The Higher Lattice Gauge Theory Model for Topological Phases of Matter

Joe Huxford

Wadham College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2022

Abstract

Topological phases of matter are characterised by long-range entanglement between their constituent degrees of freedom, which allows them to host excitations with non-trivial exchange statistics. While topological phases in 2+1d are relatively well studied, less is known about their 3+1d counterparts. In this thesis, we examine features of these higher-dimensional phases in a tractable example model, based on a generalisation of lattice gauge theory called higher lattice gauge theory. As well as hosting point-like excitations, the model supports loop-like excitations with non-trivial loop-loop and point-loop braiding statistics. We explicitly construct operators to produce and move these excitations, and use these to find the loop-loop and point-loop braiding relations. These creation operators also reveal that some of the excitations are confined, costing energy to separate, while others are condensed and can be produced locally. This is discussed in the context of condensation-confinement transitions between different cases of this model. We also explain a method for measuring topological charge and use explicit measurement operators to re-derive a relationship between the number of charges measured by a 2-torus and the ground-state degeneracy of the higher lattice gauge theory model on the 3-torus. From these measurement operators, we can see that the ground state degeneracy on the 3-torus is related to the number of types of linked loop-like excitations.

Acknowledgements

My warmest thanks go to my supervisor Steven Simon for guiding me through this project and spending many hours reading and commenting on my work, including this thesis. I am also grateful to my colleagues Yves Kwan, Michele Fava and Greg Henderson, along with many others in the department, for their support and encouragement.

This thesis is based on my pre-publications Ref. [1] and Ref. [2], with some references to results I prove in Ref. [3]. These works were made possible with support from EPSRC Grant EP/S020527/1.

Contents

Chapter 1

Introduction

The field of condensed matter physics is concerned with systems with many constituents in relatively close proximity. In these systems, collective phenomena are often more significant than the properties of the constituent particles. One interesting class of condensed matter system that exemplifies this idea includes so-called topological phases of matter [4, 5, 6], which fall outside of the usual Landau picture of phases of matter arising from spontaneous symmetry breaking described by local order parameters [7]. These topological phases, which include the celebrated fractional quantum Hall systems [8, 9, 10, 11, 12], are characterised by long-range entanglement between their local degrees of freedom [6, 13, 14]. While the fact that these phases cannot be described by symmetry breaking is itself interesting, topological phases can also possess rather unique properties as a result of this long-range entanglement. For example, these long-range entangled topological phases may have a ground state degeneracy even in the absence of additional symmetry [15, 16] (when also considering phases with enforced symmetry, the classification of topological phases becomes more rich and includes so-called symmetry protected and symmetry enriched topological phases [13]). This ground-state degeneracy depends on the topology of the manifold on which the topological phase is placed, meaning that such a phase may have a ground state degeneracy on the torus but not the sphere for example, with this degeneracy being resistant to local perturbations [15, 16]. This feature may have practical applications in

the future, because it can allow such topological phases to serve as quantum memories [16, 17, 18, 19], with information encoded in the topologically protected degenerate subspace being resistant to local noise [16, 20].

Further intriguing properties of the long-ranged entangled topological phases are revealed when we consider excitations. In 2+1d, the entanglement structure allows these phases to support *anyonic* excitations, which are generalizations of the more familiar bosons and fermions, and can undergo non-trivial transformations when two anyons are passed around each-other, even at large distances [21, 22, 23] (for the interested reader, we note that there are many works giving pedagogical introductions to anyon physics, such as Ref. [24], Ref. [25] and Appendix E in Ref. [26]). Because these transformations (called braiding relations) do not depend on local details, it is believed that these excitations can be used for fault-tolerant quantum computation [27, 28], should sufficiently stable phases and excitations be constructed. In three spatial dimensions, while any point-like excitations must be fermionic or bosonic [25, 27, 29, 30], topological phases can admit loop-like excitations that undergo non-trivial transformations when a point-like or loop-like excitation is passed *through* a loop-like excitation [31, 32] (even the Aharanov-Bohm effect [33] can be thought of as braiding a point-like electron with a loop-like or string-like magnetic flux tube).

In order to provide a setting where the unique properties of topological phases can be studied in detail, it is convenient to use exactly solvable toy models [16, 34, 35, 36]. While these models may not resemble those used to describe real materials [35], or only describe the renormalization group fixed point of their phase [37], they provide representatives for a large class of phases of matter [35]. This means that such constructions can be used to probe (and attempt to classify [36]) which kinds of phases can exist. The toy models have Hamiltonians constructed out of commuting projector operators, so that the quasiparticle excitations can be found exactly. Of these constructions for 2+1d topological phases, two of the most successful are the Levin-Wen

string-net model [36] and Kitaev's Quantum Double model [16] (see also work on discrete gauge theory related to the Quantum Double model in Refs. [38] and [39]). The Kitaev Quantum Double class of models includes the toric code, which appears to have a practical application as a robust way to store qubits [16]. Indeed one approach to building quantum computers uses so-called surface codes, which take inspiration from the toric code [40] and which have recently been experimentally realized on a small scale [41, 42]. The string-net construction is more general than Kitaev's Quantum Double model, and has been conjectured to cover all phases that can be represented by commuting projector models in 2+1d in the absence of an additional symmetry [43], when generalized appropriately from the original construction in Ref. [36] (see Refs. [34, 35, 44, 45, 46] for such generalizations). In both of these classes of models, it is well understood how to find the ground-state degeneracy [16, 47] and the properties of the excitations, such as braiding statistics [16, 35, 36] (as we will elaborate on shortly).

In the 2+1d commuting projector models, a useful way of obtaining information on the underlying topological theory is to find the operators, known as ribbon operators, that create and move the quasiparticle excitations [16]. This approach was used in Ref. [16] to study the excitations in Kitaev's Quantum Double model, and in Ref. [36] for the string-net model. As well as classifying the quasiparticles, these ribbon operators can be used to find the braiding relations of the quasiparticles, by taking appropriate commutation relations of the ribbon operators. Further, in Ref. [48] a method was developed for constructing operators to measure topological charge, which is a conserved charge that can exist without symmetry, by using closed ribbon operators. This method was used to identify confined charges and to demonstrate the condensation of certain charges during a condensation-confinement transition, in a modified version of the Quantum Double model introduced in that paper [48]. It is clear then that these ribbon operators provide a wealth of information about the topological phase under study.

The models that we have discussed so far describe topological phases in two spatial dimensions. However, there are also many toy models for topological phases in three spatial dimensions. Existing commuting projector Hamiltonian models include the twisted gauge theory model [49, 50, 51], which is a generalized version of the Quantum Double model in 3+1d and is based on the Dijkgraaf-Witten topological field theory [52]; a class of models developed from Unitary G-crossed Braided Fusion Categories (UGxBFCs) [53]; the Walker Wang models [54, 55, 56, 57], which are 3+1d generalizations of the Levin-Wen string-net models [54]; and the higher lattice gauge theory models [58, 59, 60, 61], based on a generalization of lattice gauge theory (and related to the Yetter quantum field theory [62]). For the twisted gauge theory model in particular, there has been significant study of the properties of the ground state [49] and the excitations, including their braiding properties [50, 51]. However, the general approach to studying these 3+1d models has been different to the approach used for models in two spatial dimensions. While in the 2+1d case, the use of ribbon operators to obtain the properties of the excitations is common, in the 3+1d case an explicit construction of the ribbon and membrane operators (the higher dimensional counterparts to ribbon operators, which produce loop-like excitations) can be difficult. There are some examples of such explicit constructions for the twisted gauge theory models in three spatial dimensions, such as in Refs. [63] and [64], but less so for other models. Instead, indirect methods like dimensional reduction [31, 50] and tube algebras [51, 60] are often used. These methods are certainly useful, but seem to offer a less complete picture of the excitations than a direct construction. Given the success of ribbon operator approaches in 2+1d, and these examples of membrane operators in the 3+1d twisted gauge theory model, we would like to be able to apply similar approaches to other 3+1d models. In this work we will do precisely that, with one of the models discussed above.

In this work we will study the higher lattice gauge theory models introduced in Ref.

[58], which can be defined in arbitrary dimension but which we will study in 3+1d. We also considered the 2+1d case in Ref. [3], but we will omit discussion of this case here for brevity. Prior to (or in some cases concurrently with) our work, some important properties of this model had already been explored by other researchers. In Ref. [58], where the Hamiltonian models based on higher lattice gauge theory were introduced, for example, the ground state degeneracy was given in terms of the partition function of a topological quantum field theory (TQFT), the Yetter TQFT, and also explicitly computed for some examples. Then in Ref. [60], the excitations were studied using a tube algebra approach, through which the loop-like excitations in the model were classified and the simple types were counted. Furthermore, it was shown in Ref. [60] that there is a relationship between the number of types of elementary excitation and the ground-state degeneracy of the model on a 3-torus. In addition, it was shown in Ref. [65] that higher gauge theory could lead to loop-like excitations with non-trivial loop braiding statistics, and the associated representations of the loop braid group were found (the loop braid group describes the motions of loops [66, 67]), although this was not done in the Hamiltonian model but instead from more geometric reasoning about the fluxes and gauge transforms involved. Until now, there was no explicit construction of these excitations in the Hamiltonian model using ribbon and membrane operators, and the braiding statistics of the excitations in the Hamiltonian model have not been found. We aim to address this, and describe some of the other features of the excitations, in this work. That is, in this thesis we will study these higher lattice gauge theory models [58] by explicitly constructing the membrane and ribbon operators for this model for the first time and applying them to find the other properties of the excitations. We note that these models are particularly interesting to study in this way because they share a similar structure to lattice gauge theory models, which helps with the difficult task of directly constructing ribbon and membrane operators, and yet still exhibit features not seen in ordinary (1-gauge) gauge theory models, as we elaborate on shortly.

As mentioned earlier, the higher lattice gauge theory models are based on a generalization of lattice gauge theory, known as higher lattice gauge theory. In Chapter 2, we will describe lattice gauge theory and higher lattice gauge theory and use this to motivate the model proposed by Bullivant et al. in Ref. [58]. To summarize briefly here, in higher lattice gauge theory the edges of a graph are labelled by group elements which describe the effect of parallel transport of points across the edges, while the plaquettes are labelled by elements in another group, which describe the parallel transport of paths across these plaquettes (whereas for lattice gauge theory, there are no plaquette labels). The fact that the Hamiltonian model is based on such geometrical reasoning allows us to have some intuition about the properties of the model. We use this fact in Chapter 3 and draw from ideas in gauge theory to motivate some of the results that we describe in more detail later.

After Chapter 3, the rest of this work will cover the results we have obtained during this project. In Chapter 4, we find the membrane and ribbon operators which produce the excitations for this model, in a broad subset of the higher lattice gauge theory models. We find that the basic excitations are either loop-like or point-like and that some of the point-like excitations are confined, costing energy to separate a particle from its anti-particle. In Chapter 5, we explain that this arises from a condensation-confinement transition between different higher lattice gauge theory models, during which some of the loop-like excitations condense out, becoming topologically trivial. This model therefore provides an example of condensation and confinement in a 3+1d setting, which we would not see in a lattice gauge theory model without some modification. Then in Chapter 6, using our direct construction of the ribbon and membrane operators, we find the (loop)-braiding relations of our excitations, which describe what happens as we move point-like or loop-like excitations around and through each-other. These braiding relations are described in terms of simple group-theoretic quantities. We find that the braiding is generally non-Abelian, so that our relations involve more than a simple accumulation of phase. Instead the excitations have an internal space,

which undergoes a transformation under braiding. While the internal spaces are transformed, the excitations carry a conserved topological charge which is not changed by the braiding. This topological charge is of significant interest, and so in Chapter 7 we consider the charges present in the higher lattice gauge theory model. Extending the methods of Ref. [48] to 3+1d, we construct operators that can measure the topological charge present in a region. These measurement operators are made from closed membrane and ribbon operators applied on the boundary of the region in question, and the topology of this boundary determines what types of charge we can resolve. For example, the charge associated to point-like objects is measured by putting a sphere around that charge. This is analogous to how the electric charge of an object may be measured by looking at the electric flux through a surface enclosing that object and applying Poisson's law. On the other hand, loop-like excitations require a surface with handles in order to detect their loop-like character. As far as we know, this is the first time that this method has been applied in 3+1d using explicit membrane operators, although this method has some similarity to the tube algebra methods used in Ref. [60]. In the tube algebra method, the excitations are described in terms of the allowed degrees of freedom on the boundary of a given region, which is presumably equivalent to finding the allowed measurement operators for the surface forming the boundary of the region. Indeed, just like Ref. [60] we find that the number of different charges that can be measured by a torus is equal to the ground state degeneracy of the model when placed on a 3-torus. Finally, in Chapter 8 we summarize our results and describe future avenues of research based on this work.

Chapter 2

The Higher Lattice Gauge Theory Model

2.1 Lattice Gauge Theory

As described earlier, we will consider a model for topological phases of matter based on higher lattice gauge theory. It will therefore be useful to first discuss lattice gauge theory and then higher lattice gauge theory. Gauge theory has been used extensively in the study of topological phases. For example, Chern-Simons theory [68] is used to describe fractional quantum Hall systems, and Kitaev's Quantum Double model [16] (which is based on lattice gauge theory, as we describe later) is one of the best studied commuting projector models for topological phases in 2+1d. In a continuum gauge theory there are three main ingredients: matter fields, gauge fields (which describe parallel transport of the matter fields) and gauge symmetry. One familiar example of such a gauge theory is conventional electromagnetism. In this case the matter fields describe charges, such as electrons. These matter fields couple to the usual scalar and vector gauge fields, which describe the parallel transport of the charged matter via the Aharanov-Bohm effect [33]. The defining feature of a gauge theory, however, is the presence of a gauge symmetry, which describes a redundancy in our description of the physical system. This symmetry manifests in the presence of gauge transforms, which

appear to change the values of the gauge and matter fields. However, states that are related by gauge transforms represent the same physical state and are simply different descriptions of the same physical system.

While gauge theory is commonly considered in the continuum, with a continuous gauge group, we can use the same ideas to construct a lattice gauge theory model [69]. We replace the continuous spatial dimensions with a discrete representation, namely a graph. A graph consists of a collection of vertices and (generally directed) edges (later we generalize this to explicitly include plaquettes, which are the smallest surface elements in the space). Note that we make no requirements for the graph to have a repeating structure (which would give a lattice), although it is sometimes convenient to do so and we will often refer to the graph informally as a lattice. We can then place the ingredients of gauge theory into this discrete setting. The matter field resides on the vertices of the lattice [69], although in this work we are interested only in the gauge fields and will ignore these potential fields. The gauge field is associated to the (directed) edges of the lattice and is valued in some discrete group G (note that while a discrete group may seem unnatural in the context of gauge theory, it can arise by spontaneously breaking a continuous group [38]). The label of each edge represents the result of parallel transport across that edge. We can then use the group structure of the gauge field to describe the parallel transport across more general paths, made from multiple edges. We construct larger paths by composing smaller ones, where the label of a path made from two smaller paths lying end-to-end is given by group multiplication of the labels of the constituent paths [69], as shown in Figure 2.1. The paths are oriented objects, and the result of parallel transport along a path against its orientation is described by the inverse of the path's label. This means that we can combine two paths that point in opposite directions, by first reversing the direction of one of them (and simultaneously replacing the label of that path with the inverse), so that they align. Therefore, if in Figure 2.1 the path labelled g_2 pointed in the opposite direction, the combined path would have label $g_1 g_2^{-1}$.

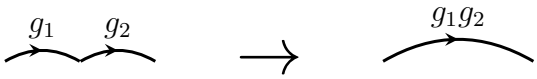

Figure 2.1: Composition of paths is described by group multiplication

2.1.1 Gauge Transforms

Having described how the gauge field fits into the lattice picture, we now look at the gauge symmetry. This symmetry is implemented by a set of local gauge transforms that each act on the degrees of freedom near a vertex. Each operator is labelled by the vertex it acts on and an element of G, so that the gauge transform for a vertex v and element $x \in G$ is denoted by A_v^x [16]. This transform affects the edges surrounding it by pre-multiplying the group element on each adjacent edge by x if the edge is outgoing, and post-multiplying the element by x^{-1} if the edge is ingoing [16, 69, 70, 71]. As we see from the example shown in Figure 2.2, this action is equivalent to adding an imaginary edge, labelled by x, to the vertex v and parallel transporting the entire vertex along it.

These vertex transforms possess two particularly useful properties. Firstly, the vertex transform only affects the labels of paths that start or end at that vertex. For example, in the top-left image in Figure 2.2 the path entering the vertex v from the lower left and exiting v through the lower right is labelled by the product $g_1 g_2$. In the bottom-left image of Figure 2.2, which represents the state after the gauge transform, the same path is labelled by $g_1 x^{-1} x g_2 = g_1 g_2$. That is, the path label is unchanged by the gauge transform, because the path does not start or terminate at the vertex v. Secondly, note that applying two gauge transforms to the same vertex is the same as parallel transporting along two edges in sequence. This should be equivalent to parallel transport of the vertex across a single path composed of the two edges, and so is the same as applying a single vertex transform with a label obtained by combining the labels of the two edges (and so combining the labels of the two original transforms). If we first apply a vertex transform A_v^g and then another transform A_v^h, the label of the

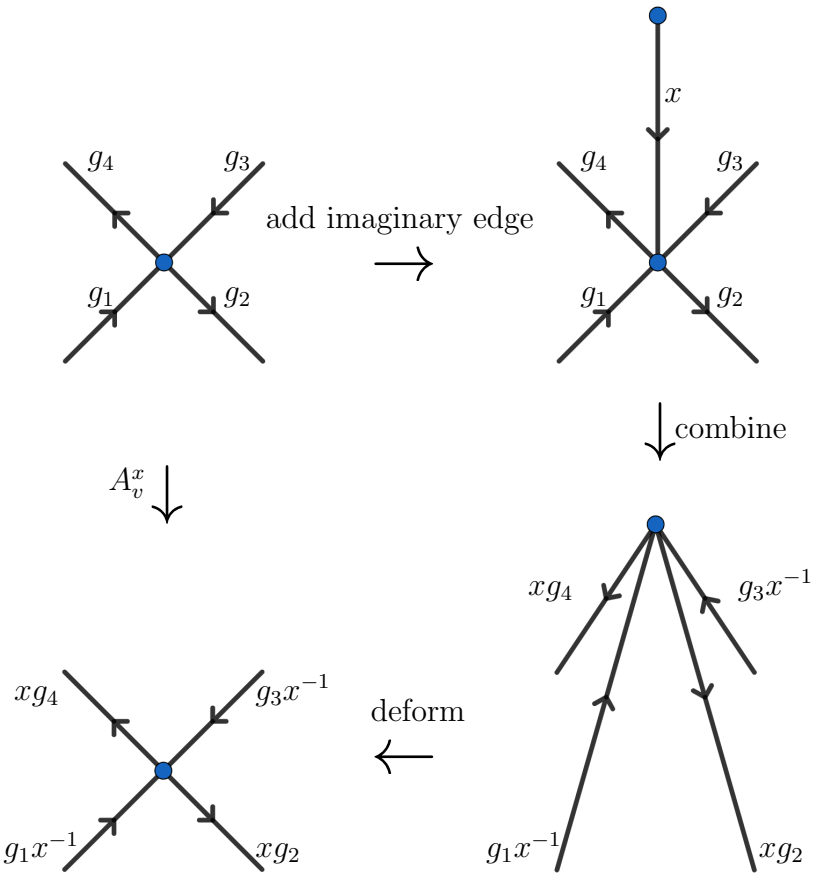

Figure 2.2: The gauge transform on a vertex is equivalent to adding an imaginary edge at that vertex and then combining this edge into the diagram, or equivalently transporting the vertex along that edge.

---

combined path introduced by the transforms is hg (it is hg rather than gh, due to the fact that the vertex is parallel transported against the direction of the edge, as seen in Figure 2.2). Therefore, we must have that $A_v^h A_v^g = A_v^{hg}$ [16].

2.1.2 Gauge-Invariants

Because states related by gauge transforms are equivalent, any physical quantity should be gauge-invariant. We can construct these gauge-invariant quantities from the closed loops of our lattice [70, 71]. Under a gauge transform, the group element assigned to a closed loop is at most conjugated by the vertex transforms [71]. Therefore, the conjugacy class of that label is a gauge-invariant quantity. For example, consider Figure 2.3 below, which shows the action of a vertex transform A_v^x on a closed loop

starting at v. The group element associated to the closed loop is initially $g_1 g_2$. After applying the vertex transform it becomes $x g_1 g_2 x^{-1}$. This indicates that the group element is not generally a gauge-invariant quantity, but its conjugacy class is.

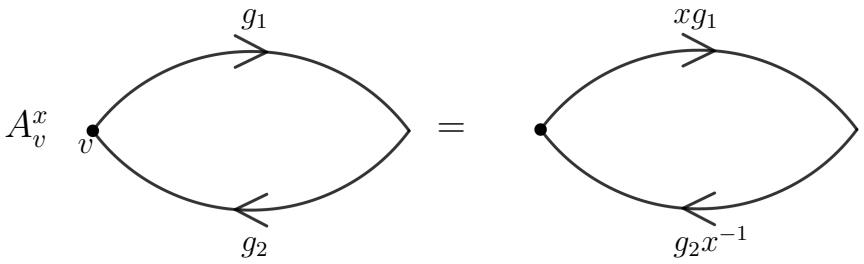

Figure 2.3: We consider the effect of a vertex transform A_v^x on a closed loop starting at the vertex v. The path label in this case goes from $g_1 g_2$ to $x g_1 g_2 x^{-1}$. That is, the path label of the closed loop is conjugated by x under the action of the transform.

As an example of the importance of such closed loops, we can consider the case of electromagnetism. Here we have a U(1) gauge symmetry, so that the edges in our lattice would be labelled by phases. There is a physical process where we take a charge q around a closed loop in the presence of a magnetic field described by a vector potential $\vec{A}$. In the continuum theory this leads to the Aharanov-Bohm effect [33], where the wavefunction accumulates a phase of $\theta = q \oint \vec{A} \cdot \vec{dl}$. The phase $e^{i\theta}$ is then the label we would give our closed loop in the lattice model. Using Stoke's theorem, the Aharanov-Bohm phase can be related to the magnetic flux through the surface enclosed by the loop. The phase is a gauge-invariant quantity, as is required by the fact that this phase can be measured in interference experiments and thus is a physical quantity.

It is important to keep track of such gauge-invariant quantities, because they allow us to differentiate between physically distinct states. For instance, many states can be reduced to the trivial configuration, where every edge is labelled by 1_G (the identity in the group G), by applying gauge transforms. The state where the edges are all labelled by the identity describes trivial parallel transport, and so the states related to this trivial state by gauge transforms must also be trivial. This indicates that in these equivalent states the apparently non-trivial edge labels only describe a change

of basis, rather than a physical change under parallel transport. On the other hand, if a state has any closed loops with non-trivial path label, then (because the conjugacy classes of closed path labels are gauge-invariant) the state cannot correspond to this trivial state. Therefore, in such a state the parallel transport across the edges must describe both a change of basis and some physical "flux", analogous to the magnetic flux in electromagnetism, which differentiates it from the trivial case.

2.1.3 The Quantum Double Model

Having now considered lattice gauge theory, we can build a model for topological phases from it, known as Kitaev's Quantum Double model [16]. The lattice represents the spatial dimensions of the models, while a Hamiltonian controls the time evolution. In order to construct the Hamiltonian, we first demote the gauge symmetry to an energetic constraint. We do this by adding an energy term to the Hamiltonian for each vertex, with the energy term enforcing the symmetry. We also energetically punish closed paths with non-trivial labels, by adding an energy term at each plaquette that punishes plaquettes with non-trivial boundary paths. The Hamiltonian is [16]

$$H = -\sum_{\text{vertices, } v} A_v - \sum_{\text{plaquettes, } p} B_p.$$

Here we have

$$A_v = \frac{1}{|G|}\sum_{g\in G} A_v^g,$$

where the A_v^g are the gauge transforms from earlier and $|G|$ is the number of elements in the discrete group G. A_v is therefore an average over all gauge transforms at vertex v. The operator A_v is a projector [16], because

$$
\begin{aligned}
A_v A_v &= \frac{1}{|G|^2} \sum_{g \in G} \sum_{h \in G} A_v^g A_v^h = \frac{1}{|G|^2} \sum_{g \in G} \sum_{h \in G} A_v^{gh} \\
&= \frac{1}{|G|^2} \sum_{g \in G} \sum_{g' = gh \in G} A_v^{g'} = \frac{1}{|G|} \sum_{g \in G} A_v \\
&= A_v. \qquad (2.1)
\end{aligned}
$$

As a projector, A_v has eigenvalues of zero and one, with the eigenvalue of one corresponding to states which are gauge symmetric at that vertex (because the gauge transforms leave such states unchanged). A_v enters the Hamiltonian with a minus sign, so the gauge-invariant states are lower in energy.

The other term in the Hamiltonian, B_p, acts on the edges around a plaquette p. It leaves states where the boundary of the plaquette is labelled by the identity unchanged and returns zero for other states. As an example, consider Figure 2.4, which illustrates the action of the plaquette term B_p on a simple plaquette made from two edges (a bigon). In this case the boundary path label is given by $g_1 g_2^{-1}$, and so the plaquette term returns the state if $g_1 g_2^{-1} = 1_G$. B_p is clearly a projector just like the vertex term, with the eigenvalue of one corresponding to states with trivial flux around the plaquette (we say the plaquette satisfies flatness in these states). Again, B_p enters the Hamiltonian with a minus sign, so that these trivial flux states are lower in energy. The trivial flux label 1_G is in a conjugacy class on its own, meaning that it is unchanged by gauge transforms. This means that the operator B_p is built out of gauge-invariant quantities and therefore commutes with the gauge transforms. All of the terms in the Hamiltonian are projectors and they all commute, so this is an example of a commuting projector model. This structure to the Hamiltonian enables the model to be solved exactly. The excitations are charge-like (excitations of the vertex term), flux-like (primarily excitations of the plaquette term, though they may also excite a vertex term), or some combination of the two [16]. These excitations are called electric if they

are charge-like, magnetic if they are flux-like and dyonic if they are a combination.

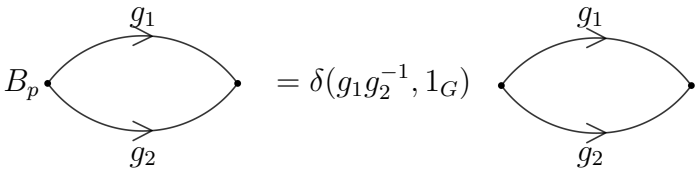

Figure 2.4: The plaquette term B_p gives 1 if the closed path forming the boundary of the plaquette p is 1_G (i.e. if it is flat) and 0 otherwise. For the example plaquette shown in this figure, where the edges are in states labelled by g_1 and g_2, that means that acting with the plaquette term gives a non-zero result only if $g_1 g_2^{-1} = 1_G$.

Even from this brief description of the Quantum Double model, it is possible to see how it uncovers the topological properties of the manifold on which we place it. For example, the plaquette terms project to the case where the smallest closed paths, the boundaries of the plaquettes, are labelled by the identity element in G. Because the larger closed paths can be constructed from products of these individual units, in the ground state space all contractible paths have trivial flux. However, if a manifold possesses non-contractible cycles, these can support non-trivial flux in the ground state.

We have seen how lattice gauge theory can be used to define a Hamiltonian model for topological phases, known as Kitaev's Quantum Double model [16]. This model is primarily defined in 2+1d, though it can also be extended to 3+1d (see e.g. Ref. [49], where Kitaev's Quantum Double model is both extended to 3+1d and generalized). The higher lattice gauge theory model that we will consider in this work is built in a similar way from a generalization of lattice gauge theory, called higher lattice gauge theory. Because of this, many of the features of Kitaev's Quantum Double model will carry over to the higher lattice gauge theory model considered in the rest of this work.

2.2 Higher Lattice Gauge Theory

We have now looked at the general principles of ordinary lattice gauge theory and its associated model for topological phases. Next we can consider higher lattice gauge

theory [72], before moving on to its associated model for topological phases [58]. As stated before, higher lattice gauge theory is a generalization of lattice gauge theory [72, 73]. In lattice gauge theory we consider parallel transport along paths, and label paths by group elements to allow composition of paths [73]. That is, we label geometric objects, the paths, with algebraic objects, the group elements. A natural generalization is to label more of the geometric objects. We still label the paths (now calling their associated gauge field the 1-gauge field), but now also consider labelling objects of a higher dimension: the surfaces. If paths describe the parallel transport of points, then surfaces describe the parallel transport of paths, that is of the 1-gauge fields [73]. We can view this pictorially as shown in Figure 2.5. The blue double arrow on the surface enclosed by the paths represents the transport of one path (the source) into another (the target) [58]. The two paths (source and target) start and end at the same points (vertices) and the start of these two paths is called the base-point of the surface, while the end of these paths is called the end-point.

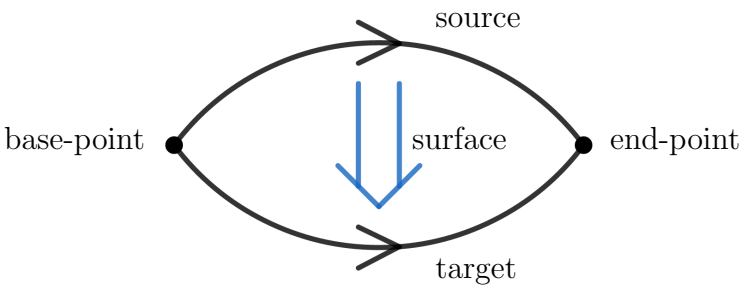

Figure 2.5: Just as a path is associated to parallel transport of points, so is a surface associated to parallel transport of a path. The initial position of the path is called the source and the final position is called the target.

In the same way that we can compose paths that lie end-to-end, so may we combine adjacent surfaces. In fact, surfaces can be composed in two ways. Firstly, they may be combined vertically [72, 73], as shown in Figure 2.6. Vertical composition corresponds to the case where we perform two parallel transportations of a path (the top path in Figure 2.6) in sequence (first moving it to the middle position in the figure and then to the bottom). We can combine these two steps to describe the two parallel transportations as parallel transport along a single, combined, surface. Note that in

order to compose the two surfaces in this way, the target of the first surface must match the source of the second one. Then the source of the combined surface is the source of the first surface and the target of the combined surface is the target of the second surface.

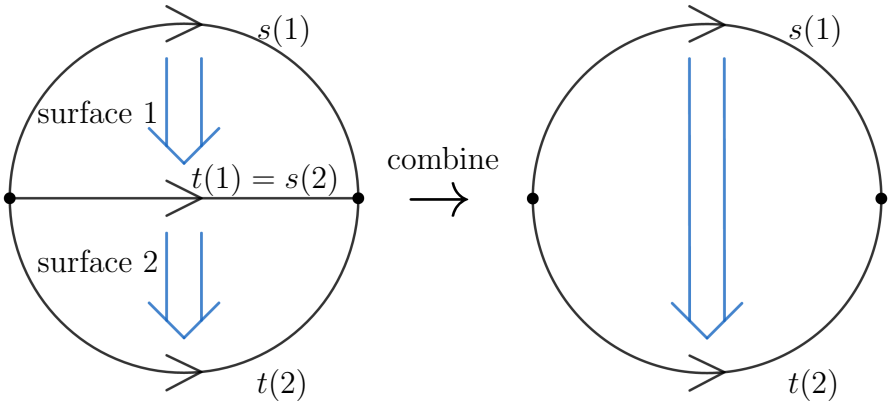

Figure 2.6: Consider two surfaces over which we can sequentially transport a path, such as the ones in the left side of the figure. In this case we first transport the top path over the upper surface (across the arrow) to the middle location (the straight path) and then over the lower surface to the bottom position. We can express the same process as transport over a single surface, made from a combination of the two individual surfaces, as shown in the right figure. This is called vertical composition of the surfaces. Note that in order to compose the surfaces, the target $t(1)$ of the first surface must match the source $s(2)$ of the second.

We may also combine the surfaces horizontally [72, 73], as shown in Figure 2.7. This horizontal combination corresponds to the case where we have two paths lying end to end, which we can parallel transport separately. However, we can also combine the two paths into one, before transporting them across a single surface. In this case the source of the combined surface is the composition of the sources of the two individual surfaces and similarly the target of the combined surface is the composition of the targets of the two surfaces.

As a special case of horizontal combination, we have the case where the first path is not parallel transported across any surface. This lets us combine a surface with a path. As an example, such a situation is shown in Figure 2.8. In Figure 2.8, we combine the edge that runs from A to B with a surface, by treating the edge and

its inverse (the inverse is the same edge, but with reversed direction) as bounding an infinitesimally thin surface and then using horizontal composition. This process of combining a surface with a path is known as whiskering [73], and can also be thought of as moving the base-point of a surface (in the case shown in Figure 2.8, the base-point of the surface is initially at B, but is moved to A, where in each case the base-point is represented by the blue dot).

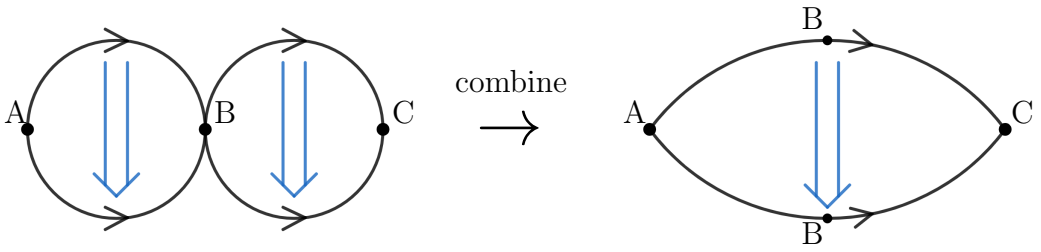

Figure 2.7: In addition to vertical composition of surfaces, as shown in Figure 2.6, we can consider horizontal composition of two surfaces that lie side by side. In this case, the individual surfaces describe parallel transport of two paths (A-B and B-C) which can be composed, while the combined surface describes the parallel transport of the paths after they have been composed (into A-C). While we show the resulting surface as a simple bigon (2-gon) for convenience, the composition does not change the shape of the constituent surfaces on the lattice, and so the point B remains on the boundary of the combined surface. The two points labelled B in the right-hand side are the same and so should be glued together.

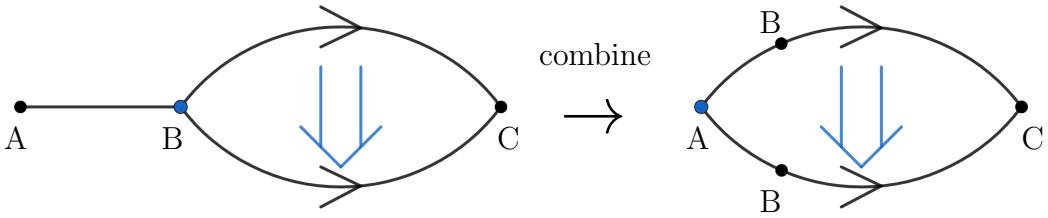

Figure 2.8: Here we consider "whiskering" a surface (the bigon with base-point B and end-point C) along a path. Just as for horizontal composition, the shape of the final surface should match the combined shape of the path and surface, though we have represented the final surface as a simple bigon. The point B remains on the boundary of the final surface, in both the source and target. Note that the base-point (which is represented by the blue dot) of the surface changes from B in the first image to A in the second.

Given this geometric structure for the paths and surfaces, we must now determine the appropriate algebraic structure with which to label it. As before, we label the

edges with elements of a group G (the 1-gauge field), to enable composition. The label assigned to a directed path is called the 1-holonomy of that path [58]. We also label the plaquettes with elements of another group, E [58, 72, 73]. Labelling every plaquette in this way gives us a field which is valued in E at each plaquette, with composition rules allowing us to find the label for an arbitrary surface. We refer to this field as the 2-gauge field and call the label associated to a surface with given source and target a 2-holonomy [58]. As we will see shortly, parallel transport will involve various mappings between the groups E and G.

Armed with the algebraic objects which describe our fields, we go back to our diagram defining parallel transport across a surface (Figure 2.5) and label it, giving the situation shown in Figure 2.9. Figure 2.9 then describes how parallel transport of a 1-gauge field across a surface depends on the 2-gauge field at that surface.

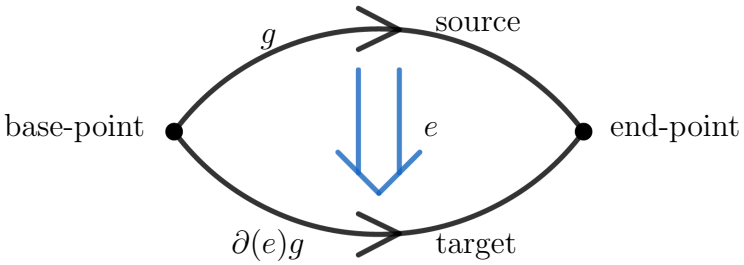

Figure 2.9: We can label the geometric objects in Figure 2.5 with algebraic quantities from the groups G and E. The paths are labelled by elements of G and the surfaces by elements of E. The parallel transport of a path over a surface labelled by $e \in E$ results in the path element g gaining a factor of $\partial(e)$, where ∂ is a group homomorphism from E to G.

In Figure 2.9, ∂ is a group homomorphism from E to G (i.e. a map that preserves the group multiplication), so that for $e \in E$, $\partial(e) \in G$ [73]. As we will now show, this is required to be a group homomorphism in order to make the vertical composition diagram consistent. We map vertical composition of two surfaces onto the group multiplication, with the multiplication of the labels in E done from right to left, following the convention in [58]. An example of this is shown in Figure 2.10. Requiring the label

of the bottom path to be the same on both sides of Figure 2.10 gives the consistency condition $\partial(e_2 \cdot e_1) = \partial(e_2)\partial(e_1)$, which is the condition for a group homomorphism. This ensures that the effect of transporting the edge along one surface, labelled by e_1, and then another surface, labelled by e_2, is the same as transporting the edge along the combined surface (labelled by $e_2 e_1$). It is worth noting that in Ref. [73] the horizontal composition, rather than the vertical composition, is described by the group multiplication.

