COUNTING DOMINO TILINGS

JOSHUA HOROWITZ

Several classic math puzzles exist in the template "Tiling X with a bunch of Ys: can you or can't you?" (E.g., the mutilated chessboard problem, Golomb's delightful theorem on trominoes.) But a point arises in one's life when "can you or can't you?" isn't enough. You need to know "how many ways?".

Take the simplest example: How many ways can you tile an $m \times n$ chessboard with 2×1 dominoes? If mn is odd, the answer is zero. If m or n is small, some fun answers like "1" or "a Fibonacci number" arise. But we should be able to come up with a general formula, right?

Turns out we can. There are precisely

$$N_{m,n} = \prod_{j=1}^{m} \prod_{k=1}^{n} \sqrt[4]{4\cos^2 \frac{\pi j}{m+1} + 4\cos^2 \frac{\pi k}{n+1}}$$

ways to tile the chessboard. This bewildering formula actually arises from fairly simply linear-algebraic techniques and showcases the remarkable power of spectral graph theory. This guide (based on Richard Kenyon's lecture notes[1]) will take you on the most straightforward path to this result.

1. From combinatorics to linear algebra

The squares of an $m \times n$ chessboard can be colored white and black in such a way that every pair of adjacent squares consists of one white square and one black square[2]. Mathematically, we say the grid graph $G_{m,n}$ is 2-colorable, or bipartite. This simplifies the structure of domino tilings – every domino goes between a white square and a black square, and the set of all dominoes in a valid coloring thus gives us a complete matching between the set of white squares and the set of black squares. Giving these two sets an arbitrary ordering, labeling them as w_i and b_i for $i = 1 \ldots \frac{mn}{2}$, this matching can be represented by a permutation σ of $\left\{1, \ldots, \frac{mn}{2}\right\}$. A given σ says that w_i is matched by a domino with $b_{\sigma(i)}$. Keep in mind, however, that not every permutation represents a valid domino tiling – some will match together black and white squares which are not adjacent. We will say that a permutation σ is "ok" if $w_i \sim b_{\sigma(i)}$ for all i (letting $\sim$ denote the white/black adjacency relation). Counting "ok" permutations is the same as counting domino tilings.

So our goal is to determine

$$N_{m,n} = \#\left\{\text{ok permutations } \sigma\right\}.$$

Date: April 10, 2013 (updated April 16, 2013).

I thank the MLS Lecture Series for hosting the talk from which these notes were composed.

[1] http://www.math.brown.edu/~rkenyon/papers/dimerlecturenotes.pdf

[2] In fact, it is difficult to find a chessboard that is not colored in this manner.

I will rewrite this in the form

$$N_{m,n} = \sum_{\sigma} \begin{cases} 1 & \text{if } \sigma \text{ is ok} \\ 0 & \text{if } \sigma \text{ is not ok} \end{cases}.$$

It's not clear what good this reformulation does you, until you see the sum over permutations and begin to hazily recollect a place where you've seen such a sum before. Searching the darker recesses of your mind, you come across a glint of light: "The determinant!" Yes! Recall that

$$\det M = \sum_{\sigma} (-1)^{\sigma} M_{1,\sigma(1)} \cdots M_{n,\sigma(n)}.$$

If we could equate the summand of $N_{m,n}$ to the summand of $\det M$ for some matrix M, we would have an expression for $N_{m,n}$ as a determinant. Since linear algebra is the most unreasonably effective field of mathematics, we'd have good reason to be excited about this development.

But what M will make this work? Forgetting for a moment about the signature $(-1)^{\sigma}$, the equality we need is

$$M_{1,\sigma(1)} \cdots M_{n,\sigma(n)} = \begin{cases} 1 & \text{if } \sigma \text{ is ok} \\ 0 & \text{if } \sigma \text{ is not ok} \end{cases}.$$

Well, σ is only ok if $w_i \sim b_{\sigma(i)}$ for all i, that is, if $\prod_i \begin{cases} 1 & w_i \sim b_{\sigma(i)} \\ 0 & w_i \not\sim b_{\sigma(i)} \end{cases} = 1$. So, if we set $M_{i,j} = \begin{cases} 1 & w_i \sim b_j \\ 0 & w_i \not\sim b_j \end{cases}$, the above equality will be satisfied! This is simply the adjacency matrix, or rather the "biadjacency" matrix which gives adjacencies between the two different colors of vertex. A permutation contributes to the determinant of the biadjacency matrix iff every matrix element it picks up is nonzero, that is, iff it represents an ok permutation!

