



Fillipo de Souza Lima Impellizieri

Domino Tilings of the Torus

Dissertação de Mestrado

Dissertation presented to the Programa de Pós-Graduação em Matemática of the Departamento de Matemática, PUC-Rio as partial fulfillment of the requirements for the degree of Mestre em Matemática.

Advisor: Prof. Nicolau Corção Saldanha

Rio de Janeiro
September 2015



Fillipo de Souza Lima Impellizieri

Domino Tilings of the Torus

Dissertation presented to the Programa de Pós-Graduação em Matemática of the Departamento de Matemática do Centro Técnico Científico da PUC-Rio, as partial fulfillment of the requirements for the degree of Mestre.

Prof. Nicolau Corção Saldanha

Advisor

Departamento de Matemática – PUC-Rio

Prof. Carlos Tomei

Departamento de Matemática – PUC-Rio

Prof. Juliana Abrantes Freire

Departamento de Matemática – PUC-Rio

Prof. Marcio da Silva Passos Telles

Instituto de Matemática e Estatística – UERJ

Prof. Robert David Morris

Instituto de Matemática Pura e Aplicada – IMPA

Prof. José Eugenio Leal

Coordinator of the Centro Técnico Científico – PUC-Rio

Rio de Janeiro, September 11th, 2015

All rights reserved

Fillipo de Souza Lima Impellizieri

The author obtained the degree of Bacharel em Matemática from PUC-Rio in July 2013

Bibliographic data

Impellizieri, Fillipo de Souza Lima

Domino Tilings of the Torus / Fillipo de Souza Lima Impellizieri ; advisor: Nicolau Corção Saldanha. — 2015.
146 f. : il. ; 30 cm

Dissertação (Mestrado em Matemática)-Pontifícia Universidade Católica do Rio de Janeiro, Rio de Janeiro, 2015.
Inclui bibliografia

1. Matemática – Teses. 2. domino. 3. tiling. 4. torus. 5. lattice. 6. flux. 7. flip. 8. height function. 9. Kasteleyn matrix. I. Saldanha, Nicolau Corção. II. Pontifícia Universidade Católica do Rio de Janeiro. Departamento de Matemática. III. Título.

CDD: 510

Acknowledgments

To CNPq, FAPERJ and PUC-Rio, for making this research possible.

To my advisor, Prof. Nicolau Corção Saldanha, for his boundless patience, eagerness to help and valuable insight.

To the Department of Mathematics, especially Creuza, Fred and Renata, for helping me get here and putting up with my fickleness.

To my jury, for their appreciation, understanding and precious feedback. In particular, I would like to thank Carlos Tomei for his spirit and enthusiasm, and Marcio Telles for his dedication.

To my professor Ricardo Earp, for his guidance and caring.

To my friends, for their positivity and reassurance. In particular, to my graduation friends Felipe, Leonardo and Gregory for holding onto this ‘academic friendship’.

To my family, for their patience, unwavering support and cherished love.

Abstract

Impellizieri, Fillipo de Souza Lima; Saldanha, Nicolau Corção (Advisor). **Domino Tilings of the Torus**. Rio de Janeiro, 2015. 146p. Msc. Dissertation — Departamento de Matemática, Pontifícia Universidade Católica do Rio de Janeiro.

We consider the problem of counting and classifying domino tilings of a quadrilateral torus. The counting problem for rectangles was studied by Kasteleyn and we use many of his ideas. Domino tilings of planar regions can be represented by height functions; for a torus given by a lattice L , these functions exhibit arithmetic L -quasiperiodicity. The additive constants determine the flux of the tiling, which can be interpreted as a vector in the dual lattice $(2L)^*$. We give a characterization of the actual flux values, and of how corresponding tilings behave. We also consider domino tilings of the infinite square lattice; tilings of tori can be seen as a particular case of those. We describe the construction and usage of Kasteleyn matrices in the counting problem, and how they can be applied to count tilings with prescribed flux values. Finally, we study the limit distribution of the number of tilings with a given flux value as a uniform scaling dilates the lattice L .

Keywords

domino; tiling; torus; lattice; flux; flip; height function; Kasteleyn matrix;

Resumo

Impellizieri, Fillipo de Souza Lima; Saldanha, Nicolau Corção. **Coberturas do Toro por Dominós**. Rio de Janeiro, 2015. 146p. Dissertação de Mestrado — Departamento de Matemática, Pontifícia Universidade Católica do Rio de Janeiro.

Consideramos o problema de contar e classificar coberturas por dominós de toros quadriculados. O problema de contagem para retângulos foi estudado por Kasteleyn e usamos muitas de suas ideias. Coberturas por dominós de regiões planares podem ser representadas por funções altura; para um toro dado por um reticulado L , estas funções exibem L -quasiperiodicidade aritmética. As constantes aditivas determinam o fluxo da cobertura, que pode ser interpretado como um vetor no reticulado dual $(2L)^*$. Damos uma caracterização dos valores de fluxo efetivamente realizados e de como coberturas correspondentes se comportam. Também consideramos coberturas por dominós do reticulado quadrado infinito; coberturas de toros podem ser vistas como um caso particular destas. Descrevemos a construção e uso de matrizes de Kasteleyn no problema de contagem, e como elas podem ser aplicadas para contar coberturas com valores de fluxo prescritos. Finalmente, estudamos a distribuição limite do número de coberturas com um dado valor de fluxo quando o reticulado L sofre uma dilatação uniforme.

Palavras-chave

dominó; cobertura; toro; reticulado; fluxo; flip; função altura; matriz de Kasteleyn;

Contents

1	Introduction	8
2	Definitions and Notation	16
3	Domino tilings on the plane	18
3.1	Flips and height functions	20
3.2	Kasteleyn matrices	26
3.3	A classical result: domino tilings of the rectangle	28
4	Domino tilings on the torus	32
4.1	Height functions on the torus	36
4.2	More general tori: valid lattices	41
5	Flux on the torus	44
5.1	The affine lattice $L^\#$	47
5.2	Characterization of flux values	50
6	Flip-connectedness on the torus	63
6.1	Tilings of the infinite square lattice	64
6.2	Back to the torus	78
7	Kasteleyn matrices for the torus	84
7.1	The structure of $\mathcal{F}(L) \cap \partial Q$	89
8	Sign distribution over $\mathcal{F}(L)$	98
8.1	Cycles and cycle flips	98
8.2	The effect of a cycle flip on the sign of a flux	101
8.3	The effect of a cycle flip on the flux itself	108
9	Kasteleyn determinants for the torus	117
9.1	The case of $M = KK^* \oplus K^*K$	117
9.2	Spaces of L -quasiperiodic functions and the case of K, K^*	124
9.3	Formulas for $\det(K_E)$, ρ and uniform scaling	132
	Glossary	139

1

Introduction

Tilings of planar regions by dominoes (and also lozenges) can be thought of as perfect matchings of a corresponding graph. In this sense, the enumeration of matchings was studied as early as 1915 by MacMahon [10], whose focus was on plane partitions. Also around the time, chemists and physicists were interested in aromatic hydrocarbons and the behavior of liquids. Hereafter, I will refer to perfect matchings simply by ‘matchings’.


Research on dimers in statistical mechanics had a major breakthrough in 1961, when Kasteleyn [7] (and, independently, Temperley and Fisher [17]) discovered a technique to count the matchings of a subgraph G of the infinite square lattice. He proved that this number is equal to the Pfaffian of a certain 0,1-matrix M associated with G . Not much later, Percus [12] showed that when G is bipartite, one can modify M so as to obtain the number from its determinant (rather than from its Pfaffian). James Propp [13] provides an interesting overview of the topic on his ‘Problems and Progress in Enumeration of Matchings’.

In the early 90s, more advances were made and gave new impetus to research. Conway [3] devised a group-theoretic argument that, in many interesting cases, may be used to show that a given region cannot be tessellated by a given set of tiles. In a related work, Thurston [18] introduced the concept of height functions: integer-valued functions that encode a tiling of a region. With them, he presented a simple algorithm that verifies the domino-tileability of simply-connected planar regions.

In 1992, Aztec diamonds were examined by Elkies, Kuperberg, Larsen and Propp [4], who gave four proofs of a very simple formula for the number of domino tilings of these regions. Later, probability gained importance with the study of random tilings, and Jockush, Propp and Shor [6, 2] proved the Arctic Circle Theorem. This framework was further generalized in the early 2000s by Kenyon, Okounkov and Sheffield [9, 8], whose work relates random tilings to Harnack curves and describes the variational problem in terms of the complex Burgers equation.


While now much is known about tilings for planar regions, higher dimensions have proven less tractable. Randall and Yngve [14] examined analogues of Aztec diamonds in three dimensions for which many of the two-dimensional results can be adapted. Hammersley [5] makes asymptotic estimates on the number of brick tilings of a d -dimensional box as all dimensions go to infinity. In his thesis, Milet [11] studied certain three-dimensional regions for which he defines an invariant that can be interpreted under knot theory.

This dissertation was motivated by the observation of a certain asymptotic behavior in the statistics of domino tilings of square tori. We elaborate: consider a quadriculated torus, represented by a square with sides of even length and whose opposite sides are identified. A *domino* is a 2×1 rectangle. Below, we have a tiling of the 4×4 torus which also happens to be a tiling of the 4×4 square.




A domino tiling of the 4×4 torus

Because in the torus opposite sides are identified, we may also consider tilings with dominoes that ‘cross over’ to the opposing side.



Tilings of the 4×4 torus featuring cross-over dominoes

The *flux* of a tiling is an algebraic construct that counts these cross-over dominoes, with a sign; one may think of it as a pair of integers. In the next figure, we assign the positive sign when a white square is to the right of the blue curve or when a black square is above the red curve (and the negative sign otherwise). Hence, their fluxes are $(0, -1)$, $(1, 0)$ and $(1, 1)$, where the first integer counts horizontal dominoes crossing the blue curve and the second integer counts vertical dominoes crossing the red curve.



From left to right, tilings with flux $(0, -1)$, $(1, 0)$ and $(1, 1)$

We may thus count tilings of tori by flux. In the 4×4 model, we have:

Flux	$(0, 0)$	$(0, \pm 1), (\pm 1, 0)$	$(\pm 1, \pm 1)$	$(0, \pm 2), (\pm 2, 0)$
Tilings	132	32	2	1

For a total of 272 tilings. Observe the proportion of total tilings by flux:

Flux	$(0, 0)$	$(0, \pm 1), (\pm 1, 0)$	$(\pm 1, \pm 1)$	$(0, \pm 2), (\pm 2, 0)$
Proportion	0.48529	0.11765	0.00735	0.00368

Now we repeat the process for different square tori:

Flux	$(0, 0)$	$(0, \pm 1), (\pm 1, 0)$	$(\pm 1, \pm 1)$	$(0, \pm 2), (\pm 2, 0)$
4×4	0.48529	0.11765	0.00735	0.00368
6×6	0.48989	0.11082	0.01416	0.00253
10×10	0.49436	0.10575	0.01820	0.00141
16×16	0.49564	0.10411	0.02053	0.00109

In each case, tilings with flux $(0, 0)$ comprise almost half of all tilings of the $2n \times 2n$ square torus. For other values of flux in the table it may not be as apparent, but as n increases the proportions stabilize.

Theorem. *As n goes to infinity, the proportions converge to a discrete gaussian distribution. More specifically, for each $i, j \in \mathbb{Z}$, as n goes to infinity*

the proportion relative to flux (i, j) tends to

$$\frac{2 \cdot \Gamma\left(\frac{3}{4}\right)^2}{\sqrt{(6 + 4\sqrt{2})} \cdot \pi} \cdot \exp\left(-\frac{1}{2}(i^2 + j^2)\right)$$

The formula for the rather curious constant can be derived from theta-function identities; see Yi [19]. For comparison, we provide the previous table, together with the limit value given by the formula above:

Flux	$(0, 0)$	$(0, \pm 1), (\pm 1, 0)$	$(\pm 1, \pm 1)$	$(0, \pm 2), (\pm 2, 0)$
4×4	0.48529	0.11765	0.00735	0.00368
6×6	0.48989	0.11082	0.01416	0.00253
10×10	0.49436	0.10575	0.01820	0.00141
16×16	0.49564	0.10411	0.02053	0.00109
Limit	0.49629	0.10317	0.02145	0.00093

We will **not** prove this theorem in this dissertation.

Nevertheless, it motivated us to study domino tilings of the torus and the underlying combinatorial and algebraic structures involved. We expect the content of this text lays the groundwork for writing a proof of this theorem in the future.

That said, results of this kind are not new to physicists, and in fact neither to mathematicians. Boutillier and de Tilière [1] derived explicit formulas for the limit proportions in the honeycomb model of the torus (in this model, a matching may be thought of as a lozenge tiling of the torus). They interpret matchings as loops (see Cycles and cycle flips, Section 8.1) and study the asymptotic behavior of corresponding winding numbers. Although their methods differ from ours, parallels can be drawn.

In Chapter 3, we examine domino tilings of quadriculated planar, simply-connected regions. We discuss how the study of domino tilings is related to the problem of determining perfect matchings of a graph, and present the idea of black-and-white colorings (so our equivalent graphs are bipartite). In Section 3.1 we explore two concepts, as well as their relations. A *flip* is a move on a

tiling that exchanges two dominoes tiling a 2×2 square by two dominoes in the only other possible configuration. A *height function* is an integer valued function on the vertices of the squares of a tiling t that encodes t . Later, these concepts will be generalized to the torus case, and many results of this section (like a characterization of height functions, or the flip-connectedness of these regions) admit adaptation.

Section 3.2 details the construction of *Kasteleyn matrices* and explains how their determinants can be used to count domino tilings of a region. Finally, Section 3.3 contains a worked, classical example: the problem of enumerating domino tilings of the rectangle.

In Chapter 4 begins our study of the torus; we initially consider the square torus \mathbb{T}_n with side length $2n$. The notion of *flux* is introduced here, and an overview of how Kasteleyn matrices can be adapted is provided. We also supply a figure with all tilings of the 4×4 square torus.

Section 4.1 extends height functions to this scenario by interpreting \mathbb{T}_n as a quotient \mathbb{R}^2 / L , where L is the lattice generated by $\{(2n, 0), (0, 2n)\}$. Moreover, we show the flux manifests in the *arithmetic quasiperiodicity* of height functions: they satisfy $h(u + v) = h(u) + k$ for some v, k and all u .

In Section 4.2, we consider more general tori \mathbb{T}_L by allowing other lattices L in the quotient. These are called *valid lattices*: their vectors have integral coordinates that are the same parity. This condition is necessary for the resulting graph to be bipartite.

Chapter 5 further investigates the flux. In Section 5.1, we describe how the flux can be thought of as an element of the *dual lattice* $(2L)^*$. More precisely, we show there is a translate of L^* in $(2L)^*$ that contains all flux values; we call this affine lattice $L^\#$.

Section 5.2 provides our first theorem. For a valid lattice L , let $\mathcal{F}(L)$ be the set of all flux values of tilings of \mathbb{T}_L ; the inner product identification allows us to regard $\mathcal{F}(L)$ as a subset of \mathbb{R}^2 . Consider also the (filled) square $Q \subset \mathbb{R}^2$ with vertices $(\pm\frac{1}{2}, 0), (0, \pm\frac{1}{2})$.

Theorem 1 (Characterization of flux values). $\mathcal{F}(L) = L^\# \cap Q$.

The proof is given by two separate propositions, each showing one inclusion. Much of the technical work here relates to the description of *maximal height functions* (given a base value at a base point).


In Chapter 6, we discuss how flip-connectedness extends to the torus. Flips preserve flux values, so of course the situation must be unlike that of

Section 3.1. It turns out that for flux values in the interior of Q , its tilings are flip-connected, but for flux values in the boundary, none of its tilings admit any flips: they are flip-isolated!

In order to show that, Section 6.1 is devoted to understanding tilings that do not admit flips, and contains our second theorem. It is a fairly independent section, requiring only that the reader be familiar with (maximal) height functions and flips; see Sections 3.1 and 5.2.

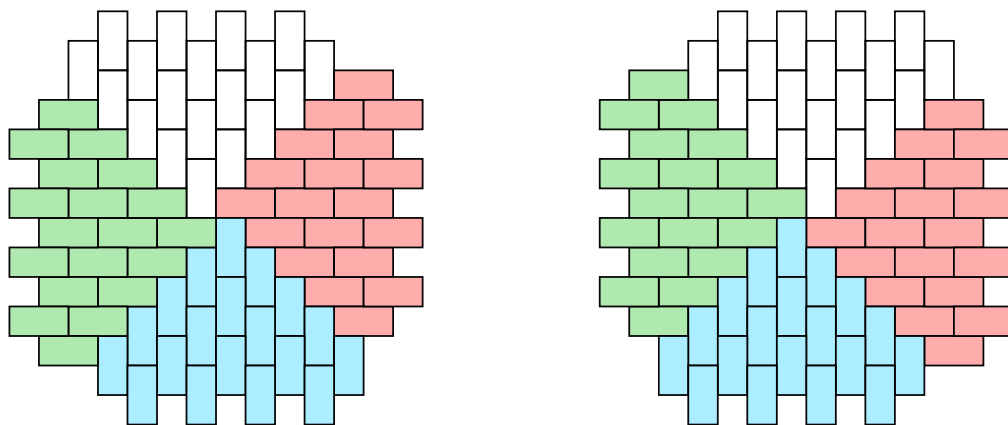
Theorem 2 (Characterization of tilings of the infinite square lattice). *Let t be a tiling of \mathbb{Z}^2 . Then exactly one of the following applies:*

1. t admits a flip;
2. t consists entirely of parallel, doubly-infinite domino staircases;



Examples of domino staircases

3. t is a windmill tiling.



Windmill tilings of \mathbb{Z}^2


The proof (and theory leading up to it) delves into properties of domino staircases and staircase edge-paths.

Tilings of the torus can be seen as periodic tilings of \mathbb{Z}^2 . Section 6.2 combines this observation with Theorem 2 to obtain relations between the

shape of a tiling and its flux. The final result is the above description of flip-connectedness on the torus. For a survey of flip-connectedness on more general surfaces, see Saldanha, Tomei, Casarin and Romualdo [16].

In Chapter 7, we go into detail about the construction of a Kasteleyn matrix for the torus. Some of its entries are monomials in $q_0, q_1, q_0^{-1}, q_1^{-1}$, or Laurent monomials in q_0, q_1 , so its determinant is a *Laurent polynomial* p_K in q_0, q_1 . We show that each monomial in p_K counts tilings with a flux given by the exponents of q_0, q_1 . Later, Chapter 9 will consider these variables as complex numbers on the unit circle. Moreover, Section 7.1 examines the structure of tilings with flux in the boundary of Q , primarily through a move called *stairflip*, that exchanges a doubly-infinite domino staircase by the only other one.

Chapter 8 elaborates on how the signs of monomials in p_K are assigned. The main tool here are *cycles* and *cycle flips*. Cycles are obtained by representing two tilings simultaneously, and cycle flips use them to go from one tiling to the other.




Cycles: one tiling has black dominoes, the other has blue dominoes

In the end of Section 8.3, we exhibit an odd-one-out pattern for signs over $\mathcal{F}(L)$, and use it to show that the total number of tilings can be given as a linear combination of $p_K(\pm 1, \pm 1)$ where each coefficient is either $\frac{1}{2}$ or $-\frac{1}{2}$.

Chapter 9 revisits techniques used in Section 3.3 and refines them for the calculation of Kasteleyn determinants of the torus. In Section 9.1, we examine the case of $M = KK^* \oplus K^*K$, for which we can compute all eigenvalues. Section 9.2 interprets K, K^* as linear maps on spaces of L -quasiperiodic functions, allowing us to exhibit bases for which they are diagonal. Studying the change of basis, we are able to relate the determinant of the original matrix to that of its diagonal version.

Finally, Section 9.3 makes explicit calculations on these determinants and investigates the effects of scaling L uniformly. This leads to our third and last theorem, which relates the Laurent polynomials (from Kasteleyn determinants) for L and nL by a simple product formula.



The odd-one-out sign pattern

Let $p_{[L,E]} : \mathbb{R}^2 \rightarrow \mathbb{C}$ be defined by $p_{[L,E]}(u_0, u_1) = \det(K_E(q_0, q_1))$, where K_E is the diagonal Kasteleyn matrix for L and $q_m = \exp(2\pi i \cdot u_m)$.

Theorem 3. For any positive integer n and reals u_0, u_1

$$p_{[nL,E]}(n \cdot u_0, n \cdot u_1) = \prod_{0 \leq i, j < n} p_{[L,E]} \left(u_0 + \frac{i}{n}, u_1 - \frac{j}{n} \right).$$

Intuitively, Theorem 3 says $\det(K[nL]_E(q_0, q_1))$ can be obtained from determinants of $K[L]_E$ by considering all n -th roots of q_0 and of q_1 . Product formulas of this kind have been encountered by Saldanha and Tomei [15] in their study of quadriculated annuli.

2

Definitions and Notation

This will be a short chapter detailing definitions and conventions used in the dissertation.

The imaginary unit will be denoted by the boldface \mathbf{i} .

A *lattice* is a subgroup of \mathbb{R}^2 that is isomorphic to \mathbb{Z}^2 and spans \mathbb{R}^2 (as a real vector space). An equivalent description is that a lattice L is the (additive) group of all integer linear combinations of a basis β of \mathbb{R}^2 ; in this case, we say L is generated by β . Notice different bases may generate the same lattice.

The *dual lattice* of L is $L^* = \text{Hom}(L, \mathbb{Z})$, the set of homomorphisms from L to \mathbb{Z} . Observe that, under addition, L^* is a group. Moreover, we may identify an element $f \in L^*$ with a unique $\tilde{f} \in \mathbb{R}^2$ via $f(v) = \langle \tilde{f}, v \rangle$ (for all $v \in L$). This allows us to see L^* as an additive subgroup of \mathbb{R}^2 , so that L^* is itself a lattice. Under this representation, it is easy to see that $(L^*)^* = L$. We will generally not make a distinction between f and \tilde{f} .

Given a basis $\beta = \{v_0, v_1\}$ of L , its *dual basis* is $\beta^* = \{v_0^*, v_1^*\}$, where $\langle v_i^*, v_j \rangle = \delta_{ij}$ ($0 \leq i, j \leq 1$). Geometrically, this means v_i^* is perpendicular to v_j and its length is determined by the equality $\langle v_i^*, v_i \rangle = 1$. It is a straightforward exercise to check that β^* generates L^* , and that $(\beta^*)^* = \beta$. We can also make explicit calculations; let $v_0 = (a, b)$ and $v_1 = (c, d)$. Then

$$\begin{aligned}v_0^* &= \frac{1}{ad - bc} \cdot (d, -c) \\v_1^* &= \frac{1}{ad - bc} \cdot (-b, a)\end{aligned}$$

Notice that because $\{v_0, v_1\}$ is a basis of \mathbb{R}^2 , $ad - bc$ is always nonzero, so the dual basis is well-defined.

A *fundamental domain* for a lattice L is a set $D \subset \mathbb{R}^2$ such that, for all $v \in \mathbb{R}^2$, the affine lattice $L + v$ intersects D exactly once. Another way to think of this is as follows: L acts on \mathbb{R}^2 by translation, so the orbit of any $v \in \mathbb{R}^2$ under L (that is, the set of images of v under L) is the affine lattice $L + v$. Hence, D contains exactly one point from each orbit: it is a visual realization

of the representatives of each orbit. It is easily seen that \mathbb{R}^2 is partitioned by the sets $D + v, v \in L$.

For a lattice generated by $\{v_0, v_1\}$, the fundamental domain is usually the parallelogram $\{s \cdot v_0 + t \cdot v_1 \mid s, t \in [0, 1)\}$, but we will generally prefer other kinds of fundamental domain (discussed in Section 4.2).

Consider the infinite square lattice \mathbb{Z}^2 . A *quadriculated region* is a union of (closed, filled) unit squares with vertices in \mathbb{Z}^2 . We say two squares are adjacent if they share an edge. A *domino* is a union of two adjacent unit squares, that is, a 2×1 rectangle with vertices in \mathbb{Z}^2 . A (*domino*) *tiling* of a quadriculated region is a collection t of dominoes on R with pairwise disjoint interiors and such that every unit square of R belongs to a domino in t .

A *torus* is a quotient \mathbb{R}^2 / L ; we may represent it by a fundamental domain of L whose boundary has appropriate identifications. If the fundamental domain is chosen to be a quadriculated region, we say the torus is a *quadriculated torus*. A tiling of a quadriculated torus is much like that of its fundamental domain, except dominoes account for boundary identifications. Alternatively, a tiling of a quadriculated torus is an L -periodic tiling of the infinite square lattice.

3

Domino tilings on the plane

Let R be a finite, simply-connected, quadriculated planar region. A *domino* is a 2×1 rectangle made of two unit squares. Is it possible to tile R entirely using only domino pieces? In how many ways can this be done?

For instance, if R has an odd number of squares, then there is no domino tiling of R . If R is the 2×3 rectangle below...




Figure 3.1: A 2×3 rectangle.

...then there are exactly three distinct domino tilings of R :




Figure 3.2: Domino tilings of the 2×3 rectangle.

The first observation is this problem can be converted to a dual problem on graph theory. This conversion associates to the region R a graph G (R 's *dual graph*) obtained by substituting each square of R by a vertex and joining neighboring vertices by an edge (horizontally and vertically, but not diagonally). On a domino tiling level, each domino corresponds to an edge on G : the edge joining the two vertices whose associated squares that are tiled by that domino.

For instance, the region R and the graph G in Figure 3.3 are dual. Likewise, the domino tiling of R and the G subgraph in Figure 3.4 are dual.

In this context, a question on domino tilings of R can be translated naturally into a question on the matchings of G . A *matching* M of a graph G




Figure 3.3: A quadriculated region R and its dual graph G .




Figure 3.4: A domino tiling of R and the corresponding G subgraph.

is a set of edges on G with no common vertex. If two vertices on G are joined by an edge of M , we say M matches those vertices. A *perfect matching* of a graph G is a matching of G that matches all vertices on G .

Now we may translate the opening questions: ‘Is it possible to tile R by dominoes?’ becomes ‘Is there a perfect matching of the dual graph G ?’; and ‘In how many ways can this be done?’ becomes ‘How many perfect matchings does the dual graph G have?’. Henceforth, unless explicitly stated, we shall use *matchings* when referring to perfect matchings. Non-perfect matchings do not interest us in this study.

The second observation is that these constructions lend themselves naturally to the concept of bipartite graphs. A graph G is *bipartite* if its vertices can be separated into two disjoint sets U and V so that every edge on G joins a vertex in U to a vertex in V . In this case, the sets U and V are called a *bipartition* of G . With this in mind, we may return to our initial problem and consider a prescribed ‘bipartition’ on R : we assign the label ‘black’ to an initial square, then assign the label ‘white’ to its neighbors, and so on in alternating fashion. Naturally, the vertices of the dual graph G inherits the labels.




Figure 3.5: A quadriculated region and its dual graph colored as above.

At this point, notice every domino in a tiling of R must be made of a single black square and a single white square. Hence, a **necessary condition** for R to admit a domino tiling is that the number of black squares and the number of white squares be equal. Observe, however, that it is not sufficient.




Figure 3.6: A region that satisfies the black-and-white condition but admits no domino tiling.

We point out that we will generally think of G as embedded on the region R , with each vertex lying on the center of its corresponding square and each edge a straight line.

3.1

Flips and height functions

We now introduce the concept of *flips*. To that end, notice a 2×2 square can be tiled by two dominoes in exactly two ways: by using both dominoes vertically, or by using both dominoes horizontally.

Consider two adjacent parallel dominoes forming a 2×2 square. A *flip* of these two dominoes consists in substituting the domino tiling of the square they form by the only other domino tiling of that same square. Naturally, the concept of flip is transferred to the graph treatment of the problem.




Figure 3.7: A flip on a region's tiling and on its corresponding graph.

Of course, given a domino tiling of a planar region R , the execution of a flip takes us to a new domino tiling of R . Following this train of thought, a natural question might be whether two given domino tilings of R can be joined by a sequence of flips. To answer this question, we will investigate the **height function** h of a domino tiling t of R .

We highlight the distinction between an edge on a graph G and an edge on a quadriculated region R : the latter refers to an edge on the boundary of a square on R . Similarly, an edge on a domino tiling t of R is an edge on R (of a square, not of a domino) that does not cross a domino (it has not been ‘erased’ to produce said domino).

We choose once and for all the clockwise orientation for black squares; the other orientation is assigned to white squares. This choice induces an orientation on each edge on R . Notice it is consistent: along an edge where two squares meet, each square will have a different orientation and thus the orientations induced on the edge will agree.

Now, choose a base vertex v on R and assign an integer value to it; we will always choose a base vertex in the boundary ∂R of the region R and we will always assign the value 0 to it. This is the value h takes on v . We now propagate that value across all vertices of R as follows. For each vertex w joined to v by an edge on t , that edge may point from v to w or from w to v , depending on its orientation as defined above. In the first case, w is assigned the integer value $h(v) + 1$; otherwise, it is assigned the integer value $h(v) - 1$.

By connectivity, this process defines the height function on each vertex of R , but it may not be clear whether or not the definition is consistent. It’s easy to verify consistency on a single domino, as the image below shows.




Figure 3.8: Height consistency for horizontal dominoes.

Consistency for a general simply-connected planar region R can be proved as follows: starting from a vertex on which h is well-defined (for instance, the base vertex v), suppose we wish to check consistency on another vertex, say w . Consider then two different edge-paths γ_0 and γ_1 on t joining those vertices; these paths can be seen as the boundary of a region tiled by dominos. The area of that region is thus well-defined. Now, incrementally deform γ_0 onto γ_1 , with

each step producing a new region with less area than the previous one through the removal of a domino. Here, consistency on a single domino ensures each step is consistent with the previous one. Finally, the simply-connectedness of R guarantees this process can fully deform γ_0 onto γ_1 .

We provide a simple example of this process below.




Figure 3.9: Deforming one edge-path into another.

With these conventions, given a domino tiling t and a base vertex v of a black-and-white quadriculated region R , the height function h of t is well-defined. An example of height function h can be seen in the following image; the marked vertex is the base vertex.




Figure 3.10: An example of height function.

This provides a constructive definition of height functions, but we highlight now some of their properties.

Proposition 3.1.1. *Let R be a black-and-white quadriculated region. Fix a base vertex $v \in \partial R$ (independent of choice of tiling). Then (1) the values a height function takes on ∂R and (2) the mod 4 values a height function takes on all of R are all independent of choice of tiling.*

Proof. Remember that, regardless of the choice of tiling t , an edge on ∂R is an edge on t . Since we have already proved consistency, (1) is automatic.

For (2), let u and w be vertices on R joined by an edge e . Notice the orientation of e depends only on the region R and not on choice of tiling; assume then that e is oriented from u to w . The constructive definition implies a change in height function along e occurs in one of the following ways:

- If e is on the tiling t , then $h(w) = h(u) + 1$.
- If e is not on the tiling t , then $h(w) = h(u) - 3$.

Notice in both cases $h(w)$ has the same mod 4 value. The same occurs when e is oriented from w to u . By connectivity, we are done. \square

Proposition 3.1.1 allows us to fully characterize height functions of tilings of a region R .

Proposition 3.1.2 (Characterization of height functions). *Let R be a black-and-white quadriculated region. Fix a base vertex $v \in \partial R$. Then an integer function h on the vertices of R is a height function (of a tiling of R) if and only if h satisfies the following properties:*

1. h has the prescribed values on ∂R .
2. h has the prescribed mod 4 values on all of R .
3. h changes by at most 3 along an edge on R .

Proof. Proposition 3.1.1 and its proof guarantee that any height function satisfies the listed properties. We will now show that if an integer function on the vertices of R satisfies those properties, it is the height function of a tiling on R . To that end, we will construct a tiling t that realizes one such function h .

On R , whenever two vertices joined by an edge have h -values that differ by 3, erase that edge (thus producing a domino). We claim the result is a

domino tiling t on R . Indeed, properties (2) and (3) ensure each square on R will have exactly one of its sides erased. Furthermore, by (1) that side will never occur on ∂R . It's easy to see this yields a domino tiling of R ; furthermore, by construction this tiling's height function is h . \square

From now on, for any black-and-white quadriculated region R , assume the base vertex v is fixed independently of choice of tiling.

Another interesting and perhaps less obvious property of height functions is that the minimum of two height functions is itself a height function.

Proposition 3.1.3. *Let R be a black-and-white quadriculated region and t_1, t_2 be two domino tilings of R with corresponding height functions h_1, h_2 . Then $h_m = \min\{h_1, h_2\}$ is a height function on R .*

Proof. Indeed, by Proposition 3.1.2, it suffices to show that h_m changes by at most 3 along an edge on R . This is trivially verified on vertices v and w joined by an edge whenever $h_m = h_1$ or $h_m = h_2$ on both v and w . Suppose this is not the case; furthermore, suppose without loss of generality $h_m(v) = h_1(v)$, $h_m(w) = h_2(w)$ and that the edge joining them points from v to w .

The edge's orientation implies $h_i(w) = h_i(v) + 1$ if the edge is on t_i and $h_i(w) = h_i(v) - 3$ otherwise ($i = 1, 2$). Since $h_1(v) < h_2(v)$, the only possibility that realizes $h_2(w) < h_1(w)$ is the edge being on t_1 and not on t_2 , so that $h_1(w) = h_1(v) + 1$ and $h_2(w) = h_2(v) - 3$. Now, because $h_1(v) - h_2(v) < 0$ and mod 4 values are prescribed, the difference must be $-4k$ for some positive integer k , so that $h_1(v) = h_2(v) - 4k$.

Finally, $h_2(w) < h_1(w)$ can now be rewritten as $h_2(v) - 3 < (h_2(v) - 4k) + 1$, or simply $-4 < -4k$. This is a contradiction, implying only the cases when $h_m = h_1$ or $h_m = h_2$ on both v and w can occur. \square

Corollary 3.1.4 (Minimal height function). *Let R be a black-and-white quadriculated region. If R can be tiled by dominoes, then there is a minimal height function.*

Along a 2×2 square tiled by dominoes, it's easy to verify that height function values are distributed so that the center vertex is a local maximum or minimum. Furthermore, applying a flip changes a local maximum vertex to a local minimum vertex, and vice-versa, leaving other values unchanged. Figure 3.11 illustrates this phenomenon.

Together with Corollary 3.1.4, an application of this technique provides the following result.




Figure 3.11: The effect of a flip on the height function.

Proposition 3.1.5. *Let R be a black-and-white quadriculated region with minimal height function h_m . Let $h \neq h_m$ be a height function associated to the domino tiling t of R . Then there is a flip on t that produces a height function $\tilde{h} \leq h$ with $\tilde{h} < h$ on one vertex of R .*

Proof. Consider the difference $h - h_m$. By Proposition 3.1.2, it is 0 along the boundary and takes nonnegative values on $4\mathbb{Z}$. Let V be the set of vertices of R on which $h - h_m$ is maximum, and choose a vertex $v \in V$ that maximizes h . Notice by hypothesis V is non-empty, and does not intersect ∂R . We assert that v is a local maximum of h .

Suppose v were not a local maximum of h , that is, suppose there were a vertex w joined to v by an edge e so that $h(w) > h(v)$. There are two cases:

1. e is on t and points from v to w , so that $h(w) = h(v) + 1$.
2. e is not on t and points from w to v , so that $h(w) = h(v) + 3$.

Remember edge orientation does not depend on choice of tiling (and thus does not depend on the height function considered).

In case (1), $h_m(w) = h_m(v) + 1$ if e is on the associated minimal tiling t_m , and $h_m(w) = h_m(v) - 3$ otherwise. Neither can occur: the first contradicts v maximizing h (since $h(w) > h(v)$), and the latter contradicts v maximizing $h - h_m$ (since $h(w) - h_m(w) > h(v) - h_m(v)$).

Case (2) is similar: $h_m(w) = h_m(v) - 1$ if e is on t_m , and $h_m(w) = h_m(v) + 3$ otherwise. The first contradicts v maximizing $h - h_m$, and the latter contradicts v maximizing h .

Whatever the situation, we derive a contradiction, implying v must indeed be a local maximum. Since v is not on ∂R , we can perform a flip round v . This makes it a local minimum while preserving the values h takes on all other vertices of R and completes the proof. \square

Because the situation is finite, Proposition 3.1.5 essentially tells us any tiling of a region R can be taken by a sequence of flips to the tiling that minimizes height functions over tilings of R . A simple but important corollary follows.


Corollary 3.1.6 (Flip-connectedness). *Let R be a black-and-white simply-connected quadriculated region tileable by dominoes. Then any two distinct tilings of R can be joined by a sequence of flips.*

3.2

Kasteleyn matrices

A *Kasteleyn matrix* ‘encodes’ a quadriculated black-and-white region R in matrix form, and its construction is similar to that of adjacency matrices.

Given one such region R , we can obtain an adjacency matrix A of R from its dual graph G as follows: enumerate each black vertex (starting from 1), and do the same to white vertices. Then $A_{ij} = 1$ if the i -th black vertex and j -th white vertex are joined by an edge, and 0 otherwise.



$$A = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix}$$

Figure 3.12: The construction of an adjacency matrix A for a quadriculated, colored region.

Consider now an $n \times n$ adjacency matrix A and the combinatorial expansion of its determinant:

$$\det(A) = \sum_{\sigma \in S_n} \text{sgn}(\sigma) \prod_{i=1}^n A_{i,\sigma(i)} \tag{3-1}$$

In the expansion above, each nonzero term of the form $\prod_{i=1}^n A_{i,\sigma(i)}$ can be seen as corresponding to a matching of G . In fact, the term is nonzero if and

only if each factor in the product is 1, in which case the i -th black vertex is joined by an edge to the $\sigma(i)$ -th white vertex. Since σ is a permutation on $\{1, \dots, n\}$, the collection of these edges is by construction a set of edges on G in which each vertex of G features exactly once. The observation follows.

Of course, the correspondence goes both ways. This means that, except for $\text{sgn}(\sigma)$, $\det(A)$ counts the number of matchings of G (and thus also the domino tilings of R). How do we get past the sign?

The obvious way would be to consider the *permanent* of A

$$\text{perm}(A) = \sum_{\sigma \in S_n} \prod_{i=1}^n A_{i, \sigma(i)},$$

but permanents lack a number of interesting properties when compared to determinants, and are also much more costly to compute.

The answer is precisely the Kasteleyn matrix K : an altered adjacency matrix in which some entries are replaced by -1 . Its construction is similar to the ordinary adjacency matrix, except some edges on G are assigned the value -1 rather than $+1$. This distribution of minus signs can be done in many ways, but the following observation explains the general principle behind it: a flip on a matching of G always changes the sign of the corresponding permutation in (3-1). This is because, on a permutation level, applying a flip amounts to multiplying the original permutation by a cycle of length 2.

With this in mind, the distribution of minus signs over edges on G is made so that the sign change in a permutation caused by a flip is always counterbalanced by a sign change on the corresponding product of entries of K . Such a distribution ensures that applying a flip does not change the ‘total’ sign of the term

$$\text{sgn}(\sigma) \prod_{i=1}^n K_{i, \sigma(i)}$$


in (3-1). And since we’ve shown that any two distinct domino tilings of R (and thus matchings of G) can be joined by a sequence flips, this means the sum in (3-1) is carried over identically signed numbers. In other words, for a Kasteleyn matrix K of a region R , $|\det(K)|$ is the number of domino tilings of R .

An easy, convenient way of distributing minus signs over edges on G is assigning them to all horizontal edges in alternating lines (say, all odd lines, or all even lines). This way, a 2×2 square in the dual graph will always contain exactly one negative edge (either the topmost or the bottommost horizontal

line), so that a flip always will always produce a sign change on the product of entries of K .

We highlight that in his original paper [7], flip-connectedness (or more generally, flips) was not a part of Kasteleyn's exposition. His methods were combinatorial but he employed Pfaffians.

Below, we show an example of construction of a Kasteleyn matrix. In the corresponding dual graph, negative edges are red and dashed.



$$K = \begin{pmatrix} -1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -1 \end{pmatrix}$$

Figure 3.13: The construction of a Kasteleyn matrix K .

3.3

A classical result: domino tilings of the rectangle

We will end this chapter by using our methods to provide a classical result: the counting of domino tilings of an $m \times n$ black-and-white rectangular region, $R_{m,n}$. Of course, if both m and n are odd, that number is 0; we assume then that m is even.

Let G be $R_{m,n}$'s dual graph with minus signs assigned to all horizontal edges in even lines, from which we obtain the corresponding Kasteleyn matrix K . Rather than compute the determinant of K , we will consider the matrix $M = KK^* \oplus K^*K$; it's clear that $\det(M) = \det(K)^4$.

M can be seen as a double adjacency matrix of G , acting as a linear map on the space of formal linear combinations of vertices. It takes a vertex v to the sum of vertices that are joined to v via an edge-path of length two on the graph. Notice edge sign and vertex multiplicity (when a vertex can be reached from v via two distinct edge-paths) are taken into account.

Because edge-paths considered have length two, M takes white vertices to white vertices and black vertices to black vertices. Another way of thinking this is as follows: when interpreting the Kasteleyn matrix as a linear map (like M above), our general construction of the Kasteleyn matrix implies it goes from the space of white vertices W to the space of black vertices B . Of course, this also means $K^* : B \rightarrow W$. It then becomes clear by the definition of M that it is a color-preserving map. This essentially means M acts independently on B and W .

Consider the grid below, so that each vertex of G is identified by a double index (i, j) .




Figure 3.14: Indexation grid.

In the obvious notation, vertices $v_{i,j}$ of G at least two units away from the boundary satisfy $Mv_{i,j} = 4v_{i,j} + v_{i+2,j} + v_{i-2,j} + v_{i,j+2} + v_{i,j-2}$. The coefficient in $v_{i,j}$ comes from moving forward then backwards in each cardinal direction; notice that a negative edge traversed this way will account for two minus signs, so the end result is always positive. Vertices of the form $v_{i\pm 1, j\pm 1}$ do not feature because each of them can be reached via exactly two distinct edge-paths with necessarily opposite signs.

The formula can be extended to all vertices of G as follows. Put $v_{i,j} = 0$ if it is immediately outside the boundary of G ; then, for each line of zero-vertices, reflect G through that line and set a vertex $v_{i,j}$ obtained this way to be **minus** the vertex from which it was reflected. We then repeat this process, so that in the end $v_{i,j}$ will be defined for all $i, j \in \mathbb{Z}$.

