

Recurrence vs. Transience in \mathbb{Z}^2 and Higher Dimensions: A Walk Through Probability

Laura Kurtz, Mentored by Dr. Fehrman

May 2, 2025

Abstract

This paper explores the recurrence and transience properties of nearest neighbor random walks on integer lattices, with a focus on \mathbb{Z}^2 and higher dimensions. We begin with foundational concepts in probability theory and Markov processes while deriving key results for random walks on \mathbb{Z} , such as passing times and time of first return. We then extend the study to \mathbb{Z}^d , where we establish recurrence criteria and demonstrate the transition between recurrence in \mathbb{Z}^2 and transience in \mathbb{Z}^3 and beyond. Through analytical techniques such as the reflection principle and functional equations, this paper will highlight the dimensional dependence of random walk behavior.

Contents

1	Introduction	3
2	What's a Random Walk?	3
2.1	The Drunkard's Step: Bernoulli Foundations	3
2.2	Nearest Neighbor Random Walks	4
2.2.1	An Elementary Probability Theory Definition	4
2.2.2	A Definition Rooted in Markov Processes	4
2.3	Passage Times	5
2.3.1	via Reflection Principle	5
2.3.2	An Important Theorem	6
2.3.3	via Functional Equations	6
2.4	Time of First Return	7
2.4.1	The Fate of the Wanderer	7
2.4.2	The Paradox of the Symmetric Walk	7
2.4.3	The Non-Symmetric Case: A Swift (or Eternal!) Farewell	7
2.4.4	A Lesson in Persistence and Impatience	8
3	Recurrence Properties of the Random Walk	8
3.1	Extending the Random Walk to \mathbb{Z}^d : We Have d Different Walkers!	8
3.2	Finally ... A Recurrence Criterion	9
3.2.1	The Total Time at the Origin	9
3.2.2	The Dichotomy	9
3.2.3	Implications and Applications	10
4	In \mathbb{Z}^2 vs. Higher Dimensions	10
4.1	Recurrence in \mathbb{Z}^2 : Why Our Walker Always Comes Home	10
4.1.1	The Key to Recurrence Criterion in \mathbb{Z}^2	10
4.1.2	Symmetry: The Origin is the Most Likely Place!	10
4.1.3	The Decisive Estimate: Squeezing the Return Probability	10
4.2	Transience In \mathbb{Z}^3	11
4.2.1	Bounding Return Probabilities	11
4.2.2	A Key Idea: Coupling with Independent Walks	11
4.2.3	Breaking Down the Probability	11
4.2.4	The Magic of Concentration Inequalities	12
4.3	Beyond \mathbb{Z}^3 ...What Then?	12
5	Conclusions	13
6	Acknowledgements	14

1 Introduction

Imagine a drunk mathematician stumbling home from a bar. Each step they take is random—moving left or right with equal probability. Will they eventually return to the bar, or wander off forever? This is a key question in the random walk problem.

When we extend this idea to higher dimensions, such as a city grid (\mathbb{Z}^2) or a 3D lattice (\mathbb{Z}^3), the drunkard's fate changes dramatically based on the dimension. In 1D and 2D, the mathematician will always return to the starting point eventually, whereas in 3D and higher, there is a non-zero probability they will never return. We call these phenomena *recurrence* and *transience*, respectively. Even before and after this threshold, we see huge differences in the drunkard's behavior based on dimension! In 1D vs. 2D, while both dimensions guarantee eventual return, the expected time it will take to return differs. In 1D, returns happen relatively quickly, whereas in 2D, the wait is much, much longer (though still certain). The jump from 2D to 3D marks a sharp transition—guaranteed return vs. possible escape.

