

Some combinatorial problems arising in the dimer model

R. Kenyon*

1 Introduction

We collect here a few diverse problems on the dimer model. For background on dimers one can read for example [6].

Let $G = (V, E)$ be a finite, connected, bipartite graph. A *dimer cover*, or *perfect matching*, of \mathcal{G} is a collection $m \subset E$ of edges that covers each vertex exactly once. We denote by $\mathcal{M} = \mathcal{M}(\mathcal{G})$ the set of dimer covers of \mathcal{G} .

Let $c : E \rightarrow \mathbb{R}_{>0}$ be a positive weight function edges of \mathcal{G} . If $m \in \mathcal{M}$ is a dimer cover, we define $c(m) = \prod_{e \in m} c_e$ to be the *weight* of m . This weight function defines a probability measure $\mu = \mu_c$ on \mathcal{M} , giving a dimer cover a probability proportional to its weight, that is, $\Pr(m) = \frac{c(m)}{Z}$, where the constant of proportionality Z is the *partition function*: $Z = \sum_{m \in \mathcal{M}} c(m)$.

Essentially the most fundamental problem in the dimer model is the following:

Problem 0. *For a random dimer cover, compute the edge probabilities, that is, compute the probability $\Pr(e)$ that any particular edge e is used.*

This is easy if we can compute Z as a function of edge weights, since

$$\Pr(e) = \frac{c_e}{Z} \frac{\partial Z}{\partial c_e} = \frac{\partial \log Z}{\partial \log c_e}. \quad (1)$$

However on general graphs computing Z is hard, even for constant edge weights: enumerating dimer covers is in fact $\#P$ -complete, by a celebrated result of Valiant [11]. Note that computing $\Pr(e)$ is in fact just as hard as computing Z , since one can reconstruct Z from integrating $\Pr(e)$ using (1). For planar graphs we can however compute Z and thus solve the above problem, using Kasteleyn's method of enumerating dimer covers with determinants, see [6] and Section 2.2 below.

One useful change of perspective is to consider the edge weights as a “ \mathbb{C}^* -local system” on \mathcal{G} , that is, as a connection on a line bundle on the graph. This point of view may seem at first just an added level of abstraction and taking us away from combinatorics, but it turns out to have a number of advantages, such as:

*Mathematics Department, Yale University, New Haven

1. It highlights a nonobvious symmetry of the problem: the gauge invariance (see below).
2. It generalizes naturally to other gauge groups such as $\mathrm{SL}_n(\mathbb{R})$, with interesting combinatorial interpretations (see Sections 6, 7 below).
3. It connects the problem to geometry.

Our choice of problems in this article is (mostly) motivated by this geometric point of view.

Acknowledgements

This research was supported by NSF grant DMS-1940932 and the Simons Foundation grant 327929.

2 Preliminaries

2.1 Vector bundles and connections on graphs

What is a bundle with connection on a graph? Fix a vector space \mathcal{V} and assign an isomorphic copy \mathcal{V}_v of \mathcal{V} to each vertex $v \in V$. For each edge $e = uv$ we associate a linear isomorphism $\phi_{uv} : \mathcal{V}_u \rightarrow \mathcal{V}_v$ of these vector spaces, with $\phi_{vu} = \phi_{uv}^{-1}$. Generally we have $\phi_{uv} \in \mathrm{GL}(\mathcal{V})$ but we can restrict if desired to other subgroups of $\mathrm{GL}(\mathcal{V})$, and even noninvertible endomorphisms [2]. In this paper we will only use $\mathrm{SL}_n(\mathbb{C})$ connections on \mathbb{C}^n or \mathbb{C}^* -connections on \mathbb{C} . (Here $\mathbb{C}^* = \mathbb{C} \setminus \{0\}$ is the group of \mathbb{C} -linear isomorphisms of \mathbb{C} .) These are called SL_n -local systems and \mathbb{C}^* -local systems respectively.

If we choose a basis for each vector space \mathcal{V}_v then ϕ_{uv} is a matrix. Changing basis by matrices A_v at vertices v results in a new matrix $\phi'_{uv} = A_u \phi_{uv} A_v^{-1}$. Such changes are called *gauge transformations*. They form a symmetry of the system.