Figure 3.15 is a visual representation of this extension; in it, each gray square is a zero vertex. More generally, one such extension can be succinctly represented by the relations:

$$\begin{cases} v_{0,j} = v_{m+1,j} = v_{i,0} = v_{i,n+1} = 0 \\ v_{-i,j} = v_{i,-j} = -v_{i,j} \\ v_{i+2m+2,j} = v_{i,j+2n+2} = v_{i,j} \end{cases} \quad (3-2)$$

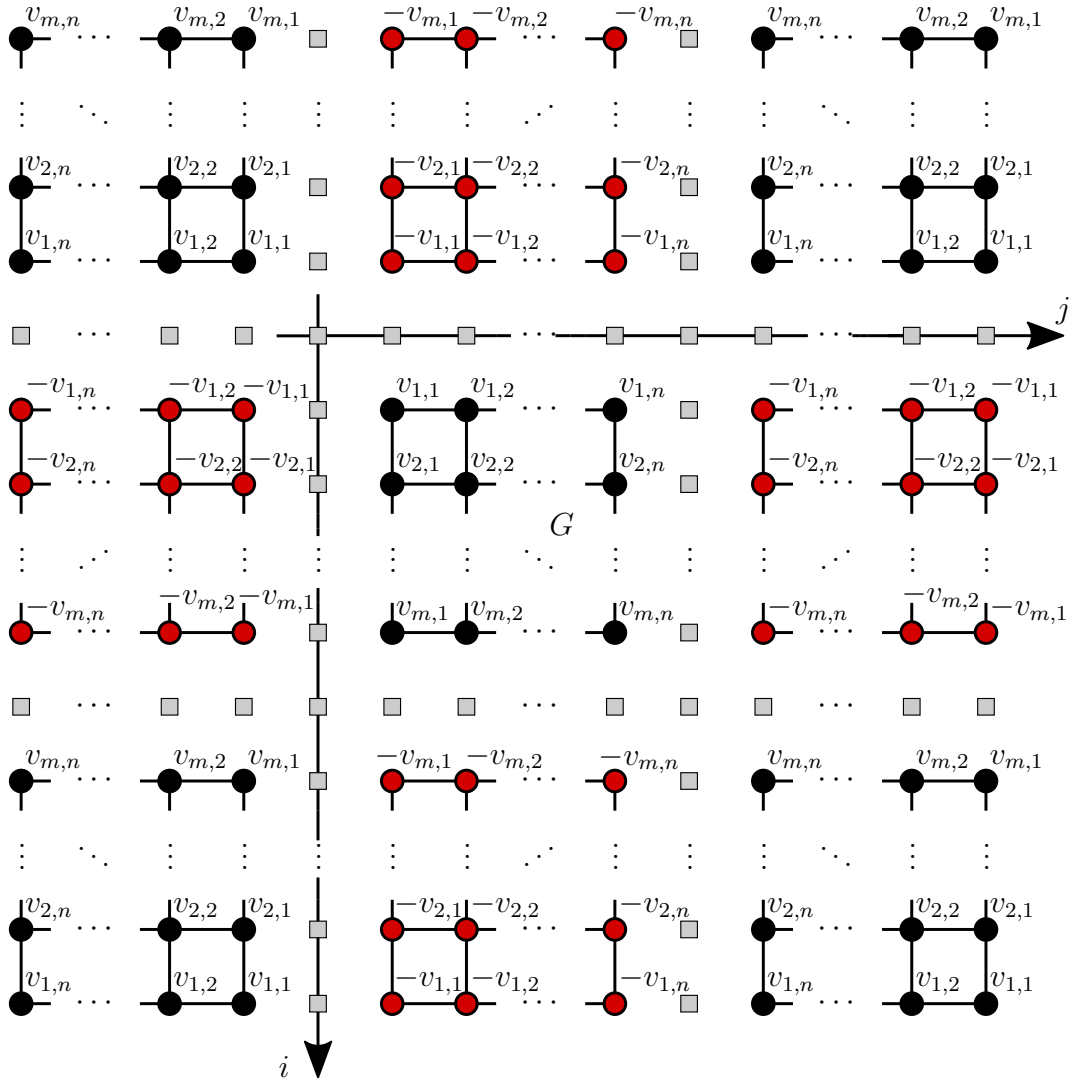


Figure 3.15: Extension through reflections.

With this, the formula for M now holds not only on all of G , but also on all of \mathbb{Z}^2 . Notice the space on which M acts is still mn -dimensional, since coordinates on the original vertices of G propagate to all vertices of \mathbb{Z}^2 via the relations above. We will now compute $\det(M)$.

We will always refer to the imaginary unit by the boldface \mathbf{i} . Let $\zeta_1 = \exp\left(\frac{\pi \mathbf{i}}{m+1}\right)$ and $\zeta_2 = \exp\left(\frac{\pi \mathbf{i}}{n+1}\right)$. For all $k, l \in \mathbb{Z}$ with $1 \leq k \leq m$ and $1 \leq l \leq n$, let

$$\begin{aligned} v(k, l) &= \sum_{i,j} 4 \sin\left(\frac{i \cdot k \cdot \pi}{m+1}\right) \sin\left(\frac{j \cdot l \cdot \pi}{n+1}\right) \cdot v_{i,j} \\ &= \sum_{i,j} (\zeta_1^{ik} - \zeta_1^{-ik}) (\zeta_2^{jl} - \zeta_2^{-jl}) \cdot v_{i,j} \end{aligned}$$

Notice $v(k, l)$ is a valid vector in that its coordinates respect the relations in (3-2). We claim each such $v(k, l)$ is an eigenvector of M . Indeed, we have

that the $v_{i,j}$ -coordinate of $Mv(k, l)$ is

$$\begin{aligned} & 4 (\zeta_1^{ik} - \zeta_1^{-ik}) (\zeta_2^{jl} - \zeta_2^{-jl}) \\ & + (\zeta_1^{(i+2)k} - \zeta_1^{-(i+2)k}) (\zeta_2^{jl} - \zeta_2^{-jl}) + (\zeta_1^{(i-2)k} - \zeta_1^{-(i-2)k}) (\zeta_2^{jl} - \zeta_2^{-jl}) \\ & + (\zeta_1^{ik} - \zeta_1^{-ik}) (\zeta_2^{(j+2)l} - \zeta_2^{-(j+2)l}) + (\zeta_1^{ik} - \zeta_1^{-ik}) (\zeta_2^{(j-2)l} - \zeta_2^{-(j-2)l}) \end{aligned}$$

The expression above may be rearranged to yield:

$$\begin{aligned} [Mv(k, l)]_{i,j} &= (\zeta_1^{ik} - \zeta_1^{-ik}) (\zeta_2^{jl} - \zeta_2^{-jl}) \left((\zeta_1^k + \zeta_1^{-k})^2 + (\zeta_2^l + \zeta_2^{-l})^2 \right) \\ &= [v(k, l)]_{i,j} \cdot \left((\zeta_1^k + \zeta_1^{-k})^2 + (\zeta_2^l + \zeta_2^{-l})^2 \right). \end{aligned}$$

In other words, $v(k, l)$ is an eigenvector of M with associated eigenvalue

$$\begin{aligned} \lambda(k, l) &= \left((\zeta_1^k + \zeta_1^{-k})^2 + (\zeta_2^l + \zeta_2^{-l})^2 \right) \\ &= 4 \cdot \left(\cos \left(\frac{k\pi}{m+1} \right) \right)^2 + 4 \cdot \left(\cos \left(\frac{l\pi}{n+1} \right) \right)^2 \end{aligned}$$

It is well known from Fourier series that the vectors $v(k, l)$ are linearly independent (and in fact orthogonal), so indeed these are all the mn eigenvalues of M . Hence, $\det(M) = \prod_{k=1}^m \prod_{l=1}^n \lambda(k, l)$.

Since m is even, for all $k \in \mathbb{Z}$ with $1 \leq k \leq m/2$ it holds that $\cos \left(\frac{k\pi}{m+1} \right) = -\cos \left(\pi - \frac{k\pi}{m+1} \right)$. Because in $\lambda(k, l)$ the cosines are squared, this implies $\lambda(k, l) = \lambda(m+1-k, l)$ whenever $1 \leq k \leq m/2$. Thus, we may write $\det(M) = \prod_{k=1}^{m/2} \prod_{l=1}^n \lambda(k, l)^2$.

Finally, remember that $\det(M) = \det(K)^4$, so the number of tilings of $R_{m,n}$ is

$$|\det(K)| = \prod_{k=1}^{m/2} \prod_{l=1}^n 2 \left(\left(\cos \left(\frac{k\pi}{m+1} \right) \right)^2 + \left(\cos \left(\frac{l\pi}{n+1} \right) \right)^2 \right)^{1/2}$$

4

Domino tilings on the torus

Domino tilings on the torus (or more generally planar surfaces) are amenable to much of what's been discussed in the plane case, but modifications are unavoidable. Let $D_n \subset [0, 2n]^2$ be a $2n \times 2n$ black-and-white quadrilaterated square region with the square $[0, 1]^2$ fixed as black: initially, this will be our model fundamental domain for studying the torus. The torus arises from identifying opposite sides on D_n ; we will refer to a torus obtained this way by \mathbb{T}_n . Notice D_n respects the black-and-white condition. Questions on the existence and number of domino tilings for \mathbb{T}_n will receive treatment similar to that of ordinary planar regions.

We first introduce the concept of *flux* of a tiling. Our torus model is very similar to the ordinary square planar region, except for the identification of opposing sides. This identification enables us to represent tilings of \mathbb{T}_n on D_n with dominoes that 'cross over' to the opposing side. For instance, the 4×4 model of the torus admits tilings such as those represented on D_2 below.




Figure 4.1: Example tilings of \mathbb{T}_2 with 'cross over' dominoes.

Notice these new tilings also introduce new 'cross-flips': we can flip 'cross over' dominoes just like ordinary dominoes, in the obvious way. In fact, cross-flips are possible in the first two tilings above.

In a loose sense, the flux of a tiling counts these 'cross-over' dominoes. We now describe this in detail.

Remember a domino is always made up of exactly one black square and one white square. We fix a positive orientation for dominoes: from their black

square to their white square. This orientation applies to cross-over dominoes in the obvious way, and is also inherited by edges on the dual graph in the natural way. We also fix positive horizontal and vertical orientations for the fundamental domain D_n itself: from left to right and from top to bottom.


Let t be a tiling of the torus \mathbb{T}_n on D_n . The *horizontal* flux of t is the number of vertical cross-over dominoes on t , counted positively if their orientation agrees with D_n 's own vertical orientation, and negatively otherwise. Similarly, the *vertical* flux of t is the number of horizontal cross-over dominoes on t , counted positively if their orientation agrees with D_n 's own horizontal orientation, and negatively otherwise. The definitions may appear to be a mismatch, but the reason for them should become clear in time.

This rather algebraic definition interacts well with flips, in a manner similar to how a signed region does: a flip always preserves the flux of a tiling. Indeed, an ordinary flip does not involve cross-over dominoes at all, while a cross-flip replaces a count of $+1$ and -1 by a count of two zeroes. On the other hand, this also implies our space of tilings is no longer flip-connected: two tilings with different flux values can never be joined by a sequence of flips! Consequently, we cannot expect to achieve good results from naïvely applying our Kasteleyn matrix method to the torus model. Accordingly, we will modify the construction of our Kasteleyn matrix K to account for the flux of a tiling.

\mathbb{T}_n 's dual graph G is the same as D_n 's dual graph, except it features new 'crossing edges'. Like in the planar case, enumerate each black vertex of G (starting from 1), and do the same to white vertices. We will assign weights to each edge on G . Let e_{ij} be the edge joining the i -th black vertex to the j -th white vertex: if no such edge exists, we assign the weight 0 to it; otherwise, it's assigned the weight $+1$. Next, distribute minus signs over edges on G like before (taking crossing edges into account); for each edge that's assigned a minus sign this way, multiply its weight by -1 .

Finally, consider vertical crossing edges. For each of those, if its orientation agrees with the graph's own vertical orientation, multiply its weight by q_0 ; otherwise, multiply its weight by q_0^{-1} . Similarly, consider horizontal crossing edges. For each of those, if its orientation agrees with the graph's own horizontal orientation, multiply its weight by q_1 ; otherwise, multiply its weight by q_1^{-1} . Notice the effect of minus signs and crossing edges is cumulative!

Now that all edges are assigned their corresponding weights, $K(i, j)$ is simply the weight of e_{ij} . Below, we have the dual graph of \mathbb{T}_2 , where dashed red edges are assigned the minus sign, and the corresponding Kasteleyn matrix.



$$K = \begin{pmatrix} 1 & q_1^{-1} & 1 & 0 & 0 & 0 & q_0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & 0 & q_0 \\ 1 & 0 & -1 & -1 & 1 & 0 & 0 & 0 \\ 0 & 1 & -q_1 & -1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & q_1^{-1} & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 1 \\ q_0^{-1} & 0 & 0 & 0 & 1 & 0 & -1 & -1 \\ 0 & q_0^{-1} & 0 & 0 & 0 & 1 & -q_1 & -1 \end{pmatrix}$$

Figure 4.2: The construction of a Kasteleyn matrix for the dual graph of \mathbb{T}_2 .


It's clear that $\det(K)$ is not a number, but rather a Laurent polynomial P_K in q_0 and q_1 . For instance, the Kasteleyn matrix above yields the polynomial

$$132 - 32 \cdot (q_0 + q_0^{-1} + q_1 + q_1^{-1}) - 2 \cdot (q_0 q_1 + q_0^{-1} q_1 + q_0 q_1^{-1} + q_0^{-1} q_1^{-1}) + q_0^2 + q_0^{-2} + q_1^2 + q_1^{-2}$$


Much further ahead, Proposition 7.0.8 will show that each monomial of the form $c_{i,j} \cdot q_0^i q_1^j$ in $\det(K)$ counts the number of tilings of \mathbb{T}_n with horizontal flux i and vertical flux j : that number is precisely the modulus of c . Notice i or j may be negative. The total number of tilings is thus $\sum |c_{i,j}|$, but Proposition 8.3.5 will show it can be computed as a suitable linear combination of $P_K(\pm 1, \pm 1)$, dependent on the assignment of negative edges. For instance, the linear combination for P_K above is $1/2 \cdot [-P_K(1, 1) + P_K(-1, 1) + P_K(1, -1) + P_K(-1, -1)]$.

We provide all the possible tilings of \mathbb{T}_2 , represented as matchings of its dual graph, and grouped according to flux values. Compare them with the polynomial P_K above. Notice we omit the vertices at the endpoints of each edge.


Flux = (0, 0)




Flux = (0, 1)



Flux = (0, -1)



Flux = (1, 0)



PUC-Rio - Certificação Digital N° 1321816/CA

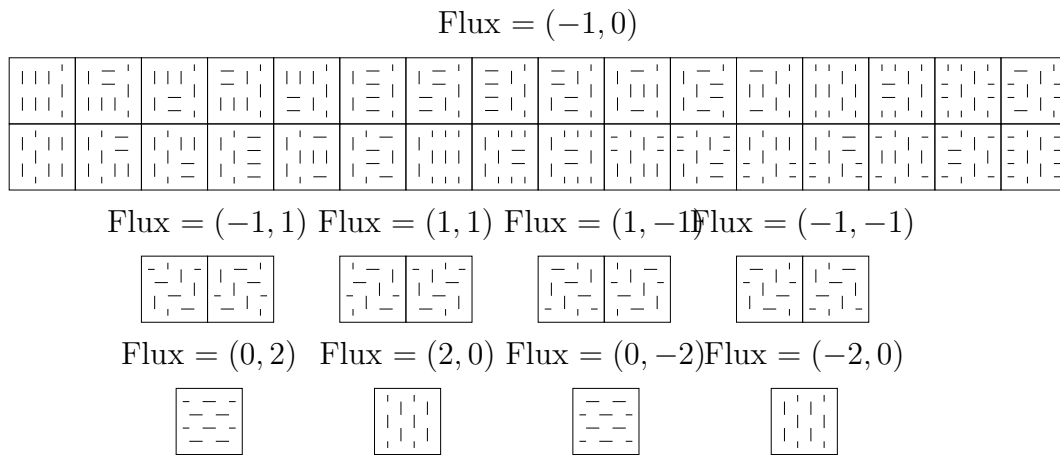


Figure 4.3: Tilings of \mathbb{T}_2 , 272 in total.

4.1

Height functions on the torus

We could attempt to define height functions constructively on \mathbb{T}_n as we did before, but since the torus is not simply-connected we will generally find the process results in inconsistencies. Remember opposite sides are identified, so that corresponding vertices on opposite sides should have the same value. The following image provides an example of a tiling of \mathbb{T}_2 on D_2 with an associated ‘naïve’ height function; notice its values on corresponding vertices do *not* agree.




Figure 4.4: A ‘wrong’ height function for a tiling of \mathbb{T}_2 .

Instead, we will change our methods. We will interpret \mathbb{T}_n as the quotient \mathbb{R}^2 / L , where $L \subset \mathbb{Z}^2$ is the lattice generated by $\{(2n, 0), (0, 2n)\}$; notice we can still take D_n as its fundamental domain. Consider the projection map $\Pi : \mathbb{R}^2 \rightarrow \mathbb{T}_n$. If base points are provided on each of \mathbb{T}_n and \mathbb{R}^2 , any tiling of

\mathbb{T}_n can be lifted by Π to a tiling of the infinite square lattice \mathbb{Z}^2 in the obvious way (lift colors too!). Given a tiling t of \mathbb{T}_n on D_n , we will always choose the point $(\frac{1}{2}, \frac{1}{2})$ as base point for both \mathbb{T}_n and \mathbb{R}^2 . Notice this choice guarantees the fiber over every vertex of a square on D_n will consist of points in \mathbb{Z}^2 , that the square $[0, 1]^2 \subset \mathbb{R}^2$ will be colored black, and also that D_n and t will be lifted to an exact copy on the square $[0, 2n]^2 \subset \mathbb{R}^2$.

Because identifications and the black-and-white condition are respected, it's easy to see the end result is an L -periodic domino tiling \tilde{t} of \mathbb{Z}^2 . Since \mathbb{Z}^2 is simply-connected, we can define the height function of \tilde{t} as before to be a function $\tilde{h} : \mathbb{Z}^2 \rightarrow \mathbb{Z}$. This will not always be the case, but unless stated otherwise, consider the base vertex as lying on the origin with base value 0 assigned to it.

Finally, we define the height function h of t to be the height function \tilde{h} of \tilde{t} . We provide an example of this construction below, using the tiling of Figure 4.4. The marked vertex is the origin.

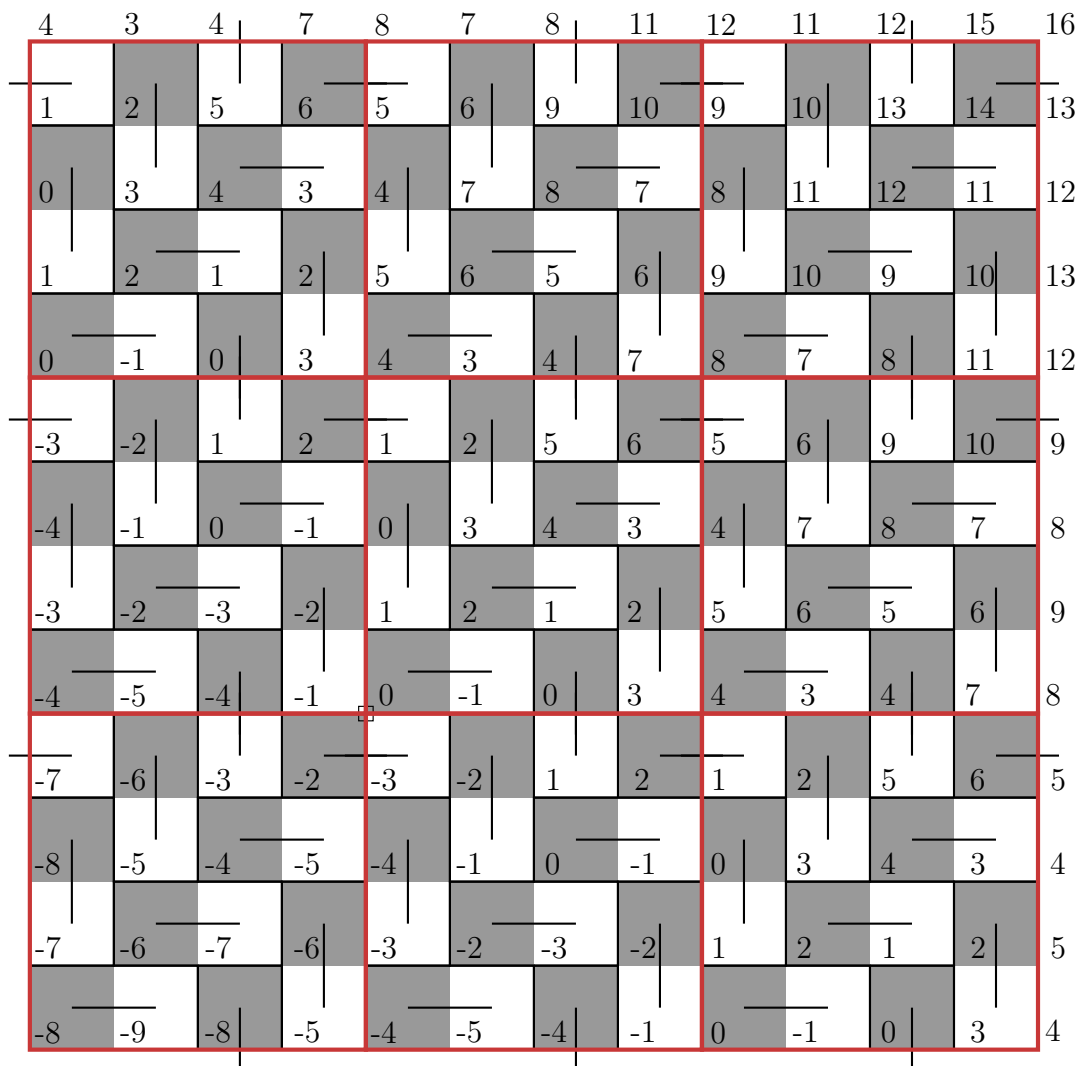


Figure 4.5: A proper height function for a tiling of \mathbb{T}_2 , defined on \mathbb{Z}^2 .

Inspecting the proofs of Propositions 3.1.1 and 3.1.2, Proposition 4.1.1 follows immediately.

Proposition 4.1.1 (Height functions on the infinite square lattice). *Let \mathbb{Z}^2 be the black-and-white infinite square lattice with $[0, 1]^2 \subset \mathbb{R}^2$ colored black. Then an integer function h on \mathbb{Z}^2 is a height function if and only if it satisfies the following properties:*

1. h has the prescribed mod 4 values on \mathbb{Z}^2 .
2. h changes by at most 3 along an edge on \mathbb{Z}^2 .

Furthermore, when h has base vertex lying on the origin with base value 0, then $h(0) = 0$.

The mod 4 prescription function $\Phi : \mathbb{Z}^2 \rightarrow \{0, 1, 2, 3\}$ is easily computed for the usual choice of base vertex $h(0) = 0$. It is given by

$$\Phi(x, y) = \begin{cases} 0 & \text{if } x \equiv 0 \text{ and } y \equiv 0 \pmod{2}; \\ 1 & \text{if } x \equiv 0 \text{ and } y \equiv 1 \pmod{2}; \\ 2 & \text{if } x \equiv 1 \text{ and } y \equiv 1 \pmod{2}; \\ 3 & \text{if } x \equiv 1 \text{ and } y \equiv 0 \pmod{2}. \end{cases}$$

It is important to point out that while Proposition 4.1.1 characterizes general height functions on the infinite square lattice \mathbb{Z}^2 , a function satisfying these properties may not be one obtained from a domino tiling of the torus. General domino tilings of the infinite square lattice need not have any kind of periodicity.

Consider a domino tiling t of \mathbb{T}_n on D_n and its height function h . The flux of t has a very concrete manifestation in h which we now describe. Follow the leftmost side of the tiled copy of $D_n \subset \mathbb{R}^2$ on $[0, 2n]^2$, edge by edge and starting from the origin, until the opposite horizontal side is reached. Notice how h changes along this path according to whether or not that edge is on t : whenever an edge with a black square on its right is traversed, h changes by +1 if that edge is on t and by -3 otherwise. Similarly, whenever an edge with a white square on its right is traversed, h changes by -1 if that edge is on t and by +3 otherwise.

Now, observe that in this situation an edge being on t means there is no horizontal cross-over domino along it, and an edge **not** being on t means there **is** a horizontal cross-over domino along it. It's then easy to see that whenever we have a -3 change on h along that path, the vertical flux changes by -1 ,

and whenever we have a +3 change, the vertical flux changes by +1. If every edge on that path were on t , the total change on h along it would be 0, and it would correspond to a vertical flux value of 0. This analysis thus makes it clear that in general the total change is a value $4k$ with $k \in \mathbb{Z}$, corresponding to a vertical flux value of k .

Let $u_y = (0, 2n) \in \mathbb{Z}^2$. The conclusion of the preceding paragraphs can be succinctly expressed as

$$h(u_y) = 4k \iff t \text{ has a vertical flux value of } k, \quad (4-1)$$

where $k \in \mathbb{Z}$. This analysis considered only the path along the leftmost side of the original copy of D_n , starting from the base vertex lying on the origin, but in fact it can be made more general. We claim that for **any** vertex v of \mathbb{Z}^2 it holds that

$$h(v + u_y) - h(v) = 4k \iff t \text{ has a vertical flux value of } k \quad (4-2)$$

Indeed, because (4-1) means that (4-2) holds when $v = 0$, it suffices to show that for any two vertices v, w of \mathbb{Z}^2 we have $h(v + u_y) - h(v) = h(w + u_y) - h(w)$, or equivalently $h(v + u_y) - h(w + u_y) = h(v) - h(w)$. Choose any edge-path γ_0 in \tilde{t} joining v to w . Because of how \tilde{t} is obtained from t (via lifting), the translated path $\gamma_1 = \gamma_0 + u_y$ is an edge-path in \tilde{t} joining $v + u_y$ to $w + u_y$, and furthermore the constructive definition of height functions implies the total change of h along γ_0 or along γ_1 is the same. This proves (4-2).

Equation (4-2) means a vertical flux value k of a domino tiling t manifests in its height function h as the (arithmetic) quasiperiodicity relation $h(v + u_y) = h(v) + 4k$. The same techniques used above show that something very similar holds for the horizontal flux.

Let $u_x = (2n, 0) \in \mathbb{Z}^2$. Then for **any** vertex v of \mathbb{Z}^2 it holds that

$$h(v + u_x) - h(v) = 4l \iff t \text{ has a horizontal flux value of } l \quad (4-3)$$

Of course, equation (4-3) means a horizontal flux value l of a domino tiling t manifests in its height function h as the quasiperiodicity relation $h(v + u_x) = h(v) + 4l$. These relations allow us to fully characterize *toroidal height functions*, that is, height functions on the infinite square lattice obtained from domino tilings on the torus.

Proposition 4.1.2 (Toroidal height functions). *Let \mathbb{Z}^2 be the black-and-white infinite square lattice, as before. Then a height function h on \mathbb{Z}^2 (see Proposition 4.1.1) is a toroidal height function of \mathbb{T}_n if and only if h satisfies the following properties:*

1. $\exists k \in \mathbb{Z}, \forall v \in \mathbb{Z}^2, h(v + u_y) = h(v) + 4k.$
2. $\exists l \in \mathbb{Z}, \forall v \in \mathbb{Z}^2, h(v + u_x) = h(v) + 4l.$

where $u_x = (2n, 0)$ and $u_y = (0, 2n)$.

Moreover, if h is a toroidal height function of \mathbb{T}_n and t is its associated domino tiling of \mathbb{T}_n , then k is t 's vertical flux value and l is t 's horizontal flux value.

Proof. From our previous discussions, it follows immediately that a toroidal height function h of \mathbb{T}_n satisfies those properties and that the integers k, l determine the flux of h 's associated domino tiling of \mathbb{T}_n . On the other hand, properties 1 and 2 ensure the domino tiling \tilde{t} of $\mathbb{Z}^2 \subset \mathbb{R}^2$ associated to the height function h is invariant under translation by u_x or by u_y ; in other words, letting L be the lattice generated by $\{u_x, u_y\}$, \tilde{t} is L -periodic. This means a quotient by L will result in a tiled torus \mathbb{T}_n , and by construction that tiling's height function is h . \square

We will use this quasiperiodic characterization of toroidal height functions to establish results similar to those we obtained in the planar case.

Proposition 4.1.3. *Let t_1, t_2 be two tilings of \mathbb{T}_n with identical flux values k, l and corresponding toroidal height functions h_1, h_2 . Then $h_m = \min\{h_1, h_2\}$ is a toroidal height function of \mathbb{T}_n with flux values k, l .*

Proof. We first use Proposition 4.1.1 to check h_m is a height function on \mathbb{Z}^2 . Like before, it suffices to show that h_m changes by at most 3 along an edge on \mathbb{Z}^2 , and the proof of Proposition 3.1.3 applies verbatim here. We then need only check the conditions on Proposition 4.1.2. Notice that

$$\begin{aligned} h_m(v + u_y) &= \min \{h_1(v + u_y), h_2(v + u_y)\} = \min \{h_1(v) + 4k, h_2(v) + 4k\} \\ &= \min \{h_1(v), h_2(v)\} + 4k = h_m(v) + 4k, \end{aligned}$$

so $h_m(v + u_y) = h_m(v) + 4k$.

Similarly, $h_m(v + u_x) = h_m(v) + 4l$, and we are done. \square

Corollary 4.1.4 (Minimal height functions on the torus). *If there is a tiling of \mathbb{T}_n with flux values $k, l \in \mathbb{Z}$, then there is a tiling of \mathbb{T}_n with flux values k, l and so that its height function is minimal over tilings of \mathbb{T}_n with flux values k, l .*

4.2

More general tori: valid lattices

Remember the torus \mathbb{T}_n may be seen as a quotient \mathbb{R}^2/L , where L is the lattice generated by $\{(2n, 0), (0, 2n)\}$. We wish to consider other tori — or equivalently, other lattices. Of course, for the quotient \mathbb{R}^2/L to be a torus, we still need L to be generated by two linearly independent vectors. Moreover, when choosing a planar region D_L to be its fundamental domain, we will avoid those whose boundary crosses an edge on the infinite square lattice; this ensures D_L consists of whole squares whenever $L \subset \mathbb{Z}^2$. See the image below.




Figure 4.6: A lattice’s usual fundamental domain, and one made up of whole squares.

We would like to ensure the black-and-white condition is respected (that is, each domino consists of exactly one white square and one black square). This is *not* equivalent to D_L having the same number of black squares and white squares. For instance, the figure below features a lattice $L \subset \mathbb{Z}^2$ whose fundamental domain D_L has the same number of black squares and white squares; however, the identifications allow us to use dominoes consisting of two white squares or two black squares when tiling it.




Figure 4.7: A tiling of a torus which does not respect the black-and-white condition.

Perhaps a better interpretation of this situation is that the copies of the fundamental domain D_L that cover \mathbb{R}^2 are not all equally colored.




Figure 4.8: An invalid lattice: fundamental domains are not all equally colored.

It's now easy to see that a necessary and sufficient condition for them to be equally colored (and thus for the black-and-white condition to hold) is that the L -periodicity in \mathbb{R}^2 preserve square color. If L is generated by vectors $v_0, v_1 \in \mathbb{Z}^2$, this is equivalent to the sum of v_i 's coordinates being even. Notice this ensures D_L has an even number of squares, because that number is the area of D_L which is given by $\det(v_0, v_1)$. Moreover, we claim in this situation D_L automatically has an equal number of black squares and white squares.

Indeed, if D_L is a rectangle, the claim clearly holds (since it has an even number of squares). Otherwise, D_L can be taken to be an L-shaped figure, like in the figure below.




Figure 4.9: D_L can always be taken to be a rectangle or an L-shaped figure.

Notice in this case D_L can be decomposed into a union of two rectangles in two different ways:

- $a_0 \times b_0$ rectangle with area R_{00} and $a_1 \times (b_0 + b_1)$ rectangle with area R_{01}
- $(a_0 + a_1) \times b_0$ rectangle with area R_{10} and $a_1 \times b_1$ rectangle with area R_{11}

We need only show that for at least one $i = 0, 1$ both R_{i0} and R_{i1} are even. Since $R_{i0} + R_{i1}$ is always even (because that is the area of D_L), R_{i0} and R_{i1} always have the same parity. Thus, the claim would fail to hold only if all R_{ij} were odd. Studying the parity for a_0, a_1, b_0 and b_1 , it's easy to see this cannot be, so we are done.

Let $\mathcal{E}, \mathcal{O} \subset \mathbb{Z}^2$ be the sets of vertices whose coordinates are respectively both even and both odd, that is $\mathcal{E} = 2\mathbb{Z}^2$ and $\mathcal{O} = 2\mathbb{Z}^2 + (1, 1)$. It is clear from this discussion that whenever a lattice $L \subset \mathbb{Z}^2$ is generated by two vertices v_0, v_1 with $\det(v_0, v_1) \neq 0$ and $v_0, v_1 \in \mathcal{E} \sqcup \mathcal{O}$, the fundamental domain D_L satisfies the black-and-white condition and can be tiled by dominoes. We say such a lattice L is a *valid lattice*. We will refer to a torus obtained from a valid lattice L by \mathbb{T}_L . Toroidal height functions for these tori are defined much in the same way as before, via lifting.

5

Flux on the torus

We now explain how the flux definition that relies on counting cross-over dominoes can be adapted to these more general tori. Let L be a valid lattice generated by $\{v_0, v_1\}$ and t be a tiling of the torus on the fundamental domain D_L . As before, consider its lift to a tiling on \mathbb{Z}^2 .

For any vertex $v \in \mathbb{Z}^2$ there are two L-shaped paths joining v to $v + v_0$; call them u_0 and u_1 . Observe that if one of v_0 's coordinates is 0, u_0 coincides with u_1 . Generally, these edge-paths form the boundary of a quadriculated rectangle $R \subset \mathbb{Z}^2$ in which v and $v + v_0$ are opposite vertices.

Remember that whenever an edge-path crosses a domino on a tiling, the height function of that tiling changes by either $+3$ or -3 along that edge-path. Were we to define the flux of t through v_0 as before, we would like to say it is the number n_i of dominoes (horizontal **or** vertical) that cross u_i , each of which is counted positively if its corresponding height change along u_i is $+3$, and negatively if it is -3 ; notice u_i 's orientation matters. However, there is no particular reason why n_0 should be used over n_1 .

When we previously defined the flux via counting dominoes (in Chapter 4, for the square torus \mathbb{T}_n), u_0 and u_1 always coincided, so the distinction was irrelevant. If $v_0 \in \mathcal{E}$, the situation is similar. In this case, R is a rectangle with an even number of squares; in particular, the number of black squares and the number of white squares in R are the same. This means n_0 and n_1 are equal, so the choice of path does not matter.

Real change occurs if $v_0 \in \mathcal{O}$. In this case, R is a rectangle with an odd number of squares, so the number of black squares and the number of white squares in it differ by 1. This means $|n_0 - n_1| = 1$. Rather than arbitrarily choosing one of n_0, n_1 , we opt for a measured approach: we take their average. Notice that applying this to previous cases yields the same result.

Of course, this means that whenever $v_0 \in \mathcal{O}$, the flux of t through v_0 will be some k in $(\mathbb{Z} + \frac{1}{2})$, rather than in \mathbb{Z} . This does not contradict our original quasiperiodicity relations, and inspecting the proof of Proposition 4.1.2, we need only show that $h(v_0) = 4k$ for this case too.

Lemma 5.0.1. *Let L be a valid lattice generated by $\{v_0, v_1\}$. Let t be a tiling of \mathbb{T}_L and h its toroidal height function. Then*

$$h(v_0) = 4k \iff k \text{ is the flux of } t \text{ through } v_0,$$

where k is defined as above.

Proof. Let u_0 and u_1 be the L-shaped edge-paths joining the origin to v_0 (and oriented from the origin to v_0). Let d_i^+ be the number of dominoes crossing u_i that are counted positively and let d_i^- be the number of dominoes crossing u_i that are counted negatively. Let e_i^+ be the number of edges on u_i whose orientation (as induced by the coloring of \mathbb{Z}^2) agrees with u_i 's own, and let e_i^- be the number of edges on u_i whose orientation reverses u_i 's own.

Suppose $v_0 \in \mathcal{E}$. In this case, $e_i^+ = e_i^-$ for each $i = 0, 1$. Thus, if no domino crosses u_i (that is, when d_i^+ and d_i^- are both 0) the constructive definition of height functions implies $h(v_0) = 0$. Each domino counted by d_i^+ crosses an edge counted by e_i^- , and contributes with a height change of +3 along that edge (rather than -1); in other words, each domino counted by d_i^+ contributes with a total change of +4 for $h(v_0)$. Similarly, each domino counted by d_i^- contributes with a total change of -4 for $h(v_0)$. All of this implies the following formula¹ holds for each $i = 0, 1$:

$$h(v_0) = 4(d_i^+ - d_i^-) \tag{5-1}$$

The lemma follows from observing that $d_i^+ - d_i^-$ is t 's flux through v_0 .

Now suppose $v_0 \in \mathcal{O}$. In this case we no longer have $e_i^+ = e_i^-$; instead, we claim $e_i^+ - e_i^- = \pm 2$, where the sign in \pm is different for each $i = 0, 1$. Indeed, let R be the quadriculated rectangle whose boundary is given by $u_0 \cup u_1$. Each u_i can be divided into three segments as follows: a middle segment of length two fitting a corner square in R , and the other two outer segments (each of which possibly has length 0); see the Figure 5.1.

For each u_i , each of the outer segments has even length and features edges that are alternately counted by e_i^+ and by e_i^- , so $e_i^+ - e_i^-$ is given entirely by the middle segment. That segment has two edges that are counted with the same sign, but for each u_i that sign is different, so the claim is proved.

Without loss of generality, say $e_0^+ - e_0^- = 2$ and $e_1^+ - e_1^- = -2$. If no domino crosses u_0 , the constructive definition of height functions implies $h(v_0) = 2$. The same technique used above implies the following formula holds:

$$h(v_0) = 2 + 4(d_0^+ - d_0^-) \tag{5-2}$$

¹Equation (5-1) provides another way to see that when v_0 's coordinates are both even, the numbers n_0 and n_1 are equal.




Figure 5.1: The paths u_0 and u_1 , each divided into three segments.

Applying this process to u_1 gives us the formula²:

$$h(v_0) = -2 + 4(d_1^+ - d_1^-) \tag{5-3}$$

Combining the two yields $h(v_0) = 4 \cdot \frac{1}{2} [(d_0^+ - d_0^-) + (d_1^+ - d_1^-)]$. Since $\frac{1}{2} [(d_0^+ - d_0^-) + (d_1^+ - d_1^-)]$ is the flux of t through v_0 , the proof is complete. \square

The reader might question the choice of L -shaped paths for the flux definition. In this regard, we note the following. For any edge-path γ , let $\mathcal{R}(\gamma)$ be the edge-path obtained from γ by reflecting it across the middle point between 0 and v_0 (in particular, notice $\mathcal{R}(u_0) = u_1$). Consider the numbers n_γ and $n_{\mathcal{R}(\gamma)}$ of crossing dominoes, as we defined n_0, n_1 for u_0, u_1 . Then $n_\gamma + n_{\mathcal{R}(\gamma)} = n_0 + n_1$, so that ‘any measured approach’ to choosing an edge-path would yield the same results.

When $v_0 \in \mathcal{O}$, the flux of a tiling through v_0 is some k in $(\mathbb{Z} + \frac{1}{2})$, so Lemma 5.0.1 implies $h(v_0) \equiv 2 \pmod{4}$, rather than the usual 0 . Observe that this is consistent with the mod 4 prescription function Φ calculated just after Proposition 4.1.1.

Also, it’s clear that the flux of t through v_1 is similarly defined, and these properties also hold for v_1 .

Because of Lemma 5.0.1, generalizations of Propositions 4.1.2 and 4.1.3 to this new scenario are automatic.

Proposition 5.0.2 (General toroidal height functions). *Let L be a valid lattice generated by $\{v_0, v_1\}$. Then a height function h on \mathbb{Z}^2 (see Proposition 4.1.1) is*

²Together, equations (5-2) and (5-3) provide another way to see that when v_0 ’s coordinates are both odd, the numbers n_0 and n_1 differ by 1.

a toroidal height function of \mathbb{T}_L if and only if h satisfies the following property for each $i = 0, 1$:

$$\begin{aligned} v_i \in \mathcal{E} &\Rightarrow \exists k_i \in \mathbb{Z}, \quad \forall v \in \mathbb{Z}^2, \quad h(v + v_i) = h(v) + 4k_i \\ v_i \in \mathcal{O} &\Rightarrow \exists k_i \in (\mathbb{Z} + \frac{1}{2}), \quad \forall v \in \mathbb{Z}^2, \quad h(v + v_i) = h(v) + 4k_i \end{aligned}$$

Furthermore, if h is a toroidal height function of \mathbb{T}_L and t is its associated domino tiling, then k_0 is t 's flux through v_0 and k_1 is t 's flux through v_1 .

Proposition 5.0.3. *Let L be a valid lattice and t_1, t_2 be two tilings of \mathbb{T}_L with identical flux values k, l and corresponding toroidal height functions h_1, h_2 . Then $h_m = \min\{h_1, h_2\}$ is a toroidal height function of \mathbb{T}_L with flux values k, l .*

Corollary 5.0.4 (Minimal height functions on general tori). *Let L be a valid lattice. If there is a tiling of \mathbb{T}_L with flux values k, l , then there is a tiling of \mathbb{T}_L with flux values k, l and so that its height function is minimal over tilings of \mathbb{T}_L with flux values k, l .*

5.1

The affine lattice $L^\#$

Let L be a valid lattice generated by $\{v_0, v_1\}$. Proposition 5.0.2 provides a new way to interpret the flux of a tiling of \mathbb{T}_L . Given one such tiling t , let h_t be its toroidal height function. The quantities

$$\begin{aligned} \varphi_t(v_0) &= \frac{1}{4}(h_t(v + v_0) - h_t(v)) \\ \varphi_t(v_1) &= \frac{1}{4}(h_t(v + v_1) - h_t(v)) \end{aligned}$$

do not depend on $v \in \mathbb{Z}^2$. By the same token, for $i, j \in \mathbb{Z}$, h_t 's quasiperiodicity implies

$$\varphi_t(i \cdot v_0 + j \cdot v_1) = i \cdot \varphi_t(v_0) + j \cdot \varphi_t(v_1),$$

so φ_t can be seen as a homomorphism on L . Additionally, since $L \subset \mathbb{Z}^2 \subset \mathbb{R}^2$ is generated by two linearly independent vectors, the usual inner product $\langle \cdot, \cdot \rangle$ provides the means to identify φ_t with $\varphi_t^* \in \mathbb{R}^2$ via $\varphi_t(u) = \langle \varphi_t^*, u \rangle$. From now on, using this identification, we will not distinguish between \mathbb{R}^2 and $(\mathbb{R}^2)^*$, and similarly we will not distinguish between φ_t and φ_t^* .

What can be said about the image of the homomorphism φ_t ? Of course, it is entirely defined by the values φ_t takes on v_0 and on v_1 . If $v_i \in \mathcal{E}$, $\varphi_t(v_i) \in \mathbb{Z}$.

