# Quantum Networking and Internetworking

Rodney Van Meter

May 16, 2012

### Abstract

Quantum networks build on *entanglement* and *quantum measurement* to bring new capabilities to communication systems. Quantum physical effects can be used to detect eavesdropping, to improve the shared sensitivity of separated astronomical instruments, or to create distributed states that will enable numerical quantum computation over a distance using teleportation. Because quantum data is fragile and some quantum operations are probabilistic, errors and distributed calculations must be managed aggressively and perhaps cooperatively among nodes. Solutions to these problems will have both similarities to and differences from purely classical networks. Architectures for large-scale quantum networking and internetworking are in development, paralleling theoretical and experimental work on physical layers and low-level error management and connection technologies. With unentangled quantum networks already deployed, entangled networks may appear within the next few years and will form a vibrant research topic in the coming decade.

**Keywords:** quantum communication, quantum networks, quantum repeaters, network protocols, distributed quantum computation

## 1 Introduction

"Teleportation" is a magic word, exotic and evocative, but it has been appearing in serious technical literature with increasing frequency. Both theoretically fascinating and experimentally demonstrated, teleportation is the key to quantum networks [1, 2]. When used in discussions about quantum information, teleportation refers not to Captain Kirk stepping into a machine on the starship Enterprise, dissolving and reappearing on a planet's surface, but to an operation in which a quantum variable dissolves *here* and reappears *there*, on a different physical device. Only the quantum *state* moves; the electron or other physical device remains where it was, and the receiver can in fact be a very different form of physical device than the sender. Classical networks communicate by physically copying data and transmitting the copy, but the rules of quantum mechanics forbid the creation of independent copies of an unknown, arbitrary quantum state. Instead of risking the loss of valuable, fragile quantum data by directly transmitting our only copy, networks will prepare generic states that are used to teleport data, or to perform teleportation-derived operations on the data.

Quantum networks, like classical networks, allow distributed computation, and support the movement of data from place to place. The motivations for doing so are the same for both quantum and classical networks: the desire to connect people, devices such as computers or sensors, or databases that are in separate locations, for technical, economic, political, logistical, or sometimes purely historical reasons. What differs is the type of data and operation involved. Quantum computers, and quantum networks, use quantum variables rather than classical ones; the analogue of the classical bit is the quantum bit, or *qubit*.

Proper use of quantum information opens up new possibilities, making feasible solutions to some problems that are computationally intractable for classical computers (most famously, factoring large numbers), and adding new physical capabilities (most famously, detection of eavesdropping, leading to new, secure, distributed cryptographic key generation mechanisms). Other applications for distributed quantum systems include long-baseline optical interferometry for telescopes [3], high-precision clock synchronization, and quantum forms of distributed tasks such as leader election, Byzantine agreement [4], and coin flipping. Quantum and classical networks and computing systems will hybridize, allowing applications to select the most efficient mechanism for accomplishing a particular function.

Modern work on quantum communications can be said to have begun with Wiesner's quantum cryptography proposal, originated around 1970 [5], followed by Bennett and Brassard's 1984 proposal for *quantum key distribution* (QKD) [6], which utilizes the new low-level quantum capability of eavesdropping detection to build a specific system function, namely the creation of shared, secret random numbers for keying of classical cryptographic systems. However, QKD in its basic form is limited in distance to a few hundred kilometers in optical fiber or perhaps more through free space, and is a single-application system.

Bennett *et al.*'s 1996 proposal for *quantum teleportation* made it possible to move data and execute simple calculations remotely, extending the feasible distance for QKD and vastly expanding the range of conceivable distributed quantum applications. Teleportation involves local quantum operations at each end, and classical messages from the sender to the receiver. It consumes a quantum state known as a *Bell pair* (explained below), shared between the two end points, so a key function of quantum networks is to replenish the supply of distributed Bell pairs as necessary. As with any physical operation, teleportation operates imperfectly, requiring an extensive system that labors to suppress errors. More than a goal in itself, teleportation serves as a building block for distributed quantum applications.

The need to deal with imperfect quantum states and to span multiple hops spurred the development of the concept of *quantum repeaters* [7, 8], which are a vibrant area of research in both experiment and theory. Quantum repeaters, unlike classical repeaters, cannot simply amplify a signal; instead, they support high-fidelity, long-distance quantum communication using teleportation-like techniques over shorter distances and forms of error correction ranging from a simple parity check on a Bell pair to extraordinarily complex, full error correction schemes based on the mathematics of topology. Some repeater architectures manage data movement using computations distributed across all of the nodes in a path between source and destination, while others are more akin to hop-by-hop packet forwarding; the best approach for a given set of physical capabilities remains an important open question. The basics of teleportation and simple forms of error correction have been experimentally demonstrated, and the race is on to build more complete repeaters.

Although QKD networks using trusted relays and optical switches are in use in medium-scale testbeds, the key architectural issues in large-scale repeater networks are only beginning to be addressed. Protocols to actually implement the repeater functionality must be developed. Path selection and resource management, both at the node level, where memory resources are precious, and the network level, including choosing who gets access to the network, will play a role in determining whether the networks actually work.

Beyond single networks lies the issue of internetworking, beginning of course with the ability to recode quantum data from one form to another. Internetworking will require classical sharing of the correct abstraction for describing quantum states or computation requests, and the ability to translate protocols for error management, as well as again the issues of resource management and path selection.

