Using Bijections and Lattice Paths to Enumerate Tilings of the Aztec Diamond

Mark: 85

2017/18

1 Introduction

In this essay, we are aiming to solve the problem of counting domino tilings of a region called the Aztec Diamond of order n, for each positive integer n. Although the result has consequences in the broader study of tilings[1], this essay aims to use the Aztec Diamond theorem as a means of exploring an identity common across combinatorics, which we introduce in Section 2.

In [1], Ardila and Stanley construct an Aztec Diamond of order n, denoted AD_n, by stacking unit squares in rows of length $2, 4, \ldots, 2n, 2n, \ldots, 4, 2$ on top of each other. Figure 1 shows the way we construct a given size of this region. We define a *domino* as a region comprised of the union of two adjacent unit squares, and a *domino tiling* of a given region as a way to overlay dominoes onto the region, with no overlaps or gaps. Two such tilings of AD_3 are given in Figure 2.

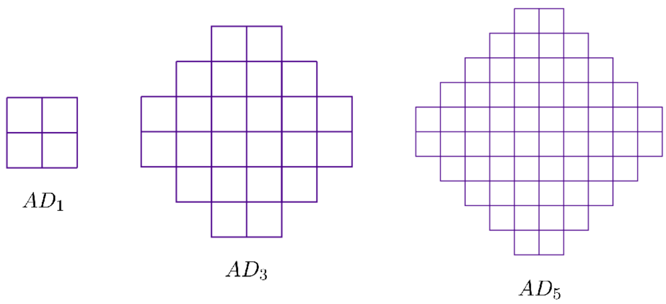

Figure 1: Some sizes of the Aztec Diamond.

Formally, we define our Aztec Diamond as following:

[1]Some examples of this are further discussed in Section 6.

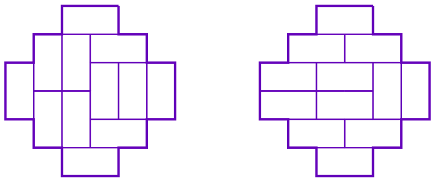

Figure 2: Two tilings of AD_3. We will show there is 32 in total.

Definition 1.1. The *Aztec Diamond* of order n, denoted by AD_n, is the union of all the unit squares with integral corners (x,y) satisfying $|x|+|y| \leq n+1$.

We will be proving in this essay that the number of domino tilings of the nth Aztec Diamond is $2^{n(n+1)/2}$. Interest in proving the formula began due to the elegance of this number. Compare this formula with the number of domino tilings of a $2m \times 2n$ rectangle[2],

$$4^{mn} \prod_{j=1}^{m} \prod_{k=1}^{n} \left(\cos^2 \frac{j\pi}{2m+1} + \cos^2 \frac{k\pi}{2n+1} \right).$$

As highlighted by Ardilla and Stanley [1], the comparison underlines the surprising simplicity of the number of Aztec Diamond domino tilings. However, no proof of this result is as simple as we would imagine. This essay will follow the method used by Sen-Peng Eu and Tung-Shan Fu in [5] in 2005, and along the way aims to communicate how else we may use such a method.

2 An Introduction to Involutions

Suppose we were given the sum $\sum_{k=1}^{2n} (-1)^k$, and asked to find its value. One natural method would be state that there are n positive terms in the sum, and n negative terms, so we can match pairs of terms up to say that $\sum_{k=1}^{2n} (-1)^k = \sum_{k=1}^{n} 0 = 0$.

Similarly, suppose we were asked to prove the following, as given by [2] and [7]:

Example 2.1. For all $n \geq 1$, find the value of

$$\sum_{k=0}^{n} (-1)^k \binom{n}{k}.$$

[2]This is a special case of Equation 13 in Kasteleyn's paper, [6], examining ways of arranging dimers (molecules made up of two identical, simpler molecules) in chemistry.

We can prove this algebraically, using the binomial theorem to say that $\sum_{k=0}^{n}(-1)^k\binom{n}{k} = \sum_{k=0}^{n}\binom{n}{k}(-1)^k 1^{n-k} = (-1+1)^n = 0$.

However, we could also use a combinatorial method. If we let X be the set of all subsets of $\{1, 2, \ldots, n\}$, then there are $\binom{n}{k}$ subsets of size k, so showing our identity sums to zero is showing that the number of subsets, X_E, with an even number of elements is equal to the number, X_O, with an odd number of elements.

Given this, we could try and find a bijection between X_E and X_O. Let us define a function ϕ which for every subset S in X,

$$\phi(S) = \begin{cases} S \setminus \{1\} & \text{if } 1 \in S \\ S \cup \{1\} & \text{if } 1 \notin S. \end{cases}$$

This means that if a subset contains 1, ϕ removes it, and if it does not contain 1, ϕ adjoins it. If we recognise that ϕ sends odd subsets to even subsets and even subsets to odd, then we have a bijection between X_E and X_O, because the function is its own inverse. This implies that $|X_E| = |X_O|$, and so since every element of X is either odd or even, $\sum_{k=0}^{n}(-1)^k\binom{n}{k} = |X_E| - |X_O| = 0$. □

In this chapter we will be formalising this idea. The method involves describing a set of objects, where each object has a positive or negative sign, and focusing on summing the signs of each object. Then, if we can find pairs comprised of one positive and one negative object, their signs will cancel, and we need only consider the objects that do not belong to a pair.

Definition 2.2. An *involution* on a finite set X is a function $I : X \to X$ such that $I \circ I = id_X$.

Such a function must be made up of transpositions and fixed points, where we define the *fixed point set* of an involution I as $\mathrm{Fix}(I) = \{x \in X : I(x) = x\}$.

Definition 2.3. An involution $I : X \to X$ is *sign-reversing* if, given a function sgn: $X \to \{+1, -1\}$ that assigns a sign to every object in X, $\mathrm{sgn}(I(x)) = -\mathrm{sgn}(x)$ for all $x \in X \setminus \mathrm{Fix}(I)$.