If $v_i \in \mathcal{O}$, $\varphi_t(v_i) \in (\mathbb{Z} + \frac{1}{2})$. This allows us to analyze each case separately. Consider the following sets:

$$\begin{aligned} L_{00}^* &= \left\{ \varphi \in \text{Hom}\left(L; \frac{1}{2}\mathbb{Z}\right) \mid \varphi(v_0), \varphi(v_1) \in \mathbb{Z} \right\} \\ L_{01}^* &= \left\{ \varphi \in \text{Hom}\left(L; \frac{1}{2}\mathbb{Z}\right) \mid \varphi(v_0) \in \mathbb{Z}, \varphi(v_1) \in \left(\mathbb{Z} + \frac{1}{2}\right) \right\} \\ L_{10}^* &= \left\{ \varphi \in \text{Hom}\left(L; \frac{1}{2}\mathbb{Z}\right) \mid \varphi(v_0) \in \left(\mathbb{Z} + \frac{1}{2}\right), \varphi(v_1) \in \mathbb{Z} \right\} \\ L_{11}^* &= \left\{ \varphi \in \text{Hom}\left(L; \frac{1}{2}\mathbb{Z}\right) \mid \varphi(v_0), \varphi(v_1) \in \left(\mathbb{Z} + \frac{1}{2}\right) \right\} \end{aligned}$$

Then it's readily checked that:

$$\begin{aligned} v_0, v_1 \in \mathcal{E} &\Rightarrow \varphi_t \in L_{00}^* \\ v_0 \in \mathcal{E}, v_1 \in \mathcal{O} &\Rightarrow \varphi_t \in L_{01}^* \\ v_0 \in \mathcal{O}, v_1 \in \mathcal{E} &\Rightarrow \varphi_t \in L_{10}^* \\ v_0, v_1 \in \mathcal{O} &\Rightarrow \varphi_t \in L_{11}^* \end{aligned}$$

Notice that $L_{00}^* = \text{Hom}(L; \mathbb{Z}) = L^*$. Furthermore, the sets L_{ij}^* decompose $\text{Hom}\left(L; \frac{1}{2}\mathbb{Z}\right)$ into four disjoint and non-empty subsets. Observe that the parities of $\varphi(2v_0)$ and of $\varphi(2v_1)$ provide a way to identify $\text{Hom}\left(L; \frac{1}{2}\mathbb{Z}\right)$ with $(2L)^* = \text{Hom}(2L; \mathbb{Z})$.

Another description of these sets can be given in terms of a basis for $(2L)^*$. For each $i, j = 0, 1$ let $\varphi_i \in (2L)^*$ be defined by $\varphi_i(v_j) = \frac{1}{2}\delta_{ij}$. The set $\{\varphi_0, \varphi_1\}$ is a basis for $(2L)^*$, and the following characterizations are immediate:

$$\begin{aligned} L_{00}^* &= \{x_0 \cdot \varphi_0 + x_1 \cdot \varphi_1 \in (2L)^* \mid x_0, x_1 \in 2\mathbb{Z}\} \\ L_{01}^* &= \{x_0 \cdot \varphi_0 + x_1 \cdot \varphi_1 \in (2L)^* \mid x_0 \in 2\mathbb{Z}, x_1 \in (2\mathbb{Z} + 1)\} \\ L_{10}^* &= \{x_0 \cdot \varphi_0 + x_1 \cdot \varphi_1 \in (2L)^* \mid x_0 \in (2\mathbb{Z} + 1), x_1 \in 2\mathbb{Z}\} \\ L_{11}^* &= \{x_0 \cdot \varphi_0 + x_1 \cdot \varphi_1 \in (2L)^* \mid x_0, x_1 \in (2\mathbb{Z} + 1)\} \end{aligned}$$

It should now be clear the sets L_{ij}^* are related by translations of φ_0 and/or φ_1 . Since L_{00}^* is itself a lattice, we can generally say the L_{ij}^* are affine, or translated, lattices. The inner product identification (like the one we did with φ_t) allows us to see this concretely, representing $(2L)^*$, and naturally also the L_{ij}^* , in \mathbb{R}^2 . Under this representation, $(2L)^* = \frac{1}{2}L^*$ and we have the chain of inclusions

$$L \subset \mathbb{Z}^2 \subset \frac{1}{2}\mathbb{Z}^2 \subset (2L)^*$$

For any given valid lattice L , all flux values of tilings of \mathbb{T}_L belong to one same L_{ij}^* and no other, depending on the parity of v_0 and v_1 's coordinates. We will call this set $L^\#$.

As an example, for the torus \mathbb{T}_n we have $\varphi_0 = (\frac{1}{4n}, 0)$ and $\varphi_1 = (0, \frac{1}{4n})$, so $(2L)^* \subset \mathbb{R}^2$ is the lattice generated by these vectors. Moreover, in this case $L^\# = L_{00}^*$, so $L^\# \subset \mathbb{R}^2$ is the lattice generated by $\{(\frac{1}{2n}, 0), (0, \frac{1}{2n})\}$.




Figure 5.2: The lattice $(2L)^*$ represented in \mathbb{R}^2 . Each L_{ij}^* corresponds to a color: $L_{00}^* = L^*$ is black, L_{10}^* is red, L_{01}^* is purple and L_{11}^* is green. The marks round black vertices indicate $L^\# = L_{00}^*$.

Proposition 5.1.1. *Let L be a valid lattice. Under the inner product identification, it holds that $\pm(\frac{1}{2}, 0)$ and $\pm(0, \frac{1}{2})$ are in $L^\#$. In particular, $L^\# = L^* + (\frac{1}{2}, 0)$.*

Proof. Let L be generated by $v_0 = (a, b)$ and $v_1 = (c, d)$. Then it's easily

checked that:

$$\begin{aligned}\varphi_0 &= \frac{1}{2} \cdot \frac{1}{ad - bc} \cdot (d, -c) \\ \varphi_1 &= \frac{1}{2} \cdot \frac{1}{ad - bc} \cdot (-b, a)\end{aligned}$$

From these, we derive

$$\begin{aligned}\pm (a \cdot \varphi_0 + c \cdot \varphi_1) &= \pm \left(\frac{1}{2}, 0\right) \\ \pm (b \cdot \varphi_0 + d \cdot \varphi_1) &= \pm \left(0, \frac{1}{2}\right),\end{aligned}$$

where choice of signs is the same across a line.

Since L is valid, a, b and c, d have the same parity, so these points are all in the same L_{ij}^* . It suffices to see this set is $L^\#$. \square

Notice the calculations in Proposition 5.1.1 also prove that when v_0 is multiplied by k_0 and v_1 is multiplied by k_1 , φ_0 is multiplied by k_0^{-1} and φ_1 is multiplied by k_1^{-1} . In other words, as the moduli of v_0 and v_1 increase (but the angle between them is kept constant), the moduli of φ_0 and φ_1 decrease, and vice-versa. Visually, this means that as L becomes more scattered, $L^\#$ becomes more cluttered.

5.2

Characterization of flux values


For a valid lattice L , let $\mathcal{F}(L)$ be the set of all flux values of tilings of \mathbb{T}_L . We know $\mathcal{F}(L) \subset L^\#$, but what more can be said about it? What elements of $L^\#$ are in $\mathcal{F}(L)$? Surely not all — $L^\#$ is infinite, and the definition of flux via counting dominoes makes it clear $\mathcal{F}(L)$ must be finite. This section is devoted to answering these questions, and does so via a full characterization of $\mathcal{F}(L)$.

For $v = (x, y) \in \mathbb{R}^2$, let $\|v\|_1 = |x| + |y|$ and $\|v\|_\infty = \max\{|x|, |y|\}$. Let $Q \subset \mathbb{R}^2$ be the set $\{v \in \mathbb{R}^2; \|v\|_1 \leq \frac{1}{2}\}$.

Theorem 1 (Characterization of flux values). $\mathcal{F}(L) = L^\# \cap Q$.

Its proof will be given by Propositions 5.2.1 and 5.2.7, each showing one of the inclusions.

Proposition 5.2.1. *For any valid lattice L , $\mathcal{F}(L) \subset L^\# \cap Q$.*

Figure 5.3: The set $Q \subset \mathbb{R}^2$.

For the proof of Proposition 5.2.1, we will need to develop new techniques.

There is a height function h_{\max} on \mathbb{Z}^2 that is maximal over height functions h on \mathbb{Z}^2 with $h(0) = 0$. Before providing a characterization, recall that a finite *edge-path* in a quadrilaterated region R is a sequence of vertices $(p_n)_{n=0}^m$ in R such that p_j is neighbor to p_{j+1} for all $j = 0, \dots, m-1$; in this case, it's clear $p_j p_{j+1}$ is an edge in R joining those two vertices. We say an edge-path $(p_n)_{n=0}^m$ joins p_0 (its starting point) to p_m (its endpoint) and has length m . We will also consider infinite edge-paths: those with no starting point, those with no endpoint, and those with neither a starting point nor an endpoint. In the last case, we say the edge-path is *doubly-infinite*. Notice the ordering of an edge-path's vertices imbues its edges with a natural orientation, and it need not agree with the natural orientation of R 's edges (induced by the coloring).

Given a tiling t of R , an edge-path in t is an edge-path in R such that each of its edges are in t (that is, none of its edges cross a domino in t).

For $v, w \in \mathbb{Z}^2$, let $\Gamma(v, w)$ be the set of all edge-paths in \mathbb{Z}^2 joining v to w that respect edge orientation (as induced by the coloring of \mathbb{Z}^2).

Finally, let $\mathcal{H}_0(R)$ be the set of height functions h on R with $h(0) = 0$. We are now ready to state the characterization.

Proposition 5.2.2 (Characterization of h_{\max}). *Consider the infinite black-and-white square lattice \mathbb{Z}^2 (with $[0, 1]^2$ black) and let $h_{\max} \in \mathcal{H}_0(\mathbb{Z}^2)$ be its maximal height function. Then*

$$h_{\max}(v) = \min_{\gamma \in \Gamma(0, v)} l(\gamma),$$

where $l(\gamma)$ is the length of γ .

Proof. Fix $v \in \mathbb{Z}^2$. We claim that for all $\gamma \in \Gamma(0, v)$ and for all $h \in \mathcal{H}_0(\mathbb{Z}^2)$ it holds that $h(v) \leq l(\gamma)$. Indeed, the constructive definition of height functions implies that whenever an edge in γ is traversed, h changes by $+1$ if that edge is on t and by -3 otherwise, so an induction on the length of γ justifies the claim.

Since $\Gamma(0, v)$ is never empty, it follows that any $h \in \mathcal{H}_0(\mathbb{Z}^2)$ satisfies $h(v) \leq \min_{\gamma \in \Gamma(0, v)} l(\gamma)$. Letting go of the requirement that $v \in \mathbb{Z}^2$ be fixed, this inequality then holds for all $v \in \mathbb{Z}^2$.

Now define $h_M : \mathbb{Z}^2 \rightarrow \mathbb{Z}$ to be the function given by $h_M(v) = \min_{\gamma \in \Gamma(0, v)} l(\gamma)$, so that $h(v) \leq h_M(v)$ for all $v \in \mathbb{Z}^2$ and for all $h \in \mathcal{H}_0(\mathbb{Z}^2)$. If we show that $h_M \in \mathcal{H}_0(\mathbb{Z}^2)$, it follows immediately that $h_M = h_{\max}$ and the proposition is proved.

By inspection, $h_M(0) = 0$. Using Proposition 4.1.1, it's easy to verify h_M is a height function. Indeed, property 1 follows from the fact that edge-paths in $\Gamma(0, v)$ respect edge orientation. For property 2, it suffices to check that any two neighboring vertices in \mathbb{Z}^2 can always be joined by an edge-path that respects edge orientation and has length at most three: either the edge joining those two vertices, or the edge-path going round a square that contains those two vertices. \square

There is elegance to the simplicity of this rather abstract proof, but it does little to shed light on the structure and properties of h_{\max} ; our next proposition addresses this. In addition, we provide an image of h_{\max} along with its associated tiling t_{\max} ; see Figure 5.4.

Proposition 5.2.3. *Let $v = (x_1, x_2) \in \mathbb{Z}^2$. If $x_1 \equiv x_2 \pmod{2}$, then*

$$h_{\max}(v) = 2 \cdot \|v\|_{\infty}. \quad (5-4)$$

More generally, it holds that

$$|h_{\max}(v) - 2 \cdot \|v\|_{\infty}| \leq 1. \quad (5-5)$$

Proof. The idea is to describe edge-paths $\gamma \in \Gamma(0, v)$ with minimal length. Because of Proposition 5.2.2, the constructive definition of height functions implies any such γ is an edge-path not only in \mathbb{Z}^2 , but also in t_{\max} . The explicit construction of these paths will allow us to derive relations (5-4) and (5-5).




Figure 5.4: The tiling t_{\max} and its associated height function h_{\max} . The marked vertex is the origin. Notice its only local extremum is the origin, a minimum, and it is not the height function of any torus (since it is not quasiperiodical).

We introduce the concept of *edge-profiles* round a vertex. When horizontal edges round a vertex point toward it and vertical edges round that vertex point away from it, we say the edge-profile round that vertex is type-0. When horizontal edges round a vertex point away from it and vertical edges round that vertex point toward it, we say the edge-profile round that vertex is type-1. It's clear those are the only possible cases, see the image below.



Figure 5.5: The two edge-profiles; type-0 to the left and type-1 to the right.

Notice the edge-profile round a vertex depends only on the region (and not on a tiling of the region). Moreover, two neighbouring vertices will always have distinct edge-profiles, so that any edge-path on a region will always feature successive vertices with alternating edge-profiles.

This means an edge-path that respects edge orientation will necessarily alternate between vertical and horizontal edges, correspondingly as that edge emanates from a vertex with edge-profile respectively type-0 and type-1. On the other hand, whenever an edge-path alternates between vertical and horizontal edges, it either always respects orientation (if vertical edges emanate from vertices with edge-profile type-0) or always reverses orientation (if vertical edges emanates from vertices with edge-profile type-1). This is the content of Corollary 5.2.4 below.

We can now characterize edge-paths in $\Gamma(0, v)$. Since the edge-profile round the origin in \mathbb{Z}^2 (as we have colored it) is type-0, any edge-path in $\Gamma(0, v)$ is an alternating sequence of vertical and horizontal edges, starting from the origin with a vertical edge and ending in v .

Consider then the vectors $e_1 = (1, 0)$ and $e_2 = (0, 1)$. By the characterization, any edge-path in $\Gamma(0, v)$ can be uniquely represented as an ordered sum of $\pm e_i$ in which the first term is either e_2 or $-e_2$ and no two consecutive terms are collinear vectors. It's that clear the length of an edge-path in this representation is simply the number of terms in the ordered sum. Furthermore, because an edge-path in $\Gamma(0, v)$ starts at the origin, if we carry out the sum of this unique representation the result **is** in fact the vector $v \in \mathbb{Z}^2$.

How does the ordered sum representation of a path $\gamma \in \Gamma(0, v)$ with minimal length look like? Let $v = x_1 \cdot e_1 + x_2 \cdot e_2$ be a vertex in \mathbb{Z}^2 and $i, j \in \{1, 2\}$ be different indices with $|x_j| \geq |x_i|$. The ordered sum representing γ will have exactly $|x_j|$ terms of the form $\pm e_j$, all of them with sign given by $\text{sgn}(x_j)$. Notice they add up to $x_j \cdot e_j$, and no smaller number of $\pm e_j$ terms does so.

Similarly, the ordered sum will feature $|x_i|$ terms of the form $\pm e_i$, all of them with sign given by $\text{sgn}(x_i)$, adding up to $x_i \cdot e_i$. Because $|x_j| \geq |x_i|$, in order for the ordered sum to fulfill the requirement that it be alternating in $\pm e_1$ and $\pm e_2$, it must have a number m (possibly zero) of additional $\pm e_i$ terms. Since the ordered sum starts with a $\pm e_2$ term and sums to v , m is uniquely defined.

It is clear that whenever γ has an ordered sum representation described as above, $\gamma \in \Gamma(0, v)$. Additionally, no path in $\Gamma(0, v)$ may have smaller length, for the unique ordered sum representation was chosen to have the smallest possible number of terms. By Proposition 5.2.2, $h_{\max}(v) = l(\gamma)$. We provide an example of this construction for $v = (4, -1)$ in Figure 5.6.




Figure 5.6: The paths in $\Gamma(0, (4, -1))$ with minimal length, along with their ordered sum representation. The marked vertex is the origin.

Notice this analysis ensures all of γ 's horizontal edges or all of γ 's vertical edges have the same orientation (possibly both); see Figure 5.7. This fact will be used in Lemma 6.1.2 later.

In the construction above, m is always even. Indeed, the number of plus signs and the number of minus signs in the additional m terms of the form $\pm e_i$ must be equal, for otherwise they would not add up to 0. When $x_1 \equiv x_2 \pmod{2}$, $|x_j| - |x_i|$ is even, and in this case it's easy to see we can take $m = |x_j| - |x_i|$. This implies the ordered sum representation has a total of $2 \cdot |x_j|$ terms, so formula (5-4) is proved.

It remains to prove inequality (5-5). Formula (5-4) means it trivially holds whenever $x_1 \equiv x_2 \pmod{2}$, so we need only check when x_1 and x_2 have different mod 2 values. In particular, we may assume $|x_j| > |x_i|$ (the inequality is strict).

Let $\gamma \in \Gamma(0, v)$ have minimal length. Consider the edge-path $\tilde{\gamma}$ obtained from γ by removing its last edge e . It is an edge-path in $\Gamma(0, v - e)$ with minimal length, for otherwise $\gamma \in \Gamma(0, v)$ would not have minimal length. This implies the equality $h_{\max}(v) = h_{\max}(v - e) + 1$.

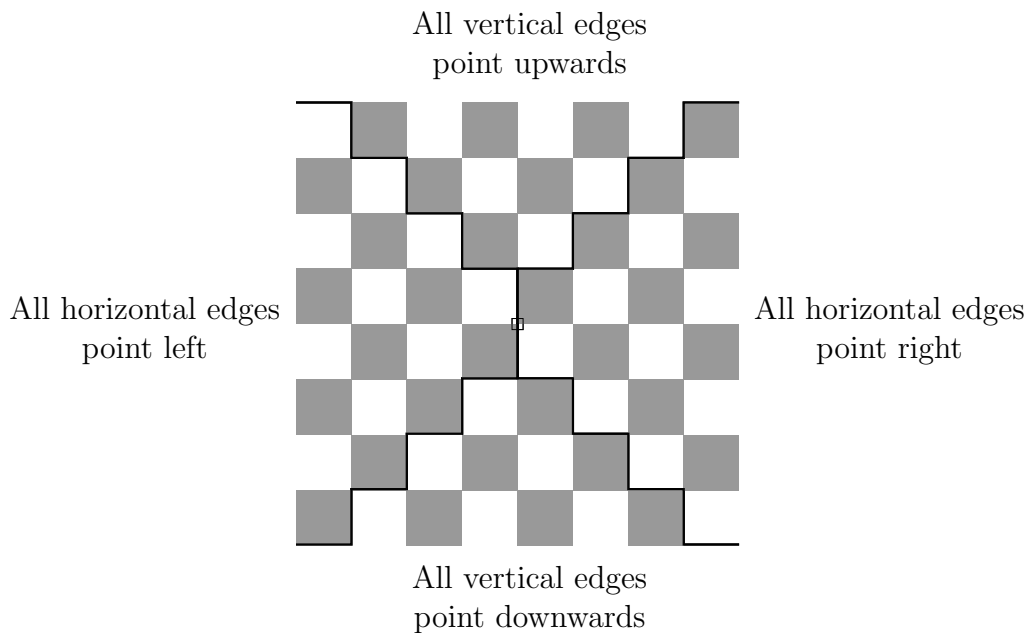


Figure 5.7: For each of the four regions above, if v belongs to that region, the edges of any $\gamma \in \Gamma(0, v)$ with minimal length satisfy the corresponding property. The marked vertex is the origin.

Write $v - e = (y_1, y_2)$. Observe that $v - e$ is obtained from v by changing one of its coordinates by ± 1 . Since x_1 and x_2 have different mod 2 values, it follows that $y_1 \equiv y_2 \pmod{2}$, and formula (5-4) applies: $h_{\max}(v - e) = 2 \cdot \max\{|y_1|, |y_2|\}$. Combining the two equalities yields

$$h_{\max}(v) - 2 \cdot \max\{|y_1|, |y_2|\} = 1 \tag{5-6}$$

There are two cases: (1) e is of the form $\pm e_i$; and (2) e is of the form $\pm e_j$.

In case (1), v 's x_i coordinate is changed by ± 1 , so $|y_j| = |x_j| \geq |y_i|$. Substituting into (5-6), the inequality holds.

In case (2), because $|x_j| > |x_i|$, all of γ 's edges of the form $\pm e_j$ have the same orientation. This implies $|y_j| = |x_j| - 1 \geq |x_i| = |y_i|$. Once again, substituting into (5-6) the inequality holds, and we are done. \square

Corollary 5.2.4. *Let R be a planar region and γ an (oriented) edge-path in R . Then the following are equivalent:*

- γ always respects or always reverses edge orientation (as induced by the coloring of R);
- γ 's edges alternate between horizontal and vertical.

Proof. See Proof of Proposition 5.2.3. \square

For $v, w \in \mathbb{Z}^2$, let $\Psi(v, w)$ be the set of all edge-paths in \mathbb{Z}^2 joining v to w that reverse edge orientation. The techniques used to obtain the characterization of h_{\max} can be very similarly employed to obtain a characterization of the minimal height function h_{\min} on $\mathcal{H}_0(\mathbb{Z}^2)$, and derive analogous results.

Corollary 5.2.5 (Characterization of h_{\min}). *Consider the infinite black-and-white square lattice \mathbb{Z}^2 (with $[0, 1]^2$ black) and let $h_{\min} \in \mathcal{H}_0(\mathbb{Z}^2)$ be its minimal height function. Then*

$$h_{\min}(v) = - \left(\min_{\gamma \in \Psi(0, v)} l(\gamma) \right),$$

where $l(\gamma)$ is the length of γ . Furthermore, if $v = (x_1, x_2) \in \mathbb{Z}^2$ and $x_1 \equiv x_2 \pmod{2}$, then

$$h_{\min}(v) = -2 \cdot \|v\|_{\infty}.$$

More generally, it holds that

$$|h_{\min}(v) + 2\|v\|_{\infty}| \leq 1.$$

Proof. Similar to the proofs of Propositions 5.2.2 and 5.2.3. \square

Before proving Proposition 5.2.1, we will need a quick lemma.

Lemma 5.2.6. *Let v, w be linearly independent vectors in \mathbb{Z}^2 . Then for each choice of signs in $(\pm x, \pm x)$, there is a nonzero integer linear combination of v, w with that form.*

Proof. Let $v = (a, b)$ and $w = (c, d)$. It suffices to prove for (x, x) and $(x, -x)$.

For the (x, x) case, take $k = -c + d$ and $l = a - b$, so that $k \cdot v + l \cdot w = (ad - bc, ad - bc)$. For the $(x, -x)$ case, take $k = c + d$ and $l = -a - b$, so that $k \cdot v + l \cdot w = (ad - bc, -(ad - bc))$. In each case, the combination uses integer coefficients, and it is nonzero because v and w are linearly independent. \square

We are now ready to prove Proposition 5.2.1.

Proof of Proposition 5.2.1. Let L be a valid lattice and t a tiling of \mathbb{T}_L with flux φ_t and height function h_t . It suffices to show that $\varphi_t \in Q$.

For any $(x, y) \in L$, we have that $\varphi_t(x, y) = \frac{1}{4}h_t(x, y)$. Of course, this means $\varphi_t(x, y) \leq \frac{1}{4}h_{\max}(x, y)$ for any $(x, y) \in L$. Because L is a valid lattice, it is generated by two linearly independent vectors. By Lemma 5.2.6, for any choice of signs in $(\pm x, \pm x)$, there is a vector in L with that form.

Proposition 5.2.3 then implies $\varphi_t(\pm x, \pm x) \leq \frac{1}{2}|x|$, so $\langle \varphi_t, (\pm 1, \pm 1) \rangle \leq \frac{1}{2}$. Writing $\varphi_t = (x_t, y_t)$, it then holds that $\pm x_t \pm y_t \leq \frac{1}{2}$.

In particular, there is a choice of signs in the previous inequality that yields $\|\varphi_t\|_1 = |x_t| + |y_t| \leq \frac{1}{2}$, so $\varphi_t \in Q$ as desired. \square

We now provide the remaining inclusion in Theorem 1.

Proposition 5.2.7. *For any valid lattice L , $\mathcal{F}(L) \supset L^\# \cap Q$.*

Before proving it, we will need a few lemmas.

Lemma 5.2.8. *Let L be a valid lattice. For all $v \in L$ and $\varphi \in L^\#$ it holds that*

$$4 \cdot \langle \varphi, v \rangle \equiv \Phi(v) \pmod{4},$$

where Φ is the mod 4 prescription function on the infinite square lattice³.

Proof. Suppose L is generated by $v_0 = (x_0, y_0)$ and $v_1 = (x_1, y_1)$. Given $v \in L$, there are unique integers a and b with $v = a \cdot v_0 + b \cdot v_1$, so that $v = (ax_0 + bx_1, ay_0 + by_1)$.

Similarly, given $\varphi \in L^\#$, there are unique integers z_0 and z_1 with $\varphi = z_0 \cdot \varphi_0 + z_1 \cdot \varphi_1$. We may then write

$$\begin{aligned} 4 \cdot \langle \varphi, v \rangle &= 4(a z_0 \cdot \langle \varphi_0, v_0 \rangle + b z_1 \cdot \langle \varphi_1, v_1 \rangle) \\ &= 2(a z_0 + b z_1) \end{aligned} \tag{5-7}$$

Notice $x_0 \equiv y_0 \equiv z_0 \pmod{2}$, because L is valid and $\varphi \in L^\#$. By the same token, $x_1 \equiv y_1 \equiv z_1 \pmod{2}$. Moreover, L being valid implies v 's coordinates have the same parity. We now analyze the mod 4 value of the expression in (5-7) for each case.

Suppose first that v 's coordinates are both even, that is, $ax_0 + bx_1 \equiv ay_0 + by_1 \equiv 0 \pmod{2}$. We must show $2(a z_0 + b z_1) \equiv 0 \pmod{4}$, or equivalently $az_0 + bz_1 \equiv 0 \pmod{2}$. This is implied by the mod 2 equivalences between x_0, y_0 and z_0 , and between x_1, y_1 and z_1 , so we are done.

Suppose now that v 's coordinates are both odd, that is, $ax_0 + bx_1 \equiv ay_0 + by_1 \equiv 1 \pmod{2}$. We must show $2(a z_0 + b z_1) \equiv 2 \pmod{4}$, or equivalently $az_0 + bz_1 \equiv 1 \pmod{2}$. Once again, this is implied by the mod 2 equivalences, and the proof is complete. \square

³ Φ is calculated just after Proposition 4.1.1.

Lemma 5.2.9. *Let L be a valid lattice and $\varphi \in L^\#$. For each $w \in \mathbb{Z}^2$, consider the expression*

$$\min_{v \in L} \left(4 \cdot \langle \varphi, v \rangle + \min_{\gamma \in \Gamma(v, w)} l(\gamma) \right).$$

The minimum exists if and only if $\varphi \in Q$.

Proof. First, observe that Lemma 5.2.8 guarantees the minimum is taken over integer-valued expressions; in other words, the existence of the minimum is equivalent to the existence of a lower bound.

On the one hand, the following estimate holds whenever $v \in L$:

$$-\|\varphi\|_1 \cdot \|v\|_\infty \leq \langle \varphi, v \rangle \leq \|\varphi\|_1 \cdot \|v\|_\infty \quad (5-8)$$

On the other, $\min_{\gamma \in \Gamma(v, w)} l(\gamma) = \min_{\gamma \in \Gamma(0, w-v)} l(\gamma) = h_{\max}(w-v)$, because v 's coordinates have the same parity (it is in L), so a translation by v preserves orientation. Proposition 5.2.3 then implies that for all $v \in L$

$$\begin{aligned} \min_{\gamma \in \Gamma(v, w)} l(\gamma) &\geq 2 \cdot \max\{|w_1 - v_1|, |w_2 - v_2|\} - 1 \\ &\geq 2 \cdot \max\{||w_1| - |v_1||, ||w_2| - |v_2||\} - 1. \end{aligned} \quad (5-9)$$

Consider the set $R(w) = \{(x, y) \in \mathbb{R}^2; |x| \leq |w_1| \text{ and } |y| \leq |w_2|\}$. Because it is bounded and L is discrete, $R(w) \cap L$ is finite. Thus, we need only show $4 \cdot \langle \varphi, v \rangle + \min_{\gamma \in \Gamma(v, w)} l(\gamma)$ has a lower bound for $v \in L$ outside $R(w)$. In this situation, inequality (5-9) allows us to write

$$\begin{aligned} \min_{\gamma \in \Gamma(v, w)} l(\gamma) &\geq 2 \cdot \max\{|v_1| - |w_1|, |v_2| - |w_2|\} - 1 \\ &\geq 2 \cdot \max\{|v_1|, |v_2|\} - 2 \cdot \max\{|w_1|, |w_2|\} - 1 \\ &= 2 \cdot \|v\|_\infty - 2 \cdot \|w\|_\infty - 1. \end{aligned}$$

Combining the two yields for all $v \in L$ the estimate

$$4 \cdot \langle \varphi, v \rangle + \min_{\gamma \in \Gamma(v, w)} l(\gamma) \geq (2 - 4 \cdot \|\varphi\|_1) \cdot \|v\|_\infty - 2 \cdot \|w\|_\infty - 1,$$

so that when $\varphi \in Q$ a lower bound outside $R(w)$ is given by $-(2 \cdot \|w\|_\infty + 1)$. Similar manipulations show that when $\varphi \in Q$, $-(2 \cdot \|w\|_\infty + 1)$ is a lower bound everywhere. In other words, when $\varphi \in Q$, the minimum exists.

To complete the proof, we show that when $\varphi \notin Q$, there is no lower bound. Observe that Lemma 5.2.6 guarantees the existence of a vertex $\tilde{v} \in L$ of the form $(\pm x, \pm x)$ and such that for all $n \in \mathbb{Z}$ it holds that $\langle \varphi, n \cdot \tilde{v} \rangle =$

$-n \cdot \|\varphi\|_1 \cdot |x|$. When $\varphi \notin Q$, there is some $\epsilon > 0$ for which $\|\varphi\|_1 > \frac{1}{2} + \frac{\epsilon}{4}$; in this case, we have $4 \cdot \langle \varphi, n \cdot \tilde{v} \rangle < -2n|x| - n\epsilon|x|$ for all positive integers n .

As before, Proposition 5.2.3 can be used to show the following estimate must hold outside $R(w)$:

$$\min_{\gamma \in \Gamma(n \cdot \tilde{v}, w)} l(\gamma) \leq 2n \cdot |x| - 2 \cdot \|w\|_\infty + 1$$

Thus, when $\varphi \notin Q$, it holds that for all positive integers n

$$4 \cdot \langle \varphi, n \cdot \tilde{v} \rangle + \min_{\gamma \in \Gamma(n \cdot \tilde{v}, w)} l(\gamma) < -n\epsilon|x| - 2 \cdot \|w\|_\infty + 1.$$

For fixed $w \in \mathbb{Z}^2$, we see the expression has no lower bound as n tends to infinity, so the proof is complete. \square

We are now ready to prove Proposition 5.2.7.

Proof of Proposition 5.2.7. Let $\varphi \in L^\# \cap Q$. We will construct the height function $h_{\max}^{L, \varphi}$ that is maximal over toroidal height functions of \mathbb{T}_L with flux φ (and base value 0 at the origin).

First, observe that if h is a toroidal height function of \mathbb{T}_L with flux φ , then $h(v) = 4 \cdot \varphi(v) = 4 \cdot \langle \varphi, v \rangle$ for all $v \in L$. Consider for each $v \in L$ the height function $h_{\max}^{v, \varphi}$, maximal over height functions that take the value $4 \cdot \langle \varphi, v \rangle$ on v (notice they need not have the value 0 on the origin). An easy adaptation of Proposition 5.2.2 yields

$$h_{\max}^{v, \varphi}(w) = 4 \cdot \langle \varphi, v \rangle + \min_{\gamma \in \Gamma(v, w)} l(\gamma).$$

As in Lemma 5.2.9, $\min_{\gamma \in \Gamma(v, w)} l(\gamma) = \min_{\gamma \in \Gamma(0, w-v)} l(\gamma) = h_{\max}(w-v)$, so we also have

$$h_{\max}^{v, \varphi}(w) = 4 \cdot \langle \varphi, v \rangle + h_{\max}(w-v).$$

If h is a toroidal height function of \mathbb{T}_L with flux φ , it follows that $h(w) \leq h_{\max}^{v, \varphi}(w)$ for all $w \in \mathbb{Z}^2$ and $v \in L$. By Lemma 5.2.9, the function given by $h_{\max}^{L, \varphi}(w) = \min_{v \in L} h_{\max}^{v, \varphi}(w)$ is well defined, which implies $h(w) \leq h_{\max}^{L, \varphi}(w)$ for all $w \in \mathbb{Z}^2$. We claim $h_{\max}^{L, \varphi}$ is a toroidal height function of \mathbb{T}_L with flux φ ; in this case, clearly it is maximal over such height functions.

We first prove it is a height function that takes the value 0 on the origin, as characterized by Proposition 4.1.1.

Item 1. $h_{\max}^{L,\varphi}(0) = 0$.

Since $0 \in L$ and by inspection $h_{\max}^{0,\varphi}(0) = 0$, we have the inequality $h_{\max}^{L,\varphi}(0) \leq 0$. We then need only show $h_{\max}^{v,\varphi}(0) \geq 0$ for all $v \in L$. Since for each $v \in L$ we have $h_{\max}^{v,\varphi}(0) = 4 \cdot \langle \varphi, v \rangle + h_{\max}(-v)$, the equivalent inequality $4 \cdot \langle \varphi, v \rangle \geq -h_{\max}(-v)$ suffices. Now, because $v \in L$ and L is valid, Proposition 5.2.5 implies $-h_{\max}(-v) = -2 \cdot \|v\|_{\infty}$. The inequality follows from applying estimate (5-8) in Lemma 5.2.9 (remember $\varphi \in Q$).

Item 2. $h_{\max}^{L,\varphi}$ has the prescribed mod 4 values on all of \mathbb{Z}^2 .

We show that $h_{\max}^{v,\varphi}$ satisfies this condition for all $v \in L$, from which the claim follows. Indeed, Lemma 5.2.8 implies $h_{\max}^{v,\varphi}$ respects the condition on L . To see this holds on all of \mathbb{Z}^2 , it suffices to note the edge-paths in $\min_{\gamma \in \Gamma(v,w)} l(\gamma)$ respect edge orientation.

Item 3. $h_{\max}^{L,\varphi}$ changes by at most 3 along an edge on \mathbb{Z}^2 .

Let e be an edge on \mathbb{Z}^2 joining w_1 to w_2 (in the orientation induced by the coloring of \mathbb{Z}^2). We claim $h_{\max}^{L,\varphi}(w_2) \leq h_{\max}^{L,\varphi}(w_1) + 1$. Indeed, there is some $v \in L$ and $\gamma \in \Gamma(v, w_1)$ with $h_{\max}^{L,\varphi}(w_1) = 4 \cdot \langle \varphi, v \rangle + l(\gamma)$. Consider the path $\tilde{\gamma} = \gamma * e$, where $*$ is edge-path concatenation. It is clear $\tilde{\gamma} \in \Gamma(v, w_2)$. Thus, it follows that $h_{\max}^{L,\varphi}(w_2) \leq h_{\max}^{v,\varphi}(w_2) \leq 4 \cdot \langle \varphi, v \rangle + l(\tilde{\gamma})$. Since $l(\tilde{\gamma}) = l(\gamma) + 1$, the claim holds.

Finally, we claim $h_{\max}^{L,\varphi}(w_1) \leq h_{\max}^{L,\varphi}(w_2) + 3$. Like before, there is some $\tilde{v} \in L$ and $\beta \in \Gamma(\tilde{v}, w_2)$ with $h_{\max}^{L,\varphi}(w_2) = 4 \cdot \langle \varphi, \tilde{v} \rangle + l(\beta)$. Consider the path $\tilde{\beta} = \beta * \tilde{e}$, where \tilde{e} is the edge-path joining w_2 to w_1 that goes round a square containing e . Observe that \tilde{e} respects edge-orientation, so $\tilde{\beta} \in \Gamma(\tilde{v}, w_1)$. It follows that $h_{\max}^{L,\varphi}(w_1) \leq h_{\max}^{\tilde{v},\varphi}(w_1) \leq 4 \cdot \langle \varphi, \tilde{v} \rangle + l(\tilde{\beta})$. Since \tilde{e} has length 3, $l(\tilde{\beta}) = l(\beta) + 3$ and the claim holds.

Together, both inequalities prove (3) above.

We have thus shown that $h_{\max}^{L,\varphi}$ is a height function; it remains to show it is L -quasiperiodic with flux φ . For the L -quasiperiodicity, we will prove that for all $v \in L$ and $w_1, w_2 \in \mathbb{Z}^2$

$$h_{\max}^{L,\varphi}(w_1 + v) - h_{\max}^{L,\varphi}(w_1) = h_{\max}^{L,\varphi}(w_2 + v) - h_{\max}^{L,\varphi}(w_2)$$

To that end, observe that any $u \in L$ can be written as $\tilde{u} + v$, so

$$\begin{aligned} h_{\max}^{L,\varphi}(w_i + v) &= \min_{(\tilde{u}+v) \in L} \left(4 \cdot \langle \varphi, \tilde{u} + v \rangle + h_{\max}(w_i - \tilde{u}) \right) \\ &= \min_{(\tilde{u}+v) \in L} \left(4 \cdot \langle \varphi, \tilde{u} \rangle + h_{\max}(w_i - \tilde{u}) \right) + 4 \cdot \langle \varphi, v \rangle \\ &= \min_{\tilde{u} \in L} \left(4 \cdot \langle \varphi, \tilde{u} \rangle + h_{\max}(w_i - \tilde{u}) \right) + 4 \cdot \langle \varphi, v \rangle \\ &= h_{\max}^{L,\varphi}(w_i) + 4 \cdot \langle \varphi, v \rangle \end{aligned}$$

Now, because $h_{\max}^{L,\varphi}(0) = 0$, this also shows that $h_{\max}^{L,\varphi}(v) = 4 \cdot \langle \varphi, v \rangle$ for all $v \in L$, so $h_{\max}^{L,\varphi}$ has flux φ and the proof is complete. \square

Combining Propositions 5.2.1 and 5.2.7, we obtain the full characterization provided by Theorem 1 at the beginning of this section.

6

Flip-connectedness on the torus

We would like to prove a flux-analogue of Proposition 3.1.5. This would allow us to use flip-connectedness in studying the properties of Kasteleyn matrices for the torus, in a manner similar to the planar case. However, it turns out that for extremal values of the flux — those lying on the boundary¹ of Q —, domino tilings of the torus with that flux value are not flip-connected. In fact, we will see that each tiling with an extremal flux value is a flip-isolated point in the space of domino tilings of the torus.

The image below features two tilings of the torus with identical flux values, but notice none of them admits any flip at all.




Figure 6.1: Two tilings of \mathbb{T}_2 that admit no flips and have identical flux values.

In attempting to reproduce the proof of Proposition 3.1.5, we can see how the situation is different. Let h, h_m be height functions of domino tilings of \mathbb{T}_L with identical flux values, h_m minimal and $h \neq h_m$. Consider the difference $g = h - h_m$. By Proposition 5.0.2, for all $v \in \mathbb{Z}^2$ we have that $g(v) = g(v + u_x) = g(v + u_y)$, so that g is L -periodic and in particular bounded. Moreover, Proposition 4.1.1 means g takes nonnegative values in $4\mathbb{Z}$. Let V be the set of vertices of \mathbb{Z}^2 on which g is maximum. Were we to proceed as before, we would now look for a vertex $v \in V$ that maximizes h , but there's generally no such vertex because of h 's quasiperiodicity.

What we truly seek, however, is a local maximum of h ; if one such vertex exists, then h 's quasiperiodicity would give rise to a copy of the local maximum

¹Notice that Proposition 5.1.1 implies $L^\#$ always intersects ∂Q .

on each copy of the fundamental domain D_L . Performing a flip on each of those would result in a new toroidal height function \tilde{h} that is less than h at these points, and identical at every other; remember flips preserve flux values.

Nonetheless, once again, there's nothing that guarantees the existence of a local maximum, and in fact for extremal values of the flux, it does not exist. In order to understand these behaviors, we will investigate properties of tilings that do not admit flips.

6.1

Tilings of the infinite square lattice

The next result will characterize domino tilings of the plane, but before stating it we need to introduce two new concepts. A *domino staircase*, or simply a staircase when the context is clear, is a sequence of neighbouring dominoes such that:

- All dominoes in the staircase are parallel, that is, either all of them are horizontal or all of them are vertical;
- Neighbouring dominoes in the staircase always touch along one edge of the longer side;
- Except for the first and last dominoes (if they exist), each domino in the staircase has exactly two neighbouring dominoes in the staircase;
- For each domino in the staircase with exactly two neighbouring dominoes in the staircase, those two neighbours touch the domino at different squares.




Figure 6.2: Examples of domino staircases.

If a staircase is finite, its length is the number of dominoes in it. A staircase may be infinite in a single direction; if it is infinite in two directions, we say it is *doubly-infinite*.

Notice that staircases have the very important property that for each domino in the staircase with exactly two neighbouring dominoes in the

staircase, there are no flips involving that domino. In particular, a doubly-infinite staircase admits no flips involving one of its dominoes.

A *windmill tiling* is a domino tiling of the plane that has one of the forms below:




Figure 6.3: Windmill tilings of \mathbb{Z}^2 .

Observe that a windmill tiling admits no flips and consists entirely of infinite staircases that are never doubly-infinite. As Theorem 2 will show, they are the only tilings of \mathbb{Z}^2 with this property.

Theorem 2 (Characterization of tilings of the infinite square lattice). *Let t be a tiling of \mathbb{Z}^2 . Then exactly one of the following applies:*

1. t admits a flip;
2. t consists entirely of parallel, doubly-infinite domino staircases;
3. t is a windmill tiling.

We will need a few lemmas for the proof, and thus it will be delayed.

A *staircase edge-path* is an edge-path whose edges alternate between vertical and horizontal and such that all of its vertical edges have the same orientation and all of its horizontal edges have the same orientation. Notice by Corollary 5.2.4 a staircase edge-path always respects or always reverses edge orientation.

There are essentially two kinds of staircase edge-paths: northeast-southwest or northwest-southeast. If orientation (as induced by colors) is taken into consideration, there are four: we will refer to them by 1-3 and 3-1 for northeast-southwest, and by 2-4 and 4-2 for northwest-southeast (think quadrants in the plane). There are thus four *types* of staircase-edge paths. Examples are provided in Figure 6.4.




Figure 6.4: From left to right: 3-1, 1-3, 4-2 and 2-4 staircase edge-paths.

Observe that a domino staircase always admits two parallel staircase edge-paths that fit it, one on each side. Furthermore, with orientation induced by the coloring of \mathbb{Z}^2 , the staircase-edge paths are always the same type. We can thus speak of types of domino staircases: it is the same type as that of the edge-paths that fit it (when those are oriented as per the coloring of \mathbb{Z}^2).

Four particularly interesting tilings of the plane related to this observation are the *brick walls*. Each of them uses only one type of domino (vertical or horizontal) and consists entirely of doubly-infinite domino staircases. Another characterization is as follows: each of them can be seen as consisting entirely of northeast-southwest doubly-infinite domino staircases **and** entirely of northwest-southeast doubly-infinite domino staircases. It's easy to see they are the only tilings of the plane with this property.

The following image shows the four different brick walls; the marked vertex is the origin.



Figure 6.5: The four brick walls.

Later, Proposition 7.0.7 will expand on the importance of brick walls.

Lemma 6.1.1. *Let t be a tiling of \mathbb{Z}^2 . Suppose there is a doubly-infinite staircase edge-path γ in t , dividing \mathbb{R}^2 into two disjoint and quadriculated connected components \mathcal{Z}_1 and \mathcal{Z}_2 . Then for each of \mathcal{Z}_1 and \mathcal{Z}_2 , its tiling by t contains a doubly-infinite domino staircase that fits γ or admits a flip.*

Proof. Let γ be a doubly-infinite staircase edge-path in t and choose any square Q in \mathcal{Z}_1 touching γ along one (and thus two) of its edges. Notice that any choice of type (either vertical or horizontal) for the domino tiling Q in t propagates

infinitely in one direction of γ , producing a domino staircase S in \mathcal{Z}_1 that fits γ and is infinite in that direction. Observe the image below.




