

EXCEPTIONAL GRAPHS FOR THE RANDOM WALK

JUHAN ARU, CARLA GROENLAND, TOM JOHNSTON, BHARGAV NARAYANAN,
ALEX ROBERTS, AND ALEX SCOTT

ABSTRACT. If \mathcal{W} is the simple random walk on the square lattice \mathbb{Z}^2 , then \mathcal{W} induces a random walk \mathcal{W}_G on any spanning subgraph $G \subset \mathbb{Z}^2$ of the lattice as follows: viewing \mathcal{W} as a uniformly random infinite word on the alphabet $\{\mathbf{x}, -\mathbf{x}, \mathbf{y}, -\mathbf{y}\}$, the walk \mathcal{W}_G starts at the origin and follows the directions specified by \mathcal{W} , only accepting steps of \mathcal{W} along which the walk \mathcal{W}_G does not exit G . For any fixed $G \subset \mathbb{Z}^2$, the walk \mathcal{W}_G is distributed as the simple random walk on G , and hence \mathcal{W}_G is almost surely recurrent in the sense that \mathcal{W}_G visits every site reachable from the origin in G infinitely often. This fact naturally leads us to ask the following: does \mathcal{W} almost surely have the property that \mathcal{W}_G is recurrent for *every* $G \subset \mathbb{Z}^2$? We answer this question negatively, demonstrating that exceptional subgraphs exist almost surely. In fact, we show more to be true: exceptional subgraphs continue to exist almost surely for a countable collection of independent simple random walks, but on the other hand, there are almost surely no exceptional subgraphs for a branching random walk.

1. INTRODUCTION

Let us say that a walk on a graph G is *recurrent* if the walk visits every site in the connected component of its starting point in G infinitely often, and *transient* otherwise. It is a classical result of Pólya [17] that a simple random walk on the square lattice \mathbb{Z}^2 is almost surely recurrent. In this paper, we shall be concerned with how ‘robust’ this property is in the following sense: do the coin tosses that determine a recurrent random walk on \mathbb{Z}^2 also determine a recurrent random walk on *every* subgraph of \mathbb{Z}^2 simultaneously? We make this question precise below in a few different ways.

Date: 20 July 2018.

2010 Mathematics Subject Classification. Primary 05C81; Secondary 60G50.

Key words and phrases. Random walks, recurrence and transience, exceptional graphs, traversal sequences.

We view a simple random walk \mathcal{W} on \mathbb{Z}^2 as a random infinite word on the four-letter alphabet $\{\mathbf{x}, -\mathbf{x}, \mathbf{y}, -\mathbf{y}\}$, where $\mathbf{x} = (1, 0)$ and $\mathbf{y} = (0, 1)$, with each letter of \mathcal{W} being chosen independently and uniformly at random. For any (spanning) subgraph $G \subset \mathbb{Z}^2$ of the lattice, the random walk \mathcal{W} then induces a (random) walk \mathcal{W}_G on G : starting at the origin, we consider the letters of \mathcal{W} one at a time, and for each letter of \mathcal{W} , we take a step in the appropriate direction in G provided the edge in question is present in G , and stand still otherwise. For any *fixed* $G \subset \mathbb{Z}^2$, it is clear that \mathcal{W}_G is distributed as the simple random walk on G , and since G is a subgraph of a recurrent graph, we conclude that \mathcal{W}_G is almost surely recurrent. It follows immediately from Fubini's theorem that, if we consider any Borel probability measure on the space of subgraphs of \mathbb{Z}^2 , then the random walk \mathcal{W} almost surely has the property that the induced walks \mathcal{W}_G are recurrent for *almost all* $G \subset \mathbb{Z}^2$. We are then naturally led to the following question: does the random walk \mathcal{W} almost surely have the property that the induced walks \mathcal{W}_G are recurrent for *all* $G \subset \mathbb{Z}^2$? Our first result answers this question negatively in the following strong sense.

Theorem 1.1. *If \mathcal{W} is a simple random walk on \mathbb{Z}^2 , then there almost surely exists a (random) exceptional subgraph $H \subset \mathbb{Z}^2$ for which the induced walk \mathcal{W}_H*

- (1) *visits each site reachable from the origin in H finitely many times, and*
- (2) *fails to visit infinitely many sites reachable from the origin in H .*

More generally, we can ask whether a countably infinite independent collection of simple random walks almost surely has the property that, for every $G \subset \mathbb{Z}^2$, at least one of the walks in this collection induces a recurrent walk on G . An extension of the proof of Theorem 1.1 allows us to prove the following result, which answers this question negatively as well.

Theorem 1.2. *If $\{\mathcal{W}_i\}_{i \in \mathbb{N}}$ is a collection of independent simple random walks on \mathbb{Z}^2 , then there almost surely exists a (random) exceptional subgraph $H \subset \mathbb{Z}^2$ so that, for every $i \in \mathbb{N}$, the induced walk $(\mathcal{W}_i)_H$*

- (1) *visits each site reachable from the origin in H finitely many times, and*
- (2) *fails to visit infinitely many sites reachable from the origin in H .*

What about an uncountable collection of simple random walks? To avoid measurability issues around independence, we need to be careful about how we phrase such a question. One natural formulation is in the language of branching

random walks. A *branching random walk* on \mathbb{Z}^d starts with a single particle at the origin. At each time step, each particle independently generates a number of additional particles at its current location according to some fixed *offspring distribution*, and we say that this offspring distribution is *nontrivial* if the number of offspring is nonzero with positive probability. Independently of the other particles and their history, all particles then take a step in a direction chosen uniformly at random, which leads to a random family of dependent simple random walks. In this language, one can then ask whether a branching random walk on \mathbb{Z}^2 almost surely has the property that, for every $G \subset \mathbb{Z}^2$, at least one of the branches of the branching random walk induces a recurrent walk on G . Our final result answers this question positively, and furthermore, shows that the same is true in dimensions greater than two as well.

