

Excitations in the higher-lattice gauge theory model for topological phases. I. OverviewJoe Huxford^{1,2} and Steven H. Simon¹¹*Rudolf Peierls Centre for Theoretical Physics, Clarendon Laboratory, Oxford OX1 3PU, United Kingdom*²*Department of Physics, University of Toronto, Ontario M5S 1A7, Canada* (Received 8 August 2022; revised 20 June 2023; accepted 25 October 2023; published 12 December 2023)

In this series of papers, we study a Hamiltonian model for (3+1)-dimensional topological phases introduced in [Bullivant *et al.*, *Phys. Rev. B* **95**, 155118 (2017)], based on a generalization of lattice gauge theory known as “higher-lattice gauge theory.” Higher-lattice gauge theory has so-called “2-gauge fields” describing the parallel transport of lines, in addition to ordinary 1-gauge fields which describe the parallel transport of points. In this series we explicitly construct the creation operators for the pointlike and looplike excitations supported by the model. We use these creation operators to examine the properties of the excitations, including their braiding statistics. These creation operators also reveal that some of the excitations are confined, costing energy to separate that grows linearly with the length of the creation operator used. This is discussed in the context of condensation-confinement transitions between different cases of this model. We also discuss the topological charges of the model and use explicit measurement operators to rederive a relationship between the number of charges measured by a 2-torus and the ground-state degeneracy of the 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 looplike excitations. This first paper provides an accessible summary of our findings, with more detailed results and proofs to be presented in the other papers in the series.

DOI: [10.1103/PhysRevB.108.245132](https://doi.org/10.1103/PhysRevB.108.245132)**I. INTRODUCTION**

Outside of the phases of matter described by Landau symmetry-breaking classification [1], there exist so-called topological phases of matter [2–4]. These topological phases, which include the celebrated fractional quantum Hall systems [5–9], are characterized by long-range entanglement between their local degrees of freedom [4,10,11]. 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 [12,13] (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 [10]). This ground-state degeneracy depends on the topology of the manifold on which the topological phase is placed (e.g., such a phase may have no ground-state degeneracy on the sphere, but have a degeneracy on the torus), with this degeneracy being resistant to local perturbations [12,13]. This feature may allow such topological phases to serve as quantum memories [13–16] because encoding information in the topologically protected degenerate subspace makes the information resistant to local noise [13,17].

Further intriguing properties of the long-ranged entangled topological phases are revealed when we consider excitations. In 2+1 dimensions [(2+1)D], the entanglement structure allows these phases to support *anyonic* excitations, which are generalizations of the more familiar bosons and fermions. Moving two anyons around each other can induce nontrivial transformations, even at large distances [18–20] (for the interested reader, we note that there are many works giving pedagogical introductions to anyon physics, such as Refs. [21,22] and Appendix E in Ref. [23]). 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 [24,25], should sufficiently stable phases and excitations be constructed. In three spatial dimensions, while any pointlike excitations must be fermionic or bosonic [22,24,26,27], topological phases can admit loop braiding, such that pointlike or looplike excitations transform nontrivially when passed *through* a looplike excitation [28,29]. This can be thought of as a generalization of the Aharonov-Bohm effect [30,31], where braiding a pointlike electron around a looplike or stringlike magnetic flux tube results in a phase depending on the magnetic field enclosed.

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 [13,32–34]. While these models may not resemble those used to describe real materials [33], or only describe the renormalization group fixed point of their phase [35], they provide representatives for a large class of phases of matter [33]. This means that such constructions can be used to probe (and attempt to classify [34]) which kinds of phases can exist. The toy models have Hamiltonians that are constructed out of commuting projector operators, which

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

allows the quasiparticle excitations to be found exactly. Of these constructions for (2+1)D topological phases, two of the most successful are the Levin-Wen string-net model [34] and Kitaev's quantum double model [13] (which is related to discrete gauge theory that had priorly been discussed in Refs. [36,37]). The Kitaev quantum double class of models includes the toric code as its simplest case, which appears to have a practical application as a robust way to store qubits [13]. Indeed, one approach to building quantum computers uses so-called surface codes, which take inspiration from the toric code [38] and which have recently been experimentally realized on a small scale [39,40]. 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+1)D in the absence of an additional symmetry [41], when generalized appropriately from the original construction in Ref. [34] (see Refs. [32,33,42–44] for such generalizations). In both of these classes of models, it is well understood how to find the ground-state degeneracy [13,45] and the properties of the excitations, such as braiding statistics [13,33,34].

In the (2+1)D 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 [13]. This approach was used in Ref. [13] to study the excitations in Kitaev's quantum double model, and in Ref. [34] 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. Furthermore, in Ref. [46] 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. By applying this method to a modified version of the quantum double model which describes a condensation-confinement transition, the charges which condense and the charges that confined during the transition were identified [46]. 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 [47–49], which is a generalized version of the quantum double model in 3+1 dimensions [(3+1)D] and is based on the Dijkgraaf-Witten topological field theory [50]; a class of models developed from unitary G-crossed braided fusion categories (UGxBFCs) [51]; the Walker-Wang models [52–55], which are (3+1)D generalizations of the Levin-Wen string-net models [52]; and the higher-lattice gauge theory models [56–59], based on a generalization of lattice gauge theory (and related to the Yetter quantum field theory [60]). For the twisted gauge theory model, in particular, there has been significant study of the properties of the ground state [47] and the excitations, including their braiding properties [48,49]. However, the general approach to studying these (3+1)D models has been different from the approach used for models in two spatial dimensions. While in the (2+1)D case, the use of ribbon operators

to obtain the properties of the excitations is common, in the (3+1)D case an explicit construction of the ribbon and membrane operators (the higher-dimensional counterparts to ribbon operators, which produce looplike 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. [61,62], but less so for other models. Instead, indirect methods like dimensional reduction [28,48] and tube algebras [49] 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+1)D, and these examples of membrane operators in the (3+1)D twisted gauge theory model, we would like to be able to apply similar approaches to other (3+1)D models. In this work we will do precisely that, with one of the models discussed above.

In this series of papers we study a model [56] based on higher-lattice gauge theory [63,64], which can be defined in arbitrary dimension but which we will study in two and three spatial dimensions. Higher-lattice gauge theory is a generalization of lattice gauge theory, where there is a second gauge field which describes the parallel transport of the ordinary 1-gauge field across surfaces. This type of higher gauge theory (both on the continuum and in the lattice) has seen significant prior study in the context of topological phases. In Refs. [65,66], related topological quantum field theory (TQFT) constructions were used to describe confinement in regular gauge theories, while in Ref. [67] the corresponding TQFT was treated as a theory in its own right. In Ref. [56], a Hamiltonian model was constructed which realizes higher-lattice gauge theory, in the same way that Kitaev's quantum double model [13] realizes lattice gauge theory. Reference [56] also explored several of the properties of these Hamiltonian models. For example, the ground-state degeneracy was given in terms of the partition function of a TQFT, the Yetter TQFT, and also explicitly computed for some examples. Then in Ref. [58], the excitations were studied using a tube algebra approach, through which the looplike excitations in the model were classified and the simple types were counted. Furthermore, it was shown in Ref. [58] 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. [68] that higher gauge theory could lead to looplike excitations with nontrivial loop braiding statistics, and the associated representations of the loop braid group were found (the loop braid group describes the motions of loops [69,70]), although this was not done in the Hamiltonian model but instead from more geometric reasoning about the fluxes and gauge transforms involved.

However, 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. To do so, we will explicitly construct the membrane and ribbon operators for the Hamiltonian model [56] and use 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.

Our main results in the (3+1)D case are as follows. We construct 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 looplike or pointlike and that some of the pointlike excitations are confined, with an energy cost to separate a particle from its antiparticle that grows linearly with the length of ribbon used to do so. This is described in terms of a condensation-confinement transition between different higher-lattice gauge theory models, during which some of the looplike excitations condense out, becoming topologically trivial. Then, using our direct construction of the ribbon and membrane operators, we find the (loop-) braiding relations of our excitations 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 generally have an internal space, which can transform nontrivially under braiding, in addition to a conserved topological charge, which is not changed by the braiding. This topological charge is of significant interest, and so we also consider the charges present in the higher-lattice gauge theory model. Extending the methods of Ref. [46] to (3+1)D, 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 pointlike objects is measured by putting a sphere around that charge, similar to Gauss's law for electric charge. On the other hand, looplike excitations require a surface with handles in order to detect their looplike character. This is similar in concept to the tube algebra methods used in Ref. [58], which classify the boundary conditions of unexcited regions of space. Indeed, just like Ref. [58] 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.

A. Structure of this series

Due to the large amount of algebra needed to fully describe and prove our results, we have divided the discussion into three parts. This work is the first of the series, so we feel that it would be valuable to provide a brief guide to the set of papers. In this work, we will provide a more informal and descriptive overview of our main results for the (3+1)D model. We suggest that a general reader consider this work, before looking through the other papers in the series if they are interested in more detail, or are specifically interested in the (2+1)D model.

In the second paper [71], we consider the (2+1)D version of the higher-lattice gauge theory model. Perhaps the most interesting feature of this model is that, despite being in (2+1)D, this model still hosts looplike excitations. In addition to studying the topological content of the model, we demonstrate that in certain cases the looplike excitations can be described as domain walls between different symmetry sectors. This

idea that the (2+1)D models can describe symmetry-enriched topological phases is further expanded on when we map a subset of these models to another construction for such phases, the symmetry-enriched string-net model from Ref. [41].

In the final paper [72], we return to the (3+1)D model to provide more detailed results. This includes an explicit presentation of the commutation relations of the ribbon and membrane operators that give us the braiding relations. We also directly construct the measurement operators for topological charge within a torus and a sphere, and find the pointlike charge of the simple excitations of the model (including the looplike ones).

B. Structure of this paper

In the rest of the Introduction, we describe the model introduced in Ref. [56] and introduce other important concepts from existing work. To introduce the model, we first discuss lattice gauge theory in Sec. IC, then higher-lattice gauge theory in Sec. ID. In Sec. IE we use these ideas to motivate the Hamiltonian model from Ref. [56], and explicitly define the model. Throughout the paper, we will be placing different conditions on the model to examine special cases, which we describe in Sec. IF. In the last section of the Introduction, Sec. IG, we describe what we mean by braiding statistics in (3+1)D.

After discussing the background for our work, we then move on to a description of our results. We start in Sec. II by using ideas from gauge theory to motivate the excitations and some of their properties. Then in the rest of the paper we will look at how these properties (and additional features) arise in the lattice model. In Sec. III we construct the operators that create and move the various excitations. Then in Sec. IV we describe how some of the excitations are confined, with a cost to separate two of these particles that grow linearly with separation. We explain how, at least in certain cases, this can arise from a "condensation-confinement" transition between different higher-lattice gauge theory models. After this, in Sec. V we present the braiding relations of the excitations, describing the result of exchanging our excitations in various ways. Finally in Sec. VI we discuss topological charge, a type of conserved charge realized by topological phases, and point out a relation between the allowed values of this charge and the ground-state degeneracy of our model.

C. Lattice gauge theory

In this section, we review lattice gauge theory and Kitaev's quantum double model. This material may be familiar to some readers, who may still wish to read it to familiarize themselves with the notation we use throughout. To describe a continuum gauge theory, the key ingredients are matter fields, gauge fields (which describe parallel transport of the matter fields), and gauge symmetry. As an example, we can consider conventional electrodynamics. In this case the matter fields describe charges, such as electrons, which couple to the usual gauge field. This gauge field describes parallel transport of the charged matter via the Aharonov-Bohm effect [30,31]. Finally, there is a gauge symmetry, which gives us gauge transforms that appear to change the values of the gauge

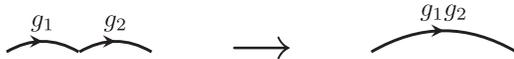


FIG. 1. Composition of paths is described by group multiplication.

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 theories are typically constructed in the continuum, the same ideas can be applied to a lattice gauge theory [73]. The first thing to consider is the physical space on which we consider the model, that is the lattice. Throughout this paper, we will use lattice in the more informal sense, referring to a collection of vertices and edges (and later plaquettes), without requiring a repeating structure (i.e., we consider a graph). Then we have to place the other ingredients of gauge theory into this discrete setting. The matter field is placed on the vertices of the lattice [73], while the gauge field is placed on the (directed) edges of the lattice and determines the result of transport of matter along the edges [73]. However, for the purposes of this paper, we will not include matter as a dynamical field, and charges are instead represented by violations of the gauge symmetry. This leaves us only with the gauge field, which is valued in some discrete group G . The group structure means that two paths that lie end to end can be composed, with the field label of the resulting path given by group multiplication of the labels of the constituent paths [73], as shown in Fig. 1. If we wish to combine two paths that point in opposite directions, we must first reverse the orientation of one of them, so that they align. The group element associated to the reversed path is then the inverse of original group element. For example, if in Fig. 1 the path labeled g_2 pointed in the opposite direction, the combined path would instead have label $g_1 g_2^{-1}$.

1. Gauge transforms

Having considered the fields present in the model, we now look at the gauge symmetry. The gauge symmetry is included through a set of local operators that each act on the degrees of freedom near a vertex. Each operator is labeled 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 [13]. This transform affects the edges surrounding it by premultiplying the group element on each adjacent edge by x if the edge is outgoing, and postmultiplying the element by x^{-1} if the edge is incoming [13,73–75]. We give an example of the action of the vertex transform in Fig. 2, from which we can see that this action is equivalent to adding an imaginary edge, labeled by x , to the vertex v and parallel transporting the entire vertex along it.

This geometric picture reveals two important properties of the vertex transform. First, the vertex transform only affects paths that start or end at that vertex because a path passing through the vertex will travel both ways along the added edge. For example, in the top-left image in Fig. 2 the path entering the vertex v from the lower left and exiting v through the lower right is labeled by the product $g_1 g_2$. In the bottom-left image of Fig. 2, which represents the state after the gauge transform, the same path is labeled by $g_1 x^{-1} x g_2 = g_1 g_2$. That

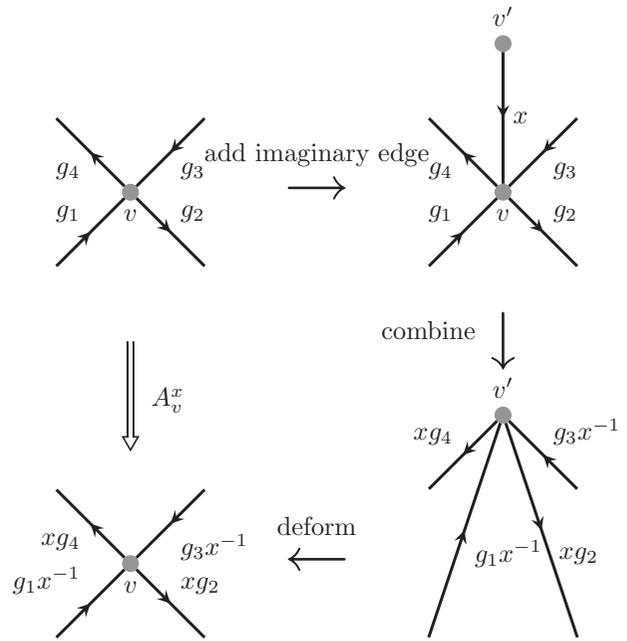


FIG. 2. The gauge transform on a vertex is equivalent to adding an imaginary edge at that vertex and then combining this edge into the diagram, or equivalently transporting the vertex along that edge.

is, the path label is unchanged by the gauge transform because the path does not start or terminate at the vertex v . Second, note that applying two gauge transforms to the same vertex is the same as parallel transporting along two edges in sequence. This is 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 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 Fig. 2). Therefore, we must have that $A_v^h A_v^g = A_v^{hg}$ [13].

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 [74,75]. Under a gauge transform, the group element assigned to a closed loop is at most conjugated by the vertex transforms [75]. Therefore, the conjugacy class of that label is a gauge-invariant quantity. For example, consider Fig. 3, which shows the action of a vertex transform A_v^x on a closed loop starting at v . Initially the group element associated to the closed loop is $g_1 g_2$. 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.

As an example of the importance of such closed loops, we can consider the case of electromagnetism. Here we have a U gauge symmetry, so that the edges in our lattice would be labeled by phases. There is a physical process where we take

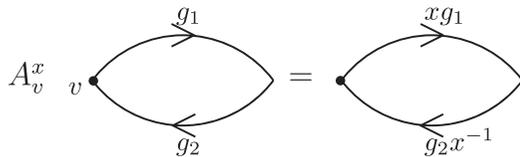


FIG. 3. We consider the effect of a vertex transform A_v^x on a closed loop starting at the vertex v . The path label in this case goes from $g_1 g_2$ to $x g_1 g_2 x^{-1}$. That is, the path label of the closed loop is conjugated by x under the action of the transform.

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 Aharonov-Bohm effect [30,31], where the wave function accumulates a phase of $\theta = q \oint \vec{A} \cdot d\vec{l}$. This phase is the label we would give our closed loop in the lattice model. Using Stoke's theorem, the Aharonov-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.

These gauge-invariant quantities allow us to differentiate between physically distinct states. For instance, many gauge configurations can be reduced to the trivial configuration, where every edge is labeled by 1_G (the identity in the group G), by applying gauge transforms. The state where the edges are all labeled 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 nontrivial 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 nontrivial 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.

3. The quantum double model

Lattice gauge theory can be used to build a model for topological phases, known as Kitaev's quantum double model [13]. 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 gauge invariance to an energetic constraint by adding an energy term to the Hamiltonian for each vertex that enforces the symmetry. We also add an energy term at each plaquette that penalizes plaquettes with nontrivial boundary paths. The Hamiltonian is [13]

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

Here we have

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

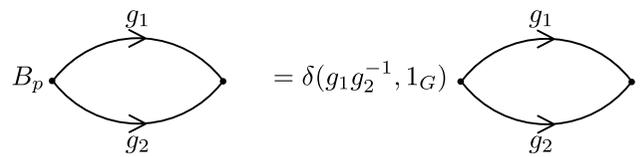


FIG. 4. The plaquette term B_p gives 1 if the closed path forming the boundary of the plaquette p is 1_G (i.e., if it is flat) and 0 otherwise. For the example plaquette shown in this figure, where the edges are in states labeled by g_1 and g_2 , that means that acting with the plaquette term gives a nonzero result only if $g_1 g_2^{-1} = 1_G$.

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 [13] because

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

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 labeled by the identity unchanged and returns zero for other states. As an example, consider Fig. 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 charginelike (excitations of the vertex term), fluxlike (primarily excitations of the plaquette term, though they may also excite a vertex term), or some combination of the two [13]. These excitations are called electric if they are charginelike, magnetic if they are fluxlike, and dyonic if they are a combination. As we will see later, some of these properties will carry over to the higher-lattice gauge theory Hamiltonian model.

D. Higher-lattice gauge theory

In lattice gauge theory we consider parallel transport along paths, and label paths by group elements to allow composition

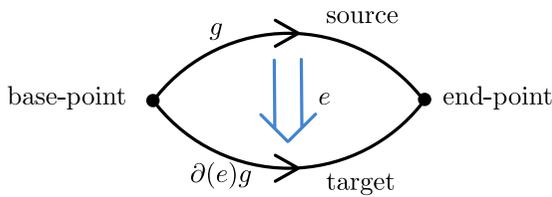


FIG. 5. Just as a path is associated to parallel transport of points, so is a surface associated to parallel transport of a path. The initial position of the path is called the source and the final position is called the target. The parallel transport of a path over a surface labeled by $e \in E$ results in the path element g gaining a factor of $\partial(e)$, where ∂ is a group homomorphism from E to G .

of paths. That is, we label geometric objects, the paths, with algebraic objects, the group elements. A natural generalization is to label more types of geometric objects. We still label the paths with elements of a group G (this is the 1-gauge field, or the 1-holonomy of that path [56]). However, we now also label the surfaces with elements in a second group, E . We refer to this field as the 2-gauge field. As we will see shortly, parallel transport will involve various mappings between the groups E and G . If paths describe the parallel transport of points, then surfaces describe the parallel transport of paths, that is of the 1-gauge fields [64]. We can view this pictorially as shown in Fig. 5. The blue double arrow on the surface enclosed by the paths represents the transport of one path (the source) into another (the target) [56]. Both of these paths must be specified in order to give the surface a label, which is called the 2-holonomy [56] for that surface. The two paths (source and target) both start at a common vertex, called the start point of the surface, and end at a common vertex, called the end point. As indicated in Fig. 5, the parallel transport over a surface labeled by e causes the source to gain a factor of $\partial(e)$, where ∂ is a group homomorphism from E to G (i.e., a map that preserves the group multiplication), so that for $e \in E$, $\partial(e) \in G$ [64]. Normally, the labels of the two paths on either side of the surface are independent variables, but if the label of the source is related to that of the target by this parallel transport rule, then the surface is called fake flat. These fake-flat surfaces play an important role in the theory. Fake flatness replaces the trivial flux condition for the quantum double model and will determine the low-energy space in the topological model. In the rest of this section, we will therefore discuss such fake-flat surfaces unless otherwise mentioned.

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. First, they may be combined vertically [63,64], as shown in Fig. 6. Vertical composition corresponds to the case where we perform two parallel transportations of a path (the top path in Fig. 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. In order to compose the two surfaces in this way, the target of the first surface must match the source of the second one. After composition, 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.

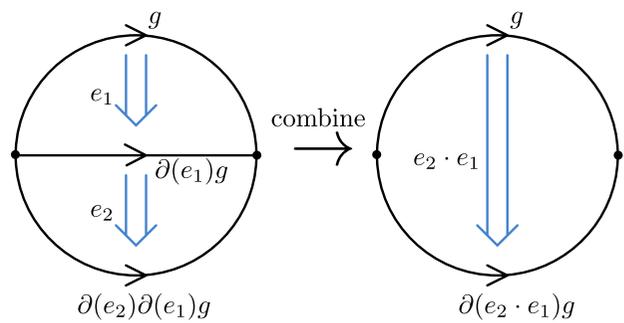


FIG. 6. Consider two surfaces over which we can sequentially transport a path, such as the ones in the left side of the figure. In this case we first transport the top path over the upper surface (across the arrow) to the middle location (the straight path) and then over the lower surface to the bottom 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.

We map vertical composition of two surfaces onto the group multiplication, with the first surface label on the right and the second on the left, following the convention in Ref. [56]. As shown in Fig. 6, requiring the label of the bottom path to be the same on both sides of Fig. 6 gives the consistency condition $\partial(e_2 e_1) = \partial(e_2) \partial(e_1)$, which is why ∂ must be a group homomorphism. This ensures that the effect of transporting the edge along one surface, labeled by e_1 , and then another surface, labeled by e_2 , is the same as transporting the edge along the combined surface (labeled by $e_2 e_1$).

We may also combine the surfaces horizontally [63,64], as shown in Fig. 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.

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 Fig. 8. In Fig. 8, we combine the edge that runs from A to B with a surface, by treating

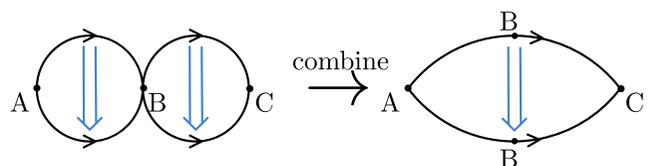


FIG. 7. In addition to vertical composition of surfaces, as shown in Fig. 6, we can consider horizontal composition of two surfaces that lie side by side. In this case, the individual surfaces describe parallel transport of two paths (A-B and B-C) which can be composed, while the combined surface describes 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 labeled B in the right-hand side are the same and so should be glued together.

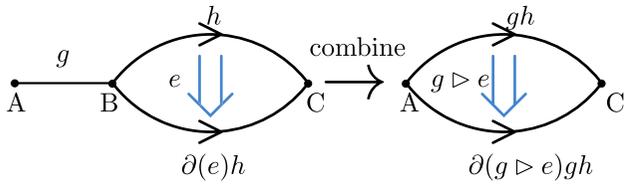


FIG. 8. Here we consider “whiskering” a surface (the bigon with base point B and end point C) along a path, which is described by the map \triangleright . Just as for horizontal composition, the shape of the final surface should match the combined shape of the path and surface, though we have represented 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 of the surface changes from B in the first image to A in the second. 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 $\partial(g \triangleright e) = g\partial(e)g^{-1}$.

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

[64], and can also be thought of as moving the base point of a surface (in the case shown in Fig. 8, the base point of the surface is initially at B, but is moved to A). Because parallel transport of objects along paths is described by the group G , the whiskering must be described by an action of G on E . This action is given by a map \triangleright from G to the endomorphisms on E [56,64] (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 $g \triangleright e$. In Fig. 8, 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 labeled 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.