Figure 6.6: Domino propagation across staircases.

Let S_1 be the maximal domino staircase in t containing the domino that tiles \mathcal{Q} ; the preceding paragraph makes it clear that S_1 is in \mathcal{Z}_1 and is infinite in at least one direction. If it is doubly-infinite, there is nothing to prove. Otherwise, S_1 has exactly one extremal domino \mathcal{D}_1 . Consider the square \mathcal{D}_2 in $\mathcal{Z}_1 \setminus S_1$ that touches γ along one of its edges and is closest to S_1 's extremal domino.




Figure 6.7: The domino tiling \mathcal{D}_2 must be of the other type.

The domino tiling \mathcal{D}_2 cannot have the same type as the dominoes of S_1 , for otherwise \mathcal{D}_1 would not be S_1 's extremal domino. The choice of type for the domino tiling \mathcal{D}_2 is thus fixed, and like before it propagates infinitely, except this time in the other direction of γ . This produces a new domino staircase S_2 in \mathcal{Z}_1 that fits γ and is infinite in that direction. Finally, consider the square in \mathcal{Z}_1 touching both S_1 and S_2 .




Figure 6.8: A flip is inevitable.

Regardless of the choice of type for the domino tiling that square in t , a flip is enabled (either with \mathcal{D}_1 or with \mathcal{D}_2), so the claim on \mathcal{Z}_1 is proved.

The argument does not rely on any particular property of \mathcal{Z}_1 , and using \mathcal{Z}_2 in its stead yields the complete proof. \square

Lemma 6.1.2. *Let t be a tiling of a planar region R and γ an edge-path in t joining v to w that respects (respectively reverses) edge orientation. Then*

1. *Any edge-path in R joining v to w that respects (respectively reverses) edge orientation and has length $l(\gamma)$ is an edge-path in t ;*
2. *γ has minimal length over edge-paths in R joining v to w that respect (respectively reverse) edge orientation;*
3. *At least one of the following properties is true:*
 - (a) *Every horizontal edge in γ has the same orientation;*
 - (b) *Every vertical edge in γ has the same orientation.*

Proof. (1) follows immediately from the constructive definition of height functions.

For (2), either $\gamma \in \Gamma(v, w)$ or $\gamma \in \Psi(v, w)$. Suppose we're in the first case; the other is analogous. Let h be t 's associated height function. Consider the auxiliary height function h_{aux} , maximal over height functions \tilde{h} on \mathbb{Z}^2 with $\tilde{h}(v) = h(v)$ (notice they need not have the value 0 on the origin). An easy adaptation of Proposition 5.2.2 yields $h_{\text{aux}}(u) = h(v) + \min_{\tilde{\gamma} \in \Gamma(v, u)} l(\tilde{\gamma})$. Then $h \leq h_{\text{aux}}$ wherever h is defined, and since the constructive definition of height functions implies $h(w) = h(v) + l(\gamma)$ (because γ is an edge-path in t), we have

$$h(v) + l(\gamma) = h(w) \leq h_{\text{aux}}(w) = h(v) + \min_{\tilde{\gamma} \in \Gamma(v, w)} l(\tilde{\gamma}).$$

In other words, $l(\gamma) \leq \min_{\tilde{\gamma} \in \Gamma(v, w)} l(\tilde{\gamma})$, so it must be an equality and (2) is proved.

Finally, consider the edge-path $\tilde{\gamma}$ in \mathbb{Z}^2 obtained from γ by a translation that takes v to the origin, that is, $\tilde{\gamma} = \gamma - v$. It is clear either $\tilde{\gamma} \in \Gamma(0, w - v)$ or $\tilde{\gamma} \in \Psi(0, w - v)$, and by (2) it has minimal length over its corresponding set. The Proof of Proposition 5.2.3 (or a trivial adaptation for Ψ and h_{min}) then implies (3), and we are done. \square

Consider the black-and-white infinite square lattice \mathbb{Z}^2 . A finite planar region $R \subset \mathbb{Z}^2$ is a *rugged rectangle* if its boundary consists of two pairs of staircase edge-paths $\{S_{00}, S_{01}\}$ and $\{S_{10}, S_{11}\}$ such that for each $i \in \{0, 1\}$, S_{i0} and S_{i1} are the same type (**including** color-induced orientation). In this case, we say the rugged rectangle has side lengths l and m , where l is the number of squares on R that fit S_{00} or S_{01} , and m is the number of squares on R that fit S_{10} or S_{11} . Observe that because R is finite, there must be two northeast-southwest staircases and two northwest-southeast staircases, so side lengths are finite and do not depend on choice of staircase.

We provide examples of rugged rectangles below.




Figure 6.9: A rugged rectangle with side lengths 7 and 11, and one with side lengths 10 and 5.

Induction shows a rugged rectangle can be tiled in a single way, either entirely by vertical dominoes or entirely by horizontal dominoes. In fact, rugged rectangles are pieces of brick walls; see Figure 6.10.

Lemma 6.1.3. *Let t be a tiling of \mathbb{Z}^2 and γ a finite edge-path in t that respects (respectively reverses) edge orientation. At least one of the following holds²:*

- (a) *Every horizontal edge in γ has the same orientation;*
- (b) *Every vertical edge in γ has the same orientation.*

If exactly one of (a), (b) holds, then γ defines a rugged rectangle R in t .

Furthermore, if (a) holds but (b) doesn't, the side lengths of R are given by the number of vertical edges in γ that point upwards and the number of

²This is provided by Lemma 6.1.2.




Figure 6.10: The unique tilings of the rugged rectangles from Figure 6.9.

vertical edges in γ that point downwards. If (b) holds but (a) doesn't, the side lengths of R are given by the number of horizontal edges in γ that point left and the number of horizontal edges in γ that point right.

Proof. Let v be γ 's starting point and w be γ 's endpoint, so that γ joins v to w . Suppose (a) holds but (b) doesn't; the opposite case is analogous.

Let γ_{up} and γ_{down} be edge-paths joining v to w that respect (respectively reverse) edge orientation and have length $l(\gamma)$; γ_{up} features all vertical edges pointing upwards before any vertical edge pointing downwards and γ_{down} features all vertical edges pointing downwards before any vertical edge pointing upwards. Notice γ , γ_{up} and γ_{down} have the same number of horizontal edges (all of which have the same orientation across all three edge-paths), the same number of vertical edges pointing upwards and the same number of vertical edges pointing downwards. Moreover, because (b) does not hold, there is at least one vertical edge of each type.

Remember that whenever an edge-path respects or reverses edge orientation, vertical and horizontal edges appear in alternating fashion (see Corollary 5.2.4). This means γ_{up} has a 'rising' staircase edge-path γ_{up}^+ followed by a 'descending' staircase edge-path γ_{up}^- , and γ_{down} has a 'descending' staircase edge-path γ_{down}^- followed by a 'rising' staircase edge-path γ_{down}^+ . Figure 6.11 provides an example of this construction.

By Lemma 6.1.2, γ_{up} and γ_{down} are edge-paths in t . We assert that the region R they enclose is a rugged rectangle. Indeed, $\{\gamma_{\text{up}}^+, \gamma_{\text{down}}^+\}$ is a pair of parallel staircase edge-paths, for in both γ_{up}^+ and γ_{down}^+ , vertical edge points upward and horizontal edges have the same orientation (because (a) holds).

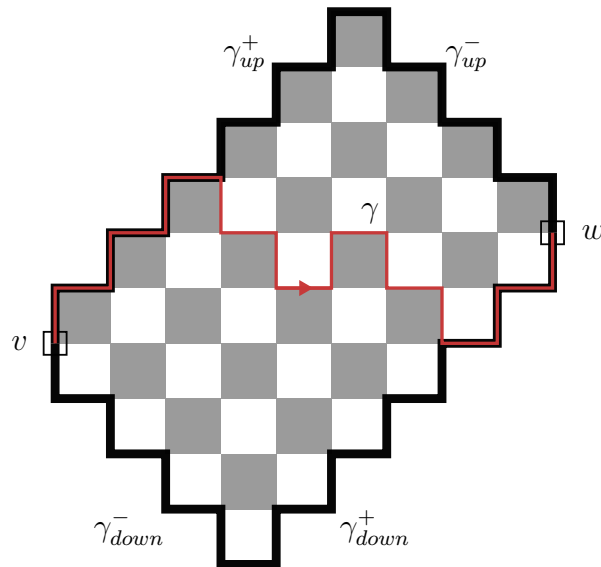


Figure 6.11: An example construction of the paths γ_{up} and γ_{down} .

Furthemore, because both respect (respectively reverse) edge orientation, this also means they are the same type. Similarly, $\{\gamma_{up}^-, \gamma_{down}^-\}$ is a pair of staircase edge-paths that are the same type, so the assertion holds.

It's also easy to see that any square on R that fits γ_{up}^+ lies beside a vertical edge on γ_{up}^+ , and any square on R that fits γ_{up}^- lies beside a vertical edge on γ_{up}^- . Since every vertical edge on γ_{up}^+ points upwards and every vertical edge on γ_{up}^- points downwards, the claim on the side lengths of R is proved. \square

Two non-parallel doubly-infinite staircase edge-paths meet along a single edge and divide \mathbb{R}^2 into four disjoint and quadriculated connected components. Each such connected component is a planar region we will call *rugged quadrant*.

There are eight kinds of rugged quadrants, four of which can be tiled in a single way: either entirely by vertical dominoes, or entirely by horizontal dominoes. These are the north, south, east and west rugged quadrants, and we will refer to them as *cardinal* rugged quadrants. Each of the other four kinds admits an infinite number of tilings, and will be of no interest to us. Figure 6.12 provides examples of rugged quadrants.

Lemma 6.1.4. *Let t be a tiling of \mathbb{Z}^2 and Q_0, Q_1 be cardinal rugged quadrants in t . If Q_0 and Q_1 are the same type (north, south, east or west), then there is a cardinal rugged quadrant \tilde{Q} in t that is their type and contains them both.*

Proof. If $Q_0 \subseteq Q_1$ or $Q_1 \subseteq Q_0$, there is nothing to prove; suppose this is not the case. Furthermore, suppose Q_0 and Q_1 are both north; the proof for other types is analagous to the one that follows.




Figure 6.12: The eight rugged quadrants, and the unique tilings of each cardinal quadrant.

Let γ_i be the edge-path that fits the border of Q_i ($i = 0, 1$). Since its edges alternate between horizontal and vertical, an orientation of γ_i always respects or always reverses edge orientation (see Corollary 5.2.4); choose the orientation that always respects it. We claim this choice is the same for both γ_0 and γ_1 , that is, either both go from left to right or both go from right to left.

Indeed, because Q_0 neither contains nor is contained in Q_1 , γ_0 and γ_1 must intersect on non-parallel segments along a single edge. Furthermore, since a vertical domino lies above every horizontal edge of each γ_i (because Q_i is north), that edge must be horizontal; see the following image.




Figure 6.13: Intersecting north quadrants meet along a single horizontal edge.

The only possible edge orientations are shown below.




Figure 6.14: Color-induced orientations on the boundary of intersecting north quadrants.

Regardless of the situation, the claim holds. Notice this implies every horizontal edge of γ_0 and every horizontal edge of γ_1 have the same orientation.

Decompose γ_i into disjoint staircase edge-paths γ_i^- and γ_i^+ , where γ_i^- has vertical edges pointing downwards and γ_i^+ has vertical edges pointing upwards. Because γ_0 and γ_1 both have the same orientation, γ_0^- and γ_1^- are parallel, and γ_0^+ and γ_1^+ are parallel. This implies either γ_0^+ intersects γ_1^- or γ_1^+ intersects γ_0^- . Suppose without loss of generality that γ_0^+ intersects γ_1^- and call e the horizontal edge along which they intersect.

Let β^+ be the maximal segment of γ_0^+ that ends with e and β^- the maximal segment of γ_1^- that starts with e . Observe that $\gamma_0^- \cup \beta^+ \cup \beta^- \cup \gamma_1^+$ is the edge-path that fits the border of $Q_0 \cup Q_1$, and with the orientation inherited from γ_0 and γ_1 it always respects edge orientation.

Consider the edge-path $\beta = \beta^+ \cup \beta^-$. Let v be β 's starting point and w be β 's endpoint. Consider the edge-path α joining v to w that respects edge orientation, has length $l(\beta)$ and features all vertical edges pointing downwards before any vertical edge pointing upwards. Notice β and α have the same number of horizontal edges (all of which have the same orientation), the same

number of vertical edges pointing upwards and the same number of vertical edges pointing downwards.

By Corollary 5.2.4, edges in α alternate between vertical and horizontal, so that α has a ‘descending’ staircase edge-path α^- followed by a ‘rising’ staircase edge-path α^+ . Furthermore, by Lemma 6.1.2, α is an edge-path in t .

We claim $\gamma_0^- \cup \alpha^-$ is a ‘descending’ staircase edge-path. Indeed, since γ_0^- and α^- both respect edge orientation, so does their union; this implies the edges on $\gamma_0^- \cup \alpha^-$ alternate between horizontal and vertical. Furthermore, by construction all of its vertical edges point downwards, and all of its horizontal edges have the same orientation. The claim thus holds. Similarly, $\alpha^+ \cup \gamma_1^+$ is a ‘rising’ staircase edge-path. Moreover, the entire union $\zeta = \gamma_0^- \cup \alpha^- \cup \alpha^+ \cup \gamma_1^+$ features edges that alternate between horizontal and vertical, because $\alpha = \alpha^- \cup \alpha^+$ does.




Figure 6.15: The edge-path $\alpha = \alpha^- \cup \alpha^+$ is in t and fits the border of a north quadrant.

This means ζ is an edge-path in t that fits the border of a north rugged quadrant \tilde{Q} . Since γ_0^- and γ_1^+ are contained in ζ (alternatively, since $\gamma_0^- \cup \beta^+ \cup \beta^- \cup \gamma_1^+$ is contained in \tilde{Q}), it is clear \tilde{Q} contains both Q_0 and Q_1 , and we are done. \square

We are now ready to prove Theorem 2.

Proof of Theorem 2. For the first part, we will show that if t contains a doubly-infinite domino staircase but does not admit a flip, then it consists entirely of parallel, doubly-infinite staircases. Indeed, suppose t contains a doubly-infinite

domino staircase S ; then there are two distinct doubly-infinite staircase edge-paths in t that fit S , one on either side of it. By Lemma 6.1.1, if t admits no flips, then on either side of S lies another doubly-infinite domino staircase that is parallel to S , so by induction t consists entirely of those.

For the second part, suppose t neither contains a doubly-infinite staircase nor admits a flip. We will show that t is a windmill tiling. Let h be t 's associated height function. Because t admits no flips at all, h cannot have local extrema.

Take any $v_0 \in \mathbb{Z}^2$. Since v_0 is not a local maximum of h , there must be a neighbouring vertex $v_1 \in \mathbb{Z}^2$ for which $h(v_1) > h(v_0)$. Of course, v_1 is not a local maximum of h , so we can repeat the process. This produces a list $(v_n)_{n \geq 0} \subset \mathbb{Z}^2$ with $h(v_{n+1}) > h(v_n)$ for all $n \geq 0$ and in which each v_n is neighbour to v_{n+1} .

Notice we may assume $h(v_{n+1}) = h(v_n) + 1$. Indeed, when $h(v_{n+1}) = h(v_n) + 3$, the edge joining v_n to v_{n+1} crosses a domino in t , so going from v_n to v_{n+1} round that domino is allowed. It's clear that each edge traversed this way increases h by $+1$, so there is no loss of generality in the assumption.

Similarly, since v_0 is not a local minimum of h , there must be a neighbouring vertex $v_{-1} \in \mathbb{Z}^2$ for which $h(v_0) > h(v_{-1})$. Repeating the process, we obtain a new list $(v_m)_{m \leq 0} \subset \mathbb{Z}^2$ with $h(v_m) > h(v_{m-1})$ for all $m \leq 0$ and in which each v_m is neighbour to v_{m-1} . Like before, we may assume $h(v_{m-1}) = h(v_m) - 1$.

The union of these two lists yields an indexed list $(v_k)_{k \in \mathbb{Z}}$ with $h(v_{k+1}) = h(v_k) + 1$ for all $k \in \mathbb{Z}$ and in which each v_k is neighbour to v_{k+1} .

For each $n \in \mathbb{N}$, consider the edge-paths $\gamma_n = (v_k)_{k=-n}^n$, $\gamma_n^+ = (v_k)_{k=0}^n$ and $\gamma_n^- = (v_k)_{k=-n}^0$. Because h always changes by $+1$ along an edge on each of these paths, they are by construction edge-paths in t that respect edge orientation, so Lemma 6.1.2 applies to them. Furthermore, because γ_n is always contained in γ_{n+1} , at least one of the statements below is true.

- (a) Each horizontal edge of $\cup_{n \in \mathbb{N}} \gamma_n$ has the same orientation;
- (b) Each vertical edge of each $\cup_{n \in \mathbb{N}} \gamma_n$ has the same orientation.

We claim exactly one of these hold. Indeed, if both (a) and (b) hold, $\cup_{n \in \mathbb{N}} \gamma_n$ is a doubly-infinite staircase edge-path, so Lemma 6.1.1 applies. This contradicts our initial assumption that t neither contains a doubly-infinite domino staircase nor admits a flip, and the claim is thus proved.

Suppose then that (a) holds but not (b).

Let u_n^+ be the number of vertical edges pointing upwards in γ_n^+ , d_n^+ be the number of vertical edges pointing downwards in γ_n^+ and similarly for u_n^- and d_n^- . Notice $(u_n^+)_{n \in \mathbb{N}}$ is a nondecreasing sequence of nonnegative integers, and the same is true for the others. Furthermore, because $(u_n^+ + d_n^+)$ is the number of vertical edges in γ_n^+ , at least one of $(u_n^+)_{n \in \mathbb{N}}$ and $(d_n^+)_{n \in \mathbb{N}}$ is unbounded, and similarly for $(u_n^-)_{n \in \mathbb{N}}$ and $(d_n^-)_{n \in \mathbb{N}}$. We assert that:

$$(d_n^-) \text{ is unbounded} \iff \begin{matrix} (u_n^+) \text{ is unbounded and} \\ (u_n^-, d_n^+) \text{ are bounded} \end{matrix} \tag{6-1}$$

The \Leftarrow implication is obvious. Suppose now that $(d_n^-)_{n \in \mathbb{N}}$ is unbounded; we will prove the \Rightarrow implication. Consider the doubly-infinite staircase edge-path S_d (respectively S_u) defined by:

- Its vertical edges all point downwards (respectively upwards);
- Its horizontal edges have the same orientation as those in $\cup_{n \in \mathbb{N}} \gamma_n$;
- It respects edge orientation;
- It contains v_0 .

Let S_d^- be the infinite segment of S_d that ends in v_0 (as per the edge-path's own orientation), S_d^+ be the infinite segment that starts at v_0 , and similarly for S_u^- and S_u^+ . Observe the image below.

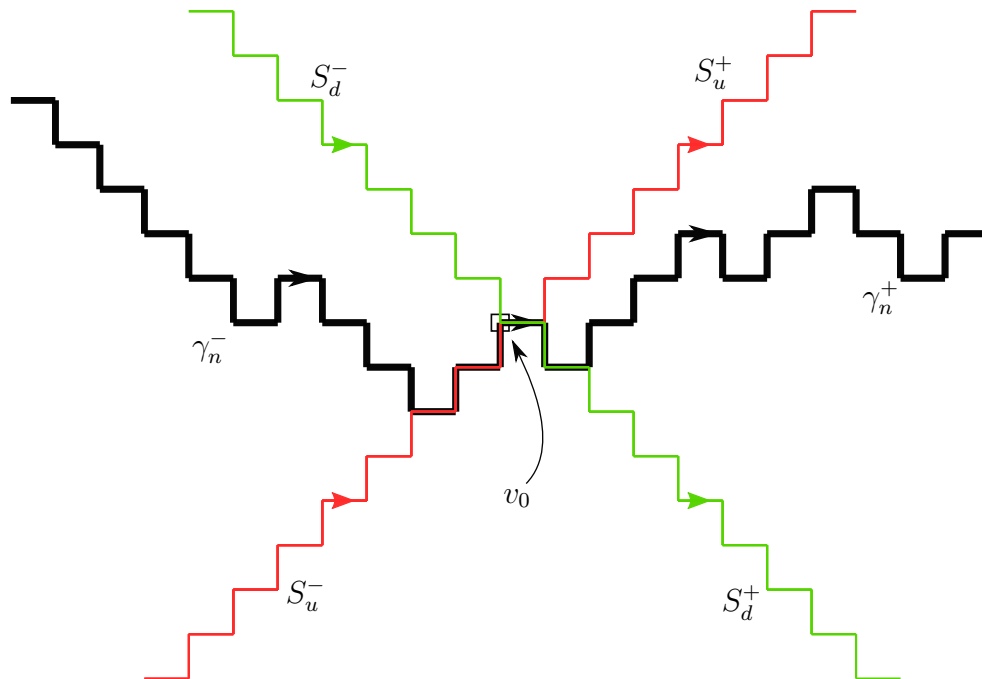


Figure 6.16: The union $\cup_{n \in \mathbb{N}} \gamma_n$ defines doubly-infinite staircase edge-paths S_d and S_u about v_0 .

We claim that no domino in t crosses S_d^- , that is, S_d^- is an edge-path in t . Indeed, if $(u_n^-)_{n \in \mathbb{N}}$ is the 0 sequence $S_d^- = \cup_{n \in \mathbb{N}} \gamma_n^-$, so the claim is obviously true. If $(u_n^-)_{n \in \mathbb{N}}$ is not the 0 sequence there is some $m \geq 0$ with the property that $u_n^- > 0$ for all $n \geq m$. Then by Lemma 6.1.3, each γ_n^- with $n \geq m$ defines a rugged rectangle R_n with side lengths d_n^- along S_d^- and u_n^- along S_u^- . Since $(d_n^-)_{n \in \mathbb{N}}$ is unbounded, the claim is proved.

We will now show that $(d_n^+)_{n \in \mathbb{N}}$ is bounded. Indeed, if it were not, then S_d^+ would be an edge-path in t , as in the preceding paragraph. Since we have shown that S_d^- is an edge-path in t , this would mean S_d is an edge-path in t , and because it is a doubly-infinite staircase edge-path, Lemma 6.1.1 applies. This contradicts our initial assumptions, so $(d_n^+)_{n \in \mathbb{N}}$ must be bounded.

Now, because $(d_n^+)_{n \in \mathbb{N}}$ is bounded, $(u_n^+)_{n \in \mathbb{N}}$ must be unbounded. Like before, this implies S_u^+ is an edge-path in t . Finally, as in the previous paragraph, it follows that $(u_n^-)_{n \in \mathbb{N}}$ is bounded, thus proving the equivalence in (6-1). There are then two cases:

- $(d_n^-)_{n \in \mathbb{N}}$, $(u_n^+)_{n \in \mathbb{N}}$ are unbounded and $(u_n^-)_{n \in \mathbb{N}}$, $(d_n^+)_{n \in \mathbb{N}}$ are bounded;
- $(d_n^-)_{n \in \mathbb{N}}$, $(u_n^+)_{n \in \mathbb{N}}$ are bounded and $(u_n^-)_{n \in \mathbb{N}}$, $(d_n^+)_{n \in \mathbb{N}}$ are unbounded.

In the first case, S_d^- and S_u^+ are edge-paths in t , so their union is the border of a north rugged quadrant in t . In the latter case, S_u^- and S_d^+ are edge-paths in t , so their union is the border of a south rugged quadrant in t .

In other words, v_0 belongs to the ‘tip’ of a north rugged quadrant in t or to the ‘tip’ of a south rugged quadrant in t .

We drew this conclusion under the supposition that (a) holds but not (b). When (b) holds but not (a), the same techniques can be used to conclude that v_0 belongs to the ‘tip’ of an east rugged quadrant in t or to the ‘tip’ of a west rugged quadrant in t .

In particular, since v_0 is free to assume any value in \mathbb{Z}^2 , we discover that every vertex $v \in \mathbb{Z}^2$ belongs to the ‘tip’ of a cardinal rugged quadrant in t .

Let $\mathcal{N}(t)$ be the set of all north rugged quadrants in t ; suppose it is non-empty. We claim the union $Q_{\mathcal{N}} = \cup_{Q \in \mathcal{N}(t)} Q$ is in $\mathcal{N}(t)$. Indeed, Lemma 6.1.4 ensures $Q_{\mathcal{N}}$ is either a north rugged quadrant, a ‘rugged half plane’, or the entire plane. Since t cannot contain doubly-infinite domino staircases, the last two possibilities are excluded and the claim holds. Hence, if $\mathcal{N}(t)$ is non-empty, there is a maximal element $Q_{\mathcal{N}} \in \mathcal{N}(t)$ that contains every $\tilde{Q} \in \mathcal{N}(t)$

This also applies to $\mathcal{S}(t)$, $\mathcal{E}(t)$ and $\mathcal{W}(t)$, respectively the set of all south, east and west rugged quadrants in t , whenever they’re non-empty.

Now, notice none of $\mathcal{N}(t)$, $\mathcal{S}(t)$, $\mathcal{E}(t)$ or $\mathcal{W}(t)$ may be empty, for we have shown that every vertex of \mathbb{Z}^2 belongs to the ‘tip’ of a cardinal rugged quadrant in t , and no fewer than four maximal cardinal rugged quadrants with different types can tile \mathbb{Z}^2 . It follows that t can be decomposed into the four pieces $Q_{\mathcal{N}}$, $Q_{\mathcal{S}}$, $Q_{\mathcal{E}}$ and $Q_{\mathcal{W}}$.

Finally, it’s easy to check the only ways to fit these pieces into a tiling of \mathbb{Z}^2 produce windmill tilings, so we are done. \square

6.2

Back to the torus

Because tilings of \mathbb{T}_L are L -periodic when lifted to \mathbb{Z}^2 , a windmill tiling can never be the tiling of a torus. Theorem 2 then implies the following corollary:

Corollary 6.2.1 (Characterization of tilings of the torus). *Let L be a valid lattice and t a tiling of \mathbb{T}_L . Then exactly one of the following applies:*

1. t admits a flip;
2. t consists entirely of parallel, doubly-infinite domino staircases.

The next proposition shows this characterization can be described in terms of the flux of a tiling.

Proposition 6.2.2. *Let L be a valid lattice and t a tiling of \mathbb{T}_L with flux $\varphi_t \in \mathcal{F}(L)$. Then t admits no flips if and only if $\varphi_t \in \partial Q$.*

Before proving it, we need a lemma.

Lemma 6.2.3. *Let L be a valid lattice, and t a tiling of \mathbb{T}_L . Suppose there is a staircase edge-path γ in t joining w to $w+v$, where $w \in \mathbb{Z}^2$ and $v \in L$. Then t consists entirely of doubly-infinite domino staircases, each parallel to γ .*

Proof. Since t is L -periodic and $v \in L$, t is v -periodic. Thus, for each $n \in \mathbb{Z}$ the translated edge-path $(\gamma + n \cdot v)$ is in t . Now, $l(\gamma)$ is the sum of v ’s coordinates, and because L is valid, that number is even. This means γ ’s first and last edge are different types (horizontal or vertical), which in turn implies the union $\cup_{n \in \mathbb{Z}} (\gamma + n \cdot v)$ is itself a staircase edge-path $\tilde{\gamma}$ in t , except now doubly-infinite.

For any doubly-infinite staircase edge-path, the choice of a single domino fitting it propagates infinitely along the staircase in one direction; the direction

is given by that domino's type (horizontal or vertical). Figure 6.6 in Lemma 6.1.1 illustrates this.

Using the v -periodicity of our tiling, this propagation can be extended infinitely to the other direction, so $\tilde{\gamma}$ has a doubly-infinite domino staircase on each side. For each of those, there is another doubly-infinite staircase edge-path that fits it and is parallel to $\tilde{\gamma}$ (and thus also to γ), so we may repeat the process. The lemma follows. \square

We now prove Proposition 6.2.2.

Proof of Proposition 6.2.2. Suppose first that t admits no flips. In this case, Corollary 6.2.1 implies t consists entirely of parallel, doubly-infinite domino staircases. Consider then a staircase edge-path γ that respects edge orientation, starts at the origin, and fits a staircase in t . Such an edge-path always exists; see the image below. The marked vertex is the origin.

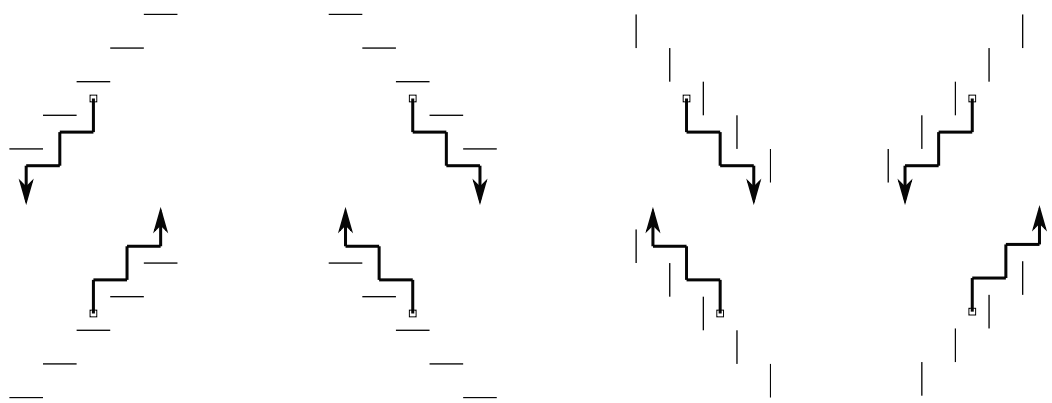


Figure 6.17: Possible domino staircases about the origin and choice of staircase edge-path γ .

Notice that for a suitable and fixed choice of signs, γ contains all vertices of the form $(\pm x, \pm x)$, $x \in \mathbb{N} \setminus \{0\}$. For these vertices, the constructive definition of height function implies $h(\pm x, \pm x) = 2x$. By Lemma 5.2.6, one of those vertices is in L . We thus have

$$\varphi_t(\pm x, \pm x) = \langle \varphi_t, (\pm x, \pm x) \rangle = \frac{1}{4}h(\pm x, \pm x) = \frac{1}{2}x.$$

On the other hand, it holds that

$$\langle \varphi_t, (\pm x, \pm x) \rangle \leq \|\varphi_t\|_1 \cdot x \leq \frac{1}{2}x,$$

where the last inequality follows from the fact that $\varphi_t \in Q$ (see Theorem 1). Combining the two yields $\|\varphi_t\|_1 = \frac{1}{2}$, so $\varphi_t \in \partial Q$ as desired.

Suppose now that $\varphi_t \in \partial Q$. Lemma 5.2.6 guarantees that for each choice of signs in $(\pm x, \pm x)$, there is a vertex in L with that form. For a suitable choice of signs then, there is a vertex in L with that form and

$$\langle \varphi_t, (\pm x, \pm x) \rangle = \|\varphi_t\|_1 \cdot x = \frac{1}{2}x,$$

which implies $h(\pm x, \pm x) = 2x$. We may assume without loss of generality x is positive (otherwise, take $-x$ instead).

Consider the staircase edge-paths that start at the origin and end in $(\pm x, \pm x)$; there are two: both have length $2x$, and one respects edge orientation while the other reverses it. Because $h(\pm x, \pm x) = 2x$, the constructive definition of height functions implies the staircase edge-path that respects edge orientation is in t . By Lemma 6.2.3, t consists entirely of parallel, doubly-infinite domino staircases, so it admits no flips and the proof is complete. \square

We now know that if $\varphi \in \mathcal{F}(L) \cap \text{int}(Q)$, every tiling of \mathbb{T}_L with flux φ admits a flip; in other words, it has a local extremum. It turns out, however, that it must have both a local minimum and a local maximum, that is, it must admit at least two flips.

Proposition 6.2.4. *Let L be a valid lattice and t a tiling of \mathbb{T}_L with height function h and flux $\varphi \in \mathcal{F}(L) \cap \text{int}(Q)$. Then h has both a local minimum and a local maximum.*

Proof. The proof is by contradiction. We will show that if h does not have one kind of local extremum, there is an infinite staircase edge-path in t . We claim in this case Lemma 6.2.3 applies. Indeed, any edge-path γ in \mathbb{Z}^2 can be projected onto an edge-path in \mathbb{R}^2/L ; since \mathbb{R}^2/L is finite, if γ is long enough the projection must self-intersect, so the claim holds. By Lemma 6.2.3, t consists entirely of parallel, doubly-infinite domino staircases, contradicting $\varphi \in \mathcal{F}(L) \cap \text{int}(Q)$.

Suppose h does not have a local maximum. The argument that follows goes similar to the proof of Theorem 2, and is analogous when h does not have a local minimum.

Since no $v_0 \in \mathbb{Z}^2$ is a local maximum of h , there is a list $(v_n)_{n \geq 0} \subset \mathbb{Z}^2$ with $h(v_{n+1}) = h(v_n) + 1$ for all $n \geq 0$ and in which each v_n is neighbour to v_{n+1} . For each $n \in \mathbb{N}$, consider the edge-path $\gamma_n^+ = (v_k)_{k=0}^n$. Because h always changes by $+1$ along an edge on each of these paths, they are by construction edge-paths in t that respect edge orientation, so Lemma 6.1.2 applies to them. Furthermore, because γ_n^+ is always contained in γ_{n+1}^+ , at least one of the statements below is true.

- (a) Each horizontal edge of $\cup_{n \in \mathbb{N}} \gamma_n^+$ has the same orientation;
- (b) Each vertical edge of each $\cup_{n \in \mathbb{N}} \gamma_n^+$ has the same orientation.

If both (a) and (b) hold, $\cup_{n \in \mathbb{N}} \gamma_n^+$ is an infinite staircase edge-path in t , and we are done.

Suppose now (a) holds but not (b); the other case is analogous. Consider the doubly-infinite staircase edge-path S_d^+ (respectively S_u^+) defined by:

- Its vertical edges all point downwards (respectively upwards);
- Its horizontal edges have the same orientation as those in $\cup_{n \in \mathbb{N}} \gamma_n$;
- It respects edge orientation;
- Its starting point is v_0 .

Let u_n^+ be the number of vertical edges pointing upwards in γ_n^+ and d_n^+ be the number of vertical edges pointing downwards in γ_n^+ . Notice $(u_n^+)_{n \in \mathbb{N}}$ and $(d_n^+)_{n \in \mathbb{N}}$ are nondecreasing sequences of nonnegative integers. Furthermore, because $(u_n^+ + d_n^+)$ is the number of vertical edges in γ_n^+ , at least one of $(u_n^+)_{n \in \mathbb{N}}$ and $(d_n^+)_{n \in \mathbb{N}}$ is unbounded.

As in the proof of Theorem 2, Lemma 6.1.3 guarantees that when $(u_n^+)_{n \in \mathbb{N}}$ is unbounded, S_u^+ is in t ; and when $(d_n^+)_{n \in \mathbb{N}}$ is unbounded, S_d^+ is in t . In other words, at least one of S_u^+ and S_d^+ is in t . Since they're both infinite staircase edge-paths, the proof is complete. \square

Let L be a valid lattice. Remember $\mathbb{T}_L = \mathbb{R}^2 / L$, so any vertex of \mathbb{T}_L has an L -equivalence class in \mathbb{Z}^2 ; we will denote v 's equivalence class by $[v]_L$.

Let t be a tiling of \mathbb{T}_L with associated height function h . Because h is L -quasiperiodic, if $v \in \mathbb{Z}^2$ is a local extremum of h , each vertex in $[v]_L$ will also be a local extremum of the same kind. In other words, we may perform a flip round each vertex in $[v]_L$. We call this process an L -flip (round v): it is how a flip on a tiling of \mathbb{T}_L manifests in the planar, L -periodic representation of t .

An L -flip round v preserves h 's quasiperiodicity. This is clear when v is not in the equivalence class of the origin $[0]_L$; in this case, the height change on each vertex in $[v]_L$ will be the same, and no height change will occur on other vertices.

When v is in $[0]_L$, the situation is different. The toroidal height functions we consider take the base value 0 at the origin, so performing an L -flip round the origin does not change the value h takes on it; instead, it changes the value

on each vertex **not** in $[0]_L$. Nonetheless, since that change is the same across all such vertices³, h 's quasiperiodicity is preserved in this case too.

This also shows an L -flip preserves a tiling's flux value, because in either case the value h takes on L does not change — observe that the equivalence class of the origin is L itself.

If the reader had any thoughts about how Proposition 6.2.4 and Corollary 5.0.4 were contradictory, the discussion above should have cleared those. There is no conflict: for each flux $\varphi \in \mathcal{F}(L) \cap \text{int}(Q)$, $h_{\min}^{L,\varphi}$ must have all of its local maxima lying on $[0]_L$, so that performing an L -flip round those vertices does not contradict the minimality of $h_{\min}^{L,\varphi}$. By the same token, $h_{\max}^{L,\varphi}$ must have all of its local minima lying on $[0]_L$.

We are now poised to prove the flux-analogue of Proposition 3.1.5 for the torus.

Proposition 6.2.5. *Let L be a valid lattice and $\varphi \in \mathcal{F}(L) \cap \text{int}(Q)$. Let $h_{\min}^{L,\varphi}$ be minimal over height functions of tilings of \mathbb{T}_L with flux φ . Let $h \neq h_{\min}^{L,\varphi}$ be a height function associated to a tiling t of \mathbb{T}_L with flux φ . Then there is an L -flip on t that produces a height function $\tilde{h} \leq h$ with $\tilde{h} < h$ on one equivalence class of vertices of \mathbb{Z}^2 .*

Proof. By Proposition 6.2.4, we know h has a local maximum, but this is not enough. Our previous consideration makes it clear we need to show h has a local maximum on a vertex that is not in the equivalence class of the origin.

Consider the difference $g = h - h_{\min}^{L,\varphi}$. By Proposition 5.0.2, g is L -periodic and in particular bounded. Moreover, Proposition 4.1.1 means g takes nonnegative values in $4\mathbb{Z}$. Let V be the set of vertices of \mathbb{Z}^2 on which g is maximum. We assert that h has a local maximum lying on V . Notice this is sufficient: since $h \neq h_{\min}^{L,\varphi}$, g necessarily assumes positive values on V , so $[0]_L$ does not intersect V (because g is 0 at the origin).

We prove the assertion by contradiction; suppose V contained no local maximum of h and choose any $v_0 \in V$. Since v_0 is not a local maximum of h there must be a neighboring vertex $v_1 \in \mathbb{Z}^2$ for which $h(v_1) > h(v_0)$. We claim $v_1 \in V$. Indeed, let e be the edge joining v_0 to v_1 . The possible height changes along e are either $+1$ and -3 , or -1 and $+3$, depending on e 's orientation as induced by the coloring of \mathbb{Z}^2 . In either case, $h_{\min}^{L,\varphi}$ and h must both increase along e , for if $h_{\min}^{L,\varphi}$ decreased along e , it would contradict the maximality of g on v_0 .

³When the origin is a local maximum, the change is $+4$; when the origin is a local minimum, the change is -4 .

It follows that v_1 is also not a local maximum of h , so we may repeat the process. This produces a list $(v_n)_{n \geq 0} \subset V$ with $h(v_{n+1}) > h(v_n)$ for all $n \geq 0$ and in which each v_n is neighbour to v_{n+1} .

Notice we may assume $h(v_{n+1}) = h(v_n) + 1$. Indeed, let $t_{\min}^{L,\varphi}$ be the tiling associated to $h_{\min}^{L,\varphi}$. When $h(v_{n+1}) = h(v_n) + 3$, the edge joining v_n to v_{n+1} crosses a domino in both t and $t_{\min}^{L,\varphi}$, so going from v_n to v_{n+1} round that domino is allowed in both tilings. It's clear that each edge traversed this way increases h by $+1$, so there is no loss of generality in the assumption.

At this point, the proof of Proposition 6.2.4 can be applied verbatim here: the existence of one such list implies the existence of an infinite staircase edge-path in both t and $t_{\min}^{L,\varphi}$. By Lemma 6.2.3, both t and $t_{\min}^{L,\varphi}$ must consist entirely of parallel, doubly-infinite domino staircases, contradicting $\varphi \in \text{int}(Q)$. \square

Like in the planar case, because the situation is finite, Proposition 6.2.5 tells us any tiling of \mathbb{T}_L with flux φ can be taken by a sequence of L -flips to $t_{\min}^{L,\varphi}$. The following corollary is immediate.

Corollary 6.2.6 (Flip-connectedness on the torus). *Let L be a valid lattice and $\varphi \in \mathcal{F}(L) \cap \text{int}(Q)$. Any two distinct tilings of \mathbb{T}_L with flux φ can be joined by a sequence of flips.*

7

Kasteleyn matrices for the torus

This chapter is devoted to adapting the construction of Kasteleyn matrices for tori.

Consider the dual graph $G(\mathbb{Z}^2)$ of \mathbb{Z}^2 . We will represent each unit square of \mathbb{Z}^2 by the vertex in its center, so each vertex of $G(\mathbb{Z}^2)$ lies in $(\mathbb{Z} + \frac{1}{2})^2$. Let L be a valid lattice. As with vertices, edges on $G(\mathbb{Z}^2)$ have L -equivalence classes: two edges belong to the same class if they're related by a translation in L . Given an edge e on $G(\mathbb{Z}^2)$, its equivalence class will be denoted by $[e]_L$.

We first tackle the problem of determining an L -Kasteleyn signing of $G(\mathbb{Z}^2)$, that is, an assignment of plus and minus signs to equivalence classes of edges on $G(\mathbb{Z}^2)$ with the following property: for every four edges on $G(\mathbb{Z}^2)$ that make up a square, the product of their signs is -1 (where of course the sign of an edge is the sign of its equivalence class).

Similar to the planar case, these conditions guarantee that whenever we perform an L -flip on a tiling of \mathbb{T}_L , the total sign on the corresponding summands of the Kasteleyn determinant does not change, but this will become clear later.

An initial observation is that our usual assignment of minus signs to alternating lines of edges on $G(\mathbb{Z}^2)$ is generally not an L -Kasteleyn signing. Indeed, if e is an edge and $v \in \mathcal{O} \cap L$, then e and $e + v$ have different signs.

Rather than show the existence of an L -Kasteleyn signing for a given valid lattice L , we will exhibit a universal Kasteleyn signing, which applies to all valid lattices.

Recall the special brick wall tilings, defined just before Lemma 6.1.1 and shown in Figure 6.5. For any $v \in \mathcal{E} \sqcup \mathcal{O}$, translation by v is color preserving and therefore a symmetry of each brick wall. In particular, for each valid lattice L and brick wall b , b is L -periodic and thus a tiling of \mathbb{T}_L . If all four brick walls are represented on $G(\mathbb{Z}^2)$, it's easy to see that for every four edges on $G(\mathbb{Z}^2)$ that make up a square, each of those edges lies in a different brick wall. We may thus use each brick wall b to define a universal Kasteleyn signing: simply

assign -1 to every edge on $G(\mathbb{Z}^2)$ that is also in b , and $+1$ to every other edge (or vice-versa, exchanging -1 with $+1$).

In light of this, we expand on the significance of Proposition 5.1.1.