Theorem 1.3. *Fix $d \in \mathbb{N}$ and let $(\mathcal{W}_i)_{i \in S}$ be a family of random walks generated by a branching random walk on \mathbb{Z}^d with a nontrivial offspring distribution. Then, almost surely, for every (spanning) subgraph $G \subset \mathbb{Z}^d$, there is some $j \in S$ for which the induced walk $(\mathcal{W}_j)_G$ is recurrent.*

Our work here fits into the broader context of attempting to understand the robustness of objects such as random walks and Brownian paths, in terms of their quasi-everywhere properties or their dynamical sensitivity; see [10, 12, 5] for example.

Let us mention two results in this general direction that are particularly close to our results in spirit. The first relevant result is a theorem of Adelman, Burdzy and Pemantle [1] on the projections of three-dimensional Brownian motion. The projection of Brownian motion in \mathbb{R}^3 onto any *fixed* plane yields Brownian motion in that plane which is neighbourhood recurrent; Adelman, Burdzy and Pemantle [1] however show that there almost surely exists a (random) exceptional plane on which the projection is not neighbourhood recurrent. The second result that is pertinent is a theorem of Hoffman [11] demonstrating that recurrence of the simple random walk on \mathbb{Z}^2 is dynamically sensitive; in other words, if the coin tosses of the random walk are refreshed continuously with Poisson clocks generating a *dynamic random walk*, then although dynamic random walk is almost surely recurrent at any *fixed* time, there almost surely exists a (random) exceptional time at which the dynamic random walk is not recurrent.

Although the two results mentioned above bear some similarities in flavour to our first two results, it is perhaps worth remarking that the methods of proof are somewhat different: while the results in both [1] and [11] are based on second moment computations, the proofs of Theorems 1.1 and 1.2, in contrast, proceed by explicitly ‘embedding drift’. In terms of proofs, our constructive approach however shares more in terms of analogy with the following two results. First, in [2], Amir, Benjamini, Gurel-Gurevich and Kozma give a construction of a transient walk by dynamically changing conductances of the underlying graph $G = \mathbb{Z}^2$ along the walk in a way not too dissimilar to our construction. Second, in [4], Barlow and Perkins construct a (deterministic) exceptional graph $G \subset \mathbb{Z}^2$, a certain binary tree embedded in \mathbb{Z}^2 , on which the random walk remains barely recurrent: along the sequence $m_n = 12(2^{2^{n+2}} - n)$ the simple random walk $(X_t)_{t \geq 0}$ on G satisfies

$$\liminf_{n \rightarrow \infty} |X_{m_n}| (m_n \log m_n)^{-1/2} \geq 1/17,$$

almost surely.

This note is organised as follows. In Section 2, we first sketch a natural approach that fails, but nonetheless motivates our construction, then we prove Theorem 1.1, and we finally sketch how the same argument extends to prove Theorem 1.2. Section 3 is devoted to the proof of Theorem 1.3. We conclude in Section 4 with a discussion of some open problems.

2. EXISTENCE OF EXCEPTIONAL SUBGRAPHS

In this section, we prove Theorems 1.1 and 1.2. Before we do so, let us sketch a construction which, while failing to prove Theorem 1.1, serves as the motivation for the construction in our proof.

Suppose that \mathcal{W} is a simple random walk on \mathbb{Z}^2 . Let us construct a (random) subgraph $P \subset \mathbb{Z}^2$ which exhibits some drift. We shall ensure that P is an infinite non-decreasing path (i.e., a north-east path) passing through the origin, and we reveal P as follows. We shall read off the letters of \mathcal{W} one at a time and follow the induced walk \mathcal{W}_P on P , revealing more of P as and when \mathcal{W}_P needs to know if a particular edge is present in P . At any finite time, it is clear that P (or rather, what has been revealed of P so far) is a finite non-decreasing path through the origin, and we are forced to reveal more of P at this time if and only if \mathcal{W}_P is at one of the leaves of P and the next letter of \mathcal{W} would cause \mathcal{W}_P to exit P in a non-decreasing fashion. Our strategy for constructing P is then as follows: if \mathcal{W}_P

is at the north-eastern leaf of P at some stage, and the next letter of \mathcal{W} causes \mathcal{W}_P to travel either north or east, we extend P so as to allow this, proceeding analogously at the south-western leaf as well.

What can we say about the induced walk \mathcal{W}_P on the path P as constructed above? It is not hard to see that if we identify P with \mathbb{Z} by ‘unrolling’ it, then \mathcal{W}_P is a random walk on \mathbb{Z} with the following law: if $[a, b] \subset \mathbb{Z}$ is the range of the walk at some time and the walk is at $x \in [a, b]$ at this time, then the walk moves to either $x - 1$ or $x + 1$ both with probability $1/2$, unless $x \in \{a, b\}$, in which case, the walk moves to $b + 1$ with probability $2/3$ and to $b - 1$ with probability $1/3$ at $x = b$, and similarly to $a - 1$ with probability $2/3$ and to $a + 1$ with probability $1/3$ at $x = a$. In other words, for the path P constructed as described above, the induced walk \mathcal{W}_P behaves like a random walk on \mathbb{Z} with a tiny amount of drift; indeed, the walk possesses some drift away from the origin when it is at the boundary of its range, but behaves like the simple random walk in the interior of its range. Unfortunately, this tiny amount of drift does not stop \mathcal{W}_P from being recurrent, but this construction nevertheless demonstrates that it is possible to construct (random) subgraphs of \mathbb{Z}^2 where the induced walk possesses some drift; below, we prove Theorem 1.1 with a more careful construction that embeds more drift into the induced walk.

We need two simple facts about the simple random walk on \mathbb{Z} . First, we require the following well-known fact.

Proposition 2.1. *The probability that the simple random walk on the interval $\{0, 1, \dots, n\}$ started at 1 visits n before it visits 0 is $1/n$.* \square

Next, we shall also make use of the following crude bound.