For the diagrams that we have considered 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 Fig. 9, which shows a diagram involving the horizontal composition of two surfaces (on the top line), where the left and right surfaces are labeled by group elements

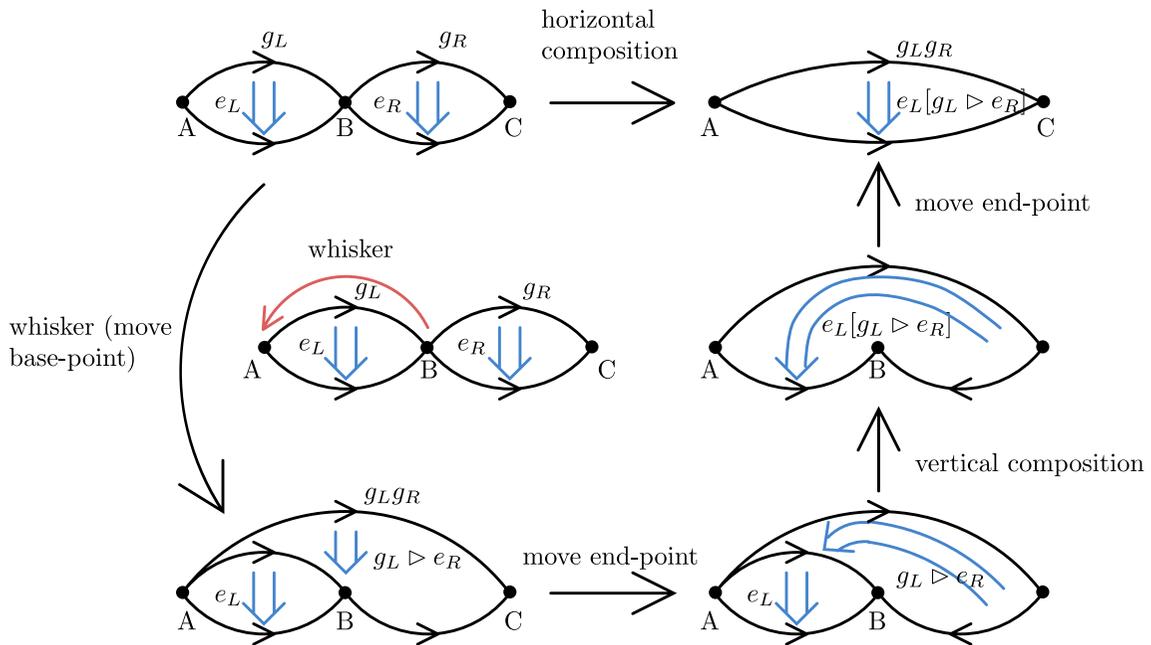


FIG. 9. The requirement that the different ways of combining a diagram must be consistent means that we can express the horizontal composition of two surfaces in terms of whiskering and vertical composition. Consider the top-left image, consisting of a left and right surface with labels e_L and e_R , 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 a 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 center-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.

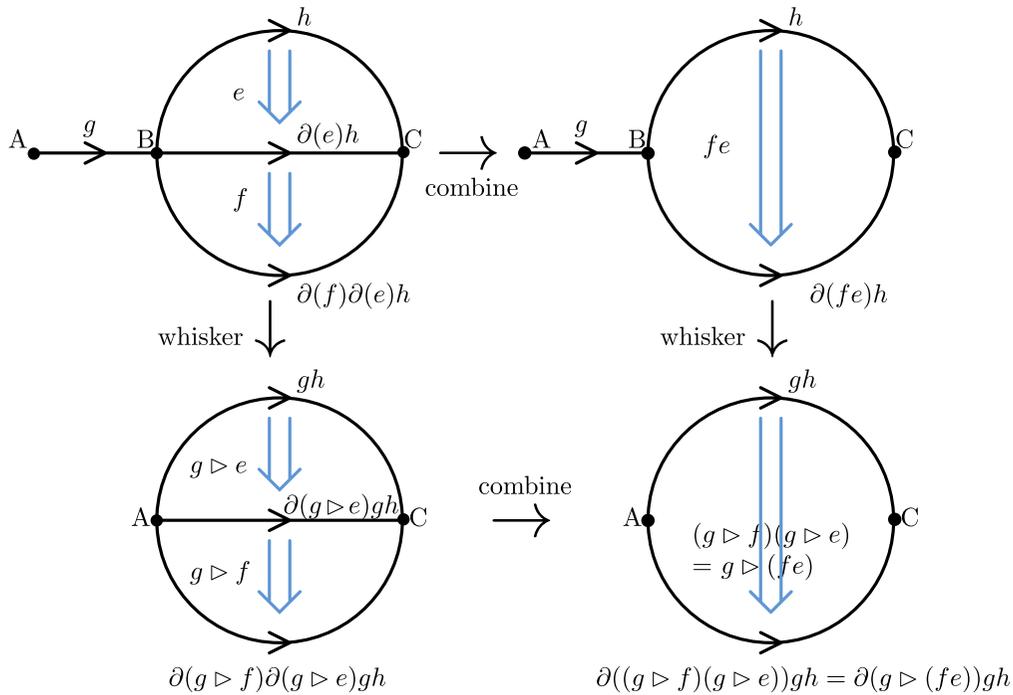


FIG. 10. Requiring consistency of whiskering with the vertical composition of surfaces demands that the map $g \triangleright$, for arbitrary $g \in G$, is a group homomorphism on E . That is $(g \triangleright f)(g \triangleright e) = g \triangleright (fe)$ for $g \in G$ and $e, f \in E$. This can be 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.

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 Fig. 9), we find that the label resulting from horizontal composition must be $e_L[g_L \triangleright e_R]$.

Requiring the consistency of various diagrams also enforces certain restrictions on the algebraic objects we have already discussed. 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 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 Figs. 10 and 11 or Ref. [64] for more detail) gives us the following conditions for all $g, h \in G$ and $e, f \in E$ [56]:

$$g \triangleright (ef) = (g \triangleright e)(g \triangleright f), \quad (2)$$

$$g \triangleright (h \triangleright e) = (gh) \triangleright e. \quad (3)$$

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 Eq. (3), $g^{-1} \triangleright$ is the inverse of $g \triangleright$], they are automorphisms.

As illustrated in Ref. [64] (though note that different conventions are used in this reference and in particular group multiplication describes horizontal composition of surfaces), consistency of whiskering with other diagrams (see Figs. 8 and 12) also demands that [63,76]

$$\partial(g \triangleright e) = g\partial(e)g^{-1}, \quad (4)$$

$$\partial(e) \triangleright f = efe^{-1}. \quad (5)$$

These two conditions are known as the Peiffer conditions [56]. The algebraic structure $(G, E, \partial, \triangleright)$ satisfying all of these conditions [Eqs. (2)–(5) in addition to ∂ being a group homomorphism] is known as a crossed module.

Definition 1. A crossed module is a collection $(G, E, \partial, \triangleright)$, where G and E are groups, and $\partial : E \rightarrow G$ and $\triangleright : G \rightarrow \text{Aut}(E)$ are group homomorphisms satisfying the Peiffer conditions (4) and (5).

In order to familiarize the reader with these crossed modules, we describe a handful of examples here.

Example 1. One example of a crossed module is $(G, G, \text{id}, \text{ad})$, where G is any finite group, id is the identity map, and ad maps $g \in G$ to conjugation by g [56]. That is, we have $\partial(e) = e$ and $g \triangleright e = geg^{-1}$. This clearly satisfies the first Peiffer condition because

$$\partial(g \triangleright e) = g \triangleright e = geg^{-1} = g\partial(e)g^{-1}.$$

It also satisfies the second condition as

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

This crossed module describes a model where all of the excitations are either confined or carry trivial charge.

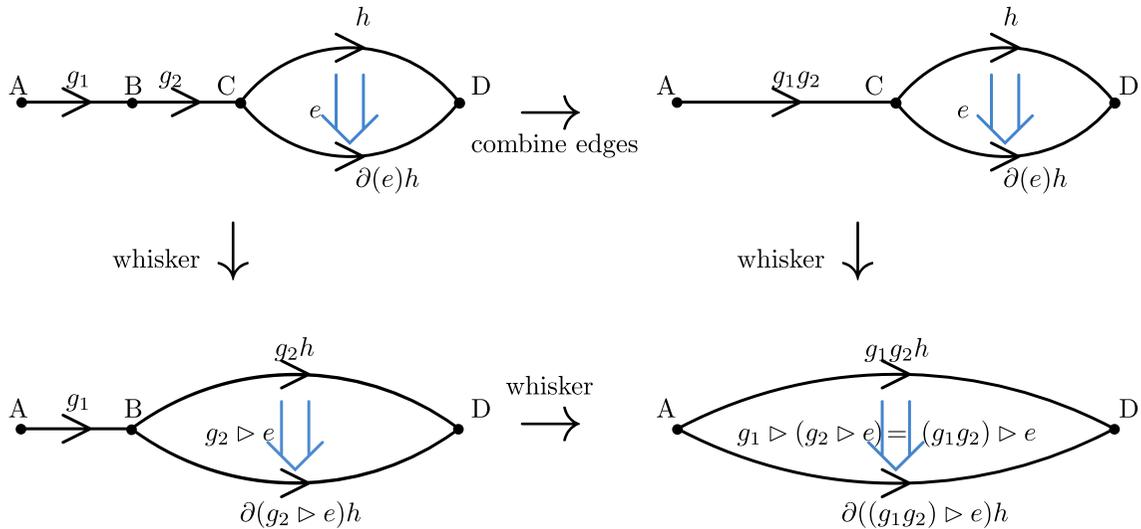


FIG. 11. We require that whiskering is consistent with the composition of edges. That is, combining two edges together and then whiskering a surface along the combined edge should give the same result as whiskering that surface by one edge and then the other in sequence. This gives us the mathematical condition $g_1 \triangleright (g_2 \triangleright e) = (g_1g_2) \triangleright e$, which is the condition that the map $\triangleright : G \rightarrow \text{End}(E)$ be a group homomorphism on G .

Example 2. Another example is $(G, \{1_E\}, \partial \rightarrow 1_G, \triangleright \rightarrow \text{id})$ [56]. That is, we take the group E to be trivial. Then ∂ maps the element of E to the identity of G and $g \triangleright$ is the identity map on E (clearly these are the only allowed ∂ and \triangleright

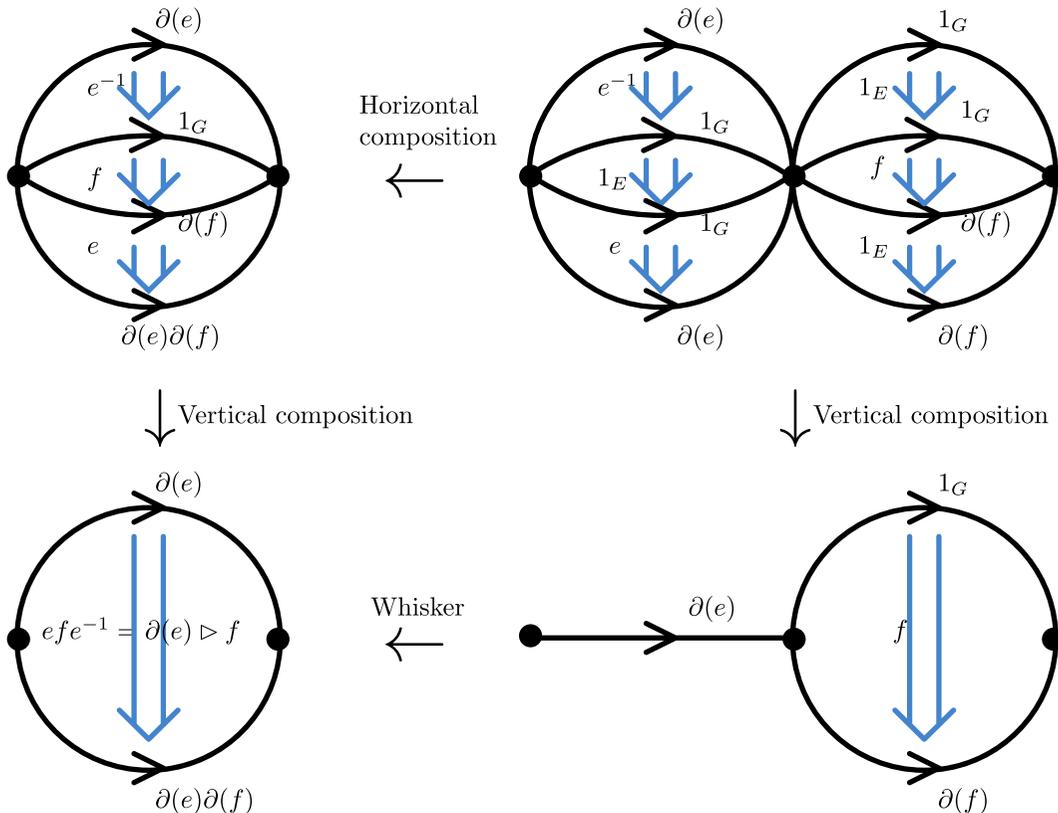


FIG. 12. The second Peiffer condition can be derived from demanding that this figure be consistent. Starting at the top-right diagram, we may combine the left and right parts of the diagram to obtain the top-left diagram, using the rules for horizontal composition given in Fig. 9. We can then use vertical 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 $efe^{-1} = \partial(e) \triangleright f$.

when E is $\{1_E\}$. We have that

$$\partial(g \triangleright e) = 1_G = gg^{-1} = g1_Gg^{-1} = g\partial(e)g^{-1}$$

so the first Peiffer condition is satisfied. Furthermore,

$$\partial(e) \triangleright f = 1_E = efe^{-1}$$

because the only element of E is the identity, so the second Peiffer condition is also satisfied. This special case recovers lattice gauge theory because the surfaces all have trivial label and so we can just neglect to label them.

Example 3. A third, more interesting, example is $(\mathbb{Z}_2, \mathbb{Z}_3, \partial \rightarrow 1_G, \triangleright)$ [56]. We take the elements of $G = \mathbb{Z}_2$ to be 1_G and -1_G and the elements of $E = \mathbb{Z}_3$ to be $1_E, \omega_E,$ and ω_E^2 . Then we define \triangleright by $1_G \triangleright e = e$ and $-1_G \triangleright e = e^{-1}$ (where $\omega_E^{-1} = \omega_E^2$). This satisfies the requirement of having a group structure on the elements of G [as described in Eq. (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. For $-1_G \triangleright$ we have

$$\begin{aligned} -1_G \triangleright (e_1e_2) &= e_2^{-1}e_1^{-1} = e_1^{-1}e_2^{-1} \\ &= (-1_G \triangleright e_1)(-1_G \triangleright e_2), \end{aligned}$$

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 $\partial(g \triangleright e) = 1_G = \partial(e)gg^{-1} = g\partial(e)g^{-1}$ [using $\partial(e) = 1_G$ and the fact that the group G is Abelian]. Finally, $\partial(e) \triangleright f = 1_G \triangleright f = f = fee^{-1} = efe^{-1}$, where we used that E is Abelian. Because all of the consistency conditions are satisfied, this is indeed a valid crossed module. This $(\mathbb{Z}_2, \mathbb{Z}_3, \partial \rightarrow 1_G, \triangleright)$ crossed module can be generalized slightly by replacing \mathbb{Z}_3 with \mathbb{Z}_n , where n is an odd integer, with $-1_G \triangleright$ still acting as inversion.

1. Composing general surfaces

So far, we have considered how we may combine surfaces when their sources and targets are compatible. We can combine two surfaces using vertical composition when the target of one surface matches the source of the other. However, we may also 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. The group element assigned to a surface corresponds to parallel transport of a particular path over that surface. However, we can also pull other paths over that same surface. For example, consider a square, with different paths denoted as the source or target, as shown in Fig. 13. In the left diagram, we consider the process where we transport the top edge (which is the source for the surface) into the bottom three (which form the target).

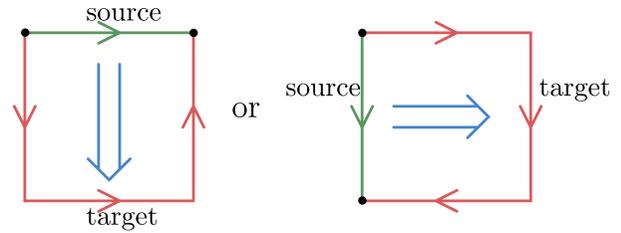


FIG. 13. The same surface can correspond to different 2-holonomies, depending on which parts of the boundary of that surface are designated as the source and target. For example, in the left image the source is the top edge and the target the bottom three edges, so the 2-holonomy corresponds to a process where we transport 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.

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. We therefore need to know how the label changes when we change the transport process. The first thing we can do is to swap the source and target [56]. The resulting plaquette label is just inverted [56], as shown in Fig. 14.

Next, we can move the base point around. We can either move it along the plaquette (as shown in Fig. 15), or away from the plaquette [56] (as shown in Fig. 16). 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 [56].

We can also move the end point, either along the plaquette, or away from it [56]. Either way, the surface label is unchanged [56]. This latter move (as shown in Fig. 17) allows us to add additional edges to the boundary of the surface, though these additional edges enclose no area. Although in Fig. 17 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 traveled in opposite directions for their two appearances, meaning that they do not contribute to the path element of the source or

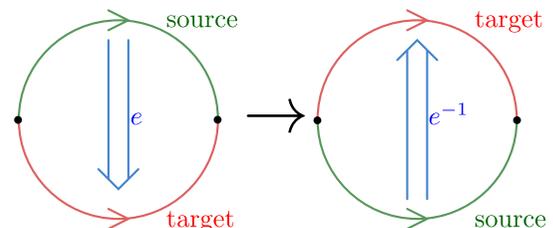


FIG. 14. We can change the source and target of a surface by following a set of rules that tells us how the label of that surface should change. The first rule allows us to swap the source and target of a surface. If the plaquette has a 2-holonomy of e , then swapping the source and target changes the 2-holonomy to e^{-1} .

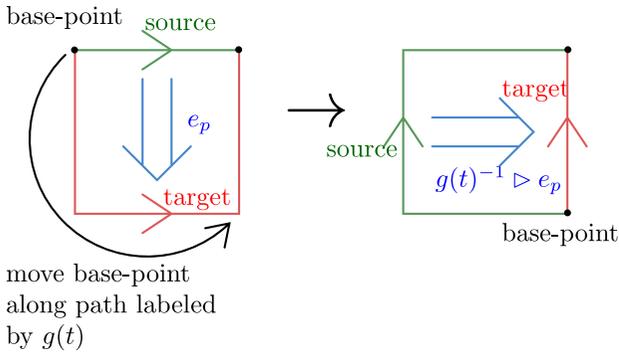


FIG. 15. We can also move the base point of our surface along the boundary of that surface, which adds or removes edges from the start of the source (and removes or adds those edges to the target). This results in a \triangleright action on the surface label.

target [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 target. Either way, their total contribution to the group element associated to the surface boundary (the 1-gauge value assigned to the path around the surface) is trivial. For example, in Fig. 17 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 combine two surfaces when their sources and targets are not immediately compatible. In Fig. 18 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 (labeled by e_2), we first move the end point of the second surface, as shown in the top right of Fig. 18. Because moving

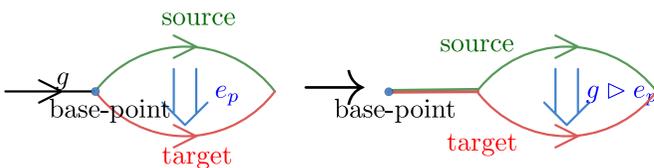


FIG. 16. We can whisker a surface by moving its base point away from the original boundary of that surface. In the right image, the red and green section (which is the black path from the left image) is part of both the source and target.

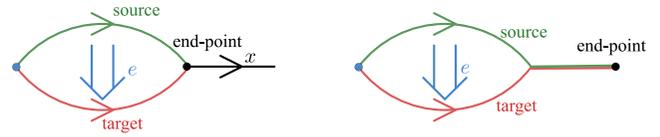


FIG. 17. We can move the end point (the black dot) away from the original boundary of the surface, thereby adding edges (in this case the black path from the left image) to the boundary. These edges do not enclose any area and appear once in the source and once in the target (in the right image, the rightmost edge 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.

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 the fake-flatness condition, meaning that each surface obeys the parallel transport rules given in Fig. 5.

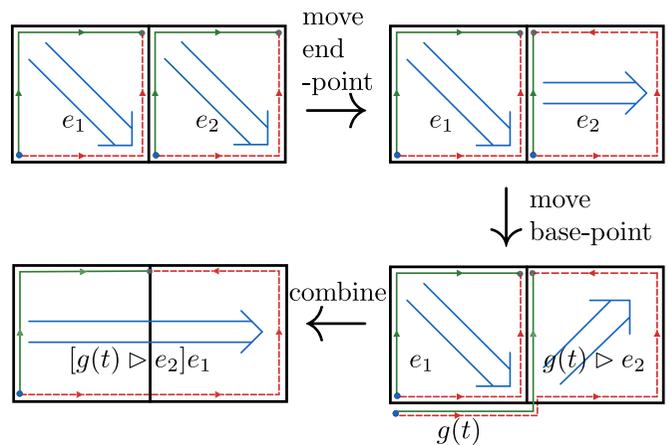


FIG. 18. In order to combine two adjacent surfaces whose sources (shown in green) and targets (shown as red dashed lines) are not compatible, we need to manipulate the end point and base points of the surfaces first. In the first step, we move the end point of the second surface (with label e_2), then 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.

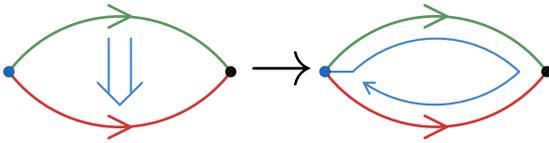


FIG. 19. Instead of illustrating the orientation of the surface as an arrow between the source and target, we can draw it as an arrow that passes clockwise or anticlockwise around the surface, starting and ending at the base point of the surface. The direction of this arrow matches the direction of the source.

2. 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 Fig. 19. Due to the convenience of this notation, we will generally use it when we do not need to indicate the source and target of a surface explicitly.

3. Gauge transforms

Now that we have considered the fields and the parallel transport rules, we can describe the gauge transforms. There are two types of gauge transforms: those associated to the more familiar 1-gauge field and those associated to the 2-gauge field [56]. We label the 1-gauge transforms associated to a vertex v by A_v^g , where there is one such transform for each element g of G . Just like for lattice gauge theory, 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 Fig. 20). The effect on the edges around the vertex is therefore the same as in the lattice gauge theory case and so only paths that start or terminate on the vertex are affected by the gauge transform (see Sec. IC 1 and Fig. 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, as illustrated in Fig. 20. In the last step we relabel the vertex v' to v in order to match the original vertex, so that the lattice is the same at the end as it was before the transform, apart from changes to the group labels. We can recognize the middle diagram in Fig. 20

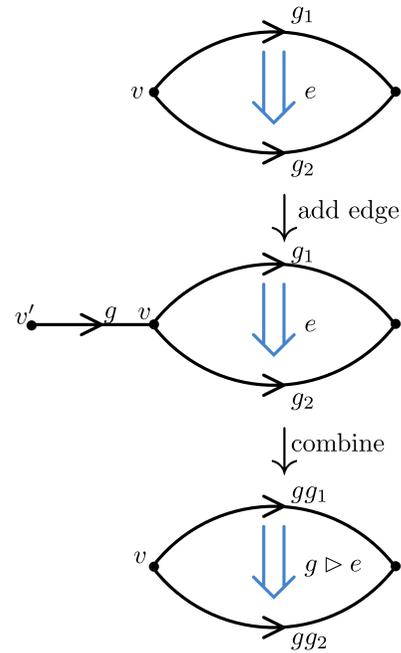


FIG. 20. If the base point of a plaquette is at v , then it is affected by the action of a gauge transform A_v^g on that plaquette. The vertex transform is equivalent to parallel transport of the base point (i.e., whiskering) and induces a $g \triangleright$ action.

as the whiskering diagram (see Fig. 8), 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 [56]. In summary, the 1-gauge transform acts on an edge i or plaquette p according to [56]

$$\begin{aligned}
 A_v^g : g_i &\rightarrow \begin{cases} gg_i & \text{if } v \text{ is the start of } i, \\ g_i g^{-1} & \text{if } v \text{ is the end of } i, \\ g_i & \text{otherwise,} \end{cases} \\
 A_v^g : e_p &\rightarrow \begin{cases} g \triangleright e_p & \text{if } v \text{ is the base-point of } p, \\ e_p & \text{otherwise.} \end{cases} \quad (6)
 \end{aligned}$$

In addition to these 1-gauge transforms, we also have 2-gauge transforms, which act on an edge and the surfaces that adjoin it [56]. The 2-gauge transform on an edge i and labeled by an element $e \in E$ (denoted by \mathcal{A}_i^e) acts like parallel transport of the edge along a surface labeled by e . That is, to find the action of a 2-gauge transform on a diagram we add a surface and combine the surface with the rest of the diagram, as shown in Fig. 21. This 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).

In Fig. 21, 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. To treat a more general case, we can use the rules for changing the base point of a surface to move them to the start of edge i . Then we can perform

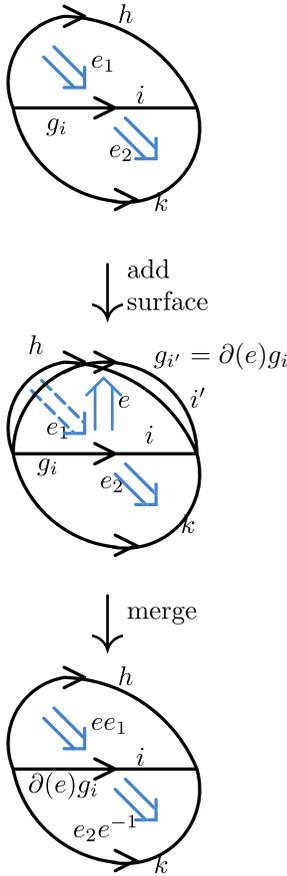


FIG. 21. A 2-gauge transform \mathcal{A}_i^e on an edge i , with initial label g_i , acts like parallel transport of that edge across an additional surface of label e . In the middle picture, the new surface points upwards, out of the plane of the other two surfaces. In the third picture, we combine this new surface with the others and then relabel the edge i' (the target of the additional surface) to i .

the gauge transform \mathcal{A}_i^e on this simple case before moving the base points back to their original positions. 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 $(g \triangleright e)e_p$ [56] 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. In order to define the edge transform, we therefore need a prescription for choosing this path.

