

Resource Costs for Fault-Tolerant Linear Optical Quantum Computing

Ying Li,¹ Peter C. Humphreys,² Gabriel J. Mendoza,³ and Simon C. Benjamin¹

¹*Department of Materials, University of Oxford, Parks Road, Oxford OX1 3PH, United Kingdom*

²*Clarendon Laboratory, Department of Physics, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom*

³*Centre for Quantum Photonics, H.H. Wills Physics Laboratory and Department of Electrical and Electronic Engineering, University of Bristol, Merchant Venturers Building, Woodland Road, Bristol BS8 1UB, United Kingdom*

(Received 26 April 2015; revised manuscript received 11 August 2015; published 14 October 2015)

Linear optical quantum computing (LOQC) seems attractively simple: Information is borne entirely by light and processed by components such as beam splitters, phase shifters, and detectors. However, this very simplicity leads to limitations, such as the lack of deterministic entangling operations, which are compensated for by using substantial hardware overheads. Here, we quantify the resource costs for full-scale LOQC by proposing a specific protocol based on the surface code. With the caveat that our protocol can be further optimized, we report that the required number of physical components is at least 5 orders of magnitude greater than in comparable matter-based systems. Moreover, the resource requirements grow further if the per-component photon-loss rate is worse than 10^{-3} or the per-component noise rate is worse than 10^{-5} . We identify the performance of switches in the network as the single most influential factor influencing resource scaling.

DOI: [10.1103/PhysRevX.5.041007](https://doi.org/10.1103/PhysRevX.5.041007)

Subject Areas: Optics, Quantum Information

I. INTRODUCTION

Numerous different physical systems have been explored as platforms for quantum information processing. Most approaches involve embodying information in matter systems such as ions or superconducting qubits, but a striking alternative is linear optical quantum computing (LOQC), where all information is encoded in electromagnetic field modes and processing is carried out using only linear optical elements. The remarkable fact that such an approach is possible was first established in 2001 by Knill, Laflamme, and Milburn [1]; for a review see Ref. [2]. Using light as the information medium takes advantage of the low decoherence suffered by optical fields and the relative ease with which quantum information can be encoded photonically. However there are drawbacks, in particular, the impossibility of deterministic entanglement and the impact of photon loss (whether due to absorption, leakage, or detector failure). Such difficulties can be solved by increasing the physical complexity of the circuitry. Thus, while LOQC may benefit from simple building blocks, conversely, it may require more complex circuits than other approaches, and the balance of these factors will determine whether the approach is a practical competitor to matter-based processors. A number of studies have explored the

benefits of hybrid systems, where modules that use matter-light interactions can enrich the available manipulations (see, e.g., Ref. [3]), but our interest in the present paper is to explore the feasibility of “all-photonic” LOQC.

The most developed method for LOQC to date is based on a discrete dual-rail encoding, in which each qubit is encoded in the field modes occupied by a single photon [4], as distinct from coherent-state approaches [5]. The dual-rail modes can be spatial, polarization, time-frequency, or any other degree of freedom supported by electromagnetic fields. Crucially, even though entangling operations between dual-rail-encoded photonic qubits cannot succeed deterministically, it is nonetheless possible to build an essentially deterministic universal quantum computer using only linear optics. This can be achieved by attempting probabilistic entangling operations (PEO) between many resource states in order to ensure that, with high probability, a sufficient number of operations will succeed to allow for quantum computing [6,7].

Techniques for mitigating photon loss have also been developed. It has been shown that if quantum information is suitably encoded in a multiphoton state, then losses of up to 50% of the photons can be tolerated before the encoded information is lost [8] (in the limit of ideal sources and detectors in the state-encoding process). However, in any realistic implementation of a quantum computer, one must account for how complex multiphoton states can be created, given that every component, at every level, will be associated with finite rates of photon loss and other forms of noise. Furthermore, the circuitry associated with

Published by the American Physical Society under the terms of the Creative Commons Attribution 3.0 License. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

overcoming nondeterministic entanglement will require many linear optical elements, including delay lines and switching networks, in order to dynamically reroute the outputs of successful operations to the next stage of processing. These elements will induce further errors and losses, and in this sense, the twin issues of nondeterministic entanglement and photon loss aggravate one another in LOQC. Fortunately, the threshold theorem assures us that, if all physical error rates are sufficiently low, then errors at the logical level can be made arbitrarily rare, and scalable fault-tolerant quantum computing can be achieved. The central challenge of quantum computing is therefore to demonstrate the operations necessary for quantum computing with error rates below these thresholds. Theoretical studies have established the required thresholds for architectures relevant to superconducting qubits [9] and to matter-optical networks [10], and experimental systems have been demonstrated at, or beyond, the required performance levels [11–13]. However, to our knowledge, no prior paper has established requirements of LOQC at the per-component level while simultaneously tracking the overall resource costs.

In this paper, we propose a protocol for LOQC that includes every step from the initial generation of entanglement primitives to the deployment of a fully fault-tolerant scalable unit for quantum computing. We consider computational errors and losses at each stage and endeavor to employ the most efficient known protocols for optical quantum information processing. In contrast to previous studies of noise thresholds in optical quantum computing [14,15], we explicitly account for the substantial resource costs of LOQC protocols. We focus on a purely linear optical network, without employing matter qubits as memories or for entanglement generation. However, we do assume the availability of on-demand sources of single photons, without concerning ourselves with the particular method with which these would be generated [16].

It is important to recognize that our results only represent an upper bound on the physical characteristics that are required of the components in an LOQC system: Our protocol may admit various further optimizations, which will make the physical requirements somewhat less stringent, and moreover, it is possible that there is an as-yet undiscovered alternative approach that can achieve significantly lower resource costs. Nevertheless, we believe the results we offer are highly relevant to the field, having been derived from protocols that represent the present “state of the art,” and moreover, it is therefore fair to compare the results here with those that have been derived for matter and hybrid matter-optical systems.

Our analysis allows us to make an estimate of the overall scale of the resources necessary to construct a fully fault-tolerant optical quantum computer. We choose the number of detectors as a metric for the device size, recognizing that the total numbers for the other kinds of component will scale roughly proportionately. We find that one would

require upwards of 10^5 detectors per physically encoded qubit in the cluster state, therefore requiring a total of at least 10^{11} detectors to build a 1000 logical qubit quantum computer [17]. Further, such a quantum computer would require loss rates per component below approximately 10^{-3} and error rates below approximately 10^{-5} per component. These numbers result from intensive numerical simulations; we present the simulation data in a series of graphs, and we discuss these results both in the main paper and in Appendix F.

II. PROTOCOL

Our protocol is based on a three-dimensional (3D) cluster state [19–21] [Fig. 1(a)]. With the cluster-state approach, all entanglement required by the quantum computation is generated ahead of the computation itself, which then proceeds purely through measurements. The 3D cluster states enable measurement-based implementations of topological quantum computing using the surface code [22–24], providing high thresholds for both qubit loss and computational errors [25]. Without qubit loss, 3D cluster states tolerate phase errors with a rate up to 3% on each qubit; conversely, without computational errors, they tolerate up to 24.9% qubit losses [25]; and with both computational errors and qubit loss, the threshold of errors decreases approximately linearly with the loss rate. Thus, cluster states are particularly well suited to LOQC, as they can be efficiently prepared with linear optics: There is no fundamental difficulty caused by a high rate of entanglement failure during the creation of the cluster state, provided that once it is created it surpasses these thresholds [6,7,26,27].

The 3D cluster state is a graph state on a 3D lattice, which can be understood by supposing that each vertex on the graph denotes a qubit initialized in the state $|+\rangle$, and each edge denotes a controlled-phase gate entangling the two linked qubits. In the particular lattice we require, each qubit is connected to four neighboring qubits; see Fig. 1(a). In order to create such a cluster state, our protocol requires one complete building-block state to be prepared for each eventual qubit in the cluster. Importantly, these building-block states contain sufficient redundant encoding that entanglement links between building blocks can be generated with a probability above that necessary for fault-tolerant computing. If suitable building blocks can be constructed, a fault-tolerant cluster state of arbitrary size can then be generated deterministically. We can therefore focus on the optimal approach to constructing these building-block states (Fig. 2), without concerning ourselves with the precise details of the 3D cluster state that will ultimately be generated. However, we note that it is only necessary to generate one 2D layer of the cluster state “at a time,” since the process of synthesizing a layer can simultaneously entangle it with the previous layer while leaving “dangling bonds” to fuse with the succeeding layer—this is shown in Fig. 1, where the darker layer is

