Discrete Spin Structures, The Dimer Model, and Dual Graphs

Samuel Edward Blyth

Summer Research Project, University of Auckland
Supervised by Pedram Hekmati

April 3, 2026

Abstract

The partition function for the dimer model on a surface graph is known to be calculable from a sum of Pfaffians of Kasteleyn matrices. We follow the derivation of the formula given in [1], and apply it to explicitly calculate the partition function for the graph $K_{3,3}$. We then examine how Kasteleyn orientations can be viewed as a discrete formulation of a spin structure, and introduce another discrete formulation of spin structures given in [2]. Finally, we describe how one can convert between these two formulations by taking the dual of the graph.

1 The Dimer Model

1.1 The Partition Function

The dimer model concerns the absorption of a diatomic gas onto a crystal lattice. It asks how many ways there are in which the vertices of the lattice can be covered by non-overlapping dimer molecules, and seeks to quantify the relative probabilities of these configurations.

Given a finite connected graph Γ, a *perfect matching* or *dimer configuration* on Γ is a subset D of the edges of Γ such that each vertex of Γ is adjacent to precisely one edge in D. If a dimer configuration exists, it is clear that Γ must have an even number of vertices, however this not a sufficient condition. Take for example a Y-shaped graph, which has 4 vertices but does not admit a dimer configuration. Unless stated otherwise, we will assume that any graph we consider admits a dimer configuration.

An example of a graph which admits a dimer configuration is given in Figure 1, which we shall refer to as the lambda graph, for it contains that shape on the right hand side. One can simply see that the lambda graph admits only two different dimer configurations, as shown in Figure 2, where the dimer edges are drawn in bold.

Given an arbitrary graph Γ, with edge set $E(\Gamma)$ and vertex set $V(\Gamma)$, we may wish to determine how many different dimer configurations are possible. To this end, we introduce an *edge weight system* on Γ, that is, a function $\nu : E(\Gamma) \to \mathbb{R}^+$ which assigns a weight to each edge of Γ. These weightings may represent how likely a dimer is to be absorbed across that edge. Hence we derive a probability distribution on the set $\mathcal{D}(\Gamma)$ of possible dimer configurations,

$$\mathsf{P}(D) = \frac{\nu(D)}{Z(\Gamma;\nu)},$$

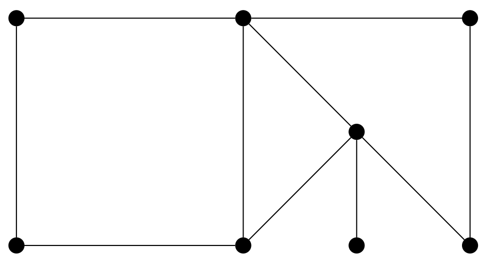

Figure 1: The lambda graph.

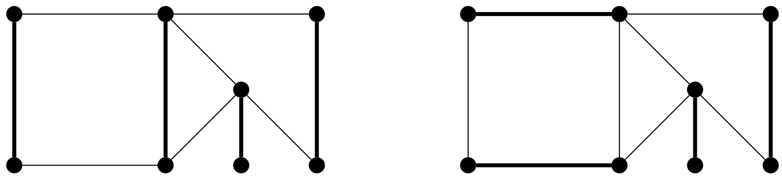

Figure 2: Dimer configurations of the lambda graph.

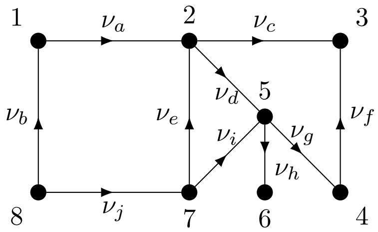

$$\mathsf{A}^K = \begin{pmatrix} 0 & \nu_a & 0 & 0 & 0 & 0 & 0 & -\nu_b \\ -\nu_a & 0 & \nu_c & 0 & \nu_d & 0 & -\nu_e & 0 \\ 0 & -\nu_c & 0 & -\nu_f & 0 & 0 & 0 & 0 \\ 0 & 0 & \nu_f & 0 & -\nu_g & 0 & 0 & 0 \\ 0 & -\nu_d & 0 & \nu_g & 0 & \nu_h & -\nu_i & 0 \\ 0 & 0 & 0 & 0 & -\nu_h & 0 & 0 & 0 \\ 0 & \nu_e & 0 & 0 & \nu_i & 0 & 0 & -\nu_j \\ \nu_b & 0 & 0 & 0 & 0 & 0 & \nu_j & 0 \end{pmatrix}$$

Figure 3: An arbitrary weighted orientation of the lambda graph, and the associated adjacency matrix.

where $\nu(D) = \prod_{e\in D} \nu(e)$ is the weight of a given dimer configuration, and

$$Z(\Gamma;\nu) = \sum_{D\in\mathcal{D}(\Gamma)} \nu(D) \tag{1}$$

is called the *partition function* of Γ under the weight system ν. In the case where the weight system is equal to the constant 1, we find that the partition function is simply the number of possible dimer configurations on Γ,

$$Z(\Gamma;1) = \sum_{D\in\mathcal{D}(\Gamma)} 1 = |\mathcal{D}(\Gamma)|\,.$$

We would therefore like to obtain an algorithmic method of calculating the partition function for arbitrary graphs. If Γ does not admit a dimer configuration, for example because $|V(\Gamma)|$ is odd, the partition function is trivially equal to 0, since the sum in equation 1 is empty. To calculate the partition function in non-trivial cases, we follow closely the method developed in [3].

1.2 Planar Graphs

To begin we restrict our consideration to planar graphs. In addition to the weight system on Γ, we fix an arbitrary orientation K on the edges of Γ, and enumerate the $2n$ vertices of Γ by $1,\ldots,2n$. Define the adjacency matrix of the weighted oriented graph Γ to be

$$\mathsf{A}^K = \left(\mathsf{a}^K_{ij}\right) = \left(\sum_{e\in E(\Gamma)} \varepsilon^K_{ij}(e)\,\nu(e)\right),$$

where ε^K_{ij} is given by

$$\varepsilon^K_{ij}(e) = \begin{cases} 1 & e \text{ connects } i \text{ to } j \text{ as oriented by } K, \\ -1 & e \text{ connects } j \text{ to } i \text{ as oriented by } K, \\ 0 & \text{otherwise.} \end{cases}$$

By this construction, A^K must be skew-symmetric. The weighting, vertex enumeration, and an arbitrary orientation of the lambda graph, as well as the corresponding adjacency matrix are shown in Figure 3.

Recall that the *Pfaffian* of a skew-symmetric matrix is the square root of its determinant, and is a polynomial in the components of the matrix given by

$$\mathrm{Pf}(\mathsf{A}) = \frac{1}{2^n n!} \sum_{\sigma\in S_{2n}} \mathrm{sgn}(\sigma) \prod_{\ell=1}^{n} \mathsf{a}_{\sigma(2\ell-1)\sigma(2\ell)}.$$

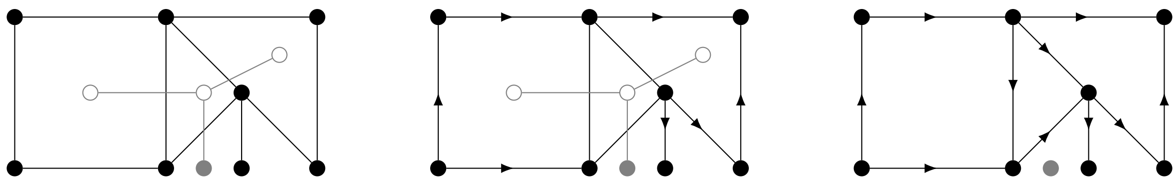

Figure 4: Construction of a Kasteleyn orientation on the lambda graph.

In the case of our adjacency matrix A^K, each permutation induces a matching of $1,\ldots,2n$ into ordered pairs of vertices, $\{\{\sigma(1),\sigma(2)\},\ldots,\{\sigma(2n-1),\sigma(2n)\}\}$. If a matching contains any pair of vertices $\{k,l\}$ which are not connected by an edge in Γ, the component a_{kl} will be zero, so that permutation's whole contribution to the Pfaffian will be zero. Hence we only need to sum over those matchings which give a valid dimer configuration. Furthermore, each matching of $1,\ldots,2n$ is represented by $2^n n!$ permutations, so only one such permutation needs to be taken into account.

Let $D\in\mathcal{D}(\Gamma)$ be a dimer configuration on Γ given by edges $e_1,\ldots,e_n$, and choose any permutation $\sigma_D\in S_{2n}$ which sends $1,\ldots,2n$ to $\{\{i_1,j_1\},\ldots,\{i_n,j_n\}\}$, where $\{i_\ell,j_\ell\}$ are the vertices adjacent to the edge e_ℓ. We define the sign

$$\varepsilon^K(D)=\operatorname{sgn}(\sigma_D)\prod_{\ell=1}^{n}\varepsilon^K_{i_\ell,j_\ell}(e_\ell),$$

and hence find that the Pfaffian of A^K is given by

$$\operatorname{Pf}\left(\mathsf{A}^K\right)=\sum_{D\in\mathcal{D}(\Gamma)}\varepsilon^K(D)\,\nu(D). \tag{2}$$

It now follows that if we can find an orientation K such that $\varepsilon^K(D)$ is independent of the dimer configuration D, then the partition function can be computed directly from the Pfaffian of the adjacency matrix. To achieve this, we define a particular class of orientations as follows.

