AQUA: Advancing Quantum Architecture
Annual Report 2015

February 7, 2017

Abstract

The AQUA (Advancing Quantum Architecture) working group continued research activities advancing quantum computing and communication, especially quantum networking and distributed quantum computing systems. Our research contributes to planning for the long-term evolution of the computing and networking industries as Moore's Law comes to an end. In 2016, AQUA members published eight papers in journals and workshops on quantum error correction and quantum repeater networks.

1 Introduction

WIDE, through the AQUA working group, is well positioned to participate in and help guide the field in this exciting area, particularly as it moves from theoretical papers and small laboratory technology demonstrations toward actual systems.

This report first discusses recent work in WIDE on quantum networks, then quantum error correction and quantum architecture. This is followed by a summary of 2017's major publications. An introduction to the AQUA group and work areas is included as Appendix A. A brief introduction to the field of quantum information is included as Appendix B.

2 Quantum Networks

A quantum repeater's work consists of three tasks: (1) generation of base-level entanglement with its nearest neighbors, using fiber or free space links; (2) managing errors (via error detection or error correction); and (3) coupling the single-hop entanglement into longer-distance entanglement, e.g. via a method known as entanglement swapping [19]. Experimental and theoretical physicists have worked hard on the physical layer mechanisms for generating entanglement, and theorists have studied means of managing errors while building entanglement along a chain of repeaters, but little energy has been invested so far in designing *networks* of quantum repeaters.

Building on the work done over the last ten years, the work done by AQUA in 2016 and early 2017 has completed our list of provisional technical proposals for almost every aspect of creating a true entanglement-based Quantum Internet above the physical layer.

2.1 Recent Accomplishments

Interoperability between different network types: Quantum repeater networks will be built on a variety of different physical technologies, but from our point of view the bigger concern is whether different types of networks can be made to interoperate at the logical level. Several different types of networks have been proposed. Some depend on *entanglement purification* [13, 12], which detects but cannot correct errors, and some depend on *quantum error correction* [20, 16]. Is it possible to bridge the different types, so that a true Quantum Internet can be built?

In this work, Nagayama *et al.* examined several methods of creating *Bell pairs* (the basic form of entangled quantum state used in quantum networks) that span two different types of networks, and proposed a system that is robust and performs well. This method is optimized for use in quantum routers that sit at the boundary between two networks, and use an optical switch to connect two "line cards" that face into separate types of networks. We believe that this method will serve as the primary workhorse for interoperability in the Quantum Internet.

Analyzing quantum network coding over repeater networks: Like classical network coding [1], quantum network coding (QNC) is a means of utilizing computation in the middle of a network to enhance throughput by alleviating the load on bottleneck links [18, 21]. However, early analyses ignored the detailed process of executing QNC on repeater networks, which Satoh *et al.* remedied in 2012 [29].

In 2016, Satoh *et al.* extended this work to incorporate mixed states (those with some error component), assessing the output fidelity and comparing it to the simpler entanglement swapping approach. They found that QNC results in lower fidelity, and is especially sensitive to local gate error rates due to the larger number of operations, and is therefore most useful when maximizing performance is more important than the output fidelity.

Assessing the assessment of quantum states: Quantum tomography is the complete evaluation of a system's ability to create a particular quantum state, done by recreating the state many times and measuring it in different fashions [4]. In quantum networks, tomography, or perhaps a simplified procedure optimized for Bell pairs, will be needed both at the link level and end to end. In operational networks, the process is complicated by the need to perform tomography on quantum states that are held in distributed fashion, and the need to conduct it in real time as end-to-end connections are created across the network, rather than in batch fashion after a laboratory data collection run. Oka *et al.* created a classical network protocol to support this distributed tomography [25].

The goal of tomography is to create a description of the quantum state known as the *density matrix*, building on knowledge of the state creation process and the experimental measurement results. The reconstruction process involves two phases: the reconstruction itself, followed by evaluation of the likelihood that the observed measurements can be explained by the reconstruction. An important observation is that the latter process is inherently stochastic, hence, two parties both attempting this evaluation may reach different conclusions. In a network, this could result in different operational decisions and failure of the networking protocols, an undesirable result. Thus, the protocol created employs a master-slave architecture.

