

From quantum multiplexing to high-performance quantum networking

W. J. Munro^{1,2*}, K. A. Harrison², A. M. Stephens¹, S. J. Devitt¹ and Kae Nemoto¹

Quantum repeaters will be critical to quantum communication and quantum computation. Here, we describe a mechanism that permits the creation of entanglement between two qubits, connected by fibre, with probability arbitrarily close to one and in constant time. We show how this mechanism may be extended to ensure that the entanglement has high fidelity without compromising these properties. Finally, we describe how it may be used to construct a quantum repeater that is capable of creating a linear quantum network connecting two distant qubits with high fidelity. The communication rate is shown to be a function of the maximum distance between any two adjacent quantum repeaters rather than of the entire length of the network.

The twentieth century saw the discovery of quantum mechanics, a set of principles describing physical reality at the atomic level of matter. These principles have been used to develop much of today's advanced technology, including, for example, microprocessors. Quantum physics also allows a new paradigm for the processing of information known as quantum information processing^{1,2}. Over the last decade there has been a huge worldwide effort to develop and explore quantum-based devices and technologies^{3,4}. Devices for quantum key distribution (QKD) are already commercially available⁵. Quantum repeaters are now an obvious next subject for research^{6,7}. Their function is to enable the creation of entangled states between remote locations. Long-distance entanglement is achieved by arranging a number of repeater nodes in a linear network and creating entangled links between adjacent nodes. Once a node has links with both neighbours, entanglement swapping within the node then creates a longer-range link. When entanglement swapping has occurred at all intermediate nodes, an end-to-end link will have formed. This entanglement can be used for QKD, quantum communication, or distributed quantum computation.

The current goal of many research groups is to produce a stream of entangled qubits over long distances, preferably with a communication rate in the megahertz range. Proposals have generally focused on the quantum components necessary to create entangled links between neighbouring nodes, purification of these links, and entanglement swapping to create longer-range links^{8–18}. Entangled links are generally created by entangling a quantum-optical signal with a qubit and then transmitting that signal over a channel to the neighbouring node. There the signal entangles with another qubit and a measurement is then made on the quantum signal, indicating success or failure^{9–11}. The probability of successfully entangling the two qubits scales as $\exp[-L/L_0]$, where L is the distance between repeater nodes and L_0 the attenuation length of the fibre.

The next step is to look at the overall design of a repeater network, considering both the quantum and classical components. A repeater network must be underpinned by experimental techniques for entanglement generation, incorporate purification or error correction to achieve high-fidelity entangled links, and be controlled by classical communication across the network. This should be done in a practical manner without compromising real-world

applicability. Typically, the communication time required for classical messages to be transmitted between nodes severely limits the performance of repeater networks. Many rounds of messaging between nodes need to occur if entanglement distribution and purification are probabilistic processes. In this Article, we show how to maintain near-determinism throughout all aspects of a repeater network. This allows us to propose an efficient, pipe-lined architecture in which it is known when the end-to-end entangled pairs are going to be available and where the requirements on quantum memories are minimized.

Results

Quantum multiplexing. The core element of any repeater network is the creation of entanglement between neighbouring nodes. The fact that this is a probabilistic process is an issue that affects the performance of a repeater network, as one cannot predict when links are going to be available. A classical signal needs to be sent between repeater nodes to confirm that a link is available, so the generation rate is ultimately limited by this roundtrip transmission time. With typical repeater nodes being separated by, for example, 40 km, this would be on the order of 400 μ s. With the probability of success for entanglement generation at such distances being less than 25%, a number of attempts will be needed before one is effectively guaranteed a link. A significant delay results if the attempts are performed sequentially. One could parallelize the operations, but this would require significantly more resources^{19–21}.

A more efficient design is shown in Fig. 1. In this design each repeater node comprises two parts: a bank of transmitters and a bank of receivers. The creation of an entangled link begins with a classical pulse initiating all the transmitters in a node to prepare individual quantum-optical signals. These signals then interact and become entangled with respective qubits. They then propagate, temporally multiplexed together with the classical heralding pulse, along a fibre to the next node. The classical pulse announces that a series of quantum signals is about to arrive. A receiver initializes a qubit into the appropriate state, and this interacts with the first signal. The signal is measured using an appropriate detection scheme (photon-counting or homodyne) to determine whether a successful entanglement-creation operation has occurred. If not, the qubit is re-initialized for the second signal, and another

