



PAPER • OPEN ACCESS

Minimally complex ion traps as modules for quantum communication and computing

To cite this article: Ramil Nigmatullin *et al* 2016 *New J. Phys.* **18** 103028

View the [article online](#) for updates and enhancements.

You may also like

- [The focusing effect of electron flow and negative refraction in three-dimensional topological insulators](#)
Kai-Tong Wang, Yanxia Xing, King Tai Cheung et al.
- [Ising exchange interaction in lanthanides and actinides](#)
Livia F Chibotaru and Naoya Iwahara
- [Electron-spin dynamics induced by photon spins](#)
Heiko Bauke, Sven Ahrens, Christoph H Keitel et al.



PAPER

OPEN ACCESS

RECEIVED
2 June 2016REVISED
17 September 2016ACCEPTED FOR PUBLICATION
21 September 2016PUBLISHED
19 October 2016

Original content from this
work may be used under
the terms of the [Creative
Commons Attribution 3.0
licence](#).

Any further distribution of
this work must maintain
attribution to the
author(s) and the title of
the work, journal citation
and DOI.



Minimally complex ion traps as modules for quantum communication and computing

Ramil Nigmatullin¹, Christopher J Ballance², Niel de Beaudrap³ and Simon C Benjamin^{1,4}¹ Department of Materials, University of Oxford, Parks Road, Oxford, UK² Department of Physics, University of Oxford, Parks Road, Oxford, UK³ Department of Computer Science, University of Oxford, Parks Road, Oxford, UK⁴ Author to whom any correspondence should be addressed.E-mail: simon.benjamin@materials.ox.ac.uk**Keywords:** quantum technologies, quantum networks, ion traps

Abstract

Optically linked ion traps are promising as components of network-based quantum technologies, including communication systems and modular computers. Experimental results achieved to date indicate that the fidelity of operations within each ion trap module will be far higher than the fidelity of operations involving the links; fortunately internal storage and processing can effectively upgrade the links through the process of purification. Here we perform the most detailed analysis to date on this purification task, using a protocol which is balanced to maximise fidelity while minimising the device complexity and the time cost of the process. Moreover we ‘compile down’ the quantum circuit to device-level operations including cooling and shuttling events. We find that a linear trap with only five ions (two of one species, three of another) can support our protocol while incorporating desirable features such as *global control*, i.e. laser control pulses need only target an entire zone rather than differentiating one ion from its neighbour. To evaluate the capabilities of such a module we consider its use both as a universal communications node for quantum key distribution, and as the basic repeating unit of a quantum computer. For the latter case we evaluate the threshold for fault tolerant quantum computing using the surface code, finding acceptable fidelities for the ‘raw’ entangling link as low as 83% (or under 75% if an additional ion is available).

1. Introduction

In order to realise the promise of quantum technologies it is highly desirable to create modular units that can interlink photonically, each having an internal storage and processing capacity. Long range quantum communication networks will require repeaters to overcome photon loss and accumulated noise [1]; generally a small quantum processor with photonic outputs can be seen as a universal communications node suitable for supporting any network-based task. Meanwhile, for quantum computing the modular approach could be used to build a large-scale machine on a single site [2–4]. Here modules may be called ‘remote’ but they might be separated by centimetres or less. Several technologies have shown the in-principle capability to serve as photonically interlinked modular cells, including ion traps [5], nitrogen-vacancy centres in diamond [6, 7], and superconducting qubits [8, 9]. Here our sole focus is on strategies for exploiting ion traps. Nevertheless we hope that elements of our study, in particular the compact three-tier purification circuits that we analyse, may also be useful in diamond or superconducting approaches. All systems of this general kind will have the same desiderata: high purifying power with low time cost and low system complexity.

Trapped ion systems are one of the most mature quantum technologies. A variety of trap devices now exist, but in all cases electromagnetic fields are configured to spatially confine ions. In high vacuum and under the action of laser cooling, the ions organise into Coulomb crystals. The electronic and spin states of the ions can be manipulated using optical and microwave techniques, and thus each ion can embody a qubit once a suitable pair levels are identified. Coupling between the ions due to the vibrational mode of the Coulomb crystal implies the

possibility of controlled manipulation of quantum states of two-or-more ions—multi-qubit gate operations. High fidelity quantum operations in ion traps have been demonstrated by a number of groups worldwide [10–13]. Proof of principle experiments have demonstrated several quantum algorithms in single crystal devices; ranging from Shor’s algorithm [14] to simulation of quantum Ising spin chains [15, 16].

Spectral crowding means that the larger the ion crystal the more difficult it is to coherently control the quantum dynamics of individual ions. For this reason, it is advantageous to adopt a modular design: decomposition of a device into a large number of interconnected ion traps. The links between the traps can be implemented either by shuttling of the ions between the traps [17] or with a hybrid system, for example, with photonic interfaces. The advantage of the hybrid ion/photon approach is that it does not involve the design and manufacture of large traps with complex electrode geometries. The disadvantage is that, at the moment, the ion/photon interface is significantly noisier and far slower than the operations within an isolated ion crystal [5]. Upcoming iterations of the technology should increase the speed of the interface (as we presently discuss), however it may be that the inherent fidelity of remotely generated entanglement will never match that of local operations. Fortunately the process of *entanglement purification* can be employed to close this gap. A substantial literature on this topic has been established over the last two decades. While we will not attempt a review here, we note that the seminal paper introducing the concept of purification was that of Bennett *et al* in 1996 [18], while the key challenge of performing compact tiered purification was addressed by Briegel and Dür in 2003 [19]. The present work derives its foundations from these papers and the broad literature which they exemplify. In all cases, the use of purification inevitably adds additional resource cost [4, 20]. Surprisingly it has been shown that the resource cost associated with the adoption of the flexible network architecture may vary little over a wide range of module sizes [21].

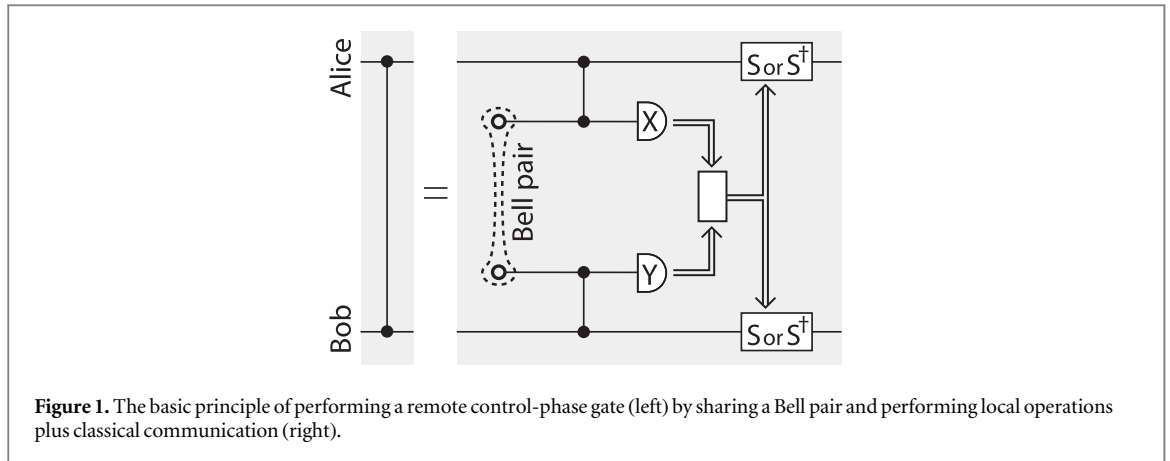
The focus of this paper is to design the simplest possible ion trap modules, suitable for photonically linking into a scalable communications network or a single-site modular quantum computer [22]. In our analysis, we assume gate fidelities that are already accessible in state of the art experiments. In section 2 we note the general requirements for photonically linking modules. In section 3, we list our design priorities and describe the extent to which they will be met. Section 4 then provides a systematic construction of the purification protocol and shows that using three ions can reduce the infidelity of a Bell pair from ϵ to $\frac{2}{9}\epsilon^2 + O(\epsilon^3)$. A suitable device structure and a device-level specification is given in section 5. In section 6, we numerically evaluate the performance of the node—the final fidelity and the running time of the protocol as a function of network fidelity. In sections 7 and 8, we indicate the performance in two practical applications: communication and fault tolerant computing. For the latter, we evaluate the fault tolerance threshold for a toric code. Sections 9 and 10 provide discussions of the practicality of optical wiring and of the expected overall performance, in both near-term and further future realisations. Finally we offer some conclusions in section 11.

2. Interlinking ion traps via photons

In any modular quantum technology it will be necessary to achieve entanglement that spans the modules. For computing applications, we will wish to be able to perform quantum gates between qubits in remote locations. An elegant and practical route to achieving this is to create a shared Bell pair between the two modules (through a process that might be probabilistic, provided that success is heralded) and then consume this Bell pair in order to implement the gate. Suppose that Alice and Bob each have a module, and within each module is a single ‘application qubit’, i.e. a qubit that is part of the overall task that Alice and Bob are performing, be it communication or computing. Suppose that they also share a Bell pair $|\Psi^+\rangle \equiv (|00\rangle + |11\rangle)/\sqrt{2}$.

If Alice and Bob would like to perform a control-phase (cPhase) gate operation between their application qubits, then they can do by the process shown in figure 1 which involves three steps: (1) they each perform local cPhase operations between their application qubit and their ‘half’ of the Bell pair, (2) Alice measures her Bell qubit in the x -basis while Bob measures his in the y -basis, and finally (3) they each apply a single qubit gate to their application qubit. The required gate depends on the measurement outcomes at the second stage; if their measurements were the same (both measured in the positive direction on their apparatus, or both negative) then Alice and Bob should both apply single qubit gate $S \equiv \text{diag}\{1, i\}$, otherwise they should apply S^\dagger . This is a simple instance of gate teleportation and the process can be verified in a few lines, see appendix 1. Adopting this approach to remote gate operations, the challenge of realising a modular quantum computer becomes ‘How can we create *high fidelity* shared Bell pairs?’.

The demand for high fidelity is crucial since any imperfection on the Bell pair will translate to noise on the remote gate operation. It is a reasonable presumption that the entanglement channel creates ‘raw’ Bell pairs with a fidelity far below that of the local gates. In ion trap experiments, all local gates have been demonstrated with fidelities of 99.9% or higher [11, 13], while entanglement between traps has been achieved at the level of about 85% [5]. While we can expect these numbers to continually improve, it may be that local operations are always of



superior fidelity to the entangling channel; whenever this is the case, we may wish to perform purification in order to effectively upgrade the channel fidelity to a level comparable to the local gates.

In the following, we will refer to the ancillas which represent the purified Bell pairs as the ‘envoy’ qubits; their role in mediating the link between modules is ‘high status’ in the sense that they embody the superior, purified entanglement and they interact directly with the crucial ‘application qubits’. Below the envoys are other ancilla qubits which process the lower grade entanglement, or create the ‘raw’ entanglement between modules. However these lower status ancillas will never interact directly with the application qubit. Thus, an ion trap module will contain several ions with different designated functions: raw entanglement ion, purification ions, the envoy ion and the application ion.

We will not discuss the means by which raw entanglement is created, except that we assume it is done optically and that the fidelity achieved is relatively poor. The details of the process are of course important since they determine the nature of the infidelity on the raw pair. Here we will assume a noise model where all imperfections are equally likely. In reality a given method for Bell state generation (e.g. [5]) will have a unique noise spectrum; generally however, structure in the noise will make it easier to purify, and therefore our assumption of structureless noise means that the performance metrics we predict tend to the conservative end.

3. Design considerations

Following the reasoning in the previous section, we now proceed to design the minimally complex ion trap that suffices as a module for communications or computing over an imperfect network. We assume that each trap contains *only one* ‘application qubit’ and that all other ions exist only in order that the sole application qubit can perform high fidelity gates with partners in other traps. Generalisations to variants with two or more application qubits per trap are straightforward, but by focusing on this minimal device we can address the question ‘What is the simplest ion trap that can suffice as a module of a quantum technology?’.

The following features are desirable for a practical module:

1. A high level of purification should be achieved (at least an order of magnitude reduction to infidelity).
2. The time cost of the purification should be modest.
3. The trap geometry should be simple, ideally linear.
4. The trap should have as few zones as possible.
5. At most two ion species should be used (ideally one).
6. Shutting/permuring of ions should be minimised.
7. Two-qubit gates are preferable to higher order gates.
8. Ideally, two-qubit gates take place in small crystals.
9. The fewest possible measurement/control systems (lasers, lenses, detectors etc) should be required.
10. The issue of cross talk, e.g. unwanted interaction of a laser or emitted photons with another ion, should be minimised by design.
11. The need for sympathetic cooling must be allowed for.