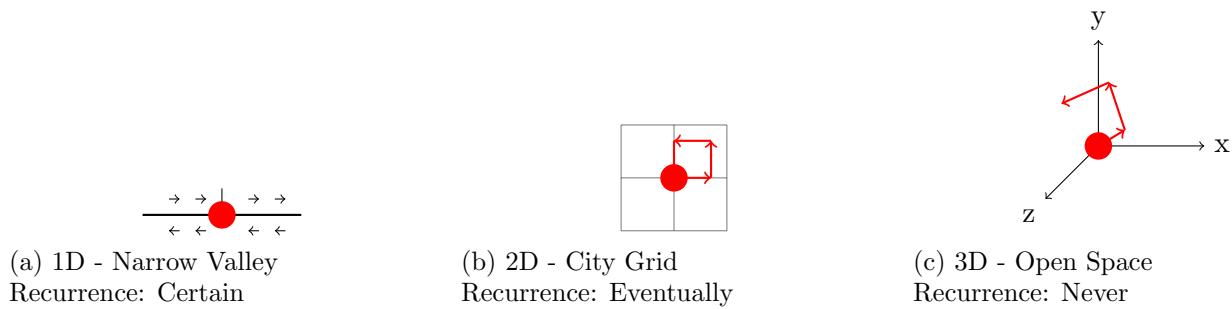


Figure 1: Illustration of random walk recurrence properties in different dimensions. In one dimension (a), the walker must return to the origin. In two dimensions (b), the walker will eventually return. In three dimensions (c), the walker will almost surely become permanently lost.

We now explore why dimension creates such a threshold.

2 What's a Random Walk?

2.1 The Drunkard's Step: Bernoulli Foundations

When we think of the word “walk,” we think of *steps*. In the case of the random walk, we're talking about either taking a step in one direction or another, with some probability. We will make this more rigorous later, but first, we need to talk about these “steps.”

Definition 1. A *Bernoulli random variable* is a random variable that takes two values.

We can think of many examples of Bernoulli random variables. A coin flip that takes values either “heads” or “tails” is a Bernoulli random variable, or a switch that has two choices, “on” or “off,” is also a Bernoulli random variable. But in the case of our walk, the two values are our step “forward” or “backward,” which we describe as “1” or “−1.” When we talk about these variables, we usually are referring to independent and identically distributed (“i.i.d.”) random variables.

Remark 1. The notation $\mathbb{P}(X \mid Y)$ means “The probability of X happening, given that Y has happened.”

Definition 2. Two random variables X and Y are said to be independent if $\mathbb{P}(X \mid Y) = \mathbb{P}(X)$, and $\mathbb{P}(Y \mid X) = \mathbb{P}(Y)$.

In the terms of our random walk, we're saying that at the fourth step we take, the probability we step forward is exactly the same as it was when we took our very first step. No matter how many steps we've taken, we always have the same probability of moving in a particular direction.

Definition 3. Two random variables X and Y are identically distributed if $\mathbb{P}(X \leq x) = \mathbb{P}(Y \leq x)$.

In our case, this just means that the odds of stepping forward is the same at any step.

2.2 Nearest Neighbor Random Walks

2.2.1 An Elementary Probability Theory Definition

We can now rigorously define our random walks in one dimension. Let $\{B_n \mid n \in \mathbb{Z}^+\}$ be a sequence of i.i.d. Bernoulli random variables taking values in $\{-1, 1\}$, with

$$P(B_n = 1) = p \quad \text{and} \quad P(B_n = -1) = q = 1 - p, \quad (1)$$

for some $p \in (0, 1)$. B_n is the “step” taken at time n : one step to the right (if it's 1), or one step to the left (if it's -1). This is the “step” we discussed in our preliminaries! Now, fix a sequence $E = (\epsilon_1, \epsilon_2, \dots, \epsilon_n) \in \{-1, 1\}^n$. The probability of taking exactly this sequence of steps is given by:

$$P(B_1 = \epsilon_1, \dots, B_n = \epsilon_n) = p^{N(E)} q^{n-N(E)}, \quad (2)$$

where $N(E)$ is the number of steps to the right in the sequence:

$$N(E) := |\{m \mid \epsilon_m = 1\}| = \frac{n + \sum_{m=1}^n \epsilon_m}{2}. \quad (3)$$

This formula captures the likelihood of any specific walk path based on how many steps move right versus left. Then we accumulate the total displacement from the origin:

$$X_0 = 0, \quad \text{and for } n \geq 1, \quad X_n = \sum_{m=1}^n B_m. \quad (4)$$

Definition 4. The collection $\{X_n \mid n \in \mathbb{Z}\}$ is called a nearest neighbor random walk on \mathbb{Z} .

