Computational Power of Symmetry-Protected Topological Phases

David T. Stephen, Dong-Sheng Wang, Abhishodh Prakash, Tzu-Chieh Wei, and Robert Raussendorf

Phys. Rev. Lett. **119**, 010504 — Published 5 July 2017

DOI: 10.1103/PhysRevLett.119.010504

Computational power of symmetry-protected topological phases

David T. Stephen,[1] Dong-Sheng Wang,[1] Abhishodh Prakash,[2] Tzu-Chieh Wei,[2] and Robert Raussendorf[1]

[1] *Department of Physics and Astronomy, University of British Columbia, Vancouver, British Columbia V6T 1Z1, Canada*

[2] *C. N. Yang Institute for Theoretical Physics and Department of Physics and Astronomy, State University of New York at Stony Brook, Stony Brook, NY 11794-3840, USA*

(Dated: June 20, 2017)

We consider ground states of quantum spin chains with symmetry-protected topological (SPT) order as resources for measurement-based quantum computation (MBQC). We show that, for a wide range of SPT phases, the computational power of ground states is uniform throughout each phase. This computational power, defined as the Lie group of executable gates in MBQC, is determined by the same algebraic information that labels the SPT phase itself. We prove that these Lie groups always contain a full set of single-qubit gates, thereby affirming the long-standing conjecture that general SPT phases can serve as computationally useful phases of matter.

PACS numbers: 03.67.Mn, 03.65.Ud, 03.67.Ac

Introduction. In many-body physics, the essential properties of a quantum state are determined by the phase of matter in which it resides. Recent years have witnessed tremendous progress in the discovery and classification of quantum phases [1–10], and it is thus pertinent to ask: what can a phase of matter be used for? A traditional example is the ubiquitous superconductor, while newly discovered phases such as topological insulators [11] and quantum spin liquids [12] have promising future applications. Quantum phases are useful in quantum information processing as well: certain topological phases allow for error-resilient topological quantum computation via the braiding and fusion of their anyonic excitations [13, 14]. These applications all operate due to properties of a phase rather than a particular quantum state, hence they enjoy passive protection against certain sources of noise and error.

In this letter, we establish a general connection between the symmetry-protected topological (SPT) phases in one dimension (1D) [3–5] and quantum computation. To do this we use the framework of measurement-based quantum computation (MBQC) [15, 16], in which universal computation is possible using only single-body measurements on an entangled many-body system. The computational power of an MBQC scheme, defined by the set of logical gates that can be performed using measurements, is related to the entanglement structure of the many-body ground state. Whether this computational power is particular to individual states, or a property of a phase as above, is a long-standing open problem [17–29]. An important early result showed that every ground state within certain SPT phases has the ability to faithfully transport quantum information along a 1D chain; however, "universal" single-qubit gates appeared to be properties only of special points in the phases [27]. Later, it was shown that, for one particular SPT phase (namely one that is protected by S_4 symmetry), universal single-qubit gates can be implemented throughout the entire phase [23]. Yet, it remains unknown whether a general SPT phase can serve as such a *computational phase of matter*.

Here, we construct a general computational scheme that harnesses the part of a ground state that is fully constrained by symmetry. This part is uniform throughout the SPT phase, and therefore the computational power in our scheme is a property of SPT phases rather than individual states. This power is determined by the same algebraic structure that is used to classify the SPT phases, namely group cohomology. This establishes a firm connection between SPT order and the computational power of many-body ground states.

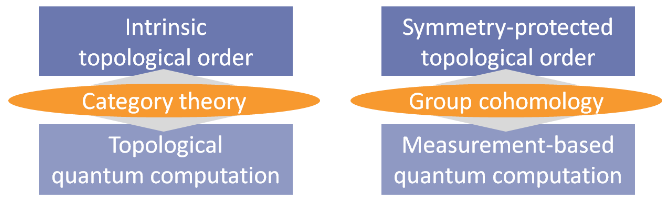

Figure 1. In the same way that the language of category theory allows us to classify gates that can be executed by braiding the anyonic excitations of topologically ordered systems in 2D [14, 30, 31], group cohomology determines the gates implementable in measurement-based quantum computation using 1D resource states with symmetry-protected topological order.