¹National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan, ²Hewlett-Packard Laboratories, Filton Road, Stoke Gifford, Bristol BS34 8QZ, UK. *e-mail: bill.munro@me.com

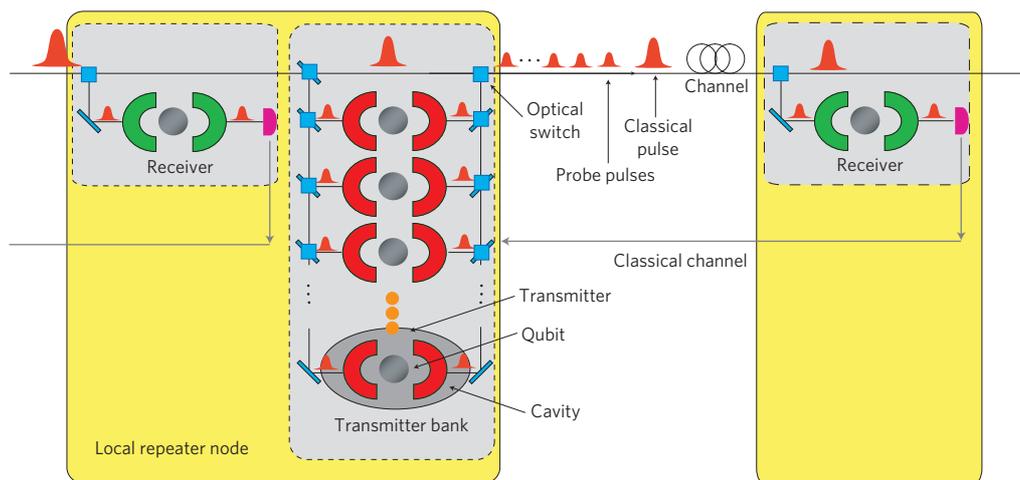


Figure 1 | Schematic of a quantum repeater node and its link to its nearest neighbour. The repeater node is composed of two fundamental components: a bank of transmitters (cavities each with a qubit within them) and a bank of receivers (receiving cavity with a qubit within it and a signal detector). There are generally more transmitters than receivers, but here we consider only a single receiver. The situation where we have a single transmitter and a single receiver in each node is equivalent to existing entanglement-distribution schemes. Two channels are required: a quantum channel to connect the adjacent nodes, which carries a heralding pulse and the quantum-optical signals in the forward direction, and a classical channel to return information about which transmitter was successful.

interaction and measurement procedure is performed. This continues until a success is reported, which triggers two operations. First, any further signals are prevented from interacting with the qubit. Second, a classical message is dispatched back to the first node, informing it as to which transmitter was successful. The time between initiating the transmitters and receiving the classical message is essentially the roundtrip time between nodes and is independent of the probability of successfully entangling a single transmitter and receiver, assuming that the quantum signals can be tightly temporally multiplexed and that local gate operations are sufficiently fast. With enough transmitters we can effectively guarantee that an entangled link is created. The probability of not establishing a link is given by $P_f = (1 - P)^n$, where n is the number of transmitters and P is the probability of successfully entangling a single transmitter and receiver.

With only one receiver, once the entangled link is created we have to discard any further signals. However, we could use a few extra receivers. Once the first receiver in a node has been entangled, the remaining signals are routed to the second receiver in that same node. When that is successful they go on to the next receiver, and so on. With n transmitters and m receivers, the failure probability that all m entangled links have not been created is

$$P_f(m) = \sum_{j=1}^m \binom{n}{j-1} P^{j-1} (1 - P)^{n-j+1}$$

Table 1 | Resources required to generate entangled links.

m	$P_f = 0.1$	$P_f = 0.01$	$P_f = 0.001$
	n	n	n
1	10	20	30
2	17	30	41
3	25	38	50
5	37	53	67
10	67	88	105
50	288	326	356
100	552	604	644

The number of transmitters n and receivers m required to generate m entangled links for various failure probabilities P_f , where we have chosen the probability of successfully entangling a single transmitter and receiver to be $P = 0.2$. In the limit of large m , n asymptotes to $n = m/P$, which implies $n = 5m$.

The numbers of transmitters and receivers needed are quite modest, as shown in Table 1 for $P = 0.20$ and $P_f = 0.1, 0.01, 0.001$. The number of extra receivers required for multiple entangled links to be created simultaneously is modest, even for low numbers of receivers. More importantly, by creating links in one roundtrip time in a near-deterministic fashion, we potentially lessen the requirements on the quantum memories.