If γ is an oriented loop in \mathcal{G} based at x_0 , the *monodromy* of the connection around γ is the composition of the isomorphisms along γ starting at x_0 . Its conjugacy class is invariant under gauge transformations. Note also that changing the basepoint along γ conjugates the monodromy.

For one-dimensional bundles, with $\mathcal{V} = \mathbb{R}$, $\phi_{uv} \in \mathbb{R}$ is a scalar. Since we are assuming \mathcal{G} is bipartite, we can canonically orient the edges from black to white. Then ϕ_{bw} is a scalar quantity associated to edge bw which we call the *edge weight*. In probability settings we usually take ϕ_{bw} positive. However it is sometimes useful to consider $\mathcal{V} = \mathbb{C}$ and $\phi_{bw} \in \mathbb{C}^*$. Gauge transformations are functions $A : V \rightarrow \mathbb{C}^*$ and transform ϕ into $\phi'_{bw} = \phi_{bw} A_b A_w^{-1}$.

For an \mathbb{R} -bundle, when edge weights ϕ_{bw} are positive, we can associate a probability measure μ as above. Then positive gauge transformations (i.e. when the $A_v > 0$) change the weights of individual dimer covers but not the probability measure μ : If we multiply the

edge weights of all edges at a single vertex by a constant $\lambda > 0$, then the weight of every dimer cover is multiplied by λ . This implies that the probability measure μ only depends on the equivalence class of edge weights under gauge transformations, that is, on the “edge weights modulo gauge”.

2.2 Planar bipartite graphs and the Kasteleyn matrix

Let \mathcal{G} be a planar bipartite graph with a \mathbb{C}^* -local system Φ (equivalently, a set of edge weights in \mathbb{C}^* , see the previous section). We define a matrix K , the *Kasteleyn matrix for \mathcal{G}* , as follows. The matrix $K = (K_{wb})$ has rows indexing white vertices and columns indexing black vertices, and

$$K_{wb} = \begin{cases} \pm\phi_{bw} & b \sim w \\ 0 & \text{otherwise} \end{cases}$$

where the signs are chosen by the “Kasteleyn rule”: a face of length ℓ has $\frac{\ell}{2} + 1 \pmod 2$ minus signs. Kasteleyn showed in this case that the determinant of K counts weighted dimer covers:

Theorem 1 ([4, 7]). $|\det K| = \sum_{m \in \mathcal{M}} c(m)$.

As an application, probabilities of one or more edges being in a dimer cover are minors of the inverse of K , see [5].

If \mathcal{G} has an SL_n -local system, one can similarly define a Kasteleyn matrix, whose entries are now $n \times n$ matrices themselves; see Sections 6.1, 7.2 below.

3 From edge weights to probabilities

We wish to study the mapping from edge weights to probabilities. This requires first defining the correct spaces.

3.1 Fractional matchings

A *fractional matching* of $G = (V, E)$ is a function $f : E \rightarrow [0, 1]$ which sums to one at each vertex: $\sum_{u \sim v} f(u) = 1$. A dimer cover m determines a fractional matching, by assigning $f(e) = 1$ for each $e \in m$ and $f(e) = 0$ otherwise. The set of fractional matchings $\Omega(\mathcal{G})$ is a polytope in \mathbb{R}^E (in fact a subset of $[0, 1]^E$) whose vertices are exactly the dimer covers of \mathcal{G} , see [8].

A probability measure on dimer covers is a probability measure on the vertices of Ω , and so has a barycenter which is a point in Ω . The coordinate of this barycenter corresponding to an edge is the probability of that edge.

So a more elegant statement of Problem 0 is

Problem 1. *Compute the expected fractional matching.*

See an example in Figure 1.




Figure 1: A graph and its expected fractional matching (for uniform edge weights). This graph has three dimer covers; two of these use the leftmost vertical edge, so its probability is $2/3$.




Figure 2: A degenerate graph. The dimension of the cycle space is 3 but Ω is of dimension 2. Note that the two diagonal edges are “unused”, that is, do not appear in any dimer cover.