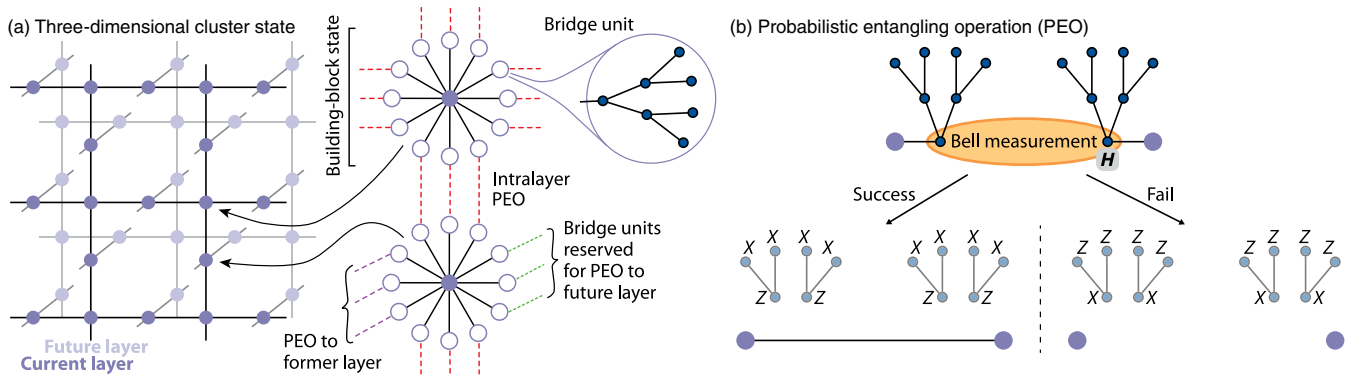


FIG. 1. Protocol for linear optical quantum computing using 3D cluster states. This figure uses the graph-state notation: Each diagram represents the multiqubit state that would result if one could prepare one qubit in state $|+\rangle$ for each node (filled circle) and then perform a controlled-phase gate for each edge (i.e., connecting line). However, since LOQC does not permit deterministic entanglement, our states must actually be created by a more lengthy process. (a) The full 3D cluster state is constructed in a near-deterministic step once we have created sufficiently complex building-block states. Each building-block state contains a photonic core qubit (solid circle) and several bridge units (empty circles), which are themselves tree graph states of photonic qubits. The core qubits of the building-block states will form the qubits in the 3D cluster state. Entanglement between these cores (i.e., edges in the eventual cluster state) is established by attempting PEOs between bridge units. (b) PEOs on two bridge units are used to make entanglement links between core qubits. The PEO is composed of a Hadamard gate on one of the bridge’s root qubits followed by a Bell measurement (BM) between the two root qubits. Other photonic qubits in the bridge qubit are then measured in bases according to the outcome of the BM. If the BM succeeds, an entanglement link is generated. Regardless of whether it succeeds or fails, the remains of the two bridge units must be removed from their building blocks.

the “current” layer and the lighter-colored qubits are part of the next layer to be formed. Therefore, the building-block factories can be reused to generate each layer of the 3D cluster.

As in other cluster-state-generation protocols [6,28–33], our building-block state is also a graph state. We employ the star graph as the basic structure of our building blocks; see Fig. 1(a). This state is composed of one core qubit and several bridge units. While the core qubit is a single photonic qubit, each bridge unit is physically encoded in a tree-structure graph state of several photonic qubits. In order to implement a PEO between two different building-block states, a Bell measurement is carried out between the root qubits at the base of each bridge unit. The treelike structure within the bridge units enables two key properties: In the case of a successful PEO between bridge units on two different building-block states, the core qubits on each building block become entangled, and the remaining qubits within each unit can be trimmed away; see the left part of Fig. 1(b). Moreover, on failure of the entangling operation, then the measurements on the remaining qubits within the bridge units allow us to

identify any necessary phase correction to the core qubit, preventing its corruption (with high probability). The right part of Fig. 1(b) summarizes this protocol; see Appendix A for further details. This method for recovering from PEO failure via measurements on ancillary qubits follows the approach introduced in Ref. [8].

With these two properties, it is possible to make multiple attempts to form links between core qubits while still ensuring that errors remain below the fault-tolerant threshold. Two building blocks can therefore be successfully connected with a high rate, provided that there are enough bridge units.

Since each core qubit must be linked to four other core qubits, the number of bridge units on the building blocks is chosen to be a multiple of four, with a quarter of the bridge units allocated for each connection. To establish an entanglement link, PEOs are performed on corresponding bridge units in parallel. If there is one successful PEO and the removal measurements are also successful, then the connection is successfully established. If there is more than one successful PEO, we keep only one link. The connection fails if there are




FIG. 2. Building-block states are constructed in a series of stages from initial three-qubit GHZ-state entanglement resources. Further details are given in Appendix B.

no successful PEOs or one of the removal operations fails. Connection failures are dealt with by treating core qubits with a failed connection as missing qubits [33], which can be tolerated by MBQC on the 3D cluster state.

The finally prepared state of core qubits is equivalent to the 3D cluster state up to some feedback single-qubit gates, depending on outcomes of single-qubit measurements and Bell measurements for preparing the state. In the MBQC algorithm on the 3D cluster state, (core) qubits are measured in four bases, which are σ^x , σ^z , and $(\sigma^x \pm \sigma^y)/\sqrt{2}$ (only for magic state injection). In our protocol, the feedback gate on a core qubit is always either the identity $\mathbb{1}$ or the phase gate σ^z . Therefore, remarkably, core qubits can be measured before the cluster state is prepared. It is beneficial to do so, in order to reduce the effect of photon loss: Core qubits are measured as early as possible, and obviously, if a measurement fails, then that particular building block is abandoned at its initial stage. Once the full feedback is known, we may update (flip) the recorded outcomes of any cores measured in σ^x or $(\sigma^x \pm \sigma^y)/\sqrt{2}$.

III. GENERATION OF BUILDING BLOCKS

Each building-block state must be generated from an initial resource of unentangled single photons. In our scheme, these single photons are first entangled into three-qubit Greenberger-Horne-Zeilinger (GHZ) states. These entanglement primitives can then be sequentially combined into larger units using Bell measurements (Fig. 2). This process is known to be efficient for loss rates of less than $1/3$ [34], since at each stage it is then possible to increase (up to double) the size of the resulting entangled states. Further details on the building-block construction process are given in Appendix B.

Regardless of the specific architecture of the building block to be generated, this process requires two primary circuit elements. The first element, a GHZ-state factory, produces GHZ states probabilistically from single-photon inputs. The second element probabilistically joins two independent graph states into a larger graph state. Along with these two processing elements, it is also necessary to construct switching networks and delay lines in order to route photons between processing states. All of these operations must be realized using only linear optical elements, e.g., single-photon sources, beam splitters, switches, delay lines, and photon detectors [Fig. 3(a)].

In our scheme, we use the same GHZ-state factory as proposed in Ref. [34]. This circuit requires six single-photon inputs and, in the lossless case, successfully generates GHZ states with probability $1/32$ [Fig. 3(b)]. Our fusion elements use Bell measurements as PEOs for joining intermediate states. These Bell measurements consume one photon from each input state [7]. A tempting alternative is to employ the type-I fusion gate, which consumes only one photon and can also connect two graph states [7]. However, a type-I fusion gate may convert

photon loss into computational errors (see Appendix B), which should be avoided, as overcoming errors is usually harder than overcoming photon loss. Therefore, we only use Bell measurements in our protocol [34]. The circuits we use for the Bell measurement are also shown in Fig. 3(c). Without any ancillary resources, a linear optical Bell measurement (often termed type-II fusion) can succeed with 50% probability in the lossless case. However, with the help of four ancillary single photons, the success probability of a Bell measurement can be boosted to 75% [35]. The same success probability can also be achieved with a Bell state as the ancillary resource [36]. With a resource state of more entangled photons, the success probability can be further boosted [35,36].

As neither GHZ-state generation nor Bell measurements can succeed deterministically, we select successful outcomes from these operations to feed into the next stage of construction. This requires a rapidly reconfigurable switchyard consisting of a network of switches. For example, to select N successful copies of the three-qubit GHZ state from M attempts in parallel, we need six M -input-to- N -output switchyards, one for each output mode of the GHZ-state-generation circuit. Before photons enter switchyards, delay lines are necessary to allow time for the switchyard to be reconfigured.

We consider two different approaches to this switching requirement. In the ideal case, this switchyard would consist of a single reconfigurable switch with multiple inputs and outputs [37], in which there is no extra cost in terms of losses or errors as N or M increases. This may prove impossible to achieve, and so we also consider the opposite limit, in which a switchyard is built out of a network of two-to-two switches. Such an M -to- N switchyard can be realized with M one-to- N switchyards and N M -to-one switchyards [Fig. 3(d)]. Each one-to- N and M -to-one switchyard is respectively composed of approximately N and approximately M two-to-two switches, as also shown in Fig. 3(d). With such a network of simple switches, each photon must go through approximately $\log_2(MN)$ switches. (Our numerical simulations account for the number and their topology exactly.) To minimize photon loss, switchyards with multiple inputs M but a single output are favorable. However, resources are not used efficiently in this case, and many successful PEO outputs will be discarded. To increase the efficiency, it is preferable to use switchyards with more output modes. In our numerical simulations, we have considered different configurations of the switch network to obtain the optimal threshold of a computer built with two-to-two switches.

IV. PHOTON LOSS AND COMPUTATIONAL ERRORS

The main source of noise in LOQC is photon loss, which may be induced by any component on the optical path of the qubit. We assume that a loss occurs at single-photon sources, beam splitters, delay lines (for the time period

