Squishing dimers on the hexagon lattice

Ben Young *

Submitted: Aug 12, 2008; Accepted: Jul 16, 2009; Published: Jul 24, 2009
Mathematics Subject Classification: 05A15

Abstract

We describe an operation on dimer configurations on the hexagon lattice, called "squishing", and use this operation to explain some of the properties of the Donaldson-Thomas partition function for the orbifold $\mathbb{C}^3/\mathbb{Z}_2 \times \mathbb{Z}_2$ (a certain four-variable generating function for plane partitions which comes from algebraic geometry).

1 Introduction

In this paper, we will describe and use a novel technique called "squishing", which one applies to the dimer model on the regular honeycomb lattice. We developed this technique as an attempt to verify a conjectured generating function which arises in algebraic geometry (specifically, in the Donaldson-Thomas theory of the orbifold $\mathbb{C}^3/\mathbb{Z}_2 \times \mathbb{Z}_2$ [1]). Our attempt was only partially successful, and we were later able to compute the generating function by other means. However, the technique is interesting in of itself, being one of relatively few dimer model techniques which exploits the self-similarity of the lattice at different scales.

We begin by describing the original motivation for this work.

Definition 1 *A* 3D Young diagram *(or* 3D diagram*)* π *is a subset of* $(\mathbb{Z}_{\geq 0})^3$ *such that if* $(x, y, z) \in \pi$*, then* $(x', y', z') \in \pi$ *whenever* $x' \leq x$*,* $y' \leq y$*, and* $z' \leq z$*.*

3D Young diagrams are called *boxed plane partitions* or *3D partitions* elsewhere in the literature. We refer to the points in π as *boxes* – the point (i, j, k) corresponds to the unit cube with vertices $\{(i \pm \frac{1}{2}, j \pm \frac{1}{2}, k \pm \frac{1}{2})\}$.

We will be discussing the following generating functions:

Definition 2 *A* weighting *of* $(\mathbb{Z}_{\geq 0})^3$ *is a map*

$$w : (\mathbb{Z}_{\geq 0})^3 \to \{p, q, r, s\},$$

*Email: byoung@math.mcgill.ca

where p, q, r, s are formal indeterminates. We say that $w(P)$ is the weight *of the lattice point P; the* weight *of a 3D diagram is the product of the weights of the lattice points at the centers of all of its boxes. The* w-partition function *is then defined to be the formal sum*

$$Z_w = \sum_{\pi \text{ 3D diagram}} w(\pi).$$

The weightings we will be concerned with are the *monochromatic* weighting $w_{\{1\}}$,

$$(i, j, k) \mapsto p$$

and the $\mathbb{Z}_2 \times \mathbb{Z}_2$ weighting $w_{\mathbb{Z}_2 \times \mathbb{Z}_2}$,

$$(i, j, k) \mapsto \begin{cases} p & \text{if } i-k \equiv 0, j-k \equiv 0 \pmod 2, \\ q & \text{if } i-k \equiv 1, j-k \equiv 0 \pmod 2, \\ r & \text{if } i-k \equiv 0, j-k \equiv 1 \pmod 2, \\ s & \text{if } i-k \equiv 1, j-k \equiv 1 \pmod 2, \end{cases}$$

along with various specializations of these weightings; we shall denote their partition functions $Z_{\{1\}}(p)$ and $Z_{\mathbb{Z}_2 \times \mathbb{Z}_2}(p, q, r, s)$, respectively.

It is a classical result [4] that

$$\sum_{\pi} w_{\{1\}}(\pi) = M(1, p) \tag{1}$$

where we define

$$M(a, z) = \prod_{n=1}^{\infty} \left(\frac{1}{1 - az^n} \right)^n.$$

Equation (1) also arises in algebraic geometry, essentially because 3D Young diagrams are the same as monomial ideals $I \subset \mathbb{C}^3[x, y, z]$ (simply read off the exponents of the elements of the coordinate ring $\mathbb{C}^3/I$; these are the boxes of π). As a result, (1) is, in a certain sense, an invariant of $\mathbb{C}^3$; specifically, it is the *Donaldson-Thomas partition function* for the space $\mathbb{C}^3$, up to a sign on p. [5].

It is possible [1] to develop Donaldson-Thomas theory for orbifolds of $\mathbb{C}^3$ under the actions of certain finite abelian groups. It turns out that for the group $\mathbb{Z}_2 \times \mathbb{Z}_2$, the partition function is given by

$$Z_{\mathbb{Z}_2 \times \mathbb{Z}_2} = \frac{M(1, Q)^4 \widetilde{M}(qr, Q) \widetilde{M}(qs, Q) \widetilde{M}(rs, Q)}{\widetilde{M}(-q, Q) \widetilde{M}(-r, Q) \widetilde{M}(-s, Q) \widetilde{M}(-qrs, Q)} \tag{2}$$

where $Q = pqrs$ and $\widetilde{M}(a, z) = M(a, z) M(a^{-1}, z)$.

The curious identity (2) was proven in [1] using vertex operators; it was conjectured earlier by Bryan (based on the behaviour of related Donaldson-Thomas partition functions) and independently by Kenyon (in an equivalent form, based on empirical computer work).

Figure 1: A matching of $H_{5,4,3}$.

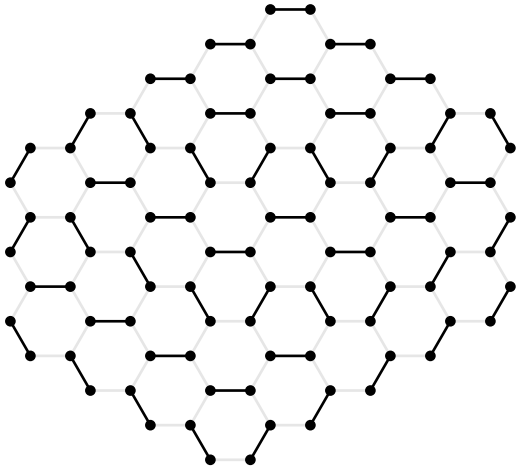

It is possible to check some of the properties of 2 in an elementary manner. For example, it should be the case that specializing $p = q = r = s$ should give

$$Z_{\mathbb{Z}_2\times\mathbb{Z}_2}(p, p, p, p) = Z_{\{1\}},$$

and indeed (2) does satisfy this relation. Another striking relation suggested by (2) is:

$$Z_{\mathbb{Z}_2\times\mathbb{Z}_2}(p, -1, -1, -1) = M(1, Q)^2; \tag{3}$$

however, the combinatorial reason for this is far less clear. This paper demonstrates (3) in an elementary manner, without relying on the theorems of [1].

2 Matchings, and Weightings

Our attack on (3) immediately requires us to to encode the "surface" of a 3D Young diagram with dimers on the hexagon lattice. Let us fix some terminology.

Definition 3 *Let $G = (V, E)$ be a graph. A* matching *or* 1-factor *of G is a subgraph $M = (V, E')$ such that the degree of every vertex $v \in V$ is 1. Equivalently, a matching of G is a partition of the vertices of G into a disjoint union of* dimers*, or pairs of vertices joined by a single edge of G.*

Note that elsewhere in the graph theory literature, these are called "perfect matchings", and "matching" means a 1-factor of some subgraph of G.