We can use this connection to prove that universal single-qubit gates are a property of all phases considered by Ref. [27], and many more. Going beyond this, we identify classes of phases that also allow operations on qudits of arbitrarily large dimension. Overall, our results highlight how the algebraic classification of quantum phases can contribute to the study of the structures responsible for quantum computational power, as outlined in Fig. 1

In the following, we begin by reviewing the virtual space picture of MBQC [32], which aids our subsequent analysis. We then introduce the three key elements of our scheme, and demonstrate their use through the examples of the AKLT state and the Haldane phase before generalizing to other SPT phases. We finish by using the algebraic classification of SPT phases to determine their computational power.

Computation in Virtual Space. We consider MBQC in the virtual space picture, where states are represented in the matrix product state (MPS) form [33]. The wave function $|\psi\rangle$ of a 1D system of N interacting sites of local dimension d (i.e. a spin chain) can be written in MPS form by introducing the square matrices A^i, $i = 0, \ldots, d-1$, such that

$$|\psi\rangle = \sum_{i_1,\ldots,i_N} \langle R|A^{i_N} A^{i_{N-1}} \ldots A^{i_1}|L\rangle |i_1 \ldots i_N\rangle, \quad (1)$$

where $|R\rangle$, $|L\rangle$ are states in the so-called "virtual space" that encode the boundary conditions of the finite chain. The MPS formalism leads to a useful interpretation of MBQC which occurs in virtual space [32]: measuring the leftmost spin in the chain with outcome $|s\rangle$ reduces chain length by one and evolves the virtual system as $|L\rangle \to \left(\sum_i \langle s|i\rangle A^i\right)|L\rangle$. With a proper choice of measurement basis, this can correspond to unitary evolution and can simulate computation up to outcome-dependent byproduct operators. Since we consider only 1D resource states, we say a state is *universal* if measurement can induce a full set of gates for a single qudit, corresponding to operators in $SU(D)$ on some D-level subspace in virtual space.

A simple example is the spin-1 AKLT state, which is well-known to be a universal resource [34]. The MPS matrices are the Pauli matrices, $A^i = \sigma^i$, with respect to the *wire basis* $\mathcal{B} = \{|x\rangle, |y\rangle, |z\rangle\}$ where $|i\rangle$ is the 0 eigenstate of the spin-1 operator S^i. To achieve a rotation by θ about the z-axis, we measure in the basis $\mathcal{B}(z,\theta) = \{|\theta_x\rangle, |\theta_y\rangle, |z\rangle\} \equiv \left\{\cos\frac{\theta}{2}|x\rangle - \sin\frac{\theta}{2}|y\rangle, \sin\frac{\theta}{2}|x\rangle + \cos\frac{\theta}{2}|y\rangle, |z\rangle\right\}$ and propagate the byproducts σ^x, σ^y, and σ^z, respectively. We enact byproduct propagation via symmetry transformations of future measurement bases, as described in Ref. [22]. With this, the first two outcomes give the desired rotation by θ while the third does nothing, so the gate is probabilistic with success probability $\frac{2}{3}$. Rotations about the x-axis can be achieved similarly, giving a full set of $SU(2)$ operations.

To extend the universality of the AKLT state and others like it to entire SPT phases, we introduce three modifications to the usual MBQC procedure, as described in Fig. 2. The purpose and justification of each are given in the following section, using the AKLT state and Haldane phase as examples.