Consider the path around one of the plaquettes affected by the transform, starting at the base point v_0 of the plaquette and traveling 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 Fig. 22. The path up to this point is denoted by $g(v_0 - v_i)$ [56]. Now consider a path starting at the base point of the plaquette, but traveling against the circulation of 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 against the circulation of the plaquette [56]. This overline notation is illustrated in Fig. 23, 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

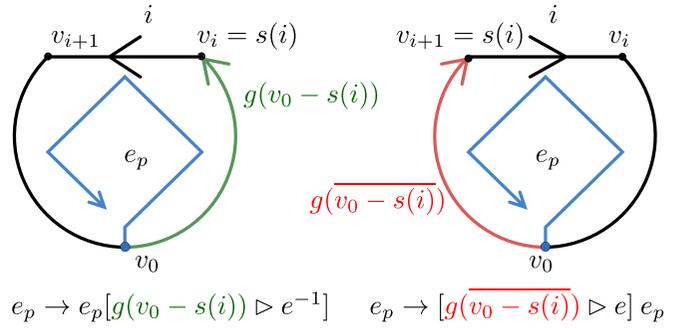


FIG. 22. The path involved in the effect of the 2-gauge transform \mathcal{A}_i^e on a plaquette p depends on whether the edge i is aligned with the p (as in the left case) or antialigned (as in the right case). If the edge is aligned with the plaquette, then the path $[v_0 - s(i)]$ in the transformation of the plaquette label is aligned with p , whereas if i is antialigned with p then the path $[\overline{v_0 - s(i)}]$ appearing in the transformation is antialigned with p . Either way, the path is aligned with the edge i .

p is [56]

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

$$\mathcal{A}_i^e : e_p \rightarrow \begin{cases} e_p[g(v_0 - v_i) \triangleright e^{-1}] & \text{if } i \text{ is on } p \text{ and} \\ & \text{aligned with } p, \\ [g(\overline{v_0 - v_{i+1}}) \triangleright e]e_p & \text{if } i \text{ is on } p \text{ and} \\ & \text{aligned against } p, \\ e_p & \text{otherwise.} \end{cases} \quad (7)$$

The paths involved in the cases where the edge is aligned or antialigned with the plaquette are indicated in Fig. 22. In either case the path terminates at the source of edge i , which is the adjacent vertex which the edge points away from (while the target is the vertex it points towards). This means that we can replace v_i (in the aligned case) or v_{i+1} (in the antialigned case) in the expression for the paths in Eq. (7) with this source $s(i)$.

4. Gauge invariants

In ordinary lattice gauge theory we could build gauge-invariant quantities out of closed loops. What are the

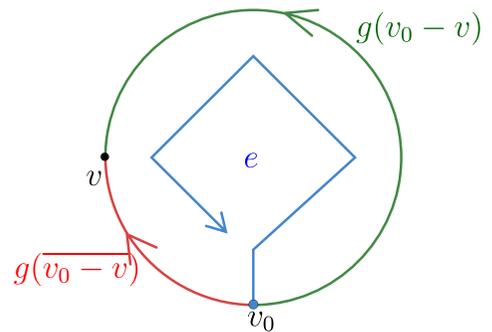


FIG. 23. The two different paths from the base point v_0 to the same vertex v on a plaquette are shown in green and red. Paths that antialign with the surface circulation (represented by the blue arrow in the center) are indicated using overline notation.

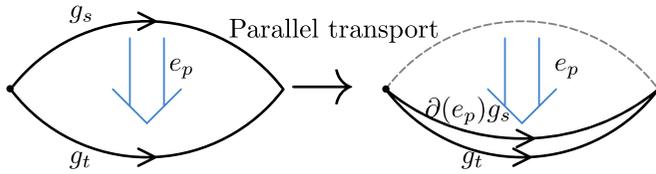


FIG. 24. The gauge invariant associated to plaquettes is modified to account for parallel transport of the paths over the plaquette.

appropriate quantities for higher-lattice gauge theory? We can build gauge invariants from the closed loops as before, but also from closed surfaces. For the closed loops, 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 Fig. 24, 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 Fig. 24.

The parallel transport modifies the group element associated to the closed loop in Fig. 24, 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 $H_1(p)$ [56]. This label can be changed only within a conjugacy class by either the vertex transforms (as in Fig. 20) or the edge transforms (as in Fig. 21), so those conjugacy classes are gauge-invariant quantities [56].

As an example, we can consider acting on the diagram in Fig. 24 with a vertex transform. This gives us the situation shown in Fig. 20. From that figure, we see that the plaquette holonomy, which is initially given by $\partial(e)g_1 g_2^{-1}$, transforms as

$$\begin{aligned} \partial(e)g_1 g_2^{-1} &\rightarrow \partial(g \triangleright e)g g_1 g_2^{-1} g^{-1} \\ &= g \partial(e)g^{-1} g g_1 g_2^{-1} g^{-1} \\ &\quad \text{(using the Peiffer conditions)} \\ &= g \partial(e)g_1 g_2^{-1} g^{-1}, \end{aligned}$$

which is only conjugation.

In addition to closed paths, closed surfaces have their own gauge invariants. The gauge invariant assigned to a closed surface can be found from 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. This label, which we call the 2-flux of that surface, is only changed within certain equivalence classes by the gauge transforms [56] (as long as the constituent plaquettes satisfy fake flatness). Again, the identity element is in a class on its own, so that trivial 2-flux is preserved by the transforms [56].

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

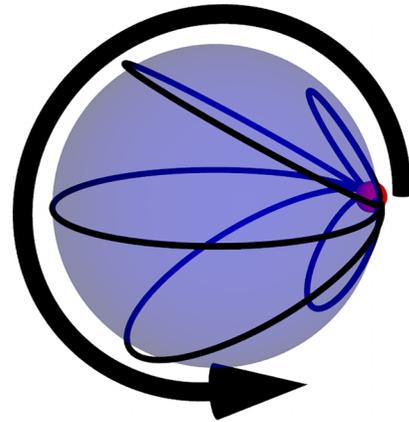


FIG. 25. The 2-holonomy of a surface (in this case a sphere) can be measured by a transport process. A small loop is created at the base point (the small red sphere), then dragged over the surface (the larger blue sphere), as indicated by the arrow.

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 Fig. 25. 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 noncontractible cycles, the corresponding transport process may not involve nucleation and collapse.

E. Hamiltonian model

Having considered higher-lattice gauge theory, we can now define the Hamiltonian model based on it (as introduced in Ref. [56]). The three spatial dimensions of the model 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 [56]. Labelling every edge and plaquette gives a configuration (or coloration). These configurations then form a basis for the Hilbert space, so that a general state is a linear combination of the different labelings 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. As described in Sec. ID 1, we can then use a set of rules to manipulate the source and target of a plaquette to find the label that would be associated to a different process. In order to have this unambiguous reference process, we define a “canonical” position for the source and target paths of every plaquette when we set up the lattice. Because the label of the plaquette is invariant under changes to the end point, it is sufficient to choose a base point and orientation for each plaquette. We also need to choose an orientation for each edge. This can be done formally via a branching structure, which assigns every vertex in the lattice a unique integer vertex. The edges and plaquettes then inherit their data from the vertices involved [56]. The details of

this are not important for our discussion, so we will directly choose the canonical data (orientation and base point) for each edge and plaquette. We will sometimes refer to this choice as the branching structure, or the decoration of the lattice. In the Appendix, we demonstrate how the energy terms change under changes to this branching structure.

To motivate the Hamiltonian considered by Bullivant *et al.* [56], we can take the same approach used for Kitaev’s quantum double model. We first demote the gauge symmetries to energetic constraints, by including them as terms in the Hamiltonian:

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

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

$$A_v = \frac{1}{|G|} \sum_{g \in G} A_v^g, \tag{8}$$

where the action of the gauge transform A_v^g is defined in Eq. (6).

As with Kitaev’s quantum double model, the vertex transforms satisfy $A_v^g A_v^h = A_v^{gh}$ for any $g, h \in G$, which follows from their interpretation in terms of parallel transport (see Fig. 20). The relation $A_v^g A_v^h = A_v^{gh}$ results in A_v being a projector [56], just as for the equivalent term in Kitaev’s quantum double model [see Eq. (1)]. 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 and then using the algebra of the vertex transforms:

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

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

$$\begin{aligned} A_v^g |\psi\rangle &= A_v^g A_v |\psi\rangle \\ &= A_v |\psi\rangle \\ &= |\psi\rangle. \end{aligned}$$

Therefore, the eigenvalue of one for the energy term corresponds to states which are gauge invariant at that vertex.

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 [56]:

$$\mathcal{A}_i = \frac{1}{|E|} \sum_{e \in E} \mathcal{A}_i^e. \tag{9}$$

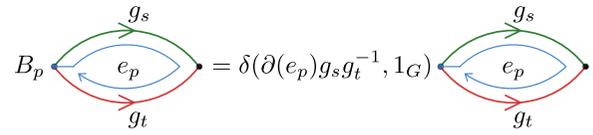


FIG. 26. The plaquette term checks that the flux through a plaquette is trivial. In this case, the boundary of the plaquette is made of two edges, labeled by g_s and g_t , and the boundary label is $g_s g_t^{-1}$. The flux through the plaquette is therefore $\partial(e_p)g_s g_t^{-1}$, where e_p is the plaquette label, and the plaquette term checks whether this expression for the flux is trivial.

The edge terms can be combined in the same way as the vertex terms [56]:

$$\mathcal{A}_i^e \mathcal{A}_i^f = \mathcal{A}_i^{ef}. \tag{10}$$

As with the vertex terms, this means that

$$\mathcal{A}_i^e \mathcal{A}_i = \mathcal{A}_i. \tag{11}$$

This leads to the energy term \mathcal{A}_i being a projector [56], with the eigenvalue of one corresponding to states that are 2-gauge symmetric at that edge. The minus sign with which this term enters the Hamiltonian ensures that the energy term favors these gauge-symmetric states.

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 Sec. ID 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. We can therefore energetically penalize cycles that have nontrivial 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 nontrivial. As explained in Sec. ID 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 Fig. 26. The plaquette terms enter the Hamiltonian with a minus sign, which ensures that the lowest-energy states have trivial flux on the plaquettes. We call plaquettes that satisfy this condition fake flat [56].

Finally, we consider the gauge-invariant quantity associated to the closed surfaces. In the same way as for closed cycles, we penalize closed surfaces with nontrivial 2-flux (2-holonomy). This is done with an energy term at each “blob” (3-cell) [56]. 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 [56], as shown in Fig. 27. We denote the blob term associated to a blob b by \mathcal{B}_b . The blob term also enters the Hamiltonian with a minus sign, so that the full Hamiltonian is given by [56]

$$H = - \sum_{\text{vertices, } v} A_v - \sum_{\text{edges, } i} \mathcal{A}_i - \sum_{\text{plaquettes, } p} B_p - \sum_{\text{blobs, } b} \mathcal{B}_b. \tag{12}$$

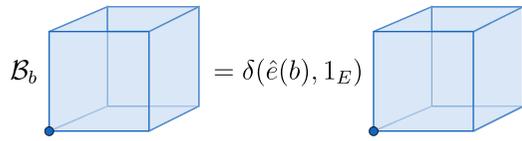


FIG. 27. The blob energy term \mathcal{B}_b checks whether the total surface label of the blob b , $\hat{e}(b)$, is the identity element or not. This surface label must be determined by using the rules for combining surface elements from Secs. ID and IE to combine the plaquettes on the boundary of the 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).

Note that the model can also be defined in (2+1)D, in which case there are no blob energy terms.

Building the Hamiltonian out of gauge transforms and gauge-invariant quantities should mean that the different energy terms commute. However, when fake flatness is not satisfied, the blob terms may not actually be gauge invariant [56]. 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. This means that the energy terms do not commute on the full Hilbert space and the model is not a commuting projector model (and so is not necessarily solvable). This problem will occur in models where \triangleright is nontrivial, 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 regardless [56]. One solution to this problem for nontrivial \triangleright is to define the blob terms to be zero when any of the nearby plaquette terms are not satisfied [56], similar to the approach taken for plaquette terms in the string-net model when the neighboring vertices are not satisfied. However, 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 Sec. IF.

F. Some special cases and consistency

As we mentioned in the previous section, for the most general crossed modules the Hamiltonian model has certain inconsistencies. As an example of this, consider the surface holonomy of a plaquette e_p . We can move the base point of the plaquette all the way around the plaquette and back to its initial position. This induces a change to the surface label of the plaquette, given by $e_p \rightarrow g_p^{-1} \triangleright e_p$ [56], where g_p is the path label of the boundary of the plaquette, as shown in Fig. 28. 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 plaquette is fake flat. In that case, the boundary label satisfies $g_p^{-1} = \partial(e_p)$ and so $g_p^{-1} \triangleright e_p = \partial(e_p) \triangleright e_p = e_p e_p e_p^{-1} = e_p$ [56], where we used the Peiffer condition (5) in the second step. However, if fake flatness is not satisfied, then we cannot guarantee that the plaquette label is unchanged.

This is not the only issue arising from violating fake flatness: as we describe in Appendix A1, the edge energy term also appears to become inconsistent with changes to the

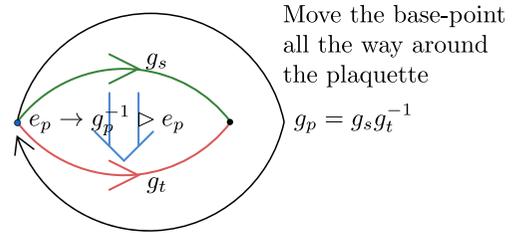


FIG. 28. Moving the base point of a plaquette all the way around the plaquette transforms the surface label from e_p to $g_p^{-1} \triangleright e_p$, where g_p is the path label of the plaquette's boundary.

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 in the paper introducing this model [56]. However, another possibility is to take \triangleright trivial, so that the base point of the plaquette loses any meaning, but allows fake-flatness violations. If we use this condition, then all of the energy terms commute naturally, with no need to restrict the Hilbert space. 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. Some additional consequences of taking \triangleright trivial are that E must be Abelian and that ∂ maps to the center of G . The first condition, that E is Abelian, comes from the second Peiffer condition [Eq. (5) in Sec. ID] because $\partial(e) \triangleright f = e f e^{-1} \Rightarrow f = e f e^{-1}$ so that any elements of E commute with each other. The second condition, that ∂ maps to the center of G , comes from the first Peiffer condition [Eq. (4) in Sec. ID], as $\partial(g \triangleright e) = g \partial(e) g^{-1}$ becomes $\partial(e) = g \partial(e) g^{-1}$, so that $\partial(e)$ commutes with all elements of G . In this paper, both of these special cases (\triangleright trivial and restricting to fake-flat configurations) are considered. We also consider a third case where E is Abelian and ∂ maps to the center of G , but we do not enforce fake flatness on the level of the Hilbert space or require \triangleright to be trivial. In this case, the inconsistencies we mentioned previously are still present, but are not as generic. For example, 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 Sec. IV). We will see other simplifications that occur due to E being Abelian and ∂ mapping to the center of G throughout the paper. 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 I.

TABLE I. A summary of the special cases of the model.

| Case | E | \triangleright | $\partial(E)$ | Full Hilbert Space |
|------|---------|------------------|----------------------------|--------------------|
| 1 | Abelian | Trivial | $\subset \text{center}(G)$ | Yes |
| 2 | Abelian | General | $\subset \text{center}(G)$ | Yes |
| 3 | General | General | General | No |

G. Braiding relations in (3+1)D

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+1)D, it may be useful to give an overview of braiding in (3+1)D, particularly where looplike excitations are involved. In (2+1)D, topological phases support excitations with exchange statistics that generalize the familiar Fermi and Bose statistics [18–20], and which are described by the (colored) braid group. On the other hand, in (3+1)D the pointlike particles can only have Fermi or Bose statistics [22,26,27]. However, the presence of looplike excitations means that we can still have interesting braiding statistics. The motion of loops can be described by the (colored) loop braid group [68–70] (considered under different names and contexts in papers such as Refs. [77–79]). The braid and loop braid groups are both examples of motion groups [80,81], which describe the motions of arbitrary objects, up to homotopy [69] (formally the objects should return to the same positions, perhaps swapping positions).

The loop braid group is generated by a few simple motions. First, consider braiding that involves only two looplike excitations, and for simplicity suppose that those two excitations are stacked vertically, as shown in Fig. 29. 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 Fig. 30, is to simply move the lower loop around and past the upper one, so that neither loop passes through the other [68,70]. We refer to this motion as a *permutation* move because such moves generate the permutation group (symmetric group) [69] that would describe the motions of point particles in (3+1)D. The second way, shown in the left side of Fig. 31, is to move the lower loop *through* the upper loop [68,70]. 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 Fig. 31, is to pull the

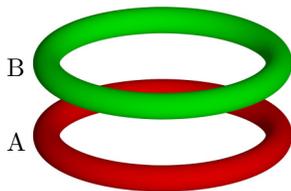


FIG. 29. We first consider the motion of a pair of loops A and B, in the frame where the upper one (B) is held fixed. The motion of a set of loops can be described by such pairwise motions, together with flips of the two loops.

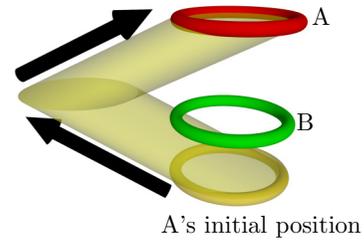


FIG. 30. In the permutation move, we move one loop around the other so that neither loop passes through the other. The motion of the (red) loop A is represented by the (yellow) sheet, with the initial position of the moving loop represented by the lowest (yellow) torus.

lower loop *over* the upper loop, which we can also think of as the upper loop passing 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 Fig. 32. 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) [68,70]. 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.

Generally, we are interested in comparing 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 Fig. 30 to the braiding move in Fig. 31, 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. As we explain in Sec. V, we will be considering processes that involve the production of the looplike excitations from the ground state, rather than just the movement of existing excitations, but the same principles discussed in this section hold regardless.

In addition to looplike particles, the lattice model supports pointlike excitations. While braiding between two pointlike particles in (3+1)D is bosonic or fermionic [26,27] as described earlier (and is exclusively bosonic in this model), the pointlike excitations may braid nontrivially with the looplike excitations of the model. In order to move a pointlike particle

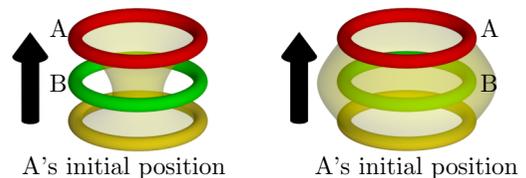


FIG. 31. Schematic of the two braid moves. The sheet swept by the motion of the red loop A is shown in yellow, with the initial position of the moving loop shown as a yellow torus. In the left image, the loop A is passed up *through* loop B. In the right image, the loop A passes over loop B, meaning that loop B passes down through loop A from its perspective. 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).

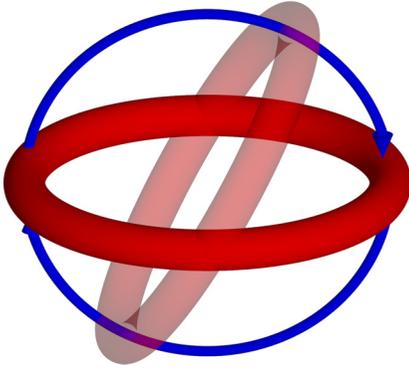


FIG. 32. Unlike point particles, loop-like excitations have an orientation. We can reverse the orientation of a loop by rotating it by π radians out of its plane, which leaves the loop in the same position but flipped over [70].

past a loop excitation, we can either move the pointlike excitation *around* the looplike excitation (analogous to Fig. 30 for two loops, if we replace the initially lower loop A with a pointlike excitation), or we can move the pointlike excitation through the looplike one (analogous to the left side of Fig. 31, if we replace the lower loop A with a pointlike excitation).

II. PROPERTIES FROM GAUGE THEORY PICTURE

A. 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, using ideas from higher gauge theory. As a starting point, we shall briefly review the excitations in ordinary gauge theory, which are electric charges and magnetic fluxes. A clear exposition on these objects and their properties, in the (2+1)D case, is given by Preskill's lecture notes on topological quantum computation [82, Chap. 9] and an early description of non-Abelian magnetic fluxes is given in Ref. [83] (see also Refs. [36,37]). Here we will instead examine the (3+1)D case, as described in (for example) Ref. [84].

Electric charges are point particles labeled by irreducible representations of the group G and excite the vertex gauge terms. 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 . The Hilbert space therefore splits into subspaces that transform as irreducible representations (irreps) of G under the action of the gauge transforms at every vertex. The trivial irrep corresponds to states that are gauge invariant at that vertex, which can be thought of as the absence of an electric charge. On the other hand, if a state transforms as some other, nontrivial, irrep at a particular vertex, then that vertex will be excited. This means that we expect to find excitations that carry some nontrivial irrep of G , with this irrep describing how the excitations transform under the gauge transforms. These excitations should be produced in pairs, with a particle and antiparticle. We denote such a pair, associated to irrep R , by (R, a, b) , where a, b are the matrix indices of the representation and describe an internal space for the pair. The irrep determines the action of the vertex operator applied on one of

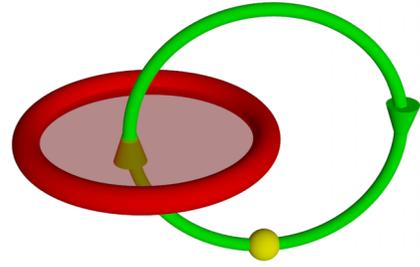


FIG. 33. The magnetic excitation (thick red torus) is measured by a nontrivial closed loop linking with it, such as the thin green line. This closed loop begins at some start point (shown as a yellow sphere) and the value of the flux that we measure depends on which start point we choose, but the conjugacy class of the flux does not.

the particles:

$$A_v^g(R, a, b) = \sum_c [D^R(g^{-1})]_{ac}(R, c, b), \quad (13)$$

where $D^R(g)$ is the matrix representation of element g . 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. In the current prescription, this right multiplication instead describes the transformation of the antiparticle under a vertex transformation at its position. This transformation under the vertex transforms can also be used to tell us something about the transport properties of the excitation. In Sec. ID3 we explained that the vertex transforms are equivalent to parallel transport. Therefore, Eq. (13) tells us how we expect these excitations to behave under parallel transport over an edge labeled by g^{-1} . Looking at Eq. (13), 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 labeled by the representation, while the matrix indices describe some nonconserved details.

In addition to the electric charges, lattice gauge theory hosts magnetic flux tubes. Recall from Sec. IC2 that fluxes are associated with closed paths that have nontrivial labels. In that section, we drew an analogy to how a magnetic field leads to a nontrivial Aharonov-Bohm effect for taking a charge around a closed loop. In a (2+1)D 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+1)D 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 Fig. 33. Then there is no way to smoothly contract the measurement path without

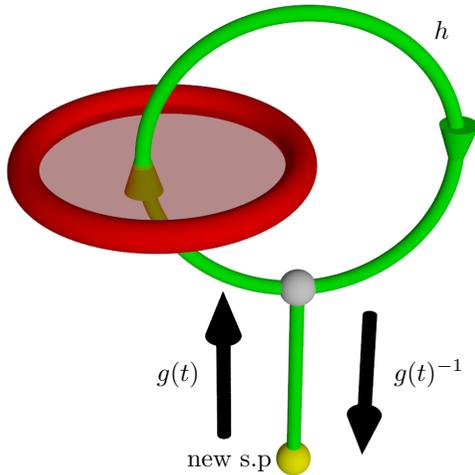


FIG. 34. With the start point (s.p.) at the higher (gray) sphere, the flux label is h , however with the s.p. at the lower (yellow) sphere the flux is $g(t)hg(t)^{-1}$.

it intersecting with the flux tube, and so the label of the path can be a topological quantity.

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, we consider taking a particular closed path that links with that excitation and then changing its start point, as shown in Fig. 34. In order to traverse the new path, we must first travel the path joining the start points, then the original closed path, and then back along the path joining the start points. If the original path has label h , then the new one will have label $g(t)hg(t)^{-1}$, where t is 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.

This geometric picture for the two types of excitations also gives us their braiding relations. We already established how an electric excitation should transform as it moves through space. We can now create a pair of electric excitations, labeled by an irrep R , and move one on a closed path that links with a flux tube labeled by h , as shown in Fig. 35. If this closed path

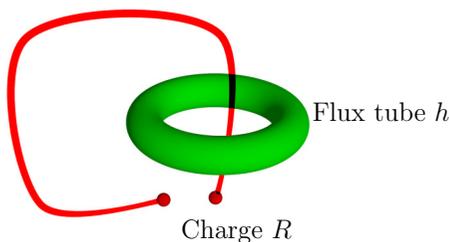


FIG. 35. Schematic view of braiding a charge through a loop. The red line tracks the motion of the charge.