For a general survey of results on the dimer model, see [2].

We will be considering matchings on the semiregular hexagonal mesh of side lengths a, b, c, a, b, c, denoted $H_{a,b,c}$ (see Figure 1 for a definition–by–picture). Matchings on this

Figure 2: A 3D Young diagram viewed as a matching on $H_{3,3,3}$.

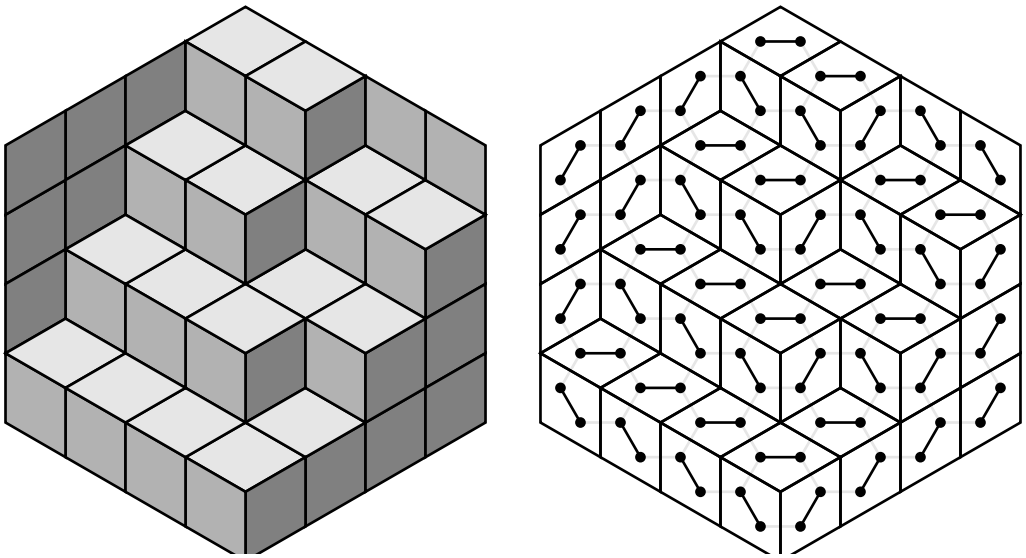

graph are in bijection with 3D Young diagram which are contained within an $a \times b \times c$ box. To see why, imagine viewing a 3D Young diagram from a distance(i.e., under isometric projection). The faces of the 3D diagram are then rhombi, each of which is composed of two equilateral triangles (see Figure 2). The centers of these triangles fall at the vertices of $H_{a,b,c}$, so we get a matching by replacing each rhombus with the corresponding edge.

A weighting of a graph assigns a monomial to each of the graph's edges. Our first example, called a *monochromatic weighting*, is the one depicted in figure 3, which assigns a weight of one to all non-horizontal edges, and weight p^i to horizontal edges which have i other horizontal edges directly below them. We will adopt the convention that if an edge has no weight written beside it, then that edge has a weight of one.

If our graph has a weighting, then the weight of a matching is the product of all of the weights of the edges of the graph.

Now, we have defined two monochromatic weightings – one for 3D Young diagrams and one for graphs. One might ask whether they are "the same": is the weight of a 3D diagram which fits inside an $i \times j \times k$ box the same as the weight of the associated matching of $H_{i,j,k}$? Strictly speaking, the answer is no, because the weight of the empty 3D Young diagram should be one, but the weight of the associated matching (see Figure 5) is not equal to one.

However, given a matching M, we can define its *normalized weight* as the weight of M divided by the weight of the empty 3D Young diagram. Now it is easy to see that the normalized weight of M is equal to the weight of the associated 3D diagram π. The proof is by induction on the number n of boxes in π, the case $n = 0$ being trivial. One only needs to check that the operation of adding a box to π has the effect of increasing both weights by p, which is easy to do.

Figure 3: A monochromatic weighting on $H_{4,4,4}$

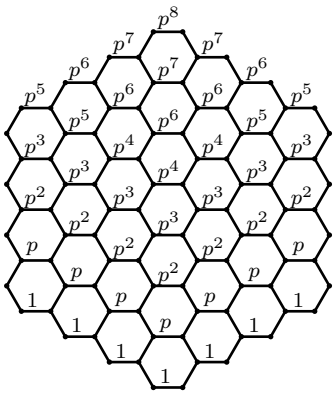

There are many other weightings on the hexagon meshes whose normalized versions are equivalent to the monochromatic weight. For example, one could rotate the weighted mesh by 120 degrees. Indeed, for our purposes, the following weighting is superior (see Figure 4): we replace p with t, and superimpose three copies of the old monochromatic weighting: one normal, one rotated 120 degrees, and one weighted 240 degrees. This weighting assigns each box the weight t^3, so substituting $p = t^3$ gives a new monochromatic weighting.

Definition 4 *The weighting described above is called w_p.*

It is cumbersome to draw diagrams with $p^{1/3}$ as an edge weight, so throughout the paper, we shall use the convention that $p = t^3$ when it is convenient.

We would next like to define a weighting of $H_{a,b,c}$ whose normalized version is equivalent to the $\mathbb{Z}_2 \times \mathbb{Z}_2$ weighting. There are several ways of doing this, but for our purposes, the best way is to first define three weightings: the qt–, rt–, and st–weightings (see Figure 6). The qt–weighting is equivalent to the $\mathbb{Z}_2 \times \mathbb{Z}_2$ weighting under the specialization $p \mapsto t; r, s \mapsto 1$, and similarly for the other two weightings.

Having done this, we construct the $\mathbb{Z}_2 \times \mathbb{Z}_2$ weighting by assigning each edge in $H_{a,b,c}$ the product of its qt,rt, and st–weights. This weights boxes colored q, r, s correctly, but each box in the p position gets the weight t^3. So specializing $t \mapsto p^{1/3}$ gives the $\mathbb{Z}_2 \times \mathbb{Z}_2$ weighting (see Figure 7). Observe that we have assigned weight 1 to all of the grey edges.

Definition 5 *We call the weighting of Figure 7 $w_{p,q,r,s}$.*

3 Overlaying pairs of matchings

We were introduced to the ideas in this section by Kuo's beautiful paper on graphical condensation [3]. In the following, G will always be a bipartite graph.

Figure 4: A "better" monochromatic weighting, w_p, where $p = t^3$.