3.2 Cycles

A *flow* on a graph \mathcal{G} is a function f on oriented edges satisfying $f(-e) = -f(e)$ where $-e$ represents the edge with the reversed orientation. The *cycle space* $H_1(\mathcal{G})$ is the space of divergence-free flows: flows f satisfying $\sum_{u \sim v} f(vu) = 0$ for all v .

Since \mathcal{G} is bipartite, we can orient edges from black to white, and then for any two fractional matchings $f_1, f_2 \in \Omega$, $f_1 - f_2$ represents an element of the cycle space of \mathcal{G} , denoted $[f_1 - f_2]$. Fixing a basepoint $f_0 \in \Omega$, we can map Ω linearly and injectively into $H_1(\mathcal{G})$ by $f \mapsto [f - f_0]$.

We say \mathcal{G} is *nondegenerate* if Ω is of dimension $d = |E| - |V| + 1$, the dimension of the cycle space of the graph. A graph is *degenerate* (that is, not nondegenerate) if there are unused edges (edges which participate in no dimer cover) whose removal does not disconnect the graph [8]. Figure 3.2 shows a degenerate graph.

3.3 Gauge equivalence

For a bipartite graph, the space of equivalence classes of edge weights modulo gauge transformations, that is, the space of \mathbb{C}^* -local systems, is isomorphic to $(\mathbb{C}^*)^d$ where d is the dimension of the cycle space of the graph. Explicitly, for a cycle γ we can define X_γ to be the monodromy of the connection around γ . In terms of edge weights, X_γ is the alternating

product of edge weights around γ : assuming we start at a black vertex, X_γ is the first weight, divided by the second, times the third, and so on. Note that X_γ is invariant under gauge transformations. Reversing the orientation inverts the monodromy. If γ ranges over a basis for the cycle space, the set of monodromies $\{X_\gamma\}$ parameterize the space of all edge weights modulo gauge, that is, all \mathbb{C}^* -local systems.

For a planar graph, a basis for the cycle space is the set of cycles around the bounded faces. Consequently we can take the “face weights” $\{X_f\}_{f \in F}$ on the bounded faces of the graph to parameterize the choices of edge weights modulo gauge.

Now a perhaps more interesting variant of Problem 0 is:

Problem 2. *For nondegenerate graphs, study the map Ψ from cycle weights $\{X_f\}$ to the expected fractional matching.*

The domain and range of Ψ have the same dimension: the dimension of the cycle space of the graph.

Theorem 2. *If \mathcal{G} is nondegenerate then Ψ is a diffeomorphism from the space of cycle weights $(\mathbb{R}_{>0})^d$ to $\Omega(\mathcal{G})$.*

Proof. Fix the gauge by choosing a spanning tree T of \mathcal{G} and setting edge weights 1 on the tree T and weights c_e for $e \notin T$. Let $Z = Z(\{c_e\})$ be the associated partition function. Let p_e be the probability of edge e for $e \notin T$. All the other edge probabilities (and therefore the fractional matching f) are uniquely determined by the probabilities $\{p_e \mid e \notin T\}$: one can determine the value of f on the leaves of the tree, then remove these and continue with the leaves of the subtrees inductively. As in (1) we have $p_e = c_e \frac{\partial \log Z}{\partial c_e}$. We can write this as $\vec{p} = \nabla \log Z$ where the gradient is taken with respect to the vector of “edge energies” $(\log c_e)_{e \notin T}$.

Since Z is a positive polynomial (that is, its coefficients are nonnegative), it is, up to scale, the probability generating function of a probability measure on \mathbb{Z}^d of finite support. The Hessian of $\log Z$ with respect to the logs of its variables is the covariance matrix of this probability measure. Any covariance matrix is positive semidefinite, and thus this Hessian is positive semidefinite: $\text{Hess} \log Z \geq 0$. In fact it is positive definite as long as the Newton polytope of Z is of positive volume, that is, if the exponents of the monomials of Z span the cycle space. But this is guaranteed by nondegeneracy. Since $\text{Hess}(\log Z)$ is positive definite, $\log Z$ is strictly convex and so $\nabla \log Z$ is a diffeomorphism onto its image. \square

Note that $\det(\nabla \Psi)$ is a rational function of the c_e . What is the degree of $\nabla \Psi$ as a rational map on \mathbb{C}^d ? If we extend Ψ to all of \mathbb{C}^d it appears that the inverse of a point $\vec{p} \in \Omega$ consists in *real* values of c_e (only one preimage consists in positive reals, but there can be multiple non-positive preimages; it is not clear that these are always real, although they are in the small examples we tried). A similar situation, for resistor networks, where reality of the other preimages is proved, can be found in [1].