This is really lovely. Let's give it a try. We expect $N_{2,2} = 2$. The 2×2 chessboard has two white squares and two black squares. Each of the white squares is connected to each of the black squares, and visa versa. So

$$M = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}.$$

And what's $\det M$? Well... 0. This is troubling. Something went wrong. Our two tilings correspond to the two possible permutations:

$$\sigma_1 : \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \qquad \sigma_2 : \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}.$$

As hoped for, the σ_1 contributes 1 to the determinant. Alas, σ_2 contributes -1, canceling out the contribution of σ_1. The signature $(-1)^{\sigma}$, which we tried to forget about, has come back to bite us.

What is to be done? One option would be to consider the *permanent* of M, a function defined identically to the determinant but leaving out the $(-1)^{\sigma}$. The problem with this approach is that the permanent is useless – it lacks all of the lovely properties the determinant has which made us so excited to use it. So we will try a second option: modify the elements of the matrix M to exactly cancel out the effects of $(-1)^{\sigma}$.

It's not immediately clear how to do this. We still want a permutation to contribute zero to the determinant if it isn't ok. So we will leave the 0 elements of M alone. All we have left are the

1 elements, which represent weights on the edges of $G_{m,n}$. We need the product of these weights to flip sign as we transpose elements in σ. Take the 2×2 case above. Suppose σ_1 represents two vertical dominoes placed next to one another. When we move to the tiling consisting of two horizontal dominoes (represented by σ_2) we swap the black-square partners of the two white squares, adding a factor of -1 to the signature. To cancel this out, the horizontal edges need weights multiplying to -1. The most natural, symmetric option is to give them each weight i. Now the matrix of consideration looks like

$$K = \begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix},$$

and $\det K$ is 2, as desired!

We can generalize this to the $m \times n$: Assign every vertical edge weight 1 and every horizontal edge weight i, to get a matrix $K_{m,n}$ (the *Kasteleyn* matrix). We know each of these little 2×2 swaps will preserve the contribution of a tiling to the determinant of K. There's a more general argument that in *every* transition between two different tilings, the change in signature will cancel out the change in weighting, and every tiling will contribute to $\det K_{m,n}$ with the same sign. Hence, $N_{m,n} = |\det K_{m,n}|$.

2. Linear algebra

OK, so we know that we are looking for $|\det K_{m,n}|$, where $K_{m,n}$ is the white-black adjacency matrix for $G_{m,n}$ with vertical edges weighted 1, horizontal edges weighted i. We face a linear-algebraic calculation with our comprehensive suite of linear-algebraic tools. Let's get to it!

First of all, we're going to look at $A_{m,n} = \begin{bmatrix} 0 & K_{m,n} \\ K_{m,n}^t & 0 \end{bmatrix}$. Turns out $\det A_{m,n} = (\det K_{m,n})^2$ and $A_{m,n}$ is our good friend the (full) adjacency matrix! (With, admittedly, some pretty weird edge weightings.) Our graph $G_{m,n}$ is the *cartesian product* of a vertical path (of length m) and a horizontal path (of length n) so we can write its adjacency matrix as $A_{m,n} = A_m \otimes 1_n + i 1_m \otimes A_n$, where A_i is the adjacency matrix of the path of length i.

To compute the determinant of this matrix, we'll look at its eigenvalues/eigenvectors. It turns out we can get $A_{m,n}$'s eigenstuff from A_m's and A_n's eigenstuff. If $A_m \vec{v}_m = \lambda_m \vec{v}_m$ and $A_n \vec{v}_n = \lambda_n \vec{v}_n$, then

$$\begin{aligned} A_{m,n} \left(\vec{v}_m \otimes \vec{v}_n \right) &= A_m \vec{v}_m \otimes 1_n \vec{v}_n + i 1_m \vec{v}_m \otimes A_n \vec{v}_n \\ &= \left(\lambda_m + i \lambda_n \right) \left(\vec{v}_m \otimes \vec{v}_n \right), \end{aligned}$$

so the eigenvalues of $A_{m,n}$ are these particular complex combinations of eigenvalues of A_m and A_n.

2.1. Eigenvalues of A_m.

What are the eigenvalues of a path adjacency matrix A_m?

$$A_m = \begin{bmatrix} 0 & 1 & & & \\ 1 & 0 & 1 & & \\ & 1 & 0 & \ddots & \\ & & \ddots & \ddots & 1 \\ & & & 1 & 0 \end{bmatrix}.$$

This matrix is *almost* nice to work with, but it has those two nasty ends. If we instead considered the adjacency matrix A_∞ of the bidirectionally infinite path graph, the math might be easier.

$$A_\infty = \begin{bmatrix} \ddots & \ddots & & \\ \ddots & 0 & 1 & \\ & 1 & 0 & \ddots \\ & & \ddots & \ddots \end{bmatrix}.$$

If we found some A_∞-eigenvector $\vec{v}$ with $v_0 = v_{m+1} = 0$, the "subvector" $\vec{v}' = \{v_j\}_{j=1\dots m}$ would form an eigenvector for A_m!