For shorthand in our essay, we shall separate a set X with associated sign function sgn: $X \to \{+1, -1\}$ into two sets: $X^+ = \{x \in X : \mathrm{sgn}(x) = +1\}$, and $X^- = \{x \in X : \mathrm{sgn}(x) = -1\}$. Then $X = X^+ \cup X^-$ and

$$\sum_{x \in X} \mathrm{sgn}(x) = |X^+| - |X^-|.$$

For example, if we let X be the set of integers from $\{1, \ldots, 2n\}$, and have associated sign function sgn: $X \to \{+1, -1\}$, such that for x in X, $\mathrm{sgn}(x) =$

$(-1)^x$, then X^+ is the set even integers between 1 and $2n$, and X^- is the set of odd integers in this range.

We can now show formally that this method works, using the method of [7, p. 153].

Theorem 2.4 (Involution Theorem)**.** *Given a finite set X of signed objects and a sign-reversing involution I on X,*

$$\sum_{x \in X} \operatorname{sgn}(x) = \sum_{x \in \operatorname{Fix}(I)} \operatorname{sgn}(x).$$

Proof. Let $I : X \to X$ be an involution. If we restrict I to $X^+ \setminus \operatorname{Fix}(I)$, since I is sign-reversing, we can obtain a function I^+ from $X^+ \setminus \operatorname{Fix}(I)$ to $X^- \setminus \operatorname{Fix}(I)$, and similarly we get $I^- : X^- \setminus \operatorname{Fix}(I) \to X^+ \setminus \operatorname{Fix}(I)$. Note that these two functions are the inverse of each other, and hence bijections. This means $|X^+ \setminus \operatorname{Fix}(I)| = |X^- \setminus \operatorname{Fix}(I)|$. Then

$$\sum_{x \in X} \operatorname{sgn}(x) = \sum_{x \in X^+ \setminus \operatorname{Fix}(I)} \operatorname{sgn}(x) + \sum_{x \in X^- \setminus \operatorname{Fix}(I)} \operatorname{sgn}(x) + \sum_{x \in \operatorname{Fix}(I)} \operatorname{sgn}(x)$$

$$= |X^+ \setminus \operatorname{Fix}(I)| - |X^- \setminus \operatorname{Fix}(I)| + \sum_{x \in \operatorname{Fix}(I)} \operatorname{sgn}(x) = \sum_{x \in \operatorname{Fix}(I)} \operatorname{sgn}(x). \qquad \square$$

This is referred to as the *Description, Involution, Exceptions (D.I.E)* method [2]. We *describe* a set of objects which is counted by the sum when we do not consider sign, and assign each object a sign. We find a sign-changing *involution* from objects in the sum counted positively to those counted negatively. Finally, we look for *exceptions*, where the involution is undefined. Counting these exceptions and noting their sign gives the value of the sum.

In the next chapter we will be applying this idea in a new context, to show how sign-changing involutions can be used to find solutions to more abstract problems.

3 Lattice Paths and The Reflection Principle

In this chapter we will be looking at how problems can be represented as lattice paths. A *lattice path* is a path formed by line segments between integer points in the plane. We restrict the possible steps of a lattice path to three types — *up steps* of $(1, 1)$, *down steps* of $(1, -1)$ and *level steps* of $(2, 0)$. We denote these steps by **U**, **D** and **L** respectively. Sometimes we restrict lattice paths to only up steps and down steps, depending on the problem being represented. In Figure 3, we show some such lattice paths.

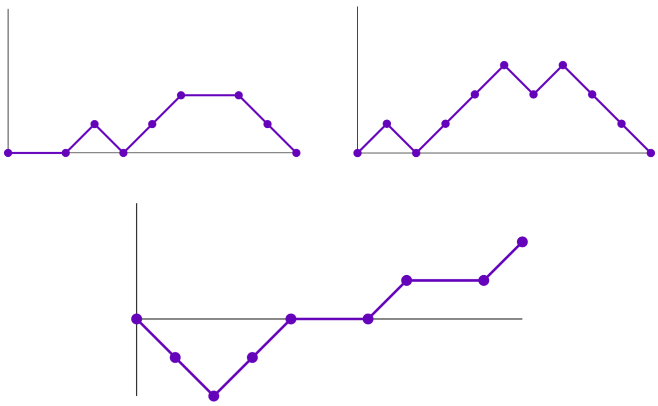

Figure 3: Three examples showing how a lattice path may look.

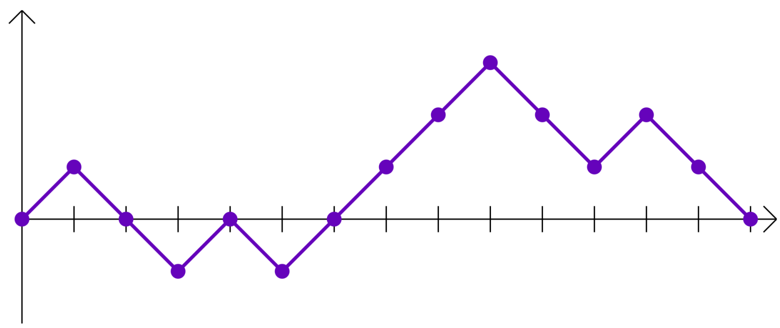

Figure 4: A lattice path representing ABBABAAAABBABB.

Let us use our new involution method and lattice path interpretation to solve a version of Bertrand's Ballot Problem, named after Joseph Bertrand who posed it in 1887 [3]. We will be following the proof found in [10, pp. 130–132].

Theorem 3.1. *Suppose A and B are the two parties running in an election, and both parties receive n votes. The number of ways to count the votes such that A is never behind B is $\frac{1}{n+1}\binom{2n}{n}$.*

Proof. We shall represent this problem in the form of a lattice path. If we let an up step of $(1,1)$ represent a vote for A, and a down step of $(1,-1)$ represent a vote for B, then any sequence of up steps and down steps starting at $(0,0)$ and ending at $(2n,0)$ represents a way of counting $2n$ votes in which A and B both receive n votes (since there is the same number of up and down steps). For example, Figure 4 shows how we would represent 14 votes being counted in the order $ABBABAAAABBABB$.