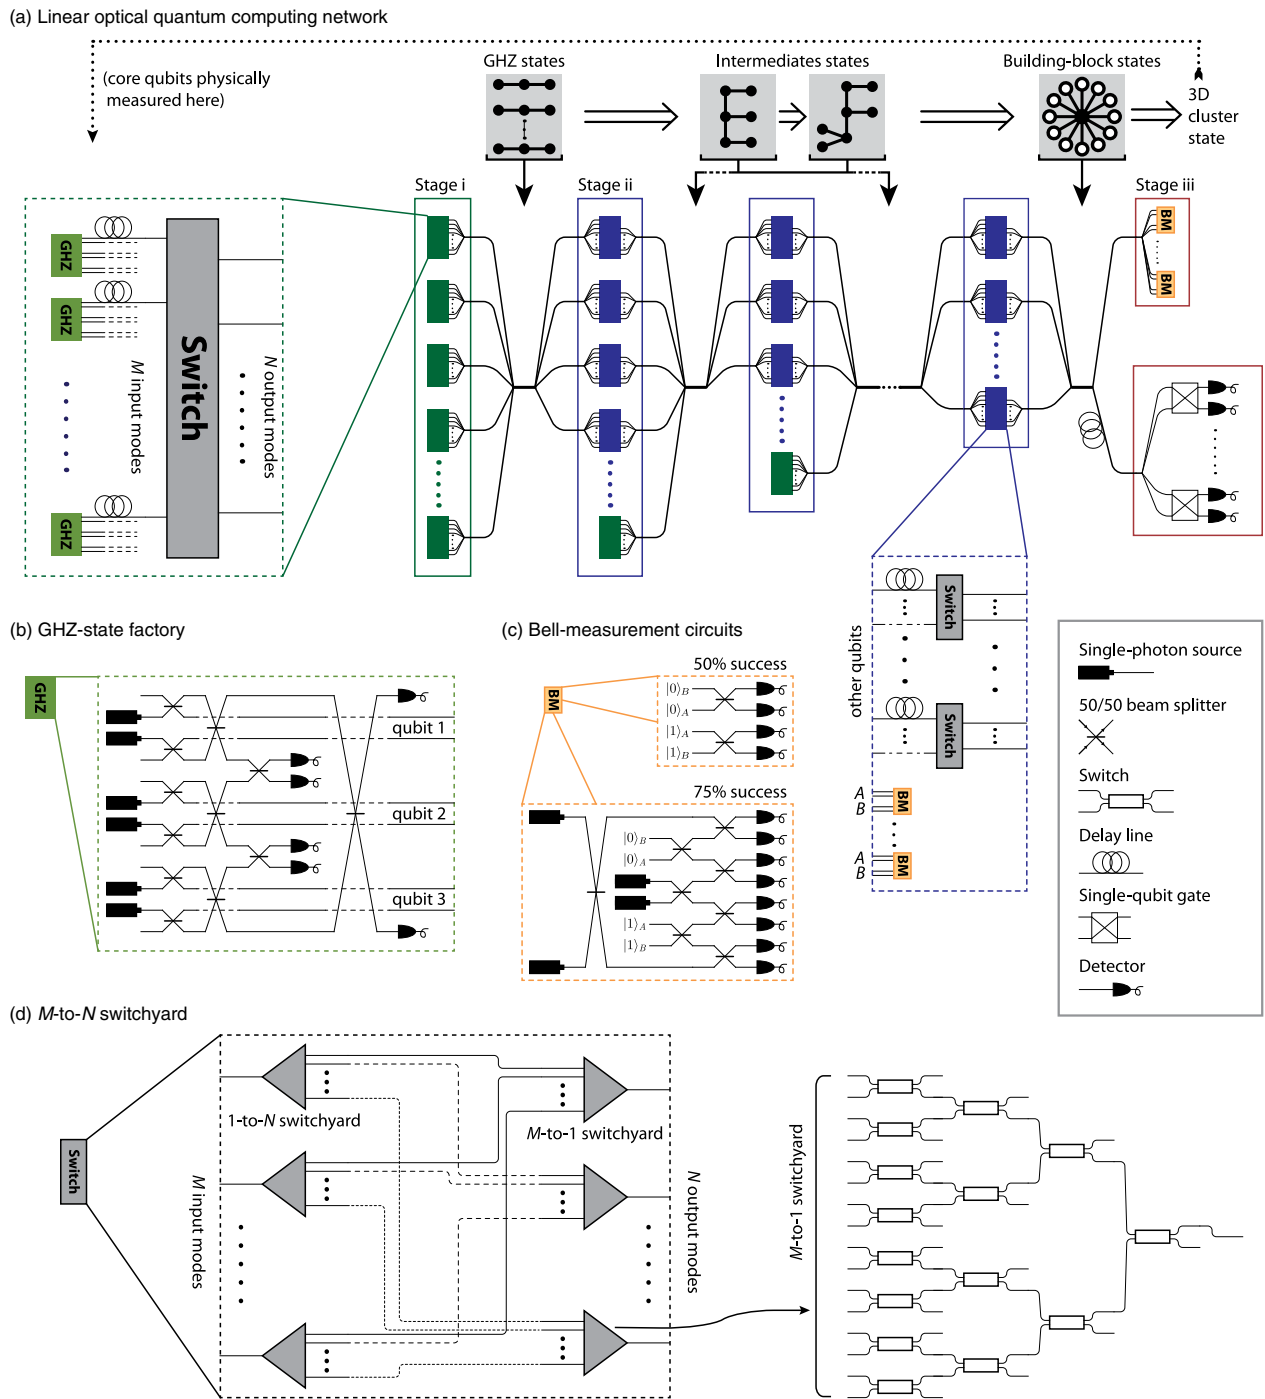


FIG. 3. (a) Circuit for realizing linear optical quantum computing using a three-dimensional cluster state. The circuit includes stages for (i) generating initial three-qubit GHZ states, (ii) synthesizing these states into building-block states, and finally (iii) constructing the 3D cluster state. Qubits on the cluster state are physically measured right after they are generated by a GHZ-state factory. (b) GHZ-state factories probabilistically generate three-qubit GHZ states from six single photons. Successful generation of the GHZ state is heralded by specific three-photon detection events at the detectors. To select at most N successful copies of the GHZ state from M attempts, we need six M -to- N switchyards. Delay lines are necessary before photons enter switchyards to allow for feed forward. (c) Bell-measurement circuits synthesize building-block states. The two different circuits depicted succeed with probabilities 50% and 75%, respectively. For the 75%-success circuit, four ancillary single photons are used. At each synthesis stage, many copies of input states are prepared, and measurements are performed on these states in parallel. Successful output states are selected with switch networks. (d) A M -to- N switchyard can be realized with one-to- N switchyards and M -to-one switchyards. For each one-to- N switchyard, every output mode is connected to an input mode of a different M -to-one switchyard. A M -to-one switchyard is composed of approximately M two-to-two switches. A one-to- N switchyard is similar.

required for one PEO stage), switches, and detectors with the rates p_e , p_b , p_d , p_s , and p_m , respectively.

In addition to photon losses, we also have to consider computational errors. Because measurements are eventually attempted on all photonic qubits in the protocol, computational errors are induced by any source of noise that can affect these measurement outcomes. For example, any asymmetry, e.g., phase difference or biased transmission, between two modes of a qubit may result in computational errors. All computational errors are equivalent to Pauli errors (see Appendix C). In this paper, we assume that depolarizing errors may happen at beam splitters, delay lines, and switches with the rates ϵ_b , ϵ_d , and ϵ_s , respectively. Imperfect mode overlap between different photon sources will lead to imperfect quantum interference at beam splitters and therefore also to Pauli errors. For simplicity, in our model, we incorporate this form of error into ϵ_b .

Other types of noise are also tolerable in our protocol. For example, a photon source may emit two photons rather than a single photon into the circuit. Similar errors can be induced by switching errors, in which a photon enters the wrong mode, and from dark counts of detectors. To first order, all of these errors will be caught during measurement, since if these

extra photons survive in the optical path, we will measure more than the expected number of photons. This requires detectors with at least the capacity to resolve one- and two-photon signals [16]. In the event of such a measurement, we can simply treat the qubit as missing, an error which can be overcome in the same way as true photon loss. However, if a two-photon error is followed by a photon-loss event, only one photon will be detected, and the measurement on such a qubit may give a wrong outcome. These computational errors are also equivalent to Pauli errors and can be corrected with our protocol. Although these errors can be corrected, we consider regimes in which they will occur at a rate much lower than the first-order error terms, and so we do not explicitly include them in our threshold study.

V. THRESHOLDS

In this approach to LOQC, the fault-tolerance threshold depends on the complexity of each building-block state. With more resources, one can prepare bigger building blocks, and thus a higher level of photon loss is tolerable. In Fig. 4, fault-tolerant thresholds are obtained numerically (see Appendix D for details; our computer code is available