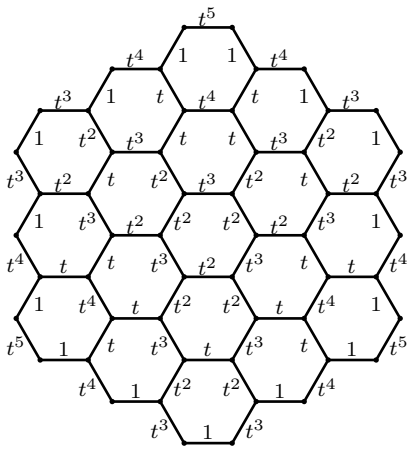

Figure 5: The empty 3D Young diagram and its associated w_p-weighted matching

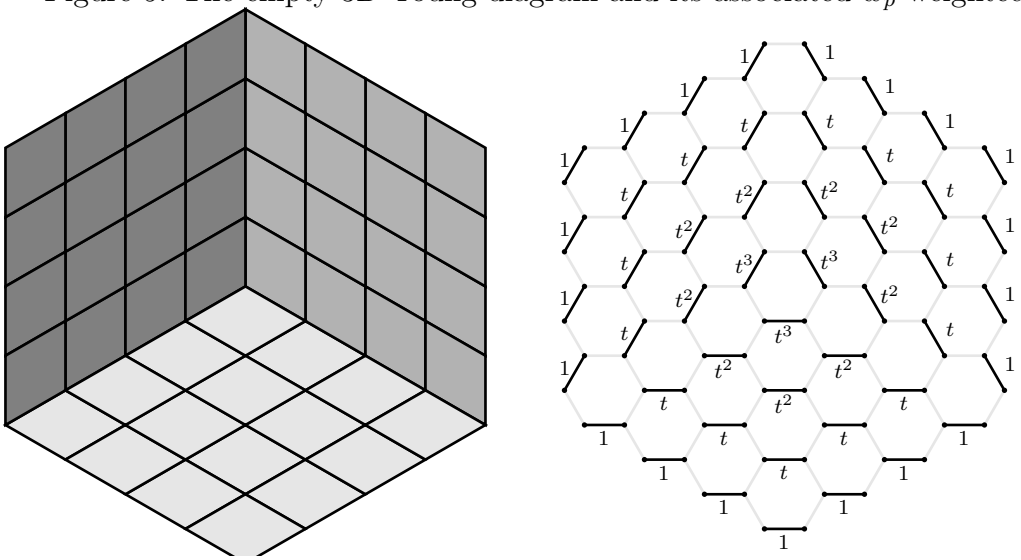

Figure 6: The qt–, rt–, and st–weightings on $H_{4,4,4}$.

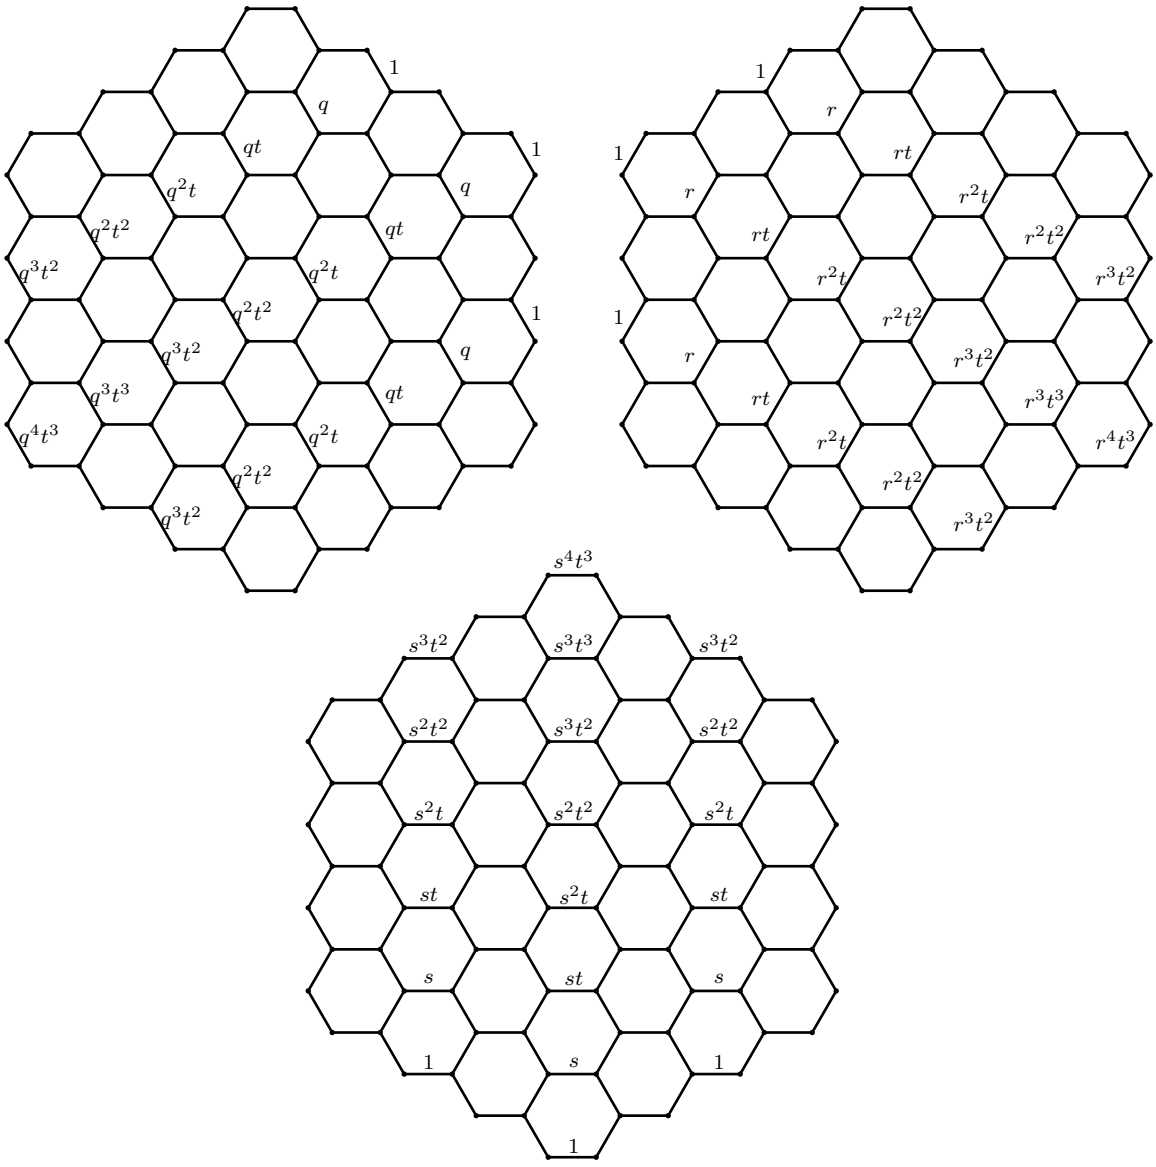

Figure 7: The $\mathbb{Z}_2 \times \mathbb{Z}_2$ weighting.

