

Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links

Naomi H. Nickerson,¹ Joseph F. Fitzsimons,² and Simon C. Benjamin³

¹*Department of Physics, Imperial College London,
Prince Consort Road, London SW7 2AZ, United Kingdom*

²*Singapore University of Technology and Design, 20 Dover Drive, Singapore 138682, Singapore
and Center for Quantum Technologies, National University of Singapore,*

Block S15, 3 Science Drive 2, Singapore 117543, Singapore

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

(Received 25 July 2014; published 9 December 2014)

Exquisite quantum control has now been achieved in small ion traps, in nitrogen-vacancy centers and in superconducting qubit clusters. We can regard such a system as a universal cell with diverse technological uses from communication to large-scale computing, provided that the cell is able to network with others and overcome any noise in the interlinks. Here, we show that loss-tolerant entanglement purification makes quantum computing feasible with the noisy and lossy links that are realistic today: With a modestly complex cell design, and using a surface code protocol with a network noise threshold of 13.3%, we find that interlinks that attempt entanglement at a rate of 2 MHz but suffer 98% photon loss can result in kilohertz computer clock speeds (i.e., rate of high-fidelity stabilizer measurements). Improved links would dramatically increase the clock speed. Our simulations employ local gates of a fidelity already achieved in ion trap devices.

DOI: [10.1103/PhysRevX.4.041041](https://doi.org/10.1103/PhysRevX.4.041041)

Subject Areas: Quantum Physics

Within the past year, there have been remarkable advances in the fidelity with which small quantum devices can be controlled. The two most mature systems are ion traps and superconducting qubits. In ion trap devices, single-qubit fidelities [1] have reached 99.9999%, with combined preparation and measurement of 99.93%. Moreover, two-qubit operations [2] have been reported with fidelities up to 99.9%. Meanwhile, a superconducting qubit device (SQD) containing five qubits [3] has been demonstrated with all qubit manipulations above 99.3%. At the same time, there has been rapid progress in the study of nitrogen-vacancy (NV) centers in diamond—single electron spin manipulation is possible with 99% fidelity [4], and it is possible to manipulate nuclei that are relatively far from the center, so that each NV center may be thought of as a group of several qubits interacting with an optically active core [5].

These prototype systems are small; none of them contain as many as 20 qubits. But importantly in each case, it is possible to bridge between small systems using photonic channels, albeit with lower entanglement fidelities and in a probabilistic way that may require many attempts. In the ion trap community, there are well-established methods for

entangling ions in separate traps, and recent progress associated with projects such as the MUSIQ initiative [6] have led to successful entanglement at a rate of hertz [7]. This can be improved by orders of magnitude by hardware advances and by loss-adapted protocols, as we describe in this paper. In SQDs, a well-established means of interfacing qubits is to exploit microwave photons in cavities [8]. This suffices for short-range bridging, and, moreover, remote entanglement of two superconducting qubits separated by more than a meter of coaxial cable has recently been demonstrated [9]. In the case of NV center research, successful optical linking of qubits can occur either within the same sample [4] or over meters of separation [10], and teleportation [11] with fidelity around 86% as been achieved.

Thus, the quantum state of the art includes well-controlled small groups of qubits, which we refer to as “cells” in this paper, together with intercell entanglement links that may be nondeterministic and, even when successful, noisy. These ingredients may already suffice to develop fully scalable technologies: although the fidelities over the links are too low to directly enable secure communication or fault-tolerant computing, crucially the fidelities *within* cells are now high enough to support entanglement purification. This process allows one to improve the fidelity of a quantum channel by combining several successive uses of the link. Thus, at the cost of lowering the effective bandwidth, we have a powerful paradigm in which small cells link to one another through a

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.

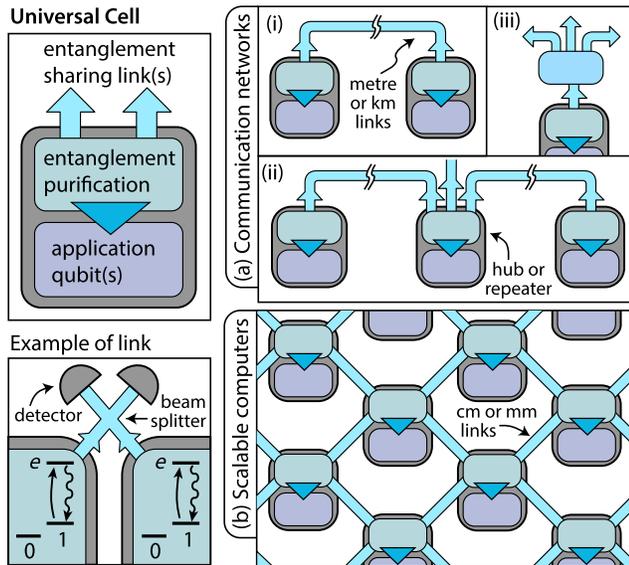


FIG. 1. A small, well-controlled quantum system interfaced to a noisy entanglement-sharing channel constitutes a universal cell if it can purify the entanglement to a high fidelity. Such cells enable secure communication; monogamy of entanglement [12] means the links need not be secure and may be either direct (i) or via repeater hubs (ii) or switches (iii). (b) Moreover, a dense array of cells bridged by short links constitutes a freely scalable computer, as we analyze here.

kind of internal digital filter where purification is performed; see Fig. 1. This paradigm universally supports quantum technologies on any scale. On the large scale, when the bridges between cells are meters or kilometers long, cellular nodes enable secure communication and other distributed information tasks. However, this paper reports on freely scalable quantum computing, where the optical bridges connect a dense array of cells with spacings on the order of centimeters or less.

Any technology based on high-performing cells bridged by very imperfect links will be practical only if entanglement purification is efficient and robust. The protocol should have frugal requirements for “work space” qubits within a cell, it should require only achievable levels of fidelity for local gates and measurements, and, most importantly, it should minimize the time cost by requiring only a few uses of the noisy quantum channel (i.e., a small number of low-fidelity Bell pairs) in order to purify a high-fidelity shared state. These desiderata are in tension with one another, and, of course, the achievable values will depend on the fidelity of the native channel as well as the target fidelity that enables the task in question (e.g., communication or computation). Furthermore, certain tasks, such as the stabilizer-based computing considered here, are best enabled by multiparty entangled states that are more complex than simple Bell pairs.

A seminal Letter in the purification literature is that of Briegel and Dür [13], which showed that by using a tiering

system, one can promote even very noisy “raw” entanglement to a fidelity that is of the order of the fidelity of the local operations. A number of authors have extended this idea, for example, through a means to purify phase noise very efficiently [14] and by introducing the idea of “double selection” [15], which was then used in a scheme [16] that can tolerate channel noise up to 30% in the context of quantum computing. Recent work has even pushed the acceptable limits of local noise to comparable levels [17]. Here, we employ a range of such techniques in order to obtain what we believe to be the most practical purification protocols yet described for the context of quantum computing. Moreover, we show how these protocols can guide specific hardware design to achieve optimum efficiency. We aim to determine whether the intracell operations and limited intercell links that are regarded as achievable today can suffice for full-scale quantum computing, assuming that the various accomplishments that have been made in different (but compatible) experiments can be engineered into a single platform. We conclude that the answer is yes.

Our approach here is an evolution of the scheme in Ref. [18], which was designed to fight network *noise*. We extend that scheme to be efficient versus severe *loss* while retaining the noise tolerance, and this allows us to analyze the “clock speed” of the resulting computer. Our approach requires a total of at least five qubits per cell: four qubits for purification of noise on the cell-cell coupling links and one that is involved in the actual quantum computation (a so-called “data qubit,” as we explain presently). Ion traps, SQDs, and NV center systems can all scale to five qubits. However, with ion traps and SQDs, we may eventually have the luxury of tens of qubits per cell. In that case, we can make good use of the additional structure: presently, we discuss a buffered ion trap design that is optimal for entanglement purification, maximizing the processing speed of a computer formed from such cells. A cell possessing tens of qubits could also embody multiple data qubits—we do not pursue this possibility here, but it is an obvious method for reducing the number of cells required for a given computational task. In any case, a useful quantum computer will require a great many basic cells, but since each cell is likely to be of subcentimeter dimensions, a machine composed of millions of cells could fit within the space allocated to a conventional supercomputer.

To support quantum computing with the cellular paradigm, we must select an approach to achieving fault tolerance—this will effectively set the target fidelity with which purified intercell operations must be performed. We opt to employ a surface code, first introduced by Kitaev and co-workers [19,20], because of its high thresholds and local structure [21]. The approach involves repeatedly measuring certain *stabilizers*—these correspond to simple parity measurements on groups of qubits; i.e., we need to learn whether the total number of 1’s in the group is odd or

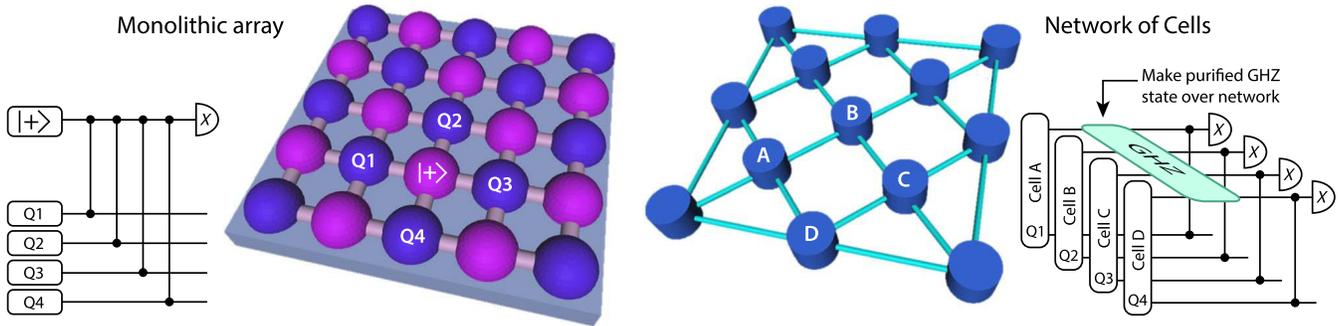


FIG. 2. Monolithic versus network quantum computing. A piece of surface code is represented in the monolithic picture on the left, and in a network architecture on the right. (Here, for clarity we depict a 2D planar code; our network simulations will actually employ the toric code, i.e., periodic boundaries. These two variants are very similar; see Appendix B.) In a monolithic structure [24,25], all qubits are contained within one physical system and two-qubit gates are performed directly between qubits via some physical interaction. In the network picture, the system is instead divided up into small cells, each of which contains a modest number of qubits that interact directly, while between cells only noisy and lossy interactions are possible. For the monolithic system, stabilizer measurements are performed by using a dedicated set of ancilla qubits (pink) interlaced with the data qubits (blue). In contrast, an efficient route to making stabilizer measurements in the network model is to first purify a shared GHZ state between the cells involved and then use this resource to evaluate the stabilizer: The parity of the four qubits measured out from the GHZ tells us the stabilizer outcome. In subsequent cycles, cells are grouped into different sets of four in order to evaluate a complementary set of stabilizers; see Appendix B 3.

even. The basic repeating cycle of the computer involves alternating patterns of parity checks separated by Hadamard rotations to switch between the x and z basis. Remarkably, this simple principle allows for far more than merely protecting quantum information from errors: certain operations between encoded logical qubits can be performed merely by altering the patterns of parity measurements [20,22], and together with a technique such as magic state distillation [23], all the operations required for universal quantum computation can be performed this way.