These desiderata are largely self-explanatory. The need for a purifying factor of at least ten follows from the fact that we can expect the ‘raw’ entanglement fidelity to be at least ten times worse than the local gate fidelity (in present experiments it is two orders worse). The need to minimise the time cost follows from the fact that it will be challenging to create entanglement before significant decoherence has occurred to the application qubit. By ‘zone’ in point (4) we refer to a region of the trap that is significantly remote from other regions, effectively forming a sub-trap; one zone may have several electrodes to define it and move ions. Point (7) is motivated by the observation that experiments to date have reported lower fidelity as the number of ions involved in a gate increases. Point (8) is advantageous given that, while two-qubit gates can be directly performed between non-adjacent ions [23], the fidelity is expected to decrease with increasing number of ions and the approach may necessitate additional hardware complexity.

We find that by permitting ourselves two species of ion, the other desiderata can be satisfied to a remarkable degree. For the variant that we analyse in greatest detail,

1. 10% raw infidelity is purified to 0.6% infidelity.
2. Average time cost is a factor of ~ 8 .
3. A linear geometry does suffice.
4. Only two zones are required.
5. Two species are employed, e.g. Ca and Sr.
6. Only one ion performs any shuttling, and ions need never be permuted*.
7. Only one-qubit and two-qubit gates are employed.
8. All two-qubit gates are on nearest-neighbours.
9. Only a single instance of each control/measurement system is required for each entire trap device.
10. By adopting a *global control* principle, laser cross talk is negligible; laser beams need not be tightly focused.
11. Cooling is efficiently integrated via a dual-role ion.

The asterisk in point (6) is present because for certain functions, such as fault tolerant surface code computing, it may be desirable to periodically exchange the roles of application qubit and the envoy. An efficient way to do this would be to physically permute the two ions so that they exchange places; however if this is not possible, a logical SWAP operation will suffice instead. The reason that this exchange may be desirable is explained presently when we appraise the module’s performance for fault tolerant computing.

Generally the desired features are quite compatible with one another, and in particular it is quite natural to support (7) and (8) because of the tiered nature of the purification process. The use of two zones proves to be very valuable in meeting the other desiderata especially point (10). In the following we will use the label ‘rowdy end’ for the zone in which raw entanglement is created and all measurements are performed. The term ‘tranquil end’ will be used for the zone where the application qubit resides. The envoy qubit will shuttle between the two zones, ultimately delivering purified entanglement to the application qubit in order to perform the remote gate. Thus at any one time each module would contain only two crystals with 2–3 ions.

We emphasize that the desiderata listed in this section are aimed at minimizing the complexity with the current available technology and as the technology improves the design criteria may become more flexible.

4. Purification protocol

The purpose of purification is to establish high-fidelity entanglement over noisy channels. This is achieved probabilistically by consuming several noisy Bell pairs to produce one Bell pair of higher fidelity. Entanglement purification has been demonstrated experimentally in photonic systems [24, 25] and internally in atomic systems [26]. To date there are no experimental demonstrations of purification in a system involving atomic species linked through a photonic channel. A number of previous theoretical works have proposed resource-efficient purification protocols for modular technologies, whether generically (as in, for example, [27, 28]) or for specific combinations of a system and an application, for example coupled spins in a communication setting [29]. The present study is, to the authors’ knowledge, the most detailed blueprint yet for purification-enabled ion trap technology.

Our purification protocol will use the following primitives: generation of raw entangled pairs, two qubit gates between adjacent qubits, and measurements. Presently we will ‘compile down’ the circuit to a set of device-level operations including cooling and shuttling operations.

The raw entangled state is assumed to be of the form of a depolarized Werner state

$$\rho_\epsilon = (1 - \epsilon)\Phi^+ + \frac{\epsilon}{3}\Phi^- + \frac{\epsilon}{3}\Psi^+ + \frac{\epsilon}{3}\Psi^-, \quad (1)$$

where $\epsilon \in [0, 0.5)$ and states Φ^+ , Φ^- , Ψ^+ and Ψ^- are the standard Bell states. Here we have chosen Φ^+ as the desired Bell state (obviously, it requires only a single-qubit rotation to transform any Bell state to another, so we are not limiting ourselves to any particular entanglement generation protocol by assuming Φ^+). The fidelity of ρ_ϵ is $(\text{Tr} \sqrt{\sqrt{\rho_\epsilon} \Phi^+ \sqrt{\rho_\epsilon}})^2 = 1 - \epsilon$. A state given by equation (1) is fully depolarized i.e. its errors are evenly distributed across the X , Y and Z error channels. As noted above, in the context of purification, a fully depolarized input state is a conservative assumption since structured noise can be beneficially exploited in purification.

The most basic purification circuit takes two pairs ρ_ϵ and produces a state $\tilde{\rho}_\epsilon$ of fidelity greater than $1 - \epsilon$. It consists of two cNOT gates and a parity measurement. This purification protocol is the first part (Level 1) of the circuit shown in figure 3. We will denote the map describing this purification protocol by $F : Q \otimes Q \rightarrow Q$, where Q is a set of two qubit density matrices. Measurements are in the standard basis and are postselected to have the same parity. The resulting state $\tilde{\rho}_\epsilon^{(1)}$ is given by

$$\tilde{\rho}_\epsilon^{(1)} \equiv F[\rho_\epsilon, \rho_\epsilon] = \left(1 - \frac{2}{3}\epsilon - \frac{2}{3}\epsilon^2\right)\Phi^+ + \left(\frac{2}{3}\epsilon + \frac{2}{9}\epsilon^2\right)\Phi^- + \frac{2}{9}\epsilon^2\Psi^+ + \frac{2}{9}\epsilon^2\Psi^- + O(\epsilon^3). \quad (2)$$

The fidelity of $\tilde{\rho}_\epsilon^{(1)}$ is $1 - \frac{2}{3}\epsilon + O(\epsilon^2)$. To achieve further improvements in fidelity, one constructs a tiered purification protocol, where a level consists of a single application of map F and the outputs from a given level may be inputs at a higher level. As we will see shortly, between each stage of the process it may be necessary to perform local rotations.

We now systematically construct such multilevel purification protocols. First, we introduce a shorthand notations where the state

$$\rho = (1 - r_1 - r_2 - r_3)\Phi^+ + r_1\Phi^- + r_2\Psi^+ + r_3\Psi^-, \quad (3)$$

is represented by a tuple

$$\rho \sim (r_1, r_2, r_3). \quad (4)$$

By direct computation one finds that (to lowest order) the effect of a single iteration of map F on states of general form (4) is

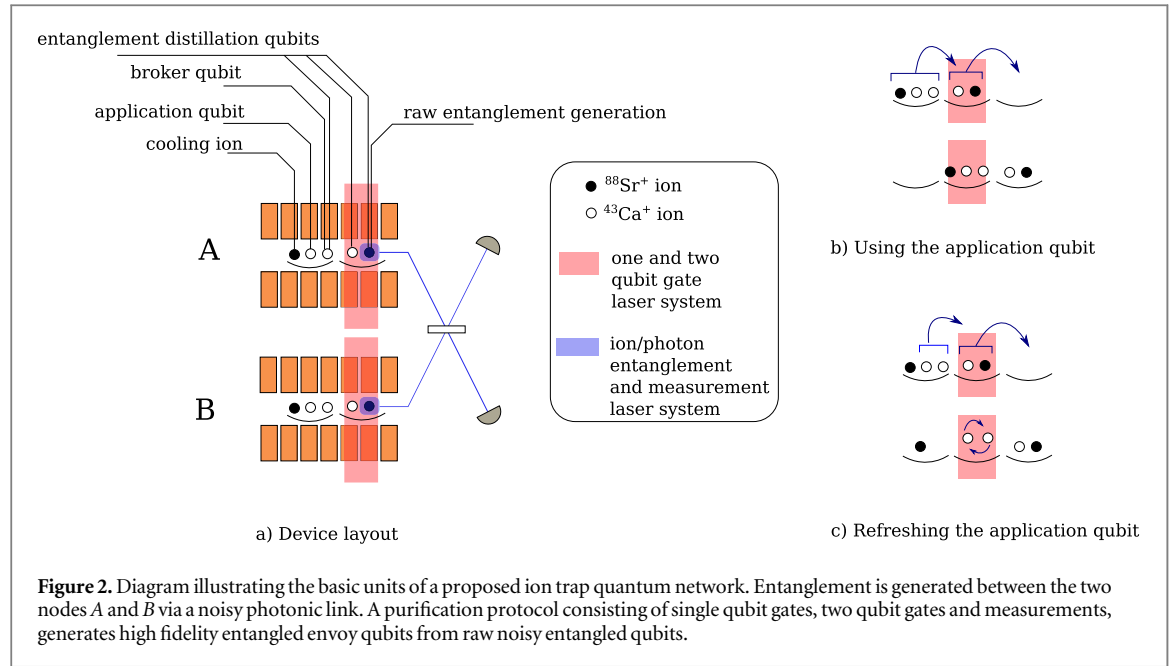
$$F[(r_1, r_2, r_3), (s_1, s_2, s_3)] \sim (r_1 + s_1, r_2s_2 + r_3s_3, r_2s_3 + r_3s_2). \quad (5)$$

Thus, for instance, applying equation (5) on a pair of states $\rho_\epsilon \sim (\frac{\epsilon}{3}, \frac{\epsilon}{3}, \frac{\epsilon}{3})$, produces a state $\tilde{\rho}_\epsilon^{(1)} \sim (\frac{2}{3}\epsilon, \frac{2}{9}\epsilon^2, \frac{2}{9}\epsilon^2)$ in agreement with equation (2). From equation (5) we can see that the map F suppresses the contributions from the Ψ^\pm modes but increases the contribution of the Φ^- channel. Further iterations of the map F will decrease the fidelity of the resulting state due to the concentration of noise in the Φ^- mode. In order to continue to improve the fidelity with successive applications of F one can permute the Φ^- , Ψ^+ and Ψ^- modes by applying local rotations. The three modes Φ^- , Ψ^+ and Ψ^- can be described in terms of single-qubit Pauli errors on the noise-free mode Φ^+

$$\begin{aligned} \Phi^- &= (I \otimes Z)\Phi^+(I \otimes Z), \\ \Psi^+ &= (I \otimes X)\Phi^+(I \otimes X), \\ \Psi^- &= (I \otimes Y)\Phi^+(I \otimes Y). \end{aligned}$$

These modes can be permuted using any single-qubit Clifford group operations which leave Φ^+ invariant. These operations form a representation of the dihedral group D_3 , and in particular form a group of order 6; they are generated by the operation $g_1 \equiv H \otimes H$ and the operation $g_2 = S^\dagger \otimes S$, where $S = \exp(-i\frac{\pi}{4}Z) \propto \text{diag}(1, i)$.

With this in mind, we can see that one can form a higher fidelity state $\tilde{\rho}^{(2)}$ by applying F on two states $\tilde{\rho}^{(1)} = F[\rho_\epsilon, \rho_\epsilon]$ if they are first locally rotated using g_1 .



$$\begin{aligned}
 \tilde{\rho}^{(2)} &= F[g_1 \tilde{\rho}^{(1)} g_1^\dagger, g_1 \tilde{\rho}^{(1)} g_1^\dagger] \\
 &\sim F\left[\left(\frac{2}{9}\epsilon^2, \frac{2}{3}\epsilon, \frac{2}{9}\epsilon^2\right), \left(\frac{2}{9}\epsilon^2, \frac{2}{3}\epsilon, \frac{2}{9}\epsilon^2\right)\right] \\
 &\sim \left(\frac{4}{9}\epsilon^2, \frac{4}{9}\epsilon^2, \frac{8}{27}\epsilon^3\right).
 \end{aligned} \tag{6}$$

The fidelity of this state is $1 - \frac{8}{9}\epsilon^2 + O(\epsilon^3)$. This purification map that produces state $\tilde{\rho}^{(2)}$ is represented by the Level 2 of the circuit shown in figure 3.