The authors also evaluated the workload imposed by tomography, and found that it may take several hours. This may be acceptable for bootstrapping a link, but is unacceptable for dynamic connection establishment, forcing us to look for other solutions. This factor will affect the design of connection establishment protocols, and ultimately important details of the entire network architecture.

Entanglement creation via container ship: Existing designs for quantum repeaters are likely to be physically large devices that require substantial maintenance and infrastructure, making them unsuitable for deployment along the ocean floor. Devitt *et al.* recognized that the power of quantum error correction allows us to create Bell pairs on shore, then keep one member of the pair in place while the other is carried via ship across the ocean [10]. Because quantum entanglement is a generic resource not incorporating important data, the latency of the ship is irrelevant, much like the sharing of one-time pads for encryption.

Assessing security and stability of quantum repeater network operations: We have estimated the impact of the hijacking of a single quantum repeater on the work executed by an entire network. Due to the fragility of quantum states, connections across quantum repeater networks are very sensitive to the presence of eavesdroppers, which is well known, but until now no one has assessed whether this fact and the distributed nature of purification and entanglement swapping make quantum networks more vulnerable to serious operational disruption.

Satoh *et al.* have taken a big step by analyzing whether the hijacking of a single repeater gives the hijacker significantly more leverage than the hijacking of a single classical Internet router would [30]. We conclude that networks that utilize purification, which already suffer a logarithmic overhead in work due to the purification, may also suffer a similar logarithmic overhead in the amount of work that can be disrupted by a single hijacker. However, we do not expect that this will make quantum repeater networks significantly less stable than classical networks. Moreover, as we apply quantum tomography throughout the network to monitor the quality of connections, the same tomography, if done in a secure fashion, can be used to scan for the presence of hijacking and assist in the isolation of the hijacked repeater.

This work was preceded by work examining hardware attacks on individual nodes, which resulted in some suggested guidelines for quantum repeater hardware architectures [31].

Optimization of performance and resource consumption for some networks: Creation of entanglement across a link is a probabilistic process, due to the loss of photons at every optical element and interface and through the channel (e.g., fiber). We can divide the success probability into several ranges; in the "high probability" range, we can expect that enough entanglements will be created every round trip across the link that we can use quantum error correction-based methods and effectively optimize the use of memory along the path [38].

Van Meter *et al.* showed that path utilization patterns that minimize memory buffering time (and the corresponding loss of state fidelity due to decoherence) always exist for two of the three major entanglement usage patterns: B class, which includes Bell inequality violation experiments and quantum key distribution (QKD); and C class, which execute a limited range of quantum calculations known as Clifford group operations. Buffering above the link level can be reduced to zero for these classes, whereas T class operations, including full quantum computation and teleportation, require additional buffering to accommodate the propagation of classical signals. T class operations have a range of Pareto optimal patterns on any type of path, but with non-zero buffering. Zero buffering is possible only for one particular link arrangement, discussed in the paper.

Resource demands of Quantum Byzantine Agreement: The demands of quantum key distribution are relatively well understood for single-photon systems, though perhaps somewhat less so for entanglement-based QKD [14, 15]. However, the performance requirements of the broader set of applications of distributed entanglement are less well understood. We can classify applications into three groups, with some overlap: distributed cryptographic functions, wide-area sensor networks, and distributed quantum computation. In earlier years, we have worked to establish performance requirements for distributed quantum computation, in which we concluded that for many applications the performance demands are very high, leaving us very far from networks capable of supporting them [9]. Likewise, some assessments of sensor network applications are very demanding [17].

AQUA's recent work therefore returns to the question of assessing the performance of cryptographic functions other than QKD. Quantum Byzantine Agreement (QBA) appears to have the potential to be an early application of quantum repeater networks [6]. QBA provides a set of security and asymptotic performance characteristics not available in a classical Byzantine agreement protocol. Taherkhani *et al.*'s analysis suggests that nodes with a few hundred qubits, capable of creating several hundred of Bell pairs and performing a few thousand operations before quantum decoherence sets in, can perform the full QBA protocol designed by Ben-Or and Hassidim [32].