We can translate this definition into the intuition we developed earlier—we're essentially collecting steps, either right or left, as we see in equation (4). However, this definition relies mainly on fundamental probability theory, but this process is a very dynamic one that could be modeled in a more Markovian way.

2.2.2 A Definition Rooted in Markov Processes

Definition 5. The Markov property, or “lack of memory property,” is the property such that

$$\mathbb{P}(X_{n+1} \mid X_n, X_{n-1}, \dots) = \mathbb{P}(X_{n+1} \mid X_n) \quad (5)$$

We can see that our process we defined in Definition 4 is a Markovian process—at a given step n , the probability we step forward or backward is the same no matter what time n we are at. Therefore, we could replace equation (4) with the following:

$$\mathbb{P}(X_0 = 0) = 1, \quad \text{and} \quad \mathbb{P}(X_n - X_{n-1} = \epsilon \mid X_0, \dots, X_{n-1}) = \begin{cases} p & \text{if } \epsilon = 1 \\ q & \text{if } \epsilon = -1 \end{cases} \quad (6)$$

Note that this way of defining our random walk captures the dynamicity of the process, since we are saying that we start at 0 from time $n = 0$ and proceeding either a step forward with probability p or a step backward with probability q .

2.3 Passage Times

To discuss recurrence, we need more information about our process. To find out when or if our process will return to 0, we first need to develop a tool to tell us when our process first passes a certain point.

Definition 6. *The first passage time to a point $a \in \mathbb{Z}$, ζ_a , is given by*

$$\zeta_a = \inf\{n \geq 1 \mid X_n = a\} \quad (7)$$

where $\zeta_a = \infty$ if $X_n \neq a$ for any $n \geq 1$.

Equivalently, note that what we're truly looking for is an expression to calculate $\mathbb{P}(\zeta_a = n)$, or the probability that we pass a at time n . We can derive a way to find this in two methods, one based on section 2.2.1, and one based in section 2.2.2.

2.3.1 via Reflection Principle

To find the first passage time based on the probabilistic definition of the random walk, let $a \in \mathbb{Z}^+$, and suppose $n \in \mathbb{Z}^+$ has the same parity as a .^a Now observe that

$$\mathbb{P}(\zeta_a = n) = \mathbb{P}(X_n = a \text{ and } \zeta_a > n - 1) = p\mathbb{P}(\zeta_a > n - 1 \text{ and } X_{n-1} = a - 1)^{(*)}. \quad (8)$$

Hence, if we compute $(*)$, we have computed $\mathbb{P}(\zeta_1 = n)$. To this end, note that for any $E \in \{-1, 1\}^{n-1}$ with $X_{n-1} = a - 1$, there exists a bijection between paths that stay strictly above a and those that stay strictly below a up to time $n - 1$. The *reflection principle* allows us to relate these probabilities through symmetry. For $a > 0$, we have:

$$\mathbb{P}(\zeta_a \leq n) = 2\mathbb{P}(X_n \geq a) - \mathbb{P}(X_n = a). \quad (9)$$

This follows by considering that for every path reaching a at time $k \leq n$ and ending above a , there's a corresponding reflected path ending below a . The correction term accounts for paths ending exactly at a . For the special case $a = 1$, we can derive an exact probability:

$$\mathbb{P}(\zeta_1 = 2m - 1) = \frac{1}{2m - 1} \binom{2m - 1}{m} 2^{-(2m-1)}. \quad (10)$$

The asymptotic behavior follows from Stirling's approximation:

$$\mathbb{P}(\zeta_1 = n) \sim \frac{1}{\sqrt{2\pi}} n^{-3/2} \quad \text{as } n \rightarrow \infty. \quad (11)$$

This $n^{-3/2}$ decay is characteristic of first passage times in one-dimensional random walks and explains why the expected return time to the origin is infinite despite the walk being recurrent.

