Markov Chain Algorithms for Planar Lattices

In this lecture we begin studying the problem of generating the following structures on planar lattices:

— domino tilings of Cartesian lattices,

— lozenge tilings of triangular lattices, and

— 3-colorings of Cartesian lattices.

First, we establish a 1-1 correspondence between these lattice structures and *routings* defined on appropriately chosen lattices. Markov Chains for the previous problems translate naturally to Markov Chains on the associated set of routings; by applying coupling arguments in the routing world we will be able to prove rapid convergence for a number of algorithms.

Domino Tilings and Routings

A domino tiling is a covering of a finite region of the Cartesian lattice with dominoes, where each domino covers two adjacent squares of the region.

The associated routings are defined in this case on a triangular lattice. To specify this lattice, first color the squares of the Cartesian lattice black and white as on a chessboard. The vertices of the triangular lattice are the centers of vertical edges that have a black square to their right. Finally, connect a vertex (x, y) to $(x+1, y+1)$, $(x+1, y-1)$ and $(x+2, y)$ to complete the triangular lattice. Sources are boundary vertices with the interior of S to their right, and sinks those with the interior to their left. Once again, we pair up sources $\{s_1, \ldots, s_k\}$ and sinks $\{t_1, \ldots, t_k\}$ in the obvious way. A *domino routing* of S is then a collection of non-intersecting shortest paths on the triangular lattice within S from s_i to t_i for each i.

The correspondence between domino tilings and domino routings is bijective; this correspondence is shown in Figure 1 for a rectangular region of the Cartesian lattice. (Note that the construction can be applied to any Cartesian lattice that can be tiled with dominoes, not only to rectangular ones.)

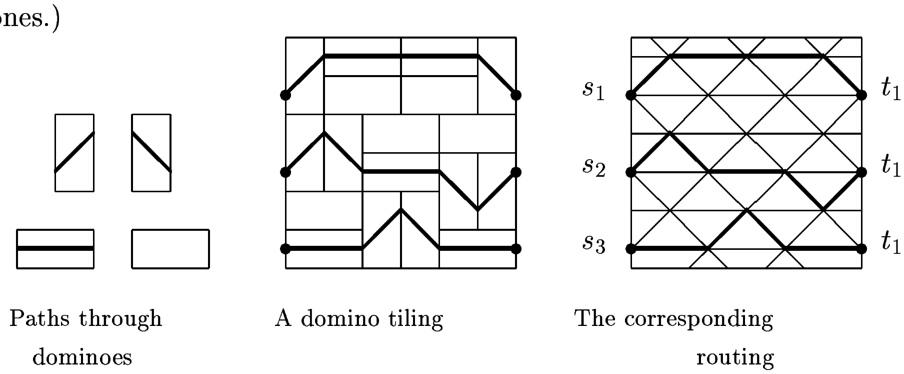

Figure 1: Domino tilings and routings

Lozenge Tilings and Routings

A *lozenge* is the analogue of a domino in the Cartesian lattice: each lozenge covers two adjacent triangles in the triangular lattice, and has three possible orientations. Again, for a finite region of the triangular lattice, we are interested in the problem of generating random lozenge tilings.

The routings corresponding to lozenge tilings are defined on a Cartesian lattice rotated by 45°. Formally, the vertices of this lattice are the midpoints of the vertical edges of the triangular lattice; two vertices are connected by an edge if they lie on adjacent triangles. For a finite, simply-connected region S of the triangular lattice, we designate each vertex of the Cartesian lattice that lies on the boundary of S as either a source or a sink: a vertex v is a *source* if the interior of S is on the right of v, and a *sink* otherwise. If a lozenge tiling of S exists, the numbers of sources and sinks are necessarily equal. Following the boundary in counterclockwise order, starting at a source, we label the sources $\{s_1, \ldots, s_k\}$ and each sink t_i, where s_i was the last unmatched source we labeled.

A *lozenge routing* of S is a set of k non-intersecting shortest paths on the Cartesian lattice within S from s_i to t_i for each i. Note that the length of the path from s_i to t_i is the same in every routing (see Figure 2). Again, the correspondence between lozenge tilings and lozenge routings is bijective.

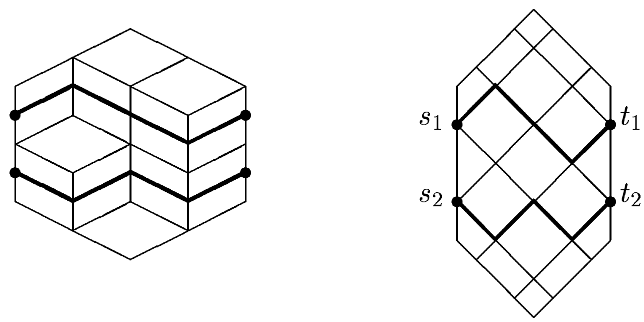

A lozenge tiling The corresponding routing

Figure 2: Lozenge tilings and routings

There is a natural 1-1 correspondence between lozenge tilings and yet another set of combinatorial objects: surfaces determined by partial fillings with unit cubes of a given box (the box is a cube in Figure 2). The disjoint paths in a lozenge routing can be interpreted as the level curves for this surface.