4 Dimer walks

To each dimer cover $m \in \mathcal{M}(\mathcal{G})$ is associated an involution $\pi_m : V \rightarrow V$. This is the involution exchanging a vertex with the vertex it is paired with.

Problem 3. Let $\{m_1, m_2, \dots\}$ be a sequence of i.i.d. dimer covers of \mathcal{G} . Study the random walks on S_V , the permutation group on the vertices, defined by the π_{m_i} .

Another way to state this is to define $g = \frac{1}{Z} \sum_{m \in \Omega_1(\mathcal{G})} \pi(m) \in \mathbb{R}[S_V]$, the element of the group algebra $\mathbb{R}[S_V]$ of S_V defined by a random choice of dimer cover. What can be said about g^n ?

As a simple example one can take $\mathcal{G} = K_4$ with vertices $[4] = \{1, 2, 3, 4\}$. The three dimer covers correspond to the three permutations $(12)(34), (13)(24), (14)(23)$ which generate a subgroup of S_4 isomorphic to $(\mathbb{Z}_2)^2$. The dimer random walk just becomes simple random walk on this group, where every element is adjacent to any other element, that is, we have simple random walk on K_4 .

For another simple example, let \mathcal{G} be the 3×2 grid. Rather than record the element in S_6 , without loss of information we can record simply the permutations of the x -coordinates of the vertices, in S_3 . There are 3 dimer covers, and the walk corresponds to multiplication by $g = \frac{1}{3}(1 + e_{(12)} + e_{(23)})$. Its eigenvalues are $1, \frac{2}{3}, \frac{2}{3}, -\frac{1}{3}, 0, 0$.

As a more interesting example, when \mathcal{G} is the $n \times n$ grid graph on a torus, each vertex undergoes a simple random walk on V , since with probability $1/4$ it is matched to any neighbor. These simple random walks are coupled to avoid each other. How quickly does the process mix? What is the expected winding of two walks around each other, as a function of time?

5 Magnetic double dimer model

5.1 Double dimers

Recall that $\mathcal{M} = \mathcal{M}(\mathcal{G})$ is the set of dimer covers of \mathcal{G} . A *double dimer cover* is a function $m : E \rightarrow \{0, 1, 2\}$ which sums to 2 at each vertex. See Figure 3. A union of two single dimer covers determines a double dimer cover, and (for bipartite graphs) every double dimer cover occurs this way. The map from pairs of dimer covers to double dimer covers is not, however, injective: each double dimer cover m arises from 2^ℓ ordered pairs of single dimer covers, where ℓ is the number of cycles in m , where the cycles are formed from the edges of multiplicity 1.

Let \mathcal{M}_2 be the set of double dimer covers of \mathcal{G} . The natural probability measure on \mathcal{M}_2 (for us) is *not* the uniform measure. It is, rather, the projection of the uniform measure on $\mathcal{M} \times \mathcal{M}$ under the standard map $\mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}_2$. Thus a double dimer cover m has a probability proportional to 2^ℓ where ℓ is the number of cycles (cycles formed from edges of




Figure 3: A double dimer cover of a grid. Edges of multiplicity 1 or 2 are shown.

multiplicity 1; edges of multiplicity 2 do not count as cycles). We call 2^ℓ the *weight* $c(m)$ of m .