Embed the graph Γ in the plane, then the faces of Γ inherit the anticlockwise orientation of $\mathbb{R}^2$. Given a face f of Γ, with boundary $\partial f\subset E(\Gamma)$, let $n^K(\partial f)$ denote the number of edges along which the anticlockwise orientation of f differs from K. An orientation is *Kasteleyn* if $n^K(\partial f)$ is odd for each face f of Γ.

We say that a face obeys the *Kasteleyn condition* if $n^K(\partial f)$ is odd, and we call the adjacency matrix A^K of a Kasteleyn orientation a *Kasteleyn matrix*. A graph admits a Kasteleyn orientation if and only if it has an even number of vertices ([1], Theorem 1).

We now describe how to construct Kasteleyn orientations on a graph Γ. First, fix a spanning tree T in Γ^*, the dual graph of Γ, rooted at the vertex of Γ^* representing the outer face of Γ. Orient the edges of Γ which do not intersect T arbitrarily. Now, starting at the leaves of T and moving towards the root, remove the branches of T one at a time, and each time fix the orientation of the edge of Γ which intersected this branch such that the Kasteleyn condition is satisfied in the newly completed face. A demonstration of this process for the lambda graph is shown in Figure 4.

It is shown in [3] (Theorem 3.1) that for planar graphs, the sign $\varepsilon^K(D)$ is independent of D when K is Kasteleyn. Hence, (2) implies that the partition function of a planar graph Γ with a Kasteleyn orientation K is given by

$$Z(\Gamma;\nu)=\left|\operatorname{Pf}\left(\mathsf{A}^K\right)\right|. \tag{3}$$

This allows us to calculate the partition function of a graph without explicitly determining all the possible dimer configurations it admits. We do this now for the lambda graph. The Kasteleyn matrix corresponding to the Kasteleyn orientation obtained in Figure 4 is given by

$$
\mathsf{A}^K = \begin{pmatrix} 0 & \nu_a & 0 & 0 & 0 & 0 & 0 & -\nu_b \\ -\nu_a & 0 & \nu_c & 0 & \nu_d & 0 & \nu_e & 0 \\ 0 & -\nu_c & 0 & -\nu_f & 0 & 0 & 0 & 0 \\ 0 & 0 & \nu_f & 0 & -\nu_g & 0 & 0 & 0 \\ 0 & -\nu_d & 0 & \nu_g & 0 & \nu_h & -\nu_i & 0 \\ 0 & 0 & 0 & 0 & -\nu_h & 0 & 0 & 0 \\ 0 & -\nu_e & 0 & 0 & \nu_i & 0 & 0 & -\nu_j \\ \nu_b & 0 & 0 & 0 & 0 & 0 & \nu_j & 0 \end{pmatrix},
$$

and its determinant is

$$
\begin{aligned}
\det \mathsf{A}^K &= -\nu_h\nu_f\left(\nu_a\nu_f\nu_j\left(-\nu_a\nu_h\nu_j - \nu_e\nu_b\nu_h\right) + \nu_b\nu_e\nu_f\left(-\nu_a\nu_h\nu_j - \nu_e\nu_b\nu_h\right)\right) \\
&= \nu_h\nu_f\left(\nu_a\nu_f\nu_j + \nu_b\nu_e\nu_f\right)\left(\nu_a\nu_h\nu_j + \nu_e\nu_b\nu_h\right) \\
&= \nu_h^2\nu_f^2\left(\nu_a\nu_j + \nu_b\nu_e\right)^2 \\
&= \left(\nu_h\nu_f\left(\nu_a\nu_j + \nu_b\nu_e\right)\right)^2.
\end{aligned}
$$

Its Pfaffian is therefore

$$
\mathrm{Pf}\left(\mathsf{A}^K\right) = \pm\nu_h\nu_f\left(\nu_a\nu_j + \nu_b\nu_e\right).
$$

Although it is not necessary for this calculation of the partition function, it will later be useful to know how to determine the correct sign of the Pfaffian. This is done by explicitly calculating the sign $\varepsilon^K(D)$ for one of the terms in the Pfaffian. In this case we take the term $\nu_h\nu_f\nu_a\nu_j$, and choose any permutation representing this matching of vertices, for example, $\sigma_D = (1,2,4,3,5,6,8,7) = (3\,4)(7\,8)$ in one-line notation and cycle notation respectively. In particular, we have chosen a permutation such that the signs $\varepsilon^K_{12}(e_a) = \varepsilon^K_{43}(e_f) = \varepsilon^K_{56}(e_h) = \varepsilon^K_{87}(e_j) = 1$, so we have that

$$
\varepsilon^K(D) = \mathrm{sgn}(\sigma_D) = (-1)^2 = 1.
$$

Thus, the term $\nu_h\nu_f\nu_a\nu_j$ takes the positive sign, and we have that

$$
\mathrm{Pf}\left(\mathsf{A}^K\right) = \nu_h\nu_f\left(\nu_a\nu_j + \nu_b\nu_e\right).
$$

Finally, we can use (3) to calculate the partition function of the lambda graph,

$$
\begin{aligned}
Z(\Gamma;\nu) &= \left|\mathrm{Pf}\left(\mathsf{A}^K\right)\right| \\
&= \nu_a\nu_f\nu_h\nu_j + \nu_b\nu_e\nu_f\nu_h,
\end{aligned}
$$

where these two terms represent the two possible dimer configurations of the graph shown in Figure 2. In particular, the number of possible dimer configurations is $Z(\Gamma;1) = 2$.

1.3 Surface Graphs

The fact that $\varepsilon^K(D)$ is independent of D for planar graphs relies on the property that all simple closed curves in the plane bound a disc. This property, and therefore (3), does not extend directly to non-planar graphs. In order to obtain a formula for the partition function of non-planar graphs, we must embed any non-planar graphs in a surface of a higher genus than the plane.

Given a graph Γ, it is always possible embed it in a closed surface Σ of a fixed genus g, such that $\Gamma \subset \Sigma$ is the 1-skeleton of a cellular decomposition of Σ. Equivalently, Γ must have no self-intersections on the surface Σ, and g must be minimal. Throughout this section, we will consider $K_{3,3}$, the Utility Graph, embedded in a genus-1 torus, as shown in Figure 5 with one possible dimer configuration indicated.

In order to construct Kasteleyn orientations on non-planar graphs, we will use a topological intuition known as *homology*. The *first homology space*, $H_1(\Sigma;\mathbb{Z}_2)$, of a genus-g surface is a vector space over $\mathbb{Z}_2$ of dimension $2g$, and its vectors are equivalence classes of curves on the surface which traverse the holes in the surface in the same way. Each of the g holes in the surface provides two generators, one going through the hole, and the other travelling around it. Two curves are in the same *homology class* if they travel around every generator the same number of times modulo 2. The four homology classes on the torus can be represented by the cycles $\alpha_0, \alpha_1, \alpha_2, \alpha_3$ on the utility graph as shown in Figure 6.

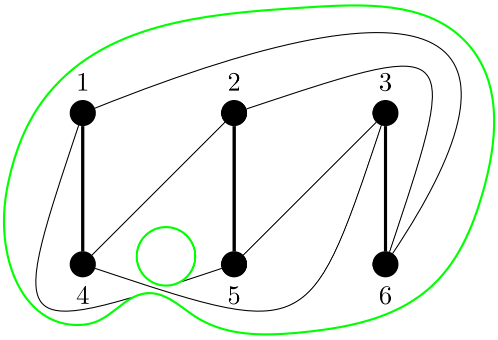

Figure 5: The utility graph embedded in a geometric torus.

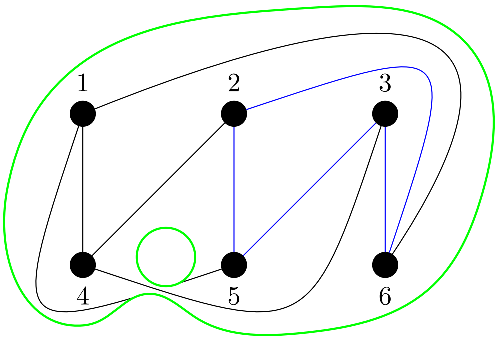

(a) The trivial loop, α_0.

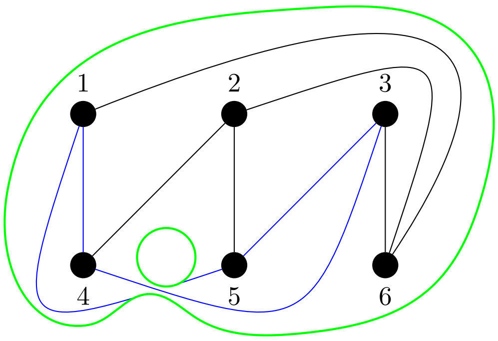

(b) The first generator, α_1.

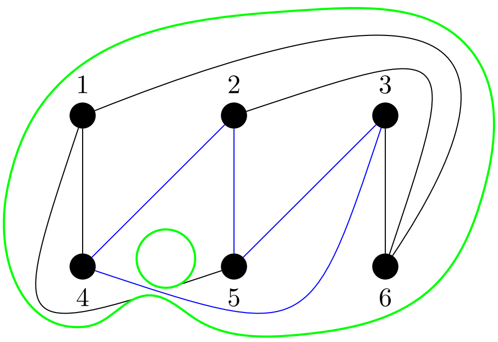

(c) The second generator, α_2.

