Cambridge University Press
978-1-108-84396-6 — Lectures on Random Lozenge Tilings
Vadim Gorin
Excerpt
More Information

1

Lecture 1: Introduction and Tileability

1.1 Preamble

The goal of the lectures is for the reader to understand the mathematics of tilings. The general setup is to take a lattice domain and tile it with elementary blocks. For the most part, we study the special case of tiling a polygonal domain on the triangular grid (of mesh size 1) by three kinds of rhombi that we call "lozenges."

Panel (a) of Figure 1.1 shows an example of a polygonal domain on the triangular grid. Panel (b) of Figure 1.1 shows the lozenges: each of them is obtained by gluing two adjacent lattice triangles. A triangle of the grid is surrounded by three other triangles; attaching one of them, we get one of the three types of lozenges. The lozenges can also be viewed as orthogonal projections onto the $x + y + z = 0$ plane of three sides of a unit cube. Figure 1.2 provides an example of a lozenge tiling of the domain of Figure 1.1.

Figure 1.3 shows a lozenge tiling of a large domain, with the three types of lozenges shown in three different colors. The tiling here is generated uniformly at random over the space of all possible tilings of this domain. More precisely, it is generated by a computer that is assumed to have access to perfectly random bits. It is certainly not clear at this stage how such "perfect sampling" may be done computationally; in fact, we address this issue in the very last lecture. Figure 1.3 is meant to capture a "typical tiling," making sense of what this means is another topic that will be covered in this book. The simulation reveals an interesting feature: there are special regions next to the boundaries of the domain, and in each such region, there is only one (rather than three) type of lozenge. These regions are typically referred to as "frozen regions," and their boundaries are "arctic curves"; their discovery and study have been one of the important driving forces for investigations of the properties of random tilings.

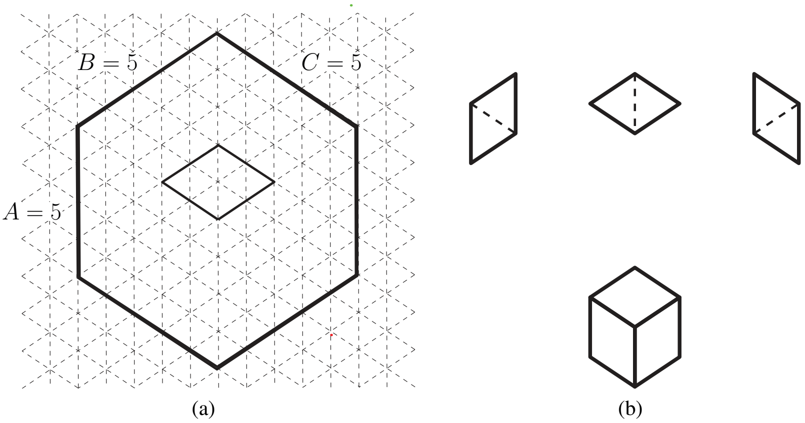

Figure 1.1 Panel (a): A $5 \times 5 \times 5$ hexagon with 2×2 rhombic hole. Panel (b): Three types of lozenges obtained by gluing two adjacent triangles of the grid.

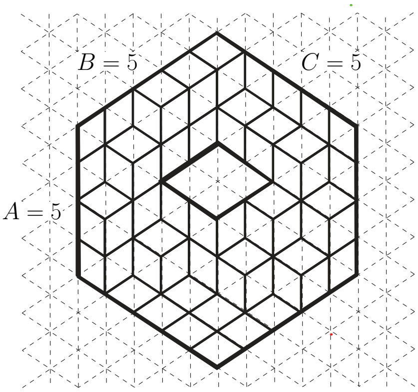

Figure 1.2 A lozenge tiling of a $5 \times 5 \times 5$ hexagon with a hole.

We often identify a tiling with a so-called height function. The idea is to think of a two-dimensional (2D) stepped surface living in a three-dimensional (3D) space and treat tiling as a projection of such surface onto $x + y + z = 0$ plane along the $(1, 1, 1)$ direction. In this way, three lozenges become projections of three elementary squares in 3D space parallel to each of the three coordinate planes. We formally define the height function later in this lecture. We refer to a web page of Borodin and Borodin[1] for a gallery of height functions in a 3D virtual-reality setting.