In a monolithic 2D device, there is a natural layout for the physical qubits such that nearest-neighbor interactions suffice to efficiently perform the parity evaluation (see Fig. 2, left). But how should one find the parity of four data qubits if they are instead incorporated into four different cells? The solution used in Ref. [18] is to employ ancilla qubits within the cells, interacting them with one another across the network so as to build up a four-qubit Greenberger-Horne-Zeilinger (GHZ) state with one qubit in each cell (see Fig. 2, right). Given this GHZ state, it is trivial to deterministically find the parity of the data qubits, as shown in the figure. The challenge is to efficiently make a high-fidelity GHZ state in an efficient manner. We now specify the protocols that we have developed (Fig. 3) and we establish their performance in terms of the fault-tolerance thresholds. The full process of deriving these performance figures is fairly involved, and is described in Appendix C.

In order to minimize the impact of photon loss on entanglement generation rates (and so ultimately maximize the computer’s clock speed), we must optimize the mechanism by which raw entanglement is achieved between cells. Typical schemes for optical entanglement generation, such

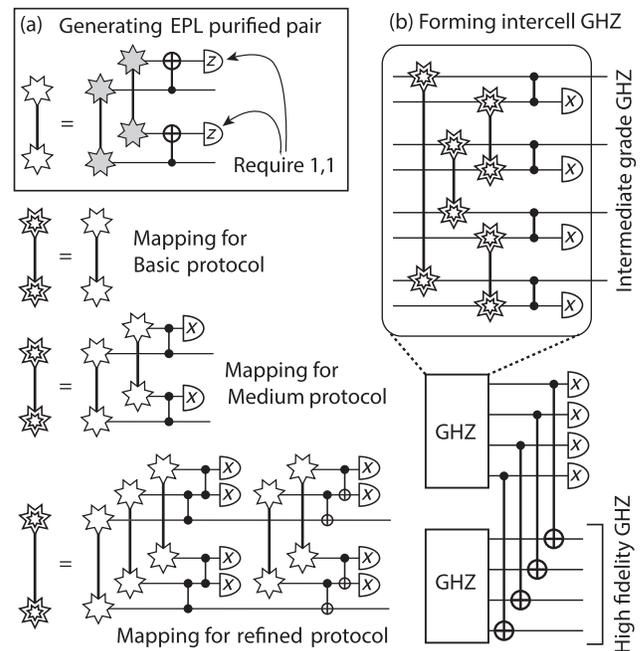


FIG. 3. Distilling a GHZ state with EPL-generated entanglement. (a) The symbols with gray-shaded stars represent the highly mixed Bell states ρ_{raw} or ρ'_{raw} obtained when a single detector “click” is seen. Following the the extreme photon loss protocol [26], we can combine two such pairs and measure one out, thus producing a single greatly improved pair. This process is represented by a symbol with an open star. (b) Circuit diagram showing the adapted GHZ distillation process for entanglement generated using the EPL protocol. Two GHZ states are produced and one is used to make a four-qubit parity projection onto the other. The three different protocols, basic, medium, and refined are shown.

as the Barrett-Kok method [27], are based on heralding by two photons so that the rate of successful entanglement has a quadratic dependence on the probability that a photon avoids loss, $R_{\text{BK}} \sim \frac{1}{2}(1 - p_{\text{loss}})^2$. Here, p_{loss} is the probability that an emitted photon fails to make it through the system and yield a detector click, whether due to loss or to detector failure. The antibunching scheme employed by Monroe and co-workers in Ref. [7] has a similar form: $R_M \sim \frac{1}{4}(1 - p_{\text{loss}})^2$. In both cases, the quadratic dependence is punishing when losses are severe and only rare photons are captured and detected. The use of cavities to enhance matter-light coupling may eventually allow more sophisticated entanglement channels (as recently demonstrated [28,29]), but here we assume that cavities are not employed, and therefore, we must minimize the impact of loss. We adapt a scheme of Campbell and Benjamin [26] called the extreme photon loss (EPL) protocol, which requires one additional qubit at each site and results in success rate $R_{\text{EPL}} \sim \frac{1}{8}(1 - p_{\text{loss}})$, i.e., linear in the photon loss rate (the precise prefactor depends on a parameter in the scheme).

In common with the approach of Barrett and Kok, this scheme requires a system capable of conditionally emitting a photon depending on its state (see lower left-hand panel in Fig. 1). Two optically active “broker” qubits [30], each in a separate cell, are initialized to a state $\sqrt{p_0}|0\rangle + \sqrt{p_1}|1\rangle$ and optically excited causing the $|1\rangle$ component to emit a photon as it immediately decays from a short-lived excited state. Any such photons emitted pass through a beam splitter before impinging on photon detectors. The EPL protocol assumes that useful entanglement is heralded by the detection of a single photon “click.” In the absence of photon loss or noise, this would produce an odd parity Bell state, say, $|\psi\rangle \equiv (|01\rangle + |10\rangle)/\sqrt{2}$. In reality, there are various sources of error in this process. The first is the consequence of general imperfections in the preparation and manipulation of the qubits, which we model by mixing the ideal Bell state with the identity, leading to

$$\rho_{\text{imperfect}} = (1 - p_n)|\psi\rangle\langle\psi| + \frac{p_n}{3} \sum_{i=1,2,3} |\phi_i\rangle\langle\phi_i|, \quad (1)$$

where the ϕ_i are the other three Bell states. We note that the noise model is actually quite general: any state can be “twirled” into this form using local operations, which (being relatively high fidelity) will not significantly degrade the entanglement. However, if the imperfections in the system lead to a state with biased noise, as, for example, if phase noise dominates, then in fact this bias may be advantageous [14]; it is typically most difficult to purify structureless “white” noise of the form assumed here.

Now, in addition to this general noise, we have the specific problems of photon loss and dark counts. With photon loss, the primary issue is that when we see a single click, as required by the protocol, it may be that that, in fact,

two photons were emitted, one from each qubit, but one photon was lost, and we thus incorrectly heralded a success. In that event, the eventual broker state is $|11\rangle$. Because of this possibly, our state is

$$\rho_{\text{raw}} = (1 - r)\rho_{\text{imperfect}} + r|11\rangle\langle 11|, \quad (2)$$

where $r = p_{\text{loss}}/(p_1^{-1} - 1 + p_{\text{loss}})$. In our simulations, we assume that p_{loss} is very severe—it approaches unity and, therefore, $r \approx p_1$. Thus, ρ_{raw} is highly mixed: its two terms will have comparable weight.

Having thus accounted for photon loss, it remains to assess the impact of dark counts. Our protocol proves to be quite robust versus this issue; a full analysis is presented in Appendix F. We show that the key parameter is $d = p_{\text{dc}}/(1 - p_{\text{loss}})$, where p_{dc} is the probability that a given detector registers a dark count in the detection window of a single entanglement attempt. Provided that $d \lesssim 10^{-2}$, then to a good approximation dark counts simply increase the network infidelity; for example, if we set $r = 1/2$, then finite dark counts result in $p_n \rightarrow p_n + 3d$.

Given that we have seen a detector click and so heralded the existence of ρ_{raw} , we now store this state and proceed to create another instance of it. Note that there will typically be many heralded failures, i.e., instances where no detector click is reported, before another success is seen. When that success occurs, we again have a ρ_{raw} , except that we apply an additional random phase shift to account for the fact that a substantial time may pass between creation of the two pairs, so that a finite unknown phase drift in the network may have occurred. This case can be modeled by saying that a π phase shift has occurred with probability p_{drift} , as follows:

$$\rho'_{\text{raw}} = (1 - p_{\text{drift}})\rho_{\text{raw}} + p_{\text{drift}}Z_1\rho_{\text{raw}}Z_1. \quad (3)$$

Note that apart from this possible drift *between* the two heralded successes, the approach is otherwise interferometrically stable: an unknown phase shift that is acquired by both ρ_{raw} and ρ'_{raw} will cancel out in the next step; see Appendix A. This step proceeds as shown in Fig. 3. Within each of the two cells, a local control-NOT operation is performed; it is controlled by the broker associated with ρ_{raw} and targets the broker associated with ρ'_{raw} . The brokers associated with ρ'_{raw} are then separately measured in the z basis. As explained in Appendix A, the measurement outcome 1,1 is inconsistent with *either* of the two pairs ρ_{raw} or ρ'_{raw} having been originally in the $|11\rangle$ state; thus, if that outcome is seen, the $|11\rangle\langle 11|$ component of the surviving entangled pair is removed. A convenient feature of this protocol is that the desired measurement outcome, $|11\rangle$, can be made to correspond to “bright” states of the matter qubits, which have a higher measurement fidelity than their “dark” counterpart in several optical systems.

Of course, the local intracell operations and measurements must themselves be treated as noisy (see Appendix C

for the noise model). Given that they are reasonably high fidelity, the result of a successful “1,1” outcome is the Bell state ρ_{EPL} which, while still imperfect, is far higher fidelity than the parent states ρ_{raw} and ρ'_{raw} . We then take these EPL-derived Bell pairs ρ_{EPL} as the basic resource for our GHZ creation. (Optionally, we could use the EPL protocol to perform some, or all, of the parity projections involved in creating the GHZ; however, this possibility is not explored here.) We introduce three new purification protocols with varying time-versus-fidelity tradeoffs. These are depicted in Fig. 3(b). The basic protocol is fast but tolerates only a limited network error rate. The refined protocol carries out several rounds of entanglement distillation, making it much more robust against noise, but also quite time consuming. The medium option sits between these two extremes.

Having established our procedure for generating shared GHZ states across the network links, we proceed to determine the performance of the quantum machine by simulating and tracking errors. This is an intensive numerical process benefiting from the use of a cluster-scale computer facility. The process is detailed in Appendix C. We summarize it here: For a given set of local error rates, we pick a network error rate p_n and a network size characterized by parameter L such that there are $2L^2$ cells in the complete toric network. We then simulate a large number of stabilizer cycles of the computer. At the end of this simulation, we inspect the state to determine whether the logical qubit was corrupted, a simple “yes or no” outcome. We repeat this numerical experiment many thousands of times (typically 3×10^4) to determine the probability that logically encoded qubits will survive these stabilizer cycles without error; this produces one data point for Fig. 4. This process is now repeated with a different network size—if the larger network has a lower logical qubit error rate, we deem that the surface code is operating successfully and, therefore, our chosen network error rate is within the threshold for fault tolerance.

The analysis is then repeated for different levels of network noise in order to determine the threshold precisely.

The results of these calculations are shown in Fig. 4, from which we find that the basic protocol leads to a threshold of 7.7%, the medium complexity protocol has a threshold of 13.3%, while the most aggressive protocol, refined, is able to tolerate very high network noise of up to 19.4%; note in each case we allow for an additional 1% phase drift between the two rounds of EPL.