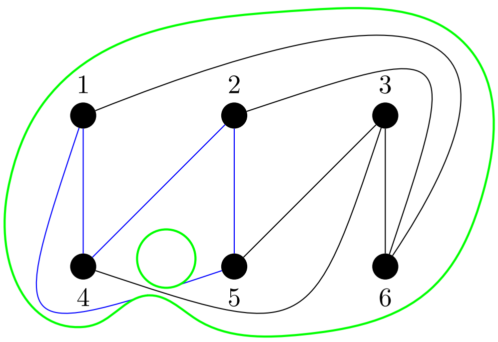

(d) The generators combined, α_3.

Figure 6: Homology classes of the utility graph.

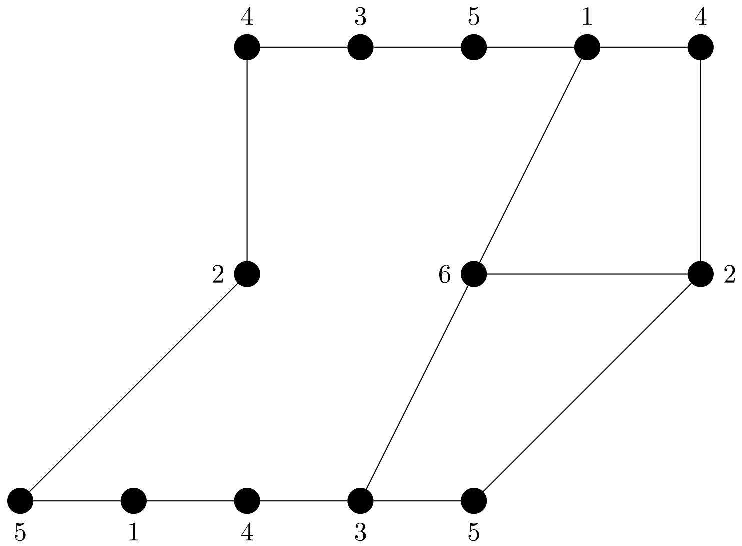

Figure 7: The utility graph cut along the curves α_1 and α_2.

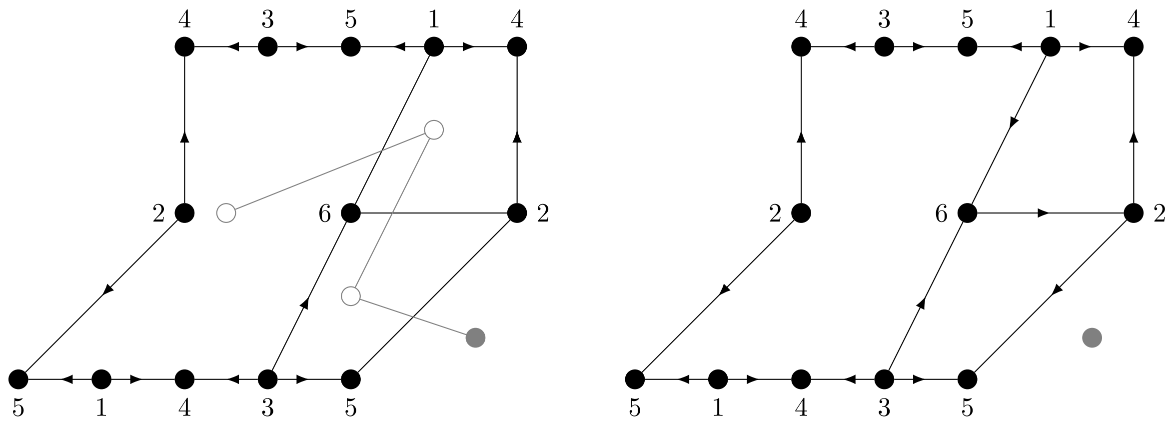

Figure 8: Construction of a Kasteleyn orientation on the utility graph.

In order to extend the idea of a Kasteleyn orientation to surface graphs, we choose a set of simple closed curves in Γ such that when Σ is cut along these curves, it unfolds to a disc $\Sigma' \subset \mathbb{R}^2$. For the utility graph, we cut along the curves $\{\alpha_1, \alpha_2\}$, and obtain the unfolded graph shown in Figure 7.

One can then repeat the process for constructing a Kasteleyn orientation described in Section 1.2 on this planar representation of Γ, with two additional conditions.

1. The spanning tree T must be rooted in the external face of the graph, $\mathbb{R}^2 \setminus \Sigma'$.
2. When choosing the arbitrary edge orientations, edges which are identified in Σ must have the same orientation.

This leaves only one edge in the boundary of Σ' unoriented, and the orientation it receives by the usual process of orienting the remaining edges will be consistent with the second condition, provided that Γ has an even number of vertices. This construction is shown on the utility graph in Figure 8.

Kasteleyn orientations on a graph are not unique. In fact, one can obtain a new Kasteleyn orientation from an existing one by flipping the orientations of all the edges adjacent to a given vertex. We shall say that two Kasteleyn orientations on the same graph are equivalent if they can be transformed to one another by a series of flips at their vertices. It is shown in [1] (Corollary 1) that a surface graph of genus g has 2^{2g} equivalence classes of Kasteleyn orientations.

In order to obtain non-equivalent Kasteleyn orientations from an existing one, we choose a family B of simple closed curves on Σ which avoid the vertices of Γ and form a basis of $H_1(\Sigma; \mathbb{Z}_2)$. For the utility

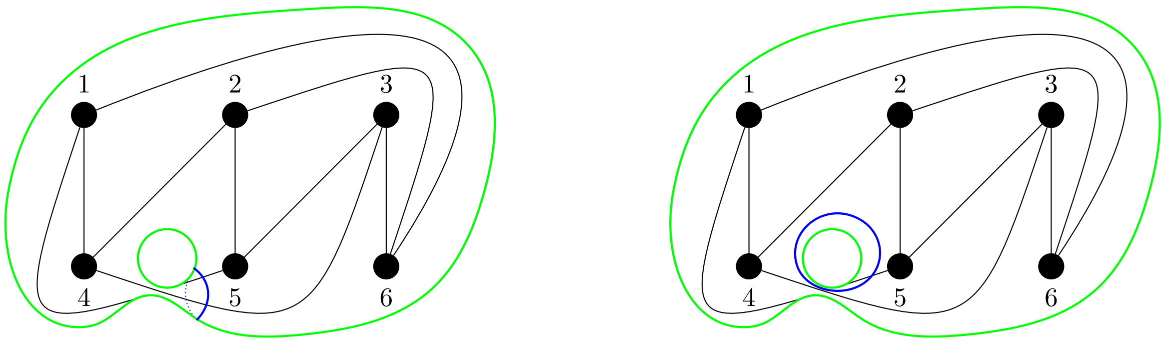

(a) The curve β_1. (b) The curve β_2.

Figure 9: A basis of the first homology space of the utility graph.

graph, we take the curves $\{\beta_1, \beta_2\}$ shown in Figure 9. Then take a subset $I \subset B$ of these curves, and flip the orientation of each edge once for every curve in I which it intersects.

On the utility graph, curve β_1 intersects only edge e_{15}, while β_2 intersects edges e_{15}, e_{34}. Hence, the four classes of Kasteleyn orientations are given by the one shown in Figure 8, and those obtained by flipping the edge sets $\{e_{15}\}$, $\{e_{15}, e_{34}\}$, and $\{e_{34}\}$. These four Kasteleyn orientations are shown in Figure 10.

The equivalence classes of Kasteleyn orientations on a surface graph allow us to define a *quadratic form* on the first homology space of the surface. A quadratic form $q : H \to \mathbb{Z}_2$ is a map on a $\mathbb{Z}_2$-vector space endowed with a bilinear form, $(H, \cdot)$, such that for all $a, b \in H$, we have $q(a+b) = q(a) + q(b) + a \cdot b$. The bilinear form on the first homology space we will use is the intersection number of curves representing

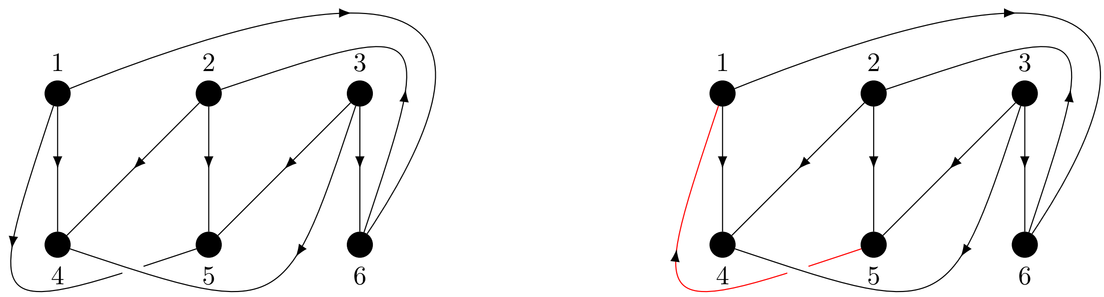

(a) K_0, the original Kasteleyn orientation. (b) K_1, given by flipping along β_1.