Proposition 2.2. *The expected number of times the simple random walk on \mathbb{Z} started at 0 visits 0 in the first N steps is at most $10\sqrt{N}$.* \square

Armed with these two facts, we are now ready to prove our main result.

Proof of Theorem 1.1. To prove the theorem, we will construct a (random) graph H based on the random walk \mathcal{W} where the induced walk \mathcal{W}_H exhibits a strong drift away from the origin.

The graph H we construct will consist of the vertical lines $L_n = \{(x, y) \in \mathbb{Z}^2 : x = 2^n - 1\}$ for all integers $n \geq 0$, and a (random) collection of finite horizontal

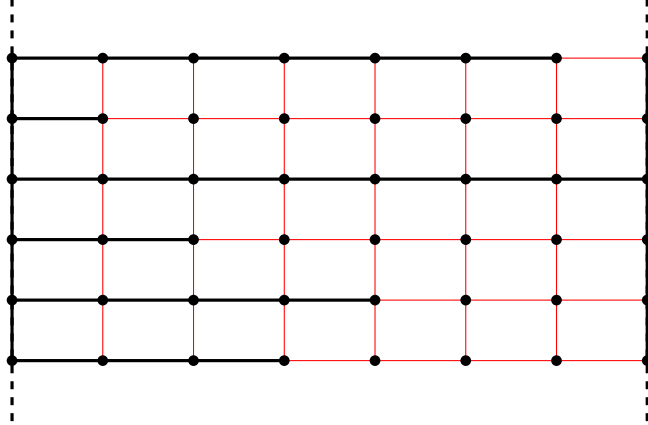


FIGURE 1. A slice of H between two consecutive vertical lines in H ; edges coloured black are present while edges coloured red are absent.

segments between any consecutive pair of vertical lines with the property that exactly one such horizontal segment connects any consecutive pair of vertical lines; here, by lines and segments, we mean the edges in the appropriate paths in \mathbb{Z}^2 , as shown in Figure 1.

Let us denote the location of the induced walk \mathcal{W}_H at a given time $t \geq 0$ by $\mathcal{W}_H(t) = (X_H(t), Y_H(t))$, with the time t tracking steps along \mathcal{W}_H (as opposed to \mathcal{W}). As before, we shall read off the letters of \mathcal{W} one at a time, and we shall reveal H by following the induced walk \mathcal{W}_H and revealing more of H as necessary.

Notice that the only vertical edges in H are (deterministically) those on a vertical line L_n for some $n \geq 0$, so at any time $t \geq 0$, the induced walk \mathcal{W}_H accepts a vertical step of \mathcal{W} if and only if $\mathcal{W}_H(t) \in L_n$ for some $n \geq 0$.

We reveal the horizontal edges of H in stages: during stage $n \geq 0$, we shall reveal all the horizontal edges of H in between the lines L_n and L_{n+1} , with the stage ending as soon as there is a horizontal path connecting these vertical lines in H . Note in particular that during stage n , we have already revealed all the horizontal edges in H between L_0 and L_n , and none of the horizontal edges in H to the right of L_{n+1} .

We begin by declaring every horizontal edge to the left of L_0 as being absent in H , and for $n \geq 0$, having completed stage $n - 1$, we reveal H in the following fashion. At some $t \geq 0$ during stage n , there are two possibilities. If $X_H(t) < 2^n - 1$, then we have nothing to do when we read off the next letter of \mathcal{W} since all the

edges of H to the left of L_n have already been revealed. If $X_H(t) \geq 2^n - 1$ on the other hand, we reveal H in such a way so as to ensure that \mathcal{W}_H always accepts a letter of \mathcal{W} that would cause the induced walk to travel to the right. The stage ends as soon as we have a single horizontal path connecting L_n and L_{n+1} , or in other words, at the first time $t \geq 0$ when we have $X_H(t) = 2^{n+1} - 1$. At the end of the stage, all horizontal edges between L_n and L_{n+1} whose presence or absence in H have not been revealed over the course of the stage, we declare as being absent in H . In particular, the line L_{n+1} is incident to precisely one horizontal edge to its left in H and it has been revealed by the end of the stage.

The above construction clearly ensures that H has the structure we promised. More is true, however; as we shall shortly see, our construction endows the induced walk \mathcal{W}_H with a strong drift to the right.

For $n \geq 0$, let τ_n be the first time t at which we have $X_H(t) = 2^n - 1$, and let E_n denote the event that the walk $(\mathcal{W}_H(t))_{t \geq \tau_n}$ hits the line L_{n-1} before hitting the line L_{n+1} ; in other words, E_n is the event that there exists a time $t \in [\tau_n, \tau_{n+1})$ at which $X_H(t) = 2^{n-1} - 1$. With these definitions in place, we have the following claim.

Claim 2.3. *There exists an absolute constant $c \in (0, 1)$ such that $\mathbb{P}(E_n) \leq c^n$ for all $n \geq 0$.*

Proof. Let α be the y -coordinate of the unique horizontal path joining the vertical lines L_{n-1} and L_n , and let T be the first time after τ_n at which the walk \mathcal{W}_H hits either the line L_{n-1} or the line L_{n+1} . Note that the time between τ_n and T naturally decomposes into *excursions*, where an excursion is a maximal interval of time during which the walk \mathcal{W}_H remains at some fixed y -coordinate.

Let us now describe the walk \mathcal{W}_H in terms of its excursions. First, note that our construction of H ensures that the y -coordinates of successive excursions of \mathcal{W}_H are determined by a simple random walk on \mathbb{Z} started at α . Also, we can describe an excursion at some y -coordinate β as follows. If $\beta \neq \alpha$, then during an excursion at β , successive x -coordinates of \mathcal{W}_H are determined by a simple random walk on the interval $\{2^n - 1, 2^n, \dots, 2^{n+1} - 1\}$ started at $2^n - 1$, with the excursion ending either, with probability $1/3$, when the x -coordinate of the walk is $2^{n+1} - 1$, or, with probability $2/3$, when the x -coordinate of the walk is $2^n - 1$. If $\beta = \alpha$ on the other hand, then successive x -coordinates of \mathcal{W}_H are determined by a simple random walk on the interval $\{2^{n-1} - 1, 2^{n-1}, \dots, 2^n - 1\}$ started at $2^{n-1} - 1$, with

the excursion ending either, with probability 1, when the x -coordinate of the walk is either $2^{n-1} - 1$ or $2^{n+1} - 1$, or, with probability $1/2$, when the x -coordinate of the walk is $2^n - 1$. Crucially, note that by the strong Markov property of the walk \mathcal{W}_H , each excursion depends on past excursions only through the endpoint of the last excursion preceding it.