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 object (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 $g(t)hg(t)^{-1}$ to describe the transport of the charge to and from the start point (along path t) as well as around the defining loop of the flux. This gives us the braiding relation

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

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

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 Fig. 36. 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 Fig. 36, which shows the situation after we perform the braiding move, we see that the measurement path for the initially lower flux tube (the blue tube) is pulled through the other (now lower) tube (the red tube). This new deformed path carries the original flux label (h in Fig. 36) 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 Fig. 36). 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 Fig. 36, 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 g of path β . 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 so the new label of the path α is g (which is now associated with the red flux tube). 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' and it keeps its label of h , as

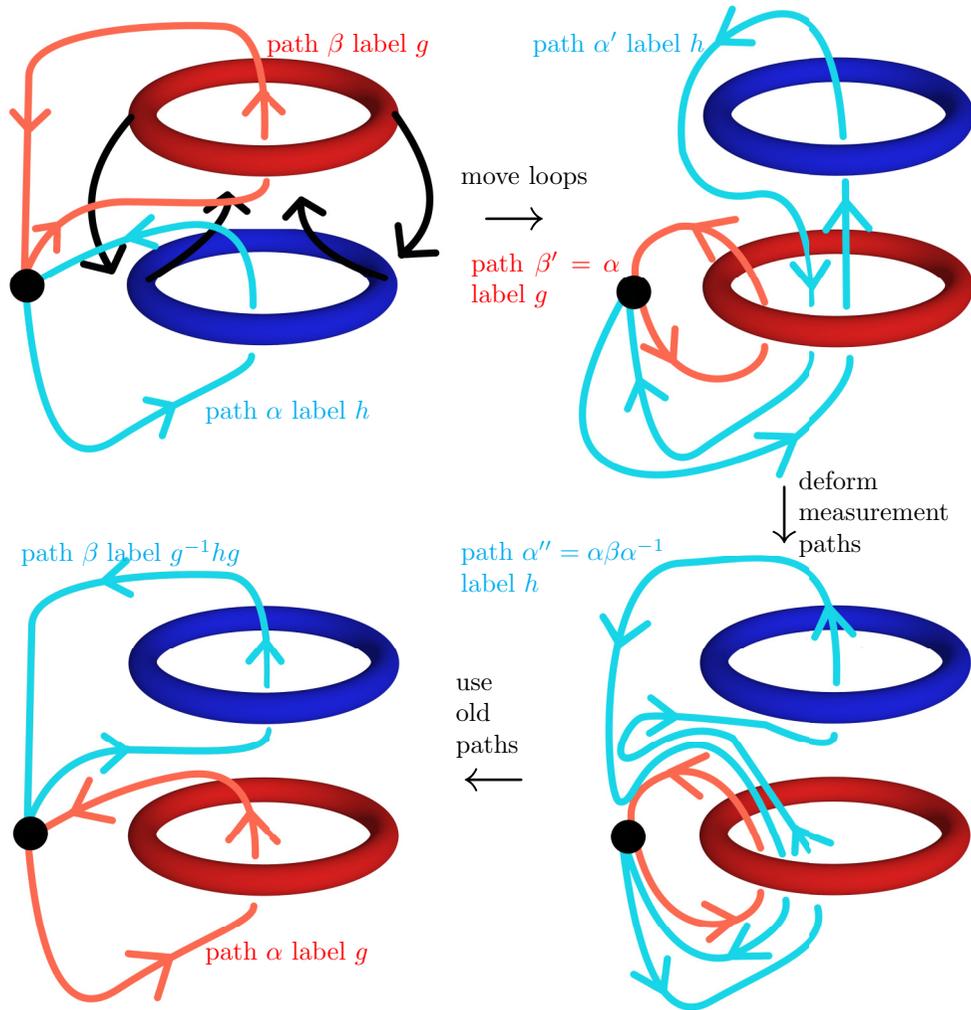


FIG. 36. Starting with two fluxes (top left) we can move the lower one through the upper one, swapping their positions. When we do so, we must also move the measurement paths associated with each flux (top right). Then we can deform these paths (bottom right) to write them as products of the original paths. Knowing the labels for these new paths (which are just the original flux labels of the two loops) allows us to find new labels for the fluxes, when measured along the original paths (bottom left).

indicated in the upper right diagram. 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 bottom-right figure. Therefore, we have $\alpha' = \alpha\beta\alpha^{-1}$ 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 $g^{-1}hg$. 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 $(h, g) \rightarrow (g, g^{-1}hg)$, 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 $h \rightarrow g^{-1}hg$ for the blue tube and $g \rightarrow g$ 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.

B. Higher gauge theory

We have so far described the excitations for ordinary lattice gauge theory, but 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 to find electric charges labeled 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 the quantum double model. Similarly, we expect to find magnetic flux tubes that are similar to those from 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 $\partial(e)$. 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 Sec. IV A, 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 nontrivial label $\partial(e)$ even in the ground state. 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 Sec. IV B.

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 labeled 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” to be extended objects. In fact, we will find that these 2-charges are looplike objects. Then when we transport a loop, labeled by an irrep μ and matrix indices a and b , over a surface labeled by e , we should obtain the transformation $(\mu, a, b) \rightarrow \sum_c [D^\mu(e^{-1})]_{ac}(\mu, c, b)$. There is some subtlety to this, however. Whenever we define a surface element we must give it a base point. If we change the base point we change the label of that surface from e to some $g \triangleright e$ 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 we must define the 2-charges with respect some start point, just like the flux excitations. 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 particularly important for these loop excitations because the action of the group G on the start point, which changes the surface label by some $g \triangleright$ map, enables G to affect the loop excitations. This action can even change the irrep labeling 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 define the classes with the equivalence relation

$$\mu_1 \sim \mu_2 \iff \exists g \in G \text{ s.t. } \mu_1(e) = \mu_2(g \triangleright e) \forall e \in E \quad (14)$$

and when E is non-Abelian we must generalize this to account for irreps related by conjugation. 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 nontrivial loops, there are also “2-fluxes” associated to nontrivial closed surfaces, labeled by elements in E . We expect excitations corresponding to the 2-flux of a sphere to be point excitations because we can shrink a sphere to enclose just a single point. Because these 2-fluxes are measured by closed surfaces, and every surface must be defined with a base point, we must choose a base point for our 2-flux excitation, which we call the start point of the excitation. Note that 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. For the 2-flux excitations the base point of our measurement surface is the start point of the excitation. Moving the base point of a surface along a path labeled 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 it moves along a path. 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

$$e \sim f \iff \exists g \in G \text{ such that } e = g \triangleright f. \quad (15)$$

These “ \triangleright classes” are then conserved under motion.

As this picture gives us the transport properties of these excitations, we can obtain their braiding relations as well. We know how our E -valued loops, the 2-charges, transform under transport over a surface, which tells us how they transform when pulled over the surface assigned to a 2-flux. 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 fluxes. When a 2-flux moves along a path t , the 2-flux label changes from e to $g(t) \triangleright e$. Therefore, when moving a 2-flux labeled by e around an ordinary magnetic flux labeled 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, as explained in Ref. [68], there may be more complicated excitations as well. We may expect to find loop particles that generate both a 2-flux and a 1-flux. The nontrivial 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. [68], the braiding relation between two such loops is established, using geometric arguments and study of the loop braid group. The authors look at the situation where they braid two such excitations labeled 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 labeled by (g, e) is pushed through the one labeled 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 Sec. V, we do

indeed find such excitations with these braiding statistics in the lattice model.

III. EXCITATIONS

The aim of this study is to find the excitations of the higher-lattice gauge theory model and their properties. In the previous section, we gave brief arguments about these characteristics from geometric arguments. However, we wish to show that the lattice model does indeed support such excitations, and give a fuller description of their properties. A significant feature of this model is that we can explicitly find the operators to produce the excitations in various broad cases. Here we will explicitly construct these operators for the higher-lattice gauge theory model, and use the operators to study the excitations directly.

Pointlike excitations are produced by *ribbon operators*. These 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 [13,34]. 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 antiparticle, 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. In (3+1)D, there are multiple types of topological charge, corresponding to different measurement surfaces, as we explain in Sec. VI (and in more detail in Ref. [72]). In this work we will consider the charge measured by a sphere, which captures the pointlike character of an excitation or set of excitations, and the charge measured by a torus, which captures the looplike character. 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.

Instead of being produced at the ends of ribbon operators, the looplike 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

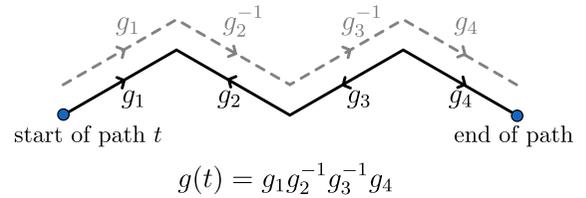


FIG. 37. An electric ribbon operator measures the value of a path and assigns a weight to each possibility, creating excitations at the two ends of the path. In this example, the edges along the path are shown in black. Some of the edges are antialigned with the path and so we must invert the elements associated to these edges to find their contribution to the path element. This is represented by the gray dashed lines, which are labeled with the contribution of each edge to the path.

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 [34]. 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.

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+1)D, two of these types of excitation are pointlike and two types are looplike.

A. 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 Sec. II. 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. An example of this is shown in Fig. 37. In order to measure each possible value of the path label, we apply a ribbon operator of the form

$$S^{\vec{\alpha}}(t) = \sum_{g \in G} \alpha_g \delta(\hat{g}(t), g), \quad (16)$$

where $\hat{g}(t)$ is the path element for path t and α_g is a coefficient (or weight) for the element g . This operator can excite the two vertex terms at the ends of the path t . We call the start of path t the start point of the operator, and it can be thought of as the position where the pair of excitations is created before the excitations are moved.

These electric excitations are equivalent to the electric excitations found in Kitaev’s quantum double model [13], even

up to the precise form of the ribbon operator that creates them. 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 Sec. IV A [and prove in Sec. S-I of Ref. [71], with the proof holding in both (2+1)D and (3+1)D], some of our electric excitations are confined, with the ribbon operator having an energy cost that scales with its length.

Any set of weights α_g that we choose in Eq. (16) 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 Sec. II. Each basis element is labeled 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

$$\hat{S}^{R,a,b}(t) = \sum_{g \in G} [D^R(g)]_{ab} \delta(\hat{g}(t), g), \quad (17)$$

where $D^R(g)$ is the matrix representation of the element g in the irrep R . The operator labeled 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 labeled by a nontrivial 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 topological charge of the excitations. As we show in Ref. [72] (in Sec. IX A 1), the pointlike topological charge for (nonconfined) pure electric excitations is labeled 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 Sec. II). 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. On the other hand, if R is nontrivial in the subgroup $\partial(E)$, then the excitation is confined.

B. Blob excitations

The next pointlike 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 nontrivial closed surfaces we discussed in Sec. ID 4 of the Introduction. That is, the blob excitations correspond to nontrivial “2-fluxes” on a sphere. In order to excite a blob term, which enforces that the surface element of that blob is trivial, we can multiply one of the plaquette labels on that 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

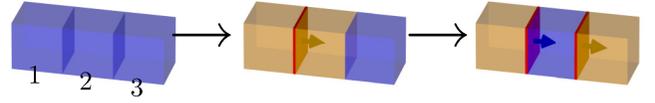


FIG. 38. We consider a series of blobs in the ground state (left-most image). In the ground state, all of the blob terms are satisfied, which we represent here by coloring the blobs blue (dark gray in grayscale). Changing the label of the plaquette between blobs 1 and 2 excites both adjacent blobs, as can be seen in the middle image (we represent excited blobs by coloring them orange, or lighter gray in grayscale). Multiplying another plaquette label on blob 2 to try to correct it just moves the right-hand excitation from blob 2 to blob 3 (rightmost image). In each step, the plaquettes whose labels we changed are indicated by the (red) squares and their orientations are indicated by an arrow.

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 Fig. 38. 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 centers of the blobs (3-cells). The operator that changes the plaquette labels is therefore another of our ribbon operators.

The precise action of our blob ribbon operator depends on which special case from Sec. IF we consider. In the simplest case, where \triangleright is trivial (case 1 in Table I), 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 centers 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 Fig. 38. 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). This results in the two blobs at the end of the path being excited, as we discussed for the example in Fig. 38. As we explain in Sec. IV A, some of the blob ribbon operators [those labeled 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.

We must modify the action of the blob ribbon operators slightly if \triangleright is nontrivial, as in cases 2 and 3 of Table I. When \triangleright is nontrivial, 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 the label of each plaquette by e or e^{-1} 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 Fig. 8 for

a reminder). The total change to e_p is therefore

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

where t is the path from the base point $v_0(p)$ of plaquette p to the start point of our operator. This gives us the action

$$\begin{aligned}
B^e(r) : e_p \\
= \begin{cases} e_p [g(\text{s.p.}(r) - v_0(p))^{-1} \triangleright e^{-1}] & \text{if } p \text{ is aligned} \\ & \text{with } r, \\ [g(\text{s.p.}(r) - v_0(p))^{-1} \triangleright e] e_p & \text{if } p \text{ is antialigned} \\ & \text{with } r, \end{cases} \quad (18)
\end{aligned}$$

where $g(\text{s.p.}(r) - v_0(p))$ is the path element from the start point of our operator to the base point of p and the plaquette is aligned with the ribbon if the circulation of the plaquette can be obtained from the local direction of the ribbon by using the right-hand rule. The path from the start point to the base point of the plaquette can be deformed over a fake-flat region without affecting the action of the operator, so its precise position is not usually important. However, the start point itself is important and the vertex term there may be excited by the ribbon operator. This is because the vertex transform affects the path element $g(\text{s.p.}(r) - v_0(p))$ and so does not commute with the ribbon operator in general. While we leave a more detailed discussion and proof of this for Ref. [72], we can understand this somewhat intuitively using the geometric interpretation discussed in Sec. II. The blob excitations correspond to nontrivial 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 Sec. ID), 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-labeled 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.

C. E -valued loops

The first looplike 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 Secs. ID and IE and apply a weight depending on the result. This is very similar to the electric ribbon operators except that the membrane operator measures 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. 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

$$L^{\mu,a,b}(m) = \sum_{e \in E} [D^\mu(e)]_{ab} \delta(\hat{e}(m), e), \quad (19)$$

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

$$\hat{e}(m) = \prod_{\text{plaquettes } p \in m} g(\text{s.p.}(m) - v_0(p)) \triangleright e_p^{\sigma_p}, \quad (20)$$

where $v_0(p)$ is the base point of the plaquette p , $\text{s.p.}(m)$ 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 depending on the plaquette’s orientation ($+1$ if it matches that of m and -1 if it is antialigned with m). Note that when E is non-Abelian, the order of the product in Eq. (20) is important and must be obtained by applying the rules for composing surfaces given in Secs. ID and IE. There are generally many ways of composing the surfaces and each way is associated to different paths $[\text{s.p.}(m) - v_0(p)]$ and different orders of multiplication, which should give the same result as long as fake flatness is satisfied on the membrane. Applying the membrane operator in Eq. (19) causes the edges along the boundary of the surface to become excited (as long as μ is nontrivial), as indicated in Fig. 39. 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 from Sec. IF that we look at. As we explained in Sec. IF, when \triangleright is trivial, the Peiffer conditions [Eqs. (4) and (5) in Sec. ID] imply that E is Abelian. In this case, the irreps are all one dimensional, so we can drop the matrix indices a and b in Eq. (19). In addition, a trivial \triangleright means that we do not need to keep track of the start point of the membrane.

On the other hand, when \triangleright is nontrivial, E may be non-Abelian and so generally we must include the matrix indices. In addition, the start point of the membrane becomes important and cannot generally be changed without affecting the action of the operator. Much as with the blob excitations, this start point can be excited by the operator, which reflects the nontrivial transformation undergone by the operator when we move the start point (due to the connection between

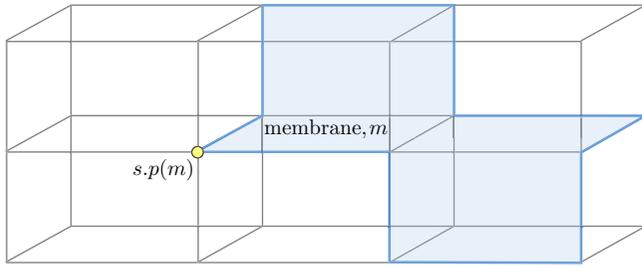


FIG. 39. We consider applying an E -valued membrane on the shaded membrane m in a fragment of the three-dimensional lattice. The membrane operator measures the surface label of the membrane, with a weight for each possible label. When measuring a surface, if \triangleright is nontrivial we must specify 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 nontrivial 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.

vertex transforms and parallel transport). The start point is not excited if the membrane operator is made of a linear combination $\sum_e \alpha_e \delta(\hat{\triangleright}(m), e)$ whose set of coefficients α_e is a function of \triangleright class (i.e., $\alpha_e = \alpha_{g \triangleright e}$). If the start point is not excited, then it can be moved without affecting the operator. On the other hand, the start point may be excited if the coefficient α_e transforms nontrivially under the \triangleright action, as we describe in more detail in Refs. [71,72]. The start-point excitation is significant because it can carry a (pointlike) topological charge, which must be balanced by a pointlike charge on the loop itself, as we prove explicitly in Ref. [72], in Sec. IX A 1. Furthermore, as we describe in Ref. [72] in Sec. V, this pointlike charge can be confined for certain membrane operators, 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 I.

D. 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 [13] and its analogy to magnetic flux. The magnetic loop excitation is primarily associated with excitation of the plaquette energy terms. Recall from Sec. IE (see Fig. 26) 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 must 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 Sec. IIIB). We therefore consider trying to excite a single 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 sur-

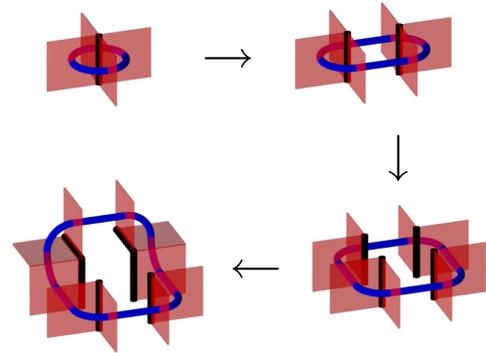


FIG. 40. In order to excite one of the plaquettes in the lattice and produce a magnetic excitation, we change the label of one of the edges (black cylinders) on the boundary of the plaquette. However, this excites all of the plaquettes adjacent to that edge, as shown in the first image (the excited plaquettes are shown in red). Note that these plaquettes lie on a closed loop (blue tube) through their centers. 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. 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).

rounding this edge. 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 always get a closed string of excited plaquettes, as shown in Fig. 40 (unless we collapse the loop to nothing). This means that the magnetic excitation is indeed looplike. 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 Fig. 41. 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.

Next we will explicitly describe 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 I). In this case the operator to produce the excitation is analogous to the (2+1)D magnetic ribbon operator from Kitaev’s quantum double model [13]. We denote the magnetic membrane operator labeled by an element $h \in G$ 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 the labels of 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. On the other hand, when the group is non-Abelian, we must multiply each edge by some element in the same conjugacy class as h (or the inverse). 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 cut by the membrane. Denoting the path to

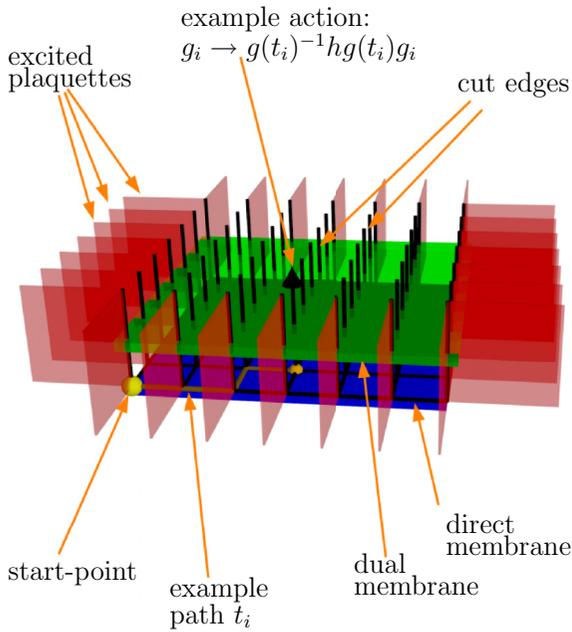


FIG. 41. Here we give an example of the membranes for the flux creation operator (magnetic membrane operator). The dual membrane (green) cuts through the edges changed by the operator. The direct membrane (blue) contains a vertex at the end of each of these cut edges (such as the orange sphere). A path 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.

edge i by t_i , the action of the membrane operator on edge i is

$$C^h(m) : g_i = \begin{cases} g(t_i)^{-1} h g(t_i) g_i & \text{if } i \text{ points away from} \\ & \text{the direct membrane,} \\ g_i g(t_i)^{-1} h^{-1} g(t_i) & \text{if } i \text{ points towards} \\ & \text{the direct membrane.} \end{cases} \quad (21)$$

This action is shown in Fig. 41. 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 [13], which is a thickened string.

The precise choice of the paths on the direct membrane is not usually significant, as we can deform them over a fake-flat region without affecting the action of the membrane operator. 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. When \triangleright is trivial, elements of $\partial(E)$ are in the center of G (see Sec. 1F) and so do not affect the expression $g(t_i)^{-1} h g(t_i)$. Because we usually apply membrane operators on regions without any other excitations, this means that we do not generally need to specify the precise positions of the paths. An important exception that we describe in Ref. [72] is when we produce two linked magnetic excitations by applying intersecting membrane operators.

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 [13], and the potential start-point excitation is also present for the magnetic ribbon operators in that model (see, for example, Ref. [85]). We can interpret the start point of the magnetic membrane operators in this model, as well as the start points of the magnetic ribbon operators in Kitaev's quantum double model, in terms of gauge theory. Recall from Sec. II that whenever we measure a flux, we must do so with respect to a certain start point. The flux created by $C^h(m)$ is only h when we measure with respect to the start point of the membrane operator (or ribbon operator for the quantum double model). Measuring the flux from a different point gives us a result of $\hat{g}(t) h \hat{g}(t)^{-1}$, 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 is then connected to whether the start point of a magnetic membrane operator is 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 conjugates 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. [72]. 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 coefficients 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 pointlike charge [or in the (2+1)D case such as the quantum double model, a pair of excitations that can be annihilated, as described in Ref. [82]].

So far this consideration of the magnetic excitations has all been in the case where \triangleright is trivial. In the case where \triangleright is nontrivial but we restrict to fake-flat configurations (case 3 in Table I), 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 case 2 from Table I, where we loosen the restrictions on the crossed module without throwing out the non-fake-flat configurations. Specifically, we require that ∂ maps to the center of G and that E is Abelian. In this case, we are allowed to keep the magnetic excitations, though their operators must be modified. We briefly describe this modification here, but give a full description in Ref. [72]. The new membrane operators act on edge elements in the same way as described above in the \triangleright trivial case, but they also act on the plaquette elements around the membrane in two ways. First, the membrane operator directly affects the plaquettes that are cut by the dual membrane. If a cut plaquette p has its base point on the direct membrane, then its label e_p is changed to $(g(t)^{-1} h g(t)) \triangleright e_p$,

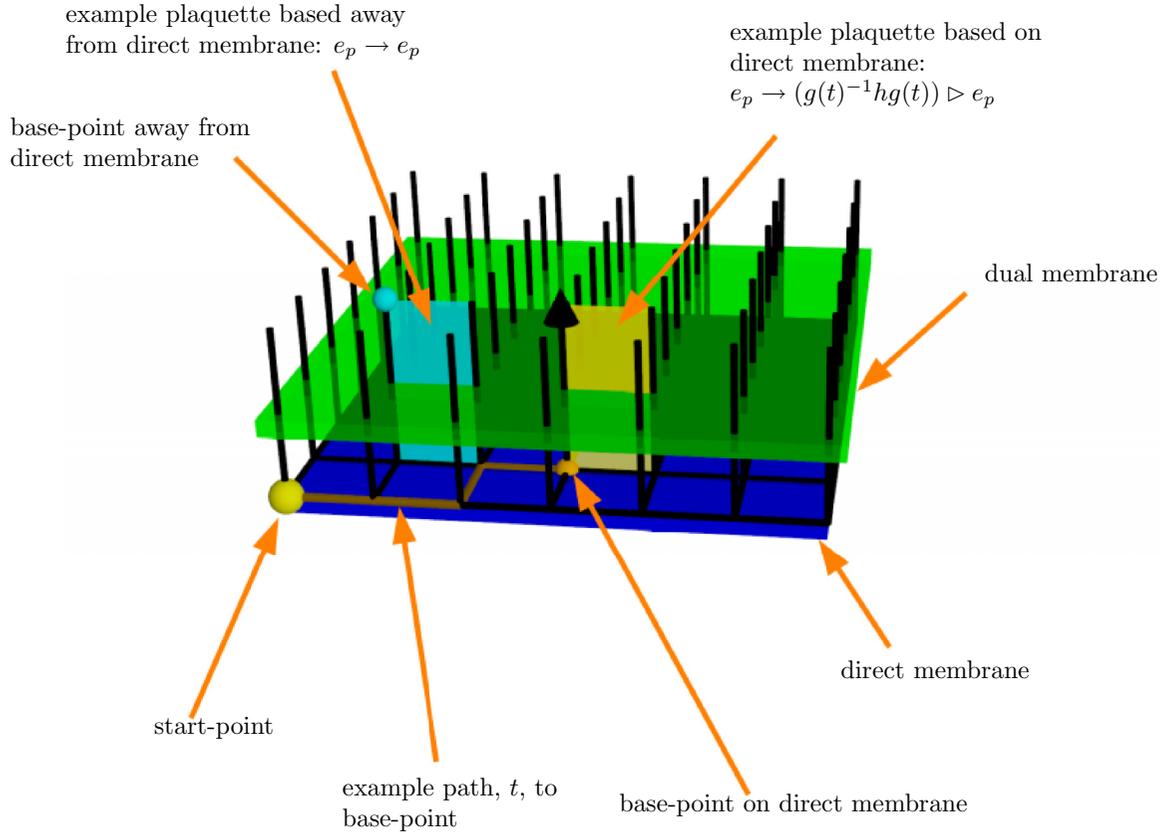


FIG. 42. In addition to changing the edges cut by the dual membrane, when \triangleright is nontrivial the magnetic membrane operator affects the plaquettes cut by the dual membrane if their base points lie on the direct membrane.

where $g(t)$ is the path from the start point of the operator to the base point of the plaquette. Any plaquette whose base point is away from the direct membrane is left unaffected. An example of this action is shown in Fig. 42. 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. Indeed, 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 (it is this equivalence that leads to the membrane operators being topological, as we explain in Ref. [72]).