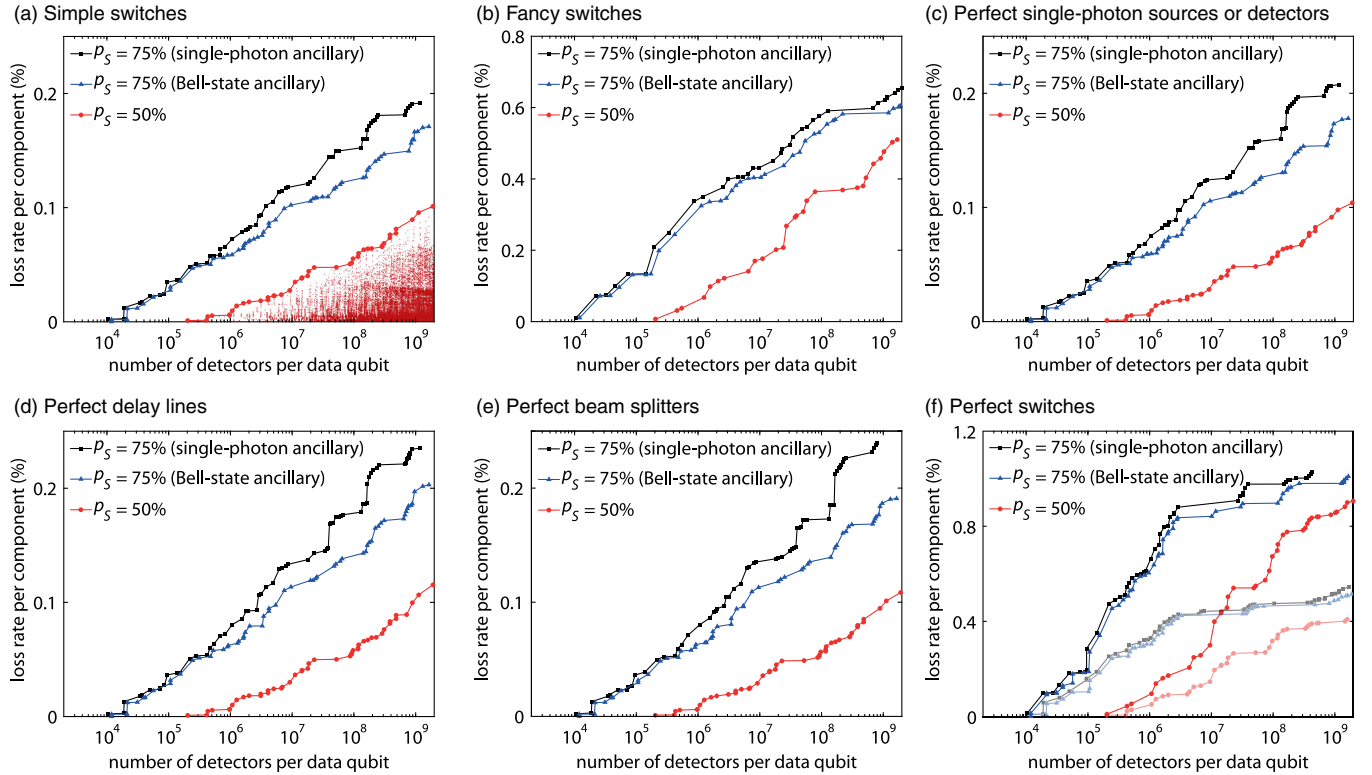


FIG. 4. Thresholds on the loss rate per component p as a function of the number of detectors per data qubit in the case of no computational errors; p_S denotes the success rate of Bell measurements. We again note that each logical qubit will be encoded in a surface code consisting of ≥ 1000 data qubits. (a) Model in which each switchyard is composed of cascaded two-to-two switches. All components have the same loss rate, i.e., $p_e = p_b = p_d = p_s = p_m = p$. (b) Each switchyard is a single switch with multiple inputs and outputs. All components have the same loss rate p . (c)–(f) In each plot, one form of the component is assumed to be perfect, while all other components have the same loss rate p . As can be seen, the most dramatic improvement is seen when the switches are assumed to be perfect. Transparent curves in subfigure (f) correspond to the case that the loss rate of switches is 10% of other components.

openly [38]). In order to provide some physical intuition for the size of such linear optical quantum computers, we choose to specify the total number of detectors needed as our metric of the resources required. It can be seen that this approximately corresponds to twice the number of single photons needed and therefore twice the number of single-photon sources. It is likely that the resource burden of the other elements, e.g., beam splitters, delay lines, and switches, will be of similar magnitudes.

Note that each curve in Fig. 4(a) is actually an envelope representing the best of a very large number of protocols that were tested. Each small dot within the red curve in the upper left figure represents the outcome of one such simulation; these dots are omitted from other curves for clarity.

It is vital to appreciate that the number of detectors shown in the figure is for a single building-block state rather than the entire computer. A building-block state corresponds to only one qubit on the cluster state, i.e., one data qubit of the surface code, which could correspond to just one ion in an ion-trap quantum computer or one superconducting qubit in a superconducting quantum computer. It is anticipated that a fault-tolerant quantum computer will need at least approximately 10^6 data qubits in order to be able to compete with state-of-the-art classical computers [18,39]. We therefore do not consider building-block states with a resource requirement of greater than 2×10^9 detectors, since at that point one finds the entire computer requires thousands of trillions of components.

In Fig. 4(a), we consider the case in which all components of the computer have the same photon-loss rate. Depending on the choice of Bell-measurement protocol, the threshold loss rate per component varies from approximately 0.1% to approximately 0.2%. In this subfigure, we consider the worst-case approach, in which each switchyard is built out of a cascade of two-to-two switches. For comparison, in Fig. 4(b), we consider a more sophisticated computer, in which each switchyard is a multiple-input–multiple-output switch with the same loss rate as the other components. In this case, the threshold is approximately 3–5 times higher than that of a simple-switch computer. We note that these thresholds approach the 1% limiting loss rate that has been discussed in the context of a computing paradigm where gates are essentially deterministic but suffer a small probability of qubit loss [40]. (We have achieved this at the cost of the additional resource overhead, of course.)

In order to further explore which components have the most significant impact on the fault-tolerance threshold, in Figs. 4(c)–(f), we modify the model in Fig. 4(a), assuming in each that one of the circuit components is lossless. These simulations confirm that it is the switching networks which most strongly impact the loss tolerance of the quantum computer. This suggests that alternative approaches in which intermediate cluster states are extensively recycled (similar to the recycling discussed in Ref. [32]) will suffer from the associated increase in complexity of the switching networks.




FIG. 5. Thresholds on the loss rate per component p for given the error rate per component ϵ and the number of detectors per data qubit. Switchyards are composed of cascaded two-to-two switches, and single-photon ancilla-assisted Bell measurements with 75%-success probability are used. All components have the same loss rate, i.e., $p_e = p_b = p_d = p_s = p_m = p$, and all error rates are equal, i.e., $\epsilon_b = \epsilon_d = \epsilon_s = \epsilon$. See Fig. 11 for the full data with more details.

The threshold changes dramatically with increased success probability Bell measurements. With a higher success probability, the size of building-block states (number of bridge units) can be smaller; hence, both the level of noise and the resource cost can be lower. Bell measurements with 75% success probability perform significantly better than those with 50% success probability. Further, the single-photon ancilla-assisted Bell measurement is slightly better than the Bell-state-assisted Bell measurement. These boosted Bell measurements do, however, require photon detectors with additional photon-number resolution. In the case of the 50%-success Bell measurement, we need detectors that can distinguish photon numbers 0,1,2, while for 75%-success Bell measurements, we need detectors that can distinguish photon numbers 0,1,2,3. With more complex ancillary states, the success probability can be further boosted. However, in this case, more resources are required for preparing these ancillary states, which can counteract the benefits of the higher success probabilities (see Appendix E).

The presence of computational errors reduces the threshold loss rate. A general study of the threshold for both loss rates and error rates is shown in Fig. 5, in which we only consider quantum computers built with two-to-two switches and single-photon ancilla-assisted Bell measurements with 75% success probability. More data for other Bell-measurement circuits and fancy switches can be found in Fig. 10. The threshold error rate per component is on the order of 10^{-5} .

VI. DISCUSSION

We have proposed a comprehensive protocol for LOQC with 3D cluster states, in which we consider the full network

of linear optical devices necessary to realize a quantum computer. We find thresholds for loss and error rates of approximately 10^{-3} and approximately 10^{-5} per component, respectively. This per-component performance is beyond the current state of the art in photonics [41–47]. Furthermore, we find that such a quantum computer would require on the order of 10^{11} detectors and similar numbers of other components including deterministic and indistinguishable single-photon sources. We note that these component counts are 5 or more orders of magnitude beyond those required for systems with matter qubits and deterministic gates. In such systems, the data qubits of the surface code correspond to the lowest level of the physical machine [18], but our LOQC machine requires a further tier of complexity beneath, where a large building-block structure must be created for each eventual data qubit. The resource costs are discussed in more detail in Appendix F.

We wish to emphasize that these stringent thresholds should be taken as a challenge to the community, aiming to stimulate further discussion and innovation in LOQC. From an experimental perspective, we have tried to determine which components will prove most critical in the development of an optical quantum computer. We found that, for our scheme, it is the performance of the optical switches that have by far the most impact on the threshold loss and error rates, while other components contribute more equally. We hope that this will help guide the priorities of future experimental work aimed toward realizing LOQC. From a theoretical perspective, we hope that our work will stimulate others to improve on our thresholds by exploring alternative schemes.

ACKNOWLEDGMENTS

This work was supported by the EPSRC platform grant “Molecular Quantum Devices” (EP/J015067/1), and the EPSRC National Quantum Technology Hub in Networked Quantum Information Processing. We would like to thank Joshua Nunn, Animesh Datta, Benjamin Metcalf, Jacques Carolan, Dan Browne, and Mercedes Gimeno-Segovia for helpful discussions. The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work [48].

APPENDIX A: LOSS AND ERROR-TOLERANT BUILDING BLOCKS

As discussed in the main text, our protocol is based on a 3D cluster state [Fig. 1(a)]. To create the cluster state, one intermediate building-block state must be prepared for each qubit in the cluster.

We employ the star graph as the basic structure of our building blocks [Fig. 1(a)]. This state is composed of one core photonic qubit and several bridge units. To tolerate photon loss and failures of PEOs, each bridge unit is

encoded as a tree-structure graph state of several photonic qubits with a root qubit connected to the core qubit (Fig. 6). The PEO for connecting two core qubits includes a Hadamard gate on one root qubit and a Bell measurement on two root qubits [Fig. 1(b)]. Because the Bell measurement can only succeed probabilistically in LOQC, the overall operation is probabilistic. If the PEO is successful (fails), qubits on first-generation branches, which are directly connected to the root, are measured in the σ^z (σ^x) basis, qubits on second-generation branches are measured in the σ^x (σ^z) basis, and so on. This measurement pattern removes redundant branches from two connected building blocks if the PEO is successful and removes entire bridge units from two independent building blocks if the PEO is failed. This removal operation is not always successful due to photon loss. When the tree graph state is large enough, the removal operation can succeed with an arbitrarily high probability if the photon-loss rate is lower than 50% [8].

APPENDIX B: CONSTRUCTING BUILDING BLOCKS

Building-block states are generated by fusing three-qubit GHZ states (see Fig. 6) with Bell measurements. There are four types of graph states occurring in the generation process, which are three-qubit GHZ states, rake-structure graph states, tree-structure graph states, and rake-tree states. In the first step, rake states are prepared from GHZ states. These states, along with further GHZ states, are the basic ingredients of rake-tree states. Using these ingredients, rake-tree states are generated and enlarged with Bell measurements. When the tree of a rake-tree state is large enough, it can be converted into a building-block state by removing the rake. As an example, the construction process for a building-block state with branching numbers (8,2,2) is shown in Fig. 6.