Computation in the Haldane phase. We begin this section by introducing the "mixed state interpretation" of MBQC that will be used throughout this letter. Here we argue its validity, with a formal proof given in the Supplemental Material [35]. We define a computation by a sequence of n measurement bases, which are fixed modulo byproduct propagation. In general, an input state $|\psi\rangle$ will be taken to a final state $|\psi_{\vec{s}}\rangle$ which depends on the measurement outcomes $\vec{s} = (s_1, \ldots, s_n)$. Then we measure some observable O on $|\psi_{\vec{s}}\rangle$, whose eigenvalues o_i appear with probability $p(o_i|\vec{s})$. To garner measurement statistics of O, we must repeat the computation, whereupon the full statistics are given by $p(o_i) = \sum_{\vec{s}} p(o_i|\vec{s}) p_{\vec{s}}$ where $p_{\vec{s}}$ is the probability of outcomes $\vec{s}$. These statistics are encoded in the mixed state $\hat{\sigma} = \sum_{\vec{s}} p_{\vec{s}} |\psi_{\vec{s}}\rangle\langle\psi_{\vec{s}}|$, for instance $\langle O\rangle = \sum_{\vec{s}} p_{\vec{s}} \langle\psi_{\vec{s}}|O|\psi_{\vec{s}}\rangle \equiv \mathrm{Tr}\,(O\hat{\sigma})$. Hence in this probabilistic scenario the computational output must be interpreted to be $\hat{\sigma}$.

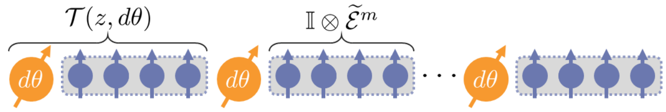

Figure 2. Illustration of the measurements needed to execute a rotation about the z-axis in the Haldane phase example. Our scheme consists of three modifications to the usual MBQC procedure: (1) In analysis of the scheme, measurement outcomes are summed over, such that the computational output is interpreted as a mixed state, (2) finite rotations are split into smaller pieces $d\theta$ that each differ only slightly from the identity, and (3) consecutive gates are separated by many applications of the identity gate.

To determine the mixed state $\hat{\sigma}$, we simply sum over all possible outcomes of each measurement. It is crucial that this sum-over-outcomes is implemented *after* byproduct propagation, making it very different from simply tracing over each spin in the chain. The byproducts accumulated at the end of the computation affect the basis of computational readout, during which we do not sum over outcomes. By analysing the computation in this way, we can design a sequence of measurement bases such that $\hat{\sigma}$ approximates the desired output. If the computation defined by this sequence of measurements is repeated many times, it will deterministically produce the desired measurement statistics of any observable O, even though each run of the algorithm may produce a different output state that is meaningless on its own.

Let us return to the AKLT state as an example. By measuring in the basis $\mathcal{B}(z,\theta)$ and summing over measurement outcomes, we find that an initial state $|L\rangle\langle L|$ becomes:

$$\hat{\sigma} = \frac{2}{3} e^{-i\theta\sigma_z/2}|L\rangle\langle L|e^{i\theta\sigma_z/2} + \frac{1}{3}|L\rangle\langle L|. \quad (2)$$

Since the original gate is probabilistic, this is a mixed state and does not represent unitary evolution. However, for small angles $d\theta$, it is unitary up to first order:

$$\hat{\sigma} = e^{-i\frac{2}{3}d\theta\sigma_z/2}|L\rangle\langle L|e^{i\frac{2}{3}d\theta\sigma_z/2} + \mathcal{O}(d\theta^2). \quad (3)$$

So for small rotation angles $d\theta$, the mixed output state is our initial state rotated by a reduced angle $\frac{2}{3}d\theta$ about the z-axis. Restriction to gates that are close to the identity is an unavoidable consequence of the mixed state interpretation, and finite rotations must be split into many infinitesimal pieces [36]. The number of measurements needed to execute a unitary gate with rotation angle θ and admissible error ϵ is $\mathcal{O}(\theta^2/\epsilon)$; details can be found in Sec. II of the Supplemental Material.