Since any such sequence can be represented as a lattice path, and any lattice path from $(0,0)$ to $(2n,0)$ can be translated to a unique sequence of votes, there is a bijection between all sequences of $2n$ votes and the lattice path representations. Note that a sequence in which A is never behind B, corresponds to a lattice path which never descends below the x-axis.

Before we define our signed set, let us displace our set of lattice path one unit vertically, so each lattice path goes from $(0,1)$ to $(2n,1)$. Now there is a bijection between sequences in which A is never behind B, and lattice paths which do not cross or touch the x-axis.

Define our set of objects with positive sign (denoted S^+) as the set of all paths from $(0,1)$ to $(2n,1)$. Note that precisely n of these steps must be up steps, so there are $\binom{2n}{n}$ lattice paths in S^+. Let our set of objects with negative sign, S^-, be the set of all lattice paths from $(-1,0)$ to $(2n,1)$. Here each member of S^- must have 2 more up steps than down steps, hence there must be precisely $n+1$ up steps, giving $\binom{2n}{n+1}$ members of S^-.

Now if we let $S = S^+ \cup S^-$, then

$$\sum_{P \in S} \operatorname{sgn}(P) = |S^+| - |S^-| = \binom{2n}{n} - \binom{2n}{n+1} = \frac{1}{n+1}\binom{2n}{n}.$$

Let us recall Theorem 2.4, in which we proved that given a sign-reversing involution I, $\sum_{P \in S} \operatorname{sgn}(P) = \sum_{P \in \mathrm{Fix}(I)} \operatorname{sgn}(P)$. What remains to show in the proof is that by finding a particular involution I, the fixed point set of I is the set of lattice paths corresponding to A never being behind B.

In order to find the appropriate involution, which we will denote ϕ, we will use the *reflection principle*. First we note that every member of S^- will cross the x-axis, so for each path P in S^-, denote by m the smallest x co-ordinate where P touches the x-axis. Now reflect the initial segment of P from $x = 0$ to $x = m$ across the x-axis, leaving the rest of P unchanged. Denote the path we obtain $\phi(P)$. An example is given in Figure 5.

For all paths P in S^+ which touch or cross the x-axis, we can define $\phi(P)$ in the same way. This new path will belong to S^-. If a path doesn't touch the x-axis, let $\phi(P) = P$. Then, it is a fixed point of ϕ, because it is unchanged by the function.

We need to check ϕ is an involution, and that it sign-reversing.

Recall that for a function to be an involution, it must be such that $\phi(\phi(P)) = P$ for every path P in S, as in Definition 2.2. If P is in the fixed point set of ϕ, ϕ does not change P, so $\phi(\phi(P)) = P$. If ϕ reflects the first section of P, then since the first point of intersection is not changed by ϕ, applying ϕ twice will give us P. Hence, ϕ is an involution. The function ϕ is also sign-reversing because provided the path P is not in the fixed point set of ϕ, then if P is in S^+, $\phi(P)$ is in S^-, and vice versa.

So for this sign-reversing involution on the set S of lattice paths, the only members of $\mathrm{Fix}(\phi)$ are those elements of S^+ which do not touch or cross the x-axis. This precisely corresponds to the sequences in which A is never behind B, so the number of such sequences is equal to $|\mathrm{Fix}(\phi)|$. Using the fact that all members of $\mathrm{Fix}(\phi)$ have positive sign, and Theorem 2.4, we get

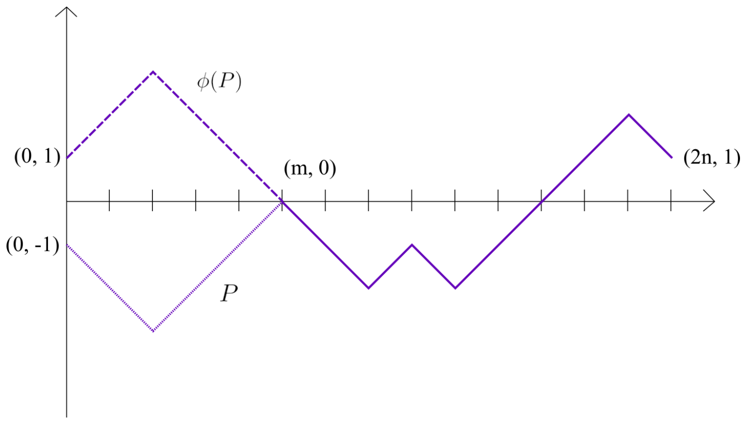

Figure 5: A lattice path P in S^- and its corresponding S^+ member, $\phi(P)$. Observe that after the first intersection point, $(m, 0)$, the two paths are identical. We can see that $\phi(P)$ will be in S^+, since it is a lattice path from $(0, 1)$ to $(2n, 1)$.

$$|\operatorname{Fix}(\phi)| = \sum_{P \in \operatorname{Fix}(I)} \operatorname{sgn}(P) = \sum_{P \in S} \operatorname{sgn}(P) = \frac{1}{n+1}\binom{2n}{n}. \qquad \square$$

In this proof, we have found a bijection between the set we are interested in and a set of lattice paths, and used a sign-changing involution to count such lattice paths. Hence, we have shown how we can use such a method to solve abstract problems that, otherwise, would be more difficult to visualise. In the remainder of this essay, we will be considering how we can apply a similar method to the Aztec Diamond.

4 Schröder Numbers and The Aztec Diamond Revisited

In this chapter we aim to define Schröder numbers and Schröder paths, and using the method of [5], show how a certain n-tuple (ordered set of n objects) of Schröder paths can be mapped in a bijection to tilings of our Aztec Diamond.