^a X_n is always odd if n is odd, and similarly if n is even.

2.3.2 An Important Theorem

Theorem 1 (Probability of Ever Hitting Level a). *For a random walk with step probabilities p, q ($p + q = 1$), the probability of eventually hitting level $a \in \mathbb{Z} - \{0\}$ is:*

$$\mathbb{P}(\zeta_a < \infty) = \begin{cases} 1 & \text{if } a \in \mathbb{Z}^+ \text{ and } p \geq q, \text{ or } -a \in \mathbb{Z}^+ \text{ and } p \leq q, \\ \left(\frac{p}{q}\right)^{|a|} & \text{otherwise.} \end{cases}$$

Proof. Base Case: For $a = 1$, the generating function yields:

$$\mathbb{E}[s^{\zeta_1}] = \frac{1 - \sqrt{1 - 4pqs^2}}{2qs} \implies \mathbb{P}(\zeta_1 < \infty) = \min\left(1, \frac{p}{q}\right). \quad (12)$$

For $a > 0$, by the strong Markov property:

$$\mathbb{P}(\zeta_{a+1} < \infty) = \mathbb{P}(\zeta_1 < \infty) \cdot \mathbb{P}(\zeta_a < \infty).$$

Induction gives $\mathbb{P}(\zeta_a < \infty) = \mathbb{P}(\zeta_1 < \infty)^a$. The $a < 0$ case follows by symmetry (swap p and q). Substituting $\mathbb{P}(\zeta_1 < \infty) = \min(1, p/q)$ yields the result. \odot

2.3.3 via Functional Equations

There has to be a better way! We continue our discussion of passage times with a less computational derivation- the reflection principle gives us a geometric insights, but to truly understand what is happening, we need generating functions. Our goal is now to derive the distribution of ζ_a in a single expression in a series of recursive relations.

For $a \in \mathbb{Z} - \{0\}$ and $s \in (-1, 1)$, define the generating function:

$$u_a(s) = \mathbb{E}[s^{\zeta_a}]. \quad (13)$$

We first make an observation: reaching $a + 1$ requires that we reach a , and then, independently of the past, make an additional step of 1. Given $a \in \mathbb{Z}^+$, we apply the strong Markov property^b to obtain the recursive relation:

$$u_{a+1}(s) = \sum_{m=1}^{\infty} s^m \mathbb{E}[s^{\zeta_1 \circ \Sigma^m}, \zeta_a = m] = u_a(s)u_1(s). \quad (14)$$

We note that this propagates upwards through the integers. For negative values, we note a symmetrical structure:

$$u_{a-1}(s) = u_a(s)u_{-1}(s). \quad (15)$$

These relations yield the general form:

$$u_a(s) = u_{\text{sgn}(a)}(s)^{|a|} \quad \text{for } a \in \mathbb{Z} - \{0\}. \quad (16)$$

So now the problem reduces itself to solving for $u_1(s)$. We can condition our first step (X_1):

$$u_1(s) = \mathbb{E}[s^{\zeta_1}, X_1 = 1] + \mathbb{E}[s^{\zeta_1}, X_1 = -1] = ps + qsu_1(s)^2. \quad (17)$$

^bThe strong Markov property differs from the Markov property only in that the strong Markov property holds at any stopping time.

where the term $u_i(s)^2$ emerges from the two possible states after the initial step. But how do we choose which root to use?