The question of how these protocols behave when well below threshold (i.e., the regime where a device would realistically operate) requires a different approach to the Monte Carlo simulations performed here, as considered in several recent works [31,32]. While this is beyond the scope of this paper, we note that using the medium protocol at half the threshold network error (7%) with a lattice size of $L = 16$ yields a logical error rate per L stabilizer rounds of fewer than 1 in 10^6 .

These simulations establish the tolerable levels of error, which are comparable to (but better than) our earlier paper, Ref. [18]. However, because this new approach is founded on the EPL protocol for entanglement generation, the overall time needed to perform a stabilizer—and, hence, the fundamental clock cycle of the quantum computer—will be much faster than in prior schemes.

The achievable computer speed depends on the kind of cell architecture that we have available; see Figs. 5 and 6. An obvious advantage is to have multiple entanglement channels connected to each cell; the example in Fig. 5(c) has eight channels (two to each of its four neighboring cells; see also Fig. 11). There is also a second, independent characteristic of the architecture, which we analyze by distinguishing two limiting cases, the minimal architecture and the buffered architecture. In a minimal system, there are just enough qubits to perform our protocols. There will then be uncertainty as to how long it takes to complete a given

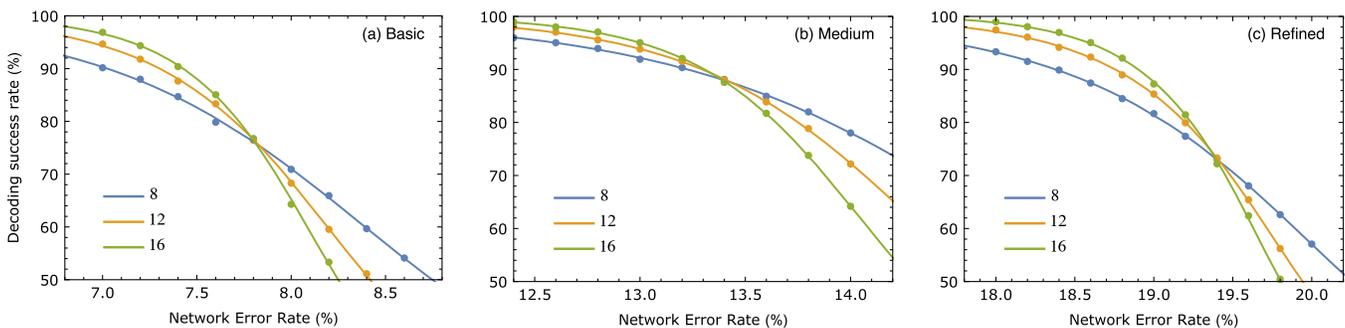


FIG. 4. Results of threshold calculations for the three protocols considered: (a) basic, (b) medium, and (c) refined. The logical error rate in the toric code is calculated for varying values of the network error rate p_n . The cells’ internal error rates are taken to be the best currently demonstrated in an ion trap system: a two-qubit gate error rate of 0.1% and a measurement error rate of 0.05% [1,2]. We select $p_1 = \frac{1}{4} \Rightarrow r \approx \frac{1}{4}$, and we take the phase shift in the EPL entanglement generation to be p_{drift} , of 1%. Details of the error model can be found in Appendix C. The three curves on each plot denote the results for increasing lattice sizes, where $L = 8, 12,$ and 16 (therefore containing $2L^2$ data qubits). The threshold is defined as the intersection of these curves from which we find the basic protocol has a threshold of $p_n = 7.7\%$, medium has a threshold of $p_n = 13.3\%$, and refined has a threshold of $p_n = 19.4\%$. In Appendix E, we further explore the dependence of the network error threshold on the error rates for intracell two-qubit gates and measurements.

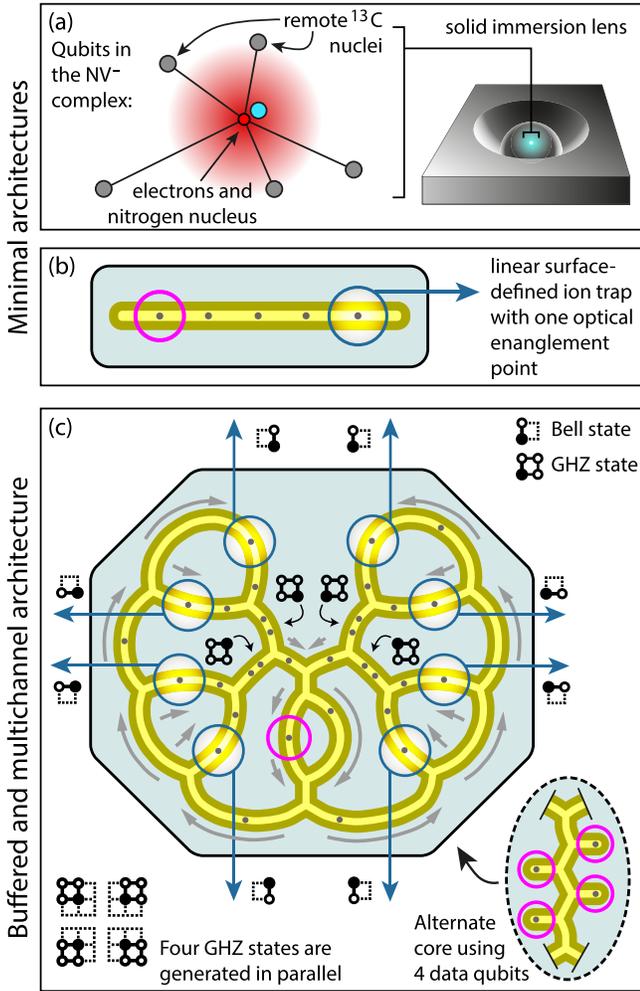


FIG. 5. Example architectures relevant to the cellular network paradigm. In (a) and (b) there are 5 qubits available, the minimum required by the protocol: (a) An NV center with several ^{13}C atoms within range of the core constitutes a five-qubit cell with one qubit (the electron spin) coupling to an optical channel. A simple ion trap (b) need only have five ions, but a more complex architecture (c) offers the advantage of temporarily storing, or “buffering,” the incomplete GHZ states. In this illustration, the eight independent entanglement sites further enhance the GHZ generation rate and thus increase the “clock speed” of the computer; this also obviates the need for optical switching, as illustrated in Fig. 6. The small square symbols indicate the generation of Bell pairs between cells and the subsequent synthesis of GHZ states out of those Bell pairs. A filled circle indicates an ion in this cell, whereas open circles are ions in neighboring traps; this may be more apparent from the multicell schematic in Fig. 11.

stabilizer measurement (since the protocols are probabilistic) and this necessitates a delay for synchronization; see Appendix D 1. In contrast, a buffered architecture has additional internal storage allowing us to queue our qubits and smooth out the timing irregularities, thus avoiding the difficulty in synchronization. Table I summarizes the time

TABLE I. The threshold of tolerable network error rates for each of the three distillation protocols considered, and the time cost for making a complete high-fidelity four-qubit GHZ assuming we operate well under threshold (3%, 5%, and 7% for the three protocols, respectively). Such a GHZ state enables a stabilizer measurement. The distinction between minimal and buffered architectures is illustrated in Fig. 5.

| Protocol | Threshold error rate ($p_n + p_{\text{drift}}$) | Time to make GHZ (units of T_0) | |
|----------|---|------------------------------------|-----------------------|
| | | Minimal architecture | Buffered architecture |
| Basic | 7.7% + 1% | 22 | 5.2 |
| Medium | 13.3% + 1% | 47 | 12.2 |
| Refined | 19.4% + 1% | 102 | 31.6 |

cost to perform a high-fidelity stabilizer measurement on four data qubits. It is quantified in terms of T_0 , the time to produce a single basic EPL Bell pair, i.e., the state ρ_{EPL} .

Given all of these contributing factors, we can now estimate an achievable rate for the clock cycle of our computer. We neglect the time for local gates and measurements; thus, our estimate is accurate only if such gates

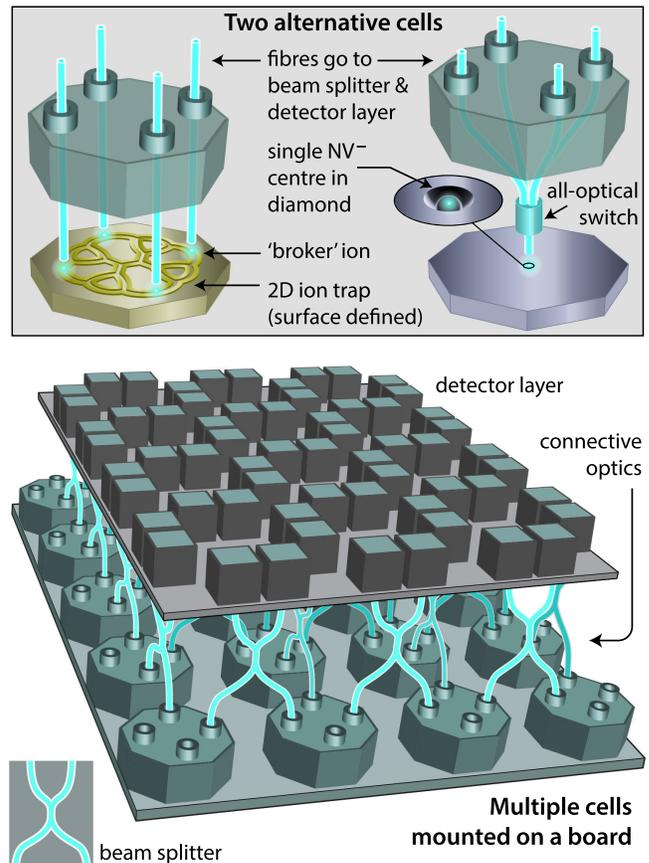


FIG. 6. Illustration of the cellular architecture in 3D. See also the detailed schematics in Figs. 10 and 11.

are performed on the scale of microseconds. This appears achievable, but we note that in established experiments, the highest fidelities are seen for longer gate times (see, e.g., Ref. [2]). One cycle of a surface code quantum computer corresponds to a set of parity measurements over all its data qubits (either in the x basis or the z basis, alternatingly). Suppose that the cells in our machine correspond to the design in Fig. 5(c), or a superconducting qubit device of equivalent complexity. Assume that each cell-cell link is an entanglement channel, which is realistic with today’s technology: the entanglement attempt rate is 2 MHz and the end-to-end photon detection probability is only 2%. We select $p_1 = \frac{1}{4}$ and find that the average time cost for an entanglement channel to create one Bell resource (ρ_{EPL}) is $T_0 = 0.27$ ms. Now, further assume that we have opted for the medium purification protocol because we have network noise at a level of 5% (well within the medium protocol’s threshold of 13.3%). According to Table I, a single-channel cell will require time $12.1T_0$ to create one high-fidelity GHZ state. Our cells have eight channels, which together generate such GHZs at a rate of 2.5 kHz; however, two GHZ states per cell are consumed in making a complete set of stabilizer measurements (either x basis or z basis), as explained in Appendix C 3. Therefore, our overall clock rate is 1.2 kHz.