The AKLT state is in the Haldane phase, which we define as the SPT phase protected by on-site $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry [37]. Every state in the Haldane phase can be viewed as an AKLT state with some additional entanglement that encodes the microscopic details of the state. This is formally expressed in terms of the MPS matrices, which factorize as $A^i = \sigma^i \otimes B^i$ in the wire basis $\mathcal{B}$ [27]. The Pauli part acts in the *logical subspace* into which information is encoded and processed. The matrices B^i act in the *junk subspace* and contain all of the microscopic details of the state. Importantly, byproduct propagation via symmetry transformations acts only within the logical subspace. This is not a problem for measurements in the wire basis, which evolve the two subsystems independently. But measurements in other bases will mix the junk and logical subspaces, which hides the logical information and introduces an unavoidable outcome dependence into the computation. We now show how the mixed state interpretation allows us to solve both of these problems in a relatively simple way.

Consider a measurement in the infinitesimally tilted basis $\mathcal{B}(z, d\theta)$. Without loss of generality, we assume that our initial state is factorized across the subspaces as $|\phi\rangle\langle\phi| \otimes \rho_{\text{fix}}$, for a particular fixed point state ρ_{fix} that will be defined later. If we get the outcome $|\theta_x\rangle$ and propagate σ^x on the logical subspace, our state becomes:

$$
\begin{aligned}
&|\phi\rangle\langle\phi| \otimes \rho_{\text{fix}} \to |\phi\rangle\langle\phi| \otimes B^x \rho_{\text{fix}} B^{x\dagger} \\
&+ i\frac{d\theta}{2}\left(|\phi\rangle\langle\phi|\sigma^z \otimes B^x\rho_{\text{fix}}B^{y\dagger} - \sigma^z|\phi\rangle\langle\phi| \otimes B^y\rho_{\text{fix}}B^{x\dagger}\right),
\end{aligned} \tag{4}
$$

up to first order in $d\theta$. We see that the two subsystems are no longer factorized, and the logical state $|\phi\rangle\langle\phi|$ is no longer accessible.

To remedy this, we will flow the junk subspace towards a fixed point. This is accomplished by simply measuring a large number of spins in the wire basis. In the mixed state interpretation, a measurement in the wire basis followed by logical byproduct propagation effects the operation $\mathbb{I} \otimes \sum_i B^i(\cdot)B^{i\dagger} \equiv \mathbb{I} \otimes \widetilde{\mathcal{E}}$. Since every state in the Haldane phase is short-range correlated, the channel $\widetilde{\mathcal{E}}$ will have a unique fixed point, which is ρ_{fix}, with all other eigenvalues of modulus less than unity [22]. Hence measuring m consecutive spins in the wire basis results in the linear channel $\mathbb{I} \otimes \widetilde{\mathcal{E}}^m$ and projects the junk subspace onto the fixed point ρ_{fix}. The projection occurs exponentially fast over the correlation length ξ of the state.

Applying this to Eq. 4, which must be summed with its counterparts for the other measurement outcomes $|\theta_y\rangle$ and $|z\rangle$, we find that for large enough m,

$$
\hat{\sigma} = \left(\nu|\phi\rangle\langle\phi| + i\frac{d\theta}{2}(\nu_{xy}+\nu_{yx})\left[|\phi\rangle\langle\phi|, \sigma^z\right]\right) \otimes \rho_{\text{fix}}, \tag{5}
$$

where we have defined $\lim_{m\to\infty} \widetilde{\mathcal{E}}^m(B^i\rho_{\text{fix}}B^{j\dagger}) = \nu_{ij}\rho_{\text{fix}}$ and $\nu = \nu_{xx} + \nu_{yy} + \nu_{zz}$. Up to first order in $d\theta$, this corresponds to a unitary rotation acting on the logical subspace:

$$
\mathcal{T}(z, d\theta) = \exp\left\{-id\theta\left(\frac{\nu_{xy}+\nu_{yx}}{2\nu}\right)\sigma^z\right\}. \tag{6}
$$

Hence, making a measurement in the rotated basis $\mathcal{B}(z, d\theta)$, followed by a series of measurements in the wire basis, produces the desired rotation of the virtual state $|\phi\rangle$ up to a scaling factor $\frac{\nu_{xy}+\nu_{yx}}{\nu}$. As long as this factor is non-zero, it can be measured on the chain prior to computation by attempting a finite rotation (split into small pieces), and measuring the reduction in rotation angle [38]. The parameters ν_{ij} contain all relevant microscopic details of our resource state $|\psi\rangle$. Since they can be measured during a calibration step, any state in the phase can be used as a resource without prior knowledge of its identity.