Since $u_1(s)$ is a probability generating function, it must satisfy $u_1(1^-) = 1$ and also remain bounded for $|s| < 1$. Thus, our hand is forced by $\sqrt{1 - 4pqs^2}$ from equation (12):

$$u_1(s) = \frac{1 - \sqrt{1 - 4pqs^2}}{2qs} \quad (\text{choosing the root with } u_1(s) \leq 1). \quad (18)$$

Combining these results, we obtain the closed-form expression for all $a \neq 0$:

$$\mathbb{E}[s^{\zeta_a}] = \begin{cases} \left(\frac{1 - \sqrt{1 - 4pqs^2}}{2qs} \right)^a & \text{if } a \in \mathbb{Z}^+ \\ \left(\frac{1 - \sqrt{1 - 4pqs^2}}{2ps} \right)^{-a} & \text{if } -a \in \mathbb{Z}^+ \end{cases} \quad \text{for } |s| < 1, \quad (19)$$

which is our universal formula for passage times to *any* $a \neq 0$.

2.4 Time of First Return

Now we ask our key question: when will our random walk come back? The time of first return, ρ_0 , tells us whether the walker will come back, and if so, how long will it take?

2.4.1 The Fate of the Wanderer

When a walker stands at the origin of \mathbb{Z} , taking steps to the right with probability p and to the left with probability $q = 1 - p$, the question of *recurrence*, whether the walker is fated to return, is dependent on the balance between p and q . We learn that the probability of eventual return is

$$\mathbb{P}(\rho_0 < \infty) = 2(p \wedge q)^c \quad (20)$$

This result reveals that the walker is certain to return only in the symmetric case, where $p = q = \frac{1}{2}$.

2.4.2 The Paradox of the Symmetric Walk

When the walk is symmetric, we have shown that our walker will come back, but there's a twist! The expected time to return, $\mathbb{E}[\rho_0]$, is infinite. This is an interesting result- the walker will always return, but on average, we'd have to wait forever to see it happen. The generating function

$$\mathbb{E}[s^{\rho_0}] = 1 - \sqrt{1 - 4pqs^2} \quad \text{for } |s| < 1 \quad (21)$$

shows us this delicate balance, as it approaches 1 (*has 100% probability*) as $s \rightarrow 1$.

2.4.3 The Non-Symmetric Case: A Swift (or Eternal!) Farewell

For biased walks ($p \neq q$), the story is much different. With probability $1 - 2(p \wedge q)$, the walker drifts away forever. But if they do return, they do so relatively quickly. Equation (22) quantifies this:

$$\mathbb{E}[\rho_0 \mid \rho_0 < \infty] = \frac{2p \vee q}{|p - q|} = 1 + \frac{1}{|p - q|} \quad (22)$$

Here, the expected return time grows as the bias $|p - q|$ shrinks, reflecting the increasing difficulty of returning when the walk is *almost* symmetric.

^c $p \wedge q$ denotes a minimum, and $p \vee q$ denotes a maximum.

2.4.4 A Lesson in Persistence and Impatience

The picture painted by these results is clear. Symmetry ensures return but demands infinite patience. Asymmetry offers a quicker resolution, which is either a swift return or a permanent departure.

3 Recurrence Properties of the Random Walk

3.1 Extending the Random Walk to \mathbb{Z}^d : We Have d Different Walkers!

To generalize nearest neighbor random walks to \mathbb{Z}^d , we first identify the set of nearest neighbors of the origin in \mathbb{Z}^d as the $2d$ points where $(d-1)$ coordinates are 0 and the remaining coordinate is in $\{-1, 1\}$. We then consider independent, identically distributed N_d -valued random variables B_1, B_2, \dots ,^d where N_d is the set of nearest neighbors in \mathbb{Z}^d . More directly— we have a choice of any of the nearest neighbors. In two dimensions, our nearest neighbors to the origin are $(1, 0)$, $(-1, 0)$, $(0, 1)$, and $(0, -1)$. In d dimensions, this would be the $2d$ points where one coordinate is ± 1 and the rest are 0. Figures 1 and 2 illustrate the $2d$ nearest neighbors for $d = 2, 3$:




Figure 2: 2D Grid (4 Nearest Neighbors)




Figure 3: 3D Space (6 Nearest Neighbors)

^dThe existence of B_n 's can be seen as a consequence: let $\{U_n \mid n \in \mathbb{Z}^+\}$ be a family of mutually independent random variables which are uniformly distributed on $[0,1)$, and let there be an $F : [0,1) \rightarrow N_d$ (defined as necessary) and then set $B_n = F(U_n)$. See [3] Chapter 6.