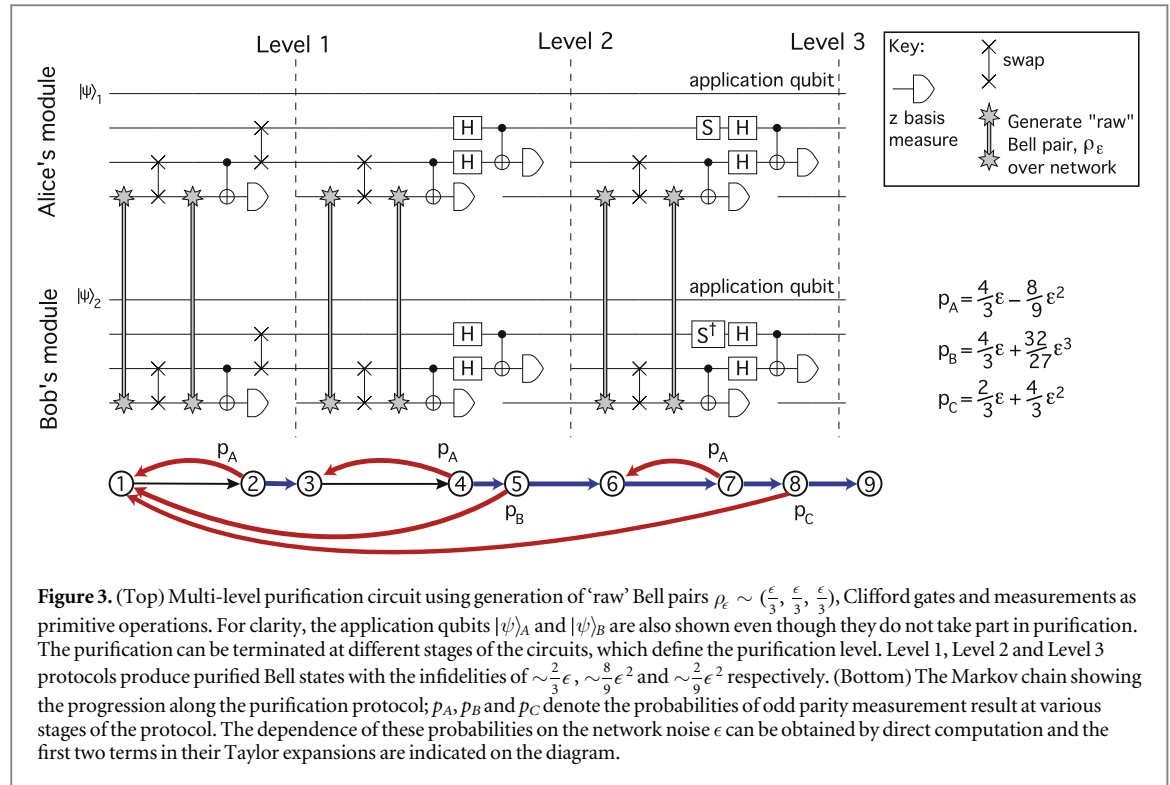
Finally, we construct a third level to our purification protocol that uses $\tilde{\rho}^{(2)}$ and $\tilde{\rho}^{(1)}$, with suitable local rotations, to produce a purified state $\tilde{\rho}^{(3)}$. The Level 3 purification protocol is represented by the whole circuit shown in figure 3. To the lowest orders in ϵ the state $\tilde{\rho}^{(3)}$ is given by

$$\begin{aligned}
 \tilde{\rho}^{(3)} &= F[g_1 g_2 \tilde{\rho}^{(2)} g_2^\dagger g_1^\dagger, g_1 \tilde{\rho}^{(1)} g_1^\dagger] \\
 &\sim F\left[g_1 g_2 \left(\frac{4}{9}\epsilon^2, \frac{4}{9}\epsilon^2, \frac{8}{27}\epsilon^3\right) g_2^\dagger g_1^\dagger, g_1 \left(\frac{2}{3}\epsilon, \frac{2}{9}\epsilon^2, \frac{2}{9}\epsilon^2\right) g_1^\dagger\right] \\
 &\sim F\left[g_1 \left(\frac{4}{9}\epsilon^2, \frac{8}{27}\epsilon^3, \frac{4}{9}\epsilon^2\right) g_1^\dagger, \left(\frac{2}{9}\epsilon^2, \frac{2}{3}\epsilon, \frac{2}{9}\epsilon^2\right)\right] \\
 &\sim F\left[\left(\frac{8}{27}\epsilon^3, \frac{4}{9}\epsilon^2, \frac{4}{9}\epsilon^2\right), \left(\frac{2}{9}\epsilon^2, \frac{2}{3}\epsilon, \frac{2}{9}\epsilon^2\right)\right] \\
 &\sim \left(\frac{2}{9}\epsilon^2, \frac{8}{27}\epsilon^3, \frac{8}{27}\epsilon^3\right).
 \end{aligned} \tag{7}$$

The complete Level 3 purification circuit uses three pairs of raw states ρ_ϵ of infidelity ϵ to produce a state $\tilde{\rho}^{(3)}$ of infidelity $\sim \frac{2}{9}\epsilon^2$. If we were to add to the resources another pair of raw states ρ_ϵ then it is possible to produce a state $\tilde{\rho}^{(4)}$ of infidelity $\sim \frac{16}{27}\epsilon^3$. However, in the current paper, we will not analyse maps that produce states of infidelity of order ϵ^3 and higher. We presently find that states produced by the Level 3 purification protocol are already of fidelity high enough for quantum communication and computing applications.

Note that each additional stage of purification requires generation of an additional two raw Bell pairs, while introducing a new *tier* to the process would double the total requirements. Since with current technology photonic entanglement is a slow operation, protocols significantly more complex than that described above would likely lead to unacceptable slowdown to the rate at which remote gates can be applied between the application qubits in separate modules.

The purification protocol is post-selective—if a measurement produces an odd parity result then the protocol has to be repeated from a particular point. Thus in practice the number of raw entangled pairs needed to complete a protocol is not fixed. The progression along the purification protocol can be represented using the Markov chain model [27, 28], which is shown below the circuit diagram in figure 2.



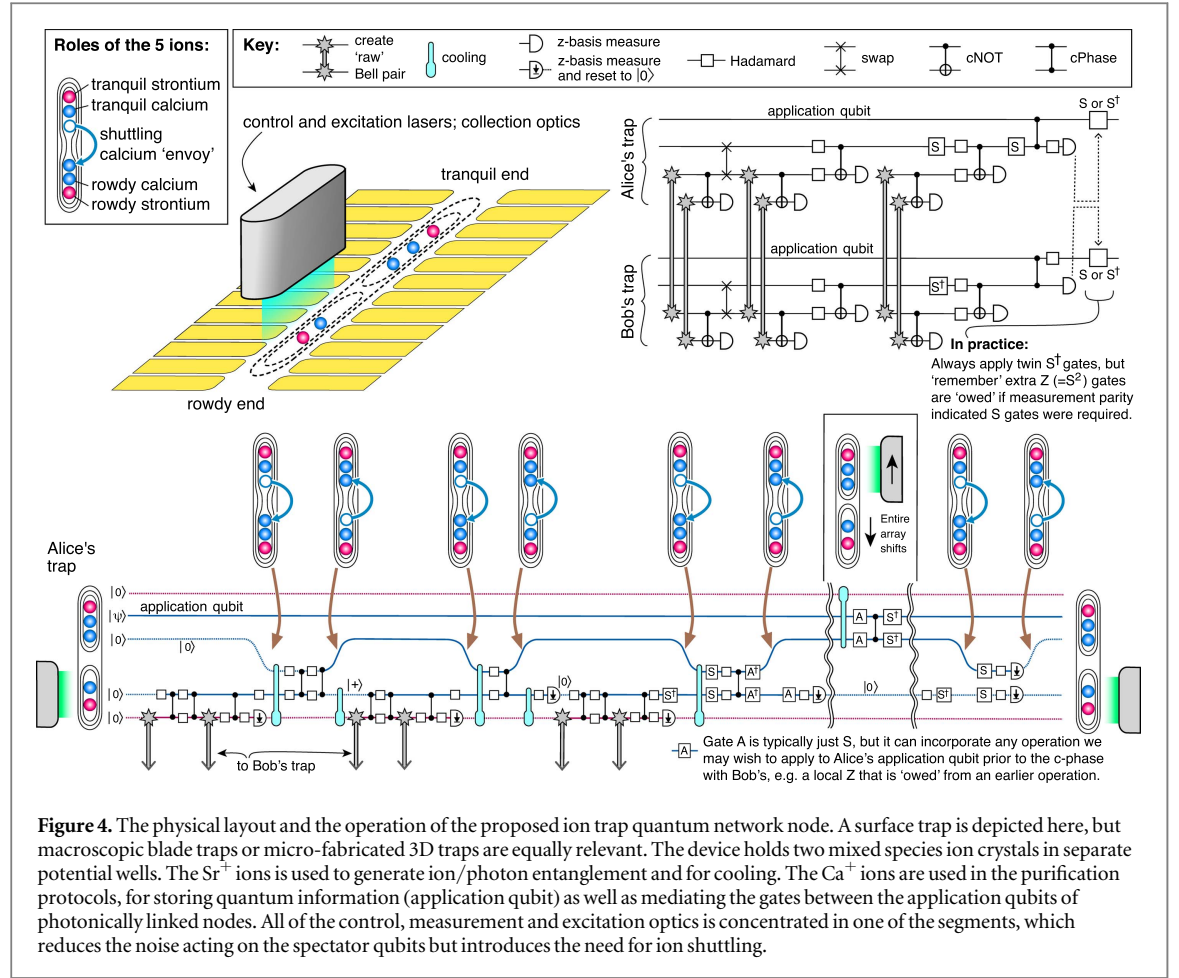
A few remarks about the optimality of the above purification circuits are in order. Our choice of the primitive 2-to-1 purification protocol is strongly motivated by the setting in which it is to be applied. As we prefer to perform nearest-neighbour operations (which themselves are noisy) within a linear array, it is important to minimise the number of operations to be performed to reduce the propagation of errors. Furthermore, as we consider traps with very few ions, it is important to limit the number of qubits to be stored simultaneously. Given these considerations, the simplicity of the 2-to-1 purification is advantageous. We then require only that the raw entangled states have little enough noise that each round of purification may succeed with high probability.

It remains to consider whether we may obtain further improvements in the noise reduction, given the same device complexity but using maps other than the particular arrangements of Hadamard gates and $\pi/4$ gates which we describe. Our choices of transformations are optimal over the set of Clifford gates: operations from the Clifford group only serve to permute the Pauli noise channels, and our operations are chosen to optimise the rate at which these noise channels are suppressed upon success. Any choice of non-Clifford gates will at best mix the Pauli noise channels prior to purification, and at worst introduce more noise channels which are not described by Pauli operators, reducing the (admittedly small) probability of cancelling the noise from different noisy entangled states, without increasing the probability of success in purification. This leads us to conclude that the approach we have taken is the best choice for purification in our setting.

5. Physical layout of the device

We now combine the designed purification circuit shown in figure 2 with the general ion trap consideration of section 2 to produce a detailed blueprint for the ion trap quantum network node and its operational steps. This is presented in figure 4.

The allowed primitive operations in the device are ion shuttling operations (splitting, joining and moving ion crystals in a linear array), raw photonic entanglement generation, local qubit rotations, the symmetric two-qubit phase gate cPhase, measurements, and crystal cooling. (Note that the basic two-qubit operation may not in practice be the cPhase, but rather e.g. a Mølmer–Sørensen gate; however by suitably replacing the adjacent pairs of Hadamard gates by another symmetric pair of single-qubit gates, they are made equivalent.) Figure 4 shows the suggested device layout with the explicit purification and application sequence. At any time the module contains two mixed-species ion crystals; there will be two ions in one potential well, and three in the other. In this paper we assume the species are $^{43}\text{Ca}^+$ and Sr^+ , however there are of course other suitable possibilities. The two Sr^+ ions are used for cooling and/or photonic entanglement generation. The three $^{43}\text{Ca}^+$ ions are used for purification, storing quantum information and mediating the gates between the application qubits of separate



modules. The control, excitation and collection optics are all focused on one trap region only, i.e. it targets only one of the ion crystals. Typically the targeted region is the rowdy end; when the time comes for the envoy qubit to entangle with the application qubit, the entire trap potential shifts (without any change to the relative positions of the ions) so that the laser control systems now target the tranquil end; once the gates there are implemented, the potentials shift back.

Note that in the fully compiled circuit it is never necessary to differentially target one ion over another of the same species in the same zone; in fact zones are under *global control* in the sense that control pulses target entire zones and, where there is more than one ion of the relevant species, both respond: we therefore restrict ourselves to symmetric gates cPhase gate and $G \otimes G$, where G is any single qubit rotation. This negates the issue of cross talk *within* a given zone, leaving us only concerned with the possibility of accidental excitation (by scattered laser light or emitted photons) of ions in the other zone; given that the zone separation could be of the order of a centimetre if need be, this source of error should be easily made negligible.