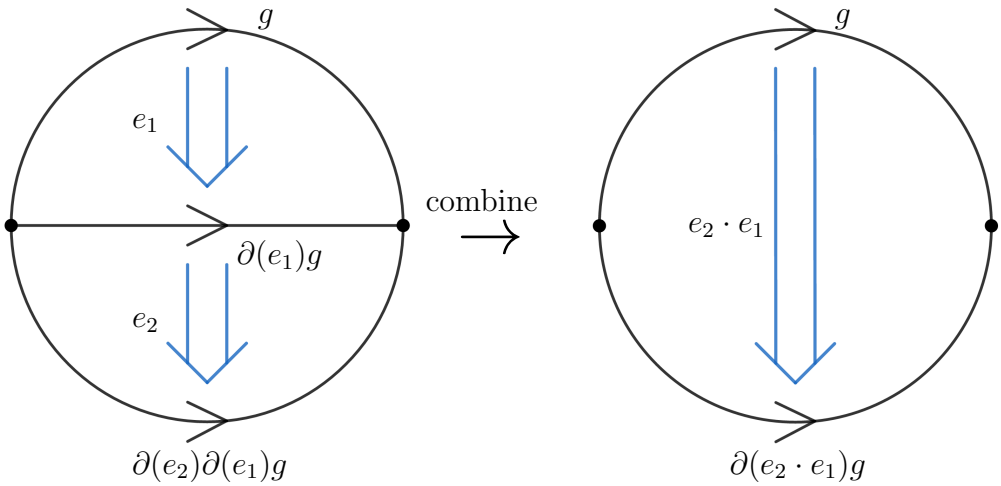

Figure 2.10: Vertical composition of two 2-holonomies is described by the group multiplication in E. Note that requiring this composition diagram to be consistent means that ∂ is a group homomorphism (because we must have $\partial(e_2 \cdot e_1) = \partial(e_2)\partial(e_1)$ to give a consistent label for the lower path in the left and right sides of the figure).

We need a second group-theoretic quantity to describe whiskering [58, 73] (see Figure 2.8). Whiskering can be thought of as parallel transport of the base-point of a surface along an edge. Because parallel transport of objects along paths is described by the group G, we need some way for G to act on E. This is handled by a map $\triangleright$ from G to the endomorphisms on E [58, 73] (endomorphisms are homomorphisms from a group to itself). That is, given an element g of G, the object $g \triangleright$ is then a map from E to itself. We write $g \triangleright$ acting on an element e of E as $g \triangleright e$, so that $g \triangleright e$ is an element of E. In Figure 2.11, we see how this map $\triangleright$ is involved in whiskering. If we move the base-point of a surface (initially B in the Figure) across an edge labelled by g, against the direction of the edge, then the surface label changes from e to $g \triangleright e$. On the other

hand, moving the end-point (C in the figure), rather than the base-point, has no effect on the surface label.

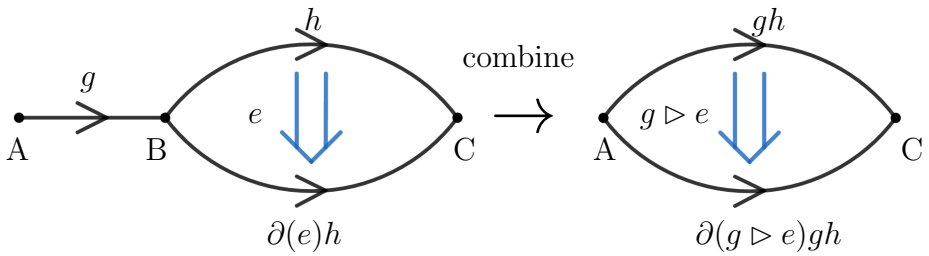

Figure 2.11: Whiskering is described by the map $\triangleright$. Requiring the lower path from A to C to have the same label after combining the edge with the surface via whiskering gives the condition that $g\partial(e)h = \partial(g \triangleright e)gh$, so that $\partial(g \triangleright e) = g\partial(e)g^{-1}$.

For our diagrams to give a consistent theory, the different ways of combining the elements of the diagram must be consistent. One consequence of this is that we can find the result of horizontal composition of two surfaces, by combining the rules for vertical composition and whiskering. Consider Figure 2.12, which shows a diagram involving the horizontal composition of two surfaces (on the top line), where the left and right surfaces are labelled by group elements e_L and e_R respectively. We can reproduce this horizontal composition with a series of other manipulations, which takes us the other way around the diagram. These other processes involve changing the base-point and end-point of surfaces, as well as vertical composition, all of which we already know how to perform. Applying these manipulations (as explained in Figure 2.12), we find that the label resulting from horizontal composition must be $e_L[g_L \triangleright e_R]$.

In addition to allowing us to define horizontal composition without introducing additional algebraic objects, requiring the consistency of various diagrams enforces certain restrictions on the algebraic objects we already have. For example, if we have a diagram with three surfaces to combine, the order in which we combine the surfaces should not matter. This restricts our multiplication of surface labels to be associative. Because our vertical composition is described by group multiplication in the group E, this associativity is immediately guaranteed by the group properties without any

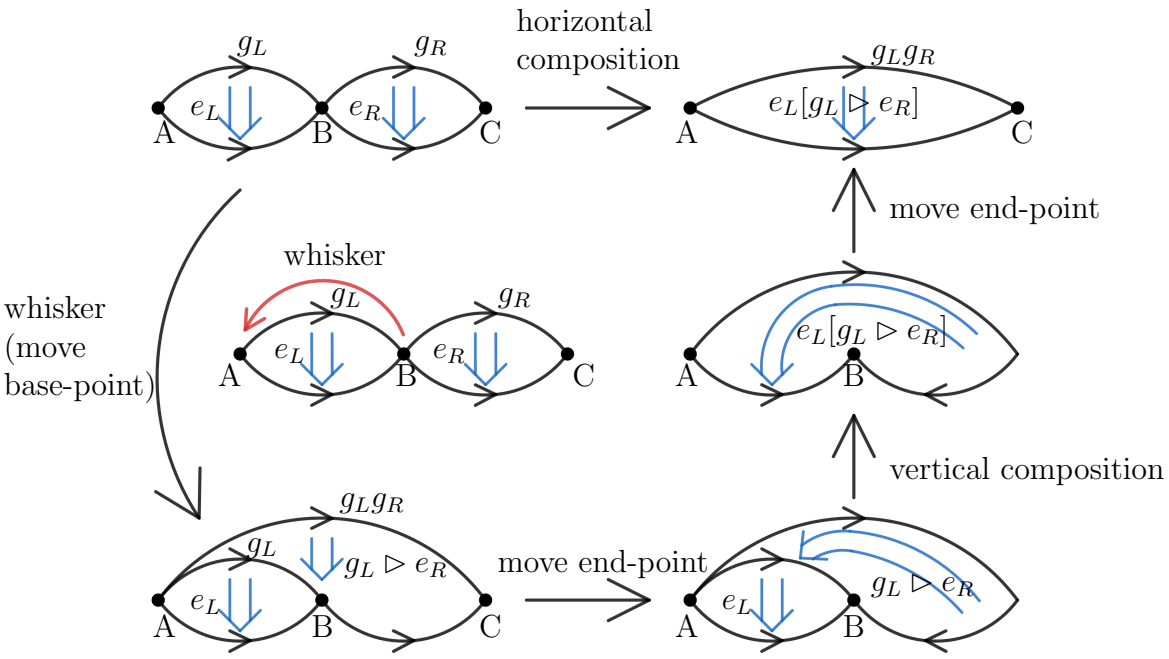

Figure 2.12: The requirement that the different ways of combining a diagram must be consistent means that we can express the horizontal composition of two surfaces in terms of whiskering and vertical composition. Consider the top-left image, consisting of a left and right surface with label e_L and e_R respectively. We wish to combine these two surfaces to obtain the surface in the top-right image. There are two ways around the diagram that lead from the top-left to the top-right image, and these are required to give the same result for the label of the final surface. The first way is horizontal composition of the two surfaces, which gives us an unknown label for the final surface that we wish to find. The second way around the diagram involves the other processes that we do have algebraic expressions for. We can therefore use this to find the label resulting from horizontal composition. The first step (represented by the downwards arrow) is to whisker the right-hand surface so that it has the same base-point (A) as the left surface. This gives the right surface the label $g_L \triangleright e_R$, where g_L is the label of the path from A to the original base-point B, as shown in the bottom-left image. The next step is to move the end-point of the right-surface from C to B, to match the end-point of the left surface, as shown in the bottom-right image. This has no effect on the label of the surface. Then the target of the right surface is the same as the source of the left surface (the path from A to B with label g_L) and so the two surfaces can be combined via vertical composition, giving a surface with label $e_L[g_L \triangleright e_R]$ as shown in the centre-right image. Finally, moving the end-point from B to C (the original end-point of the right surface) gives us the same surface that we would have from horizontal composition, with label $e_L[g_L \triangleright e_R]$. This is therefore the label resulting from horizontal composition.

---

additional conditions on the group. However, there are additional constraints that must be satisfied by the maps ∂ and $\triangleright$. Requiring the consistency of whiskering with vertical composition of surfaces and composition of paths (see Figures 2.13 and 2.14

or Ref. [73] for more detail) gives us the following conditions for all g, $h \in G$ and e, $f \in E$ [58]:

$$g \triangleright (ef) = (g \triangleright e)\,(g \triangleright f) \tag{2.2}$$

$$g \triangleright (h \triangleright e) = (gh) \triangleright e. \tag{2.3}$$

These are the conditions for a group action of G on E. That is, these conditions mean that $\triangleright$ is a homomorphism from G to the endomorphisms on E, where endomorphisms are group homomorphisms from E to itself. Furthermore, because these endomorphisms are invertible (from Equation 2.3, $g^{-1} \triangleright$ is the inverse of $g \triangleright$), they are automorphisms.

As illustrated in Ref. [73] (though note the different conventions used in this reference), consistency of whiskering with other diagrams (see Figures 2.11 and 2.15) also demands that [72, 74]

$$\partial(g \triangleright e) = g\partial(e)g^{-1} \tag{2.4}$$

$$\partial(e) \triangleright f = efe^{-1}. \tag{2.5}$$

These two conditions are known as the Peiffer conditions [58]. The algebraic structure $(G, E, \partial, \triangleright)$ satisfying all of these conditions (Equations 2.2, 2.3, 2.4 and 2.5 in addition to ∂ being a group homomorphism) is known as a crossed module. That is:

Definition 1: A crossed module is a collection $(G, E, \partial, \triangleright)$, where G and E are groups, and $\partial : E \to G$ and $\triangleright : G \to \mathrm{Aut}(E)$ are group homomorphisms satisfying the Peiffer conditions Equations 2.4 and 2.5.

As an example, consider the crossed module $(Z_2, Z_3, \partial \to 1_G, \triangleright)$ [58]. We take the elements of $G = Z_2$ to be 1_G and -1_G and the elements of $E = Z_3$ to be 1_E, ω_E and ω_E^2. Then we define $\triangleright$ by $1_G \triangleright e = e$ and $-1_G \triangleright e = e^{-1}$ (where $\omega_E^{-1} = \omega_E^2$).

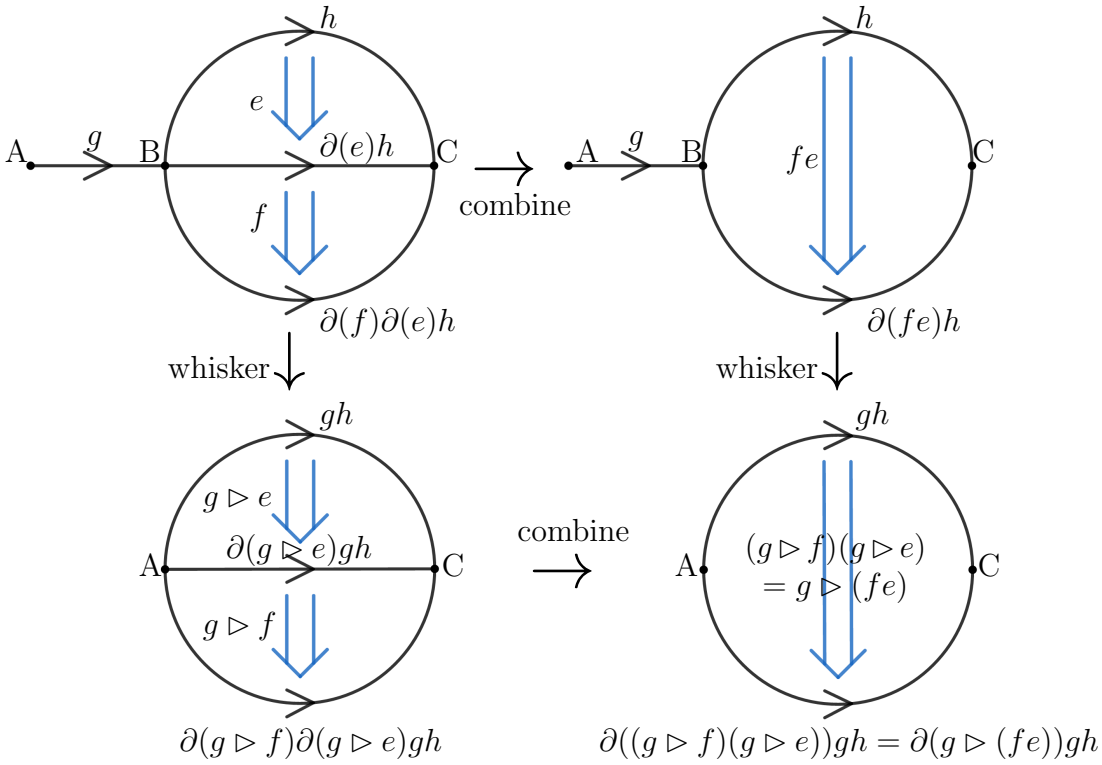

Figure 2.13: Requiring consistency of whiskering with the vertical composition of surfaces demands that the map $g \triangleright$, for arbitrary $g \in G$, is a group homomorphism on E. That is $(g \triangleright f)(g \triangleright e) = g \triangleright (fe)$ for $g \in G$ and $e, f \in E$. This can be seen from the figure, because consistency demands that the diagram commute, i.e. the two routes from the top-left image to the bottom-right image should give the same result.

This satisfies the requirement of having a group structure on the elements of G (as described in Equation 2.2) because applying two $-1_G \triangleright$ maps in sequence gives

$$\begin{aligned} -1_G \triangleright (-1_G \triangleright e) &= -1_G \triangleright e^{-1} = e \\ &= 1_G \triangleright e \\ &= (-1_G \cdot -1_G) \triangleright e, \end{aligned}$$

while the other conditions for the group structure involve $1_G \triangleright$ and are satisfied because $1_G \triangleright$ is the identity map. The individual maps $g \triangleright$ are also endomorphisms as required.

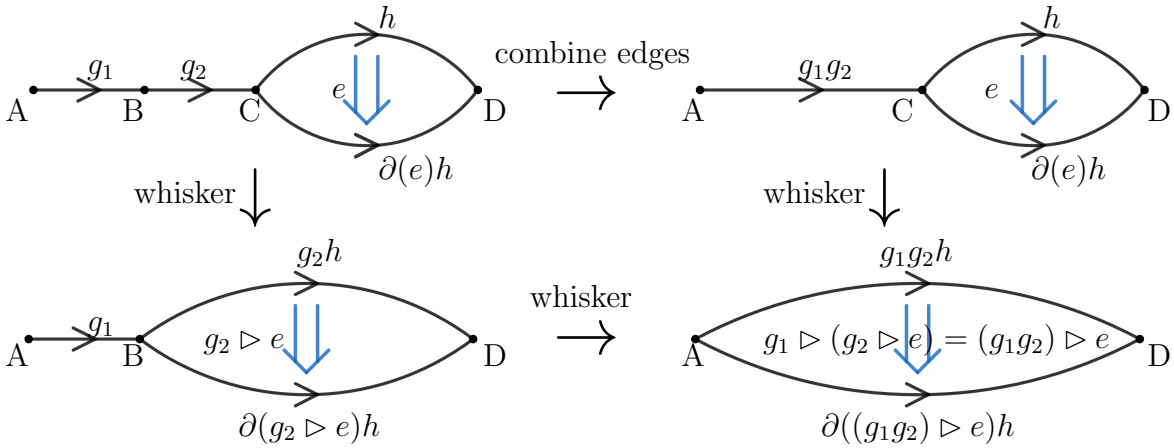

Figure 2.14: We require that whiskering is consistent with the composition of edges. That is, combining two edges together and then whiskering a surface along the combined edge should give the same result as whiskering that surface by one edge and then the other in sequence. This gives us the mathematical condition $g_1 \triangleright (g_2 \triangleright e) = (g_1 g_2) \triangleright e$, which is the condition that the map $\triangleright : G \to \mathrm{End}(E)$ be a group homomorphism on G.

For $-1_G \triangleright$ we have

$$\begin{aligned} -1_G \triangleright (e_1 e_2) &= e_2^{-1} e_1^{-1} \\ &= e_1^{-1} e_2^{-1} \\ &= (-1_G \triangleright e_1)(-1_G \triangleright e_2), \end{aligned}$$

where we used the fact that E is Abelian to swap the order of multiplication in the second line. This indicates that $-1_G \triangleright$ is a group homomorphism on E. $1_G \triangleright$ is also a homomorphism because it is the identity map. Therefore, $\triangleright$ is indeed a group action of G on E. Next, we will check that the Peiffer conditions are satisfied. We have $\partial(g \triangleright e) = 1_G = \partial(e) g g^{-1} = g \partial(e) g^{-1}$ (using $\partial(e) = 1_G$ and the fact that the group G is Abelian). Finally $\partial(e) \triangleright f = 1_G \triangleright f = f = f e e^{-1} = e f e^{-1}$, where we used that E is Abelian. Because all of the consistency conditions are satisfied, this is indeed a valid crossed module.

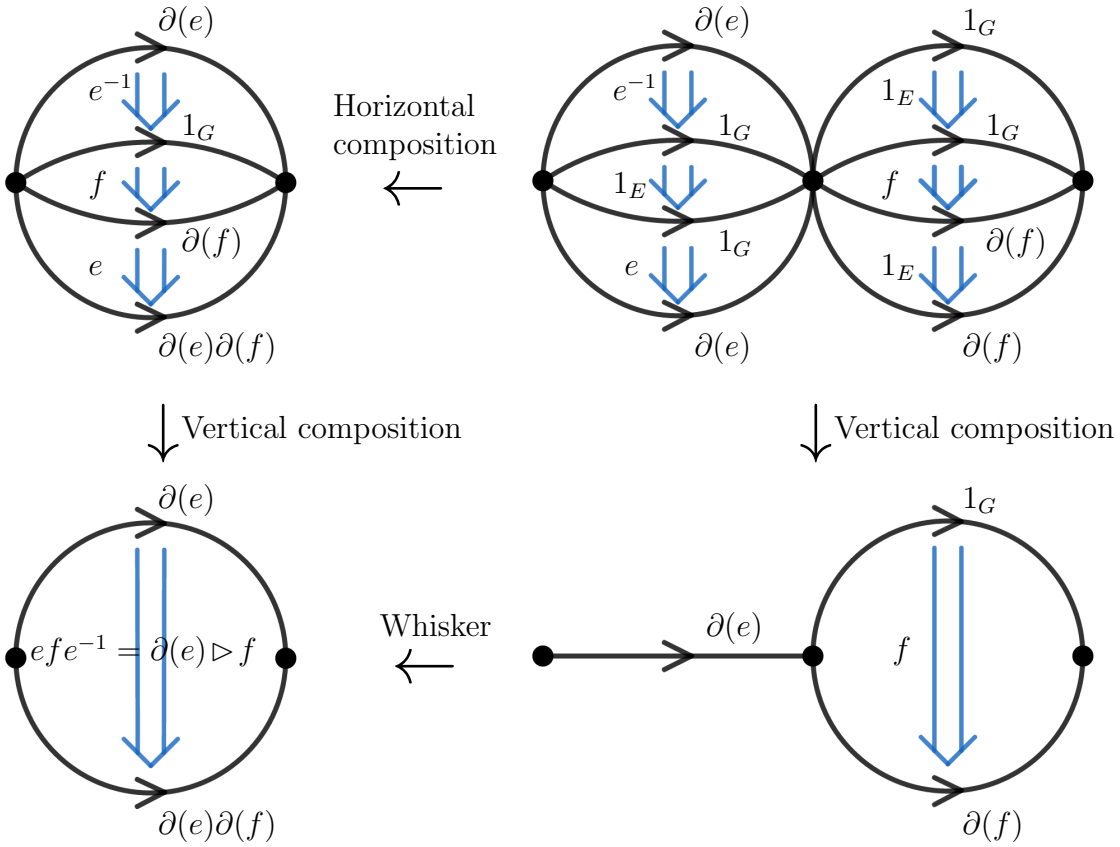

Figure 2.15: The second Peiffer condition can be derived from demanding that this figure be consistent. Starting at the top-right diagram, we may combine the left and right parts of the diagram to obtain the top-left diagram, using the rules for horizontal composition given in Figure 2.12. We can then use vertical composition to obtain the bottom-left diagram, which should have a label of efe^{-1}. However, we could also have performed vertical composition on the top-right diagram to obtain the bottom-right diagram, before whiskering to obtain the bottom-left diagram. In that case the surface label is $\partial(e) \triangleright f$. Consistency therefore demands that $efe^{-1} = \partial(e) \triangleright f$.

2.2.1 Composing General Surfaces

So far, we have considered how we may combine surfaces when their sources and targets are compatible. That is, we know that we can combine two surfaces using vertical composition when the target of one surface matches the source of the other. However when we consider a lattice model, we may need to combine adjacent surfaces for which the sources and targets are not compatible. To understand this, we should first look in more detail at how we interpret the 2-holonomy in the case of a fixed lattice. When we label a surface with a group element, that element corresponds

to parallel transport of a particular path along that surface. It does not immediately describe the parallel transport of a different path across the same surface. For example, consider a square, with different paths denoted as the source or target, as shown in Figure 2.16. In the left diagram, the surface corresponds to the case where we deform the top edge (which is the source for the surface) into the bottom three (which form the target). However, as indicated in the right diagram, we could also transport the left edge into the right three over the same surface. Despite corresponding to the same square in space, the label in E associated with these two parallel transports is different in general. If we wish to relate the label of one to the label of the other, we use certain rules that let us change the source and target and give us the appropriate label for the new parallel transport process [58]. The first thing we can do is to swap the source and target. The resulting plaquette label is just inverted [58], as shown in Figure 2.17.

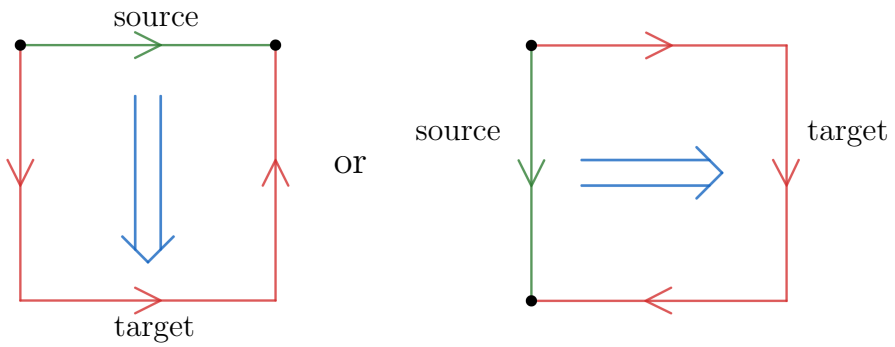

Figure 2.16: The same surface can correspond to different 2-holonomies, depending on which parts of the boundary of that surface are designated as the source and target. For example, in the left image the source is the top edge and the target the bottom three edges, so the 2-holonomy corresponds to a process where we deform the upper edge into the bottom three. On the other hand, in the right-hand figure the 2-holonomy corresponds to the process where we transport the left edge into the other three. We expect the labels assigned to these processes to be different but related, as we describe shortly.

Next, we can move the base-point around. We can either move it along the plaquette (as shown in Figure 2.18), or away from the plaquette (as shown in Figure 2.19) [58]. In either case, the surface label changes from its original label e_p to $g(t)^{-1} \triangleright e_p$, where t is the path along which we move the base-point and $g(t)$ is the group element assigned to that path [58].

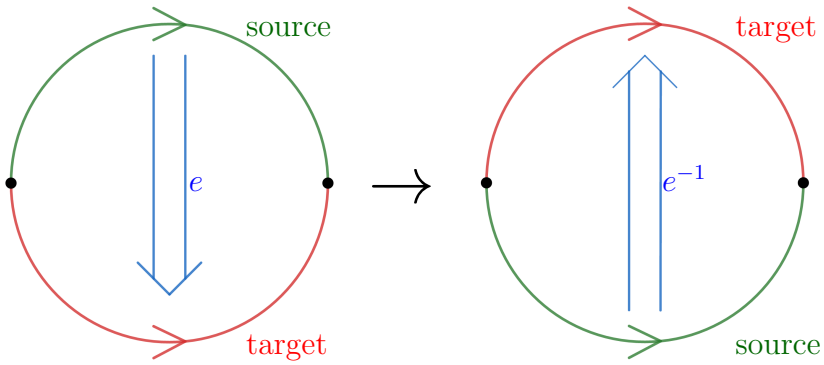

Figure 2.17: We can change the source and target of a surface by following a set of rules that tells us how the label of that surface should change. The first rule allows us to swap the source and target of a surface. If the plaquette has a 2-holonomy of e, then swapping the source and target changes the 2-holonomy to e^{-1}.

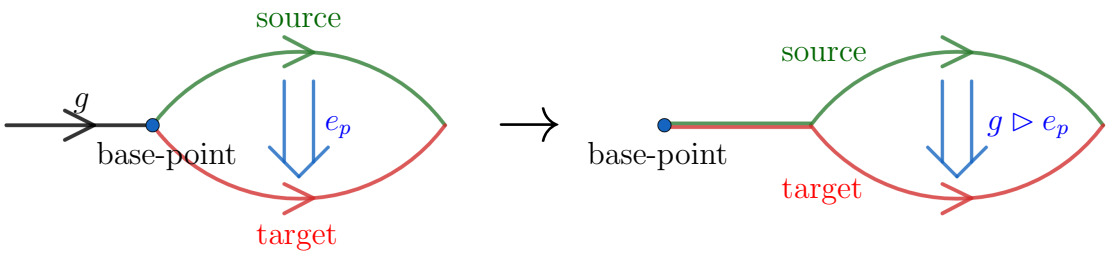

Figure 2.19: We can whisker a surface by moving its base-point away from the original boundary of that surface. In the right image, the red and green section (which is the black path from the left image) is part of both the source and target.

We can also move the end-point, either along the plaquette, or away from it [58]. An example of the latter case is shown in Figure 2.20. Either way, the plaquette label is unchanged [58]. This latter move (as shown in Figure 2.20) allows us to add additional edges to the boundary of the surface, though these additional edges enclose no area. Though in Figure 2.20 the edges are added near the end-point, we can add these additional edges anywhere on the surface's boundary. If these edges are not added at the end-point, the added edges appear twice consecutively in the source or target and are travelled in opposite directions for their two appearances, meaning that they do not contribute to the path element of the source or target (because adding a path t to the surface in this way contributes $g(t)g(t)^{-1} = 1_G$ to the source or target). If the edges are added at the end-point, they contribute equally to the end of the source and

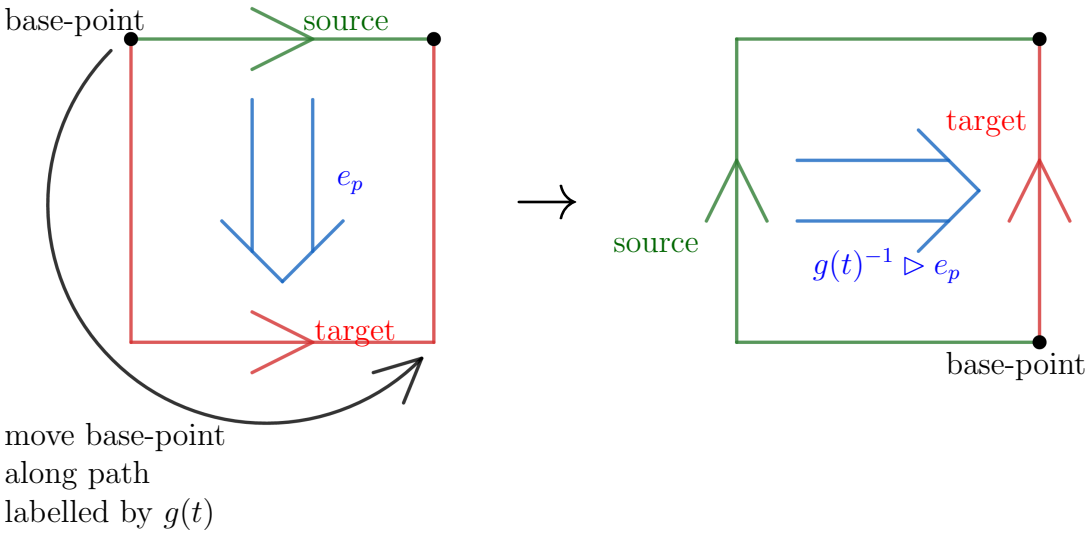

Figure 2.18: Another useful manipulation of the source and target is to move the base-point of our surface along the boundary of that surface, which adds or removes edges from the start of the source (and removes or adds those edges to the target). This results in a $\triangleright$ action on the surface label.

---

target. Either way, their contribution to the group element associated to the surface boundary (the 1-gauge value assigned to the path around the surface) cancels. For example, in Figure 2.20 adding the edge of label x to the end-point takes the path label of the boundary from

$$g(\text{boundary}) = g(\text{source})g(\text{target})^{-1}$$

to

$$\begin{aligned} g(\text{source})x(g(\text{target})x)^{-1} &= g(\text{source})xx^{-1}g(\text{target})^{-1} \\ &= g(\text{source})g(\text{target})^{-1} \\ &= g(\text{boundary}), \end{aligned}$$

whereas adding such an edge in the middle of the source or target would lead to similar cancellation within $g(\text{source})$ or $g(\text{target})$.

Now we consider an example of how we can use the rules we have discussed so far to

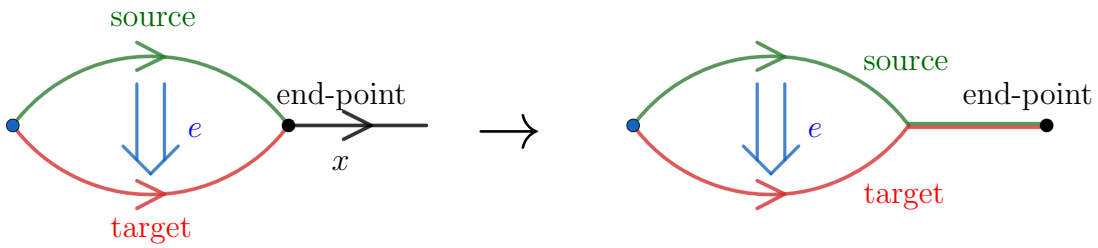

Figure 2.20: We can move the end-point (the black dot) away from the original boundary of the surface, thereby adding edges (in this case the black path from the left image) to the boundary. These edges do not enclose any area and appear once in the source and once in the target (in the right image, the green and red path is part of both the source and target). Moving the end-point in this way does not change the label of the surface, unlike moving the base-point.

---

combine two surfaces when their sources and targets are not immediately compatible. In Figure 2.21 we show two such adjacent surfaces. For each surface, the source is represented by the solid green line and the target by the dashed red one, and we have displaced the source and target slightly away from the edges of the graph (shown in black) for clarity. In order to match the target of the first surface (with surface label e_1) to the source of the second (labelled by e_2), we first move the end-point of the second surface, as shown in the top-right of Figure 2.21. Because moving the end-point of the surface does not affect its label, the second surface still carries a label of e_2. Next we move the base-point of the second surface to match that of the first, as shown in the bottom-right image. When we do this, we must whisker the second surface, so that the path t (the edge at the bottom of the first surface) appears in both the source and target of the second surface (represented by the parallel red and green arrows below that edge). Upon doing so, the label of the second edge is changed to $g(t) \triangleright e_2$, because t is the path from the new base-point of the surface to the old one. By moving the base-point and end-points in this specific way, we ensure that the target of the first surface matches the source of the second (consisting of the bottom edge of the first surface and the edge separating the two surfaces), so we can compose the surfaces. This gives us a combined surface with label $[g(t) \triangleright e_2]e_1$. In general, there may be many ways to combine a given set of surfaces into the same final surface (i.e., a final surface with the same source and target). These are guaranteed to be consistent

only when the surfaces that we are combining satisfy an additional condition called fake-flatness, meaning that each surface obeys the parallel transport rules given in Figure 2.9, which we will discuss further in Section 2.2.5.

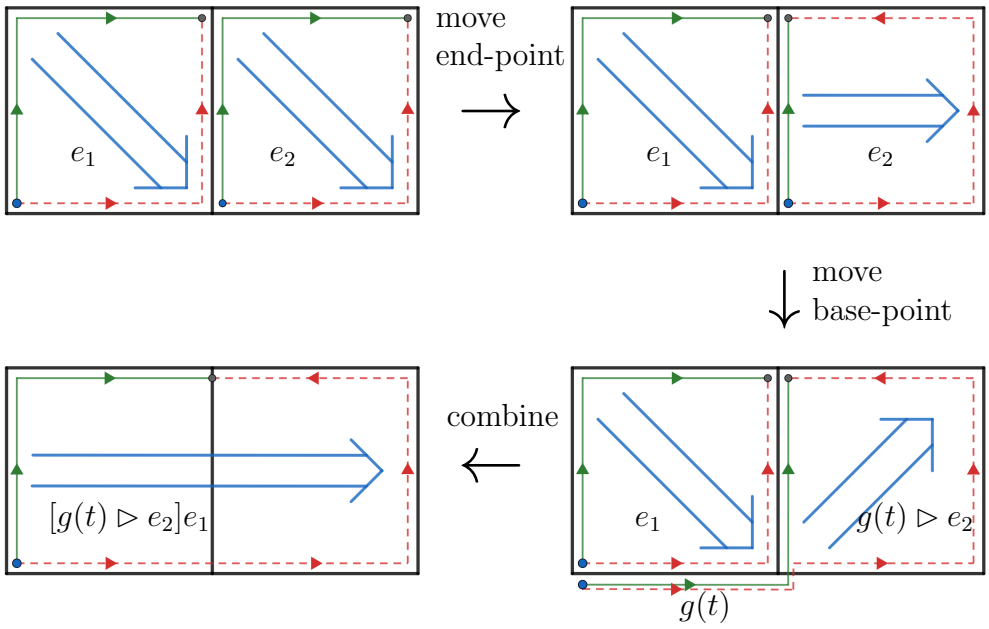

Figure 2.21: In order to combine two adjacent surfaces whose sources (shown in green) and targets (shown as red dashed lines) are not compatible, we need to manipulate the end-point and base-points of the surfaces first. In the first step, we move the end-point of the second surface (with label e_2), then in the second step we move the base-point of that surface. After doing this, the target of the first surface matches the source of the second, so we can combine them (as shown in the final step).

2.2.2 A Note About Notation

So far, when describing surfaces we have specified both the source and target of the surface. However, the fact that the label of a surface is unchanged when we move the end-point of the source and target means that we do not need to keep track of all of the information specifying a surface in order to be able to assign that surface a group label. This motivates us to consider a change of notation. Rather than specify the source and the target as two paths, with an arrow between them to highlight the parallel transport, we simply combine the source with the target by moving the end-point all the way along the target (so that the new source is now the original source composed

with the inverse of the original target, and the new target is an empty path). This means that we now just have one path all the way around the surface. To specify this, we only need the start of that path (the base-point) and its orientation. Rather than draw an arrow, we indicate this as a circulation, as shown in Figure 2.22. We note that this circulation labelling scheme is not used in the original paper on the higher lattice gauge theory model [58], but we will often use it in this thesis for its convenience, at least when we do not need to indicate the source and target of a surface explicitly.

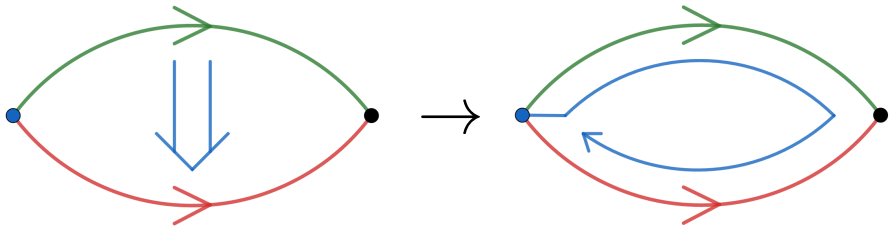

Figure 2.22: Instead of illustrating the orientation of the surface as an arrow between the source and target, we can draw it as an arrow that passes clockwise or anticlockwise around the surface, starting and ending at the base-point of the surface. The direction of this arrow matches the direction of the source.