[1] A. Borodin and M. Borodin, A 3D representation for lozenge tilings of a hexagon, http://math.mit.edu/~borodin/hexagon.html.

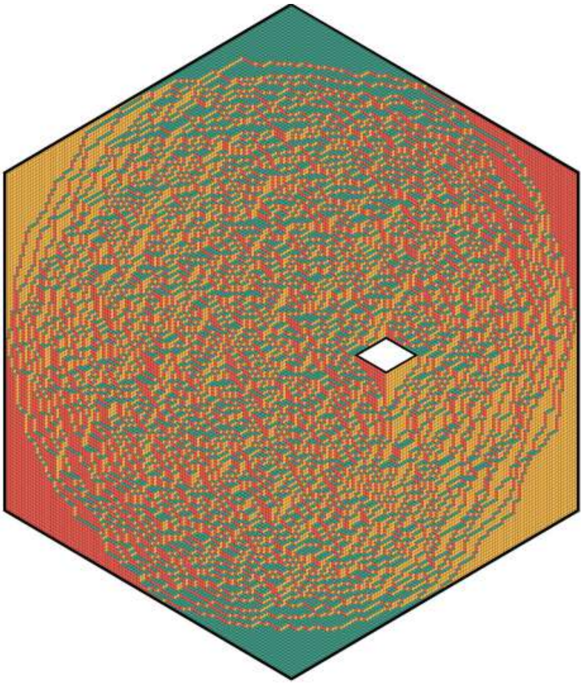

Figure 1.3 A perfect sample of a uniformly random tiling of a large hexagon with a hole. (I thank Leonid Petrov for this simulation.)

1.2 Motivation

Figure 1.3 is beautiful, and for mathematicians, this suffices for probing more deeply into it and trying to explain the various features that we observe in the simulation.

There are also some motivations from theoretical physics and statistical mechanics. Lozenge tilings serve as "toy models" that help in understanding the 3D Ising model (a standard model for magnetism). Configurations of the 3D Ising model are assignments of ⊕ and ⊖ spins to lattice points of a domain in $\mathbb{Z}^3$. A parameter (usually called "temperature") controls how much the adjacent spins are inclined to be oriented in the same direction. The zero-temperature limit leads to spins piling into as large as possible groups of the same orientation; these groups are separated with stepped surfaces whose projections in $(1, 1, 1)$ direction are lozenge tilings. For instance, if we start from the Ising model in a cube and fix boundary conditions to be ⊕ along three faces (sharing a single vertex) of this cube and ⊖ along the other three faces, then we end up with lozenge tilings of a hexagon in the zero-temperature limit.[2]

Another deformation of lozenge tilings is the six-vertex or square-ice model, whose state space consists of configurations of the molecules H_2O on the grid. There are six weights in this model (corresponding to the six ways to match an oxygen with two out of the neighboring four hydrogens), and for particular choices of the weights one discovers weighted bijections with tilings.

We refer the reader to Baxter (2007) for more information about the Ising model and the six-vertex model, further motivations to study them, and approaches to the analysis. In general, both the Ising and six-vertex models are more complicated objects than lozenge tilings, and they are much less understood. From this point of view, the theory of random tilings that we develop in these lectures can be treated as the first step toward the understanding of more complicated models of statistical mechanics.

For yet another motivation, we notice that the 2D stepped surfaces of our study have flat faces (these are frozen regions consisting of lozenges of one type, cf. Figures 1.3 and 1.5) and, thus, are relevant for modeling facets of crystals. One example from everyday life is a corner of a large box of salt. For a particular (nonuniform) random tiling model leading to the shapes reminiscent of such a corner, we refer the reader to Figure 10.1 in Lecture 10.

1.3 Mathematical Questions

We now turn to describing the basic questions that drive the mathematical study of tilings.