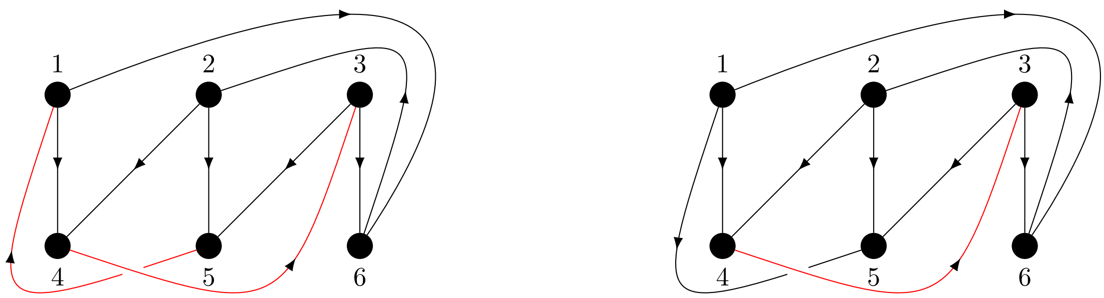

(c) K_2, given by flipping along β_2. (d) K_3, given by flipping along β_1 and β_2.

Figure 10: Four non-equivalent Kasteleyn orientations on the utility graph.

K	α_0	α_1	α_2	α_3	Arf
K_0	0	1	1	1	-1
K_1	0	0	1	0	1
K_2	0	1	0	0	1
K_3	0	0	0	1	1

Table 1: The quadratic forms $q^K_{D_0}$ evaluated on the utility graph, and their Arf invariants.

the homology classes.

Specifically, we fix a dimer configuration D in Γ, and let $\ell_D(C)$ denote the number of vertices in the oriented curve C at which the adjacent dimer in D sticks out strictly to the left of C in Σ. Recall also that $n^K(C)$ denotes the number of edges in C where the orientation of C differs from K.

Theorem 1.1 ([1], Theorem 3 & Corollary 2)**.** *Given a class $\alpha \in H_1(\Sigma;\mathbb{Z}_2)$ represented by oriented simple closed curves $C_1,\dots,C_m$, a Kasteleyn orientation K, and a dimer configuration D on $\Gamma \subset \Sigma$, the function $q^K_D : H_1(\Sigma;\mathbb{Z}_2) \to \mathbb{Z}_2$ given by*

$$q^K_D(\alpha) = \sum_{i<j} C_i \cdot C_j + \sum_{i=1}^{m} \left(n^K(C_i) + \ell_D(C_i) + 1\right) \pmod 2,$$

is a well-defined quadratic form on $(H_1(\Sigma;\mathbb{Z}_2),\cdot)$, where $\cdot$ denotes the intersection form. Furthermore, the map $[K] \mapsto q^K_D$ from the set of equivalence classes of Kasteleyn orientations on Γ, to the set of quadratic forms on $(H_1(\Sigma;\mathbb{Z}_2),\cdot)$, is an isomorphism.

Notice that this implies that there are also 2^{2g} distinct quadratic forms on a surface of genus g. In the cases where a homology class can be represented by a single curve, such as those in Figure 6, this calculation can be simplified.

Corollary 1.2. *If C is an oriented simple closed curve in Γ, then*

$$q^K_D([C]) = n^K(C) + \ell_D(C) + 1 \pmod 2.$$

We also require an additional property of quadratic forms known as their *Arf invariant*, $\mathrm{Arf}(q) \in \{-1,1\}$, which is given by

$$\mathrm{Arf}\left(q^K_D\right) = \frac{1}{\sqrt{|H_1|}} \sum_{\alpha \in H_1} (-1)^{q^K_D(\alpha)}.$$

Equivalently, if $k \in \mathbb{Z}_2$ is the value which q^K_D takes most often over $H_1(\Sigma;\mathbb{Z}_2)$, then $\mathrm{Arf}\left(q^K_D\right) = (-1)^k$.

Denote the dimer configuration shown in Figure 5 as D_0. We hence provide an example calculation of the values and Arf invariant of the form $q^{K_0}_{D_0}$ on the utility graph, and we then summarise these results, and those for the remaining quadratic forms in Table 1. We first need an arbitrary orientation on each of the curves $\alpha_0, \alpha_1, \alpha_2, \alpha_3$ representing the homology classes, so we choose the orientations which agree with K_0 along each of the edges in D_0. We thus obtain

$$\begin{aligned}
q^{K_0}_D(\alpha_0) &= n^K(\alpha_0) + \ell_D(\alpha_0) + 1 = 1 + 0 + 1 = 0,\\
q^{K_0}_D(\alpha_1) &= n^K(\alpha_1) + \ell_D(\alpha_1) + 1 = 2 + 2 + 1 = 1,\\
q^{K_0}_D(\alpha_2) &= n^K(\alpha_2) + \ell_D(\alpha_2) + 1 = 2 + 2 + 1 = 1,\\
q^{K_0}_D(\alpha_3) &= n^K(\alpha_3) + \ell_D(\alpha_3) + 1 = 2 + 0 + 1 = 1,\\
\mathrm{Arf}\left(q^{K_0}_{D_0}\right) &= \frac{1}{\sqrt{|H_1|}} \sum_{\alpha \in H_1} (-1)^{q^{K_0}_{D_0}(\alpha)} = \frac{1}{\sqrt{4}}(1 - 1 - 1 - 1) = -1.
\end{aligned}$$

The existence of non-equivalent Kasteleyn orientations on a surface graph means that the partition function cannot be calculated from a single Pfaffian. In fact, it requires a contribution from the Pfaffian associated with all 2^{2g} non-equivalent Kasteleyn orientations.

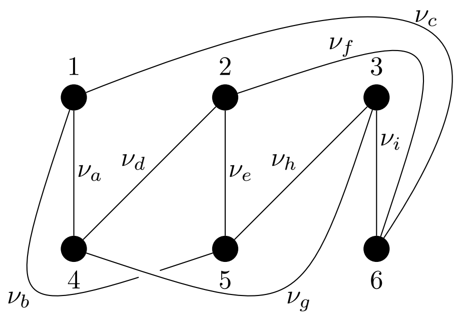

Figure 11: An edge weight system on the utility graph.

Theorem 1.3 ([1], Theorem 5)**.** *The partition function of a dimer model on a surface graph $\Gamma \subset \Sigma$ for a closed surface of genus g is given by*

$$Z(\Gamma;\nu) = \frac{1}{2^g}\sum_{[K]} \operatorname{Arf}\left(q^K_{D_0}\right)\varepsilon^K(D_0)\operatorname{Pf}\left(\mathsf{A}^K\right),$$

where the sum is taken over all equivalence classes of Kasteleyn orientations. Furthermore, the sign $\operatorname{Arf}\left(q^K_{D_0}\right)\varepsilon^K(D_0)$ *is independent of* D_0*.*

We can now calculate the partition function for the utility graph, beginning with assigning the edge weights shown in Figure 11.

The first Kasteleyn matrix is given by,

$$\mathsf{A}^{K_0} = \begin{pmatrix} 0 & 0 & 0 & \nu_a & \nu_b & \nu_c \\ 0 & 0 & 0 & \nu_d & \nu_e & -\nu_f \\ 0 & 0 & 0 & \nu_g & \nu_h & \nu_i \\ -\nu_a & -\nu_d & -\nu_g & 0 & 0 & 0 \\ -\nu_b & -\nu_e & -\nu_h & 0 & 0 & 0 \\ -\nu_c & \nu_f & -\nu_i & 0 & 0 & 0 \end{pmatrix}.$$

Since the utility graph is bipartite, its Kasteleyn matrix is of the block form

$$\mathsf{A}^{K_0} = \begin{pmatrix} 0 & \mathsf{B} \\ -\mathsf{B}^\mathsf{T} & 0 \end{pmatrix},$$

so its Pfaffian is simplified to $\operatorname{Pf}\left(\mathsf{A}^{K_0}\right) = (-1)^{\frac{n(n-1)}{2}}\det\mathsf{B}$, where B is of size $n \times n$. Hence,

$$\begin{aligned}\operatorname{Pf}\left(\mathsf{A}^{K_0}\right) &= -\det\mathsf{B} \\ &= -\nu_a(\nu_e\nu_i + \nu_f\nu_h) + \nu_b(\nu_d\nu_i + \nu_f\nu_g) - \nu_c(\nu_d\nu_h - \nu_e\nu_g) \\ &= -\nu_a\nu_e\nu_i - \nu_a\nu_f\nu_h + \nu_b\nu_d\nu_i + \nu_b\nu_f\nu_g - \nu_c\nu_d\nu_h + \nu_c\nu_e\nu_g.\end{aligned}$$

The Pfaffians of the remaining Kasteleyn orientations can be calculated from $\operatorname{Pf}\left(\mathsf{A}^{K_0}\right)$ simply by changing the signs of the edges whose orientations were flipped as highlighted in Figure 10. We must also calculate the signs induced by the dimer configuration, $\varepsilon^K(D_0)$. Since all four Kasteleyn orientations have the same orientations on these dimer edges, we only need to calculate this sign once, which we will do using the permutation $\sigma : (1,2,3,4,5,6) \mapsto (1,4,2,5,3,6)$. In cycle notation, this is $\sigma = (2\;4\;5\;3) = (2\;4)\,(4\;5)\,(5\;3)$, so $\operatorname{sgn}(\sigma) = -1$, and we have

$$\varepsilon^K(D_0) = \operatorname{sgn}(\sigma)\prod_{\ell=1}^{3}\varepsilon^K_{i_\ell,j_\ell}(e_\ell) = -1 \times 1 \times 1 \times 1 = -1.$$

