$$

The first inequality follows from the fact that, when $\text{Timely}_{\ell-1}$ occurs, $\mathcal{T}_m - \mathcal{T}_{m-1} \leq t_m$ for $m \leq \ell - 1$, and $\mathcal{T}_\ell^- - \mathcal{T}_{\ell-1} \leq t_\ell$. The second inequality holds because $t_m \leq (\log \rho_m)^2$ by definition.

Next, assume w.l.o.g. that cluster i is least separated at time $\mathcal{T}_{\ell-1}$, meaning $\rho_\ell = \text{sep}(U_{\mathcal{T}_{\ell-1}}^i)$. Since $\text{Timely}_{\ell-1}$ occurs, Lemma 6.8 applies and with its repeated use we establish (6.17):

$$\sum_{m=1}^{\ell} (\log \rho_m)^2 \leq \sum_{m=1}^{\ell} (\log \text{sep}(U_{\mathcal{T}_{m-1}}^i))^2 \leq \sum_{m=1}^{\ell} \left(\log \left(\left(1 + \frac{e^{-n}}{1-e^{-n}}\right)^{\ell-m} \rho_\ell \right) \right)^2 \leq \ell (\log(2\rho_\ell))^2 \leq 2n(\log \rho_\ell)^2.$$

The first inequality is due to the definition of ρ_m as the least separation at time \mathcal{T}_{m-1} . This step is helpful because it replaces each summand with one concerning the i^{th} cluster. The second inequality holds because, by Lemma 6.8,

$$\rho_\ell = \text{sep}(U_{\mathcal{T}_{\ell-1}}^i) \geq (1 - e^{-n})^{\ell-m} \text{sep}(U_{\mathcal{T}_{m-1}}^i) \implies \text{sep}(U_{\mathcal{T}_{m-1}}^i) \leq \left(1 + \frac{e^{-n}}{1-e^{-n}}\right)^{\ell-m} \rho_\ell.$$

The third inequality follows from $\ell \leq n$ and $(1 + \frac{e^{-n}}{1-e^{-n}})^n \leq 2$ when $n \geq 2$. The fourth inequality is due to $\ell \leq n$ and $\rho_\ell \geq e^{\theta 2n}$ from (6.15). (The factor of 2 could be replaced by $1 + o_n(1)$.) Combining the displays proves (6.17). \square

When cluster collapse is timely, we can bound cluster diameter at time $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_\ell^-]$ from above, in terms of its separation at time $\mathcal{T}_{\ell-1}$ or at time t .

Lemma 6.10. *For any cluster i , when $\text{Timely}_{\ell-1}$ occurs and when t is at most \mathcal{T}_ℓ^- ,*

$$\text{diam}(U_t^i) \leq (\log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^2. \quad (6.18)$$

Additionally, if x_i is the center of the i^{th} cluster resulting from the exponential clustering of U_0 , then when t is at most \mathcal{T}_ℓ^- ,

$$U_t^i \subseteq D_{x_i} \left((\log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^2 \right) \quad \text{and} \quad U_t \setminus U_t^i \subseteq D_{x_i} (0.99 \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^c. \quad (6.19)$$

Lastly, if i, j label any two clusters which are nonempty at time $\mathcal{T}_{\ell-1}$, then when t is at most \mathcal{T}_ℓ^- ,

$$\frac{\log \text{diam}(U_t^i)}{\log \text{dist}(U_t^i, U_t^j)} \leq \frac{2.1 \log \log \rho_\ell}{\log \rho_\ell}. \quad (6.20)$$

We use factors of 0.99 and 2.1 for concreteness; they could be replaced by $1 - o_n(1)$ and $2 + o_n(1)$. Lemma 6.10 implements the diameter and separation bounds we discussed in Section 6.2.2 (there, we used ρ in the place of $\text{sep}(U_{\mathcal{T}_{\ell-1}}^i)$). Before proving the lemma, we discuss some heuristics which explain (6.18) through (6.20).

If a cluster is initially separated by a distance ρ , then it has a diameter of at most $2 \log \rho$ by (5.1), which is negligible relative to an expiry time of order $(\log \rho)^2$. Diameter increases at most linearly in time by (6.12), so when cluster collapse is timely the diameter of U_t^i is at most $(\log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^2$. In fact, the definition of the expiry time subtracts the lower order terms, so the bound will be exactly this quantity. Moreover, since $(\log \rho)^2$ is negligible relative to the separation ρ , and as separation decreases at most linearly in time by (6.13), the separation of U_t^i should be at least $\text{sep}(U_{\mathcal{T}_{\ell-1}}^i)$, up to a constant which is nearly one.

Combining these bounds on diameter and separation suggests that the ratio of the diameter of U_t^i to its separation from another cluster U_t^j should be roughly the ratio of $(\log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^2$ to $\text{sep}(U_{\mathcal{T}_{\ell-1}}^i)$, up to a

constant factor. Because this ratio is decreasing in the separation (for separation exceeding, say, e^2) and because the separation at time $\mathcal{T}_{\ell-1}$ is at least ρ_ℓ , the ratio $\frac{(\log \rho_\ell)^2}{\rho_\ell}$ should provide a further upper bound, again up to a constant factor. These three observations correspond to (6.18) through (6.20).

Proof of Lemma 6.10. We first address (6.18) and use it to prove (6.19). We then combine the results to prove (6.20). We bound $\text{diam}(U_t^i)$ from above in terms of $\text{diam}(U_0^i)$ as

$$\text{diam}(U_t^i) \leq \text{diam}(U_0^i) + \mathcal{T}_\ell^- \leq \text{diam}(U_0^i) + \mathcal{T}_{\ell-1} + \mathfrak{t}_\ell. \quad (6.21)$$

The first inequality holds because diameter grows at most linearly in time (6.12) and because t is at most \mathcal{T}_ℓ^- . The second inequality is due to the definition of \mathcal{T}_ℓ^- . We then bound $\text{diam}(U_0^i)$ from above in terms of $\text{sep}(U_{\mathcal{T}_{\ell-1}}^i)$ as

$$\text{diam}(U_0^i) \leq 2 \log \text{sep}(U_0^i) \leq 2 \log (\text{sep}(U_{\mathcal{T}_{\ell-1}}^i) + \mathcal{T}_{\ell-1}). \quad (6.22)$$

The exponential separation property (5.1) implies the first inequality and (6.13) implies the second.

Combining the two preceding displays, we find

$$\text{diam}(U_t^i) \leq 2 \log (\text{sep}(U_{\mathcal{T}_{\ell-1}}^i) + \mathcal{T}_{\ell-1}) + \mathcal{T}_{\ell-1} + \mathfrak{t}_\ell.$$

Substituting the definition of \mathfrak{t}_ℓ , the right-hand side becomes

$$2 \log (\text{sep}(U_{\mathcal{T}_{\ell-1}}^i) + \mathcal{T}_{\ell-1}) + (\log \rho_\ell)^2 - 4 \log (\rho_\ell + \mathcal{T}_{\ell-1}).$$

By definition, ρ_ℓ is the least separation at time $\mathcal{T}_{\ell-1}$, so we can further bound $\text{diam}(U_t^i)$ from above by substituting $\text{sep}(U_{\mathcal{T}_{\ell-1}}^i)$ for ρ_ℓ :

$$\text{diam}(U_t^i) \leq (\log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^2 - 2 \log (\text{sep}(U_{\mathcal{T}_{\ell-1}}^i) + \mathcal{T}_{\ell-1}). \quad (6.23)$$

Dropping the negative term gives (6.18).

We turn our attention to (6.19). To obtain the first inclusion of (6.19), we observe that U_t^i is contained in the disk $D_{x_i}(\text{diam}(U_0^i) + \mathcal{T}_{\ell-1} + \mathfrak{t}_\ell)$, the radius of which is the quantity in (6.21) that we ultimately bounded above by $(\log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^2$.

Concerning the second inclusion of (6.19), we observe that for any y in $U_t \setminus U_t^i$, there is some y' in $U_{\mathcal{T}_{\ell-1}} \setminus U_{\mathcal{T}_{\ell-1}}^i$ such that $|y - y'|$ is at most \mathfrak{t}_ℓ , because t is at most \mathcal{T}_ℓ^- . By the triangle inequality and the bound on $|y - y'|$,

$$|x_i - y| \geq |x_i - y'| - |y - y'| \geq |x_i - y'| - \mathfrak{t}_\ell.$$

Next, we observe that the distance between x_i and y' is at least

$$|x_i - y'| \geq \text{sep}(U_{\mathcal{T}_{\ell-1}}^i) - \text{diam}(U_0^i).$$

The two preceding displays and (6.22) imply

$$|x_i - y| \geq \text{sep}(U_{\mathcal{T}_{\ell-1}}^i) - 2 \log (\text{sep}(U_{\mathcal{T}_{\ell-1}}^i) + \mathcal{T}_{\ell-1}) - \mathfrak{t}_\ell. \quad (6.24)$$

We continue (6.24) with

$$|x_i - y| \geq \text{sep}(U_{\mathcal{T}_{\ell-1}}^i) - (\log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^2 \geq \left(1 - \frac{(\log \rho_\ell)^2}{\rho_\ell}\right) \text{sep}(U_{\mathcal{T}_{\ell-1}}^i) \geq 0.99 \text{sep}(U_{\mathcal{T}_{\ell-1}}^i). \quad (6.25)$$

The first inequality follows from substituting the definition of \mathfrak{t}_ℓ into (6.24) and from $\text{sep}(U_{\mathcal{T}_{\ell-1}}^i) \geq \rho_\ell$. The second inequality holds because the ratio of $(\log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^2$ to $\text{sep}(U_{\mathcal{T}_{\ell-1}}^i)$ decreases as $\text{sep}(U_{\mathcal{T}_{\ell-1}}^i)$ increases and because $\text{sep}(U_{\mathcal{T}_{\ell-1}}^i) \geq \rho_\ell$. The fact (6.15) that ρ_ℓ is at least $e^{\theta_{2n}} \geq e^{\theta_2(\gamma)}$ when $\text{Timely}_{\ell-1}$ occurs implies that the ratio in (6.25) is at most 0.01 because $\theta_2(\gamma)^2 \leq 0.01 e^{\theta_2(\gamma)}$, which justifies the third inequality. (6.25) proves the second inclusion of (6.19).

Lastly, to address (6.20), we observe that any element x in U_t^i is within a distance $(\log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^2$ of x_i by (6.23). So, by (6.25) and simplifying with $\rho_\ell \geq e^{\theta_{2n}}$, the distance between U_t^i and U_t^j is at least

$$\text{sep}(U_{\mathcal{T}_{\ell-1}}^i) - 2(\log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^2 \geq 0.99 \text{sep}(U_{\mathcal{T}_{\ell-1}}^i).$$

Combining this with (6.18), and then using the fact that $\text{sep}(U_{\mathcal{T}_{\ell-1}}^i)$ is at least ρ_ℓ , gives

$$\frac{\log \text{diam}(U_t^i)}{\log \text{dist}(U_t^i, U_t^j)} \leq \frac{2 \log \log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i)}{\log(0.99 \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))} \leq \frac{2.1 \log \log \rho_\ell}{\log \rho_\ell}.$$

□

The next lemma concerns two properties of the midway point introduced in Section 6.2. We recall that the midway point (for the period beginning at time $\mathcal{T}_{\ell-1}$ and continuing until \mathcal{T}_ℓ^-) is a circle of radius $\frac{1}{2}\rho_\ell$, centered on the center x_i (given by the initial exponential clustering of U_0) of a cluster i which is least separated at time $\mathcal{T}_{\ell-1}$. The first property is the simple fact that, when collapse is timely, the midway point separates U_t^i from the rest of U_t until time \mathcal{T}_ℓ^- . This is clear because the midway point is a distance of $\frac{1}{2}\rho_\ell$ from $U_{\mathcal{T}_{\ell-1}}$ and \mathcal{T}_ℓ^- is no more than $(\log \rho_\ell)^2$ steps away from $\mathcal{T}_{\ell-1}$ when collapse is timely. The second property is the fact that a random walk from anywhere in the midway point hits U_t^i before the rest of U_t (excluding the site of the activated particle) with a probability of at most 0.51, which is reasonable because the random walk begins effectively halfway between U_t^i and the rest of U_t . In terms of notation, when activation occurs at u , the bound applies to the probability of the event

$$\{\tau_{U_t^i \setminus \{u\}} < \tau_{U_t \setminus (U_t^i \cup \{u\})}\}.$$

We will stipulate that u belongs to a cluster in U_t which is not a singleton as, otherwise, its activation at time t necessitates $t = \mathcal{T}_\ell$.

Lemma 6.11. *Suppose cluster i is least separated at time $\mathcal{T}_{\ell-1}$ and recall that x_i denotes the center of the i^{th} cluster, determined by the exponential clustering of U_0 . When $\text{Timely}_{\ell-1}$ occurs and when t is at most \mathcal{T}_ℓ^- :*

- (1) *the midway point $C(i; \ell) = C_{x_i}(\frac{1}{2}\rho_\ell)$ separates U_t^i from $U_t \setminus U_t^i$, and*
- (2) *for any u in U_t which does not belong to a singleton cluster and any y in $C(i; \ell)$,*

$$\mathbb{P}_y \left(\tau_{U_t^i \setminus \{u\}} < \tau_{U_t \setminus (U_t^i \cup \{u\})} \right) \leq 0.51. \quad (6.26)$$

Proof. Property (1) is an immediate consequence of (6.19) of Lemma 6.10, since $\frac{1}{2}\rho_\ell$ is at least $(\log \rho_\ell)^2$ and less than $0.99\rho_\ell$.

Now let u and y satisfy the hypotheses from #2, denote the center of the i^{th} cluster by x_i , and denote $C((\log \rho_\ell)^2)$ by B . To prove property (2), we will establish

$$\mathbb{P}_{y-x_i} (\tau_B < \tau_{z-x_i}) \leq 0.51, \quad (6.27)$$

for some $z \in U_t \setminus (U_t^i \cup \{u\})$. This bound implies (6.26) because, by (6.19), B separates $U_t^i - x_i$ from the rest of $U_t - x_i$.

We can express the probability in (6.27) in terms of hitting probabilities involving only three points:

$$\begin{aligned} \mathbb{P}_{y-x_i} (\tau_B < \tau_{z-x_i}) &= \mathbb{P}_{y-x_i} (\tau_o < \tau_{z-x_i}) + \mathbb{E}_{y-x_i} \left[\mathbb{P}_{S_{\tau_B}} (\tau_{z-x_i} < \tau_o) \mathbf{1}(\tau_B < \tau_{z-x_i}) \right] \\ &\leq \mathbb{P}_{y-x_i} (\tau_o < \tau_{z-x_i}) + \max_{v \in B} \mathbb{P}_v (\tau_{z-x_i} < \tau_o) \mathbb{P}_{y-x_i} (\tau_B < \tau_{z-x_i}). \end{aligned}$$

Rearranging, we find

$$\mathbb{P}_{y-x_i}(\tau_B < \tau_{z-x_i}) \leq \left(1 - \max_{v \in B} \mathbb{P}_v(\tau_{z-x_i} < \tau_o)\right)^{-1} \mathbb{P}_{y-x_i}(\tau_o < \tau_{z-x_i}). \quad (6.28)$$

We will choose z so that the points $y - x_i$ and $z - x_i$ will be at comparable distances from the origin and, consequently, $\mathbb{P}_{y-x_i}(\tau_o < \tau_{z-x_i})$ will be nearly $1/2$. In contrast, every element of B will be far nearer to the origin than to $z - x_i$, so $\mathbb{P}_v(\tau_{z-x_i} < \tau_o)$ will be nearly zero for every v in B . We will write these probabilities in terms of the potential kernel using Lemma 3.9. We will need bounds on the distances $|z - x_i|$ and $|z - y|$ to simplify the potential kernel terms; we take care of this now.

Suppose cluster j was nearest to cluster i at time $\mathcal{T}_{\ell-1}$. We then choose z to be the element of U_t^j nearest to U_t^i . Note that such an element exists because, when t is at most $\mathcal{T}_{\ell-1}^-$, every cluster surviving until time $\mathcal{T}_{\ell-1}$ survives until time t . By (6.19) of Lemma 6.10,

$$|z - x_i| \geq 0.99\rho_\ell.$$

Part (2) of Lemma A.1 then gives the lower bound

$$\mathbf{a}(z - x_i) \geq \frac{2}{\pi} \log(0.99\rho_\ell). \quad (6.29)$$

In the inter-collapse period before $\mathcal{T}_{\ell-1}^-$, the separation between z and y (initially $\frac{1}{2}\rho_\ell$) can grow by at most $\mathbf{t}_\ell + \text{diam}(U_{\mathcal{T}_{\ell-1}}^j)$:

$$|z - y| \leq \frac{1}{2}\rho_\ell + \mathbf{t}_\ell + \text{diam}(U_{\mathcal{T}_{\ell-1}}^j).$$

By (6.18), the diameter of cluster j at time $\mathcal{T}_{\ell-1}$ is at most $(\log \rho_\ell)^2$; this upper bound applies to \mathbf{t}_ℓ as well, so

$$|z - y| \leq \frac{1}{2}\rho_\ell + 2(\log \rho_\ell)^2 \leq 0.51\rho_\ell.$$

We obtained the second inequality using the fact (6.15) that, when Timely($\ell - 1$) occurs, ρ_ℓ is at least $e^{\theta_{2n}}$. (In what follows, we will use this fact without restating it.)