Quantum error correction. There are various possibilities regarding how multiple entangled links between adjacent nodes can be used. The simplest is just to use them in parallel to improve the rate of the overall network; however, as the entangled links are unlikely to be perfect, they need to be purified. Normal purification protocols are problematic, however, because they are probabilistic and require two-way communication to determine if one has succeeded or failed^{22–25}. Upon failure the entangled links are destroyed and one must start the link generation again, negating the benefit of our protocol. This can be overcome using quantum error correction^{26,27}. Error correction can be used to purify entangled links while only requiring classical information to be sent one way²⁸.

Which error-correction code is best to use will depend on the type of errors induced during the process of creating entangled links and on the failure rate of the quantum gates at each node. If, for simplicity, we assume that the predominant error, excluding loss, is a bit-flip (X) error and we have perfect local gates, then our entangled link can be represented by

$$\rho(F) = \frac{F}{2} |gg + ee\rangle\langle gg + ee| + \frac{1 - F}{2} |ge + eg\rangle\langle ge + eg| \quad (1)$$

where F measures the fidelity (quality) of the entangled link one is trying to create and $|g\rangle$ and $|e\rangle$ are the two states of the relevant qubits. In this case, to create an entangled link of fidelity $F' > F$ we can use a three-qubit repetition code, which corrects a single X error (depicted in Fig. 2). This conditions three links of fidelity F to a single link of fidelity $F' = F^3 + 3F^2(1 - F)$, up to a known X correction. Because corrections simply update the so-called Pauli frame, they need only be noted and communicated to one end of the network²⁸. This means we do not need to wait, and the transmitters and receivers can be processed and reused immediately.

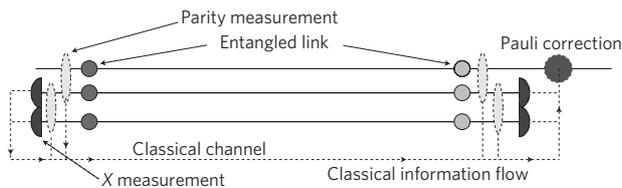


Figure 2 | Schematic of a quantum circuit for performing one-way error correction on imperfect links. In this example, three entangled links of fidelity $\rho(F)$ are conditioned to a single link of fidelity $\rho[F' = F^3 + 3F^2(1 - F)]$. On both the transmitter and receiver sides, non-destructive, two-qubit Z-parity measurements are carried out on the first and second qubits and then on the second and third ones. The second and third qubits on each side are then measured in the X basis. A message is sent from the transmitter side to the receiver side to update the Pauli frame. This simple protocol is quite effective at increasing the fidelity of the remaining link relative to that of the initial link. For instance, with three links of $F = 0.95$ one can create a link of $F' \geq 0.99$ (with five links of $F = 0.9$ one can create a link of $F' = 0.99$ using a five-qubit repetition code).

The general features of this protocol can be retained under more challenging error models. In the case of a general channel error and faulty local gates, to achieve fault tolerance while keeping within the spirit of our design one can replace physical qubits with logical qubits encoded using a Calderbank–Shor–Steane (CSS) quantum code^{29–31}. It is a property of CSS codes that an error syndrome can be obtained by teleporting a logical qubit using an already prepared logical Bell state, where the two logical measurements in the teleportation circuit reveal the syndrome of the code, which is decoded to determine the logical operation required to complete the teleportation³². Error correction by teleportation is particularly suited to communication with quantum repeaters, because teleportation of logical qubits to perform X- and Z-error correction is achieved using the same physical operations that perform entanglement swapping in the repeater network²⁸.

Our protocol creates entangled links in the presence of loss without negatively affecting fidelity. This means it can effectively underly any scheme for fault-tolerant quantum communication. In addition to schemes based on the abovementioned ideas, recent schemes based on surface codes³³ and cluster states may be of interest^{34,35}. Most importantly, by using error correction we can avoid the non-determinism inherent in schemes based on purification, allowing for pipe-lining of the overall repeater network.

Quasi-asynchronous design. A repeater network involves quantum resources, both between and within nodes, and classical resources such as communication between nodes and clocking. Two choices for how such a network could operate are a synchronous design and an asynchronous design. We will focus on the quasi-asynchronous network (Fig. 3a). An advantage of this design is that distances between adjacent nodes need not be the same.