Hence, Theorem 1.3 allows us to calculate the partition function of the utility graph,

$$\begin{aligned}
Z(\Gamma;\nu) &= \frac{1}{2^1}\sum_{i=0}^{3}\operatorname{Arf}\left(q_{D_0}^{K_i}\right)\varepsilon^{K_i}(D_0)\operatorname{Pf}\left(\mathsf{A}^{K_i}\right) \\
&= -\frac{1}{2}\left(-\operatorname{Pf}\left(\mathsf{A}^{K_0}\right)+\operatorname{Pf}\left(\mathsf{A}^{K_1}\right)+\operatorname{Pf}\left(\mathsf{A}^{K_2}\right)+\operatorname{Pf}\left(\mathsf{A}^{K_3}\right)\right) \\
&= \frac{1}{2}(-\nu_a\nu_e\nu_i-\nu_a\nu_f\nu_h+\nu_b\nu_d\nu_i+\nu_b\nu_f\nu_g-\nu_c\nu_d\nu_h+\nu_c\nu_e\nu_g \\
&\qquad +\nu_a\nu_e\nu_i+\nu_a\nu_f\nu_h+\nu_b\nu_d\nu_i+\nu_b\nu_f\nu_g+\nu_c\nu_d\nu_h-\nu_c\nu_e\nu_g \\
&\qquad +\nu_a\nu_e\nu_i+\nu_a\nu_f\nu_h+\nu_b\nu_d\nu_i-\nu_b\nu_f\nu_g+\nu_c\nu_d\nu_h+\nu_c\nu_e\nu_g \\
&\qquad +\nu_a\nu_e\nu_i+\nu_a\nu_f\nu_h-\nu_b\nu_d\nu_i+\nu_b\nu_f\nu_g+\nu_c\nu_d\nu_h+\nu_c\nu_e\nu_g) \\
&= \nu_a\nu_e\nu_i+\nu_a\nu_f\nu_h+\nu_b\nu_d\nu_i+\nu_b\nu_f\nu_g+\nu_c\nu_d\nu_h+\nu_c\nu_e\nu_g,
\end{aligned}$$

showing that there are $Z(\Gamma;1)=6$ possible dimer configurations on the utility graph.

2 Discrete Spin Structures

2.1 Kasteleyn Formulation

We have seen already in Theorem 1.1 that the set of equivalence classes of Kasteleyn orientations on a surface graph $\Gamma\subset\Sigma$ is isomorphic to the set of quadratic forms on $(H_1(\Sigma;\mathbb{Z}_2),\cdot)$. It is a well known result of Johnson [4] that the set of quadratic forms on Σ is in fact also isomorphic to the set of *spin structures* on Σ. This implies the following result.

Corollary 2.1 ([1], Corollary 3)**.** *Any dimer configuration D on a surface graph $\Gamma\subset\Sigma$ induces an isomorphism from the set of equivalence classes of Kasteleyn orientations on Γ to the set of spin structures on Σ.*

While we shall not go into their formal definition (see [1], Section 5.2), a spin structure is a higher notation of orientation which allows for the definition of *spinors* (such as fermions) on a surface. Spinors are the vector-like objects on which the Dirac operator (the square-root of the Laplacian) acts. They have the property that when parallel transported around a 2π spacial rotation, they pick up a sign change, and in fact require a 4π rotation to return to their original state. The sign change picked up by a spinor parallel transported in a closed loop γ is given by its *holonomy*, $\zeta(\gamma)\in\mathbb{Z}_2$.

In order to calculate the holonomy of curves on a surface endowed with a Kasteleyn orientation of a surface graph and a fixed dimer configuration, we deform γ via a homeomorphism to a curve $\alpha\subset\Gamma$ which lies along the edges of the surface graph. The holonomy of γ is then given by[1]

$$\zeta(\gamma)=q_D^K(\alpha)-1 \pmod 2.$$

In particular, the boundary of a disc α_0, is the identity element of $H_1(\Sigma;\mathbb{Z}_2)$, and will intersect any curve α an even number of times. We therefore have that

$$q(\alpha)=q(\alpha+\alpha_0)=q(\alpha)+q(\alpha_0)+\alpha\cdot\alpha_0=q(\alpha)+q(\alpha_0) \pmod 2,$$

so we must have that $q(\alpha_0)=0$ for any quadratic form. Hence, the holonomy of the boundary of a disc is $\zeta(\alpha_0)=1$, reproducing the sign-change property of spinors transported around a 2π rotation.

In this way, the pairing of Kasteleyn orientations on a surface graph with a fixed dimer configuration is a discrete representation of a spin structure on an Riemann surface.

2.2 PLCW Formulation

An alternative discrete representation of spin structures on a Riemann surface is given in [2], which is built upon piecewise-linear CW cell decompositions of surfaces. We shall first show that our surface

[1] This result will be proved in Theorem 2.6.

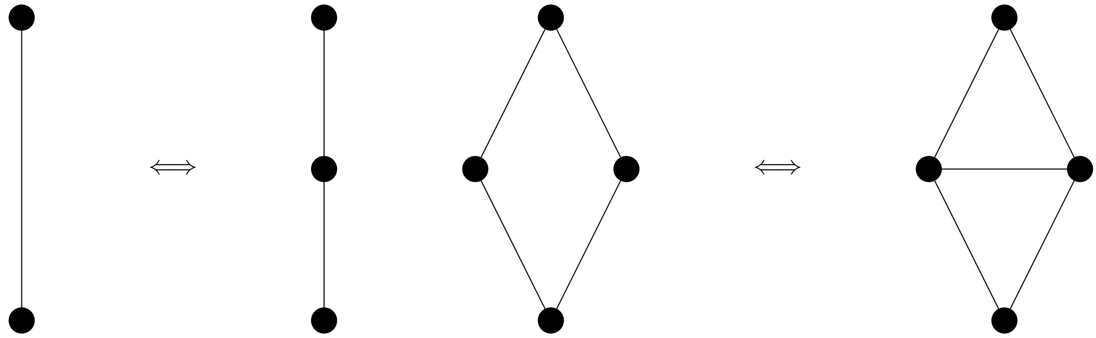

(a) Inserting a vertex in the middle of an edge. (b) Inserting an edge in the middle of a face.

Figure 12: The two possible elementary moves on PLCW decompositions of a surface.

graphs, and their dual graphs, are PLCW decompositions in the sense of [2]. Recall that the dual Γ^* of a graph Γ is given by placing a vertex of Γ^* in each face of Γ, then for each edge $e \in \Gamma$, draw an edge $e^* \in \Gamma^*$ between the vertices of the faces separated by e. We shall use the facts that any triangulation of a surface is a PLCW decomposition, and that an *elementary move* as defined below takes one PLCW decomposition to another.

Generally, an elementary move involves the insertion of an m-cell in order to split one $m+1$-cell into two, or the removal of an m-cell in order to combine two $m+1$-cells into one. For 2-dimensional surfaces, this means either inserting/removing a vertex along an edge (Figure 12a), or inserting/removing an edge, between two existing vertices, in the middle of a face (Figure 12b).

Theorem 2.2. *A surface graph $\Gamma \subset \Sigma$ and its dual surface graph $\Gamma^* \subset \Sigma$ are both PLCW decompositions.*

Proof. Denote the vertices, edge, and faces in the planar representation of Γ by V, E, and F respectively. Beginning with the graph Γ, we will first describe the series of elementary moves to transform it into a triangulation.

1. Insert a vertex in the middle of every edge in E. Denote the collection of these vertices $\tilde{V}$.
2. Insert an edge in the middle of each face in F, between any two distinct vertices in V bounding this face. Denote the collection of these edges $\tilde{E}$
   (a) If a face contains only one vertex in V, place this edge between that vertex and one from $\tilde{V}$, and note sub-point 5a.
3. Insert a vertex in the middle of each edge $\tilde{E}$. Denote the collection of these vertices V^*.
4. For each vertex $v \in V^*$, insert an edge between v and each vertex in V bounding the face containing v to which v is not already connected. These new edges are added to $\tilde{E}$.
   (a) Any vertex visited multiple times when walking the boundary of a face should receive an edge for each time it is visited.
5. For each vertex $v \in V^*$, insert an edge between v and each vertex in $\tilde{V}$ bounding the face containing v. Denote the collection of these edges E^*.
   (a) For faces containing only one vertex in V, do not insert a new edge between v and the vertex in $\tilde{V}$ to which it is already connected. Instead, move the edge connecting these vertices from $\tilde{V}$ to V^*.

This completes a triangulation of Γ using elementary moves, showing that Γ is a PLCW decomposition. We can now continue the process in order to obtain the dual graph Γ^*. Let F^* denote the faces induced by the edge set E^*.

6. Remove[2] all edges in E.

[2]Before performing this step graphically, it is useful to tessellate parts of the graph so that the outer boundary is now defined by edges in E^*, rather than edges in E.

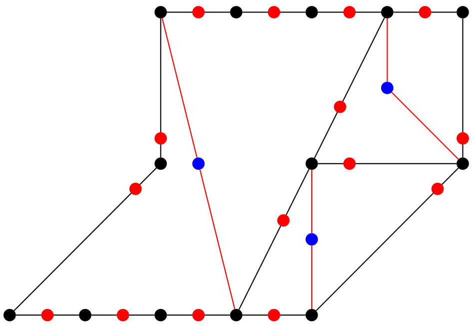

(a) The graph after steps 1-3.