2.2.3 Gauge Transforms

Now that we have considered the fields and the parallel transport rules, it is time to look at the gauge transforms. We have gauge transforms associated to the more familiar 1-gauge field in addition to transforms associated to the 2-gauge field [58]. The 1-gauge transforms are similar to the gauge transforms in ordinary lattice gauge theory. As before, we label the 1-gauge transforms associated to a vertex v by A_v^g, where we have one such transform for each element g of G. A_v^g acts on the degrees of freedom near the vertex v in a way equivalent to parallel transport of the vertex v along an edge of label g^{-1} (or g if we transport the vertex against the direction of the edge, as in Figure 2.23). The effect on the edges around the vertex is therefore the same as in the lattice gauge theory case and, just as in the lattice gauge theory case, only paths that start or terminate on the vertex are affected by the gauge transform (see Section 2.1.1 and Figure 2.2 in particular). The only difference is that now we

must also consider parallel transport of surfaces along the edge, so that the vertex transform also affects the surface labels. This parallel transport can be performed by adding a new edge and vertex, which we proceed to combine with the rest of the lattice (so that the lattice is the same at the end as it was before the transform, apart from changes to the group labels), as illustrated in Figure 2.23. In the last step we relabel the vertex v' to v in order to match the original vertex. We can recognise the middle diagram in Figure 2.23 as the whiskering diagram (see Figure 2.11), so combining the edge with the plaquette gives us a $g \triangleright$ action on the plaquette label. This tells us that any surface with base-point at the vertex on which we apply the transform must be acted on by $g \triangleright$. On the other hand, surfaces not based at that vertex are left unaffected [58]. In summary, the 1-gauge transform acts on an edge i or plaquette p according to [58]

$$
\begin{aligned}
A_v^g : g_i \to & \begin{cases} g g_i & \text{if } v \text{ is the start of } i \\ g_i g^{-1} & \text{if } v \text{ is the end of } i \\ g_i & \text{otherwise} \end{cases} \\
A_v^g : e_p \to & \begin{cases} g \triangleright e_p & \text{if } v \text{ is the base-point of } p \\ e_p & \text{otherwise.} \end{cases}
\end{aligned}
\tag{2.6}
$$

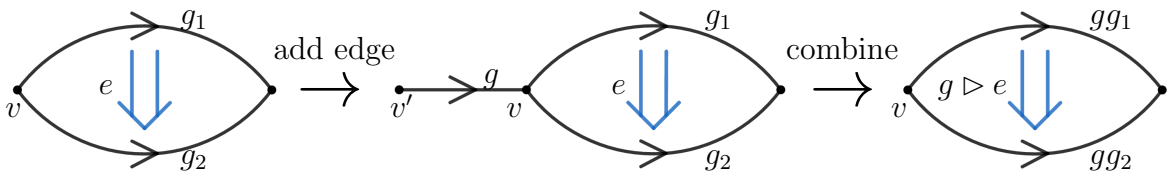

Figure 2.23: The effect of a vertex transform A_v^g on a plaquette attached to the vertex v depends on the base-point of that plaquette. If the base-point is at v, as in this case, the plaquette label obtains a $g\triangleright$ action, because the vertex transform is equivalent to parallel transport of the vertex v, and so of the base-point (i.e. the vertex transform is like whiskering the plaquette, as we see in the middle diagram). If the base-point of the plaquette is not at v, then the plaquette is not whiskered and its label is left unchanged.

In addition to these 1-gauge transforms, we also have 2-gauge transforms, which act

on an edge and the surfaces that adjoin it [58]. The 2-gauge transform on an edge i and labelled by an element $e \in E$ (denoted by $\mathcal{A}_i^e$) acts like parallel transport of the edge along a surface labelled by e (this follows from the action defined in Ref. [58]). Recall that for the ordinary gauge transforms we added an edge and combined it with the diagram, as shown in Figure 2.23. For the 2-gauge transforms, we instead add a surface and combine the surface with the rest of the diagram, as shown in Figure 2.24.

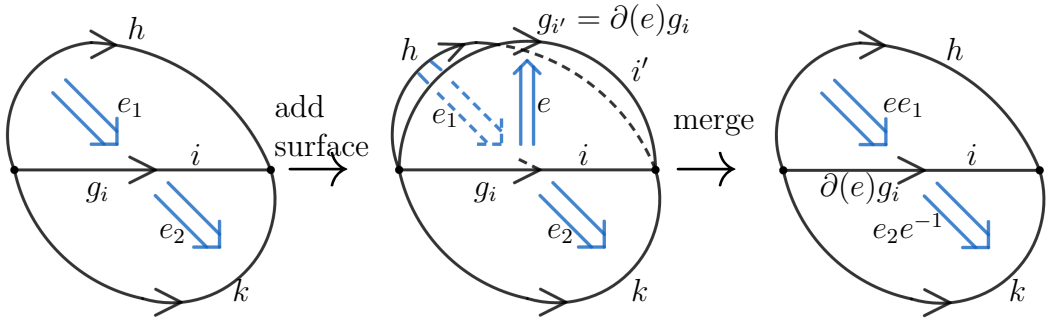

Figure 2.24: A 2-gauge transform $\mathcal{A}_i^e$ on an edge i, with initial label g_i, acts like parallel transport of that edge across an additional surface of label e. In the middle picture, the new surface points upwards, out of the plane of the other two surfaces. In the third picture, we combine this new surface with the others and then relabel the edge i' (the target of the additional surface) to i.

Note that in Figure 2.24, the base-point of each surface is also the start of edge i, which results in the simple expression for the edge transform given in that figure. If the base-points are not at the start of edge i, we can use the rules for changing base-points to find the appropriate transformation for the surface label. Given such a general case, we can move the base-points of the surfaces to the start of edge i using the procedure described earlier. Then we can perform the gauge transform $\mathcal{A}_i^e$ on this simple case before moving the base-point back. Because moving the base-point has an $\triangleright$ action on the plaquette label, this results in the plaquette label e_p becoming $e_p(g \triangleright e^{-1})$ or $(g \triangleright e)e_p$ [58] rather than just $e_p e^{-1}$ or ee_p, where g is the label of the path on which we had to move the base-point (we will give a more concrete description of this path shortly). While this action can look fairly complicated, the essence is that the edge transform fluctuates the plaquette labels surrounding an edge, as well as changing the edge label itself. This is similar to how the 1-gauge transform at a vertex fluctuates

the edges around the vertex (along with any plaquettes based at that vertex).

Now we want to describe the action of the edge operator resulting from this procedure algebraically. The edge transform acts on an edge i and neighbouring plaquettes. Consider the path around one of the plaquettes, starting at the base-point v_0 of the plaquette and travelling along its boundary, aligned with its orientation. This path reaches the edge at a vertex that we call v_i, as shown in the left picture of Figure 2.25. The path up to this point is denoted by $g(v_0 - v_i)$ [58]. Now consider the path starting at the base-point of the plaquette, but travelling the opposite way around the plaquette. At some point this path will reach the other vertex on the edge, which we call v_{i+1}. This path is denoted by $g(\overline{v_0 - v_{i+1}})$, where the overline is used to indicate that this path travels the "wrong way" around the plaquette [58], meaning against its circulation. The meaning of the overline notation is further illustrated in Figure 2.26, where we look at different paths around the plaquette to the same vertex. Then the action of the edge transform on each edge i' and plaquette p is [58]:

$$\mathcal{A}_i^e : g_{i'} \to \begin{cases} \partial(e) g_{i'} & \text{if } i = i' \\ g_{i'} & \text{otherwise} \end{cases}$$

$$\mathcal{A}_i^e : e_p \to \begin{cases} e_p[g(v_0 - v_i) \triangleright e^{-1}] & \text{if } i \text{ is on } p \text{ and aligned with } p \\ [g(\overline{v_0 - v_{i+1}}) \triangleright e] e_p & \text{if } i \text{ is on } p \text{ and aligned against } p \\ e_p & \text{otherwise.} \end{cases} \tag{2.7}$$

The paths involved in the cases where the edge is aligned or anti-aligned with the plaquette are indicated in Figure 2.25. Note that in either case the path terminates at the source of edge i, where the source of an edge is the vertex attached to that edge and which the edge points away from (while the target of the edge is the vertex attached to the edge and which the edge points towards). This means that we can replace v_i (in the aligned case) or v_{i+1} (in the anti-aligned case) in the expression for the paths in Equation 2.7 with this source, $s(i)$.

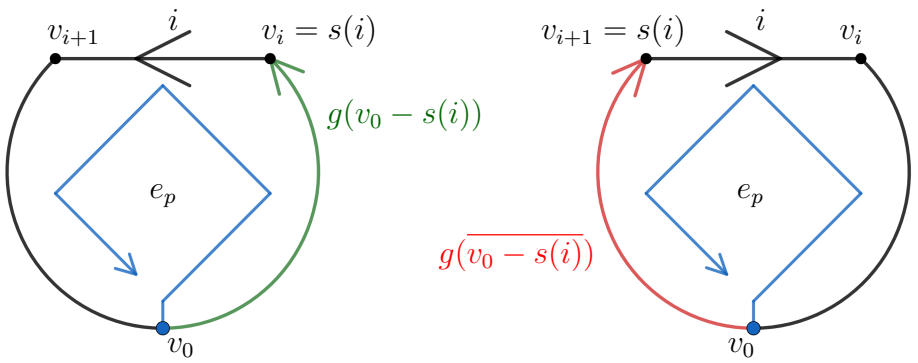

Figure 2.25: The path involved in the effect of the 2-gauge transform $\mathcal{A}_i^e$ on a plaquette p depends on whether the edge i is aligned with the p (as in the left case) or anti-aligned (as in the right case). If the edge is aligned with the plaquette, then the path $(v_0 - s(i))$ in the transformation of the plaquette label is aligned with p, whereas if i is anti-aligned with p then the path $(\overline{v_0 - s(i)})$ appearing in the transformation is anti-aligned with p. Either way, the path is aligned with the edge i.

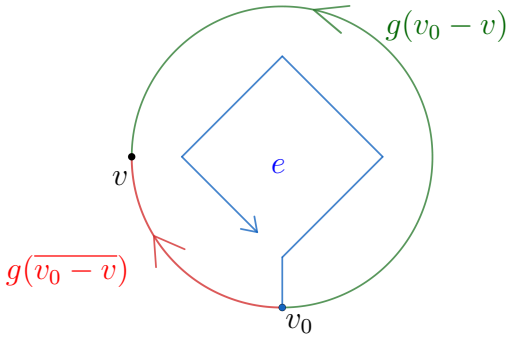

Figure 2.26: The two different paths from the base-point v_0 to the same vertex v on a plaquette are shown in green and red. Paths that anti-align with the surface circulation (represented by the blue arrow in the centre) are indicated using overline notation.

2.2.4 Gauge-Invariants

In ordinary lattice gauge theory we could build gauge-invariant quantities out of closed loops. What are the appropriate quantities for higher lattice gauge theory? We can build gauge-invariants from the closed loops as before, but also from closed surfaces. The closed loops work much as before, except that we need to modify the group element that labels them to account for parallel transport of paths over surfaces. Given a closed

loop made of two paths, as shown in Figure 2.27, to work out the group element for the loop, we need to transport the paths so that they are in the same location. This is necessary because the two paths may be defined with different gauge choices, with the conversion between the gauge choices performed by parallel transport. To obtain a gauge-invariant, we will need to ensure that the two paths are described in the same gauge. The relevant transport is shown in Figure 2.27.

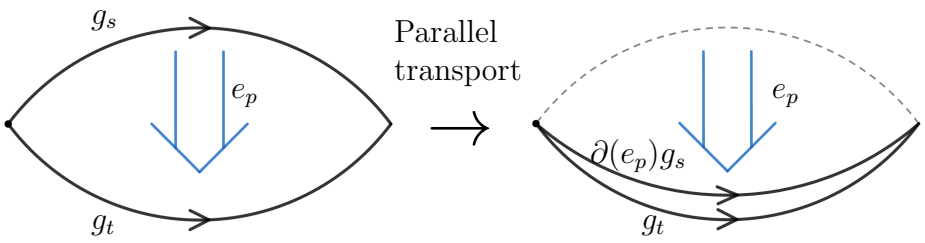

Figure 2.27: The gauge-invariant associated to plaquettes is modified to account for parallel transport of the paths over the plaquette.

The parallel transport modifies the group element associated to the closed loop in Figure 2.27, from $g_s g_t^{-1}$ to $\partial(e_p)g_s g_t^{-1}$. This quantity is the 1-flux or 1-holonomy for the closed loop. For a general surface, we replace $g_s g_t^{-1}$ with the label of the boundary of the surface. For a plaquette p, with boundary label g_p, the 1-flux is given by $\partial(e_p)g_p$ and we refer to this quantity as $H_1(p)$ [58]. This label can be changed only within a conjugacy class by either the vertex transforms (as in Figure 2.23) or the edge transforms (as in Figure 2.24), so those conjugacy classes are gauge-invariant quantities [58]. As an example, we can consider acting on the diagram in Figure 2.27 with a vertex transform. This gives us the situation shown in Figure 2.23. From that figure, we see that the plaquette holonomy, which is initially given by $\partial(e)g_1 g_2^{-1}$, transforms as

$$\begin{aligned}\partial(e)g_1 g_2^{-1} &\to \partial(g \triangleright e)g g_1 g_2^{-1} g^{-1} \\ &= g\partial(e)g^{-1} g g_1 g_2^{-1} g^{-1} \\ &= g\partial(e)g_1 g_2^{-1} g^{-1},\end{aligned}$$

where we used the Peiffer condition Equation 2.4 to write $\partial(g \triangleright e)$ as $g\partial(e)g^{-1}$. We see that the plaquette holonomy is only conjugated by the vertex transform. This indicates that the conjugacy class of this group element is a gauge-invariant quantity, in a similar way to the closed cycle in lattice gauge theory.

In addition to closed paths, we have closed surfaces, with their own gauge-invariants. In this case, the gauge-invariant assigned to a closed surface is simply the group label (2-gauge label) assigned to that closed surface, which may be obtained by using the rules for composing surfaces if that closed surface is comprised of multiple plaquettes. The 2-gauge label for a closed surface, which we call the 2-flux of that surface, is only changed within certain equivalence classes by the gauge transforms [58]. Again, the identity element is in a class on its own, so that trivial 2-flux is preserved by the transforms [58]. An important thing to note is that this is only strictly true when the system satisfies certain additional constraints [58], which we will discuss in Section 2.2.6.

In the same way that the 1-flux on a closed loop determines the result of a process where we move a charge around the loop, the 2-flux of a closed surface corresponds to a transport process. For a sphere at least, we can measure this 2-flux by nucleating a small loop at the base-point of that surface, before passing it over the surface and then contracting it again, as indicated in Figure 2.28. This reflects the fact that a spherical closed surface (which can be built from a series of open surfaces) can have empty source and target, and so can represent a transport process where we nucleate the loop at the start and collapse it at the end. For a surface such as a torus, with non-contractible cycles, the corresponding transport process may not involve nucleation and collapse.

2.2.5 Hamiltonian Model

Having considered higher lattice gauge theory, we are now in a position to consider the Hamiltonian model based on it (as described in Ref. [58]). The spatial dimensions

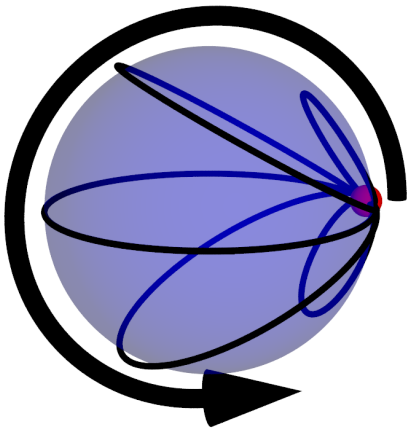

Figure 2.28: The 2-holonomy of a surface (in this case a sphere) can be measured by a transport process. A small loop is created at the base-point (the small red sphere), then dragged over the surface (the larger blue sphere), as indicated by the arrow.

---

of the model (we will consider the model in three spatial dimensions in this work) are represented by a lattice, while the temporal dimension is continuous and time evolution is controlled by the Hamiltonian. As already alluded to, we label each edge of the lattice with an element of group G and each plaquette with an element of group E [58]. Labelling every edge and plaquette gives a configuration (or colouration). These configurations then form a basis for the Hilbert space, so that a general state is a linear combination of the different labellings of the lattice. However, we have seen that a given plaquette can correspond to different transport processes depending on the source and target, so we need a way of specifying which transport process the assigned label corresponds to. In order to do this, we define a "canonical" position for the source and target paths of every plaquette when we set up the lattice. This is done via a so-called branching structure. Every vertex in our lattice is assigned a unique integer index [58]. Edges point from the lower indexed vertex to the higher indexed one. Then for a given plaquette, the lowest indexed vertex on that surface is the base-point for that surface (recall that the base-point is the start of the source and target paths for the surface) [58], meaning that the vertex with index 0 is the base-point in the example shown in Figure 2.29. This base-point is adjacent to two other vertices on that surface. The target is the edge from the base-point to the higher indexed of these

two neighbours. In the example in Figure 2.29, the base-point is adjacent to vertices with indices 2 and 3, so the target is the edge 0-3 from the base-point to the higher indexed neighbour. On the other hand, the source is the path from the base-point to the same vertex, but the other way around the plaquette. It first passes from the base-point to the lower indexed neighbour, then around the plaquette to the higher indexed neighbour. This means that in Figure 2.29, the source is the path 0-2-1-3.

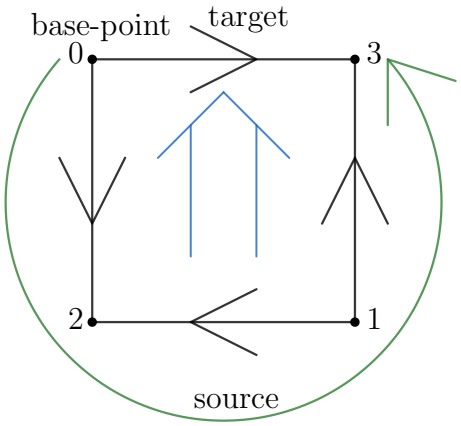

Figure 2.29: For each plaquette (the square in the figure being an example), we define a reference source and target in order to define the 2-holonomy of the plaquette. This is done by indexing the vertices in the lattice. In this example, the base-point of the plaquette is indexed by zero (which is the lowest index on this plaquette). The target of the plaquette is then 0-3, the direct path from the base-point to the higher indexed of its two neighbours, while the source is 0-2-1-3, the path from the base-point to this higher indexed neighbour via the lower indexed one.

The canonical sources and targets, which are fixed by the branching structure, form a reference that we can use to unambiguously define a state (because we need to specify the sources and targets whenever we define a 2-holonomy). The branching structure defines the precise parallel transport process associated to the plaquette by default, and the group label given to the plaquette corresponds to this process. As described in Section 2.2.1, we can then use a set of rules to manipulate the source and target of a plaquette and find the appropriate label for the plaquette afterwards.

To motivate the Hamiltonian considered by Bullivant et al. [58], we can take the same approach as when we considered a model based on lattice gauge theory. We

first demote the gauge symmetries to energetic constraints, by including energy terms which are averages over the associated gauge transforms:

$$H = -\sum_{\text{vertices, } v} A_v - \sum_{\text{edges, } i} \mathcal{A}_i + ...$$

Here A_v is the average over gauge transforms at the vertex v:

$$A_v = \frac{1}{|G|} \sum_{g \in G} A_v^g. \tag{2.8}$$

We saw how the gauge transform A_v^g acts on the degrees of freedom in the lattice in Figure 2.23, and again in Equation 2.6. As with Kitaev's Quantum Double model, the vertex transforms satisfy $A_v^g A_v^h = A_v^{gh}$ for any g, $h \in G$. The action of the transforms on the edges satisfies this relationship because the higher lattice gauge theory vertex transforms act on the edges in the same way as for the lattice gauge theory case. The relationship also holds for the action on the plaquettes, because if v is the base-point of a plaquette p then we have

$$A_v^g A_v^h : e_p = g \triangleright (h \triangleright e_p) = (gh) \triangleright e_p,$$

where in the last step we used Equation 2.3 from Section 2.2, while if v is not the base-point then the plaquette is unaffected. Either way $A_v^g A_v^h : e_p = A_v^{gh} : e_p$. The fact that the vertex transforms satisfy the algebra $A_v^g A_v^h = A_v^{gh}$ means that A_v is a projector [58], just as for the equivalent term in Kitaev's Quantum Double model (see Equation 2.1), where the eigenvalue of one corresponds to states which are gauge-symmetric at that vertex. The algebra also means that we can absorb vertex transforms into the corresponding vertex energy term, by which we mean that $A_v^x A_v = A_v$ for any $x \in G$, as we can demonstrate by expanding the vertex term:

$$\begin{aligned} A_v^x A_v &= A_v^x \frac{1}{|G|} \sum_{g \in G} A_v^g = \frac{1}{|G|} \sum_{g \in G} A_v^{xg} \\ &= \frac{1}{|G|} \sum_{g' = xg \in G} A_v^{g'} = A_v. \end{aligned}$$

This means that a state $|\psi\rangle$ (such as a ground state) which satisfies $A_v |\psi\rangle = |\psi\rangle$ is invariant under the individual vertex transforms, rather than just the energy term:

$$A_v^x |\psi\rangle = A_v^x A_v |\psi\rangle = A_v |\psi\rangle = |\psi\rangle . \tag{2.9}$$

Therefore, such states are gauge-invariant, as we claimed earlier.

In a similar way to the vertex terms, the edge term $\mathcal{A}_i$ is the average over 2-gauge transforms at the edge i [58]:

$$\mathcal{A}_i = \frac{1}{|E|} \sum_{e \in E} \mathcal{A}_i^e. \tag{2.10}$$

These edge transforms can be combined in the same way as the vertex transforms [58]: $\mathcal{A}_i^e \mathcal{A}_i^f = \mathcal{A}_i^{ef}$. As with the vertex terms, this means that an individual transform can be absorbed into the energy term:

$$\mathcal{A}_i^e \mathcal{A}_i = \mathcal{A}_i. \tag{2.11}$$

This leads to the energy term $\mathcal{A}_i$ being a projector [58], with the eigenvalue of one corresponding to states that are 2-gauge-symmetric at that edge, as we will see shortly. The minus sign with which this term enters the Hamiltonian ensures that the energy term favours these gauge-symmetric states. To see that these lower energy states (such as the ground states) are invariant not just under the edge energy terms, but also under the individual 2-gauge transforms $\mathcal{A}_i^e$, note that the lower energy states satisfy $\mathcal{A}_i |\psi\rangle = |\psi\rangle$ by definition and so

$$\mathcal{A}_i^e |\psi\rangle = \mathcal{A}_i^e \mathcal{A}_i |\psi\rangle = \mathcal{A}_i |\psi\rangle = |\psi\rangle .$$

So far we have considered energy terms that enforce the 1-gauge symmetry and 2-gauge symmetry. Now we add terms that depend on quantities that are invariant under the two types of gauge transform. By building these terms from gauge-invariant quantities, we guarantee that the new terms commute with the gauge transforms. Recall from

Section 2.2.4 that there are gauge-invariant quantities associated to the closed cycles of the lattice. In particular, whether a cycle has a trivial group element or not is invariant under gauge transforms (once we account for the effect of parallel transport over the surface bounded by the cycle). We can therefore energetically punish cycles that have non-trivial 1-flux. We do this with an energy term at each plaquette, which gives one if the plaquette has trivial flux and zero if the flux is non-trivial. As explained in Section 2.2.4, the 1-flux for a plaquette with label e_p and path label g_p for its boundary is given by $\partial(e_p)g_p$. The plaquette term therefore acts as $\delta(\partial(e_p)g_p, 1_G)$. An example of the plaquette energy term is shown in Figure 2.30.

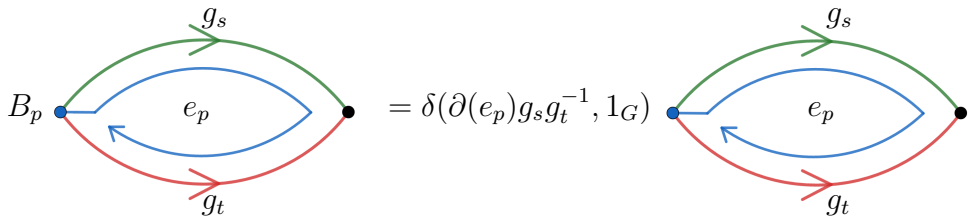

Figure 2.30: The plaquette term checks that the flux through a plaquette is trivial. In this case, the boundary of the plaquette is made of two edges, labelled by g_s and g_t, and the boundary label is $g_s g_t^{-1}$. The flux through the plaquette is therefore $\partial(e_p)g_s g_t^{-1}$, where e_p is the plaquette label, and the plaquette term checks whether this expression for the flux is trivial.

The plaquette terms enter the Hamiltonian with a minus sign, so that

$$H = -\sum_{\text{vertices, } v} A_v - \sum_{\text{edges, } i} \mathcal{A}_i - \sum_{\text{plaquettes, } p} B_p + ...$$

Including these terms with a minus sign ensures that the lowest energy states have trivial flux on the plaquettes. This trivial flux condition is called fake-flatness [58], due to its analogy with the flatness condition of ordinary lattice gauge theory (which enforces that all cycles are trivial). Plaquettes that satisfy the fake-flatness condition are then called fake-flat.

Finally, we consider the gauge-invariant quantity associated to the closed surfaces. In the same way as for closed cycles, we punish closed surfaces with non-trivial 2-

flux (2-holonomy). This is done with an energy term at each "blob" (3-cell) [58]. The blobs are the smallest three-dimensional volumes, such as the smallest cubes in a cubic lattice. For each blob, we have an energy term that checks the value of the surface of that blob, leaving it unchanged if that value is 1_E and giving zero otherwise [58], as shown in Figure 2.31. We denote the blob term associated to a blob b by $\mathcal{B}_b$. In this work we will consider the case where the lattice is in three spatial dimensions (so the overall model is 3+1d), although in Ref. [3] we also looked at the 2+1d case, for which the lattice has no blobs and there are no blob energy terms.

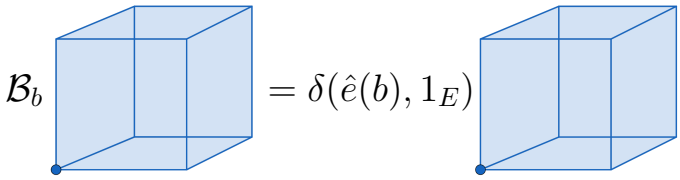

Figure 2.31: The blob energy term $\mathcal{B}_b$ checks whether the total surface label of the blob b, $\hat{e}(b)$, is the identity element or not. This surface label must be determined by using the rules for combining surface elements from Sections 2.2 and 2.2.5 to combine the plaquettes on the boundary of the blob. For example, when $\triangleright$ is trivial the surface label is a product of the plaquette labels (with inverses if the orientation of the plaquette needs to be reversed to match the overall surface).

The blob term also enters the Hamiltonian with a minus sign, so that the full Hamiltonian is given by [58]

$$H = -\sum_{\text{vertices, } v} A_v - \sum_{\text{edges, } i} \mathcal{A}_i - \sum_{\text{plaquettes, } p} B_p - \sum_{\text{blobs, } b} \mathcal{B}_b. \tag{2.12}$$

Building the Hamiltonian out of gauge transforms and gauge-invariant quantities would mean that the different energy terms commute (given that the two types of gauge transform and two types of gauge-invariant commute). However, when fake-flatness is not satisfied (that is the plaquette terms are not minimised), the blob terms may not actually be 2-gauge-invariant [58] (in fact, the rules for combining surfaces become inconsistent and so the blob terms are ill-defined without some convention for how combination should be done). Therefore, the blob terms may not commute with the

edge terms in the Hamiltonian. This in turn means that the model is not a commuting projector model (and so not necessarily solvable), although this can be fixed by defining the blob terms to be zero when any of the plaquettes are not satisfied [58]. This problem will occur in models where $\triangleright$ is non-trivial, i.e. models for which $g \triangleright e \neq e$ in general. When $\triangleright$ is trivial this complication does not occur and we have a commuting projector Hamiltonian (using the results from Ref. [58]). There are some further complications when $\triangleright$ is not trivial in the general case, and so we make some restrictions to the model in order to make it more manageable, as we discuss in Section 2.2.6.

We note that when we take E to be the trivial group (so that we consider lattice gauge theory), the terms $\mathcal{A}_i$ and $\mathcal{B}_b$ become trivial. In addition, the vertex transforms and plaquette terms become equivalent to the corresponding terms from Kitaev's Quantum Double model. Therefore, we recover that model in this case, indicating that the higher lattice gauge theory model is a generalization of the Quantum Double model [58].

2.2.6 Some Special Cases and Consistency

We now wish to consider the conceptual issues that we mentioned in the previous section. For the most general crossed modules, the Hamiltonian model has certain inconsistencies. As an example of what we mean, consider the surface holonomy of a plaquette, e_p. We can move the base-point of the plaquette around the plaquette, so consider moving it all the way around the plaquette. This induces a change to the surface label of the plaquette, given by $e_p \to g_p^{-1} \triangleright e_p$ [58], where g_p is the path label of the boundary of the plaquette, as shown in Figure 2.32. The base-point is back to the same position, and the surface appears to be the same, yet the label may have changed. The label does stay constant if the configuration is fake-flat, meaning that it satisfies the plaquette energy term defined in Section 2.2.5). In that case, the boundary label satisfies $g_p^{-1} = \partial(e_p)$ and so $g_p^{-1} \triangleright e_p = \partial(e_p) \triangleright e_p = e_p e_p e_p^{-1} = e_p$ [58], where we used the Peiffer condition Equation 2.5 in the second step. However, if fake-flatness is not satisfied, then we cannot guarantee that the plaquette label is unchanged.

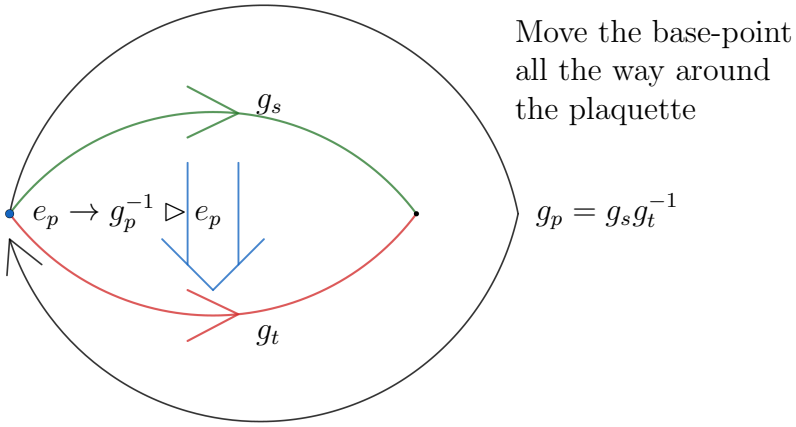

Figure 2.32: Moving the base-point of a plaquette all the way around the plaquette transforms the surface label from e_p to $g_p^{-1} \triangleright e_p$, where g_p is the path label of the plaquette's boundary.

This is not the only issue arising from violating fake-flatness: as we describe in the Appendix of Ref. [1], the edge energy term also appears to become inconsistent with changes to the branching structure of the lattice. One approach for dealing with this problem is to enforce fake-flatness on the level of the Hilbert space, as a hard constraint rather than an energy term. This is the case most closely considered by the authors of the original paper on this model [58]. However, another possibility is to take $\triangleright$ to be trivial, so that the base-point of the plaquette loses any meaning, but allow fake-flatness violations. In this case the model loses some of its complexity, due to the 1-gauge field having no way to act on the 2-gauge field. If we use this condition, then all of the energy terms commute naturally, with no need to restrict the Hilbert space. Some additional consequences of taking $\triangleright$ trivial are that E must be Abelian and that ∂ maps to the centre of G. The first condition, that E is Abelian, comes from the second Peiffer condition (Equation 2.5 in Section 2.2), because $\partial(e) \triangleright f = efe^{-1} \implies f = efe^{-1}$ so that any elements of E commute with each-other. The second condition, that ∂ maps to the centre of G, comes from the first Peiffer condition (Equation 2.4 in Section 2.2), as $\partial(g \triangleright e) = g\partial(e)g^{-1} \implies \partial(e) = g\partial(e)g^{-1}$ so that $\partial(e)$ commutes with all elements of G. In this work we consider both of these special cases ($\triangleright$ trivial and restricting

to fake-flat configurations). We also examine another case, where E is Abelian and ∂ maps to the centre of G, but we do not enforce fake-flatness on the level of the Hilbert space or require $\triangleright$ to be trivial. This is the most interesting case that we consider, but also the most difficult to deal with due to the non-commutativity of energy terms. This case is still more tractable than the fully general case (i.e. the case where we do not make any restrictions on either the crossed module or the Hilbert space), not just because the groups are simpler algebraically, but also because the inconsistency under changes to the base-point that we considered above is not as generic. If a plaquette p with label e_p violates fake-flatness because the boundary label g_p of the plaquette differs from $\partial(e_p)^{-1}$ only by an element $\partial(e) \in \partial(E)$, then moving the base-point of the plaquette around p results in the plaquette label e_p transforming to

$$g_p^{-1} \triangleright e_p = (\partial(e)^{-1}\partial(e_p)) \triangleright e_p = (e^{-1}e_p)e_p(e^{-1}e_p)^{-1},$$

which is just e_p because E is Abelian. This case, where the boundary label differs from $\partial(e_p^{-1})$ by an element in $\partial(E)$, is significant because it occurs when the fake-flatness violation is caused by a change to the plaquette label e_p, rather than changes to the edge labels. Such flatness-violating changes to the plaquette labels occur for a whole class of ribbon operators (the confined blob ribbon operators we describe in Chapter 5). This means that the inconsistencies considered previously are less prevalent in the case where E is Abelian. We will see other simplifications that occur due to E being Abelian and ∂ mapping to the centre of G throughout this work. Because we will refer to these restrictions, along with the other cases that we have considered in this section, many times in the following text, we summarize all of them in Table 2.1.

Case	E	$\triangleright$	$\partial(E)$	Full Hilbert Space
1	Abelian	Trivial	$\subset$ centre(G)	Yes
2	Abelian	General	$\subset$ centre(G)	Yes
3	General	General	General	No

Table 2.1: A summary of the special cases of the model

Chapter 3

Properties From Gauge Theory

Before we discuss the excitations that we find in the model in great mathematical detail, it will be instructive to give a more qualitative description of the excitations that we expect to find, using ideas from higher gauge theory. As a starting point, we shall give a brief description of the picture in ordinary gauge theory. In ordinary gauge theory, we find objects which we call "electric charges" and "magnetic fluxes", which we will describe shortly. A clear exposition on these objects and their properties, in the 2+1d case, is given by Preskill's lecture notes on topological quantum computation [75, Chapter 9] and an early description of non-Abelian magnetic fluxes is given in Ref. [76] (see also Refs. [38] and [39]). Here we will instead examine the 3+1d case, as described in (for example) Ref. [77].

Electric charges are point particles labelled by irreducible representations of the group G. To see how these charges arise from our model, we look at our vertex gauge transforms, as defined in Figure 2.2. The gauge transforms A_v^g at a particular vertex form a group, with the product $A_v^g A_v^h = A_v^{gh}$, which is isomorphic to G. Our Hilbert space then splits into subspaces that transform as irreducible representations (irreps) of G under the action of the gauge transforms at every vertex. These irreps are related to the energy of the states. The states that are unexcited (lower energy) at each vertex are those states for which $A_v = 1$ at each vertex v (a vertex is excited in a state if

$A_v = 0$ for that state). Such states are gauge-invariant, meaning that they are invariant under the individual vertex transforms, as discussed in Section 2.2.5. These gauge-invariant states are therefore exactly those described by the trivial irrep of G. From this we see that the states described by the trivial irrep at each vertex have unexcited vertex terms. On the other hand, if a state transforms as some other, non-trivial, irrep at a particular vertex, then that vertex will be excited. This means that we expect to find excitations that carry some non-trivial irrep of G, with this irrep describing how the excitations transform under the gauge transforms. We denote an excitation which transforms under the irrep R by (R, a, b), where a and b are the matrix indices of the representation and describe an internal space for the class of excitations labelled by R. We define such an excitation by the action of the vertex operator on it:

$$A_v^g \cdot (R, a, b) = \sum_c [D^R(g^{-1})]_{ac}(R, c, b), \tag{3.1}$$

where $D^R(x)$ is the matrix representation of element $x \in G$ in irrep R. The label g^{-1}, rather than g, is used to ensure that the action of the A_v^g satisfies the composition rule $A_v^g A_v^h = A_v^{gh}$. We could equally have defined the action of A_v^g to be right multiplication by $D^R(g)$ instead, which would also satisfy the composition law. This transformation under the vertex transforms can also be used to tell us something about the transport properties of the excitation. In Section 2.2.3 we explained that the vertex transforms are equivalent to parallel transport. Therefore, Equation 3.1 tells us how we expect these excitations to behave under parallel transport over an edge labelled by g^{-1}. Looking at Equation 3.1, we see that there is mixing between states defined by different matrix indices, while the irrep is unchanged. This suggests that the electric charges carry some conserved charge labelled by the representation, while the matrix indices describe some non-conserved details.

In addition to the electric charges, we anticipate magnetic fluxes. Recall from Section 2.1.2 that fluxes are associated with closed paths that have non-trivial labels. In that section, we drew an analogy to how a magnetic field leads to a non-trivial Aharanov-

Bohm effect for taking a charge around a closed loop. In a 2+1d model, flux can be created by a point particle, which we can think of as being similar to a magnetic field penetrating our surface at a point. If this point particle generates a flux, then this flux should be measured by a closed loop that encloses the particle. Therefore, we can describe this particle with the closed loop that measures the flux. However, a point particle cannot be sensibly described by a closed loop in a 3+1d topological theory. This is because any closed loop around a point particle can be smoothly deformed away from that particle and contracted to nothing, i.e. to a path with a trivial label. Therefore, the label of the path cannot be a topological quantity if the excitation generating the flux is a point particle. Instead the magnetic flux particles should be closed loops (flux tubes). The flux generated by such a tube can be measured by a closed path that links with the flux tube, as shown in Figure 3.1. Then there is no way to smoothly contract the measurement path without it intersecting with the flux tube, and so the label of the path can be a topological quantity.