Definition 4.1. By first fixing an integer a, the *large Schröder numbers* $(r_n)_{n\geq 0} = 1, 2, 6, 22, 90, 394, 1806, \ldots$ count the number of lattice paths from $(a, 0)$ to $(a + 2n, 0)$, with up steps of $(1, 1)$, down steps of $(-1, 1)$ and level steps of $(2, 0)$, which do not pass below the x-axis. We call such a path a *large Schröder path of size* n, or a *large n-Schröder path*.

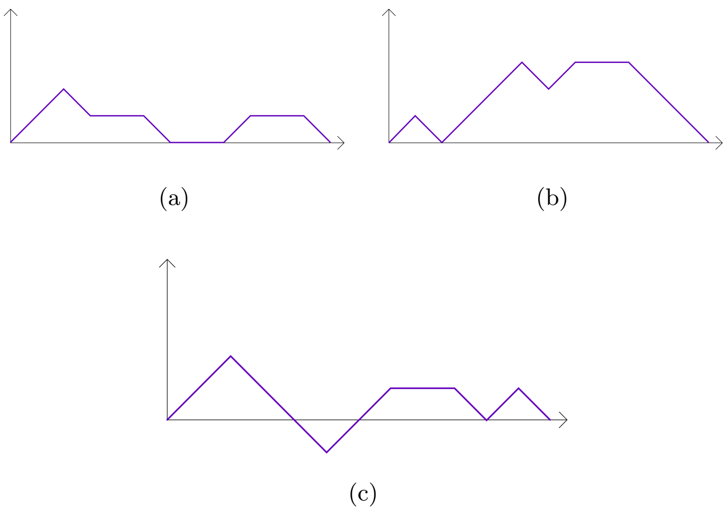

Figure 6: Three examples of lattice paths. Path (a) is a large but not small Schröder path because it has a level step on the x-axis, path (b) is both, and path (c) is neither, since it descends below the x-axis.

Definition 4.2. Again fixing an integer a, the *small Schröder numbers* $(s_n)_{n\geq 0} = 1, 1, 3, 11, 45, 197, 903, \ldots$ count the number of large n-Schröder paths from $(a, 0)$ to $(a+2n, 0)$ which have no level steps on the x-axis. Such a path is called a *small n-Schröder path.*

Figure 6 shows the conditions lattice paths must fulfill to be a large or small Schröder path.

Proposition 4.3. *For $n \geq 1$, the number of large n-Schröder paths is twice the number of small n-Schröder paths, or in other words, $r_n = 2s_n$.*

Remark. We can see this fact explicitly from the first terms in each sequence above. A proof shall not be provided in this essay, but one such proof using the generating functions of the two sequences can be found in [4, pp. 310–313].

Let us define a set, Π_n, of n-tuples of large Schröder paths $(\pi_1, \pi_2, \ldots, \pi_n)$ where π_i is a large i-Schröder path satisfying the following two conditions:

C1 For $0 \leq i \leq n$, the path π_i begins at $(-2i+1, 0)$ and ends at $(2i-1, 0)$.

C2 No two paths π_i and π_j intersect with one another.

We will similarly define Ω_n a the set of n-tuples of *small* Schröder paths $(\omega_1, \omega_2, \ldots, \omega_n)$ which satisfy **C1** and **C2**. In less formal words, Π_n is the set of n-tuples of lattice paths where each lattice path never descends below

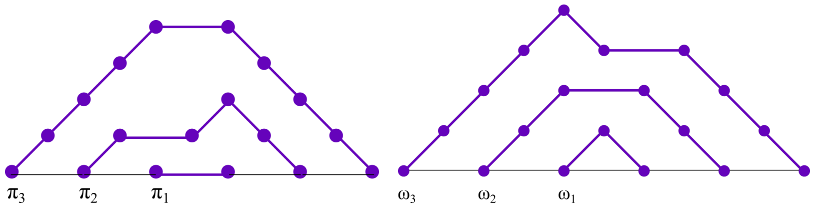

Figure 7: Figure (a) shows a member of Π_3, and Figure (b) a member of Ω_3. Note that $(\omega_1, \omega_2, \omega_3)$ belongs to Π_3 as well.

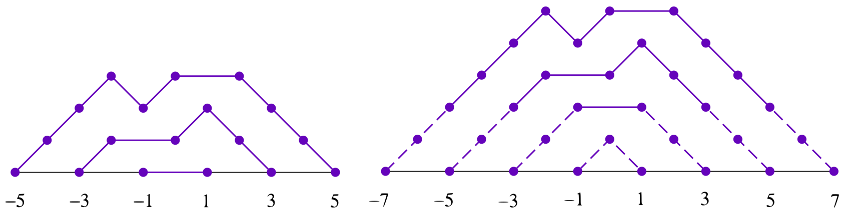

Figure 8: A member of Π_3 and its corresponding member of Ω_4.

the x-axis, and the lattice paths do not intersect with one another. The set Ω_n has the same conditions, along with the additional condition that each path cannot have level steps on the x-axis. Figure 7 shows what a member of these two sets might look like.

We are following Eu and Fu's method as in [5].

Lemma 4.4. *For $n \geq 2$, $|\Pi_{n-1}| = |\Omega_n|$.*

Proof. We want to show there is a bijection between Π_{n-1} and Ω_n. Define a function $\phi : \Pi_{n-1} \to \Omega_n$. For a member $(\pi_1, \ldots, \pi_{n-1})$ of Π_{n-1}, let $\phi((\pi_1, \ldots, \pi_{n-1})) = (\omega_1, \ldots, \omega_n)$, where $\omega_1 = \mathbf{UD}$, and $\omega_i = \mathbf{UU}\pi_{i-1}\mathbf{DD}$, for $2 \leq i \leq n$.

Figure 8 shows how the bijection works. Note that since a member of Ω_n must have no level steps on the x-axis, all members of Ω_n must have $\omega_1 = \mathbf{UD}$, and for $i > 1$, ω_i must start with an up step, and end with a down step. Once we know this, for $i > 1$, ω_i must start with two up steps and end with two down steps, in order to avoid intersections with ω_{i-1}.