Let us say that an excursion is *positively successful* if it ends on account of the walk \mathcal{W}_H reaching the line L_{n+1} , and *negatively successful* if it ends on account of the walk \mathcal{W}_H reaching the line L_{n-1} . In this language, we see that E_n is precisely the event that we witness a negatively successful excursion before a positively successful one.

We would like to show that in the first 3^n excursions, there is a positively successful excursion, but no negatively successful one. A minor technicality arises from the fact that excursions are only defined until the first successful one occurs. To circumvent this issue, consider a modified process coupled with \mathcal{W}_H which, after a successful excursion with y -coordinate β , teleports to the line L_n , taking y -coordinate $\beta + 1$ with probability $1/2$, and $\beta - 1$ otherwise, before again moving according to \mathcal{W}_H . Let F_1 to be the event that at least one of the first 3^n excursions in this modified process is positively successful, and let F_2 be the event that none of the first 3^n excursions in the modified process are negatively successful. Since \mathcal{W}_H and our modified process behave identically until the first successful excursion, if F_1 and F_2 both occur, then E_n cannot occur. Therefore, it suffices to show that both F_1 and F_2 are overwhelmingly likely.

First, we deal with the event F_1 . It is easy to see from Proposition 2.1 that an excursion is positively successful with probability at least $(1/100)2^{-n}$. Using the strong Markov property, we conclude that

$$\mathbb{P}(F_1^c) \leq (1 - 2^{-n}/100)^{3^n} \leq c_1^n,$$

where $c_1 \in (0, 1)$ is an absolute constant.

Next, we handle the event F_2 . Notice that we may only witness a negatively successful excursion at y -coordinate α ; with this in mind, let Z be the number of excursions in the first 3^n excursions at y -coordinate α . Since the y -coordinates of successive excursions are determined by a simple random walk on \mathbb{Z} started at α , we conclude from Proposition 2.2 that

$$\mathbb{P}(Z \geq (7/4)^n) \leq \frac{10\sqrt{3^n}}{(7/4)^n} \leq c_2^n,$$

where $c_2 \in (0, 1)$ is an absolute constant. As before, we know from Proposition 2.1 that an excursion at y -coordinate α is negatively successful with probability at most $100 \cdot 2^{-n}$, so we may again use the strong Markov property to conclude that

$$\begin{aligned} \mathbb{P}(F_2^c) &\leq \mathbb{P}(Z \geq (7/4)^n) + \mathbb{P}(F_2^c \mid Z < (7/4)^n) \\ &\leq c_2^n + (1 - (1 - 100 \cdot 2^{-n})^{(7/4)^n}) \\ &\leq c_2^n + c_3^n, \end{aligned}$$

where $c_3 \in (0, 1)$ is an absolute constant.

The result follows from the estimates above since $\mathbb{P}(E_n) \leq \mathbb{P}(F_1^c) + \mathbb{P}(F_2^c)$. \square

It follows from the above claim, by the Borel–Cantelli lemma, that the walk \mathcal{W}_H almost surely visits the line L_n only finitely many times for each $n \geq 0$, thus proving the result. \square

The proof of Theorem 1.2 follows from a simple modification of the proof of Theorem 1.1; therefore, we only provide a sketch highlighting the main differences.

Proof of Theorem 1.2. As in the argument above, we begin with the vertical lines L_1, L_2, \dots , where L_n has x -coordinate $2^n - 1$. Suppose that we have defined the exceptional graph H up to L_i by using the walks $\mathcal{W}_1, \dots, \mathcal{W}_i$. To define the portion of H between L_i and L_{i+1} , we run \mathcal{W}_{i+1} on the already defined portion of H until it first hits L_i . At that point, we run the algorithm used in the proof above on all of the $i + 1$ walks $\mathcal{W}_1, \dots, \mathcal{W}_{i+1}$ (so that they can move freely to the right of L_i but not up or down in that region) in, say, numerical order, stopping each walk when it first hits L_{i+1} . We continue this process, introducing one new walk at each step. The analysis goes through essentially as before: if $E_{n,i}$ is the event that, after hitting L_n , the walk \mathcal{W}_i hits L_{n-1} before L_{n+1} , then as in Claim 2.3, we have $\mathbb{P}(E_{n,i}) \leq nc^n$ for all $n \geq i$ for some absolute constant $c \in (0, 1)$. Notice that we rely on independence here: we use the fact that the walk \mathcal{W}_i , conditioned on its trajectory up to hitting L_n and on the exceptional graph built up to L_n , still proceeds as a simple random walk. The result again follows from the Borel–Cantelli lemma. \square

3. NON-EXISTENCE OF EXCEPTIONAL SUBGRAPHS

We now consider branching random walks on \mathbb{Z}^d in all dimensions $d \geq 2$. Recall that we start with a single particle at the origin, and that at each time step, every

particle independently generates a number of additional particles at its current location according to some fixed nontrivial distribution, where at least one new particle is generated with positive probability; independently of the other particles and their history, all particles then take a step in a direction chosen uniformly at random. The result is a random family $(\mathcal{W}_i)_{i \in S}$ of dependent simple random walks, where the branches \mathcal{W}_i with $i \in S$ each correspond to a walk on \mathbb{Z}^d obtained by starting with the original particle at the origin and then following either the particle presently under consideration or one of its children at each time.

We will use the following Chernoff-type bound; see [15] for a proof.