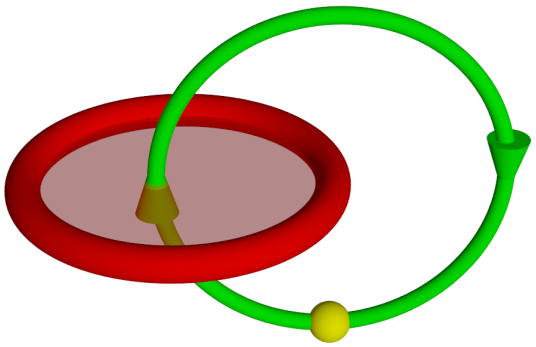

Figure 3.1: The magnetic excitation (thick red torus) is measured by a non-trivial closed loop linking with it, such as the thin green torus. This closed loop begins at some start-point (shown as a yellow sphere) and the value of the flux that we measure depends on which start-point we choose, but the conjugacy class of the flux does not.

We can label a flux tube by the group element of the closed path that measures the flux. This label describes how a charge would evolve as it travels along that closed path. However, the value we assign to a path depends on the start-point of that path. To see this, consider taking a particular closed path that links with that excitation and then changing its start-point, as shown in Figure 3.2. In order to traverse the

new path, a charge must first travel along the path from the new start-point to the old one, then along the original closed path, and then back along the path joining the start-points. This means that the label of the new path is given by conjugating the group element of the old path by the element associated to the path between the start-points. We therefore see that this new path has a label that is different from the label of the original path, but which lies in the same conjugacy class. This means that the conjugacy class describes the excitation in a robust way, but to obtain a full description of the flux tube we must also specify its element within that conjugacy class and the start-point from which we measure that value.

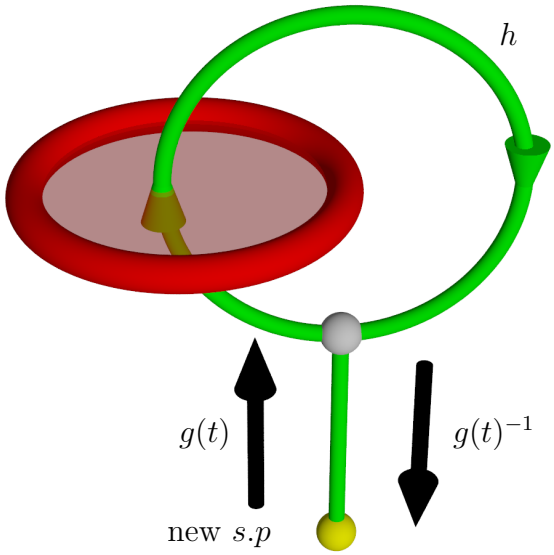

Figure 3.2: With the start-point ($s.p$) at the grey (higher) sphere, the flux label is h, however with the $s.p$ at the yellow (lower) sphere the flux is $g(t)hg(t)^{-1}$

In addition to giving us the excitations that we would expect for a lattice gauge theory model, this picture will give us the braiding relations. We already established how an electric excitation should transform as it moves through space. We can now consider taking an electric excitation, labelled by an irrep R, on a closed path that links with a flux tube labelled by h, as shown in Figure 3.3. If this closed path is the one used to define the flux, then the path label is given by the flux label of the flux tube, h. Therefore, the excitation (R, a, b) should become $\sum_c [D^R(h)]_{ac}(R, c, b)$ after the motion. On the other hand, if the path has a different start-point to the defining path

of the flux, we should replace h with some other element in the conjugacy class of h, in order to describe the transport of the charge to and from the start-point as well as around the defining loop of the flux. For instance, if we were to first take the charge on some path t to get from a new start-point to that of the flux, then take the charge around the defining path of the flux before finally returning it along t^{-1} to its initial position, then we would have

$$\begin{aligned}(R,a,b) \to & \sum_c \sum_d \sum_e [D^R(g(t))]_{ac}[D^R(h)]_{cd}[D^R(g(t)^{-1})]_{de}(R,e,b) \\ & = \sum_e [D^R(g(t)hg(t)^{-1})]_{ae}(R,e,b),\end{aligned}$$

from which we see that the charge experiences the same transform as if it travelled around a flux of $g(t)hg(t)^{-1}$, as expected from our rules for changing the start-points of fluxes.

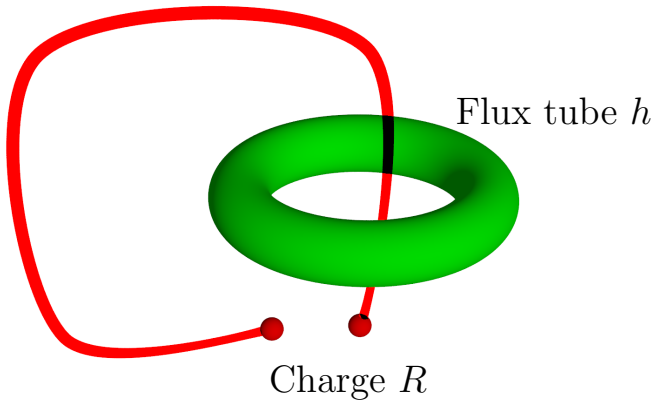

Figure 3.3: Schematic view of braiding a charge through a loop. The red line tracks the motion of the charge.

We can also obtain the braiding relations of the magnetic fluxes using this picture. Consider the case where we have two flux tubes, which we define with the same start-point. We want to keep track of the measurement paths and the flux labels, as we move the fluxes around. We consider exchanging the two flux loops by pushing one through the other, as shown in Figure 3.4. When we move a flux loop, the measurement path (which is associated with the flux label) moves with it so that the measurement path

and flux tube remain linked (we can imagine the flux tube dragging the measurement path with it). For example, in the top-right part of Figure 3.4, which shows the situation after we perform the braiding move, we see that the measurement path for the blue flux tube (which is originally lower than the red flux tube) is pulled through the (now lower) red flux tube. This new deformed path carries the original flux label (h in Figure 3.4) and so this flux label is now associated to a process where we pull a charge through the lower loop, then around the upper one and then back through the lower loop, rather than a process where we simply braid the charge around the upper loop. We want to define our fluxes with respect to our original measurement paths, in order to find the labels associated to the original measurement processes and so to find the change to the system under braiding. That is, we want to find the labels associated to the original measurement paths (α and β in Figure 3.4). To do this, we need to write the original paths in terms of the new deformed ones, for which we know the path labels. This will allow us to obtain the labels of the original paths and so tell us the result of braiding our fluxes.

Looking at Figure 3.4, we see that β is the path originally associated to the upper (red) flux, with label g. When we deform space to push the lower (blue) flux tube through the upper (red) flux tube, this path β is deformed to β'. This means that the label of this new path, β', is equal to the original label of path β, i.e. g. However, this path is equivalent to (i.e. can be smoothly deformed into) the original path α around the old lower flux. So we have $\beta' = \alpha$ and because β' is labelled by g, this is the new label of the path α (which now links with the red flux tube rather than the blue one) after braiding. On the other hand, consider the path α originally associated with the blue (initially lower) flux. When we move the fluxes, this path is deformed into α', as indicated in the upper right diagram, meaning that α' is labelled by h after the braiding. We want to write α' in terms of our old paths α and β. To do this we note that α' can be smoothly deformed into another path α'', which is equal to the path $\alpha\beta\alpha^{-1}$ obtained by traversing α then β and then α in reverse, as shown in the

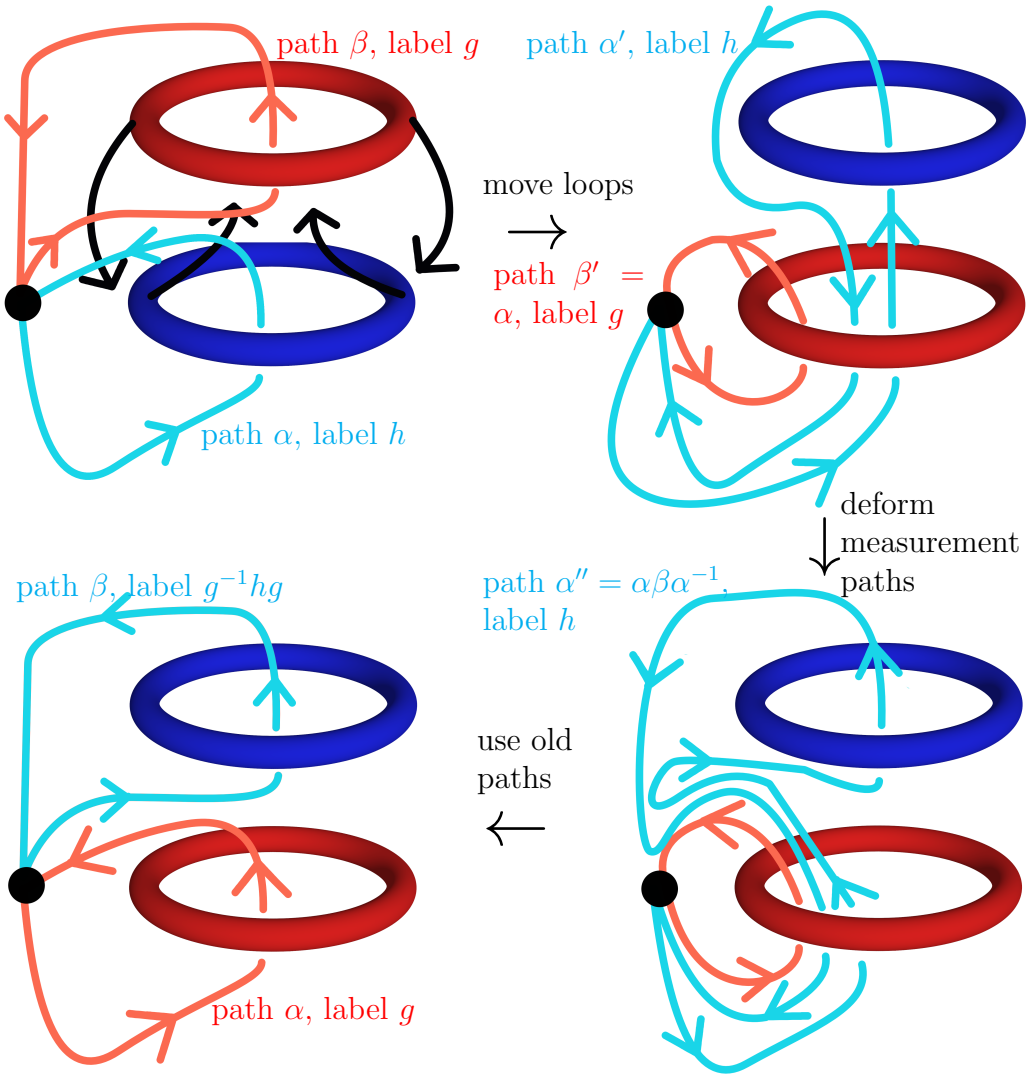

Figure 3.4: Starting with two fluxes (top-left) we can move the lower one through the upper one, swapping their positions. When we do so, we must also move the measurement paths associated with each flux (top-right) to stay linked with the flux. We can deform these paths (bottom-right) in order to express them as products of the original paths. Knowing the labels for these new paths (which are just the original flux labels of the two loops) allows us to find the labels of the original measurement paths (bottom-left), giving us the result of braiding the two flux tubes.

---

bottom-right figure. Therefore, we have that $\alpha' = \alpha\beta\alpha^{-1}$ and so $\beta = \alpha^{-1}\alpha'\alpha$. Using the fact that α now has the label g and α' has the label h, we see that β has the label $g^{-1}hg$ after braiding. We can write this braiding relation in the following way. We start with (h, g), where the first symbol in brackets is the label given to path α and the second is the one given to path β. Then under braiding we have $(h, g) \to (g, g^{-1}hg)$, where on both sides the first symbol refers to the value of path α and the second to the

value of path β, rather than giving the label of a particular one of the excitations (blue or red). If we instead keep track of the labels of each tube, we see that $h \to g^{-1}hg$ for the blue tube and $g \to g$ for the red tube. Therefore, we see that the label of one of our flux tubes is conjugated by the label of the other one under our braiding.

We have so far described the results we expect from ordinary lattice gauge theory, but not higher lattice gauge theory. However, we can use very similar ideas to explore higher lattice gauge theory. Our vertex terms still have the same algebra as in ordinary gauge theory. Namely, we have a group isomorphic to G at each vertex. Therefore, we expect electric charges labelled by irreps of G, just as with lattice gauge theory. Again, under parallel transport the charges will transform according to this irrep. This suggests that our electric excitations will be largely unchanged when compared to those from ordinary lattice gauge theory. Similarly, we expect magnetic flux tubes, just as in lattice gauge theory. However, there is some subtlety in considering the braiding between these two types of excitation. This is because the lattice does not satisfy flatness in the ground state, but instead fake-flatness. This means that deforming a path over an unexcited region causes the path label to pick up a factor of the form $\partial(e)$, rather than remain unchanged. As we discussed earlier, moving an electric excitation through space, such as when we braid it around a magnetic flux tube, causes it to transform according to the label of the path traversed. This suggests that the result of braiding an electric charge around such a tube depends on the precise path chosen for the braiding, not just its homotopy class, implying that the braiding relation is not topological. As we shall discuss further in Section 5.3, the resolution to this is that any electric excitation that is sensitive to such factors of $\partial(e)$ must be confined (i.e., cost energy to separate from its antiparticle), and so is not topological. In addition, the fact that we have fake-flatness rather than flatness indicates that a closed path may have a non-trivial label of the form $\partial(e)$ even in the ground state (and we would have to measure the surface enclosed by that path to check that fake-flatness is satisfied). This implies that fluxes with label in $\partial(E)$ cannot be distinguished from trivial fluxes

just by measuring the closed path. Furthermore, deforming the measurement path for a magnetic flux tube will change the label measured by an element of $\partial(E)$. Therefore, when we talk about the flux of a magnetic excitation, we should only define it up to elements in $\partial(E)$. This leads to magnetic excitations with label in $\partial(E)$ becoming topologically trivial, in a sense that we explain in Section 5.2.

In addition to the vertex gauge transforms, we also have the 2-gauge (or edge) transforms, which again have a group structure. These operators form a group isomorphic to E for each edge: $\mathcal{A}_i^e \cdot \mathcal{A}_i^f = \mathcal{A}_i^{ef}$. Therefore, we expect to find edge excitations that are labelled by irreps of the group E. Recall that the 2-gauge transform is equivalent to parallel transport of an extended object (a line object) over a surface. Therefore, an object which transforms as a particular irrep under the 2-gauge transform should also transform as that irrep under parallel transport over a surface. Because of the fact that this transport is over a surface rather than a path, we expect our "2-charges" for the 2-gauge field to be extended objects. In fact, we will find that these 2-charges are loop-like objects. Then when we transport a loop, labelled by an irrep μ and matrix indices a and b, over a surface labelled by e, we should obtain the transformation $(\mu, a, b) \to \sum_c [D^\mu(e^{-1})]_{ac}(\mu, c, b)$. However, there is some subtlety to this. Recall from Section 2.2 that every surface element has a base-point, which we can change by whiskering. If we change the base-point, we change the label of that surface from e to $g \triangleright e$ for some $g \in G$. Then how do we know if we should have $[D^\mu(e)]$ or $[D^\mu(g \triangleright e)]$ in our transformation when we transport the loop? That is, where should we take the base-point of our surfaces? The answer is that, just as our flux excitations are defined with respect to some start-point, so must our 2-charges. When we define our 2-charges, as well as giving the irrep and its matrix indices, we give a special point on the lattice, which we call the start-point. When we move the loop over a surface, we always take the label of that surface with respect to the start-point of our loop excitation. This start-point is important for these loop excitations, because the action of the group G on the start-point via a vertex transform, which changes the surface

label by some $g\triangleright$ map, enables G to affect the loop excitations. This action can even change the irrep labelling a 2-charge loop, suggesting that the irrep is not a conserved quantity. Instead there is some mixing within certain classes of irreps, which we term $\triangleright$-Rep classes of irreps of E, with the irreps in a particular class being related by the action of $\triangleright$. When E is Abelian we can define the classes with the equivalence relation

$$\mu_1 \sim \mu_2 \iff \mu_1(e) = \mu_2(g \triangleright e) \text{ for some } g \in G, \tag{3.2}$$

although when E is non-Abelian we must generalize this to account for another equivalence between irreps, namely the fact that irreps related by conjugation by a constant matrix should not be regarded as independent. We therefore see that $\triangleright$ plays a significant role in determining how the excitations behave.

Just as we have magnetic fluxes that are associated with non-trivial loops, we established in Section 2.2 that we have "2-fluxes" associated to non-trivial closed surfaces. We expect excitations corresponding to the 2-flux of a sphere (as the blob terms in our Hamiltonian correspond to spheres on the surface of each blob) to be point excitations, because we can shrink a sphere to enclose just a single point. Just as the fluxes are labelled by elements of G, the 2-fluxes (or blob excitations) should be labelled by elements of E. In the same way that we must define a closed path to specify a flux, we must specify a closed surface to specify our 2-flux. Recalling that every surface must be defined with a base-point, we see that we must choose a base-point for our 2-flux excitation, which we call the start-point of the excitation. Note that we have now used the term start-point in relation to several excitations, and we deliberately use the same term here. From here on we use the term *start-point* to refer to privileged vertices related to the excitations, while *base-point* is used to refer to the base-point of a surface. In the case of the 2-flux excitations (as well as the 2-charges), the start-point of the excitation is also the base-point of a measurement surface (whereas for the 1-flux tubes, the start-point is the start of a measurement path). Moving the base-point of a surface along a path labelled by g changes the surface label from e to

$g^{-1} \triangleright e$, so similarly changing the start-point of our 2-flux changes its label by this $g^{-1} \triangleright$ action. We also expect the 2-flux excitation to transform in this way as the excitation moves along a path (for example, as it braids through a 1-flux tube). This tells us that the 2-flux label is not conserved, but rather each group element belongs in a class of elements related by the $\triangleright$ action. The equivalence relation defining such a class is that two elements $e, f \in E$ satisfy

$$e \sim f \iff \exists\, g \in G \text{ such that } e = g \triangleright f. \tag{3.3}$$

These "$\triangleright$-classes" are then the quantities that are conserved under motion.

As this picture gives us the transport properties of these excitations, we can obtain the braiding relations as well. We know how our 2-gauge charges transform under transport over a surface. Therefore, we know how they will transform when we pull them over one of our 2-fluxes, as a 2-flux corresponds to a particular closed surface with a particular label. For a loop excitation (μ, a, b) and a 2-flux e, defined with the same start-point, the loop excitation becomes $\sum_c [D^\mu(e^{-1})]_{ac}(\mu, c, b)$ when it is pulled over the 2-flux. We can also work out how the 2-fluxes braid with ordinary 1-fluxes. We know how the 2-fluxes transform when they move along edges. The 2-flux label changes from e to $g(t) \triangleright e$ from moving along a path t. Therefore, when moving a 2-flux labelled by e around an ordinary magnetic flux labelled by h, the 2-flux becomes $h \triangleright e$ (or $h^{-1} \triangleright e$, depending on the orientation of the magnetic flux).

This picture therefore tidily describes several types of simple excitation that we expect to find. However, there are more complicated excitations as well. We may expect loop particles that generate both a 2-flux and a 1-flux. The non-trivial magnetic 1-flux is associated to a closed path that links with the excitation, while the 2-flux corresponds to a spherical surface enclosing that excitation. In Ref. [65], a possible braiding relation between two such loops is described, motivated using geometric arguments. The authors look at the situation where they braid two such excitations labelled by

(g, e) and (h, f), where the first label of each pair gives the magnetic flux and the second the 2-flux. When the excitation labelled by (g, e) is pushed through the one labelled by (h, f), the excitations should transform under braiding to become $(h^{-1}gh, h^{-1} \triangleright e)$ and $(h, ef[h^{-1} \triangleright e^{-1}])$. As will be explained in Section 6.4, we do indeed find such excitations with these braiding statistics.

Having discussed the excitations we will find in terms of electric charges, magnetic fluxes and their 2-gauge equivalents, we can now look in more detail at how these arise in our Hamiltonian model. We will describe the operators that produce and move these excitations and will also see features that are not apparent from this brief discussion, including which energy terms the fluxes and charges excite and which of the excitations are confined (cost energy to separate from their antiparticle).

Chapter 4

Ribbon and Membrane Operators

4.1 Introduction

One of the central objects in commuting projector models for topological phases is the so-called ribbon operator [16, 36]. Ribbon operators act on a linearly extended region (often with some finite width), called a ribbon, and produce excitations at the two ends of the ribbon. One of the defining properties of the ribbon operators that produce the topological excitations is that they commute with the Hamiltonian everywhere except at the start and end of the ribbon and so act to produce a pair of anyons [16, 36]. Because the bulk of the ribbon does not produce any excitations, the ribbon itself is largely invisible apart from its end-points. Indeed, ribbon operators are *topological*, in the sense that they can be smoothly deformed through the ground state (or any unexcited region of the lattice) without affecting the action of the operator. Important exceptions to this are the ribbon operators which produce confined excitations. These excitations cost energy to separate from their anti-particle, and so the corresponding ribbon operators have an energy cost associated to the length of the ribbon and are therefore not topological (the location of the ribbon can be detected by the energy terms along the length of the ribbon).

The fact that the topological excitations must be produced in pairs at the ends of

ribbon operators, rather than locally, is no accident. These excitations carry a conserved charge, known as topological charge. Because this charge is conserved, it can only be produced in one region by moving it out of another region. Ribbon operators do exactly this, and the charge carried by the excitation produced at one end of the ribbon operator must be balanced by the charge carried by the excitation at the other end.

While in 2+1d phases the anyons are point-like, and so are all produced by ribbon operators, in 3+1d we can have loop-like excitations. Instead of being produced at the ends of ribbon operators, these loop-like excitations are created at the boundary of so-called membrane operators. As the name suggests, membrane operators act across some extended surface in the 3d spatial lattice, often with some finite thickness. These membrane operators must be applied on unexcited regions of the lattice, otherwise they may produce additional excitations or in some cases become ill-defined. The membrane operators, just like the ribbon operators, are topological, meaning that the membrane is largely invisible and deforming it through an unexcited region without changing its boundary leaves the action of the corresponding operator unchanged.

These ribbon and membrane operators carry significant information about topological phases. They can be used to obtain the fusion rules for the associated particles, which describe how two anyons can be combined into a single particle [36]. Furthermore, ribbon and membrane operators allow us to find the braiding statistics of the topological excitations, because these operators encode the creation and motion of the excitations. As we describe in more detail in Section 6.1, this allows us to calculate braiding relations from the commutation relations of the ribbon and membrane operators.

In the following sections we construct the ribbon and membrane operators for the higher lattice gauge theory model and use the operators to study the properties of the excitations directly. We find four types of excitation in our model, with a rough

correspondence to the four energy terms of the Hamiltonian. In 3+1d, two of these types of excitation are point-like and two types are loop-like.

4.2 Electric Excitations

The first type of excitation we construct is called the electric excitation and is primarily associated to the vertex terms of our Hamiltonian. These electric excitations therefore correspond to the "electric charges" that we described at the start of Chapter 3. In order to create these electric excitations, we measure the group element associated to some path on our lattice and apply weights depending on the result. In order to measure a path element, we take the product of edge elements along the path, with inverses if the orientation of the edge is against the orientation of the path. For example, considering the path shown in Figure 4.1, the edges labelled by g_2 and g_3 are aligned against the path, and so those group elements appear with an inverse in the path element. In order to measure each possible value of the path label, we apply an operator of the form

$$S^{\vec{\alpha}}(t) = \sum_{g \in G} \alpha_g \delta(\hat{g}(t), g), \tag{4.1}$$

where $\hat{g}(t)$ is the path element for path t and α_g is a coefficient (or weight) for the element g. The operator in Equation 4.1, which measures the path element and gives a weight to each possibility, is our first example of a ribbon operator and can produce excitations at the vertices at the two ends of the ribbon.

At this point, we would like to note that the ribbon operators that produce the electric excitations in the higher lattice gauge theory model are equivalent to the corresponding operators in Kitaev's Quantum Double model [16]. As we go on, we will find that several features of the Quantum Double model (which is based on lattice gauge theory) carry over to the higher lattice gauge theory model. However, we will also see important distinctions between the two models. For instance, as we describe in Section 5.3 (and prove in Ref. [3], in Section S-I A of the Supplementary Material), some

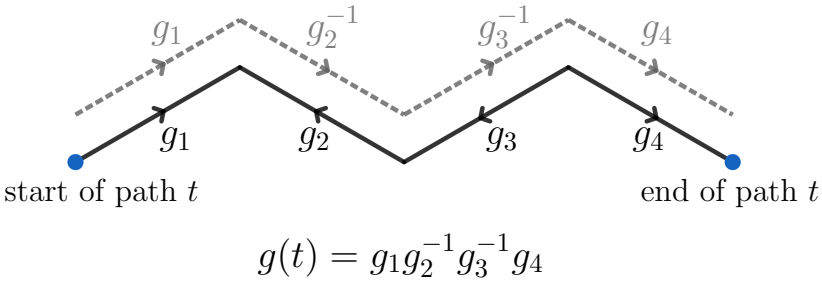

Figure 4.1: An electric ribbon operator measures the value of a path and assigns a weight to each possibility, creating excitations at the two ends of the path. In order to find the group element associated to the path, we must first find the contribution of each edge to the path. In this example, the edges along the path are shown in black. Some of the edges are anti-aligned with the path and so we must invert the elements associated to these edges to find their contribution to the path. This is represented by the grey dotted lines, which are labelled with the contribution of each edge to the path.

---

of our electric excitations are confined, meaning that the two excitations at the end of the electric ribbon operator cost energy to separate. This manifests in the ribbon operators exciting the edges along the ribbon in addition to the vertices at the ends of the ribbon.

As discussed in Section 4.1, one of the important properties of unconfined ribbon operators is that they are topological, meaning that the action of the ribbon operator on an unexcited region is unchanged if we deform the ribbon on which we apply the operator, provided that this deformation does not cause the ribbon to cross another excitation. In the case of the electric ribbon operators, this topological property is a result of the fake-flatness condition of the ground state. If we deform the ribbon from an initial path t_1 to a final path t_2, then the two paths enclose a surface which is fake-flat (provided no excitations are present on that surface). This means that the path elements of the two paths satisfy $g(t_1)g(t_2)^{-1} = \partial(e)$ for some element $e \in E$. The two path elements therefore differ only by an element of $\partial(E)$. If the ribbon operator is not sensitive to elements in $\partial(E)$ (i.e., if the weights α_g in Equation 4.1 satisfy $\alpha_g = \alpha_{g\partial(e)}$ for all $e \in E$ and $g \in G$) then the ribbon operator is therefore left

invariant under such a change in path. As we discuss in Section 5.3, electric ribbon operators that are sensitive to elements in $\partial(E)$ are confined (or contain a confined component), so the unconfined ribbon operators (which are not sensitive to elements in $\partial(E)$) satisfy the condition to be topological, as we require.

Any set of weights α_g that we choose in Equation 4.1 will give a valid ribbon operator. Varying the weights therefore takes us through a space of these electric ribbon operators. A particularly useful basis for this space has the weights described using representations of G, as we anticipated in Chapter 3. Each basis element is labelled by an irreducible representation (irrep) of the group G, along with two matrix indices. Then for the irrep R and the indices a and b, the corresponding basis ribbon operator is given by

$$\hat{S}^{R,a,b}(t) = \sum_{g \in G} [D^R(g)]_{ab} \delta(\hat{g}(t), g), \tag{4.2}$$

where $D^R(g)$ is the matrix representation of the element g in the irrep R. The operator labelled by the identity irrep is then the identity operator

$$1 = \sum_{g \in G} \delta(g, \hat{g}(t))$$

and so does not produce any excitations. Any other basis operator (i.e. an operator labelled by a non-trivial irrep) does produce excitations at the two ends of the ribbon.

In addition to determining which operators excite the vertices, this basis is a good choice for examining the properties of the electric excitations. The excitations of the model carry a set of conserved charges, called *topological charge*. The excitations are then split into sectors that have different values of topological charge. In 3+1d, there are multiple types of topological charge, corresponding to different measurement surfaces, as we explain in Chapter 7 (and in more detail in Ref. [2]). In this work we will consider the charge measured by a sphere, which captures the point-like character of an excitation or set of excitations, and the charge measured by a torus, which

captures the loop-like character. The electric excitations are point-like and so we are interested in the point-like charge, or the charge measured by a sphere. As we show in Ref. [2] (in Section IX A1), the point-like topological charge for (non-confined) pure electric excitations is labelled by the irreps of the group $G/\partial(E)$ (the quotient removes the confined excitations) and the basis operators given above transport definite values of topological charge. That is, in the operator above R labels a conserved charge, while the matrix indices a and b describe some internal space to the sector (as we expect from our discussion of gauge theory in Chapter 3). In particular, when R is the trivial irrep the ribbon operator transports the vacuum charge, as we require from the fact that the ribbon operator is just the identity.

4.3 Blob Excitations

The next point-like excitations that we find are called blob excitations, because they primarily correspond to violations of the blob terms of the Hamiltonian. These excitations are therefore associated to the non-trivial closed surfaces we discussed in Section 2.2.4. That is, the blob excitations correspond to non-trivial "2-fluxes" on a sphere. The Hamiltonian includes a blob term for each blob (3-cell), which enforces that a particular product of the surface elements around the blob gives the identity element, 1_E. To violate this condition, we can therefore consider multiplying one of the plaquette labels on a blob by an element e of E. However, each plaquette is shared by two blobs, one on either side of the plaquette. Changing a plaquette's label will therefore excite both of the adjacent blobs. We can try to correct this by changing another plaquette on this second blob. However, that plaquette will in turn be connected to a third blob, which will become excited, as illustrated in Figure 4.2. This means that we simply move the second excitation from the second to the third blob. We therefore see that we produce these blob excitations in pairs, just as with the electric excitations. The series of plaquettes that we have to change to produce the excitations forms a string that passes through the centres of the blobs (3-cells). The operator that changes the

plaquette labels is therefore another of our ribbon operators.

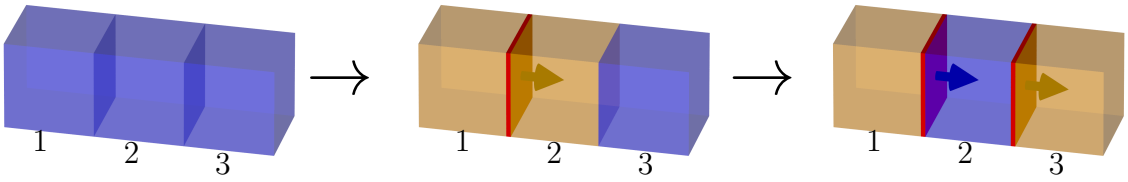

Figure 4.2: We consider a series of blobs on the ground state (left-most image). In the ground state, all of the blob terms are satisfied, which we represent here by colouring the blobs blue. Changing the label of the plaquette between blobs 1 and 2 excites both adjacent blobs, as can be seen in the middle image (we represent excited blobs by colouring them orange). Multiplying another plaquette label on blob 2 to try to correct it just moves the right-hand excitation from blob 2 to blob 3 (right-most image). In each step, the plaquettes whose labels we changed are shown in red and their orientations are indicated by an arrow.

4.3.1 Blob Operators in the $\triangleright$ Trivial Case

We have so far been somewhat vague about the precise action of our blob ribbon operator. This is because the action depends on which special case (as described in Section 2.2.6) that we take. In the simplest case, where $\triangleright$ is trivial (Case 1 in Table 2.1), the action is fairly simple. We choose an element of E to label the operator. We also choose a path on the dual lattice, which passes between the centres of blobs just as a path on the direct lattice passes between vertices. This path cuts through the plaquettes that separate the blobs (analogous to direct paths passing along edges), such as the red plaquettes in the example in Figure 4.2. The choice of element $e \in E$ and path r gives us a blob ribbon operator $B^e(r)$. The action of this operator is just to multiply the labels of all of the plaquettes pierced by this dual path by e or e^{-1}, depending on the orientation of these plaquettes relative to the ribbon (where the orientation of the plaquette is obtained from its circulation using the right-hand rule). If the plaquette is aligned with the ribbon, its label is multiplied by e^{-1} and if it is anti-aligned it is multiplied by e. This results in the two blobs at the end of the path being excited, as we discussed for the example in Figure 4.2. As we explain in Section 5.3, some of the blob ribbon operators (those labelled by an element $e \in E$ for which $\partial(e)$ is not 1_G) also excite the plaquettes pierced by the ribbon operator, resulting in

the corresponding blob excitations being confined.

4.3.2 Blob Operators in the $\triangleright$ Non-Trivial Case

We must modify the action of the blob ribbon operators slightly if $\triangleright$ is non-trivial, as in Cases 2 and 3 of Table 2.1. When $\triangleright$ is non-trivial, we must keep track of the base-points of the plaquettes that we want to change. We first move all of their base-points to a common location at the start of our operator (for example, the base-point of the first plaquette that we want to change), which we call the start-point. Then we multiply each plaquette label by e or e^{-1} (according to the orientation of the plaquette as before) before moving their base-point back to its original location. Recall from the introduction that moving the base-point of a plaquette along path t changes its label from e_p to $g(t)^{-1} \triangleright e_p$ (see Figure 2.11 for a reminder). Assuming that the plaquette p is aligned with the ribbon, the total change to e_p is therefore

$$\begin{aligned} e_p &\to g(t)^{-1} \triangleright e_p \text{ (move base-point)} \\ &\to [g(t)^{-1} \triangleright e_p]\, e^{-1} (\text{ postmultiply by } e^{-1}) \\ &\to g(t) \triangleright ([g(t)^{-1} \triangleright e_p]\, e^{-1}) \text{ (move base-point back)} \\ &= e_p[g(t) \triangleright e^{-1}], \end{aligned}$$

where t is the path from the base-point $v_0(p)$ of plaquette p to the start-point of our operator. Generally, we can write the action of the blob ribbon operator on a plaquette p pierced by the ribbon as

$$B^e(r) : e_p = \begin{cases} e_p[g(s.p - v_0(p))^{-1} \triangleright e^{-1}] & \text{if } p \text{ aligns with } t \\ [g(s.p - v_0(p))^{-1} \triangleright e]e_p & \text{if } p \text{ anti-aligns with } t, \end{cases} \tag{4.3}$$

where $g(s.p - v_0(p))$ is the label of the path from the start-point of our operator to the base-point of p. Provided that the ribbon operator acts on a fake-flat region of the lattice, the path $(s.p - v_0(p))$ only needs to be defined up to smooth deformations (as we show in Ref. [2]) and it is convenient to choose a single path that runs from

the start-point through the base-point of every plaquette pierced by the ribbon, with the particular path $(s.p - v_0(p))$ then being the appropriate section of this larger path. This dependence of the blob ribbon operator on paths from the start-point to the base-points of the plaquettes means that the operator acts on a ribbon, rather than a path. We call the path on the dual lattice, which pierces the affected plaquettes, the dual path and call the path from the start-point to the base-points the direct path (because it lies on the direct lattice).

This more complicated action means that the start-point vertex may be excited by the ribbon operator, because vertex transforms at the start-point affect the path elements $g(s.p - v_0(p))$ that appear in the ribbon operator and so do not commute with the ribbon operator. While we leave a more detailed discussion and proof of this for Ref. [2], we can understand this somewhat intuitively using the geometric interpretation discussed in Chapter 3. The blob excitations correspond to non-trivial 2-fluxes, and these 2-fluxes must be defined with respect to a base-point, which is the start-point of the ribbon operator. Moving this base-point induces a $g \triangleright$ action on the label of the 2-flux. Then, because applying a vertex transform at a vertex is analogous to parallel transport of that vertex (see Section 2.2), the vertex transform also has a similar $g \triangleright$ action on the label of the 2-flux produced by the ribbon operator. The ribbon operators which are invariant under this action therefore commute with the vertex transforms and so with the vertex energy term at the start-point. However, generic ribbon operators are not invariant under this action, which leads to some of them exciting the start-point. Specifically, given a linear combination $\sum_e \alpha_e B^e(r)$ of group-labelled blob ribbon operators, the start-point is not excited if the coefficients α_e are a function of $\triangleright$-class, so that $\alpha_f = \alpha_{g \triangleright f}$ for all pairs $g \in G$ and $f \in E$. On the other hand, the start-point is definitely excited if the sum of the coefficients in each $\triangleright$-class is zero (i.e. $\sum_{g \in G} \alpha_{g \triangleright e} = 0$ for all $e \in E$). Ribbon operators that do not satisfy either of these conditions do not produce eigenstates of the vertex term when acting on the ground state, but can be written as a linear combination of operators

that excite the vertex and operators that do not.

Just like the electric ribbon operators, the unconfined blob ribbon operators (we discuss the pattern of confinement in Section 5.3) are topological, meaning that they can be deformed through an unexcited region without changing their action. In this case, the reason that the ribbon operators are topological is that deforming the ribbon operator is equivalent to applying edge transforms, which act trivially when applied on unexcited edges (due to the fact that edge transforms can be absorbed into the edge energy terms, as described in Equation 2.11). We give a more detailed proof of this in Ref. [2].