We can repeat the above procedure for rotations about the x-axis to generate all of $SU(2)$. Hence every state in the Haldane phase, with the exception of a null subset in which some of the constants ν_{ij} are 0, has the same computational power as the AKLT state (which satisfies $\nu_{ij} = \frac{1}{3}\ \forall i, j$). To complete the scheme, we would require a method to read out and initialize the virtual state which also works throughout the phase. This can be done without the need of ancillary systems on the boundaries [38].

Generalization to Other Phases. Our scheme does not depend on any properties that are particular to the Haldane phase, so it can be generalised to a large class of other SPT phases. A general 1D SPT phase without symmetry breaking is defined with respect to an on-site symmetry group G such that $u(g)^{\otimes n}|\psi\rangle = |\psi\rangle$ for some unitary representation u of G. The phase is then labelled by a cohomology class $[\omega] \in H^2(G, U(1))$ in the second cohomology group of G which describes how this symmetry acts in the virtual space [3].

The Haldane phase is an example of a maximally non-commutative SPT phase, as defined in Ref. [27]. Such phases satisfy all conditions needed to apply our methods, namely the existence of a logical subspace and the ability to propagate byproduct operators within it. Indeed, suppose that G is finite abelian and $[\omega]$ is maximally non-commutative, meaning $\{g \in G | \omega(g, g') = \omega(g', g)\ \forall g' \in G\} = \{e\}$. By diagonalizing the representation u, we obtain the wire basis $\mathcal{B} = \{|0\rangle, \ldots, |d-1\rangle\}$ such that $u(g)|i\rangle = \chi_i(g)|i\rangle\ \forall g \in G$ where $\chi_i(g)$ are linear characters of G. Maximal non-commutativity then implies the MPS tensor A^i can be written in the wire basis as [27]:

$$
A^i = C^i \otimes B^i, \tag{7}
$$

where C^i are $D \times D$ unitary and trace-orthogonal matrices and $D = \sqrt{|G|}$ is the dimension of our logical subspace [39]. C^i can be determined uniquely from G, $[\omega]$, and χ_i as described in the Supplemental Material. In general, if some group G has a finite abelian subgroup H such that $[\omega|_H]$ is maximally non-commutative, we can make the exact same argument with H taking the place of G everywhere. This means the following results also apply to certain non-abelian groups and Lie groups.

Now we follow the same steps used to perform computation in the Haldane phase. Measurement in the slightly tilted basis $\mathcal{B}(i, j; d\theta, \varphi) = \{|0\rangle, \ldots, |i\rangle + d\theta e^{i\varphi}|j\rangle, |j\rangle - d\theta e^{-i\varphi}|i\rangle, \ldots, |d-1\rangle\}$, followed by measurements in $\mathcal{B}$ to drive the junk subspace to a fixed-point state, induces an infinitesimal rotation in the logical subspace:

$$
\begin{aligned}
&\mathcal{T}(i, j; d\theta, \varphi) = \\
&\exp\left\{d\theta\frac{|\nu_{ij}|}{\nu}\left(e^{i(\varphi+\delta_{ij})}C^{i\dagger}C^j - e^{-i(\varphi+\delta_{ij})}C^{j\dagger}C^i\right)\right\},
\end{aligned} \tag{8}
$$