In addition to this \triangleright action, we have to multiply the membrane operator by blob ribbon operators, 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. In the \triangleright trivial case (case 1 of Table I), the plaquettes around the membrane were excited and the start point could also be excited. In this more general case (case 2 of Table I) the special blob, blob 0, may also be excited, as may the edges and blobs surrounding the membrane. The word “may” is important for the edges and blobs at the boundary of the membrane because, 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. 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, which 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 Sec. II. Recall that a nontrivial 2-flux indicates a closed surface with a nontrivial label. We can see that the loop excitation must be associated to a nontrivial surface when blob 0 of the membrane operator is excited. This is because blob 0 being excited indicates that this blob carries a nontrivial 2-flux, which must be balanced by a 2-flux belonging to the looplike excitation itself.

IV. CONDENSATION AND CONFINEMENT

A. Confinement

As we alluded to in Sec. III, we found that some of the excitations in the higher-lattice gauge theory model are confined, meaning that there is an energetic cost to separating particles (or growing and moving looplike excitations) which grows at least linearly with separation (or with the area swept by the loop). Specifically, some of the pointlike (electric and blob) excitations are confined.

For the blob excitations, the mechanism for this confinement is the plaquette terms. Recall from Sec. III B that the

blob ribbon operator $B^e(t)$ multiplies each plaquette p pierced by the ribbon by an element $[g(\text{s.p.}(t) - v_0(p))^{-1} \triangleright e^{\pm 1}]$ 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. This is because 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 $\partial(e_p)g_p = 1_G$. If the ribbon operator changes $\partial(e_p)$ for the plaquettes that it pierces, it will excite those plaquettes. This occurs when $\partial\{[g(\text{s.p.}(t) - v_0(p))^{-1} \triangleright e]\}$ is nontrivial. The Peiffer condition (4) states that $\partial(g \triangleright e) = g\partial(e)g^{-1}$, and so $\partial\{[g(\text{s.p.}(t) - v_0(p))^{-1} \triangleright e]\}$ is the identity element when $\partial(e) = 1_G$ regardless of the value of $g(\text{s.p.}(t) - v_0(p))$. That is, the blob ribbon operator $B^e(t)$ excites every plaquette it pierces when e is outside the kernel of ∂ . In this case, the ribbon operator has an energetic cost that grows linearly with the length of the ribbon and so the associated blob excitations are confined. Note that for case 3 in Table I, where we exclude states that excite the plaquette terms from the Hilbert space, we must not allow the confined blob excitations.

We also find that some of the electric excitations are confined. The mechanism of this confinement is the edge terms along the ribbon. This is because the edge transforms change the path element measured by the electric ribbon operator by an element in $\partial(E)$ (the image of ∂) and so may fail to commute with the ribbon operator. To determine which electric ribbon operators are confined, we use the irrep basis for the ribbon operators. Given an electric ribbon operator labeled by irrep R of G , we can determine whether the corresponding excitations are confined by evaluating the irrep R on elements of the normal subgroup $\partial(E)$ and treating this as a representation of the subgroup. Restricting the irrep R to the subgroup in this way produces a generally reducible representation of the 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 center of G]. If R branches to the trivial irrep in this way, then the ribbon operator will transform trivially when we alter the path element by an element in $\partial(E)$. This means that the ribbon operator commutes with the edge transforms and the electric excitations are not confined. On the other hand, if R does not branch to the trivial irrep then the ribbon operator does not commute with the edge transforms. In particular, it transforms as a nontrivial irrep of $\partial(E)$ under them and so gives zero when we act with the edge energy term, which is an average over all labels of the edge transform, due to the grand orthogonality theorem. This means that all of the edges along the ribbon are excited and the corresponding excitations are confined. We note that this is equivalent to the confinement for the field theory discussed in Ref. [65], where the electric operators are confined if they can detect factors in a subgroup $\pi_1(H)$ [equivalent to $\partial(E)$ here].

Some of the E -valued looplike excitations can also be confined in a certain sense when E is non-Abelian, as described in Sec. III C. This confinement does not give an energy cost to growing the looplike excitation, but instead it costs energy to move the excitation away from the start point of the membrane. As we discussed in Sec. III C, it seems like

the pointlike charge carried by the loop excitation is confined, rather than the looplike charge. This is also reflected in the topological properties of the membrane operator. Normally, the creation operator (ribbon or membrane operator) for a confined excitation is not topological because the position of the operator can be detected by the energy terms it excites. In the case of the membrane operators producing the particular confined looplike excitations in this model, however, the membrane operator is still partially topological: we can deform the membrane without affecting the action of the operator, but we must keep the location of the excited edges fixed (and so the location of the confining string is fixed). This is again because it is the pointlike charge that is confined (and so the motion of the pointlike charge is not topological).

B. Condensation

The phenomenon of confinement is closely related to a process known as condensation. Consider a topological model with some set of topological charges. By deforming the Hamiltonian, we may find that some of the nontrivial topological charges from the old Hamiltonian are present in the ground state of the deformed Hamiltonian. Because the ground state corresponds to the topological vacuum, this means that the excitations which carried the previously nontrivial charges now carry the trivial charge, although they may still remain energetic. We say that those excitations *condense* [87–91]. When this occurs, any excitations that braided nontrivially with those excitations in the original become confined in the deformed model [87,88] (this is a bit of a simplification, but is sufficient to describe this model). This process is known as a condensation-confinement transition. While this is fairly well understood in (2+1)D, there has been comparatively little study of such transitions in the (3+1)D case (examples of work in this area include Refs. [92,93]), and so it is interesting to see how condensation and confinement arise in this (3+1)D model.

In the higher-lattice gauge theory models, we can consider the process of condensation by constructing two models related by such a transition. That is, we consider cases where we can turn the confinement on and off by changing a parameter in the Hamiltonian. We cannot change the groups G and E , which also fix the Hilbert space, but we can change the map ∂ . When E is Abelian we can construct a model with no confinement, described by the crossed module $(G, E, \partial \rightarrow 1_G, \triangleright)$, and called the uncondensed model. Here $\partial \rightarrow 1_G$ indicates that ∂ maps to the identity of G . The condition that E be Abelian is required from the second Peiffer condition (5), which 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 also trivial: $efe^{-1} = 1_G \triangleright f = f$. Starting from this uncondensed model we can “turn on” ∂ (while keeping the groups fixed), by moving into a model described by a crossed module where ∂ maps to some nontrivial subgroup of G . This can be done by interpolating between the two Hamiltonians because the two models have the same Hilbert space. During this process, some of the excitations condense. When this occurs, the condensing

excitations may still cost energy in the new model, but they carry trivial charge in this condensed phase. We therefore refer to such excitations in the condensed phase as condensed excitations. These condensed excitations can be produced by an operator local to the excitation. For looplike excitations this means that a condensed looplike excitation can be produced with an operator that only acts near the loop, rather than on an entire membrane. Indeed, finding such an operator to produce the loop excitation is one way to show that the excitation is condensed.

In this model, we find that some of the looplike excitations are indeed condensed. Recall from Sec. III D that the magnetic excitations are labeled by elements of the group G . If that element is in the image of ∂ , then the excitation is condensed. This can be seen from considering the magnetic excitations in light of the plaquette term. We discussed in Sec. II that the magnetic excitations are associated with closed loops that have nontrivial 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 just being the identity element. 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 ∂). However, note that this does not mean that these magnetic excitations are not excitations at all. Changing the path label by an element of $\partial(E)$ still leads to plaquette excitations because the path labels do not match the surface enclosed by the loop. In Ref. [72] (in Sec. S-III) 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 looplike excitation is condensed. Like the confinement of the electric excitations, this condensation is equivalent to that for the field theory discussed in Ref. [65], where fluxes in the subgroup $\pi_1(H)$ [equivalent to $\partial(E)$] are condensed.

Some of the E -valued loops are also condensed. This is because some of 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 [labeled 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

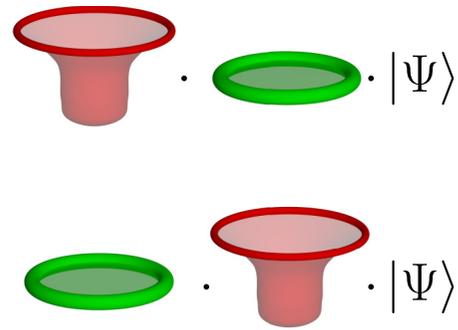


FIG. 43. The commutation of operators used to calculate the braiding. The partially transparent surfaces indicate the membranes for the operators, while the opaque loops indicate the excited regions, which are the boundaries of the membranes.

label. 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. In this case, the corresponding loop excitation can be produced by an electric ribbon operator that only acts near the loop itself, indicating that the loop cannot carry nontrivial loop topological charge and is condensed.

We can be more precise about this notion of sensitivity to the kernel by using the irrep basis for the membrane operators. Recall from Sec. III C that the E -valued loops are labeled by irreps of E , as seen in Eq. (19). We can construct an irrep of the kernel of ∂ (which is a subgroup of E) by restricting the irrep of E to the kernel. This results in a (generally reducible) representation of $\ker(\partial)$. The kernel is always central in E due to the second Peiffer condition (5), so Schur's Lemma applies. This means that the matrix representation of any group element in the kernel must be a scalar multiple of the identity, with the scalar being an irrep of the kernel. If this irrep of the kernel is trivial then the excitation is condensed, otherwise it is not condensed. That is, for an excitation produced by a membrane operator labeled by an irrep μ of E , if the matrix representation satisfies $D^\mu(e_K) = \mathbf{1}$ for all e_K in the kernel then the excitation is condensed.

V. BRAIDING

In this section we will first explain how braiding relations can be obtained from the ribbon and membrane operators, then describe our results for this model. 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 across the membrane. Therefore, we should be able to find the result of braiding by applying successive membrane operators. In particular, the braiding is related to the commutation relations of the membrane operators. For instance, consider the commutation relation shown in Fig. 43, 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 Fig. 31. In the second line of Fig. 43, 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.

There are some subtleties when determining the braiding relations, which mean that we have to take care when interpreting the commutation relations between membrane operators. The formulas given in this section will refer to “same-site” braiding. This refers to the case where the start points of the operators (as defined when discussing the excitations in Sec. III) are in the same location. In this case the excitations involved have a definite fusion product (i.e., well-defined combined topological charge). In a non-Abelian anyon theory, where two anyons may fuse to multiple different types of anyon, the braiding of two particles depends not only on the charges of the two particles, but also which charge they fuse to [21, Preliminaries]. To find the braiding relations we therefore want to consider the case where they have definite total charge, otherwise the braiding relations are not well defined (which in practice is reflected in this model by the presence of operator labels in the braiding relations). In simple cases, where the fusion is Abelian (i.e., where two charges only have one fusion channel), this requirement for same-site braiding is lifted (and this is usually reflected by the start point of the corresponding membrane or ribbon operators being irrelevant to the commutation relation). While we only discuss the same-site case here, we will give more general results in Ref. [72].

Having discussed the method for finding the braiding, we now discuss our results. First, we note that any braiding of two pointlike excitations (which we term permutation) is trivial. 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+1)D, the ribbon operators that permute the two excitations do not need to intersect and so commute, which leads to trivial (bosonic) exchange statistics. That is, the result of moving a pointlike excitation past another one is always the same as moving that excitation through empty space (and then producing the other pointlike particle), as long as the excitations stay well separated. Even if the ribbons do intersect, they can be deformed away from one another using the topological property without affecting their action. This contrasts with the (2+1)D case where crossings cannot always be removed without pulling a ribbon over an excitation. For the same reason, exchange involving looplike 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.

Next, we consider the braiding where we move a pointlike or looplike excitation through a looplike excitation, as shown in Fig. 35 for the pointlike case. First consider the case where \triangleright is trivial. In this case, the excitations split into two separate sets, with nontrivial braiding only within each set. First, we have the excitations that are labeled by objects related to G , namely, the electric and magnetic excitations. We have nontrivial point-loop braiding between the electric and magnetic excitations, where the magnetic excitation labeled by a group element h acts on the electric excitation, labeled by an irrep R and indices a and b , by multiplication by the matrix $D^R(h)$. That is, given an electric ribbon operator $S^{R,a,b}(t) = \sum_g [D^R(g)]_{ab} \delta(\hat{g}(t), g)$ and a magnetic membrane operator $C^h(m)$, commuting the operators gives the following braiding relation:

Flux-charge braiding relation

$$S^{R,a,b}(t)C^h(m)|GS\rangle = C^h(m) \sum_{c=1}^{|R|} [D^R(h)]_{ac} S^{R,c,b}(t)|GS\rangle. \quad (22)$$

This braiding mixes ribbon operators labeled by different indices, but not operators labeled by different irreps. This is because R labels a conserved charge, whereas the matrix indices only represent some internal space to the topological sector, as we mentioned in Sec. III A. When the irrep is one dimensional, such as when the group G is Abelian, the braiding relation results in the accumulation of a phase of $R(h)$ because the matrices belonging to a 1D unitary representation are just phases. This braiding relation holds for a specific orientation of the looplike excitation and direction with which the pointlike particle is moved through the loop. If either orientation was reversed, then we would replace h with its inverse in the braiding relation. This braiding relation is natural from the gauge theory perspective and indeed is equivalent to the one that we predicted in Sec. II.

We also find nontrivial loop-loop braiding between the magnetic excitations when the group G is non-Abelian. When one loop, labeled by g , is passed through another, labeled by h , the first loop has its label conjugated, while the other label is left unchanged, as indicated by Eq. (23). This braiding is just as we would expect for the braiding of the magnetic excitations in ordinary lattice gauge theory, as explained in Sec. II:

Flux-flux braiding relation (\triangleright trivial)

$$C^g(m_1)C^h(m_2)|GS\rangle = C^h(m_2)C^{h^{-1}gh}(m_1)|GS\rangle. \quad (23)$$

Next we consider the other set of excitations, those labeled by objects corresponding to the group E . In this set the only nontrivial braiding, whether point-loop or loop-loop, is the braiding between the pointlike blob excitations and the E -valued loop excitations. When a blob excitation labeled by an element $e \in E$ passes through a loop excitation labeled by the 1D irrep α of E , a phase of $\alpha(e)$ (or the inverse phase, depending on the orientation of the loop and direction of braiding) is accumulated. When \triangleright is trivial, the group E is Abelian, so all of the irreps are 1D and so the braiding transformation is only a phase, as shown in Eq. (24):

Loop-blob braiding relation

$$\begin{aligned}
 B^e(t) \sum_{f \in E} \alpha(f) \delta(f, \hat{e}(m)) |GS\rangle \\
 = \alpha(e) \sum_{f \in E} \alpha(f) \delta(f, \hat{e}(m)) B^e(t) |GS\rangle. \quad (24)
 \end{aligned}$$

The braiding relations are a little different when \triangleright is non-trivial. If we restrict to fake-flat configurations (case 3 from Table I), we have to throw out the magnetic excitations (and the blob excitations labeled by elements outside the kernel of ∂) and the only nontrivial braiding is between the blob excitations and the E -valued loops. When we pass a blob excitation labeled by e through a loop labeled by the irrep μ and indices a and b , the loop transforms by multiplication by the matrix $D^\mu(e)$ or the inverse. On the other hand, if we take our other special case (case 2 in Table I), where ∂ maps from an Abelian E to the center of G , then the braiding is richer than in the \triangleright trivial case. While braiding not involving the magnetic excitations is the same as in the \triangleright trivial case, the magnetic excitations now braid nontrivially with all of the types of excitation. To obtain the braiding relations we have to combine the magnetic excitation with the E -valued loop, giving us an excitation we call a higher-flux loop excitation. This is because the magnetic excitation now carries a 2-flux, but this flux is not well defined unless we also apply an E -valued membrane operator $\delta(e, \hat{e}(m))$ on the same region of space (which fixes the 2-flux). Then the combined membrane operator produces an excitation labeled by a pair of group elements (g, \tilde{e}) , where g is in G and gives the 1-flux, while \tilde{e} is in E and gives the 2-flux. Here g is simply the label of the magnetic membrane operator, but \tilde{e} is related to the label e of the additional E -valued membrane operator by $\tilde{e} = e[g^{-1} \triangleright e^{-1}]$, as we explain in more detail in Ref. [72]. When a pair of these higher-flux excitations braid, we have the following:

Higher-flux–higher-flux braiding relation

$$((g, \tilde{e}_2), (h, \tilde{e}_1)) \rightarrow ((h, \tilde{e}_1 \tilde{e}_2 [h \triangleright \tilde{e}_2^{-1}]), (hgh^{-1}, h \triangleright \tilde{e}_2)). \quad (25)$$

This braiding then matches the prediction of Ref. [68], where the result is argued from geometric grounds. Note that the conjugation of the 1-flux label is slightly different from Eq. (23) due to a different convention we use for the orientation of the higher-flux excitations.

Finally, when a blob excitation labeled by e is passed through a higher-flux excitation labeled by (h, \tilde{f}) , the braiding relation is given by the following:

Higher-flux–blob braiding relation

$$B^e(t) C_T^{h, \tilde{f}}(m) |GS\rangle = C_T^{h, \tilde{f}e}(m) B^e(t'_1) B^{h^{-1} \triangleright e}(t'_2) |GS\rangle, \quad (26)$$

where t'_1 and t'_2 are the parts of the ribbon before and after the intersection with the membrane. Note that the labels of the excitations change in two ways. First, the label f of the pinned E -valued loop is changed by multiplication by e or by $h \triangleright e^{-1}$ (depending on orientation), which induces a change in the 2-flux $\tilde{f} = f[h^{-1} \triangleright e^{-1}]$ of the excitation by multiplication by $e[h^{-1} \triangleright e^{-1}]$ or $e[h \triangleright e^{-1}]$. Second, the blob ribbon operator labeled by e is acted on by the magnetic operator, so that $e \rightarrow h^{-1} \triangleright e$ or $h \triangleright e$ after the intersection (again depending on orientation). From this we can see that the product $\tilde{f}e$ of the 2-fluxes of the two excitations is preserved by the braiding.

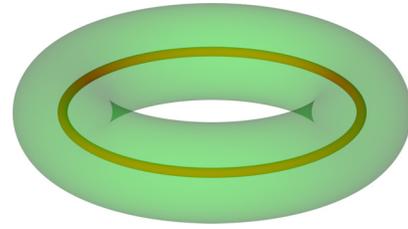


FIG. 44. Given a loop excitation (thin red torus), we can measure its topological charge with a toroidal surface (larger green torus) enclosing it.

VI. TOPOLOGICAL CHARGE

In this section we will explain in more detail what we mean by 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 carries the trivial charge, also called the 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 flux of the electric field through 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 Fig. 44. 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.

The choice of measurement 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 surface. A spherical measurement surface measures a different set of topological charges from a toroidal surface, for example, because a spherical measurement surface cannot distinguish a looplike excitation from a pointlike one. For this reason we say that a spherical surface measures the pointlike charge of an excitation (or set of excitations). In order to determine the looplike charge of an excitation, we must use a toroidal surface, such as the one shown in Fig. 44 (or a surface of higher genus, although we will not consider these in this work).

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 and so

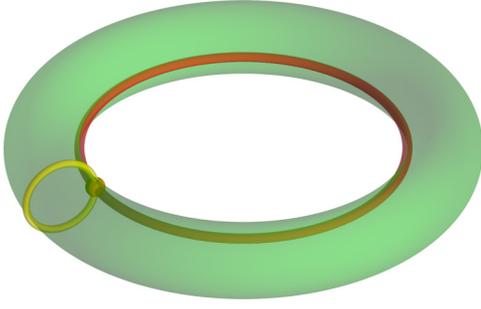


FIG. 45. We apply closed ribbon operators on the two cycles of the torus (the thin red and yellow loops) and a membrane operator over the surface itself (the thicker green torus).

must commute with the Hamiltonian. In addition, 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 [46], we construct such measurement operators using closed ribbon and closed membrane operators [the latter because we consider the (3+1)D case, whereas Ref. [46] considered a (2+1)D model]. This may seem 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.

As an example, consider the case of a torus, shown in Fig. 45. We apply ribbon operators around the two noncontractible cycles of the torus, and membrane operators on the torus itself. Any other closed ribbon operators that we could apply would not be independent (i.e., could be reduced to operators of the types already considered by deformation or other means), or would leave excitations on the surface. This means that the only operators we can apply are an electric and a blob ribbon operator running around each independent noncontractible cycle and a magnetic and E -valued membrane operator over the torus itself. This gives us six labels, one for each operator. We then construct a linear combination over all possible labels, with coefficients chosen so that the sum commutes with the Hamiltonian:

$$\sum_{e_{c_1}, e_{c_2}, e_m \in E} \sum_{g_{c_1}, g_{c_2}, h \in G} \alpha_{(g_{c_1}, g_{c_2}, h, e_{c_1}, e_{c_2}, e_m)} B^{e_{c_1}}(c_1) \times B^{e_{c_2}}(c_2) C_T^h(m) \delta(\hat{e}(m), e_m) \delta(\hat{g}(c_1), g_{c_1}) \delta(\hat{g}(c_2), g_{c_2}), \quad (27)$$

where c_1 and c_2 are the two cycles of the torus, while m is its surface. Unfortunately, not being able to construct the magnetic excitations for a general crossed module (case 3 in Table I) prevents us from constructing all of our charge measurement operators in every case. However, when E is Abelian and ∂ maps from an Abelian group to the center of G (case 2 in Table I, of which case 1 is a subset), we are able to explicitly construct the measurement operators. As explained

in Ref. [72], we find that the following restrictions are necessary for the operator to commute with the Hamiltonian:

Conditions for the measurement operators

First, the coefficients $\alpha_{(g_{c_1}, g_{c_2}, h, e_{c_1}, e_{c_2}, e_m)}$ are only nonzero when

$$\partial(e_m) = [g_{c_2}, g_{c_1}], \quad (28)$$

$$\partial(e_{c_2}) = [g_{c_1}, h], \quad (29)$$

$$\partial(e_{c_1}) = [h, g_{c_2}], \quad (30)$$

$$1_E = [h \triangleright e_m^{-1}] e_m e_{c_1}^{-1} [g_{c_1}^{-1} \triangleright e_{c_1}] e_{c_2}^{-1} [g_{c_2}^{-1} \triangleright e_{c_2}]. \quad (31)$$

Second, the coefficients must satisfy

$$\alpha_{(g_{c_1}, g_{c_2}, h, e_{c_1}, e_{c_2}, e_m)} = \alpha_{(gg_{c_1}g^{-1}, gg_{c_2}g^{-1}, ghg^{-1}, g \triangleright e_{c_1}, g \triangleright e_{c_2}, g \triangleright e_m)} \forall g \in G, \quad (32)$$

$$\alpha_{(g_{c_1}, g_{c_2}, h, e_{c_1}, e_{c_2}, e_m)} = \alpha_{(\partial(e)^{-1} g_{c_1}, g_{c_2}, h, e_{c_1}, e_{c_2}, h \triangleright e)} e^{-1}, e_m e^{-1} [g_{c_2}^{-1} \triangleright e]} \forall e \in E, \quad (33)$$

$$\alpha_{(g_{c_1}, g_{c_2}, h, e_{c_1}, e_{c_2}, e_m)} = \alpha_{(g_{c_1}, \partial(r)g_{c_2}, h, e_{c_1}, [h \triangleright r] r^{-1}, e_{c_2}, e_m r^{-1} [g_{c_1}^{-1} \triangleright r])} \forall r \in E, \quad (34)$$

$$\alpha_{(g_{c_1}, g_{c_2}, h, e_{c_1}, e_{c_2}, e_m)} = \alpha_{(g_{c_1}, g_{c_2}, \partial(e)h, e_{c_1}, [g_{c_2}^{-1} \triangleright e] e^{-1}, e_{c_2}, [g_{c_1}^{-1} \triangleright e^{-1}] e, e_m)} \forall e \in E. \quad (35)$$

The number of linearly independent operators satisfying these conditions is then the number of topological charges measured by the torus. In Ref. [72] we explicitly construct a basis for this space, consisting of operators that project to definite charge. In this case we find that the number of charges matches the ground-state degeneracy of the 3-torus. 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) [94]. 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 [56,60] 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 the 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. [58], using a different method (employing tube algebra). Indeed, the projection operators we construct in Ref. [72] are labeled by the same objects labeling the simple modules of the tube algebra found in Ref [58].