Higher rates could be achieved simply by introducing more entanglement channels [the branched design in Fig. 5(c) obviously generalizes from 8 channels to 2^N]. This is consistent with ideas in the MUSIQC project [6]. Alternatively, if we look to the medium-term future and assume that the use of integrated cavities [29] (or other advances) can reduce the photon loss rate to $\sim 50\%$, and that the network noise can be taken well below the 7.5% threshold of our basic protocol, then the same device design in Fig. 5(c) should begin to approach megahertz rates for stabilizer measurement. At this point, the local gate speeds may be the limiting factor.

In conclusion, we consider an architecture for quantum technologies that is motivated by the recent achievements in ion traps, superconducting devices, and NV centers. We consider small quantum “cells” composed of a few (5-to-50) qubits under high-fidelity deterministic control, together with intercell network links that are both noisy and lossy and therefore nondeterministic. This architecture is relevant to communication when the links are long (e.g., kilometers), but our focus here is on quantum computing using a large number of such cells with short (e.g., centimeter) bridges. We find that the exquisitely high levels of control recently achieved in small systems can enable very compact and efficient entanglement purification, allowing one to use relatively poor photonic links between cells. We consider three different purification protocols and derive the corresponding thresholds for fault-tolerant quantum computing, finding that threshold network fidelity can be as low as 80%. Moreover, we study

the time cost of the purification process, and thus the time to evaluate a set of stabilizer measurements across the network—effectively, the “clock speed” of the quantum computer. We relate this speed to the complexity of the cell. Given cells that are sufficiently complex to incorporate parallel operations and buffering (temporary storage), we find that even the highly lossy links that are realistic today should support kilohertz-rate, freely scalable quantum computing.

ACKNOWLEDGMENTS

We acknowledge very useful conversations with Chris Ballance, Earl Campbell, Austin Fowler, Ronald Hanson, Tom Harty, Winfried Hensinger, and David Lucas. Computing facilities were provided by the Advanced Research Computing service of the University of Oxford, the Imperial College High Performance Computing Service, and the Dieter Jaksch group. S. C. B. acknowledges EPSRC platform Grant No. EP/J015067/1. This material is based on research funded in part by the Singapore National Research Foundation under NRF Award No. NRF-NRFF2013-01.

APPENDIX A: EXPLANATION OF THE EPL METHOD

Here, we describe the EPL protocol’s handling of photon loss and its inherent tolerance of systematic phase errors. We neglect all other imperfections, but of course these are accounted for elsewhere in our analysis and in our simulations (cf. Appendixes C 1 and F).

For simplicity, we initialize each of our two remote “broker” qubits into the state $(|0\rangle + |1\rangle)/\sqrt{2}$. (Note that, in fact, we need not start from an equal superposition; one selects an optimal level of excitation.) The $|1\rangle$ state is optically active; we excite both brokers and route the collected light through a beam splitter prior to detection as in the standard picture. If we see a photon, then we have

$$\rho_{\text{simple}} = (1 - r)|\Theta\rangle\langle\Theta| + r|11\rangle\langle 11|,$$

where $|\Theta\rangle = (|01\rangle + e^{i\phi}|10\rangle)/\sqrt{2}$ and ϕ is some phase introduced due to the optical apparatus. The ratio r depends on the severity of photon loss; in the limit of loss tending to unity, we find $r \rightarrow \frac{1}{2}$. An intuitive explanation of this is as follows: the state $|\Theta\rangle$ generates only one photon whereas $|11\rangle$ generates two; however, the state $|\Theta\rangle$ has twice the probability of $|11\rangle$ in the original broker-broker product state, thus these factors compensate and the two states have equal weight in the mixture. Now we store ρ_{simple} and attempt to create another such pair. We may fail a number of times before we again succeed. Provided that the apparatus does not suffer phase “drift” *between* the two successful events, the second entangled pair will have the exact same form, ρ_{simple} . (Recall that our full analysis in the

main text does account for the possibility that such a drift occurs.)

Having obtained two noisy entangled pairs, each of the form ρ_{simple} , we now apply CNOT gates locally within each cell according to the Fig. 3(a) circuit. These CNOT gates map each of the pure states within our $\rho_{\text{simple}} \otimes \rho_{\text{simple}}$ mixture as follows, where we take the left-hand side to be the controlling qubits and the right-hand side as the target qubits:

$$\begin{aligned}
(|01\rangle + e^{i\phi}|10\rangle)(|01\rangle + e^{i\phi}|10\rangle) &\rightarrow \\
|01\rangle(|00\rangle + e^{i\phi}|11\rangle) + e^{i\phi}|10\rangle(|11\rangle + e^{i\phi}|00\rangle) & \\
|11\rangle(|01\rangle + e^{i\phi}|10\rangle) &\rightarrow |11\rangle(|10\rangle + e^{i\phi}|01\rangle) \\
(|01\rangle + e^{i\phi}|10\rangle)|11\rangle &\rightarrow |01\rangle|10\rangle + e^{i\phi}|10\rangle|01\rangle \\
|11\rangle|11\rangle &\rightarrow |11\rangle|00\rangle. \tag{A1}
\end{aligned}$$

The protocol then calls for the one of the qubit pairs, the right-hand side pair in the present notation, to be measured in the z basis. Any result other than “11” is rejected. We see that only the first of the four possibilities listed above can pass this filter. Revisiting Eq. (A1) and collecting terms, we have

$$\begin{aligned}
|01\rangle(|00\rangle + e^{i\phi}|11\rangle) + e^{i\phi}|10\rangle(|11\rangle + e^{i\phi}|00\rangle) \\
= (|01\rangle + e^{2i\phi}|10\rangle)|00\rangle + e^{i\phi}(|01\rangle + |10\rangle)|11\rangle.
\end{aligned}$$

Thus, measuring “11” implies that the remaining qubit pair is state $(|01\rangle + |10\rangle)/\sqrt{2}$ and we have eliminated both the $|11\rangle$ components due to photon loss and the unwanted phase ϕ (which, therefore, we need not know). This occurs with probability $\frac{1}{2}(1-r)^2$, i.e., $\frac{1}{8}$ when $r = \frac{1}{2}$.

In practice, of course, there are other sources of error, both in the network and in the local gates and measurements, as discussed in the error model below. But for all levels of noise relevant to the devices we are considering, the result of this process is to generate a state ρ_{EPL} , which is far higher fidelity than the two parent states.

APPENDIX B: PLANAR VERSUS TORIC NETWORK TOPOLOGIES

The main text presents results for the toric code. In this version of the surface code the boundaries of the surface wrap to form a torus. While this would be difficult to realize if we were employing a monolithic structure (cf. Fig. 2, left), in a network paradigm there is no in-principle difficulty. However, it might be that there are reasons to prefer to layout the network in 2D and maintain all links of the same physical length—in this case, one would adopt the planar version of the surface code. Our threshold-finding numerical programs can simulate either the toric or the planar variant. We find that while the toric code exhibits a slightly sharper threshold because it has no “edge effect,” in

fact, the value of a threshold obtained from the two approaches (given all other factors are held the same) varies only slightly.

APPENDIX C: SIMULATION METHODS

Here, we give an overview of the methods used in calculating threshold values for our system. The approach can be divided into two distinct sections: In the first part, we derive superoperators representing the net effect on the data qubits of our stabilizer measurement protocols with all their various errors. In the second part, we use a classical algorithm to track the effect of these superoperators as we simulate a surface code embodying logical qubits.

1. Error model

All of the protocols we consider are composed of a small number of low-level basic operations, each with an associated noise model. 1. Network error model. This is described in the main text. States ρ_{raw} and ρ'_{raw} are defined in Eqs. (2) and (3). They involve the network noise rate p_n and the photon loss rate p_{loss} , as well as a parameter p_{drift} accounting for phase drift in the entanglement channel between the creation of the two states. The dark-count rate d can be subsumed into the network noise p_n as we explain in Appendix F. 2. Local (intracell) controlled- Z and controlled- X gates. For a gate error rate p_g , the noise is modeled as a perfect gate operation, which with probability p_g is followed at random with one of the 15 nontrivial two-qubit Pauli errors $\sigma_i \otimes \sigma_j$, where $i = 0, 1, 2, 3$ and σ_0 is the identity. If ρ represents the ideal state after gate operation, this noise map can be written as

$$N_{\text{gate}}(\rho) = (1 - p_g)\rho + \frac{p_g}{15} \sum_{(i,j) \neq (0,0)} (\sigma_i \otimes \sigma_j)\rho(\sigma_i \otimes \sigma_j)^\dagger. \tag{C1}$$

3. Single-qubit measurement in the X and Z bases. Given a measurement error rate p_m , a particular outcome of the measurement, $q \in \{0, 1\}$, corresponds to the intended projection P_q applied to the state with probability $(1 - p_m)$ and the opposite projection $P_{\bar{q}}$ applied with probability p_m . This noisy projector can be written as

$$\mathcal{P}_q(p_m) = (1 - p_m)|q\rangle\langle q| + p_m|\bar{q}\rangle\langle \bar{q}|. \tag{C2}$$

2. Stabilizers as superoperators

To characterise the entire process of the stabilizer measurement, we carry out a full simulation of the measurement procedure including all sources of noise and use the Choi-Jamiolkowski isomorphism [33] to generate a superoperator from the result. Thus, we completely describe the action of the stabilizer measurement procedure with

$$\mathcal{S}(\rho) = \sum_{i=0} p_i K_i \rho K_i^\dagger. \tag{C3}$$

This probabilistic decomposition describes the operation as a series of Kraus operators K_i applied to the initial state with probabilities p_i , which depend on the chosen protocol, noise model, and the error rates. The leading term $i = 0$ will have corresponding K_0 representing the reported parity projection, and large p_0 . For the protocols considered here, the other Kraus operations can be decomposed and expressed as a parity projection with additional erroneous operations applied. For example, if a noisy stabilizer measurement is made that returns an “even” outcome, we find $K_0 = P_{\text{even}}$, the reported even parity projection, and $K_1 = P_{\text{odd}}$, which implies that a perfect odd parity projection was applied, but the wrong outcome was recorded—a “lying” stabilizer measurement. All of the other K_i can be represented as K_0 or K_1 followed by single-qubit Pauli errors. This decomposition then involves two distinct types of error: lies, where an incorrect outcome is recorded, and qubit errors, where a physical error occurs on a data qubit. The probability of each combination of events can be calculated from the values of the p_i . This information on stabilizer performance then enables classical simulation of a full planar code array, and its fault-tolerance threshold can be assessed.

3. Scheduling stabilizer measurement

Each qubit in the body of the lattice is part of four different stabilizer groups. Therefore, in a physical implementation, the measurement of a full cycle of stabilizers is divided into four distinct rounds, two of plaquette measurement and two of star measurement, as shown in Fig. 7.

For the purpose of simulation, it is desirable to break down the evolution of the lattice into complete rounds of perfect plaquette or star measurement separated by rounds of errors. This can be achieved by making use of the fact that each Kraus operator can be decomposed in different but equivalent ways, namely, that each K_i can be written with the Pauli errors either preceding or following a parity projection.

4. Decoding

Decoding is performed using Kolmogorov’s Blossom V implementation of Edmonds’ minimum weight-matching algorithm [34,35] to generate a “perfect” matching between stabilizer violations. To do this, the syndrome on the lattice must be formulated into a weighted graph. In the case of perfect measurements, each “-1” outcome in the syndrome becomes a node of a completely connected graph, where the weight of each edge is given by the distance between the corresponding two nodes on the lattice.