Proposition 7.0.7. *Let L be a valid lattice. Each of the points $\pm(\frac{1}{2}, 0)$ and $\pm(0, \frac{1}{2})$ is in $\mathcal{F}(L)$. For each of those, there is only one tiling of \mathbb{T}_L which realizes that flux, and it is a brick wall.*

Proof. That the points are in $\mathcal{F}(L)$ is provided by Proposition 5.1.1 and Theorem 1. We first show that the fluxes of the four brick walls are given by these four points.

Let L be a valid lattice and b a brick wall with flux φ_b and associated height function h_b . Each marked edge-path on the image below is a staircase that respects edge orientation and by Lemma 5.2.6 intercepts L . The marked vertex is the origin.




Figure 7.1: Choice of staircase edge-paths for each brick wall.

This means there are nonzero $x, y \in \mathbb{Z}$ with $(x, x), (y, -y) \in L$ and

$$\begin{cases} h_b(x, x) = 4 \cdot \langle \varphi_b, (x, x) \rangle = 2|x| \\ h_b(y, -y) = 4 \cdot \langle \varphi_b, (y, -y) \rangle = 2|y| \end{cases} \quad (7-1)$$

Each of the four brick walls corresponds to a choice of signs in $x, y \in \mathbb{Z}$ (positive or negative), and the corresponding solution of system (7-1) for the

flux φ_b yields one of $\pm(\frac{1}{2}, 0), \pm(0, \frac{1}{2})$. By inspection, $b_{\mathcal{E}}$ has flux $(\frac{1}{2}, 0)$, $b_{\mathcal{N}}$ has flux $(0, \frac{1}{2})$, $b_{\mathcal{W}}$ has flux $(-\frac{1}{2}, 0)$ and $b_{\mathcal{S}}$ has flux $(0, -\frac{1}{2})$.

It remains to show that a tiling with flux given by one of the four points is a brick wall. We will do this for the flux $(\frac{1}{2}, 0)$, but for other points the reasoning is analogous. Let L be a valid lattice and t a tiling of \mathbb{T}_L with flux $(\frac{1}{2}, 0)$ and associated height function h . By Lemma 5.2.6, there are positive integers x_0, x_1 with $(x_0, x_0), (x_1, -x_1) \in L$. Let γ_0 be a 3-1 staircase edge-path joining the origin to (x_0, x_0) and γ_1 be a 2-4 staircase edge-path joining the origin to $(x_1, -x_1)$, like the edge-paths in Figure 7.1a; notice they respect edge orientation.

Clearly, $l(\gamma_i) = 2x_i$. Since $h(x_0, x_0) = 4 \cdot \langle (\frac{1}{2}, 0), (x_0, x_0) \rangle = 2x_0$ and similarly $h(x_1, -x_1) = 2x_1$, the constructive definition of height functions implies γ_0 and γ_1 are both edge-paths in t . Applying Lemma 6.2.3 to each of γ_0 and γ_1 , it follows that t consists entirely of 3-1 doubly-infinite domino staircases **and** entirely of 2-4 doubly-infinite domino staircases, so it must be a brick wall (in particular, since it has flux $(\frac{1}{2}, 0)$, it must be $b_{\mathcal{E}}$). \square

Fix once and for all a choice of universal Kasteleyn signing: assign -1 to every edge on $G(\mathbb{Z}^2)$ that is also in $b_{\mathcal{N}}$. We now describe how to construct a Kasteleyn matrix K for a torus \mathbb{T}_L .

Let x_0 be the smallest positive integer with $v_0 = (x_0, 0) \in L$. Now let y_1 be the smallest positive integer for which the set $Y = \{(x, y_1) \mid x \in \mathbb{R}\}$ intersects L . Choose the vertex $v_1 = (x_1, y_1) \in Y \cap L$ with $0 \leq x_1 < x_0$; notice there is always exactly one such v_1 . Indeed, if $v = (x, y_1) \in Y \cap L$, there is a unique integer m with $m \cdot x_0 \leq x < (m+1) \cdot x_0$, so we may take $v_1 = v - m \cdot v_0$; if there were more than one, their difference would contradict the minimality of x_0 . Arguments similar to this show that $\{v_0, v_1\}$ generates L .

We now take the following fundamental domain: let $u(0)$ be the straight line edge-path in \mathbb{Z}^2 joining the origin to v_0 . Let $u(1)$ be the L-shaped edge-path in \mathbb{Z}^2 joining the origin to v_1 that never coincides with u_0 away from the origin. Our fundamental domain $D_L \subset \mathbb{R}^2$ is the rectangle with vertices $(0, 0), (x_0, 0), (0, y_1)$ and (x_0, y_1) . Notice $u(0)$ and $u(1)$ always lie in the boundary of D_L , and these edge-paths are used in the flux definition that counts cross-over dominoes.

Enumerate each of D_L 's black squares (starting from 1) and do the same to white squares. There is an obvious correspondence between squares on D_L and equivalence classes of vertices on $G(\mathbb{Z}^2)$. Also, observe that any equivalence class of edges on $G(\mathbb{Z}^2)$ is given by two equivalence classes of vertices on $G(\mathbb{Z}^2)$,

its endpoints. With this in mind, we now assign weights to each equivalence class of edges on $G(\mathbb{Z}^2)$. Let e_{ij} be the edge joining the i -th back vertex to the j -th white vertex: if no such edge exists, we assign the weight 0 to $[e_{ij}]_L$; otherwise, we assign its corresponding Kasteleyn sign (either $+1$ or -1) to it.

Next, if there is an edge in $[e_{ij}]_L$ that crosses $u(0)$, we will multiply the weight of $[e_{ij}]_L$ by either q_0 or q_0^{-1} ; we now explain how the exponent is chosen. Remember that edges on the dual graph represent dominoes, and whenever an edge-path crosses a domino on a tiling, the height function of that tiling changes by either $+3$ or -3 along that edge-path. When the edge in $[e_{ij}]_L$ that crosses $u(0)$ corresponds to a height change of $+3$ along $u(0)$, we choose q_0 ; when it corresponds to a height change of -3 , we choose q_0^{-1} . Observe that because $u(0)$ joins the origin to v_0 (and v_0 is in a basis of L), there is at most one edge in $[e_{ij}]_L$ that crosses $u(0)$. Of course, there may be none, and in that case this step does not change the previously assigned weight of $[e_{ij}]_L$.

Finally, we repeat the last step for $u(1)$ and q_1 or q_1^{-1} . Notice the effect of Kasteleyn signs and each of the q_k 's is cumulative!

Now that all equivalence classes of edges on $G(\mathbb{Z}^2)$ are assigned their corresponding weights, the matrix entry $K(i, j)$ is simply the weight of $[e_{ij}]_L$. We provide an example of this construction in the next page.


Similarly to the planar case, each nonzero term in the combinatorial expansion of $\det(K)$ can be seen as an L -periodic matching of $G(\mathbb{Z}^2)$. In other words, it can be seen as an L -periodic tiling of \mathbb{Z}^2 , or a tiling of \mathbb{T}_L . For each tiling t of \mathbb{T}_L , let K_t be its corresponding nonzero term in the combinatorial expansion of $\det(K)$. From the construction of K and the flux definition via counting cross-over dominoes, the following is clear:

$$\forall \varphi \in \mathcal{F}(L), \exists n_0, n_1 \in \mathbb{Z}, \forall \text{ tiling } t \text{ of } \mathbb{T}_L \text{ with flux } \varphi, K_t = \pm q_0^{n_0} q_1^{n_1}.$$


In fact, whenever $v_k \in \mathcal{E}$, the flux through v_k is precisely n_k ; whenever $v_k \in \mathcal{O}$, the flux through v_k is $n_k + \frac{1}{2}$. Furthermore, the use of a Kasteleyn signing in the construction of K implies via Corollary 6.2.6 that these terms are all identically signed whenever $\varphi \in \mathcal{F}(L) \cap \text{int}(Q)$. Thus, for each monomial of the form $c_{ij} \cdot q_0^i q_1^j$ in the full expansion of $\det(K)$, if $q_0^i q_1^j$ corresponds to a flux value $\varphi \in \mathcal{F}(L) \cap \text{int}(Q)$, then $|c_{ij}|$ is the number of tilings of \mathbb{T}_L with flux φ .

It turns out this is also true for fluxes in $\mathcal{F}(L) \cap \partial Q$.

Proposition 7.0.8. *Let L be a valid lattice and K a Kasteleyn matrix for \mathbb{T}_L . Let $\varphi \in \mathcal{F}(L)$ and $c_{ij} \cdot q_0^i q_1^j$ be the monomial in the full expansion of $\det(K)$ that corresponds to φ . Then $|c_{ij}|$ is the number of tilings of \mathbb{T}_L with flux φ .*



(a) A lattice L with $v_0 = (14, 0)$ and $v_1 = (4, 2)$. The fundamental domain D_L is represented by the red rectangle, and its squares are enumerated.



(b) The enumeration applied to vertices on $G(\mathbb{Z}^2)$. Negative edges are dashed.

$$\begin{pmatrix}
 -1 & 0 & 0 & 0 & 0 & 0 & q_1^{-1} & 0 & q_0 & 0 & 0 & 0 & 0 & 1 \\
 1 & -1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & q_0 & 0 & 0 & 0 & 0 \\
 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & q_0 & 0 & 0 & 0 \\
 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & q_0 & 0 & 0 \\
 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & q_0 & 0 \\
 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & q_0 q_1 \\
 0 & 0 & 0 & 0 & 0 & 1 & -1 & q_0 q_1 & 0 & 0 & 0 & 0 & 1 & 0 \\
 1 & 0 & 0 & 0 & 0 & \frac{1}{q_0 q_1} & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 1 \\
 0 & 1 & 0 & 0 & 0 & 0 & \frac{1}{q_0 q_1} & 1 & -1 & 0 & 0 & 0 & 0 & 0 \\
 q_0^{-1} & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 \\
 0 & q_0^{-1} & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 \\
 0 & 0 & q_0^{-1} & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\
 0 & 0 & 0 & q_0^{-1} & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 \\
 0 & 0 & 0 & 0 & q_0^{-1} & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & -q_1
 \end{pmatrix}$$

(c) The resulting Kasteleyn matrix K .

Figure 7.2: An example construction of a Kasteleyn matrix for a torus.

We know Proposition 7.0.8 needs to be proved only for $\varphi \in \mathcal{F}(L) \cap \partial Q$, but before doing it we will study the structure of $\mathcal{F}(L) \cap \partial Q$.

7.1

The structure of $\mathcal{F}(L) \cap \partial Q$

Observe that in Theorem 2, when t consists entirely of parallel, doubly-infinite domino staircases, these staircases are in fact the same type. This is because every doubly-infinite staircase edge-path that fits one of these domino staircases actually fits two of them (one on each side), so they are both the same type. By induction, this applies to them all.

Consider the boundary of $Q \subset \mathbb{R}^2$; it is a square. Let $Q_1 \subset \partial Q$ be the side of the square lying on the first quadrant of \mathbb{R}^2 , and similarly for Q_2, Q_3 and Q_4 . We will use the observation above to classify the tilings in each $\mathcal{F}(L) \cap Q_k$.

Proposition 7.1.1. *Let L be a valid lattice. For each $k \in \{1, 2, 3, 4\}$, there is a type of domino staircase such that each tiling of \mathbb{T}_L with flux in $\mathcal{F}(L) \cap Q_k$ consists entirely of doubly-infinite domino staircases which are all that type.*

Proof. We will prove this for $k = 1$, but for other values of k the proof is analogous. Let t be a tiling of \mathbb{T}_L with associated height function h and flux $\varphi \in \mathcal{F}(L) \cap Q_1$. We will show t consists entirely of 3-1 doubly-infinite domino staircases.

Write $\varphi = (a, b)$. Because $\varphi \in Q_1$, $a + b = \frac{1}{2}$. By Lemma 5.2.6, there is some positive integer x_0 with $(x_0, x_0) \in L$. We thus have $h(x_0, x_0) = 4 \cdot \langle \varphi, (x_0, x_0) \rangle = 2x_0$. Consider the edge-path γ_0 below. As in the proof of Proposition 7.0.7, it respects edge orientation, and is a 3-1 staircase edge-path joining the origin to (x_0, x_0) .





Figure 7.3: The edge-path γ_0 ; the marked vertex is the origin.

Like before, it's clear $l(\gamma_0) = 2x_0$, so by the constructive definition of height functions it must be an edge-path in t . We may thus apply Lemma 6.2.3 to it, from which we conclude t consists entirely of 3-1 doubly-infinite domino staircases, as desired. \square

Let L be a valid lattice and t be a tiling of \mathbb{T}_L with flux $\varphi \in \mathcal{F}(L)$. By inspection, we have that:

- $\varphi \in Q_1 \Rightarrow t$ consists entirely of 3-1 doubly-infinite domino staircases
- $\varphi \in Q_2 \Rightarrow t$ consists entirely of 4-2 doubly-infinite domino staircases
- $\varphi \in Q_3 \Rightarrow t$ consists entirely of 1-3 doubly-infinite domino staircases
- $\varphi \in Q_4 \Rightarrow t$ consists entirely of 2-4 doubly-infinite domino staircases

Of course, the converse is also true. This can be schematically represented by the diagram below:



PUC-Rio - Certificação Digital Nº 1321816/CA

Figure 7.4: Schematic representation of the behavior of tilings with flux in ∂Q .

Notice how each brick wall belongs to two distinct Q_k 's: they are as transition tilings between their respective Q_k 's. This will become clearer with the concept of stairflips.

Observe that for any doubly-infinite domino staircase, the dominoes in it are all the same type (horizontal or vertical). A *stairflip* on a doubly-infinite domino staircase S is the process of exchanging all dominoes in S by dominoes of the other type (horizontal or vertical). It is clear a stairflip on

S produces a new doubly-infinite domino staircase \tilde{S} ; furthermore, because the doubly-infinite staircase edge-paths that fit S and \tilde{S} are the same, a stairflip preserves the type of a doubly-infinite domino staircase. Of course, performing two successive stairflips produces no change.




Figure 7.5: An example of stairflip. Remember the staircases are doubly-infinite!

Let L be a valid lattice and S a doubly-infinite domino staircase in \mathbb{Z}^2 . We define its L -equivalence class $[S]_L$ as the set of all doubly-infinite domino staircases \tilde{S} in \mathbb{Z}^2 that are the same type as S and satisfy $\Pi(\tilde{S}) = \Pi(S)$, where $\Pi : \mathbb{R}^2 \rightarrow \mathbb{T}_L$ is the projection map. Notice all dominoes in all staircases of $[S]_L$ must be the same type (horizontal or vertical).

Like with flips, we may define an L -stairflip on a doubly-infinite domino staircase S : simply apply a stairflip to each staircase in $[S]_L$.

Henceforth, we will use interchangeably the terms type-1 staircase and 3-1 staircase, and similarly for types 2, 3 and 4. Let $\text{Stair}(L) = \{ [S]_L \mid S \text{ is a doubly-infinite domino staircase in } \mathbb{Z}^2 \}$ and for $k \in \{1, 2, 3, 4\}$ let $\text{Stair}(L; k) = \{ [S]_L \mid S \text{ is a type-}k \text{ doubly-infinite domino staircase in } \mathbb{Z}^2 \}$. Clearly, $\text{Stair}(L)$ is finite. A perhaps less obvious observation is that $\text{Stair}(L; 1)$ and $\text{Stair}(L; 3)$ have the same cardinality. Indeed, translation by $e_1 = (1, 0)$ is a bijection between 1-3 and 3-1 doubly infinite domino staircases, which extends into a bijection between corresponding equivalence classes. By the same token, $\text{Stair}(L; 2)$ and $\text{Stair}(L; 4)$ have the same cardinality.

We may further decompose $\text{Stair}(L; k)$ into two disjoint subsets. Let $\text{Stair}(L; k; \text{vert})$ be the set of equivalence classes in $\text{Stair}(L; k)$ whose domino staircases are all made up of vertical dominoes. Define $\text{Stair}(L; k; \text{hor})$ similarly for horizontal dominoes. The L -stairflip is an obvious bijection between them.

Define the L -stairflip operator $\xi_L : \text{Stair}(L) \rightarrow \text{Stair}(L)$. One may restrict it to $\text{Stair}(L; k) \rightarrow \text{Stair}(L; k)$, and furthermore to $\text{Stair}(L; k; \text{vert}) \rightarrow \text{Stair}(L; k; \text{hor})$. Notice that in each case, ξ_L is an involution.

We may use these sets to describe tilings of \mathbb{T}_L with flux in ∂Q . Let t be a tiling of \mathbb{T}_L with flux in Q_k . We know t consists entirely type- k doubly-infinite

domino staircases, each of which corresponds to an element of $\text{Stair}(L; k)$. We may thus identify t with a subset $C(t)$ of $\text{Stair}(L; k)$. What can we say about $C(t)$? The crucial observation is that for each $[S]_L \in \text{Stair}(L; k)$ we have $[S]_L \in C(t)$ if and only if $\xi_L([S]_L) \notin C(t)$. Clearly, $[S]_L$ and $\xi_L([S]_L)$ cannot both be in $C(t)$, for their lifts overlap¹. On the other hand, one of them must be in $C(t)$, for otherwise t would not consist entirely of type- k doubly-infinite domino staircases. This property immediately implies $\text{card}(C(t)) = \text{card}(\text{Stair}(L; k; \text{vert})) = \text{card}(\text{Stair}(L; k; \text{hor}))$, but there's more.

We say a set $C \subset \text{Stair}(L; k)$ is ξ_L - k -exclusive if it satisfies

$$\forall [S]_L \in \text{Stair}(L; k), [S]_L \in C \iff \xi_L([S]_L) \notin C.$$

By the preceding paragraph, every tiling t of \mathbb{T}_L with flux in Q_k corresponds to a ξ_L - k -exclusive set $C(t)$. Nonetheless, by Corollary 6.2.1 and Propositions 6.2.2 and 7.1.1, every ξ_L - k -exclusive set C also corresponds to a tiling t_C of \mathbb{T}_L with flux in Q_k . The correspondence is obvious: simply lift the elements in C , producing an L -periodic tiling t_C of \mathbb{Z}^2 that consists entirely of type- k doubly-infinite domino staircases.

Notice $\text{Stair}(L; k; \text{vert})$ is a ξ_L - k -exclusive set, and it corresponds to a brick wall b . It is clear any such set can be obtained from $\text{Stair}(L; k; \text{vert})$ by choosing a number of its elements and applying ξ_L to them. In other words, any tiling of \mathbb{T}_L with flux in Q_k can be obtained from b by applying an L -stairflip to each of a number of equivalence classes of doubly-infinite domino staircases in b . Moreover, this also means the number of tilings of T_L with flux in Q_k is 2^{c_k} , where $c_k = \frac{1}{2} \text{card}(\text{Stair}(L; k)) = \text{card}(\text{Stair}(L; k; \text{vert}))$: for each element of $\text{Stair}(L; k; \text{vert})$, choose whether or not to apply ξ_L to it.

Observe that c_k depends only on the parity of k , i.e. $c_1 = c_3$ and $c_2 = c_4$. Moreover, they provide a way to count the number of tilings with flux in ∂Q :

$$2^{c_1} + 2^{c_2} + 2^{c_3} + 2^{c_4} - 4 = 2 \cdot (2^{c_1} + 2^{c_2} - 2).$$

Here, we subtract 4 because each brick wall is counted twice — each belongs to two Q_k 's.

Proposition 7.1.2. *Let L be a valid lattice and $k \in \{1, 2, 3, 4\}$. As above, each tiling t of \mathbb{T}_L with flux in $\mathcal{F}(L) \cap Q_k$ corresponds to a unique ξ_L - k -exclusive set; call it $C(t)$. For each such tiling, let $n_k^{\text{vert}}(t) = \text{card}(C(t) \cap \text{Stair}(L; k; \text{vert}))$. If t is a tiling of \mathbb{T}_L with flux in $\mathcal{F}(L) \cap Q_k$, the flux of t depends only on $n_k^{\text{vert}}(t)$.*

¹Recall the projection map $\Pi : \mathbb{R}^2 \rightarrow \mathbb{T}_L$.

Proof. We will prove for $k = 1$, but for other values of k the proof is analogous.

Let t be a tiling of \mathbb{T}_L with flux $\varphi \in \mathcal{F}(L) \cap Q_1$ and associated height function h_t . By Proposition 7.1.1, t consists entirely of 3-1 doubly-infinite domino staircases, so each 3-1 doubly-infinite staircase edge-path in \mathbb{Z}^2 is in t . In particular, there is one such path through the origin, so for each $x \in \mathbb{Z}$, $h_t(x, x) = 2x$. Notice this implies the value h_t takes on a vertex of the form (x, x) in L is the same for all tilings of \mathbb{T}_L with flux in $\mathcal{F}(L) \cap Q_1$.

For $s \in \mathbb{R}$, let $\alpha(s) = (s, \frac{1}{2})$ be a continuous path in \mathbb{R}^2 . As α is traversed, it crosses each 3-1 doubly-infinite domino staircase in t . Let S_0 be the staircase containing $\alpha(\frac{1}{2})$, and for each $i \in \mathbb{Z}$ let S_{i+1} be the first staircase α crosses after S_i . See the image below.




Figure 7.6: Enumerating the S_i 's with α . The marked vertex is the origin.

Observe that $[S]_L = [\tilde{S}]_L$ if and only if there is some $v \in L$ with $S+v = \tilde{S}$. This means $[S_i]_L = [S_i + v]_L$ for all $i \in \mathbb{Z}$ and $v \in L$. Let c be the smallest positive integer for which $[S_0]_L = [S_c]_L$. Letting $u \in L$ be such that $S_0+u = S_c$, it's easy to see that $S_i + u = S_{i+c}$ for all $i \in \mathbb{Z}$. In particular, $[S_i]_L = [S_{i+c}]_L$ for all $i \in \mathbb{Z}$, and no smaller positive integer may have this property. Notice this also implies $[S_i]_L = \{S_{i+k \cdot c} \mid k \in \mathbb{Z}\}$ for all $i \in \mathbb{Z}$. Indeed, suppose it were $\{S_{i+k \cdot c} \mid k \in \mathbb{Z}\} \subsetneq [S_i]_L$. In this case, there is some integer j not of the form $i + k \cdot c$ with $S_j \in [S_i]_L$. There must be an integer m such that S_j lies between $S_{i+m \cdot c}$ and $S_{i+(m+1) \cdot c}$. Let $w \in L$ have the property that $S_{i+m \cdot c} + w = S_j$. Then $[S_0 + w]_L = [S_0]_L$, and letting \tilde{c} be such that $S_0 + w = S_{\tilde{c}}$ it suffices to note that $0 < \tilde{c} < c$, contradicting the minimality of c .

We claim $C(t) = \{[S_0]_L, [S_1]_L, \dots, [S_{c-1}]_L\}$. Indeed, since $[S_i]_L = [S_{i+c}]_L$ for all $i \in \mathbb{Z}$, $C(t) \subseteq \{[S_0]_L, [S_1]_L, \dots, [S_{c-1}]_L\}$. On the other hand, because $[S_i]_L = \{S_{i+k \cdot c} \mid k \in \mathbb{Z}\}$ for all $i \in \mathbb{Z}$, the elements in $\{[S_0]_L, [S_1]_L, \dots, [S_{c-1}]_L\}$ are all distinct, so the inclusion must be an equality. Notice this also implies $c = \text{card}(C(t)) = \text{card}(\text{Stair}(L; 1; \text{vert}))$, which does not depend on t .

Let γ be the doubly-infinite staircase edge-path through the origin that fits S_0 . Let $R \subset \mathbb{R}^2$ be the connected region enclosed by γ and $\gamma + u$. We claim $\text{int}(R) \cap L = \emptyset$. Indeed, $\gamma + u$ fits S_c on the same side that γ fits S_0 . If there were some $v \in \text{int}(R) \cap L$, then there would be some integer $0 < j < c$ with $S_j = S_0 + v \neq S_c$, contradicting the minimality of c .

Observe that for any tiling of \mathbb{T}_L , its flux is entirely determined by the value its height function takes on two linearly independent vectors on L . By Lemma 5.2.6, γ intersects L away from the origin, say at w ; and because $S_c = S_0 + u \neq S_0$, u is not of the form (x, x) . In other words, u and w are linearly independent, so φ is entirely defined by the value h_t takes on them. Now, we've shown the value h_t takes on vertices of the form (x, x) in L is the same for all tilings of \mathbb{T}_L with flux in $\mathcal{F}(L) \cap Q_1$, so φ is entirely defined by the value h_t takes on u .

Let β_0 be the horizontal edge-path in \mathbb{Z}^2 of length $2c$ joining the origin to $\gamma + u$; its endpoint is $(2c, 0)$. Let β_1 be the edge-path in $\gamma + u$ joining $(2c, 0)$ to u , so $\beta_0 * \beta_1$ is an edge-path in \mathbb{Z}^2 joining the origin to $u \in L$. Notice γ is a 3-1 staircase, so it is an edge-path in each tiling of \mathbb{T}_L with flux in $\mathcal{F}(L) \cap Q_1$. Of course, this means $\gamma + u$ and β_1 also have this property, so each such tiling has the same height change along β_1 . Thus, $h_t(u)$ depends only on the height change along β_0 .

Each edge joining vertices $(2k - 1, 0)$ and $(2k, 0)$ is an edge on a 3-1 staircase, so it is an edge on each tiling of \mathbb{T}_L with flux in $\mathcal{F}(L) \cap Q_1$. Each edge joining vertices $(2k, 0)$ and $(2k + 1, 0)$ is an edge that crosses S_k . It follows that the height change along β_0 depends only on these latter edges, and thus depends only on how many dominoes β_0 crosses along those. Now, it's easy to check that for each horizontal edge crossing S_k , that edge crosses a domino if and only if S_k is made up of vertical dominoes, that is, if and only if $[S_k]_L \in \text{Stair}(L; 1; \text{vert})$. The Proposition follows from the fact that β_0 crosses each equivalence class in $C(t)$. \square

We saw that any ξ_L - k -exclusive set can be obtained from $\text{Stair}(L; k; \text{vert})$ by choosing a number of its elements and applying ξ_L to them. By Proposition 7.1.2, this number entirely determines the flux of the corresponding tiling of \mathbb{T}_L , regardless of the choice of elements themselves. We also knew how to count the total tilings of \mathbb{T}_L with flux in $\mathcal{F}(L) \cap Q_k$, but now we have a way to count the tilings of \mathbb{T}_L with a given flux $\varphi \in \mathcal{F}(L) \cap Q_k$. Indeed, let $c_k = \text{card}(\text{Stair}(L; k; \text{vert}))$. By Proposition 7.1.2, $n_k^{\text{vert}}(t)$ is the same for each tiling t of \mathbb{T}_L with flux φ ; we may thus speak of $n_k^{\text{vert}}(\varphi)$. The number of tilings

of \mathbb{T}_L with flux φ is then simply $(n_k^{\text{vert}}(\varphi)^{c_k})$, corresponding to a choice of $n_k^{\text{vert}}(\varphi)$ elements in $\text{Stair}(L; k; \text{vert})$ to keep, and applying ξ_L to the others.

Let L be a valid lattice and fix $k \in \{1, 2, 3, 4\}$. Let t be any tiling of \mathbb{T}_L with flux $\varphi \in \mathcal{F} \cap Q_k$ and ξ_L - k -exclusive set $C(t)$. If $C(t) \neq \text{Stair}(L; k; \text{hor})$, then $C(t) \cap \text{Stair}(L; k; \text{vert})$ is nonempty. Let \tilde{C} be any set obtained from $C(t)$ by choosing an element in $C(t) \cap \text{Stair}(L; k; \text{vert})$ and applying ξ_L to it. By Proposition 7.1.2, the flux of the tiling that corresponds to \tilde{C} does not depend on the choice of element; let $\xi_{L;k}^{\text{vert}}(\varphi)$ be that flux.

Proposition 7.1.3. *The vector $\varphi - \xi_{L;k}^{\text{vert}}(\varphi)$ is constant and nonzero across all $\varphi \in (\mathcal{F}(L) \cap Q_k) \setminus \{(0, \frac{1}{2}), (0, -\frac{1}{2})\}$.*

Proof. Notice b_N and b_S are the only tilings whose ξ_L - k -exclusive sets may be given by $\text{Stair}(L; k; \text{hor})$, and their fluxes are respectively $(0, \frac{1}{2})$ and $(0, -\frac{1}{2})$. Since these are excluded in the statement, $\xi_{L;k}^{\text{vert}}(\varphi)$ is always well-defined.

Let t, \tilde{t} be tilings of \mathbb{T}_L with fluxes respectively $\varphi, \xi_{L;k}^{\text{vert}}(\varphi)$ and associated height functions respectively h, \tilde{h} . Observe that, given two linearly independent vectors $w_0, w_1 \in \mathbb{R}^2$, any vector $w \in \mathbb{R}^2$ is entirely determined by $\langle w, w_0 \rangle$ and $\langle w, w_1 \rangle$. For vertices on L , h and \tilde{h} are given by inner product formulas with respectively φ and $\xi_{L;k}^{\text{vert}}(\varphi)$, so $\varphi - \xi_{L;k}^{\text{vert}}(\varphi)$ is entirely defined by the value $h - \tilde{h}$ takes on two linearly independent vertices of L . The idea of the proof is to show $h - \tilde{h}$ is constant on two linearly independent vertices of L (across all φ as in the statement), so $\varphi - \xi_{L;k}^{\text{vert}}(\varphi)$ is always the same, and nonzero if $h - \tilde{h}$ is nonzero on at least one of those vectors.

Recall the proof of Proposition 7.1.2; in what follows, we will use its notation and ideas. Once again, we will show only the case $k = 1$, but for other values of k the proof is analogous.

Because t, \tilde{t} have flux in $\mathcal{F}(L) \cap Q_1$, their height functions coincide on (x, x) for all $x \in \mathbb{Z}^2$; in other words, $h - \tilde{h}$ is always 0 on those vertices. Since by Lemma 5.2.6 there is a nonzero $v \in L$ with that form, we need only find a vertex in L not of that form on which $h - \tilde{h}$ is constant and nonzero.

Let $C(t)$ be t 's ξ_L -1-exclusive set and similarly for $C(\tilde{t})$. Choose \tilde{t} so that $C(t)$ and $C(\tilde{t})$ differ in only one element; it's clear this is always possible (as in the paragraph just before the statement of Proposition 7.1.3). Analyzing how h and \tilde{h} change along β_0 , it is clear they differ only along the edge e crossing the staircase S whose equivalence class is different in $C(t)$ and $C(\tilde{t})$. In t , $[S]_L \in \text{Stair}(L; 1; \text{vert})$, so e crosses a domino; in \tilde{t} , $[S]_L \in \text{Stair}(L; 1; \text{hor})$, so e does not cross a domino. Therefore, the coloring of \mathbb{Z}^2 implies h changes

by $+3$ along e while \tilde{h} changes by -1 along e . It follows that $h(u) - \tilde{h}(u)$ is always 4, and we are done. \square

Proposition 7.1.3 provides a visual way to interpret the counting of tilings of \mathbb{T}_L with flux in $\mathcal{F}(L) \cap \partial Q$. For each k , orient Q_k (as a line segment in \mathbb{R}^2) from $\text{Stair}(L; k; \text{vert})$ to $\text{Stair}(L; k; \text{hor})$. Let $\mathcal{F}(L) \cap Q_k = \{p_0, p_1, \dots, p_{c_k}\}$, where the order respects Q_k 's orientation; in particular, p_0 corresponds to $\text{Stair}(L; k; \text{vert})$ and p_{c_k} corresponds to $\text{Stair}(L; k; \text{hor})$. Notice $c_k = \text{card}(\text{Stair}(L; k; \text{vert}))$.

For each $0 \leq j < c_k$, Proposition 7.1.2 implies $\xi_{L;k}^{\text{vert}}(p_j) \in Q_k \setminus \{p_j\}$. In particular, if $p_i = \xi_{L;k}^{\text{vert}}(p_0)$, then $i > 0$. By Proposition 7.1.3, if we define $p_{i_j} = \xi_{L;k}^{\text{vert}}(p_j)$ it follows that $i_j > j$, for $p_{i_j} - p_j$ is constant. In particular, it must be $i_{c_k-1} = c_k$, so by induction we have $i_j = j + 1$. The following formula thus holds: $(\xi_{L;k}^{\text{vert}})^i(p_0) = p_i$.

This means that for all $0 \leq i \leq c_k$, a tiling of \mathbb{T}_L with flux p_i is obtained from the only tiling of \mathbb{T}_L with flux p_0 — the brick wall that corresponds $\text{Stair}(L; k; \text{vert})$ — by choosing i elements in its ξ_L - k -exclusive set that are also in $\text{Stair}(L; k; \text{vert})$ and applying ξ_L to them. In other words, $n_k^{\text{vert}}(p_i) = c_k - i$, so the number of tilings of \mathbb{T}_L with flux p_i is simply $\binom{c_k}{c_k - i} = \binom{c_k}{i}$.

Notice this means that regardless of L or $\mathcal{F}(L)$'s behaviour, if we know $\mathcal{F}(L) \cap Q_k$, then for each $\varphi \in \mathcal{F}(L) \cap Q_k$ we know the number of tilings of \mathbb{T}_L with flux φ .

Moreover, the effect of L -stairflips on a tilings with flux in $\mathcal{F}(L) \cap \partial Q$ can now be better understood: it navigates between adjacent fluxes. We explain it: let t be one such tiling and suppose its flux φ lies in the interior of Q_k . Then it consists entirely of type- k doubly-infinite domino staircases, some of which are made up of vertical dominoes and some of which are made of horizontal ones. Each extremal point of Q_k corresponds to a different tiling that consists entirely of type- k doubly-infinite staircases and uses dominoes of only one type (vertical or horizontal). Applying an L -stairflip to a staircase in t takes us to a tiling of \mathbb{T}_L whose flux is closest to φ in Q_k : if that stairflip is applied to a vertical staircase, the new flux is closest to the extremal point of horizontal staircases, and vice-versa. Now, if the flux φ lies in the boundary of Q_k , it is one of the brick walls. In this case, t can be seen as consisting entirely of type- k_0 domino staircases and entirely of type- k_1 domino staircases: it is in $\partial Q_{k_0} \cap \partial Q_{k_1}$. Applying an L -stairflip to a type- k_i staircase in t takes us to a tiling of \mathbb{T}_L whose flux is closest to φ in Q_{k_i} — and now there's only one such

flux! This also makes it clear how brick walls are as transition tilings between different Q_k 's.

We are now ready to prove Proposition 7.0.8.

Proof of Proposition 7.0.8. The discussion before the statement of Proposition 7.0.8 makes it clear we need only prove the case $\varphi \in \mathcal{F}(L) \cap \partial Q$. Furthermore, it suffices to show that, for each tiling t of \mathbb{T}_L with flux $\varphi \in \mathcal{F}(L) \cap \partial Q$, the sign in $K_t = \pm q_0^i q_1^j$ depends only on the flux (and not on t).

By Proposition 7.1.1, given $\varphi \in \mathcal{F}(L) \cap \partial Q$, there is some $k \in \{1, 2, 3, 4\}$ such that each tiling of \mathbb{T}_L with flux φ consists entirely of type- k doubly-infinite domino staircases. In particular, $\varphi \in Q_k$. Let $c_k = \text{Stair}(L; k; \text{vert})$. For each tiling of \mathbb{T}_L with flux in Q_k , let $C(t)$ be its ξ_L - k -exclusive set. By Proposition 7.1.2, each tiling t of \mathbb{T}_L with flux φ satisfies $\text{card}(C(t) \cap \text{Stair}(L; k; \text{vert})) = n_k^{\text{vert}}(\varphi)$, and $\text{card}(C(t) \cap \text{Stair}(L; k; \text{hor})) = c_k - n_k^{\text{vert}}(\varphi)$; in particular, neither depends on t .

Remember our universal Kasteleyn signing, based on $b_{\mathcal{N}}$. An equivalence class of edges on $G(\mathbb{Z}^2)$ is on $b_{\mathcal{N}}$, and thus is negatively signed, if and only if it lies in an element of $\text{Stair}(L; 1; \text{hor})$ or in an element of $\text{Stair}(L; 2; \text{hor})$. Now,

$$\begin{aligned} & \text{card}\left(C(t) \cap \left(\text{Stair}(L; 1; \text{hor}) \sqcup \text{Stair}(L; 2; \text{hor})\right)\right) = \\ & \text{card}\left(C(t) \cap \text{Stair}(L; 1; \text{hor})\right) + \text{card}\left(C(t) \cap \text{Stair}(L; 2; \text{hor})\right), \end{aligned}$$

and the preceding paragraph then implies that for each tiling t of \mathbb{T}_L with flux φ , the expression above depends only on φ .

Notice that the number of equivalence classes of edges on $G(\mathbb{Z}^2)$ in any two elements of $\text{Stair}(L; 1) \sqcup \text{Stair}(L; 3)$ is the same, and similarly for any two elements of $\text{Stair}(L; 2) \sqcup \text{Stair}(L; 4)$: the former is given by the smallest $x > 0$ for which $(x, x) \in L$, and the latter is given by the smallest $x > 0$ for which $(x, -x) \in L$. Together with the previous paragraph, this means that for each tiling of \mathbb{T}_L with flux φ , the number of equivalence classes of edges on $b_{\mathcal{N}}$ is the same. Moreover, this also shows that for each such tiling t the permutation associated to K_t is given by c_k cycles of equal length; in particular, the signs of these permutation are all the same.

It follows that the sign of K_t depends only on φ , as desired. \square

8

Sign distribution over $\mathcal{F}(L)$

We now know that each tiling of \mathbb{T}_L with flux $\varphi \in \mathcal{F}(L)$ is assigned the same sign in the combinatorial expansion of the Kasteleyn determinant, so we may speak of the sign of φ . Our next results will work to describe how these signs are distributed.

8.1

Cycles and cycle flips

Let t_0, t_1 be two tilings of \mathbb{T}_L . Represent both simultaneously on a fundamental domain D_L and orient dominoes of t_0 from black to white and dominoes of t_1 from white to black. With this orientation, D_L is decomposed into disjoint domino cycles whose dominoes belong alternatingly to t_0 and t_1 ; call the collection of these cycles $C(t_0, t_1)$. This set provides a way to go from t_0 to t_1 : for each cycle c in $C(t_0, t_1)$, simply replace each domino in c that is also in t_0 with the domino that follows it in c . We call this process a *cycle flip*.

The image below provides an example of this construction. Notice each cycle in $C(t_0, t_1)$ defines edge-paths that are in both t_0 and t_1 (the edge-paths that fit that cycle).




Figure 8.1: An example of cycle construction. The dominoes of t_0 are represented by the black, thicker edges while the dominoes of t_1 are represented by the thinner, blue edges. Marked vertices are in the lattice L , and the red rectangle is the fundamental domain D_L .

Of course, we may lift this representation to \mathbb{Z}^2 , decomposing it into disjoint, L -periodic domino paths whose dominoes belong alternatingly to t_0

and t_1 . Under this representation, each cycle is a collection of either finite, closed domino paths or infinite domino paths. We will refer to the former by *closed cycles* and to the latter by *open cycles*.




Figure 8.2: Lifting our previous example to \mathbb{Z}^2 . We can see $C(t_0, t_1)$ has two trivial cycles, one closed cycle, and two open cycles.

Remember dominoes may be seen as edges on $G(\mathbb{Z}^2)$, and vertices of $G(\mathbb{Z}^2)$ lie on $(\mathbb{Z} + \frac{1}{2})^2$. Thus, each domino path of a cycle can be seen as an edge-path in $(\mathbb{Z} + \frac{1}{2})^2$, which decomposes \mathbb{R}^2 into two disjoint, connected components. For open cycles these components are both unbounded, while for closed cycles one is unbounded and one is bounded. In the latter case, we call the unbounded component the exterior of the path, and the bounded component the interior of the path.

Define the interior $\text{int}(c)$ of a closed cycle c to be the union of the interior of domino paths in c , and the exterior $\text{ext}(c)$ of c to be the intersection of the exterior of domino paths in c . Notice $\text{int}(c)$ is never connected, and $\text{ext}(c)$ is always connected; moreover, $\text{ext}(c) \cap \mathbb{Z}^2$ is always connected by edge-paths. We provide an example of these constructions below.

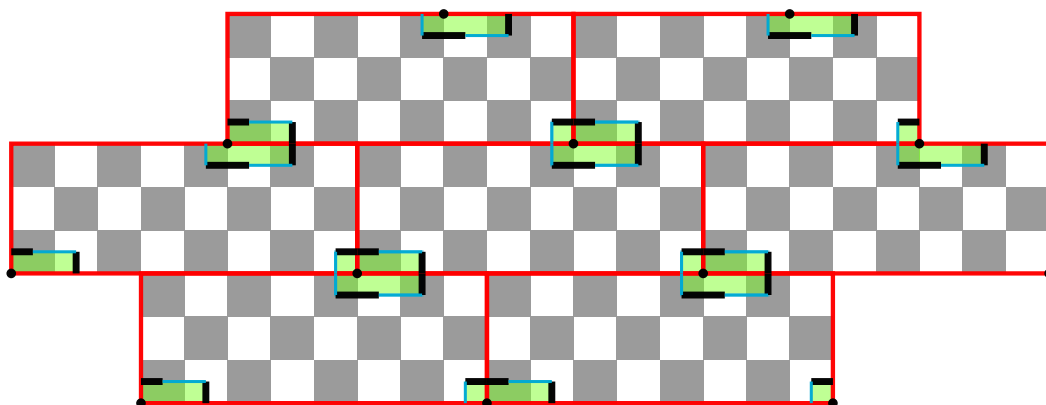


Figure 8.3: The closed cycle in our previous example. Its interior is tinted green.

Proposition 8.1.1. *Let L be a valid lattice and t_0, t_1 be two tilings of \mathbb{T}_L with fluxes respectively $\varphi_0, \varphi_1 \in \mathcal{F}(L)$. If t_0 and t_1 differ by a single closed cycle c , then $\varphi_0 = \varphi_1$.*

Proof. Let h_i be t_i 's associated height function. We will show h_0 and h_1 agree on L , from which the proposition follows. There are two cases:

1. $\text{int}(c)$ does not intersect L ;
2. $\text{int}(c)$ does intersect L .

In the first case, $L \subset \text{ext}(c)$, so any two points of L can be joined by an edge-path contained entirely in $\text{ext}(c)$. Since t_0 and t_1 coincide along these edge-paths, the height change along them is the same for both h_0 and h_1 , and they agree on L .

In the second case, because domino paths of c are L -periodic, each $v \in L$ must belong to the interior of a single domino path of c , and each domino path of c must contain a single $v \in L$. Let $v_a, v_b \in L$ and δ_a, δ_b be their respective domino paths of c . Let $u_a \in \mathbb{Z}^2$ be the first point to v_a 's right in the exterior of δ_a , and similarly for u_b . Let γ_a be the horizontal edge-path in \mathbb{Z}^2 joining v_a to u_a and γ_b be the horizontal edge-path in \mathbb{Z}^2 joining u_b to v_b . Finally, let γ_{ab} be any edge-path in $\text{ext}(c) \cap \mathbb{Z}^2$ joining u_a to u_b , so $\gamma = \gamma_a * \gamma_{ab} * \gamma_b$ is an edge-path in \mathbb{Z}^2 joining v_a to v_b . Notice γ has a single edge e_a that crosses δ_a , and a single edge e_b that crosses δ_b .