Accordingly, the difference between $\mathbf{a}(z - y)$ and $\mathbf{a}(y - x_i)$ satisfies

$$\mathbf{a}(z - y) - \mathbf{a}(y - x_i) \leq \frac{2}{\pi} \log(2 \cdot 0.51) + 4\lambda\rho_\ell^{-2} \leq \frac{2}{\pi}. \quad (6.30)$$

By Lemma 3.9, the first term of (6.28) equals

$$\mathbb{P}_{y-x_i}(\tau_o < \tau_{z-x_i}) = \frac{1}{2} + \frac{\mathbf{a}(z - y) - \mathbf{a}(y - x_i)}{2\mathbf{a}(z - x_i)}. \quad (6.31)$$

Substituting (6.29) and (6.30) into (6.31), we find

$$\mathbb{P}_{y-x_i}(\tau_o < \tau_{z-x_i}) \leq \frac{1}{2} + \frac{1}{\log \rho_\ell} \leq 0.501. \quad (6.32)$$

We turn our attention to bounding above the maximum of $\mathbb{P}_v(\tau_{z-x_i} < \tau_o)$ over v in B . This quantity should be close to 0 for every $v \in B$ because the elements of B are only a distance of roughly $(\log \rho_\ell)^2$ from the origin, but they are nearly ρ_ℓ from $z - x_i$. For any such v , Lemma 3.9 gives

$$\mathbb{P}_v(\tau_{z-x_i} < \tau_o) = \frac{1}{2} + \frac{\mathbf{a}(v) - \mathbf{a}(z - x_i - v)}{2\mathbf{a}(z - x_i)}. \quad (6.33)$$

By Lemma A.2, $\mathbf{a}(v)$ is at most $\mathbf{a}'((\log \rho_\ell)^2) + 2(\log \rho_\ell)^{-2}$. Then, since

$$|z - x_i - v| \geq 0.99\rho_\ell - (\log \rho_\ell)^2 \geq 0.98\rho_\ell,$$

we have

$$\mathbf{a}(z - x_i - v) - \mathbf{a}(v) \geq \frac{2}{\pi} \log(0.98\rho_\ell) - \frac{4}{\pi} \log \log \rho_\ell - 4(\log \rho_\ell)^{-2} \geq \frac{2 \cdot 0.99}{\pi} \log(0.99\rho_\ell). \quad (6.34)$$

Substituting (6.29) and (6.34) into (6.33), we find

$$\mathbb{P}_v(\tau_{z-x_i} < \tau_o) \leq \frac{1}{2} - \frac{0.99}{2} \leq 0.005.$$

This bound holds uniformly over v in B . Applying it and (6.32) to (6.28), we find

$$\mathbb{P}_{y-x_i}(\tau_B < \tau_{z-x_i}) \leq (1 - 0.005)^{-1} 0.501 \leq 0.51.$$

□

Combined with the separation lower bound (6.15) of Lemma 6.8, the inclusions (6.19) of Lemma 6.10 ensure that, when $\text{Timely}_{\ell-1}$ occurs, nonempty clusters at time $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_\ell^-]$ are contained in well separated disks. A natural consequence is that, when $\text{Timely}_{\ell-1}$ occurs, every nonempty cluster has positive harmonic measure in U_t . Later, we will use this fact in conjunction with Theorem 4 to control the activation step of the HAT dynamics.

Lemma 6.12. *Let I_ℓ be the set of indices of nonempty clusters at time $\mathcal{T}_{\ell-1}$. When $\text{Timely}_{\ell-1}$ occurs and when t is at most \mathcal{T}_ℓ^- , $\mathbb{H}_{U_t}(U_t^i) > 0$ for every $i \in I_\ell$.*

The proof shares some ideas with the proof of Lemma 3.13. Recall the definition of the *-exterior boundary (3.18) and define the disk D^i to be the one from (6.19)

$$D^i = D_{x_i}((\log \text{sep}(U_{\mathcal{T}_{\ell-1}}^i))^2) \quad \text{for each } i \in I_\ell. \quad (6.35)$$

For simplicity, assume $1 \in I_\ell$. Most of the proof is devoted to showing that there is a path Γ from $\partial_{\text{ext}}^* D^1$ to a large circle C about U_t , which avoids $\cup_{i \in I_\ell} D^i$ and thus avoids U_t . To do so, we will specify a candidate path from $\partial_{\text{ext}}^* D^1$ to C , and modify it as follows. If the path encounters a disk D^i , then we will reroute the path around $\partial_{\text{ext}}^* D^i$ (which will be connected and will not intersect another disk). The fact that $\partial_{\text{ext}}^* D^i$ will not intersect another disk is a consequence of the separation between clusters. The modified path encounters one fewer disk. We will iterate this argument until the path avoids every disk and therefore never returns to U_t .

Proof of Lemma 6.12. Suppose $\text{Timely}_{\ell-1}$ occurs and $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_\ell^-]$, and assume w.l.o.g. that $1 \in I_\ell$. Let $y \in U_t^1$ satisfy $\mathbb{H}_{U_t^1}(y) > 0$. (Note that there must be such a y because we are considering U_t^1 , not U_t .) For each $i \in I_\ell$, let D^i be the disk defined in (6.35). As $\mathbb{H}_{U_t^1}(y)$ is positive, there is a path from y to $\partial_{\text{ext}}^* D^1$ which does not return to U_t^1 . In a moment, we will show that $\partial_{\text{ext}}^* D^1$ is connected, so it will suffice to prove that there is a subsequent path from $\partial_{\text{ext}}^* D^1$ to $C = C_{x_1}(2 \text{diam}(U_t))$ which does not return to $E = \cup_{i \in I_\ell} D^i$. This suffices when $\text{Timely}_{\ell-1}$ occurs because then, by Lemma 6.10, $U_t^i \subseteq D^i$ for each $i \in I_\ell$, so $U_t \subseteq E$.

We make two observations. First, because each D^i is finite and *-connected, Lemma 2.23 of [Kes86] (alternatively, Theorem 4 of [Tim13]) states that each $\partial_{\text{ext}}^* D^i$ is connected. Second, $\partial_{\text{ext}}^* D^i$ is disjoint from E when $\text{Timely}_{\ell-1}$ occurs; this is an easy consequence of (6.19) and the separation lower bound (6.15).

We now specify a candidate path from $\partial_{\text{ext}}^* D^1$ to C and, if necessary, modify it to ensure that it does not return to E . Because $\mathbb{H}_{U_t^1}(y)$ is positive, there is a shortest path Γ from $\partial_{\text{ext}}^* D^1$ to C , which does not return to U_t^1 . Let L be the set of labels of disks encountered by Γ . If L is empty, then we are done. Otherwise, let i be the label of the first disk encountered by Γ , and let Γ_a and Γ_b be the first and last elements of Γ which

intersect $\partial_{\text{ext}}^* D^i$. By our first observation, $\partial_{\text{ext}}^* D^i$ is connected, so there is a shortest path Λ in $\partial_{\text{ext}}^* D^i$ from Γ_a to Γ_b . When edit Γ to form Γ' as

$$\Gamma' = (\Gamma_1, \dots, \Gamma_{a-1}, \Lambda_1, \dots, \Lambda_{|\Lambda|}, \Gamma_{b+1}, \dots, \Gamma_{|\Gamma|}).$$

Because Γ_b was the last element of Γ which intersected $\partial_{\text{ext}}^* D^i$, Γ' avoids D^i . Additionally, by our second observation, Λ avoids E , so if L' is the set of labels of disks encountered by Γ' , then $|L'| \leq |L| - 1$. If L' is empty, then we are done. Otherwise, we can relabel Γ to Γ' and L to L' in the preceding argument to continue inductively, obtaining Γ'' and $|L''| \leq |L| - 2$, and so on. Because $|L| \leq n$, we need to modify the path at most n times before the resulting path from y to C does not return to E . \square

The last result of this section bounds above escape probabilities; we will shortly specialize it for our setting. Note that ∂A_ρ denotes the exterior boundary of the ρ -fattening of A , not the ρ -fattening of ∂A .

Lemma 6.13. *If A is a subset of \mathbb{Z}^2 with at least two elements and if ρ is at least twice the diameter of A , then, for x in A ,*

$$\mathbb{P}_x \left(\tau_{\partial(A \setminus \{x\})_\rho} < \tau_{A \setminus \{x\}} \right) \leq \frac{\log \text{diam}(A) + 2}{\log \rho}. \quad (6.36)$$

The added 2 in (6.36) is unimportant. If $A \setminus \{x\} = \{o\}$, then a standard result (e.g., [Law13, Proposition 1.6.7]) states that the probability in question is approximately $\frac{\log |x|}{\log \rho} \leq \frac{\log \text{diam}(A)}{\log \rho}$. If A has more than two elements, then it can only be more difficult to escape $A \setminus \{x\}$, hence this approximate bound continues to hold. However, it is more convenient for us to directly prove this bound with explicit constants.

Proof of Lemma 6.13. We will replace the event in (6.36) with a more probable but simpler event and bound above its probability instead. By hypothesis, A has at least two elements, so for any x in A , there is some y in $A \setminus \{x\}$ nearest to x . To escape to $\partial(A \setminus \{x\})_\rho$ without hitting $A \setminus \{x\}$ it is necessary to escape to $C_y(\rho)$ without hitting y . Accordingly, for a random walk from x , the following inclusion holds

$$\{\tau_{\partial(A \setminus \{x\})_\rho} < \tau_{A \setminus \{x\}}\} \subseteq \{\tau_{C_y(\rho)} < \tau_y\}. \quad (6.37)$$

To prove (6.36) it therefore suffices to obtain the same bound for the larger event.

The hypothesis $\rho \geq 2 \text{diam}(A)$ ensures that $x - y$ lies in $D(\rho)$, so we can apply the optional stopping theorem to the martingale $\mathbf{a}(S_{j \wedge \tau_o})$ at the stopping time $\tau_{C(\rho)}$. Doing so, we find

$$\mathbb{P}_x(\tau_{C_y(\rho)} < \tau_y) = \mathbb{P}_{x-y}(\tau_{C(\rho)} < \tau_o) = \frac{\mathbf{a}(x-y)}{\mathbb{E}_{x-y}[\mathbf{a}(S_{\tau_{C(\rho)}}) \mid \tau_{C(\rho)} < \tau_o]}. \quad (6.38)$$

We apply Lemma A.2 with $r = \rho$ and $x = o$ to find

$$\mathbb{E}_{x-y}[\mathbf{a}(S_{\tau_{C(\rho)}}) \mid \tau_{C(\rho)} < \tau_o] \geq \mathbf{a}'(\rho) - \rho^{-1} \geq \frac{2}{\pi} \log \rho. \quad (6.39)$$

By (3.9) and the facts that $1 \leq |x - y| \leq \text{diam}(A)$ and $\kappa + \lambda \leq 1.1$, the numerator of (6.38), is at most

$$\mathbf{a}(x-y) \leq \frac{2}{\pi} \log |x-y| + \kappa + \lambda |x-y|^{-2} \leq \frac{2}{\pi} \log \text{diam}(A) + 1.1. \quad (6.40)$$

Substituting (6.39) and (6.40) into (6.38), and simplifying with $\frac{1.1\pi}{2} \leq 2$, we find

$$\mathbb{P}_x(\tau_{C_y(\rho)} < \tau_y) \leq \frac{\log \text{diam}(A) + 2}{\log \rho}.$$

Due to the inclusion (6.37), this implies (6.36). \square

6.4. Proof of Proposition 6.3. Recall that, for $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_{\ell}^-]$, the midway point is a circle which surrounds one of the clusters which is least separated at time $\mathcal{T}_{\ell-1}$. (We choose this cluster arbitrarily from the least separated clusters.) We call this cluster the watched cluster, to distinguish it from other clusters which are least separated at $\mathcal{T}_{\ell-1}$. The results of this section are phrased in these terms and through the following events.

Definition 6.14. For any $x \in \mathbb{Z}^2$, time $t \geq 0$, and any $1 \leq i \leq k$, define the activation events

$$\text{Act}(x, t) = \{x \text{ is activated at time } t\} \quad \text{and} \quad \text{Act}(i, t) = \bigcup_{x \in U_t^i} \text{Act}(x, t).$$

Additionally, define the deposition event

$$\text{Dep}(i, t) = \bigcup_{x \in U_t} \text{Act}(x, t) \cap \left\{ \tau_{U_t^i \setminus \{x\}} < \tau_{U_t \setminus (U_t^i \cup \{x\})} \right\}.$$

In words, the deposition event requires that, at time t , the activated particle deposits at the i^{th} cluster.

When $\text{Timely}_{\ell-1}$ occurs, if the i^{th} cluster is the watched cluster at time $\mathcal{T}_{\ell-1}$, then for any time $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_{\ell}^-]$, define the “midway” event as

$$\text{Mid}(i, t; \ell) = \bigcup_{x \in U_t} \text{Act}(x, t) \cap \left\{ \tau_{C(i; \ell)} < \tau_{U_t \setminus \{x\}} \right\}.$$

In words, the midway event specifies that, at time t , the activated particle reaches $C(i; \ell)$ before deposition.

We will now use the results of the preceding subsection to bound below the probability that activation occurs at the watched cluster and that the activated particle subsequently reaches the midway point. Essentially, Theorem 4 addresses the former probability and Theorem 5 addresses the latter. However, it is necessary to first ensure that the watched cluster has positive harmonic measure, so that at least one of its particles can be activated and the lower bound (1.2) of Theorem 4 can apply. This is handled by Lemma 6.12, the hypotheses of which are satisfied whenever $\text{Timely}_{\ell-1}$ occurs and $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_{\ell}^-]$. The hypotheses of Theorem 5 will be satisfied in this context so long as we estimate the probability of escape to a distance ρ which is at least twice the cluster diameter. The distance from the watched cluster to the midway point is roughly ρ_{ℓ} , while the cluster diameter is at most $(\log \rho_{\ell})^2$ by (6.19) of Lemma 6.10, so this will be the case.

The lower bounds from Theorems 4 and 5 will imply that a particle with positive harmonic measure is activated and reaches the midway point with a probability of at least

$$\exp(-c_1 n \log n + \log(c_2 n^{-2})) \cdot (\log \rho_{\ell})^{-1}$$

for constants c_1, c_2 . By our choice of γ (6.10), the first factor in the preceding display is at least α_n^{-1} , where

$$\alpha_n = e^{\gamma n \log n}. \tag{6.41}$$

Proposition 6.15. Let cluster i be least separated at time $\mathcal{T}_{\ell-1}$. When $\text{Timely}_{\ell-1}$ occurs and when $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_{\ell}^-]$, we have

$$\mathbf{P}(\text{Mid}(i, t; \ell) \cap \text{Act}(i, t) \mid \mathcal{F}_t) \geq (\alpha_n \log \rho_{\ell})^{-1}. \tag{6.42}$$

Proof. Fix ℓ , suppose the i^{th} cluster is least separated at time $\mathcal{T}_{\ell-1}$ and $\text{Timely}_{\ell-1}$ occurs, and let $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_{\ell}^-]$. For any $x \in U_t^i$, we have

$$\mathbf{P}(\text{Mid}(i, t; \ell) \cap \text{Act}(i, t) \mid \mathcal{F}_t) \geq \mathbf{P}(\text{Mid}(i, t; \ell) \mid \text{Act}(x, t), \mathcal{F}_t) \mathbf{P}(\text{Act}(x, t) \mid \mathcal{F}_t). \tag{6.43}$$

Let B denote the set of all points within distance ρ_ℓ of U_t^i . We have the following inclusion when $\text{Act}(x, t)$ occurs:

$$\left\{ \tau_{\partial B} < \tau_{U_t^i} \right\} \subseteq \left\{ \tau_{C(i;\ell)} < \tau_{U_t \setminus \{x\}} \right\} = \text{Mid}(i, t; \ell). \quad (6.44)$$

From (6.44), we have

$$\mathbf{P}(\text{Mid}(i, t; \ell) \mid \text{Act}(x, t), \mathcal{F}_t) \geq \mathbf{P}\left(\tau_{\partial B} < \tau_{U_t^i} \mid \text{Act}(x, t), \mathcal{F}_t\right) = \mathbb{P}_x\left(\tau_{\partial B} < \tau_{U_t^i}\right). \quad (6.45)$$

Now let x be an element of U_t^i which is exposed and which maximizes the right-hand side of (6.45). Such an element must exist because, by Lemma 6.12, when $\text{Timely}_{\ell-1}$ occurs and when $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_\ell^-]$, $\mathbb{H}_{U_t}(U_t^i)$ is positive. We aim to apply Theorem 5 to bound below the probability in (6.45). The hypotheses of Theorem 5 require $|U_t^i| \geq 2$ and $\rho_\ell \geq 2 \text{diam}(U_t^i)$. First, the cluster U_t^i must contain at least two elements as, otherwise, activation at x would necessitate $t = \mathcal{T}_\ell$. Second, ρ_ℓ is indeed at least twice the diameter of U_t^i because, when $\text{Timely}_{\ell-1}$ occurs, U_t^i is contained in a disk of radius $(\log \rho_\ell)^2$ by (6.19). Theorem 5 therefore applies to (6.45), giving

$$\mathbf{P}(\text{Mid}(i, t; \ell) \mid \text{Act}(x, t), \mathcal{F}_t) \geq c_2(n^2 \log \rho_\ell)^{-1}. \quad (6.46)$$

The harmonic measure lower bound (1.2) of Theorem 4 applies because x has positive harmonic measure. According to (1.2), the harmonic measure of x is at least

$$\mathbf{P}(\text{Act}(x, t) \mid \mathcal{F}_t) = \mathbb{H}_{U_t}(x) \geq e^{-c_1 n \log n}. \quad (6.47)$$

Combining (6.46) and (6.47), we find

$$\mathbf{P}(\text{Mid}(i, t; \ell) \cap \text{Act}(i, t) \mid \mathcal{F}_t) \geq c_2(n^2 \log \rho_\ell)^{-1} \cdot e^{-c_1 n \log n} \geq (\alpha_n \log \rho_\ell)^{-1}.$$

The second inequality is due to the definition of α_n (6.41). \square

Next, we will bound below the conditional probability that activation occurs at the watched cluster, given that the activated particle reaches the midway point.

Proposition 6.16. *Let cluster i be the watched cluster at time $\mathcal{T}_{\ell-1}$. When $\text{Timely}_{\ell-1}$ occurs and when $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_\ell^-]$, we have*

$$\mathbf{P}(\text{Act}(i, t) \mid \text{Mid}(i, t; \ell), \mathcal{F}_t) \geq (3\alpha_n \log \log \rho_\ell)^{-1}. \quad (6.48)$$

Proof. Suppose $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_\ell^-]$ and $\text{Timely}_{\ell-1}$ occurs. If we obtain a lower bound p_1 on $\mathbf{P}(\text{Mid}(i, t; \ell) \cap \text{Act}(i, t) \mid \mathcal{F}_t)$ and an upper bound p_2 on $\mathbf{P}\left(\bigcup_{j \neq i} \text{Mid}(i, t; \ell) \cap \text{Act}(j, t) \mid \mathcal{F}_t\right)$, then

$$\mathbf{P}(\text{Act}(i, t) \mid \text{Mid}(i, t; \ell), \mathcal{F}_t) \geq \frac{p_1}{p_1 + p_2}. \quad (6.49)$$

First, the probability $\mathbf{P}(\text{Mid}(i, t; \ell) \cap \text{Act}(i, t) \mid \mathcal{F}_t)$ is precisely the one we used to establish (6.43) in the proof of Proposition 6.15; p_1 is therefore at least $(\alpha_n \log \rho_\ell)^{-1}$.