The quasi-asynchronous network begins with the clock in the leftmost node initiating a classical heralding pulse that will propagate along the entire network from left to right. As it goes it will initiate the transmitters to transmit signals to the receivers in the adjacent nodes to the right. The transmitters will be initiated in a temporal progression from the left-hand side of the network to the right-hand side. Each node reports to its left neighbour, via a classical message, which transmitted signals were successful. To extend entanglement beyond neighbouring nodes, whenever a transmitter is entangled to the right and a receiver to the left, entanglement swapping is performed. This creates a longer-range link and frees the transmitters and receivers in that node to participate in creating the next link. While the results of the measurements made during entanglement swapping are available at each node,

no local corrections need to be applied as we propagate this information to the very end of the network along with any information required for error correction. It is important that the heralding pulse associated with the next link arrives at a repeater node after entanglement swapping has been performed, as the herald must be updated with this information. This design allows us to know exactly when the entangled links are ready to use and so we have an efficient, pipe-lined design.

Butterfly design. As entanglement generation is flowing from left to right, the leftmost transmitter and the rightmost receiver become entangled at different times; effectively, the left endpoint of the entanglement link is stationary and the right endpoint is moved to the right by entanglement swapping until it reaches the rightmost node. If this is an issue, a simple solution is to split the network into two halves (Fig. 3b). The actual location of the split depends on the topology of the network, but should be chosen to maximize throughput and to balance the availability of qubits in the leftmost and rightmost nodes. Each side will see a generalized parity for its half of the network. The two halves can be simply connected by entanglement swapping and this information propagated to either the left or right by the next heralding pulse. Here, effectively, the endpoints of the entangled link are moved from the middle node to the leftmost and rightmost nodes by entanglement swapping. Because the transmitters and receivers in the intermediate nodes are disentangled by entanglement swapping, these nodes are free to participate in creating the next link, and we do not need exceptionally long-lived qubits anywhere in the repeater network. This may significantly lessen the technological challenge inherent in distributed quantum-information processing, as the quantum memories now have to be good on timescales associated with the longest roundtrip time between any two adjacent nodes and not the propagation time over the entire network.

Discussion

In the butterfly network the rate of generation of entangled links is a function of the roundtrip time between adjacent nodes and not the propagation time over the entire network. If the distances between nodes vary, then the rate will be limited by the longest roundtrip time between any two adjacent nodes. The rate will be given by $R_f = c/2L$, with c being the speed of light in fibre. This suggests that better rates will be achieved for shorter L , but at the expense of more repeater stations. For practical reasons we want to space the repeaters as far apart as possible. Here we will assume $L = 40$ km, implying that the roundtrip time between adjacent nodes is ~ 400 μ s, giving a rate of $\sim 2,500$ entangled pairs per second. The number of qubits required at each node to maintain this rate without compromising fidelity scales polylogarithmically with the number of nodes. Continuing on from our example of a three-qubit repetition code, a five-node network (total distance, 200 km) requires ~ 40 qubits per node to create pairs of $F' \geq 0.99$ from pairs of $F = 0.95$, where $P = 0.2$ as usual. Increasing the number of nodes to 25 (1,000 km total distance) requires ~ 90 qubits to achieve the same fidelity using a larger repetition code. A more significant, but still polylogarithmic overhead is required to tolerate faulty local gates. For example, for one level of the $[[9,1,3]]$ subsystem code³⁶ we require 109 qubits at each node. The rate can be improved by increasing the resources at each node; increasing the number of qubits by an order of magnitude may increase the rate by approximately two orders of magnitude (see Table 1).

Comparing the performance of our network with others in the literature^{8–18} is difficult, because assumptions about the channel properties, error models and the number of nodes in the network vary. Although networks use different node separations, most aim for a total communication distance of 1,000 km. We will present