Multiple rounds of stabilization are performed, producing a three-dimensional syndrome cube where the third dimension represents time. Each point where one stabilizer measurement differs from its value in the previous round gives rise to a node in the graph. Matching in the spatial dimensions of the cube correct for physical errors on the lattice, while timelike matchings correct for lying stabilizers. The rate of lie-type errors and physical errors will not generally be the same; to account for this, timelike and spacelike paths are weighted differently. The ratio of these weights is chosen to optimize performance.

In the toric code, error chains on the surface will always result in two stabilizer violations, one at each end of the

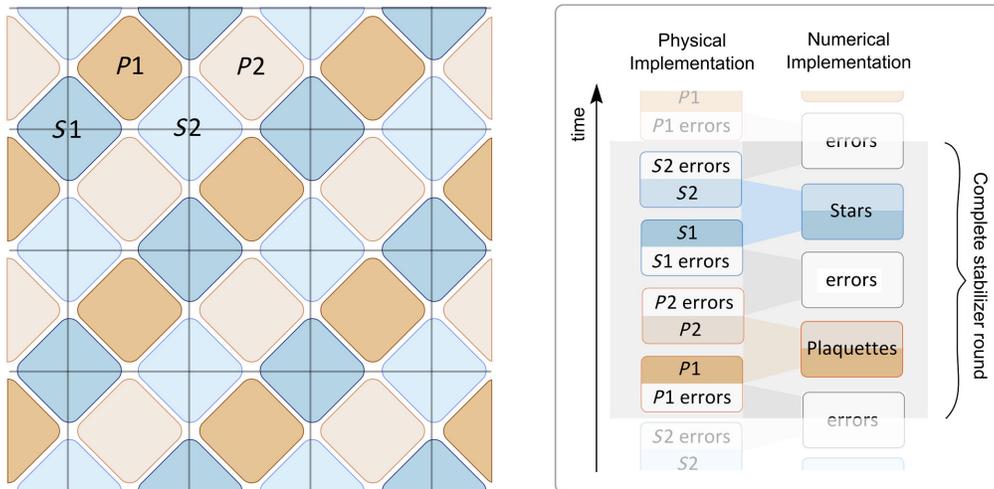


FIG. 7. Scheduling stabilizer measurements. Stabilizer measurements are split into four rounds, two of each type, plaquettes (orange) and stars (blue). The underlying lattice and qubit structure is shown here in black for a planar code of lattice dimension $L = 4$. Decomposition of the stabilizer measurement procedure allows each round to be described as a round of perfect measurements either followed or preceded by errors. The stabilizer implementation including these errors is shown on the right.

chain. In the planar code, however, if an error chain reaches an edge of the lattice, only one stabilizer violation will be seen. To account for this, each node in the original graph is uniquely connected to a new node located at the nearest boundary position, following the method described in Ref. [36]. This gives the possibility for each -1 stabilizer to match to a boundary as well as any other node on the lattice itself.

The process described above is a “vanilla” implementation of the perfect-matching decoder—there are many possibilities for optimizing the decoder that are not pursued in this paper. For example, it is well understood that correlating the X and Z errors reveals information about Y errors. Moreover, there are opportunities to exploit the classical information that occurs during a stabilizer evaluation: Most importantly, in the case where we are using a simple, serial architecture with no buffering available (a NV-center-based technology, for example), we need to impose a cutoff time after which an attempt to measure a stabilizer is abandoned; see Appendix D 1. At present, our decoder makes no use of the information that a given stabilizer has not been evaluated, and instead simply replaces the missing information with a copy of the previous result—this is obviously not optimal.

We emphasize that these limitations in our decoder do not undermine the accuracy of the simulations in this paper, since the operator of a quantum computer is free to use any classical decoder they wish when they run the machine. Thus, the thresholds that we find should be considered a lower bound on the achievable thresholds—a better decoder should boost the performance.

5. Calculating the threshold

We wish to determine the logical error rate of the code as a function of the physical error rates. To do this, we perform Monte Carlo simulations of a lattice under noisy stabilizer measurement. For each instance, the evolution of the lattice is simulated by applying random errors drawn from the distributions specified by the derived superoperators. A total of $3L$ complete stabilizer rounds are performed, where L is the lattice dimension, before a decoding attempt is made and the result analyzed to test whether a logical error has occurred. For each physical error rate, the logical error rate is calculated for three different lattice sizes, $L = 8, 12, 16$. For each data point, a minimum of 30 000 instances are simulated, and error bars are calculated by treating each result as a sample drawn from a Bernoulli distribution.

If the error rate is below the threshold, then increasing the lattice size will improve the performance of the code; that is, the logical error rate will become smaller. So, to find the threshold, we must find the point at which the curves from the different lattice sizes intersect. To estimate the threshold error rate, we use the method

described by Wang *et al.* [36] to model the behavior of the logical error rate close to the crossing. This tells us that for a large enough lattice size L , the decoding failure probability is given by

$$P_{\text{fail}} = (p - p_{\text{th}})L^{1/\nu_0}. \quad (\text{C4})$$

The threshold data are fitted to a quadratic function, to account for small system-size effects, and the threshold crossing value drawn from the resultant fit parameters:

$$P_{\text{fail}} = a + b(p - p_{\text{th}})L^{1/\nu_0} + c(p - p_{\text{th}})^2L^{2/\nu_0}. \quad (\text{C5})$$

APPENDIX D: COMPUTER OPERATIONAL SPEED

The entanglement generation described in the previous section is only one aspect that determines the overall speed of the device. Figure 8 shows a summary of all factors contributing to the final clock cycle of the quantum computer. The contributors are divided into three categories. Entanglement generation has already been discussed. We now consider the factors affecting GHZ state distillation time and stabilizer measurement.

1. Handling probabilistic stabilizer evaluation

The generation and distillation of entanglement using optical links between cells is a probabilistic process, and steps must sometimes be attempted repeatedly. If the cells do not have sufficient internal complexity to queue up, or buffer, the results of the different stages of the purification, then necessarily the time taken to complete a stabilizer measurement will also be probabilistic. In that case, there is a potential difficulty in performing a complete set of stabilizers (or a complete subset, cf. Fig. 7) over the entire computer—should we wait until the very last stabilizer has been successfully performed, before moving to the next set? In fact, to do so would require a time cost that scales with the computer size. Fortunately, this is not necessary; instead, we can simply wait a fixed time and then abandon any stabilizers that have not yet been measured. Figure 9 shows the results from recalculating thresholds for the case where each stabilizer has a 1% chance of not being evaluated; this shows a minimal difference to the threshold. Limiting the number of stabilizers that must be completed to a fixed fraction of the whole also keeps the mean evaluation time constant in the lattice size. We note that our GHZ-based approach to stabilizer measurement is particularly “friendly” to this process of abandoning the “slowest” 1% of measurements in each round, because when we abandon an attempt to create a GHZ, the data qubits in those cells have not been involved in any gate operations. Operations on data

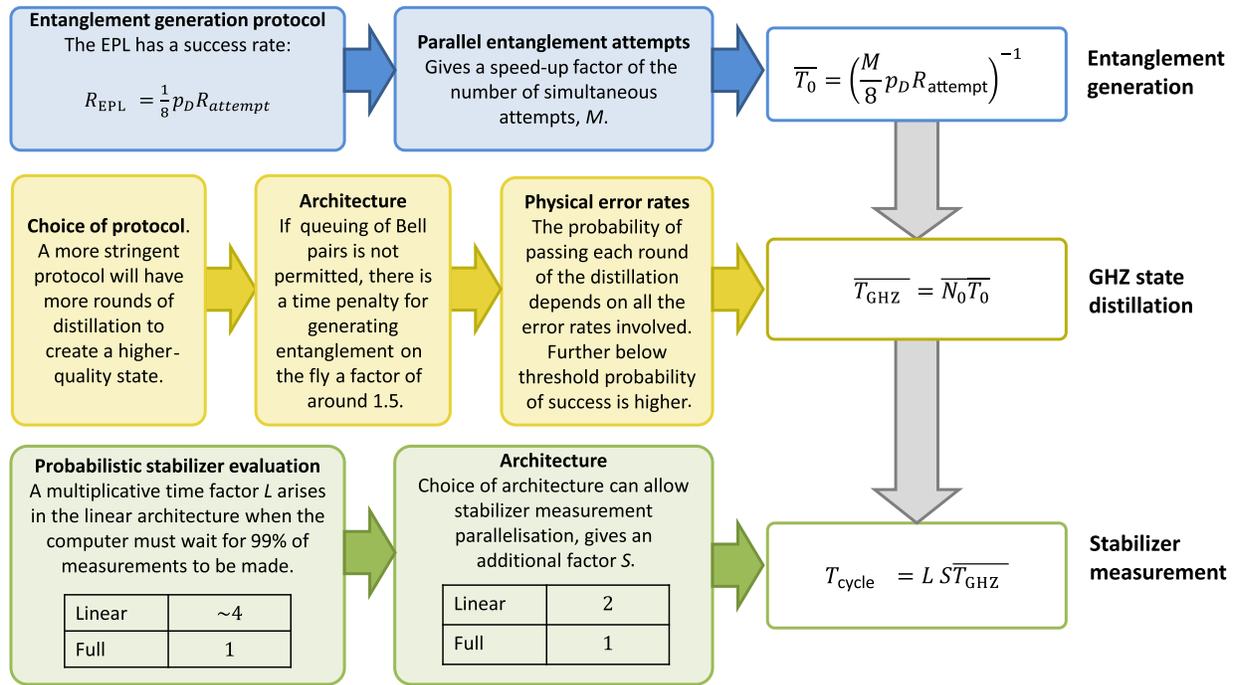


FIG. 8. The seven major factors contributing to the overall potential running speed of the device.

qubits take place only after successful completion of the high-fidelity GHZ state.

This approach is still less than ideally efficient since on average cells will be inactive for a significant portion of their time, having finished well before the “cutoff time.” Since some potential technologies for the cell, including, for example, NV centers, do not have the internal complexity to act as a sophisticated device such as Fig. 5(c), it is interesting to ask whether there is another route with higher efficiency. A possibility is an asynchronous stabilizer measurement, where measurements are made not in discrete rounds but as soon as the necessary entanglement has been generated.

2. Cell design to support parallelization

The design of the cell can lead to a number of ways to parallelize the protocol. The full architecture shown in Fig. 5(c) is designed to exploit all of these possibilities.