Second, for any $j \neq i$, we use the trivial upper bound $\mathbf{P}(\text{Act}(j, t) \mid \mathcal{F}_t) \leq 1$ and address the midway component by writing

$$\mathbf{P}(\text{Mid}(i, t; \ell) \mid \text{Act}(j, t), \mathcal{F}_t) = \mathbf{E}\left[\mathbb{P}_X(\tau_{C(i;\ell)} < \tau_{U_t \setminus \{X\}}) \mid \text{Act}(j, t), \mathcal{F}_t\right]. \quad (6.50)$$

Use ρ to denote $\text{dist}(U_{\mathcal{T}_{\ell-1}}^i, U_{\mathcal{T}_{\ell-1}}^j)$ and B to denote the set of all points within a distance $\rho/3$ of $U_t^j \setminus \{X\}$. We can use Lemma 6.13 to bound the probability in (6.50) because the following inclusion holds:

$$\text{Act}(j, t) \cap \{\tau_{C(i;\ell)} < \tau_{U_t \setminus \{X\}}\} \subseteq \text{Act}(j, t) \cap \left\{ \tau_B < \tau_{U_t^j \setminus \{X\}} \right\}.$$

If $\text{Act}(j, t)$ occurs, then U_t^j must have least two elements as, otherwise, t would equal \mathcal{T}_ℓ . Then, because $\rho/3$ is at least twice its diameter, an application of Lemma 6.13 with $A = U_t^j$ and ρ yields

$$\mathbb{P}_x \left(\tau_B < \tau_{U_t^j \setminus \{x\}} \right) \leq \frac{\log \text{diam}(U_t^j) + 2}{\log \rho} \leq \frac{2.2 \log \log \rho_\ell}{\log \rho_\ell},$$

uniformly for x in U_t^j . The second inequality follows from (6.20), which bounds the ratio of $\log \text{diam}(U_t^j)$ to $\log \rho$ by $\frac{2.1 \log \log \rho_\ell}{\log \rho_\ell}$.

Applying the preceding bound to (6.50), we find

$$\mathbf{P}(\text{Mid}(i, t; \ell) \mid \text{Act}(j, t), \mathcal{F}_t) \leq \frac{2.2 \log \log \rho_\ell}{\log \rho_\ell} =: p_2.$$

Then, substituting p_1 and p_2 in (6.49), we conclude

$$\mathbf{P}(\text{Act}(i, t) \mid \text{Mid}(i, t; \ell), \mathcal{F}_t) \geq (1 + 2.2\alpha_n \log \log \rho_\ell)^{-1} \geq (3\alpha_n \log \log \rho_\ell)^{-1}.$$

□

We now use Lemma 6.11 to establish that an activated particle, upon reaching the midway point, deposits at the watched cluster with a probability of no more than 0.51.

Proposition 6.17. *Let cluster i be the watched cluster at time $\mathcal{T}_{\ell-1}$. When $\text{Timely}_{\ell-1}$ occurs and when $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_\ell^-]$, for x in U_t , we have*

$$\mathbf{P}(\text{Dep}(i, t) \mid \text{Mid}(i, t; \ell), \text{Act}(x, t), \mathcal{F}_t) \leq 0.51. \quad (6.51)$$

Proof. Using the definitions of $\text{Dep}(i, t)$ and $\text{Mid}(i, t; \ell)$, we write

$$\begin{aligned} & \mathbf{P}(\text{Dep}(i, t) \mid \text{Mid}(i, t; \ell), \text{Act}(x, t), \mathcal{F}_t) \\ &= \mathbf{P} \left(\bigcup_{y \in U_t} \text{Act}(y, t) \cap \left\{ \tau_{U_t^i \setminus \{y\}} < \tau_{U_t \setminus (U_t^i \cup \{y\})} \right\} \mid \bigcup_{y \in U_t} \text{Act}(y, t) \cap \left\{ \tau_{C(i; \ell)} < \tau_{U_t \setminus \{y\}} \right\}, \text{Act}(x, t), \mathcal{F}_t \right). \end{aligned}$$

Because $\text{Act}(y, t)$ only occurs for one particle y in U_t^i at any given time t , the right-hand side simplifies to

$$\mathbf{E} \left[\mathbb{P}_x \left(\tau_{U_t^i \setminus \{x\}} < \tau_{U_t \setminus (U_t^i \cup \{x\})} \mid \tau_{C(i; \ell)} < \tau_{U_t \setminus \{x\}} \right) \mid \text{Act}(x, t), \mathcal{F}_t \right].$$

We then apply the strong Markov property to $\tau_{C(i; \ell)}$ to find that the previous display equals

$$\mathbf{E} \left[\mathbb{P}_{S_{\tau_{C(i; \ell)}}} \left(\tau_{U_t^i \setminus \{x\}} < \tau_{U_t \setminus (U_t^i \cup \{x\})} \right) \mid \text{Act}(x, t), \mathcal{F}_t \right] \leq 0.51,$$

where the inequality follows from the estimate (6.26). □

The preceding three propositions realize the strategy of Section 6.2.2. We proceed to implement the strategy of Section 6.2.3. In brief, we will compare the number of particles in the watched cluster to a random walk and bound the collapse time using the hitting time of zero of the walk.

Let cluster i be the watched cluster at time $\mathcal{T}_{\ell-1}$ and denote by $(\eta_{\ell, m})_{m \geq 1}$ the consecutive times at which the midway event $\text{Mid}(i, \cdot; \ell)$ occurs. Set $\eta_{\ell, 0} \equiv \mathcal{T}_{\ell-1}$ and for all $m \geq 1$ define

$$\eta_{\ell, m} = \inf \{ t > \eta_{\ell, m-1} : \text{Mid}(i, t; \ell) \text{ occurs} \}.$$

Additionally, we denote the number of midway event occurrences by time t as

$$N_\ell(t) = \sum_{m=1}^{\infty} \mathbf{1}(\eta_{\ell, m} \leq t).$$

The number of elements in cluster i viewed at these times can be coupled to a lazy random walk $(W_m)_{m \geq 0}$ on $\{0, \dots, n\}$ from $W_0 \equiv |U_{\eta_{\ell,0}}^i|$, which takes down-steps with probability $q = (7\alpha_n \log \log \rho_\ell)^{-1}$ and up-steps with probability $1 - q$, unless it attempts to take a down-step at $W_m = 0$ or an up-step at $W_m = n$, in which case it remains where it is.

When $\text{Timely}_{\ell-1}$ occurs, at each time $\eta_{\ell, \cdot}$, the watched cluster has a chance of losing a particle of at least $0.49(3\alpha_n \log \log \rho_\ell)^{-1} \geq q$ (Propositions 6.15 and 6.16). The standard coupling of $|U_{\eta_{\ell,m}}^i|$ and W_m will then guarantee that $|U_{\eta_{\ell,m}}^i| \leq W_m$. However, this inequality will only hold when $N_\ell(\mathcal{T}_\ell^-) \geq m$.

Lemma 6.18. *Let cluster i be the watched cluster at time $\mathcal{T}_{\ell-1}$. There is a coupling of $(|U_{\eta_{\ell,m}}^i|)_{m \geq 0}$ and $(W_m)_{m \geq 0}$ such that, when $\text{Timely}_{\ell-1}$ and $\{N_\ell(\mathcal{T}_\ell^-) \geq M\}$ occur, $|U_{\eta_{\ell,m}}^i| \leq W_m$ for all $m \leq M$.*

Proof. Define

$$q_{\ell,m} = \mathbf{P} \left(\text{Act}(i, \eta_{\ell,m}) \cap \bigcup_{j \neq i} \text{Dep}(j, \eta_{\ell,m}) \mid \mathcal{F}_{\eta_{\ell,m}} \right).$$

In words, the event in the preceding display is the occurrence of $\text{Mid}(i, \eta_{\ell,m}; \ell)$, preceded by activation at cluster i and followed by deposition at cluster $j \neq i$; the watched cluster loses a particle when this event occurs.

Couple $(|U_{\eta_{\ell,m}}^i|)_{m \geq 0}$ and $(W_m)_{m \geq 0}$ using the standard monotone coupling. When $\text{Timely}_{\ell-1}$ occurs, by Propositions 6.16 and 6.17, the estimates (6.48) and (6.51) hold for all $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_\ell^-]$. In particular, these estimates hold at time $\eta_{\ell,m}$ for any $m \leq M$ when $\{N_\ell(\mathcal{T}_\ell^-) \geq M\}$ occurs. Accordingly, for any such m , we have $q_{\ell,m} \geq 0.49(3\alpha_n \log \log \rho_\ell)^{-1} \geq q$. \square

Denote by τ_0^U and τ_0^W the first hitting times of zero for $(|U_{\eta_{\ell,m}}^i|)_{m \geq 0}$ and $(W_m)_{m \geq 0}$. Under the coupling, τ_0^W cannot precede τ_0^U . So an upper bound on τ_0^W of m implies $\tau_0^U \leq m$ and therefore it takes no more than m occurrences of the midway event after $\mathcal{T}_{\ell-1}$ for the collapse of the watched cluster to occur. In other words, $\mathcal{T}_\ell - \mathcal{T}_{\ell-1}$ is at most $\eta_{\ell,m}$.

Lemma 6.19. *Let cluster i be the watched cluster at time $\mathcal{T}_{\ell-1}$. When $\text{Timely}_{\ell-1}$ and $\{N_\ell(\mathcal{T}_\ell^-) \geq M\}$ occur,*

$$\{\mathcal{T}_\ell - \mathcal{T}_{\ell-1} > \eta_{\ell,M}\} \subseteq \{\tau_0^W > M\}. \quad (6.52)$$

Proof. From Lemma 6.18, there is a coupling of $(|U_{\eta_{\ell,m}}^i|)_{m \geq 0}$ and $(W_m)_{m \geq 0}$ such that, when $\text{Timely}_{\ell-1}$ and $\{N_\ell(\mathcal{T}_\ell^-) \geq M\}$ occur, $|U_{\eta_{\ell,m}}^i| \leq W_m$ for all $m \leq M$. In particular,

$$\{\tau_0^U > M\} \subseteq \{\tau_0^W > M\}.$$

The inclusion (6.52) then follows from the fact that $\{\mathcal{T}_\ell - \mathcal{T}_{\ell-1} > \eta_{\ell,M}\} \subseteq \{\tau_0^U > M\}$. \square

We now show that τ_0^W , the hitting time of zero for W_m , is at most $(\log \log \rho_\ell)^n$, up to a factor depending on n , with high probability. With more effort, we could prove a much better bound (in terms of its dependence on n), but this improvement would not affect the conclusion of Proposition 6.3. By Lemma 6.19, the bound on τ_0^W will imply a bound on $\mathcal{T}_\ell - \mathcal{T}_{\ell-1}$. For brevity, denote $\beta_n = (8\alpha_n)^n$.

Lemma 6.20. *Let cluster i be the watched cluster at time $\mathcal{T}_{\ell-1}$ and let $K \geq 0$ be such that $M = \lfloor K\beta_n(\log \log \rho_\ell)^n \rfloor$ satisfies $\mathbf{P}(N_\ell(\mathcal{T}_\ell^-) \geq M \mid \mathcal{F}_{\mathcal{T}_{\ell-1}}) > 0$. Then*

$$\mathbf{P}(\mathcal{T}_\ell - \mathcal{T}_{\ell-1} > \eta_{\ell,M} \mid N_\ell(\mathcal{T}_\ell^-) \geq M, \mathcal{F}_{\mathcal{T}_{\ell-1}}) \mathbf{1}_{\text{Timely}_{\ell-1}} \leq e^{-\lfloor K \rfloor}. \quad (6.53)$$

The factor $(\log \log \rho_\ell)^n$ appears because $(W_m)_{m \geq 0}$ takes down-steps with a probability which is the reciprocal of $O_n(\log \log \rho_\ell)$, and we will require it to take n consecutive down-steps. Note that $\{N_\ell(\mathcal{T}_\ell^-) \geq M\}$ cannot occur if K is large enough, because $N_\ell(\mathcal{T}_\ell^-)$ cannot exceed $t_\ell \leq (\log \rho_\ell)^2$ (i.e., there can be no more occurrences of the midway event than there are HAT steps). The implicit bound on K is $(\log \rho_\ell)^{2-O_n(1)}$. We will apply the lemma with a K of approximately $(\log \rho_\ell)^\delta$ for some $\delta \in (0, 1)$.

Proof of Lemma 6.20. Denote the distribution of $(W_m)_{m \geq 0}$ by \mathbb{P}^W . If $\text{Timely}_{\ell-1}$ and $\{N_\ell(\mathcal{T}_\ell^-) \geq M\}$ occur, then by Lemma 6.19, we have the inclusion (6.52):

$$\{\mathcal{T}_\ell - \mathcal{T}_{\ell-1} > \eta_{\ell, M}\} \subseteq \{\tau_0^W > M\}.$$

Since $(W_m)_{m \geq 0}$ is never greater than n , it never takes more than n down-steps for W_m to hit zero. Since $W_{m+1} = W_m - 1$ with a probability of q_W whenever $m \leq M - n$, we have

$$\mathbb{P}^W(\tau_0^W > m + n \mid \tau_0^W > m) \leq 1 - q_W^n.$$

Applying this fact $\lfloor M/n \rfloor$ times, we find that

$$\mathbb{P}^W(\tau_0^W > M) \leq (1 - q_W^n)^{\lfloor \frac{M}{n} \rfloor} \leq e^{-\lfloor K \rfloor}.$$

For the second inequality, we used the fact that $\lfloor \beta_n(\log \log \rho_\ell)^n/n \rfloor$ is at least q_W^{-n} and therefore $\lfloor M/n \rfloor$ is at least $\lfloor K \rfloor \cdot q_W^{-n}$. Combining this with (6.52) gives (6.53). \square

For convenience, we will treat $(\log \rho_\ell)^{1+2\delta}$ as an integer, as the distinction will be unimportant. Define the time

$$\mathfrak{s}_\ell(\delta) = 5n\alpha_n(\log \rho_\ell)^{1+2\delta}$$

and the event that occurrences of the midway event are frequent:

$$\text{FreqMid}_\ell(\delta) = \{\eta_{\ell, m} - \eta_{\ell, m-1} \leq \mathfrak{s}_\ell(\delta), 1 \leq m \leq N_\ell(\mathcal{T}_\ell^-)\}.$$

When $\text{Timely}_{\ell-1}$ occurs, the midway event occurs frequently, with high probability.

Proposition 6.21. *Let cluster i be the watched cluster at time $\mathcal{T}_{\ell-1}$ and let $\delta = (2n)^{-4}$. Then*

$$\mathbf{P}(\text{FreqMid}_\ell(\delta) \mid \mathcal{F}_{\mathcal{T}_{\ell-1}}) \geq \left(1 - e^{-4n(\log \rho_\ell)^{2\delta}}\right) \mathbf{1}_{\text{Timely}_{\ell-1}}. \quad (6.54)$$

Proof. Proposition 6.15 states that

$$\mathbf{P}(\text{Mid}(i, t; \ell) \mid \mathcal{F}_t) \geq (\alpha_n \log \rho_\ell)^{-1}, \quad (6.55)$$

when $\text{Timely}_{\ell-1}$ and $t \in [\mathcal{T}_{\ell-1}, \mathcal{T}_\ell^-]$ occur. The probability that the time between consecutive midway events exceeds $\mathfrak{s}_\ell(\delta)$ therefore satisfies

$$\mathbf{P}\left(\eta_{\ell, N_\ell(t)} - \eta_{\ell, N_\ell(t)-1} > \mathfrak{s}_\ell(\delta) \mid \mathcal{F}_{\eta_{\ell, N_\ell(t)-1}}\right) \leq \left(1 - \frac{1}{\alpha_n \log \rho_\ell}\right)^{\mathfrak{s}_\ell(\delta)} \leq e^{-5n(\log \rho_\ell)^{2\delta}}. \quad (6.56)$$

Since $\text{FreqMid}_\ell(\delta)^c$ is a union of $N_\ell(\mathcal{T}_\ell^-) \leq (\log \rho_\ell)^2$ such events, a union bound and (6.56) imply that

$$\mathbf{P}(\text{FreqMid}_\ell(\delta)^c \mid \mathcal{F}_t) \leq (\log \rho_\ell)^2 e^{-5n(\log \rho_\ell)^{2\delta}} \leq e^{-4n(\log \rho_\ell)^{2\delta}}.$$

The second inequality follows from the fact that $(\log \rho_\ell)^2 \leq e^{(\log \rho_\ell)^{2\delta}}$ since $\log \rho_\ell \geq (2\delta)^{-2}$. \square

Proof of Proposition 6.3. Let $\delta = (2n)^{-4}$ and define the event that the ℓ^{th} collapse is fast by

$$\text{FastCol}_\ell(\delta) = \{\mathcal{T}_\ell - \mathcal{T}_{\ell-1} \leq (\log \rho_\ell)^{1+6\delta}\}, \quad 1 \leq \ell < k.$$

We will use Lemma 6.20 and Proposition 6.21 to prove the bound

$$\mathbf{P}(\text{FastCol}_\ell(\delta) \mid \mathcal{F}_{\mathcal{T}_{\ell-1}}) \geq \left(1 - e^{-3n(\log \rho_1)^{2\delta}}\right) \mathbf{1}_{\text{Timely}_{\ell-1}}. \quad (6.57)$$

This implies the proposition because \mathcal{T}_{k-1} is at most $(\log d)^{1+7\delta}$ when every collapse is fast, i.e., when $\text{FastCol}(\delta) = \cap_{\ell=1}^{k-1} \text{FastCol}_\ell(\delta)$ occurs, and because (6.57) implies

$$\mathbf{P}(\text{FastCol}(\delta) \mid \mathcal{F}_0) \geq 1 - e^{-2n(\log \rho_1)^{2\delta}}. \quad (6.58)$$

Note that this lower bound is at least the one in (6.2) because the clustering parameter r (6.11) satisfies $r \leq \log \rho_1$.

Indeed, $\text{FastCol}(\delta) \subseteq \{\mathcal{T}_{k-1} \leq (\log d)^{1+7\delta}\}$ because, when $\text{FastCol}(\delta)$ occurs, the collapse time satisfies

$$\mathcal{T}_{k-1} = \sum_{\ell=1}^{k-1} (\mathcal{T}_\ell - \mathcal{T}_{\ell-1}) \leq \sum_{\ell=1}^{k-1} (\log \rho_\ell)^{1+6\delta} \leq 2n(\log d)^{1+6\delta} \leq (\log d)^{1+7\delta}.$$

The second inequality holds because ρ_ℓ is never more than twice the diameter d of the initial configuration U_0 when $\text{FastCol}(\delta)$ occurs. The third follows from the fact that $d \geq \theta_{4n}$ by assumption, so $\log d \geq (2n)^{\frac{1}{\delta}}$ in particular.