Since t_0 and t_1 agree except on c , h_0 and h_1 have the same change along γ except along e_a and e_b . Now, because of c 's L -periodicity, e_a and e_b have the same color-induced orientation, but are traversed by γ in opposite directions. Moreover, the L -periodicity also implies the domino e_a crosses in δ_a and the domino e_b crosses in δ_b belong to the same t_i . In other words, the changes of h_i along e_a and along e_b have equal magnitude but opposite signs, so they cancel each other out. This means the total change of h_0 and h_1 along γ is in fact the same. Since they agree on the origin (and the origin is in L), they must agree everywhere on L , and we are done. \square

Corollary 8.1.2. *Let L be a valid lattice and t_0, t_1 be two tilings of \mathbb{T}_L with fluxes respectively $\varphi_0, \varphi_1 \in \mathcal{F}(L)$. If each cycle in $C(t_0, t_1)$ is closed, then $\varphi_0 = \varphi_1$.*

8.2

The effect of a cycle flip on the sign of a flux

We know any two tilings of \mathbb{T}_L can be joined by cycle flips, but by Corollary 8.1.2 only flips on open cycles can affect the flux. Our attention now turns to studying how flips on open cycles affect the Kasteleyn sign of a tiling, and thus of their respective fluxes.

Let L be a valid lattice. We say $v \in L$ is *short* if $s \cdot v \notin L$ for all $s \in [0, 1)$.

Let t_0, t_1 be two tilings of \mathbb{T}_L . Let c be an open cycle in $C(t_0, t_1)$. Let γ_c be any infinite domino path of c . Since these paths are L -periodic, for each $v \in L$ $\gamma_c + v$ is a domino path of c . Moreover, there is a short $u \in L$ for which $\gamma_c = \gamma_c + u$, and it is clear u is unique up to multiplication by -1 . Now, if $\tilde{\gamma}_c$ is another infinite domino path of c , L -periodicity implies this unique vector is the same. We will then say $u = u(c)$ is c 's *parameter*.

Let \tilde{c} be any other open cycle in $C(t_0, t_1)$. We claim $u(c) = u(\tilde{c})$. Indeed, if there were some $k \in \mathbb{R} \setminus \{-1, 1\}$ with $u(c) = k \cdot u(\tilde{c})$, it would contradict the shortness of $u(c)$ or of $u(\tilde{c})$. On the other hand, if $u(c)$ and $u(\tilde{c})$ were linearly independent, c and \tilde{c} would intersect, contradicting their disjointness. It follows that whenever $C(t_0, t_1)$ contains an open cycle, it has a well-defined parameter $u = u(t_0, t_1)$, a short vector in L unique up to multiplication by -1 .

For any short $v \in L$, a v -*quasicycle* is a function $\gamma : \mathbb{Z} \rightarrow (\mathbb{Z} + \frac{1}{2})^2$ with $\|\gamma(t+1) - \gamma(t)\| = 1$ for all $t \in \mathbb{Z}$ and such that there is a positive integer T with $\gamma(t+T) = \gamma(t) + v$ for all $t \in \mathbb{Z}$. We say T is γ 's quasiperiod, and notice it is always even (because v is in a valid lattice). We say $\gamma : \mathbb{Z} \rightarrow (\mathbb{Z} + \frac{1}{2})^2$ is *simple* if it is injective; in other words, if it does not self-intersect in the plane.

Observe that any quasicycle γ can be interpreted as a domino-path in the infinite square lattice. With this in mind, let $C(t_0, t_1)$ contain an open cycle and v be its parameter. It's easy to see that for any open cycle $c \in C(t_0, t_1)$ and any infinite domino path γ_c of c , γ_c is a simple v -quasicycle.

For any v -quasicycle γ , define its sign by

$$\text{sgn}(\gamma) = (-1)^{\left(\frac{T}{2} + 1\right)} \cdot \prod_{0 \leq t < T} \text{sgn}([\gamma(t), \gamma(t+1)]),$$

where T is γ 's quasiperiod and $\text{sgn}(e)$ is the Kasteleyn sign of the edge e in $G(\mathbb{Z}^2) = (\mathbb{Z} + \frac{1}{2})^2$.

Notice that for any u in L and any edge e in $G(\mathbb{Z}^2)$, $\text{sgn}(e) = \text{sgn}(e + u)$. In particular, because $\gamma(t + T) = \gamma(t) + v$ and $v \in L$, we always have that

$$\text{sgn}([\gamma(t), \gamma(t+1)]) = \text{sgn}([\gamma(t), \gamma(t+1)] + v) = \text{sgn}([\gamma(t+T), \gamma(t+1+T)]).$$

This means the sign of a quasicycle γ obtained from any infinite domino path of an open cycle in $C(t_0, t_1)$ does not depend on a particular choice of parametrization (a choice of edge to be $\gamma(0)$). Moreover, any two infinite domino paths of one same open cycle in $C(t_0, t_1)$ are related by a translation in L , which preserves the quasiperiod and each Kasteleyn sign. We may thus define the sign of an open cycle $c \in C(t_0, t_1)$ to be the sign of any quasicycle obtained from any infinite domino path of c .

Let t_0, t_1 be tilings of \mathbb{T}_L with fluxes respectively φ_0 and φ_1 . Suppose $\varphi_0 \neq \varphi_1$. By Corollary 8.1.2, $C(t_0, t_1)$ contains an open cycle c . Let \tilde{t} be obtained from t_0 by a cycle flip on c , and let $\tilde{\varphi}$ be its flux. Notice $\varphi_0 \neq \tilde{\varphi}$, and by construction $\text{sgn}(c)$ is the sign change produced by the cycle flip on c , so

$$\frac{\text{sgn}(\varphi_0)}{\text{sgn}(\tilde{\varphi})} = \text{sgn}(c).$$

Indeed, $\frac{T}{2}$ is the length of the permutation cycle that represents the cycle flip, so its sign is $(-1)^{\binom{T}{2}+1}$. The product accounts for the sign changes from the Kasteleyn signing.

When γ is simple, it divides \mathbb{R}^2 into two unbounded connected components: $\mathbb{R}_+(\gamma)$ to the left of γ , and $\mathbb{R}_-(\gamma)$ to the right of γ . For any simple v -quasicycle γ , let γ_+ be the edge-path in $\mathbb{R}_+(\gamma) \cap \mathbb{Z}^2$ that fits it, and similarly for γ_- . The height changes along γ_+ and γ_- are well-defined. For any vertex w in γ_+ , there is an edge-path along γ_+ joining w to $w + v$; call it γ_+^w . We claim the height change along γ_+^w does not depend on w . Indeed, if w_1, w_2 are vertices in γ_+ , there is an edge-path $\gamma_+^{w_1, w_2}$ along γ_+ joining w_1 to w_2 . Because γ is a v -quasicycle, $\gamma_+^{w_1, w_2} + v$ is also an edge-path in γ_+ ; it joins $w_1 + v$ to $w_2 + v$. It's easy to see the height change along $\gamma_+^{w_1, w_2}$ and $\gamma_+^{w_1, w_2} + v$ is the same, from which the claim follows. Define then $h(\gamma_+)$ to be the height change along any γ_+^w . Of course, the same applies to γ_- , and $h(\gamma_-)$ is well-defined. We claim $h(\gamma_+) = h(\gamma_-)$.

Indeed, let u_+ and u_- be adjacent vertices with u_+ in γ_+ and u_- in γ_- . Let e be the edge joining u_+ to u_- . Of course, $u_+ + v$ and $u_- + v$ are also adjacent vertices in their respective edge-paths, and $e + v$ joins them. Consider

then the edge-path $\beta = \gamma_+^{u_+} * (e + v) * (\gamma_-^{u_-})^{-1} * e^{-1}$, where e is oriented from u_+ to u_- and \cdot^{-1} indicates a reversal of orientation.




Figure 8.4: The simple v -quasicycle γ in black, the marked vertices u_+ , $u_+ + v$, u_- and $u_- + v$, and the edge-paths $\gamma_+^{u_+}$ and $\gamma_-^{u_-}$ in red.

Since it is closed, the height change along β is 0, so $h(\gamma_+) + h(e + v) - h(\gamma_-) - h(e) = 0$, where $h(e)$ is the height change along e . The claim follows from noting that $h(e + v) = h(e)$, because $v \in L$.

We say a tiling of \mathbb{T}_L is *compatible with* γ if t contains every other domino in γ . In this case, it's clear γ_+ and γ_- are edge-paths in t , so by definition

$$h(\gamma_+) = 4 \cdot \langle \varphi_t, v \rangle = h(\gamma_-), \tag{8-1}$$

where φ_t is the flux of t .

Define then γ 's *pseudo-flux* by $\phi(\gamma) = \frac{1}{4}h(\gamma_+) = \frac{1}{4}h(\gamma_-)$, so $\phi(\gamma) = \langle \varphi_t, v \rangle$ whenever t is a tiling that's compatible with γ . Here, φ_t is t 's flux and v is γ 's parameter.

We are now ready to state Proposition 8.2.1.

Proposition 8.2.1. *Let L be a valid lattice and $v \in L$ be short. If γ_0, γ_1 are simple v -quasicycles, then*

$$\frac{\text{sgn}(\gamma_0)}{\text{sgn}(\gamma_1)} = (-1)^{\phi(\gamma_0) - \phi(\gamma_1)}.$$

For the proof of Proposition 8.2.1, we will need two lemmas.

For any simple quasicycle γ , define $V_+(\gamma) = \mathbb{R}_+(\gamma) \cap \mathbb{Z}^2$, the set of all vertices of \mathbb{Z}^2 that lie in $\mathbb{R}_+(\gamma)$, and similarly for $V_-(\gamma)$. Notice that when γ is a v -quasicycle, $V_+(\gamma)$ and $V_-(\gamma)$ are invariant under translation by v .

Now, suppose γ_0, γ_1 are both simple v -quasicycles¹. Consider the set $V(\gamma_0, \gamma_1)$ of vertices in \mathbb{Z}^2 that are to the left of one γ but to the right of the other; in other words, the set $V(\gamma_0, \gamma_1) = V_+(\gamma_0) \Delta V_+(\gamma_1) = V_-(\gamma_0) \Delta V_-(\gamma_1)$. Since the $V_{\pm}(\gamma_i)$ are invariant under translation by v , so too is $V(\gamma_0, \gamma_1)$; this means that for each $u \in V(\gamma_0, \gamma_1)$, there are infinitely many copies of u in $V(\gamma_0, \gamma_1)$. We may thus take the quotient $\tilde{V}(\gamma_0, \gamma_1) = V(\gamma_0, \gamma_1) / \langle v \rangle$.

Lemma 8.2.2. *Let L be a valid lattice and $v \in L$ be short. If γ_0, γ_1 are simple v -quasicycles and $\text{card}(\tilde{V}(\gamma_0, \gamma_1)) = 1$, then*

$$\frac{\text{sgn}(\gamma_0)}{\text{sgn}(\gamma_1)} = (-1)^{\varphi(\gamma_0) - \varphi(\gamma_1)}.$$

In other words, when $\text{card}(\tilde{V}(\gamma_0, \gamma_1)) = 1$, the statement of Proposition 8.2.1 holds.

Notice $\text{card}(\tilde{V}(\gamma_0, \gamma_1)) = 0$ if and only if $\gamma_0 = \gamma_1$.

Proof. Choose a vertex $w \in (\mathbb{Z} + \frac{1}{2})^2$ that belongs to both γ_0 and γ_1 . For each $i \in \{0, 1\}$, consider the segment γ_i^w of γ_i that joins w to $w + v$. They coincide except round the boundary of a square S with vertices in $(\mathbb{Z} + \frac{1}{2})^2$, which represents the equivalence class in $\tilde{V}(\gamma_0, \gamma_1)$. There are two cases:

1. Three of the edges of S belong to a γ_i^w , the other edge belongs to the other γ_j^w ;
2. Two of the edges of S belong to a γ_i^w , the other two edges belong to the other γ_j^w .



Figure 8.5: An example of each case. γ_i^w is represented by black edges and γ_j^w by dashed, red edges.

¹Note that the same v applies to both γ 's, and this also implies they are oriented the same way.

Recall that either one or three of the edges of S are negatively signed, so the product of Kasteleyn signs in each $\text{sgn}(\gamma_i)$ is always different. It follows that $\frac{\text{sgn}(\gamma_0)}{\text{sgn}(\gamma_1)}$ is defined entirely by T_0 and T_1 .

In case (1), suppose without loss of generality three of the edges of S belong to γ_0^w . Then in the obvious notation, the quasiperiods satisfy $T_0 = T_1 + 2$. This means the signs $(-1)^{\frac{1}{2}T_i+1}$ in each $\text{sgn}(\gamma_i)$ are also different, so $\text{sgn}(\gamma_0) = \text{sgn}(\gamma_1)$.

On the other hand, in this case there is a sign $x \in \{+, -\}$ such that $(\gamma_0)_x$ and $(\gamma_1)_x$ coincide. For instance, in our previous example $(\gamma_0)_- = (\gamma_1)_-$. It follows that $\phi(\gamma_0) = \frac{1}{4}h((\gamma_0)_x) = \frac{1}{4}h((\gamma_1)_x) = \phi(\gamma_1)$, so $(-1)^{\phi(\gamma_0)-\phi(\gamma_1)} = 1$ and the proof for this case is complete.

Suppose now we are in case (2). Clearly $T_0 = T_1$, implying the signs $(-1)^{\frac{1}{2}T_i+1}$ in each $\text{sgn}(\gamma_i)$ are the same, so $\text{sgn}(\gamma_0) = -\text{sgn}(\gamma_1)$.

On the other hand, in this case for each sign $x \in \{+, -\}$, $(\gamma_0)_x$ and $(\gamma_1)_x$ coincide except round the boundary of a square S_x and its translations by v . Along the segments on S_x , the height change for a $(\gamma_i)_x$ is $+2$ while the height change for the other $(\gamma_j)_x$ is -2 , so $\phi(\gamma_0) - \phi(\gamma_1) = \frac{1}{4} \cdot (h(\gamma_{0x}) - h(\gamma_{1x}))$ is either $+1$ or -1 . Regardless of the situation, $(-1)^{\phi(\gamma_0)-\phi(\gamma_1)} = -1$ as desired. \square

Let γ be an oriented edge-path, and suppose no two of its consecutive edges coincide (except for orientation). For such a path γ , we now describe the construction of an *argument function* defined over its edges.

We choose γ 's initial edge e_0 as a base edge, and assign the choice of base value 0 to it; in other words, $\arg(e_0) = 0$. For other edges, \arg is given recursively by $\arg(e_{j+1}) = \arg(e_j) + \alpha_j$, where $\alpha_j \in \mathbb{R}$ is the angle with smallest modulus such that rotation by α_j round e_j 's starting point results in an edge that is parallel and identically oriented to e_{j+1} . Notice the condition we imposed guarantees α_j is always well-defined (there is never a choice between π or $-\pi$), so \arg is too.

Lemma 8.2.3. *Let L be a valid lattice and $v \in L$ be short. Let γ_0, γ_1 be distinct simple v -quasicycles. There is a finite sequence of simple v -quasicycles $(\beta_k)_{k=0}^n$ with $\beta_0 = \gamma_0$, $\beta_n = \gamma_1$, and such that for all $0 \leq k < n$ it holds that $\text{card}(\tilde{V}(\beta_k, \beta_{k+1})) = 1$.*

Proof. Observe that γ 's quasiperiod is the length of any segment in γ joining a vertex to its translation by v . The reader may find it easier to follow the proof with this interpretation, and over the course of this proof, we will refer to γ 's quasiperiod as γ 's length.

We need only prove for γ_1 with minimal length.² Indeed, let γ_0 and γ_1 be any two simple v -quasicycles, and suppose $\tilde{\gamma}$ is a v -quasicycle with minimal length; clearly it is simple. If Lemma 8.2.3 holds for $\gamma_0, \tilde{\gamma}$ and for $\tilde{\gamma}, \gamma_1$, we may combine both sequences obtained this way, so Lemma 8.2.3 applies to γ_0 and γ_1 . We thus assume without loss of generality that γ_1 has minimal length.

Let γ be any simple v -quasicycle. Because it is simple, no two of its consecutive edges coincide, so we may define an argument function \arg_γ over its edges. Now, it's easy to see that the following are equivalent:

- γ has minimal length;
- \arg_γ assumes at most two values.

We divide the proof in two cases.

Case 1. γ_0 also has minimal length.

Observe that \arg_{γ_0} assumes a single value if and only if one of v 's coordinates is 0, that is, if and only if γ_0 and γ_1 are parallel straight lines. In this situation, it is obvious the lemma holds.

Suppose then that \arg_{γ_0} assumes two values, let $\beta_0 = \gamma_0$ and consider $\tilde{V}(\beta_0, \gamma_1)$. Make correspond to each vertex in \mathbb{Z}^2 the square in \mathbb{R}^2 with side length 1 centered on that vertex, so each edge of β_0 is the side of one such square. Choose a class $w \in \tilde{V}(\beta_0, \gamma_1)$ such that β_0 fits each of w 's corresponding squares on two of its sides. Notice $\tilde{V}(\beta_0, \gamma_1)$ is non-empty (since $\beta_0 = \gamma_0$ and γ_1 are distinct), and there's always one such w because \arg_{β_0} assumes two values.

Now, because \arg_{β_0} does not assume three or more values, all horizontal edges of β_0 have the same orientation and all of its vertical edges also do. This means no edge of β_0 touches any of the other two sides of each of w 's squares. We may thus consider the path β_1 that coincides with β_0 except on w 's squares; it fits each of these squares on the other two sides.

By construction, β_1 is a simple v -quasicycle, it preserves the minimal length, and its argument function also assumes two values. Moreover, it is clear $\tilde{V}(\beta_0, \beta_1) = \{w\}$, so $\text{card}(\tilde{V}(\beta_0, \beta_1)) = 1$. Finally, because w lies on different sides of β_0 and of β_1 , and also on different sides of β_0 and of γ_1 , w lies on the same side of β_1 and of γ_1 . Since $\text{card}(\tilde{V}(\beta_0, \gamma_1)) = \text{card}(\tilde{V}(\gamma_0, \gamma_1))$ is finite, this means $\text{card}(\tilde{V}(\beta_1, \gamma_1)) = \text{card}(\tilde{V}(\beta_0, \gamma_1)) - 1$.

If $\text{card}(\tilde{V}(\beta_1, \gamma_1)) > 0$ we may repeat the process, and in fact as long as $\text{card}(\tilde{V}(\beta_k, \gamma_1)) > 0$ we may do so. It's easy to verify that $\text{card}(\tilde{V}(\beta_k, \gamma_1)) =$

²The minimal length is $\|v\|_1$.

$\text{card}(\tilde{V}(\beta_0, \gamma_1)) - k$, so after a finite number $n = \text{card}(\tilde{V}(\beta_0, \gamma_1))$ of steps, we have $\text{card}(\tilde{V}(\beta_n, \gamma_1)) = 0$, that is, $\beta_n = \gamma_1$ as desired.

Case 2. γ_0 does not have minimal length.

If γ_0 does not have minimal length, arg_{γ_0} assumes at least three values. This means γ_0 contains at least one segment as in the figure below:




Figure 8.6: A segment with three consecutive argument values.

Each of these segments is a sequence of consecutive edges $(e_k)_{k=1}^n$ satisfying one of the following:

$$\begin{aligned} \text{arg}_{\gamma_0}(e_1) < \text{arg}_{\gamma_0}(e_2) = \dots = \text{arg}_{\gamma_0}(e_{n-1}) < \text{arg}_{\gamma_0}(e_n) \\ \text{arg}_{\gamma_0}(e_1) > \text{arg}_{\gamma_0}(e_2) = \dots = \text{arg}_{\gamma_0}(e_{n-1}) > \text{arg}_{\gamma_0}(e_n) \end{aligned}$$

For one such segment, we say its length is the number $n \geq 3$. Now, let δ be one such segment with minimal length. Because its length is minimal, γ_0 does not touch any of the ‘inner vertices’ near δ , indicated by a square in Figure 8.6. We may thus consider the simple v -quasicycles $\beta_1, \beta_2, \dots, \beta_{n-2}$ obtained from $\beta_0 = \gamma_0$ by changing the edges (and its translations by v) as shown in the following image:




Figure 8.7: The simple v -quasicycles β_k obtained from $\beta_0 = \gamma_0$.

It is clear $\text{card}(\tilde{V}(\beta_k, \beta_{k+1})) = 1$. Moreover, in the obvious notation the lengths satisfy $l(\beta_0) = l(\beta_1) = \dots = l(\beta_{n-3}) = l(\beta_{n-2}) + 2$, so the length has decreased. If $l(\beta_{n-2})$ is not minimal, then $\text{arg}_{\beta_{n-2}}$ assumes at least three values, and we may repeat the process. Since $l(\beta_0)$ is finite, this must end in a finite

number of steps, producing a simple v -quasicycle β_k with minimal length. We have thus reduced it to the previous case, and the proof is complete. \square

We are now ready to prove Proposition 8.2.1.

Proof of Proposition 8.2.1. By Lemma 8.2.2, the Proposition holds when $\text{card}(\tilde{V}(\gamma_0, \gamma_1)) \leq 1$. When $\text{card}(\tilde{V}(\gamma_0, \gamma_1)) > 1$, we may use Lemma 8.2.3 to obtain a sequence of simple v -quasicycles $(\beta_k)_{k=0}^n$ with $\beta_0 = \gamma_0$, $\beta_n = \gamma_1$, and such that for all $0 \leq k < n$ it holds that $\text{card}(\tilde{V}(\beta_k, \beta_{k+1})) = 1$. Then

$$\frac{\text{sgn}(\gamma_0)}{\text{sgn}(\gamma_1)} = \frac{\text{sgn}(\beta_0)}{\text{sgn}(\beta_n)} = \prod_{0 \leq k < n} \frac{\text{sgn}(\beta_k)}{\text{sgn}(\beta_{k+1})}$$

Applying Lemma 8.2.2 to each β_k, β_{k+1} yields

$$\begin{aligned} \frac{\text{sgn}(\gamma_0)}{\text{sgn}(\gamma_1)} &= \prod_{k=0}^n (-1)^{\phi(\beta_k) - \phi(\beta_{k+1})} = (-1)^{\sum_{k=0}^n \phi(\beta_k) - \phi(\beta_{k+1})} \\ &= (-1)^{\phi(\beta_0) - \phi(\beta_n)} = (-1)^{\phi(\gamma_0) - \phi(\gamma_1)}, \end{aligned}$$

so we are done. \square

The following corollary is automatic.

Corollary 8.2.4 (Sign formula for quasicycles). *Let $v \in L$ be short. There is a sign $C_v \in \{-1, +1\}$ such that for any v -quasicycle γ*

$$\text{sgn}(\gamma) = C_v \cdot (-1)^{\phi(\gamma)}.$$

8.3

The effect of a cycle flip on the flux itself

We know each open cycle in $C(t, t_0)$ has the same effect on $\text{sgn}(\varphi_t)$: as per Corollary 8.2.4, it is given by $C_v \cdot (-1)^{\langle \varphi_t, v \rangle}$, where v is $C(t, t_0)$'s parameter. But what about the effect of on the flux itself?

Proposition 8.3.1. *Let L be a valid lattice. Consider two different fluxes $\varphi_0, \varphi_1 \in \mathcal{F}(L)$ and let t_i be any tiling of \mathbb{T}_L with flux φ_i . Let v be $C(t_0, t_1)$'s parameter. Then $\varphi_1 - \varphi_0 \perp v$.*

Choose any open cycle $c \in C(t_0, t_1)$ and let t_c be obtained from t_0 by a cycle flip on c . Let φ_c be its flux. Then, in addition to $(\varphi_c - \varphi_0) \perp v$, $\varphi_c - \varphi_0$ is short in L^ . In particular, $\varphi_c - \varphi_0$ is uniquely defined up to sign.*

Proof. That $C(t_0, t_1)$ contains an open cycle (and thus its parameter is well-defined) is provided by Corollary 8.1.2.

For any infinite domino path of c , of course t_0 and t_c are both compatible with that path, so $\langle \varphi_0, v \rangle = \langle \varphi_c, v \rangle$, as in equation (8-1). This implies $\varphi_c - \varphi_0 \perp v$; notice this argument also applies to t_0 and t_1 . It remains to show that $\varphi_c - \varphi_0$ is short in L^* .

Consider the set $\Gamma(c) = \{\gamma \mid \gamma \text{ is an infinite domino path of } c\}$. Let $u \in L$ be such that $\{u, v\}$ is a basis for L ; the existence of such a vertex is guaranteed by Lemma 8.3.2 below. Choose any $\gamma \in \Gamma(c)$, and for each $k \in \mathbb{Z}$ let $\gamma_k = \gamma + k \cdot u$. Notice that because c has parameter v and $\{u, v\}$ is linearly independent, $\gamma_j = \gamma_k$ if and only if $j = k$. Clearly, each $\gamma_k \in \Gamma(c)$. We claim $\Gamma(c) = \{\gamma_k \mid k \in \mathbb{Z}\}$. Indeed, L acts transitively on $\Gamma(c)$ by translation, so for any $\tilde{\gamma} \in \Gamma(c)$ there is some $w = a \cdot u + b \cdot v \in L$ with $\tilde{\gamma} = \gamma + w$. It follows that $\tilde{\gamma} = \gamma_a + b \cdot v = \gamma_a$, because elements of $\Gamma(c)$ are invariant under translation by v (they are v -quasicycles).

Now, as edge-paths in $G(\mathbb{Z}^2) = (\mathbb{Z} + \frac{1}{2})^2$, two distinct γ_k 's are disjoint and 'parallel' with respect to v . Consider then $R = \mathbb{R}^2 \setminus \Gamma(c)$. Each connected component of R has boundary given by exactly two consecutive γ_k 's, and each γ_k is in the boundary of exactly two adjacent connected components of R . Moreover, it is clear $\mathbb{Z}^2 \subset R$.

Let R_0 be the connected component of R that contains the origin. Let $\partial R_0 = \gamma_j \sqcup \gamma_{j+1}$. Relabel the paths $\gamma_k := \gamma_{k-j}$, so the boundary of R_0 is given by $\gamma_0 \sqcup \gamma_1$. More generally, let R_k be the connected component of R whose boundary is given by $\gamma_k \sqcup \gamma_{k+1}$; clearly, R_k and R_{k+1} are adjacent.

We now compare the height functions h_0 of t_0 and h_c of t_c on each $R_k \cap \mathbb{Z}^2$. Remember t_0 and t_c differ only by the γ_k 's, so they coincide on each R_k . Now, because h_0 and h_c are both 0 at the origin, they coincide everywhere on R_0 ; in particular, $h_c - h_0$ is always 0 on $\langle v \rangle$.

We now compute $h_c - h_0$ on $u \in L \cap R_1$. Choose any edge-path β in \mathbb{Z}^2 joining the origin to u and that crosses $\gamma_1 = \partial R_0 \cap \partial R_1$ only once. It is clear such a path exists, because each $R_k \cap \mathbb{Z}^2$ is connected by edge-paths in \mathbb{Z}^2 . The height change along β for each of h_0 and h_c is the same except on the edge that crosses γ_1 . Either for one that change $+3$ and for the other that change is -1 , or these changes are -3 and $+1$, depending on whether the domino of γ_1 that's crossed by β lies in t_0 or in t_c , and on the orientation (as induced by the coloring of \mathbb{Z}^2) of the crossing edge on β . This means $h_c - h_0$ is either $+4$ or -4 on u (and thus on all of $R_1 \cap \mathbb{Z}^2$). By the same token, it's easy to see

that when it is $+4$, $h_c - h_0$ is $4k$ on all of $R_k \cap \mathbb{Z}^2$; when it is -4 , $h_c - h_0$ is $-4k$ on all of $R_k \cap \mathbb{Z}^2$.

Recall that evaluating a height function on L yields information about the corresponding flux via the inner product identification. In particular, we now know $\langle \varphi_c - \varphi_0, v \rangle = 0$ and $\langle \varphi_c - \varphi_0, u \rangle = \pm 1$. This entirely defines $\varphi_c - \varphi_0 \in L^*$. Moreover, for any $\varphi \in L^*$ we have $\langle \varphi, u \rangle \in \mathbb{Z}$, so $\varphi_c - \varphi_0$ must be short, completing the proof. \square

The following lemma was used in the proof of Proposition 8.3.1 above.

Lemma 8.3.2. *Let L be a lattice generated by linearly independent vectors $v_0, v_1 \in \mathbb{Z}^2$. Let $v \in L$ be short. Then there is some $u \in L$ such that $\{u, v\}$ is a basis for L .*

Proof. Observe that $v = a \cdot v_0 + b \cdot v_1 \in L$ is short if and only if $\gcd(a, b) = 1$. In this case, there are integers $k_a, k_b \in \mathbb{Z}$ with $k_a \cdot a + k_b \cdot b = 1$. Let $u \in L$ be the vector $-k_b \cdot v_0 + k_a \cdot v_1$. It's easy to check that $k_a \cdot v - b \cdot u = v_0$ and $k_b \cdot v + a \cdot u = v_1$, so $\{u, v\}$ generates L and is thus a basis for it. \square

We now show there is a set of two ‘short moves’ that connects $\mathcal{F}(L)$. We will use it to describe a ‘sign pattern’ for fluxes in $\mathcal{F}(L)$ via Proposition 8.3.1.

Recall that for any two fluxes in $\mathcal{F}(L)$, their difference is an element of L^* (indeed, as per Proposition 5.1.1, $L^\#$ is a translation of L^*). We say a basis $\{v_0^*, v_1^*\}$ for L^* is *flux-connecting* if given any two fluxes $\varphi, \tilde{\varphi} \in \mathcal{F}(L)$ there is a sequence of fluxes $(\varphi_k)_{k=0}^n$ with $\varphi_k \in \mathcal{F}(L)$ for all $0 \leq k \leq n$, $\varphi_0 = \varphi$, $\varphi_n = \tilde{\varphi}$ and such that $\varphi_{k+1} - \varphi_k = \pm v_i^*$ for all $0 \leq k < n$. In other words, the moves $\pm v_0^*$ and $\pm v_1^*$ connect $\mathcal{F}(L)$.

Lemma 8.3.3. *For any valid lattice L , L^* admits a flux-connecting basis.*

Proof. Consider $Q_2 \subset Q$, the side of ∂Q contained in the second quadrant. We know $b_{\mathcal{N}} = (0, \frac{1}{2})$ and $b_{\mathcal{W}} = (-\frac{1}{2}, 0)$ are in $\mathcal{F}(L) \cap Q_2$, so every flux in $\mathcal{F}(L)$ belongs to a unique line that is parallel to Q_2 . We say $l_0 \supset Q_2$ is the first such line, and l_{k+1} is the line just below l_k . See Figure 8.8.

Consider the brick wall $b_{\mathcal{W}}$, and let f_0 be the flux in l_0 that is closest to it (but different from it). Consider the line l_1 , and let f_1 be the flux in it that is closest to $b_{\mathcal{W}}$. Let $v_0^* = f_0 - f_1$ and $v_1^* = b_{\mathcal{W}} - f_1$. Figure 8.9 illustrates this construction. We claim $\{v_0^*, v_1^*\}$ is a flux-connecting basis of L^* .

First, we show the moves $\pm v_i^*$ connect any two fluxes in one same line l_k . Let $\varphi, \tilde{\varphi} \in \mathcal{F}(L) \cap l_k$ and suppose without loss of generality $\tilde{\varphi}$ is to the right of




Figure 8.8: Enumerating the lines l_k . Marked vertices are elements of $\mathcal{F}(L)$.




Figure 8.9: The vectors v_0^* and v_1^* form a flux-connecting basis of L^* .

φ . There must be a line l_{k-1} above l_k or a line l_{k+1} below it (possibly both, but at least one). In the first case, v_0^* takes $\varphi_0 = \varphi$ to a flux φ_1 in l_{k-1} ; and $-v_1^*$ takes φ_1 to a flux φ_2 back in l_k . Notice φ_2 is to the right of φ , and because of how v_0^*, v_1^* were chosen, there can be no flux in l_k between φ and φ_2 . If $\varphi_2 \neq \tilde{\varphi}$, we may repeat the process, and because $l_k \cap \mathcal{F}(L)$ is finite, it must end after a finite number of steps.

The latter case is similar: $-v_1^*$ takes $\varphi_0 = \varphi$ to a flux φ_1 in l_{k+1} ; and v_0^* takes φ_1 to a flux φ_2 back in l_k . Once again, induction shows φ and $\tilde{\varphi}$ are connected by the moves $\pm v_i^*$.

To complete the proof, we show the moves $\pm v_i^*$ connect any two adjacent lines. Consider the lines l_k, l_{k+1} . Let φ_k be the first flux in l_k , that is, the flux in l_k that is closest to Q_3 ; and similarly for φ_{k+1} . When φ_k is to the right of

φ_{k+1} , v_0^* takes φ_k to φ_{k+1} ; when φ_k is to the left of φ_{k+1} , v_1^* takes φ_{k+1} to φ_k . The image below exemplifies the two cases.




Figure 8.10: Each v_i^* connects adjacent lines in each case.

Regardless of the situation, the moves $\pm v_i^*$ connect two adjacent lines. Similar line arguments also show $\{v_0^*, v_1^*\}$ generates L^* , so we are done. \square

We are now ready to describe the aforementioned sign pattern.

Proposition 8.3.4 (Sign patterns in $\mathcal{F}(L)$). *Let L be a valid lattice and $\{v_0^*, v_1^*\}$ be a flux-connecting basis of L^* . Decompose $\mathcal{F}(L)$ in lines parallel to v_0^* and in lines parallel to v_1^* . Then, along any given one of those lines, the sign change between two adjacent fluxes is always the same. Moreover, for parallel and adjacent lines, the sign change along each line is different.*

Proof. Let $i, j \in \{0, 1\}$ be distinct. We will show that if $\varphi \in \mathcal{F}(L)$ is such that $\varphi + v_i^*$ and $\varphi - v_i^*$ are in $\mathcal{F}(L)$, then

$$\frac{\text{sgn}(\varphi)}{\text{sgn}(\varphi + v_i^*)} = \frac{\text{sgn}(\varphi - v_i^*)}{\text{sgn}(\varphi)}. \tag{8-2}$$

This proves the claim on each given line. Additionally, we will show that if $\varphi + v_j^*$ and $\varphi + v_i^* + v_j^*$ are in $\mathcal{F}(L)$, then

$$\frac{\text{sgn}(\varphi)}{\text{sgn}(\varphi + v_i^*)} = -\frac{\text{sgn}(\varphi + v_j^*)}{\text{sgn}(\varphi + v_j^* + v_i^*)}. \tag{8-3}$$

This proves the claim on parallel and adjacent lines.

Let t be a tiling of \mathbb{T}_L with flux φ , and similarly for t^\pm with $\varphi \pm v_i^*$, for t_j with $\varphi + v_j^*$ and for t_j^+ with $\varphi + v_j^* + v_i^*$. Suppose all relevant vectors are in $\mathcal{F}(L)$. Let v^+ and v^- be respectively the parameters of $C(t, t^+)$ and $C(t^-, t)$. Observe that $(\varphi + v_i^*) - \varphi = \varphi - (\varphi - v_i^*) = v_i^*$, so by Proposition 8.3.1 $v_i^* \perp v^+, v^-$. Since v^+ and v^- are both short, this implies $v^+ = v^-$ or $v^+ = -v^-$; in other words, they're equal up to multiplication by -1 .

Inspecting the proof of Proposition 8.3.1, we see that any cycle flip on a cycle with parameter u changes the flux by a vector in L^* that is short and

perpendicular to u . In particular, for any open cycle c in $C(t, t^+) \cup C(t^-, t)$, a cycle flip on c changes the flux by either v_i^* or $-v_i^*$ (recall that v_i^* is in a basis for L^* , so it must be short). This implies the number of open cycles in $C(t, t^+)$ is odd. Indeed, the total change (from performing a cycle flip on each open cycle) is v_i^* , so the number n^- of cycles with a $-v_i^*$ change and the number n^+ of cycles with a v_i^* change must satisfy $n^+ = n^- + 1$. Of course, the same holds for $C(t^-, t)$.

Now, any open cycle in $C(t, t^+) \cup C(t^-, t)$ is compatible with t , so by Corollary 8.2.4 there is a sign C_{v^+} such that for each open cycle $c \in C(t, t^+) \cup C(t^-, t)$

$$\text{sgn}(c) = C_{v^+} \cdot (-1)^{\langle \varphi, v^+ \rangle}.$$

Since we've shown each of $C(t, t^+)$ and $C(t^-, t)$ has an odd number of open cycles, the total sign change in each case is precisely $C_{v^+} \cdot (-1)^{\langle \varphi, v^+ \rangle}$. In particular, in each case the total sign change is the same, so we've proved equation (8-2).

For equation (8-3), observe that $(\varphi + v_j^* + v_i^*) - (\varphi + v_j^*) = v_i^*$, so once again the parameter w^+ of $C(t_j, t_j^+, t_i^+)$ is perpendicular to v_i^* and thus equal to v^+ up to multiplication by -1 (because they're both short). The same argument used above also shows $C(t_j, t_j^+)$ has an odd number of open cycles, and since each such cycle is compatible with t_j , its total sign change is given by

$$C_{v^+} \cdot (-1)^{\langle \varphi + v_j^*, v^+ \rangle} = C_{v^+} \cdot (-1)^{\langle \varphi, v^+ \rangle} \cdot (-1)^{\langle v_j^*, v^+ \rangle}.$$

We will show $\langle v_j^*, v^+ \rangle$ is either $+1$ or -1 , from which equation (8-3) follows. To that end, notice v_j^* is a short element of L^* (because it is in a basis). Since v^+ is short in L , the only possible integer values for $\langle v_j^*, v^+ \rangle$ that respect v_j^* 's shortness are -1 , 0 , and $+1$. Now, $\langle v_i^*, v^+ \rangle = 0$ (because $v_i^* \perp v^+$), so if $\langle v_j^*, v^+ \rangle$ were also 0 it would contradict $\{v_i^*, v_j^*\}$ being a basis for L^* , for all its elements would be 0 on the sublattice generated by v^+ . It follows that $\langle v_j^*, v^+ \rangle$ is either $+1$ or -1 , as desired. \square

Notice in the proof above that since v^+ is unique up to multiplication by -1 , we may choose it with $\langle v_j^*, v^+ \rangle = 1$. In particular, since $\langle v_i^*, v^+ \rangle = 0$, we have $\langle v_k^*, v^+ \rangle = \delta_{ik}$. This implies the following: let $\{v_0, v_1\}$ be the basis of L that is dual to the flux-connecting basis of L^* $\{v_0^*, v_1^*\}$, that is, they satisfy $\langle v_k^*, v_l \rangle = \delta_{kl}$. Let t, t_0 and t_1 be tilings of \mathbb{T}_L with fluxes respectively $\varphi, \varphi + v_0^*$ and $\varphi + v_1^*$. Then v_1 is $C(t, t_0)$'s parameter and v_0 is $C(t, t_1)$'s parameter.

The sign pattern in Proposition 8.3.4 may also be described as a pattern in which one odd-one-out sign is always surrounded by different signs.




Figure 8.11: In this example signing, the minus sign is the odd-one-out.

Of course, which precise sign (+1 or -1) is the odd-one-out depends on the choice of Kasteleyn signing for edges of $G(\mathbb{Z}^2)$ and the enumeration of D_L 's squares, but the pattern is always the same.

Proposition 8.3.4 is also instrumental in showing we can always obtain the total number of tilings of \mathbb{T}_L with a linear combination of $\det(K(\pm 1, \pm 1))$.

Proposition 8.3.5. *Let L be a valid lattice and K a Kasteleyn matrix for \mathbb{T}_L . Let $p_K(q_0, q_1)$ be the Laurent polynomial given by $\det(K)$. Then there is a choice of constants $s_{00}, s_{01}, s_{10}, s_{11} = \pm \frac{1}{2}$ such that the total number of tilings of \mathbb{T}_L is*

$$s_{00} \cdot p_K(1, 1) + s_{01} \cdot p_K(1, -1) + s_{10} \cdot p_K(-1, 1) + s_{11} \cdot p_K(-1, -1).$$

Proof. For the first part, we will show that for each monomial of the form $c_{ij} \cdot q_0^i q_1^j$ in the full expansion of $\det(K)$, the sign of c_{ij} depends only on the parity of i, j . To that end, observe that by Proposition 8.3.4 whenever $\varphi_0, \varphi_1 \in \mathcal{F}(L)$ are such that $\varphi_0 - \varphi_1 \in 2L^*$, then $\text{sgn}(\varphi_0) = \text{sgn}(\varphi_1)$.

Recall our construction of the Kasteleyn matrix, using the rectangular fundamental domain D_L defined by $v_0 = (x_0, 0)$ and $v_1 = (x_1, y_1)$. Here, v_0 and v_1 generate L , $x_0, y_1 > 0$ and $0 \leq x_1 < x_0$. We know that by construction

$$\forall \varphi \in \mathcal{F}(L), \exists i_0, i_1 \in \mathbb{Z}, \forall \text{ tiling } t \text{ of } \mathbb{T}_L \text{ with flux } \varphi, K_t = \pm q_0^{i_0} q_1^{i_1},$$

so we may speak of the q -exponents (i_0, i_1) of φ . Also by construction, we know that if y_1 is even and φ has q -exponents (i_0, i_1) , then $\langle \varphi, v_0 \rangle = i_0$ and $\langle \varphi, v_1 \rangle = i_1$. When y_1 is odd, the latter changes to $\langle \varphi, v_1 \rangle = i_1 + \frac{1}{2}$.

Let $q_0^{i_0} q_1^{i_1}$ be a monomial in p_K and consider any monomial $q_0^{j_0} q_1^{j_1}$ in p_K such that $i_0 \equiv j_0 \pmod{2}$ and $i_1 \equiv j_1 \pmod{2}$. Let $\varphi(i_0, i_1) \in \mathcal{F}(L)$ correspond to $q_0^{i_0} q_1^{i_1}$ and similarly for $\varphi(j_0, j_1)$. Consider the difference $\varphi(i_0, i_1) - \varphi(j_0, j_1) \in L^*$. Regardless of the parity of y_1 , it satisfies

$$\begin{aligned} \langle \varphi(i_0, i_1) - \varphi(j_0, j_1), v_0 \rangle &= i_0 - j_0 \\ \langle \varphi(i_0, i_1) - \varphi(j_0, j_1), v_1 \rangle &= i_1 - j_1 \end{aligned}$$

Because i_k and j_k have the same parity, these numbers are both even integers, so $\varphi(i_0, i_1) - \varphi(j_0, j_1) \in 2L^*$. It follows that $\varphi(i_0, i_1)$ and $\varphi(j_0, j_1)$ have the same sign, that is, the coefficients of $q_0^{i_0} q_1^{i_1}$ and $q_0^{j_0} q_1^{j_1}$ in p_K have the same sign, so the first part is complete.

Now, for the second part, write

$$\begin{aligned} & s_{00} \cdot p_K(1, 1) + s_{01} \cdot p_K(1, -1) + s_{10} \cdot p_K(-1, 1) + s_{11} \cdot p_K(-1, -1) \\ &= \sum_{i,j} |c_{ij}| \cdot \text{sgn}(c_{ij}) \cdot (s_{00} + (-1)^j \cdot s_{01} + (-1)^i \cdot s_{10} + (-1)^{i+j} \cdot s_{11}), \end{aligned}$$

where from the first part we know $\text{sgn}(c_{ij})$ depends only on the parity of i, j . We thus need to check the system below admits a solution:

$$\begin{cases} s_{00} + s_{01} + s_{10} + s_{11} = \text{sign of } i, j \text{ even} \\ s_{00} - s_{01} + s_{10} - s_{11} = \text{sign of } i \text{ even, } j \text{ odd} \\ s_{00} + s_{01} - s_{10} - s_{11} = \text{sign of } i \text{ odd, } j \text{ even} \\ s_{00} - s_{01} - s_{10} + s_{11} = \text{sign of } i, j \text{ odd} \end{cases}$$

Notice Proposition 8.3.4 ensures exactly three of the signs are equal. The claim follows from observing that the matrix

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}$$

is invertible³, and that applying A^{-1} to any of the eight possible sign configurations yields s_{ij} 's as in the statement. \square

³We have $\det(A) = 16$.