Problem 4. *What can be said about the distribution of loops in the double dimer model on \mathbb{Z}^2 ?*

Let \mathcal{G} be a bipartite planar graph, and $q \in \mathbb{C}^*$ be a variable. One can define a \mathbb{C}^* -connection on \mathcal{G} so that each face has counterclockwise monodromy q , see for example Figure 4. Let $K(q)$ be an associated Kasteleyn matrix. For a double-dimer configuration m_2 define

$$c(m_2) = \prod_{\text{loops } \gamma \text{ of } m_2} (q^{A_\gamma} + q^{-A_\gamma}) \quad (2)$$

where A_γ is the area (number of faces) enclosed by γ . We then have

Theorem 3.

$$\det K(q) \det K(1/q) = \sum_{m_2 \in \mathcal{M}_2} \prod_{\text{loops } \gamma \text{ of } m_2} (q^{A_\gamma} + q^{-A_\gamma}). \quad (3)$$

Proof. The quantity $\det K(q) \det K(1/q)$ counts pairs of dimer configurations (m, m') , where m' has the inverted q -weight. In the superposition of m and m' , we orient m 's dimers from black to white and m' 's dimers from white to black, reversing the parallel transport, so that the orientations of loops are consistent. In this way the weight of an oriented loop is given by its monodromy, that is, q^A where A is the area enclosed by the loop. When we sum over both orientations of each loop, we find (3). \square




Figure 4: Connections with monodromy q per face.

It is tempting to use (3) to answer Problem 4. For example one can expand $\det K(q) \det K(1/q)$ near $q = 1$ to get various moments. However extracting information about the areas from this does not appear easy.

An alternative approach is to invert the standard Kasteleyn matrix (when $q = 1$) and use it to compute the probability of any given shape of loop. For the square grid \mathbb{Z}^2 , this can be used to compute the expected density of loops of fixed small area, as follows. The “inverse Kasteleyn matrix” for \mathbb{Z}^2 can be obtained as a limit as $n \rightarrow \infty$ of $K_{\mathcal{G}_n}^{-1}$ for \mathcal{G}_n the $2n \times 2n$ grid, see [6]. There is an explicit formula: when $b - w = (x, y) \in \mathbb{Z}^2$,

$$K^{-1}((0, 0), (x, y)) = \frac{1}{(2\pi i)^2} \iint_{|z|=|w|=1} \frac{z^{(-x+y+1)/2} w^{(-x-y+1)/2} dz dw}{1 + z + w - zw} \frac{dz}{z} \frac{dw}{w}.$$

Minors of this matrix compute edge probabilities for random dimer covers of \mathbb{Z}^2 , see [5]:

$$Pr((w_1, b_1), \dots, (w_k, b_k)) = |\det(K^{-1}(w_i, b_j)_{1 \leq i, j \leq n})|.$$

The expected density of loops of area 1 (the probability that a given face is in such a loop) is the probability that for a pair of single dimer models (m, m') on \mathbb{Z}^2 , m contains edges $(0, 0)(1, 0)$ and $(0, 1)(1, 1)$ and m' contains edges $(0, 0)(0, 1)$ and $(1, 0)(1, 1)$, or the reverse. The probability that m contains both edges $(0, 0)(1, 0)$ and $(0, 1)(1, 1)$ is, by a short computation, $\frac{1}{8}$. So the expected density of loops of area 1 in the double dimer model is $\frac{2}{8^2} = \frac{1}{32}$. Similarly the density of loops of area 2 can be computed to be $\frac{(\pi-1)^2}{2\pi^4}$. The density for area-3 loops is already more complicated:

$$\frac{3(64 - 192\pi^2 + 192\pi^3 - 32\pi^4 - 32\pi^5 + 24\pi^6 - 8\pi^7 + \pi^8)}{32\pi^8}.$$

6 SL_2 connections

6.1 Kasteleyn matrix

Consider \mathcal{G} a planar bipartite graph with an SL_2 -local system Φ . Let $K = (K_{wb})$ be an associated Kasteleyn matrix for \mathcal{G} , that is a matrix with rows indexing white vertices and columns indexing black vertices, and

$$K_{wb} = \begin{cases} \pm\phi_{bw} & b \sim w \\ 0 & \text{otherwise} \end{cases}$$

where the signs are chosen by the same Kasteleyn rule as in the scalar case: a face of length ℓ has $\frac{\ell}{2} + 1 \pmod 2$ minus signs. Here by 0 we mean the zero matrix in $M_2(\mathbb{R})$. Then K is an $N \times N$ matrix with entries in $M_2(\mathbb{R})$. Let \tilde{K} be the $2N \times 2N$ matrix obtained from K by replacing each entry with its 2×2 array of real numbers.