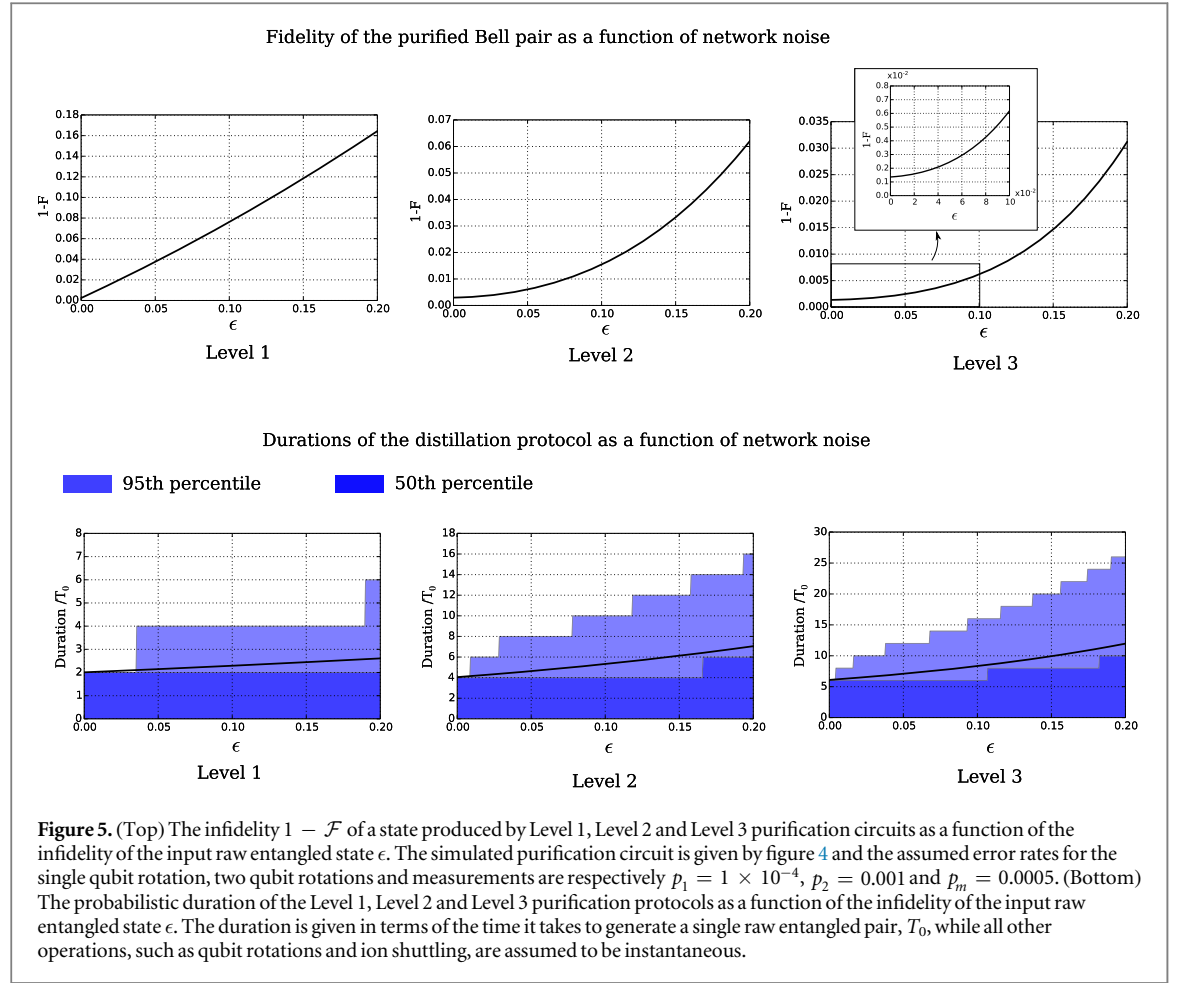
6. Performance of two connected modules

In this section, we numerically evaluate the performance of the designed purification protocols assuming realistic level of gate noise.

We model single-qubit and two-qubit gate noise by a depolarizing noise channel. In practice, rates of bit flip and phase noise can differ considerably, but using a depolarized channel is a good theoretical approach since it gives bounds on the expected fidelity of the purified Bell pairs—any asymmetry in the noise channel can be exploited favourably in the purification protocol. Single qubit noise is modelled by a perfect gate followed by a trace preserving noise process

$$\mathcal{N}_1(\rho) = (1 - p_1)\rho + \frac{p_1}{3}(X\rho X + Y\rho Y + Z\rho Z), \quad (8)$$

where X , Y and Z denote the Pauli matrices.



Two qubit noise is modelled by perfect gate followed by a noise process

$$\mathcal{N}_2(\rho) = (1 - p_2)\rho + \frac{p_2}{15} \sum_{A,B} (A \otimes B)\rho(A \otimes B)^\dagger, \quad (9)$$

where operator $A \in \{I, X, Y, Z\}$ acts on the first qubit, and similarly B acts on the second qubit but the case $I \otimes I$ is excluded from the sum. Note that a given experimental system will have noise that deviates from an even distribution of errors over all channels (see e.g. [11]), but by making this assumption we ensure that all error types are corrected.

Given the measurement error rate p_m , a particular outcome $q \in \{0, 1\}$ 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 . The superoperator describing the measurement is thus

$$\mathcal{P}_m(\rho) = (1 - p_m)|q\rangle\langle q| + p_m|\bar{q}\rangle\langle\bar{q}|. \quad (10)$$

The fidelity of the purified state as a function of ϵ is shown in figure 5 for Level 1, Level 2 and Level 3 purification protocols. The values chosen for the intra-module error rates correspond to the values reported in recent ion trap experiments [11, 12], $p_1 = 1 \times 10^{-4}$, $p_2 = 0.001$ and $p_m = 0.0005$. Note that both the single-qubit and two-qubit gates are laser driven, which can be readily localized to the scale of one of our trap regions (rather than encompassing both). Figure 5 also displays the probabilistic running times of the purification protocols calculated by numerically simulating the Markov chain in figure 3.

7. Application: communication

In this section we consider how a purification module as described above can be used to distribute high quality Bell pairs between remote locations in a network. A module capable of purifying, storing, and processing entangled states is an enabling device for communications applications in general, and can be thought of as a universal communications node. For the sake of demonstration, we select one specific scenario and calculate the resulting rate and fidelity of the long-range entanglement that can be achieved. Beyond the obvious relevance to

Table 1. Creating long range entanglement. Note that the initial (short range) purified fidelity at stage (i) is recovered at stages (iii) and (v). These data are obtained using a network error rate of $\epsilon = 0.1\%$, local two-qubit gate fidelity of 99.9%, and perfect single qubit gates (which have been realised at far higher fidelity). During purification processes (iii) and (v), local operations are used to permute the three erroneous Bell states so that the state which will escape the next purification has the lowest probability. Fidelity of state at final stage (v) obtained when starting at different network error rates ϵ are shown in figure 7.

Stage	Φ^+ (fidelity)	Error channels (decreasing order)
i	0.993817	0.00443, 0.000957, 0.000796
ii	0.922	0.052, 0.01396, 0.0123
iii	0.994154	0.00433, 0.00103, 0.000487
iv	0.925	0.0511, 0.0146, 0.00888
v	0.99450	0.0044, 0.000681, 0.000422

quantum key distribution (QKD), a range of applications for shared entanglement exist [30], including for example secure remote use of a quantum computer via ‘blind’ QIP [31, 32], secret sharing [33] and fingerprinting [34]. The degree to which these emerging applications tolerate network noise is not yet fully established in the literature; we will target a small infidelity of 0.5% (and in a known channel) since this is well within threshold for fault tolerant QIP. We note that while QKD tolerates rather higher infidelities, nevertheless low infidelity is very beneficial especially when the need to communicate relatively short strings is considered [35].

We will consider a relatively naive use of the purification module to act as a repeater, so that a chain of such modules spans the distance between two remote parties. This will be a ‘first generation’ approach in the sense of [1]. We estimate some performance characteristics in this simple scenario, and indicate where a more sophisticated approach based on code states may become preferable.

In our scenario Alice and Bob are at two remote points in a quantum network, and they wish to use the network to generate Bell states so as to create a shared secret key known only to themselves. Being part of a network they do not have a direct connection between them, but they can identify a path (or paths) involving a number of intermediate nodes. For simplicity we will assume that they identify a single path and make exclusive use of the nodes along it until they have succeeded in their task, see figure 6(a). Obviously, generalisations are possible involving using multiple paths and/or sharing node functionality with other network users. Alice and Bob will use the nodes along the path in order to generate Bell pairs that they alone share; they need not trust the operators of those intervening nodes.

Suppose that each link along the path is an ion trap device as outlined previously. We assume that a frequency conversion technology is used in order to translate single photons from the ion’s native emission frequency to a telecoms frequency 1550 nm, so that transmission over distance through fibre is feasible. We also assume that the ion trap modules alternate with measurement modules along the chain as in figure 6(d) (although it is possible that a measurement system can be co-located with each ion trap figure 6(b)).

Each module could be as simple as the five-ion device detailed above, in which case entanglement would be created and purified first with the neighbour to one side, then stored in the ‘application qubit’ (which acts simply as a memory) while entanglement is created and purified with the neighbour to the other side (figure 6(c)). This could suffice for a smaller scale network. However in figure 6(d) a series of grey qubits are indicated; these are additional memory ions which are not essential but have the effect of increasing efficiency: they store purified Bell pairs so as to provide a ‘buffer’ between Bell creation and consumption. Provided that the rate of consuming the Bell pairs is set to be somewhat slower than the average creation rate, the buffers will tend to replenish and the system can operate in a near deterministic way, rather than having the entire device wait for the slowest link. The result will be an increase in bit rate by a factor that is logarithmic in the chain length; this is likely to be a worthwhile enhancement and will require a third zone in addition to the two indicated in figure 4. Since this zone is purely for storage, it should not add greatly to the complexity.

All nodes are continuously seeking to create purified Bell pairs with their immediate neighbours, figure 6(d) (i). Periodically, certain modules will act to fuse a Bell qubit that is shared with a left neighbour with a Bell qubit shared to the right. This is achieved by performing a Bell-basis measurement on the local pair, recording the measurement outcome and communicating it classically to one of the neighbouring modules; that module can apply a local single-qubit rotation to complete the process. (Note that a complete Bell-basis measurement is easily achieved using, for example, a control-NOT between the two qubits, followed by a Hadamard on the controlling qubit and then measurement of each qubit separately in the z -basis.) As shown in figure 6(d)(ii) these

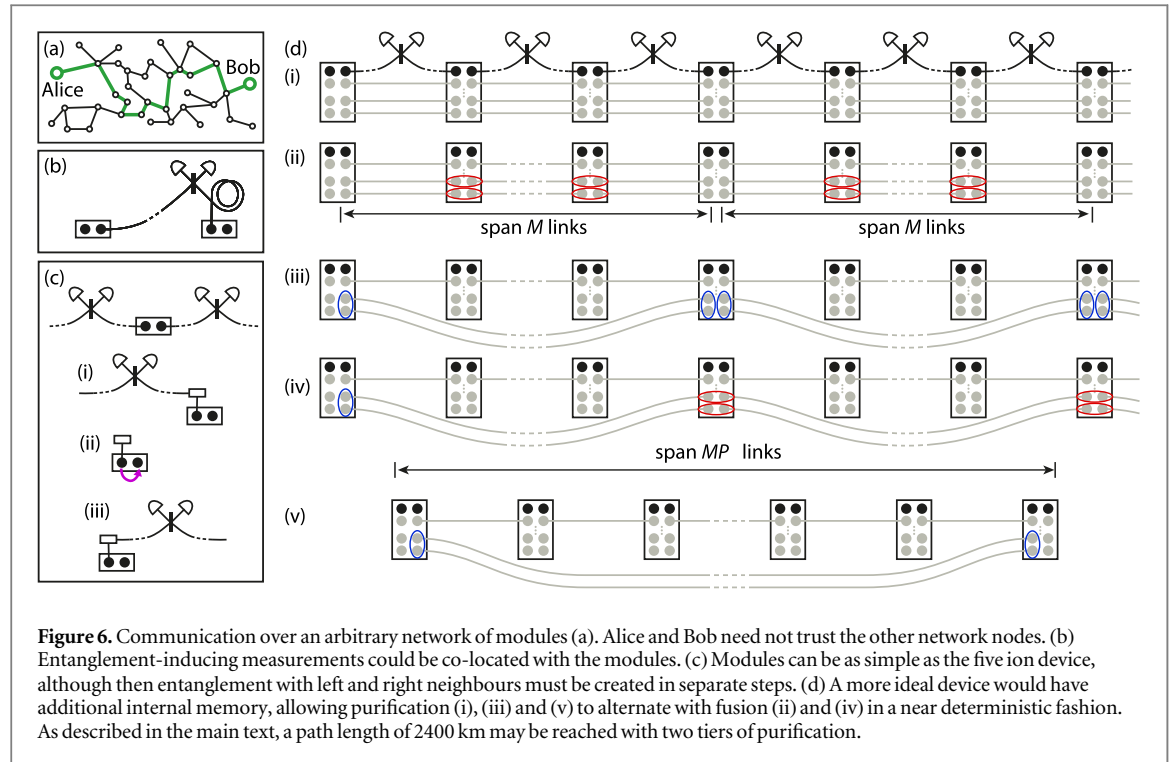


Figure 6. Communication over an arbitrary network of modules (a). Alice and Bob need not trust the other network nodes. (b) Entanglement-inducing measurements could be co-located with the modules. (c) Modules can be as simple as the five ion device, although then entanglement with left and right neighbours must be created in separate steps. (d) A more ideal device would have additional internal memory, allowing purification (i), (iii) and (v) to alternate with fusion (ii) and (iv) in a near deterministic fashion. As described in the main text, a path length of 2400 km may be reached with two tiers of purification.

fusion operations can happen in $M - 1$ consecutive modules so as to fuse together M Bell pairs in a single step. The result is a long-range Bell pair shared between the end-point nodes. Any noise present in the original Bell pairs will contribute to the noise in the newly fused, long-range pair. Therefore these long-range pairs are purified, before the same process is repeated to fuse them into very long-range pairs.