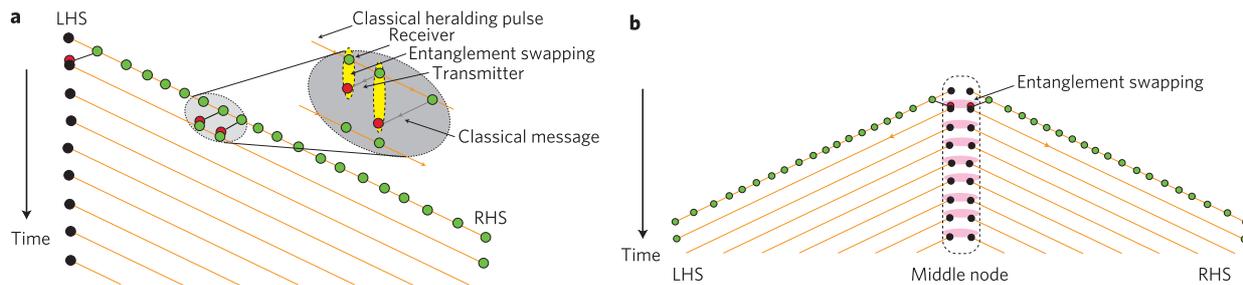


Figure 3 | Quasi-asynchronous repeater network. **a**, A quasi-asynchronous repeater network in which entanglement generation is initiated at the left-hand side where the system clock is located. The classical heralding pulse from the leftmost node propagates to the rightmost node, initiating all of the transmitters as it propagates. Entanglement swapping within local nodes occurs when the local transmitters and receivers have links to their respective neighbours. The leftmost node can start a new entanglement generation cycle after a time equal to the longest roundtrip between any two adjacent nodes has elapsed. The classical heralding pulse for this next cycle picks up the Pauli-frame information from the last cycle as it propagates through the network, and makes it available to the rightmost node as it arrives. **b**, A butterfly network, which creates entangled qubits at the leftmost and rightmost nodes at approximately the same time. This design can be thought of as two quasi-asynchronous repeater schemes (one flowing from the middle to the right and one from the middle to the left). Once links have been established to the left and the right in this middle node, entanglement swapping can be performed. The resulting classical correction information can be sent with the next heralding pulse.

entanglement generation rates as well as the number of qubits used in each node to attempt to give a fair comparison. The DLCZ scheme¹⁵ achieves approximately one entangled pair per 4,000 s (ref. 8) using two qubits (memories) per node. This scheme can be improved by an order of magnitude by performing entanglement swapping by two-photon detection³⁷. Another scheme based on solid-state emitters is that of Childress and colleagues¹³, which uses two qubits (one electron spin and one nuclear spin) per node. It potentially achieves a few pairs per second. Hybrid approaches using bright coherent light^{9–12} generate between 10 and 100 pairs per second using 100 qubits per node. The performance of each of these schemes suffers due to the time delays necessary for the propagation of the classical messages required for purification between nodes. This time increases as the total communication distance increases. In each of these schemes quantum memories must be good for the roundtrip time of the entire network (40 ms) rather than the roundtrip time between adjacent nodes (0.4 ms). In fact, for many, the memories must be good for the time required to generate an entangled link between the end nodes of the network, which could be many times longer. Finally, a recent scheme incorporating error correction²⁸ achieves 100 entangled pairs per second using ~100–150 qubits per node, depending on the error-correction code used.

In summary, we have presented an optimized design for a quantum repeater and its associated use in a network. The key element is a scheme for generating a high-fidelity entangled link between adjacent nodes in constant time. By using error correction instead of purification, the near-deterministic nature of this scheme can be maintained, even with faulty local gates. This allows us to lessen significantly the requirements on the quantum memories, which only need to be sufficient to preserve qubits for the roundtrip time between any adjacent nodes. Then, by using a butterfly design, the end nodes become entangled at roughly the same time, with the generalized parity results arriving one adjacent node roundtrip time later. This allows for an efficient, pipe-lined architecture. For repeater nodes separated by 40 km we could achieve a rate of 2,500 entangled pairs per second with a number of qubits at each node that scales polylogarithmically with the communication distance. With more qubits per repeater node one can achieve megahertz rates. Finally, although we have considered only a linear design, the network topology can be generalized easily.

Methods

Probabilistic entanglement distribution. The core element of any repeater network is the creation of entanglement between neighbouring nodes. This entanglement can

be created between two electron spins placed in cavities at neighbouring nodes with nuclear spins available for quantum memory. The electron- and nuclear-spin systems may be achieved, for example, by single electrons trapped in quantum dots, by neutral donor impurities in semiconductors, or by nitrogen-vacancy diamond centres. For a sufficient interaction between the electron and the light field, the system should be placed in a cavity resonant with the light. Physical mechanisms for entanglement generation between nodes generally fall into one of two categories: first, the heralded creation of high-fidelity entangled links with a low probability of success using single photons or weak coherent sources^{13–17} and, second, the heralded creation of moderate-fidelity entangled links with a moderate to high probability of success using strong coherent fields and homodyne detection^{9–12}. Which approach is better depends on the physical system, but the latter can use the same qubit-photon interaction for local gate operations necessary for error correction and entanglement swapping.