By a theorem of [2],

$$\det \tilde{K}(\Phi) = \sum_{m \in \mathcal{M}_2} \prod_{\text{loops in } m} \mathrm{Tr}(\phi_\gamma). \quad (4)$$

Here ϕ_γ is the monodromy of the connection Φ around the loop γ , and $\mathrm{Tr}(\phi_\gamma)$ is its trace. Even though ϕ_γ depends on a starting vertex and the orientation of the loop, the trace $\mathrm{Tr}(\phi_\gamma)$ is independent of starting vertex (since the trace of a matrix only depends on its conjugacy class) and orientation (since for matrices $A \in \mathrm{SL}_2$ we have $\mathrm{Tr}(A) = \mathrm{Tr}(A^{-1})$).

6.2 Flat connections and simple laminations

Suppose graph \mathcal{G} is drawn on a multiply-connected planar domain Σ , and Φ is a *flat* SL_2 connection. This means the monodromy of Φ around any cycle in \mathcal{G} which is contractible as a cycle in Σ , is the identity. In this case the trace of the monodromy around a loop γ only depends on its isotopy class as a loop in Σ . In other words two loops in \mathcal{G} which are isotopic as loops in Σ , have monodromies with the same trace.

A *simple lamination* is an isotopy class of finite collections of pairwise disjoint simple closed loops on Σ . Let Λ_2 be the collection of simple laminations.

When Φ is flat we can group the terms in the sum (4) according to their isotopy classes:

$$\det \tilde{K}(\Phi) = \sum_{\lambda \in \Lambda_2} C_\lambda \mathrm{Tr}(\phi_\lambda) \quad (5)$$

with $C_\lambda \in \mathbb{Z}_{\geq 0}$.

By a Theorem of Fock and Goncharov [3], the functions $\mathrm{Tr}(\phi_\lambda)$ as λ runs over Λ_2 , considered as functions of Φ , are linearly independent and in fact form a linear basis for the regular (i.e. polynomial) functions on the character variety (the space of flat SL_2 -local systems Φ modulo gauge). As a consequence C_λ can be in principle determined from $\det \tilde{K}(\Phi)$




Figure 5: Simple laminations on a pair of pants consist in three kinds of loops: those surrounding p_1, p_2 or both.

as Φ varies over flat connections. Mysteriously, even though for a finite graph $\det \tilde{K}(\Phi)$ is a polynomial function of the matrix entries, and the sum (5) is a finite sum, it is not at all clear how to extract the individual C_λ from it.

Problem 5. *How can one extract C_λ from $\det \tilde{K}(\Phi)$ as in (5)?*

For an example where we know how to extract these coefficients, consider the simplest nontrivial case when Σ is an annulus. Suppose graph \mathcal{G} is drawn on an annulus with flat connection having monodromy A around the nontrivial isotopy class γ , the generator of $\pi_1(\Sigma)$. Then a simple lamination on \mathcal{G} just consists in some number n of copies of loops each isotopic to γ (we are ignoring orientation). The isotopy classes Λ_2 are thus in correspondence with the set of nonnegative integers. We can write

$$\det \tilde{K}(A) = \sum_{j=0}^{\infty} C_j \text{Tr}(A)^j.$$

From this expression it is easy to extract C_j : let $\text{Tr}(A) = z$, then the C_j are power series coefficients of $\det \tilde{K}(A)$.

The next simplest planar surface is a disk with two small disks p_1, p_2 , removed (also known as a “pair of pants”). In this case the space Λ_2 has a simple parameterization by $(\mathbb{Z}_+)^3$ where (i, j, k) corresponds to the lamination with i, j , and k loops surrounding respectively p_1 only, p_2 only or both p_1 and p_2 , see Figure 5. If A, B are the monodromies around p_1, p_2 respectively then given any $x, y, z \in \mathbb{C}^*$ we can choose A, B in $\text{SL}_2(\mathbb{C})$ so that $x = \text{Tr}(A), y = \text{Tr}(B), z = \text{Tr}(AB)$. Then one can extract $C_{i,j,k}$ by a contour integral

$$C_{i,j,k} = \frac{1}{(2\pi i)^3} \int_{(S^1)^3} \frac{\det K(\Phi)}{x^i y^j z^k} \frac{dx}{x} \frac{dy}{y} \frac{dz}{z}.$$