3.2 Finally ... A Recurrence Criterion

We now attempt to discover a rigorous criterion for whether or not a walk will return to its starting point.

3.2.1 The Total Time at the Origin

Consider a random walk $\{X_n : n \geq 0\}$ starting at the origin ($X_0 = 0$). Let T_0 denote the total time the walk spends at the origin:

$$T_0 = \sum_{n=0}^{\infty} \mathbf{1}_{\{0\}}(X_n). \quad (23)$$

Since the walk begins at 0, T_0 is at least 1. For $n \geq 1$, the event $\{T_0 > n\}$ corresponds to the walk returning to the origin at least n times. This connects T_0 to the return times $\rho_0^{(n)}$, defined as the time of the n -th return to 0. Specifically,

$$T_0 > n \iff \rho_0^{(n)} < \infty. \quad (24)$$

Using this relationship, we derive the expected value of T_0 :

$$\mathbb{E}[T_0] = \sum_{n=0}^{\infty} \mathbb{P}(T_0 > n) = 1 + \sum_{n=1}^{\infty} \mathbb{P}(\rho_0^{(n)} < \infty). \quad (25)$$

A key observation is that the probability of returning n times factors into the n -th power of the probability of returning once:

$$\mathbb{P}(\rho_0^{(n)} < \infty) = \mathbb{P}(\rho_0 < \infty)^n. \quad (26)$$

This multiplicative property arises because each return time $\rho_0^{(n)}$ depends only on the steps *after* the previous return, and these segments are independent. Substituting this into the expectation yields:

$$\mathbb{E}[T_0] = 1 + \sum_{n=1}^{\infty} \mathbb{P}(\rho_0 < \infty)^n = \frac{1}{1 - \mathbb{P}(\rho_0 < \infty)} = \frac{1}{\mathbb{P}(\rho_0 = \infty)}. \quad (27)$$

This formula tells us that the expected time spent at the origin is finite *if and only if* there is a positive chance the walk never returns.

3.2.2 The Dichotomy

The behavior of T_0 is interesting:

1. **Finite Returns:** If $\mathbb{P}(T_0 < \infty) > 0$, then $\mathbb{E}[T_0] < \infty$. This happens when $\mathbb{P}(\rho_0 < \infty) < 1$, meaning the walk has a chance to escape forever.
2. **Infinite Returns:** If $\mathbb{E}[T_0] = \infty$, then $\mathbb{P}(T_0 = \infty) = 1$. Here, $\mathbb{P}(\rho_0 < \infty) = 1$, so the walk returns infinitely often with certainty.

This dichotomy mirrors the broader classification of random walks as *transient* or *recurrent*.

Definition 7. A walk is *transient* if it escapes to infinity. More precisely, if $\mathbb{E}[T_0] < \infty$ (finite expected visits).

Definition 8. A walk is *recurrent* if it revisits the origin infinitely often. More precisely, if $\mathbb{E}[T_0] = \infty$ (positive probability of infinite visits).

3.2.3 Implications and Applications

If the walk has a directional bias (e.g., $\mathbb{E}[B_1] \neq 0$), the Strong Law of Large Numbers^e ensures $|X_n| \rightarrow \infty$ almost surely. This forces $\mathbb{P}(\rho_0 = \infty) > 0$, so $\mathbb{E}[T_0] < \infty$ —the walk is transient. For unbiased walks (e.g., simple symmetric random walk), the return probability $\mathbb{P}(\rho_0 < \infty)$ often equals 1, leading to recurrence.