The costs of this process can be estimated as follows: suppose that the infidelity in the raw entanglement process is 10% (and this must include the effects of the frequency conversion technology). Take the initial purification to be a Level 3 process according to figure 3. This will produce Bell pairs with an infidelity of 0.6%, and moreover the noise will be largely in a single channel (see equation (7) and table 1). The average time cost will correspond to the creation of 8.34 ‘raw’ Bell pairs. When a chain of M such Bell pairs are now fused together by high grade local operations (figure 6(d)(ii)), the result will be a Bell pair whose infidelity is greater by a factor of approximately M . For $M = 12$, numerical simulation produces the numbers shown in table 1. Now suppose that a Level 2 purification is performed on these pairs. Because of the structure in the noise, level two suffices to reduce the infidelity back below 0.6% as shown in the table. This necessarily consumes at least 4 of the long range pairs, and because of the failure possibility the average cost is in fact 4.77. Finally the long range pairs are again combined, with P of them being fused into very-long-range pairs. If $P = 12$, then we have a total range of $12 \times 12 = 144$ modules. Finally performing another Level 2 purification yields final Bell pairs with infidelity 0.55%, at an average cost 4.78 input pairs. We see that we can characterise this process by saying that we suffer a reduction in the rate of pair distribution by a factor of ~ 4.8 every time we increase the range by a factor of 12. The process could be continued to reach longer chains, using the same rule.

Suppose that the modules are spaced apart by d kilometres, and that standard silica telecoms fibre (Corning SMF-28) is used with a photon loss rate of order 0.17 dB per kilometre. Then the loss in reaching the measuring station mid-distance between modules, figure 6(d)(i) will be $(0.085)d$ dB. However, it is likely that a procedure involving detecting two photons would be employed in order to alleviate the demands for interferometric stability. Therefore the photon loss probability will impact the success rate quadratically, so that it falls as $(0.17)d$ dB. We might reasonably assume that d is chosen so that the total success rate only falls by a factor of 2 (on top of all other loss mechanisms including collection and detector inefficiency, losses in links and within the frequency conversion system). Then $d = 17$ km is an appropriate spacing, and our chain of $12 \times 12 = 144$ modules spans more than 2 400 km. But to reach this range we have used several tiers of purification: a Level 3 purification at the initial stage of entanglement generation between modules, costing 8.34 raw pairs per success, and then two additional purification phases (figures 6(c)(iii) and (c)(v)) each of Level 2, taking 4.77 and 4.78 input pairs, respectively, to produce one upgraded pair. The total factor between the raw Bell pairs and the purified remote pairs is therefore 190. (Note that a further factor of 4.8 would allow us another factor of twelve in separation, reaching 29 000 km and so exceeding the distance between any two points on Earth’s surface.)

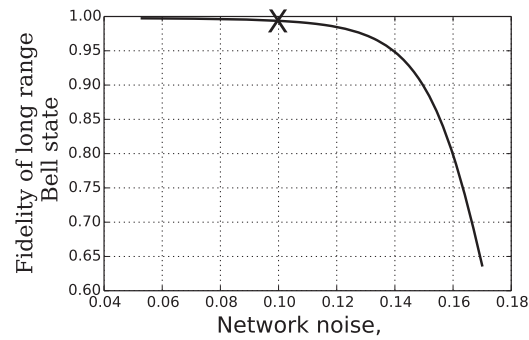


Figure 7. Distributed entanglement fidelity as a function of raw network noise ϵ , calculated for the communication scenario described in the main text. The instance considered in the main text and table 1, where network noise ϵ is taken to be 10%, is marked on the graph by a cross.

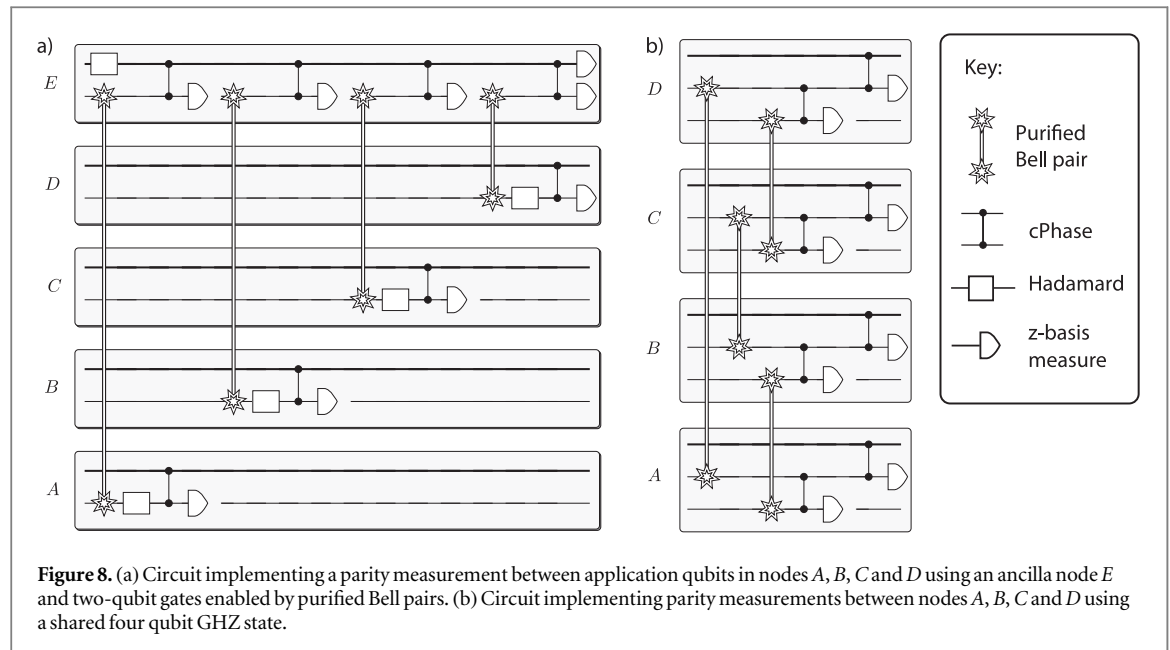
To the factor of 190 we must add the cost of translating the remote Bell pairs into shared secret bits, i.e. the QKD protocol itself. The Bell pairs we have created have higher fidelity than those considered in the relevant literature, and moreover our noise is concentrated into specific channels. Therefore while it is not possible to identify an exact cost for this stage without further study, it is reasonable to hope that this cost can be somewhat less than those that have been computed to date: Inspecting table 1 in [35] one might expect the ratio between rates of shared key creation and Bell pair distribution to be between 0.2 and 0.6, depending on how large a key is generated (larger keys being more efficient to generate).

Ultimately, then, we may see nearly three orders of magnitude reduction between the rate of ‘raw’ entanglement generation between neighbouring modules, and the rate of generating secure bits between users 2400 km apart. Given that key lengths of 10^4 bits may be necessary for practical QKD, one would require a raw entanglement generation rate of 100 kHz to create such a key in 100 s. This is very demanding given that raw entanglement rates *in the lab* are presently four orders of magnitude slower. However, it is reasonable to suppose that multiple ion trap devices can and would be implemented within each repeater station; since each would be independent from the others, this would not increase the technical sophistication and indeed the expensive components such as control lasers could be used as a common resource for multiple traps. Through such an approach, together with anticipated improvements in the efficiency of collecting light from ions, it may be possible to reach the communication rates described here.

Note that once one considers having multiple ion trap devices within each repeater station, the possibility arises that one could interlink those devices locally; then it would become possible to use an encoding such as Raussendorf’s 3D cluster state [36] to fault tolerantly transmit entangled pairs [37, 38], with each repeater forming a ‘sheet’ of the structure. One could employ a variant of the approach described in [39] to achieve transmission rates that do not deteriorate with distance (at the cost of considerably greater complexity in the repeaters). Moreover such a system, being a ‘third generation’ repeater [1], need not be limited by classical communication times—we conclude by assessing the significance of this limit for our naive approach.

The finite nature of the speed of light leads to a bound on the entanglement rate: in order to know whether an entanglement attempt has succeeded or failed one must wait for light to travel from the ion to the detector system, and the return of a classical signal. If the attempt has failed, the ion cannot be reset for another attempt until this information is received; this therefore puts a limit on the attempt rate of c/d where c is the speed of light. In dense urban networks this might not be an issue, since d may be less than a kilometre. But for our long range network where we have assumed $d = 17$ km, the implied maximum cycle rate is 18 kHz, a factor of 26 below the 470 kHz rates that have been used in the lab in entanglements [5]. The obvious solution is to reduce the separation d between repeater stations, but this is expensive and will imply that more purification is needed for a given total distance. A more advanced solution would be to have multiple ions available for ‘raw’ entanglement, such that each acts briefly as memory while the results of its latest entanglement attempt are awaited. Within this time several other ions would begin the process of entanglement generation, sequentially.

Thus for long range networks an ideal device might employ the same central two zones as described in section 5, but peripherally there would be one zone for passive memory ions (the grey circles in figure 7) and another region containing several ions that are dedicated to obtaining raw entanglement.



8. Application: fault tolerant computing

We now consider the use of our modules in the context of scalable, fault tolerant quantum computing. We will evaluate the performance of our five-ion module as described earlier, as a building block of an architecture that uses the toric quantum error correcting code. The application qubit within each module now represents one ‘data qubit’ of the toric code. We numerically simulate the code and determine the error correction thresholds in terms of the network noise.

The toric code involves repeatedly measuring *stabilisers*, which correspond in practice to parity measurements on groups of four data qubits. The basic repeating cycle of the computer involves alternating these measurements with Hadamard rotations to switch between the x and z basis. In addition to preserving the logical quantum information, one can implement all operations required for universal computation by varying these parity-checking measurements. The full process is complex and involves magic state purification [40, 41].

There are different ways in which one can carry out a stabilizer measurement on qubits in separate traps that share some entangled states. We consider two methods whose circuit diagrams are shown in figures 8(a) and (b). Method (a) uses five nodes, each of which is a module of the kind described in section 5. The circuit shown in figure 8(a) effectively induces a parity measurement on the application qubits in nodes A, B, C and D. Node E is an ancilla—the application qubit of node E will be measured and that measurement result determines whether the parity of A, B, C and D is even or odd.

A second way of inducing a parity measurement follows the approach in [20] and is shown in figure 8(b). This method uses four nodes, each of which contain two envoy qubits and one application qubit. The idea behind this method is to create a shared 4-qubit GHZ state between the four nodes, which can then be used to generate the non-local parity measurement of the application qubits in node A, B, C and D. The advantages of method (b) over method (a) are that method (b) uses only four nodes, that the operations can be performed parallel and (as we will see) it has higher error correction thresholds. The disadvantage of method (b) is that it requires an extra ion per trap to store an additional envoy qubit and would involve a longer purification protocol. Thus this is not strictly compatible with the ion trap layout shown earlier; we would need to introduce an additional Ca ion and recompile the purification process to this different layout; however, the purification process (denoted by the linked double star symbol in figure 8) would remain the same, therefore we can include this case for comparison without generating an explicit low-level blueprint.

The fault-tolerance threshold is an important characteristic of a quantum error correcting code. The concept here is that applying the error-correcting process will actually make things worse, i.e. accelerate the degradation of the logical qubit, if the process of parity measurement is too noisy. Then enlarging the code will actually increase the logical error rate. However, if in fact we find that enlarging the code suppresses errors on the logical qubit, then indeed we are successfully operating within the fault-tolerant threshold. We will set the local gate fidelities to a constant level (corresponding to state of the art numbers) and then consider different levels of network noise, in order to identify the threshold.