Proposition 3.1. *For $n \in \mathbb{N}$, $p \in (0, 1)$ and $\varepsilon > 0$, we have*

$$\mathbb{P}(\text{Bin}(n, p) \leq np(1 - \varepsilon)) \leq \exp\left(-\frac{\varepsilon^2 np}{2}\right). \quad \square$$

We are now ready to give the proof of our final result.

Proof of Theorem 1.3. Note that removing particles reduces the set of walks and makes the problem of finding a recurrent walk harder. Hence, if at some point a particle has more than one child, we may discard all but one of its children. We may therefore assume that, for some $\varepsilon > 0$, each particle either has one child with probability ε , or no children at all with probability $1 - \varepsilon$. It will be helpful to have some notation: in what follows, we write $B_G(x, \rho)$ to denote the set of vertices at graph distance at most ρ from x in a graph G , write $B(x, \rho)$ for $B_{\mathbb{Z}^d}(x, \rho)$, and abbreviate $B(0, \rho)$ by $B(\rho)$.

In this proof, we will need to show that, with very high probability, certain particles have exponentially many children at a given future time and ensure that most of these children do not wander very far from the origin. To this end, we need two results which we state and prove below.

First, we need the following estimate for the probability of a simple random walk on \mathbb{Z}^d getting unexpectedly far from its starting point after taking some finite number of steps.

Claim 3.2. *Fix $\delta > 0$ and let S_n be an n -step simple random walk on some subgraph $G \subset \mathbb{Z}^d$ starting from the origin. Then we have*

$$\mathbb{P}(S_n \notin B_G(0, \delta n)) \leq c_d n^d \exp\left(-\frac{\delta^2 n}{2}\right),$$

where $c_d > 0$ is a constant depending only on the dimension d .

Proof. We make use of an old bound on the transition probabilities of a Markov chain due to Varopoulos [19] and Carne [6], although we use it in a more recent form due to Peyre [16]. To state this bound, we need a little set up. Given a graph G and a pair of vertices x and y of G in the same connected component, let $p_t(x, y)$ denote the probability that a simple random walk on G starting at x is at y after t steps, and write $\rho(x, y)$ for the graph distance between x and y in G ; in this language, we have

$$p_t(x, y) \leq 2\sqrt{\frac{\deg(y)}{\deg(x)}} \exp\left(-\frac{\rho(x, y)^2}{2t}\right).$$

With the above bound in hand, we conclude that the probability that an n -step simple random walk on some subgraph $G \subset \mathbb{Z}^d$ starting at the origin ends up outside the ball $B_G(0, \delta n)$ is at most

$$\sum_{y \in B(n) \setminus B_G(0, \delta n)} p_n(0, y) \leq \sum_{y \in B(n)} 2\sqrt{2d} \exp\left(-\frac{\delta^2 n^2}{2n}\right) \leq \sqrt{8d}(2n+1)^d \exp\left(-\frac{\delta^2 n}{2}\right);$$

it is clear that the bound above is of the required form, proving the claim. \square

Second, we need the following estimate for the rate of growth of a Galton–Watson branching process.

Claim 3.3. *Let $(N_j)_{j \geq 0}$ be the number of descendants at time j of a Galton–Watson branching process started with a single particle and with an offspring distribution that takes the value 2 with probability ε and 1 with probability $1 - \varepsilon$. Then there exists constants $c, c' > 0$ such that*

$$\mathbb{P}(N_j \geq (1+c)^j) \geq 1 - e^{-c'j}$$

for all $j \geq 1$.

Proof. Conditioned on N_{j-1} , the random variable N_j is distributed as $\text{Bin}(N_{j-1}, \varepsilon) + N_{j-1}$, independent of everything else in the past. For $j \geq 1$, define $Y_j = N_j/N_{j-1}$

and note that $Y_j \in [1, 2]$ and $\mathbb{E}[Y_j] = 1 + \varepsilon$. From Proposition 3.1, we deduce that

$$\begin{aligned} \mathbb{P}(Y_j \geq 1 + \varepsilon/2) &= \sum_{i=1}^{\infty} \mathbb{P}(\text{Bin}(i, \varepsilon) \geq i\varepsilon/2 \mid N_{j-1} = i) \mathbb{P}(N_{j-1} = i) \\ &\geq \sum_{i=1}^{\infty} (1 - e^{-i\varepsilon/8}) \mathbb{P}(N_{j-1} = i) \\ &\geq (1 - e^{-\varepsilon/8}). \end{aligned}$$

Thus, setting $p = (1 - e^{-\varepsilon/8})$, there exist independent random variables $(X_j)_{j \geq 1}$ dominated by the Y_j such that X_j takes the value $1 + \varepsilon/2$ with probability p and 1 otherwise with probability $1 - p$. Since $N_j = Y_j Y_{j-1} \cdots Y_1 \geq X_j X_{j-1} \cdots X_1$, by Proposition 3.1, we have

$$\begin{aligned} \mathbb{P}(N_j \geq (1 + \varepsilon/2)^{jp/2}) &\geq \mathbb{P}\left(|\{i \in [j] : X_i > 1\}| \geq \frac{jp}{2}\right) \\ &\geq \mathbb{P}\left(\text{Bin}(j, p) \geq \frac{jp}{2}\right) \\ &\geq 1 - e^{-\frac{jp}{8}}, \end{aligned}$$

as required. \square

Before diving into the details of the proof, let us sketch our plan of attack. Our argument will proceed in *stages* bookended by a rapidly-increasing sequence of times $(T_i)_{i \in \mathbb{N}}$. At each time T_i , and for every possible finite subgraph $G \subset \mathbb{Z}^d$ contained in the box $[-T_i, T_i]^d$, we arbitrarily choose a representative particle p_G . We then show that, with very high probability, for each such representative particle p_G and every possible extension of the corresponding graph G onto $[-T_{i+1}, T_{i+1}]^d$, some descendant of p_G visits every vertex of G with distance at most i (in G) to the origin, and ends up exactly at or one step away from the origin at time T_{i+1} . We shall show that the failure probabilities decay rapidly enough so that we may finish by applying the Borel–Cantelli lemma. To avoid clutter, we will not worry about making sure all the appropriate values are integers.