Figure 10 shows the relatively modest cells needed for basic functionality, while Fig. 11 shows cells that boast a fully parallel architecture. At the lowest level, entanglement generation can be parallelized. If M simultaneous attempts at entanglement are made, the effective rate of production is, of course, increased by the same factor. Further to this, a cell may be able to support entanglement “queuing” such that Bell states are continuously created and stored for later use. In such an architecture, entanglement can be treated as being deterministically generated at the mean rate (or slightly less, to maintain buffers). On the other hand, if entanglement must be generated “on the fly,” as required, then this must be treated as a stochastic process. At a higher level, GHZ states can be stored as and when they are created, leaving other qubits free to generate more. This removes the inefficiency discussed in the previous section, where most cells must wait for the slowest stabilizer measurements to complete. Instead,

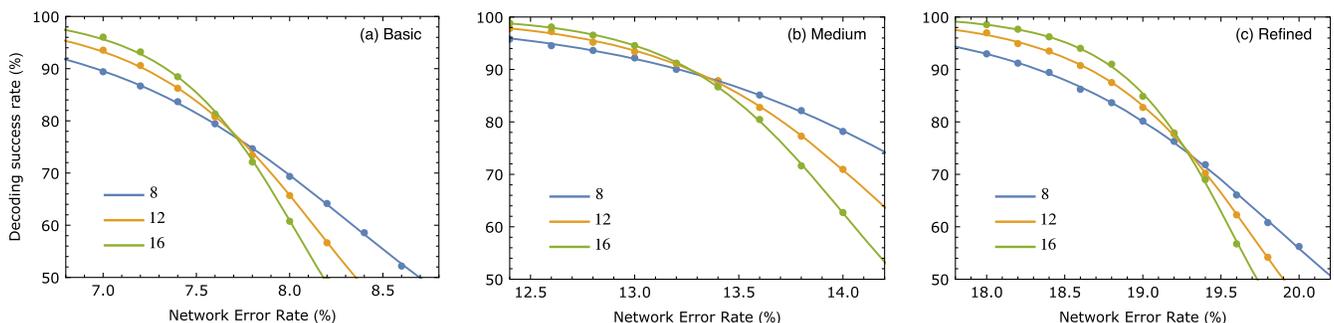


FIG. 9. Thresholds with 1% of stabilizer outcomes missing: (a) Basic, (b) medium, (c) refined.

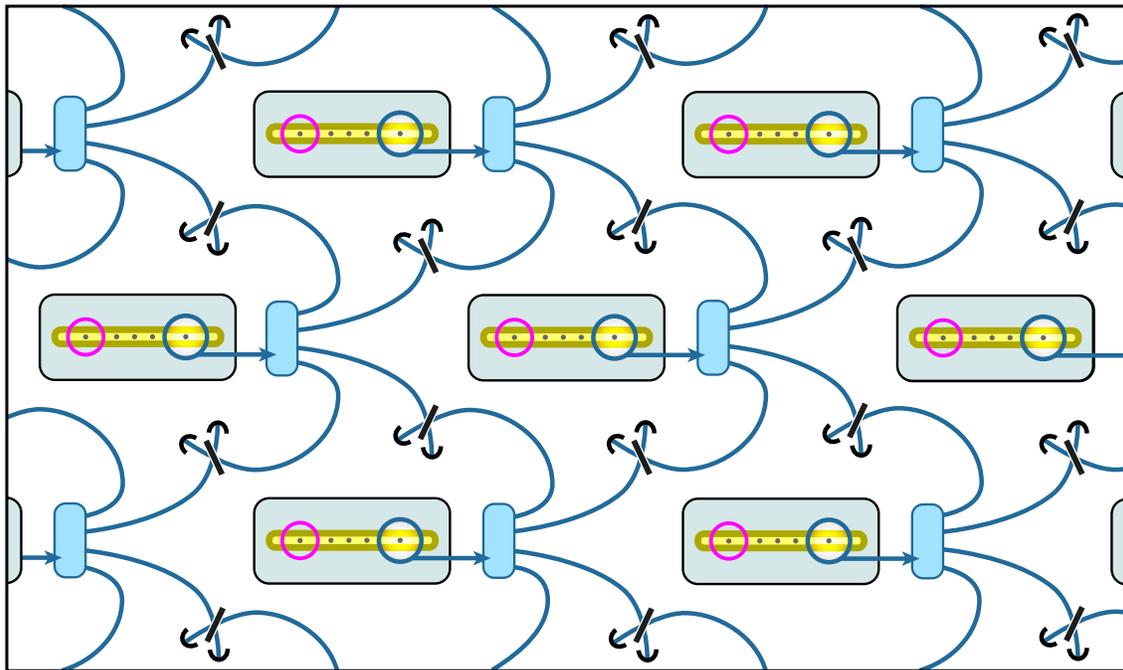


FIG. 10. Four cells of the design shown in Fig. 5(b) with the connections achieved by optical switching. Note that the ion trap elements in this figure could equivalently be any other few-qubit, optically active system such as a NV center in diamond. The system shown here is equivalent to that shown in Fig. 11, which has more complex cells and dedicated (nonswitched) cell-cell links; those features increase the speed, but the error thresholds are the same.

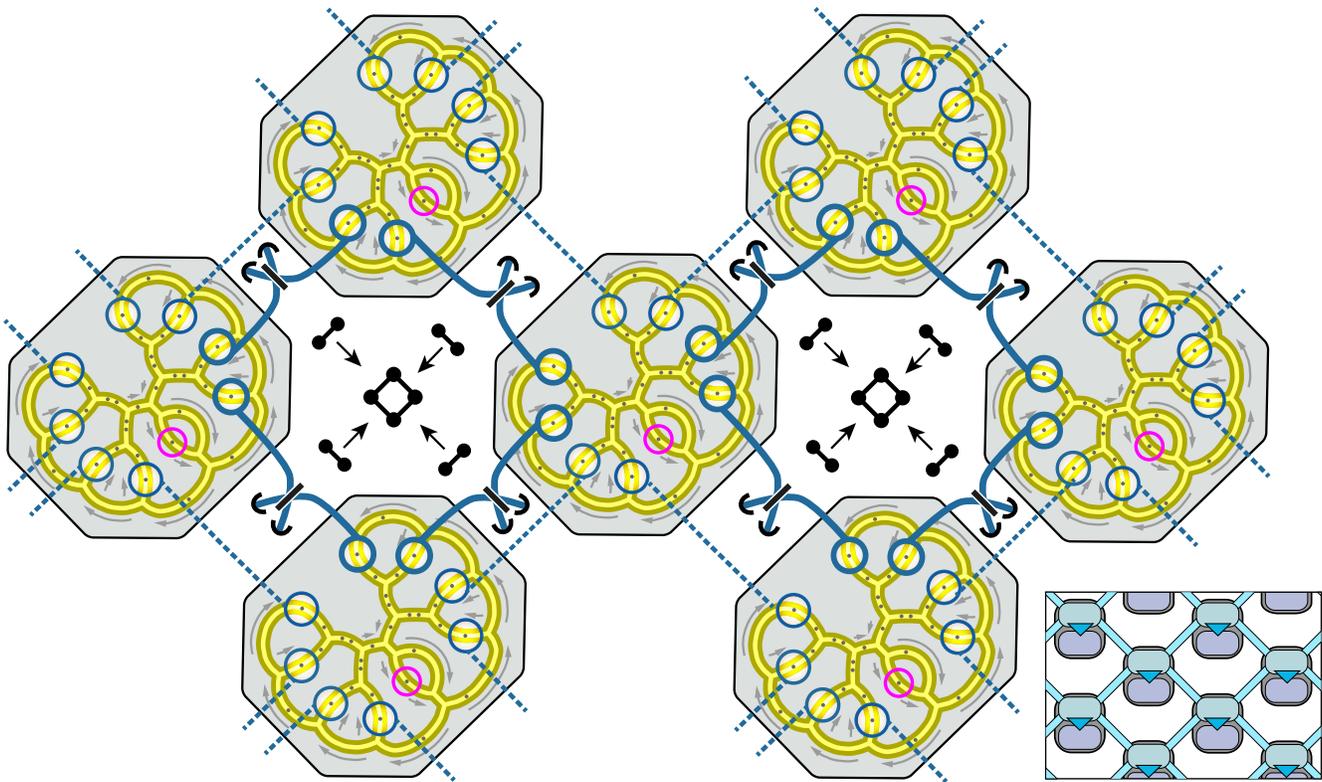


FIG. 11. Four cells of the design shown in Fig. 5(c) with the connections relevant to building their mutual GHZ states highlighted. The inset shows the abstract concept of the networked computer from Fig. 1; this ion trap design and the design in Fig. 10 are specific realizations.

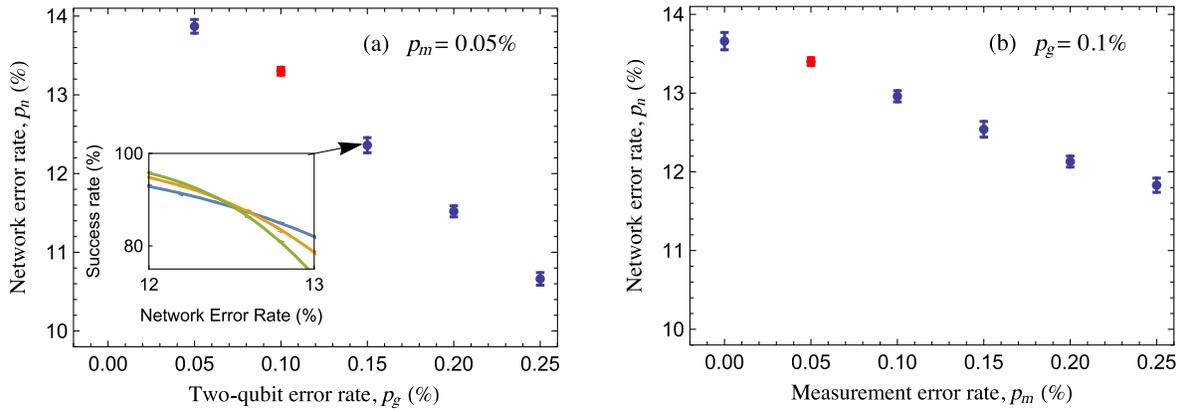


FIG. 12. Threshold dependence on local error rates for the medium protocol. The threshold for fault tolerance is dependent on all the physical error rates in the system. In Fig. 4, we show the threshold in the network error rate p_n when the local measurement and gate error rates p_m and p_g are fixed to 0.05% and 0.1%, respectively. Here, we demonstrate the robustness of the protocol to changes in these intracell error rates. In (a), we show the threshold's dependence on p_g , while (b) reveals the dependence on p_m . Each plotted point represents an entire threshold calculation as shown in the inset of (a). The red data points correspond to the values discussed in the main text. We stress that as the errors deviate further from this point, it is likely that our protocol becomes suboptimal; thus, the threshold line should be regarded as a lower bound on the possible threshold.

stabilizers can be measured at the mean rate of GHZ distillation, providing a further level of speed-up. Finally, the greater connectivity of the full architecture allows all the GHZ states required for a stabilizer round (plaquettes or stars) to be simultaneously distilled. This gives an additional speed-up factor of 2 over a minimal cell where this process must be broken into two rounds, as shown in Fig. 7.

APPENDIX E: Probing the Threshold

In the main text, we fix the measurement and gate error rates to $p_m = 0.05\%$ and $p_g = 0.1\%$ [1,2] and then vary only the network noise to identify a threshold, reporting the results in Fig. 4. The threshold is, of course, a function of all the input error rates. We now probe the behavior of the threshold when these local error parameters are varied. The threshold is evaluated for several values of p_g using the medium protocol and the method described in Appendix C 5. The results are shown in Fig. 12. The threshold exhibits a roughly linear dependence on both p_g and p_m , but a stronger dependence on the gate errors is seen. It is important to note, however, that all threshold calculations here are made using the same protocol, which we design to optimize the threshold for the parameters considered in the main text. For any significantly different combination of error values, a new protocol designed for this regime would be likely to outperform the results we find here.