The thresholds are calculated using Monte-Carlo simulations [42]. The procedure is the same as that described in [4, 43] and indeed the same base numerical code was employed (please see ‘/naominickerson/fault_tolerance_simulations/releases’ on github.com for the base code). To summarise the procedure: for given local error rates ($p_1 = 1 \times 10^{-4}$, $p_2 = 0.001$ and $p_m = 0.0005$) we select a network size characterized by parameter L such that there are $2L^2$ cells in the toric network. The stabilizer measurement cycles are then simulated. Each stabilizer measurement may introduce error(s), which in turn may induce changes in the stabilizer outcomes in the next cycle. This syndrome information is recorded over $4L$ cycles and then Edmonds’ minimum weight perfect matching algorithm is used to attempt to infer the appropriate corrective operations that would recover the ideal state. Consequently the logical qubit either does, or does not, receive an error. This numerical experiment is repeated many times (typically ~ 16000) to find the probability that the logically encoded qubit avoids an error, and this is plotted as the y -axis of the graphs in figure 9. The entire process is then repeated for a larger L , to establish whether this raises or lowers the probability of logical error. We note that our simulations use a conventional decoder that assumes measurements are available synchronously; in the physical machine, one would require an asynchronous decoder or else a (modest) time cost is incurred; this is explained in appendix section B.

The procedure outlined above requires as input a set of error rates. For example there will be a specific probability that the correlated error $XYII$ will occur on the four data-qubits involved in a stabilizer measurement; similarly there are specific probabilities for all other error combinations including measurement error on the ancilla qubits. These are pre-calculated by finding the superoperator that describes the effect of the measurement protocol on the input qubits. The superoperator can be obtained by simulating the circuit and making use of the Choi-Jamilkowsky isomorphism. The superoperator can be written as

$$\mathcal{S}(\rho) = \sum_{i=0} p_i K_i \rho K_i^\dagger. \quad (11)$$

The map $\mathcal{S}(\rho)$ describes the operation as Kraus operators K_i applied to the input state ρ 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. All of K_i can be expressed as one of the ideal parity projectors followed by (one or more) 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 an application 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 toric code array.

Figure 9 shows the results of the toric code simulations using methods (a) and (b) of measuring stabilizers with modules containing Level 1, Level 2 and Level 3 purification protocols. We see that the effect of the purification process has been to tolerate network noise at a very high level; for the five-ions-per-module approach the threshold is 17% noise, and the addition of one further ion can boost this to $>25\%$ noise through the GHZ protocol.

In order to achieve the highest possible thresholds, we exploited the fact that our Level 3 purification results in most of the error probability being associated with a specific one of the three incorrect Bell states (see equation (7)). Single qubit gates suffice to move this probability to whichever erroneous Bell state we wish; we moved it to that Bell state which, when employed in the remote gating process figure 8(a), gives rise to a Z error on the ancilla but no error on the data qubit. Ultimately this leads to an incorrect stabiliser result being recorded when the ancilla is measured; this pure ‘lie’ is the most well tolerated type of error in the surface code approach.

Finally we briefly return to the idea of exchanging the roles of the application qubit and the envoy qubit, which we alluded to in section 3. In the present context the application qubit is a single ‘data qubit’ of the surface code. The reason to exchange roles is to limit the impact of leakage errors, i.e. errors where an ion leaves the qubit subspace and enters some other state that is not computationally meaningful. Such events are tolerated by the surface code (without explicitly identifying the errors) only if the data qubit is subsequently returned to the computational subspace, so that only a very small proportion of data qubits are in an invalid state at a given time. Since the act of measuring an ion will return it to the correct subspace, an elegant solution is to measure out the ion bearing the data qubit when it is entangled with the envoy, thus teleporting the data qubit onto the envoy which now *becomes* the new data qubit. This ‘passing of the torch’ means that no single ion will remain unmeasured for very long; however it does leave the new application qubit in the incorrect physical position. Thus one is motivated to apply physical permutation prior to the exchange. Such a permutation also has the benefit that the ‘old’ data qubit can be shuttled to the ‘rowdy end’ of the trap prior to measurement, maintaining the principle that measurement always occurs at that end.

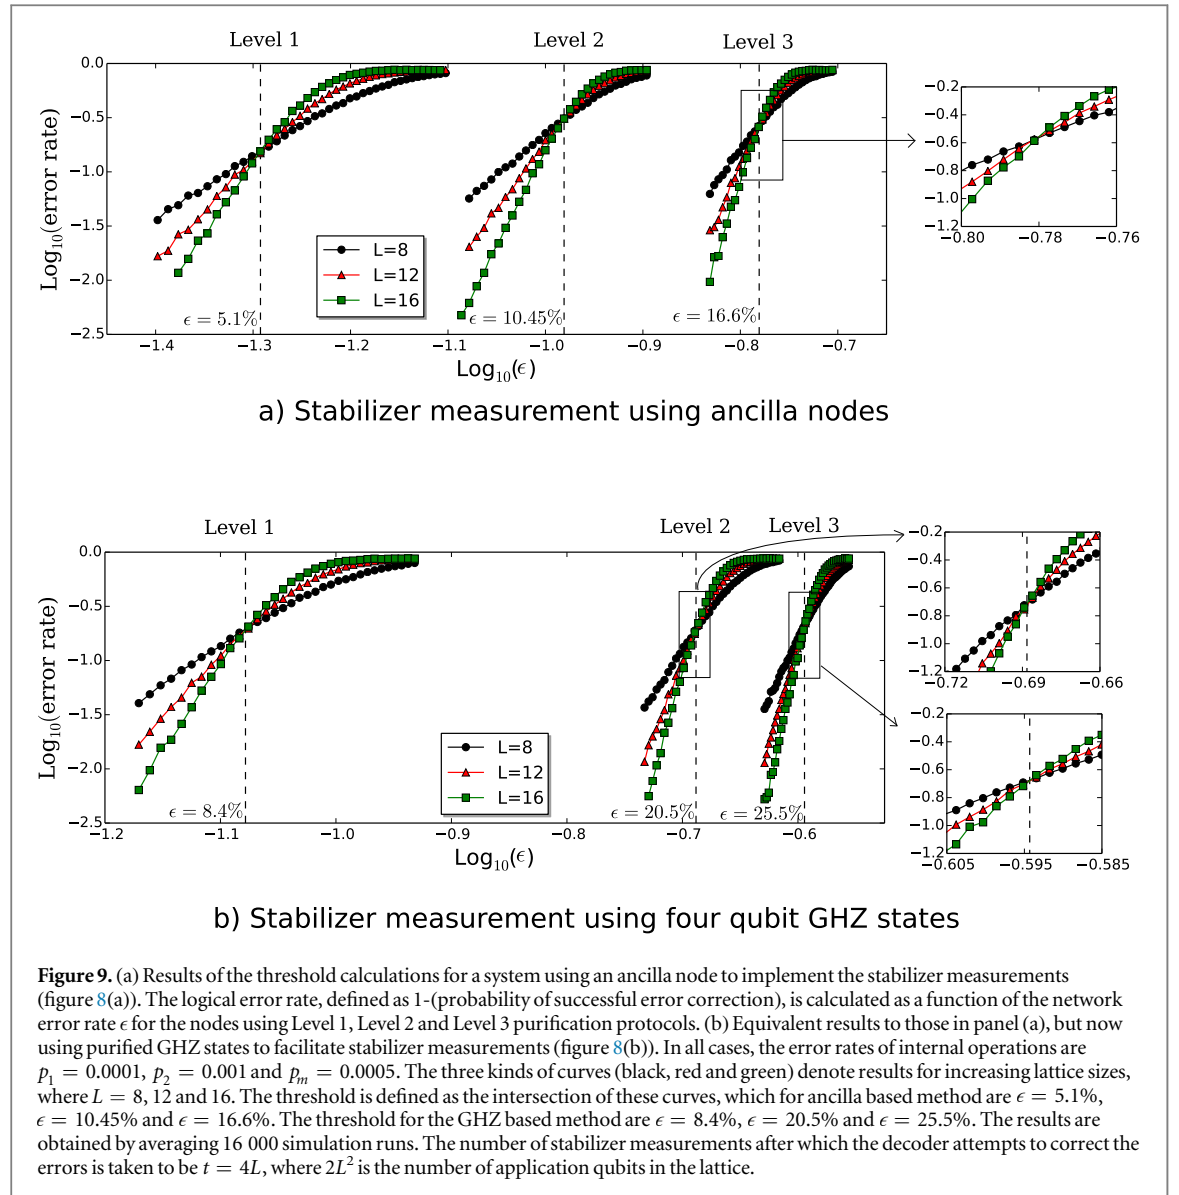


Figure 9. (a) Results of the threshold calculations for a system using an ancilla node to implement the stabilizer measurements (figure 8(a)). The logical error rate, defined as $1 - (\text{probability of successful error correction})$, is calculated as a function of the network error rate ϵ for the nodes using Level 1, Level 2 and Level 3 purification protocols. (b) Equivalent results to those in panel (a), but now using purified GHZ states to facilitate stabilizer measurements (figure 8(b)). In all cases, the error rates of internal operations are $p_1 = 0.0001$, $p_2 = 0.001$ and $p_m = 0.0005$. The three kinds of curves (black, red and green) denote results for increasing lattice sizes, where $L = 8, 12$ and 16 . The threshold is defined as the intersection of these curves, which for ancilla based method are $\epsilon = 5.1\%$, $\epsilon = 10.45\%$ and $\epsilon = 16.6\%$. The threshold for the GHZ based method are $\epsilon = 8.4\%$, $\epsilon = 20.5\%$ and $\epsilon = 25.5\%$. The results are obtained by averaging 16 000 simulation runs. The number of stabilizer measurements after which the decoder attempts to correct the errors is taken to be $t = 4L$, where $2L^2$ is the number of application qubits in the lattice.

9. Optical switching

Prior experimental demonstrations of remote entanglement do not include optical switches, and one might be concerned with the photon loss that will inevitably occur with such elements. Fortunately this loss need not be very significant: there are already commercially available switches⁵ which offer a one-to-twelve branching with a typical loss of only 0.9 dB, i.e. absorption of about 19% of photons. We note that a 12-fold branching is more than sufficient to support a surface code implementation: each module would require only a four-fold branch in order to link to the north, south, east, west neighbouring modules. The eight remaining links could fruitfully be used to enable each logical qubit (i.e. each torus in the toric surface code) to be transversally linked to eight others. If still higher levels of connectivity were desired, then micromechanical mirror arrays could be employed; here, well established devices exist which allow impressive levels of switching at only modestly higher loss. For example, a 2003 report describes a 238-to-238 switch with a typical loss of only 1.4 dB [44].

10. Performance

We now discuss the question of how quickly a network device could operate, i.e. the ‘clock speed’ for operations between modules. The fundamental requirement is that this speed should be much faster than the rate of environmentally induced decoherence on each module’s application qubit, so that the fidelity of inter-module

⁵ See e.g. leoni ‘fiber optics’ catalog, p 331 (available at leoni-fiber-optics.com).

operations is not significantly degraded by such decoherence. We will conclude that this is possible even with today's technology, and indeed such a machine would surpass threshold for operation in fault tolerant mode; however it would be slow in such a mode. In the subsequent section we estimate the higher clock speeds that should be achievable as technology matures in the foreseeable future.

A machine with today's performance. The performance figures reported in the literature to date indicate that it is feasible to achieve high fidelity inter-module gates on a timescale fast compared to decoherence. As noted above, this is the key criterion for any kind of quantum processing (whether it is fault tolerant, code based computing or simply using the physical qubits directly as algorithmic qubits).

Trapped ion qubits can have long decoherence times; recent experiments report a dephasing time $T_2 = 50$ s and negligible spontaneous decay rates [45]. If we assume a 50 s dephasing time for the application qubit, then the time it takes for the fidelity of the application qubit to drop from 1 to 0.99 is ~ 0.72 s. The Level 3 purification protocol requires on average eight raw entangled pairs (assuming 10% network noise). For the decoherence to be negligible we must have $8T_0 \ll 0.72$ s i.e. $T_0 \ll 0.09$ s, where T_0 is the time required to generate one raw entangled pair. Thus the entanglement rate should be greater than 11 Hz. This condition is within a factor of two of the rate reported in [5], and one can expect that the generation of systems now under testing will comfortably meet this condition. Additionally, the target would fall if the application qubit's decoherence rate were reduced by either technical improvements or by encoding that qubit over two ions, $|0\rangle \rightarrow |01\rangle, |1\rangle \rightarrow |10\rangle$, so as to negate collective phase noise.