A rake with r branches can be prepared with $2(r-1)$ GHZ states (assuming all Bell measurements are successful) in $\text{ceil}[\log_2(r-1)] + 1$ steps. In the first step, each pair of GHZ states is fused by a controlled-phase (CP) operation [see Fig. 7(a)] to obtain a four-qubit linear cluster state, which is also a rake with two branches. Two rakes can be combined into a bigger rake by a parity-projection (PP) operation [see Fig. 7(b)]: If two input rakes, respectively, have r_1 and r_2 branches, the output rake has $r_1 + r_2 - 2$ branches. Since in each step the number of branches can be nearly doubled, $r-1$ two-branch rakes can be combined into the r -branch rake in $\text{ceil}[\log_2(r-1)]$ steps.

A rake itself is a rake-tree graph state in which the tree is one level but the branching number is 0. A GHZ state itself is also a rake-tree graph state in which the rake has only one branch and the tree is one level with the branching number 1. From these two kinds of graph states, rake-tree states can be generated and enlarged with two basic processes: increasing the branching number of the tree and increasing

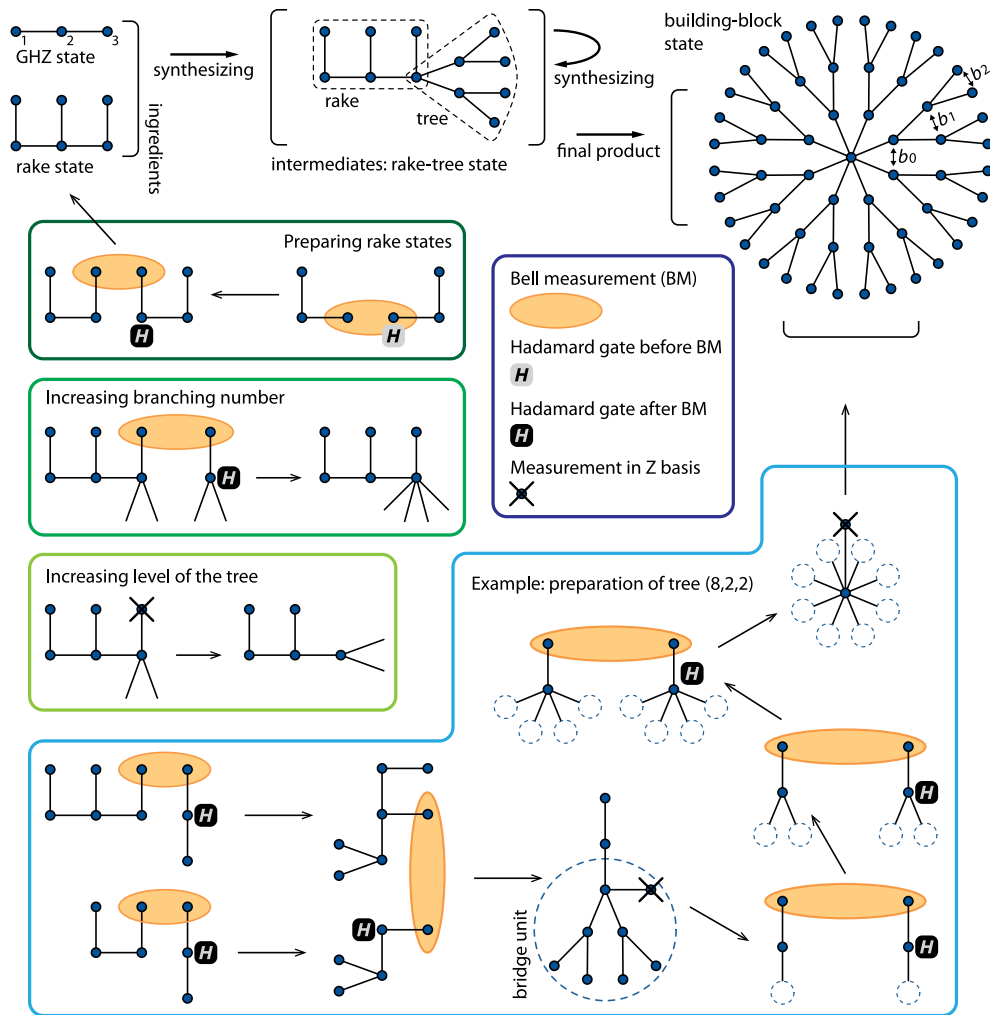


FIG. 6. Scheme for generating building-block states with three-qubit GHZ states. The building-block state is a tree graph state of photonic qubits. Such a tree can be characterized by branching numbers (b_0, b_1, b_2, \dots) , where b_0 is the number of bridge units and b_i is the number of i th-generation branches of a bridge unit. In our protocol, b_0 is always a multiple of 4. Graph states occurring in the generation process include three-qubit GHZ states, rake states, tree states, and rake-tree states. On the top of the figure, the example rake state has three branches, the rake-tree state is composed of a rake with three branches and a two-level tree with branching numbers (2,2), and the building-block state is a three-level tree state with branching numbers (8,2,2).

the level of the tree (see Fig. 6). The branching number is increased by fusing two rake-tree graph states with a PP operation, where one of the rakes always has only one rake branch. The level of the tree is increased by removing, i.e., measuring in the σ^z basis, one branch of the rake. We would like to remark that, when the level number is increased from 1 to 2, the branch of the rake (which is supposed to be measured) can be kept as a branch of the tree.

In our protocol of generating photonic tree-structure graph states, the rake structure allows us to increase the level of the tree with a single-qubit measurement (which is physically performed at the GHZ-state-generation stage due to the same reason of measuring core qubits as early as possible). This process is efficient; the largest rake state can be prepared in $\text{ceil}[\log_2(R-1)] + 1$ steps, where R is the level of the final tree. As a comparison, the previous

protocol reported in Ref. [34] requires probabilistic Bell measurements for increasing the level of the tree. Therefore, the number of construction stages is reduced in our protocol for high-level trees. Minimizing the number of construction stages can reduce noise induced by delay lines and switchyards and also reduce the resource cost. If the success probability of a Bell measurement is p_S , for each successful output state, roughly speaking, $1/p_S$ input states need to be prepared. For a GHZ state going through n Bell-measurement stages, $1/p_S^n$ copies will be required to ensure that each stage is successful. Reducing n is therefore critical to reducing the resource costs.

As we have discussed in the main text, we choose Bell measurement rather than type-I fusion gate as the operation of entangling photons because a type-I fusion gate may convert photon loss into computational errors. For a type-I




FIG. 7. Fusion operations on graph states using Bell measurements in the basis $|\text{BS}_{\mu,\nu}\rangle = Z_{a'}^\mu X_{b'}^\nu (|00\rangle_{a',b'} + |11\rangle_{a',b'})/\sqrt{2}$, where $\mu, \nu = 0, 1$. (a) A CP operation includes a Hadamard gate on one of two measured qubits and then the Bell measurement. Depending on the measurement outcome, an operation $Z_a^\nu Z_b^\mu$ needs to be performed. (b) A PP operation includes the Bell measurement and a Hadamard gate on the nearest neighboring qubit of one of two measured qubits. Depending on the measurement outcome, an operation $Z_a^\nu Z_b^\mu \prod_{i \in N'(b)} Z_i^\mu$ need to be performed. Here, $N'(c)$ denotes neighboring qubits of the qubit c except the qubit c' .

fusion gate, if there is no photon loss, two input photonic qubits are projected into the subspace of HH and VV (for the polarization encoding) when only one photon is detected or the state HV (VH) if zero photons (two photons) are detected. With photon loss, the input qubits may be in the state VH rather than the subspace of HH and VV if only one photon is detected and the other is missing. Therefore, photon loss may result in computational errors in a type-I fusion gate.

APPENDIX C: THE MODEL OF COMPUTATIONAL ERRORS

In our protocol of LOQC, operations performed on three-qubit GHZ states include Hadamard gates, Bell measurements, and single-qubit measurements in the σ^x and σ^z bases. Single-qubit measurements in bases $(\sigma^x \pm \sigma^y)/\sqrt{2}$ (for magic state injection) do not affect the fault-tolerance threshold. Therefore, any computational error is equivalent to a bit-flip error $[\sigma^x]$, a phase-flip error $[\sigma^z]$, or a combined error $[\sigma^y]$ on a single qubit or a combination of these three types of errors on a group of qubits. Because GHZ states are prepared separately, there is no correlation between them right after they are generated. Correlations between GHZ states may occur when two qubits of different GHZ states are measured by a Bell measurement. However, for the Bell measurement on qubits A and B , all errors are equivalent to three types of Pauli errors $[\sigma_A^x]$, $[\sigma_A^y]$, and $[\sigma_A^z]$ (or, equivalently, $[\sigma_B^x]$, $[\sigma_B^y]$, and $[\sigma_B^z]$, corresponding to three possible incorrect outcomes), which are all single-qubit errors. Therefore, all computational errors are equivalent to Pauli errors within GHZ states. For the GHZ state shown in Fig. 6, these Pauli errors could be $[\sigma_1^y]$, $[\sigma_1^z]$, $[\sigma_2^x]$, $[\sigma_2^y]$, $[\sigma_2^z]$, $[\sigma_3^y]$, and $[\sigma_3^z]$, and all other errors are equivalent to these seven types of errors.

Correlations may also occur in switchyards. To deal with these correlations, we can ensure that states from the same switchyard are utilized in different cluster-state qubits that

are separated by distances much larger than the dimension of logical qubits. In this way, these correlations induced by switchyards never form a correlated error that can risk a logical qubit.

The depolarizing error reads

$$\mathcal{E} = (1 - \epsilon)[1] + \frac{\epsilon}{3}([\sigma^x] + [\sigma^y] + [\sigma^z]). \quad (\text{C1})$$

Here, ϵ is the error rate, and the superoperator $[U]\rho = U\rho U^\dagger$.