APPENDIX F: ANALYSIS OF DARK COUNTS

We begin by considering the simple case of $r = 1/2$, where it is straightforward to show that the effect of dark

counts is to produce an adjusted level of network noise p_n . From the main text, we have the following expressions for the case where the dark counts are neglected:

$$\rho_{\text{imperfect}} = (1 - p_n)|\psi\rangle\langle\psi| + \frac{p_n}{3}(|\phi\rangle\langle\phi| + |00\rangle\langle 00| + |11\rangle\langle 11|) \quad (\text{F1})$$

and

$$\rho_{\text{raw}} = (1 - r)\rho_{\text{imperfect}} + r|11\rangle\langle 11|. \quad (\text{F2})$$

Let us take the simple case that $r = \frac{1}{2}$, and introduce $d = p_{\text{dc}}/(1 - p_{\text{loss}})$ as above, where p_{dc} is the probability that the system will experience a dark count in a given *one* of the two detectors during a given attempt at entanglement. We can write a new expression ρ'_{raw} , which has the same form above, but where p'_n and r' replace the unprimed parameters and have absorbed the dark-count parameter d . Roughly speaking, $p'_n \approx p_n + 3d$.

Consider the limit of high photon loss, where almost all dark-count events occur on occasions when all emitted photons have been lost. When we see a dark count on such an occasion, we wrongly conclude that we have heralded the creation of ρ_{raw} . In fact, the state of the two optically active qubits is simply the completely mixed state, because they were prepared in an equal superposition and then (effectively) measured by the environment. Thus, we can write

$$\begin{aligned}
\rho'_{\text{raw}} &= \left(\rho_{\text{raw}} + 2d \frac{\mathbb{1}}{4} \right) / (1 + 2d) \\
&= \left(\frac{1}{2} \rho_{\text{imperfect}} + \frac{1}{2} |11\rangle\langle 11| + \frac{d\mathbb{1}}{2} \right) / (1 + 2d) \\
&= (1 + 2d)^{-1} \left[\left(\frac{1 - p_n}{2} + \frac{d}{2} \right) |\psi\rangle\langle\psi| + \left(\frac{p_n}{6} + \frac{d}{2} \right) (|\phi\rangle\langle\phi| + |00\rangle\langle 00|) + \left(\frac{p_n}{6} + \frac{d}{2} + \frac{1}{2} \right) |11\rangle\langle 11| \right] \\
&= (1 - r') \rho'_{\text{imperfect}} + r' |11\rangle\langle 11|, \quad \text{with} \quad \rho'_{\text{imperfect}} = (1 - p'_n) |\psi\rangle\langle\psi| + \frac{p'_n}{3} (|\phi\rangle\langle\phi| + |00\rangle\langle 00| + |11\rangle\langle 11|), \quad (\text{F3})
\end{aligned}$$

where the last line introduces

$$r' = \frac{1}{2 + 4d} \quad \text{and} \quad p'_n = \frac{p_n + 3d}{1 + 4d}. \quad (\text{F4})$$

We now present a more general analysis of dark counts for arbitrary r , again demonstrating that they serve to additively increase the effective network error p_n . To accomplish this, we first note that ρ_{raw} (defined, as in the main text, as the state heralded by a single click, with dark counts assumed impossible) is always diagonal in the basis $\{|\psi\rangle, |\phi\rangle, |00\rangle, |11\rangle\}$, where $|\phi\rangle$ is the antisymmetric state; see Eq. (F2). If a nonzero dark-count probability p_{dc} is taken into account, then the state of the system (given it has been *postselected* due to a single detector clicking as required) will be altered due to three additional ways in which that single click can be produced. The first is that the state $|00\rangle$ may survive postselection due to a single dark count occurring, and this occurs with an absolute probability $p_{00} = 2p_{\text{dc}}(1 - p_{\text{dc}})(1 - p_1)^2$. The second way in which a dark count can lead to an effect on the postselected state is that both a single-photon loss and a single dark count can occur, resulting in the state $\frac{1}{2}(|\psi\rangle\langle\psi| + |\phi\rangle\langle\phi|)$ with absolute probability

$p_{\psi+\phi} = 4p_{\text{dc}}(1 - p_{\text{dc}})p_{\text{loss}}p_1(1 - p_1)$. Finally, a combination of both the loss of two photons and a single dark count can lead to the erroneous inclusion of the $|11\rangle$ state after postselection, which occurs with absolute probability $p_{11} = 2p_{\text{dc}}(1 - p_{\text{dc}})p_{\text{loss}}^2p_1^2$. As we wish to examine the regime where p_{dc} is comparable to or less than the probability of photon loss not occurring, it is convenient to introduce the constant $d = p_{\text{dc}}/1 - p_{\text{loss}}$. In the regime of high loss, where $p_{\text{loss}} \rightarrow 1$, the probabilities become

$$p_{00} = 2d(1 - p_{\text{loss}})(1 - p_{\text{dc}})(1 - r)^2, \quad (\text{F5})$$

$$p_{\psi+\phi} = 4d(1 - p_{\text{loss}})(1 - p_{\text{dc}})r(1 - r), \quad (\text{F6})$$

$$p_{11} = 2d(1 - p_{\text{loss}})(1 - p_{\text{dc}})r^2. \quad (\text{F7})$$

Note that in the absence of dark counts the state after postselection will be ρ_{raw} , which occurs with probability $p_{\text{raw}} = 2(1 - p_{\text{dc}})^2(1 - p_{\text{loss}})p_1(1 - p_1 + p_{\text{loss}}p_1)$, which in the high-loss regime becomes $p_{\text{raw}} = 2(1 - p_{\text{dc}})^2(1 - p_{\text{loss}})r$. Thus, the state of the system after dark counts are included is given by

$$\rho_{\text{dc}} = \frac{p_{\text{raw}}\rho_{\text{raw}} + p_{00}|00\rangle\langle 00| + p_{11}|11\rangle\langle 11| + \frac{p_{\psi+\phi}}{2}|\psi\rangle\langle\psi| + \frac{p_{\psi+\phi}}{2}|\phi\rangle\langle\phi|}{p_{\text{raw}} + p_{00} + p_{11} + p_{\psi+\phi}} \quad (\text{F8})$$

$$= \frac{(1 - p_{\text{dc}})r\rho_{\text{raw}} + d(1 - r)^2|00\rangle\langle 00| + dr^2|11\rangle\langle 11| + dr(1 - r)|\psi\rangle\langle\psi| + dr(1 - r)|\phi\rangle\langle\phi|}{d + (1 - p_{\text{dc}})r}. \quad (\text{F9})$$

As the dark-count rate in many of the current generation of experiments is already very low, we can consider this expression in the case of small p_{dc} , in which case

$$\rho_{\text{dc}} \approx \frac{r\rho_{\text{raw}} + d(1 - r)^2|00\rangle\langle 00| + dr^2|11\rangle\langle 11| + dr(1 - r)|\psi\rangle\langle\psi| + dr(1 - r)|\phi\rangle\langle\phi|}{d + r} \quad (\text{F10})$$

$$= \frac{(1 - r)[d(1 - r) + r\frac{p_n}{3}]}{d + r}|00\rangle\langle 00| + \frac{r[(d + 1)r + (1 - r)\frac{p_n}{3}]}{d + r}|11\rangle\langle 11| \quad (\text{F11})$$

$$+ \frac{r(1 - r)(d + 1 - p_n)}{d + r}|\psi\rangle\langle\psi| + \frac{r(1 - r)(d + \frac{p_n}{3})}{d + r}|\phi\rangle\langle\phi|. \quad (\text{F12})$$

When phase drift between the creation of this step and the application of the EPL pair distillation step is taken into account, the state of the system is given by $\rho'_{dc} = (1 - p_{\text{drift}})\rho_{dc} + p_{\text{drift}}Z\rho_{dc}Z$. Hence, we have

$$\rho'_{dc} \approx \frac{(1-r)[d(1-r) + r\frac{p_n}{3}]}{d+r} |00\rangle\langle 00| + \frac{r[(d+1)r + (1-r)\frac{p_n}{3}]}{d+r} |11\rangle\langle 11| \quad (\text{F13})$$

$$+ \frac{r(1-r)[d+1 - p_n - p_{\text{drift}}(1 - \frac{4}{3}p_n)]}{d+r} |\psi\rangle\langle\psi| \quad (\text{F14})$$

$$+ \frac{r(1-r)[d + \frac{p_n}{3} + p_{\text{drift}}(1 - \frac{4}{3}p_n)]}{d+r} |\phi\rangle\langle\phi|. \quad (\text{F15})$$

Note that the application of the EPL protocol conditioned on a (1,1) outcome is not sufficient to ensure that the output state is in subspace spanned by $|\psi\rangle$ and $|\phi\rangle$, since this outcome also occurs when one pair is in the state $|00\rangle$ and the other state is $|11\rangle$. Note, however, that the (1,1) outcome cannot occur when both pairs are in state $|00\rangle$ or both pairs are in state $|11\rangle$. Thus, the pair produced by an application of the EPL protocol to two noisy pairs is

$$\rho_{\text{EPL}} \approx f(r, d, p_n)(|00\rangle\langle 00| + |11\rangle\langle 11|) + g(r, d, p_n, p_{\text{drift}})|\psi\rangle\langle\psi| + h(r, d, p_n, p_{\text{drift}})|\phi\rangle\langle\phi|, \quad (\text{F16})$$

where

$$f(r, d, p_n) = \frac{p_{\text{EPL}}^{-1}(1-r)(d(1-r) + r\frac{p_n}{3})r[(d+1)r + (1-r)\frac{p_n}{3}]}{(d+r)^2} (|00\rangle\langle 00| + |11\rangle\langle 11|), \quad (\text{F17})$$

$$g(r, d, p_n, p_{\text{drift}}) = \frac{p_{\text{EPL}}^{-1}r^2(1-r)^2\{[d+1 - p_n - p_{\text{drift}}(1 - \frac{4}{3}p_n)]^2 + [d + \frac{p_n}{3} + p_{\text{drift}}(1 - \frac{4}{3}p_n)]^2\}}{2(d+r)^2} |\psi\rangle\langle\psi|, \quad (\text{F18})$$

$$h(r, d, p_n, p_{\text{drift}}) = \frac{p_{\text{EPL}}^{-1}r^2(1-r)^2[d+1 - p_n - p_{\text{drift}}(1 - \frac{4}{3}p_n)][d + \frac{p_n}{3} + p_{\text{drift}}(1 - \frac{4}{3}p_n)]}{(d+r)^2} |\phi\rangle\langle\phi|, \quad (\text{F19})$$

and

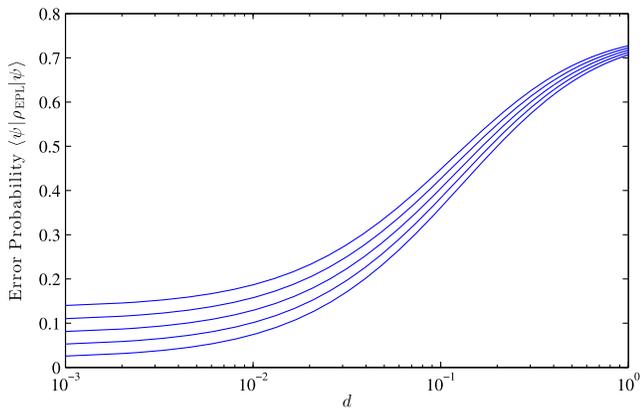


FIG. 13. Total error probability. Here, we take $r = \frac{1}{4}$ and $p_{\text{drift}} = 0.01$. The five lines represent the total error probability corresponding, from bottom to top, to preparation error probabilities $p_n = \{0, 0.025, 0.05, 0.075, 0.1\}$. The error probability is close to the error rate without dark counts while d is below approximately 0.01 but increases rapidly thereafter, passing through the distillation threshold even for the case $p_n = 0$.