9

Kasteleyn determinants for the torus

In Section 3.3, we calculated the Kasteleyn determinant for rectangles with integral sides by finding a suitable basis of eigenvectors, from which we derived the eigenvalues of the Kasteleyn matrix itself. We will now employ similar techniques for calculating Kasteleyn determinants of tori.

Let L be a valid lattice and K be a Kasteleyn matrix for \mathbb{T}_L . For all $i, j \in \mathbb{Z}$, let $v_{i,j} = (i, j) + (\frac{1}{2}, \frac{1}{2})$, so the $v_{i,j}$ enumerate the vertices of $G(\mathbb{Z}^2)$. $v_{i,j}$ is black whenever $i \equiv j \pmod{2}$, and it is white otherwise. As we defined it, K is a linear map from the space of equivalence classes of black vertices in $G(\mathbb{Z}^2)$ to the space of white ones, and it acts as follows:

$$\begin{aligned} K v_{i,j} &= \text{Kast}(v_{i,j} v_{i,j+1}) \cdot v_{i,j+1} + \text{Kast}(v_{i,j} v_{i,j-1}) \cdot v_{i,j-1} \\ &\quad + \text{Kast}(v_{i,j} v_{i+1,j}) \cdot v_{i+1,j} + \text{Kast}(v_{i,j} v_{i-1,j}) \cdot v_{i-1,j} \end{aligned}$$

Here, $\text{Kast}(e)$ is the Kasteleyn weighting of the edge e , and if p_0, p_1 are two adjacent vertices, $p_0 p_1$ indicates the edge joining them. Observe that $\text{Kast}(e)$ does not depend on a particular choice of representative of $[e]_L$, so our dropping the braces round each $v_{i,j}$ is justified.

9.1

The case of $M = KK^* \oplus K^*K$

Like before, we will consider the matrix $M = KK^* \oplus K^*K$ rather than K . Notice K^* goes from the space of equivalence classes of white vertices in $G(\mathbb{Z}^2)$ to the space of black ones, via:

$$\begin{aligned} K^* v_{i,j} &= \overline{\text{Kast}(v_{i,j} v_{i,j+1})} \cdot v_{i,j+1} + \overline{\text{Kast}(v_{i,j} v_{i,j-1})} \cdot v_{i,j-1} \\ &\quad + \overline{\text{Kast}(v_{i,j} v_{i+1,j})} \cdot v_{i+1,j} + \overline{\text{Kast}(v_{i,j} v_{i-1,j})} \cdot v_{i-1,j} \end{aligned}$$

Here, $\overline{\cdot}$ is the complex conjugation.

The entries of K are Laurent monomials in q_0 and q_1 (their exponents indicate flux value), so we may think of K as a matrix $K(q_0, q_1)$. Of course,

this also means $\det(K)$ is a Laurent polynomial p_K in q_0 and q_1 . Rather than tackle the problem of finding a formula for $p_K(q_0, q_1) = \det(K)$, we will find the values of p_K when $q_0, q_1 \in \mathbb{S}^1 \subset \mathbb{C}$. Observe this is enough to describe p_K .

Notice $\bar{z} = z^{-1}$ whenever $z \in \mathbb{S}^1$, so for $q_0, q_1 \in \mathbb{S}^1$ we have $K_{ij}^* = (K_{ji})^{-1}$. In other words, for $q_0, q_1 \in \mathbb{S}^1$ we have that

$$K^*v_{i,j} = \text{Kast}(v_{i,j}v_{i,j+1})^{-1} \cdot v_{i,j+1} + \text{Kast}(v_{i,j}v_{i,j-1})^{-1} \cdot v_{i,j-1} \\ + \text{Kast}(v_{i,j}v_{i+1,j})^{-1} \cdot v_{i+1,j} + \text{Kast}(v_{i,j}v_{i-1,j})^{-1} \cdot v_{i-1,j}$$

Recall the rectangular fundamental domain $D_L \subset \mathbb{R}^2$ we used in the construction of K ; its vertices are $(0, 0)$, $(x_0, 0)$, $(0, y_1)$ and (x_0, y_1) , where $v_0 = (x_0, 0)$ and $v_1 = (x_1, y_1)$ generate L (and $0 \leq x_1 < x_0$). Because translations of D_L by v_0 and v_1 tile \mathbb{R}^2 , we may use it to partition $(\mathbb{Z} + \frac{1}{2})^2$, the set of vertices on $G(\mathbb{Z}^2)$. For $a, b \in \mathbb{Z}$, consider the sets $D(a, b) = (D_L + a \cdot v_0 + b \cdot v_1) \cap (\mathbb{Z} + \frac{1}{2})^2$. It's easy to see that $P = \{D(a, b) \mid a, b \in \mathbb{Z}\}$ is a partition of $(\mathbb{Z} + \frac{1}{2})^2$.

We now study transitions between adjacent $D(a, b)$'s via edge-paths in $(\mathbb{Z} + \frac{1}{2})^2$. The diagram below represents this behavior schematically; remember the orientation induced by the coloring of \mathbb{Z}^2 on edges of its unit squares when determining the Kasteleyn weights of crossing dominoes.




Figure 9.1: Schematic representation of q -weights for crossing edges. Red rectangles are fundamental domains, square vertices are elements of L , and round black or white vertices are elements of $G(\mathbb{Z}^2)$.

More precisely, we mean that:

- Edges joining a black vertex in $D(a, b)$ (respectively white) to a white vertex in $D(a, b + 1)$ (respectively black) have Kasteleyn weight $\pm q_0^{-1}$ (respectively $\pm q_0$);
- Edges joining a black vertex in $D(a, b)$ (respectively white) to a white vertex in $D(a + 1, b)$ (respectively black) have Kasteleyn weight $\pm q_1$ (respectively $\pm q_1^{-1}$);

- Edges joining a black vertex in $D(a, b)$ (respectively white) to a white vertex in $D(a - 1, b + 1)$ (respectively black) have Kasteleyn weight $\pm q_0^{-1} q_1^{-1}$ (respectively $\pm q_0 q_1$).

Remember our choice of positive orientation for dominoes: from their black square to their white square. When we assigned Kasteleyn weights to dominoes it was done irrespective of the domino's own orientation, but K naturally maps black vertices to white vertices, so we may think of it as the weight being assigned to dominoes with positive orientation. For dominoes with negative orientation, we assign it the inverse of its weight with positive orientation. Thus, each oriented domino has an *oriented weight*: its Kasteleyn weight for dominoes oriented positively, and the inverse of its Kasteleyn weight for dominoes oriented negatively.

For adjacent vertices $p_0, p_1 \in (\mathbb{Z} + \frac{1}{2})^2$, let $p_0 p_1$ the edge that joins them and is oriented from p_0 to p_1 . For an oriented edge e , let $o(e)$ denote its oriented weight. We thus have that:

- $\text{Kast}(p_0 p_1) = \text{Kast}(p_1 p_0)$, that is, Kasteleyn weight does not depend on orientation;
- $o(p_0 p_1) = o(p_1 p_0)^{-1}$, that is, reversing edge orientation inverts oriented weight;
- If p_0 is black, then $o(p_0 p_1) = \text{Kast}(p_0 p_1)$.

Furthermore, our previous observation may be simplified. Let u, v be adjacent vertices with $u \in D(a, b)$.

- If $v \in D(a, b + 1)$, then $o(uv) = \pm q_0^{-1}$;
- If $v \in D(a + 1, b)$, then $o(uv) = \pm q_1$;
- If $v \in D(a - 1, b + 1)$, then $o(uv) = \pm q_0^{-1} q_1^{-1}$.

Let b_i be D_L 's i -th black vertex and w_j be its j -th white vertex, as we enumerated them. With these conventions, $K_{ij} = o(b_i w_j)$. Moreover, observe that when $q_0, q_1 \in \mathbb{S}^1$, $K_{ij}^* = (K_{ji})^{-1}$, so $K_{ij}^* = o(b_j w_i)^{-1} = o(w_i b_j)$. In this case, the actions of K and K^* can be described by essentially the same formula, below. The '=' symbol draws attention to the fact that each of K, K^* acts on vertices of different colors, so for any given $v_{i,j}$ only one of $K(v_{i,j}), K^*(v_{i,j})$ actually makes sense.

$$\begin{aligned}
 K(v_{i,j}) \text{ '=' } K^*(v_{i,j}) \text{ '=' } & o(v_{i,j} v_{i,j+1}) \cdot v_{i,j+1} + o(v_{i,j} v_{i,j-1}) \cdot v_{i,j-1} \\
 & + o(v_{i,j} v_{i+1,j}) \cdot v_{i+1,j} + o(v_{i,j} v_{i-1,j}) \cdot v_{i-1,j}
 \end{aligned} \tag{9-1}$$

For $a, b \in \mathbb{Z}$, let $\mathcal{Q}(a, b) = q_0^{-b} q_1^a$. For $v \in (\mathbb{Z} + \frac{1}{2})^2$, let $\mathcal{Q}(v) = \mathcal{Q}(a_v, b_v)$, where $v \in D(a_v, b_v)$. Recall that because $(\mathbb{Z} + \frac{1}{2})^2$ is partitioned by the $D(a, b)$'s, $\mathcal{Q}(v)$ is always well-defined.

Proposition 9.1.1. *Let $(p_k)_{k=0}^n$ be a sequence of adjacent vertices in $(\mathbb{Z} + \frac{1}{2})^2$.*

Then

$$\prod_{k=0}^{n-1} |o(p_k p_{k+1})| = \frac{\mathcal{Q}(p_n)}{\mathcal{Q}(p_0)}.$$

In particular, the product depends only on the initial and final $D(a, b)$'s.

Proof. Notice that whenever u and v are adjacent vertices in the same $D(a, b)$, then $|o(uv)| = 1$. Combined with our previous (simplified) observation, it's easy to check that for any two adjacent vertices $u, v \in (\mathbb{Z} + \frac{1}{2})^2$ we have

$$|o(uv)| = \frac{\mathcal{Q}(v)}{\mathcal{Q}(u)}.$$

From this, the proposition follows. \square

With Proposition 9.1.1 and formula (9-1), we can make a good description of the action of $M = KK^* \oplus K^*K$. Notice (9-2) applies irrespective of $v_{i,j}$'s color.

$$\begin{aligned} Mv_{i,j} = & 4 \cdot v_{i,j} - \frac{\mathcal{Q}(v_{i+2,j})}{\mathcal{Q}(v_{i,j})} \cdot v_{i+2,j} - \frac{\mathcal{Q}(v_{i-2,j})}{\mathcal{Q}(v_{i,j})} \cdot v_{i-2,j} \\ & + \frac{\mathcal{Q}(v_{i,j+2})}{\mathcal{Q}(v_{i,j})} \cdot v_{i,j+2} + \frac{\mathcal{Q}(v_{i,j-2})}{\mathcal{Q}(v_{i,j})} \cdot v_{i,j-2} \end{aligned} \quad (9-2)$$

The coefficient in $v_{i,j}$ comes from moving forward then backwards in each cardinal direction, so any negative edge traversed this way will account for two minus signs, and the end result is always 1, as per Proposition 9.1.1. Vertices of the form $v_{i\pm 1, j\pm 1}$ do not feature: each such vertex can be reached from $v_{i,j}$ via exactly two distinct edge-paths, which enclose a square with vertices in $(\mathbb{Z} + \frac{1}{2})^2$. By Proposition 9.1.1, these paths contribute with coefficients that are equal in absolute value, but the definition of Kasteleyn signing ensures they have opposite signs, so they cancel out. The vertices $v_{i,j\pm 2}$ are reached from $v_{i,j}$ via vertical edge-paths, so they never contain a negative edge for our choice of Kasteleyn signing. On the other hand, the vertices $v_{i\pm 2, j}$ are reached from $v_{i,j}$ via horizontal edge-paths, so they always contain exactly one negative edge for our fixed Kasteleyn signing.

We will use formula (9-2) to find M 's eigenvectors. M acts on the space of equivalence classes of vertices of $G(\mathbb{Z}^2)$, so a vector in that space may be thought of as a weight on each such equivalence class. Alternatively, we may

think of it as a function on the vertices of $G(\mathbb{Z}^2)$ that is L -periodic, so it coincides on each vertex of a given equivalence class.

For each $x \in (\mathbb{R}^2)^*$, consider the function $\zeta_x : \mathbb{R}^2 \rightarrow \mathbb{S}^1$ given by $\zeta_x(v) = \exp(2\pi\mathbf{i} \cdot x(v))$. Now, let $q_0, q_1 \in \mathbb{S}^1$ be fixed, and take any $z \in (\mathbb{R}^2)^*$ with $\zeta_z(v_0) = q_1^{-1}$ and $\zeta_z(v_1) = q_0$, where $v_0, v_1 \in L$ are the vectors used in the definition of D_L . Notice this implies $\zeta_z(v + a \cdot v_0 + b \cdot v_1) = q_0^b \cdot q_1^{-a} \cdot \zeta_z(v)$ for all $v \in \mathbb{R}^2$. Finally, consider the function f_z defined on $(\mathbb{Z} + \frac{1}{2})^2$ by $f_z(v) = \zeta_z(v) \cdot \mathcal{Q}(v)$; observe it is L -periodic. We claim f_z is an eigenvector of $M = M(q_0, q_1)$, i.e.,

$$v_z = \sum_{i,j} f_z(v_{i,j}) \cdot v_{i,j}$$

is an eigenvector of M . Indeed, notice the $v_{i,j}$ -coordinate of Mv_z is given by

$$\begin{aligned} [Mv_z]_{i,j} &= 4 \cdot f_z(v_{i,j}) - f_z(v_{i+2,j}) \cdot \frac{\mathcal{Q}(v_{i,j})}{\mathcal{Q}(v_{i+2,j})} - f_z(v_{i-2,j}) \cdot \frac{\mathcal{Q}(v_{i,j})}{\mathcal{Q}(v_{i-2,j})} \\ &\quad + f_z(v_{i,j+2}) \cdot \frac{\mathcal{Q}(v_{i,j})}{\mathcal{Q}(v_{i,j+2})} + f_z(v_{i,j-2}) \cdot \frac{\mathcal{Q}(v_{i,j})}{\mathcal{Q}(v_{i,j-2})}, \end{aligned}$$

which we may rearrange to

$$f_z(v_{i,j}) \cdot \left(4 - \frac{\zeta_z(v_{i+2,j})}{\zeta_z(v_{i,j})} - \frac{\zeta_z(v_{i-2,j})}{\zeta_z(v_{i,j})} + \frac{\zeta_z(v_{i,j+2})}{\zeta_z(v_{i,j})} + \frac{\zeta_z(v_{i,j-2})}{\zeta_z(v_{i,j})} \right).$$

If $x = (x_0, x_1)$, let $\zeta_{x,0} = \exp(2\pi\mathbf{i} \cdot x_0)$ and $\zeta_{x,1} = \exp(2\pi\mathbf{i} \cdot x_1)$. Then we may write

$$\begin{aligned} [Mv_z]_{i,j} &= f_z(v_{i,j}) \cdot (4 - \zeta_{z,0}^2 - \zeta_{z,0}^{-2} + \zeta_{z,1}^2 + \zeta_{z,1}^{-2}) \\ &= f_z(v_{i,j}) \cdot \left(-(\zeta_{z,0} - \zeta_{z,0}^{-1})^2 + (\zeta_{z,1} + \zeta_{z,1}^{-1})^2 \right). \end{aligned}$$

In other words, for $z = (z_0, z_1)$, v_z is an eigenvector of $M = M(q_0, q_1)$ with associated eigenvalue

$$\begin{aligned} \lambda_z &= \left(-(\zeta_{z,0} - \zeta_{z,0}^{-1})^2 + (\zeta_{z,1} + \zeta_{z,1}^{-1})^2 \right) \\ &= 4 \cdot \left((\sin(2\pi z_0))^2 + (\cos(2\pi z_1))^2 \right). \end{aligned}$$

We will think of each pair $q_0, q_1 \in \mathbb{S}^1$ as a homomorphism $q : L \rightarrow \mathbb{S}^1$ with $q(v_0) = q_1^{-1}$ and $q(v_1) = q_0$, so the condition that $\zeta_z(v_0) = q_1^{-1}$ and $\zeta_z(v_1) = q_0$ reduces to $\zeta_z|_L = q$. Thus, for each $z \in (\mathbb{R}^2)^*$ with $\zeta_z|_L = q$, f_z is an eigenvector of $M = M(q)$. Are all its eigenvectors of that form?

First note that $\zeta_z|_L = \zeta_{\tilde{z}}|_L \iff \forall v \in L, (z - \tilde{z})(v) \in \mathbb{Z} \iff z - \tilde{z} \in L^*$.

We may thus identify $\text{Hom}(L, \mathbb{S}^1)$ with $(\mathbb{R}^2)^* / L^*$. On the other hand, $z - \tilde{z} \in (\mathbb{Z}^2)^* \iff \exists c \in \{\pm 1\}, \forall v \in (\mathbb{Z} + \frac{1}{2})^2, \zeta_z(v) = c \cdot \zeta_{\tilde{z}}(v)$.

In particular, if z and \tilde{z} are in the same equivalence class $q \in (\mathbb{R}^2)^* / L^*$ and $z - \tilde{z}$ is in $(\mathbb{Z}^2)^*$ — i.e., z and \tilde{z} represent the same element in $q / (\mathbb{Z}^2)^*$ —, then f_z and $f_{\tilde{z}}$ are essentially the same eigenvector of $M(q)$.

Proposition 9.1.2. *Let z, \tilde{z} be in the same equivalence class $q \in (\mathbb{R}^2)^* / L^*$, but in different equivalence classes of $q / (\mathbb{Z}^2)^*$. Then v_z and $v_{\tilde{z}}$ are orthogonal.*

Proof. We will show that $\langle v_z, v_{\tilde{z}} \rangle = 0$. Notice f_z and $f_{\tilde{z}}$ are both L -periodic, so we need only verify $\sum_{i,j} f_z(v_{i,j}) \cdot \overline{f_{\tilde{z}}(v_{i,j})} = 0$, where the sum is carried over the $v_{i,j}$ in a single $D(a, b)$, say $D(0, 0)$.

Since $\overline{\zeta_x} = \zeta_{-x}$ for all $x \in (\mathbb{R}^2)^*$, it holds that

$$\sum_{i,j} f_z(v_{i,j}) \cdot \overline{f_{\tilde{z}}(v_{i,j})} = \sum_{i,j} \zeta_z(v_{i,j}) \cdot \overline{\zeta_{\tilde{z}}(v_{i,j})} = \sum_{i,j} \exp\left(2\pi i \cdot (z - \tilde{z})(v_{i,j})\right).$$

The $v_{i,j}$ in $D(0, 0)$ are vertices $(i, j) + (\frac{1}{2}, \frac{1}{2})$ in $(\mathbb{Z} + \frac{1}{2})^2$ with $0 \leq i < x_0$ and $0 \leq j < y_1$, where $v_0 = (x_0, 0)$ and $v_1 = (x_1, y_1)$ are the generators of L used in the construction of D_L . Letting $C = \exp\left(2\pi i \cdot \langle z - \tilde{z}, (\frac{1}{2}, \frac{1}{2}) \rangle\right)$, we have that $\sum_{i,j} f_z(v_{i,j}) \cdot \overline{f_{\tilde{z}}(v_{i,j})}$ is

$$C \cdot \left(\sum_{0 \leq i < x_0} \exp\left(2\pi i \cdot \langle z - \tilde{z}, (i, 0) \rangle\right) \right) \cdot \left(\sum_{0 \leq j < y_1} \exp\left(2\pi i \cdot \langle z - \tilde{z}, (0, j) \rangle\right) \right)$$

Rewrite the expression above as:

$$\sum_{i,j} f_z(v_{i,j}) \cdot \overline{f_{\tilde{z}}(v_{i,j})} = C \cdot \left(\sum_{0 \leq i < x_0} (\zeta_{z-\tilde{z},0})^i \right) \cdot \left(\sum_{0 \leq j < y_1} (\zeta_{z-\tilde{z},1})^j \right) \quad (9-3)$$

Observe that $(\zeta_{z-\tilde{z},0})^{x_0} = \zeta_{z-\tilde{z}}(v_0) = 1$, because z and \tilde{z} are both in the same equivalence class $q \in (\mathbb{R}^2)^* / L^*$ (they satisfy $\zeta_z|_L = \zeta_{\tilde{z}}|_L$). In similar fashion, $\zeta_{z-\tilde{z}}(v_1) = (\zeta_{z-\tilde{z},0})^{x_1} \cdot (\zeta_{z-\tilde{z},1})^{y_1} = 1$, so $(\zeta_{z-\tilde{z},1})^{y_1} = (\zeta_{z-\tilde{z},0})^{-x_1}$. Now, $\zeta_{z-\tilde{z},0}$ and $\zeta_{z-\tilde{z},1}$ may not both be 1, for z and \tilde{z} are in different equivalence classes of $q / (\mathbb{Z}^2)^*$, so their difference does not lie in $(\mathbb{Z}^2)^*$.

If $\zeta_{z-\tilde{z},0} \neq 1$, then $(\zeta_{z-\tilde{z},0})^{x_0} = 1$ implies that $\sum_{0 \leq i < x_0} (\zeta_{z-\tilde{z},0})^i$ is a symmetrical sum in \mathbb{S}^1 , and hence 0. If $\zeta_{z-\tilde{z},0} = 1$, then $(\zeta_{z-\tilde{z},1})^{y_1} = 1$, so similarly $\sum_{0 \leq j < y_1} (\zeta_{z-\tilde{z},1})^j = 0$. Regardless of the situation, we see from (9-3) that $\sum_{i,j} f_z(v_{i,j}) \cdot \overline{f_{\tilde{z}}(v_{i,j})} = 0$, as desired. \square

Proposition 9.1.2 tells us each equivalence class of $q / (\mathbb{Z}^2)^*$ corresponds to a different eigenvector of $M(q)$. The next proposition shows these are ‘all’ of its eigenvectors.

Proposition 9.1.3. *The groups $L^* / (\mathbb{Z}^2)^*$, $(\mathbb{Z}^2 / L)^*$ and \mathbb{Z}^2 / L are all isomorphic.*

Proof. We first show $L^* / (\mathbb{Z}^2)^*$ and $(\mathbb{Z}^2 / L)^*$ are isomorphic (naturally so, in fact). Any $f \in L^*$ admits a unique extension by linearity to a functional in $(\mathbb{R}^2)^*$, which we still call f . Given any such f , define $\zeta_f : \mathbb{R}^2 \rightarrow \mathbb{S}^1$ by $\zeta_f(v) = \exp(2\pi i \cdot f(v))$; notice it is a homomorphism. Consider $\zeta : (\mathbb{R}^2)^* \rightarrow \text{Hom}(\mathbb{R}^2, \mathbb{S}^1)$ given by $\zeta(f) = \zeta_f$. We claim ζ defines an isomorphism between $L^* / (\mathbb{Z}^2)^*$ and $(\mathbb{Z}^2 / L)^*$. It’s easy to check that it is an homomorphism; we prove it is a bijection. Observe that the following are equivalent:

- f and \tilde{f} are in L^* and belong to the same equivalence class of $L^* / (\mathbb{Z}^2)^*$;
- $\zeta_f|_L = \zeta_{\tilde{f}}|_L = \mathbb{1}$ and $\frac{\zeta_f}{\zeta_{\tilde{f}}}|_{\mathbb{Z}^2} = \mathbb{1}$ (i.e., ζ_f and $\zeta_{\tilde{f}}$ coincide on \mathbb{Z}^2).

In addition, each $g \in (\mathbb{Z}^2 / L)^* = \text{Hom}(\mathbb{Z}^2 / L, \mathbb{S}^1)$ corresponds to a unique $\tilde{g} \in \text{Hom}(\mathbb{Z}^2, \mathbb{S}^1)$ that is L -periodic (it satisfies $\tilde{g}|_L = \mathbb{1}$). It follows that ζ injectively maps $L^* / (\mathbb{Z}^2)^*$ to $(\mathbb{Z}^2 / L)^*$.

To see ζ also does so surjectively, observe that any $\tilde{g} \in \text{Hom}(\mathbb{Z}^2, \mathbb{S}^1)$ satisfies

$$\tilde{g}(a_1, a_2) = \tilde{g}(1, 0)^{a_1} \cdot \tilde{g}(0, 1)^{a_2}.$$

Choose any $\alpha_1, \alpha_2 \in \mathbb{R}$ with $\tilde{g}(1, 0) = \exp(2\pi i \cdot \alpha_1)$ and $\tilde{g}(0, 1) = \exp(2\pi i \cdot \alpha_2)$. Then

$$\tilde{g}(a_1, a_2) = \exp\left(2\pi i \cdot (\alpha_1 \cdot a_1 + \alpha_2 \cdot a_2)\right),$$

so \tilde{g} is realized by ζ_f , where $f(a_1, a_2) = \alpha_1 \cdot a_1 + \alpha_2 \cdot a_2$. Now it’s easy to check that $\tilde{g}|_L = \mathbb{1}$ if and only if $\forall v \in L, f(v) \in \mathbb{Z}$ — that is, if and only if $f \in L^*$.

We now show $(\mathbb{Z}^2 / L)^*$ and \mathbb{Z}^2 / L are isomorphic. Since \mathbb{Z}^2 / L is abelian and finite, there are prime powers m_1, \dots, m_k with $\mathbb{Z}^2 / L \simeq \mathbb{Z} / (m_1) \oplus \dots \oplus \mathbb{Z} / (m_k)$. For each $1 \leq i \leq k$, let e_i be a generator of $\mathbb{Z} / (m_i)$, so each $f \in (\mathbb{Z}^2 / L)^*$ is uniquely defined by the value it takes on the e_i ’s. Each $1 \leq i \leq k$ satisfies $f(e_i)^{m_i} = 1$, so $f(e_i)$ must be a m_i -th root of unity. Under multiplication, the n -th roots of unity form a cyclic group of order n ,

which is isomorphic to $\mathbb{Z}/(n)$. It follows that $(\mathbb{Z}^2/L)^*$ is also isomorphic to $\mathbb{Z}/(m_1) \oplus \dots \oplus \mathbb{Z}/(m_k)$, which completes the proof. \square

Each class $q \in (\mathbb{R}^2)^*/L^*$ is isomorphic to L^* , so $q/(\mathbb{Z}^2)^* \simeq L^*/(\mathbb{Z}^2)^*$. Since $n = \text{card}(\mathbb{Z}^2/L)$ is the number of squares on the fundamental domain D_L , M is $n \times n$. It follows from Proposition 9.1.3 that $q/(\mathbb{Z}^2)^*$ has n elements, so Proposition 9.1.2 implies they provide all n of $M(q)$'s eigenvectors.

Corollary 9.1.4. *For each $q \in \text{Hom}(L, \mathbb{S}^1) \simeq (\mathbb{R}^2)^*/L^*$ we have*

$$\det(M(q)) = \prod_{z \in q/(\mathbb{Z}^2)^*} \lambda_z \quad ,$$

where for $z = (z_0, z_1)$ it holds that $\lambda_z = 4 \cdot \left((\sin(2\pi z_0))^2 + (\cos(2\pi z_1))^2 \right)$.

Notice λ_z does not depend on choice of representative for $z \in q/(\mathbb{Z}^2)^*$. Moreover, $M = KK^* \oplus K^*K$ implies $\det(M) = |\det(K)|^4$; in particular, $\det(M)$ is always a nonnegative real and $|\det(K)| = \sqrt[4]{\det(M)}$.

9.2

Spaces of L -quasiperiodic functions and the case of K, K^*

We saw that M acted on the space of L -periodic functions on $(\mathbb{Z} + \frac{1}{2})^2$ — the weights on squares of D_L —, but we now present another interpretation of the situation. For each $q \in \text{Hom}(L, \mathbb{S}^1)$, consider the space $\mathcal{F}(L, q)$ of complex functions on $(\mathbb{Z} + \frac{1}{2})^2$ that are L -quasiperiodic with parameter q , meaning

$$f \in \mathcal{F}(L, q) \iff \forall u \in (\mathbb{Z} + \frac{1}{2})^2, \forall v \in L, f(u + v) = q(v) \cdot f(u).$$

Of course, q is defined by the values it takes on a basis of L , v_0 and v_1 for instance. If $q(v_0) = q_1^{-1}$ and $q(v_1) = q_0$, the figure in the next page is a representation of some generic $f \in \mathcal{F}(L, q)$, where the number next to a vertex indicates the value f takes on it.

There is a host of obvious isomorphisms between the space of L -periodic functions on $(\mathbb{Z} + \frac{1}{2})^2$ and $\mathcal{F}(L, q)$: simply make two functions correspond if they agree on any particular fundamental domain of L . We will use the isomorphism ψ that checks for functions agreeing on $D_L = D(0, 0)$; under ψ , each $v_{i,j} \in D_L$

PUC-Rio - Certificação Digital Nº 1321816/CA

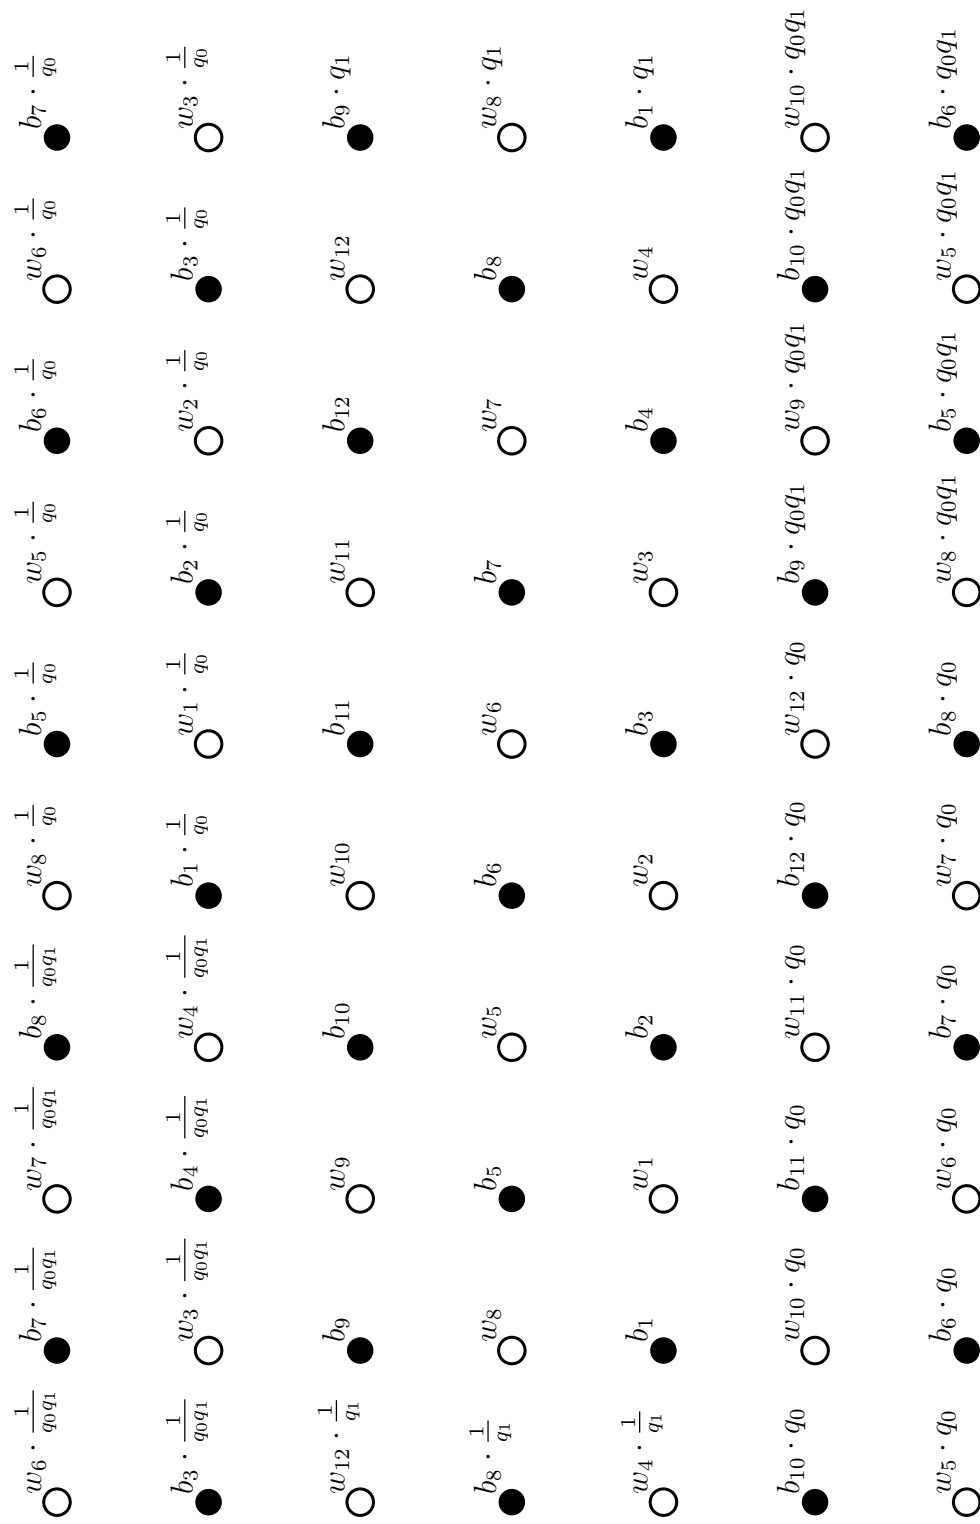


Figure 9.2: Representation of some generic $f \in \mathcal{F}(L, q)$, where L is generated by $v_0 = (8, 0)$ and $v_1 = (3, 3)$. Vertices are in $G(\mathbb{Z}^2)$.

corresponds to the function $g_{i,j}$ defined by taking the value 1 on $v_{i,j}$ and 0 on every other vertex of D_L (and by extension through L -quasiperiodicity).

This definition serves to take the $\mathcal{Q}(v)$'s away from formula (9-2) and into the space itself, which in turn simplifies the expression of $M(q)$. Indeed, let \widetilde{M} be given by

$$\begin{aligned} (\widetilde{M}g)(x, y) = & 4 \cdot g(x, y) - g(x + 2, y) - g(x - 2, y) \\ & + g(x, y + 2) + g(x, y - 2) \end{aligned} \tag{9-4}$$

With our choice isomorphism and convention for values of $q(v_0)$ and $q(v_1)$ above, it's easily checked that $\widetilde{M}\psi(v_{i,j}) = \psi(Mv_{i,j})$, so $\widetilde{M}(q)$ emulates $M(q)$'s action on $\mathcal{F}(L, q)$. We will thus refer to \widetilde{M} also by M . Moreover, arguments similar to those used before will show that for formula (9-4), functions of the form ζ_z are eigenvectors of M (rather than $f_z = \zeta_z \cdot \mathcal{Q}$), and they are associated to the same eigenvalues λ_z . Perhaps a yet better benefit of this approach is that it defines the linear map $M(q)$ without resorting to a particular choice or construction of fundamental domain for L .

We may use this interpretation to describe K and K^* similarly. Let $L_0 \subset \mathbb{Z}^2$ be the lattice spanned by $\{(2, 0), (1, 1)\}$; observe that any valid lattice is a sublattice of L_0 . Consider the affine lattices $L_b = L_0 + (\frac{1}{2}, \frac{1}{2})$ and $L_w = L_0 + (\frac{1}{2}, -\frac{1}{2})$; notice L_b is the set of black vertices of $G(\mathbb{Z}^2)$ and L_w is the set of white ones. Now let $\mathcal{B}(L, q)$ be the space of complex functions on L_b that are L -quasiperiodic with parameter q , and similarly for $\mathcal{W}(L, q)$.

Like with $M(q)$, we may interpret $K(q)$ as a linear map $\mathcal{B}(L, q) \rightarrow \mathcal{W}(L, q)$ and $K(q)^*$ as a linear map $\mathcal{W}(L, q) \rightarrow \mathcal{B}(L, q)$ (note that for $z \in \mathbb{S}^1$, $\overline{z^{-1}} = z$). We provide the analogues of formula (9-4):

$$(Kg)(x, y) = g(x + 1, y) - g(x - 1, y) + g(x, y + 1) + g(x, y - 1) \tag{9-5a}$$

$$(K^*g)(x, y) = -g(x + 1, y) + g(x - 1, y) + g(x, y + 1) + g(x, y - 1) \tag{9-5b}$$

In each case, which term is negatively signed is justified by our fixed choice of Kasteleyn signing: each domino on the brick wall $b_{\mathcal{N}}$ has a black square on the left and a white square on the right.




Figure 9.3: The brick wall $b_{\mathcal{N}}$. The marked vertex is the origin.

The matrices K and K^* we previously constructed represent the linear maps above in bases we now describe. Let b_i be D_L 's i -th black vertex and w_j its j -th white vertex, as we enumerated them. Let $g[b_i] \in \mathcal{B}(L, q)$ be defined by taking the value 1 on b_i and 0 on each other black vertex of D_L , and let $g[w_j] \in \mathcal{W}(L, q)$ take the value 1 on w_j and 0 on each other white vertex of D_L . Then the $g[b_i]$ form an ordered basis D_b of $\mathcal{B}(L, q)$, and the $g[w_j]$ form an ordered basis D_w of $\mathcal{W}(L, q)$. The matrix $K = K_D$ represents the linear map of (9-5a) from D_b to D_w , and in similar fashion K_D^* goes from D_w to D_b .

We will now choose different bases for $\mathcal{B}(L, q)$ and $\mathcal{W}(L, q)$, given by ζ_z 's. Observe that one such function satisfies $\zeta_z(u + v) = \zeta_z(u) \cdot \zeta_z(v)$, so it¹ is an element of $\mathcal{B}(L, q)$ or of $\mathcal{W}(L, q)$ if and only if $\zeta_z|_L = q$. Recall our identification of $\text{Hom}(L, \mathbb{S}^1)$ with $(\mathbb{R}^2)^* / L^*$: $\zeta_z|_L = \zeta_{\tilde{z}}|_L$ if and only if z and \tilde{z} are in the same equivalence class.

Granted, for z and \tilde{z} in one same equivalence class $q \in (\mathbb{R}^2)^* / L^*$, ζ_z and $\zeta_{\tilde{z}}$ need not agree on L_b . We assert that $\exists c \in \mathbb{C}$, $\zeta_z|_{L_b} = c \cdot \zeta_{\tilde{z}}|_{L_b} \iff z - \tilde{z} \in L_0^*$. Indeed, it's easy to check that

$$\zeta_z|_{L_b} = c \cdot \zeta_{\tilde{z}}|_{L_b} \iff \forall u \in L_0, \exp\left(2\pi\mathbf{i} \cdot (z - \tilde{z})(u)\right) \cdot \exp\left(2\pi\mathbf{i} \cdot (z - \tilde{z})\left(\frac{1}{2}, \frac{1}{2}\right)\right) = c,$$

so when there is one such c we have $\exp\left(2\pi\mathbf{i} \cdot (z - \tilde{z})\left(\frac{1}{2}, \frac{1}{2}\right)\right) = c$, because $u = 0 \in L_0$. It follows that $\exp\left(2\pi\mathbf{i} \cdot (z - \tilde{z})(u)\right) = 1$ whenever $u \in L_0$, so $z - \tilde{z} \in L_0^*$. On the other hand, when $z - \tilde{z} \in L_0^*$ we may take $c = \exp\left(2\pi\mathbf{i} \cdot (z - \tilde{z})\left(\frac{1}{2}, \frac{1}{2}\right)\right)$, and the assertion is proved.

By the same token, $\exists c \in \mathbb{C}$, $\zeta_z|_{L_w} = c \cdot \zeta_{\tilde{z}}|_{L_w} \iff z - \tilde{z} \in L_0^*$.

Proposition 9.2.1. *Let z, \tilde{z} be in the same equivalence class $q \in (\mathbb{R}^2)^* / L^*$, but in two different equivalence classes of q / L_0^* . Then ζ_z and $\zeta_{\tilde{z}}$ are orthogonal in each of $\mathcal{B}(L, q)$ and $\mathcal{W}(L, q)$.*

Proof. Observe that if $L = L_0^*$, there are no such z, \tilde{z} as in the statement: q / L_0^* has a single class. We thus suppose without loss of generality that $L \subsetneq L_0$.

¹Its restriction to the relevant subset of $(\mathbb{Z} + \frac{1}{2})^2$

The idea of the proof goes similar to that of Proposition 9.1.2. We will show the expression below is always 0 whenever $v = z - \tilde{z} \in L^* \setminus L_0^*$.

$$C \cdot \left[\sum_{0 \leq i < x_0} (\zeta_{v,0})^i \cdot \left(\sum_{\substack{0 \leq j < y_1 \\ j \equiv i \pmod{2}}} (\zeta_{v,1})^j \right) \right] \tag{9-6a}$$

$$= C \cdot \left[\sum_{0 \leq j < y_1} (\zeta_{v,1})^j \cdot \left(\sum_{\substack{0 \leq i < x_0 \\ i \equiv j \pmod{2}}} (\zeta_{v,0})^i \right) \right], \tag{9-6b}$$

where $C = \exp(2\pi\mathbf{i} \cdot v(\frac{1}{2}, \frac{1}{2}))$ for $\mathcal{B}(L, q)$ and $C = \exp(2\pi\mathbf{i} \cdot v(\frac{1}{2}, -\frac{1}{2}))$ for $\mathcal{W}(L, q)$. Notice that because $v \in L^*$ we have

$$\zeta_v(v_0) = (\zeta_{v,0})^{x_0} = 1 \tag{9-7a}$$

$$\zeta_v(v_1) = (\zeta_{v,0})^{x_1} \cdot (\zeta_{v,1})^{y_1} = 1. \tag{9-7b}$$

On the other hand, $v \notin L_0^*$ implies $(\zeta_{v,0}, \zeta_{v,1})$ is not in the span of $\{(\frac{1}{2}, -\frac{1}{2}), (0, 1)\}$. We will divide the proof in cases.

Case 1. $\zeta_{v,0} \neq 1$ and $x_0 > 2$.

Remember x_0 is always positive and even, so $x_0 = 2k$ for some integer $k > 1$. Equation (9-7a) implies $((\zeta_{v,0})^2)^k = 1$, so that

$$\sum_{\substack{0 \leq i < x_0 \\ i \text{ even}}} (\zeta_{v,0})^i = \sum_{0 \leq i < k} ((\zeta_{v,0})^2)^i$$

is a symmetrical sum in \mathbb{S}^1 , and hence 0. Of course,

$$\sum_{\substack{0 \leq i < x_0 \\ i \text{ odd}}} (\zeta_{v,0})^i = \zeta_{v,0} \cdot \sum_{\substack{0 \leq i < x_0 \\ i \text{ even}}} (\zeta_{v,0})^i,$$

and thus it is also 0. It follows that expression (9-6b) is 0.

Case 2. $\zeta_{v,0} \neq 1$ and $x_0 = 2$.