Let $T_0 = 0$, and for the singleton graph containing just the origin, we choose the initial particle at the origin. Let δ be sufficiently small such that $2^{-4d\delta}(1+c)^{(1-4\delta)} > 1$ and $c\delta - c'(1-5\delta) < 0$, where c, c' are the absolute constants promised by Claim 3.3.

Suppose now that we have run the branching walk over the course of $\ell - 1$ stages until time $n = T_{\ell-1}$. For each of the at most $2^{d(2n+1)^d}$ possible finite

subgraphs $G \subset \mathbb{Z}^d$ contained in the box $[-n, n]^d$, we now arbitrarily choose a representative particle p_G at some location in $[-n, n]^d$. Note that these particles need not be distinct, and that a single particle may be at different locations in different subgraphs. We now describe how we construct the ℓ th stage of the branching walk.

Let $N > n/\delta$ be a suitably large integer. We specify what conditions N needs to satisfy in what follows, and shall then take $T_\ell = N$. We shall further divide the ℓ th stage, which consists of the interval of time $[n + 1, N]$, into three smaller blocks of time as follows.

In the first block, we run the branching walk for δN more steps (after time n) so that for each $G \subset \mathbb{Z}^d$ contained in the box $[-n, n]^d$, the representative particle p_G has a set of descendants P_G after these δN steps. By Claim 3.3, any fixed representative particle p_G has at least $(1 + c)^{\delta N}$ descendants with probability at least $1 - e^{-c'\delta N}$. If even one of the representative particles has not branched this much, we declare that this step has *failed*. Note that we are free to discard particles, so we may assume (as long as this step has not failed) that each representative particle p_G has a set P_G of exactly $(1 + c)^{\delta N}$ descendants. Note that all the particles under consideration are at a graph distance of at most $n + \delta N < 2\delta N$ from the origin in all the graphs under consideration.

In the second block, we run the branching walk for another $N - 5\delta N$ steps. For each $G \subset \mathbb{Z}^d$ contained in the box $[-n, n]^d$, we now count the number of descendants of each $q \in P_G$. We assume that every such $q \in P_G$ has at least $(1 + c)^{N - 5\delta N}$ descendants, and if this ever fails to hold, we again declare that the step has *failed*. By Claim 3.3, the probability of failure for any particular particle q as above is at most $\exp(-c'(N - 5\delta N))$.

Now, fix $H \subset \mathbb{Z}^d$ contained in the box $[-N, N]^d$, and suppose that it induces a graph H' on $[-n, n]^d$. We say that a particle $q \in P_{H'}$ is *H-good* if at least half of its descendants are no further than $3\delta N$ from the origin in H after the $n + N - 4\delta N$ steps taken so far. Let us estimate the probability λ_H that any given particle $q \in P_{H'}$ is *H-good*. Fix $q \in P_{H'}$ and denote the position of q at time $n + \delta N$ by L_q . By Claim 3.2, the probability that a given descendant of q is within distance δN of L_q in H , and hence within distance $3\delta N$ of the origin, is at least

$$1 - c_d N^d \exp\left(-\frac{\delta^2 N}{2}\right).$$

Let Y be the proportion of descendants of q which are within $3\delta N$ of the origin. By the linearity of expectation,

$$\mathbb{E}[Y] \geq 1 - c_d N^d \exp\left(-\frac{\delta^2 N}{2}\right).$$

If q is not good, then $Y \leq 1/2$, and since $Y \leq 1$, we find that $\lambda_H + \frac{1}{2}(1 - \lambda_H) \geq \mathbb{E}[Y]$. Hence,

$$\lambda_H \geq 1 - 2c_d N^d \exp\left(-\frac{\delta^2 N}{2}\right)$$

and, by choosing N sufficiently large, we may certainly assume that $\lambda_H \geq 2/3$. We now also declare the step to have *failed* if, for some H as above, at most $\frac{1}{3}(1+c)^{\delta N}$ elements of P_H are H -good, and deduce from Proposition 3.1 that the probability of failing in this fashion for any fixed H is at most

$$\begin{aligned} \mathbb{P}\left(\text{Bin}((1+c)^{\delta N}, \lambda) \leq \frac{1}{3}(1+c)^{\delta N}\right) &\leq \mathbb{P}\left(\text{Bin}((1+c)^{\delta N}, \lambda) \leq \frac{1}{2}(1+c)^{\delta N} \lambda\right) \\ &\leq \exp\left(-\frac{1}{8}(1+c)^{\delta N} \lambda\right) \\ &\leq \exp\left(-\frac{1}{16}(1+c)^{\delta N}\right). \end{aligned}$$

Hence, if we have not already failed, then for any H as above, counting the descendants of H -good particles that don't stray too far from the origin, we have at least $\frac{1}{6}(1+c)^{N-4\delta N}$ such descendants at graph distance at most $3\delta N$ from the origin in H after $n + N - 4\delta N$ steps; we call these descendants *H-counters*.

In the third and final block of the stage, our goal is to visit all vertices at distance at most $\ell - 1$ from the origin in every $H \subset \mathbb{Z}^d$ contained in the box $[-N, N]^d$. For any fixed H as above, there are at most $(2\ell - 1)^d$ such vertices close to the origin. We enumerate these vertices and pick a path of length at most $\ell - 1$ from the origin to each such vertex in H ; putting these together, we get a walk W_H of length at most $(2\ell - 2)(2\ell - 1)^d$ in H which starts at the origin, ends at the origin, and visits every vertex at most distance $\ell - 1$ from the origin; furthermore, by choosing N sufficiently large, we may assume that $(2\ell - 2)(2\ell - 1)^d \leq \delta N - n$.