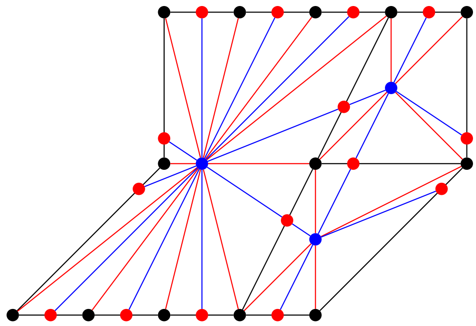

(b) The graph after steps 4–5.

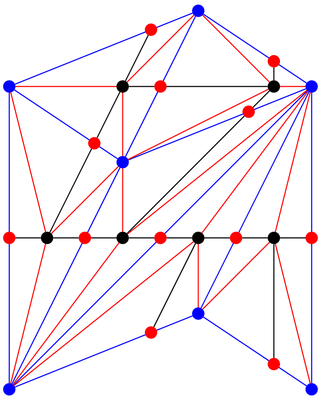

(c) The graph before step 6.

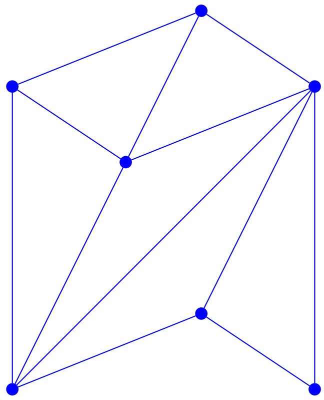

(d) The graph after steps 6–10.

Figure 13: The triangulation process on the utiltiy graph.

   (a) When a face in F^* contains only one edge in E (and therefore one edge in $\tilde{E}$, do not remove this edge, instead move it from E to $\tilde{E}$

7. For each vertex in V, remove all but two edges in $\tilde{E}$ connected to it.
8. Remove all vertices in V.
9. Remove the remaining edges in $\tilde{E}$.
10. Remove all vertices in $\tilde{V}$.

This transforms the triangulation into the dual graph Γ^* using elementary moves, showing that Γ^* is also a PLCW decomposition. □

The process of applying these steps to the utility graph is shown in Figure 13. Elements in V and E belonging to Γ are shown in black. Elements in $\tilde{V}$ and $\tilde{E}$ belonging to the triangulation only are shown in red. Elements in V^* and E^* belonging to Γ^* are shown in blue.

We can now describe the encoding of a spin structure onto a PLCW decomposition as given in [2], and in doing so provide a dictionary that can be used to convert between the Kasteleyn representation of a spin structure on a graph Γ, and the equivalent PLCW representation on its dual graph Γ^*. We will do so on the utility graph shown in Figure 14, representing Γ in black, and Γ^* in blue below.

The discrete spin data on PLCW decompositions is encoded by a *marking*, which consists of three aspects:

- A choice of a marked edge for each face.
- An orientation of each edge.
- An edge index $s_{e^*} \in \mathbb{Z}_2$ for each edge e^*.

The first two can be chosen arbitrarily, while the edge indices must obey a condition around each vertex which will be described later. We can obtain the first two on Γ^* from our choice of a dimer configuration D and Kasteleyn orientation K on Γ. We will use the dimer configuration D shown in Figure 5, and the Kasteleyn orientation K shown in 8.

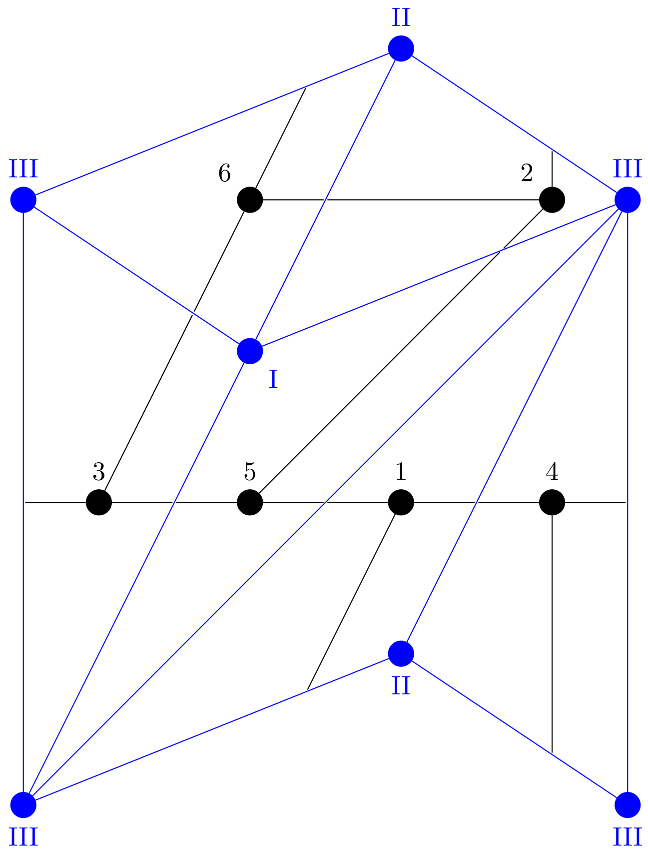

Figure 14: The utility graph and its dual.

For each face $f^* \in F^*$ on the dual graph, there is precisely one dimer edge $e \in D$ which goes through the vertex $v \in V$ associated with f^*. Mark the edge $e^* \in \partial(f^*)$ which intersects e. Although it is not necessary for marked PLCW decompositions in general, this construction ensures that each edge e^* intersecting a dimer edge e will be the marked edge for both faces it borders, so we can refer to these edges simply as *marked edges*, without having to specify the face on which they are marked. We represent the dimer edges in bold, and the marked edge of each face with a red semicircle pointing into that face, on Figure 15.

Each edge $e^* \in E^*$ intersects one edge $e \in E$, which has an orientation given by K. If e^* is a marked edge, rotate the orientation of e clockwise about the edges' intersection until it lines up with e^*, thus giving the orientation of e^*. If e^* is not a marked edge, instead rotate it anticlockwise. We represent the Kasteleyn orientation on the edges $e \in E$ in black, and the derived orientation on the edges $e^* \in E^*$ in blue, on Figure 15. Let us call this method of obtaining marked edges and edge orientations from D and K the *natural* marking of the dual graph Γ^*.

We can now consider the edge indices. In order to refer to the edges on Γ^*, we give each of them the same letter as the intersecting edge on Γ had in Section 1. These labels are also shown in Figure 15.

For each vertex $v^* \in V^*$, let D_{v^*} be the number of faces whose marked edge has v^* as its boundary in the clockwise direction. For example, vertex III is the clockwise boundary of the marked edges in faces 4, 5, and 6, so $D_{\text{III}} = 3$.

Next, the orientation of each edge gives a starting and ending vertex for each edge (these may be the same if an edge connects a vertex to itself). Let $N_{v^*}^{\text{start}}$ and $N_{v^*}^{\text{end}}$ be the number of edges starting and ending at v^* respectively, and let

$$N_{v^*} = N_{v^*}^{\text{start}} + N_{v^*}^{\text{end}}.$$

We can see for example that edge f starts at vertex I, and edges e, h, i end there, so $N_{\text{I}}^{\text{start}} = 1$, $N_{\text{I}}^{\text{end}} = 3$, and $N_{\text{I}} = 4$. Finally, we let

$$\partial^{-1}(v^*) := \{e^* \in E^* | v^* \in \partial(e^*)\}$$

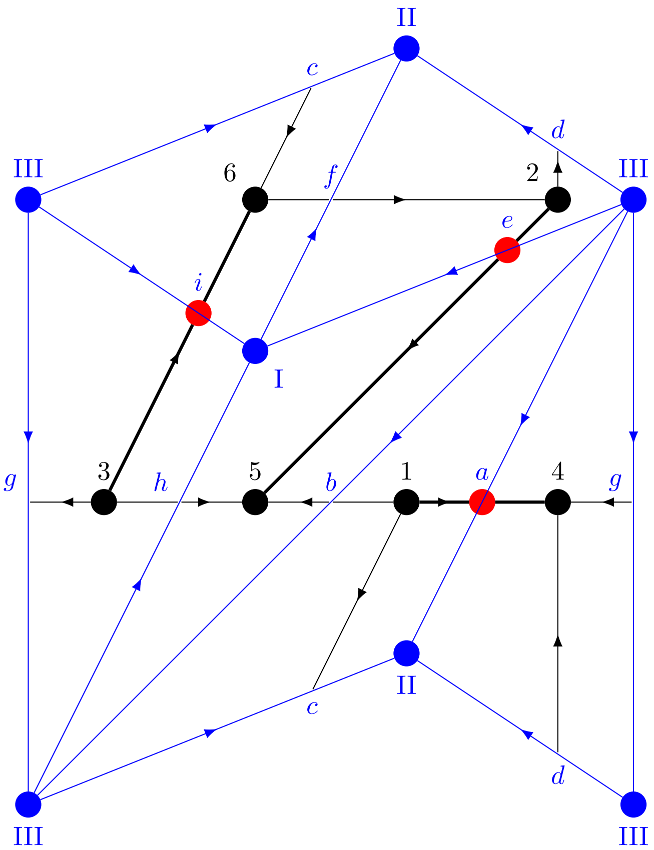

Figure 15: Marked edges and orientations on the dual of the utility graph.

denote the edges whose boundary contains v^*, and have the following condition on the edge indices s_{e^*} around each vertex $v^* \in V^*$,

$$\sum_{e^* \in \partial^{-1}(v^*)} s_{e^*} = D_{v^*} + N_{v^*}^{\text{start}} + 1 \pmod 2. \tag{4}$$

Any assignment of edge indices satisfying these conditions defines a spin structure, however with the natural marking of Γ^* described above, there is a trivial choice of indices.