APPENDIX D: NUMERICAL SIMULATIONS

The threshold of fault-tolerant quantum computing is determined by evaluating p_L and p_P , which, respectively, are the loss rate and phase-error rate of a cluster-state qubit, for the given loss rate and error rate per component. Our threshold-finding code is available openly [38]. In Ref. [25], the inferred critical threshold is almost a straight line in the (loss rate, computational error rate) parameter space. For the 3D cluster state, the phase-error-rate threshold without loss is 2.93% [19], and the loss-rate threshold without error is 24.9%. Therefore, the threshold of (p_L, p_P) is estimated as

$$\frac{p_L}{24.9\%} + \frac{p_P}{2.93\%} = 1.$$

To obtain thresholds of the loss rate per component without computational error, we have considered two-level trees and three-level trees (see Fig. 8) with branching numbers not larger than 20 as building-block states, which includes 8400 different tree structures in total. In the case that switchyards are composed by two-input–two-output switches, we have considered configurations of the switch network for which the number of outputs for all of the switchyards is the same. We simulate output numbers $N = 2^n$ with $n = 0, 1, \dots, 10$. For switchyards at the outputs of GHZ-state factories, the input number is $M = m \times N$ with $m = 32, 36, \dots, 256$. (The success rate of generating GHZ states is $1/32$ [34].) For switchyards for selecting successful Bell measurements, the input number is determined by the input number of GHZ-state switchyards, which is $\text{ceil}[M/(32p_S)]$, where $p_S = 50\%, 75\%, 87.5\%$ is the success rate of Bell measurements without photon loss. Similarly, for switchyards for selecting successfully generated Bell states, the input number is $\text{ceil}(M/4)$, where we have used the circuit for generating Bell states with the success rate $1/8$, which can be boosted to $3/16$ if a switch is introduced [7]. Therefore, we have in total considered 627 different configurations of the switch network composed by two-input–two-output switches. In the case that each switchyard is a fancy switch with arbitrarily large input and output numbers, we have assumed that the ratio of output number to input number equals the actual success rate (including the effect of photon loss) of corresponding operations.

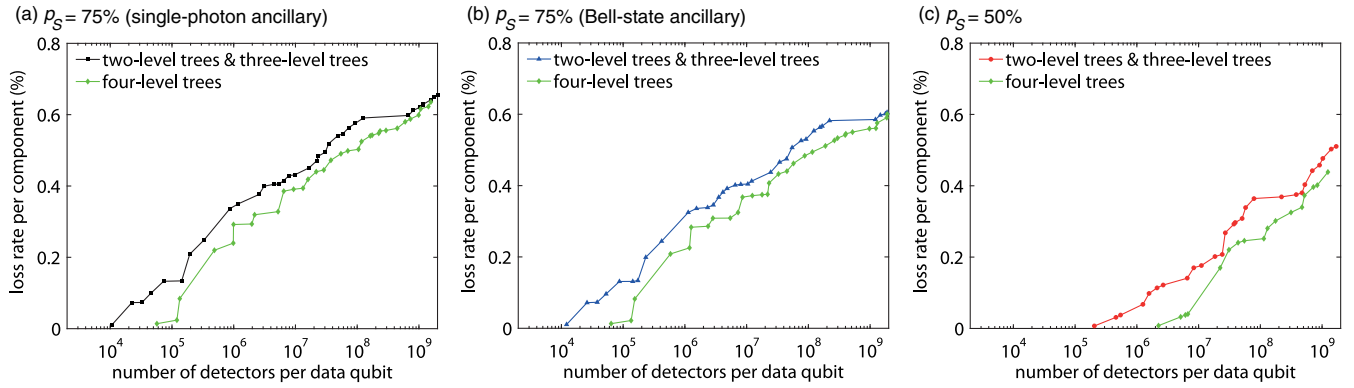


FIG. 8. Thresholds of the loss rate per component without computational errors for a computer built with fancy switches. Panels (a)–(c) correspond to different values of p_s , as labelled in the figure. All components have the same loss rate. In addition to two-level and three-level trees, we have also considered four-level trees. One can find that four-level trees tolerate less photon loss per component within the regime in which we are interested.

For each curve in Figs. 4(a) and 4(c)–(f), thresholds of the loss rate per component are evaluated for 8400×627 different protocols. Each protocol includes the building-block structure and the configuration of the switch network. Each curve is obtained as the envelope of these thresholds. For each curve in Fig. 4(b), thresholds of the loss rate per component are evaluated for 8400 different protocols, which are only determined by building-block structures.

To obtain thresholds of the loss rate per component with computational errors, we have selected about 500 protocols from protocols that require not more than 2×10^9 detectors for each case. These selected protocols are all close to the envelope, i.e., have the best performance of tolerating photon loss. Specifically, we have drawn a straight line connecting the highest point (corresponding to the protocol that tolerates the highest loss rate per component) and the lowest point (corresponding to the protocol with the smallest number of detectors) on the envelope. This line is then shifted downward until there are about 500 protocols whose thresholds of the loss rate per component are above it. Thresholds of the loss rate per component with computational errors in Fig. 5

are obtained from these selected protocols. Computational errors are evaluated using Monte Carlo methods. In each protocol, for each value of the loss rate and the error rate, the phase-error rate on a cluster-state qubit is obtained with 100 000 samples.

APPENDIX E: BELL MEASUREMENTS WITH ENTANGLED ANCILLARY STATES

In addition to the Bell measurement assisted by a Bell state (see Fig. 4), which has the success probability 75%, we also have considered the Bell measurement assisted by a four-qubit GHZ state (see Fig. 9), which has the success probability 87.5% [36]. The four-qubit GHZ state is prepared with two three-qubit GHZ states generated with the circuit in Fig. 3(b). By using a PP operation, in which the Bell measurement is assisted by a Bell state, i.e., the success probability is 75%, two three-qubit GHZ states can be fused into a four-qubit GHZ state. We find that further boosting the success probability of Bell measurements with more entangled ancillary photons is not helpful.

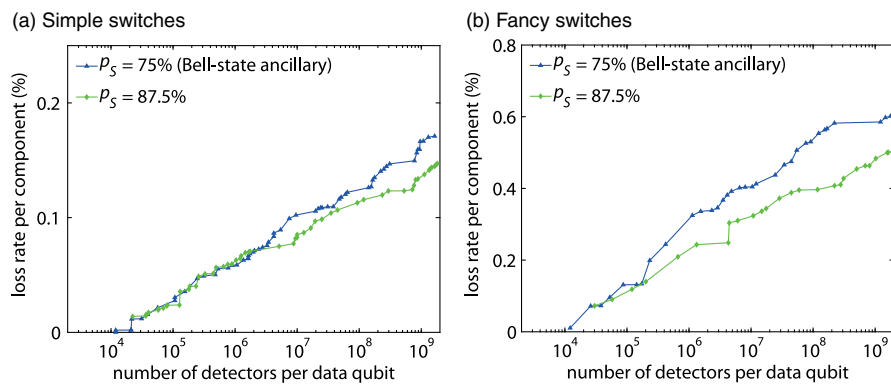


FIG. 9. Thresholds of the loss rate per component without computational errors using Bell measurements with the success probability 75% (assisted by a Bell state) and the success probability 87.5% (assisted by a four-qubit GHZ state). Panel (a) corresponds to the use of simple switches while panel (b) corresponds to fancy switches, i.e. switches with arbitrarily large numbers of inputs and outputs. All components have the same loss rate.

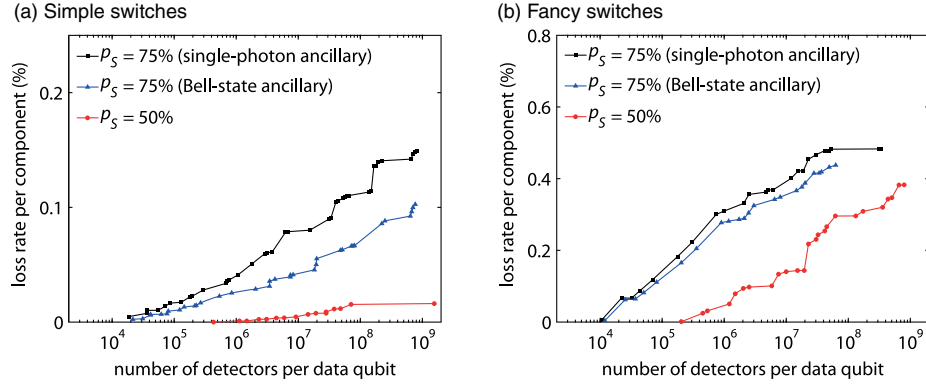


FIG. 10. Thresholds of the loss rate per component p as a function of the number of detectors per data qubit in the case with computational errors. Panel (a) corresponds to the use of simple switches while panel (b) corresponds to fancy switches. All components have the same loss rate, i.e., $p_e = p_b = p_d = p_s = p_m = p$, and the rate of computational errors per component is $\epsilon_b = \epsilon_d = \epsilon_s = 3 \times 10^{-6}$.

APPENDIX F: FURTHER DISCUSSION OF RESOURCE COSTS

Suppose that we have in mind a specific task (e.g., code breaking, optimization, etc.) together with a quantum algorithm that can perform the task. From that, one can find the required number of ideal, “logical” qubits that will suffice to perform the algorithm, as well as the number of operations that will need to be executed (and, in particular, the number of “non-Clifford” operations, which require additional resources to execute [49]). Moreover, we could then specify a logical error rate that is low enough to make it unlikely that any error occurs in the course of the calculation.

Each logical qubit will correspond to a large number (at least hundreds) of more basic qubits, broadly divided between “data qubits,” which collectively represent the encoded logical qubit, and helper qubits for measuring stabilizers. Let us refer to these as code-level qubits. As explained in Appendix F 1, the number of code-level qubits depends on the error rate ϵ that they suffer but otherwise varies only moderately between different architectures and realizations.