3-colorings and Routings

The correspondence between 3-colorings[1] of a Cartesian lattice and routings can be obtained by using as an intermediate step either BCSOS (body-centered, solid-on-solid) surfaces or Eulerian orientation.

Consider height functions defined over the cells of the lattice. Two functions are said to be equivalent if they differ at all cells by the same constant. *BCSOS surfaces* are the equivalence classes of height functions for which adjacent cells differ by ± 1 (for our purposes, two cells are adjacent if they share an edge). Given a 3-coloring of the lattice, we can obtain a BCSOS surface by "unwinding" the coloring. Indeed, since the colors of two adjacent cells differ by $\pm 1 (\mathrm{mod} 3)$, we obtain valid cell heights by fixing the height of one arbitrary cell to 0 and propagating the ± 1 differences to all other cells. Conversely, from a BCSOS surface we get a 3-coloring by coloring each cell with its height (*mod* 3). Hence, the correspondence between 3-colorings and BCSOS surfaces is 1-1.

For a given BCSOS surface consider the surface tilted by 45° around the diagonal of an arbitrary cell (the new surface is obtained by proportionally increasing/decreasing cell heights, see Figure 3).

[1]For convenience, we color the centers of lattice cells.

On the tilted surface, the boundaries between regions of different heights define a (possibly empty) set of non-crossing edge-disjoint paths (note that the paths can touch each other); this is the desired routing. It can be proved that the correspondence between BCSOS surfaces and routings is again 1-1.

2	0	1	2
1	2	0	1
2	0	1	2
0	1	0	1

3-coloring

-1	0	1	2
-2	-1	0	1
-1	0	1	2
0	1	0	1

BCSOS surface

-7	-5	-3	-1
-7	-5	-3	-1
-5	-3	-1	1
-3	-1	-1	1

Tilted surface and routings

Figure 3: BCSOS surfaces and tilted BCSOS surfaces

An *Eulerian orientation* of a finite region of the Cartesian lattice is an orientation of the edges incident to interior vertices (vertices of degree 4) such that each interior vertex has two incoming and two outcoming edges. Edges incident to only one interior vertex are called boundary edges (edges adjacent to no interior vertex are ignored). The rules represented in Figure 4 define a 1-1 mapping between 3-colorings and Eulerian orientations (3-colorings that are obtained by cyclic permutations are considered identical).

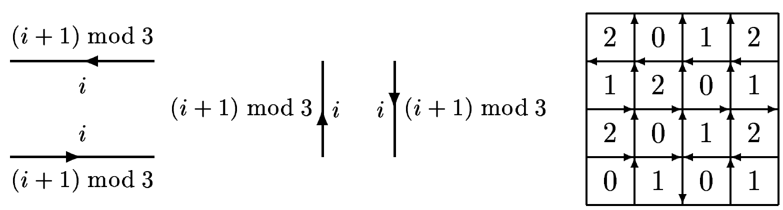

Figure 4: 3-coloring and associated Eulerian orientation

Eulerian routings are defined as follows. For a given orientation of the boundary, the *sources* of the orientation are boundary edges pointing toward an interior vertex from the left or from below, and the *sinks* are those pointing out from an interior vertex to the right or up. We will restrict ourselves to orientations of the boundary in which the number of edges pointing toward interior vertices is the same with the number of edges pointing out from them; this implies that the number of sources is equal with the number of sinks (this number can be 0 for some orientations). An *Eulerian routing* is a set of edge-disjoint directed shortest paths from sources to sinks.

For a given Eulerian orientation, define the sources and the sinks as above. We can obtain an Eulerian routing from the Eulerian orientation with the following algorithm. The paths are constructed one by one in increasing order of the source index. Initially, all edges are unmarked, and each path is started from its source. If there is an unmarked edge pointing up from the current vertex on the path, we mark it and move along. If not, we do the same for the edge pointing to the right (see Figure 5).

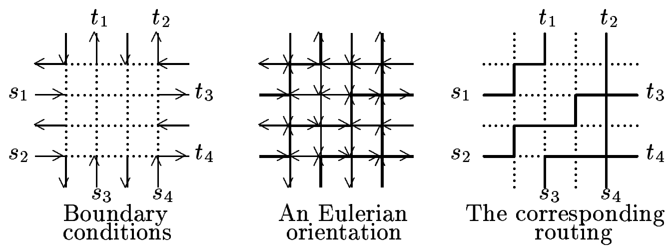

Figure 5: Eulerian orientations and routings

The above algorithm defines a mapping from Eulerian orientations to Eulerian routings. The Eulerian routings obtained by this algorithm have the property that consecutive paths do not cross each other (except the common points, one is situated entirely above or entirely below the other). Conversely, given an Eulerian routing that satisfies this path monotony property, we orient boundary and path edges according with the routing. Then, this partial orientation extends in a unique way to an Eulerian orientation, and the initial routing is the Eulerian routing associated with the orientation. It follows that the correspondence between Eulerian orientations and monotone Eulerian routings is 1-1.