4 In \mathbb{Z}^2 vs. Higher Dimensions

4.1 Recurrence in \mathbb{Z}^2 : Why Our Walker Always Comes Home

We have a surprising result in \mathbb{Z}^2 : recurrence. This means that, no matter how far the walker strays, they will always return to the origin infinitely often, with certainty. The one-dimensional case (\mathbb{Z}) is intuitive, but the two-dimensional result is not: how does the walker, with so many more directions to escape, never truly get away? The answer lies in the delicate balance of symmetry and the slow decay of return probabilities.

4.1.1 The Key to Recurrence Criterion in \mathbb{Z}^2

We begin by examining the connection between recurrence and return probabilities. For a random walk $\{\mathbf{X}_n\}$ on \mathbb{Z}^d , let T_0 count the number of times the walk visits the origin. Then:

$$\mathbb{E}[T_0] = \sum_{n=0}^{\infty} \mathbb{P}(\mathbf{X}_n = 0). \quad (28)$$

This is deceptively straightforward: the expected number of returns is just the sum of the probabilities of being at the origin at each step. But combined with our recurrence criterion ($\{\mathbf{X}_n\}$ is recurrent *if and only if* $\mathbb{E}[T_0] = \infty$), it becomes much more useful: the walk is recurrent precisely when this sum diverges.

4.1.2 Symmetry: The Origin is the Most Likely Place!

Let $\mathbf{k} \in \mathbb{Z}^2$ and observe that:

$$\mathbb{P}(\mathbf{X}_{2n} = \mathbf{k}) \leq \sum_{\ell \in \mathbb{Z}^d} \mathbb{P}(\mathbf{X}_n = \ell)^2 = \mathbb{P}(\mathbf{X}_{2n} = 0). \quad (29)$$

The inequality arises from Schwarz's inequality^f, but because of symmetry, we actually get equality: $\mathbb{P}(\mathbf{X}_n = \ell) = \mathbb{P}(\mathbf{X}_n = -\ell)$. This means the walker, no matter how far they roam, is always most likely to be at the origin at even times.

4.1.3 The Decisive Estimate: Squeezing the Return Probability

The crux of the argument lies in bounding $\mathbb{P}(\mathbf{X}_{2n} = 0)$. For \mathbb{Z}^2 , symmetry and the walk's second-moment properties yield:

$$\mathbb{P}(\mathbf{X}_{2n} = 0) \geq \frac{1}{8n + 2}. \quad (30)$$

^e“A mathematical law that states that the average of the results obtained from a large number of independent random samples converges to the true value, if it exists.” [1]

^fsee [1]

This estimate comes from two observations:

1. The walk doesn't spread too fast: Markov's inequality shows $\mathbb{P}(|\mathbf{X}_{2n}| \geq 2\sqrt{n}) \leq \frac{1}{2}$, so the walker stays within $2\sqrt{n}$ of the origin half the time.
2. Counting lattice points: There are roughly $4n$ points within this range, and symmetry ensures the origin's probability dominates.

Summing these probabilities gives a harmonic series $\sum_n \frac{1}{n}$, which diverges. Thus, $\mathbb{E}[T_0] = \infty$, and the walk is recurrent.

4.2 Transience In \mathbb{Z}^3

4.2.1 Bounding Return Probabilities

Our goal is to prove that in dimensions $d \geq 3$, the probability of returning to the origin after $2n$ steps decays quickly enough that the total expected number of returns is finite. Mathematically, we want to show:

$$\sum_{n=0}^{\infty} \mathbb{P}(\mathbf{X}_{2n} = \mathbf{0}) < \infty. \quad (31)$$

To do this, we need an upper bound on $\mathbb{P}(\mathbf{X}_{2n} = \mathbf{0})$ that captures the correct decay rate.

4.2.2 A Key Idea: Coupling with Independent Walks

The challenge is that the coordinates of \mathbf{X}_n (the random walk in \mathbb{Z}^d) are not independent—each step changes only one coordinate. However, we can cleverly relate \mathbf{X}_n to a collection of independent one-dimensional random walks $\{X_{i,n}\}$ using a random index selection process:

1. Auxiliary randomness: Introduce a sequence $\{I_n\}$ of uniform random indices in $\{1, \dots, d\}$, independent of everything else.
2. Counting coordinate updates: For each coordinate i , let $N_{i,n}$ count how many times it has been selected up to step n .
3. Constructing \mathbf{Y}_n : Define a new process $\mathbf{Y}_n = (X_{1,N_{1,n}}, \dots, X_{d,N_{d,n}})$.