To see why (6.57) implies (6.58), note that $\cap_{i < j} \text{FastCol}_i(\delta) \subseteq \text{Timely}_{j-1}$ for each $j < k$, so

$$\mathbf{P}\left(\cap_{i < j} \text{FastCol}_i(\delta) \cap \text{FastCol}_j(\delta)^c \mid \mathcal{F}_0\right) \leq \mathbf{E}\left[\mathbf{P}\left(\text{FastCol}_j(\delta)^c \mid \mathcal{F}_{\mathcal{T}_{j-1}}\right) \mathbf{1}_{\text{Timely}_{j-1}} \mid \mathcal{F}_0\right].$$

By (6.57), the right-hand side is at most $e^{-3n(\log \rho_1)^{2\delta}}$. We use this bound and the fact that $k \leq n$ to conclude (6.58) as

$$\begin{aligned} \mathbf{P}(\text{FastCol}(\delta)^c \mid \mathcal{F}_0) &= \sum_{j=1}^{k-1} \mathbf{P}\left(\cap_{i < j} \text{FastCol}_i(\delta) \cap \text{FastCol}_j(\delta)^c \mid \mathcal{F}_0\right) \\ &\leq \sum_{j=1}^{k-1} e^{-3n(\log \rho_1)^{2\delta}} \leq e^{-2n(\log \rho_1)^{2\delta}}. \end{aligned}$$

It remains to prove (6.57). Consider the bound

$$\begin{aligned} \mathbf{P}(\text{FastCol}_\ell(\delta)^c \mid \mathcal{F}_{\mathcal{T}_{\ell-1}}) \\ \leq \mathbf{P}(\text{FastCol}_\ell(\delta)^c \cap \text{FreqMid}_\ell(\delta) \mid \mathcal{F}_{\mathcal{T}_{\ell-1}}) + \mathbf{P}(\text{FreqMid}_\ell(\delta)^c \mid \mathcal{F}_{\mathcal{T}_{\ell-1}}). \end{aligned} \quad (6.59)$$

Proposition 6.21 states that the second term on the right-hand side of (6.59) is at most $e^{-4n(\log \rho_\ell)^{2\delta}}$ when $\text{Timely}_{\ell-1}$ occurs. We claim that the same is true of the first term. Note that when the ℓ^{th} collapse is slow and when occurrences of the midway event are frequent, there must be many such occurrences. In other words, $\text{FastCol}_\ell(\delta)^c \cap \text{FreqMid}_\ell(\delta) \subseteq \{N_\ell(\mathcal{T}_\ell^-) \geq M_\ell(\delta)\}$ where $M_\ell(\delta) = \lceil \frac{(\log \rho_\ell)^{1+6\delta}}{s_\ell(\delta)} \rceil > (\log \rho_\ell)^{3\delta}$. Consequently, this term satisfies

$$\mathbf{P}(\text{FastCol}_\ell(\delta)^c \cap \text{FreqMid}_\ell(\delta) \mid \mathcal{F}_{\mathcal{T}_{\ell-1}}) \leq \mathbf{P}(\mathcal{T}_\ell - \mathcal{T}_{\ell-1} > \eta_{\ell, M_\ell(\delta)} \mid N_\ell(\mathcal{T}_\ell^-) \geq M_\ell(\delta), \mathcal{F}_{\mathcal{T}_{\ell-1}}).$$

It is easy to check that $M_\ell(\delta)$ is at least $[K_\ell(\delta)\beta_n(\log \log \rho_\ell)^n]$ for $K_\ell(\delta) = 4n(\log \rho_\ell)^{2\delta}$ when $\text{Timely}_{\ell-1}$ occurs, because $\rho_\ell \geq e^{\theta_{2n}}$ by Lemma 6.8. Hence, by Lemma 6.20, the preceding bound is at most $e^{-4n(\log \rho_\ell)^{2\delta}}$ when $\text{Timely}_{\ell-1}$ occurs. We conclude that, when $\text{Timely}_{\ell-1}$ occurs,

$$\mathbf{P}(\text{FastCol}_\ell(\delta)^c \mid \mathcal{F}_{\mathcal{T}_{\ell-1}}) \leq 2e^{-4n(\log \rho_\ell)^{2\delta}} \leq 2e^{-4n(\log(\rho_1/2))^{2\delta}} \leq e^{-3n(\log \rho_1)^{2\delta}}.$$

The first inequality follows from the upper bound of $e^{-4n(\log \rho \ell)^{2\delta}}$ on both of the terms on the right-hand side of (6.59); the second and third inequalities follow from Lemma 6.8. \square

7. EXISTENCE OF THE STATIONARY DISTRIBUTION

In this section, we will prove Theorem 2, which has two parts. The first part states the existence of a unique stationary distribution, π_n , supported on the equivalence classes of non-isolated configurations, $\widehat{\text{NonIso}}(n)$, to which the HAT dynamics converges from any n -element configuration. The second part provides a tail bound on the diameter of configurations under π_n . We will prove these parts separately, as the following two propositions.

Proposition 7.1. *For all $n \geq 1$, from any n -element subset U , HAT converges to a unique stationary distribution π_n on $\widehat{\text{NonIso}}(n)$, given by*

$$\pi_n(\widehat{U}) = \frac{1}{\mathbf{E}_{\widehat{U}} \mathcal{T}_{\widehat{U}}}, \text{ for } \widehat{U} \in \widehat{\text{NonIso}}(n), \quad (7.1)$$

in terms of the return time $\mathcal{T}_{\widehat{U}} = \inf\{t \geq 1 : \widehat{U}_t = \widehat{U}\}$.

Proposition 7.2. *For any $d \geq 2\theta_{4n}$,*

$$\pi_n(\text{diam}(\widehat{U}) \geq d) \leq \exp\left(-\frac{d}{(\log d)^{1+o_n(1)}}\right). \quad (7.2)$$

For the sake of concreteness, this is true with $6n^{-4}$ in the place of $o_n(1)$.

Proof of Theorem 2. Combine Propositions 7.1 and 7.2. \square

It will be relatively easy to establish Proposition 7.2 using the inputs to the proof of Proposition 7.1 and Corollary 6.4, so we focus on presenting the key components of the proof of Proposition 7.1.

By standard theory for countable state space Markov chains, to prove Proposition 7.1, we must prove that the HAT dynamics is positive recurrent, irreducible, and aperiodic. We address each of these in turn.

Proposition 7.3 (Positive recurrent). *For any $U \in \text{NonIso}(n)$, $\mathbf{E}_{\widehat{U}} \mathcal{T}_{\widehat{U}} < \infty$.*

To prove Proposition 7.3, we will estimate the return time to an arbitrary n -element configuration \widehat{U} by separately estimating the time it takes to reach the line segment \widehat{L}_n from \widehat{U} , where $L_n = \{y e_2 : y \in \{0, 1, \dots, n-1\}\}$ in terms of $e_2 = (0, 1)$, and the time it takes to hit \widehat{U} from \widehat{L}_n . The first estimate is the content of the following result.

Proposition 7.4. *There is a constant c such that, if U is a configuration in $\text{NonIso}(n)$ with a diameter of R , then, for all $K \geq \max\{R, \theta_{5n}(cn)\}$,*

$$\mathbf{P}_U\left(\mathcal{T}_{\widehat{L}_n} \leq K^4\right) \geq 1 - e^{-2K}. \quad (7.3)$$

The second estimate is provided by the next proposition.

Proposition 7.5. *There is a constant c such that, if U is a configuration in $\text{NonIso}(n)$ with a diameter of R , then, for all $K \geq \max\{e^{R^{2.1}}, \theta_{5n}(cn)\}$,*

$$\mathbf{P}_{\widehat{L}_n}\left(\mathcal{T}_{\widehat{U}} \leq K^5\right) \geq 1 - e^{-K}. \quad (7.4)$$

The proof of Proposition 7.3 applies (7.3) and (7.4) to the tail sum formula for $\mathbf{E}_{\widehat{U}}\mathcal{T}_{\widehat{U}}$.

Proof of Proposition 7.3. Let $U \in \text{NonIso}(n)$. We have

$$\mathbf{E}_{\widehat{U}}\mathcal{T}_{\widehat{U}} = \sum_{t=0}^{\infty} \mathbf{P}_{\widehat{U}}(\mathcal{T}_{\widehat{U}} > t) \leq \sum_{t=0}^{\infty} \left(\mathbf{P}_{\widehat{U}}(\mathcal{T}_{\widehat{L}_n} > \frac{t}{2}) + \mathbf{P}_{\widehat{L}_n}(\mathcal{T}_{\widehat{U}} > \frac{t}{2}) \right). \quad (7.5)$$

Suppose U has a diameter of at most R and let $J = \max\{e^{R^{2.1}}, \theta_{5n}(cn)\}$, where c is the larger of the constants from Propositions 7.4 and 7.5. We group the sum (7.5) over t into blocks:

$$\mathbf{E}_{\widehat{U}}\mathcal{T}_{\widehat{U}} \leq O(J^5) + \sum_{K=J}^{\infty} \sum_{t=2K^5}^{2(K+1)^5} \left(\mathbf{P}_{\widehat{U}}(\mathcal{T}_{\widehat{L}_n} > \frac{t}{2}) + \mathbf{P}_{\widehat{L}_n}(\mathcal{T}_{\widehat{U}} > \frac{t}{2}) \right).$$

By (7.3) and (7.4) of Propositions 7.4 and 7.5, each of the $O(K^4)$ summands in the K^{th} block is at most

$$\mathbf{P}_{\widehat{U}}(\mathcal{T}_{\widehat{L}_n} > K^5) + \mathbf{P}_{\widehat{L}_n}(\mathcal{T}_{\widehat{U}} > K^5) \leq 2e^{-K}. \quad (7.6)$$

Substituting (7.6) into (7.5), we find

$$\mathbf{E}_{\widehat{U}}\mathcal{T}_{\widehat{U}} \leq O(J^5) + O(1) \sum_{K=J}^{\infty} K^4 e^{-K} < \infty.$$

□

Propositions 7.4 and 7.5 also imply irreducibility.

Proposition 7.6 (Irreducible). *For any $n \geq 1$, HAT is irreducible on $\widehat{\text{NonIso}}(n)$.*

Proof. Let $\widehat{U}, \widehat{V} \in \widehat{\text{NonIso}}(n)$. It suffices to show that HAT reaches \widehat{V} from \widehat{U} in a finite number of steps with positive probability. By Propositions 7.4 and 7.5, there is a finite number of steps $K = K(U, V)$ such that

$$\mathbf{P}_U(\mathcal{T}_{\widehat{L}_n} < K) > 0 \quad \text{and} \quad \mathbf{P}_{L_n}(\mathcal{T}_{\widehat{V}} < K) > 0.$$

By the Markov property applied to $\mathcal{T}_{\widehat{L}_n}$, the preceding bounds imply that $\mathbf{P}_{\widehat{U}}(\mathcal{T}_{\widehat{V}} < 2K) > 0$. □

Lastly, because aperiodicity is a class property, it follows from irreducibility and the simple fact that \widehat{L}_n is aperiodic.

Proposition 7.7 (Aperiodic). *\widehat{L}_n is aperiodic.*

Proof. We claim that $\mathbf{P}_{L_n}(U_1 = L_n) \geq \frac{1}{4}$, which implies that $\mathbf{P}_{\widehat{L}_n}(\widehat{U}_1 = \widehat{L}_n) \geq \frac{1}{4} > 0$. Indeed, every element of L_n neighbors another, so, regardless of which one is activated, we can dictate one random walk step which results in transport to the site of activation and $U_1 = L_n$. □

The preceding results constitute a proof of Proposition 7.1.

Proof of Proposition 7.1. Combine Propositions 7.3, 7.6, and 7.7. □

The subsections are organized as follows. In Section 7.1, we prove some preliminary results, including a key lemma which states that it is possible to reach any configuration $U \in \text{NonIso}(n)$ from L_n , in a number of steps depending only on n and $\text{diam}(U)$. These results support the proofs of Propositions 7.4 and 7.5 in Sections 7.2 and Sections 7.3, respectively. In Section 7.4, we prove Proposition 7.2.

7.1. Preliminaries of hitting estimates for configurations. The purpose of this section is to prove that if HAT can reach \widehat{V} from \widehat{L}_n , then it does so in at most $O_n(\text{rad}(V))$ steps with a probability of at least $e^{-O_n(\text{rad}(V)^2)}$. (Recall that the radius of $A \subseteq \mathbb{Z}^2$ is defined as $\text{rad}(A) = \sup\{|x| : x \in A\}$.) We split the proof into two lemmas. The first states that if HAT can form V_1 from V_0 in one step, then it does so with a probability of at least $e^{-O_n(\text{rad}(V_0))}$.

Lemma 7.8. *There is a constant c such that, if V_0 and V_1 are subsets of \mathbb{Z}^2 with $n \geq 2$ elements that satisfy $\mathbf{P}_{V_0}(U_1 = V_1) > 0$, then*

$$\mathbf{P}_{V_0}(U_1 = V_1) \geq e^{-cn^2 \text{rad}(V_0)}. \quad (7.7)$$

With more work, we could improve the lower bound in (7.7) to $\Omega_n(\frac{1}{\log \text{rad}(V_0)})$, but this would make no difference in our application of the lemma.

Proof of Lemma 7.8. Since $\mathbf{P}_{V_0}(U_1 = V_1)$ is positive, the transition from V_0 to V_1 can be realized by activation at some x and transport to some y , where x is exposed in V_0 and there is a path of length $O(n\text{rad}(V_0))$ from x to y that lies outside of $V_0 \setminus \{x\}$ (Lemma 3.13). The former implies that $\mathbb{H}_{V_0}(x) \geq e^{-O(n \log n)}$ (Theorem 4), while the latter implies that $\mathbb{P}_x(S_{\tau_{V_0 \setminus \{x\}} - 1} = y) \geq 4^{-O(n\text{rad}(V_0))}$. The claimed bound (7.7) then follows from

$$\mathbf{P}_{V_0}(U_1 = V_1) \geq \mathbb{H}_{V_0}(x) \mathbb{P}_x(S_{\tau_{V_0 \setminus \{x\}} - 1} = y) \geq e^{-O(n \log n)} 4^{-O(n\text{rad}(V_0))}.$$

□

The second lemma proves that every $V \in \text{NonIso}(n)$ can be reached from L_n in $O_n(\text{rad}(V))$ steps, through a sequence of configurations with diameters of $O_n(\text{rad}(V))$.

Lemma 7.9. *For any number of elements $n \geq 2$ and configuration V in $\text{NonIso}(n)$, if the radius of V is at most an integer $r \geq 10n$, then there is a sequence of $k \leq 100nr$ activation sites x_1, \dots, x_k and transport sites y_1, \dots, y_k which can be “realized” by HAT from $V_0 = L_n$ to $V_k = V$ in the following sense: if we set $V_i = (V_{i-1} \setminus \{x_i\}) \cup \{y_i\}$ for each $i \in \{1, \dots, k\}$, then each transition probability $\mathbf{P}_{V_{i-1}}(U_i = V_i)$ is positive. Additionally, each V_i is contained in $D(r + 10n)$.*

The factors of 10 and 100 in the lemma statement are for convenience and have no further significance. We will prove Lemma 7.9 by induction on n . Informally, we will remove one element of L_n to facilitate the use of the induction hypothesis, forming most of V before returning the removed element. There is a complication in this step, as we cannot allow the induction hypothesis to “interact” with the removed element. We will resolve this problem by proving a slightly stronger claim than the lemma requires.

The proof will overcome two main challenges. First, removing an element from a configuration V in $\text{NonIso}(n)$ can produce a configuration in $\text{Iso}(n - 1)$, in which case the induction hypothesis will not apply. Indeed, there are configurations of $\text{NonIso}(n)$ for which the removal of any exposed, non-isolated element produces a configuration of $\text{Iso}(n - 1)$ (such a V is depicted in Figure 12). Second, if an isolated element is removed alone, it cannot be returned to form V by a single step of the HAT dynamics. To see how these difficulties interact, suppose that $\partial_{\text{exp}} V$ (defined as $\{x \in V : \mathbb{H}_V(x) > 0\}$) contains only one non-isolated element (say, at v), which is part of a two-element connected component of V . We cannot remove it and still apply the induction hypothesis, as $V \setminus \{v\}$ belongs to $\text{Iso}(n - 1)$. We then have no choice but to remove an isolated element.

When we are forced to remove an isolated element, we will apply the induction hypothesis to form a configuration for which the removed element can be “treadmilled” to its proper location, chaperoned by a element which is non-isolated in the final configuration and so can be returned once the removed element reaches its destination.

We briefly explain what we mean by treading a pair of elements. Consider elements v_1 and $v_1 + e_2$ of a configuration V . If $\mathbb{H}_V(v_1)$ is positive and if there is a path from v_1 to $v_1 + 2e_2$ which lies outside of $V \setminus \{v_1\}$, then we can activate at v_1 and transport to $v_1 + 2e_2$. The result is that the pair $\{v_1, v_1 + 2e_2\}$ has shifted by e_2 . Call the new configuration V' . If $v_1 + e_2$ is exposed in V' and if there is a path from $v_1 + e_2$ to $v_1 + 3e_2$ in $V' \setminus \{v_1 + e_2\}$, we can analogously shift the pair $\{v_1 + e_2, v_1 + 2e_2\}$ by another e_2 .

Proof of Lemma 7.9. The proof is by induction on $n \geq 2$. We will actually prove a stronger claim, because it facilitates the induction step. To state the claim, we denote by $W_i = V_{i-1} \setminus \{x_i\}$ the HAT configuration “in between” V_{i-1} and V_i and by E_i the event that, during the transition from V_{i-1} to V_i , the transport step takes place inside of $B_i = D(r + 10n) \setminus W_i$:

$$E_i = \{\{S_0, \dots, S_{\tau_{W_i}}\} \subseteq B_i\}.$$

We claim that Lemma 7.9 is true even if the conclusion $\mathbf{P}_{V_{i-1}}(U_i = V_i) > 0$ is replaced by $\mathbf{P}_{V_{i-1}}(U_i = V_i, E_i) > 0$.

To prove this claim, we will show that, for any V satisfying the hypotheses, there are sequences of at most $100nr$ activation sites x_1, \dots, x_k , transport sites y_1, \dots, y_k , and random walk paths $\Gamma^1, \dots, \Gamma^k$ such that the activation and transport sites can be realized by HAT from $V_0 = L_n$ to $V_k = V$, and such that each Γ^i is a finite random walk path from x_i to y_i which lies in B_i . While it is possible to explicitly list these sequences of sites and paths in the proof which follows, the depictions in upcoming Figures 12 and 13 are easier to understand and so we omit some cumbersome details regarding them.