For some approaches to QIP, this is the “end of the story”: The code level directly corresponds to the physical hardware, and therefore the error level ϵ is the level of physical error. However, in our LOQC machine, these code-level qubits are realized as a photonic cluster state that must itself be created through an ongoing process with significant complexity and cost. This is, of course, the main topic of the present paper and is responsible for the significant gap between the resource needs of LOQC versus monolithic matter systems. This is discussed further below in Appendix F 2.

1. Size of the surface code

If the error rate ϵ associated with the data qubits is lower than the threshold value ϵ_{th} , then errors occurring on a logical qubit can be suppressed arbitrarily by sufficiently

increasing the size of the surface code. For each round of stabilizer measurements, logical qubit errors occur with a rate that has been estimated [18] as

$$\epsilon_L \sim 0.03(\epsilon/\epsilon_{\text{th}})^{(d+1)/2}, \quad (\text{F1})$$

where d is the dimension of the surface code, and the number of qubits for encoding the logical qubit is $d^2 + (d-1)^2$. The desired ϵ_L will vary depending on the computation task but must certainly be very small, e.g., $\epsilon_L < 10^{-12}$. From the equation, it is apparent that if the physical error rate ϵ were to approach the threshold ϵ_{th} , then the device would require a very high value of d in order to achieve this target ϵ_L . Practically, one would wish to be at least an order of magnitude within threshold, i.e., $\epsilon \sim (0.1)\epsilon_{\text{th}}$ or lower, so that the logical qubit error rate decreases rapidly with the dimension of the surface code. In this case, an acceptable logical qubit error rate may be achieved using of order 1000 data qubits per logical qubit.

In addition to these data qubits, which form the surface code itself, we require additional qubits to realize stabilizer measurements; the number required will vary somewhat with the architecture in question but is roughly comparable to the number of data qubits. One finds that the total number of code-level qubits is of order $4d^2$ for a monolithic matter-qubit array and of order $6d^2$ for photonic qubits in a 3D cluster-state realization [58]. Therefore, at this level, the approaches are similar.

2. The resource gap: LOQC versus matter qubits

In a monolithic matter-based device, the approximately $4d^2$ code-level qubits that comprise one logical qubit are in fact the lowest-level physical qubits, and therefore, this is also an indicator of the number of ancillary components such as detectors, etc., that the machine requires [18]. In contrast, in our LOQC approach, each code-level qubit in the 3D cluster state is the result of successfully making and

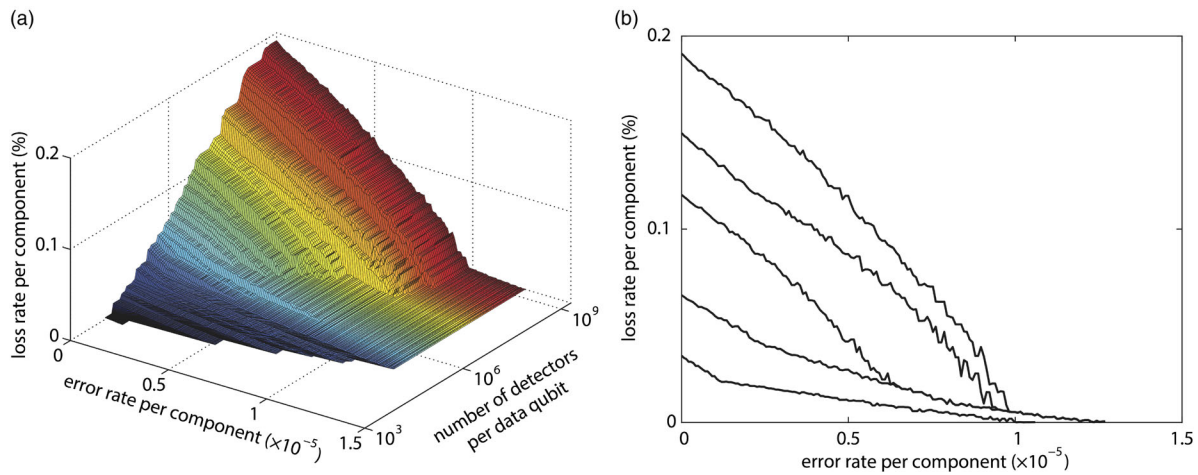


FIG. 11. (a) The full data of Fig. 5. (b) Thresholds on the loss rate per component p given the error rate per component ϵ with $10^5, 10^6, \dots, 10^9$ detectors per data qubit from bottom to top. Fluctuations are due to the uncertainty of the Monte Carlo method.

fusing a building-block object involving many photonic qubits. Our analysis in the main paper found that each building-block structure would require 10^4 detectors in the limit of small photon-loss rates, rising by orders of magnitude as loss increases (see, e.g., Fig. 4). This “detectors-per-data qubit” is a therefore a measure of the difference between the two approaches. (Of course, the detectors themselves are very different technologies—our intention here is to give a measure of the complexity.) The panels in Fig. 4 reveal that the exact cost is a complex function of the quality of the components, but it appears fair to select 10^5 as characteristic of the regions with low-but-nonzero loss.

We reiterate that our results correspond to a specific fault-tolerance strategy, albeit one that we have searched over exhaustively to find the best specific implementations. (“Rejected” variants are depicted in Fig. 4 as small dots; we estimate that in total there were over 1×10^7 such variants assessed by our numerical code.) It remains possible that a fundamentally different alternative approach could reduce this overhead significantly. Having said that, the approach we used is a synthesis of the most established ideas in the literature, i.e., the Kitaev surface code [22–24] implemented via a 3D cluster state adapted from Raussendorf and co-workers [19–21], exploiting the inherent high loss tolerance [25] and with resource-preparation strategies that adopt the recent ideas for boosted success probabilities [35,36]. Therefore, we believe it is a nontrivial task to significantly outperform this approach, and indeed it may not be possible.

Intuitively, one should expect a very substantial gap between the hardware demands of LOQC and those of matter-based approaches. It is the consequence of the mutually antagonistic relationship between the twin burdens of LOQC, i.e., nondeterministic entanglement and finite qubit-loss rates. Either issue alone could be tolerated by straightforward measures, but to mitigate both, we found

it necessary to use complex multilevel building-block objects. Although it is not possible to remove either burden completely from a pure LOQC system, it is worth noting that there are benefits to making loss very low: In Ref. [40], numerical results suggest that if the loss rate is lower than $1/5$ of the loss-rate threshold, the overhead is less than twice the zero-loss overhead. This is also reflected in the nearly flat “skirt” at the base of the “mountain” in Fig. 5.

-
- [1] E. Knill, R. Laflamme, and G. Milburn, *A Scheme for Efficient Quantum Computation with Linear Optics*, *Nature (London)* **409**, 46 (2001).
 - [2] T. C. Ralph and G. J. Pryde, *Progress in Optics* (Elsevier, New York, 2010).
 - [3] S. J. Devitt, A. G. Fowler, A. M. Stephens, A. D. Greentree, L. C. L. Hollenberg, W. J. Munro, and K. Nemoto, *Architectural Design for a Topological Cluster State Quantum Computer*, *New J. Phys.* **11**, 083032 (2009).
 - [4] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn, *Linear Optical Quantum Computing with Photonic Qubits*, *Rev. Mod. Phys.* **79**, 135 (2007).
 - [5] A. P. Lund, T. C. Ralph, and H. L. Haselgrove, *Fault-Tolerant Linear Optical Quantum Computing with Small-Amplitude Coherent States*, *Phys. Rev. Lett.* **100**, 030503 (2008).
 - [6] M. A. Nielsen, *Optical Quantum Computation Using Cluster States*, *Phys. Rev. Lett.* **93**, 040503 (2004).
 - [7] D. E. Browne and T. Rudolph, *Resource-Efficient Linear Optical Quantum Computation*, *Phys. Rev. Lett.* **95**, 010501 (2005).
 - [8] M. Varnava, D. E. Browne, and T. Rudolph, *Loss Tolerance in One-Way Quantum Computation via Counterfactual Error Correction*, *Phys. Rev. Lett.* **97**, 120501 (2006).
 - [9] J. Ghosh and A. G. Fowler, *Leakage-Resilient Approach to Fault-Tolerant Quantum Computing with Superconducting Elements*, *Phys. Rev. A* **91**, 020302(R) (2015).