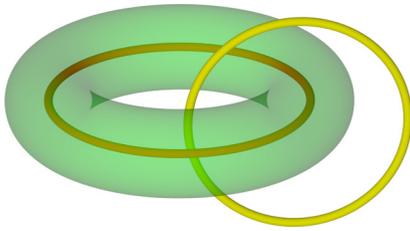


FIG. 46. The torus can measure loops that are held within it (e.g., the red loop shown) or that link with it but are outside (e.g., the yellow loop shown here). More generally, the surface can measure two linked loops, or collections of loops that link together in this way, together with point particles within the torus.

One important thing to note is that the torus surface can measure loops that link with either of the cycles of the torus, as shown in Fig. 46. This means that the general object measured by the surface is a pair of linked loops (or a set of objects that can be fused into such a pair). In particular, we note that the number of charges measured by the torus (and so the ground-state degeneracy) is not equal to the number of distinct looplike excitations, but instead the number of linklike excitations (counting those obtained by fusion with pointlike excitations as well). This suggests that it is important to consider the linking of looplike excitations when studying (3+1)D topological phases. This reinforces work by other authors [28,62] that shows that so-called three-loop (or necklace) braiding, which is braiding of two loops while both are linked to a third loop, is important for characterizing (3+1)D topological phases.

VII. CONCLUSION

In this paper 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 (the 3D version of) the quantum double model 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 the quantum double model, 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 nontrivial braiding only within the sectors (i.e., between the G -valued excitations and between the E -valued excitations). However, unlike in the quantum double model, some of the pointlike excitations are confined, costing additional energy to separate pairs of excitations after producing them, while some of the looplike 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 nontrivial. In particular, we looked at the case where the group E labeling the 2-flux is Abelian and where the map ∂ maps to the center 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 nontrivial 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 nontrivial 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 looplike excitations described in Ref. [68].

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+1)D there are both pointlike and looplike excitations, with associated pointlike and looplike charge. While looplike excitations may possess a pointlike charge, which may be measured by enclosing the excitation in a sphere, they also possess looplike charge, which can only be measured by surfaces with noncontractible handles such as the torus. This approach of explicitly constructing the charge measurement operators (which we present more fully in Ref. [72]) could allow us to measure the charge of collections of excitations and so provide a tool to consider fusion of topological charge.

There are many interesting avenues of research that would build on the results we have obtained in this work, either directly or indirectly. First, 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 I) and while there are some conceptual issues associated with taking the most general case (as described in Sec. IF), 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+1)D topological phases, perhaps in a more general and conceptual 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, either in this model or more generally. 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 nontrivial three-loop braiding statistics [28,48]. 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 pointlike particles) [95]. Because these models have a similar structure to the model considered in this paper, 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+1)D 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+1)D is perhaps incomplete, particularly when it comes to looplike excitations condensing, and the examples considered in this work may be useful in studying this process.

In compliance with EPSRC policy framework on research data, this publication is theoretical work that does not require supporting research data.

ACKNOWLEDGMENTS

The authors thank J. Faria Martins and A. Bullivant for helpful discussions about the higher-lattice gauge theory model. We also thank P. Fendley for advice on preparing our paper. We acknowledge support from EPSRC Grant No. EP/S020527/1.

APPENDIX: CONSISTENCY OF THE HIGHER-LATTICE GAUGE THEORY MODEL UNDER CHANGES TO THE BRANCHING STRUCTURE

As we discussed in Sec. IE, when we define a lattice for the higher-lattice gauge theory model, we must give the lattice a *branching structure*. This determines the orientations of the edges of the lattice as well as the orientation and base point of each plaquette, which we call the “*dressing*” of the lattice. In order for the model to be topological, it must somehow be

resilient to changes to these details. In this Appendix we will show that there is a certain sense in which the energy terms are invariant under changes to the branching structure (to be more precise, a branching structure forbids local oriented loops, and has an ordering of vertices, but we will not make these restrictions for the altered dressing). We note that this issue of consistency under changes to the branching structure is also addressed by Ref. [96], but we will need the more explicit treatment used here for future results. To understand what we mean by consistency under changes to the branching structure, let us first consider an example of how we can change the dressing of the lattice. In Sec. IC, we described the notion of parallel transport across an edge of the lattice. The label of the edge describes the result of parallel transport and the orientation of the edge describes the direction of the parallel transport that gives this result. If we were to perform parallel transport in the opposite direction, we would expect the inverse transformation. This suggests that there is a natural operation in which we can take an edge with a given orientation and a label g and reverse the orientation of the edge, while simultaneously changing the label to g^{-1} , without changing the physical meaning of the label. In a similar way, in Sec. IE we described how reversing the orientation of a surface should change its label from e to e^{-1} . We can also move the base point of a plaquette from v_1 to v_2 , which should change its label from e to $g(v_1 - v_2)^{-1} \triangleright e$, where $g(v_1 - v_2)$ is the group element assigned to the path along which we move the base point. We can treat these transformations as maps between two copies of our Hilbert space, where each Hilbert space corresponds to different dressings of the lattice. We extend the action of these transformations to states where the edges or plaquettes are labeled by linear combinations of the group elements (and to states where the different degrees of freedom are entangled) in the sensible way, to make these transformations linear. Then the transformations, which involve changes to the dressing of a lattice (orientations and base points) as well as to the labels attached to the different cells of the lattice, are maps that connect states from different Hilbert spaces that should have the same physical content, but have different descriptions. Each of the Hilbert spaces is equipped with a Hamiltonian of the form given by Eq. (12) from Sec. IE. If the states related by the transformations do indeed contain the same physical content, then the energy of a state should be preserved under these transformations. That is, if we first apply the Hamiltonian on a state in one of the Hilbert spaces and then apply a transformation to change the dressing of the lattice, this should give us the same result as applying the transformation then the Hamiltonian on the new Hilbert space. Equivalently, given a transformation \hat{T} that changes the dressing, then the Hamiltonian H_1 in the original space should be related to the Hamiltonian H_2 in the second space by $\hat{T}^{-1}H_2\hat{T} = H_1$. If this is true, then we have a way to move between the different dressed lattices, while preserving the structure endowed on the space by the Hamiltonian. This demonstrates the desired resilience of the higher-lattice gauge theory model to changes to the dressing of the lattice. In this section we will show that the Hamiltonian does indeed have this property, subject to certain caveats related to fake flatness and the map \triangleright .

1. Reversing the orientation of an edge

Let us start by considering the procedure where we reverse the orientation of an edge and simultaneously invert the label of that edge. We denote this edge-flipping transformation on an edge i by \hat{P}_i . We wish to show that this map preserves each energy term individually (which is a slightly stronger condition than just preserving the Hamiltonian). Suppose that we start in one copy of our Hilbert space \mathcal{H}_1 , and P_i maps to another copy \mathcal{H}_2 of the Hilbert space. Then for any vertex, edge, plaquette, or blob in our lattice we have an energy term \hat{O}_1 acting on the first Hilbert space, and a corresponding energy term \hat{O}_2 acting on the second Hilbert space, where these energy terms are defined in Sec. IE. We wish to show that the energy is preserved under the transformation \hat{P}_i , by which we mean that $\hat{P}_i \hat{O}_1 = \hat{O}_2 \hat{P}_i$. Thinking in terms of the eigenstates, if this is true then for a state $|\psi\rangle$ which is an eigenstate of the energy term \hat{O}_1 , with eigenvalue λ , the state $\hat{P}_i|\psi\rangle$ reached by flipping the edge i is an eigenstate of the equivalent energy term \hat{O}_2 with the same eigenvalue λ . Because the map \hat{P}_i is invertible (we can undo P_i by flipping the edge back and inverting the edge element again, so P_i is in fact its own inverse), we can also write the condition as $\hat{O}_1 = \hat{P}_i^{-1} \hat{O}_2 \hat{P}_i$. Therefore, when we say that the energy is preserved by the

transformation, we mean that changing the orientation of the edge, applying an energy term and then changing the orientation back has the same net effect as simply applying the corresponding energy term without flipping the edge.

We first wish to show that the edge-flipping procedure is consistent with the vertex terms A_v . Each vertex term is an equal sum of the vertex transforms: $A_v = \frac{1}{|G|} \sum_{g \in G} A_v^g$. Therefore, if we can show that the vertex transforms A_v^g are invariant under P_i , this will also be true for the vertex terms A_v . The vertex transform on a vertex v only has support on the edges that are adjacent to that vertex (and neighboring plaquettes, although the action on plaquettes is independent of edge orientation). Therefore, the vertex transform is only sensitive to the edge-flipping procedure on adjacent edges. Recall that the action of the vertex transform on an adjacent edge i , labeled by g_i , is

$$A_v^x : g_i = \begin{cases} xg_i & i \text{ points away from } v, \\ g_i x^{-1} & i \text{ points towards } v. \end{cases}$$

The action of the vertex transform on edge i is only sensitive to the orientation of edge i itself and not to the orientation of any other edges. Therefore, we only need to consider the effect of flipping edge i itself. We have

$$\begin{aligned} \hat{P}_i^{-1} A_v^x \hat{P}_i : g_i &= \hat{P}_i^{-1} A_v^x : g_i^{-1} \\ &= \hat{P}_i^{-1} : \begin{cases} g_i^{-1} x^{-1} & i \text{ originally pointed away from (now towards) } v, \\ x g_i^{-1} & i \text{ originally pointed towards } v \end{cases} \\ &= \begin{cases} x g_i & i \text{ points away from } v, \\ g_i x^{-1} & i \text{ points towards } v, \end{cases} \end{aligned}$$

from which we see that the action of the vertex transform is preserved under the edge-flipping procedure P_i .

Next, we consider the plaquette and blob energy terms. These terms only involve the edges through path elements (the path around the boundary of the plaquette for the plaquette term, and paths between the base points of surfaces for the blob term). When we construct a path element from the edge elements, each edge contributes the group element it would have if it pointed along the path. That is, if the edge is aligned with the path the edge contributes its group element, but if the edge is antialigned with the path, then it instead contributes the inverse of its group element. This means that path elements are invariant under the flipping procedure. If an edge with label g_i already points along the path, then it contributes g_i to the path. If we flip the edge, then we change its label to g_i^{-1} . However, the flipped edge then points against the path, so it contributes $(g_i^{-1})^{-1} = g_i$ to the path due to the way that we calculate path elements. Similarly, if the edge originally points against the path, then it contributes g_i^{-1} . If we flip it, then we change its label to g_i^{-1} . However, the edge then points along the path, so it still contributes g_i^{-1} . The path elements being invariant under the edge-flipping procedure in this way means that both the plaquette and blob energy terms are similarly invariant under this procedure.

This leaves us to consider the edge energy terms. The energy term for an edge i is \mathcal{A}_i , which is made of a sum of edge transforms: $\mathcal{A}_i = \frac{1}{|E|} \sum_{e \in E} \mathcal{A}_i^e$. Let us consider how an edge transform is affected by the edge-flipping procedures. Recall that the edge transform \mathcal{A}_i^e acts on the edge i and on neighboring plaquettes. The support of \mathcal{A}_i^e also includes some of the other edges along the boundaries of the adjacent plaquettes because the action of \mathcal{A}_i^e on an adjacent plaquette p is

$$\mathcal{A}_i^e : e_p = \begin{cases} e_p [g(v_0(p) - v_i) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(v_0(p) - v_{i+1}) \triangleright e] e_p & i \text{ points against the boundary,} \end{cases} \tag{A1}$$

which depends on the path element $g[v_0(p) - v_i]$ or $g[v_0(p) - v_{i+1}]$. Flipping the orientation of the edges along these paths will not affect the action on the edge transform because the path elements are invariant under the flipping

operation, as we just saw. However, the action of \mathcal{A}_i^e directly depends on the orientation of i itself. Therefore, we need to consider whether \mathcal{A}_i^e is invariant under the edge-flipping operation on edge i . First consider the action of \mathcal{A}_i^e on the

edge i itself. If the edge initially has label g_i , then the edge transform acts as $\mathcal{A}_i^e : g_i = \partial(e)g_i$. Therefore,

$$\begin{aligned} P_i^{-1} \mathcal{A}_i^e P_i : g_i &= P_i^{-1} \mathcal{A}_i^e : g_i^{-1} \\ &= P_i^{-1} : \partial(e)g_i^{-1} \\ &= g_i \partial(e)^{-1} \\ &= g_i \partial(e)^{-1} g_i^{-1} g_i \\ &= \partial(g_i \triangleright e^{-1}) g_i \\ &= \mathcal{A}_i^{g_i \triangleright e^{-1}} : g_i, \end{aligned}$$

where we used the Peiffer condition (4) to write $\partial(g_i \triangleright e^{-1}) = g_i \partial(e) g_i^{-1}$. We therefore see that, unlike the vertex transforms, the individual edge transforms are not generally invariant under the edge-flipping procedure. However, this does not mean that the edge energy term itself is not invariant. We have, at least for the action on the edge,

$$\begin{aligned} P_i^{-1} \mathcal{A}_i P_i : g_i &= \frac{1}{|E|} \sum_{e \in E} P_i^{-1} \mathcal{A}_i^e P_i : g_i \\ &= \frac{1}{|E|} \sum_{e \in E} \mathcal{A}_i^{g_i \triangleright e^{-1}} : g_i \\ &= \frac{1}{|E|} \sum_{e' = g_i \triangleright e^{-1}} \mathcal{A}_i^{e'} : g_i = \mathcal{A}_i : g_i, \end{aligned}$$

$$\mathcal{A}_i^e : e_p = \begin{cases} e_p [g(v_0(p) - s(i)) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(v_0(p) - s(i)) \triangleright e] e_p & i \text{ points against the boundary.} \end{cases}$$

On the other hand if we act with $P_i^{-1} \mathcal{A}_i^e P_i$, we have

$$P_i^{-1} \mathcal{A}_i^e P_i : e_p = \begin{cases} [g(v_0(p) - t(i)) \triangleright e] e_p & i \text{ originally points along the boundary of } p, \\ e_p [g(v_0(p) - t(i)) \triangleright e^{-1}] & i \text{ originally points against the boundary.} \end{cases}$$

Note that if \triangleright is trivial, we can forget the path elements and so we get $P_i^{-1} \mathcal{A}_i^e P_i = \mathcal{A}_i^{e^{-1}}$ and the overall edge term will commute with the flipping procedure. On the other hand, if \triangleright is nontrivial we need to consider the path elements more closely. Looking at Fig. 47, we see that when i points along the boundary of p , we have

$$g(v_0(p) - t(i)) g_i^{-1} g(v_0(p) - s(i))^{-1} = g(\text{boundary}(p))^{-1}.$$

Provided that the plaquette satisfies fake flatness, so that $g[\text{boundary}(p)]^{-1} = \partial(e_p)$, this means that

$$\begin{aligned} g(v_0(p) - t(i)) &= g(\text{boundary}(p))^{-1} g(v_0(p) - s(i)) g_i \\ &= \partial(e_p) g(v_0(p) - s(i)) g_i \end{aligned}$$

and so

$$\begin{aligned} P_i^{-1} \mathcal{A}_i^e P_i : e_p &= [g(v_0(p) - t(i)) \triangleright e] e_p \\ &= [(\partial(e_p) g(v_0(p) - s(i)) g_i) \triangleright e] e_p. \end{aligned}$$

Using the Peiffer condition (5), this becomes

$$\begin{aligned} P_i^{-1} \mathcal{A}_i^e P_i : e_p &= e_p [g(v_0(p) - s(i)) g_i \triangleright e] e_p^{-1} e_p \\ &= e_p [g(v_0(p) - s(i)) \triangleright (g_i \triangleright e^{-1})^{-1}] \\ &= \mathcal{A}_i^{g_i \triangleright e^{-1}} : e_p. \end{aligned}$$

which suggests that the energy term is invariant under the procedure, even if the individual transforms are not. However, we have only shown this for the action on the edges, and it must also hold for the plaquettes in order to obtain the operator relation $P_i^{-1} \mathcal{A}_i^e P_i = \mathcal{A}_i^{g_i \triangleright e^{-1}}$ (i.e., we must also satisfy the relationship $P_i^{-1} \mathcal{A}_i^e P_i : e_p = \mathcal{A}_i^{g_i \triangleright e^{-1}} : e_p$ for all plaquettes p).

In order to determine whether the action on the plaquettes also satisfies this relationship, consider Eq. (A1) which defines this action on the adjacent plaquettes. When we flip the edge, we reverse the relative orientation of i and the plaquette, and so change which path element $g(v_0(p) - v_i)$ or $g(v_0(p) - v_i)$ appears in the expression for the action on the plaquette (i.e., we change whether the path travels along the circulation of the plaquette or against it). Furthermore, note that the vertex v_i or v_{i+1} which is the end-point of the path is always the source of the edge i , the vertex which the edge points away from. When we flip the edge, we exchange the source and target of the edge (where the target is the vertex the edge points towards), which therefore changes the end point of the path that appears in the edge transform. These changes are shown in Fig. 47. Denoting the original source of the edge i by $s(i)$ and the original target of i by $t(i)$, the action of the edge transform on an adjacent plaquette (without flipping the edge) is

In a similar way, if i points against the boundary of the plaquette, we see from Fig. 47 that

$$g(v_0(p) - t(i)) g_i^{-1} g(v_0(p) - s_i)^{-1} = g(\text{boundary}(p)).$$

Again, if the plaquette satisfies fake flatness, so that $g(\text{boundary}(p)) = \partial(e_p)^{-1}$, then (following the same steps as before)

$$\begin{aligned} P_i^{-1} \mathcal{A}_i^e P_i : e_p &= e_p g(v_0(p) - t(i)) \triangleright e^{-1} \\ &= e_p ([\partial(e_p)^{-1} g(v_0(p) - s_i) g_i] \triangleright e^{-1}) \\ &= e_p e_p^{-1} ([g(v_0(p) - s_i) g_i] \triangleright e^{-1}) e_p \\ &= (g(v_0(p) - s_i) \triangleright (g_i \triangleright e)^{-1}) e_p \\ &= \mathcal{A}_i^{g_i \triangleright e^{-1}} : e_p. \end{aligned}$$

We can therefore see that, for the action on every degree of freedom (and so for the operators themselves), $P_i^{-1} \mathcal{A}_i^e P_i = \mathcal{A}_i^{g_i \triangleright e^{-1}}$ and so $P_i^{-1} \mathcal{A}_i P_i = \mathcal{A}_i$. However, note that (unless \triangleright is trivial) this relies on the plaquettes around the edge satisfying fake flatness. If fake flatness is not satisfied near the edge and \triangleright is not trivial, then we cannot say that the edge energy term is invariant under the edge-flipping procedure.

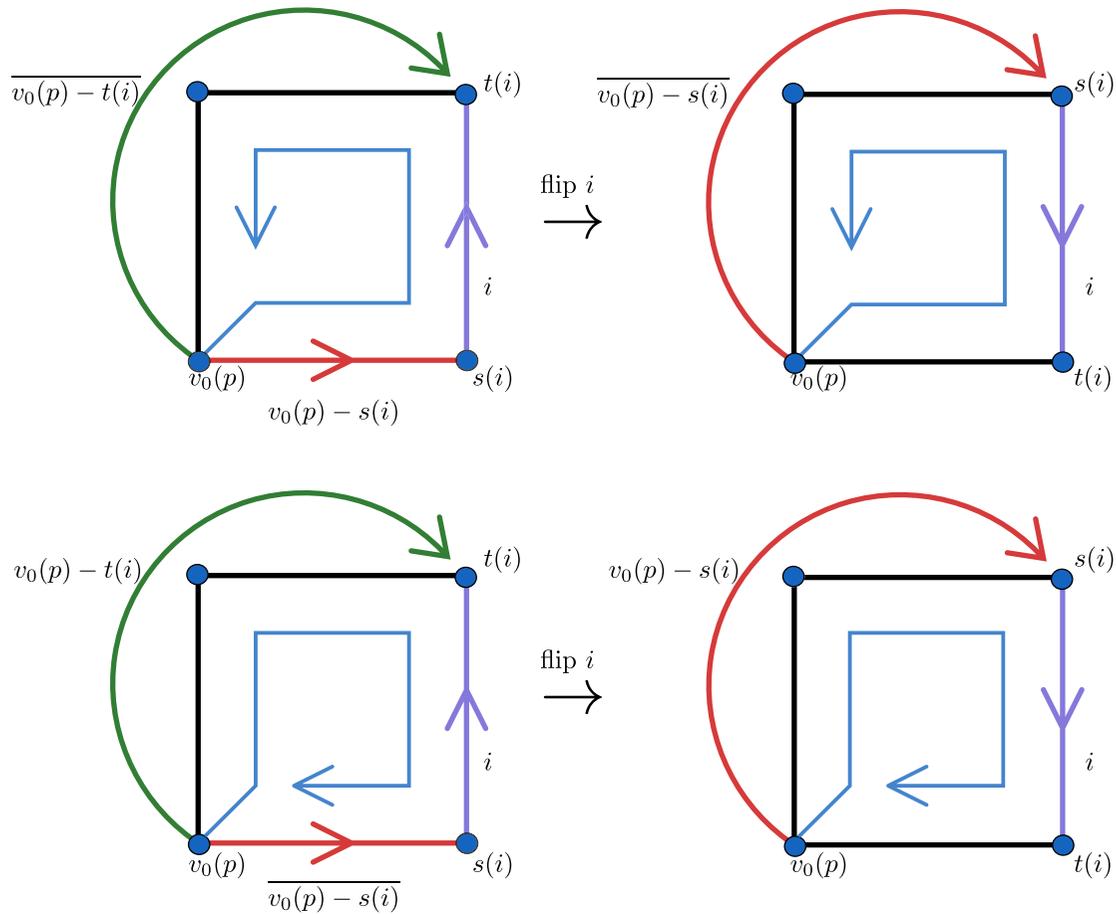


FIG. 47. The action of an edge transform applied on edge i on an adjacent plaquette p depends on the orientation of that edge. If we flip the edge, then the paths which appear in the action of the edge transform change. In the top line, we consider the case where the edge i initially points along the plaquette's orientation. The path which appears in the action of the edge transform is then $v_0(p) - s(i)$. If we flip the edge and then apply the edge transform, then instead the path $\overline{v_0(p) - s(i)}$ from the right side appears in the edge transform. Because we have flipped the edge, the source of the edge after it is flipped is the target of the edge before the flip. Therefore, the path $\overline{v_0(p) - s(i)}$ that appears in the edge transform after the flip is the same path as $v_0(p) - t(i)$ (green path on the left side) before the flip. In the lower line, we show the analogous situation when the edge initially points against the plaquette's orientation.

This indicates that the energy of the edge energy term is not independent of the supposedly arbitrary dressing of the lattice and so we do not consider the energy of the edge to be well defined (although note that if we fix a branching structure and never consider changing it, the energy can still be defined). However, in the ground state, or more generally in regions where fake flatness is satisfied, the energy of the edge is well defined. We are typically interested in the ground state, and states with a few excitations. For such states, the edge energy is only ill defined near these excitations. This does not affect the topological quantities of the theory, such as the braiding relations or topological charges, because these can be measured far away from the affected regions, where the theory is still fully consistent.

2. Reversing the orientation of a plaquette

Having considered the procedure of flipping an edge and inverting the corresponding edge label, and having shown that this is consistent with the various energy terms, we now consider the analogous procedure where we reverse the ori-

entation of a plaquette and invert its label. We denote this operation by Q_p for a plaquette p . First consider the vertex transform A_v^g . Apart from the edges, which are unaffected by the orientation of plaquettes, this transform acts only on plaquettes with base point at v . We have

$$A_v^g : e_p = \begin{cases} g \triangleright e_p & \text{if } v_0(p) = v, \\ e_p & \text{otherwise,} \end{cases}$$

where $v_0(p)$ is the base point of plaquette p . The action of the vertex transform on plaquette p can only be affected by flipping the plaquette p itself (it is not affected by flipping other plaquettes). We have

$$\begin{aligned} Q_p^{-1} A_v^g Q_p &= Q_p^{-1} A_v^g : e_p^{-1} \\ &= Q_p^{-1} : \begin{cases} g \triangleright e_p^{-1} & \text{if } v_0(p) = v, \\ e_p^{-1} & \text{otherwise} \end{cases} \\ &= \begin{cases} g \triangleright e_p & \text{if } v_0(p) = v, \\ e_p & \text{otherwise} \end{cases} \\ &= A_v^g : e_p, \end{aligned}$$

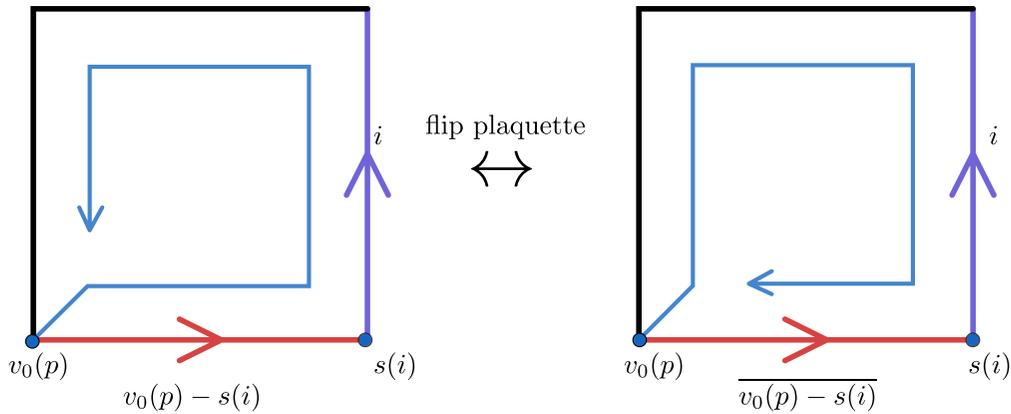


FIG. 48. The paths used for the action of an edge transform on a plaquette depend on the orientation of the plaquette with respect to the edge. We may therefore think that flipping the orientation of the plaquette will change the path used for the action of the edge transform. However, we can see that while we do swap whether the path $v_0(p) - s(i)$ or $\overline{v_0(p) - s(i)}$ is used, we also swap the definition of these paths. If the plaquette is initially aligned with the edge and we swap the orientation of the plaquette, then the path $v_0(p) - s(i)$ before the flip is the same as the path $\overline{v_0(p) - s(i)}$ after the flip (and vice versa).

from which we see that the vertex transforms (and so the vertex energy term) are invariant under the plaquette-flipping procedure.