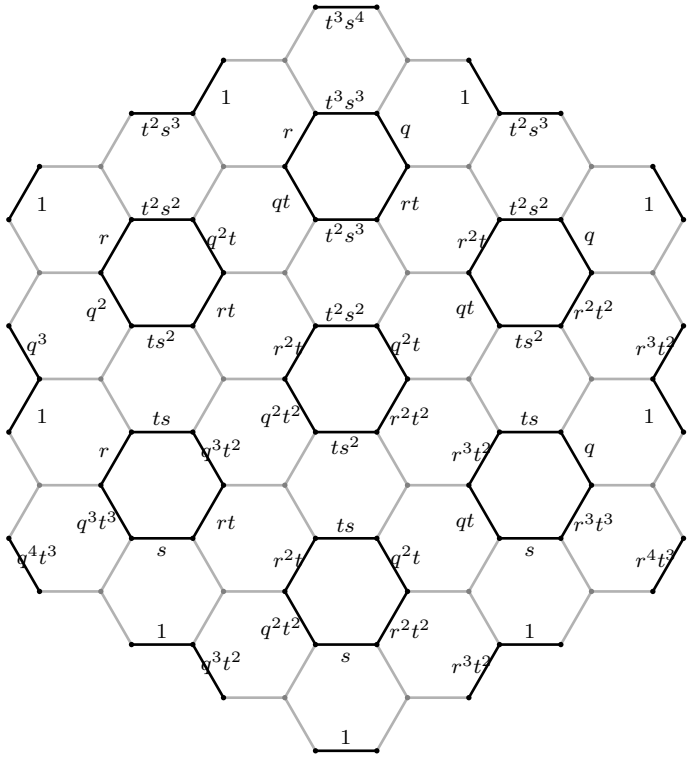

Figure 8: Overlaying two matchings on $H_{3,3,3}$.

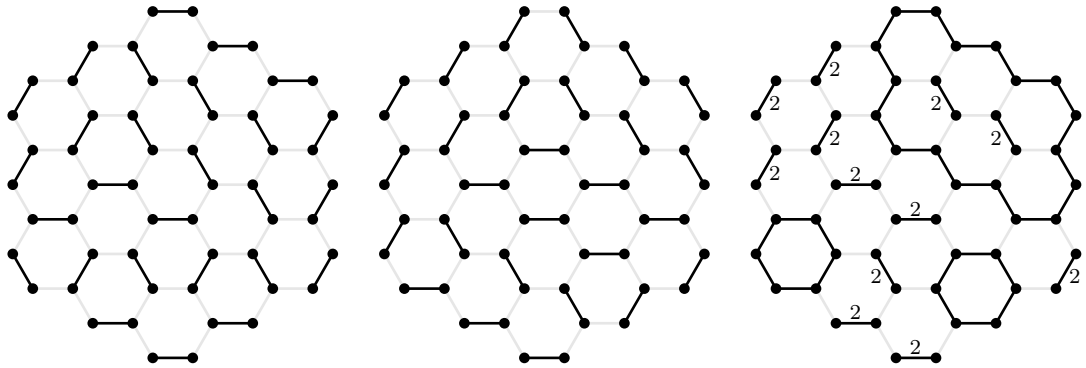

Suppose that we have two matchings M_1, M_2 of G. If we overlay these two matchings on the vertex set V of G, we have a multigraph N in which each vertex has degree two, called a *2-factor* of G. This terminology is slightly nonstandard in graph theory – elsewhere in the literature, a 2-factor is usually a collection of closed loops and isolated edges (not doubled).

If the edge e occurs in both M_1 and M_2, then e occurs as a doubled edge in N. In this case, since the degree of both endpoints of e is two, e is a connected component in N. Conversely, all doubled edges in N must occur in both M_1 and M_2. If we disregard the doubled edges, the rest of N decomposes into a collection of disjoint closed paths.

Conversely, one can split a 2–factor into two one-factors:

Lemma 6 *A 2-factor N may be partitioned into an ordered pair of matchings (M_1, M_2) in precisely $2^{\#\{\text{closed paths in } N\}}$ distinct ways.*

Proof. Suppose we have a 2–factor N of a bipartite graph G. We may obtain two matchings M_1 and M_2 of G as follows: If e is doubled in N, then place e into both M_1 and M_2. If P is a closed path in N, select one of the edges in P and place it into M_1. Place the next edge in the path into M_2, and so forth. Since G is bipartite, the path P is of even length, so each vertex in P has degree 1 in both M_1 and M_2.

There are 2 ways to divide P between M_1 and M_2, so there are $2^{\#\{\text{closed paths in } N\}}$ pairs of matchings M_1, M_2 which correspond to N. □

4 Squishing

Consider the hexagonal mesh with even side lengths, $H_{2a,2b,2c}$. The leftmost diagram in Figure 9 shows a picture of $H_{4,4,4}$, with some of the edges colored grey. The grey edges come in sets of three, all of which are incident to one central vertex. I shall call these three edges a *propeller*.

Figure 9: The hexagonal mesh $H_{4,4,4}$ being squished.

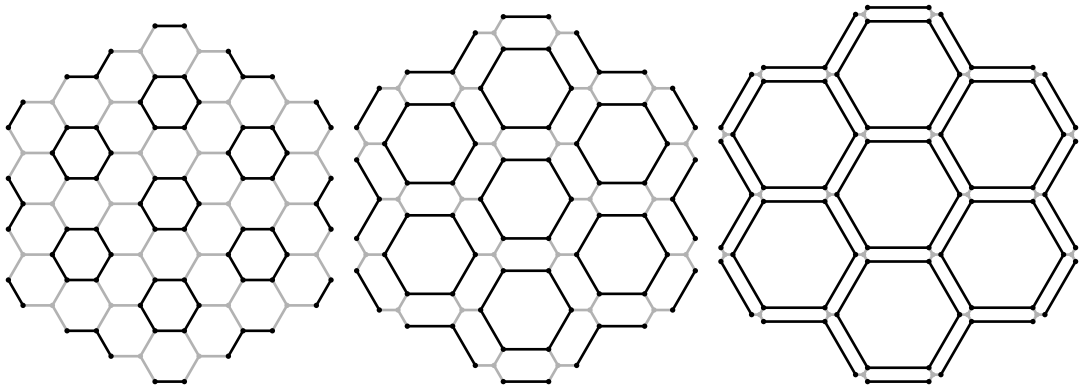

The other two diagrams in Figure 9 show what happens when the length of the edges in each propeller is decreased, while the other edges remain long. We call this procedure *squishing*. Of course, the length that we choose to draw the edges in the graph has no bearing of the structure of the graph, but when the propellers are quite small, the graph of $H_{4,4,4}$ "looks like" the graph of $H_{2,2,2}$ with each edge doubled.

We will denote this squishing operation by the symbol ψ :

$$\psi : H_{2a,2b,2c} \setminus \{\text{propellers}\} \longrightarrow H_{a,b,c}$$

where ψ sends each edge to its image after squishing. It is ofted useful to speak of squishing a set E of edges of $H_{2a,2b,2c}$, and we shall also denote this operation by ψ:

$$\psi(E) = \bigcup_{e \in E \setminus \{\text{propellers}\}} \psi(e).$$

Sometimes, given $E' \subset H_{a,b,c}$, we will need to look for sets of edges E for which $\psi(E) = E'$. Naturally, there are many such E, since ψ ignores the propellers and is two-to-one on all other edges. However, for a given $E' \subset H_{a,b,c}$, there is a "most relevant" preimage of E' under ψ, which we shall call $\varphi(E') \subset H_{2a,2b,2c}$, defined as follows:

$$E_0 := \bigcup_{e \in E'} \phi^{-1}(e),$$
$$\varphi(E') := E_0 \cup \{\text{all propellers adjacent to } E_0\}$$

In other words, $\varphi(E')$ contains all edges which squish to edges of E', as well as all propellers incident to those edges (See Figure 10). It is clear that $\psi \circ \varphi(E') = E'$.