For each H -counter v , let A_v be the event that, in the next $4\delta N - n$ steps, v visits every vertex at most distance $\ell - 1$ from the origin. If no A_v occurs, for any H as above, we again declare the stage to have *failed*. Conditional on the positions of the H -counters at the start of this block, the events A_v are independent of one another. Now, fix these 'starting positions' of the $\frac{1}{6}(1+c)^{N-4\delta N}$ H -counters.

For each H -counter v , there is a path of length at most $3\delta N$ from its starting position to the origin. Let B_v be the event that v strictly follows this path, then the walk W_H and then never leaves $B_H(0, 1)$. Clearly B_v implies A_v , and so $\mathbb{P}[A_v] \geq \mathbb{P}[B_v] \geq (2d)^{-4\delta N}$. Hence, conditional on the starting positions of the H -counters, the probability that no A_v occurs is at most

$$(1 - (2d)^{-4\delta N})^{\frac{1}{6}(1+c)^{N-4\delta N}} \leq \exp\left(-\frac{1}{6}(2d)^{-4\delta N}(1+c)^{N-4\delta N}\right).$$

Since the bound above is independent of the starting positions of the H -counters, it also holds without conditioning for any fixed H as above.

In summary, we declare the ℓ th stage to have failed if one of the following happens.

- (1) In the first block, there is some G contained in $[-n, n]^d$ whose representative particle p_G does not branch sufficiently. This happens with probability at most

$$2^{d(2n+1)^d} \exp(-c'\delta N). \quad (1)$$

- (2) In the second block, there is some G contained in $[-n, n]^d$ for which some $q \in P_G$ does not branch sufficiently. This happens with probability at most

$$2^{d(2n+1)^d} (1+c)^{\delta N} \exp(-c'(N-5\delta N)). \quad (2)$$

- (3) In the second block, there is some H contained in $[-N, N]^d$ for which too few particles in the corresponding set $P_{H'}$ are H -good. This happens with probability at most

$$2^{d(2N+1)^d} \exp\left(-\frac{1}{16}(1+c)^{\delta N}\right). \quad (3)$$

- (4) In the third block, there is some H contained in $[-N, N]^d$ for which no H -counter visits every vertex at graph distance most $\ell - 1$ from the origin during steps $n + N - 4\delta N$ through N in H . This happens with probability at most

$$2^{d(2N+1)^d} \exp\left(-\frac{1}{6}(2d)^{-4\delta N}(1+c)^{N-4\delta N}\right). \quad (4)$$

We now finish the proof as follows. By the union bound, we conclude that the probability of the ℓ th stage being declared a failure is at most the sum of the estimates in (1), (2), (3) and (4); by the choice of δ , for fixed n , this sum tends to 0 as $N \rightarrow \infty$. We choose N large enough to both ensure that the bounds above hold and to make the probability of the ℓ th stage being declared a failure at most

$2^{-\ell}$, and as mentioned earlier, we set $T_\ell = N$. In the case of success, for each $H \subset \mathbb{Z}^d$ contained in the box $[-T_\ell, T_\ell]^d$, we pick an arbitrary particle that walked the chosen path above and take that particle to be p_H , whereas in the case of failure, we choose representative particles arbitrarily. The Borel–Cantelli lemma now implies that almost surely, there are only finitely many stages that fail, so for every subgraph, there is a particle that visits every reachable vertex infinitely often. \square

4. CONCLUSION

We have shown that a countable collection of independent simple random walks in two dimensions can almost surely be made transient by dropping to a suitable random two-dimensional subgraph; on the other hand, in any number of dimensions, a branching random walk is almost surely recurrent on every subgraph, in the sense that some branch is recurrent on each subgraph. Natural intermediate questions arise from considering the dynamic random walk $\mathcal{W}_d^{(t)}$ mentioned earlier. This object was introduced by Benjamini, Häggström, Peres and Steif [5], who showed that in three or four dimensions (i.e., when $d \in \{3, 4\}$), there is almost surely some time T such that $\mathcal{W}_d^{(T)}$ is recurrent; while, in five or more dimensions, almost surely the walk is transitive at all times. In two dimensions, Hoffman [11] showed that there is almost surely a time T such that the walk is $\mathcal{W}_2^{(T)}$ transitive; see also [3]. In the light of these facts, the following question seems of interest.

Problem 4.1. *Let $d \in \{2, 3, 4\}$, and let $\mathcal{W}_d^{(t)}$ be a dynamic random walk on \mathbb{Z}^d . Is there almost surely a (random) subgraph $H \subset \mathbb{Z}^d$ on which $\mathcal{W}_d^{(t)}$ is transitive for every time $t \geq 0$?*

It seems plausible that the answer is positive in four dimensions and negative in two dimensions; we are not prepared to offer a guess in three dimensions, however.

Another interesting question concerns paths. It is clear that for any path P in \mathbb{Z}^d through the origin, a random walk on P is almost surely recurrent. In Section 2, we noted that when one attempts to build a (random) path in two dimensions that greedily forces the random walk north and east, the resulting induced walk is almost surely recurrent; the induced walk exhibits some drift (compared to a fixed path) but only at the end points, and this is not enough to make it transient. This suggests the following question.

Problem 4.2. *Let \mathcal{W} be a simple random walk on \mathbb{Z}^d . Is it almost surely the case that \mathcal{W}_P is recurrent for every path $P \subset \mathbb{Z}^d$ through the origin?*

In a similar vein, as pointed out to us by one of the referees, one could also ask what happens if we replace the underlying graph \mathbb{Z}^2 by a strongly-recurrent graph in the sense of [13], such as, for example, the the Sierpinski carpet.

In the cases where we can force a simple random walk to be transient, what can we say about its escape velocity? In dimensions $d \geq 3$, it was shown in [8, 9] that a simple random walk on \mathbb{Z}^d escapes at a rate of about $\sqrt{n}/\log^{c_d+o(1)} n$. What can be said in our context?