2.2 Prior Years

2016's results build on years of prior work quantum network and quantum internetwork architecture [35]. Almost every major network design issue has now been touched upon, though the design choices will evolve continually even after initial deployments. Our overall goal, therefore, is to build a future-proof, flexible architecture that allows indefinite innovation. A few key results are highlighted here.

Internetwork architecture: The Quantum Recursive Network Architecture (QRNA), building on the classical Recursive Network Architecture [33], will scale management of connections, work across heterogeneous networks, and retain autonomy and privacy of network operations [39].

Routing: A modified form of Dijkstra's algorithm will work for intermediate-sized networks [37].

Multiplexing: Assignment of resources to quantum network connections will likely follow circuit switching principles due to the need to continually execute distributed operations, but moment-by-moment allocation of entangled states may work acceptably even when performed in Internet-style best-effort, first-come, first-served fashion [5].

3 Quantum Error Correction and Quantum Computer Architecture

In addition to the work on quantum networks, AQUA members have conducted research on error correction for quantum computers and quantum computer architecture. In our opinion, as well as the opinion of a number of others, the *surface code* represents the most attractive method, encoding a logical qubit in the parity of chains of qubits on a surface [11, 26, 27].

A road map toward scalable distributed architectures: Van Meter and Devitt, in a recent article in a special issue of *IEEE Computer*, discussed surface code quantum computation and the prospects for several quantum information technologies.

Designing a million-qubit quantum computer: One of the hurdles to implementation of quantum computers is the enormous resources required, encouraging distributed architectures, as above. In collaboration with researchers from Duke University, AQUA member Van Meter studied the application and error correction behavior and hardware quality requirements assuming a distributed system of optically interconnected ion traps. Using more traditional error correcting codes, Ahsan *et al.* found that a million-qubit computer may be both achievable and useful [2].

Surface code on a defective lattice: Unfortunately, fabrication of solid-state qubits is expected to be far from perfect for the foreseeable future. Nagayama *et al.* recognized this difficulty, and developed methods to allow the surface code to work around physically defective qubits, then invested millions of hours of CPU time in simulating systems to determine the effectiveness [23]. They found that chips with 90% of the qubits functioning properly would allow construction of large-scale quantum computers.

A more compact form of the surface code: One of the drawbacks of the surface code is the high resource consumption compared to the code distance. Nagayama *et al.* found a new representation on the 2-D surface code lattice that packs logical qubits in about half the space of prior work [24].

4 Future Work

Quantum network security: Work on understanding how to make quantum networks operationally robust against misbehavior (whether or not intentional) is needed as we begin to plan quantum repeater networks.

Impact of tomography on network operations: The tomography work and hijacking framework described above represent the first steps in understanding how the network is monitored in real time operation. A key concern is doing so in Internet-scale interconnected systems, where latencies are high and heterogeneity forces technology-independent data representations and communications.

Connection establishment methods: We now understand reasonably well how quantum connections want to use the network and at the abstract level how to select the resources to use. However, actually identifying and reserving the resources requires a good deal of engineering. A key concern is doing so in Internet-scale interconnected systems, where privacy and autonomy of operation are paramount.

Application analysis: Many more prospective applications of distributed quantum entanglement exist in the literature, but the operational demands they make of quantum networks is still poorly understood.