It is interesting to infer that technically, near-current technology could support fault tolerant computing in the sense that gate fidelities would surpass the required threshold. However, with a gate time of order tens of Hertz, large scale calculations (i.e. those for which one would require fault tolerance) may be very slow. Instead we envisage first generation machines operating at this speed and using a modest number of application qubits (perhaps from fifty to several thousand) in order to tackle interesting tasks using *directly encoded* qubits, i.e. without incurring the significant resource overheads needed for fault tolerant QIP. Operating in such a domain, even low clock speeds may suffice to achieve performance that is superior to classical systems.

For context, a recent study [46] performed exact simulation of 40 qubits using optimised software running on a leading supercomputer; the authors found that a single-qubit gate operations took over a second to implement. Moreover they concluded that a classical simulation of 50 qubits will not be achieved within a decade. Thus, it is an interesting proposition to build a universal but *non-fault-tolerant* quantum computer with a number of well-controlled qubits exceeding 50 and a gate times of a second or less.

Foreseeable performance as systems mature. In order to reach the era of practical fault tolerant computing, one would wish to increase the speed at which high fidelity inter-module gates can be performed. In this section we consider the speeds achievable within the paradigm of simple five-ion traps.

The rate at which one can *attempt* to generate entanglement ultimately upper-bounds the rate of inter-module operations. One anticipates that the main limiting step will be the optical pumping. The duration will need to be at least five times the decay constant, and once one includes some modulator latency and switching times, together with processing latency, practical repetitions may require > 200 ns. We should also recall that one cannot begin a new attempt at entanglement until the success or failure of the prior attempt is known, and this is limited by the speed of light 'round trip' from the modules to the detectors, however this does not become a constraint unless the components are tens or hundreds of metres apart.

We can therefore suppose that entanglement between a given pair of devices is attempted at rate of a few megahertz. This would be an acceptable clock speed, even allowing for the factor of 8 reduction when one accounts for purification from a presumed initial infidelity of 10% (as shown in figure 5). However, unfortunately the rate of raw entanglement generation is typically reduced by the square of the probability that an emitted photon successfully yields a detector click, and this probability is low in experiments to date. The total probability of photon retention is the product of the retention probability at each stage: collection from the ion into the fibre mode, transfer through the network, and detector efficiency. Fortunately all stages *except* the initial collection can be engineered to have rather low loss rates, so that a photon that is successfully captured into the network will indeed yield a detector click with probability greater than 50%; we note that suitable room temperature switches exceeding 60% efficiency are available commercially⁶. However the efficiency of collection from an ion cannot approach unity unless integrated cavities are employed. This possibility is being investigated by multiple groups (see e.g. [47, 48]) but it appears very challenging at present. Meanwhile although deep parabolic mirrors have been employed to achieve collection efficiency above 50% [49], these are impractical for the small and simple structures we envisage in the present paper. Recently however a 5.8% collection efficiency has been achieved by exploiting diffractive mirrors integrated into the trap substrate [50], and the authors predict a theoretical potential of 8% efficiency from the same prototype design. We might reasonably assume that efficiency can reach 10% in a mature device of this type. Then assuming that downstream losses from

⁶ See e.g. the 'count blue' device available from laser components gmbh (www.lasercomponents.com).

switches (discussed above) and detectors in our mature system will account for a net loss of only 30%, one would have a 7% probability that an emitted photon is detected. Typical entanglement protocols require that two photons are retained, thus this probability must be squared, and also suffer an inherent 50% failure rate (they project an initial $|+\rangle|+\rangle$ state onto the odd parity subspace [51]). Finally we must account for our factor of 8 for purification (from a presumed infidelity of 10%). We conclude that an *attempt rate* of three megahertz can translate to purified high fidelity entanglement generation at about one kilohertz. This is therefore the achievable average rate for remote, high-fidelity operations between application qubits in different traps, since each such gate (for example, a c-Phase gate) consumes one Bell pair, see figure 1.

It is worth observing that there are also variants to the usual entanglement protocol that can increase the rate. For example [52] describes a protocol which, at a cost of a factor of two in the inherent success rate, avoids squaring the photon retention probability. Given the above numbers this boosts the entanglement rate 7-fold. The challenge of adopting such a scheme is that the system must be interferometrically stable over the timescale of two successive entanglement events.

What is the utility of a one kilohertz machine? A very recent study [53] has estimated the time required to execute Shor's algorithm on classically-intractable problem sizes, using the surface code (as evaluated here in section 8). Balancing the tradeoff between space and time resources, the authors estimate that a machine with a clock speed of 1 kHz would require about a fortnight to perform a classically infeasible task, i.e a Shor factorisation of a 1000 bit number. They note that further optimisations will likely reduce this number. A kilohertz clock speed would therefore be sufficient for a general fault-tolerant computer to achieve meaningful, classically-infeasible tasks.

We note that the goal of [53] is to optimise the vital low-level process of magic state distillation; Shor's algorithm is considered as an example of a high level task. A full optimisation of the software hierarchy required for Shor factorisation would no doubt identify additional costs and efficiencies. Fortunately the author's implicit assumption that there are no connectivity restrictions between the physical qubits is broadly supported by our approach: the modular architecture enables traversal two-qubit gates as an alternative to local approaches such as braiding or lattice surgery. A transversal CNOT would require a single physical CNOT to be performed between each corresponding pair of physical qubits (in parallel), possibly followed by several rounds of stabiliser evaluation according to the sophistication of the decoder.

Longer-term enhancements to modules. Our goal here has been to show that even very simple modules can be entirely adequate as building blocks of universal communication and computation systems. The numbers we discuss above are in this context; but of course complexity of the modules would offer further improvements. For example a branched 2D trap containing dozens of ions could support multiple sites for the generation of 'raw' entanglement and could queue or 'buffer' the entangled ions in order to make the device quasi-deterministic. Moreover, the introduction of cavities [47] (see footnote 5) into the ion trap design could boost the optical collection efficiencies to very high levels, dramatically boosting the entanglement rates. Thus, there are a number of avenues to pursue which could ultimately realise an entanglement rate comparable to, or greater than, the speed of high fidelity local operations within the device. The gate speed of such a modular ion trap system would then match the gate speed of any other (e.g. monolithic) ion trap system, while maximising the connectivity and scalability.

11. Conclusion

We have designed a simple ion trap device for general use as a building block of optically-linked quantum technologies. The unit is capable of interfacing with other similar units over a noisy optical channel, and purifying that channel to enable high fidelity quantum operations between the units. We simultaneously designed a novel purification protocol alongside the device layout in order to meet a number of desiderata for the system. Notably, the proposed system has only five ions in a linear arrangement. These ions (which are of two species) suffice for all entanglement generation, processing, storage and cooling operations. Laser control systems need not differentiate a given ion from its neighbour.

We evaluated this device in the context of communication over a network of untrusted nodes. Using gate fidelities already reported in the literature we found that each twelve-fold increase in range between the parties leads to a reduction in the rate of communication by a modest factor of 4.8. This shows the efficiency of the five-ion purification process, however we noted that additional peripheral ions would be desirable as a 'buffered' memory and for rapid sequential entanglement attempts.

The same five-ion device was assessed as a component for scalable fault-tolerant computing. Again using reported gate fidelities, we concluded that very high noise can be tolerated in the links of such a system: the threshold was 17%. In contrast to the communication scenario, here the five-ion core device will suffice as a

building block without additional memory or entangling qubits (although permitting one additional ion will raise the noise threshold further, to exceed 24%).

We conclude that this relatively simple system is a powerful and general device, suitable to be the core of a universal communications node or the building block of a scalable computer (in both the pre- and post-fault tolerance eras). For all applications, a key challenge is to increase the rate of entanglement between modules.

Acknowledgments

The authors are grateful for helpful discussions with Tom Harty and David Lucas. The authors acknowledge support from the EPSRC National Hub in Networked Quantum Information Technologies EP/M013243/1.

Appendix

A.1. Two qubit gates between remote modules

As noted in the main text, there is a simple three-step process by which Alice and Bob, having previously created a high fidelity Bell pair shared between them, can consume this pair in order to perform a cPhase gate between their ‘application qubits’. First they each perform a cPhase gate between their Bell qubit and their application qubit. The result can be written

$$\frac{1}{\sqrt{2}}(|00\rangle \mathbb{I} + |11\rangle Z_A Z_B)|S\rangle$$

where Z_A is a single-qubit phase gate $\text{diag}\{1, -1\}$ acting on Alice’s application qubit, Z_B is analogous, and $|S\rangle$ represents both parties’ application qubits as well as any other qubits entangled with them, i.e. the ‘rest of the system’.

Now Alice measures her Bell qubit in the x -basis and Bob does the same but in the y -basis. Each receives either a $+1$ or -1 on their measuring device, resulting in

$$\begin{aligned} & \frac{1}{\sqrt{2}} \langle X \pm_A | \langle Y \pm_B | (|00\rangle \mathbb{I} + |11\rangle Z_A Z_B) |S\rangle \\ &= \frac{1}{2\sqrt{2}} (\mathbb{I} + i(\pm 1_A)(\pm 1_B) Z_A Z_B) |S\rangle \\ &= G |S\rangle \quad \text{for } ++ \text{ or } --, \\ & \quad G^\dagger |S\rangle \quad \text{otherwise.} \end{aligned}$$

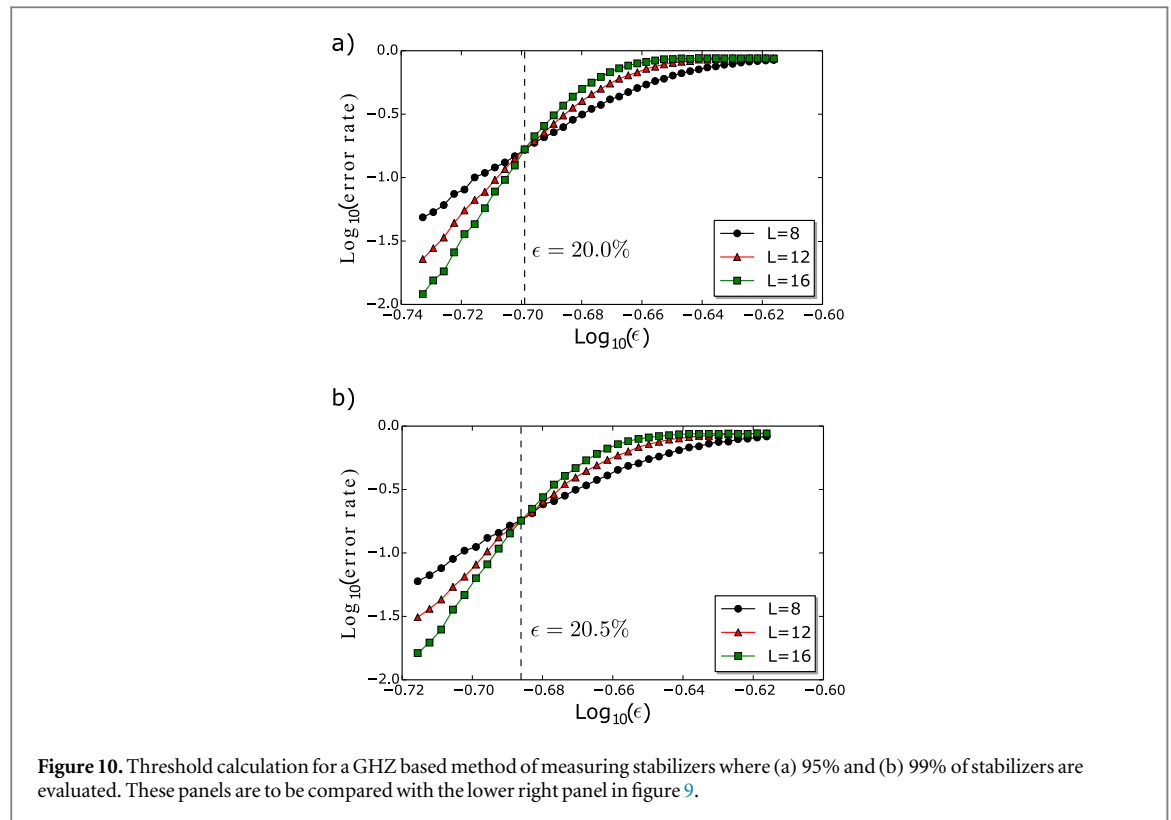
Here G is a diagonal matrix with elements $\{1, -i, -i, 1\}$, and we are neglecting global phases $\exp(\pm i\pi/4)$. Matrix G is transformed to the desired cPhase operation, i.e. $\text{diag}\{1, 1, 1, -1\}$, by local gates $S = \text{diag}\{1, i\}$ performed by both Alice and Bob. Meanwhile G^\dagger instead requires them to apply S^\dagger .