By construction, \mathbf{Y}_n has the same distribution as \mathbf{X}_n , but it is now expressed in terms of independent walks, each moving only when their coordinate is chosen.

4.2.3 Breaking Down the Probability

The probability $\mathbb{P}(\mathbf{X}_{2n} = \mathbf{0})$ splits into two cases:

1. All coordinates are updated “enough” times: If each $N_{i,2n} \geq \frac{n}{d}$, we can use known bounds for independent walks.
2. Some coordinates are neglected: If any $N_{i,2n} < \frac{n}{d}$, we control this scenario using concentration inequalities.

The dominant term comes from the first case, where we exploit the independence to get a decay like $n^{-d/2}$. The second case turns out to be negligible thanks to sharp tail bounds on the $N_{i,n}$.

4.2.4 The Magic of Concentration Inequalities

The counts $N_{i,n}$ are Binomial random variables with mean $\frac{n}{d}$. Using moment-generating functions and Taylor expansions, we derive an exponential bound:

$$\mathbb{P}\left(N_{i,2n} \leq \frac{n}{d}\right) \leq e^{-2nR^2}, \quad (32)$$

which decays much faster than $n^{-d/2}$. This ensures that the “neglected coordinates” scenario doesn’t ruin our estimate. Combining, we prove:

$$\mathbb{P}(\mathbf{X}_{2n} = \mathbf{0}) \leq A(d) n^{-d/2}, \quad (33)$$

which is summable for $d \geq 3$. Thus, the random walk is transient: it has a finite expected number of returns to the origin, meaning it eventually escapes forever.

4.3 Beyond \mathbb{Z}^3 ...What Then?

The pattern continues for higher dimensions: return probabilities satisfy $\mathbb{P}(\mathbf{X}_{2n} = \mathbf{0}) \asymp n^{-d/2}$.^g The series converges iff $d \geq 3$:

$$\sum_{n=1}^{\infty} n^{-d/2} \begin{cases} = \infty & \text{for } d \leq 2 \\ < \infty & \text{for } d \geq 3 \end{cases}. \quad (34)$$

^g \asymp is notation for “converging asymptotically.”

5 Conclusions

Imagine our tipsy mathematician, stumbling through the infinite grid of \mathbb{Z}^2 , doomed to forever retrace their steps back to the bar no matter how far they wander. In one and two dimensions, they always come home, even if it takes an eternity. But the moment they stumble into the third dimension (maybe our friend has learned to fly!), freedom awaits! With a whole new axis to explore, they might just vanish into the void, never to be seen again. Two dimensions: You're stuck in an endless loop. Three: The world is your oyster.

But, maybe we aren't stumbling through a city grid. Maybe we're a stock price, or a particle moving through space. The question of whether our not our walk will return home is a useful one to answer. And as we've shown, whether or not the walk is symmetric, or how many dimensions you're moving through, is a key question in whether or not we'll return home.

References

- [1] M. Dekking. *A Modern Introduction to Probability and Statistics*. Graduate Texts in Mathematics. Springer, 2005.
- [2] W. Monroe. Bernoulli and binomial random variables. *Lecture Notes whilst at Stanford University*, 2017.
- [3] D. W. Stroock. *An Introduction to Markov Processes*. Graduate Texts in Mathematics. Springer, 2005.

6 Acknowledgements

Thank you to:

- Dr. Ben Fehrman for mentoring this project and providing significant help and advice, and for providing edits
- Dr. Stephen Shipman for providing edits
- Chris Bunting for being a helpful first reader
- The late Dr. Daniel W. Stroock for reminding me that math writing can and should be fun.