Quantum error correction. As with all error-correction schemes, performance depends on the number and fidelity of the entangled links that are available, the number of qubits at each node, and the target fidelity, which in turn depends on the distance over which we want to communicate. For error correction by teleportation, we require at least enough qubits at each node to fault-tolerantly prepare a logical Bell state as well as enough transmitters and receivers to simultaneously and reliably send and receive a logical qubit. Because logical Bell pairs are required to perform error correction, one approach is to produce many logical Bell pairs at each node, rejecting pairs when errors are detected, so that a high-quality pair is always available when required³². This will yield a scheme that has a high threshold (>1%) for channel and gate errors while retaining the near-deterministic nature of the protocol, but at the expense of a large resource overhead. If resources are limited so that only one logical Bell pair can be prepared and stored at each node at any time, then gate errors will need to occur with a probability of less than $\sim 1 \times 10^{-4}$ (ref. 38).

Received 25 December 2009; accepted 27 July 2010;
published online 29 August 2010

References

- Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge Univ. Press, 2000).
- Spiller, T. P., Munro, W. J., Barrett, S. D. & Kok, P. An introduction to quantum information processing: applications and realizations. *Contemp. Phys.* **46**, 407–436 (2005).
- Dowling, J. P. & Milburn, G. J. Quantum technology: the second quantum revolution. Preprint at <http://arxiv.org/abs/quant-ph/0206091> (2002).
- Spiller, T. P. & Munro, W. J. Towards a quantum information technology industry. *J. Phys.: Condens. Matter* **18**, 1–10 (2006).
- Gisin, N., Ribordy, G., Tittel, W. & Zbinden, H. Quantum cryptography. *Rev. Mod. Phys.* **74**, 145–195 (2002).
- Briegel, H.-J., Dür, W., Cirac, J. I. & Zoller, P. Quantum repeaters: the role of imperfect local operations in quantum communication. *Phys. Rev. Lett.* **81**, 5932–5935 (1998).
- Dür, W., Briegel, H.-J., Cirac, J. I. & Zoller, P. Quantum repeaters based on entanglement purification. *Phys. Rev. A* **59**, 169–181 (1999).
- Sangouard, N., Simon, C., de Riedmatten, H. & Gisin, N. Quantum repeaters based on atomic ensembles and linear optics. Preprint at <http://arxiv.org/abs/0906.2699> (2009) and references within.