Hence we can define an inverse function $\phi^{-1} : \Omega_n \to \Pi_{n-1}$ which given $(\omega_1, \ldots, \omega_n)$ in Ω_n, ignores ω_1 and lets ϕ_i be ω_{i+1} without the first two up steps or last two down steps, for $i \in 1, \ldots, n-1$. This obtains a sequence $(\pi_1, \ldots, \pi_{n-1}) \in \Pi_{n-1}$, and applying ϕ and ϕ^{-1} to a member of Π_{n-1} returns the original sequence. Hence ϕ is a bijection, and so the lemma holds. □

To see how a member of Π_n relates to a tiling of the Aztec Diamond,

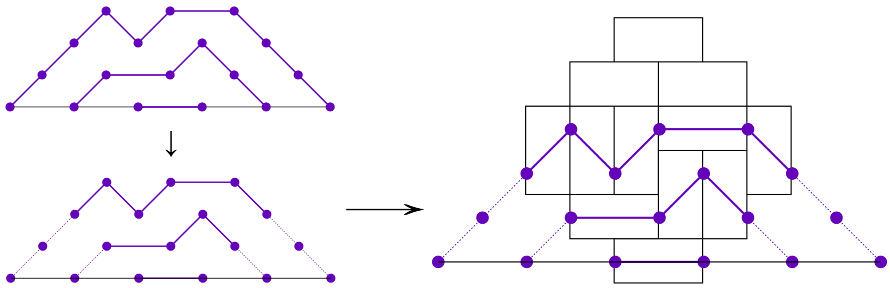

Figure 9: This is an example of the the bijection we will be defining in Proposition 4.5. Note in particular that there is only one possibility for the top 3 dominoes once we have defined the rest of the tiling.

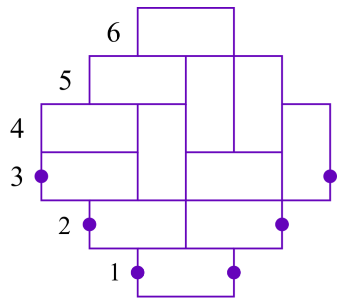

Figure 10: A tiling of AD_3 with its rows indexed. The marked points on rows 1 to 3 indicate where we will begin and end our lattice paths.

we may first simply add some dashes to the previous diagram, and overlay a corresponding tiling of the Aztec Diamond, as in Figure 9.

Hence, what we want to show is that for every tiling of AD_n, we can find a corresponding member of Π_n, and that this mapping is a bijection. We will again be going through a sketch of the proof used by [5].

Proposition 4.5. *There is a bijection between the tilings of the Aztec Diamond of size n, and the set of n-tuples* $(\pi_1, \pi_2, \ldots, \pi_n)$ *of large Schröder paths which satisfy the conditions* ***C1*** *and* ***C2****.*

Sketch of Proof. First we will find a way to encode a tiling T of AD_n as a sequence of non-intersecting lattice paths, and check that such an n-tuple of lattice paths $(\tau_1, \tau_2, \ldots, \tau_n)$ is unique, that is, we can recover a single tiling of AD_n. Note at this stage we do not require that each path is a large Schröder path.

For a given tiling of AD_n, index its rows by $1, 2, \ldots, 2n$ from bottom to top, as in Figure 10. Then, for each i for $1 \leq i \leq n$, define a path τ_i which starts at the centre of the left hand edge of the ith row, and finishes at the

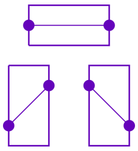

(a)

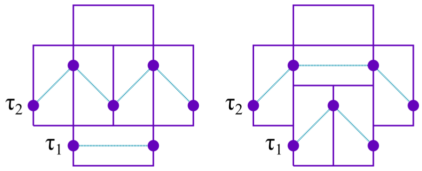

(b)

Figure 11: Figure (a) shows the rules we define for a domino. Figure (b) shows how we can apply these rules to two tilings of AD_2.

centre of the right hand edge of the ith row. In between, we define τ_i as follows:

Define a domino D as having 6 edges, each of unit edge. At a given step of τ_i, that is with the point at the left side of the domino fixed,

- if D is vertical and the fixed point is at the bottom left edge, set the step to go to the top right (an up step).
- if D is vertical and the fixed point is at the top left edge, set the step to go to the bottom right (a down step).
- if D is horizontal, translate this to a level step.

Figure 11 shows how we can apply these domino rules to obtain a set of n lattice paths.

We can then check that each tiling T has a unique n-tuple of paths, and that for all $1 \leq i, j \leq n$, the paths τ_i and τ_j do not intersect. We can also reverse this procedure — for an n-tuple of non-intersecting paths $(\tau_1, \tau_2, \ldots, \tau_n)$ with start and finish points defined as above, we can recover a unique tiling T of AD_n. We would go about this by overlaying dominoes on each path in the same manner as before, and as the paths are non-intersecting, we can show these dominoes will not overlap, and that any gaps in the produced regions can be filled precisely by dominoes to map to a tiling of AD_n which is unique. Let us denote by Λ_n the set of paths we have just created. We will show there is a bijection between the set of tilings and the set Π_n of non-intersecting lattice paths satisfying **C1** and **C2**, by using Λ_n. We have shown there is a bijection between the set of tilings and Λ_n, so we need to create a bijection between Λ_n and Π_n. Let us define a function $\psi : \{\text{Tilings of } AD_n\} \to \Pi_n$ as follows:

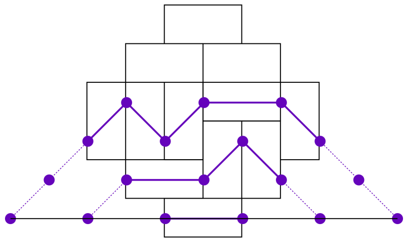

Figure 12: A set of paths for a tiling of AD_3, with the dashed lines indicating how we add up and down steps under ψ.