A.2. Probabilistic stabilizer evaluation in toric code

In the main paper we remarked that a novel decoder would be required in order to achieve the thresholds which we have derived from our models. Here we expand on this point, and note that alternatives exist if such a decoder where not forthcoming.

The network paradigm implies that the timescale for completing a stabiliser measurement is not fixed. The raw entanglement events occur at random, and moreover the process of purification can involve repeated stages (see the Markov chain in figure 3). Consequently the process of evaluating a stabiliser, whether via the method of a common ancilla or via GHZ state generation, will take an uncertain amount of time. But the surface code, in common with other fault tolerance schemes, canonically involves synchronising the process of gathering syndrome information (the stabiliser measurements) and inferring the corrections needed to maintain the logical information. Ideally one would derive a decoder whose performance matches that of this simple scenario, and yet is capable of utilising asynchronous stabiliser information. The thresholds presented in figure 9 are derived from a conventional decoder based on minimum-weight perfect matching using synchronous sheets of information, and thus implicitly we are assuming that an asynchronous decoder can be created which matches the performance of this standard approach. Decoders with a suitable structure have been created [54] but to the authors’ knowledge no such approach has matured to the point of reaching the near-ideal thresholds of the standard synchronous approach. This is therefore an area for ongoing research.

In lieu of such a decoder, there is a simple solution that makes the asynchronous measurements pseudo-synchronous, so that established decoders can be used: one can simply abandon the attempt to evaluate any stabiliser when it has taken ‘too long’. In effect one introduces a ‘deadline’ for stabiliser evaluation, such that some portion of all stabiliser measurements will have failed to report. The impact is that the decoder now lacks



any information about each such outcome. Fortunately this is not too severe an issue provided that the density of missing data is sparse (clearly, it is less severe than having a false measurement, for example). To clarify this point we have performed two threshold simulations where stabiliser evaluation is aborted at time such that 99%, or 95%, of the stabilisers have reported. The results are shown in figure 10. Comparing with the corresponding panel in the lower right of figure 9, one sees that the impact of the 99% cut-off is not visible, while the impact of the 95% cutoff is a slight reduction in threshold from 20.5% to 20%. Indeed the practical cost is not this small reduction, but rather the increase time needed: a factor of 2 for the 95% case as can be read off from the lower right panel of figure 5. Nevertheless, if one were concerned that an asynchronous decoder could not be created, or that its performance would be significantly inferior to the synchronous decoders that presently exist, then one can envisage adopting this approach.

References

- [1] Muralidharan S, Li L, Kim J, Lütkenhaus N, Lukin M D and Jiang L 2016 Optimal architectures for long distance quantum communication *Sci. Rep.* **6** 20463
- [2] Monroe C, Raussendorf R, Ruthven A, Brown K R, Maunz P, Duan L- M and Kim J 2014 Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects *Phys. Rev. A* **89** 022317
- [3] Nemoto K, Trupke M, Devitt S J, Stephens A M, Scharfenberger B, Buczak K, Nöbauer T, Everitt M S, Schmiedmayer J and Munro W J 2014 Photonic architecture for scalable quantum information processing in diamond *Phys. Rev. X* **4** 031022
- [4] Nickerson N H, Fitzsimons J F and Benjamin S C 2014 Freely scalable quantum technologies using cells of 5-to-50 qubits with very lossy and noisy photonic links *Phys. Rev. X* **4** 041041
- [5] Hucul D, Inlek I V, Vittorini G, Crocker C, Debnath S, Clark S M and Monroe C 2015 Modular entanglement of atomic qubits using photons and phonons *Nat. Phys.* **11** 37–42
- [6] Pfaff W *et al* 2014 Unconditional quantum teleportation between distant solid-state quantum bits *Science* **345** 532–5
- [7] Reiserer A, Kalb N, Blok M S, van Bemmelen K J M, Twichien D J, Markham M, Taminiau T H and Hanson R 2016 Robust quantum-network memory using decoherence-protected subspaces of nuclear spins *Phys. Rev. X* **6** 021040
- [8] Roch N, Schwartz M E, Motzoi F, Macklin C, Vijay R, Eddins A W, Korotkov A N, Whaley K B, Sarovar M and Siddiqi I 2014 Observation of measurement-induced entanglement and quantum trajectories of remote superconducting qubits *Phys. Rev. Lett.* **112** 170501
- [9] Narla A *et al* 2016 Robust concurrent remote entanglement between two superconducting qubits *Phys. Rev. X* **6** 031036
- [10] Benhelm J, Kirchmair G, Roos C F and Blatt R 2008 Towards fault-tolerant quantum computing with trapped ions *Nat. Phys.* **4** 463–6
- [11] Ballance C J, Harty T P, Linke N M, Sepiol M A and Lucas D M 2015 Laser-driven quantum logic gates with precision beyond the fault-tolerant threshold *Phys. Rev. Lett.* **117** 060504
- [12] Harty T P, Allcock D T C, Ballance C J, Guidoni L, Janacek H A, Linke N M, Stacey D N and Lucas D M 2014 High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit *Phys. Rev. Lett.* **113** 220501
- [13] Gaebler J P *et al* 2016 High-fidelity universal gate set for ${}^9\text{Be}^+$ ion qubits *Phys. Rev. Lett.* **117** 0605505

- [14] Monz T, Nigg D, Martinez E A, Brandl M F, Schindler P, Rines R, Wang S X, Chuang I L and Blatt R 2016 Realization of a scalable shor algorithm *Science* **351** 1068–70
- [15] Richerme P, Gong Z-X, Lee A, Senko C, Smith J, Foss-Feig M, Michalakakis S, Gorshkov A V and Monroe C 2014 Non-local propagation of correlations in quantum systems with long-range interactions *Nature* **511** 198–201
- [16] Britton J W, Sawyer B C, Keith A C, Wang C-C J, Freericks J K, Uys H, Biercuk M J and Bollinger J J 2012 Engineered two-dimensional ising interactions in a trapped-ion quantum simulator with hundreds of spins *Nature* **484** 489–92
- [17] Lekitsch B, Weidt S, Fowler A G, Mølmer K, Devitt S J, Wunderlich C and Hensinger W K 2015 Blueprint for a microwave trapped-ion quantum computer (arXiv:1508.00420)
- [18] Bennett C H, Brassard G, Popescu S, Schumacher B, Smolin J A and Wootters W K 1996 Purification of noisy entanglement and faithful teleportation via noisy channels *Phys. Rev. Lett.* **76** 722–5
- [19] Dür W and Briegel H-J 2003 Entanglement purification for quantum computation *Phys. Rev. Lett.* **90** 067901
- [20] Nickerson N H, Li Y and Benjamin S C 2013 Topological quantum computing with a very noisy network and local error rates approaching one percent *Nat. Commun.* **4** 1756
- [21] Li Y and Benjamin S C 2015 Hierarchical surface code for network quantum computing with modules of arbitrary size *Phys. Rev. A* **94** 042303
- [22] Brown K R, Monroe C and Kim J 2016 Co-designing a scalable quantum computer with trapped atomic ions *Quant. Inf.* accepted (arXiv:1602.02840)
- [23] Debnath S, Linke N M, Figgatt C, Landsman K A, Wright K and Monroe C 2016 Demonstration of a small programmable quantum computer module with atomic qubits *Nature* **536** 63–6
- [24] Kwiat P G, Barraza-Lopez S, Stefanov A and Gisin N 2001 Experimental entanglement distillation and /‘hidden/’ non-locality *Nature* **409** 1014–7
- [25] Pan J-W, Gasparoni S, Ursin R, Weihs G and Zeilinger A 2003 Experimental entanglement purification of arbitrary unknown states *Nature* **423** 417–22
- [26] Reichle R, Leibfried D, Knill E, Britton J, Blakestad R B, Jost J D, Langer C, Ozeri R, Seidelin S and Wineland D J 2006 Experimental purification of two-atom entanglement *Nature* **443** 838–41
- [27] Jiang L, Taylor J M, Sørensen A S and Lukin M D 2007 Distributed quantum computation based on small quantum registers *Phys. Rev. A* **76** 062323
- [28] Campbell E T 2007 Distributed quantum-information processing with minimal local resources *Phys. Rev. A* **76** 040302
- [29] Childress L, Taylor J M, Sørensen A S and Lukin M D 2005 Fault-tolerant quantum repeaters with minimal physical resources and implementations based on single-photon emitters *Phys. Rev. A* **72** 052330
- [30] Broadbent A and Schaffner C 2016 Quantum cryptography beyond quantum key distribution *Des. Codes Cryptogr.* **78** 351–82
- [31] Barz S, Kashefi E, Broadbent A, Fitzsimons J F, Zeilinger A and Walther P 2012 Demonstration of blind quantum computing *Science* **335** 303–8
- [32] Morimae T and Fujii K 2013 Secure entanglement distillation for double-server blind quantum computation *Phys. Rev. Lett.* **111** 020502
- [33] Bell B A, Markham D, Herrera-Marti D A, Marin A, Wadsworth W J, Rarity J G and Tame M S 2014 Experimental demonstration of graph-state quantum secret sharing *Nat. Commun.* **5** 587–9
- [34] Guan J-Y *et al* 2016 Observation of quantum fingerprinting beating the classical limit *Phys. Rev. Lett.* **116** 240502
- [35] Tomamichel M, Lim C C W, Gisin N and Renner R 2012 Tight finite-key analysis for quantum cryptography *Nat. Commun.* **3** 634 01
- [36] Raussendorf R, Harrington J and Goyal K 2006 A fault-tolerant one-way quantum computer *Ann. Phys., NY* **321** 2242–70
- [37] Fowler A G, Wang D S, Hill C D, Ladd T D, Van Meter R and Hollenberg L C L 2010 Surface code quantum communication *Phys. Rev. Lett.* **104** 180503
- [38] Stephens A M, Huang J, Nemoto K and Munro W J 2013 Hybrid-system approach to fault-tolerant quantum communication *Phys. Rev. A* **87** 052333
- [39] Li Y, Barrett S D, Stace T M and Benjamin S C 2013 Long range failure-tolerant entanglement distribution *New J. Phys.* **15** 023012
- [40] Bravyi S and Kitaev A 2005 universal quantum computation with ideal clifford gates and noisy ancillas *Phys. Rev. A* **71** 022316
- [41] Li Y 2015 A magic state’s fidelity can be superior to the operations that created it *New J. Phys.* **17** 023037
- [42] Fowler A G, Mariantoni M, Martinis J M and Cleland A N 2012 Surface codes: towards practical large-scale quantum computation *Phys. Rev. A* **86** 032324
- [43] O’Gorman J, Nickerson N H, Ross P, Morton J J L and Benjamin S C 2016 A silicon-based surface code quantum computer *Njp: Quantum Inf.* **2** 15019
- [44] Aksyuk V A *et al* 2003 238 x 238 micromechanical optical cross connect *IEEE Photonics Technol. Lett.* **15** 587–9
- [45] Harty T P, Allcock D T C, Ballance C J, Guidoni L, Janacek H A, Linke N M, Stacey D N and Lucas D M 2014 High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit *Phys. Rev. Lett.* **113** 220501
- [46] Smelyanskiy M, Sawaya N P D and Aspuru-Guzik A 2016 qhipster: the quantum high performance software testing environment (arXiv:1601.07195)
- [47] Stute A, Casabone B, Schindler P, Monz T, Schmidt P O, Brandstätter B, Northup T E and Blatt R 2012 Tunable ion–photon entanglement in an optical cavity *Nature* **485** 482–5
- [48] Steiner M, Meyer H M, Deutsch C, Reichel J and Köhl M 2013 Single ion coupled to an optical fiber cavity *Phys. Rev. Lett.* **110** 043003
- [49] Maiwald R, Golla A, Fischer M, Bader M, Heugel S, Chalopin B, Sondermann M and Leuchs G 2012 Collecting more than half the fluorescence photons from a single ion *Phys. Rev. A* **86** 043431
- [50] Ghadimi M *et al* 2016 Scalable ion–photon quantum interface based on integrated diffractive mirrors (arXiv:1607.00100)
- [51] Barrett S D and Kok P 2005 Efficient high-fidelity quantum computation using matter qubits and linear optics *Phys. Rev. A* **71** 060310
- [52] Campbell E T and Benjamin S C 2008 Measurement-based entanglement under conditions of extreme photon loss *Phys. Rev. Lett.* **101** 130502
- [53] O’Gorman J and Campbell E T 2016 Quantum computation with realistic magic state factories (arXiv:1605.07197)
- [54] Herold M, Campbell E T, Eisert J and Kastoryano M J 2015 Cellular-automaton decoders for topological quantum memories *Njp: Quantum Inf.* **1** 15010