Next, consider the plaquette term. Recall that the plaquette term for a plaquette p is

$$B_p = \delta(\partial(e_p)g(\text{boundary}(p)), 1_G).$$

If we flip the orientation of the plaquette, then we also reverse the orientation of its boundary, so that $g[\text{boundary}(p)]$ is inverted. Therefore,

$$\begin{aligned} Q_p^{-1} B_p Q_p &= \delta(\partial(e_p^{-1})g(\text{boundary}(p))^{-1}, 1_G) \\ &= \delta(g(\text{boundary}(p))^{-1}, \partial(e_p)) \\ &= \delta(1_G, \partial(e_p)g(\text{boundary}(p))) \\ &= B_p, \end{aligned}$$

so that the plaquette energy term is preserved by the plaquette-flipping procedure.

The blob energy term is similarly preserved by Q_p . Recall that the blob energy term checks that the blob 2-holonomy is

equal to the identity of group E , where the blob 2-holonomy is a product of the plaquette elements around the boundary of the blob. For a blob B , the 2-holonomy is

$$H_2(B) = \prod_{p \in \text{Bd}(B)} g(v_0(B) - v_0(p)) \triangleright e_p^{\sigma_p},$$

where σ_p is 1 or -1 depending on the orientation of the plaquette p , $v_0(p)$ is the base point of plaquette p , and $v_0(B)$ is the base point of the blob B . Flipping the orientation of plaquette p swaps σ_p between the two, which cancels with the inverse from the flipping procedure itself (in a similar way to how the contribution of edge elements to a path was invariant under the edge-flipping procedure), so that the blob 2-holonomy, and thus the blob energy term, is preserved by the plaquette-flipping orientation.

This just leaves the edge energy terms to consider. Recall that the action of the edge transform \mathcal{A}_i^e on an adjacent plaquette p depends on the relative orientation of the edge i and the plaquette p . We have

$$\mathcal{A}_i^e : e_p = \begin{cases} e_p [g(v_0(p) - s(i)) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(\overline{v_0(p) - s(i)}) \triangleright e] e_p & i \text{ points against the boundary.} \end{cases}$$

Swapping the orientation of the plaquette does not change the source of i , unlike the edge-flipping procedure. However, it does change the relative orientation of the plaquette and the edge. The orientation of the plaquette determines which path between the base point of the plaquette and source of the edge we use. When the edge aligns with the orientation of the plaquette we use the path $(v_0(p) - s(i))$ that also aligns with the orientation of the plaquette, and when the edge is antialigned with the plaquette we use the path $(\overline{v_0(p) - s(i)})$ that travels against the orientation of the plaquette. Therefore, flipping the plaquette changes which of these paths we use. However, when we flip the plaquette, the path that previously was aligned with the plaquette is now antialigned with the plaquette. This means that $(v_0(p) - s(i))$ before the flip is equal to $(\overline{v_0(p) - s(i)})$ afterwards (and vice versa), as we show in Fig. 48. These two effects cancel, and therefore flipping the orientation of the plaquette has no net effect on the path that appears in the edge transform (we swap whether we should use the aligned or antialigned path, but also swap which one is which). In addition, reversing the orientation of the plaquette changes whether we use premultiplication (by

$g(\overline{v_0(p) - s(i)}) \triangleright e$) or postmultiplication (by $g(v_0(p) - s(i)) \triangleright e^{-1}$). Using this, we see that

$$\begin{aligned} Q_p^{-1} \mathcal{A}_i^e Q_p : e_p &= Q_p^{-1} \mathcal{A}_i^e : e_p^{-1} \\ &= Q_p^{-1} : \begin{cases} [g(v_0(p) - s(i)) \triangleright e] e_p^{-1} & i \text{ originally points along the boundary of } p, \\ e_p^{-1} [g(\overline{v_0(p) - s(i)}) \triangleright e^{-1}] & i \text{ originally points against the boundary of } p \end{cases} \\ &= \begin{cases} e_p [g(v_0(p) - s(i)) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(\overline{v_0(p) - s(i)}) \triangleright e] e_p & i \text{ points against the boundary} \end{cases} \\ &= \mathcal{A}_i^e : e_p, \end{aligned}$$

so the edge transforms are invariant under the plaquette-flipping procedure. This means that every energy term is invariant under this procedure.

3. Moving the base point of a plaquette

The final procedure to consider is changing the base point of a plaquette. We denote the procedure that moves the base point of plaquette p from a vertex v_1 to a vertex v_2 by $E_p(v_1 \rightarrow v_2)$. In cases where the precise path by which we move the base point (rather than just the end points) is important, we will state that this is the case. Under the operation $E_p(v_1 \rightarrow v_2)$, as well as changing the base point of the plaquette, we must change the plaquette's label from e_p to $g(v_1 - v_2)^{-1} \triangleright e_p$, where $g(v_1 - v_2)$ is the group element assigned to the path along which we move the base point. In the case where \triangleright is trivial, the base point of a plaquette is irrelevant, but in the more general cases the base point of a plaquette affects the action of all of the energy terms. We first consider a vertex transform. The vertex transform at a vertex v affects any plaquette whose base point is at the vertex v . Furthermore, it affects path elements which start or end at the vertex v . This is relevant because the transformation of the plaquette label under $E_p(v_1 \rightarrow v_2)$ depends on the path element $g(v_1 - v_2)$, which is affected by vertex transforms at v_1 and v_2 . This means that the vertex transforms at v_1 and v_2 might not commute with the procedure for moving the base point. First consider the vertex transform $A_{v_1}^x$. We have $A_{v_1}^x : e_p = x \triangleright e_p$. On the other hand, we have

$$E_p^{-1}(v_1 \rightarrow v_2) A_{v_1}^x E_p(v_1 \rightarrow v_2) : e_p = E_p^{-1}(v_1 \rightarrow v_2) A_{v_1}^x : g(v_1 - v_2)^{-1} \triangleright e_p = E_p^{-1}(v_1 \rightarrow v_2) : g(v_1 - v_2)^{-1} \triangleright e_p,$$

where in the last line the vertex transform leaves the plaquette element unchanged because the base point of p is no longer at v_1 . However, the path element for $v_1 - v_2$ has been changed by the action of the vertex transform from its original value of $g(v_1 - v_2)$ to $xg(v_1 - v_2)$. This means that when we move the base point back along the path, we pick up a factor of $xg(v_1 - v_2)$ acting on the plaquette, rather than just a factor of $g(v_1 - v_2)$. This gives us

$$\begin{aligned} E_p^{-1}(v_1 \rightarrow v_2) : g(v_1 - v_2)^{-1} \triangleright e_p &= [xg(v_1 - v_2)] \triangleright [g(v_1 - v_2)^{-1} \triangleright e_p] \\ &= x \triangleright e_p, \end{aligned}$$

so that

$$E_p^{-1}(v_1 \rightarrow v_2) A_{v_1}^x E_p(v_1 \rightarrow v_2) : e_p = A_{v_1}^x : e_p.$$

We therefore see that the vertex transforms (and so the vertex energy term) at v_1 are invariant under changing the base point. Now consider the vertex transforms at v_2 , the position of the base point after we move it. If we do not move the base point, we have $A_{v_2}^x : e_p = e_p$. On the other hand, if we do move the base point we have

$$\begin{aligned} E_p^{-1}(v_1 \rightarrow v_2) A_{v_2}^x E_p(v_1 \rightarrow v_2) : e_p &= E_p^{-1}(v_1 \rightarrow v_2) A_{v_2}^x : g(v_1 \rightarrow v_2)^{-1} \triangleright e_p \\ &= E_p^{-1}(v_1 \rightarrow v_2) : x \triangleright [g(v_1 \rightarrow v_2)^{-1} \triangleright e_p], \end{aligned}$$

where the vertex transform affects the plaquette label because the base point is at v_2 after the action of $E_p(v_1 \rightarrow v_2)$. However, the vertex transform also changes the path element $g(v_1 \rightarrow v_2)$ to $g(v_1 \rightarrow v_2)x^{-1}$. Therefore,

$$E_p^{-1}(v_1 \rightarrow v_2) : x \triangleright [g(v_1 \rightarrow v_2)^{-1} \triangleright e_p] = [g(v_1 \rightarrow v_2)x^{-1}] \triangleright (x \triangleright [g(v_1 \rightarrow v_2)^{-1} \triangleright e_p]) = e_p,$$

from which we see that

$$E_p^{-1}(v_1 \rightarrow v_2) A_{v_2}^x E_p(v_1 \rightarrow v_2) : e_p = A_{v_2}^x : e_p.$$

This means that the vertex energy term at v_2 is also unaffected by our procedure for changing the base points of plaquettes, and so all of the vertex transforms are preserved by this procedure.

Next, consider the plaquette energy terms. Moving the base point of a plaquette also affects the boundary of that plaquette, as shown in Fig. 49. We see that under $E_p(v_1 \rightarrow v_2)$, the boundary of plaquette p transforms as

$$g(\text{boundary}(p)) \rightarrow g(v_1 \rightarrow v_2)^{-1} g(\text{boundary}(p)) g(v_1 \rightarrow v_2).$$

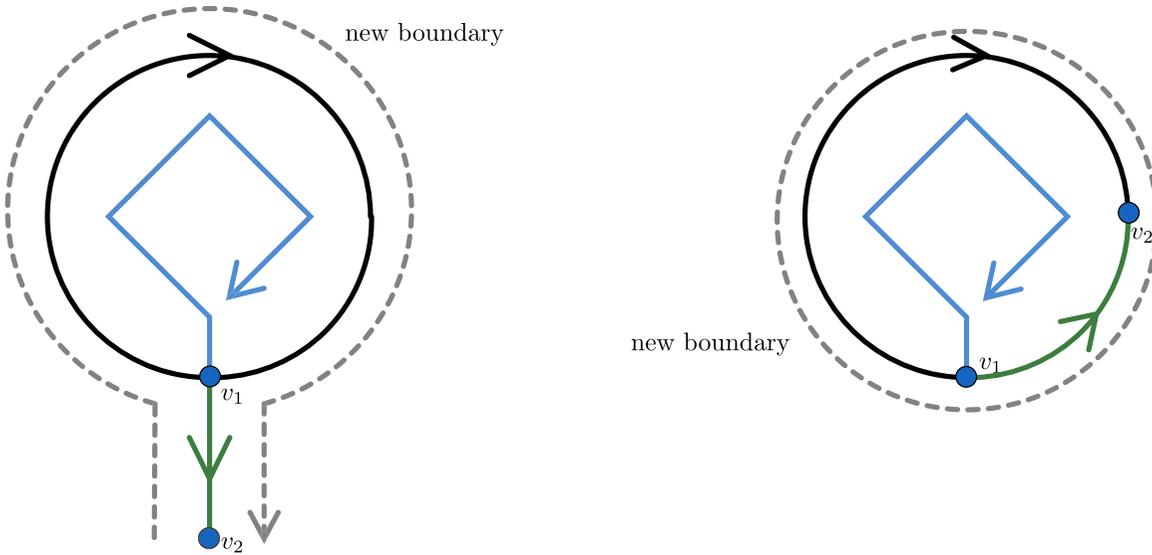


FIG. 49. When we change the base point of a plaquette, we also change the boundary of that plaquette. In the left image we show a case where the base point is moved from a vertex v_1 to v_2 , along a path $v_1 - v_2$ (green) that leaves the plaquette. That is, we whisker the plaquette. We can see that the new boundary (gray dotted line) is $(v_1 - v_2)^{-1}(\text{boundary})(v_1 - v_2)$, where (boundary) is the original boundary of the plaquette (black). In the right image, we instead show a case where the base point of the plaquette is moved along the boundary of the plaquette. In this case, it is still true that the boundary after moving the base point is $(v_1 - v_2)^{-1}(\text{boundary})(v_1 - v_2)$, although this time there is some cancellation between sections of the original boundary and $(v_1 - v_2)$.

Therefore, the plaquette holonomy $H_1(p) = \partial(e_p)g(\text{boundary}(p))$ becomes

$$\begin{aligned} & \partial[g(v_1 \rightarrow v_2)^{-1} \triangleright e_p]g(v_1 \rightarrow v_2)^{-1}g(\text{boundary}(p))g(v_1 \rightarrow v_2) \\ &= g(v_1 \rightarrow v_2)^{-1}\partial(e_p)g(v_1 \rightarrow v_2)g(v_1 \rightarrow v_2)^{-1}g(\text{boundary}(p))g(v_1 \rightarrow v_2) \\ &= g(v_1 \rightarrow v_2)^{-1}\partial(e_p)g(\text{boundary}(p))g(v_1 \rightarrow v_2), \end{aligned}$$

where we used the Peiffer condition (4) to write $\partial[g(v_1 \rightarrow v_2)^{-1} \triangleright e_p] = g(v_1 \rightarrow v_2)^{-1}\partial(e_p)g(v_1 \rightarrow v_2)$. We see that the plaquette holonomy is merely conjugated by a path element, which preserves the identity element. Therefore, the energy term (which checks if the plaquette holonomy is equal to the identity) is unaffected by the base-point changing procedure.

The next energy term to consider is the blob term. This checks that the blob 2-holonomy is the identity, where the blob 2-holonomy is

$$H_2(B) = \prod_{p \in \text{Bd}(B)} g(v_0(B) - v_0(p)) \triangleright e_p^{\sigma_p},$$

where σ_p depends on the orientation of the plaquette. The contribution of a particular plaquette p to the blob 2-holonomy is $g(v_0(B) - v_0(p)) \triangleright e_p^{\sigma_p}$. If we change the base point of this plaquette from v_1 to v_2 , then we must change the label e_p to $g(v_1 \rightarrow v_2)^{-1} \triangleright e_p$ due to the effect of E_p , but we must also change the base point $v_0(p)$ that appears in the expression $g(v_0(B) - v_0(p)) \triangleright e_p^{\sigma_p}$ along the same path (the same path so that the resulting surface is the same as the original). This means that when we change the base point of plaquette p , its contribution to the blob 2-holonomy transforms as

$$\begin{aligned} g(v_0(B) - v_0(p)) \triangleright e_p^{\sigma_p} &\rightarrow g(v_0(B) - v_2) \triangleright [g(v_1 \rightarrow v_2)^{-1} \triangleright e_p^{\sigma_p}] \\ &= (g(v_0(B) - v_1)g(v_1 \rightarrow v_2)) \triangleright [g(v_1 \rightarrow v_2)^{-1} \triangleright e_p^{\sigma_p}] \\ &= g(v_0(B) - v_1) \triangleright e_p^{\sigma_p}, \end{aligned}$$

so that the contribution of the plaquette to the blob 2-holonomy is unchanged by moving the base point of the plaquette.

The final energy terms to consider are the edge energy terms. There are several cases to consider. We need to consider edge terms that directly affect the plaquette whose base point is being moved (i.e., the edges on the boundary of that plaquette), but also any edge transforms along the path on which we move the plaquette's base point. First consider the latter kind of edge, which is not on the original boundary of the plaquette, but is on the path along which we move the base point. If we move the base point of the plaquette before acting with the edge transform, then the path becomes part of the boundary of the plaquette. That is, the plaquette is whiskered along the path through which we move the base point. The action of the edge transform, as defined in Eq. (7), does not account for the possibility of a whiskered plaquette. We give a more general definition of the edge

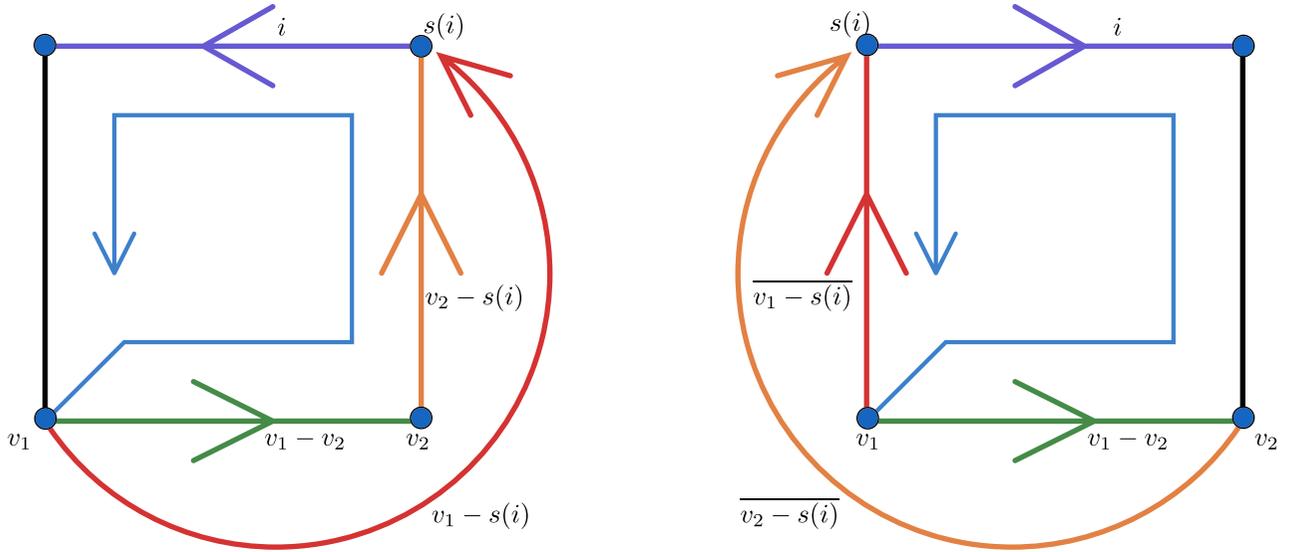


FIG. 50. The action of the edge transform on a plaquette p depends on various paths, which change when we change the base point of the plaquette. In this figure, we show two examples. In the left case, the orientation of the edge on which we apply the transform (i) matches the orientation of the plaquette (shown in blue). Because these orientations match, the path which appears in the edge transform is $v_0(p) - s(i)$, which has the same orientation as the plaquette. We see that if the original base point is v_1 and we move it to v_2 along the path $v_1 - v_2$, where the edge i is not on this path, then $v_1 - s(i) = (v_1 - v_2)[v_2 - s(i)]$. Similarly, in the right of the figure we show a case where the edge points against the orientation of the plaquette, so that the path $v_0(p) - s(i)$, which has the opposite orientation to the plaquette, appears in the edge transform. In this case $v_1 - s(i) = (v_1 - v_2)[v_0(p) - s(i)]$. Note that there are two other cases, where we flip the orientation of the plaquette in each of these images, where the paths obey similar relations.

transform in Ref. [71] (in Sec. S-I C), which describes the action on whiskered plaquettes and other special cases. We define this generalized edge transform in a way that is consistent with the procedure for changing the base point by whiskering (at least if fake flatness is satisfied), as we show in Ref. [71] when we define the transform. We will therefore not prove this consistency here.

Next, we consider the case where we move the base point along the boundary of the plaquette and i is one of the edges on the plaquette. Recall that the action of the edge transform \mathcal{A}_i^e on an adjacent plaquette is

$$\mathcal{A}_i^e : e_p = \begin{cases} e_p [g(v_0(p) - s(i)) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(v_0(p) - s(i)) \triangleright e] e_p & i \text{ points against the boundary.} \end{cases}$$

Now consider $E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2)$, where v_1 is the initial base point of the plaquette p , $v_0(p)$. Then

$$\begin{aligned} & E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2) : e_p \\ &= E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e : g(v_1 - v_2)^{-1} \triangleright e_p \\ &= E_p(v_1 \rightarrow v_2)^{-1} : \begin{cases} [g(v_1 - v_2)^{-1} \triangleright e_p] \{g[v_2 - s(i)] \triangleright e^{-1}\} & i \text{ points along the boundary of } p, \\ [g(\overline{v_2 - s(i)}) \triangleright e] [g(v_1 - v_2)^{-1} \triangleright e_p] & i \text{ points against the boundary.} \end{cases} \end{aligned}$$

Then, if the edge i is not on the path $v_1 - v_2$, so that the label of the path is unaffected by the edge transform, moving the base point of the plaquette back gives

$$\begin{aligned} & E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2) : e_p \\ &= \begin{cases} g(v_1 - v_2) \triangleright ([g(v_1 - v_2)^{-1} \triangleright e_p] \{g[v_2 - s(i)] \triangleright e^{-1}\}) & i \text{ points along the boundary of } p, \\ g(v_1 - v_2) \triangleright ([g(\overline{v_2 - s(i)}) \triangleright e] [g(v_1 - v_2)^{-1} \triangleright e_p]) & i \text{ points against the boundary} \end{cases} \\ &= \begin{cases} e_p [(g(v_1 - v_2)g(v_2 - s(i))) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [(g(v_1 - v_2)g(\overline{v_2 - s(i)})) \triangleright e] e_p & i \text{ points against the boundary.} \end{cases} \end{aligned}$$

We then need to consider the path elements involved in the expression above. There are various cases, as shown in Fig. 50. We see that $g(v_1 - v_2)g(v_2 - s(i)) = g(v_1 - s(i))$ if the edge is aligned with the boundary of p and $g(v_1 - v_2)g(\overline{v_2 - s(i)}) =$

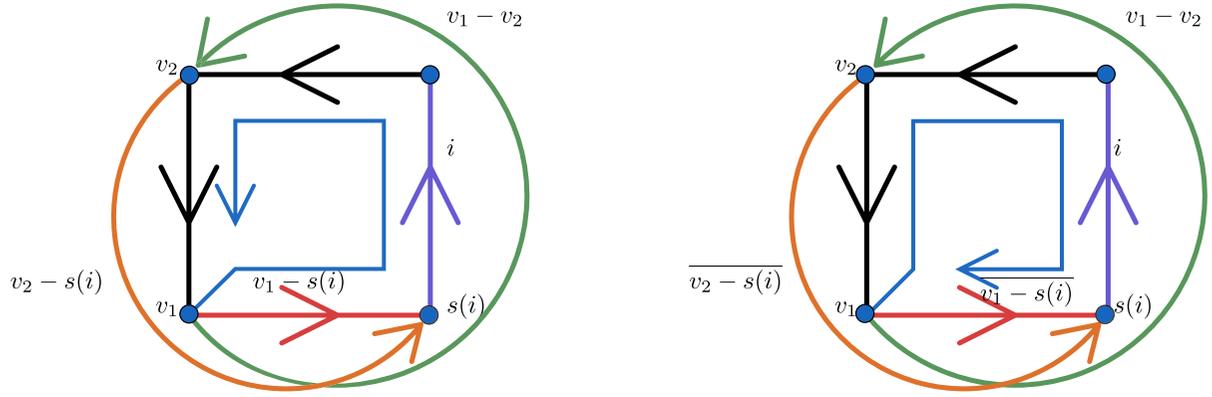


FIG. 51. In this figure, we consider the case where the path on which we move the base point includes the edge i on which we apply the edge transform. In the left case, the edge i has the same orientation as the plaquette. In this case, the path $(v_1 - v_2)[v_2 - s(i)]$ includes the entire boundary of the plaquette, so $(v_1 - v_2)[v_2 - s(i)] = (\text{boundary})[v_1 - s(i)]$. In the right image, we consider the case where the edge and plaquette have opposite orientations. In this case $(v_1 - v_2)[v_2 - s(i)] = (\text{boundary})^{-1}[v_1 - s(i)]$.

$\overline{g(v_1 - s(i))}$ if it is antialigned. Therefore,

$$\begin{aligned} E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2) : e_p &= \begin{cases} e_p ([g(v_1 - v_2)g(v_2 - s(i))] \triangleright e^{-1}) & i \text{ points along the boundary of } p, \\ ([g(v_1 - v_2)\overline{g(v_2 - s(i))}] \triangleright e) e_p & i \text{ points against the boundary} \end{cases} \\ &= \begin{cases} e_p [g(v_0(p) - s(i)) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(v_0(p) - s(i)) \triangleright e] e_p & i \text{ points against the boundary} \end{cases} \\ &= \mathcal{A}_i^e : e_p, \end{aligned}$$

and so the action of the edge transform is preserved.