The matrix A_∞ can be written as the sum of the left-shift matrix S and its inverse, the right-shift matrix S^{-1}.

$$S = \begin{bmatrix} \ddots & \ddots & & \\ \ddots & 0 & 1 & \\ & 0 & 0 & \ddots \\ & & \ddots & \ddots \end{bmatrix}, \quad S^{-1} = \begin{bmatrix} \ddots & \ddots & & \\ \ddots & 0 & 0 & \\ & 1 & 0 & \ddots \\ & & \ddots & \ddots \end{bmatrix}.$$

An eigenvector $\vec{v}$ of S with eigenvector μ must satisfy $v_{j+1} = \mu v_j$ for all j, so we can write it normalized as $\vec{w}_\mu = (\dots, \mu^{-2}, \mu^{-1}, \underline{1}, \mu, \mu^2, \dots)$ (with v_0 underlined, to clarify alignment). This $\vec{w}_\mu$ also serves as an eigenvector for S^{-1}, but with eigenvalue μ^{-1}. So, since $A_\infty = S + S^{-1}$, $A_\infty \vec{w}_\mu = \left(\mu + \mu^{-1}\right) \vec{w}_\mu$.

So far, we don't have eigenvectors for A_∞ with $v_0 = v_{m+1} = 0$. But the eigenvectors we *do* have come in degenerate pairs: $\vec{w}_\mu$ and $\vec{w}_{\mu^{-1}}$ both have eigenvalue $\lambda = \mu + \mu^{-1}$. So we can form linear combinations of these and find new eigenvectors with this same eigenvalue. To get $v_0 = 0$, we just subtract: $\vec{v} = \vec{w}_\mu - \vec{w}_{\mu^{-1}} = (\dots, \mu^{-2} - \mu^2, \mu^{-1} - \mu, \underline{0}, \mu - \mu^{-1}, \mu^2 - \mu^{-2}, \dots)$. To get $v_{m+1} = 0$, we need $\mu^{m+1} - \mu^{-m-1} = 0$, that is, $\mu^{2m+2} = 0$. Just set μ to $\mu_j = e^{i\pi j/(m+1)}$, one of the $(2m+2)^{\text{th}}$ roots of unity! Then, the eigenvector will become $\vec{v}_j = (\dots, 2\sin\frac{-2\pi j}{m+1}, 2\sin\frac{-\pi j}{m+1}, \underline{0}, 2\sin\frac{\pi j}{m+1}, 2\sin\frac{2\pi j}{m+1}, \dots)$, with eigenvalue $\lambda_j = 2\cos\frac{\pi j}{m+1}$. Since this vanishes in the right places, we find an eigenvector for A_m: $\vec{v}'_j = \left(2\sin\frac{\pi j}{m+1}, 2\sin\frac{2\pi j}{m+1}, \dots, 2\sin\frac{m\pi j}{m+1}\right)$, with the same λ_js. Note, however, that $\vec{v}_j = -\vec{v}_{(2m+2)-j}$, so to get a complete set of independent eigenvectors, we should have j range from 1 to m. We have our answer: the eigenvalues of A_m are $\lambda_j = 2\cos\frac{\pi j}{m+1}$, for $j = 1 \dots m$.

[Note: If you don't like using infinitely large graphs / matrices, you can use roughly the same argument on the cycle graph C_{2m+2}. It's just harder to motivate that particular choice of graph.]

2.2. **Bringing it together.** By our discussion of $A_{m,n}$ above, we know the eigenvalues of $A_{m,n}$ must be $2\cos\frac{\pi j}{m+1} + 2i\cos\frac{\pi k}{n+1}$, for $j = 1 \ldots m$, $k = 1 \ldots n$. Hence,

$$\begin{aligned}
|\det K_{m,n}| &= \sqrt{\det A_{m,n}} \\
&= \prod_{j=1}^{m}\prod_{k=1}^{n}\sqrt{2\cos\frac{\pi j}{m+1} + 2i\cos\frac{\pi k}{n+1}} \\
&= \prod_{j=1}^{m}\prod_{k=1}^{n}\sqrt[4]{4\cos^2\frac{\pi j}{m+1} + 4\cos^2\frac{\pi k}{n+1}}.
\end{aligned}$$

Remarkable!

E-mail address: `joshuah@alum.mit.edu`
URL: `http://joshuahhh.com/`