Figure 10: The "unsquishing" map φ

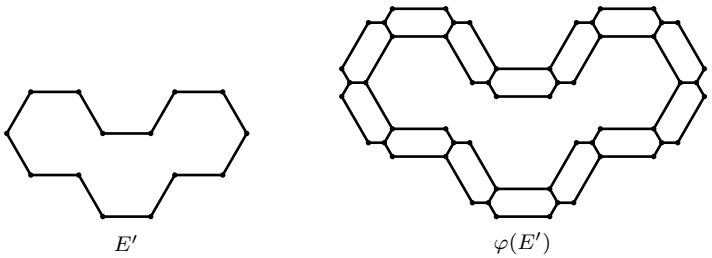

Figure 11: A matching on $H_{4,4,4}$ being squished to a 2–factor on $H_{2,2,2}$.

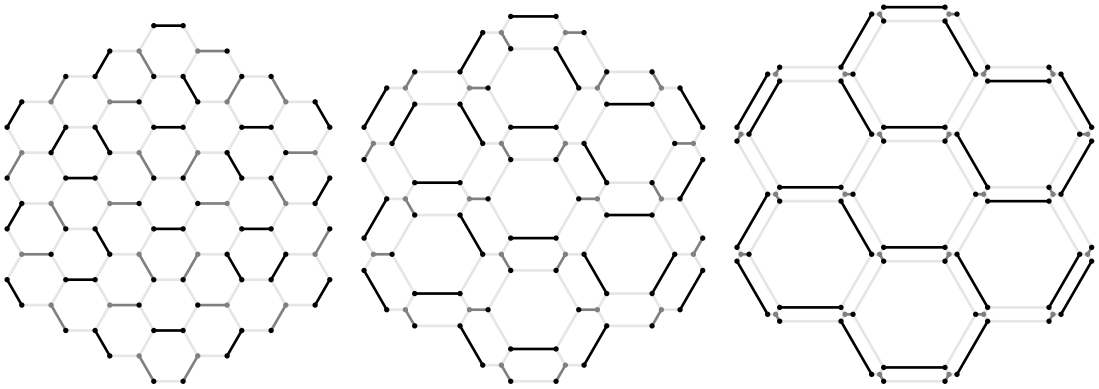

5 Three lemmas on weightings of squished graphs

When one draws a matching on the graph of $H_{4,4,4}$ before squishing, we get what looks very much like a 2-factor of $H_{2,2,2}$ (See Figure 11). We will quantify precisely how this occurs, and use the effect to prove several facts about certain specializations of the $\mathbb{Z}_2 \times \mathbb{Z}_2$ partition function.

Let us examine one of the propellers more closely, and determine what possible configurations of its edges can appear in a perfect matching M, up to symmetry. (see Figure 12). We label the central vertex D, and the other vertices A, B, C. First, observe that preciesly one of the three "short" edges must be included in M, to give the central vertex degree one. Suppose that it is the short horizontal edge CD. In order for vertices A and B to have degree 1 as well, they must each have one incident long edge. Up to symmetry, there are three ways for this to occur:

1. All three edges are horizontal;

2. The edge incident to A is horizontal; the edge incident to B is diagonal; or

Figure 12: All possible configurations of edges around a propeller in a perfect matching.

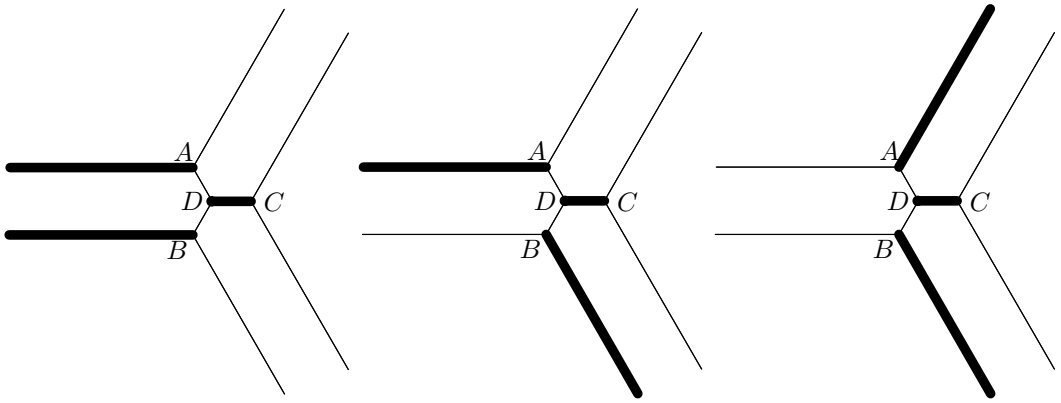

3. The edges incident to both A and B are diagonal.

If we now identify the vertices A, B, C, D and ignore the short edges, we can see that case 1 gives rise to a doubled edge, and cases 2 and 3 give rise to a path through the vertex. Therefore, collapsing all propellers in $H_{2a,2b,2c}$ to points transforms a matching on $H_{2a,2b,2c}$ into a 2-factor on $H_{a,b,c}$. In other words, the squishing operation ψ induces a map, which we will call Ψ:

$$\Psi : \{\text{1-factors of } H_{2a,2b,2c}\} \longrightarrow \{\text{2-factors of } H_{a,b,c}\}.$$

It is easy to see that Ψ is surjective but not injective. For example, $H_{2,2,2}$ has 20 perfect matchings, whereas $H_{1,1,1}$ has only 3 2-factors.

It is often necessary to find the set of all matchings on $\varphi(E')$. Strictly speaking, this set is $\left(\Psi|_{\varphi(E')}\right)^{-1}(E')$, but for simplicity of notation we shall just write this as $\Psi^{-1}(E')$.

We will need to use the $\mathbb{Z}_2 \times \mathbb{Z}_2$ weighting under the specialization $t, q, r, s \mapsto -1$, which we shall call w_{-1} for short (see Figure 13).