For a general punctured surface Fock and Goncharov [3] describe a different (and orthogonal) basis on the space of laminations, the “Peter-Weyl basis”, and show that the basis

$\{\mathrm{Tr}(\phi_\lambda)\}_{\lambda \in \Lambda_2}$ is obtained from the Peter-Weyl basis by an upper triangular transformation. One route to solving Problem 5 would be to write down this base-change matrix.

7 n -fold dimer model

7.1 n -multiwebs

An n -multiweb, or n -fold dimer cover in \mathcal{G} is a function $m : E \rightarrow \{0, 1, 2, \dots, n\}$ which sums to n at each vertex. So a 1-multiweb is a dimer cover and a 2-multiweb is a double dimer cover. A union of n single dimer covers is an n -multiweb, and every n -multiweb occurs this way (although not uniquely).

Let \mathcal{M}_n be the set of n -multiwebs of \mathcal{G} . The natural probability measure on \mathcal{M}_n is, like in the double dimer case, the projection of the uniform measure on \mathcal{M}^n under the standard map $\mathcal{M}^n \rightarrow \mathcal{M}_n$. Unlike the $n = 2$ case, the size of the preimage is difficult to compute in general, in fact #P-complete: if $n = 3$ and a 3-multiweb is a trivalent graph (that is, all edges have multiplicity 1), then the triples of dimer covers in the preimage are in bijection with edge 3-colorings, or Tait colorings. Tait colorings are colorings of the edges with three colors so that each color appears at each vertex. For planar graphs Tait colorings are dual to proper 4-colorings of the vertices of the dual graph.

7.2 Kasteleyn matrix

Consider \mathcal{G} a planar bipartite graph with an SL_n -local system Φ . Let $K = (K_{wb})$ be a Kasteleyn matrix for \mathcal{G} ; K is an $N \times N$ matrix with entries in $M_n(\mathbb{R})$. Let \tilde{K} be the $nN \times nN$ matrix obtained from K by replacing each entry with its $n \times n$ array of reals.

By a theorem of [2],

$$\det \tilde{K}(\Phi) = \sum_{m \in \mathcal{M}_n} \prod_{m \in \mathcal{M}_n} \mathrm{Tr}(\phi_m) \quad (6)$$

where we still need to define the trace of an n -multiweb with an SL_n -connection; see the next section.

7.3 Trace of a 3-multiweb

The trace of an n -multiweb m is not simple to define; it is a contraction of n -tensors defined at each vertex. We refer to [2] for the general definition and give here a working definition for $n = 3$. We cut each edge of \mathcal{G} into two half-edges, one associated with each of its vertices. We consider colorings c of the half-edges of m with three colors $\{1, 2, 3\}$, so that an edge of multiplicity $j \in \{0, 1, 2, 3\}$ gets a set of j colors, and so that all colors appear at each vertex of \mathcal{G} . See Figure 6.




Figure 6: 3-coloring of the half-edges of \mathcal{G} .

Then on an edge e of multiplicity j we have two sets of colors S, T of size j , with S located near the white vertex and T near the black vertex. We define

$$\text{Tr}(c) = \sum_{\text{colorings } c} \text{sgn}(c) \prod_e (\phi_{bw})_S^T \quad (7)$$

where $(\phi)_S^T$ denotes the S, T minor of ϕ . Here $\text{sgn}(c)$ is a sign depending on the cyclic order of colors at each vertex: at a black vertex if the colors are in counterclockwise order we get sign $+$ and otherwise sign $-$, and this convention is reversed at a white vertex; the product of these signs over all vertices is $\text{sgn}(c)$. (For edges with multiplicity the colors are assumed in natural order within that set.)

When Φ is the identity connection, and m is planar, $|\text{Tr}(m)|$ is the number of Tait colorings of m (see [2]). In other words, all nonzero terms of (7) have the same sign. The global sign of $\text{Tr}(m)$ is unimportant since it depends on an artificial choice of vertex ordering.