where $\nu_{ij} = |\nu_{ij}|e^{i\delta_{ij}}$ is as defined earlier and $\nu = \sum_{i=0}^{d-1} \nu_{ii}$. As before, the microscopic details of the state enter only as these measurable constants. Computation can only proceed if these constants are non-zero, which is satisfied for all but a null set of states. With knowledge of these constants, $\mathcal{B}(i, j; d\theta, \varphi)$ can be chosen such that the primitive gates are generated by elements of the set of anti-hermitian operators:

$$\mathcal{O} = \{\alpha C^{i\dagger} C^j - \alpha^* C^{j\dagger} C^i\} \tag{9}$$

with $i, j = 0 \dots d-1$, $i \neq j$, $|\alpha| \ll 1$. Furthermore, we have $e^{d\theta A} e^{d\theta B} e^{-d\theta A} e^{-d\theta B} \approx e^{(d\theta)^2 [A,B]}$, so that our infinitesimal generators form a real Lie algebra which in turn generates a Lie group $\mathcal{L}[\mathcal{O}]$ of executable gates.

From the above, we can see the main strength of our methods. Given only the algebraic quantities G, u, and $[\omega]$ which describe the SPT phase of our resource state, we are able to define a complete MBQC scheme, including the set of gates and the measurements needed to execute them. The computational power of each state in the phase is uniformly defined as the Lie group $\mathcal{L}[\mathcal{O}]$, which is completely determined by the same algebraic quantities. This signifies the existence of a deep connection between SPT order and MBQC via the language of group cohomology.

Determining Computational Power. To determine the computational power of a phase, we must identify the Lie group $\mathcal{L}[\mathcal{O}]$. We will do this by taking advantage of the algebraic structure inherited from the SPT phase classification. Consider first the case where the representation $u|_H$ contains all non-trivial characters of the subgroup H. This means that $\mathcal{O}$ contains $D^2 - 1$ trace-orthogonal, antihermitian operators, so $\mathcal{L}[\mathcal{O}] \cong SU(D)$. If the Hilbert space dimension of our physical sites is smaller than $D^2 - 1$, or certain characters χ_i do not appear in $u|_H$, $\mathcal{L}[\mathcal{O}]$ may be some Lie subgroup of $SU(D)$. However, with the condition of maximal non-commutativity, this subgroup is always universal on a qudit system, as stated in the following theorem:

Theorem 1. *Consider an SPT phase defined by an on-site symmetry group G and cohomology class $[\omega]$. Suppose there exists a finite abelian subgroup $H \subset G$ such that $[\omega|_H]$ is maximally non-commutative, and let p^n be a prime power dividing $\sqrt{|H|}$. Then $\mathcal{L}[\mathcal{O}] \supset SU(p^n)$.*

This result, proven in the Supplemental Material, determines the *minimal* computational power of the phase, which is *independent* of u and hence uniform amongst the phase. This shows that 1D ground states with SPT order are generically useful as MBQC resources.

Beyond this minimal case, $\mathcal{L}[\mathcal{O}]$ can often be expanded to gain additional computational power. For example when $H = (\mathbb{Z}_2)^4$, our theorem guarantees that $SU(2) \subset \mathcal{L}[\mathcal{O}]$, but this can be expanded to either $SU(4)$ or $SU(2) \times SU(2)$ depending on the on-site symmetry representation u. So, while changing u is generally considered to not change the SPT phase of a system [4], it remains an important label for total computational power in our scheme. If, however, we allow ourselves to redefine the locality of measurements by blocking neighbouring sites, $\mathcal{L}[\mathcal{O}]$ will always equal $SU(D)$ after sufficient blocking.