Lemma 7 *Let λ be a 2-factor of $H_{a,b,c}$. Then*

$$\sum_{\mu \in \Psi^{-1}(\lambda)} w_{-1}(\mu) = (-1)^{ab+bc+ca} \cdot 2^{\#\{\textit{closed paths in } \lambda\}}.$$

Proof. Let us suppose that λ decomposes as the disjoint union of doubled edges $e_1, \ldots, e_n$ and closed loops $\ell_1, \ldots, \ell_m$. Because the unsquishings $\varphi(e_i)$ and $\varphi(\ell_j)$ are all pairwise non-adjacent in $H_{2a,2b,2c}$, it is clear that

$$\sum_{\mu \in \Psi^{-1}(\lambda)} w_{-1}(\mu) = \prod_{i=1}^{n} \left(\sum_{\mu \text{ matching on } \varphi(e_i)} w_{-1}(\mu) \right) \prod_{j=1}^{m} \left(\sum_{\mu \text{ matching on } \varphi(\ell_j)} w_{-1}(\mu) \right). \qquad (4)$$

Figure 13: The $\mathbb{Z}_2 \times \mathbb{Z}_2$ weighting, under the specialization $t, q, r, s \mapsto -1$

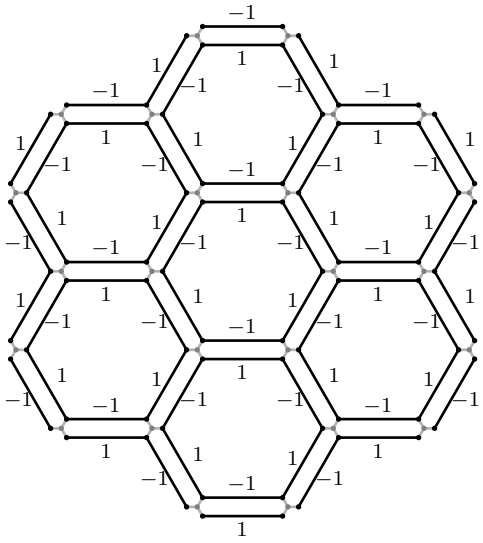

Figure 14: The unique configuration of edges which squishes to a doubled edge.

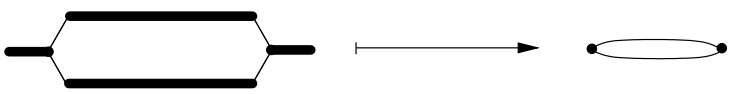

Figure 15: Lifting a closed path ℓ_j to $\varphi(\ell_j)$, and completing it to a matching.

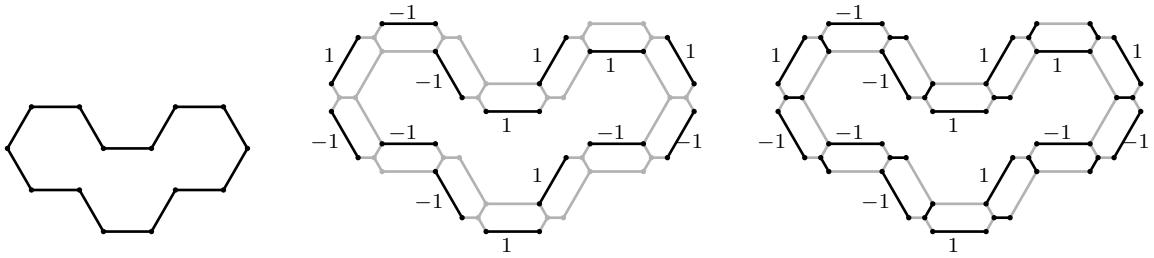

Let us first consider $\Psi^{-1}(e_i)$, the preimage of a doubled edge. There is only one possible matching on $\phi(e_i)$ (see Figure 14); it has weight -1.

Next, let us consider $\Psi^{-1}(\ell_j)$, the preimage of a closed loop. Every term in $\Psi^{-1}(\ell_j)$ is a matching on $\varphi(\ell_j)$ which maps to ℓ_j under ψ. We shall "lift" ℓ_j up to $\varphi(\ell_j)$, so that at each step we choose an edge $e' \in \varphi(\ell_j)$ which maps to e under ψ. If we do this in such a way that no two of the e' are adjacent, then there is a unique way to complete the lifted walk to a matching on $\varphi(\ell_j)$, by choosing exactly one edge of each propeller to be in the matching (see Figure 15). Note that the w_{-1}-weight of the lifting (second image) is the same as the weight of the matching (third image) since all of the propellers have weight 1.

Suppose that we walk along the lifted path in $\varphi(\ell_j)$, in the counterclockwise (positive) direction, and we step on the edges $e'_1, e'_2, \cdots, e'_k, e'_1$, in order. We keep track of the product of the weights of the edges we have stepped on.

At any point in our walk, there are essentially only four states that we can be in. We can be standing on an edge weighted 1 or -1, and we can be facing one of two types of propeller: $\succ\!\!-$ or $-\!\!\prec$. Let us label the states 1 through 4, as follows:

1. Standing on a weight 1 edge, facing a $\succ\!\!-$ propeller
2. Standing on a weight -1 edge, facing a $\succ\!\!-$ propeller
3. Standing on a weight 1 edge, facing a $-\!\!\prec$ propeller
4. Standing on a weight -1 edge, facing a $-\!\!\prec$ propeller

To get from one edge to another in the path, we must turn either left or right at the propeller we are facing, so we shall think of our walk as a sequence of Rs and Ls, read *right-to-left*. For example, the walk in Figure 15, starting at the upper leftmost horizontal edge and proceeding counterclockwise, is $LRRLLLLRLLRLLL$.

Our trick is to interpret R and L as the following *state–transition matrices*:

$$L = \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 0 \\ -1 & -1 & 0 & 0 \end{bmatrix} \qquad R = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & -1 & -1 \\ 1 & 1 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{bmatrix}$$

These matrices encode which states can be reached from which, and keep track of the total weight of all paths leading from one state to another. For example, at the upper-left most horizontal edge in Figure 15, we are in state 4, and we are about to turn to the left. We could go into state 1, picking up a weight of 1 (as in the diagram); alternately, we could go into state 2, picking up a weight of -1. This is modeled by our state-transition matrices, since

$$L \cdot \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 0 \\ -1 & -1 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ -1 \end{bmatrix}$$

By the same token, the coefficients of the matrix LR are the sums of the weights of all paths which start in state i , turn first right and then left, and end in state j. If we want to know the sum of the weights of the lifts of the closed loop in Figure 15, we should compute the matrix product $LRRLLLLRLLRLLL$. We must end in the same state as we start since our path is closed. Furthermore, our starting state must be either 3 or 4, otherwise we would be walking clockwise rather than counterclockwise. It follows that the total number of paths is the sum of the (3,3) and (4,4) entries of this matrix product.

Now, here is the punchline: It is easy to check that $L = R^{-1}$. In any closed loop on the hexagonal lattice, there must be 6 more left turns than right turns, so the combined weights of all lifts of any closed loop is the sum of the $(3, 3)$ and $(4, 4)$ entries of L^6. Now, it is also easy to check that $L^6 = -I$, so this combined weight is -2, regardless of the shape or orientation of the loop ℓ_j.

Going back to Equation 4, and recalling that there are i doubled edges and j closed loops in the 2-factor λ, we can now state that

$$\sum_{\mu \in \Psi^1(\lambda)} w_{-1}(\mu) = (-1)^i \cdot (-2^j) = (-1)^{i+j} \cdot 2^j$$

Observe that $i + j$ is the total number of connected components in λ, so the proof is complete modulo the following lemma. □

Lemma 8 *Let λ be a 2-factor of $H_{a,b,c}$. Let $C(\lambda)$ denote the number of connected components of λ. Then $C(\lambda) \equiv ab + bc + ca \pmod 2$.*

Proof. Decompose λ into two matchings, λ_1 and λ_2. Let us consider λ_1 for the moment.