1. *Existence of tilings*: Given a domain $\mathcal{R}$ drawn on the triangular grid (and thus consisting of a finite family of triangles), does there exist a tiling of it? For example, a unit-sided hexagon is trivially tileable in two different ways, and the bottom part of Panel (b) in Figure 1.1 shows one of these tilings. On the other hand, if we take the equilateral triangle of side length 3 as our domain $\mathcal{R}$, then it is not tileable. This can be seen directly because the corner lozenges are fixed and immediately cause obstruction. Another way to prove nontileability is by coloring the unit triangles inside $\mathcal{R}$ in white and black colors in an alternating fashion. Each lozenge covers one black and one white triangle, but there is an unequal number of black and white triangles

[2] See Shlosman (2001), Cerf and Kenyon (2001), and Bodineau et al. (2005) for a discussion of the common features in low-temperature and zero-temperature 3D Ising models, as well as the interplay between the topics of this book and more classical statistical mechanics.

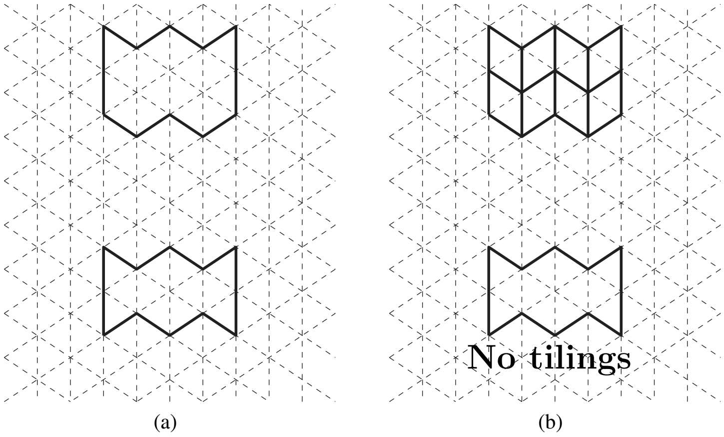

Figure 1.4 Panel (a): The top domain is tileable and the bottom one is not. Panel (b): A possible tiling.

in the $\mathcal{R}$: the equilateral triangle of side length three has six triangles of one color and three triangles of another color.

Another example is shown in Figure 1.4. In panel (a), we see two domains. The top one is tileable, whereas the bottom one is not.

More generally, is there a simple (and efficient, from the computational point of view) criterion for the tileability of $\mathcal{R}$? The answer is yes by a theorem developed by Thurston (1990). We discuss this theorem in more detail in Section 1.4 later in this lecture.

2. *How many tilings does a given domain $\mathcal{R}$ have?* The quality of the answer depends on one's taste and perhaps what one means by a "closed-form"/"explicit" answer. Here is one case where a "good" answer is known by the following theorem due to MacMahon (who studied this problem in the context of plane partitions, conjectured a formula in MacMahon (1896) and proved it in the 1915 book, see Article 495 of MacMahon (1960)). Let $\mathcal{R}$ be a hexagon with side lengths A, B, C, A, B, C in cyclic order. We denote this henceforth by the $A \times B \times C$ hexagon (in particular, panel (a) of Figure 1.1 shows $5 \times 5 \times 5$ hexagon with a rhombic hole).

Theorem 1.1 (MacMahon, 1896) *The number of lozenge tilings of $A \times B \times C$ hexagon equals*

$$\prod_{a=1}^{A} \prod_{b=1}^{B} \prod_{c=1}^{C} \frac{a+b+c-1}{a+b+c-2}. \tag{1.1}$$

As a sanity check, one can take $A = B = C = 1$, yielding the answer 2, and indeed, one readily checks that there are precisely two tilings of $1 \times 1 \times 1$

hexagon. A proof of this theorem is given in Section 2.2. Another situation where the number of tilings is somewhat explicit is for the torus, and we discuss this in Lectures 3 and 4. For general $\mathcal{R}$, one cannot hope for such nice answers, yet certain determinantal formulas (involving large matrices whose size is proportional to the area of the domain) exist, as we discuss in Lecture 2.