Concerning the base case of $n = 2$, note that $\text{NonIso}(2)$ has the same elements as the equivalence class \widehat{L}_2 , so $x_1 = e_2, y_1 = e_2, \Gamma^1 = \emptyset$ works. Suppose the claim holds up to $n - 1$ for $n \geq 3$. There are two cases:

- (1) There is a non-isolated v in $\partial_{\text{exp}} V$ such that $V \setminus \{v\}$ belongs to $\text{NonIso}(n - 1)$.
- (2) For every non-isolated v in $\partial_{\text{exp}} V$, $V \setminus \{v\}$ belongs to $\text{Iso}(n - 1)$.

It will be easy to form V using the induction hypothesis in Case 1. In Case 2, we will need to use the induction hypothesis to form a set related to V , and subsequently form V from this related set. An instance of Case 2 is depicted in Figure 12.




FIGURE 12. An instance of Case 2. If any non-isolated element of $\partial_{\text{exp}} V$ is removed, the resulting set is isolated. We use the induction hypothesis to form $V' = (V \setminus \{v_{ne}, u\}) \cup \{v_{sw} - e_2\}$. The subsequent steps to obtain V from V' are depicted in Figure 13.

Case 1. Let r be an integer exceeding $10n$ and the radius of V and denote $R = r + 10(n - 1)$. Recall that $V_0 = L_n$. Our strategy is to place one element of L_n outside of $D(R)$ and then apply the induction hypothesis to L_{n-1} to form most of V . This explains the role of the event E_i —it ensures that the element outside of the disk does not interfere with our use of the induction hypothesis.

The first $R - n + 2$ steps, which produce $L_{n-1} \cup \{a\}$ from L_n , are identical to those of Case 1. We apply the induction hypothesis to form the set

$$V' = (V \setminus \{v_{ne}, u\}) \cup \{v_{sw} - e_2\},$$

which is depicted in Figure 12. Here, v_{ne} is the easternmost of the northernmost elements of V , v_{sw} is the westernmost of the southernmost elements of V , and u is any non-isolated element of $\partial_{\text{exp}}V$ (e.g., $u = v_{ne}$ is allowed if v_{ne} is non-isolated).

The remaining steps are depicted in Figure 13. By the induction hypothesis and because the radius of V' is at most $r + 1$, there are sequences of at most $100(n - 1)(r + 1)$ activation and transport sites, which can be realized by HAT from $L_{n-1} \cup \{a\}$ to $V' \cup \{a\}$, and a corresponding sequence of finite random walk paths which lie in $D(R + 1)$.

Next, we activate at a and transport to $v_{sw} - 2e_2$, which is possible because $v_{sw} - 2e_2$ is exposed and non-isolated in V' . Like in Case 1, the existence of a finite random walk path from a to $v_{sw} - 2e_2$ which lies in $D(R + 3) \setminus V' \subseteq D(r + 10n)$ is implied by Lemma 3.13. Denote the resulting configuration by V'' .

The choice of v_{sw} ensures that $v_{sw} - e_2$ and $v_{sw} - 2e_2$ are the only elements of V'' which lie in the quadrant defined by

$$Q_{sw} = (v_{sw} - e_2) + \{v \in \mathbb{Z}^2 : v \cdot e_1 \leq 0, v \cdot e_2 \leq 0\}.$$

Additionally, the quadrant defined by

$$Q_{ne} = v_{ne} + \{v \in \mathbb{Z}^2 : v \cdot e_1 \geq 0, v \cdot e_2 \geq 0\}$$

contains no elements of V'' . As depicted in Figure 13, this enables us to treadmill the pair $\{v_{sw} - e_2, v_{sw} - 2e_2\}$ from Q_{sw} to $D(R + 3)^c$ and then to $\{v_{ne}, v_{ne} + e_2\}$ in Q_{ne} , without the pair encountering the remaining elements of V'' . It is clear that this can be accomplished by fewer than $10(R + 3)$ activation and transport sites, with corresponding finite random walk paths which lie in $D(R + 6)$. The resulting configuration is $V''' = V \cup \{v_{ne} + e_2\} \setminus \{u\}$.

Lastly, we activate at $v_{ne} + e_2$ and transport to u , which is possible because the former is exposed in V''' and the latter is exposed and non-isolated in V . As before, the fact that there is a finite random walk path in $D(r + 10n)$ which accomplishes the transport step is a consequence of Lemma 3.13. The resulting configuration is V .

In summary, there are sequences of fewer than $(R - n + 2) + 100(n - 1)(r + 1) + 10(R + 3) + 2 \leq 100nr$ (the inequality follows from the assumption that $r \geq 10n$) activation and transport sites which can be realized by HAT from L_n to V , as well as corresponding finite random walk paths which remain in $D(r + 10n)$. This proves the claim in Case 2. \square

We can combine Lemma 7.8 and Lemma 7.9 to bound below the probability of forming a configuration from a line.

Lemma 7.10. *There is a constant c such that, if V is a configuration in $\text{NonIso}(n)$ with $n \geq 2$ and a diameter of at most $R \geq 10n$, then*

$$\mathbf{P}_{\widehat{L}_n}(\mathcal{T}_{\widehat{V}} \leq 200nR) \geq e^{-cn^3R^2}.$$

Proof. The hypotheses of Lemma 7.9 require an integer upper bound r on the radius of V of at least $10n$. We are free to assume that V contains the origin, in which case a choice of $r = \lfloor R \rfloor + 1$ works, due to the assumption $R \geq 10n$. We apply Lemma 7.9 with this r to obtain a sequence of configurations $L_n = V_0, V_1, \dots, V_{k-1}, V_k = V$ such that $k \leq 100nr$, and such that $V_i \subseteq D(r + 10n)$ and $\mathbf{P}_{V_{i-1}}(U_i = V_i) > 0$ for each i .

Because the transition probabilities are positive and because they concern sets V_{i-1} in the disk of radius $r + 10n \leq 3R$, Lemma 7.8 implies that each transition probability is at least $e^{-c_1 n^2 R}$ for a constant c_1 . We use this fact in the following string of inequalities:

$$\mathbf{P}_{L_n}(\mathcal{T}_{\widehat{V}} \leq 200nR) \geq \mathbf{P}_{L_n}(\mathcal{T}_{\widehat{V}} \leq k) \geq \mathbf{P}_{L_n}(\widehat{U}_k = \widehat{V}) \geq e^{-100nr \cdot c_1 n^2 R} \geq e^{-c_2 n^3 R^2}.$$

The first inequality holds because $k \leq 100nr \leq 200nR$; the second because $\{\widehat{U}_k = \widehat{V}\} \subseteq \{\mathcal{T}_{\widehat{V}} \leq k\}$; the third follows from the Markov property, $k \leq 100nr$, and the preceding bound from Lemma 7.8; the fourth from $r \leq 2R$. \square

7.2. Proof of Proposition 7.4. We now use Lemma 7.8 to obtain a tail bound on the time it takes for a given configuration to reach \widehat{L}_n . Our strategy is to repeatedly attempt to observe the formation of \widehat{L}_n in n consecutive steps. If the attempt fails then, because the diameter of the resulting set may be larger—worsening the estimate (7.7)—we will wait until the diameter becomes smaller before the next attempt.

Proof of Proposition 7.4. To avoid confusion of U and U_t , we will use V_0 instead of U . We introduce a sequence of times, with consecutive times separated by at least n steps (which is enough time to attempt to form \widehat{L}_n) and at which the diameter of the configuration is at most $\theta_1 = \theta_{4n}(c_1 n)$ (where c_1 is the constant in Corollary 6.4). These will be the times at which we attempt to observe the formation of \widehat{L}_n . Define $\eta_0 = \inf\{t \geq 0 : \text{diam}(U_t) \leq \theta_1\}$ and, for all $i \geq 1$, the times

$$\eta_i = \inf\{t \geq \eta_{i-1} + n : \text{diam}(U_t) \leq \theta_1\}.$$

We use these times to define three events. Two of the events involve a parameter K which we assume is at least the maximum of R and θ_2 , where θ_2 equals $\theta_{5n}(cn)$ with $c = c_1 + 2c_2$ and c_2 is the constant guaranteed by Lemma 7.8. (The constant c is the one which appears in the statement of the proposition.) In particular, K is at least the maximum diameter $\theta_1 + n$ of a configuration at time $\eta_{i-1} + n$.

According to Corollary 6.4, it takes no longer than $3K(\log(3K))^{1+2n^{-4}} \leq K^2$ steps for the diameter to fall below θ_1 , except with a probability of at most e^{-3K} . The first event is the event that it takes an unusually long time for the diameter to fall below θ_1 for the first time:

$$E_1(K) = \{\eta_0 > K^2\}.$$

The second is the event that an unusually long time elapses between $\eta_{i-1} + n$ and η_i for some $1 \leq i \leq m$:

$$E_2(m, K) = \bigcup_{i=1}^m \{\eta_i - (\eta_{i-1} + n) > K^2\}.$$

The third is the event that we do not observe the formation of \widehat{L}_n in $m \geq 1$ attempts:

$$E_3(m) = \bigcap_{i=1}^m \{\mathcal{T}_{\widehat{L}_n} > \eta_{i-1} + n\}.$$

Call $E(m, K) = E_1(K) \cup E_2(m, K) \cup E_3(m)$. When none of these events occur, we can bound $\mathcal{T}_{\widehat{L}_n}$:

$$\mathcal{T}_{\widehat{L}_n} \mathbf{1}_{E(m, K)^c} \leq \left(\eta_0 + \sum_{i=1}^m (\eta_i - (\eta_{i-1} + n)) \right) \mathbf{1}_{E(m, K)^c} + (m+1)n \leq (m+1)(K^2 + n). \quad (7.8)$$

We will show that if $m = 3K \log \theta_2$, then $\mathbf{P}_{V_0}(E(m, K))$ is at most e^{-K} . Substituting this choice of m into (7.8) and simplifying with $K \geq \theta_2$, we obtain a further upper bound of

$$\mathcal{T}_{\widehat{L}_n} \mathbf{1}_{E(m, K)^c} \leq K^4. \quad (7.9)$$

By (7.9), if we show that $\mathbf{P}_{V_0}(E(m, K)) \leq e^{-2K}$, then we are done. We start with a bound on $\mathbf{P}_{V_0}(E_1(K))$. Applying Corollary 6.4 with $3K$ in the place of t , r in the place of d , and $3K = \max\{3K, R\}$ in the place of $\max\{t, d\}$, gives

$$\mathbf{P}_{V_0}(E_1(K)) \leq e^{-3K}. \quad (7.10)$$

We will use Corollary 6.4 and a union bound to bound $\mathbf{P}_{V_0}(E_2(m, K))$. Because diameter grows at most linearly in time, the diameter of $U_{\eta_{i-1}+n} \in \mathcal{F}_{\eta_{i-1}}$ is at most $\theta_1 + n \leq 3K$. Consequently, Corollary 6.4 implies

$$\mathbf{P}_{V_0}\left(\eta_i - (\eta_{i-1} + n) > K^2 \mid \mathcal{F}_{\eta_{i-1}+n}\right) \leq e^{-3K}. \quad (7.11)$$

A union bound over the constituent events of $E_2(m, K)$ and (7.11) give

$$\mathbf{P}_{V_0}(E_2(m, K)) \leq me^{-3K}. \quad (7.12)$$

To bound the probability of $E_3(m)$, we will use Lemma 7.8. First, we need to identify a suitable sequence of HAT transitions. For any $0 \leq j \leq m-1$, given \mathcal{F}_{η_j} , set $V'_0 = U_{\eta_j} \in \mathcal{F}_{\eta_j}$. There are pairs $\{(x_i, y_i) : 1 \leq i \leq n\}$ such that, setting $V'_i = (V'_{i-1} \setminus \{x_i\}) \cup \{y_i\}$ for $1 \leq i \leq n$, each transition probability $\mathbf{P}_{V'_{i-1}}(U_i = V'_i)$ is positive and $V'_n \in \widehat{L}_n$. The diameter of V'_i is at most $\theta_1 + n \leq 2\theta_1$ for each i , so Lemma 7.8 implies that the transition probabilities satisfy

$$\mathbf{P}_{V'_{i-1}}(U_i = V'_i) \geq e^{-2c_2 n^2 \theta_1}. \quad (7.13)$$

By the strong Markov property and (7.13),

$$\mathbf{P}_{V_0}\left(\mathcal{T}_{\widehat{L}_n} \leq \eta_j + n \mid \mathcal{F}_{\eta_j}\right) \geq \mathbf{P}_{V_0}\left(U_{\eta_j+1} = V'_1, \dots, U_{\eta_j+n} = V'_n \mid \mathcal{F}_{\eta_j}\right) \geq e^{-2c_2 n^3 \theta_1}. \quad (7.14)$$

Because $E_3(j) \in \mathcal{F}_{\eta_j}$, (7.14) implies

$$\mathbf{P}_{V_0}\left(\mathcal{T}_{\widehat{L}_n} \leq \eta_j + n \mid E_3(j)\right) \geq e^{-2c_2 n^3 \theta_1}. \quad (7.15)$$

Using (7.15) and observing that $\log \theta_2 \geq e^{2c_2 n^3 \theta_1}$, hence $m \geq 3K e^{2c_2 n^3 \theta_1}$, we calculate

$$\mathbf{P}_{V_0}(E_3(m)) = \prod_{j=0}^{m-1} \mathbf{P}_{V_0}\left(\mathcal{T}_{\widehat{L}_n} > \eta_j + n \mid E_3(j)\right) \leq \left(1 - e^{-2c_2 n^3 \theta_1}\right)^m \leq e^{-3K}. \quad (7.16)$$

Combining (7.10), (7.12), and (7.16), and simplifying using the fact that $m \leq 3K^2$, we find

$$\mathbf{P}_{V_0}(E(m, K)) \leq (m+2)e^{-3K} \leq e^{-2K}.$$

□

7.3. Proof of Proposition 7.5. To prove this proposition, we will attempt to observe the formation of \widehat{U} from \widehat{L}_n and wait for the set to collapse if its diameter becomes too large, as we did in proving Proposition 7.4. Note that, at the time that the set collapses, it does not necessarily form \widehat{L}_n , so we will need to use Proposition 7.4 to return to \widehat{L}_n before another attempt at forming \widehat{U} . For convenience, we package these steps together in the following lemma.

Lemma 7.11. *There is a constant c such that, if V_0 is a configuration in $\text{NonIso}(n)$ with a diameter of R , then for any $K \geq \max\{R, \theta_{5n}(cn)\}$,*

$$\mathbf{P}_{V_0}\left(\mathcal{T}_{\widehat{L}_n} \leq 2K^4\right) \geq 1 - e^{-K}. \quad (7.17)$$

Proof. Call $\theta = \theta_{5n}(cn)$ where c is the constant guaranteed by Proposition 7.4. First, we wait until the diameter falls to θ . Note that $2K(\log(2K))^{1+2n^{-4}} \leq K^2$, so Corollary 6.4 implies that

$$\mathbf{P}_{V_0} \left(\mathcal{T}(\theta) \leq K^2 \right) \geq 1 - e^{-2K}. \quad (7.18)$$

Second, from $U_{\mathcal{T}(\theta)}$, we wait until the configuration forms a line. By Proposition 7.4,

$$\mathbf{P}_{U_{\mathcal{T}(\theta)}} \left(\mathcal{T}_{\widehat{L}_n} \leq K^4 \right) \geq 1 - e^{-2K}. \quad (7.19)$$

Combining these bounds gives (7.17). \square

Proof of Proposition 7.5. We will use V to denote the target configuration instead of U , to avoid confusion with U_t . Recall that, for any configuration V in $\text{NonIso}(n)$ with a diameter upper bound of $r \geq 10n$, Lemma 7.10 gives a constant c_1 such that

$$\mathbf{P}_{\widehat{L}_n} (\mathcal{T}_{\widehat{V}} \leq 200nr) \geq e^{-c_1 n^3 r^2}.$$

Since $10nR \geq 10n$ is a diameter upper bound on V , we can apply the preceding inequality with $r = 10nR$:

$$\mathbf{P}_{\widehat{L}_n} (\mathcal{T}_{\widehat{V}} \leq 2000n^2 R) \geq e^{-c_1 n^5 R^2}. \quad (7.20)$$

With this result in mind, we denote $k = 2000n^2 R$ and define a sequence of times by

$$\zeta_0 \equiv 0 \quad \text{and} \quad \zeta_i = \inf\{t \geq \zeta_{i-1} + k : \widehat{U}_t = \widehat{L}_n\} \quad \text{for all } i \geq 1.$$

Here, the buffer of k steps is the period during which we attempt to observe the formation of V . After each failed attempt, because the diameter increases by at most 1 with each step, the diameter of U_{ζ_i+k} can be no larger than $k+n$.

We define two rare events in terms of these times and a parameter K , which we assume to be at least $\max\{e^{R^{2.1}}, \theta_{5n}(c_2 n)\}$, where c_2 is the greater of c_1 and the constant from Lemma 7.11. In particular, under this assumption, K is greater than $e^{4c_1 n^5 R^2}$ and $k+n$ —a fact we will use later.