When one removes (or adds) a cube to the 3D diagram corresponding to λ_1, one obtains a new matching. Graph-theoretically, this is equivalent to performing the change depicted in Figure 16. Let us call such an operation a *τ-transformation*. Since it is possible to change any 3D diagram into any other by adding and removing boxes one at a time, repeated τ-transformations will allow us to change any matching into any other matching.

Let us alter the matchings λ_1 and λ_2 in this fashion, so that after finitely many steps, we have obtained two copies of the empty 3D Young diagram. At each stage, superimposing

Figure 16: A τ-transformaion

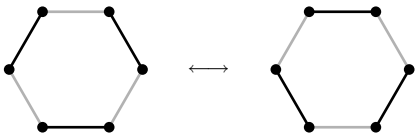

Figure 17: Changing one 2-factor into another with τ-transformations

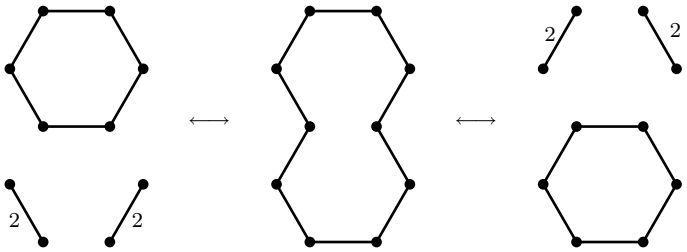

the two matchings gives a 2-factor, so we may think of τ-transformations as acting on the set of 2-factors as well (see Figure 17).

Next, we show that τ-transformations leave $C(\lambda)$ invariant mod 2. Suppose we are applying the τ transformation to λ_1, and we are affecting the edges around a particular hexagon H. Remove all edges of H from λ_1 to obtain a new subgraph λ_1° of $H_{a,b,c}$. Now superimpose λ_1° and λ_2. The resulting graph has degree 2 everywhere except at vertices of H, all of which have degree 1. As a result, every vertex of H is connected to precisely one other vertex of H by some path (possibly of length one). Up to symmetry, there are two possible ways in which this may occur (see Figure 18). When we add the edges of $\lambda_1 \cap H$ back into the graph and perform the τ-transformation, it is easy to see that the parity of the number of connected components is preserved. In the upper case, there are an odd number of path components; in the lower case, there are an even number.

Up to this point, we have shown that $C(\lambda)$ is an invariant of $H_{a,b,c}$. To calculate it, we compute the number of edges in the matching corresponding to the empty 3D Young diagram. This is easy: there are ab horizontal edges, ac edges of slope $\sqrt{3}$, and bc edges of slope $-\sqrt{3}$, and so there are $ab + bc + ca$ edges in all. □

Lemma 9 *If λ is a 2-factor on $H_{a,b,c}$ and $\mu \in \Psi^{-1}(\lambda)$, then $w_{p,1,1,1}(\lambda) = w_p(\mu)$.*

Proof. Comparing the weightings $w_{p,1,1,1}$ and w_p in Figure 19, it is clear that if e is an edge of $H_{a,b,c}$ and e' is either of the two edges in $\psi^{-1}(e)$, then $w_{p,1,1,1}(e') = w_p(e)$. The statement of the lemma follows immediately. □.

Figure 18: The two possible connectivities of $\lambda_1^\circ \cup \lambda_2$, and their behaviour under τ-transformations

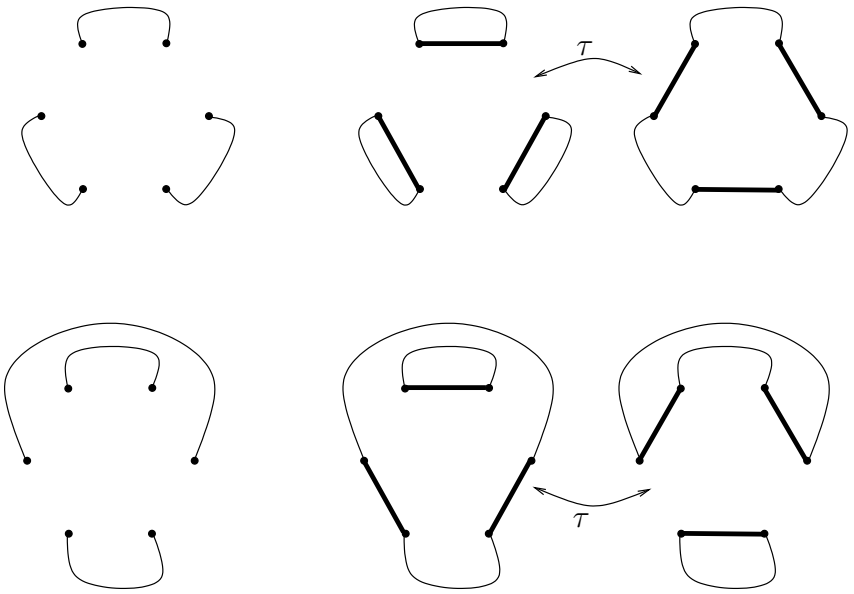

Figure 19: The weightings $w_{p,1,1,1}$ on $H_{2a,2b,2c}$, compared with the weighting w_p on $H_{a,b,c}$, where $t = p^3$.

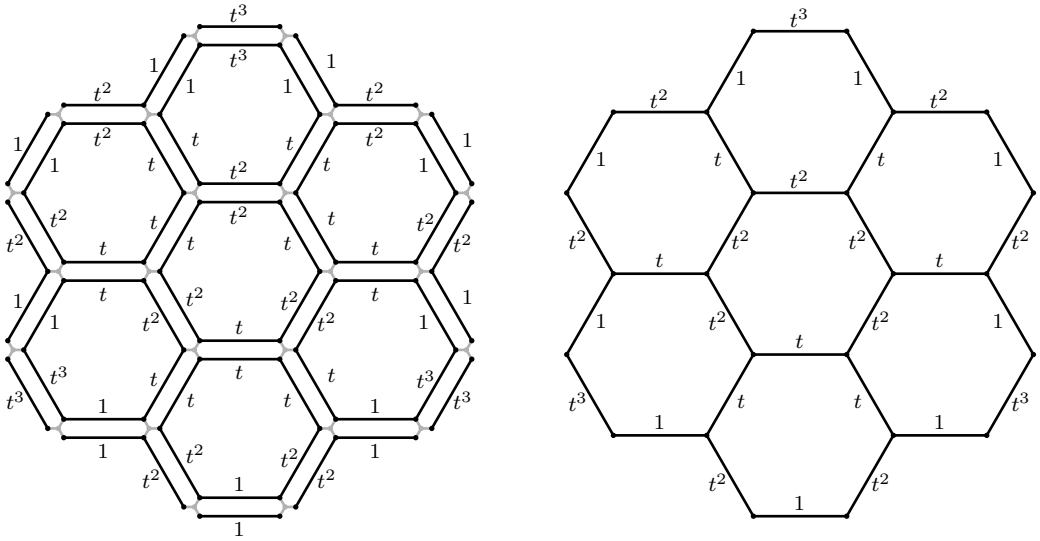

6 A specialization of the $\mathbb{Z}_2 \times \mathbb{Z}_2$ partition function

As before, $w_{p,q,r,s}$ denotes the $\mathbb{Z}_2 \times \mathbb{Z}_2$ weighting, and $w_p = w_{p,p,p,p}$ denotes the monochromatic weighting. We will be working with various specializations of these weightings, and it will be convenient to write, for example, $w_{-p,-1,-1,-1}$ for the specialization $p \mapsto -p$, $q, r, s \mapsto -1$. /usr/X11R6/bin/gvim