Next we consider the case where i is not only on the plaquette, but is also one of the edges on the path along which we move the base point, as shown in Fig. 51. In this case, the edge transform affects the path element before we return it to its original position. If the edge i points along $v_1 - v_2$, then the edge transform acts on the path element $g(v_1 - v_2) = g(v_1 - s(i))g_i g(t(i) - v_2)$ as

$$\begin{aligned} \mathcal{A}_i^e : g(v_1 - v_2) &= \mathcal{A}_i^e : \begin{cases} g(v_1 - s(i))g_i g(t(i) - v_2) & i \text{ points along the boundary of } p, \\ \overline{g(v_1 - s(i))}g_i g(t(i) - v_2) & i \text{ points against the boundary of } p \end{cases} \\ &= \begin{cases} g(v_1 - s(i))\partial(e)g_i g(t(i) - v_2) & i \text{ points along the boundary of } p, \\ \overline{g(v_1 - s(i))}\partial(e)g_i g(t(i) - v_2) & i \text{ points against the boundary of } p, \end{cases} \end{aligned}$$

where we split the path $g(v_1 - v_2)$ into parts and then used the action of the edge transform on the edge i . Then we want to write this again in terms of $g(v_1 - v_2)$, so we introduce a factor of the identity in the form of $g(v_1 - s(i))^{-1}g(v_1 - s(i))$ (or the equivalent with $\overline{g(v_1 - s(i))}$) to obtain

$$\begin{aligned} \mathcal{A}_i^e : g(v_1 - v_2) &= \begin{cases} g(v_1 - s(i))\partial(e)g(v_1 - s(i))^{-1}g(v_1 - s(i))g_i g(t(i) - v_2) & i \text{ points along the boundary of } p, \\ \overline{g(v_1 - s(i))}\partial(e)g(v_1 - s(i))^{-1}\overline{g(v_1 - s(i))}g_i g(t(i) - v_2) & i \text{ points against the boundary of } p \end{cases} \\ &= \begin{cases} \partial(g(v_1 - s(i)) \triangleright e)g(v_1 - v_2) & i \text{ points along the boundary of } p, \\ \partial(\overline{g(v_1 - s(i))} \triangleright e)g(v_1 - v_2) & i \text{ points against the boundary of } p. \end{cases} \end{aligned} \quad (\text{A2})$$

In addition, when i is on the path $v_1 - v_2$, then the path $v_1 - s(i)$ [or $\overline{v_1 - s(i)}$] along the boundary is not $(v_1 - v_2)[v_2 - s(i)]$ (or $(v_1 - v_2)[\overline{v_2 - s(i)}]$), as we can see from Fig. 51. Instead $(v_1 - v_2)[v_2 - s(i)] = [\text{boundary}(p)][v_1 - s(i)]$ if i points along the boundary and $(v_1 - v_2)[\overline{v_2 - s(i)}] = [\text{boundary}(p)]^{-1}[\overline{v_1 - s(i)}]$ if the edge i points against the boundary. If the plaquette satisfies fake flatness, then $g[\text{boundary}(p)]^{-1} = \partial(e_p)$. We can then use this to find the action of $E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2)$ on the plaquette label. We have

$$\begin{aligned} E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2) : e_p &= E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e : g(v_1 - v_2)^{-1} \triangleright e_p \\ &= E_p(v_1 \rightarrow v_2)^{-1} : \begin{cases} [g(v_1 - v_2)^{-1} \triangleright e_p] [g(v_2 - s(i)) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(v_2 - s(i)) \triangleright e] [g(v_1 - v_2)^{-1} \triangleright e_p] & i \text{ points against the boundary of } p. \end{cases} \end{aligned}$$

Then because the edge i is on the path $v_1 - v_2$, so that the label of this path is altered by the edge transform according to Eq. (A2), moving the base point of the plaquette back gives

$$E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2) : e_p = \begin{cases} (\partial[g(v_1 - s(i)) \triangleright e]g(v_1 - v_2)) \triangleright ([g(v_1 - v_2)^{-1} \triangleright e_p] [g(v_2 - s(i)) \triangleright e^{-1}]) & i \text{ points along the boundary of } p, \\ (\partial[\overline{g(v_1 - s(i))} \triangleright e]g(v_1 - v_2)) \triangleright ([\overline{g(v_2 - s(i))} \triangleright e] [g(v_1 - v_2)^{-1} \triangleright e_p]) & i \text{ points against the boundary of } p. \end{cases}$$

We can then use the Peiffer condition (5) to remove the factor of $\partial[g(v_1 - s(i)) \triangleright e]$ in $[\partial\{g(v_1 - s(i)) \triangleright e\}g(v_1 - v_2)]$ in favor of conjugation of the entire expression by $g(v_1 - s(i)) \triangleright e$ (and similar for the other orientation):

$$E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2) : e_p = \begin{cases} [g(v_1 - s(i)) \triangleright e][g(v_1 - v_2) \triangleright ([g(v_1 - v_2)^{-1} \triangleright e_p][g(v_2 - s(i)) \triangleright e^{-1}])[g(v_1 - s(i)) \triangleright e^{-1}] & i \text{ points along} \\ & \text{the boundary of } p, \\ [g(\overline{v_1 - s(i)}) \triangleright e][g(v_1 - v_2) \triangleright ([\overline{g(v_2 - s(i))} \triangleright e] [g(v_1 - v_2)^{-1} \triangleright e_p])[g(\overline{v_1 - s(i)}) \triangleright e^{-1}] & i \text{ points against} \\ & \text{the boundary of } p \end{cases}$$

$$= \begin{cases} [g(v_1 - s(i)) \triangleright e][e_p\{g(v_1 - v_2)g(v_2 - s(i)) \triangleright e^{-1}\}][g(v_1 - s(i)) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(\overline{v_1 - s(i)}) \triangleright e][\{[\overline{g(v_1 - v_2)g(v_2 - s(i))} \triangleright e]e_p\}[g(\overline{v_1 - s(i)}) \triangleright e^{-1}] & i \text{ points against the boundary of } p, \end{cases}$$

where in the last line we used the group homomorphism property $g \triangleright (e_1 e_2) = (g \triangleright e_1)(g \triangleright e_2)$ to distribute the action $g(v_1 - v_2) \triangleright$ across the terms in the curved brackets. Next we use the relationships

$$g(v_1 - s(i)) = g(\text{boundary}(p))^{-1} g(v_1 - v_2) g(v_2 - s(i)) = \partial(e_p) g(v_1 - v_2) g(v_2 - s(i))$$

and

$$g(\overline{v_1 - s(i)}) = g(\text{boundary}(p)) g(v_1 - v_2) g(\overline{v_2 - s(i)}) = \partial(e_p^{-1}) g(v_1 - v_2) g(\overline{v_2 - s(i)})$$

(where in each case the latter equality is obtained from fake flatness of the plaquette) to write

$$E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2) : e_p = \begin{cases} [g(v_1 - s(i)) \triangleright e][e_p\{[\partial(e_p^{-1})g(v_1 - s(i))]\} \triangleright e^{-1}][g(v_1 - s(i)) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(\overline{v_1 - s(i)}) \triangleright e][[\partial(e_p)g(v_1 - s(i))] \triangleright e]e_p[g(\overline{v_1 - s(i)}) \triangleright e^{-1}] & i \text{ points against the boundary} \end{cases}$$

$$= \begin{cases} [g(v_1 - s(i)) \triangleright e][e_p e_p^{-1} [g(v_1 - s(i)) \triangleright e^{-1}] e_p][g(v_1 - s(i)) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(\overline{v_1 - s(i)}) \triangleright e][e_p [g(\overline{v_1 - s(i)}) \triangleright e] e_p^{-1}][g(\overline{v_1 - s(i)}) \triangleright e^{-1}] & i \text{ points against the boundary,} \end{cases}$$

where we again used the Peiffer condition (5), this time on the factor $\partial(e_p)$. Then we can simplify the expression to give

$$E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2) : e_p = \begin{cases} [(v_1 - s(i)) \triangleright e][g(v_1 - s(i)) \triangleright e^{-1}]e_p[g(v_1 - s(i)) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(\overline{v_1 - s(i)}) \triangleright e]e_p[g(\overline{v_1 - s(i)}) \triangleright e][g(\overline{v_1 - s(i)}) \triangleright e^{-1}] & i \text{ points against the boundary} \end{cases}$$

$$= \begin{cases} e_p [g(v_1 - s(i)) \triangleright e^{-1}] & i \text{ points along the boundary of } p, \\ [g(\overline{v_1 - s(i)}) \triangleright e] e_p & i \text{ points against the boundary} \end{cases}$$

$$= \mathcal{A}_i^e : e_p.$$

This means that again the action of the edge transform, and therefore the edge energy term, is invariant under moving the base point. We also need to consider the case where i points against the path $v_1 - v_2$. The calculation for this case is very similar. Rather than go through it, we will justify the fact that this case works by noting that the edge energy term is invariant under flipping the edge (from previous calculations in this section). Similarly, the base-point moving procedure is invariant under such a flip. Therefore, we can always flip the edge so that it points along $v_1 - v_2$ before we apply either the edge transform or move the base point of the plaquette. This

means that if i points in the wrong way, we can apply P_i first. Then, using the above argument for the case where i points along the path, we have

$$\mathcal{A}_i^e P_i = E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2) P_i$$

which implies that $\mathcal{A}_i^e = E_p(v_1 \rightarrow v_2)^{-1} \mathcal{A}_i^e E_p(v_1 \rightarrow v_2)$.

We have therefore shown that each of the procedures for changing the structure of the lattice (flipping the orientations of edges and plaquettes or moving the base point of a plaquette) is consistent with each of the energy terms. It is important to note that to show this for some energy terms we had to

require that fake flatness was satisfied in the region where the energy operator has support, at least when \triangleright is nontrivial. This further shows the importance of fake flatness in ensuring the consistency of the higher-lattice gauge theory model.

4. Use of the rebranching procedures for other proofs

In addition to demonstrating the consistency of the higher-lattice gauge theory model, these procedures for changing the dressing of the lattice will be useful when considering the ribbon or membrane operators and their commutation relations with the energy terms, as we do in Refs. [71,72]. When we consider such commutation relations, the dressing of the lattice determines the action of both the energy terms and the ribbon operators, which means we need to consider several different cases for each commutation relation, depending on the branching structure in the region of the membrane operator and energy term. However, we can avoid this by instead showing that the membrane and ribbon operators are invariant under changing the dressing of the lattice in the same way as the energy terms are. In this case, we can demonstrate

the commutation relations for one choice of the branching structure. Then, because the operators are invariant under changes to the dressing, the commutation relation will hold for all choices of branching structure. Let \hat{X} be a series of these rebranching operations, and suppose a commutation relation between two operators \hat{O}_1 and \hat{O}_2 holds when acting on states that are defined when the lattice has a particular structure. Then the relation also holds for the branching structure produced by acting with \hat{X} on these states. For example, suppose the operators \hat{O}_1 and \hat{O}_2 commute for one choice of branching structure. Then for any state $|\psi\rangle$ defined on the original lattice, $\hat{O}_1\hat{O}_2|\psi\rangle = \hat{O}_2\hat{O}_1|\psi\rangle$. If the operators are consistent with the rebranching then $X^{-1}\hat{O}_iX = \hat{O}_i$. Therefore

$$\begin{aligned} X^{-1}\hat{O}_1XX^{-1}\hat{O}_2X|\psi\rangle &= X^{-1}\hat{O}_2XX^{-1}\hat{O}_1X|\psi\rangle \\ \Rightarrow X^{-1}\hat{O}_1\hat{O}_2X|\psi\rangle &= X^{-1}\hat{O}_2\hat{O}_1X|\psi\rangle \\ \Rightarrow \hat{O}_1\hat{O}_2X|\psi\rangle &= \hat{O}_2\hat{O}_1X|\psi\rangle, \end{aligned}$$

so that the commutation relation also holds for the state $X|\psi\rangle$, that is for states defined on the lattice with the altered branching structure.

-
- [1] L. D. Landau, On the theory of phase transitions, *Zh. Eksp. Teor. Fiz.* **7**, 19 (1937), in *Collected Papers of L. D. Landau*, edited by D. T. Haar (Pergamon, 1965), Chap. 29, pp. 193–216.
- [2] X.-G. Wen, Vacuum degeneracy of chiral spin states in compactified space, *Phys. Rev. B* **40**, 7387 (1989).
- [3] X.-G. Wen, Topological orders in rigid states, *Int. J. Mod. Phys. B* **04**, 239 (1990).
- [4] X.-G. Wen, Topological order: From long-range entangled quantum matter to a unified origin of light and electrons, *ISRN Condens. Matter Phys.* (2013).
- [5] D. C. Tsui, H. L. Stormer, and A. C. Gossard, Two-dimensional magnetotransport in the extreme quantum limit, *Phys. Rev. Lett.* **48**, 1559 (1982).
- [6] 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).
- [7] A. Stern, Anyons and the quantum Hall effect - A pedagogical review, *Ann. Phys.* **323**, 204 (2008).
- [8] T. Chakraborty and P. Pietlinen, *The Quantum Hall Effects: Integral and Fractional*, 2nd ed., Springer Series in Solid-State Sciences (Springer, New York, 1995).
- [9] S. D. Sarma and A. Pinczuk, *Perspectives in Quantum Hall Effects: Novel Quantum Liquids in Low-dimensional Semiconductor Structures* (Wiley, New York, 1997).
- [10] 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).
- [11] 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).
- [12] A. Mesaros and Y. Ran, Classification of symmetry enriched topological phases with exactly solvable models, *Phys. Rev. B* **87**, 155115 (2013).
- [13] A. Y. Kitaev, Fault-tolerant quantum computation by anyons, *Ann. Phys.* **303**, 2 (2003).
- [14] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, *J. Math. Phys.* **43**, 4452 (2002).
- [15] B. M. Terhal, Quantum error correction for quantum memories, *Rev. Mod. Phys.* **87**, 307 (2015).
- [16] B. J. Brown, D. Loss, J. K. Pachos, C. N. Self, and J. R. Wootton, Quantum memories at finite temperature, *Rev. Mod. Phys.* **88**, 045005 (2016).
- [17] J. C. Bridgeman, S. T. Flammia, and D. Poulin, Detecting topological order with ribbon operators, *Phys. Rev. B* **94**, 205123 (2016).
- [18] J. M. Leinaas and J. Myrheim, On the theory of identical particles, *Nuovo Cimento B* **37**, 1 (1977).
- [19] F. Wilczek, Magnetic flux, angular momentum, and statistics, *Phys. Rev. Lett.* **48**, 1144 (1982).
- [20] D. Arovas, J. R. Schrieffer, and F. Wilczek, Fractional statistics and the quantum Hall effect, *Phys. Rev. Lett.* **53**, 722 (1984).
- [21] J. K. Pachos, *Introduction to Topological Quantum Computation* (Cambridge University Press, Cambridge, 2012).
- [22] S. Rao, An anyon primer, [arXiv:hep-th/9209066](https://arxiv.org/abs/hep-th/9209066).
- [23] A. Y. Kitaev, Anyons in an exactly solved model and beyond, *Ann. Phys.* **321**, 2 (2006).
- [24] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. D. Sarma, Non-Abelian anyons and topological quantum computation, *Rev. Mod. Phys.* **80**, 1083 (2008).
- [25] V. T. Lahtinen and J. K. Pachos, A short introduction to topological quantum computation, *SciPost Phys.* **3**, 021 (2017).
- [26] S. Doplicher, R. Haag, and J. E. Roberts, Local observables and particle statistics I, *Commun. Math. Phys.* **23**, 199 (1971).
- [27] S. Doplicher, R. Haag, and J. E. Roberts, Local observables and particle statistics II, *Commun. Math. Phys.* **35**, 49 (1974).
- [28] C. Wang and M. Levin, Braiding statistics of loop excitations in three dimensions, *Phys. Rev. Lett.* **113**, 080403 (2014).

- [29] 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**, 251 (1992).
- [30] W. Ehrenberg and R. E. Siday, The refractive index in electron optics and the principles of dynamics, *Proc. Phys. Soc. B* **62**, 8 (1949).
- [31] Y. Aharonov and D. Bohm, Significance of electromagnetic potentials in the quantum theory, *Phys. Rev.* **115**, 485 (1959).
- [32] C.-H. Lin and M. Levin, Generalizations and limitations of string-net models, *Phys. Rev. B* **89**, 195130 (2014).
- [33] C.-H. Lin, M. Levin, and F. J. Burnell, Generalized string-net models: A thorough exposition, *Phys. Rev. B* **103**, 195155 (2021).
- [34] M. A. Levin and X.-G. Wen, String-net condensation: A physical mechanism for topological phases, *Phys. Rev. B* **71**, 045110 (2005).
- [35] M. Cheng, Z.-C. Gu, S. Jiang, and Y. Qi, Exactly solvable models for symmetry-enriched topological phases, *Phys. Rev. B* **96**, 115107 (2017).
- [36] F. A. Bais, P. van Driel, and M. de Wild Propitius, Quantum symmetries in discrete gauge theories, *Phys. Lett. B* **280**, 63 (1992).
- [37] M. de Wild Propitius and F. A. Bais, Discrete gauge theories, in *Particles and Fields*, CRM Series in Mathematical Physics, edited by G. Semenoff and L. Vinet (Springer, New York, 1999), pp. 353–439.
- [38] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A* **86**, 032324 (2012).
- [39] 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 (2020).
- [40] K. J. Satzinger, Y. Liu, A. Smith *et al.*, Realizing topologically ordered states on a quantum processor, *Science* **374**, 1237 (2021).
- [41] C. Heinrich, F. Burnell, L. Fidkowski, and M. Levin, Symmetry enriched string-nets: Exactly solvable models for SET phases, *Phys. Rev. B* **94**, 235136 (2016).
- [42] A. Hahn and R. Wolf, Generalized string-net model for unitary fusion categories without tetrahedral symmetry, *Phys. Rev. B* **102**, 115154 (2020).
- [43] E. Lake and Y.-S. Wu, Signatures of broken parity and time-reversal symmetry in generalized string-net models, *Phys. Rev. B* **94**, 115139 (2016).
- [44] I. Runkel, String-net models for nonspherical pivotal fusion categories, *J. Knot Theory Ramif.* **29**, 2050035 (2020).
- [45] Y. Hu, S. D. Stirling, and Y.-S. Wu, Ground state degeneracy in the Levin-Wen model for topological phases, *Phys. Rev. B* **85**, 075107 (2012).
- [46] 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**, 115421 (2008).
- [47] Y. Wan, J. C. Wang, and H. He, Twisted gauge theory model of topological phases in three dimensions, *Phys. Rev. B* **92**, 045101 (2015).
- [48] J. C. Wang and X.-G. Wen, Non-Abelian string and particle braiding in topological order: Modular $SL(3, \mathbb{Z})$ representation and 3+1D twisted gauge theory, *Phys. Rev. B* **91**, 035134 (2015).
- [49] A. Bullivant and C. Delcamp, Tube algebras, excitations statistics and compactification in gauge models of topological phases, *J. High Energy Phys.* **10** (2019) 216.
- [50] R. Dijkgraaf and E. Witten, Topological gauge theories and group cohomology, *Commun. Math. Phys.* **129**, 393 (1990).
- [51] D. J. Williamson and Z. Wang, Hamiltonian models for topological phases of matter in three spatial dimensions, *Ann. Phys.* **377**, 311 (2017).
- [52] K. Walker and Z. Wang, (3+1)-TQFTs and topological insulators, *Front. Phys.* **7**, 150 (2012).
- [53] C. W. von Keyserlingk, F. J. Burnell, and S. H. Simon, Three-dimensional topological lattice models with surface anyons, *Phys. Rev. B* **87**, 045107 (2013).
- [54] X. Chen, F. J. Burnell, A. Vishwanath, and L. Fidkowski, Anomalous symmetry fractionalization and surface topological order, *Phys. Rev. X* **5**, 041013 (2015).
- [55] Z. Wang and X. Chen, Twisted gauge theories in three-dimensional Walker-Wang models, *Phys. Rev. B* **95**, 115142 (2017).
- [56] A. Bullivant, M. Calcada, Z. Kadar, P. Martin, and J. F. Martins, Topological phases from higher gauge symmetry in 3+1D, *Phys. Rev. B* **95**, 155118 (2017).
- [57] C. Delcamp and A. Tiwari, From gauge to higher gauge models of topological phases, *J. High Energy Phys.* **10** (2018) 049.
- [58] A. Bullivant and C. Delcamp, Excitations in strict 2-group higher gauge models of topological phases, *J. High Energy Phys.* **01** (2020) 107.
- [59] 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**, 2050011 (2020).
- [60] D. Yetter, TQFT's from homotopy 2-types, *J. Knot Theory Ramif.* **02**, 113 (1993).
- [61] D. V. Else and C. Nayak, Cheshire charge in (3+1)-D topological phases, *Phys. Rev. B* **96**, 045136 (2017).
- [62] 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**, 031048 (2014).
- [63] H. Pfeiffer, Higher gauge theory and a non-Abelian generalization of 2-form electrodynamics, *Ann. Phys.* **308**, 447 (2003).
- [64] J. C. Baez and J. Huerta, An invitation to higher gauge theory, *Gen. Relativ. Gravitat.* **43**, 2335 (2011).
- [65] S. Gukov and A. Kapustin, Topological quantum field theory, nonlocal operators, and gapped phases of gauge theories, [arXiv:1307.4793](https://arxiv.org/abs/1307.4793).
- [66] A. Kapustin and R. Thorngren, Topological field theory on a lattice, discrete theta-angles and confinement, *Adv. Theor. Math. Phys.* **18**, 1233 (2014).
- [67] A. Kapustin and R. Thorngren, Higher symmetry and gapped phases of gauge theories, in *Algebra, Geometry, and Physics in the 21st Century: Kontsevich Festschrift*, edited by D. Auroux, L. Katzarkov, T. Pantev, Y. Soibelman, and Y. Tschinkel (Springer, Cham, 2017), pp. 177–202.
- [68] A. Bullivant, J. F. Martins, and P. Martin, Representations of the loop braid group and Aharonov-Bohm like effects in discrete (3+1)-dimensional higher gauge theory, *Adv. Theor. Math. Phys.* **23**, 1685 (2019).
- [69] J. C. Baez, A. S. Crans, and D. K. Wise, Exotic statistics for strings in 4d BF theory, *Adv. Theor. Math. Phys.* **11**, 707 (2007).

- [70] C. Damiani, A journey through loop braid groups, *Expo. Math.* **35**, 252 (2017).
- [71] J. Huxford and S. H. Simon, following paper, Excitations in the higher-lattice gauge theory model for topological phases. II. The (2+1)-dimensional case, *Phys. Rev. B* **108**, 245133 (2023).
- [72] 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](https://arxiv.org/abs/2206.09941).
- [73] I. Montvay and G. Münster, *Quantum Fields on a Lattice*, Cambridge Monographs on Mathematical Physics (Cambridge University Press, Cambridge, 1994).
- [74] K. G. Wilson, Confinement of quarks, *Phys. Rev. D* **10**, 2445 (1974).
- [75] B. Durhuus, On the structure of gauge invariant classical observables in lattice gauge theories, *Lett. Math. Phys.* **4**, 515 (1980).
- [76] J. C. Baez, Higher Yang-Mills theory, [arXiv:hep-th/0206130](https://arxiv.org/abs/hep-th/0206130).
- [77] J. McCool, On basis-conjugating automorphisms of free groups, *Can. J. Math.* **38**, 1525 (1986).
- [78] A. G. Savushkina, On the group of conjugating automorphisms of a free group, *Math Notes* **60**, 68 (1996).
- [79] R. Fenn, R. Rimanyi, and C. Rourke, The braid-permutation group, *Topology* **36**, 123 (1997).
- [80] D. M. Dahm, A generalization of braid theory, Ph.D. thesis, Princeton University, 1962.
- [81] D. L. Goldsmith, The theory of motion groups, *Michigan Math. J.* **28** (1981).
- [82] J. Preskill, Lecture notes for physics 219: Quantum computation, California Institute of Technology, 2004, <https://www.theory.caltech.edu/>.
- [83] F. A. Bais, Flux metamorphosis, *Nucl. Phys. B* **170**, 32 (1980).
- [84] M. Bucher, K.-M. Lee, and J. Preskill, On detecting discrete Cheshire charge, *Nucl. Phys. B* **386**, 27 (1992).
- [85] A. Komar and O. Landon-Cardinal, Anyons are not energy eigenspaces of quantum double Hamiltonians, *Phys. Rev. B* **96**, 195150 (2017).
- [86] A. H. Clifford, Representations induced in an invariant subgroup, *Ann. Math.* **38**, 533 (1937).
- [87] F. A. Bais and J. K. Slingerland, Condensate induced transitions between topologically ordered phases, *Phys. Rev. B* **79**, 045316 (2009).
- [88] F. J. Burnell, Anyon condensation and its applications, *Annu. Rev. Condens. Matter Phys.* **9**, 307 (2018).
- [89] T. Neupert, H. He, C. von Keyserlingk, G. Sierra, and B. A. Bernevig, Boson condensation in topologically ordered quantum liquids, *Phys. Rev. B* **93**, 115103 (2016).
- [90] F. A. Bais, B. J. Schroers, and J. K. Slingerland, Hopf symmetry breaking and confinement in (2+1)-dimensional gauge theory, *J. High Energy Phys.* **05** (2003) 068.
- [91] I. S. Eliëns, J. C. Romers, and F. A. Bais, Diagrammatics for Bose condensation in anyon theories, *Phys. Rev. B* **90**, 195130 (2014).
- [92] 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**, 235120 (2013).
- [93] 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**, 115104 (2016).
- [94] M. Atiyah, Topological quantum field theory, *Publ. Math. IHÉS* **68**, 175 (1988).
- [95] 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**, 021074 (2018).
- [96] V. Koppen, J. F. Martins, and P. P. Martin, Exactly solvable models for 2+1D topological phases derived from crossed modules of semisimple Hopf algebras, [arXiv:2104.02766](https://arxiv.org/abs/2104.02766).