7.4 Flat connections and reduced 3-webs

Suppose \mathcal{G} is drawn on a multiply-connected planar domain Σ , and Φ is a *flat* SL_3 connection.

A vertex in a 3-multiweb is *trivalent* if it has three adjacent edges of multiplicity 1. A 3-multiweb consists in a collection of trivalent vertices connected in pairs by chains of edges with multiplicities alternating between 1 and 2 along the chains.

We say a 3-multiweb is *reduced* or *nonelliptic* if each component, considered as a planar graph by itself, has no contractible faces with 0, 2 or 4 trivalent vertices.

There are certain *skein* relations by which the trace of any 3-multiweb can be written as a linear combination of traces of reduced 3-multiwebs (see Figure 7). The set of isotopy classes of the reduced 3-multiwebs which arise from reducing a given web are well defined.

As a consequence for a flat SL_3 -connection we can group the terms in the sum (6) according to isotopy classes of reduced 3-multiwebs:

$$\det \tilde{K}(\Phi) = \sum_{\lambda \in \Lambda_3} C_\lambda \text{Tr}(\phi_\lambda). \quad (8)$$

By a Theorem of Sikora and Westbury [10], the functions $\text{Tr}(\lambda)$ as λ runs over Λ_3 form a linear basis for the regular (i.e. polynomial) functions on the character variety of flat




Figure 7: Relations between 3-multiwebs: a loop can be removed, increasing the weight of the remaining web by a factor of 3; a bigon can be replaced with a single edge as shown multiplying the weight by 2; a doubled edge can be replaced with a single edge as shown, a square face can be resolved into a linear combination of two webs.

SL_3 -local systems modulo gauge. As a consequence C_λ can be in principle determined from $\det \tilde{K}(\Phi)$ as Φ varies over flat connections.

Problem 6. *How can one extract C_λ from $\det \tilde{K}(\Phi)$?*

Applying the skein relations to a non-reduced 3-multiweb results in a collection of reduced webs which depends on the order in which the skein relations are applied. Even though the isotopy classes of the reduced web are well-defined, the individual webs themselves will depend on the order. Is there a way to make a canonical choice in this reduction process, so that starting from a random 3-multiweb we arrive at a well-defined, canonical probability measure on reduced 3-multiwebs?

References

- [1] A. Abrams, R. Kenyon, Fixed-energy harmonic functions. *Discrete Anal.* 2017, Paper No. 18, 21 pp.
- [2] D. Douglas, R. Kenyon, H. Shi, Dimers, webs, and local systems, arxiv:2205.05139
- [3] V. Fock, A. B. Goncharov, Moduli spaces of local systems and higher Teichmüller theory, *Publ. Math. Inst. Hautes Études Sci.*,103, (2006), pp 1–211
- [4] P. Kasteleyn, Dimer statistics and phase transitions. *J. Mathematical Phys.* 4 (1963), 287–293.
- [5] R. Kenyon, Local statistics of lattice dimers. *Ann. Inst. H. Poincaré Probab. Statist.* 33 (1997), no. 5, 591–618.

- [6] R. Kenyon, Lectures on dimers. Statistical mechanics, 191–230, IAS/Park City Math. Ser., 16, Amer. Math. Soc., Providence, RI, 2009.
- [7] G. Kuperberg, An exploration of the permanent-determinant method. *Electron. J. Combin.* 5 (1998), no. 1, Research Paper 46, 34 pp.
- [8] L. Lovasz, M. Plummer, Matching theory. North-Holland Mathematics Studies, 121. *Annals of Discrete Mathematics*, 29. North-Holland Publishing Co., Amsterdam.
- [9] A. S. Sikora. SL_n -character varieties as spaces of graphs. *Trans. Amer. Math. Soc.*, 353 (2001): 2773–2804, .
- [10] A. S. Sikora and B. W. Westbury. Confluence theory for graphs. *Algebr. Geom. Topol.*, 7:439–478, 2007.
- [11] L. G. Valiant, The Complexity of Computing the Permanent. *Theoretical Computer Science*. Elsevier. 8 (2): 189–201 (1979)