3. *Law of large numbers*: Each lozenge tilings is a projection of a 2D surface and therefore can be represented as a graph of a function of two variables, which we call the "height function" (its construction is discussed in more detail in Section 1.4). If we take a *uniformly random* tiling of a given domain, then we obtain a random height function $h(x, y)$ encoding a random stepped surface. What is happening with the random height function of a domain of linear size L as $L \to \infty$? As we will see in Lectures 5–10 and in Lecture 23, the rescaled height function has a deterministic limit

$$\lim_{L\to\infty} \frac{1}{L} h(Lx, Ly) = \hat{h}(x, y).$$

An important question is how to compute and describe the *limit shape* $\hat{h}(x, y)$. One feature of the limit shapes of tilings is the presence of regions where the limiting height function is linear. In terms of random tilings, these are "frozen" regions, which contain only one type of lozenge. In particular, in Figure 1.3, there is a clear outer frozen region near each of the six vertices of the hexagon; another four frozen regions surround the hole in the middle.

Which regions are "liquid," that is, contain all three types of lozenges? What is the shape of the "arctic curve," that is, the boundary between frozen and liquid regions? For example, with the $aL \times bL \times cL$ hexagon setup, one can visually see from Figure 1.5 that the boundary appears to be an inscribed ellipse:

Theorem 1.2 (Baik et al., 2003; Cohn et al., 1998; Gorin, 2008; Petrov, 2014a) *For $aL \times bL \times cL$ hexagon, a uniformly random tiling is with high probability asymptotically frozen outside the inscribed ellipse as $L \to \infty$. In more detail, for each (x, y) outside the ellipse, with probability tending to 1 as $L \to \infty$, all the lozenges that we observe in a finite neighborhood of (xL, yL) are of the same type.*

The inscribed ellipse of Theorem 1.2 is the unique degree 2 curve tangent to the hexagon's sides. This characterization in terms of algebraic curves extends to other polygonal domains, where one picks the degree such that there is a unique algebraic curve tangent (in the interior) of $\mathcal{R}$. Various approaches to Theorem 1.2, its relatives, and generalizations are discussed in Lectures 7, 10, 16, 21, 23.

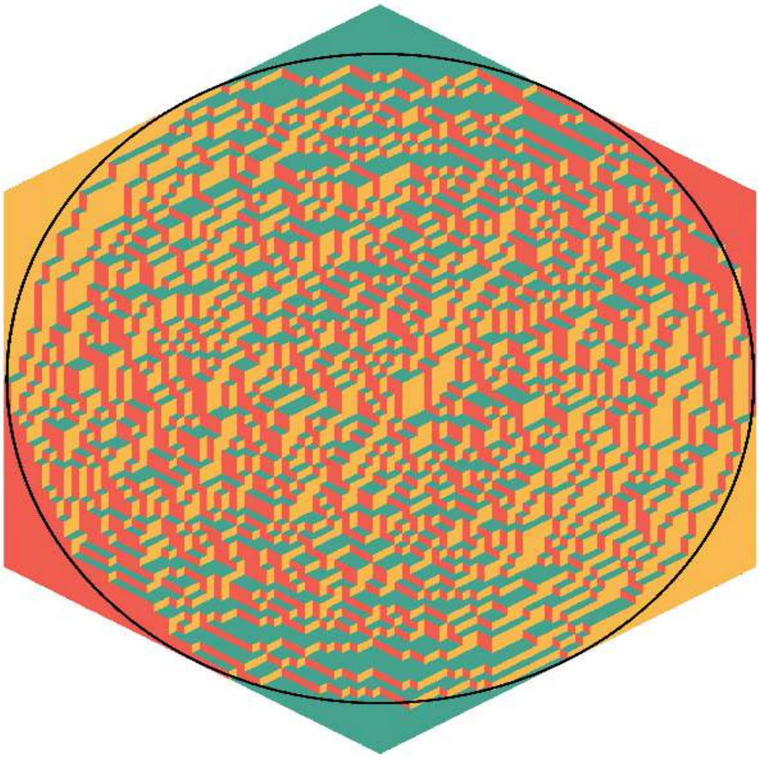

Figure 1.5 Arctic circle of a lozenge tiling.