Now we must ask: which symmetry groups protect phases that satisfy our theorem? To answer this in general is a difficult problem of group cohomology, but we can identify some particularly relevant examples. When G is a classical Lie group (except $Spin(4n)$), there is a subgroup of the form $\mathbb{Z}_N \times \mathbb{Z}_N \subset G$ such that $H^2(G, U(1)) \cong H^2(\mathbb{Z}_N \times \mathbb{Z}_N, U(1))$ [40, 41]. Since $\mathbb{Z}_N \times \mathbb{Z}_N$ protects a maximally non-commutative phase [27, 42], G must protect a phase which satisfies our theorem. The same can be said for any subgroup G' such that $\mathbb{Z}_N \times \mathbb{Z}_N \subset G' \subset G$. This has already been observed in Ref. [24] for the groups $D_4, A_4, S_4 \subset SO(3)$, which each contain $\mathbb{Z}_2 \times \mathbb{Z}_2$. Another example is the class of groups for which the subgroup H specified in Theorem 1 appears as a (semi)direct factor, that is $G = H' \rtimes H$ for some subgroup H' which could represent eg. time reversal symmetry [43].

Conclusion. By introducing three simple modifications to the usual MBQC procedure, we showed that the MBQC power of an SPT-ordered ground state of a spin chain is determined solely by the cohomological information that labels the corresponding SPT phase, and that this power is always sufficient for universal computation on a single qudit. Regarding the algebraic classification of phases of matter and its role in quantum computation, our results show that group cohomology links SPT order and MBQC in 1D, in the same way that modular tensor categories link topological order and topological quantum computation in 2D [14, 30, 31]. In each case, the algebraic framework that classifies the phases of matter also classifies their computational properties. Whether this extends to higher dimensions and other types of quantum phases is an intriguing question at the intersection of quantum information and condensed matter physics. There is already evidence that SPT order in higher dimensions can lead to unique computational properties [21, 26, 29, 44–47]. It would also be interesting to see whether the mathematical frameworks that unify topological order and SPT order, such as G-crossed braided tensor categories [6], could also describe computation with systems that have both types of order.

This work is supported by NSERC, the Canadian Institute for Advanced Research (Cifar), and the National Science Foundation under Grant No. PHY 1620252. R.R. is a fellow of the Cifar Quantum Information Science program.

[1] M. A. Levin and X.-G. Wen, Phys. Rev. B **71**, 045110 (2005).
[2] L. Fidkowski and A. Kitaev, Phys. Rev. B **83**, 075103 (2011).
[3] X. Chen, Z.-C. Gu, and X.-G. Wen, Phys. Rev. B **84**, 235128 (2011).
[4] N. Schuch, D. Pérez-García, and I. Cirac, Phys. Rev. B **84**, 165139 (2011).
[5] X. Chen, Z.-C. Gu, Z.-X. Liu, and X.-G. Wen, Phys. Rev. B **87**, 155114 (2013).