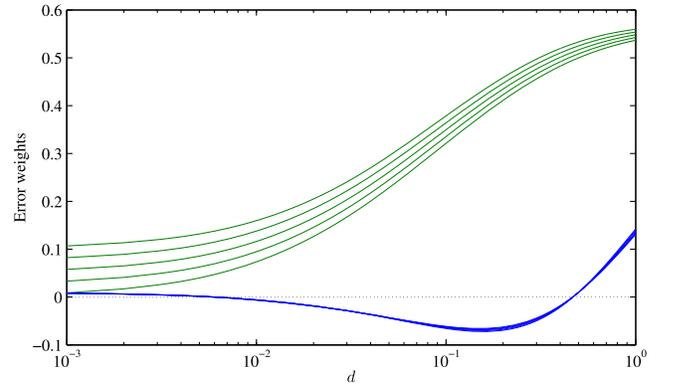


FIG. 14. Influence of dark counts on effective preparation and drift error probabilities for $r = \frac{1}{4}$ and $p_{\text{drift}} = 0.01$. The five upper lines (green) represent the effective preparation error weight (p'_n) corresponding, from bottom to top, to preparation error probabilities $p_n = \{0, 0.025, 0.05, 0.075, 0.1\}$. The lower lines (blue), indistinguishable at this scale, correspond to the effective drift error weight (p_{drift}) for the same values of p_n . In the regions $d \leq 0.007$ and $d \geq 0.469$, both p'_n and p_{drift} correspond to valid probabilities, and hence, in these regimes dark counts are indistinguishable from other noise sources.

$$p_{\text{EPL}} = \frac{r(1-r)}{2(d+r)^2} \left\{ 4 \left[d(1-r) + r \frac{p_n}{3} \right] \left[(d+1)r + (1-r) \frac{p_n}{3} \right] + r(1-r) \left(2d + 1 - \frac{2p_n}{3} \right)^2 \right\} \quad (\text{F20})$$

is the probability of obtaining the (1,1) result during the EPL protocol. The total error probability is then $\epsilon = 1 - \langle \psi | \rho_{\text{EPL}} | \psi \rangle = 1 - g(r, d, p_n, p_{\text{drift}})$. As can be seen from Fig. 13, dark counts begin to contribute significantly once d exceeds about 0.01.

It is natural to ask whether the errors present in this state due to dark counts are fundamentally different from those due to preparation and drift errors. However, by comparing ρ_{EPL} to the case of no dark counts, it is possible to find modified preparation and drift error weights (p'_n and p'_{drift}) such that the two states match. As f is independent of p_{drift} , the value of p'_n can be obtained by solving $f(r, d, p_n) = f(r, 0, p_n)$. The value for p'_{drift} can then be obtained by solving $g(r, d, p_n, p_{\text{drift}}) = g(r, d, p'_n, p'_{\text{drift}})$. Although p'_n and p'_{drift} are not guaranteed to correspond to valid probabilities, for many experimentally relevant parameter ranges they do indeed take on values between zero and one, and hence, the effect of dark counts within these parameter ranges are indistinguishable from drift and preparation errors. As can be seen from Fig. 14, as long as the value of d is kept far below 1 (in this case around 0.01), the effective modification of the drift and preparation error rates due to dark counts is relatively small. However, as expected, when d approaches unity the effective error rate rapidly increases.

-
- [1] 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).
 - [2] C. J. Ballance, T. P. Harty, N. M. Linke, and D. M. Lucas, *High-Fidelity Two-Qubit Quantum Logic Gates Using Trapped Calcium-43 Ions*, [arXiv:1406.5473](https://arxiv.org/abs/1406.5473).
 - [3] R. Barends *et al.*, *Logic Gates at the Surface Code Threshold: Superconducting Qubits Poised for Fault-Tolerant Quantum Computing*, [arXiv:1402.4848](https://arxiv.org/abs/1402.4848).
 - [4] F. Dolde, V. Bergholm, Y. Wang, I. Jakobi, B. Naydenov, S. Pezzagna, J. Meijer, F. Jelezko, P. Neumann, T. Schulte-Herbrüggen, J. Biamonte, and J. Wrachtrup, *High-Fidelity Spin Entanglement Using Optimal Control*, *Nat. Commun.* **5**, 3371 (2014).
 - [5] T. H. Taminiau, J. Cramer, T. van der Sar, V. V. Dobrovitski, and R. Hanson, *Universal Control and Error Correction in Multi-Qubit Spin Registers in Diamond*, *Nat. Nanotechnol.* **9**, 171 (2014).
 - [6] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim, *Large Scale Modular Quantum Computer Architecture with Atomic Memory and Photonic Interconnects*, *Phys. Rev. A* **89**, 022317 (2014).

- [7] D. Hucul, I. V. Inlek, G. Vittorini, C. Crocker, S. Debnath, S. M. Clark, and C. I. Monroe, *Modular Entanglement of Atomic Qubits Using Both Photons and Phonons*, [arXiv:1403.3696](https://arxiv.org/abs/1403.3696).
- [8] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, *Demonstration of Two-Qubit Algorithms with a Superconducting Quantum Processor*, *Nature (London)* **460**, 240 (2009).
- [9] N. Roch, M. E. Schwartz, F. Motzoi, C. Macklin, R. Vijay, A. W. Eddins, A. N. Korotkov, K. B. Whaley, M. Sarovar, and I. Siddiqi, *Observation of Measurement-Induced Entanglement and Quantum Trajectories of Remote Superconducting Qubits*, *Phys. Rev. Lett.* **112**, 170501 (2014).
- [10] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, *Heralded Entanglement between Solid-State Qubits Separated by 3 Meters*, *Nature (London)* **497**, 86 (2013).
- [11] W. Pfaff *et al.*, *Unconditional Quantum Teleportation between Distant Solid-State Quantum Bits*, *Science* **345**, 532 (2014).
- [12] D. Yang, *A Simple Proof of Monogamy of Entanglement*, *Phys. Lett. A* **360**, 249 (2006).
- [13] W. Dür and H.-J. Briegel, *Entanglement Purification for Quantum Computation*, *Phys. Rev. Lett.* **90**, 067901 (2003).
- [14] E. T. Campbell, *Distributed Quantum-Information Processing with Minimal Local Resources*, *Phys. Rev. A* **76**, 040302 (2007).
- [15] K. Fujii and K. Yamamoto, *Entanglement Purification with Double Selection*, *Phys. Rev. A* **80**, 042308 (2009).
- [16] K. Fujii, T. Yamamoto, M. Koashi, and N. Imoto, *A Distributed Architecture for Scalable Quantum Computation with Realistically Noisy Devices*, [arXiv:1202.6588v1](https://arxiv.org/abs/1202.6588v1).
- [17] M. Zwerger, H. J. Briegel, and W. Dür, *Hybrid Architecture for Encoded Measurement-Based Quantum Computation*, *Sci. Rep.* **4**, 5364 (2014).
- [18] N. H. Nickerson, Y. Li, and S. C. Benjamin, *Topological Quantum Computing with a Very Noisy Network and Error Rates Approaching One Percent*, *Nat. Commun.* **4**, 1756 (2013).
- [19] A. Kitaev, in *Proceedings of the Third International Conference on Quantum Communication and Measurement* (Plenum Press, New York, 1997).
- [20] A. Y. Kitaev, *Fault-Tolerant Quantum Computation by Anyons*, *Ann. Phys. (Amsterdam)* **303**, 2 (2003).
- [21] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, *Topological Quantum Memory*, *J. Math. Phys. (N.Y.)* **43**, 4452 (2002).
- [22] C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, *Surface Code Quantum Computing by Lattice Surgery*, *New J. Phys.* **14**, 123011 (2012).
- [23] S. Bravyi and A. Kitaev, *Universal Quantum Computation with Ideal Clifford Gates and Noisy Ancillas*, *Phys. Rev. A* **71**, 022316 (2005).

- [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] D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg, *Quantum Computing with Nearest Neighbor Interactions and Error Rates over 1%*, *Phys. Rev. A* **83**, 020302(R) (2011).
- [26] E. T. Campbell and S. C. Benjamin, *Measurement-Based Entanglement under Conditions of Extreme Photon Loss*, *Phys. Rev. Lett.* **101**, 130502 (2008).
- [27] S. D. Barrett and P. Kok, *Efficient High-Fidelity Quantum Computation Using Matter Qubits and Linear Optics*, *Phys. Rev. A* **71**, 060310 (2005).
- [28] A. Reiserer, N. Kalb, G. Rempe, and S. Ritter, *A Quantum Gate between a Flying Optical Photon and a Single Trapped Atom*, *Nature (London)* **508**, 237 (2014).
- [29] M. M. Steiner, H. M. Meyer, J. Reichel, and M. Köhl, *Photon Emission and Absorption of a Single Ion coupled to an Optical Fiber-Cavity*, [arXiv:1407.6036](https://arxiv.org/abs/1407.6036).
- [30] S. C. Benjamin, D. E. Browne, J. Fitzsimons, and J. J. L. Morton, *Brokered Graph-State Quantum Computation*, *New J. Phys.* **8**, 141 (2006).
- [31] S. Bravyi and A. Vargo, *Simulation of Rare Events in Quantum Error Correction*, *Phys. Rev. A* **88**, 062308 (2013).
- [32] F. H. E. Watson and S. D. Barrett, *Logical Error Rate Scaling of the Toric Code*, [arXiv:1312.5213v2](https://arxiv.org/abs/1312.5213v2).
- [33] A. Jamiolkowski, *Linear Transformations which Preserve Trace and Positive Semidefiniteness of Operators*, *Rep. Math. Phys.* **3**, 275 (1972).
- [34] J. Edmonds, *Paths, Trees, and Flowers*, *Can. J. Math.* **17**, 449 (1965).
- [35] V. Kolmogorov, *Blossom V: A New Implementation of a Minimum Cost Perfect Matching Algorithm*, *Math. Program. Comput.* **1**, 43 (2009).
- [36] D. S. Wang, A. G. Fowler, A. M. Stephens, and L. C. L. Hollenberg, *Threshold Error Rates for the Toric and Planar Codes*, *Quantum Inf. Comput.* **10**, 456 (2010).