4.4 E-Valued Loops

In this model, we find two main types of loop-like excitation, in addition to the point-like excitations that we already discussed. Whereas point-like excitations are produced by ribbon operators, which act on a path or ribbon, loop-like excitations are produced by so-called *membrane operators*, which act across a general two-dimensional surface. The first loop-like excitations that we find, which we call E-valued loop excitations, are produced by membrane operators which act primarily on the surface labels. These membrane operators measure the label of a surface in our lattice, using the rules for combining surface elements given in Sections 2.2 and 2.2.5, and apply a weight for each possible label. This procedure is very similar to that used for the electric excitations (as described by Figure 4.1 and Equation 4.2), except that we measure a surface rather than a path. Indeed, if the electric excitations are charges for the 1-gauge field, then the E-valued loops are "2-charges" corresponding to the higher gauge (2-gauge) field. This means that they can interact with our 2-fluxes, in a way somewhat analogous to the interaction of ordinary charges and fluxes, as will be discussed more closely when we examine braiding. In the same way as with the electric excitations, the weights describe a space of operators and an appropriate basis is given using irreps, this time

of the group E. For an irrep μ and matrix indices a and b, the operator acting on a membrane m is given by

$$L^{\mu,a,b}(m) = \sum_{e \in E} [D^{\mu}(e)]_{ab} \delta(\hat{e}(m), e), \tag{4.4}$$

where $\hat{e}(m)$ is the total surface element of membrane m. This surface element can be written in terms of the labels of the plaquettes making up the membrane as

$$\hat{e}(m) = \prod_{\text{plaquettes } p \in m} g(s.p(m) - v_0(p)) \triangleright e_p^{\sigma_p}, \tag{4.5}$$

where $v_0(p)$ is the base-point of the plaquette p, $s.p(m)$ is the base-point with respect to which we measure the surface label (and which we call the start-point of the membrane) and σ_p is $+1$ if the plaquette's orientation matches that of m and -1 if it is anti-aligned with m. The particular paths $(s.p(m) - v_0(p)$ in this expression (along with the order of the product if E is non-Abelian) are chosen to be consistent with the rules for combining surfaces given in Section 2.2. There are generally many ways of composing the surfaces and each way is associated to different paths $(s.p(m) - v_0(p))$ and different orders of multiplication. These will give the same result, provided that fake-flatness is satisfied on the membrane. If the membrane is not fake-flat (which occurs if certain excitations are present on the membrane), then the different ways of combining the plaquettes into the membrane are not consistent and we must have a convention for how we do so, or specify the specific composition we use when we define the membrane. Generally, however, we use the membrane operator in regions where no other excitations (especially fake-flatness violations) are present. Applying the membrane operator in Equation 4.4 on a fake-flat region causes the edges bordering the surface to become excited (as long as μ is non-trivial), as indicated in Figure 4.3. If μ is the trivial irrep, then the membrane operator is instead the identity operator, so the operator does not produce any excitations.

As with the blob excitations, there are some features that depend on which special case

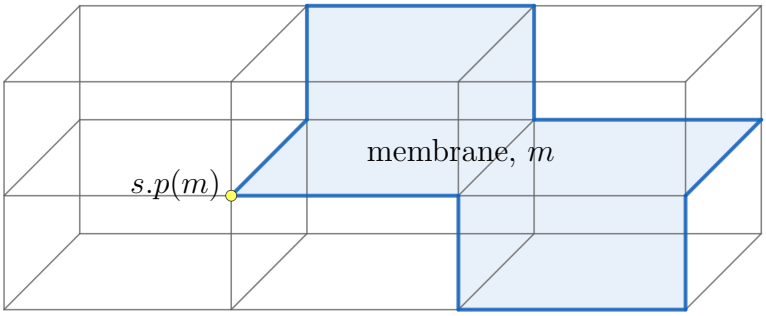

Figure 4.3: We consider applying an E-valued membrane on the shaded membrane m in a fragment of the three-dimensional lattice. The membrane operator measures the surface label of the membrane, with a weight for each possible label. When measuring a surface, if $\triangleright$ is non-trivial we must specify the base-point of that surface. The base-point of the surface measured by the membrane operator (the yellow dot) is called the start-point of the membrane. A non-trivial E-valued membrane operator excites the edges (solid blue lines) on the boundary of the membrane, and may also excite the start-point of the membrane.

---

from Section 2.2.6 we look at. As we explained in Section 2.2.6, the Peiffer conditions (Equations 2.4 and 2.5 in Section 2.2) imply that E is Abelian when $\triangleright$ is trivial. This is because we require $\partial(e) \triangleright f = efe^{-1}$ for all $e, f \in E$ from the second Peiffer condition (Equation 2.5). When $\triangleright$ is trivial, this condition becomes $f = efe^{-1}$, which implies that E is Abelian. In this case, the irreps are all one dimensional, so we can drop the matrix indices a and b in Equation 4.4.

On the other hand, when $\triangleright$ is non-trivial E may be non-Abelian and so generally we must include the matrix indices. In addition, even if E is Abelian the membrane operator has an additional feature when $\triangleright$ is non-trivial. When we measure the surface label, we must choose a base-point for that surface. We call this base-point the start-point of our membrane operator. Much as with the blob excitations, this start-point can be excited by the operator, which reflects the non-trivial transformation undergone by the operator when we move the start-point (due to the connection between vertex transforms and parallel transport). Similarly to the blob ribbon operator, the start-point is not excited if the membrane operator is made of a linear combination $\sum_e \alpha_e \delta(\hat{e}(m), e)$ whose set of coefficients α_e is a function of $\triangleright$-class (i.e. $\alpha_e = \alpha_{g \triangleright e}$). On the other

hand, the start-point may be excited if the coefficient transforms non-trivially under the $\triangleright$ action, as we describe in more detail in Refs. [2] and [3]. The start-point excitation is significant because it can carry a (point-like) topological charge, which must be balanced by a point-like charge on the loop itself, as we prove explicitly in Ref. [2] in Section IX A1. Furthermore, as we describe in Ref. [2] in Section V, it seems that this point-like charge can be confined for certain membrane operators (meaning it costs energy to separate the start-point charge from the point-like charge carried by the loop), in which case the charge drags a line of excited edges between the start-point and the loop excitation. We note that this confinement can only occur when E is non-Abelian, so it does not occur for Cases 1 and 2 in Table 2.1.

In a similar way to ribbon operators, unconfined membrane operators should be topological. This means that if we deform the membrane through an unexcited region of the lattice, while keeping the position of the excitations produced by the membrane operator fixed, the action of the membrane operator should be left unchanged. In the case of the E-valued membrane operators, this property is a result of the blob energy terms being satisfied. If we start with the membrane operator applied on a membrane m_1 and then deform the membrane into a new position m_2, then the surface $m_1 \cdot m_2^{-1}$ is a contractible closed surface (where m_2^{-1} is the surface obtained from m_1 by reversing its orientation). Then, provided that the closed surface does not enclose any excitations, the blob energy terms ensure that the label of this closed surface is just the identity 1_E. This means that the surface labels $e(m_1)$ and $e(m_2)$ of the two surfaces satisfy $e(m_1)e(m_2)^{-1} = 1_E$ and so $e(m_1) = e(m_2)$, meaning that the membrane operator is unaffected by the change of membrane.

4.5 Magnetic Excitations

The final type of elementary excitation is the magnetic excitation, named so due to its correspondence to the magnetic excitations in Kitaev's Quantum Double model

[16] and its analogy to magnetic flux. In Chapter 3, we discussed how the presence of magnetic excitations results in closed loops linking with the magnetic flux tubes having non-trivial path label. We therefore expect these magnetic excitations to be created by changing the edge labels of the lattice in order to introduce this non-trivial flux. In this section, we will see in more detail how this is done.

The magnetic loop excitation is primarily associated with excitation of the plaquette energy terms, due to its nature as a flux tube. Recall from Section 2.2.5 (see Figure 2.30) that the plaquette term checks that the 1-flux of the plaquette is trivial, where the 1-flux is given by the product of edge elements around the boundary multiplied by the image under ∂ of the plaquette's surface element. In order to create an elementary flux excitation, we want to excite the fewest number of plaquettes by changing edge labels (changing the plaquette label could also excite the plaquette, but this results in blob excitations, as we saw in Section 4.3). We therefore consider trying to excite a plaquette by changing one of the edges on that plaquette. However, in three spatial dimensions the edge will generally be shared by multiple plaquettes. Therefore, changing the edge label will excite all of the plaquettes surrounding this edge, not just the one we wished to excite. We can try to fix one of these additional excited plaquettes by changing the label of another edge on that plaquette, but this will in turn excite the other plaquettes surrounding that edge. We can repeat this process, but we can never get rid of all of the additional excitations. When we change a series of edge labels in this way, the result is a set of excited plaquettes that lie on a closed loop which intersects the centres of these plaquettes. For example, in the first step when we change just one edge, we excite all of the plaquettes on a closed loop around that edge, as shown in Figure 4.4. Upon changing additional edge labels, we deform this loop but always get a closed string of excited plaquettes (unless we collapse the loop to nothing). This means that the magnetic excitation is indeed loop-like, as we predicted earlier. The edges that we have to change to produce the magnetic excitation are bisected by a membrane bounded by the excited loop, as shown in Figure 4.5. We call the membrane cutting these edges

the dual membrane, because the membrane bisects the edges of the lattice rather than lying on the lattice itself. The fact that we must change degrees of freedom across a membrane in order to produce a general magnetic excitation means that the creation operator is a membrane operator, which we call the magnetic membrane operator.

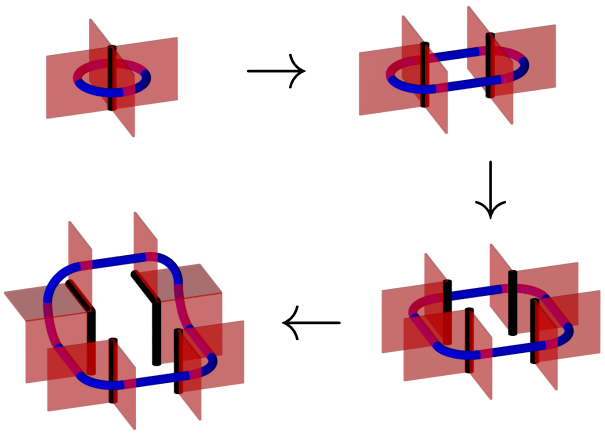

Figure 4.4: In order to excite one of the plaquettes in the lattice and produce a magnetic excitation, we change the label of one of the edges (black cylinders) on the boundary of the plaquette. However, this excites all of the plaquettes adjacent to that edge, as shown in the first image (the excited plaquettes are shown in red). Note that these plaquettes lie on a closed loop (blue tube) through their centres. If we change another edge label to try to prevent some of the plaquette excitations, we will excite the other plaquettes adjacent to that edge, as shown in the second image. This enlarges the loop of excitations. Repeating the process, by changing the additional edges shown in black in each step, simply changes the shape of this loop (unless we change all of the edges bisected by a closed membrane and shrink the loop to nothing). This tells us that the magnetic excitations are loop-like.

4.5.1 Magnetic Membrane Operators When $\triangleright$ Is Trivial

Now we should consider more concretely the action of the membrane operator. The features of the magnetic excitation depend strongly on the special case that we take. We first consider the $\triangleright$ trivial case (Case 1 in Table 2.1), in which case the membrane operator is analogous to the 2+1d magnetic ribbon operator from Kitaev's Quantum Double model [16]. Just like the ribbon operator from the Quantum Double model, a magnetic membrane operator is labelled by a group element h of G. We denote the magnetic membrane operator labelled by h and acting on a membrane m by $C^h(m)$. When the group G is Abelian, the action of this membrane operator is simple. We just

multiply each of the affected edges (those cut by the dual membrane) by the element h or its inverse, depending on the orientation of the edge relative to the normal of the membrane. On the other hand, when the group is non-Abelian the action is slightly more complicated. Instead of multiplying each edge by h, we must multiply by some element in the same conjugacy class as h. To determine which element this is, we must first endow our operator with a privileged point, called the start-point. Furthermore, we must specify a path from this start-point to each edge we want to change. Denoting the path to edge i by t_i, the edge label g_i of i transforms as $g_i \to g(t_i)^{-1}hg(t_i)g_i$ or $g_i \to g_i g(t_i)^{-1}h^{-1}g(t_i)$ depending on the orientation of that edge. This action is shown in Figure 4.5. The paths involved in this action lie on a second membrane, which we call the direct membrane, so the support of the membrane operator actually lies on both the direct and dual membranes, which we sometimes refer to together as the thickened membrane. This is analogous to how the ribbon operators in Kitaev's Quantum Double model act on a ribbon [16], which is a thickened string. We can then write the action of a magnetic membrane operator $C^h(m)$ on an edge i cut by the dual membrane as

$$C^h(m) : g_i = \begin{cases} g(s.p - v_i)^{-1}hg(s.p - v_i)g_i & \text{if } i \text{ points away from} \\ & \text{the direct membrane} \\ g_i g(s.p - v_i)^{-1}h^{-1}g(s.p - v_i) & \text{if } i \text{ points towards the} \\ & \text{direct membrane,} \end{cases} \tag{4.6}$$

where v_i is the end of edge i that lies on the direct membrane and the path $(s.p - v_i)$ runs from the start-point to this vertex. Note that while it may seem that the action of the membrane operator depends strongly on the particular choices for the paths $(s.p - v_i)$ and not just their end-points, the action is invariant under deformation of these paths provided that they are deformed over a region satisfying fake-flatness. This is because the group elements assigned to two paths differing by such a deformation only differ by an element of $\partial(E)$, due to the fake-flatness condition on the surface over which

we deform the path. When $\triangleright$ is trivial, elements of $\partial(E)$ are in the centre of G (see Section 2.2.6) and so do not affect the expression $g(s.p-v_i)^{-1}hg(s.p-v_i)$ (conjugation by an element in the centre is trivial). Because we usually apply membrane operators on regions without any other excitations (although in Ref. [2] we describe one situation where we do apply a membrane operator on an exited region, in order to produce two linked magnetic excitations), this means that we do not generally need to specify the precise positions of the paths. In addition, just like the other non-confined ribbon and membrane operators, the magnetic membrane operator is topological when applied on an unexcited region, meaning that we can deform the membrane on which we apply it without affecting the action of the operator, provided that we deform the membrane over a region containing no excitations and we keep the positions of any excitations produced by the membrane operator fixed. This arises from the fact that a contractible closed magnetic membrane operator acts in the same way as a series of vertex transforms in the region enclosed by the membrane operator (as we show in Ref. [2]) and these transforms act trivially on the ground state (as described by Equation 2.9). We can then deform an open magnetic membrane operator by applying these closed magnetic membrane operators, which act trivially, so the action of the open membrane operator must be the same before and after the deformation.

As with the blob and E-valued loop excitations, the start-point of the magnetic membrane operator may be excited. As mentioned previously, the magnetic excitations in this model are analogous to the magnetic excitations from Kitaev's Quantum Double model [16], and the potential start-point excitation is also present for the magnetic ribbon operators in that model (see for example Ref. [78]). We can interpret the start-point of the magnetic membrane operators in terms of gauge theory. Recall from Chapter 3 that whenever we measure a flux, we must do so with respect to a certain start-point. The start-point of the membrane operator is the start-point with respect to which the flux label of the corresponding excitation is well-defined. Suppose that we have a flux h when measured with respect a certain start-point. As we described in

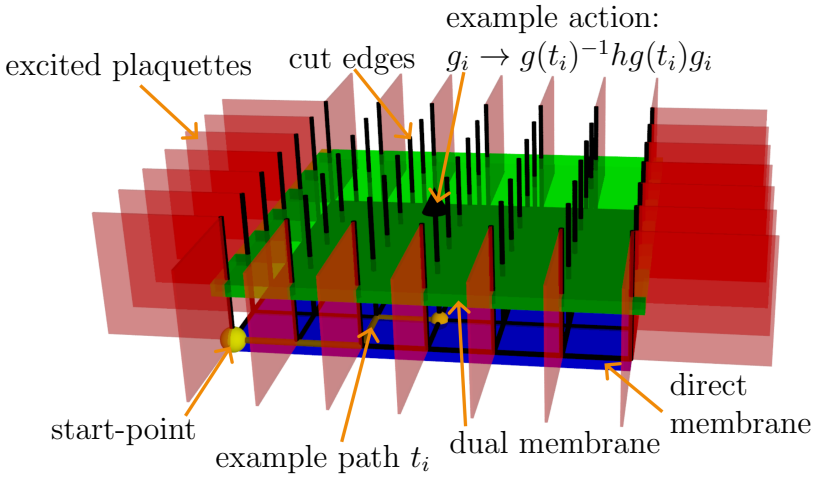

Figure 4.5: Here we give an example of the membranes for the flux creation operator (magnetic membrane operator). The dual membrane (green) cuts through the edges changed by the operator. The direct membrane (blue) contains a vertex at the end of each of these cut edges (such as the orange sphere). A path from a privileged start-point to the end of the edge (such as the example path, t_i) determines the action on the edge. This action leads to the plaquettes around the boundary of the membrane being excited.

Chapter 3, if we were to measure the flux from a different start-point, the flux that we measure would have the form $\hat{g}(t)h\hat{g}(t)^{-1}$, where $\hat{g}(t)$ is a path element operator for which we are not generally in an eigenstate (even if there are no excitations present other than the magnetic flux). Note that the element $\hat{g}(t)h\hat{g}(t)^{-1}$ is still in the conjugacy class of h, indicating that this conjugacy class is independent of the start-point even if the flux element is not. This interpretation of the start-point can help us to understand why the start-point of a magnetic membrane operator may be excited. The vertex transform at a vertex acts like parallel transport of that vertex, and so we can think of the vertex transform at the start-point as moving that start-point, which results in conjugation of the flux label. In order to diagonalize the vertex term at the start-point, we must therefore take a linear combination of magnetic membrane operators with different labels in the conjugacy class of h (so that we are considering a state in a superposition of different flux labels), as we prove in Ref. [2]. If the start-point is unexcited, it means that the membrane operator is not sensitive to changes to the start-point from which we measure the flux, which occurs when the membrane operator produces an equal combination of fluxes in the conjugacy class (i.e. the coef-

ficients of the linear combination are the same for each element in the conjugacy class). This equal combination of elements in the conjugacy class produces a flux tube with a trivial point-like charge (or in the 2+1d case such as the Quantum Double model, a pair of excitations that can be annihilated, as described in Ref. [75]).

4.5.2 Magnetic Membrane Operators When $\partial(E)$ Is in the Centre of G and E is Abelian

So far this consideration of the magnetic membrane operators has all been in the case where $\triangleright$ is trivial. As mentioned earlier, the form of the membrane operator depends quite strongly on which case we look at. In the case where $\triangleright$ is non-trivial but we restrict to fake-flat configurations (Case 3 in Table 2.1), we cannot include the magnetic excitations at all, because the magnetic excitations violate the plaquette terms and hence break fake-flatness. The most interesting case is the remaining special case we mentioned in Section 2.2.6 (Case 2 in Table 2.1), where we loosen the restrictions on the crossed module without throwing out the non-fake-flat configurations. Specifically, we require that ∂ maps to the centre of G and that E is Abelian. In this case, we are allowed to keep the magnetic excitations, although their operators must be modified. We briefly describe this modification here, but for a full and detailed description, please see Ref. [2]. The new membrane operators act on edge elements in the same way as described above for the $\triangleright$ trivial case, but they also act on the plaquette elements around the membrane. The dual membrane that cuts edges also cuts through plaquettes. These plaquettes have base-points that can be either on the direct membrane or away from it (on the other side of the dual membrane). If a cut plaquette has its base-point on the membrane, then it is acted on with a $\triangleright$ action. This action on a plaquette p, with label e_p, is given by $e_p \to (g(t)^{-1}hg(t)) \triangleright e_p$, where $g(t)$ is the path from the start-point of the operator to the base-point of the plaquette being changed (note that, just like the paths in the magnetic membrane operator when $\triangleright$ is trivial, deforming these paths over an unexcited region of the lattice leaves this action invariant). On the other hand, any plaquette whose base-point is away from

the direct membrane is left unaffected. An example of this action is shown in Figure 4.6. This action on the plaquettes is analogous to how the vertex transform affects plaquettes based at that vertex, but not plaquettes that are not based at the vertex. The difference between the action of the membrane operator and the vertex transforms is that the vertex transform acts on all plaquettes whose base-points are that vertex, but the magnetic membrane operator acts only on plaquettes whose base-points lie on the direct membrane and which are also cut by the dual membrane. In this sense, the vertex transform is like a closed magnetic membrane operator, whose dual membrane encloses that vertex and whose direct membrane is just the vertex itself.

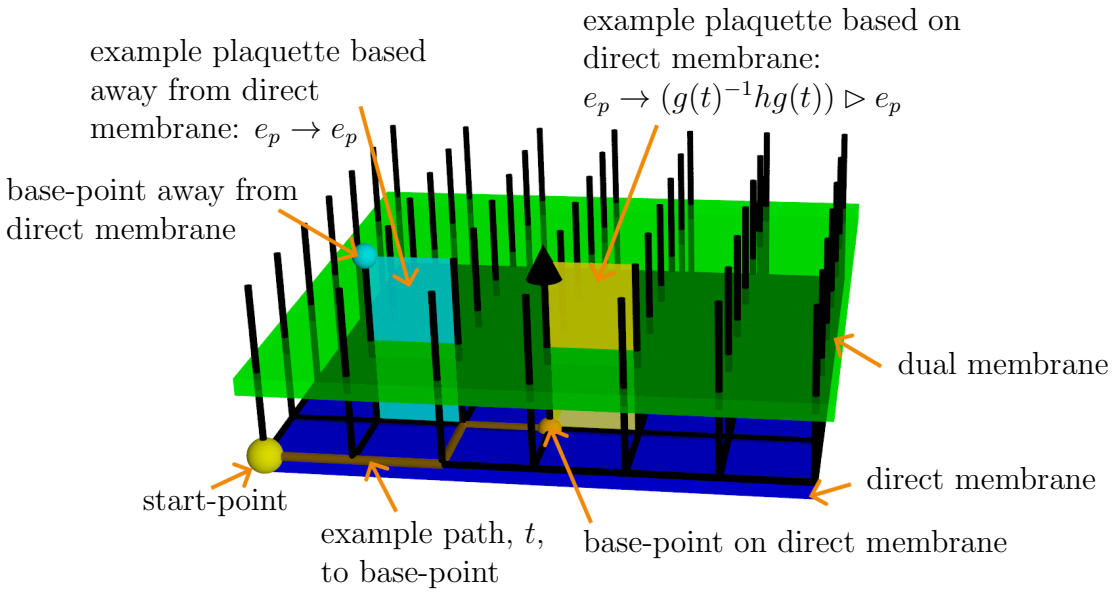

Figure 4.6: In addition to changing the edges that are cut by the dual membrane, when $\triangleright$ is non-trivial the magnetic membrane operator affects the plaquettes that are cut by the dual membrane if their base-points lie on the direct membrane

---

In addition to this $\triangleright$ action, we have to include certain blob ribbon operators in the membrane operator, with these ribbons running from a special blob defined in the membrane operator, which we call blob 0, to the plaquettes of the direct membrane. The label of these blob ribbon operators depends on the label of the corresponding plaquette on the direct membrane. For a plaquette p with label e_p and base-point $v_0(p)$, the label of the associated blob ribbon operator is

$$f(p) = [g(s.p - v_0(p)) \triangleright e_p^{\sigma_p}]\,[(h^{-1}g(s.p - v_0(p))) \triangleright e_p^{-\sigma_p}], \tag{4.7}$$

where e_p is the label of the plaquette p, and σ_p is $+1$ if the orientation of the plaquette, determined from its circulation using the right-hand rule, points away from the dual membrane and is -1 if the orientation is towards the dual membrane. Then the associated blob ribbon operator is $B^{f(p)}(\text{blob } 0 \to \text{blob } p)$, where $v_0(p)$ is the base-point of plaquette p and blob $0 \to$ blob p is shorthand for the ribbon that starts at blob 0 and terminates in the blob attached to plaquette p and cut by the dual membrane, with the start-point of the ribbon being the start-point of the membrane. Note that the blob ribbon operator is topological, so the position of the ribbon only needs to be specified up to deformations (and similarly for the path $(s.p - v_0(p))$ in the label of the ribbon operator). The full membrane operator is then given by

$$C_T^h(m) = C_\triangleright^h(m) \prod_{\substack{\text{plaquette } p \\ \text{on membrane}}} B^{f(p)}(\text{blob } 0 \to \text{blob } p), \tag{4.8}$$

where $C_\triangleright^h(m)$ includes the action of the membrane operator on the edges from the $\triangleright$ trivial case (see Equation 4.6), as well as the $\triangleright$ action on the plaquettes that are cut by the dual membrane and based on the direct membrane. That is, the action of $C_\triangleright^h(m)$ on an edge i cut by the dual membrane is given by

$$C_\triangleright^h(m) : g_i = \begin{cases} g(s.p - v_i)^{-1} h g(s.p - v_i) g_i & \text{if } i \text{ points away from} \\ & \text{the direct membrane} \\ g_i g(s.p - v_i)^{-1} h^{-1} g(s.p - v_i) & \text{if } i \text{ points towards the} \\ & \text{direct membrane,} \end{cases} \tag{4.9}$$

and the action on a plaquette p cut by the dual membrane is given by

$$C_\triangleright^h(m) : e_p = \begin{cases} (g(s.p - v_0(p))^{-1} h g(s.p - v_0(p))) \triangleright e_p & \text{if } v_0(p) \text{ lies on the} \\ & \text{direct membrane} \\ e_p & \text{otherwise.} \end{cases} \tag{4.10}$$

In the $\triangleright$ trivial case (Case 1 of Table 2.1), provided that we apply the membrane operator on a fake-flat region, only the plaquettes around the membrane and potentially the start-point of the membrane operator are excited by the action of the membrane operator. In this more general case (Case 2 of Table 2.1), in addition to these energy terms the special blob, blob 0, may also be excited, as may the edges and blobs surrounding the membrane (which are adjacent to the excited plaquettes). Near an excited plaquette, the edge and blob terms cease to commute and also become inconsistent with changes to the branching structure of the lattice (in fact, the blob terms are no longer independent of the method used to combine the surface elements on the blob), meaning that they are ill-defined. This effect would only become worse if we lifted the condition on ∂. In that case, the blob ribbon operators that we add would generally be confined, which could lead to plaquette excitations away from the boundary (although we expect this to be true only for certain labels of the membrane operator) and this is the reason that we do not consider the fully general case. In Case 2 the problematic plaquette excitations are restricted to the boundary of the membrane, and any topological quantities can be measured far from this boundary (if the membrane is sufficiently large), so the inconsistencies from the plaquette excitations are not important.

Due to the extra features of the magnetic membrane operator in this case, the magnetic excitation may carry both an ordinary 1-flux and a 2-flux, as we discussed in Chapter 3. Recall that a non-trivial 2-flux indicates a closed surface with a non-trivial label. We can see that the loop excitation must be associated to a non-trivial surface when blob 0 of the membrane operator is excited. This is because blob 0 being excited indicates that this blob carries a non-trivial 2-flux, which must be balanced by a 2-flux belonging to the loop-like excitation itself.

Chapter 5

Condensation and Confinement

5.1 Introduction

In the study of topological phases of matter, two important and related concepts are condensation and confinement. These allow us to discuss different, but related, phases of matter. Given one topological phase of matter, with its own set of conserved topological charges, it is possible to deform the Hamiltonian to *condense* some of the charges, by which we mean that these charges proliferate and become part of the ground state of the new Hamiltonian. When this occurs, the excitations that carried the condensed charges now carry the vacuum charge in the new model. We say that those excitations *condense* [79, 80, 81, 82, 83]. Any excitations that braided non-trivially with those condensing excitations in the original phase then become confined in the new phase [79, 80]. By this, we mean that it costs energy to separate a confined particle from its antiparticle, or to move and grow loop-like excitations. This process is known as a condensation-confinement transition. While this is fairly well understood in 2+1d, there has been comparatively little study of such transitions in the 3+1d case (examples of work in this area include Refs. [84] and [85]), and so it is interesting to see how condensation and confinement arise in this 3+1d model.

In the higher lattice gauge theory class of models, there exist families of models related

by condensation-confinement transitions. The models within each family have the same groups G and E (which determine the Hilbert space), but different maps ∂. In particular, when E is Abelian the family has an "uncondensed" model, where none of the excitations are confined or condensed. This uncondensed phase is described by the crossed module $(G, E, \partial \to 1_G, \triangleright)$, where ∂ maps to the identity element of G. Changing ∂ to map to a larger subgroup of G then causes a condensation-confinement transition, with the confined and condensed excitations controlled by ∂. When E is non-Abelian, however, it is not possible to construct this uncondensed model. This is because the second Peiffer condition, Equation 2.5, enforces that $\partial(e) \triangleright f = efe^{-1}$ for every pair of elements e and f in E. Because $\partial(e) = 1_G$ for the uncondensed model and $1_G \triangleright$ is always the trivial map, the second Peiffer condition ensures that conjugation is trivial: $efe^{-1} = 1_G \triangleright f = f$, which is incompatible with an Abelian group E. However, we can still discuss the pattern of condensation and confinement when E is non-Abelian, we just have a less clear picture of the uncondensed version of such models.

5.2 Condensation

When an excitation condenses and no longer carries a topological charge, it is no longer topologically protected. This means that it can be produced locally, rather than just at the boundary of a ribbon or membrane operator. We can therefore demonstrate that an excitation is condensed by showing that its ribbon or membrane operator acts equivalently (at least on the ground state) to an operator that just acts in a small region around the excitation. For a point-like excitation, this means that the corresponding ribbon operator acts in the same way as a pair of operators acting near the two ends of the ribbon (these operators are local in the usual sense). On the other hand, a loop-like excitation is itself extended, so the operator that produces them must at least be linearly extended. The excitation is condensed if the membrane operator acts on the ground state in the same way as an operator with support with only linear extent

that acts near the excitation (we say it acts locally to the excitation, even though the operator need not be local in the usual sense). By using this diagnosis, we were able to find the pattern of condensation in the higher lattice gauge theory model.

In the higher lattice gauge theory model, some of the loop-like excitations become condensed. Firstly, the magnetic excitations labelled by elements in $\partial(E)$ condense. This can be seen from considering the magnetic excitations in light of the plaquette term. As we discussed in Chapter 3, the magnetic excitations are associated with closed loops that have non-trivial label. However, the plaquette energy term enforces that the label of closed loops in our lattice match the image under ∂ of the surface element bounded by the loop, rather than enforcing that the path label is the identity element. That is, for a closed loop c bounding a surface s, the labels $g(c)$ of the loop and $e(s)$ of the surface satisfy $\partial(e(s))g(c) = 1_G$. Therefore, the ground state contains closed loops with all values in the image of ∂. A magnetic membrane operator with label in the image of ∂ modifies the labels of closed loops that link with the excitation only by multiplication by another element of $\partial(E)$, and so results in closed loop values already found in the ground state. Therefore, the topological charges of the corresponding magnetic excitations (which we measure with closed paths) belong in the ground state, and so these charges have been condensed (in the uncondensed model, only 1_G is in the image of ∂). In Ref. [2], we explicitly show that the action of a condensed membrane operator on the ground state is equivalent to the action of a (confined) blob ribbon operator around the boundary of the membrane operator, which is local to the excitation, confirming that the corresponding loop-like excitation is condensed. However, note that this does not mean that these magnetic excitations are not excitations at all. Changing the path label by an element of $\partial(E)$ (as done by these magnetic membrane operators) still leads to an excitation, because the path labels do not match the surface enclosed by the loop (and so do not satisfy fake-flatness, implying the presence of a plaquette excitation).

In a similar way, some of the E-valued loops are condensed. This is because the membrane operators that produce them are equivalent (when acting on the ground state) to an electric ribbon operator acting around the boundary of that membrane. The membrane operators for the E-valued loops measure the surface element of the membrane that they are placed on. However, in the ground state the group element e_m assigned to a surface m is related to the path around the boundary of the surface (labelled by $g(\text{boundary})$) by $\partial(e_m)g(\text{boundary}) = 1_G$, due to the plaquette terms. This suggests that we can measure the surface element just by examining the boundary, but because this expression involves only $\partial(e_m)$, this correspondence between the surface and boundary does not fully fix the value e_m. The surface element can be split into a part that describes the image under ∂ of that element and a part in the kernel of ∂, with the former part fixed by the boundary label. This means that the image under ∂ of the surface element can be measured by looking at the boundary of the surface, but not the part in the kernel. Therefore, if a membrane operator is only sensitive to the former part of the surface element, and not to the part in the kernel, then it is equivalent to an operator that simply measures the boundary path element and so the corresponding excitation is condensed.

5.3 Confinement

In the context of commuting projector models, non-confined excitations are produced by ribbon operators (in the case of point particles) with a fixed energy cost independent of the length of the ribbon and membrane operators (in the case of loop-like excitations) with only an energy cost associated to the length of the loop they create (and possibly a constant contribution as well). By contrast, confined excitations are produced by ribbon operators with a cost associated to the length of a ribbon or membrane operators with a cost associated to their area. This means that we can find which excitations are confined from their creation operators.

In the higher lattice gauge theory model, we discovered that some of the point particles are confined. This means that when we create a pair of point excitations it costs energy to separate them. Firstly, some of the blob excitations are confined. Recall from Section 4.3 that the blob ribbon operator multiplies the labels of each plaquette pierced by the ribbon by an element of E. In addition to exciting the blobs at the ends of the ribbon, this action may excite the plaquettes that are pierced by the ribbon. As discussed in Section 2.2.5, the plaquette term (see Figure 2.30) does not just depend on edge elements, but also on the surface label of that plaquette. In particular, it depends on the surface label through the map ∂, as illustrated by Figure 2.30. For a plaquette p, the plaquette term projects onto states for which the plaquette label e_p and the path label g_p of the boundary of the plaquette are related by $\partial(e_p)g_p = 1_G$. Therefore, if our ribbon operator changes $\partial(e_p)$ for the plaquettes that it pierces, it will excite those plaquettes. Because the blob ribbon operator labelled by an element $e \in E$ multiplies the plaquette label e_p of any plaquette p pierced by the ribbon by an element of the form $g \triangleright e$ (or the inverse), the ribbon operator changes $\partial(e_p)$ when $\partial(g \triangleright e)$ is non-trivial. Because of the Peiffer condition Equation 2.4, $\partial(g \triangleright e) = g\partial(e)g^{-1}$ and so $\partial(g \triangleright e)$ is the identity element when $\partial(e) = 1_G$ regardless of the value of g. That means that the blob ribbon operator will change $\partial(e_p)$ for each plaquette p pierced by the ribbon (and so will excite these plaquettes) precisely when e is outside the kernel of ∂. This means that the blob ribbon operators whose labels are outside the kernel will excite a series of plaquettes along the ribbon, with the energy cost associated to this growing with the length of the ribbon. That is, the blob excitations produced by ribbon operators with label outside the kernel of ∂ are confined, while those produced by ribbon operators with label in the kernel are not. Because the mechanism of this confinement is via the plaquette terms (which enforce fake-flatness), if we restrict to fake-flat configurations (Case 3 in Table 2.1) we are not allowed to include the confined blob excitations, which violate fake-flatness.

We also found that some of the electric excitations are confined, with the ribbon oper-

ators exciting the edges measured by the electric ribbon operator. This occurs because the edge energy terms change the path element that the electric ribbon operator measures, by an element of $\partial(E)$ (the image of ∂). Therefore, if the ribbon operator is sensitive to the path element changing by a factor in $\partial(E)$, then the associated electric excitation is confined (or at least contains a confined component). We can be more precise by using the basis for the electric ribbon operators in which the basis operators are labelled by irreps of G. These basis elements are guaranteed to be either purely unconfined or purely confined (i.e. they diagonalize the edge energy terms along the ribbon), rather than containing components of both types. To see which basis operators are confined, first note that $\partial(E)$ is a normal subgroup of G from the first Peiffer condition, Equation 2.4. Then given an irrep R of G labelling one of the basis operators, we can determine whether the corresponding excitations are confined by finding the restriction of R to the subgroup $\partial(E)$ (i.e. we only evaluate the irrep R on elements of $\partial(E)$, and treat this as a representation of the subgroup). When we restrict the irrep R to this subgroup, we obtain a generally reducible representation of this subgroup. When we decompose this representation into irreps of $\partial(E)$, by Clifford's theorem [86] the constituent irreps will all be related by conjugation. This means that if one of the irreps of $\partial(E)$ found by restricting R is the trivial irrep, then they all are (this result can also be obtained by Schur's Lemma in the case where $\partial(E)$ is in the centre of G). If R branches to the trivial irrep in this way, the electric excitation is not confined. On the other hand, if R does not branch to the trivial irrep, then the excitation is confined, with all of the edges along the ribbon being excited.

In addition, some of the E-valued loop-like excitations can be confined in a certain sense when E is non-Abelian, as described in Section 4.4. This confinement does not give an energy cost to growing the loop-like excitation, but instead it costs energy to move the excitation away from the start-point of the membrane. As we discussed in Section 4.4, it seems like the point-like charge carried by the loop excitation is confined, rather than the loop-like charge.