5 Publications

AQUA members had seven journal papers published or accepted for publication in 2016 and one peer-reviewed workshop paper, several international conference poster presentations, and three additional submissions. The published and accepted papers are:

1. Shota Nagayama, Austin G. Fowler, Dominic Horsman, Simon J. Devitt and Rodney Van Meter, "Surface Code Error Correction on a Defective Lattice," *New Journal of Physics*, to appear, 2017, [23].
2. Shota Nagayama, Takahiko Satoh and Rodney Van Meter, "State Injection, Lattice Surgery and Dense Packing of the Defermation-Based Surface Code," *Physical Review A*, 2017, to appear, [24].
3. Rodney Van Meter and Simon Devitt, "The Path to Scalable Distributed Quantum Computing," *IEEE Computer* 49(9), 31–42, Sept. 2016, [36].
4. Takahiko Satoh, Kaori Ishizaki, Shota Nagayama and Rodney Van Meter, "Analysis of quantum network coding for realistic repeater networks," *Physical Review A* 93(3), 032302, 2016, [28].
5. Shota Nagayama, Byung-Soo Choi, Simon Devitt, Shigeya Suzuki and Rodney Van Meter, "Interoperability in encoded quantum repeater networks," *Physical Review A* 93(4), 042338, 2016, [22].
6. Simon J. Devitt, Andrew D. Greentree, Ashley M. Stephens and Rodney Van Meter, "High-speed quantum networking by ship," *Scientific Reports* 6, 36163, 2016, [10].
7. Takafumi Oka and Takahiko Satoh and Rodney Van Meter, "A Classical Network Protocol to Support Distributed Quantum State Tomography," *Proc. Quantum Communications and Information Technology*, Dec. 2016, [25].
8. Muhammad Ahsan, Rodney Van Meter and Jungsang Kim, "Designing a Million-Qubit Quantum Computer Using a Resource Performance Simulator," *J. Emerg. Technol. Comput. Syst.* 12(4), 39, 2016, [3].

AQUA submitted three more papers on quantum repeater network engineering in January 2017:

1. Rodney Van Meter, Takahiko Satoh, Shota Nagayama, Takaaki Matsuo and Shigeya Suzuki, "Optimizing Timing of High-Success-Probability Quantum Repeaters," preprint `arXiv:1701.04586`.
2. Takahiko Satoh, Shota Nagayama, and Rodney Van Meter, "The Network Impact of Hijacking a Quantum Repeater," preprint `arXiv:1701.04587`.
3. M. Amin Taherkhani, Keivan Navi, Rodney Van Meter, "Resource-aware architecture for implementation of quantum aided Byzantine agreement on quantum repeater networks," preprint `arXiv:1701.04588`.

References

[1] R. Ahlswede, N. Cai, S. Li, and R. Yeung. Network information flow. *Information Theory, IEEE Transactions on*, 46(4):1204–1216, 2000.

[2] M. Ahsan, R. Van Meter, and J. Kim. Designing a million-qubit quantum computer using a resource performance simulator. *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 12(4):39, 2015.

[3] M. Ahsan, R. Van Meter, and J. Kim. Designing a million-qubit quantum computer using a resource performance simulator. *J.*

Emerg. Technol. Comput. Syst., 12(4):39:1–39:25, Dec. 2015.

[4] J. B. Altepeter, E. R. Jeffrey, and P. G. Kwiat. Photonic state tomography. *Advances in Atomic, Molecular, and Optical Physics*, 52:105–159, 2005.

[5] L. Aparicio and R. Van Meter. Multiplexing schemes for quantum repeater networks. In *Proc. SPIE*, volume 8163, page 816308, Aug. 2011.

[6] M. Ben-Or and A. Hassidim. Fast quantum Byzantine agreement. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 481–485. ACM, 2005.

[7] C. H. Bennett, G. Brassard, C. Crépeau, R. Josza, A. Peres, and W. Wootters. Teleporting an unknown quantum state via dual classical and EPR channels. *Physical Review Letters*, 70:1895–1899, 1993.

[8] H.-J. Briegel, W. Dür, J. Cirac, and P. Zoller. Quantum repeaters: the role of imperfect local operations in quantum communication. *Physical Review Letters*, 81:5932–5935, 1998.

[9] C.-H. Chien, R. Van Meter, and S.-Y. Kuo. Fault-tolerant operations for universal blind quantum computation. *J. Emerg. Technol. Comput. Syst.*, 12(1):9:1–9:26, Aug. 2015.