Given a tiling T of AD_n, let $(\tau_1, \tau_2, \ldots, \tau_n)$ in Λ_n be the n-tuple of paths we associated with T above. Then, for a given tiling T, $\psi(T) = (\pi_1, \pi_2, \ldots, \pi_n)$, where

$$\pi_i = \underbrace{\mathsf{UU}\ldots\mathsf{U}}_{i-1 \text{ times}} \tau_i \underbrace{\mathsf{D}\ldots\mathsf{D}}_{i-1 \text{ times}}$$

That is to say, the large Schröder path π_i is obtained from τ_i with $i-1$ up steps at the beginning, and $i-1$ down steps at the end, for each i with $1 \leq i \leq n$. Observe an example of this in Figure 12.

From this we can check that the resulting n-tuple of paths is in Π_n. We can see that each path starts and ends at the required points, and since each n-tuple in Λ_n is non-intersecting, so is the resulting n-tuple. Hence, it satisfies **C1** and **C2**, and so belongs to Π_n.

The final thing we can check in this proof is that given an n-tuple of lattice paths in Π_n, we can find a corresponding unique set of paths in Λ_n. This can be done by showing that for a given n-tuple of paths $(\pi_1, \pi_2, \ldots, \pi_n)$ in Π_n, the path π_i must begin with $i-1$ up steps and end with $i-1$ down steps in order for the paths to not intersect — then, we can map this path directly to a path τ_i by removing the first $i-1$ steps and the last $i-1$ steps.

The result is that this function ψ is a bijection, and hence a bijection exists between the set of tilings of AD_n and n-tuples of paths in Π_n, as required.

□

Hence, we have discovered a way of representing tilings as a set of n lattice paths, and shown there is a one-to-one mapping between them — therefore, if we can count our set of lattice paths Π_n, we have found a formula for the number of domino tilings of the Aztec Diamond of size n.

5 Determining the size of Π_n

In this chapter we will bring together some concepts from earlier in the essay, by using a signed set and a sign-reversing involution, together with the determinant of a certain matrix, to show that for a given n, there are are $2^{n(n+1)/2}$ n-tuples of paths satisfying **C1** and **C2**, and hence using our last proposition, this is the number of domino tilings of AD_n.

Before we begin, we need one more definition:

Definition 5.1. The n^{th} *Hankel matrix* $A_n = (\alpha_{ij})$ *of a sequence* (β_n) is an $n \times n$ symmetric matrix such that for $1 \leq i, j \leq n$, $\alpha_{ij} = \beta_{i+j-1}$.

So in general,

$$A_n = \begin{pmatrix} \beta_1 & \beta_2 & \beta_3 & \cdots & \beta_n \\ \beta_2 & \beta_3 & \beta_4 & \cdots & \beta_{n+1} \\ \beta_3 & \beta_4 & \beta_5 & \cdots & \beta_{n+2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \beta_n & \beta_{n+1} & \beta_{n+2} & \cdots & \beta_{2n-1} \end{pmatrix}.$$

For example, if we were given the sequence $(a_n) = (n^2)$, the 4^{th} Hankel matrix of (a_n) would be

$$\begin{pmatrix} 1 & 4 & 9 & 16 \\ 4 & 9 & 16 & 25 \\ 9 & 16 & 25 & 36 \\ 16 & 25 & 36 & 49 \end{pmatrix}.$$

Hence, let us represent our large and small Schröder numbers from Definitions 4.1 and 4.2 in this form. Denote by H_n the nth Hankel matrix of the sequence of large Schröder numbers (r_n), and by G_n the nth Hankel matrix of the small Schröder numbers (s_n), as shown below.

$$H_n = \begin{pmatrix} r_1 & r_2 & r_3 & \cdots & r_n \\ r_2 & r_3 & r_4 & \cdots & r_{n+1} \\ r_3 & r_4 & r_5 & \cdots & r_{n+2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ r_n & r_{n+1} & r_{n+2} & \cdots & r_{2n-1} \end{pmatrix} \quad G_n = \begin{pmatrix} s_1 & s_2 & s_3 & \cdots & s_n \\ s_2 & s_3 & s_4 & \cdots & s_{n+1} \\ s_3 & s_4 & s_5 & \cdots & s_{n+2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ s_n & s_{n+1} & s_{n+2} & \cdots & s_{2n-1} \end{pmatrix}$$

Lemma 5.2. *The Hankel matrix of the large and small Schröder numbers satisfies* $\det(H_n) = 2^n \det(G_n)$*, for each* $n \in \mathbb{N}$.

Proof. Let us recall Proposition 4.3, which tells us that $r_n = 2s_n$ for $n \geq 1$, that is each large Schröder number is twice the corresponding small Schröder number. Hence,

$$H_n = 2G_n = 2I_nG_n.$$

Therefore, examining the determinants of both sides, we obtain that

$$\det(H_n) = \det(2I_n)\det(G_n) = 2^n \det(G_n).$$

□

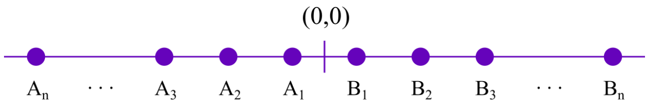

Figure 13: The points we are labeling by A_i and B_i.

We are ready for the final proof in this essay, which will again follow the steps of Eu and Fu [5].

Proposition 5.3. *For $n \geq 1$, we have*

(i) $|\Pi_n| = \det(H_n) = 2^{n(n+1)/2}$, *and*

(ii) $|\Omega_n| = \det(G_n)$.

Proof. Let us begin by defining, for $1 \leq i \leq n$, the point A_i to be $(-2i+1, 0)$, and the point B_i to be $(2i-1, 0)$, as in Figure 13.

Then for any given A_i and B_j, note that since they are $(2i-1)-(-2j+1) = 2(i+j-1)$ units apart, using Definition 4.1 of the large Schröder numbers, the number of possible large Schröder paths from A_i to B_j is given by r_{i+j-1}, which is the (i,j)-th entry of our Hankel matrix $H_n = (h_{ij})$ of the large Schröder numbers.