The first rare event is the event that an unusually long time elapses between $\zeta_{i-1} + k$ and ζ_i , for some $i \leq m$:

$$F_1(m, K) = \bigcup_{i=1}^m \{\zeta_i - (\zeta_{i-1} + k) > 72K^3\}.$$

The second is the event that we do not observe the formation of \widehat{V} in $m \geq 1$ attempts:

$$F_2(m) = \bigcap_{i=1}^m \{\mathcal{T}_{\widehat{V}} > \zeta_{i-1} + k\}.$$

Call $F(m, K) = F_1(m, K) \cup F_2(m)$. When $F(m, K)^c$ occurs, we can bound $\mathcal{T}_{\widehat{V}}$ as

$$\mathcal{T}_{\widehat{V}} \mathbf{1}_{F(m, K)^c} = \sum_{i=0}^{m-1} (\zeta_i - (\zeta_{i-1} + k)) \mathbf{1}_{E(m, K)^c} + mk \leq 72mK^3 + mk. \quad (7.21)$$

We will show that if m is taken to be $2Ke^{c_1 n^5 R^2}$, then $\mathbf{P}_{\widehat{L}_n}(F(m, K))$ is at most e^{-K} . Substituting this value of m into (7.21) and simplifying with $K \geq k$ and then $K \geq e^{4c_1 n^5 R^2}$ gives

$$\mathcal{T}_{\widehat{V}} \mathbf{1}_{F(m, K)^c} \leq K^4 e^{2c_1 n^5 R^2} \leq K^5. \quad (7.22)$$

By (7.22), if we prove $\mathbf{P}_{\widehat{L}_n}(F(m, K)^c) \leq e^{-K}$, then we are done. We start with a bound on $\mathbf{P}_{\widehat{L}_n}(F_1(m, K))$. By the strong Markov property applied to the stopping time $\zeta_{i-1} + k$,

$$\mathbf{P}_{\widehat{L}_n} \left(\zeta_i - (\zeta_{i-1} + k) > 72K^3 \mid \mathcal{F}_{\zeta_{i-1}+k} \right) = \mathbf{P}_{U_{\zeta_{i-1}+k}}(\zeta_1 > 72K^3) \leq e^{-2K}. \quad (7.23)$$

The inequality is due to Lemma 7.11, which applies to $U_{\zeta_{i-1}+k}$ and K because $U_{\zeta_{i-1}+k}$ is a non-isolated configuration with a diameter of at most $k+n$ and because $K \geq \max\{k+n, \theta_{5n}(c_2n)\}$. From a union bound over the events which comprise $F_1(m, K)$ and (7.23), we find

$$\mathbf{P}_{\widehat{L}_n}(F_1(m, K)) \leq me^{-2K}. \quad (7.24)$$

To bound $\mathbf{P}_{\widehat{L}_n}(F_2(m))$, we apply the strong Markov property to ζ_j and use (7.20):

$$\mathbf{P}_{\widehat{L}_n}(\mathcal{T}_{\widehat{V}} \leq \zeta_j + k \mid \mathcal{F}_{\zeta_j}) \geq \mathbf{P}_{\widehat{L}_n}(\mathcal{T}_{\widehat{V}} \leq k) \geq 1 - e^{-c_1 n^5 R^2}. \quad (7.25)$$

Then, because $F_2(j) \in \mathcal{F}_{\zeta_j}$ and by (7.25),

$$\mathbf{P}_{\widehat{L}_n}(\mathcal{T}_{\widehat{V}} \leq \zeta_j + k \mid F_2(j)) \geq 1 - e^{-c_1 n^5 R^2}. \quad (7.26)$$

We use (7.26) to calculate

$$\mathbf{P}_{\widehat{L}_n}(F_2(m)) = \prod_{j=0}^{m-1} \mathbf{P}_{\widehat{L}_n}(\mathcal{T}_{\widehat{V}} > \zeta_j + k \mid F_2(j)) \leq \prod_{j=0}^{m-1} (1 - e^{-c_1 n^5 R^2}) \leq e^{-2K}. \quad (7.27)$$

The second inequality is due to the choice $m = 2Ke^{c_1 n^5 R^2}$.

Recall that $F(m, K)$ is the union of $F_1(m, K)$ and $F_2(m)$. We have

$$\mathbf{P}_{\widehat{L}_n}(F(m, K)) \leq \mathbf{P}_{\widehat{L}_n}(F_1(m, K)) + \mathbf{P}_{\widehat{L}_n}(F_2(m)) \leq me^{-2K} + e^{-2K} \leq e^{-K}.$$

The first inequality is a union bound; the second is due to (7.24) and (7.27); the third holds because $m+1 \leq e^K$. \square

7.4. Proof of Proposition 7.2. We now prove a tightness estimate for the stationary distribution—that is, an upper bound on $\pi_n(\text{diam}(\widehat{U}) \geq d)$. By Proposition 7.1, the stationary probability $\pi_n(\widehat{U})$ of any non-isolated, n -element configuration \widehat{U} is the reciprocal of $\mathbf{E}_{\widehat{U}} \mathcal{T}_{\widehat{U}}$. When d is large (relative to θ_{4n}), this expected return time will be at least exponentially large in $\frac{d}{(\log d)^{1+o_n(1)}}$. This exponent arises from the consideration that, for a configuration with a diameter below θ_{4n} to increase its diameter to d , it must avoid collapse over the timescale for which it is typical (i.e., $(\log d)^{1+o_n(1)}$) approximately $\frac{d}{(\log d)^{1+o_n(1)}}$ times consecutively. Because the number of n -element configurations with a diameter of approximately d is negligible relative to their expected return times, the collective weight under π_n of such configurations will be exponentially small in $\frac{d}{(\log d)^{1+o_n(1)}}$.

We note that, while there are abstract results which relate hitting times to the stationary distribution (e.g., [GLPP17, Lemma 4]), we cannot directly apply results that require bounds on hitting times which hold uniformly for any initial configuration. This is because hitting times from \widehat{V} depend on its diameter. We could apply such results after partitioning $\widehat{\text{NonIso}}(n)$ by diameter, but we would then save little effort from their use.

Proof of Proposition 7.2. Let d be at least $2\theta_{4n}$ and take $\varepsilon = 2n^{-4}$. We claim that, for any configuration \widehat{U} with a diameter in $[2^j d, 2^{j+1} d)$ for an integer $j \geq 0$, the expected return time to \widehat{U} satisfies

$$\mathbf{E}_{\widehat{U}} \mathcal{T}_{\widehat{U}} \geq \exp\left(\frac{2^j d}{(\log(2^j d))^{1+2\varepsilon}}\right). \quad (7.28)$$

We can use (7.28) to prove (7.2) in the following way. We write $\{\text{diam}(\widehat{U}) \geq d\}$ as a disjoint union of events of the form $H_j = \{2^j \leq \text{diam}(\widehat{U}) < 2^{j+1} d\}$ for $j \geq 0$. Because a disk with a diameter of at most $2^{j+1} d$

contains fewer than $\lfloor 4^{j+1}d^2 \rfloor$ elements of \mathbb{Z}^2 , the number of non-isolated, n -element configurations with a diameter of at most $2^{j+1}d$ satisfies

$$|\{\widehat{U} \text{ in } \widehat{\text{NonIso}}(n) \text{ with } 2^j d \leq \text{diam}(\widehat{U}) < 2^{j+1}d\}| \leq \binom{\lfloor 4^{j+1}d^2 \rfloor}{n} \leq (4^{j+1}d^2)^n. \quad (7.29)$$

We use (7.1) with (7.28) and (7.29) to estimate

$$\pi_n(\text{diam}(\widehat{U}) \geq d) = \sum_{j=0}^{\infty} \pi_n(H_j) = \sum_{j=0}^{\infty} \sum_{\widehat{U} \in H_j} \pi_n(\widehat{U}) \leq \sum_{j=0}^{\infty} (4^{j+1}d^2)^n e^{-\frac{2^j d}{(\log(2^j d))^{1+2\varepsilon}}}. \quad (7.30)$$

Using the fact that $d \geq 2\theta_{4n}$, it is easy to check that the ratio of the $(j+1)$ st summand to the j th summand in (7.30) is at most e^{-j-1} , for all $j \geq 0$. Accordingly, we have

$$\pi_n(\text{diam}(\widehat{U}) \geq d) \leq (4d^2)^n e^{-\frac{d}{(\log d)^{1+2\varepsilon}}} \sum_{j=0}^{\infty} e^{-j} \leq e^{-\frac{d}{(\log d)^{1+3\varepsilon}}},$$

where the second inequality is justified by the fact that $d \geq 2\theta_{4n}$. This proves (7.2) when the claimed bound (7.28) holds.

We will prove (7.28) by making a comparison with a geometric random variable on $\{0, 1, \dots\}$ with a ‘‘success’’ probability of $e^{-\frac{d}{(\log d)^{1+\varepsilon}}}$ (or with $2^j d$ in place of d). This geometric random variable will model the number of visits to configurations with diameters below θ_{4n} before reaching a diameter of d , and the success probability arises from the fact that, for a configuration to increase its diameter to d from θ_{4n} , it must avoid collapse over $d - \theta_{4n}$ steps. By Corollary 6.4, this happens with a probability which is exponentially small in $\frac{d}{(\log d)^{1+\varepsilon}}$.

Let \widehat{U} be a non-isolated, n -element configuration with a diameter in $[2^j d, 2^{j+1}d)$ and let $\widehat{\mathcal{W}}$ be the set of configurations in $\text{NonIso}(n)$ with a diameter of at most θ_{4n} . Define N to be the number of returns to $\widehat{\mathcal{W}}$ before time $\mathcal{T}_{\widehat{U}}$, i.e., $N = \sum_{t=1}^{\mathcal{T}_{\widehat{U}}} \mathbf{1}(\widehat{U}_t \in \widehat{\mathcal{W}})$, and let \widehat{V} minimize $\mathbf{E}_{\widehat{V}} N$ among $\widehat{\mathcal{W}}$. We claim that

$$\mathbf{E}_{\widehat{U}} \mathcal{T}_{\widehat{U}} \geq (\log(2^{j+1}d))^{-2n} \mathbf{E}_{\widehat{V}} N. \quad (7.31)$$

Indeed, the factor in front of $\mathbf{E}_{\widehat{V}} N$ comes from (7.13), which implies that

$$\mathbf{P}_{\widehat{U}}(\mathcal{T}_{\widehat{L}_n} < \mathcal{T}_{\widehat{U}}) \geq (\log(2^{j+1}d))^{-2n}.$$

This bound justifies the first of the following inequalities

$$\mathbf{E}_{\widehat{U}} \mathcal{T}_{\widehat{U}} \geq (\log(2^{j+1}d))^{-2n} \mathbf{E}_{\widehat{L}_n} \mathcal{T}_{\widehat{U}} \geq (\log(2^{j+1}d))^{-2n} \mathbf{E}_{\widehat{L}_n} N \geq (\log(2^{j+1}d))^{-2n} \mathbf{E}_{\widehat{V}} N.$$

The second inequality holds because the time it takes to reach \widehat{U} is at least the number of returns to $\widehat{\mathcal{W}}$ before $\mathcal{T}_{\widehat{U}}$; the third is due to our choice of \widehat{V} . Hence, (7.31) holds.

The virtue of the lower bound (7.31) is that we can bound below $\mathbf{E}_{\widehat{V}} N$ by

$$\mathbf{E}_{\widehat{V}} N = \mathbf{P}_{\widehat{V}}(\mathcal{T}_{\widehat{\mathcal{W}}} < \mathcal{T}_{\widehat{U}}) \left(1 + \mathbf{E}_{\widehat{V}}[N - 1 \mid \mathcal{T}_{\widehat{\mathcal{W}}} < \mathcal{T}_{\widehat{U}}]\right) \geq \mathbf{P}_{\widehat{V}}(\mathcal{T}_{\widehat{\mathcal{W}}} < \mathcal{T}_{\widehat{U}}) (1 + \mathbf{E}_{\widehat{V}} N).$$

The inequality holds by the Markov property applied to $\mathcal{T}_{\widehat{\mathcal{W}}}$ and because \widehat{V} minimizes the expected number of returns to $\widehat{\mathcal{W}}$. This implies that $\mathbf{E}_{\widehat{V}} N$ is at least the expected value of a geometric random variable on $\{0, 1, \dots\}$ with success parameter p of $\mathbf{P}_{\widehat{V}}(\mathcal{T}_{\widehat{U}} < \mathcal{T}_{\widehat{\mathcal{W}}})$:

$$\mathbf{E}_{\widehat{V}} N \geq (1 - p)p^{-1}. \quad (7.32)$$

It remains to bound above p .

Because diameter increases at most linearly in time, $\mathcal{T}_{\widehat{U}}$ is at least $2^j d - \theta_{4n}$ under $\mathbf{P}_{\widehat{V}}$. Consequently,

$$\mathbf{P}_{\widehat{V}}(\mathcal{T}_{\widehat{U}} < \mathcal{T}_{\widehat{W}}) \leq \mathbf{P}_{\widehat{V}}(\mathcal{T}(\theta_{4n}) > 2^j d - \theta_{4n}). \quad (7.33)$$

We apply Corollary 6.4 with t equal to $\frac{2^j d - \theta_{4n}}{(\log(2^j d))^{1+\varepsilon}}$, finding

$$\mathbf{P}_{\widehat{V}}(\mathcal{T}(\theta_{4n}) > 2^j d - \theta_{4n}) \leq \exp\left(-\frac{2^j d - \theta_{4n}}{(\log(2^j d))^{1+\varepsilon}}\right).$$

By (7.33), this is also an upper bound on $p < \frac{1}{2}$ and so, by (7.32), $\mathbf{E}_{\widehat{V}} N$ is at least $(2p)^{-1}$. Substituting these bounds into (7.31) and simplifying with the fact that $d \geq 2\theta_{4n}$, we find that the expected return time to \widehat{U} satisfies (7.28):

$$\mathbf{E}_{\widehat{U}} \mathcal{T}_{\widehat{U}} \geq \frac{1}{2} (\log(2^{j+1} d))^{-2n} \exp\left(\frac{2^j d - \theta_{4n}}{(\log(2^j d))^{1+\varepsilon}}\right) \geq \exp\left(\frac{2^j d}{(\log(2^j d))^{1+2\varepsilon}}\right).$$

□

8. MOTION OF THE CENTER OF MASS

As a consequence of the results of Section 7 and standard renewal theory, the center of mass process $(\mathcal{M}_t)_{t \geq 0}$, after linear interpolation and rescaling $(t^{-1/2} \mathcal{M}_{st})_{s \in [0,1]}$, and when viewed as a measure on $\mathcal{E}([0,1])$, converges weakly to two-dimensional Brownian motion as $t \rightarrow \infty$. This is the content of Theorem 3.

We will use the following lemma to bound the coordinate variances of the Brownian motion limit. To state it, we denote by $\tau_i = \inf\{t > \tau_{i-1} : \widehat{U}_t = \widehat{L}_n\}$ the i^{th} return time to \widehat{L}_n .

Lemma 8.1. *Let c be the constant from Proposition 7.4 and abbreviate $\theta_{5n}(cn)$ by θ . If, for some $i \geq 0$, X is one of the random variables*

$$\tau_{i+1} - \tau_i, \quad |\mathcal{M}_{\tau_{i+1}} - \mathcal{M}_{\tau_i}|, \quad \text{or} \quad |\mathcal{M}_t - \mathcal{M}_{\tau_i}| \mathbf{1}(\tau_i \leq t \leq \tau_{i+1}),$$

then the distribution of X satisfies the following tail bound

$$\mathbf{P}_{\widehat{L}_n}(X > K^4) \leq e^{-K}, \quad K \geq \theta. \quad (8.1)$$

Consequently,

$$\mathbf{E}_{\widehat{L}_n} X \leq 2\theta^4 \quad \text{and} \quad \text{Var}_{\widehat{L}_n} X \leq 2\theta^8. \quad (8.2)$$

Proof. Because the diameter of \widehat{L}_n is at most n , for any $K \geq \theta$, Proposition 7.4 implies

$$\mathbf{P}_{\widehat{L}_n}(\tau_1 > K^4) \leq e^{-K}.$$

Applying the strong Markov property to τ_i , we find (8.1) for $X = \tau_{i+1} - \tau_i$. Using (8.1) with the tail sum formulas for the first and second moments gives (8.2) for this X . The other cases of X then follow from

$$|\mathcal{M}_{\tau_{i+1}} - \mathcal{M}_{\tau_i}| \leq \tau_{i+1} - \tau_i.$$

□

Proof of Theorem 3. Standard arguments (e.g., Section 8 of [Bil99]) combined with the renewal theorem show that $(t^{-1/2} \mathcal{M}_{st})_{t \geq 1}$ is a tight sequence of functions. We claim that the finite-dimensional distributions of the rescaled process converge as $t \rightarrow \infty$ to those of two-dimensional Brownian motion.

For any $m \geq 1$ and times $0 = s_0 \leq s_1 < s_2 < \dots < s_m \leq 1$, form the random vector

$$t^{-1/2} (\mathcal{M}_{s_1 t}, \mathcal{M}_{s_2 t} - \mathcal{M}_{s_1 t}, \dots, \mathcal{M}_{s_m t} - \mathcal{M}_{s_{m-1} t}). \quad (8.3)$$

For s in $[0, 1]$, we denote by $I(s)$ the number of returns to \widehat{L}_n by time st . Lemma 8.1 and Markov's inequality imply that $|\mathcal{M}_{s_1 t} - \mathcal{M}_{\tau_{I(s_1)}}| \rightarrow 0$ in probability as $t \rightarrow \infty$, hence, by Slutsky's theorem, the distributions of (8.3) and

$$t^{-1/2} \left(\mathcal{M}_{\tau_{I(s_1)}}, \mathcal{M}_{\tau_{I(s_2)}} - \mathcal{M}_{\tau_{I(s_1)+1}}, \dots, \mathcal{M}_{\tau_{I(s_m)}} - \mathcal{M}_{\tau_{I(s_{m-1})+1}} \right) \quad (8.4)$$

have the same $t \rightarrow \infty$ limit. By the renewal theorem, $I(s_1) < I(s_2) < \dots < I(s_m)$ for all sufficiently large t , so the strong Markov property implies the independence of the entries in (8.4) for all such t .

A generic entry in (8.4) is a sum of independent increments of the form $\mathcal{M}_{\tau_{i+1}} - \mathcal{M}_{\tau_i}$. As noted in Section 1, the transition probabilities are unchanged when configurations are multiplied by elements of the symmetry group \mathcal{G} of \mathbb{Z}^2 . This implies

$$\mathbf{E}_{\widehat{L}_n} [\mathcal{M}_{\tau_{i+1}} - \mathcal{M}_{\tau_i}] = o \quad \text{and} \quad \Sigma = \nu^2 \mathbf{I},$$

where Σ is the variance-covariance matrix of $\mathcal{M}_{\tau_{i+1}} - \mathcal{M}_{\tau_i}$ and ν is a constant which, by Lemma 8.1, is finite. The renewal theorem implies that the scaled variance $t^{-1} \nu^2 (I(s_i) - I(s_{i-1}))$ of the i^{th} entry converges almost surely to $(s_i - s_{i-1}) \chi^2$ where $\chi^2 = \nu^2 / \mathbf{E}_{\widehat{L}_n} [\tau_1]$, hence, by Slutsky's theorem, we can replace the scaled variance of each entry in (8.4) with this limit, without affecting the limiting distribution of the vector.