[10] S. J. Devitt, A. D. Greentree, A. M. Stephens, and R. Van Meter. High-speed quantum networking by ship. *Scientific Reports*, 6:36163, 2016.

[11] S. J. Devitt, W. J. Munro, and K. Nemoto. Quantum error correction for beginners. *Reports on Progress in Physics*, 76(7):076001, 2013.

[12] W. Dür and H. Briegel. Entanglement purification and quantum error correction. *Rep. Prog. Phys.*, 70:1381–1424, 2007.

[13] W. Dür, H.-J. Briegel, J. I. Cirac, and P. Zoller. Quantum repeaters based on entanglement purification. *Physical Review A*, 59(1):169–181, Jan 1999.

[14] A. Ekert. Quantum cryptography based on Bell's theorem. *Physical Review Letters*, 67(6):661–663, 1991.

[15] A. Ekert and R. Renner. The ultimate physical limits of privacy. *Nature*, 507(7493):443–447, 2014.

[16] A. G. Fowler, D. S. Wang, C. D. Hill, T. D. Ladd, R. Van Meter, and L. C. L. Hollenberg. Surface code quantum communication. *Phys. Rev. Lett.*, 104(18):180503, May 2010.

[17] D. Gottesman, T. Jennewein, and S. Croke. Longer-baseline telescopes using quantum repeaters. *Phys. Rev. Lett.*, 109:070503, Aug 2012.

[18] M. Hayashi, K. Iwama, H. Nishimura, R. Raymond, and S. Yamashita. Quantum network coding. In *24th International Symposium on Theoretical Aspects of Computer Science (STACS 2007)*, 2007. Arxiv preprint quant-ph/0601088.

[19] M. Żukowski, A. Zeilinger, M. A. Horne, and A. K. Ekert. "Event-ready-detectors" Bell experiment via entanglement swapping. *Phys. Rev. Lett.*, 71:4287–4290, Dec 1993.

[20] L. Jiang, J. M. Taylor, K. Nemoto, W. J. Munro, R. Van Meter, and M. D. Lukin. Quantum repeater with encoding. *Phys. Rev. A*, 79(3):032325, Mar 2009.

[21] D. Leung, J. Oppenheim, and A. Winter. Quantum network communication: the butterfly and beyond. *IEEE Trans. Inf. Theor.*, 56(7):3478–3490, 2010.

[22] S. Nagayama, B.-S. Choi, S. Devitt, S. Suzuki, and R. Van Meter. Interoperability in encoded quantum repeater networks. *Phys. Rev. A*, 93:042338, Apr 2016.

[23] S. Nagayama, A. G. Fowler, D. Horsman, S. J. Devitt, and R. V. Meter. Surface code error correction on a defective lattice. *New Journal of Physics*, 2017. to appear.

[24] S. Nagayama, T. Satoh, and R. V. Meter. State injection, lattice surgery and dense packing of the defermation-based surface code. *Physical Review A*, 2017. to appear.

[25] T. Oka, T. Satoh, and R. Van Meter. A classical network protocol to support distributed quantum state tomography. In *Proc. Quantum Communications and Information Technology*, Dec. 2016.

[26] R. Raussendorf and J. Harrington. Fault-tolerant quantum computation with high threshold in two dimensions. *Physical Review Letters*, 98:190504, 2007.

[27] R. Raussendorf, J. Harrington, and K. Goyal. Topological fault-tolerance in cluster state quantum computation. *New Journal of Physics*, 9:199, 2007.

[28] T. Satoh, K. Ishizaki, S. Nagayama, and R. Van Meter. Analysis of quantum network coding for realistic repeater networks. *Phys. Rev. A*, 93:032302, Mar 2016.

[29] T. Satoh, F. Le Gall, and H. Imai. Quantum network coding for quantum repeaters. *Phys. Rev. A*, 86:032331, Sep 2012.

[30] T. Satoh, S. Nagayama, and R. Van Meter. The network impact of hijacking a quantum repeater. *arXiv preprint arXiv:1701.04587*, 2017.