We want to first create a sum that we can analyse using this method. Define a set P of ordered pairs $(\sigma, (\tau_1, \dots, \tau_n))$, made up of a permutation σ of $\{1, \dots, n\}$ and a set of n large Schröder paths, where τ_i starts at A_i and ends at $B_{\sigma(i)}$. We will give each ordered pair an associated sign defined by the sign of the permutation σ. Now, we can partition the elements of P into P^+ and P^-.

For a given permutation σ, there are $h_{1\sigma(1)}h_{2\sigma(2)}...h_{n\sigma(n)}$ possible n-tuples of lattice paths, so in total we find

$$|P^+| - |P^-| = \sum_{\sigma \in S_n} \operatorname{sgn}(\sigma) \prod_{i=1}^{n} h_{i\sigma(i)}. \tag{1}$$

By recalling the definition of the determinant of a matrix, we then realise that

$$\det(H_n) = \sum_{\sigma \in S_n} \operatorname{sgn}(\sigma) \prod_{i=1}^{n} h_{i\sigma(i)} = |P^+| - |P^-|. \tag{2}$$

Now we want to find a sign-changing involution on this set of paths. Since Proposition 4.5 considers n-tuples of lattice paths which do not intersect, let us attempt to find a map that matches up n-tuples of paths which have an intersection.

Let $(\sigma, (\tau_1, \dots, \tau_n))$ be a pair in P such that at at least two of the paths have an intersection. If there are more than two intersecting paths, take

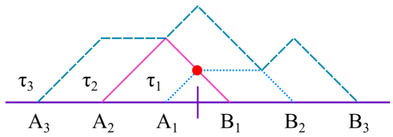

Figure 14: This figure shows an ordered pair $((1, 2), (\tau_1, \tau_2, \tau_3))$ where each path intersects with the other two. In this case, when defining the involution, we will only consider the intersection of τ_1 and τ_2.

the first two intersecting paths numerically, denoting them τ_i and τ_j. For example, if we have $(\sigma, (\tau_1, \tau_2, \tau_3))$ such that all three paths intersect, we will consider only the intersection of τ_1 and τ_2. An example of this is given in Figure 14.

Recalling the reflection principle method used in Theorem 3.1, we will use a similar method to define our involution. Given our chosen τ_i and τ_j, construct new paths τ_i' and τ_j' as follows: Denote the last intersection point of τ_i and τ_j by m. Then define τ_i' to follow the same path as τ_i up until the point m. After m, let τ_i' follow the path of τ_j. Hence τ_i' will start at A_i and end at $B_{\sigma(j)}$. Define τ_j' similarly, then this path will start at B_i and end at $A_{\sigma(i)}$. Hence, the associated permutation of the n-tuple of paths $(\tau_1, \ldots, \tau_i', \ldots, \tau_j', \ldots, \tau_n)$ is $\sigma \circ (i, j)$.

So, to set up our required sign-changing involution, let us define a function $\phi : P \to P$ such that

$$\phi((\sigma, (\tau_1, \ldots, \tau_n))) = (\sigma \circ (i, j), (\tau_1, \ldots, \tau_i', \ldots, \tau_j', \ldots, \tau_n),$$

with τ_i' and τ_j' defined as above if there is an intersection of the paths (and note this changes the sign of the ordered pair). If the n-tuple of paths contains no intersections, chose ϕ not to change the pair. Figure 15 provides an example of the result of applying ϕ to an element of P in which two paths intersect.

Hence, we only need to check that ϕ is an involution, that is that if we apply ϕ twice, we recover the original element of P. If there are no intersections in the paths of an element, applying ϕ gives the same element, so let us focus on an ordered pair with intersecting paths.

Let $p = (\sigma, (\tau_1, \ldots, \tau_n))$ be an element in P in which at least two of the lattice paths intersect. After applying ϕ, the same lattice paths will intersect as before, as the function introduces no new intersections, and if τ_i and τ_j were the paths affected by ϕ, then τ_i' and τ_j' will still intersect in $\phi(p)$. Hence, applying ϕ to $\phi(p)$ will affect only the paths τ_i' and τ_j' — note this takes them back to their original positions, returning the lattice path n-tuple $(\tau_1, \ldots, \tau_n))$. The permutation will again be composed with (i, j), giving $\sigma \circ (i, j) \circ (i, j) = \sigma$. Figure 15 illustrates how this works for a given ordered pair. So, in all cases, $\phi \circ \phi$ acts as the identity, and so ϕ is a

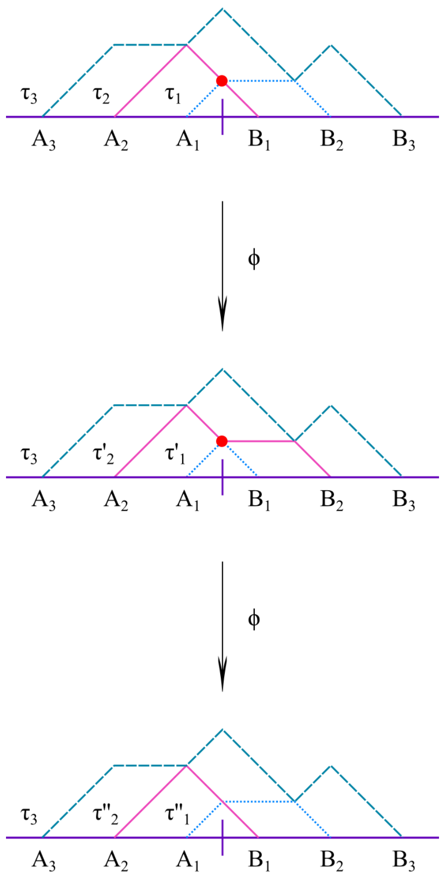

Figure 15: This is an example of how we would apply the function ϕ to a member of P, $((1, 2), (\tau_1, \tau_2, \tau_3))$. The red dot denotes the last point of intersection of the first two intersecting paths — notice that it does not change under ϕ. The diagram also shows how applying ϕ twice in this case returns the original ordered pair.

sign-changing involution.