Problem 4.3. *Fix $d \geq 2$. What is the supremum of α such that for a random walk \mathcal{W} on \mathbb{Z}^d , we can almost surely choose a subgraph $H \subset \mathbb{Z}^d$ such that the walk escapes to infinity at rate at least n^α ?*

As a first step towards this problem, it would already be interesting to know if we can get a random walk to escape with linear velocity in high dimensions.

Finally, a fundamental problem in this context, and one of our original motivations for treating the problem considered here, comes from the theory of universal traversal sequences. Call an infinite word \mathcal{Z} on the alphabet $\{\mathbf{x}, -\mathbf{x}, \mathbf{y}, -\mathbf{y}\}$ a *universal traversal sequence* for \mathbb{Z}^2 if \mathcal{Z}_G is recurrent for every $G \subset \mathbb{Z}^2$. The following basic question raised by Spink [18] remains wide open.

Problem 4.4. *Does there exist a universal traversal sequence for \mathbb{Z}^2 ?*

David and Tiba [7] recently found deterministic constructions of traversal sequences handling a reasonably large class of (but not all) subgraphs of \mathbb{Z}^2 . However, in general, the most efficient methods that we know of to construct universal traversal sequences all involve choosing a long enough traversal sequence at random; our main result rules out this standard construction on the square lattice. Either answer to the above existence question, positive or negative, would be very interesting.

ACKNOWLEDGEMENTS

The first author was supported by SNF grant 175505, the fourth author was partially supported by NSF Grant DMS-1800521, and the sixth author was supported by a Leverhulme Trust Research Fellowship.

The fourth author would like to thank Yuval Peres, Perla Sousi and Peter Winkler for interesting discussions.

Finally, we would like to thank the referees for their helpful comments, and for bringing some of the references mentioned in the introduction to our attention.

NOTE ADDED IN PROOF

After this manuscript was completed and circulated, it was brought to our attention that the existence of exceptional subgraphs for a single simple random walk has independently been established by Balister, Bollobás, Leader and Walters [14].

REFERENCES

1. O. Adelman, K. Burdzy, and R. Pemantle, *Sets avoided by Brownian motion*, Ann. Probab. **26** (1998), 429–464.
2. G. Amir, I. Benjamini, O. Gurel-Gurevich, and G. Kozma, *Random walk in changing environment*, Preprint, 1504.04870.
3. G. Amir and C. Hoffmann, *A special set of exceptional times for dynamical random walk on \mathbb{Z}^2* , Electron. J. Probab. **13** (2008), 1927–1951.
4. M. T. Barlow and E. A. Perkins, *Symmetric Markov chains in \mathbb{Z}^d : how fast can they move?*, Probab. Theory Related Fields **82** (1989), 95–108.
5. I. Benjamini, O. Häggström, Yuval Peres, and J. E. Steif, *Which properties of a random sequence are dynamically sensitive?*, Ann. Probab. **31** (2003), 1–34.
6. T. K. Carne, *A transmutation formula for Markov chains*, Bull. Sci. Math. **109** (1985), 399–405.
7. S. David and M. Tiba, *Solvability of mazes by blind robots*, Preprint, arXiv:1804.05439.
8. A. Dvoretzky and P. Erdős, *Some problems on random walk in space*, Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, University of California Press, 1951, pp. 353–367.
9. P. Erdős and S. J. Taylor, *Some problems concerning the structure of random walk paths*, Acta Sci. Hung. **11** (1960), 137–162.
10. M. Fukushima, *Basic properties of Brownian motion and a capacity on the Wiener space*, J. Math. Soc. Japan **36** (1984), 161–176.
11. C. Hoffman, *Recurrence of simple random walks on \mathbb{Z}^2 is dynamically sensitive*, ALEA **1** (2006), 35–45.

12. N. Kôno, *4-dimensional Brownian motion is recurrent with positive capacity*, Proc. Japan Acad. Ser. A Math. Sci. **60** (1984), 57–59.
13. T. Kumagai, *Random walks on disordered media and their scaling limits*, Lecture Notes in Mathematics, Springer, 2014, Lecture notes from the 40th Probability Summer School held in Saint-Flour, 2010, École d’Été de Probabilités de Saint-Flour.
14. I. Leader, Personal communication, July 2018.
15. M. Mitzenmacher and E. Upfal, *Probability and computing*, Cambridge University Press, 2017.
16. R. Peyre, *A probabilistic approach to Carne’s bound*, Potential Anal. **29** (2008), 17–36.
17. G. Pólya, *Über eine Aufgabe der Wahrscheinlichkeitsrechnung betreffend die Irrfahrt im Straßennetz*, Math. Ann. **84** (1921), 149–160.
18. H. Spink, Personal communication, December 2014.
19. N. Th. Varopoulos, *Long range estimates for Markov chains*, Bull. Sci. Math. **109** (1985), 225–252.

INSTITUTE OF MATHEMATICS, EPFL, STATION 8, 1015, LAUSANNE, SWITZERLAND

Email address: juhan.aru@epfl.ch

MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, ANDREW WILES BUILDING, RADCLIFFE OBSERVATORY QUARTER, WOODSTOCK ROAD, OXFORD OX2 6GG, UK

Email address: groenland@maths.ox.ac.uk

MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, ANDREW WILES BUILDING, RADCLIFFE OBSERVATORY QUARTER, WOODSTOCK ROAD, OXFORD OX2 6GG, UK

Email address: johnston@maths.ox.ac.uk

DEPARTMENT OF MATHEMATICS, RUTGERS UNIVERSITY, PISCATAWAY NJ 08854, USA

Email address: narayanan@math.rutgers.edu

MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, ANDREW WILES BUILDING, RADCLIFFE OBSERVATORY QUARTER, WOODSTOCK ROAD, OXFORD OX2 6GG, UK

Email address: roberts@maths.ox.ac.uk

MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, ANDREW WILES BUILDING, RADCLIFFE OBSERVATORY QUARTER, WOODSTOCK ROAD, OXFORD OX2 6GG, UK

Email address: scott@maths.ox.ac.uk