[31] S. Suzuki and R. Van Meter. Classification of quantum repeater attacks. In *Proc. NDSS Workshop on Security of Emerging Technologies*, 2015.

[32] M. A. Taherkhani, K. Navi, and R. Van Meter. Resource-aware architecture for implementation of quantum aided byzantine agreement on quantum repeater networks. *arXiv preprint arXiv:1701.04588*, 2017.

[33] J. Touch, Y. Wang, and V. Pingali. A recursive network architecture. Technical report, ISI, Oct. 2006.

[34] R. Van Meter. Quantum networking and internetworking. *IEEE Network*, 26(4):59–64, July/August 2012.

[35] R. Van Meter. *Quantum Networking*. Wiley-ISTE, Apr. 2014.

[36] R. Van Meter and S. Devitt. The path to scalable distributed quantum computing. *IEEE Computer*, 49(9):31–42, Sept. 2016.

[37] R. Van Meter, T. Satoh, T. D. Ladd, W. J. Munro, and K. Nemoto. Path selection for quantum repeater networks. *Networking Science*, 3(1):82–95, 2013.

[38] R. Van Meter, T. Satoh, S. Nagayama, T. Matsuo, and S. Suzuki. Optimizing timing of high-success-probability quantum repeaters. *arXiv preprint arXiv:1701.04586*, 2017.

[39] R. Van Meter, J. Touch, and C. Horsman. Recursive quantum repeater networks. *Progress in Informatics*, (8):65–79, Mar. 2011.

[40] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, Oct. 1982.

A What is AQUA?

A.1 Goals

The primary goal of AQUA is to advance the deployment of quantum technologies in the real world, principally by applying known techniques from classical computer architecture, networking and distributed systems to the problems of scalability in quantum systems. This work will both bring new computational capabilities and help ensure that the progress of information technology does not end when the size of transistors can no longer be reduced.

The physical technology on which modern computing systems are built will change dramatically over the course of the next several decades. Beyond the research goals, AQUA also aims to expose the current generation of students to the principles that drive the evolution of computing technology, and the underlying physics of computation, preparing the students for forty-year careers in which they will work with applied physicists and electrical engineers to drive the coming technological revolutions.

A.2 Work Areas

AQUA has current, active work in five areas contributing to distributed quantum computing systems:

- Devices: In conjunction with researchers at Stanford University, RIKEN, and the University of Tokyo we are designing semiconductor-based chips using optically-controlled *quantum dots* and *superconducting flux qubits*.
- Workloads: Although AQUA does not focus on the creation of new quantum algorithms, we do work on how to implement known quantum algorithms efficiently on realizable architectures. We also perform the reverse analysis: to implement a given algorithm, how large and how accurate a quantum system is required?
- Tools: Proper analysis of new ideas in architecture and networks requires software tools for compiling programs and optimizing their mapping to particular systems, as well as physical simulation of quantum devices and effects.
- Principles: We are searching for new principles in quantum architecture and networking, as well as applications of known principles.
- Networks: Large systems must combine multiple devices into one system that can compute collaboratively, as well as share information; we are investigating both system-area and wide-area quantum networks.

Underlying all of these is the critical issue of error management in quantum systems; quantum data is far too fragile to store or compute upon without continuous, active correction. Our primary focus is on the promising surface code error correction, looking for ways to makes its implementation resource-friendly and robust in the face of various system constraints.

B Quantum Concepts

The following is a brief summary of the key aspects of quantum communication and computation that impact network and system architecture.

Qubits. Quantum information is most often discussed in terms of *qubits*. A qubit, like a classical bit, is something with two possible values that we can label zero and one. Unlike a classical bit, a qubit can occupy both values simultaneously, known as *superposition*.

Superposition and measurement. A qubit can represent multiple values in different proportions at the same time, e.g., two-thirds of a "one" and one-third of a "zero". This *superposition* determines the relative probability of finding each value when we *measure* the state. When we measure the qubit, we get only a single classical bit of information (the "one" or "zero") with 100% probability, and the superposition *collapses*.