This means we can use the involution theorem (Theorem 2.4) to reduce this problem. Since ϕ is a sign-changing involution, we can say that

$$\sum_{\sigma \in S_n} \operatorname{sgn}(\sigma) \prod_{i=1}^{n} h_{i\sigma(i)} = \sum_{\sigma \in \operatorname{Fix}(\phi)} \operatorname{sgn}(\sigma) \prod_{i=1}^{n} h_{i\sigma(i)}.$$

Our set $\operatorname{Fix}(\phi)$ is the set of ordered pairs $(\sigma, (\tau_1, \tau_2, \ldots, \tau_n))$ in P in which $(\tau_1, \tau_2, \ldots, \tau_n)$ are non-intersecting.

From this, we can first deduce that each member of $\operatorname{Fix}(\phi)$ must have the identity (Id) as its associated permutation. If there are two paths τ_i, τ_j such that $i < j$ but $\sigma(i) > \sigma(j)$, then there will be an intersection of τ_i and τ_j, so we must have $\sigma(i) < \sigma(j)$ for all $i < j$, giving that σ is the identity permutation.

Using this, we can deduce that for each ordered pair $(Id, (\tau_1, \tau_2, \ldots, \tau_n))$ in $\operatorname{Fix}(\phi)$, $(\tau_1, \tau_2, \ldots, \tau_n)$ is in Π_n. We can show this because each path is a large Schröder path, and satisfies **C1** and **C2** — it satisfies **C1** because for each i, $1 \leq i \leq n$, τ_i goes from A_i to B_i, and it satisfies **C2** because if the ordered pair belongs to the fixed point set of ϕ, then there must be no intersections in the paths. Hence, we know that

$$\sum_{\sigma \in \operatorname{Fix}(\phi)} \operatorname{sgn}(\sigma) \prod_{i=1}^{n} h_{i\sigma(i)} = |\Pi_n|,$$

and so using Equation 2,

$$\det(H_n) = |\Pi_n|.$$

Using a similar argument, we can consider the determinant of the Hankel matrix of the small Schröder numbers G_n, and show that $\det(G_n) = |\Omega_n|$, where Ω_n is the number of n-tuples of small Schröder paths satisfying **C1** and **C2**.

Finally, we use Lemma 5.2 that $\det(H_n) = 2^n \cdot \det(G_n)$, and Lemma 4.4 that $|\Pi_{n-1}| = |\Omega_n|$ to prove our result. Since $|\Pi_n| = \det(H_n)$, $|\Omega_n| = \det(G_n)$, we therefore find that

$$|\Pi_n| = 2^n |\Omega_n| = 2^n |\Pi_{n-1}|.$$

Hence we have a recursion relation for $|\Pi_n|$. Since for $n = 1$, we have that the number of large Schröder paths from $(-1, 0)$ to $(1, 0)$ is precisely two — one level step, or an up step followed by a down step — we find that $|\Pi_1| = 2$, and so

$$|\Pi_n| = 2^n \cdot 2^{n-1} \cdots 2^2 \cdot 2 = 2^{n(n+1)/2}.$$

□

Now, we have discovered two facts:

1. There is a bijection between the set of tilings of the nth Aztec Diamond and Π_n, our set of n-tuples of lattice paths satisfying **C1** and **C2**. In particular, the number of tilings of AD_n is $|\Pi_n|$.

2. A formula for the size of Π_n, namely that $|\Pi_n| = 2^{n(n+1)/2}$.

Therefore, we can bring together these facts to prove the final theorem of this essay.

Theorem 5.4 (Aztec Diamond Theorem)**.** *The number of tilings of the Aztec Diamond of order n is $2^{n(n+1)/2}$.*

6 Conclusion

Having proven this formula, it turns out many related regions can have their tilings enumerated as variations of this problem, and hence Aztec Diamonds play a fundamental role in not only domino tilings, but many more general problems in tiling theory. Any examples of such uses would go beyond the scope of this essay, but for those interested, James Propp has written a summary on recent progress in this field [9]. He reflects on how future research may allow us to discover regions and associated dominoes that have surprisingly simple formulas for their tiling enumerations, by whether they can be reduced to considering Aztec Diamonds, or other tilings of a similar nature. Furthermore, Propp has also written an article [8] explaining how one can use a series of algorithms to reduce certain tilings of more complex regions, through the use of graph theory, to an Aztec Diamond region with certain properties[3].

Hence, the result we have proven in this essay provides us with a fundamental building block in our consideration of tilings.

References

[1] Federico Ardila and Richard P Stanley. Tilings. *The Mathematical Intelligencer*, 32(4):32–43, 2010.

[2] Arthur T Benjamin and Jennifer J Quinn. An alternate approach to alternating sums: a method to die for. *The College Mathematics Journal*, 39(3):191–202, 2008.

[3] Joseph Bertrand. Solution d un probleme. *Comptes Rendus de l'Académie des Sciences, Paris*, 105:369, 1887.

[3]This article features the section heading "DIABOLO-TILINGS OF FORTRESSES", which in itself should persuade the reader to seek it out.

[4] Richard A. Brualdi. *Introductory Combinatorics*. Pearson, 5th edition, 2009.

[5] Sen-Peng Eu and Tung-Shan Fu. A simple proof of the aztec diamond theorem. *The Electronic Journal of Combinatorics*, 12:R18, 2005.

[6] Pieter W Kasteleyn. The statistics of dimers on a lattice: I. the number of dimer arrangements on a quadratic lattice. *Physica*, 27(12):1209–1225, 1961.

[7] Nicholas Loehr. *Bijective combinatorics*. CRC Press, 2011.

[8] James Propp. Generalized domino-shuffling. *Theoretical Computer Science*, 303(2-3):267–301, 2003.

[9] James Propp et al. Enumeration of matchings: problems and progress. *New Perspectives in Geometric Combinatorics*, pages 255–291, 1999.

[10] Dennis Stanton and Dennis White. *Constructive Combinatorics*. Springer Science & Business Media, 1986.