Proposition 2.3. *If a PLCW decomposition obtains edge markings and orientations from a dimer configuration and Kasteleyn orientation in the natural way, then the choice of edge indices given by $s_{e^*} = 0$ for all $e^* \in E^*$ satisfies equation (4) at every vertex $v^* \in V^*$.*

Proof. Consider the face $f \in F$ associated with v^*, and assume that none of the edges $e^* \in \partial^{-1}(v^*)$ are marked. Then $D_{v^*} = 0$, and the orientations of the edges e^* are given by rotating the Kasteleyn orientations of edges e anticlockwise. Thus, each edge e^* starts at v^* if and only if e points clockwise around f. Since the orientations of the edges e are Kasteleyn, $\partial(f)$ has an odd number of such edges, and therefore, $N_{v^*}^{\text{start}}$ is odd. Hence, we have

$$D_{v^*} + N_{v^*}^{\text{start}} + 1 = 0 \pmod 2.$$

Now assume that one of the edges $e^* \in \partial^{-1}(v^*)$ is marked. This effectively flips the orientation of e^* compared to what it would have been if it was not marked, so $N_{v^*}^{\text{start}}$ changes parity. However, since e^* is marked, v^* will now be the clockwise boundary vertex for exactly one of the two faces which e^* borders, so D_{v^*} increases by 1. Since both D_{v^*} and $N_{v^*}^{\text{start}}$ change parity when the marking of e^* is introduced, the property $D_{v^*} + N_{v^*}^{\text{start}} + 1 = 0 \pmod 2$ is preserved, and so holds for all vertices. Equation 4 then becomes

$$\sum_{e^* \in \partial^{-1}(v^*)} s_{e^*} = 0 \pmod 2,$$

and so is satisfied by the edge index assignment $s_{e^*} = 0$ for all $e^* \in E^*$. □

We now have a complete description of how to construct a marked PLCW decomposition from a dimer configuration and Kasteleyn orientation. That they do in fact both describe isomorphic spin structures will be shown in Theorem 2.6, but first let us explore the notion of equivalent markings of PLCW decomposition as it relates to the Kasteleyn formation of spin structure.

Lemma 2.11 of [2] describes three possible changes to the marking of a PLCW decomposition which preserve the spin structure (up to isomorphism). These are

1. Flip the orientation of an edge e^*, and change its edge index $s_{e^*} \mapsto 1 - s_{e^*}$.
2. Move the marked edge of a face to the following edge anticlockwise, and change the index of the previously marked edge $s_{e^*} \mapsto 1 - s_{e^*}$.
3. Change the edge index of every edge around a face $s_{e^*} \mapsto 1 - s_{e^*}$. If an edge appears twice in the boundary of a face, it will undergo this change twice, and therefore not change overall.

Moves 1 and 2 allow one to convert an arbitrary marking of a PLCW decomposition into one which can be converted into a dimer configuration and Kasteleyn orientation describing the same spin structure (that is, the marked edges occur in a dimer configuration, and all edge indices are zero). Incidentally, the ability to move the dimer edges around allows one to convert a Kasteleyn orientation with respect to one dimer configuration to one with respect to another configuration while preserving the underlying spin structure, but we shall not explore this further.

Move 3, which is called a deck transformation, provides the exact same notion of equivalence on PLCW markings as our equivalence on Kasteleyn orientations did. Recall that two Kasteleyn orientations were equivalent if they could be related by a series of moves in which the orientations of all edges connecting a given vertex to any other vertex were flipped. Starting with a Kasteleyn orientation, let us recreate one of these moves, say at vertex v, in the PLCW formulation.

First choose any dimer configuration on Γ. Using this configuration and the given Kasteleyn orientation, construct a marking of the dual graph Γ^* as usual. Now let f^* be the face corresponding to v, and apply move 3 to this face. This sets the edge index of each edge which appears once in the boundary of f^* to $s_{e^*} = 1$. Use move 1 to set these indices back to 0, thereby flipping the orientation of all of these edges. Now apply the marking construction in reverse, to convert back to a Kasteleyn orientation. Each edge e^* which appears once in the boundary of f^* corresponds to an edge e which connects v to a distinct vertex. These are precisely the edges whose orientations have been flipped between the original Kasteleyn orientation and this new one, and this is therefore exactly the equivalence move relating these orientations.

We call the combination of a deck transformation (move 3) with the appropriate application of move 1 to preserve edge indices a *patio transformation*. Let a series of patio transformations on PLCW markings define an equivalence relation. Then, given a dimer configuration (and therefore a set of marked edges), and fixing all edge indices at zero, we have an isomorphism between Kasteleyn orientations on Γ and PLCW markings on Γ^* which preserves the equivalence classes on both structures. In particular, the number of Kasteleyn orientations on a graph is equal to the number of PLCW decomposition markings with fixed marked edges and indices on its dual.

2.3 Holonomy and Quadratic Forms

Recall that the holonomy, ζ, of a closed curve measures the rotation picked up by a vector (or spinor) which is parallel transported around the curve. In the context of spin structures, this holonomy is simply a number in $\mathbb{Z}_2$ which determines whether the spinor undergoes a sign change ($\zeta = 1$) or not ($\zeta = 0$). Exploring holonomy will demonstrate how both the Kasteleyn and PLCW formulations encode the actual spin data of the structure.

Holonomies of curves can be calculated from a marked PLCW decomposition as follows. Let $\gamma \subset \Sigma$ be an oriented simple closed curve on the surface in which Γ and Γ^* are embedded. Holonomy is preserved by homeomorphisms (indeed it is even preserved by homotopies, [2] Proposition 2.15), so apply a homeomorphism to γ so that it avoids the vertices of Γ^*.

In fact, let us assume that γ now lies completely along the edges of the graph Γ. Then, the edges of Γ^* partition γ into a series of oriented arcs $A(\gamma)$. Each arc $a \in A(\gamma)$ is contained within a face $f_a^* \in F^*$, and leaves that face through an edge $e_a^* \in E^*$. We say that e_a^* and a cross positively if the smallest rotation

#	O	D	$\hat{s}$	$\hat{\delta}$	$\hat{n}$	$\hat{\ell}$	$\hat{\zeta}$	$\hat{q}$
a	C	N	0	1	0	0	1	0
b	C	E	0	1	1	0	1	1
c	C	S	0	0	1	0	0	1
d	C	W	0	0	1	1	0	0
e	A	N	1	1	1	0	0	1
f	A	E	1	1	0	0	0	0
g	A	S	1	0	0	0	1	0
h	A	W	1	0	0	1	1	1

Table 2: Calculation of the partial holonomies and quadratic forms on each type of face.

of the orientation of e_a^* onto the orientation of a is anticlockwise. Hence, we define

$$\hat{s}^a_{e_a^*} = \begin{cases} s_{e_a^*} & e_a^* \text{ and } a \text{ cross positively,} \\ -1 - s_{e_a^*} & \text{otherwise.} \end{cases}$$

We also let v_a^* denote the vertex at the boundary of the marked edge of f_a^* in the clockwise direction, and define

$$\hat{\delta}^a_{f_a^*} = \begin{cases} 1 & v_a^* \text{ is to the right of } a, \\ 0 & \text{otherwise.} \end{cases}$$

Calculations of $\hat{s}^a_{e_a^*}$ and $\hat{\delta}^a_{f_a^*}$ in various cases can be found in Figure 16 and Table 2. We can hence use these two variables to calculate the holonomies of curves.

Proposition 2.4 ([2], Proposition 2.15). *The holonomy of an oriented simple closed curve $\gamma \subset \Sigma$ with respect to a marked PLCW decomposition Γ^* is given by*

$$\zeta(\gamma) = \sum_{a \in A(\gamma)} \hat{s}^a_{e_a^*} + \hat{\delta}^a_{f_a^*} \pmod 2.$$

We also have the following consequence for curves bounding discs.

Corollary 2.5 ([2], Proposition 2.15). *If γ bounds a disc $\mathbb{D}$ embedded in Σ, then*

$$\zeta(\gamma) = \operatorname{sgn}(\gamma) = \begin{cases} 1 & \gamma \text{ is oriented anticlockwise around the boundary of } \mathbb{D}, \\ -1 & \text{otherwise.} \end{cases}$$

As an example, let us calculate the holonomy of the curve α_1 shown in Figure 6b with respect to the marking of Γ^* in Figure 15. We orient α_1 to visit the vertices of Γ in the order 3 5 1 4.