Thurston's tiling group

Suppose that R is a Cartesian lattice bounded by an arbitrary polygon. When can R be tiled with 2×1 dominoes? The question can be answered in $O(n^2)$ using bipartite matching (here n is half the number of cells in the region). Thurston has developed a linear time algorithm using the so called *tiling group*.

We begin with a simple example showing how groups can be used to characterize geometrical properties. The free group $F(a,b)$ on $\{a,b\}$ in which a and b commute with each other is denoted by $< a,b|ab = ba >$. We can represent a path in the Cartesian lattice starting at a given point by an element of $F(a,b)$ as follows: a (a^{-1}) is identified with the traversal of a horizontal edge from left to right (right to left), b (b^{-1}) represents a vertical edge traversed from bottom to top (top to bottom, respectively). The resulting element is the product (in path order) of all group elements associated with path edges. Clearly, the group element corresponding to a path is the identity if and only if the path is closed.

The tiling group is the free group on $\{a,b\}$ satisfying the relations $a^2b = ba^2$ and $ab^2 = b^2a$ (these relations come from identifying the words formed by following the boundary of a horizontal or vertical tile with the identity). The simple characterization that worked for closed paths does not hold for tiling with dominoes: the lattice in Figure 6 has no domino tiling, but the element of $< a,b|\ a^2b = ba^2,\ ab^2 = b^2a >$ associated with its border is the identity. Fortunately, this can be dealt with.

Note that a^2 commutes with every element of $< a,b|\ a^2b = ba^2,\ ab^2 = b^2a >$, and the same is true about b^2. Fixing P_0 on the boundary of R, let α_P be the group element associated with a path starting at P_0 and ending at P. The *canonical representation* of α_P is obtained by using the commutativity of a^2 and b^2 to move them in the rightmost position. That is, the canonical form is composed of an alternating sequence followed by the product $a^{2k}b^{2l}$, for appropriate k and l (where the alternating sequence is of the form *abab...b*, *abab...ba*, *bab...ba*, or *bab...b*). The canonical representation can be reconstructed if we know the following quantities: X_P, the total number of

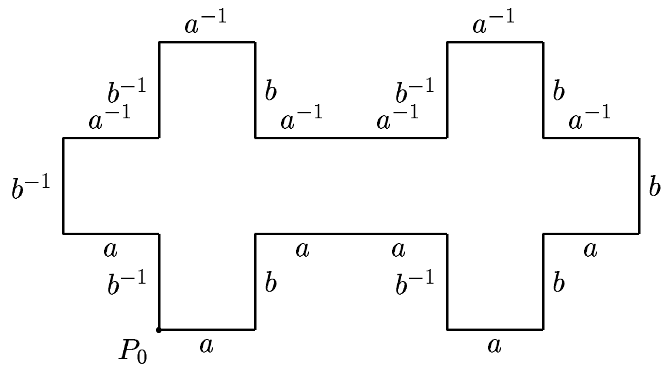

Figure 6: A lattice that cannot be tiled with dominoes

a's, Y_P, the total number of b's, and Z_P, the length of the alternating sequence, with the convention that $Z_P > 0$ for alternating sequences starting with a, and $Z_P < 0$ otherwise.

The quantity Z_P, that can be interpreted as a "height", provides the key to the linear algorithm. In the first stage we compute the values Z_P for all border vertices. The heights associated with a domino can only be in one of the two configurations shown in Figure 7 (this follows from the properties of the tiling group). So, we identify a border vertex of maximum height. If the region can be covered by dominoes, this vertex can only be the middle vertex of a domino. If a tile in this position can be placed over R, we reduce the problem to a smaller region, if not, the initial region cannot be covered with dominoes.

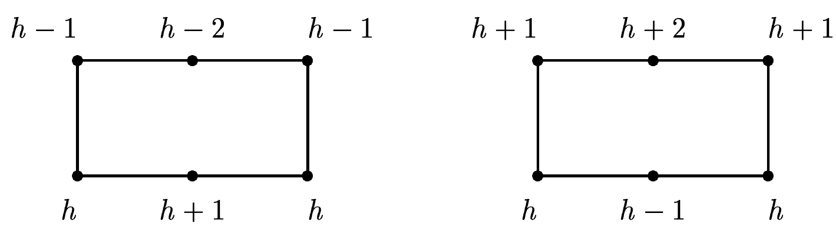

Figure 7: Heights associated with a tile