4. *Analogs of the central limit theorem*: The next goal is to understand the random field of fluctuations of the height function around the asymptotic limit shape, that is, to identify the limit

$$\lim_{L\to\infty} \left(h(Lx, Ly) - \mathbb{E}[h(Lx, Ly)]\right) = \xi(x, y). \tag{1.2}$$

Note the unusual scaling; one may naively expect a need for dividing by $\sqrt{L}$ to account for fluctuations, as in the classical central limit theorem for sums of independent random variables and many similar statements. But there turns out to be some "rigidity" in tilings, and the fluctuations are much smaller. $\xi(x, y)$ denotes the limiting random field; in this case, it can be identified with the so-called "Gaussian free field." The Gaussian free field is related to conformal geometry because it turns out to be invariant under conformal transformations. This topic will be explored in Lectures 11,

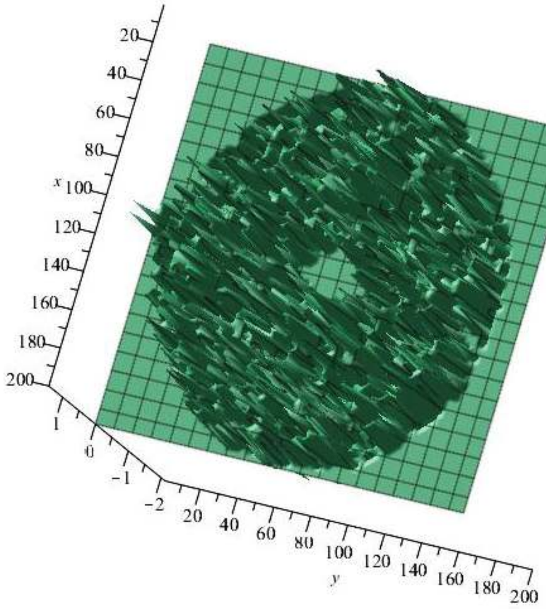

Figure 1.6 Fluctuations of the centered height function for lozenge tilings of a hexagon with a hole. Another drawing of the same system is shown in Figure 24.2 later in the text.

12, 21, and 23. For now, we confine ourselves to yet another picture, given in Figure 1.6.

5. *Height of a plateau/hole*: Consider Figures 1.2 and 1.3. The central hole has an integer-valued random height. What is the limiting distribution of this height? Note that comparing with (1.2), we expect that no normalization is necessary as $L \to \infty$, and therefore the distribution remains discrete as $L \to \infty$. Hence, the limit cannot be Gaussian. You can make a guess now or proceed to Lecture 24 for the detailed discussion.

6. *Local limit*: Suppose we "zoom in" at a particular location inside a random tiling of a huge domain. What are the characteristics of the tiling there? For example, consider a certain finite pattern of lozenges; call it $\mathcal{P}$, and see Figure 1.7 for an example. Asymptotically, what is the probability that $\mathcal{P}$ appears in the vicinity of (Lx, Ly)? Note that if $\mathcal{P}$ consists of a single lozenge, then we are just counting the local proportions for the lozenges of three types; hence, one can expect that they are reconstructed from the gradients of the limit shape $\tilde{h}$. However, for more general $\mathcal{P}$, it is not clear

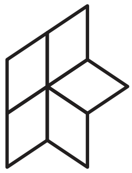

Figure 1.7 An example of a local pattern $\mathcal{P}$ of lozenges. The bulk-limit question asks about the probability of observing such (or any other) pattern in a vicinity of a given point (Lx, Ly) in a random tiling of a domain of linear scale $L \to \infty$.

what to expect. This is called a "bulk-limit" problem, and we return to it in Lectures 16 and 17.