By the central limit theorem,

$$\frac{1}{\chi \sqrt{t}} \left(\mathcal{M}_{\tau_{I(s_i)}} - \mathcal{M}_{\tau_{I(s_{i-1})+1}} \right) \xrightarrow{d} \mathcal{N}(o, (s_i - s_{i-1}) \mathbf{I}),$$

which, by the independence of the entries in (8.4) for all sufficiently large t , implies

$$\begin{aligned} \frac{1}{\chi \sqrt{t}} \left(\mathcal{M}_{s_1 t}, \mathcal{M}_{s_2 t} - \mathcal{M}_{s_1 t}, \dots, \mathcal{M}_{s_m t} - \mathcal{M}_{s_{m-1} t} \right) \\ \xrightarrow{d} (\mathbf{B}(s_1), \mathbf{B}(s_2 - s_1), \dots, \mathbf{B}(s_m - s_{m-1})), \end{aligned} \quad (8.5)$$

as $t \rightarrow \infty$. Because m and the $\{s_i\}_{i=1}^m$ were arbitrary, the continuous mapping theorem and (8.5) imply the convergence of the finite-dimensional distributions of $\left(\frac{1}{\chi \sqrt{t}} \mathcal{M}_{st}, 0 \leq s \leq 1 \right)$ to those of $(\mathbf{B}(s), 0 \leq s \leq 1)$. This proves the weak convergence component of Theorem 3.

It remains to bound χ^2 , which we do by estimating $\mathbf{E}_{\widehat{L}_n} [\tau_1]$ and ν^2 . The former is bounded above by $2\theta^4$, due to Lemma 8.1, and below by 1. Here, $\theta = \theta_{5n}(c_1 n)$ and c_1 is the constant from Proposition 7.4. To bound below ν^2 , denote the e_2 component of $\mathcal{M}_{\tau_{i+1}} - \mathcal{M}_{\tau_i}$ by X and observe that $\mathbf{P}_{\widehat{L}_n}(X = n^{-1})$ is at least the probability that, from L_n , the element at o is activated and subsequently deposited at $(0, n)$ (recall that L_n is the segment from o to $(0, n-1)$), resulting in $\tau_1 = 1$ and $\mathcal{M}_{\tau_1} = \mathcal{M}_0 + n^{-1} e_2$. This probability is at least $e^{-c_2 n}$ for a constant c_2 . Markov's inequality applied to X^2 then gives

$$\text{Var}_{\widehat{L}_n} X \geq \mathbf{P}_{\widehat{L}_n}(X^2 \geq n^{-2}) \geq n^{-2} e^{-c_2 n} \geq e^{-c_3 n}.$$

By Lemma 8.1, ν^2 is at most $2\theta^8$. In summary,

$$1 \leq \mathbf{E}_{\widehat{L}_n} [\tau_1] \leq 2\theta^4 \quad \text{and} \quad e^{-c_3 n} \leq \nu^2 \leq 2\theta^8,$$

which implies

$$\theta_{6n}(cn)^{-1} \leq e^{-c_3 n} (2\theta^4)^{-1} \leq \chi^2 \leq 2\theta^8 \leq \theta_{6n}(cn),$$

with $c = \max\{c_1, c_3\}$. □

APPENDIX A. AUXILIARY LEMMAS

A.1. Potential kernel bounds. The following lemma collects several facts about the potential kernel which are used in Section 3. As each fact is a simple consequence of (3.9), we omit its proof.

Lemma A.1. *In what follows, x, y, z, z' are elements of \mathbb{Z}^2 .*

(1) *For $\mathbf{a}(y)$ to be at least $\mathbf{a}(x)$, it suffices to have*

$$|y| \geq |x|(1 + \pi\lambda|x|^{-2} + (\pi\lambda)^2|x|^{-4}).$$

In particular, if $|x| \geq 2$, then $|y| \geq 1.06|x|$ suffices.

(2) *When $|x| \geq 1$, $\mathbf{a}(x)$ is at least $\frac{2}{\pi} \log |x|$. When $|x| \geq 2$, $\mathbf{a}(x)$ is at most $4 \log |x|$.*

(3) *If $z, z' \in C(r)$ and $y \in D(R)^c$ for $r \leq \frac{1}{100}R$ and $R \geq 100$, then*

$$|\mathbf{a}(y - z) - \mathbf{a}(y - z')| \leq \frac{4}{\pi}.$$

(4) *If x and y satisfy $|x|, |y| \geq 1$ and $K^{-1} \leq \frac{|y|}{|x|} \leq K$ for some $K \geq 2$, then*

$$\mathbf{a}(y) - \mathbf{a}(x) \leq \log K.$$

(5) *If $|x| \geq 8|y|$ and $|y| \geq 10$, then*

$$|\mathbf{a}(x + y) - \mathbf{a}(x)| \leq 0.7 \frac{|y|}{|x|}.$$

(6) *Let $R \geq 10r$ and $r \geq 10$. Then, uniformly for $x \in C(R)$ and $y \in C(r)$, we have*

$$0.56 \log(R/r) \leq \mathbf{a}(x) - \mathbf{a}(y) \leq \log(R/r).$$

In the next section, we will need the following comparison of \mathbf{a} and \mathbf{a}' .

Lemma A.2. *Let μ be any probability measure on $C_x(r)$. Suppose $r \geq 2(|x| + 1)$. Then*

$$\left| \sum_{y \in C_x(r)} \mu(y) \mathbf{a}(y) - \mathbf{a}'(r) \right| \leq \left(\frac{5}{2\pi} + 2\lambda \right) \left(\frac{|x| + 1}{r} \right).$$

In particular, if $x = o$, then $|\mathbf{a}(y) - \mathbf{a}'(r)| \leq r^{-1}$ for every $y \in C(r)$.

Proof. We recall that, for any $x \in \mathbb{Z}^2$, the potential kernel has the form specified in (3.9) where the error term conceals a constant of λ , which is no more than 0.07 [KS04]. That is,

$$\left| \mathbf{a}(x) - \frac{2}{\pi} \log |x| - \kappa \right| \leq \lambda |x|^{-2}.$$

For $y \in C_x(r)$, we have $r - |x| - 1 \leq |y| \leq r + |x| + 1$. Accordingly,

$$\begin{aligned} \mathbf{a}(y) &\leq \frac{2}{\pi} \log |r + |x| + 1| + \kappa + O(|r - |x| - 1|^{-2}) \\ &= \frac{2}{\pi} \log r + \kappa + \frac{2}{\pi} \log \left(1 + \frac{|x| + 1}{r} \right) + O(|r - |x| - 1|^{-2}). \end{aligned}$$

Using the assumption $(|x| + 1)/r \in (0, 1/2)$ with Taylor's remainder theorem gives

$$\mathbf{a}(y) \leq \mathbf{a}'(r) + \frac{2}{\pi} \left(\frac{|x| + 1}{r} + \frac{1}{2} \left(\frac{|x| + 1}{r} \right)^2 \right) + O(|r - |x| - 1|^{-2}).$$

Simplifying with $r \geq 2(|x| + 1)$ and $r \geq 2$ leads to

$$\mathbf{a}(y) \leq \mathbf{a}'(r) + \frac{2}{\pi} \left(\frac{5}{4} + \pi\lambda \right) \left(\frac{|x| + 1}{r} \right) = \mathbf{a}'(r) + \left(\frac{5}{2\pi} + 2\lambda \right) \left(\frac{|x| + 1}{r} \right).$$

The lower bound is similar. Because this holds for any $y \in C_x(r)$, for any probability measure μ on $C_x(r)$, we have

$$\left| \sum_{y \in C_x(r)} \mu(y) \mathbf{a}(y) - \mathbf{a}'(r) \right| \leq \left(\frac{5}{2\pi} + 2\lambda \right) \left(\frac{|x| + 1}{r} \right).$$

□

A.2. Comparison between harmonic measure and hitting probabilities. To prove Lemma 3.2, we require a comparison (Lemma A.3) between certain values of harmonic measure and hitting probabilities. In fact, we need additional quantification of an error term which appears in standard versions of this result (e.g. [Law13, Theorem 2.1.3]). Effectively, this additional quantification comes from a bound on λ , the implicit constant in (3.9). The proof is similar to that of Theorem 3.17 in [Pop21].

Lemma A.3. *Let $x \in D(R)^c$ for $R \geq 100r$ and $r \geq 10$. Then*

$$0.93\mathbb{H}_{C(r)}(y) \leq \mathbb{H}_{C(r)}(x, y) \leq 1.04\mathbb{H}_{C(r)}(y). \quad (\text{A.1})$$

Proof. We have

$$\mathbb{H}_{C(r)}(x, y) - \mathbb{H}_{C(r)}(y) = -\mathbf{a}(y - x) + \sum_{z \in C(r)} \mathbb{P}_y \left(S_{\tau_{C(r)}} = z \right) \mathbf{a}(z - x). \quad (\text{A.2})$$

Since $C(10r)$ separates x from $C(r)$, the optional stopping theorem applied to $\sigma_{C(10r)} \wedge \tau_{C(r)}$ and the martingale $(S_{t \wedge \tau_x} - x)$ gives

$$\begin{aligned} \mathbf{a}(y - x) &= \sum_{z \in C(r)} \mathbb{P}_y \left(S_{\tau_{C(r)}} = z \right) \mathbf{a}(z - x) \\ &\quad + \mathbb{E}_y \left[\mathbf{a} \left(S_{\sigma_{C(10r)}} - x \right) - \mathbf{a} \left(S_{\tau_{C(r)}} - x \right) \mid \sigma_{C(10r)} < \tau_{C(r)} \right] \mathbb{P}_y \left(\sigma_{C(10r)} < \tau_{C(r)} \right). \end{aligned} \quad (\text{A.3})$$

In the second term of (A.3), we analyze the difference in potentials by observing

$$S_{\sigma_{C(10r)}} - x - \left(S_{\tau_{C(r)}} - x \right) = S_{\sigma_{C(10r)}} - S_{\tau_{C(r)}}.$$

Accordingly, letting $u = S_{\tau_{C(r)}} - x$ and $v = S_{\sigma_{C(10r)}} - S_{\tau_{C(r)}}$,

$$\mathbf{a} \left(S_{\sigma_{C(10r)}} - x \right) - \mathbf{a} \left(S_{\tau_{C(r)}} - x \right) = \mathbf{a}(u + v) - \mathbf{a}(u).$$

We observe that $|v| \leq 11r + 2$ and $|u| \geq 99r - 2$, so $|u| \geq 8|v|$. Since we also have $|v| \geq 9r - 2 \geq 10$, (5) of Lemma A.1 applies to give

$$\mathbf{a}(u + v) - \mathbf{a}(u) \leq 0.7 \frac{|v|}{|u|} \leq \frac{2}{25}.$$

We analyze the other factor of (A.3) as

$$\begin{aligned} \mathbb{P}_y \left(\sigma_{C(10r)} < \tau_{C(r)} \right) &= \frac{1}{4} \sum_{z \notin C(r): z \sim y} \mathbb{P}_z \left(\sigma_{C(10r)} < \tau_{C(r)} \right) \\ &= \frac{1}{4} \sum_{z \notin C(r): z \sim y} \frac{\mathbf{a}(z - z_0) - \mathbb{E}_z \mathbf{a} \left(S_{\tau_{C(r)}} - z_0 \right)}{\mathbb{E}_z \left[\mathbf{a} \left(S_{\sigma_{C(10r)}} \right) - \mathbf{a} \left(S_{\tau_{C(r)}} \right) \mid \sigma_{C(10r)} < \tau_{C(r)} \right]}, \end{aligned}$$

where $z_0 \in A$. To obtain an upper bound on the potential difference in the denominator, we apply (6) of Lemma A.1, which gives

$$\mathbb{P}_y (\sigma_{C(10r)} < \tau_{C(r)}) \leq \frac{1}{0.6 \log 10} \mathbb{H}_{C(r)}(y).$$

Combining this with the other estimate for the second term of (A.3), we find

$$\mathbf{a}(y-x) \leq \sum_{z \in C(r)} \mathbb{P}_y (S_{\tau_{C(r)}} = z) \mathbf{a}(z-x) + \underbrace{\frac{2}{25} \cdot \frac{1}{0.56 \log 10}}_{\leq 0.063} \mathbb{H}_{C(r)}(y).$$

Substituting this into (A.2), we have

$$\mathbb{H}_{C(r)}(x, y) - \mathbb{H}_{C(r)}(y) \geq -0.063 \mathbb{H}_{C(r)}(y) \implies \mathbb{H}_{C(r)}(x, y) \geq 0.93 \mathbb{H}_{C(r)}(y).$$

We again apply (5) and (6) of Lemma A.1 to bound the factors in the second term of A.3 as

$$\mathbf{a}(u+v) - \mathbf{a}(u) \geq -0.0875 \quad \text{and} \quad \mathbb{P}_y (\sigma_{C(10r)} < \tau_{C(r)}) \geq \frac{1}{\log 10} \mathbb{H}_{C(r)}(y).$$

Substituting these into (A.3), we find

$$\mathbf{a}(y-x) \geq \sum_{z \in C(r)} \mathbb{P}_y (S_{\tau_{C(r)}} = z) \mathbf{a}(z-x) - 0.0875 \cdot \frac{1}{\log 10} \mathbb{H}_{C(r)}(y).$$

Consequently, (A.2) becomes

$$\mathbb{H}_{C(r)}(x, y) - \mathbb{H}_{C(r)}(y) \leq \frac{0.0875}{\log 10} \mathbb{H}_{C(r)}(y) \leq \frac{1}{25} \mathbb{H}_{C(r)}(y).$$

Rearranging, we find

$$\mathbb{H}_{C(r)}(x, y) \leq 1.04 \mathbb{H}_{C(r)}(y).$$

□

A.3. Uniform lower bound on a conditional entrance measure. Recall that Lemma 3.1 bounds below the conditional hitting distribution of $C(\varepsilon^2 R)$ from $C(\varepsilon R)$, given that $C(\varepsilon^2 R)$ is hit before $C(R)$, in terms of the uniform distribution on $C(\varepsilon^2 R)$. The idea of the proof is to use Lemma A.3 to approximate the hitting distribution of $C(\varepsilon^2 R)$ with the corresponding harmonic measure, which is comparable to the uniform distribution on $C(\varepsilon^2 R)$. The proof is similar to that of Lemma 2.1 in [DPRZ06].

Proof of Lemma 3.1. Fix ε and R which satisfy the hypotheses. Let $x \in C(\varepsilon R)$ and $y \in C(\varepsilon^2 R)$. We have

$$\mathbb{P}_x (S_{\tau_{C(\varepsilon^2 R)}} = y, \tau_{C(\varepsilon^2 R)} < \tau_{C(R)}) = \mathbb{H}_{C(\varepsilon^2 R)}(x, y) - \mathbb{P}_x (S_{\tau_{C(\varepsilon^2 R)}} = y, \tau_{C(\varepsilon^2 R)} > \tau_{C(R)}). \quad (\text{A.4})$$

By the strong Markov property applied to $\tau_{C(R)}$,

$$\mathbb{P}_x (S_{\tau_{C(\varepsilon^2 R)}} = y, \tau_{C(\varepsilon^2 R)} > \tau_{C(R)}) = \mathbb{E}_x [\mathbb{H}_{C(\varepsilon^2 R)}(S_{\tau_{C(R)}}, y); \tau_{C(\varepsilon^2 R)} > \tau_{C(R)}]. \quad (\text{A.5})$$

We use Lemma A.3 to uniformly bound the terms of the form $\mathbb{H}_{C(\varepsilon^2 R)}(\cdot, y)$ appearing in (A.4) and (A.5). For any $w \in C(R)$, the hypotheses of Lemma A.3 are satisfied with $\varepsilon^2 R$ in the place of r , and R as presently defined, because then $r \geq 10$ and $R \geq 100r$. Therefore, by (A.1), uniformly for $w \in C(R)$,

$$\mathbb{H}_{C(\varepsilon^2 R)}(w, y) \leq 1.04 \mathbb{H}_{C(\varepsilon^2 R)}(y). \quad (\text{A.6})$$

Now, for any $x \in C(\varepsilon R)$, the hypotheses of Lemma A.3 are again satisfied with the same r and with εR in the place of R , since $\varepsilon R \geq 100r = 100\varepsilon^2 R$ by the assumption $\varepsilon \leq \frac{1}{100}$. We apply (A.1) to find

$$\mathbb{H}_{C(\varepsilon^2 R)}(x, y) \geq 0.93 \mathbb{H}_{C(\varepsilon^2 R)}(y). \quad (\text{A.7})$$

Substituting (A.6) into (A.5), we find

$$\mathbb{P}_x \left(S_{\tau_{C(\varepsilon^2 R)}} = y, \tau_{C(\varepsilon^2 R)} > \tau_{C(R)} \right) \leq 1.04 \mathbb{H}_{C(\varepsilon^2 R)}(y) \mathbb{P}_x \left(\tau_{C(\varepsilon^2 R)} > \tau_{C(R)} \right).$$

Similarly, substituting (A.7) into (A.4) and using the previous display, we find

$$\begin{aligned} \mathbb{P}_x \left(S_{\tau_{C(\varepsilon^2 R)}} = y, \tau_{C(\varepsilon^2 R)} < \tau_{C(R)} \right) &\geq 0.93 \mathbb{H}_{C(\varepsilon^2 R)}(y) \mathbb{P}_x \left(\tau_{C(\varepsilon^2 R)} < \tau_{C(R)} \right) \\ &\quad - (1.04 - 0.93) \mathbb{H}_{C(\varepsilon^2 R)}(y) \mathbb{P}_x \left(\tau_{C(\varepsilon^2 R)} > \tau_{C(R)} \right). \end{aligned}$$

Applying hypothesis (3.3), we find that the right-hand side is at least

$$c_1 \mathbb{H}_{C(\varepsilon^2 R)}(y) \mathbb{P}_x \left(\tau_{C(\varepsilon^2 R)} < \tau_{C(R)} \right),$$

for a positive constant c_1 . The result then follows the existence of a positive constant c_2 such that $\mathbb{H}_{C(\varepsilon^2 R)}(y) \geq c_2 \mu_{\varepsilon^2 R}(y)$ for any $y \in C(\varepsilon^2 R)$. \square

The proof of Lemma 3.2 is a simple application of Lemma 3.1. A short calculation is needed to verify that the hypotheses of Lemma 3.1 are met.

Proof of Lemma 3.2. Under the conditioning, the random walk must reach $C(\delta R_{J+1})$ before $C(R_J)$. It therefore suffices to prove that there exists a positive constant c_1 such that, uniformly for all $x \in C(\delta R_{J+1})$ and $z \in C(R_J)$,

$$\mathbb{P}_x(S_\eta = z \mid \tau_{C(R_J)} < \tau_A) \geq c_1 \mu_{R_J}(z), \quad (\text{A.8})$$

where $\eta = \tau_{C(R_J)} \wedge \tau_A$. Because $\partial \mathcal{A}_J$ separates x from A , the conditional probability in (A.8) is at least

$$\mathbb{P}_x(S_\eta = z \mid \tau_{C(R_J)} < \tau_{C(R_{J+1})}, \tau_{C(R_J)} < \tau_A) \mathbb{P}_x(\tau_{C(R_J)} < \tau_{C(R_{J+1})}). \quad (\text{A.9})$$

The first factor of (A.9) simplifies to

$$\mathbb{P}_x(S_{\tau_{C(R_J)}} = z \mid \tau_{C(R_J)} < \tau_{C(R_{J+1})}), \quad (\text{A.10})$$

which we will bound below using Lemma 3.1.