[6] M. Barkeshli, M. Bonderson, M. Cheng, and Z. Wang, (2014), arXiv:1410.4540.
[7] Z.-C. Gu and X.-G. Wen, Phys. Rev. B **90**, 115141 (2014).
[8] A. Kapustin, R. Thorngren, A. Turzillo, and Z. Wang, Journal of High Energy Physics **2015**, 52 (2015).
[9] H. Song, S.-J. Huang, L. Fu, and M. Hermele, Phys. Rev. X **7**, 011020 (2017).
[10] T. Lan, L. Kong, and X.-G. Wen, (2016), arXiv:1602.05946.
[11] X.-L. Qi and S.-C. Zhang, Rev. Mod. Phys. **83**, 1057 (2011).
[12] L. Balents, Nature **464**, 199 (2010).
[13] A. Kitaev, Annals of Physics **303**, 2 (2003).
[14] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma, Rev. Mod. Phys. **80**, 1083 (2008).
[15] R. Raussendorf and H. J. Briegel, Phys. Rev. Lett. **86**, 5188 (2001).
[16] R. Raussendorf, D. E. Browne, and H. J. Briegel, Phys. Rev. A **68**, 022312 (2003).
[17] S. D. Barrett, S. D. Bartlett, A. C. Doherty, D. Jennings, and T. Rudolph, Phys. Rev. A **80**, 062328 (2009).
[18] A. C. Doherty and S. D. Bartlett, Phys. Rev. Lett. **103**, 020506 (2009).
[19] A. Miyake, Phys. Rev. Lett. **105**, 040501 (2010).
[20] S. D. Bartlett, G. K. Brennen, A. Miyake, and J. M. Renes, Phys. Rev. Lett. **105**, 110502 (2010).
[21] A. S. Darmawan, G. K. Brennen, and S. D. Bartlett, New. J. Phys. **14**, 013023 (2012).
[22] D. V. Else, S. D. Bartlett, and A. C. Doherty, New Journal of Physics **14**, 113016 (2012).
[23] J. Miller and A. Miyake, Phys. Rev. Lett. **114**, 120506 (2015).
[24] A. Prakash and T.-C. Wei, Phys. Rev. A **92**, 022310 (2015).
[25] T.-C. Wei, Y. Li, and L. C. Kwek, Phys. Rev. A **89**, 052315 (2014).
[26] H. P. Nautrup and T.-C. Wei, Phys. Rev. A **92**, 052309 (2015).
[27] D. V. Else, I. Schwarz, S. D. Bartlett, and A. C. Doherty, Phys. Rev. Lett. **108**, 240505 (2012).
[28] D.-S. Wang, D. T. Stephen, and R. Raussendorf, Phys. Rev. A **95**, 032312 (2017).
[29] T.-C. Wei and C.-Y. Huang, (2017), arXiv:1705.06833.
[30] M. Freedman, A. Kitaev, M. Larsen, and Z. Wang, Bull. Amer. Math. Soc. **40**, 31 (2003).
[31] E. Rowell, R. Stong, and Z. Wang, Communications in Mathematical Physics **292**, 343 (2009).
[32] D. Gross and J. Eisert, Phys. Rev. Lett. **98**, 220503 (2007); D. Gross, J. Eisert, N. Schuch, and D. Perez-Garcia, Phys. Rev. A **76**, 052315 (2007).
[33] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, Quant. Inf. Comp. **7**, 401 (2007), arXiv:0608197.
[34] G. K. Brennen and A. Miyake, Phys. Rev. Lett. **101**, 010502 (2008).
[35] See Supplemental Material [url], which includes Refs. [48–50], for more information on the mixed state interpretation, the efficiency and overhead of our scheme, and the proof of the main theorem.
[36] Each of the infinitesimal rotations that compose a finite gate must be implemented one after another while propagating byproducts.
[37] F. Pollmann, E. Berg, A. M. Turner, and M. Oshikawa, Phys. Rev. B **85**, 075125 (2012).
[38] R. Raussendorf, D.-S. Wang, A. Prakash, T.-C. Wei, and D. T. Stephen, (2016), arXiv:1609.07549.
[39] D is always defined since any finite abelian group which supports a maximally non-commutative factor system must have the form $G \cong G' \times G'$ for some subgroup G' [42].
[40] K. Duivenvoorden and T. Quella, Phys. Rev. B **88**, 125115 (2013).
[41] K. Duivenvoorden and T. Quella, Phys. Rev. B **87**, 125145 (2013).
[42] I. G. Berkovich and E. Zhmud, *Characters of finite groups*, Vol. 1 (American Mathematical Soc., 1998).
[43] Z. Xiong, (2016), arXiv:1701.00004.
[44] J. Miller and A. Miyake, Npj Quant. Inf. **2**, 16036 (2016).
[45] J. Miller and A. Miyake, (2016), arXiv:1612.08135.
[46] B. Yoshida, (2015), arXiv:1509.03626.
[47] B. Yoshida, Phys. Rev. B **93**, 155131 (2016).
[48] G. Karpilovsky, *Group representations*, Vol. 3 (Elsevier, 1994).
[49] M. Sanz, D. Pérez-García, M. M. Wolf, and J. I. Cirac, IEEE Trans. Inf. Theor. **56**, 4668 (2010).
[50] D. L. Zhou, B. Zeng, Z. Xu, and C. P. Sun, Phys. Rev. A **68**, 062303 (2003).