Entanglement and Bell pairs. Some groups of qubits exhibit strong correlation between the qubits that cannot be explained by independent probabilities for individual qubits. Instead, the group must be considered as a whole, with interdependent probabilities. This phenomenon is known as *quantum entanglement*. A special entangled state known as a *Bell pair* or *EPR pair*, consisting of two quantum bits, figures prominently in quantum communication. Each qubit in the pair has a 50% probability of having a value of 1 and a 50% probability of having a value of 0 when we measure it. Although we cannot predict which will be found, when we measure one member of the pair, the value of the other is immediately determined. This happens independent of the distance between the two members of the Bell pair.

No cloning. As mentioned above, a key restriction of quantum systems is that we cannot make *independent* copies of an unknown state [40]. This makes error correction exceedingly difficult.

Fidelity. The quality of a quantum state is described by its *fidelity*, which is, roughly, the probability that we correctly understand the state – if we ran the same experiment many times and measured the results, how close to our desired statistics would we be? Unfortunately, any physical operation results in a loss of fidelity, gradually degrading the state as we manipulate or even store it. We can counter this by using a form of error correction or detection.

Purification. The form of error detection historically favored in quantum repeater networks is *purification*, which uses minimal resources [8]. It sacrifices some quantum states to test the fidelity of others. There are various purification mechanisms, with different purification algorithms and different methods for determining which states are sacrificed, each with particular tradeoffs.

Quantum error correction (QEC). QEC may be based on classical codes or purely quantum concepts. The primary difficulties are extraction of errors without damaging quantum state, avoiding error propagation, and the increased resources required. (See references contained in [34], [20] and [16].)

Teleportation. Teleportation destroys the state of a qubit at the sender and recreates that state at the destination, teleporting information rather than matter, as explained in Figure 1 [7]. The process uses a Bell pair's long-distance correlation, followed by transmission of a pair of classical bits.

With these basic concepts, we can begin to construct networks. Bell pairs are consumed by teleportation, so one way to organize a network is to create a continuous stream of Bell pairs between source and destination – as long as we identify those sources and destinations, choose paths to get there, and manage the resources along the way.

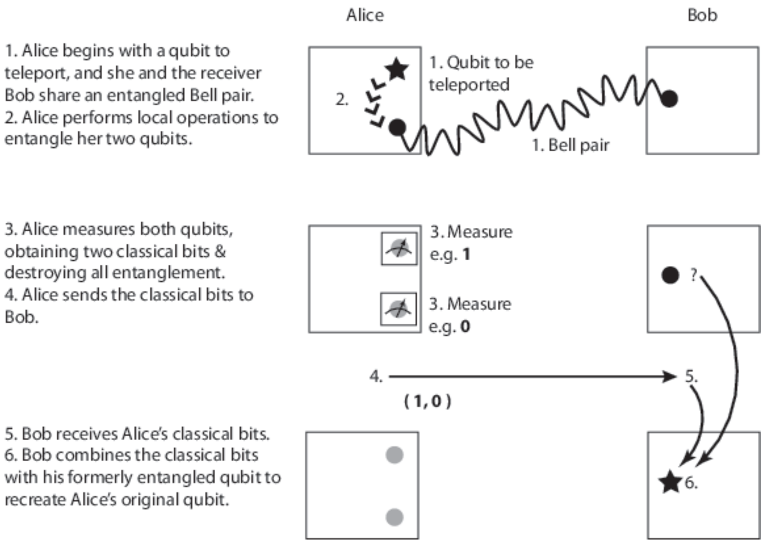

1. Alice begins with a qubit to teleport, and she and the receiver Bob share an entangled Bell pair.
2. Alice performs local operations to entangle her two qubits.

3. Alice measures both qubits, obtaining two classical bits & destroying all entanglement.
4. Alice sends the classical bits to Bob.

5. Bob receives Alice's classical bits.
6. Bob combines the classical bits with his formerly entangled qubit to recreate Alice's original qubit.

Figure 1: Operations in teleporting a qubit from Alice to Bob.