We will verify the hypotheses of Lemma 3.1 with $\varepsilon = \delta$ and $R = R_{J+1}$. The first hypothesis is $R \geq 10\varepsilon^{-2}$, which is satisfied because $R_{J+1} \geq R_1 = 10\delta^{-2}$. The second hypothesis is (3.3) which, in our case, can be written as

$$\max_{x \in C(\delta R_{J+1})} \mathbb{P}_x(\tau_{C(R_{J+1})} < \tau_{C(R_J)}) < \frac{9}{10}. \quad (\text{A.11})$$

Exercise 1.6.8 of [Law13] states that

$$\mathbb{P}_x(\tau_{C(R_{J+1})} < \tau_{C(R_J)}) = \frac{\log\left(\frac{|x|}{R_J}\right) + O(R_J^{-1})}{\log\left(\frac{R_{J+1}}{R_J}\right) + O(R_J^{-1} + R_{J+1}^{-1})}, \quad (\text{A.12})$$

where the implicit constants are at most 2 (i.e., the $O(R_J^{-1})$ term is at most $2R_J^{-1}$). For the moment, ignore the error terms and assume $|x| = \delta R_{J+1}$, in which case (A.12) evaluates to $\frac{5 \log 10 - \log 25}{5 \log 10} < 0.73$. Because $R_J \geq 10^5$, even after allowing $|x|$ up to $\delta R_{J+1} + 1$ and accounting for the error terms, (A.12) is less than $\frac{9}{10}$, which implies (A.11).

Applying Lemma 3.1 to (A.10), we obtain a constant c_2 such that

$$\mathbb{P}_x(S_{\tau_{C(R_J)}} = z \mid \tau_{C(R_J)} < \tau_{C(R_{J+1})}) \geq c_2 \mu_{R_J}(z). \quad (\text{A.13})$$

By (A.11), the second factor of (A.9) is bounded below by $\frac{1}{10}$. We conclude the claim of (A.8) by combining this bound and (A.13) with (A.9), and by setting $c_1 = \frac{1}{10} c_2$. \square

A.4. Estimate for the exit distribution of a rectangle. Informally, Lemma A.4 says that the probability a walk from one end of a rectangle (which may not be aligned with the coordinate axes) exits through the opposite end is bounded below by a quantity depending upon the aspect ratio of the rectangle. We believe this estimate is known but, as we are unable to find a reference for it, we provide one here. In brief, the proof uses an adaptive algorithm for constructing a sequence of squares which remain inside the rectangle and the sides of which are aligned with the axes. We then bound below the probability that the walk follows the path determined by the squares until exiting the opposite end of the rectangle.

Recall that $\text{Rec}(\phi, w, \ell)$ denotes the rectangle of width w , centered along the line segment from $-e^{i\phi}w$ to $e^{i\phi}\ell$, intersected with \mathbb{Z}^2 (see Figure 14).




FIGURE 14. On the left, we depict the rectangles $\text{Rec} = \text{Rec}(\phi, w, \ell)$ (shaded blue) and $\text{Rec}^+ = \text{Rec}(\phi, w, \ell + w)$ (union of blue- and red-shaded regions) for $\phi = \pi/4$, $w = 4\sqrt{2}$, and $\ell = 11\sqrt{2}$. \mathcal{I} denotes $\text{Rec} \cap \partial(\text{Rec}^+ \setminus \text{Rec})$.

Lemma A.4. For any $24 \leq w \leq \ell$ and any ϕ , let $\text{Rec} = \text{Rec}(\phi, w, \ell)$ and $\text{Rec}^+ = \text{Rec}(\phi, w, \ell + w)$. Then,

$$\mathbb{P}_o(\tau_{\partial\text{Rec}} < \tau_{\partial\text{Rec}^+}) \geq c^{\ell/w},$$

for a universal positive constant $c < 1$.

We use the hypothesis $w \geq 24$ to deal with the effects of discreteness; the constant 24 is otherwise unimportant, and many choices would work in its place.

Proof of Lemma A.4. We will first define a square, centered at the origin and with each corner in \mathbb{Z}^2 , which lies in Rec^+ . We will then translate it to form a sequence of squares through which we will guide the walk to $\text{Rec}^+ \setminus \text{Rec}$ without leaving Rec^+ (see Figure 14). We split the proof into three steps: (1) constructing the squares; (2) proving that they lie in Rec^+ ; and (3) establishing a lower bound on the probability that the walk hits ∂Rec before hitting the interior boundary of Rec^+ .

Step 1: Construction of the squares. Without loss of generality, assume $0 \leq \phi < \pi/2$. For $x \in \mathbb{Z}^2$, we will denote its first coordinate by x^1 and its second coordinate by x^2 . We will use this convention only for this proof. Let \mathcal{I} be equal to $\lfloor \frac{w}{8} \rfloor$ if it is even and equal to $\lfloor \frac{w}{8} \rfloor - 1$ otherwise. With this choice, we define

$$Q = \{x \in \mathbb{Z}^2 : \max\{|x^1|, |x^2|\} \leq \frac{1}{2}\mathcal{I}\}.$$

Since \mathcal{I} is even, the translates of Q by integer multiples of $\frac{1}{2}\mathcal{I}$ are also subsets of \mathbb{Z}^2 .




FIGURE 15. Two steps in the construction of squares. Respectively on the left and right, $y_{i+1} \in M_i$ and $y_{i+2} \in M_{i+1}$ (indicated by the \times symbols) lie above L_ϕ^∞ , so M_{i+1} and M_{i+2} are situated on the eastern sides of Q_{i+1} and Q_{i+2} . However, on the left, as Q_i was translated north to form Q_{i+1} , the relative orientation of M_i and M_{i+1} is perpendicular. In contrast, as Q_{i+1} is translated east to form Q_{i+2} , the right-hand side has parallel M_{i+1} and M_{i+2} .

We construct a sequence of squares Q_i in the following way, where we make reference to the line $L_\phi^\infty = e^{i\phi}\mathbb{R}$. Let $y_1 = o$ and $Q_1 = y_1 + Q$. For $i \geq 1$, let

$$y_{i+1} = \begin{cases} y_i + \frac{1}{2}\mathbf{I}(0, 1) & \text{if } y_i \text{ lies on or below } L_\phi^\infty \\ y_i + \frac{1}{2}\mathbf{I}(1, 0) & \text{if } y_i \text{ lies above } L_\phi^\infty \end{cases} \quad \text{and} \quad Q_{i+1} = y_{i+1} + Q.$$

In words, if the center of the present square lies on or below the line L_ϕ^∞ , then we translate the center north by $\frac{1}{2}\mathbf{I}$ to obtain the next square. Otherwise, we translate the center to the east by $\frac{1}{2}\mathbf{I}$.

We further define, for $i \geq 1$,

$$M_i = \begin{cases} \{x \in Q_i : x^2 - y_i^2 = \frac{1}{2}\mathbf{I} \text{ and } |y_i^1 - x^1| \leq \frac{1}{2}\mathbf{I} - 1\} & \text{if } y_i \text{ lies on or below } L_\phi^\infty \\ \{x \in Q_i : x^1 - y_i^1 = \frac{1}{2}\mathbf{I} \text{ and } |y_i^2 - x^2| \leq \frac{1}{2}\mathbf{I} - 1\} & \text{if } y_i \text{ lies above } L_\phi^\infty. \end{cases}$$

In words, if y_i lies on or below the line L_ϕ^∞ , we choose M_i to be the northernmost edge of Q_i , excluding the corners. Otherwise, we choose it to be the easternmost edge, excluding the corners (Figure 15). We exclude the corners to ensure that $\mathbb{P}_\omega(\tau_{M_{i+1}} \leq \tau_{\partial^{\text{int}} Q_{i+1}})$ is harmonic for all $\omega \in M_i$; we will shortly need this to apply the Harnack inequality. Upcoming Figure 16 provides an illustration of M_i in this context.

We will guide the walk to ∂Rec without leaving Rec^+ by requiring that it exit each square Q_i through M_i for $1 \leq i \leq J$, where we define

$$J = \min\{i \geq 1 : M_i \subseteq \text{Rec}^c\}.$$

That is, J is the first index for which M_i is fully outside Rec . It is clear that J is finite.

Step 2: Proof that $\cup_{i=1}^J Q_i$ is a subset of Rec^+ . Let v be the northeastern endpoint of L_ϕ , where L_ϕ is the segment of L_ϕ^∞ from o to $e^{i\phi}(\ell + w/2)$ and define k to be the first index for which y_k satisfies

$$y_k^1 > v^1 \quad \text{or} \quad y_k^2 > v^2.$$

It will also be convenient to denote by \mathcal{I} the interface between Rec and $\text{Rec}^+ \setminus \text{Rec}$ (the dashed line in Figure 14), given by

$$\mathcal{I} = \text{Rec} \cap \partial(\text{Rec}^+ \setminus \text{Rec}).$$

By construction, we have $|y_k - y_{k-1}| = \frac{1}{2}\mathbf{I}$ and $|y_k - v| \leq \frac{1}{2}\mathbf{I}$. By the triangle inequality, $|y_{k-1} - v| \leq \mathbf{I}$. As $|v| = \ell + w/2$ and because $\text{dist}(o, \mathcal{I}) \leq \ell + 1$, we must have—again by the triangle inequality—that

$\text{dist}(v, \mathcal{I}) \geq w/2 - 1$. From a third use of the triangle inequality and the hypothesized lower bound on w , we conclude

$$\text{dist}(y_{k-1}, \mathcal{I}) \geq \frac{w}{2} - 1 - 1 \geq \frac{w}{2} - 1 - \frac{w}{8} \geq \frac{w}{3} > 2\mathbb{I}. \quad (\text{A.14})$$

To summarize in words, y_{k-1} is not in Rec and it is separated from Rec by a distance strictly greater than $2\mathbb{I}$.

Because the sides of Q_{k-1} have length \mathbb{I} , (A.14) implies $Q_{k-1} \subseteq \text{Rec}^c$. Since M_{k-1} is a subset of Q_{k-1} , we must also have $M_{k-1} \subseteq \text{Rec}^c$, which implies $J \leq k - 1$. As k was the first index for which $y_k^1 > v^1$ or $y_k^2 > v^2$, y_J satisfies $y_J^1 \leq v^1$ and $y_J^2 \leq v^2$. Then, by construction, for all $1 \leq i \leq J$, the centers satisfy

$$y^1 \leq y_i^1 \leq v^1 \quad \text{and} \quad y^2 \leq y_i^2 \leq v^2. \quad (\text{A.15})$$

From (A.15) and the fact that $\text{dist}(y_i, L_\phi^\infty) \leq \frac{1}{2}\mathbb{I}$, we have

$$\text{dist}(y_i, L_\phi) = \text{dist}(y_i, L_\phi^\infty) \leq \frac{1}{2}\mathbb{I} \quad \forall 1 \leq i \leq J.$$

As the diagonals of the Q_i have length $\sqrt{2}\mathbb{I}$, (A.15) and the triangle inequality imply

$$\text{dist}(x, L_\phi) \leq \text{dist}(y_i, L_\phi) + \frac{1}{2}\sqrt{2}\mathbb{I} = \frac{1}{2}(1 + \sqrt{2})\mathbb{I} < \frac{w}{4} \quad \forall x \in \bigcup_{i=1}^J Q_i.$$

To summarize, any element of Q_i for some $1 \leq i \leq J$ is within a distance $w/4$ of L_ϕ . As Rec^+ contains all points x within a distance $\frac{w}{2}$ of L_ϕ , we conclude

$$\bigcup_{i=1}^J Q_i \subseteq \text{Rec}^+.$$

Step 3: Lower bound for $\mathbb{P}_o(\tau_{\partial\text{Rec}} < \tau_{\partial\text{Rec}^+})$. From the previous step, to obtain a lower bound on the probability that the walk exits Rec before Rec^+ , it suffices to obtain an upper bound J^* on J and a lower bound $c < 1$ on

$$\mathbb{P}_\omega(\tau_{M_{i+1}} \leq \tau_{\partial^{\text{int}}Q_{i+1}}),$$

uniformly for $\omega \in M_i$, for $0 \leq i \leq J - 1$. This way, if we denote $Y_0 \equiv y$ and $Y_i = S_{\tau_{\partial^{\text{int}}Q_i}}$ for $1 \leq i \leq J - 1$, we can apply the strong Markov property to each τ_{M_i} and use the lower bound for each factor to obtain the lower bound

$$\mathbb{P}_o(\tau_{\partial\text{Rec}} < \tau_{\partial\text{Rec}^+}) \geq c^{J^*}. \quad (\text{A.16})$$

To obtain an upper bound on J , we first recall that L_ϕ has a length of $\ell + w/2$, which satisfies

$$\ell + w/2 = \frac{1}{2} \left(\frac{2\ell}{\mathbb{I}} + \frac{w}{\mathbb{I}} \right) \leq \frac{1}{2} \left(\frac{2\ell}{w/8 - 1} + \frac{w}{w/8 - 1} \right) \leq \frac{1}{2} \left(48 \frac{\ell}{w} + 24 \right), \quad (\text{A.17})$$

due to the fact that $\mathbb{I} \geq \lfloor w/8 \rfloor - 1 \geq w/8 - 2$ and the hypothesis of $w \geq 24$. The number of steps to reach J is no more than twice the ratio $(\ell + w/2)/(1/2)$. Accordingly, using the bound in (A.17) and the hypothesis that $\ell/w \geq 1$, we have

$$J \leq 2 \left(48 \frac{\ell}{w} + 24 \right) \leq 144 \frac{\ell}{w} =: J^*. \quad (\text{A.18})$$

We now turn to the hitting probability lower bounds.

From the construction, there are only two possible orientations of M_i relative to M_{i+1} (Figure 16). Either M_i and M_{i+1} have parallel orientation or they do not. Consider the former case. The hitting probability $\mathbb{P}_\omega(\tau_{M_{i+1}} \leq \tau_{\partial^{\text{int}}Q_{i+1}})$ is a harmonic function of ω for all ω in $Q_{i+1} \setminus \partial^{\text{int}}Q_{i+1}$ and M_{i+1} in particular. Therefore, by the Harnack inequality [Law13, Theorem 1.7.6], there is a constant a_1 such that

$$\mathbb{P}_\omega(\tau_{M_{i+1}} \leq \tau_{\partial^{\text{int}}Q_{i+1}}) \geq a_1 \mathbb{P}_{y_{i+1}}(\tau_{M_{i+1}} \leq \tau_{\partial^{\text{int}}Q_{i+1}}) \quad \forall \omega \in M_{i+1}. \quad (\text{A.19})$$




FIGURE 16. The two cases for lower-bounding M_{i+1} hitting probabilities.

The same argument applies to the case when M_i and M_{i+1} do not have parallel orientation and we find there is a constant a_2 such that (A.19) holds with a_2 in place of a_1 . Setting $a = \min\{a_1, a_2\}$, we conclude that, for all $0 \leq i \leq J - 1$ and any $\omega \in M_i$,

$$\mathbb{P}_\omega \left(\tau_{M_{i+1}} \leq \tau_{\partial^{\text{int}} Q_{i+1}} \right) \geq a \mathbb{P}_{y_{i+1}} \left(\tau_{M_{i+1}} \leq \tau_{\partial^{\text{int}} Q_{i+1}} \right). \quad (\text{A.20})$$

We have reduced the lower bound for any $\omega \in M_i$ and either of the two relative orientations of M_i and M_{i+1} to a lower bound on the hitting probability of one side of Q_{i+1} from the center. By symmetry, the walk hits M_{i+1} first with a probability of exactly $1/4$. We emphasize that the probability on the left-hand side of (A.20) is exactly $1/4$ as although M_{i+1} does not include the adjacent corners of Q_{i+1} , which are elements of $\partial^{\text{int}} Q_{i+1}$, the corners are separated from y_{i+1} by the other elements of $\partial^{\text{int}} Q_{i+1}$.

Calling $b = a/4$ and combining (A.18) and (A.20) with (A.16), we have

$$\mathbb{P}_o \left(\tau_{\partial^{\text{Rec}}} < \tau_{\partial^{\text{Rec}^+} } \right) \geq b^{J^*} = b^{144\ell/w} = c^{\ell/w}$$

for a positive constant $c < 1$. □

REFERENCES

- [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.
- [DPRZ06] Amir Dembo, Yuval Peres, Jay Rosen, and Ofer Zeitouni. Late points for random walks in two dimensions. *Ann. Probab.*, 34(1):219–263, 2006.
- [GLPP17] Shirshendu Ganguly, Lionel Levine, Yuval Peres, and James Propp. Formation of an interface by competitive erosion. *Probab. Theory Related Fields*, 168(1-2):455–509, 2017.
- [Kes86] Harry Kesten. Aspects of first passage percolation. In *École d’été de probabilités de Saint-Flour, XIV—1984*, volume 1180 of *Lecture Notes in Math.*, pages 125–264. Springer, Berlin, 1986.
- [Kes87] Harry Kesten. Hitting probabilities of random walks on \mathbf{Z}^d . *Stochastic Process. Appl.*, 25(2):165–184, 1987.
- [KS04] Gady Kozma and Ehud Schreiber. An asymptotic expansion for the discrete harmonic potential. *Electron. J. Probab.*, 9:no. 1, 1–17, 2004.
- [Law93] Gregory F Lawler. A discrete analogue of a theorem of Makarov. *Combinatorics, Probability and Computing*, 2(2):181–199, 1993.
- [Law13] Gregory F. Lawler. *Intersections of random walks*. Modern Birkhäuser Classics. Birkhäuser/Springer, New York, 2013. Reprint of the 1996 edition.
- [Pop21] Serguei Popov. *Two-dimensional Random Walk: From Path Counting to Random Interlacements*, volume 13. Cambridge University Press, 2021.

[Tim13] Ádám Timár. Boundary-connectivity via graph theory. *Proc. Amer. Math. Soc.*, 141(2):475–480, 2013.

DEPARTMENTS OF MATHEMATICS AND STATISTICS, U.C. BERKELEY, EVANS HALL, BERKELEY, CA, 94720-3840, U.S.A.

Email address: jacob_calvert@berkeley.edu, sganguly@berkeley.edu, alanmh@berkeley.edu