The first arc of α_1 travels through face 3, and leaves through edge h. The smallest rotation of the orientation of edge h onto α_1 is clockwise, so $\hat{s}^a_h = -1 - s_h$. The marked edge of face 3 is edge i, and its clockwise boundary vertex is vertex I, which lies to the left of α_1, so $\hat{\delta}^a_3 = 0$. Repeating this calculation for the remaining three faces, we obtain the homology of α_1.

$$\zeta(\alpha_1) = (-1 - s_h + 0) + (s_b + 0) + (s_a + 1) + (s_g + 0) = s_a + s_b + s_g + s_h \pmod 2.$$

With the natural choice sending all edge indices to zero, we thus have $\zeta(\alpha_1) = 0$. Since we calculated the corresponding quadratic form to be $q^{K_0}_{D_0}(\alpha_1) = 1$ in Table 1, this agrees with our conjecture relating these two properties. Indeed, we are now ready to prove this relation.

Theorem 2.6. *Let γ be a simple closed curve embedded in a surface graph $\Gamma \subset \Sigma$ endowed with a dimer configuration D and a Kasteleyn orientation K. Let the dual graph Γ^* inherit the natural PLCW marking from D and K with all edge indices equal to zero. Then the holonomy of γ given by Γ^* satisfies the relation*

$$\zeta(\gamma) = q^K_D(\gamma) - 1 \pmod 2.$$

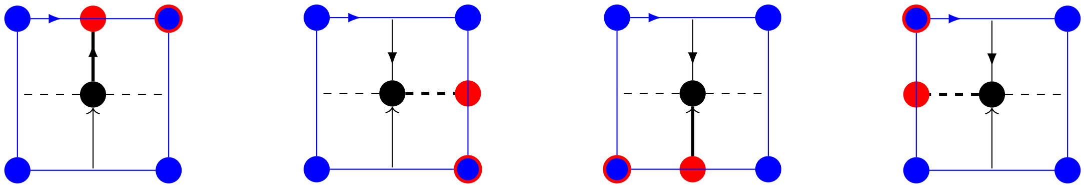

(a) Clockwise, north. (b) Clockwise, east. (c) Clockwise, south. (d) Clockwise, west.

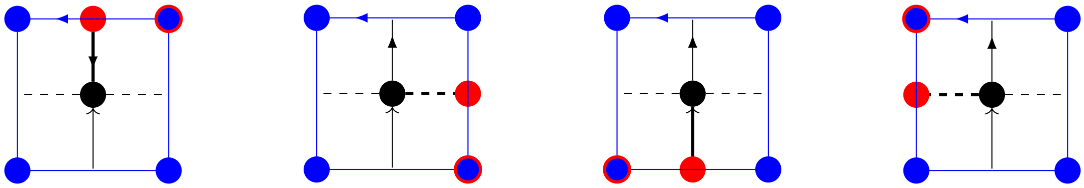

(e) Anticlockwise, north. (f) Anticlockwise, east. (g) Anticlockwise, south. (h) Anticlockwise, west.

Figure 16: The possible types of faces an arc of a curve may pass through, classified by the orientation of the edge through which the arc leaves the face, and the position of the dimer edge relative to the arc.

Proof. Like the calculation of holonomy, we shall first split the calculation of quadratic forms for a curve into arcs, with one arc travelling through each face of the dual graph Γ^*. We can classify the possible types of faces an arc may travel through by the orientation of the edge e_a^* through which the arc leaves the face, and the position of the dimer edge associated with the vertex v_a corresponding to this face f_a^*. This gives eight possibilities, as shown in Figure 16.

Shown in blue are the vertices and edges of the dual graph Γ^*, as well as the orientation of the edge e_a^*. Without loss of generality, we have chosen to represent each face with four edges, although it is possible that there may in fact be fewer or more. Shown in black are the vertices and edges corresponding to the graph Γ. The arc travels along the solid edges from south to north, as represented by the wedge-shaped arrow. The Kasteleyn orientation of the northern edge e_a is represented by the triangular arrow. The dotted edges are those that do exist in the graph Γ, but which the arc does not travel along. The dimer edge e_D is drawn with a thicker black line, and the corresponding marked edge e_D^* is indicated by a pair of red semicircles. The clockwise vertex v_D^* of this marked edge is outlined in red.

The *partial holonomy* of each arc a is given by

$$\hat{\zeta}\,(a) := \hat{s}\,(a) + \hat{\delta}\,(a) \pmod 2,$$

where we have dropped the indices of $\hat{s}^a_{e_a^*}$ and $\hat{\delta}^a_{f_a^*}$ for simplicity. Notice that since we have assumed all edge indices are zero, $\hat{s}\,(a) = 0$ if e_a^* and a cross positively, and 1 otherwise. Hence, the holonomy of the whole curve is given by summing the partial holonomies of each arc,

$$\zeta\,(\gamma) = \sum_{a\in A(\gamma)} \hat{s}\,(a) + \hat{\delta}\,(a) = \sum_{a\in A(\gamma)} \hat{\zeta}\,(a) \pmod 2.$$

Similarly, the *partial quadratic form* on each arc a is given by

$$\hat{q}\,(a) := \hat{n}\,(a) + \hat{\ell}\,(a) \pmod 2,$$

where we define $\hat{n}$ and $\hat{\ell}$ as follows. Let $\hat{n}\,(a) = 0$ if the Kasteleyn orientation of the northern edge e_a and the orientation of the arc agree, and 1 otherwise, and let $\hat{\ell}\,(a) = 1$ if the dimer edge is strictly to the left of a, and 0 otherwise. Hence applying Corollary 1.2, the quadratic form on the whole curve is given by

$$q_D^K\,(\gamma) = 1 + n^K\,(\gamma) + \ell_D\,(\gamma) = 1 + \sum_{a\in A(\gamma)} \hat{n}\,(a) + \sum_{a\in A(\gamma)} \hat{\ell}\,(a) = 1 + \sum_{a\in A(\gamma)} \hat{q}\,(a) \pmod 2.$$

Let us calculate explicitly the partial holonomy and partial quadratic form for face (a). The blue orientation of edge e_a^* aligns with the orientation of the arc when rotated anticlockwise, so these cross positively, and $\hat{s}\,(a) = 0$. The clockwise vertex of the marked edge is the northeastern edge, which is to the right of the arc, so $\hat{\delta}\,(a) = 1$. Hence, the partial holonomy of a is

$$\hat{\zeta}\,(a) = 0 + 1 = 1 \pmod 2.$$

The Kasteleyn orientation of the northern edge e_a agrees with the orientation of the arc, so $\hat{n}(a) = 0$. The dimer edge is the northern one, which is not strictly to the left of the arc, so $\hat{\ell}(a) = 0$. Hence, the partial quadratic form on a is

$$\hat{q}(a) = 0 + 0 = 0 \pmod{2}.$$

These calculations, and those for the remaining faces are summarised in Table 2, where O indicates whether the orientation of e_a^* is clockwise (C) or anticlockwise (A), and D indicates the cardinal direction (N/E/S/W) of the dimer edge containing v_a relative to the orientation of the arc.

We now show that the sum of partial holonomies and the sum of partial quadratic forms are equal. Looking at Table 2, we can see that for faces where the dimer edge is to the east or west, $\hat{\zeta}$ and $\hat{q}$ are already the same.

For the north- and south-dimer faces, $\hat{\zeta}$ and $\hat{q}$ have opposite parity. However, notice that if an arc leaves a face along a dimer edge, it must enter the next face along a dimer edge, so north- and south-dimer faces must always occur in pairs. This means that across these pairs of faces, the sums of $\hat{\zeta}$ and $\hat{q}$ will again be the same.

Bringing this all together, we indeed find that the relationship between the holonomy and quadratic form on a curve is indeed given by

$$\zeta(\gamma) = \sum_{a \in A(\gamma)} \hat{\zeta}(a) = \sum_{a \in A(\gamma)} \hat{q}(a) = q_D^K(\gamma) - 1 \pmod{2}.$$

□

Finally, holonomy provides an alternative method for calculating the Arf invariant of a spin structure. Enumerate the g holes of the genus-g surface Σ by $1, \ldots, g$. Let $a_i := \beta_1$ and $b_i := \beta_2$ be a basis of the i^{th} hole, where β_1 and β_2 are the curves shown in Figure 9. Then, from Definition 4.5 of [2], the Arf invariant of the spin structure S_Σ on Σ is given by

$$\operatorname{Arf}(S_\Sigma) = \sum_{i=1}^{g} (\zeta(a_i) + 1)(\zeta(b_i) + 1) = \sum_{i=1}^{g} q_D^K(a_i)\, q_D^K(b_i) \pmod{2}.$$

It is demonstrated in [2] that the PLCW formulation can in fact be generalised to represent r-spin structures, which might allow for the construction of higher (r^{th}) roots of the Laplacian, by considering edge indices coming from $\mathbb{Z}_r$. While these higher spin structures could allow for a generalisation of Kasteleyn orientations, they are not relevant to the dimer model, and so we shall not explore them further here. This completes our conversion of the Kasteleyn formulation of spin structure to the PLCW formulation, and with it our treatment of the dimer model.

References

[1] David Cimasoni and Nicolai Reshetikhin. Dimers on surface graphs and spin structures. i. *Communications in Mathematical Physics*, 275(1):187–208, Jul 2007.

[2] Ingo Runkel and Lóránt Szegedy. Topological field theory on r-spin surfaces and the arf-invariant. *Journal of Mathematical Physics*, 62(10), Oct 2021.

[3] David Cimasoni. The geometry of dimer models. *Winter Braids Lecture Notes*, 1:1–14, Dec 2015.

[4] Dennis Johnson. Spin structures and quadratic forms on surfaces. *Journal of the London Mathematical Society*, s2-22(2):365–373, Oct 1980.