In this case, equation (9-7a) implies $\zeta_{v,0} = -1$, so we may rewrite expression (9-6a) as

$$C \cdot \left[\left(\sum_{\substack{0 \leq j < y_1 \\ j \text{ even}}} (\zeta_{v,1})^j \right) - \left(\sum_{\substack{0 \leq j < y_1 \\ j \text{ odd}}} (\zeta_{v,1})^j \right) \right] = C \cdot \left[\sum_{0 \leq j < y_1} (-\zeta_{v,1})^j \right].$$

Additionally, because $v \notin L_0^*$ it may not be $\zeta_{v,1} = \pm 1$.

Suppose first that x_1, y_1 are both even. Then $(\zeta_{v,0})^{x_1} = 1$, and equation (9-7b) implies $(\zeta_{v,1})^{y_1} = 1$. Since $\zeta_{v,1} \neq \pm 1$, we must have $y_1 \geq 4$, so $\sum_{0 \leq j < y_1} (-\zeta_{v,1})^j$ is a symmetrical sum in \mathbb{S}^1 and hence 0.

Suppose now that x_1, y_1 are both odd. Then $(\zeta_{v,0})^{x_1} = -1$, and equation (9-7b) implies $(\zeta_{v,1})^{y_1} = -1$, or $(-\zeta_{v,1})^{y_1} = 1$. Since $\zeta_{v,1} \neq \pm 1$, we must have $y_1 \geq 3$, so $\sum_{0 \leq j < y_1} (-\zeta_{v,1})^j$ is a symmetrical sum in \mathbb{S}^1 and hence 0.

Case 3. $\zeta_{v,0} = 1$.

In this case, equation (9-7b) implies $(\zeta_{v,1})^{y_1} = 1$. Because $v \notin L_0^*$, it may not be $\zeta_{v,1} = \pm 1$, so $y_1 \geq 3$. Additionally, we may rewrite expression (9-6a) as

$$C \cdot \frac{x_0}{2} \cdot \left[\sum_{0 \leq j < y_1} (\zeta_{v,1})^j \right],$$

and $\sum_{0 \leq j < y_1} (\zeta_{v,1})^j$ is a symmetrical sum in \mathbb{S}^1 . □

Proposition 9.2.1 guarantees that if we choose a z out of each class in q/L_0^* , the ζ_z 's are linearly independent in each of $\mathcal{B}(L, q)$ and $\mathcal{W}(L, q)$. Moreover, as in Proposition 9.1.3, L^*/L_0^* and L_0/L are isomorphic, so these vectors generate their respective spaces — they are bases for them.

Now, observe that applying formulas (9-5a) and (9-5b) to ζ_z yield very simple results. Indeed, we have that $(K\zeta_z|_{L_b})(x, y)$ is

$$\begin{aligned} & \zeta_z|_{L_b}(x+1, y) - \zeta_z|_{L_b}(x-1, y) + \zeta_z|_{L_b}(x, y+1) + \zeta_z|_{L_b}(x, y-1) \\ &= \left((\zeta_{z,0} - \zeta_{z,0}^{-1}) + (\zeta_{z,1} + \zeta_{z,1}^{-1}) \right) \cdot \zeta_z|_{L_w}(x, y). \end{aligned}$$

In other words, $(K\zeta_z|_{L_b}) = \lambda(K, z) \cdot \zeta_z|_{L_w}$, and by the same token $(K^*\zeta_z|_{L_w}) = \lambda(K^*, z) \cdot \zeta_z|_{L_b}$, where for $z = (z_0, z_1)$

$$\begin{aligned} \lambda(K, z) &= (\zeta_{z,1} + \zeta_{z,1}^{-1}) + (\zeta_{z,0} - \zeta_{z,0}^{-1}) \\ &= 2 \cos(2\pi z_1) + 2i \sin(2\pi z_0) \end{aligned} \tag{9-8}$$

$$\begin{aligned} \lambda(K^*, z) &= (\zeta_{z,1} + \zeta_{z,1}^{-1}) - (\zeta_{z,0} - \zeta_{z,0}^{-1}) \\ &= 2 \cos(2\pi z_1) - 2i \sin(2\pi z_0) \end{aligned} \tag{9-9}$$

Let $n = 2m$ be $\text{card}(\mathbb{Z}^2/L)$, so $\mathcal{B}(L, q)$ and $\mathcal{W}(L, q)$ are both m dimensional. Choose z_1, \dots, z_m in different classes of q/L_0^* . For each $1 \leq i \leq m$, let $\zeta_i = \zeta_{z_i}$, $\zeta_i[b] = \zeta_i|_{L_b}$ and $\zeta_i[w] = \zeta_i|_{L_w}$. We will denote by E_b the ordered basis for $\mathcal{B}(L, q)$ given by the $\zeta_i[b]$, and by E_w the ordered basis for $\mathcal{W}(L, q)$ given by the $\zeta_i[w]$. We say these are *exponential* bases.

The discussion leading up to this point should make it clear that the matrix K_E representing K from E_b to E_w is diagonal with entries $\lambda(K, z_i)$, and in similar fashion K_E^* is diagonal with entries $\lambda(K^*, z_i)$. We can thus calculate $\det(K_E)$ and $\det(K_E^*)$ with a simple product. How do these determinants relate to that of the Kasteleyn matrix K_D we had previously constructed?

Let $X(D_b, E_b)$ be the matrix that changes basis from D_b to E_b , and similarly define $X(D_w, E_w)$. Then it's clear that $K_D = X(D_w, E_w)^{-1} \cdot K_E \cdot X(D_b, E_b)$. We will study the matrices X in order to understand the relation between $\det(K_E)$ and $\det(K_D)$.

The j -th column of $X(D_b, E_b)$ is $\zeta_j[b]$ written in the ordered basis D_b . It is easy to do so:

$$\zeta_j[b] = \sum_{1 \leq i \leq m} \zeta_j[b](b_i) \cdot g[b_i].$$

Notice $\zeta_j[b](b_i)$ is simply $\zeta_j(b_i)$. Thus, $X(D_b, E_b)_{i,j} = \zeta_j(b_i)$ and similarly $X(D_w, E_w)_{i,j} = \zeta_j(w_i)$.

Is there a relation between $\zeta_j(b_i)$ and $\zeta_j(w_i)$? Not in principle — they depend on our enumeration of colored vertices. In Figure 9.4, we present a choice of vertex enumeration on D_L that makes one such relation apparent.

We explain it in words. D_L has y_1 lines with x_0 squares each. There are two types of lines: black lines are lines whose first square (from left to right) is black, and similarly for white lines. Notice line types alternate, and the first line (whose squares touch the horizontal line through the origin) is always black. We assign 1 to the leftmost black square in that line, and the same number to the white square on its right. If there's a black square to the right of that white square, we assign the next number, and so on until the line's squares are all labeled; we then proceed to the next line. We repeat the process above, except this line is white, so we skip its first white square. Because x_0 is always even,




Figure 9.4: Enumeration of D_L 's vertices. b_1 is the black square $[0, 1]^2$. Notice this was used on the example of Figure 7.2.

a white line always ends with a black square, and as we enumerated them, no white square lies to the right of it. We then assign the first white square on that line (the one we skipped) the same number as its last, black square. Repeating this procedure enumerates all squares on D_L .

With this enumeration, each white vertex w_i satisfies $w_i = b_i + (1, 0)$ **except** for white vertices in the beginning of a white line: these satisfy $w_i + v_0 = b_i + (1, 0)$. It's easy to see there are $\lfloor \frac{y_1}{2} \rfloor$ such vertices. For vertices of the first kind, we have

$$\zeta_j(w_i) = \zeta_j(b_i + (1, 0)) = \zeta_j(b_i) \cdot \zeta_{j,0},$$

while for vertices of the second kind we have

$$\zeta_j(w_i) = \zeta_j(b_i + (1, 0) - v_0) = \zeta_j(b_i) \cdot \zeta_{j,0} \cdot q_1,$$

where the equality $\zeta_j(-v_0) = q_1$ comes from $\zeta_j|_L = q$.

Thus, $X(D_w, E_w)$ is obtained from $X(D_b, E_b)$ by multiplying each column j by $\zeta_{j,0}$ and $\lfloor \frac{y_1}{2} \rfloor$ of its lines — the ones that correspond to indices i for which w_i is in the beginning of a white line — by q_1 . It follows that

$$\det \left(X(D_w, E_w) \right) = q_1^{\lfloor \frac{y_1}{2} \rfloor} \cdot \left(\prod_{j=1}^m \zeta_{j,0} \right) \cdot \det \left(X(D_b, E_b) \right),$$

so $\det(K_D) = \rho \cdot \det(K_E)$, where $\rho = q_1^{-\lfloor \frac{y_1}{2} \rfloor} \cdot \prod_{j=1}^m (\zeta_{j,0})^{-1}$. Notice $\rho \in \mathbb{S}^1$.

9.3

Formulas for $\det(K_E)$, ρ and uniform scaling

We will now make choices for our exponential bases from which we'll derive explicit formulas for $\det(K_E)$ and ρ . Observe L_0^* may be alternatively described (under inner product identification) as the lattice $\mathbb{Z}^2 \cup (\mathbb{Z} + \frac{1}{2})^2$, so by equation (9-8) each diagonal entry of K_E is unique up to sign.

Because they're L -quasiperiodic with parameter q , the ζ_z 's must satisfy:

$$\begin{cases} \zeta_z(v_0) = (\zeta_{z,0})^{x_0} & = q_1^{-1} \\ \zeta_z(v_1) = (\zeta_{z,0})^{x_1} \cdot (\zeta_{z,1})^{y_1} & = q_0 \end{cases}$$

There are $2m = x_0 \cdot y_1$ solutions — twice the number of elements in a basis for $\mathcal{B}(L, q)$ or $\mathcal{W}(L, q)$. If $q_0 = \exp(2\pi\mathbf{i} \cdot u_0)$ and $q_1 = \exp(2\pi\mathbf{i} \cdot u_1)$, these can be written as

$$\begin{cases} \zeta_{z,0} = \zeta[k_0, k_1]_0 = \exp\left(2\pi\mathbf{i} \cdot \frac{1}{x_0} \cdot (k_0 - u_1)\right) \\ \zeta_{z,1} = \zeta[k_0, k_1]_1 = \exp\left(2\pi\mathbf{i} \cdot \frac{1}{x_0 \cdot y_1} \cdot \left((u_0 + k_1) \cdot x_0 + (u_1 - k_0) \cdot x_1\right)\right), \end{cases}$$

where $0 \leq k_0 < x_0$ and $0 \leq k_1 < y_1$.

Let $E = \{\zeta[k_0, k_1] \mid 0 \leq k_0 < \frac{1}{2}x_0 \text{ and } 0 \leq k_1 < y_1\}$.

Proposition 9.3.1. *E defines a basis for each of $\mathcal{B}(L, q)$ and $\mathcal{W}(L, q)$.*

Proof. By Proposition 9.2.1 and the discussion preceding it, if for all $0 \leq k_0, l_0 < \frac{1}{2}x_0$ and for all $0 \leq k_1, l_1 < y_1$ it holds that

$$\left(\exists c \in \mathbb{C}, \zeta[k_0, k_1]|_{L_b} = c \cdot \zeta[l_0, l_1]|_{L_b}\right) \longrightarrow k_0 = l_0 \text{ and } k_1 = l_1,$$

then E defines a basis for $\mathcal{B}(L, q)$, and similarly for $\mathcal{W}(L, q)$. As before, in both cases there is one such c if and only if $\zeta[k_0, k_1] \cdot (\zeta[l_0, l_1])^{-1} = \mathbb{1}$ on $L_0 = \text{span}\{(2, 0), (1, 1)\}$. On the other hand, we have

$$\left(\zeta[k_0, k_1] \cdot (\zeta[l_0, l_1])^{-1}\right)(2, 0) = \exp\left(2\pi\mathbf{i} \cdot \frac{2}{x_0} \cdot (k_0 - l_0)\right) \quad (9-10a)$$

$$\begin{aligned} &\left(\zeta[k_0, k_1] \cdot (\zeta[l_0, l_1])^{-1}\right)(1, 1) = \\ &\exp\left(2\pi\mathbf{i} \cdot \frac{1}{x_0 \cdot y_1} \cdot \left((k_0 - l_0) \cdot (y_1 - x_1) + (k_1 - l_1) \cdot x_0\right)\right) \quad (9-10b) \end{aligned}$$

Now, $-\frac{1}{2}x_0 < k_0 - l_0 < \frac{1}{2}x_0$, so if (9-10a) is 1 then $k_0 - l_0 = 0$. In this case, (9-10b) reduces to $\exp\left(2\pi\mathbf{i} \cdot \frac{1}{y_1} \cdot (k_1 - l_1)\right)$. Since $-y_1 < k_1 - l_1 < y_1$, if (9-10a) and (9-10b) are both 1, then $k_0 = l_0$ and $k_1 = l_1$, as desired. \square

With these bases and in the obvious notation, we have that

$$\lambda(k_0, k_1) = 2 \cos\left(2\pi \cdot \frac{1}{x_0 \cdot y_1} \cdot \left((u_0 + k_1) \cdot x_0 + (u_1 - k_0) \cdot x_1\right)\right) + 2\mathbf{i} \sin\left(2\pi \cdot \frac{1}{x_0} \cdot (k_0 - u_1)\right),$$

so $\det(K_E) = \prod_{0 \leq k_0 < \frac{1}{2}x_0} \prod_{0 \leq k_1 < y_1} \lambda(k_0, k_1)$. The term $\prod_{j=1}^m (\zeta_{j,0})^{-1}$ in the complex phase ρ also admits a simple formula:

$$\prod_{0 \leq k_0 < \frac{1}{2}x_0} \prod_{0 \leq k_1 < y_1} (\zeta[k_0, k_1]_0)^{-1} = \exp\left(2\pi\mathbf{i} \cdot \frac{y_1}{2} \cdot \left(u_1 + \frac{1}{2} - \frac{x_0}{4}\right)\right)$$

Using this, we may write

$$\rho = \underbrace{\exp\left(2\pi\mathbf{i} \cdot \frac{y_1}{2} \cdot \left(\frac{1}{2} - \frac{x_0}{4}\right)\right)}_{\rho_1} \cdot \underbrace{\exp\left(2\pi\mathbf{i} \cdot u_1 \cdot \left(\frac{y_1}{2} - \left\lfloor \frac{y_1}{2} \right\rfloor\right)\right)}_{\rho_2}$$

Notice $\rho_2 = 1$ whenever y_1 is even. When y_1 is odd, ρ_2 is a square root of $q_1 = \exp(2\pi\mathbf{i} \cdot u_1)$. If we restrict u_1 to lie on the interval $[0, 1)$ — that is, if $\arg(q_1) \in [0, 2\pi)$ —, then ρ_2 is the square root of q_1 in the upper half-plane that is not -1 .

When y_1 is even, $\rho_1 = 1$ except when $x_0 \equiv 0 \pmod{4}$ and $y_1 \equiv 2 \pmod{4}$; in this case $\rho_1 = -1$. In particular, $\rho = \pm 1$ whenever y_1 is even.

When y_1 is odd, there are more cases for ρ_1 :

$$\left\{ \begin{array}{ll} x_0 \equiv 2 \pmod{8} & \implies \rho_1 = +1 \\ x_0 \equiv 6 \pmod{8} & \implies \rho_1 = -1 \\ x_0 \equiv 0 \pmod{8}, y_1 \equiv 1 \pmod{4} & \implies \rho_1 = +\mathbf{i} \\ x_0 \equiv 0 \pmod{8}, y_1 \equiv 3 \pmod{4} & \implies \rho_1 = -\mathbf{i} \\ x_0 \equiv 4 \pmod{8}, y_1 \equiv 1 \pmod{4} & \implies \rho_1 = -\mathbf{i} \\ x_0 \equiv 4 \pmod{8}, y_1 \equiv 3 \pmod{4} & \implies \rho_1 = +\mathbf{i} \end{array} \right.$$

An interesting fact is that when y_1 is even, $\det(K_E)$ is always real, regardless of the values of $q_0, q_1 \in \mathbb{S}^1$. Indeed, it's easy to check that for each

$0 \leq k_1 < \frac{1}{2}y_1$ it holds that

$$\lambda(k_0, k_1) = -\overline{\lambda(k_0, k_1 + \frac{1}{2}y_1)},$$

so in this case $\det(K_E) = \prod_{0 \leq k_0 < \frac{1}{2}x_0} \prod_{0 \leq k_1 < \frac{1}{2}y_1} -|\lambda(k_0, k_1)|^2$, with sign given by $(-1)^{\frac{1}{4}x_0 \cdot y_1}$. In particular, taking ρ into account we conclude $\det(K_D) \leq 0$ whenever $y_1 \equiv 0 \pmod{4}$ and $\det(K_D) \geq 0$ whenever $y_1 \equiv 2 \pmod{4}$.

These formulas allow us to better understand what happens as some uniform scaling dilates the lattice L , and this will be the content of our next result. Observe that for any valid lattice L , the numbers x_0, x_1 and y_1 are uniquely defined, and vice-versa. Let $p_{[L,E]} : \mathbb{R}^2 \rightarrow \mathbb{C}$ be defined by $p_{[L,E]}(u_0, u_1) = \det(K_E(u_0, u_1))$, where K_E is the Kasteleyn matrix for L represented in our choice exponential bases and u_0, u_1 are the arguments of q_0, q_1 as above.

Proposition 9.3.2. *For any valid lattice L , $p_{[L,E]}$ satisfies the following periodicity relations:*

$$\begin{aligned} p_{[L,E]}(u_0 + 1, u_1) &= p_{[L,E]}(u_0, u_1) \\ p_{[L,E]}(u_0, u_1 + 1) &= (-1)^{y_1} \cdot p_{[L,E]}(u_0, u_1) \end{aligned}$$

In particular, we always have $p_{[L,E]}(u_0, u_1 + 2) = p_{[L,E]}(u_0, u_1)$.

Proof. In the obvious notation, we have that $p_{[L,E]}(u_0, u_1) = \det(K_D(u_0, u_1)) \cdot \rho^{-1}$, and it's clear $K_D(u_0 + a, u_1 + b) = K_D(u_0, u_1)$ whenever $a, b \in \mathbb{Z}$. We thus need only study how ρ varies with u_0, u_1 .

Inspecting the formula $\rho = \rho_1 \cdot \rho_2$ above, we see that ρ_1 depends only on L (and not on u_0, u_1) and ρ_2 depends only on u_1 , so the periodicity in u_0 is proved. Now $\rho_2(u_1)$ is always 1 when y_1 is even, so the relation holds in this case. When y_1 is odd, $\rho_2 = \exp(2\pi i \cdot \frac{1}{2}u_1)$ and the relation also is true. \square

Theorem 3. *Let L be a valid lattice. For any positive integer n and reals u_0, u_1*

$$p_{[nL,E]}(n \cdot u_0, n \cdot u_1) = \prod_{0 \leq i, j < n} p_{[L,E]} \left(u_0 + \frac{i}{n}, u_1 - \frac{j}{n} \right).$$

Proof. Let L be generated by $v_0 = (x_0, 0)$ and $v_1 = (x_1, y_1)$, with $x_0, y_1 > 0$ and $0 \leq x_1 < x_0$. Of course, in this case nL is generated by $n \cdot v_0$ and $n \cdot v_1$.

Applying our formulas to nL yields

$$p_{[nL,E]}(n \cdot u_0, n \cdot u_1) = \prod_{\substack{0 \leq k_0 < \frac{1}{2}nx_0 \\ 0 \leq k_1 < ny_1}} \lambda_n(k_0, k_1),$$

where $\lambda_n(k_0, k_1)$ is given by

$$\begin{aligned} & 2 \cos \left(2\pi \cdot \frac{1}{nx_0 \cdot ny_1} \cdot \left((nu_0 + k_1) \cdot nx_0 + (nu_1 - k_0) \cdot nx_1 \right) \right) \\ & + 2i \sin \left(2\pi \cdot \frac{1}{nx_0} \cdot (k_0 - nu_1) \right) \\ = & 2 \cos \left(2\pi \cdot \frac{1}{x_0 \cdot y_1} \cdot \left(\left(u_0 + \frac{k_1}{n} \right) \cdot x_0 + \left(u_1 - \frac{k_0}{n} \right) \cdot x_1 \right) \right) \\ & + 2i \sin \left(2\pi \cdot \frac{1}{x_0} \cdot \left(\frac{k_0}{n} - u_1 \right) \right). \end{aligned}$$

Now, for each $0 \leq k_0 < \frac{1}{2}nx_0$ there are unique integers $0 \leq l_0 < \frac{1}{2}x_0$ and $0 \leq j < n$ with $k_0 = n \cdot l_0 + j$ (division with remainder). In similar fashion, $k_1 = n \cdot l_1 + i$, where $0 \leq l_1 < y_1$ and $0 \leq i < n$ are unique integers. We may then rewrite $\lambda_n(k_0, k_1)$ as $\lambda(l_0, l_1, i, j)$:

$$\begin{aligned} & 2 \cos \left(\frac{2\pi}{x_0 \cdot y_1} \cdot \left(\left[\left(u_0 + \frac{i}{n} \right) + l_1 \right] \cdot x_0 + \left[\left(u_1 - \frac{j}{n} \right) - l_0 \right] \cdot x_1 \right) \right) \\ & + 2i \sin \left(\frac{2\pi}{x_0} \cdot \left[l_0 - \left(u_1 - \frac{j}{n} \right) \right] \right). \end{aligned}$$

It's easy to see the l_0 's and j 's are in bijection with the k_0 's, and similarly for the l_1 's and i 's with the k_1 's. It follows that:

$$p_{[nL,E]}(n \cdot u_0, n \cdot u_1) = \prod_{\substack{0 \leq i, j < n \\ 0 \leq l_0 < \frac{1}{2}x_0 \\ 0 \leq l_1 < y_1}} \lambda(l_0, l_1, i, j).$$

The theorem follows from observing that

$$\prod_{\substack{0 \leq l_0 < \frac{1}{2}x_0 \\ 0 \leq l_1 < y_1}} \lambda(l_0, l_1, i, j) = p_{[L,E]} \left(u_0 + \frac{i}{n}, u_1 - \frac{j}{n} \right).$$

□

Intuitively, Theorem 3 says $\det(K[nL]_E(q_0, q_1))$ can be obtained from determinants of $K[L]_E$ by considering all n -th roots of q_0 and of q_1 . We can make this more precise.

We will say L is odd if y_1 is odd, and L is even if y_1 is even. Let $P_{[L,E]} : \mathbb{S}^1 \times \mathbb{S}^1 \rightarrow \mathbb{C}$ be given by $P_{[L,E]}(q_0, q_1) = p_{[L,E]}(u_0, u_1) \cdot \rho_2(L, u_1)$, where $q_i = \exp(2\pi\mathbf{i} \cdot u_i)$. In other words:

$$P_{[L,E]}(q_0, q_1) = \begin{cases} p_{[L,E]}(u_0, u_1) & \text{if } L \text{ is even;} \\ p_{[L,E]}(u_0, u_1) \cdot \exp(\pi\mathbf{i} \cdot u_1) & \text{if } L \text{ is odd.} \end{cases}$$

Notice that Proposition 9.3.2 ensures $P_{[L,E]}$ is well-defined. Because $p_{[L,E]}(u_0, u_1) = \det(K_D(u_0, u_1)) \cdot \rho^{-1}$, we have that:

$$\begin{aligned} P_{[L,E]}(q_0, q_1) &= p_{[L,E]}(u_0, u_1) \cdot \rho_2(L, u_1) = \det(K_D(u_0, u_1)) \cdot \rho_1(L)^{-1} \\ &= \det(K_D(q_0, q_1)) \cdot \rho_1(L)^{-1} \end{aligned} \quad (9-11)$$

This means that, except for the complex phase $\rho_1(L)^{-1}$ — which does not depend on q_0 or q_1 —, $P_{[L,E]}$ is in fact the initial Laurent polynomial P_K we calculated from our matrix K_D . In other words, the coefficients of $P_{[L,E]}$ are the Fourier coefficients of $p_{[L,E]}$, so that in particular $p_{[L,E]}$ has finitely many nonzero Fourier coefficients. We hope to further consider this point of view in future work.

Of course, $P_{[L,E]}$ also admits its own version of Theorem 3. Observe that $P_{[nL,E]}(q_0^n, q_1^n)$ equals:

$$\begin{aligned} &\rho_2(nL, n \cdot u_1) \cdot p_{[nL,E]}(n \cdot u_0, n \cdot u_1) \\ &= \rho_2(nL, n \cdot u_1) \cdot \prod_{0 \leq i, j < n} p_{[L,E]} \left(u_0 + \frac{i}{n}, u_1 - \frac{j}{n} \right) \\ &= \rho_2(nL, n \cdot u_1) \cdot \prod_{0 \leq i, j < n} P_{[L,E]}(q_0 \cdot \zeta^i, q_1 \cdot \zeta^{-j}) \cdot \rho_2 \left(L, u_1 - \frac{j}{n} \right)^{-1}, \end{aligned}$$

where $\zeta = \exp(2\pi\mathbf{i}\frac{1}{n})$.

When L is even, $\rho_2(nL, n \cdot u_1)$ and the product $\prod_{0 \leq i, j < n} \rho_2(L, u_1 - \frac{j}{n})^{-1}$ are both trivially 1. When L is odd, it holds that

$$\begin{aligned} \prod_{0 \leq i, j < n} \rho_2 \left(L, u_1 - \frac{j}{n} \right)^{-1} &= \prod_{0 \leq i, j < n} \exp(-\pi\mathbf{i} \cdot u_1) \cdot \exp(\pi\mathbf{i} \cdot \frac{j}{n}) \\ &= \exp(-n^2\pi\mathbf{i} \cdot u_1) \cdot \exp(\pi\mathbf{i} \cdot \frac{1}{2}n(n-1)). \end{aligned}$$

If n is even, $\rho_2(nL, n \cdot u_1) = 1$. When both n and L are odd, $\rho_2(nL, n \cdot u_1) = \exp(n\pi \mathbf{i} \cdot u_1)$ and we have that

$$\begin{aligned} & \rho_2(nL, n \cdot u_1) \cdot \prod_{0 \leq i, j < n} \rho_2 \left(L, u_1 - \frac{j}{n} \right)^{-1} \\ &= \exp(-\pi \mathbf{i} \cdot n(n-1) \cdot u_1) \cdot \exp \left(\pi \mathbf{i} \cdot \frac{1}{2} n(n-1) \right) \end{aligned}$$

Observe how in each case the value of the product $\rho_2(nL, n \cdot u_1) \cdot \prod_{0 \leq i, j < n} \rho_2 \left(L, u_1 - \frac{j}{n} \right)^{-1}$ does not depend on the choice of u_1 for which $q_1^n = \exp(2\pi \mathbf{i} \cdot nu_1)$.

We can summarize these findings with

$$P_{[nL, E]}(q_0^n, q_1^n) = \mu_0(q_1, n, L) \cdot \prod_{0 \leq i, j < n} P_{[L, E]}(q_0 \cdot \zeta^i, q_1 \cdot \zeta^{-j}),$$

$$\text{where } \mu_0(q_1, n, L) = \begin{cases} 1 & \text{if } L \text{ is even;} \\ q_1^{-\frac{1}{2}n^2} & \text{if } L \text{ is odd and } n \equiv 0 \pmod{4}; \\ -q_1^{-\frac{1}{2}n^2} & \text{if } L \text{ is odd and } n \equiv 2 \pmod{4}; \\ q_1^{-\frac{1}{2}n(n-1)} & \text{if } L \text{ is odd and } n \equiv 1 \pmod{4}; \\ -q_1^{-\frac{1}{2}n(n-1)} & \text{if } L \text{ is odd and } n \equiv 3 \pmod{4}. \end{cases}$$

A more elegant expression can be given. Using equation (9-11) we write:

$$\begin{aligned} \det \left(K_D[nL](q_0^n, q_1^n) \right) &= \\ & \mu_0(q_1, n, L) \cdot \rho_1(nL) \cdot \rho_1(L)^{-n^2} \cdot \prod_{0 \leq i, j < n} \det \left(K_D[L](q_0 \cdot \zeta^i, q_1 \cdot \zeta^{-j}) \right) \end{aligned}$$

Notice we always have $\rho_1(L)^4 = 1$, so $\rho_1(L)^{-n^2}$ is 1 whenever n is even, and it is $\rho_1(L)^{-1}$ whenever n is odd. Letting $\mu_1(q_1, n, L) = \mu_0(q_1, n, L) \cdot \rho_1(nL) \cdot \rho_1(L)^{-n^2}$, it is not hard to check that

$$\mu_1(q_1, n, L) = \begin{cases} 1 & \text{if } L \text{ is even;} \\ q_1^{-\frac{1}{2}n^2} & \text{if } L \text{ is odd and } n \text{ is even;} \\ q_1^{-\frac{1}{2}n(n-1)} & \text{if } L \text{ is odd and } n \text{ is odd.} \end{cases}$$

Note how in each case the exponent of q_1 is an integer — there is no ambiguity with choosing square roots.

Corollary 9.3.3. *Let μ_1 be as above. Then for each $q_0, q_1 \in \mathbb{S}^1$*

$$\det \left(K_D[nL] (q_0^n, q_1^n) \right) = \mu_1(q_1, n, L) \cdot \prod_{0 \leq i, j < n} \det \left(K_D[L] (q_0 \cdot \zeta^i, q_1 \cdot \zeta^{-j}) \right),$$

where $\zeta = \exp \left(2\pi \mathbf{i} \frac{1}{n} \right)$ is an n -th root of unity.

Glossary

argument (function) A function that can be defined for certain edge-paths.

See page 105.

bipartite A graph whose vertices can be separated into two disjoint sets U, V such that every edge joins a vertex from U to V . See page 19.

black and white (condition) The requirement that a region's dual graph be bipartite. Alternatively, the requirement that every domino on a tiling consist of one square of each color. Implies an equal number of black squares and white squares. See pages 20, 41 and 42.

$\mathcal{B}(L, q)$ The space of complex functions on L_b that are L -periodic with parameter q . See page 126.

brick wall A tiling that uses only one type of domino (vertical or horizontal) and consists entirely of doubly-infinite domino staircases. See page 66.

compatible (with γ) A tiling of a torus is compatible with a quasicycle γ if it contains every other domino of γ ; see page 103.

cross-over domino On a fundamental domain of a torus \mathbb{T}_L , a domino that crosses a side that's been identified with another side. See page 32 .

cross-flip A flip involving a cross-over domino. See page 32.

$C(t_0, t_1)$ The set of cycles formed by the superposition of tilings t_0 and t_1 . See page 98.

cycle, cyle flip Simultaneously representing two tilings of \mathbb{T}_L on the same fundamental domain induces a decomposition of the domain into disjoint domino cycles. A cycle flip uses these cycles to go from one tiling to the other. For details, see page 98.

$D(a, b)$ A set of points in a translate of D_L . See page 118.

D_b, D_w Ordered bases of $\mathcal{B}(L, q)$ and $\mathcal{W}(L, q)$, associated to the construction of a Kasteleyn matrix. See page 127.

- D_L A specific choice of rectangular fundamental domain for a lattice L . See page 86.
- D_n A $2n \times 2n$ square fundamental domain for the $2n \times 2n$ square torus \mathbb{T}_n . See page 32.
- doubly-infinite** A staircase edge-path or domino staircase that is infinite in both directions.
- dual graph** The graph G obtained from a quadriculated region R by substituting each of its squares by a vertex and joining neighboring vertices by an edge; see page 18.
- E_b, E_w Ordered bases of $\mathcal{B}(L, q)$ and $\mathcal{B}(L, q)$, called exponential bases, associated to a diagonalization of K_D . See page 130.
- edge** An edge on a graph, or an edge on the boundary of a square of a quadriculated region.
- edge-profile (type)** The color-induced orientation of edges round a vertex, one of two types. See page 53.
- edge-path** A sequence of neighboring vertices (either on a graph or on a quadriculated region).
- \mathcal{E}, \mathcal{O} The lattices $2\mathbb{Z}^2$ and $2\mathbb{Z}^2 + (1, 1)$ of vectors whose coordinates are respectively both even and both odd.
- even lattice** A valid lattice that does not contain points with odd coordinates. See page 136.
- exterior of a cycle** The intersection of the exterior of domino paths of a closed cycle; see page 99.
- $\mathcal{F}(L)$ The set of all flux values of tilings of \mathbb{T}_L . See Section 5.2.
- flip** A move on a tiling that exchanges two dominoes tiling a 2×2 square by two dominoes in the only other possible configuration.
- flip-connected** A quadriculated region or set of tilings such that any two distinct tilings can be joined by a sequence of flips. For flip-connectedness on simply-connected planar regions, see Corollary 3.1.6 in Section 3.1. For a discussion of flip-connectedness on tori, refer to Chapter 6.
- flip-isolated** A tiling which admits no flips, or a set of tilings in which no two distinct tilings can be joined by a sequence of flips.

$\mathcal{F}(L, q)$ The space of complex functions on $(\mathbb{Z} + \frac{1}{2})^2$ that are L -periodic with parameter q . See page 124.

flux A flux of a tiling counts cross-over dominoes with a sign; see page 33 for an overview, and Chapter 5 for a detailed exposition. For a torus \mathbb{T}_L , the flux may also be thought as an element of the affine lattice $L^\#$; see Section 5.1.

flux-connecting (basis) A basis $\{v_0^*, v_1^*\}$ of L^* such that the moves $\pm v_i^*$ connect $\mathcal{F}(L)$. See page 110.

$G(\mathbb{Z}^2)$ The dual graph of the infinite square lattice. See page 84.

$\mathcal{H}_0(R)$ The set of height functions on R that are 0 at the origin. See page 51.

height function An integer function that encodes a tiling. For a tiling of a quadriculated planar region R , it is defined on the vertices of R ; see Section 3.1. For a tiling of a quadriculated torus \mathbb{R}^2/L , it is L -quasiperiodic and defined on \mathbb{Z}^2 (sometimes called its *toroidal* height function); see Section 4.1.

h_{\max} The height function that is maximal over height functions on \mathbb{Z}^2 that are 0 at the origin. See page 51.

$h_{\max}^{L, \varphi}$ The height function that is maximal over toroidal height functions of \mathbb{T}_L with flux φ . See page 60.

$h_{\max}^{v, \varphi}$ The height function that is maximal over height functions on \mathbb{Z}^2 that take the value $4 \cdot \langle \varphi, v \rangle$ at v . See 60.

h_{\min} The height function that is minimal over height functions on \mathbb{Z}^2 that are 0 at the origin. See page 57.

i The imaginary unit.

interior of a cycle The union of the interior of domino paths of a closed cycle; see page 99.

Kast(e) The Kasteleyn weight of the edge e . For Kasteleyn weights, refer to the construction of Kasteleyn matrices (Section 3.2 and Chapter 7).

Kasteleyn matrix A modified adjacency matrix whose determinant counts tilings of a region. For Kasteleyn matrices of planar regions, see Section 3.2. For an overview of Kasteleyn matrices on the square torus, see page 33; for a detailed exposition on more general tori, refer to Chapter 7.

K, K_D A Kasteleyn matrix. The subindex in K_D refers to the matrix constructed from Kasteleyn weights in a fundamental domain D_L ; see page 127.

K_E A diagonal Kasteleyn matrix, obtained from representing K_D in exponential bases. See page 130.

L A lattice. See page 16.

L^* The dual lattice of L , given by $\text{Hom}(L, \mathbb{Z})$. See page 16 .

$L^\#$ The translate of L^* in $(2L)^*$ that contains all flux values of tilings of \mathbb{T}_L . See Section 5.1.

L_0 The lattice spanned by $\{(2, 0), (1, 1)\}$; the lattice that contains every valid lattice. See page 126.

L_b, L_w The affine lattices $L_0 + (\frac{1}{2}, \frac{1}{2})$ and $L_0 + (\frac{1}{2}, -\frac{1}{2})$, respectively. Equivalently, respectively the set of black vertices and the set of white vertices of $G(\mathbb{Z}^2)$. See page 126.

Laurent polynomial A Laurent series with finitely many nonzero coefficients.

L -flip (round v) A flip round each vertex in $[v]_L$. See page 81.

$l(\gamma)$ The length of an edge-path γ .

L -Kasteleyn signing An assignment of plus and minus signs to L -equivalence classes of edges. See page 84.

L -stairflip (on S) A stairflip on each doubly-infinite domino staircase in $[S]_L$. See page 91.

(perfect) matching A perfect matching is a set of edges on a graph G in which each vertex features exactly once. Corresponds to a tiling. See page 19.

odd lattice A valid lattice that contains points with odd coordinates. See page 136.

$o(p_0p_1)$ The oriented weight of the edge p_0p_1 . See page 119.

ordered sum (representation) A unique representation of edge-paths in Γ or Ψ sets. See 54.

(cycle) parameter A short vector in L , unique up to sign, associated to an open cycle. See page 101.

Pfaffian The determinant of a skew-symmetric matrix A can be written as the square of a polynomial in A 's entries. This polynomial is A 's Pfaffian.

Q_k A component of ∂Q . See pages 89 and 90.

(v -)quasicycle A type of (v -)quasiperiodic function. See page 101.

(arithmetic) quasiperiodicity A function satisfying $f(u + v) = f(u) + C$ for some v, C and all u .

(cardinal) rugged quadrant A type of quadriculated region. See page 71.

rugged rectangle A type of quadriculated region. See page 69.

short A vector v of a lattice L such that $s \cdot v \notin L$ for all $s \in [0, 1)$.

staircase For domino staircases, see page 64. For staircase edge-paths, see page 65. For types of staircases, see pages 65 and 91.

stairflip A move on a tiling that exchanges a doubly-infinite domino staircase by a doubly-infinite domino staircase in the only other possible configuration. See page 90.

Stair(L) The set of L -equivalence classes of doubly-infinite domino staircases in \mathbb{Z}^2 . See page 91.

Stair($L; k$) The set of L -equivalence classes of type- k doubly-infinite domino staircases in \mathbb{Z}^2 . See page 91.

Stair($L; k; \text{vert}$), Stair($L; k; \text{hor}$) The sets of L -equivalence classes of type- k doubly-infinite domino staircases in \mathbb{Z}^2 whose dominoes are all vertical, or all horizontal. See page 91.

t A tiling.

\mathbb{T}_L The torus \mathbb{R}^2 / L , where L is a lattice. See page 43.

\mathbb{T}_n The $2n \times 2n$ square torus \mathbb{T}_n . See pages 32.

Universal Kasteleyn signing A Kasteleyn signing that applies to any valid lattice. See page 84.

valid lattice A lattice whose vectors have integral coordinates that are the same parity. See Section 4.2.

windmill (tiling) A tiling of the infinite square lattice that admits no flips and consists entirely of infinite domino staircases that are never doubly-infinite. See page 65.

$\mathcal{W}(L, q)$ The space of complex functions on L_w that are L -periodic with parameter q . See page 126.

$X(A, B)$ The change of basis from A to B . See page 130.

$f|_A$ The restriction of f to A .

$\Gamma(v, w), \Psi(v, w)$ The set of edge-paths in \mathbb{Z}^2 joining v to w and that respectively respect and reverse color-induced edge orientation. See pages 51 and 57.

φ A flux value.

Φ The mod 4 prescription function on the infinite square lattice \mathbb{Z}^2 . See page 38.

$\phi(\gamma)$ The pseudo-flux of the quasicycle γ ; see page 103.

ψ An isomorphism between a space of L -periodic functions and $\mathcal{F}(L, q)$. See page 124.

A/B The quotient of A by B .

$(\mathbb{Z}^2/L)^*$ The group $\text{Hom}(\mathbb{Z}^2/L, \mathbb{S}^1)$.

ρ A complex phase related to the determinants of K_D and K_E . See page 131.

ρ_1, ρ_2 Complex phases associated to ρ . See page 133.

$A \sqcup B$ The union of two disjoint sets A and B .

$[e]_L, [v]_L, [S]_L$ An L -equivalence class of a vertex v , edge e or doubly-infinite domino staircase S . See pages 81, 84 and 91.

$\|v\|_1$ For $v = (x, y)$, the 1-norm $|x| + |y|$.

$\|v\|_\infty$ For $v = (x, y)$, the infinity norm $\max\{|x|, |y|\}$.

ξ_L The L -stairflip operator. See page 91.

ξ_L - k -**exclusive** A type of subset of $\text{Stair}(L; k)$. See page 92.

Bibliography

- [1] Cédric Boutillier, Béatrice De Tilière, et al. Loop statistics in the toroidal honeycomb dimer model. *The Annals of Probability*, 37(5):1747–1777, 2009.
- [2] Henry Cohn, Noam Elkies, James Propp, et al. Local statistics for random domino tilings of the Aztec diamond. *Duke Mathematical Journal*, 85(1):117–166, 1996.
- [3] John H. Conway and Jeffrey C. Lagarias. Tiling with polyominoes and combinatorial group theory. *Journal of combinatorial theory, Series A*, 53(2):183–208, 1990.
- [4] Noam Elkies, Greg Kuperberg, Michael Larsen, and James Propp. Alternating-sign matrices and domino tilings (Part II). *Journal of Algebraic Combinatorics*, 1(3):219–234, 1992.
- [5] John M. Hammersley. Existence theorems and Monte Carlo methods for the monomer-dimer problem. In F.N. David, editor, *Research Papers in Statistics: Festschrift for J. Neyman*, pages 125–146. Wiley, 1966.
- [6] William Jockusch, James Propp, and Peter Shor. Random domino tilings and the arctic circle theorem. *arXiv:math/9801068*, 1995.
- [7] 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.
- [8] Richard Kenyon, Andrei Okounkov, et al. Planar dimers and Harnack curves. *Duke Mathematical Journal*, 131(3):499–524, 2006.
- [9] Richard Kenyon, Andrei Okounkov, and Scott Sheffield. Dimers and amoebae. *Annals of mathematics*, pages 1019–1056, 2006.
- [10] Percy A. MacMahon. *Combinatory Analysis*. Cambridge University Press, 1915-16 (reprinted by Dover Publications, New York, 2004).

- [11] Pedro H. Milet. Domino tilings of three-dimensional regions. *arXiv preprint arXiv:1503.04617*, 2015.
- [12] Jerome K. Percus. One more technique for the dimer problem. *Journal of Mathematical Physics*, 10(10):1881–1884, 1969.
- [13] James Propp. Enumeration of matchings: problems and progress. *New Perspectives in Geometric Combinatorics* (eds. L. Billera, A. Björner, C. Greene, R. Simeon, and RP Stanley), Cambridge University Press, Cambridge, pages 255–291, 1999 (updated version available at <http://faculty.uml.edu/jpropp/eom.pdf>).
- [14] Dana Randall and Gary Yngve. Random three-dimensional tilings of Aztec octahedra and tetrahedra: an extension of domino tilings. In *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*, pages 636–645. Society for Industrial and Applied Mathematics, 2000.
- [15] Nicolau C. Saldanha and Carlos Tomei. Tilings of quadriculated annuli. *Journal of Combinatorial Theory, Series B*, 88(1):153–183, 2003.
- [16] Nicolau C. Saldanha, Carlos Tomei, Mario A. Casarin Jr., and Domingos Romualdo. Spaces of domino tilings. *Discrete & Computational Geometry*, 14(1):207–233, 1995.
- [17] Harold N. V. Temperley and Michael E. Fisher. Dimer problem in statistical mechanics-an exact result. *Philosophical Magazine*, 6(68):1061–1063, 1961.
- [18] William P. Thurston. Conway’s Tiling Groups. *The American Mathematical Monthly*, 97(8):pp. 757–773, 1990.
- [19] Jinhee Yi. Theta-function identities and the explicit formulas for theta-function and their applications. *Journal of mathematical analysis and applications*, 292(2):381–400, 2004.