Chapter 6

Braiding

6.1 Introduction

One of the important features that we are interested in is the braiding of the various excitations that we find in the higher lattice gauge theory model. While we anticipate that most readers will have at least some familiarity with the concept of braiding in 2+1d, it may be useful to give an overview of braiding in 3+1d, particularly where loop-like excitations are involved. In 2+1d, topological phases support excitations with exchange statistics that generalize the familiar Fermi and Bose statistics [21, 22, 23], and which are described by the (coloured) braid group. On the other hand, in 3+1d the point-like particles can only have Fermi or Bose statistics [25, 29, 30]. However, the presence of loop-like excitations means that we can still have interesting braiding statistics. The motion of loops can be described by the (coloured) loop braid group [65, 66, 67] (considered under different names and contexts in papers such as Refs. [87, 88, 89]). The braid and loop braid groups are both examples of motion groups [90, 91], which describe the motions of arbitrary objects, up to homotopy [66] (formally the objects should return to the same positions, perhaps swapping positions). The loop braid group is generated by a few simple motions. Firstly consider braiding that involves only two loop-like excitations, and for simplicity suppose that those two excitations are stacked vertically, as shown in Figure 6.1. Then if we want to move

the lower loop up past the upper loop, there are several ways to do this. The first way, shown in Figure 6.2, is to simply move the lower loop around and past the upper one, so that neither loop passes through the other [65, 67]. We refer to this motion as a *permutation* move, because such moves generate the permutation group (also called the symmetric group) [66] that would describe the motions of point particles in 3+1d. The second way, shown in the left side of Figure 6.3, is to move the lower loop *through* the upper loop [65, 67]. We call this a braiding move and say that the lower loop has braided through the upper loop. The third way, shown in the right side of Figure 6.3, is to pull the lower loop *over* the upper loop, which we can also think of as the upper loop passing down through the lower loop (and so is the inverse of the previous motion). Another difference from point particles is that loops have an orientation, and so we must also allow a move that flips this orientation, as shown in Figure 6.4. Then any motion of the two loops can be performed by a series of such moves (we say that these generate the loop braid group for the two excitations) [65, 67]. More generally, we can have any number of loops, and the different motions of this set of loops can be performed using these pairwise moves. It is often useful to compare two motions that result in the same final position of all of the excitations. For example, we could compare the result of the permutation move in Figure 6.2 to the braiding move in Figure 6.3, or we could compare a motion that returns all particles to their initial positions to the trivial motion. Making a comparison between states where the excitations have the same final position is useful because it separates the topological content of the model from the geometric details.

In addition to loop-like particles, we have point-like particles in the lattice model. While braiding between two point-like particles in 3+1d is bosonic or fermionic [29, 30] as described earlier (and is exclusively bosonic in this model), the point-like excitations may braid non-trivially with the loop-like excitations of the model. In order to move a point-like particle past a loop excitation, we can either move the point-like excitation *around* the loop-like excitation (analogous to Figure 6.2 for two loops, if we replace

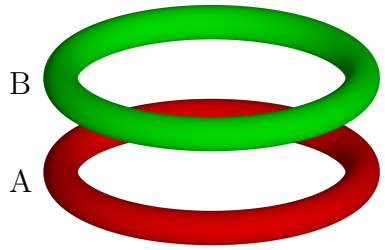

Figure 6.1: We first consider the motion of a pair of loops A and B, in the frame where the upper one (B) is held fixed. The motion of a set of loops can be described by such pairwise motions, together with flips of the two loops.

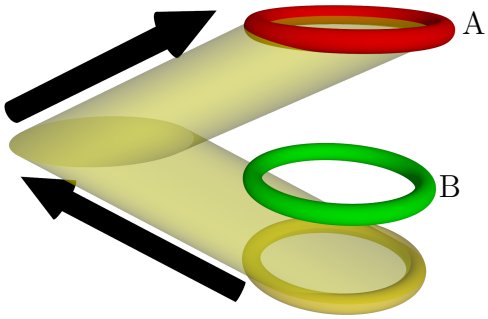

Figure 6.2: In the permutation move, we move one loop around the other so that neither loop passes through the other. The motion of the red loop is represented by the yellow sheet, with the initial position of the moving loop shown as a yellow torus.

---

the initially lower loop A with a point-like excitation), or we can move the point-like excitation through the loop-like one (analogous to the left-side of Figure 6.3, if we replace the lower loop A with a point-like excitation).

Having discussed the idea of braiding statistics in 3+1d, we will now explain how we can obtain these braiding relations in a commuting projector model using the ribbon and membrane operators. Ribbon operators can be thought of as creating a pair of excitations and then separating them along the ribbon. The ribbon operator therefore encodes the result of moving an excitation. Similarly a membrane operator can be thought of as nucleating a small loop and moving it into the final position of the excitation, as well as changing its size and shape. Therefore, we should be able to find the result of braiding by applying successive membrane operators. In particular, the braid-

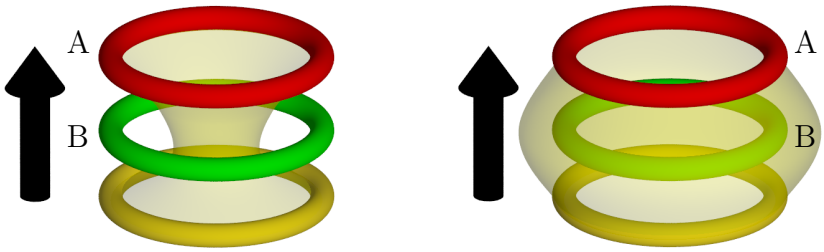

Figure 6.3: Schematic of the two braid moves. The sheet swept by the motion of the red loop is shown in yellow, with the initial position of the moving loop shown as a yellow torus. In the left image, the initially lower loop is passed through the upper one. In the right image, the initially lower loop passes over the upper loop, which is equivalent to the upper loop moving down through the lower one. The right image can therefore be thought of as the inverse of the left one (analogous to the two types of crossing in the normal braid group)

---

ing is related to the commutation relations of the membrane operators. For instance, consider the commutation relation shown in Figure 6.5, which we will see relates to loop-loop braiding. The images represent membrane operators, which are displaced horizontally to indicate an order of operators (although the membranes intersect in space). In the first line, we first act with an operator that produces a loop excitation (the green loop) and then act with an operator that creates another loop (shown in red) and moves it through that first loop. This performs a braiding move, specifically the one from the left side of Figure 6.3. In the second line of Figure 6.5, we first act with the operator that creates and moves the red loop through empty space, before creating the other loop excitation. In this case no braiding move occurs. Comparing these two lines (that is, working out the commutation relation for the two membrane operators) therefore lets us compare the situation with braiding to the one without, but where the excitations have the same final positions in each case. This latter point is important because ensuring that the excitations have the same final positions in the two cases isolates the effect of braiding from any other effects.

Having discussed the method for finding the braiding, we now discuss our results. Firstly, we note that any braiding of two point-like excitations (which we term permu-

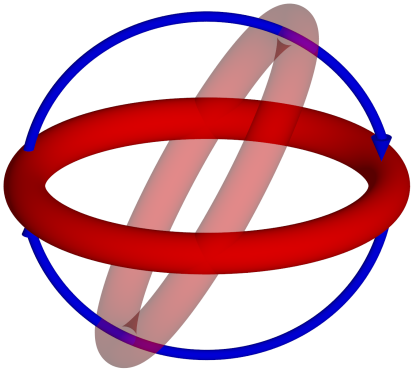

Figure 6.4: Unlike point particles, loop-like excitations have an orientation. We can reverse the orientation of a loop by rotating it by π radians out of its plane, which leaves the loop in the same position but flipped over [67].

---

tation) must be trivial in the 3+1d model. This is because in this model the permutation is implemented by ribbon operators that connect the initial and final positions of the excitations. In 3+1d, the ribbon operators that permute the two excitations do not need to intersect and so do not interact in any way (i.e. they commute), which leads to trivial (bosonic) exchange statistics. That is, the result of moving a point-like excitation around another one is the same as moving that excitation through empty space (and then producing the other point-like particle). In a similar way, exchange involving loop-like excitations moving past each-other (which we refer to as permutation), rather than through each-other (which we will refer to as loop-braiding or just braiding), is trivial for the same reason. That is, under permutation each excitation transforms just as if it were moving through empty space. This means that we only need to consider the braiding where we move a point-like or loop-like excitation through a loop-like excitation, as shown in Figure 3.3 for the point-like case and the left side of Figure 6.3 for the loop-like case. The result of this braiding move depends heavily on which special case from Table 2.1 that we consider, so we will examine these cases separately, starting with the $\triangleright$ trivial case.

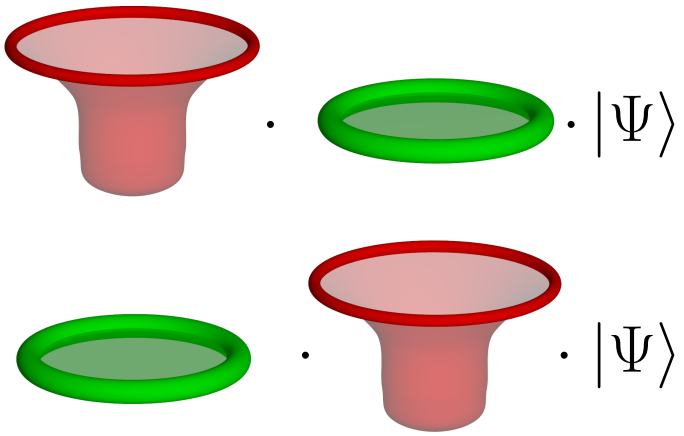

Figure 6.5: The commutation of operators used to calculate the braiding. The partially transparent surfaces indicate the membranes for the operators, while the opaque loops indicate the excited regions, which are the boundaries of the membranes.

6.2 Braiding in the $\triangleright$ Trivial Case

When $\triangleright$ is trivial, the braiding relations of the excitations are relatively simple. The elementary excitations that we have considered split into two sets, corresponding to the groups G and E, with non-trivial braiding only occurring within each set. For the group G, we have the electric and magnetic excitations. The electric excitations undergo non-trivial transformations when braided through the magnetic flux tubes, and the magnetic excitations can also braid non-trivially with each-other. For the group E, we have the point-like blob excitations and loop-like E-valued loop excitations, and the only non-trivial braiding is between these two types (the E-valued loop excitations do not braid non-trivially with each-other). Let us now consider how this non-trivial braiding arises in each case.

6.2.1 Flux-Charge Braiding

As we may expect from our discussion of gauge theory in Chapter 3, the electric and magnetic excitations can braid non-trivially. To see why this is the case, consider applying first a magnetic membrane operator $C^h(m)$ on membrane m to produce a loop-like flux tube at the boundary of m, and then applying an electric ribbon operator

on a path t that passes through that membrane in order to move an electric excitation through the flux tube. As described in Section 4.5, the magnetic membrane operator affects the edges that are cut by the dual membrane. Because the path t pierces the membrane, one of the edges on that path will have its label altered by the membrane operator. Denoting this edge by i, the label g_i of that edge becomes:

$$g_i \to \hat{g}(t_s)^{-1} h \hat{g}(t_s) g_i,$$

where t_s is the path from the start-point of the membrane to the crossing point, assuming that the edge i points away from the direct membrane. If we write the path t as $t_1 \cdot i \cdot t_2$, where t_1 and t_2 are the parts of t before and after edge i respectively, then this means that the overall path element satisfies

$$\begin{aligned}
\hat{g}(t) C^h(m) &= \hat{g}(t_1)\hat{g}_i\hat{g}(t_2) C^h(m) \\
&= C^h(m)\hat{g}(t_1)\hat{g}(t_s)^{-1} h \hat{g}(t_s)\hat{g}_i\hat{g}(t_2) \\
&= C^h(m)\hat{g}(t_1)\hat{g}(t_s)^{-1} h \hat{g}(t_s)\hat{g}(t_1)^{-1}\hat{g}(t_1)\hat{g}_i\hat{g}(t_2) \\
&= C^h(m)(\hat{g}(t_s)\hat{g}(t_1)^{-1})^{-1} h (\hat{g}(t_s)\hat{g}(t_1)^{-1})\hat{g}(t),
\end{aligned}$$

where we inserted the identity in the form of $\hat{g}(t_1)^{-1}\hat{g}(t_1)$ in order to reconstitute the path element $\hat{g}(t)$ on the right. We see that the path element gains a factor of

$$(\hat{g}(t_s)\hat{g}(t_1)^{-1})^{-1} h (\hat{g}(t_s)\hat{g}(t_1)^{-1})$$

under the braiding. This is generally an operator, so the braiding relation is not well-defined. However, there are two important cases where the factor becomes a constant instead. The first of these is when G is Abelian, in which case the relation simply becomes

$$\hat{g}(t) C^h(m) = C^h(m) h \hat{g}(t),$$

so that the path element gains a factor of h. This means that for a basis electric ribbon operator labelled by an irrep R of G, we have

$$\begin{aligned}\sum_{g\in G} R(g)\delta(g,\hat{g}(t))C^h(m) &= C^h(m)\sum_{g\in G} R(g)\delta(g,h\hat{g}(t)) \\ &= C^h(m)\sum_{g\in G} R(g)\delta(h^{-1}g,\hat{g}(t)) \\ &= C^h(m)\sum_{g'=h^{-1}g\in G} R(hg')\delta(g',\hat{g}(t)) \\ &= C^h(m)R(h)\sum_{g'\in G} R(g')\delta(g',\hat{g}(t)),\end{aligned}$$

so the result of the braiding is the accumulation of a phase $R(h)$.

The other important case where the braiding is well-defined is the "same start-point" case. When the path t for the electric ribbon operator starts at the start-point of the magnetic membrane, then the path sections t_s and t_1 start and end at the same points as each-other. Provided that these path sections can be deformed into one-another without crossing over any excitations, the fake-flatness condition imposed by the plaquette energy terms ensures that $\hat{g}(t_s) = \hat{g}(t_1)$ up to a potential factor of $\partial(e)$ for some $e \in E$. Such factors of $\partial(e)$ do not affect

$$(\hat{g}(t_s)\hat{g}(t_1)^{-1})^{-1}h(\hat{g}(t_s)\hat{g}(t_1)^{-1})$$

because $\partial(e)$ is in the centre of G, so the factor of $\partial(e)$ and $\partial(e)^{-1}$ from $\hat{g}(t_s)$ and $\hat{g}(t_s)^{-1}$ cancel. Therefore,

$$\begin{aligned}(\hat{g}(t_s)\hat{g}(t_1)^{-1})^{-1}h(\hat{g}(t_s)\hat{g}(t_1)^{-1})\hat{g}(t) &= \partial(e)^{-1}h\partial(e)\hat{g}(t) \\ &= h\hat{g}(t). \qquad (6.1)\end{aligned}$$

Then when we consider a basis electric ribbon operator labelled by a representation R and matrix indices a and b (note that when G is non-Abelian the irreps are generally

higher-dimensional), we find the braiding relation:

Flux-Charge Braiding Relation:

$$\sum_{g\in G}[D^R(g)]_{ab}\delta(g,\hat{g}(t))C^h(m)\,|GS\rangle$$
$$= C^h(m)\sum_{c=1}^{|R|}[D^R(h)]_{ac}\sum_{g'\in G}[D^R(g')]_{cb}\delta(g',\hat{g}(t))\,|GS\rangle\,. \qquad (6.2)$$

We see from Equation 6.2 that the braiding mixes electric ribbon operators labelled by different matrix indices but the same representation. The fact that the representation is left invariant suggests that the representations label the purely electric topological sectors, given that braiding cannot mix different sectors. Note that we have implicitly assumed a certain relative orientation for the path t and membrane m in finding Equation 6.2, and reversing this orientation would lead to h being replaced by h^{-1} (it may also appear that we made an assumption about the orientation of the particular edge i, but this is not necessary and reversing this edge does not affect the overall result).

The importance of the start-points of the membrane and ribbon operators when it comes to braiding can be interpreted in terms of gauge theory using the ideas we described in Chapter 3. The start-point of the magnetic membrane can be seen as a unique point in that the flux tube produced by the magnetic membrane operator has a definite flux label with respect to this point even within the conjugacy class. Moving this start-point can change the flux label that we would assign to the flux tube (see e.g. Ref. [32]). Suppose that the flux label measured with respect to a particular start-point is h. Then if we measure the flux label of the flux tube starting from a different point, the result is related to h by conjugation by a path element for a path between the two start-points. In our model this path element is not usually well-defined, because the energy eigenstates are usually linear combinations of states with different group elements assigned to the path. We say that the path element is generally operator valued. Therefore, the flux tube does not generally have definite

flux with respect to points other than the start-point (with an exception if the flux label is in the centre of G). This non-definite flux is then reflected in the flux-charge braiding.

A second interpretation of the start-point dependent braiding comes from anyon theory. In 2+1d anyon theory we only have definite braiding when the two excitations undergoing braiding have a definite fusion channel, meaning that the two excitations have a well-defined total topological charge (see e.g. Refs. [24, 26] for more information about fusion). In this model, we only have definite fusion in the case where the membrane operators share a start-point. Then, as in the 2+1d case, we may expect that we only have definite braiding when the fusion channel is fixed: that is, when the operators involved have a common start-point.

6.2.2 Flux-Flux Braiding

Next we consider the braiding between two magnetic excitations, i.e. between two flux tubes. Just as with the flux-charge braiding described previously, the braiding relation is well-defined when we use two magnetic membrane operators with the same start-point. In this case, the flux label of each excitation is defined with respect to the same start-point (i.e. in the same gauge), so we expect a well-defined braiding relation.

We consider the braiding operation where one flux tube (which we will call the inner loop) is passed through another loop (which we call the outer loop). Considering Figure 6.5, which demonstrates the appropriate commutation relation of membrane operators, the inner loop is the red one and the outer loop is the green one. A more detailed depiction of the membrane operators is shown in the left side of Figure 6.6. Note that we choose the start-points of the two membrane to be the same. Let the inner (red) membrane operator be $C^k(m_1)$ and the outer (green) membrane operator be $C^h(m_2)$, as shown in Figure 6.6. In order to calculate the braiding relation, the

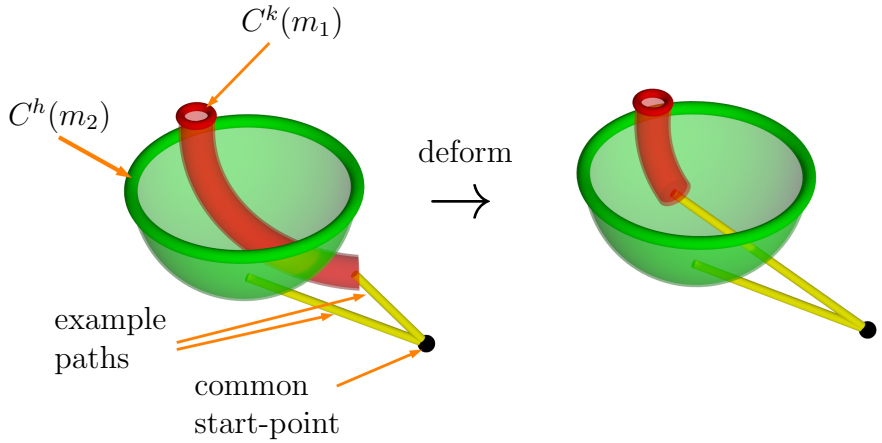

Figure 6.6: To calculate the braiding of two loops, we first apply a magnetic membrane operator $C^h(m_2)$, then apply a membrane operator $C^k(m_1)$ which intersects with the first membrane and pushes a loop excitation through that first membrane. We choose the start-points of these two membranes to be the same. Example paths from the common start-point to the two membranes are shown as yellow cylinders. We can use the topological property of the membrane operators to deform the inner (red) membrane, m_1, and pull it through the green membrane m_2 to obtain the image on the right-hand side. However, when we do so we must leave the start-point fixed. Therefore, the paths from the start-point to each point on the membrane m_1, such as the example path shown here, must pass through m_2. This leads to a non-trivial braiding relation in general.

---

simplest method is to use the topological nature of the membrane operators, which allows us to deform the membranes, pulling the inner membrane m_1 through m_2, to give us the situation shown in the right side of Figure 6.6. When we do so, the action of the membrane operators on the ground state is unchanged, provided that we keep fixed the position of any excitations, including the potentially excited start-point of the membrane operators. Recall that the action of a magnetic membrane operator on an edge i cut by its dual membrane depends on a path from the start-point to that edge:

$$C^k(m_1) : g_i = \begin{cases} g(s.p - v_i)^{-1} k g(s.p - v_i) g_i & \text{if } i \text{ points away from} \\ & \text{the direct membrane} \\ g_i g(s.p - v_i)^{-1} k^{-1} g(s.p - v_i) & \text{if } i \text{ points towards the} \\ & \text{direct membrane,} \end{cases} \tag{6.3}$$

where v_i is the end of i that lies on the direct membrane of m_1. Because the start-point is fixed and lies outside the outer membrane m_2, whereas the deformed membrane m_1 lies on the other side of the membrane, the paths from the start-point to m_1 all pass through m_2. As we described in Section 6.2.1 when considering the flux-charge braiding, paths that pass through a magnetic membrane operator $C^h(m_2)$ have their labels multiplied by h (for the same start-point case at least). This means that if we act first with $C^h(m_2)$, then with $C^k(m_1)$, the factor $g(s.p - v_i)$ in Equation 6.3 should be replaced with $hg(s.p - v_i)$. This means that

$$C^k(m_1)C^h(m_2) : g_i = \begin{cases} g(s.p - v_i)^{-1}h^{-1}khg(s.p - v_i)g_i & \text{if } i \text{ points away from the direct membrane} \\ g_i g(s.p - v_i)^{-1}h^{-1}k^{-1}jg(s.p - v_i) & \text{if } i \text{ points towards the direct membrane} \end{cases}$$
$$= C^h(m_2)C^{h^{-1}kh}(m_1) : g_i. \tag{6.4}$$

We see that under the commutation relation, and therefore under braiding, the label k of the inner loop becomes $h^{-1}kh$. This gives us the braiding relation:

Flux-Flux Braiding Relation:

$$C^k(m_1)C^h(m_2)\,|GS\rangle = C^h(m_2)C^{h^{-1}kh}(m_1)\,|GS\rangle \tag{6.5}$$

Therefore, we see that the label of one of the flux tubes is conjugated by braiding. As we expect, the conjugacy class of the flux is invariant under braiding, but the flux element within the conjugacy class can be changed. This means that if we consider a linear combination of magnetic membrane operators with label in that conjugacy class, such as $\sum_{x\in[k]} \alpha_x C^x(m_1)$, the membrane operator is unaffected by the braiding if α_x is the same for each element in the conjugacy class (because the braiding only permutes the elements within the conjugacy class). Note that this condition, that the coefficients are a function of conjugacy class, is the same requirement for the start-point of the membrane operator to be unexcited. This means that the loop transforms

trivially under the flux-flux braiding if it has an unexcited start-point. This is because if the start-point is excited the magnetic excitation carries a point-like charge, which enables it to braid non-trivially with other magnetic excitations when passed through them, but if the start-point is unexcited it does not carry such a charge. Note that the same condition does not hold for the outer membrane operator, which can affect the inner membrane operator even if it does not have an excited start-point. The point-like charge is important for the inner loop because we can shrink the inner loop to a point before braiding it without affecting the braiding relation, whereas the loop-like character of the outer loop is essential for the braiding.

It is important to note that the precise form of the braiding relation depends on the orientation of the loops involved. Flipping a magnetic excitation is equivalent to changing its label from h to h^{-1}. Therefore, if we were to reverse the orientation of the outer membrane from our earlier calculation, then the label of the inner membrane would change from k to hkh^{-1} under braiding rather than changing to $h^{-1}kh$. Flipping the orientation of the inner one does not change the expression, because the transformation is the same when we invert both sides: $k^{-1} \to h^{-1}k^{-1}h \implies k \to h^{-1}kh$. Therefore, if the orientation of both loops is reversed, the braiding transformation is

$$k \to hkh^{-1}. \tag{6.6}$$

6.2.3 Loop-Blob Braiding

So far we have considered the excitations that are described by the group G. The final non-trivial braiding is between the two types of excitation that are associated to the group E, the E-valued loops and blob excitations. In this case it is easy to find the braiding relations by looking at the effect of the blob ribbon operator on the surface measured by the E-valued membrane operator. We consider a situation where the blob excitation passes through the loop excitation. To implement this situation on our lattice, we apply both an E-valued membrane operator and a blob ribbon

operator whose path intersects with that membrane, as shown in Figure 6.7. In order to compute the braiding relation, we compare the situation where we first create the loop and then push the blob through, thus performing the braiding move, with the one where we push the blob through empty space before producing the loop. The relevant commutation relation is shown in Figure 6.8.

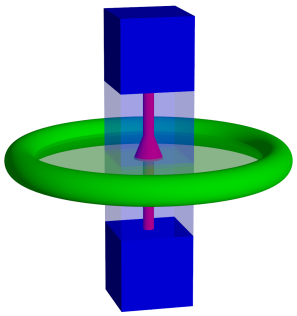

Figure 6.7: Schematic of blob-loop braiding. The blue cubes represent the blob excitations at the ends of the ribbon operator (whose ribbon is represented by the translucent blue cuboid and red arrow). The ribbon operator moves one of the blob excitations through the green loop-like excitation produced by an E-valued membrane operator on the translucent green membrane.

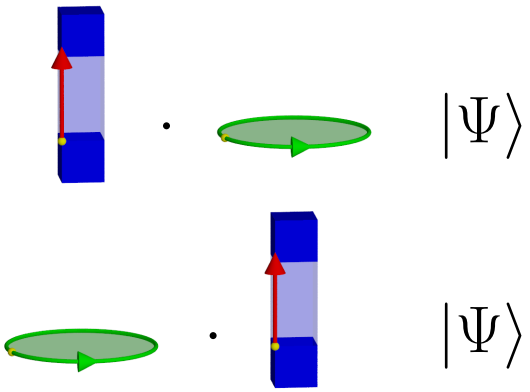

Figure 6.8: In order to determine the blob-loop braiding relations, we compare the situation shown in the top line, where we first apply an E-valued membrane operator and then a blob ribbon operator that intersects with that membrane, to the situation shown in the bottom line, where we apply the operators in the opposite order. Here $|\psi\rangle$ is any state with no other excitations near the support of the two operators (e.g. a ground state).

As we saw in Section 4.3, the blob ribbon operator with label e multiplies the labels of

the plaquettes pierced by its path by e or e^{-1}, depending on the relative orientation of the plaquette and ribbon. Recall that the E-valued membrane operator measures the surface element of the membrane on which it is applied. This surface element, $\hat{e}(m)$, is a product of the elements of individual plaquettes: $\hat{e}(m) = \prod_{\text{plaquettes } p \in m} e_p^{\sigma_p}$, where σ_p is $+1$ or -1 to account for the relative orientation of the surface and plaquette ($+1$ if they align, -1 if they anti-align). Then because the blob ribbon operator intersects with this membrane, it will affect the label of one of the plaquettes in this product. If the blob ribbon operator pierces the membrane m through a plaquette q, then the label e_q of that plaquette is multiplied by e or e^{-1}, depending on the relative orientation of the plaquette and the ribbon. This in turn means that the contribution $e_q^{\sigma_q}$ of the plaquette to the surface m will be multiplied by e or e^{-1}, this time depending on the relative orientation of the membrane and the ribbon and independent of the orientation of the plaquette. If the orientation of the membrane matches the orientation of the blob ribbon operator, $e_q^{\sigma_q}$ will be multiplied by e^{-1}. This indicates that $\hat{e}(m)B^e(t) = B^e(t)e^{-1}\hat{e}(m)$ in this case. Then, considering the basis operator for our space of E-valued membrane operators labelled by an irrep γ of E, the commutation relation with the blob ribbon operator $B^e(t)$ is given by

$$\begin{aligned} B^e(t) \sum_{e' \in E} \gamma(e')\delta(e', \hat{e}(m)) \left|GS\right\rangle &= \sum_{e' \in E} \gamma(e')\delta(e', e\hat{e}(m)) B^e(t) \left|GS\right\rangle \\ &= \sum_{e' \in E} \gamma(e')\delta(e^{-1}e', \hat{e}(m)) B^e(t) \left|GS\right\rangle \\ &= \sum_{e'' = e^{-1}e' \in E} \gamma(ee'')\delta(e'', \hat{e}(m)) B^e(t) \left|GS\right\rangle \\ &= \gamma(e) \sum_{e'' \in E} \gamma(e'')\delta(e'', \hat{e}(m)) B^e(t) \left|GS\right\rangle . \end{aligned}$$

Relabelling the dummy index e'' to e' for clarity, this gives us the braiding relation:

Loop-Blob Braiding Relation:

$$B^e(t) \sum_{e' \in E} \gamma(e')\delta(e', \hat{e}(m)) \left|GS\right\rangle = \gamma(e) \sum_{e' \in E} \gamma(e')\delta(e'', \hat{e}(m)) B^e(t) \left|GS\right\rangle \quad (6.7)$$

We therefore see that the braiding of our two excitations results in accumulating a phase of $\gamma(e)$. A similar argument holds in the case where the membrane operator and blob ribbon operator are anti-aligned, except that we should replace e by its inverse.

It is worth noting that the blob excitations with label not in the kernel of ∂ (that is, the confined blob excitations) braid non-trivially with the condensed E-valued loop excitations (those with trivial representation of the kernel), while those with label in the kernel braid trivially with them. This is because the condensed E-valued loop excitations have trivial representation of the kernel: $\gamma(e_K) = 1$ for e_K in the kernel. Therefore, the phase gained is 1 when a condensed loop braids with an unconfined blob excitation (which carries a label in the kernel). This matches our expectation that only the confined excitations can braid non-trivially with the condensed excitations.

6.2.4 Summary of Braiding When $\triangleright$ Is Trivial

For convenience, we summarize the excitations that braid non-trivially with each-other in Table 6.1. We can see that the excitations split into two sets. The electric and magnetic excitations (the excitations corresponding to the group G) braid non-trivially with each-other and the blob and E-valued loop excitations (the excitations corresponding to E) braid non-trivially with each-other, but there is no non-trivial braiding between the two sets.

Non-Trivial Braiding?	Electric	Magnetic flux	Blob	E-valued loop
Electric	✗	✓	✗	✗
Magnetic flux	✓	✓	✗	✗
Blob	✗	✗	✗	✓
E-valued loop	✗	✗	✓	✗

Table 6.1: A summary of which excitations braid non-trivially in Case 1, where $\triangleright$ is trivial. A tick indicates that at least some of the excitations of each type braid non-trivially with each-other, while a cross indicates that there is no non-trivial braiding between the two types.

6.3 Braiding in the Fake-Flat Case

Next we discuss the braiding relations in the special case in which we restrict to fake-flat configurations (Case 3 from Table 2.1). The fact that we are unable to include the magnetic excitations (because they violate fake-flatness) means that the braiding relations are rather simple. Indeed, the only remaining non-trivial braiding is between the E-valued loop excitations and the blob excitations. However, the signatures of the magnetic excitations are still present in the ground states of manifolds with non-contractible cycles, which can have labels outside of $\partial(E)$ (just like closed loops enclosing magnetic excitations). Before we discuss the braiding proper, we will briefly describe how the excitations transform as they are moved around such non-contractible cycles.

6.3.1 Moving Excitations Around Non-Contractible Cycles

The first type of excitation that we wish to consider moving around a non-contractible cycle is the electric excitation. If we compare an electric ribbon operator applied on a path s to one that is applied on the path $s \cdot t$ obtained by concatenating the original path with a non-contractible closed path t, we have

$$S^{R,a,b}(s \cdot t) = \sum_{c=1}^{|R|} [D^R(\hat{g}(t))]_{cb} S^{R,a,c}(s),$$

where $S^{R,a,b}(s \cdot t)$ is the electric ribbon operator labelled by irrep R of G and matrix indices a and b. We see that there is mixing between different electric ribbon operators labelled by the same irrep R, with this mixing controlled by the matrix $D^R(\hat{g}(t))$ representing the path element $\hat{g}(t)$ in irrep R. The path element $\hat{g}(t)$ is an operator, and the ground states are not typically eigenstates of this operator even for closed paths t.

In a similar way, we can find how the E-valued loop-like excitations transform as they are moved around a non-contractible cycle. In order to do so, we compare E-

valued membrane operators applied on two membranes m and m' which are the same except that m' is whiskered around the non-contractible cycle t. This is because the membrane operators are topological, and so a membrane travelling around the cycle can be deformed by shrinking the section of the membrane around the cycle down to nothing, so that only a whiskering path t remains, as shown in Figure 6.9. Then the E-valued membrane operator $L^{\mu,a,b}(m')$ applied on this whiskered membrane is given by

$$L^{\mu,a,b}(m') = \sum_{e \in E} [D^{\mu}(e)]_{ab} \delta(e, \hat{e}(m')),$$

where μ is the irrep of E labelling the membrane operator (and a and b are the matrix indices labelling the operator). The surface element $\hat{e}(m')$ can be written in terms of the surface element of the unwhiskered membrane m using the rules for whiskering surfaces given in Ref. [58]. We have

$$\hat{e}(m') = \hat{g}(t) \triangleright \hat{e}(m),$$

where t is the closed cycle, which is also the path $(s.p(m') - s.p(m))$ between the start-points of the two membranes. Then we have

$$L^{\mu,a,b}(m') = \sum_{e \in E} [D^{\mu}(e)]_{ab} \delta(e, \hat{e}(m')) = \sum_{e' = \hat{g}(t)^{-1} \triangleright e \in E} [D^{\mu}(\hat{g}(t) \triangleright e')]_{ab} \delta(e', \hat{e}(m)).$$

The matrix $D^{\mu}(\hat{g}(t) \triangleright e')$ can be considered as the matrix representation for element e' in a new irrep, $\hat{g}(t) \triangleright \mu$. Therefore, we see that moving the E-valued loop excitation around the path t mixes the irreps related by this $\triangleright$ action. We say that irreps related by the action of $g\triangleright$ for some $g \in G$ belong to the same $\triangleright$-Rep class of irreps.

The final type of excitation to consider passing around a non-contractible cycle is the blob excitation. Recall from Section 4.3 that the action of a blob ribbon operator $B^{e}(r)$ on a plaquette p pierced by the ribbon r is (choosing the plaquette to be aligned

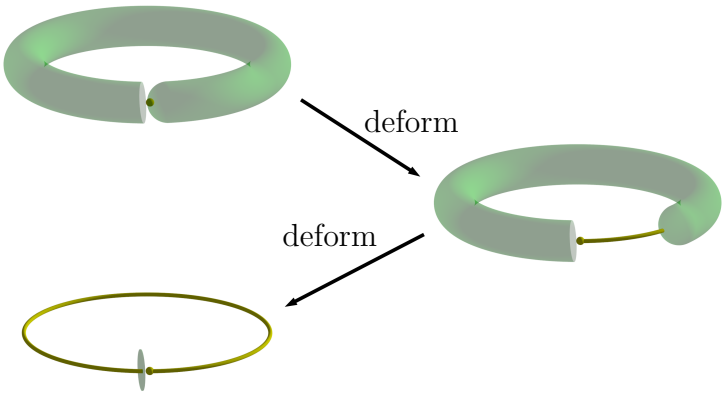

Figure 6.9: Given an E-valued membrane operator (green) wrapping around a closed cycle, we can deform the section wrapping around the cycle and shrink it down to nothing. This just leaves a whiskering string (yellow) connecting the start-point (yellow sphere) to the small part of the membrane remaining (the green disk in the final image).

---

with r for simplicity)

$$B^e(r) : e_p = e_p[g(s.p(r) - v_0(p))^{-1} \triangleright e^{-1}].$$

Taking the path $(s.p(r) - v_0(p))$ to be $t \cdot s$, where t is a closed non-contractible cycle, we can write this action as

$$\begin{aligned} B^e(r) : e_p &= e_p[g(t \cdot s)^{-1} \triangleright e^{-1}] \\ &= e_p[(g(s)^{-1} g(t)^{-1}) \triangleright e^{-1}] \\ &= e_p[g(s)^{-1} \triangleright (g(t)^{-1} \triangleright e^{-1})], \end{aligned}$$

which is the same as the action of a blob ribbon operator that does not wrap around the cycle (so that the path $(s.p(r) - v_0(p))$ is just s, rather than $t \cdot s$), except with e replaced by $g(t)^{-1} \triangleright e$. Note that here we have taken the direct path of the ribbon operator to wrap around the non-contractible cycle, but not the dual path. If we also let the dual path wrap around the cycle, then the plaquettes pierced by s obtain two factors acting on the plaquette label, one from the ribbon before it wraps the cycle (corresponding to the label e) and one from the ribbon after it wraps the cycle

(corresponding to the label $g(t)^{-1} \triangleright e$ as above).

6.3.2 Loop-Blob Braiding

Compared to the loop-blob braiding that we saw in Section 6.2.3, the loop-blob braiding in this special case is slightly more complicated. This is because the action of the ribbon and membrane operators involved now depends on the values of various paths on the lattice. For example, as we saw in Section 4.3, the blob ribbon operator multiplies plaquette elements by a group element $\hat{g}(s.p - v_0(p))^{-1} \triangleright e$ which depends on the label of a path. This label is really an operator, because the value of the path label depends on what state we are acting on. In particular, the ground state does not have a definite value of this label, instead being made up of a linear combination of states with different path labels. Because of this, we may expect that the braiding does not generally give us a definite result and that the braiding relation may depend on such operator-valued labels. However, as with the braiding of the flux tubes that we saw in Section 6.2.2, the braiding relations are simple for particular cases where the start-points of the operators match. Therefore, we are most interested in these same start-point commutation relations. To understand these, it will be useful to first discuss the interpretation of the blob and loop excitations in 2-gauge theory.

In this model, the blob excitations are associated to non-trivial 2-flux on a surface enclosing the blob excitation. The boundary of the excited blob is itself a surface with non-trivial 2-flux, because an excited blob by definition has a non-trivial surface label on its boundary. When we braid the blob excitation through an E-valued loop excitation, we measure the value of this 2-flux. However, the blob ribbon operator produces excitations that have definite 2-flux only with respect to the start-point of the blob ribbon operator. Similarly, the E-valued loop measures 2-flux with respect to the start-point of the membrane operator that creates the loop excitation. Therefore, we expect a definite braiding relation when the start-point for our blob ribbon operator matches the start-point for our E-valued loop membrane operator. Note that when

$\triangleright$ is trivial, the start-points lose meaning (we do not need a direct path for the blob ribbon operators and the surface does not need a base-point for the E-valued loop) and the loop can be nucleated at any point before being passed over the blob excitation (rather than at the specified start-point of the blob ribbon operator). This is why the braiding is simple when $\triangleright$ is trivial (as discussed in Section 6.2.3) and it is not necessary to fix the positions of the start-points in that case.