Although our tool of choice is weighted matchings on hexagon meshes, our object of study is the $\mathbb{Z}_2 \times \mathbb{Z}_2$ weighting for 3D Young diagrams. We shall use $w_{p,q,r,s}(\pi)$ to denote the $\mathbb{Z}_2 \times \mathbb{Z}_2$–weight of the 3D diagram π, and w_p to denote its monochromatic weight. Recall that if λ is the matching corresponding to π, then

$$w_{p,q,r,s}(\pi) = \frac{w_{p,q,r,s}(\lambda)}{w_{p,q,r,s}(\text{empty 3D diagram})}.$$

and we can normalize the monochromatic weight in the same fashion.

Definition 10 *The $\mathbb{Z}_2 \times \mathbb{Z}_2$ partition function for $H_{a,b,c}$ is*

$$Z^{a,b,c}(p,q,r,s) = \sum_{\substack{\pi \text{ 3d diagram} \\ \text{inside } a\times b\times c \text{ box}}} w_{p,q,r,s}(\pi).$$

The monochromatic partition function for $H_{a,b,c}$ is

$$Z^{a,b,c}(p) = \sum_{\substack{\pi \text{ 3d diagram} \\ \text{inside } a\times b\times c \text{ box}}} w_p(\pi)$$

Note that $\lim_{a,b,c\to\infty} Z^{a,b,c} = \lim_{a,b,c\to\infty} Z^{2a,2b,2c} = Z_{\mathbb{Z}_2\times\mathbb{Z}_2}$. Thus the following implies (3):

Theorem 11

$$Z^{2a,2b,2c}(p,-1,-1,-1) = \left(Z^{a,b,c}(-p)\right)^2.$$

Proof. To begin with, let us work with matchings. Observe that for any matching μ, we have $w_{p,-1,-1,-1}(\mu) = w_{-p,1,1,1}(\mu) \cdot w_{-1,-1,-1,-1}(\mu)$. Therefore,

$$\sum_{\substack{\mu \text{ matching} \\ \text{on } H_{2a,2b,2c}}} w_{p,-1,-1,-1}(\mu) = \sum_{\substack{\lambda \text{ 2-factor} \\ \text{on } H_{a,b,c}}} \sum_{\mu \in \Psi^{-1}(\lambda)} w_{-p,1,1,1}(\mu) w_{-1,-1,-1,-1}(\mu)$$

$$= \sum_{\substack{\lambda \text{ 2-factor} \\ \text{on } H_{a,b,c}}} w_{-p}(\lambda) \sum_{\mu \in \Psi^{-1}(\lambda)} w_{-1,-1,-1,-1}(\mu) \tag{5}$$

$$= (-1)^{ab+bc+ca} \sum_{\substack{\lambda \text{ 2-factor} \\ \text{on } H_{a,b,c}}} 2^{\#\{\text{closed loops in } \lambda\}} w_{-p}(\lambda) \tag{6}$$

$$= (-1)^{ab+bc+ca} \sum_{\substack{\lambda,\eta \text{ 1-factors} \\ on H_{a,b,c}}} w_{-p}(\lambda) w_{-p}(\eta) \tag{7}$$

$$= (-1)^{ab+bc+ca} \left(\sum_{\substack{\mu \text{ matching} \\ \text{on } H_{a,b,c}}} w_{-p}(\mu) \right)^2$$

In the above sequence of steps, Equation (5) uses Lemma 9 and Equation (6) uses Lemma 7.

Now, let us normalize both sides by dividing by the $w_{p,-1,-1,-1}$ weight of the matching associated to the $(2a, 2b, 2c)$ empty 3D diagram. Let us call this matching M. Call the (a, b, c) empty 3D diagram N.

On the left side, we get $Z^{a,b,c}(p, q, r, s)$. For the right side, we observe that

$$\begin{aligned} w_{p,-1,-1,-1}(M) &= w_{-1,-1,-1,-1}(M) \cdot w_{-p,1,1,1}(M) \\ &= (-1)^{ab+bc+ca} \cdot (w_{-p}(N))^2, \end{aligned}$$

again using Lemma 9. Therefore, the right-hand side normalizes to

$$\left(Z^{a,b,c}(-p)\right)^2.$$

and we are done. □

7 Conclusions

It is obvious to ask whether this method might be used for its intended purpose, namely to prove Equation 2. Unfortunately, the answer is no. Our method is clearly agnostic as to the large-scale shape of the graph; it works on any suitable subset of the hexagon lattice. However, computer calculations show that the analogue of (2) on $H_{2a,2b,2c}$ does not admit a product formula; rather, if one counts 3D Young diagrams fitting inside a $2a \times 2b \times 2c$ box, according to the $\mathbb{Z}_2 \times \mathbb{Z}_2$ weighting, one typically gets $(1 - p) \times$ (a large irreducible polynomial); the product formula only appears as $a, b, c \to \infty$.

It should be possible to apply this method to other graphs and other weights. We have not investigated this possibility in any great depth, but in order for the method to work without serious modification, it seems that the graph ought to have propeller-like vertices, at which several hexagonal faces meet; furthermore, one might expect the graph to have some degree of self-similarity at different scales.

Acknowledgements

We would like to thank Jim Bryan for editing and proofreading, and Rick Kenyon for the idea of the matrices L and R in section 5.

References

[1] Benjamin Young, with an appendix by Jim Bryan, *Generating functions for colored 3D Young diagrams and the Donaldson-Thomas invariants of orbifolds*, Submitted. arXiv:math/0802.3948.

[2] Richard Kenyon, *The planar dimer model with boundary: a survey*, Directions in mathematical quasicrystals, CRM Mathematical Monographs, vol. 13, American Math. Soc., 2004, pp. 29 – 57.

[3] Eric H. Kuo, *Applications of graphical condensation for enumerating matchings and tilings*, Theoretical Computer Science **319** (2004), no. 1, 29 – 57.

[4] Percy A. MacMahon, *Combinatory analysis*, Cambridge University Press, The Edinburgh Building, Cambridge, UK, 1915-16.

[5] Andrei Okounkov, Nikolai Reshetikhin, and Cumrun Vafa, *Quantum Calabi-Yau and classical crystals*, Progress in Mathematics **244** (2006), 597–618.