9. Van Loock, P. *et al.* Hybrid quantum repeater using bright coherent light. *Phys. Rev. Lett.* **96**, 240501 (2006).
10. Ladd, T. D., van Loock P., Nemoto K., Munro W. J. & Yamamoto Y. Hybrid quantum repeater based on dispersive CQED interactions between matter qubits and bright coherent light. *New J. Phys.* **8**, 184 (2006).
11. Munro, W. J., Van Meter, R., Louis, S. G. R. & Nemoto, K. High-bandwidth hybrid quantum repeater. *Phys. Rev. Lett.* **101**, 040502 (2008).
12. Van Meter R., Ladd T. D., Munro W. J. & Nemoto K., System design for a long-line quantum Repeater. *IEEE/ACM Transactions on Networking* **17**, 1002–1013 (2009).
13. Childress, L., Taylor, J. M., Sørensen, A. S. & Lukin, M. D. Fault-tolerant quantum communication based on solid-state photon emitters. *Phys. Rev. Lett.* **96**, 070504 (2006).
14. Enk, S. J., Cirac, J. I. & Zoller, P. Photonic channels for quantum communication. *Science* **279**, 205–208 (1998).
15. Duan, L.-M., Lukin, M. D., Cirac, J. I. & Zoller, P. Long-distance quantum communication with atomic ensembles and linear optics. *Nature* **414**, 413–418 (2001).
16. Zhao, B., Chen, Z.-B., Chen, Y.-A., Schmiedmayer, J. & Pan, J.-W. Robust creation of entanglement between remote memory qubits. *Phys. Rev. Lett.* **98**, 240502 (2007).
17. Yuan, Z. *et al.* Experimental demonstration of a BDCZ quantum repeater node. *Nature* **454**, 1098–1101 (2008).
18. Goebel, A. M. *et al.* Multistage entanglement swapping. *Phys. Rev. Lett.* **101**, 080403 (2008).
19. Simon, C. *et al.* Quantum repeaters with photon pair sources and multimode memories. *Phys. Rev. Lett.* **98**, 190503 (2007).
20. Tittel, W. *et al.* Photon-echo quantum memory in solid state systems. *Laser Photon. Rev.* **4**, 244–267 (2009).
21. Sangouard, N., Dubessy, R. & Simon, C. Quantum repeaters based on single trapped ions. *Phys. Rev. A* **79**, 042340 (2009).
22. Dür, W. & Briegel, H. J. Entanglement purification and quantum error correction. *Rep. Prog. Phys.* **70**, 1381–1424 (2007).
23. Pan, J.-W., Simon, S., Brukner, C. & Zeilinger, A. Entanglement purification for quantum communication. *Nature* **410**, 1067–1070 (2001).
24. Bennett, C. H. *et al.* Teleporting an unknown quantum state via dual classical and Einstein–Podolsky–Rosen channels. *Phys. Rev. Lett.* **70**, 1895–1899 (1993).
25. Collins, O. A., Jenkins, S. D., Kuzmich, A. & Kennedy, T. A. Multiplexed memory-insensitive quantum repeaters. *Phys. Rev. Lett.* **98**, 060502 (2007).
26. Gottesman, D. An introduction to quantum error correction and fault-tolerant quantum computation. Preprint at <http://arxiv.org/abs/0904.2557> (2009).
27. Devitt, S. J., Nemoto, K. & Munro, W. J. The idiots guide to quantum error correction. Preprint at <http://arxiv.org/abs/0905.2794> (2009).
28. Jiang, L. *et al.* Quantum repeater with encoding. *Phys. Rev. A* **79**, 032325 (2009).
29. Steane, A. M. Error correcting codes in quantum theory. *Phys. Rev. Lett.* **77**, 793–797 (1996).
30. Calderbank, A. R. & Shor, P. W. Good quantum error-correcting codes exist. *Phys. Rev. A* **54**, 1098–1105 (1996).
31. Steane, A. M. Multiple particle interference and quantum error correction. *Proc. R. Soc. Lond. A* **452**, 2551–2577 (1996).
32. Knill, E. Quantum computing with realistically noisy devices. *Nature* **434**, 39–44 (2005).
33. Fowler, A. G. *et al.* Surface code quantum communication. *Phys. Rev. Lett.* **104**, 180503 (2010).
34. Perseguers, S. *et al.* One-shot entanglement generation over long distances in noisy quantum networks. *Phys. Rev. A* **78**, 062324 (2008).
35. Perseguers, S. Fidelity threshold for long-range entanglement in quantum networks. *Phys. Rev. A* **81**, 012310 (2010).
36. Bacon, D. Operator quantum error-correcting subsystems for self-correcting quantum memories. *Phys. Rev. A* **73**, 012340 (2006).
37. Jiang, L., Taylor, J. M. & Lukin, M. D. Fast and robust approach to long-distance quantum communication with atomic ensembles. *Phys. Rev. A* **76**, 012301 (2007).
38. Aliferis, P. & Cross, A. W. Subsystem fault tolerance with the Bacon-Shor code. *Phys. Rev. Lett.* **98**, 220502 (2007).

Acknowledgements

This work was supported in part by European Union seventh framework projects Hybrid Information Processing (HIP), Quantum Interfaces, Sensors, and Communication based on Entanglement (Q-ESSENCE), a Hewlett Packard Laboratories Innovation Research Grant and Japanese grants from the Specially Promoted Research in Grants-in-Aid for Scientific Research funded by the Ministry of Education, Culture, Sports, Science and Technology (MEXT), the National Institute of Information and Communications Technology (NICT) and the Funding Program for World-Leading Innovative R&D on Science and Technology (FIRST).

Author contributions

W.J.M., K.A.H. and K.N. conceived the original entanglement-distribution concept. All authors contributed to the final design of the network. W.J.M. and A.M.S. prepared the manuscript with input from S.J.D., K.A.H. and K.N.

Additional information

The authors declare no competing financial interests. Supplementary information accompanies this paper at www.nature.com/naturephotonics. Reprints and permission information is available online at <http://npg.nature.com/reprintsandpermissions/>. Correspondence and requests for materials should be addressed to W.J.M.