In order to calculate the braiding relation, we once again relate the braiding process to a commutation relation between ribbon and membrane operators. By first producing the loop-like excitation with a membrane operator and then moving the blob excitation through the loop with a ribbon operator, we perform the braiding move. We compare this to the case where we first move the blob excitation and then produce the loop excitation. This means that the relevant commutation relation to calculate this braiding relation is the one shown in Figure 6.8 in Section 6.2.3.

The result of this same-site braiding, where we pass a ribbon operator $B^e(t)$ through an E-valued membrane operator $\delta(e_m, \hat{e}(m))$, is then

$$B^e(t)\delta(e_m, \hat{e}(m)) \left|GS\right\rangle = \delta(e_m e^{-1}, \hat{e}(m))B^e(t) \left|GS\right\rangle , \tag{6.8}$$

as we prove in the Supplementary Material of Ref. [2]. The operators $\delta(e_m, \hat{e}(m))$ for each label $e_m \in E$ form a basis for our space of E-valued membrane operators, but we want to consider the commutation of one of the basis operators labelled by an irrep of E instead. For such a basis operator, labelled by irrep α and matrix indices a and b, we have that

$$\begin{aligned} B^e(t) \sum_{e_m \in E} [D^\alpha(e_m)]_{ab} \delta(\hat{e}(m), e_m) \left|GS\right\rangle \\ = \sum_{c=1}^{|\alpha|} [D^\alpha(e)]_{cb} \sum_{e' \in E} [D^\alpha(e')]_{ac} \delta(\hat{e}(m), e') B^e(t) \left|GS\right\rangle . \end{aligned} \tag{6.9}$$

If α is a 1D representation, the braiding therefore results in the accumulation of a

phase of $\alpha(e)$, just like in the $\triangleright$-trivial case (although in the $\triangleright$-trivial case, we did not require the ribbon and membrane operators to have the same start-point). If α is higher-dimensional, then it would seem that there is mixing between operators labelled by different matrix indices. However, as we mentioned in Section 5.3, in order to ensure fake-flatness we only allow blob ribbon operators with label in the kernel of ∂. Elements in the kernel of ∂ are also in the centre of E. This follows from the second Peiffer condition, which states that $\partial(e) \triangleright f = efe^{-1}$. When e is in the kernel of ∂, $\partial(e) \triangleright f$ is just f, so e must commute with f. Because f is an arbitrary element of E, this means that e is in the centre of E. The matrix representation of an element of the centre is a scalar multiple of the identity from Schur's Lemma, so we can write $[D^\alpha(e)]_{cb} = \delta_{cb}[D^\alpha(e)]_{11}$ (where the index 1 could be replaced with any index). This means that the braiding relation Equation 6.9 simplifies to

Loop-Blob Braiding Relation:

$$B^e(t) \sum_{e_m \in E} [D^\alpha(e_m)]_{ab} \delta(\hat{e}(m), e_m) \, |GS\rangle$$
$$= [D^\alpha(e)]_{11} \sum_{e' \in E} [D^\alpha(e')]_{ab} \delta(\hat{e}(m), e') B^e(t) \, |GS\rangle \,, \quad (6.10)$$

so again we only accumulate a phase $[D^\alpha(e)]_{11}$ (this matrix element must be a phase because the matrix is diagonal and unitary). We see that the braiding relation between the E-valued loops and blob excitations is similar to the result we found for the braiding of magnetic fluxes and charges that we discussed in Section 6.2.1, except that in this case the irrep labels the loop-like, rather than point-like, excitation.

6.3.3 Summary of Braiding in the Fake-Flat Case

In Table 6.2, we summarize the braiding in the fake-flat case by indicating which excitations can braid non-trivially. Note that when we enforce fake-flatness on the level of the Hilbert space, there are no magnetic excitations, although we can still have non-trivial flux labels around non-contractible loops on manifolds that are not simply connected (e.g. the 3-torus).

Non-trivial Braiding?	Electric	Blob	E-valued loop	Around handle
Electric	✗	✗	✗	✓
Blob	✗	✗	✓	✓
E-valued loop	✗	✓	✗	✓

Table 6.2: A summary of the non-trivial braiding in the fake-flat case. In the fake-flat case, there are no magnetic excitations and the non-trivial braiding between excitations only involves the blob excitations and E-valued loops. However, there are non-trivial results from moving excitations around handles (non-contractible cycles), which can support non-trivial 1-flux.

6.4 Braiding in the Case Where $\partial \to$ Centre(G) and E Is Abelian

Finally, we consider the braiding in the case where ∂ maps to the centre of G and E is Abelian (Case 2 from Table 2.1). Any braiding not involving the magnetic excitations is the same as in the fake-flat case described in Section 6.3. Namely, there is non-trivial braiding between the blob excitations and the E-valued loops, with the same start-point braiding resulting in an accumulation of phase. The result is a phase, rather than the more general transformation given in Equation 6.9, because the irreps of E are 1D when E is Abelian.

Unlike for the fake-flat case, we can find the magnetic excitations and so describe their braiding relations. However, rather than using the magnetic membrane operator directly, it is convenient when considering braiding to combine the magnetic membrane operator with an E-valued membrane operator. We multiply the magnetic membrane operator by an E-valued membrane operator such as $\delta(e_m, \hat{e}(m))$, acting before the magnetic membrane operator. That is, we construct membrane operators of the form $C_T^h(m)\delta(e_m, \hat{e}(m))$, which we denote by $C_T^{h,e_m}(m)$. We note that combining the magnetic membrane operator with this E-valued membrane operator in this way does not excite regions of the lattice not already excited by the magnetic membrane operator, because both membrane operators only cause excitations near the boundary of the membrane, and possibly at blob 0 and the start-point of the membrane. We also

note that there are some important details to consider in order to combine the two membrane operators properly, which we explain in Ref. [2].

The reason that this combination of operators is useful is that when $\triangleright$ is non-trivial the loop excitation labelled by h and e_m acquires a 2-flux of

$$\tilde{e}_m = e_m[h^{-1} \triangleright e_m^{-1}], \tag{6.11}$$

as we show in Ref. [2]. If we did not measure the surface label $\hat{e}(m)$ with an E-valued membrane operator, this 2-flux would not be well-defined, which would lead to ill-defined braiding relations. Due to the fact that the excitations produced by this combined membrane operator carry both ordinary magnetic flux and this 2-flux, we call the combined membrane operator a "higher-flux membrane operator" and call the excitations higher-flux excitations. If we wish to instead consider the original magnetic membrane operator (without the attached E-valued membrane operator), we can simply sum over each value of e_m, because this gives us a complete sum of projectors $\delta(e_m, \hat{e}(m))$. That is $\sum_{e_m \in E} \delta(e_m, \hat{e}(m)) = 1$ and so

$$\sum_{e_m \in E} C_T^{h,e_m}(m) = C_T^h(m). \tag{6.12}$$

Having constructed this higher-flux membrane operator, we can now use it to find the braiding relations involving the higher-flux excitations. Because the action of the higher-flux membrane operator on the edges is the same as that of the magnetic membrane operator from the $\triangleright$ trivial case, the braiding relation between the magnetic and electric excitations is the same as in that case (which is described in Section 6.2.1). However, as we will see shortly, the braiding between the higher-flux loop excitation and the other excitations is significantly altered. In particular, because the higher-flux excitations can carry a non-trivial 2-flux, we expect non-trivial braiding relations with the E-valued loops, which measure 2-flux.

6.4.1 Braiding of the Higher-Flux Excitations With Blob Excitations

The first braiding relation we examine is between the higher-flux excitations and the blob excitations. We consider a blob ribbon operator $B^e(t)$, applied on a ribbon t, piercing the membrane of a higher-flux membrane operator $C_T^{h,e_m}(m)$, applied on a membrane m. The ribbon t intersects the membrane m at a plaquette q, as shown in Figure 6.10. Note that the orientation of the operators is significant. We first look at the case where the blob ribbon operator pierces the direct membrane of the magnetic membrane operator before the dual membrane. If the membrane is oriented downwards, as in Figure 6.10 (note that the point-like excitation is above the loop) then the ribbon is oriented upwards (at the point of intersection at least).

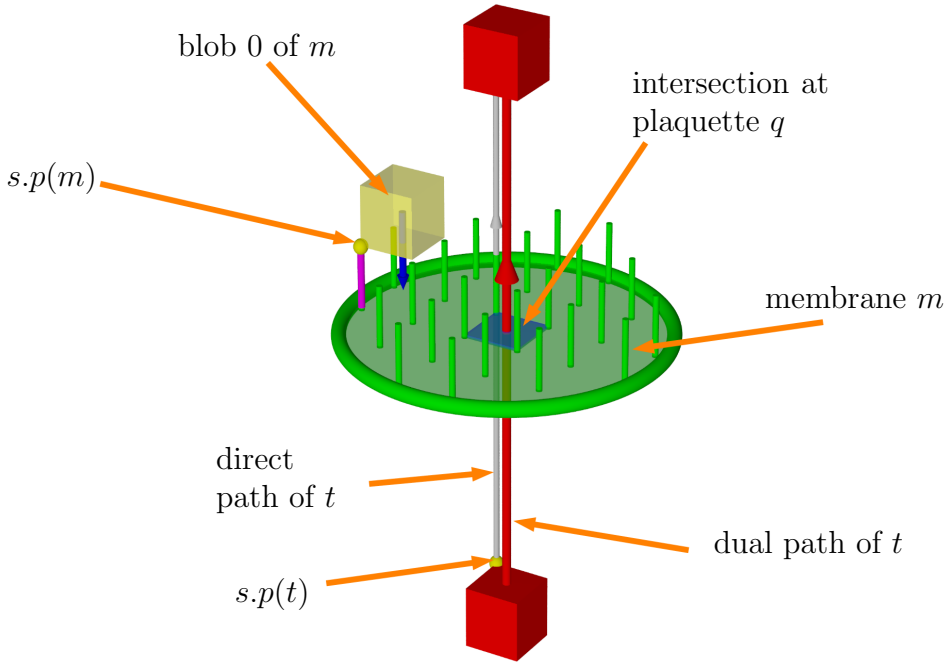

Figure 6.10: We consider a blob ribbon operator $B^e(t)$ (between the two red blobs) that passes through a higher-flux membrane operator, $C_T^{h,e_m}(m)$ (where m is shown in green). The ribbon t pierces the membrane m through a plaquette q (blue square).

We have seen in previous cases that braiding is frequently well-defined only when the start-points of the membrane and ribbon operators match. However, we shall

first examine the general case where the start-points are arbitrary. As usual, we can relate the braiding relation to a commutation relation between the two operators. We compare the case where the magnetic membrane is produced first, and then the blob excitation moved through it, to the reverse case. We find that, as we prove in the Supplementary Material of Ref. [2],

$$\begin{aligned} &B^e(t)C_T^{h,e_m}(m)\,|GS\rangle \\ &\quad = C_T^{h,e_m[\hat{g}(s.p(m)-s.p(t))\triangleright e]}(m)B^e(t_1')B^{(\hat{g}(s.p(t)-s.p(m))h^{-1}\hat{g}(s.p(t)-s.p(m))^{-1})\triangleright e}(t_2')\,|GS\rangle\,. \end{aligned} \tag{6.13}$$

In this expression, we note that the original ribbon operator is split into two parts, on ribbons t_1' and t_2', which transform differently under the braiding. Here the dual path of t_1' starts at the original origin of ribbon t and ends at blob 0 of the membrane m (corresponding to the part of the ribbon before the intersection with the membrane, except that it is diverted to end at blob 0 of m), while the dual path of t_2' starts at blob 0 and ends at the original end of ribbon t (corresponding to the part of the ribbon after the intersection), as shown in Figure 6.11 (although note that the start-point of each part of the ribbon, which is the beginning of the direct path, is still the original start-point). We therefore see that, under commutation, the ribbon t is diverted to pass through blob 0 of the membrane, and after passing through this blob it changes label from e to $(\hat{g}(s.p(t)-s.p(m))h^{-1}\hat{g}(s.p(t)-s.p(m))^{-1})\triangleright e$.

The fact that the label of the blob ribbon operator before the intersection is unaffected by the commutation relation is perhaps unsurprising, because this part of the operator corresponds to the motion of the blob excitation before it braids with the magnetic excitation and so before it has undergone its transformation. This can be seen from the fact that the blob ribbon operator actually creates two excitations and the one which is not moved should not be affected by the other one moving through the loop excitation. Another thing to note is that, as long as the blob ribbon operator is not confined, we can deform the ribbons of the blob ribbon operators without changing their action,

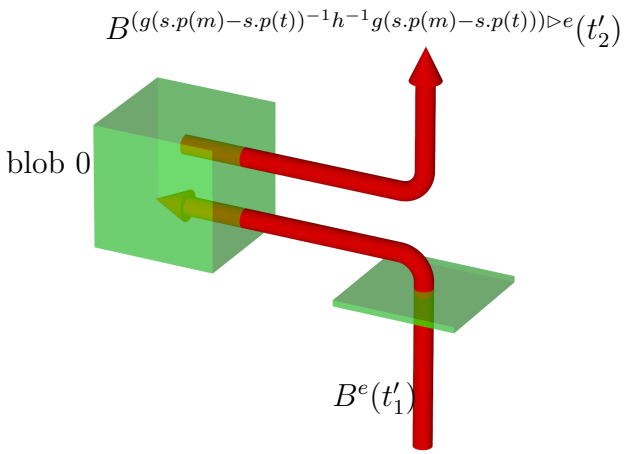

Figure 6.11: Under commutation with the membrane operator, the blob ribbon operator is diverted to pass through blob 0 of the membrane and splits into two parts. The part after the intersection has its label changed by the membrane operator. This figure shows the dual paths of the resulting ribbon operators (though note that the direct paths for each section still start at the start-point of the ribbon operator.)

---

as long as we keep the end-points fixed. Because of this, it does not matter at which plaquette q our magnetic membrane and blob ribbon operators intersect (which is reflected in the commutation relation Equation 6.13 having no explicit mention of q).

In addition to the transformation undergone by the blob ribbon operator, the E label of the higher-flux excitation changes from e_m to $e_m[\hat{g}(s.p(m) - s.p(t)) \triangleright e]$. This transformation of the E-label of the membrane operator is simply the standard braiding relation between a blob excitation and an E-valued membrane, analogous to the one we saw in Section 6.3 (see Equation 6.8, although note that there are differences due to the start-points of the ribbon and membrane being different and the orientation of the membrane being reversed).

If we give the magnetic membrane operator and the blob ribbon operator the same start-point, the braiding relation that we explained above simplifies and we are able to remove the operator $\hat{g}(s.p(m) - s.p(t))$ from the relation. We move the start-points together, without moving them through the higher-flux membrane (which would alter the commutation relations found so far). In this case the blob label goes from e before

the braiding to $h^{-1} \triangleright e$ afterwards (at least in the part after the intersection) and the E label of the higher-flux membrane operator goes from e_m beforehand to $e_m e$. That is, we obtain the braiding relation:

Higher-Flux–Blob Braiding Relation:

$$B^e(t) C_T^{h,e_m}(m) \left|GS\right\rangle = C_T^{h,e_m e}(m) B^e(t_1') B^{h^{-1} \triangleright e}(t_2') \left|GS\right\rangle . \tag{6.14}$$

If we also use the opposite relative orientation of ribbon and membrane, obtained by reversing the direction of the blob ribbon operator compared to the situation from Figure 6.10, the blob label e becomes $h \triangleright e$ and the membrane label e_m becomes $e_m[h \triangleright e^{-1}]$ under the braiding.

If we want to consider the braiding of the original magnetic excitation, produced by the membrane operator $C_T^h(m) = \sum_{e_m \in E} C_T^{h,e_m}(m)$, we simply need to sum over the E-valued label e_m of the higher-flux membrane operator. Then Equation 6.13 becomes

$$B^e(t) C_T^h(m) \left|GS\right\rangle = C_T^h(m) B^e(t_1) B^{(\hat{g}(s.p(t) - s.p(m)) h^{-1} \hat{g}(s.p(t) - s.p(m))^{-1}) \triangleright e}(t_2') \left|GS\right\rangle , \tag{6.15}$$

from which we see that the magnetic excitation is unchanged by the braiding, whereas the blob excitation is affected in the same way as in the braiding with the higher-flux excitation.

6.4.2 Braiding With Other Higher-Flux Excitations

Next, we consider the braiding between two higher-flux excitations. As shown in Figure 6.12, one of the loop-like excitations (indicated by a small red ring) is moved along the red surface and through another loop (indicated by a large green loop attached to a large green surface). To calculate the braiding relation, we apply membrane operators on these surfaces and examine the commutation relations between the membrane operators.

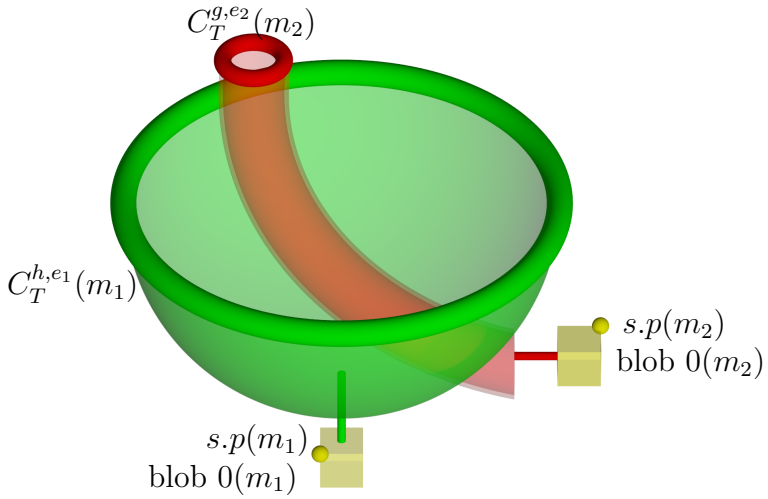

Figure 6.12: We consider the braiding move where we pull one higher-flux loop excitation (red torus) through another (green torus). This can be implemented using higher-flux membranes applied on the green and red membranes in the figure. If we first apply the membrane operator $C_T^{h,e_1}(m_1)$ on the green membrane, then $C_T^{g,e_2}(m_2)$ on the red membrane, then we are considering the case where we first produce the green loop excitation then move the red one through it. Comparing this to the opposite order of operators gives us the braiding relation.

We define the membrane operators as indicated in Figure 6.12, but then we use the topological nature of the magnetic membrane operators to pull m_2 through m_1 (as we did in the $\triangleright$ trivial case in Section 6.2.2), while keeping the start-point and blob 0 fixed. In Ref. [2], we show that this leads to the commutation relation

$$
\begin{aligned}
&C_T^{g,e_2}(m_2)C_T^{h,e_1}(m_1)\,|GS\rangle \\
&\quad = C_T^{h,e_1\left[g((1)-(2))\triangleright([h_{[2-1]}\triangleright e_2^{-1}][(h_{[2-1]}g^{-1})\triangleright e_2])\right]}(m_1)C_{\triangleright}^{h_{[2-1]}gh_{[2-1]}^{-1}}(m_2) \\
&\qquad \left(\prod_{\substack{\text{plaquette}\\ p\in m_2}} B^{[h_{[2-1]}^{-1}\triangleright e_{p|2}][(g^{-1}h_{[2-1]}^{-1})\triangleright e_{p|2}^{-1}]}((2)-(1))B^{e_{p|2}[(h_{[2-1]}g^{-1}h_{[2-1]}^{-1})\triangleright e_{p|2}^{-1}]}((1)-p)\right) \\
&\qquad \delta(h_{[2-1]}\triangleright e_2, \hat{e}(m_2))\,|GS\rangle\,,
\end{aligned}
\tag{6.16}
$$

where $g((1)-(2))$ is the path element for the path between the two start-points of the membranes and

$$
\begin{aligned}
h_{[2-1]} &= g((1)-(2))^{-1}hg((1)-(2)) \\
&= g((2)-(1))hg((2)-(1))^{-1}.
\end{aligned}
\tag{6.17}
$$

To simplify the expression, we used $e_{p|2}$ to denote the label of the plaquette p when we move its base-point to the start-point of m_2. This quantity is equivalent to $g(s.p(m_2) - v_0(p)) \triangleright e_p$. Furthermore, we used $B^{\cdots}((2)-(1))$ to denote a blob ribbon operator that runs from blob 0 of m_2 to blob 0 of m_1 and $B^{\cdots}((1)-p)$ to denote a blob ribbon operator running from blob 0 of m_1 to the blob on m_2 that is attached to plaquette p. These blob ribbon operators may seem complicated, but the situation is analogous to the braiding of blob ribbon operators with the magnetic membranes. Each of the blob ribbon operators that we added to the magnetic membrane operator has a similar commutation relation with the magnetic membrane operator as an ordinary blob ribbon operator. Namely, the blob ribbon operator splits into two parts, one that runs from blob 0 of m_2 to blob 0 of m_1 and one which runs from blob 0 of m_1 to the final destination of the original blob ribbon operator. The only difference from the ordinary blob ribbon operator braiding is that the label of the blob ribbon operator is an operator (depending on the label of a plaquette on m_2), which causes an additional apparent change to the label even for the part of the blob ribbon operator before the intersection. Note that the label of the part after the intersection reflects the change in the label of the rest of the associated magnetic membrane operator, namely that g becomes $h_{[2-1]}gh_{[2-1]}^{-1}$.

Apart from this splitting of the blob ribbon operators at the intersection of the membranes, we see that the labels of the two membrane operators change (as we mentioned previously for g). We have

$$\begin{aligned}
h &\to h,\\
e_1 &\to e_1\big[g((1)-(2)) \triangleright ([h_{[2-1]} \triangleright e_2^{-1}][(h_{[2-1]}g^{-1}) \triangleright e_2])\big],\\
g &\to h_{[2-1]}gh_{[2-1]}^{-1},\\
e_2 &\to h_{[2-1]} \triangleright e_2.
\end{aligned}$$

As usual for our braiding, when the start-points of the operators are not the same,

we have operators in our braiding relations. When we take the start-points to be the same, these relations simplify to

$$\begin{aligned} h &\to h, \\ e_1 &\to e_1[h \triangleright e_2^{-1}][(hg^{-1}) \triangleright e_2], \\ g &\to hgh^{-1}, \\ e_2 &\to h \triangleright e_2. \end{aligned} \tag{6.18}$$

This removes any operators from the labels, so that we have definite braiding. Note that the transformation of the 1-flux labels (h and g) is the same as for the braiding of two magnetic excitations in the $\triangleright$-trivial case, as given in Equation 6.6 in Section 6.2.2, replacing k with g (that equation in particular because we used the specific orientation of the loops also used to find that equation). Note that the product of the 1-fluxes of the two excitations is preserved by this relation, going from hg before the braiding to $(hgh^{-1})h = hg$ afterwards (swapping the order of multiplication for the product after braiding to reflect the change of positions of the two loops), reflecting the conservation of total 1-flux. However, the expression for the change of the E labels is not so easy to interpret. It is easier to understand these results if we change variables, from the surface labels of the membranes to the 2-fluxes possessed by the loop-like excitations, as we discussed at the start of Section 6.4. The 2-flux of the loop excitation, $\tilde{e}_1$, is related to the 1-flux label h and the surface label e_1 of the membrane operator by $\tilde{e}_1 = e_1[h^{-1} \triangleright e_1^{-1}]$. Therefore, we define

$$\tilde{e}_1 = e_1[h^{-1} \triangleright e_1^{-1}], \tag{6.19}$$

$$\tilde{e}_2 = e_2[g^{-1} \triangleright e_2^{-1}]. \tag{6.20}$$

Then, from our braiding relations in Equation 6.18, under braiding these 2-fluxes

transform according to

$$\begin{aligned}
\tilde{e}_1 &\to e_1[h \rhd \tilde{e}_2^{-1}][h^{-1} \rhd e_1^{-1}]\tilde{e}_2 \\
&= \tilde{e}_1[h \rhd \tilde{e}_2^{-1}]\tilde{e}_2, \\
\tilde{e}_2 &\to [h \rhd e_2][(hg^{-1}h^{-1}) \rhd (h \rhd e_2^{-1})] \\
&= h \rhd (e_2[g^{-1} \rhd e_2^{-1}]) \\
&= h \rhd \tilde{e}_2. \qquad (6.21)
\end{aligned}$$

This means that the product of the two fluxes transforms as

$$\tilde{e}_1\tilde{e}_2 \to \tilde{e}_1[h \rhd \tilde{e}_2^{-1}]\tilde{e}_2[h \rhd \tilde{e}_2] = \tilde{e}_1\tilde{e}_2$$

under the braiding, which indicates that the product of these 2-fluxes is conserved, reflecting the conservation of 2-flux. Putting this together, we can see that the 1-fluxes and 2-fluxes of the loop-like excitations transform under braiding as:

Higher-Flux–Higher-Flux Braiding Relation:

$$((g, \tilde{e}_2), (h, \tilde{e}_1)) \to ((h, \tilde{e}_1\tilde{e}_2[h \rhd \tilde{e}_2^{-1}]), (hgh^{-1}, h \rhd \tilde{e}_2)) \qquad (6.22)$$

Here the fact that one loop is moved past the other during the braiding is represented by swapping the order of their labels in the brackets. As usual, the braiding relation depends on the fusion channel, so this is the braiding when the total 1-flux of the two excitations is hg and the total 2-flux is $\tilde{e}_1\tilde{e}_2$. This is particularly important to note for the higher-flux excitations, because there are many ways of implementing the braiding with different fusion products.

We also wish to work out the inverse transformation, which describes the reversed braiding process. Denoting the result of the forwards transformations as primed ver-

sions, we have from Equation 6.21:

$$
\begin{aligned}
&h = h', \\
&\tilde{e}'_2 = h \triangleright \tilde{e}_2 \implies \tilde{e}_2 = h'^{-1} \triangleright \tilde{e}'_2, \\
&g' = hgh^{-1} \implies g = h'^{-1} g' h', \\
&\tilde{e}'_1 = \tilde{e}_1 [h \triangleright \tilde{e}_2^{-1}] \tilde{e}_2 \implies \tilde{e}_1 = \tilde{e}'_1 \tilde{e}'_2 [h'^{-1} \triangleright \tilde{e}_2'^{-1}].
\end{aligned}
$$

This means that the inverse braiding transformation is

$$
\big((h', \tilde{e}'_1), (g', \tilde{e}'_2)\big) \to \big((h'^{-1} g' h', h'^{-1} \triangleright \tilde{e}'_2), (h', \tilde{e}'_1 \tilde{e}'_2 [h'^{-1} \triangleright \tilde{e}_2'^{-1}])\big). \tag{6.23}
$$

This matches the braiding proposed in Ref. [65] for higher gauge theory based on discussions of the loop braid group.

As we did when considering the braiding of the higher-flux with the blob excitation, we can also consider the braiding of our original magnetic excitation, before we pinned an additional E-valued loop to it. As described by Equation 6.12, we can obtain the original magnetic membrane operators from the higher-flux membrane operators by summing over all possible elements of e for the E label. That is, we consider

$$
C_T^g(m_2) C_T^h(m_1) \left|GS\right\rangle = \sum_{e_2 \in E} \sum_{e_1 \in E} C_T^{g,e_2}(m_2) C_T^{h,e_1}(m_1) \left|GS\right\rangle .
$$

Using Equation 6.16, after some algebraic manipulation we obtain

$$
\begin{aligned}
&C_T^g(m_2) C_T^h(m_1) \left|GS\right\rangle \\
&\quad = \sum_{e_1 \in E} \sum_{e_2 \in E} C_T^h(m_1) C_\triangleright^{h_{[2-1]} g h_{[2-1]}^{-1}}(m_2) \Bigg(\prod_{\substack{\text{plaquette} \\ p \in m_2}} B^{[h_{[2-1]}^{-1} \triangleright e_{p|2}][(g^{-1} h_{[2-1]}^{-1}) \triangleright e_{p|2}^{-1}]}((2) - (1)) \\
&\qquad B^{e_{p|2}[(h_{[2-1]} g^{-1} h_{[2-1]}^{-1}) \triangleright e_{p|2}^{-1}]}((1) - p) \Bigg) \left|GS\right\rangle .
\end{aligned}
$$

Then looking at the effect of braiding on the G-valued label, we see that the result is

simply conjugation of one of the magnetic flux labels by the other. The labels of the blob ribbons corresponding to m_2 are also changed before and after the intersection of the two membranes. Just as we discussed for the higher-flux membrane operators earlier in this section, the label after the intersection reflects the change to the label g of the membrane operator applied on m_2, while the change to the label before the intersection is only an apparent change due to the operator label.

6.4.3 Braiding With E-Valued Loops

The final braiding relation to consider is the braiding between these higher-flux excitations and the E-valued loops. We can obtain this relation from the calculation for two higher-flux membrane operators in the previous section, because the E-valued loops are simply higher-flux excitations with trivial G-label. We therefore simply need to take the special case of that calculation when one of the G elements is 1_G. Rather than repeating the full equations, we will only present the results in the same-start-point braiding case.

First we consider the case where we pull the higher-flux excitation through the E-valued loop. In this case, the red excitation shown in Figure 6.12 is a higher-flux excitation, produced by a membrane operator $C_T^{g,e_{\text{mag}}}(m_2)$, while the green excitation is a pure E-valued loop, labelled by e_m. Then, from Equation 6.18, the label of the E-valued loop transforms as $e_m \to e_m e_{\text{mag}}^{-1}[g^{-1} \triangleright e_{\text{mag}}]$ under the braiding, while the labels of the magnetic membrane operator are unaffected by the braiding. When we change to consider our irrep basis for the E-valued loops (given in Equation 4.4, although when E is Abelian we can drop the matrix indices), this transformation gives us a phase of $\gamma(e_{\text{mag}}^{-1}[g^{-1} \triangleright e_{\text{mag}}])^{-1}$, where γ is the irrep of E labelling the E-valued loop. Note that if we consider the ordinary magnetic excitation by averaging over e_{mag}, the braiding relation is different for each value of e_{mag}, so the different terms in the sum accumulate different transformations. This means that even if the excitation is initially an ordinary magnetic excitation, with an equal superposition of the different E-labels, it will not

necessarily remain so after braiding. Instead, the magnetic excitation becomes an uneven superposition of the different higher-flux excitations, with labels coupled to the state of the E-valued loop.

Now we consider the opposite case where the E-valued loop passes through the higher-flux excitation. In this case, the red excitation from Figure 6.12 is a pure E-valued loop excitation produced by the membrane operator $\delta(\hat{e}(m_2), e_m)$, while the green excitation is a higher-flux loop excitation produced by the membrane operator $C_T^{h,e_1}(m_1)$. Then the label e_m of the E-valued loop transforms as $e_m \to h \triangleright e_m$ under the braiding, while the labels of the higher-flux operator are again unaffected by the braiding move. In our irrep basis for the E-valued membrane operators, this transformation actually changes the irrep γ labelling the membrane operator to a different irrep $h^{-1} \triangleright \gamma$ in the same $\triangleright$-Rep class of irreps (where two irreps of E, α and β, are in the same $\triangleright$-Rep class if there exists a $g \in G$ such that $\alpha(g \triangleright e) = \beta(e)$ for all $e \in E$). This suggests that the irreps of E are not by themselves good labels for the topological charge, because the irreps are not invariant under braiding.

6.4.4 Summary of Braiding in This Case

Table 6.3 summarizes which types of excitation can have non-trivial braiding relations in the case where E is Abelian and ∂ maps to the centre of G, where non-trivial braiding between the types of excitation is indicated by ticks. Note that the higher-flux excitations have potentially non-trivial braiding relations with every class of excitation.

Non-Trivial Braiding?	Electric	Higher-flux	Blob	E-valued loop
Electric	✗	✓	✗	✗
Higher-flux	✓	✓	✓	✓
Blob	✗	✓	✗	✓
E-valued loop	✗	✓	✓	✗

Table 6.3: A summary of which excitations braid non-trivially when the group E is Abelian and ∂ maps onto the centre of G. A tick indicates that at least some of the excitations of each type braid non-trivially with each-other, while a cross indicates that there is no non-trivial braiding between the two types.

Chapter 7

Topological Charge

7.1 Introduction

So far, we have alluded to the idea of the excitations carrying a conserved "topological" charge, without spending much time actually discussing this charge and the values it can take. In this chapter, we will explain in more detail what we mean by a topological charge and how we can measure it. Topological charge is a conserved quantity associated with the excitations of the model (the anyons), while the ground state is said to have the trivial charge, or vacuum charge. The charge held within a region can only be changed by moving that charge from inside the region to outside or moving charge from the outside in. This means that the charge in a region can only be changed by operators that connect the inside of the region to the outside. We note that there is no need for a symmetry to enforce this conservation. Topological charge is conserved on the level of the Hilbert space and can be defined without reference to any Hamiltonian (although in this case we must define the vacuum charge in another way). We can measure the topological charge associated to a region using operators on the boundary of that region, which is reminiscent of the way that we can determine the electric charge in a region by measuring the divergence of the electric field on the boundary of that region. For example, we can measure the charge associated to a loop excitation using a torus that encloses that loop, as shown in Figure 7.1. Any operator

that would move topological charge from inside a region to outside it must cross the boundary of that region, and so can be detected by the measurement operator on the surface.

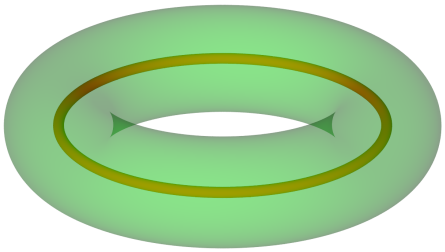

Figure 7.1: Given a loop excitation (thin red torus), we can measure its topological charge with a toroidal surface (green torus) enclosing it

---

In order to identify the operators that measure the topological charge, we first consider the characteristics that we require such operators to have. While the topological charges are properties of the Hilbert space, the Hamiltonian picks out a certain set of charges, such that the ground state has the trivial charge. Then, because a measurement operator should not change the charge in any region, we require that the measurement operators do not create any excitations. Therefore, we enforce that the measurement operator commutes with the Hamiltonian. Another property we require is that smoothly deforming the measurement operator without crossing any excitations should preserve the measured charge, because the ground state has trivial charge. Following the method of Bombin and Martin-Delgado [48], we construct such measurement operators using closed ribbon and closed membrane operators (the latter because we consider the 3+1d case, whereas Ref. [48] considered a 2+1d model). This may sound restrictive, but all operators in this model can be expressed in terms of the ribbon and membrane operators, and the ribbon operators that commute with the energy terms must be closed. However, we may need to take additional steps to guarantee that the closed operators commute with all of the energy terms (because there may be some obstruction to closing the operator without producing some excitations). Our approach is as follows. We take the surface on which we want to measure the charge. Then we apply as many independent closed membrane and ribbon operators

that act on that surface as possible, and construct all possible sums of these operators that commute with the Hamiltonian. The space of such operators corresponds to the space of possible topological charges. We note that the step of choosing the surface is important, not just because it determines where we want to measure the charge, but also because the set of charges to be measured depends on the topology of the measurement surface. A spherical measurement surface measures a different set of topological charges from a toroidal surface for example. A spherical measurement surface cannot distinguish a loop-like excitation from a point-like one, and for this reason we say that a spherical surface measures the point-like charge of an excitation (or set of excitations). In order to determine the loop-like charge of an excitation, we must use a toroidal surface, such as the one shown in Figure 7.1 (or a surface of higher genus, although we will not consider these in this work).

One subtlety with measuring the charge of a loop-like excitation is that, because the loops are extended, the excitations may pierce the measurement surface and not be wholly contained within or without the surface. As an example, consider the situation shown in Figure 7.2. As part of the measurement procedure, we must choose closed paths on which to measure any magnetic flux enclosed by the torus. However, in the presence of excitations on the surface itself, two choices of loops to measure on (for example, the blue or yellow paths in the figure) would give different results. This is because different loops may or may not link with the excitation (the thicker red torus). Because both choices are supposed to measure the charge within the torus (the partially transparent green torus), this leads to a contradiction. If we want to measure the charge held within a surface without knowing what excitations are present and where they are, this presents a difficulty. Therefore, we include in our charge measurement operators a projector to the space where the surface has no excitations. This sidesteps the above issue, but it does mean that we cannot measure the charge of confined excitations (which always cause excitations on a surface enclosing them) using this procedure.

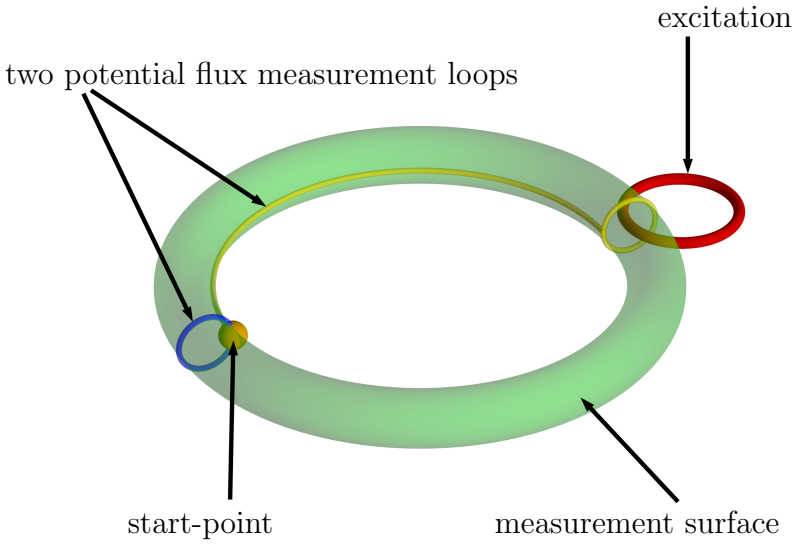

Figure 7.2: If an excitation pierces the measurement surface, then the charge within the surface is ill-defined. In the case shown in this figure, measuring the 1-flux along the two potential loops may give different results (to the point of not giving the same topological charge).