7. *Edge limit*: How does the arctic curve (border of the frozen region) fluctuate? What is the correct scaling? It turns out to be $L^{\frac{1}{3}}$ here, something that is certainly not obvious at all right now. The asymptotic law of rescaled fluctuations turns out to be given by the celebrated Tracy–Widom distribution from random matrix theory, as we discuss in Lectures 18 and 19.
8. *Sampling*: How does one sample from the uniform distribution over tilings? The number of tilings grows extremely fast (see, e.g., the MacMahon formula (1.1)), so one can not simply exhaustively enumerate the tilings on a computer, and a smarter procedure is needed. We discuss several approaches to sampling in Lecture 25.
9. *Open problem*: Can we extend the theory to 3D tiles?

1.4 Thurston's Theorem on Tileability

We begin our study from the first question: Given a domain $\mathcal{R}$, is there at least one tiling? The material here is essentially based on Thurston (1990).

Without loss of generality, we may assume $\mathcal{R}$ is a connected domain; the question of tileability of a domain is equivalent to that of its connected components. We start by assuming that $\mathcal{R}$ is simply connected, and then remove this restriction.

We first discuss the notion of a height function in more detail, and how it relates to the question of the tileability of a domain. There are six directions on the triangular grid, and the unit vectors in those directions are as follows:

$$a = (0, 1), \quad b = \left(-\frac{\sqrt{3}}{2}, -\frac{1}{2}\right), \quad c = \left(\frac{\sqrt{3}}{2}, -\frac{1}{2}\right), \quad -a, \quad -b, \quad -c.$$

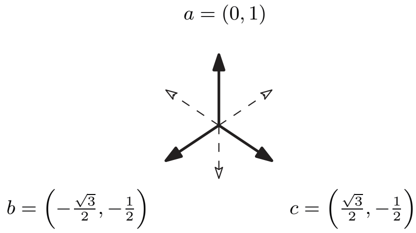

Figure 1.8 Out of the six lattice directions, three are chosen to be positive (in bold).

We call a, b, c positive directions, and their negations are negative directions, as in Figure 1.8.

We now define an *asymmetric* nonnegative distance function $d(u, v)$ for any two vertices u, v on the triangular grid (which is the lattice spanned by a and b) and a domain $\mathcal{R}$. $d(u, v)$ is the minimal number of edges in a positively oriented path from u to v staying within (or on the boundary of) $\mathcal{R}$. This is well defined as we assumed $\mathcal{R}$ to be connected. The asymmetry is clear: Consider $\mathcal{R}$ consisting of a single triangle, and let u, v be two vertices of it. Then $d(u, v) = 1$, and $d(v, u) = 2$ (or vice versa).

We now formally define a height function $h(v)$ for each vertex $v \in \mathcal{R}$, given a tiling of $\mathcal{R}$. This is given by a local rule: if $u \to v$ is a positive direction, then

$$h(v) - h(u) = \begin{cases} 1, & \text{if we follow an edge of a lozenge,} \\ -2, & \text{if we cross a lozenge diagonally.} \end{cases} \tag{1.3}$$

It may be easily checked that this height function is defined consistently. This is because the rules are consistent for a single lozenge (take, e.g., the lozenge $\{0, a, b, a + b\}$), and the definition extends consistently across unions. Note that h is determined up to a constant shift. We may assume without loss of generality that our favorite vertex v_0 has $h(v_0) = 0$.

Let us check that our definition matches the intuitive notion of the height. For that, we treat the positive directions a, b, c in Figure 1.8 as projections of coordinate axes Ox, Oy, Oz, respectively. Take one of the lozenges, say $\{0, a, b, a+b\}$. Up to rescaling by the factor $\sqrt{2/3}$, it can be treated as a projection of the square $\{(0, 0, 0), (1, 0, 0), (0, 1, 0), (1, 1, 0)\}$ onto the plane $x + y + z = 0$. Hence, our locally defined height function becomes the value of $x + y + z$. A similar observation is valid for two other types of lozenges. The conclusion is that if we identify a lozenge tiling with a stepped surface in 3D space, then our height is the (signed and rescaled) distance from the surface point to its projection onto the $x + y + z = \text{const}$ plane in the $(1, 1, 1)$ direction.