- [10] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, *Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links*, *Phys. Rev. X* **4**, 041041 (2014).
- [11] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, *State Preservation by Repetitive Error Detection in a Superconducting Quantum Circuit*, *Nature (London)* **519**, 66 (2015).
- [12] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, *High-Fidelity Preparation, Gates, Memory, and Readout of a Trapped-Ion Quantum Bit*, *Phys. Rev. Lett.* **113**, 220501 (2014).
- [13] C. J. Ballance, T. P. Harty, and N. M. Linke, D. M. Lucas, *High-Fidelity Two-Qubit Quantum Logic Gates Using Trapped Calcium-43 Ions*, [arXiv:1406.5473](https://arxiv.org/abs/1406.5473).
- [14] C. M. Dawson, H. L. Haselgrove, and M. A. Nielsen, *Noise Thresholds for Optical Quantum Computers*, *Phys. Rev. Lett.* **96**, 020501 (2006).
- [15] D. A. Herrera-Martí, A. G. Fowler, D. Jennings, and T. Rudolph, *Photonic Implementation for the Topological Cluster-State Quantum Computer*, *Phys. Rev. A* **82**, 032332 (2010).
- [16] M. D. Eisaman, J. Fan, A. Migdall, and S. V. Polyakov, *Invited Review Article: Single-Photon Sources and Detectors*, *Rev. Sci. Instrum.* **82**, 071101 (2011).
- [17] This assumes that each logical qubit will be encoded in a surface code consisting of approximately 1000 physical qubits [18]; see Appendix F.
- [18] 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).
- [19] R. Raussendorf, J. Harrington, and K. Goyal, *A Fault-Tolerant One-Way Quantum Computer*, *Ann. Phys. (N.Y.)* **321**, 2242 (2006).
- [20] R. Raussendorf and J. Harrington, *Fault-Tolerant Quantum Computation with High Threshold in Two Dimensions*, *Phys. Rev. Lett.* **98**, 190504 (2007).
- [21] R. Raussendorf, J. Harrington, and K. Goyal, *Topological Fault-Tolerance in Cluster State Quantum Computation*, *New J. Phys.* **9**, 199 (2007).
- [22] A. Kitaev, *Fault-Tolerant Quantum Computation by Anyons*, *Ann. Phys. (N.Y.)* **303**, 2 (2003).
- [23] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, *Topological Quantum Memory*, *J. Math. Phys. (N.Y.)* **43**, 4452 (2002).
- [24] A. G. Fowler, A. M. Stephens, and P. Groszkowski, *High-Threshold Universal Quantum Computation on the Surface Code*, *Phys. Rev. A* **80**, 052312 (2009).
- [25] S. D. Barrett and T. M. Stace, *Fault Tolerant Quantum Computation with Very High Threshold for Loss Errors*, *Phys. Rev. Lett.* **105**, 200502 (2010).
- [26] M. Gimeno-Segovia, P. Shadbolt, D. E. Browne, and T. Rudolph, *From Three-Photon GHZ States to Ballistic Universal Quantum Computation*, *Phys. Rev. Lett.* **115**, 020502 (2015).
- [27] H. A. Zaidi, C. Dawson, P. van Loock, and T. Rudolph, *Near-Deterministic Creation of Universal Cluster States with Probabilistic Bell Measurements and Three-Qubit Resource States*, *Phys. Rev. A* **91**, 042301 (2015).
- [28] S. D. Barrett and P. Kok, *Efficient High-Fidelity Quantum Computation Using Matter Qubits and Linear Optics*, *Phys. Rev. A* **71**, 060310 (2005).
- [29] S. C. Benjamin, *Comment on "Efficient High-Fidelity Quantum Computation Using Matter Qubits and Linear Optics"*, *Phys. Rev. A* **72**, 056302 (2005).
- [30] L.-M. Duan and R. Raussendorf, *Efficient Quantum Computation with Probabilistic Quantum Gates*, *Phys. Rev. Lett.* **95**, 080503 (2005).
- [31] T. P. Bodiya and L.-M. Duan, *Scalable Generation of Graph-State Entanglement through Realistic Linear Optics*, *Phys. Rev. Lett.* **97**, 143601 (2006).
- [32] Y. Matsuzaki, S. C. Benjamin, and J. Fitzsimons, *Probabilistic Growth of Large Entangled States with Low Error Accumulation*, *Phys. Rev. Lett.* **104**, 050501 (2010).
- [33] Y. Li, S. D. Barrett, T. M. Stace, and S. C. Benjamin, *Fault Tolerant Quantum Computation with Nondeterministic Gates*, *Phys. Rev. Lett.* **105**, 250502 (2010).
- [34] M. Varnava, D. E. Browne, and T. Rudolph, *How Good Must Single Photon Sources and Detectors Be for Efficient Linear Optical Quantum Computation?*, *Phys. Rev. Lett.* **100**, 060502 (2008).
- [35] F. Ewert and P. van Loock, *3/4-Efficient Bell Measurement with Passive Linear Optics and Unentangled Ancillae*, *Phys. Rev. Lett.* **113**, 140403 (2014).
- [36] W. P. Grice, *Arbitrarily Complete Bell-State Measurement Using Only Linear Optical Elements*, *Phys. Rev. A* **84**, 042331 (2011).
- [37] J. Kim, C. J. Nuzman, B. Kumar, D. F. Liewwen, J. S. Kraus, A. Weiss, C. P. Lichtenwalner, A. R. Papazian, R. E. Frahm, N. R. Basavanahally, D. A. Ramsey, V. A. Aksyuk, F. Pardo, M. E. Simon, V. Lifton, H. B. Chan, M. Haueis, A. Gasparyan, H. R. Shea, S. Arney, C. A. Bolle, P. R. Kolodner, R. Ryf, D. T. Neilson, and J. V. Gates, *1100 × 1100 Port MEMS-Based Optical Crossconnect with 4-dB Maximum Loss*, *IEEE Photonics Technol. Lett.* **15**, 1537 (2003).
- [38] The computer code used to generate the results in this paper is openly available online http://figshare.com/articles/Fault_tolerant_LOQC/1557857.
- [39] S. M. Hamdi, S. T. Zuhori, F. Mahmud, and B. Pal, *A Compare between Shor's Quantum Factoring Algorithm and General Number Field Sieve*, in *Proceedings of the International Conference on Electrical Engineering and Information & Communication Technology (ICEEICT), 2014*, http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6919115.
- [40] A. C. Whiteside and A. G. Fowler, *Practical Topological Cluster State Quantum Computing Requires Loss Below 1%*, [arXiv:1409.4880](https://arxiv.org/abs/1409.4880).
- [41] T. Goh, M. Yasu, K. Hattori, A. Himeno, M. Okuno, and Y. Ohmori, *Low Loss and High Extinction Ratio Strictly Nonblocking 16 × 16 Thermo-optic Matrix Switch on 6-in Wafer Using Silica-Based Planar Lightwave Circuit Technology*, *J. Lightwave Technol.* **19**, 371 (2001).

- [42] N. C. Harris, Y. Ma, J. Mower, T. Baehr-Jones, D. Englund, M. Hochberg, and C. Galland, *Efficient, Compact and Low Loss Thermo-Optic Phase Shifter in Silicon*, *Opt. Express* **22**, 10487 (2014).
- [43] M. Ringbauer, D. N. Biggerstaff, M. A. Broome, A. Fedrizzi, C. Branciard, and A. G. White, *Experimental Joint Quantum Measurements with Minimum Uncertainty*, *Phys. Rev. Lett.* **112**, 020401 (2014).
- [44] P. J. Shadbolt, M. R. Verde, A. Peruzzo, A. Politi, A. Laing, M. Lobino, J. C. F. Matthews, and J. L. O'Brien, *Generating, Manipulating and Measuring Entanglement and Mixture with a Reconfigurable Photonic Circuit*, *Nat. Photonics* **6**, 45 (2011).
- [45] S. Miki, T. Yamashita, Z. Wang, and H. Terai, *A 64-Pixel NbTiN Superconducting Nanowire Single-Photon Detector Array for Spatially Resolved Photon Detection*, *Opt. Express* **22**, 7811 (2014).
- [46] C. Niclass, C. Favi, T. Kluter, M. Gersbach, and E. Charbon, *A 128 × 128 Single-Photon Image Sensor with Column-Level 10-bit Time-to-Digital Converter Array*, *IEEE J. Solid-State Circuits* **43**, 2977 (2008).
- [47] X. C. Yao, T. X. Wang, P. Xu, H. Lu, G. S. Pan, X. H. Bao, C. Z. Peng, C. Y. Lu, Y. A. Chen, and J. W. Pan, *Observation of Eight-Photon Entanglement*, *Nat. Photonics* **6**, 225 (2012).
- [48] <http://dx.doi.org/10.5281/zenodo.22558>.
- [49] Creating the required resources through magic state injection and distillation is an important source of the overhead cost in surface-code-based quantum computing [18]. There is an ongoing effort to reduce this overhead [50–57], but for the present purposes, it suffices to remark that the costs at the surface-code level are approximately equal for both matter qubits and photonic qubits; therefore, we do not discuss this complex topic further here.
- [50] S. Bravyi and J. Haah, *Magic-State Distillation with Low Overhead*, *Phys. Rev. A* **86**, 052329 (2012).
- [51] C. Jones, *Multilevel Distillation of Magic States for Quantum Computing*, *Phys. Rev. A* **87**, 042305 (2013).
- [52] A. G. Fowler and S. J. Devitt, *A Bridge to Lower Overhead Quantum Computation*, [arXiv:1209.0510](https://arxiv.org/abs/1209.0510).
- [53] A. G. Fowler, S. J. Devitt, and C. Jones, *Surface Code Implementation of Block Code State Distillation*, *Sci. Rep.* **3**, 1939 (2013).
- [54] P. Mishra and A. G. Fowler, *Resource Comparison of Two Surface Code Implementations of Small Angle Z Rotations*, [arXiv:1406.4948](https://arxiv.org/abs/1406.4948).
- [55] E. T. Campbell, H. Anwar, and D. E. Browne, *Magic-State Distillation in All Prime Dimensions Using Quantum Reed-Muller Codes*, *Phys. Rev. X* **2**, 041021 (2012).
- [56] E. T. Campbell, *Enhanced Fault-Tolerant Quantum Computing in d-Level Systems*, *Phys. Rev. Lett.* **113**, 230501 (2014).
- [57] Y. Li, *A Magic State's Fidelity Can Be Superior to the Operations That Created It*, *New J. Phys.* **17**, 023037 (2015).
- [58] In a monolithic array of deterministically controlled matter qubits [18], typically $2d(d-1)$ ancillary qubits suffice for a total of $(2d-1)^2 \sim 4d^2$ qubits. For photonic qubits, two layers of the cluster state correspond to one round of stabilizer measurements [20], and each layer contains $d^2 + (d-1)^2 + d(d-1)$ qubits (per logical qubit). Therefore, in our protocol, each round of stabilizer measurements requires $(2d-1)^2 + d^2 + (d-1)^2 \sim 6d^2$ cluster-state qubits.