7.2 Topological Charge Within a Sphere

Before we look at the charge measured by a torus, which is sensitive to both loop-like and point-like charge, we will first examine the charge measured by a sphere. We will do this in the case where E is Abelian and ∂ maps onto the centre of G (Case 2 in Table 2.1), which includes the $\triangleright$ trivial case (Case 1 in Table 2.1) as a sub-case. Because a sphere has no non-contractible cycles, the sphere should only be sensitive to point-like charge. Nonetheless, the sphere charge is interesting, not only because it lets us look at the properties of point particles, but because loop excitations also possess point-like charge. As we explained in the previous section, to measure the charge within a sphere we first project to the case where there are no excitations on the measurement surface. Then we consider which independent closed ribbon and membrane operators we can apply on this surface.

While it may seem that we can independently apply ribbon operators around any closed

loop on the surface of the sphere, this is not the case. Any ribbon operators are either topological or confined (or can be written as a linear combination of ribbon operators of the two types). If a ribbon operator is confined, then applying it leads to excitations on the measurement surface, which we do not allow. On the other hand, if a ribbon operator is topological, then because all closed paths on the sphere are contractible on the spherical surface, the ribbon can be contracted to nothing without affecting the action of the ribbon operator. This means that applying a closed topological ribbon operator on the surface of the sphere is equivalent to applying the identity operator (at least in the subspace on which we apply measurement operators). Therefore, any ribbon operators that we are allowed to apply (the non-confined ones) act trivially.

This leaves us only with the membrane operators $C_T^h(m)$ and $L^e(m)$, where $C_T^h(m)$ is the total magnetic membrane operator defined in Section 4.5 (see Equation 4.8) and $L^e(m)$ is the E-valued membrane operator $\delta(e, \hat{e}(m))$. We consider applying these two operators over the sphere. We take the membrane m to be oriented inwards on the sphere, but taking the opposite orientation would just be equivalent to using e^{-1} instead of e so it would not give us a different set of operators. Although we apply both operators on the same sphere, when we define the membrane m for each operator to act on we need to define a start-point for the membrane. It would seem that we could choose the start-points of the membranes to be different for the two membrane operators, giving us many potential measurement operators. However, this is not the case because of the requirement that the total measurement operator commute with the energy terms on the sphere. As we have discussed previously, both the magnetic membrane operator and E-valued membrane operator commute with the vertex transforms except those at the start-points. If the two operators have different start-points, then each must individually commute with their specific start-point vertex transforms (rather than their combination having to commute with transforms at their mutual start-point). However, when a membrane operator commutes with the start-point vertex transforms, the start-point of the operator becomes arbitrary. That is,

if the start-point is not excited we can move the start-point without affecting the action of the membrane operator, because parallel transport of a vertex is equivalent to applying a vertex transform (as described in Section 2.2.3). This means that we can move the start-points to be in the same location anyway, without affecting the action of the two membrane operators. Therefore, without loss of generality, we can consider the two start-points of the membrane operators to be in the same location.

The most general operator we can apply is then a linear combination of terms with the form $C_T^h(m)L^e(m)$ for different labels h and e, where m is the spherical membrane that we are measuring the charge within. That is, a general measurement operator has the form

$$\sum_{h\in G}\sum_{e\in E}\alpha_{h,e}C_T^h(m)L^e(m),$$

where $\alpha_{h,e}$ are a set of coefficients. We now wish to find which coefficients lead to the operator commuting with the energy terms on the sphere and so give us a valid measurement operator. As we show in Ref. [2], requiring commutation with the energy terms leads to two types of restrictions for the coefficients. Some of these restrictions enforce that the coefficient $\alpha_{h,e}$ must be zero for certain labels (i.e. certain pairs of label (h,e) for $C_T^h(m)L^e(m)$ are disallowed), while other conditions mean that the coefficients of two pairs (h_1,e_1) and (h_2,e_2) must be the same (i.e. $\alpha_{h_1,e_1}=\alpha_{h_2,e_2}$). As a shorthand, we write $(h_1,e_1)\overset{\text{S}}{\sim}(h_2,e_2)$ for two pairs that are subject to this latter type of restriction (must have equal coefficients). Then we find the following restrictions for the pairs (h,e):

Spherical Charge Conditions:

$$h\triangleright e=e, \tag{7.1}$$

$$\partial(e)=1_G, \tag{7.2}$$

$$h\overset{\text{S1}}{\sim}\partial(f)h\ \forall f\in E, \tag{7.3}$$

$$(h,e)\overset{\text{S2}}{\sim}(ghg^{-1},g\triangleright e)\ \forall g\in G. \tag{7.4}$$

The first of these conditions comes from the requirement that blob 0 of the magnetic membrane operator, which can obtain a non-trivial 2-holonomy under the action of the membrane operator, remains unexcited. The second condition comes from the overall fake-flatness of the sphere, which is required if there are no plaquette excitations on the sphere, meaning that an E-valued membrane operator that does not satisfy this condition will always give zero when applied on the subspace where the sphere surface is unexcited. The third condition, Equation 7.3, is a result of the edges perpendicular to the sphere surface being unexcited: because the edges are unexcited, we can freely apply edge transforms of label f on each edge, which is equivalent to the action of a magnetic membrane operator of label $\partial(f)$, without affecting the state. This means that the measurement operator differing by such a membrane operator should not be regarded as distinct. The last condition, Equation 7.4 , comes from requiring the vertex energy term at the start-point of the membrane to remain unexcited. We give a detailed derivation of these conditions in the Supplementary material for Ref. [2].

The two equivalence relations $\overset{\text{S1}}{\sim}$ and $\overset{\text{S2}}{\sim}$ together form the equivalence relation $\overset{\text{S}}{\sim}$, where any pairs of labels (h_1, e_1) and (h_2, e_2) related by $\overset{\text{S}}{\sim}$ must have equal coefficients. We can write $\overset{\text{S}}{\sim}$ explicitly as

$$(h, e) \overset{\text{S}}{\sim} (\partial(f)ghg^{-1}, g \triangleright e), \tag{7.5}$$

for each $g \in G$ and $f \in E$ (i.e. (h, e) is in the same equivalence class as (h', e') if there exists any $g \in G$ and $f \in E$ such that $(h', e') = (\partial(f)ghg^{-1}, g \triangleright e)$). Given all of these conditions for our measurement operators, we can construct a basis for the space of allowed measurement operators. As we show in Ref. [2], one such basis is given by a set of operators labelled by two objects. The first object, C, is a $\triangleright$-class of the kernel of ∂. A $\triangleright$-class of the kernel is a subset of the kernel consisting of elements related by the equivalence relation $e \overset{\triangleright}{\sim} f$ if there exists a $g \in G$ such that $g \triangleright e = f$. It is convenient to pick a representative element r_C for each such class C. Then we define the centralizer of the class C as $Z_{\triangleright, r_C} = \{\, h \in G \mid h \triangleright r_C = r_C \,\}$. The second object

that labels our basis operators is a class within this centralizer, this time described by the equivalence relation

$$h \overset{Z_{\triangleright,r_C}}{\sim} xhx^{-1}\partial(w) \tag{7.6}$$

for any elements $x \in Z_{\triangleright,r_C}$ and $w \in E$. Note that this equivalence relation is similar to $\overset{S}{\sim}$, in that it gives the same form of relation, but only for elements $x \in Z_{\triangleright,r_C}$ (for which $x \triangleright r_C = r_C$) rather than general elements $g \in G$. The basis operator corresponding to a particular $\triangleright$-class C of the kernel and equivalence class D (defined by Equation 7.6) of the associated centralizer is

$$T^{D,C}(m) = \sum_{q \in Q_C} \sum_{d \in D} C_T^{qdq^{-1}}(m) L^{q \triangleright r_C}(m),$$

where Q_C is a set of elements of G that move us between the elements of the $\triangleright$-class C, so that each element $e_i \in C$ has a unique $q_i \in Q_C$ such that $e_i = q_i \triangleright r_C$. As we show in Ref. [2], any element $g \in G$ can be uniquely decomposed as a product of an element of Q_C and an element of $Z_{\triangleright,r_C}$. We can use this basis of operators to construct the projectors to definite topological charge within the sphere. These projectors are given by

$$T^{R,C}(m) = \frac{|R|}{|Z_{\triangleright,r_C}|} \sum_{D \in (Z_{\triangleright,r_C})_{cl}} \chi_R(D) T^{D,C}(m), \tag{7.7}$$

where R is an irrep of the quotient group $Z_{\triangleright,r_C}/\partial(E)$ with dimension $|R|$ and $(Z_{\triangleright,r_C})_{cl}$ is the set of classes in the centralizer defined by the equivalence relation Equation 7.6. Note that the character χ_R of irrep R is independent of the element $d \in D$ (because characters are a function of conjugacy class, and R being an irrep of the quotient group means that it is also insensitive to factors of $\partial(w)$ from Equation 7.6). In the Supplementary Material of Ref. [2], we prove that the operators defined by Equation 7.7 are indeed projectors and are orthogonal and complete in our space.

7.3 Topological Charge Within a Torus

Now we consider measuring the topological charge using a toroidal surface. The topological charges measured by the torus are more numerous and mathematically complicated to derive than those measured by the sphere. As illustrated in Figure 7.3, we represent the torus surface as a rectangle with opposite sides identified. These sides are then one particular choice for the two independent cycles of the torus. We will choose to apply any membrane operators on this rectangle with the boundary at the cycles, before closing the rectangle by gluing the opposite edges. This leaves "seams" at the two cycles, which may have special properties because the action of the membrane operators on either side of the seam may not match. We can understand this by imagining taking a membrane operator applied on a rectangle and folding it up to glue the opposite edges together. There is no guarantee that a membrane operator acts the same on opposite sides of the rectangle, and this disparity may remain when we glue the sides together. This leads to additional joining conditions required to prevent additional excitations being present at these seams.

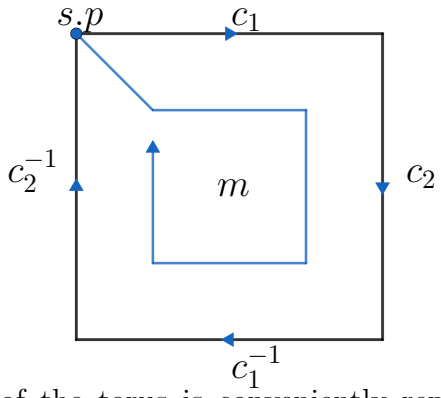

Figure 7.3: The surface of the torus is conveniently represented by a square with periodic boundary conditions. The edges of this square (which are glued due to the periodic boundary conditions) are referred to as the seams of the torus. We apply electric ribbon operators along these seams to measure the non-contractible cycles of the torus and apply an E-valued membrane operator on the surface. We will also apply blob ribbon operators around the two cycles and a magnetic membrane operator over the surface. The edges cut by the dual membrane of the magnetic membrane point outwards from the page.

To find the measurement operators, we have to first project onto the case where the surface itself is not excited, then we see what degrees of freedom are left over. After projecting onto all of the plaquettes on the surface being flat, these two cycles of the torus are still left undetermined. We therefore apply two closed electric operators $\delta(\hat{g}(c_1), g_{c_1})$ and $\delta(\hat{g}(c_2), g_{c_2})$, where c_1 and c_2 are the two cycles of the torus. We also apply a closed membrane operator $\delta(\hat{e}(m), e_m)$ on the torus, with the glued boundary of this torus being $c_1 c_2 c_1^{-1} c_2^{-1}$.

Next we apply our magnetic membrane operator $C_T^h(m)$. Because we already projected to the subspace where the torus satisfies fake-flatness, some of the details of the operator are arbitrary; in particular the set of paths on the direct membrane, which affect the action on the edges, can be freely chosen, as long as these paths do not cross the seams of the membrane (we take the convention that they do not cross the seams because two choices of path that differ by a non-contractible cycle may give different results and this gives us a consistent way of choosing the paths, but we show in Ref. [2] that choosing different paths does not give us new measurement operators).

Finally, we apply blob ribbon operators around the cycles, so that our measurement operator so far is given by

$$B^{e_{c_1}}(c_1) B^{e_{c_2}}(c_2) C_T^h(m) \delta(\hat{e}(m), e_m) \delta(\hat{g}(c_1), g_{c_1}) \delta(\hat{g}(c_2), g_{c_2}).$$

In principle we could put closed blob ribbon operators anywhere on the membrane, rather than just on the cycles c_1 and c_2. However, any blob ribbon operators with label in the kernel of ∂ can be freely deformed on the surface without affecting the action of the ribbon operators, because these operators are topological and there are no excitations on the surface. This means that any such blob ribbon operator that does not wrap around a non-contractible cycle may be contracted into nothing, while an operator that does wrap around a non-contractible cycle on the torus may be deformed to wrap around the chosen cycles c_1 and c_2 (if the ribbon operator wraps both cycles,

or wraps one multiple times, we split it into multiple ribbon operators on the cycles). This is not true for the other blob ribbon operators (those with label outside the kernel), because they are confined and so are not topological. Instead we find that this confinement leads to their position being fixed along the seams and their labels being restricted (as described in Ref. [2]), in order not to create any excitations (the labels of the blob ribbon operators are chosen so that the ribbon operators remove excitations along the seam of the magnetic membrane operator).

Much like with the sphere charge discussed in Section 7.2, requiring the measurement operator to project onto the subspace where the measurement surface is unexcited gives us several restrictions. These take the forms of conditions on the allowed labels for the ribbon and membrane operators in the measurement operator, as well as conditions on which linear combinations of these operators are allowed (represented by equivalence relations, with equivalent sets of labels appearing with equal weight in the linear combination). As we show in Ref. [2] (see the Supplementary Material of that paper for a detailed derivation), this leads to the following restrictions on the sextuples $(g_{c_1}, g_{c_2}, h, e_{c_1}, e_{c_2}, e_m)$:

Toroidal Charge Conditions:

$$\partial(e_m) = [g_{c_2}, g_{c_1}], \tag{7.8}$$

$$\partial(e_{c_2}) = [g_{c_1}, h], \tag{7.9}$$

$$\partial(e_{c_1}) = [h, g_{c_2}], \tag{7.10}$$

$$1_E = [h \triangleright e_m^{-1}]\, e_m e_{c_1}^{-1} [g_{c_1}^{-1} \triangleright e_{c_1}] e_{c_2}^{-1} [g_{c_2}^{-1} \triangleright e_{c_2}], \tag{7.11}$$

$$(g_{c_1}, g_{c_2}, h, e_{c_1}, e_{c_2}, e_m)$$

$$\sim (g g_{c_1} g^{-1}, g g_{c_2} g^{-1}, g h g^{-1}, g \triangleright e_{c_1}, g \triangleright e_{c_2}, g \triangleright e_m)\ \forall g \in G, \tag{7.12}$$

$$(g_{c_1}, e_{c_2}, e_m) \sim (\partial(e)^{-1} g_{c_1}, e_{c_2} [h \triangleright e]\, e^{-1}, e_m e^{-1} [g_{c_2}^{-1} \triangleright e])\ \forall e \in E, \tag{7.13}$$

$$(g_{c_2}, e_{c_1}, e_m) \sim (\partial(r) g_{c_2}, e_{c_1} [h \triangleright r]\, r^{-1}, e_m r^{-1} [g_{c_1}^{-1} \triangleright r])\ \forall r \in E, \tag{7.14}$$

$$(h, e_{c_1}, e_{c_2}) \sim (\partial(e) h, e_{c_1} [g_{c_2}^{-1} \triangleright e]\, e^{-1}, e_{c_2} [g_{c_1}^{-1} \triangleright e^{-1}]\, e)\ \forall e \in E. \tag{7.15}$$

Of these conditions, Equation 7.8 arises from the fake-flatness condition on the torus surface (analogous to Equation 7.2 for the sphere). Equations 7.9 and 7.10 are more subtle, and ensure that the plaquettes near the seams of the torus (the region where the opposite edges of the square in Figure 7.3 connect together) are unexcited by the action of the magnetic membrane operator and blob ribbon operators. Equation 7.11 states the condition for blob 0 of the membrane operator to remain unexcited by the measurement operator, with a contribution from both the magnetic membrane operator and the blob ribbon operators (which originate in blob 0). Then Equation 7.12 gives the condition for the start-point vertex to remain unexcited, while Equations 7.13 and 7.14 ensure that the edges along the seams of torus are unexcited. Finally, Equation 7.15 arises from the fact that we project onto the case where the edges perpendicular to the torus surface are unexcited, so we can freely apply edge transforms on these edges without affecting the state (meaning that measurement operators related by this equivalence condition should not be regarded as distinct).

These conditions show a striking resemblance to the relations that appear in the calculation of the ground state degeneracy of the 3-torus in Ref. [58] and indeed they map perfectly onto them, as we demonstrate in Ref. [2]. This indicates that the number of topological charges we can measure within a 2-torus is the same as the ground state degeneracy of the 3-torus, as also found in Ref. [60]. This is perhaps to be expected for a topological phase; in topological quantum field theories (TQFTs) there is a correspondence between the partition function associated to a closed (four-dimensional) manifold $M \times S^1$ and the dimension of the Hilbert space associated to the open manifold $M \times I$ (where I is the interval) [92]. Then the ground state degeneracy of the higher lattice gauge theory model on a (three-dimensional) manifold M is equal to the partition function of the Yetter TQFT [58, 62] on a manifold $M \times S^1$. Taking M to be the 3-torus $T^3 = S^1 \times T^2$, we may therefore perhaps expect a relationship between the ground-state degeneracy of the 3-torus (which matches the partition function of the TQFT on $T^3 \times S^1 = S^1 \times T^2 \times S^1 = S^1 \times T^3$) and dimension of the space for

the degrees of freedom on thickened 2-torus $I \times T^2$ (which should match the size of the space associated to $I \times T^2 \times S^1 = I \times T^3$ in the TQFT). This in turn should match the number of independent measurement operators that we can apply on the toroidal measurement surface (which is really a thickened torus). This correspondence between ground-state degeneracy and the charges for this model was also discovered in Ref. [60], using a different method (employing tube algebra). Indeed, in Ref. [2], we explicitly construct the projection operators to the different values of topological charge and find that these projection operators are labelled by the same objects labelling the simple modules of the tube algebra found in Ref. [60].

Chapter 8

Conclusion

In this thesis we constructed the membrane and ribbon operators that produce the excitations of the higher lattice gauge theory model in three broad cases and used these operators to obtain various properties of the excitations. The model has two classes of excitations: those present in ordinary lattice gauge theory and those that involve the surface holonomy, which only appear when we consider higher lattice gauge theory. The character of the excitations depends on which of the special cases we consider. When the map $\triangleright$ from our crossed module is trivial, the first class of excitation, those familiar from ordinary lattice gauge theory, are largely unchanged by the move to higher lattice gauge theory. The excitations related to the surface holonomy are also simple, with Abelian fusion rules. Furthermore the braiding between the two sectors is trivial, with non-trivial braiding only within the sectors (i.e. between the G-valued excitations and between the E-valued excitations). However, unlike in ordinary lattice gauge theory, some of the point-like excitations are confined, costing additional energy to separate pairs of excitations after producing them, while some of the loop-like excitations are condensed and carry trivial topological charge. We discussed how this arises from a condensation-confinement transition between higher lattice gauge theory models described by crossed modules with the same groups G and E but different maps ∂ between the groups.

We saw that the excitations are more interesting when $\triangleright$ is non-trivial. In particular, we looked at the case where the group E labelling the 2-flux is Abelian and where the map ∂ maps to the centre of the group G. In this case, even though the group E is Abelian, the excitations related to 2-flux have an internal space controlled by $\triangleright$. In addition, there is non-trivial braiding between excitations from the two different sectors. In particular, it is sensible to consider loop excitations, which we called higher-flux loops, built from a combination of excitations from the two sectors. These excitations carry both an ordinary magnetic flux along paths that link with the loop and a 2-flux associated to a surface enclosing the loop. The higher-flux loops have potentially non-trivial braiding with excitations from both sectors. We found the braiding relations of these higher-flux loops and discovered that the braiding between two higher-flux loops matches a braiding scheme for loop-like excitations described in Ref. [65].

In addition to using the membrane operators to find the braiding properties of the various excitations, we have discussed the explicit construction of the membrane operators to construct operators to measure the topological charge in a region. The charge measurement operators are operators with support on or near the surface bounding a region, with different surfaces having different potential topological charges. A related idea is the fact that in 3+1d there are both point-like and loop-like excitations, with associated point-like and loop-like charge. While loop-like excitations may possess a point-like charge, which may be measured by enclosing the excitation in a sphere, they also possess loop-like charge, which can only be measured by surfaces with non-contractible handles such as the torus. This approach of explicitly constructing the charge measurement operators (which we present more fully in Ref. [2]) could allow us to measure the charge of collections of excitations and so provide a tool to consider fusion of topological charge in the higher-flux model and beyond.

There are many interesting avenues of research that would build on the results we've

obtained in this work, either directly or indirectly. Firstly, and most obviously, there are some results for the higher lattice gauge theory model that we have so far not been able to obtain. We have considered certain special cases (described by Table 2.1) and while there are some conceptual issues associated with taking the most general case (as described in Section 2.2.6), there may still be a way to work around these issues. It is possible that there are some features exhibited by the general case that we have not been able to study in our special cases.

It would also be interesting to further study the different types of topological charge possible in 3+1d topological phases, perhaps in a more general setting than this specific model. In this work we have mostly considered the simple excitations, and only considered measurement operators for the topological charge within a sphere and torus. We believe that it would be useful to do more work with the charges themselves. There are several questions that could be explored in this direction. For example, which surfaces do we need to consider to obtain all unique charges? How do we consider the fusion of charges that are measured by surfaces other than simple spheres? For instance, considering two tori, we could fuse the enclosed charges by bringing the tori together and stacking them on top of each-other, so that a single torus can enclose both. After some preliminary calculations, we found that, in the higher lattice gauge theory model at least, this leads to consistency conditions for the "threading flux" passing up through the two tori, which must be satisfied in order to be able to fuse them, but a complete calculation is left for future work. Another way to combine two torus charges would be to bring the two tori side-by-side, so that a 2-handled torus would be needed to enclose both. For more general surfaces, there could be even more ways of fusing charges.

Another sensible direction for future study would be to extend our approach of utilizing explicit construction of membrane operators, to more general models for topological phases. In particular, it is known that twisted gauge theory models can have non-trivial

three-loop braiding statistics [31, 50]. It has been claimed that these models cover all phases that can be produced from bosonic degrees of freedom and that result in bosonic point excitations (i.e. phases with trivial braiding between point-like particles) [93]. Because these models have a similar structure to the model considered in this work, the application of our methods to the twisted gauge theory model seems feasible. By looking at other models, we may be able to see if the relation between the ground state degeneracy of the 3-torus and the number of charges measured by a torus surface holds more generally than just in the higher lattice gauge theory model, as we may expect.

Finally, it would be interesting to study the condensation-confinement transitions in 3+1d topological phases in more detail. We have seen some examples in the higher lattice gauge theory model, but in the cases where E is Abelian (which we studied in more detail) the pattern of condensation is rather simple. The current understanding of condensation in 3+1d is perhaps incomplete, particularly when it comes to loop-like excitations condensing, and the examples considered in this work may be useful in studying this process.

Bibliography

[1] J. Huxford and S. H. Simon. Excitations in the higher lattice gauge theory model for topological phases I: Overview. *arXiv:2202.08294*, 2022.

[2] J. Huxford and S. H. Simon. Excitations in the higher lattice gauge theory model for topological phases III: the 3+1d case. *arXiv:2206.09941*, 2022.

[3] J. Huxford and S. H. Simon. Excitations in the higher lattice gauge theory model for topological phases II: the 2+1d case. *arXiv:2204.05341*, 2022.

[4] X.-G. Wen. Vacuum degeneracy of chiral spin states in compactified space. *Phys. Rev. B*, 40:7387–7390, 1989.

[5] X.-G. Wen. Topological orders in rigid states. *Int. J. Mod. Phys. B*, 04(02):239 – 271, 1990.

[6] X.-G. Wen. Topological order: From long-range entangled quantum matter to a unified origin of light and electrons. *ISRN Condens. Matter Phys.*, 2013.

[7] L. D. Landau. On the theory of phase transitions. I. *Phys. Z. Sowjetunion*, 11(16), 1937.

[8] D. C. Tsui, H. L. Stormer, and A. C. Gossard. Two-dimensional magnetotransport in the extreme quantum limit. *Phys. Rev. Lett.*, 48:1559, 1982.

[9] X.-G. Wen and Q. Niu. Ground state degeneracy of the FQH states in presence of random potential and on high genus Riemann surfaces. *Phys. Rev. B*, 41:9377, 1990.

[10] A. Stern. Anyons and the quantum Hall effect - A pedagogical review. *Ann. Phys. (N. Y.)*, 323:204–249, 2008.

[11] T. Chakraborty and P. Pietlinen. *The Quantum Hall Effects: Integral and Fractional*. Springer Ser. in Solid-State Sciences. Springer, New York, 2nd edition, 1995.

[12] S. Das Sarma and A. Pinczuk. *Perspectives in Quantum Hall Effects : Novel Quantum Liquids in Low-dimensional Semiconductor Structures*. Wiley, New York, 1997.

[13] X. Chen, Z.-C. Gu, Z.-X. Liu, and X.-G. Wen. Symmetry protected topological orders and the group cohomology of their symmetry group. *Phys. Rev. B*, 87:155114, 2013.

[14] X. Chen, Z.-C. Gu, and X.-G. Wen. Local unitary transformation, long-range quantum entanglement, wave function renormalization, and topological order. *Phys. Rev. B*, 82:155138, 2010.

[15] A. Mesaros and Y. Ran. Classification of symmetry enriched topological phases with exactly solvable models. *Phys. Rev. B*, 87:155115, 2013.

[16] A. Y. Kitaev. Fault-tolerant quantum computation by anyons. *Ann. Phys. (N. Y.)*, 303:2–30, 2003.

[17] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill. Topological quantum memory. *J. Math. Phys.*, 43:4452–4505, 2002.

[18] B. M. Terhal. Quantum error correction for quantum memories. *Rev. Mod. Phys.*, 87, 2015.

[19] B. J. Brown, D. Loss, J. K. Pachos, C. N. Self, and J. R. Wootton. Quantum memories at finite temperature. *Rev. Mod. Phys.*, 88, 2016.

[20] J. C. Bridgeman, S. T. Flammia, and D. Poulin. Detecting topological order with ribbon operators. *Phys. Rev. B*, 94:205123, 2016.

[21] J. M. Leinaas and J. Myrheim. On the theory of identical particles. *Nuovo Cim. B*, 37:1–23, 1977.

[22] F. Wilczek. Magnetic flux, angular momentum, and statistics. *Phys. Rev. Lett.*, 48:1144, 1982.

[23] D. Arovas, J. R. Schrieffer, and F. Wilczek. Fractional statistics and the quantum Hall effect. *Phys. Rev. Lett.*, 53:722, 1984.

[24] J. K. Pachos. *Introduction to Topological Quantum Computation.* Cambridge University Press, Cambridge, 2012.

[25] S. Rao. An anyon primer. *arXiv:hep-th/9209066v3*, 1992.

[26] A. Y. Kitaev. Anyons in an exactly solved model and beyond. *Ann. Phys. (N. Y.)*, 321, 2006.

[27] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. D. Sarma. Non-Abelian anyons and topological quantum computation. *Rev. Mod. Phys.*, 80, 2008.

[28] V. T. Lahtinen and J. K. Pachos. A short introduction to topological quantum computation. *SciPost Phys.*, 3, 2017.

[29] S. Doplicher, R. Haag, and J. E. Roberts. Local observables and particle statistics I. *Commun. Math. Phys.*, 23, 1971.

[30] S. Doplicher, R. Haag, and J. E. Roberts. Local observables and particle statistics II. *Commun. Math. Phys.*, 35, 1974.

[31] C. Wang and M. Levin. Braiding statistics of loop excitations in three dimensions. *Phys. Rev. Lett.*, 113, 2014.

[32] M. G. Alford, K.-M. Lee, J. March-Russell, and J. Preskill. Quantum field theory of non-abelian strings and vortices. *Nucl. Phys. B*, 384, 1992.

[33] Y. Aharanov and D. Bohm. Significance of electromagnetic potentials in the quantum theory. *Phys. Rev.*, 115, 1959.

[34] C.-H. Lin and M. Levin. Generalizations and limitations of string-net models. *Phys. Rev. B*, 89, 2014.

[35] C.-H. Lin, M. Levin, and F. J. Burnell. Generalized string-net models: A thorough exposition. *Phys. Rev. B*, 103, 2021.

[36] M. Levin and X.-G. Wen. String-net condensation: A physical mechanism for topological phases. *Phys. Rev. B*, 71:045110, 2005.

[37] M. Cheng, Z.-C. Gu, S. Jiang, and Y. Qi. Exactly solvable models for symmetry-enriched topological phases. *Phys. Rev. B*, 96:115107, 2017.

[38] F. A. Bais, P. van Driel, and M. de Wild Propitius. Quantum symmetries in discrete gauge theories. *Phys. Lett. B*, 280, 1992.

[39] M. de Wild Propitius and F. A. Bais. Discrete gauge theories. *arXiv:hep-th/9511201*, 1995.

[40] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86, 2012.

[41] C. K. Andersen, A. Remm, S. Lazar, S. Krinner, N. Lacroix, G. J. Norris, M. Gabureac, C. Eichler, and A. Wallraff. Repeated quantum error detection in a surface code. *Nat. Phys.*, 16:875–880, 2020.

[42] K. J. Satzinger, Y. Liu, A. Smith, et al. Realizing topologically ordered states on a quantum processor. *Science*, 374, 2021.

[43] C. Heinrich, F. Burnell, L. Fidkowski, and M. Levin. Symmetry enriched string-nets: Exactly solvable models for SET phases. *Phys. Rev. B*, 94, 2016.

[44] A. Hahn and R. Wolf. Generalized string-net model for unitary fusion categories without tetrahedral symmetry. *Phys. Rev. B*, 102, 2020.

[45] E. Lake and Y.-S. Wu. Signatures of broken parity and time-reversal symmetry in generalized string-net models. *Phys. Rev. B*, 94, 2016.

[46] Ingo Runkel. String-net models for nonspherical pivotal fusion categories. *J. Knot Theory Ramif.*, 29(06):2050035, 2020.

[47] Y. Hu, S. D. Stirling, and Y.-S. Wu. Ground state degeneracy in the Levin-Wen model for topological phases. *Phys. Rev. B*, 85, 2012.

[48] H. Bombin and M. A. Martin-Delgado. A family of non-Abelian Kitaev models on a lattice: Topological condensation and confinement. *Phys. Rev. B*, 78, 2008.

[49] Y. Wan, J. C. Wang, and H. He. Twisted gauge theory model of topological phases in three dimensions. *Phys. Rev. B*, 92, 2015.

[50] J. C. Wang and X.-G. Wen. Non-Abelian string and particle braiding in topological order: Modular SL(3,Z) representation and 3+1D twisted gauge theory. *Phys. Rev. B*, 91, 2015.

[51] A. Bullivant and C. Delcamp. Tube algebras, excitations statistics and compactification in gauge models of topological phases. *JHEP*, 10, 2019.

[52] R. Dijkgraaf and E. Witten. Topological gauge theories and group cohomology. *Commun. Math. Phys.*, 129, 1990.

[53] D. J. Williamson and Z. Wang. Hamiltonian models for topological phases of matter in three spatial dimensions. *Ann. Phys. (N. Y.)*, 377, 2017.

[54] K. Walker and Z. Wang. (3+1)-TQFTs and topological insulators. *Front. Phys.*, 7, 2012.

[55] C. W. von Keyserlingk, F. J. Burnell, and S. H. Simon. Three-dimensional topological lattice models with surface anyons. *Phys. Rev. B*, 87, 2013.

[56] X. Chen, F. J. Burnell, A. Vishwanath, and L. Fidkowski. Anomalous symmetry fractionalization and surface topological order. *Phys. Rev. X*, 5, 2015.

[57] Z. Wang and X. Chen. Twisted gauge theories in three-dimensional Walker-Wang models. *Phys. Rev. B*, 95, 2017.

[58] A. Bullivant, M. Calcada, Z. Kadar, P. Martin, and J. Faria Martins. Topological phases from higher gauge symmetry in 3+1D. *Phys. Rev. B*, 95, 2017.

[59] C. Delcamp and A. Tiwari. From gauge to higher gauge models of topological phases. *JHEP*, 10, 2018.

[60] A. Bullivant and C. Delcamp. Excitations in strict 2-group higher gauge models of topological phases. *JHEP*, 01, 2020.

[61] A. Bullivant, M. Calcada, Z. Kadar, P. Martin, and J. Faria Martins. Higher lattices, discrete two-dimensional holonomy and topological phases in (3+1)D with higher gauge symmetry. *Rev. Math. Phys.*, 32, 2020.

[62] D. Yetter. TQFT's from homotopy 2-types. *J. Knot Theory Ramif.*, 2, 1993.

[63] D. V. Else and C. Nayak. Cheshire charge in (3+1)-D topological phases. *Phys. Rev. B*, 96, 2017.

[64] S. Jiang, A. Mesaros, and Y. Ran. Generalized modular transformations in 3+1D topologically ordered phases and triple linking invariant of loop braiding. *Phys. Rev. X*, 4, 2014.

[65] A. Bullivant, J. Faria Martins, and P. Martin. Representations of the loop braid group and Aharonov-Bohm like effects in discrete (3+1)-dimensional higher gauge theory. *arXiv:1807.09551*, 2018.

[66] J. C. Baez, D. K. Wise, and A. S. Crans. Exotic statistics for strings in 4d BF theory. *Adv. Theor. Math. Phys.*, 11, 2007.

[67] C. Damiani. A journey through loop braid groups. *Expo. Math.*, 35, 2017.

[68] E. Witten. Topological quantum field theory. *Commun. Math. Phys.*, 117, 1988.

[69] I. Montvay and G. Münster. *Quantum Fields on a Lattice*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, Cambridge, 1994.

[70] K. G. Wilson. Confinement of quarks. *Phys. Rev. D*, 10, 1964.

[71] B. Durhuus. On the structure of gauge invariant classical observables in lattice gauge theories. *Lett. Math. Phys.*, 4, 1980.

[72] H. Pfeiffer. Higher gauge theory and a non-Abelian generalization of 2-form electrodynamics. *arXiv:hep-th/0304074*, 2003.

[73] J. C. Baez and J. Huerta. An invitation to higher gauge theory. *arXiv:1003.4485v2*, 2010.

[74] J. C. Baez. Higher Yang–Mills theory. *arXiv:hep-th/0206130v2*, 2002.

[75] J. Preskill. Lecture notes for physics 219: Quantum computation. Accessed 08/07/2022.

[76] F. A. Bais. Flux metamorphosis. *Nucl. Phys. B*, 170, 1980.

[77] M. Bucher, K.-M. Lee, and J. Preskill. On detecting discrete Cheshire charge. *Nucl. Phys. B*, 386, 1992.

[78] A. Komar and O. Landon-Cardinal. Anyons are not energy eigenspaces of quantum double Hamiltonians. *Phys. Rev. B*, 96, 2017.

[79] F. A. Bais and J. K. Slingerland. Condensate induced transitions between topologically ordered phases. *Phys. Rev. B*, 79, 2009.

[80] F. J. Burnell. Anyon condensation and its applications. *Annu. Rev. Condens. Matter Phys.*, 9, 2018.

[81] T. Neupert, H. He, C. von Keyserlingk, G. Sierra, and B. A. Bernevig. Boson condensation in topologically ordered quantum liquids. *Phys. Rev. B*, 93, 2016.

[82] F. A. Bais, B. J. Schroers, and J. K. Slingerland. Hopf symmetry breaking and confinement in (2+1)-dimensional gauge theory. *JHEP*, 05, 2003.

[83] I. S. Eliëns, J. C. Romers, and F. A. Bais. Diagrammatics for Bose condensation in anyon theories. *Phys. Rev. B*, 90, 2014.

[84] F. J. Burnell, C. W. von Keyserlingk, and S. H. Simon. Phase transitions in three-dimensional topological lattice models with surface anyons. *Phys. Rev. B*, 88, 2013.

[85] P. Ye, T. L. Hughes, J. Maciejko, and E. Fradkin. Composite particle theory of three-dimensional gapped fermionic phases: Fractional topological insulators and charge-loop excitation symmetry. *Phys. Rev. B*, 94, 2016.

[86] A. H. Clifford. Representations induced in an invariant subgroup. *Ann. Math.*, 38, 1937.

[87] J. McCool. On basis-conjugating automorphisms of free groups. *Can. J. Math.*, XXXVIII, 1986.

[88] A. G. Savushkina. On the group of conjugating automorphisms of a free group. *Math Notes+*, 1996.

[89] R. Fenn, R. Rimanyi, and C. Rourke. The braid-permutation group. *Topology*, 36, 1997.

[90] D. M. Dahm. A generalization of braid theory. Ph.d. thesis, Princeton University, 1962.

[91] D. L. Goldsmith. The theory of motion groups. *Michigan Math. J.*, 28, 1981.

[92] M. F. Atiyah. Topological quantum field theory. *Publ. Math. IHÉS*, 68, 1988.

[93] T. Lan, L. Kong, and X.-G. Wen. Classification of (3+1)D bosonic topological orders: The case when pointlike excitations are all bosons. *Phys. Rev. X*, 8, 2018.