This article provides a tutorial overview of the current state of quantum networking, focusing on *entangled* networks and plans to use them to create global-scale internetworked quantum systems. The basic principles of quantum computation and communication are presented in layman's terms, followed by a discussion of how run-time errors and noise are suppressed in imperfect quantum systems. The major differences between the physical elements of entangled quantum networks – nodes and links – and their classical equivalents are articulated, followed by a brief survey of current experimental work. This article concludes with open problems in large-scale quantum internetworking.

## 2 Quantum Information

To understand teleportation and quantum networks, only a few concepts are required: superposition, measurement, interference, entanglement, and no-cloning.

Quantum computers have attracted interest because they are expected to asymptotically outperform classical computers on some important real-world problems [9, 10]. These gains in capability arise from the differences in storing and manipulating information using quantum states; here, we will restrict our discussion to qubits, though other forms of quantum information are possible. A qubit may be e.g. the direction of spin of a single electron, the direction of polarization of a single photon, or any of a large number of other proposed state variables. Like a classical bit, a qubit has two states, but unlike a classical bit, a qubit may be in a *weighted superposition* of the two states, allowing certain functions to be evaluated for *both* input values at the same time. A register of $n$ qubits can, like a classical register, hold any of $2^n$ possible values. The quantum register can in fact hold a superposition of *all* of these values, and can, in principle, be used to compute on all $2^n$ possible states at the same time.

The difficulty lies in extracting useful answers from a quantum computer. To read the results of a computation, dedicated hardware components *measure* the state of the system. The state of the quantum register *collapses* when the system is measured. It randomly picks one state out of the states that are part of the superposition, based on their relative weights. The other states go away, and it is as if they never existed.

A quantum algorithm manipulates the system to *reduce* the probability of undesirable states, and *increase* the probability of desirable states, until the system has a high probability of measuring the quantum register and getting an answer to our problem, ideally in substantially fewer computational steps than a classical system would require. This is done by creating *interference* on the quantum states to reinforce good answers.

The concept of *entanglement,* in which the states of two or more quantum subsystems are correlated in a fashion that is not possible in classical systems, is the most difficult quantum concept to grasp. Two qubits can be entangled in a continuous spectrum of possible states; four types of entangled states known as *Bell states* are commonly used. One such Bell state is a superposition of the state where both qubits are 0, and the state where both qubits are 1. In this state, when measured, each qubit has a 50% probability of being found in a 0 state and a 50% probability of being found in a 1 state. However, their probabilities are not independent; both values will be found to be the same. It is as if we had two roulette wheels, and while we can never predict what number the ball will fall on, the balls in the two wheels always land on the *same* color. This correlation extends beyond what would be possible simply by synchronizing the roulette wheels; if instead of betting on color we choose to bet on odd/even, we might find that repeated experiments always are even on both wheels, but the color is random and independent!

The Bell states are entangled two-qubit states, and form the generic communication and computation components for most distributed quantum computation. Bell states can be generalized into multi-party *graph states*. Most distributed quantum computing algorithms will build around one or more key flavors of graph state, so the network must be able to create them efficiently.

The final concept required to understand both quantum computation and communication is the *no cloning* theorem. Perfect *independent* copies of an unknown quantum state cannot be made. "Copies" of some states remain entangled with the original state. This entanglement is actually useful in many quantum algorithms, but an unentangled copy would be wildly more useful, allowing faster-than-light communication. It would be, and is, too good to be true.

A major consequence of the no-cloning theorem is that the system cannot copy and send precious quantum data when there is risk of losing the data; loss of the in-transit copy would destroy even the copy kept due to the effects of entanglement and inadvertent measurement. This fact drives the usual quantum networking approach of first building a high-quality, generic entangled state, then using that state to teleport or compute on our valuable data.

Basic teleportation is accomplished by first creating a generic Bell pair between the source and destination. The source entangles the qubit to be teleported with the source's half of the Bell pair, then both qubits are measured, destroying the entanglement of the Bell pair and any superposition state of the data qubit. The measurement results in two *random* classical bits, uncorrelated with the state of the data qubit, which must be transmitted to the destination. Local quantum operations at the destination determined by those classical bits then recreate the original data qubit's state on the remaining Bell pair member. The latency of the classical information transmission prevents information from being transferred faster than the speed of light.

## 3 Imperfect Quantum Systems

The state of a quantum system is exceedingly fragile; both the amplitude of the wave function for a particular value and its phase are subject to both systematic errors and random ones. These errors result in continuous degradation of our knowledge about the state of the quantum register, known as the *fidelity* of the state. As the state drifts from its assigned value, the probabilities of the zero and one states change and the desired effects of inteference may become muted or even incorrect. Beyond these errors that quickly accumulate, isolation of qubits from the environment is difficult, and qubits may be *accidentally* measured, destroying the valuable quantum state.

Various techniques for managing these errors have been developed, some based on classical error correction and erasure correction techniques, others on uniquely quantum approaches [11]. *Purification*, in which two or more multi-qubit states are manipulated to form one higher-fidelity state, uses few quantum memory resources and simple quantum operations, but operates only on generic states such as Bell states. More complete protection of an arbitrary quantum state requires a large number of physical qubits to represent a single logical qubit. This number can range from tens to possibly thousands, depending on the physical memory lifetime, quantum operation error rates, and the performance required to successfully execute a given algorithm.

Besides errors involving the drift of the state, quantum communication systems are also subject to loss in the channel; for those systems expecting to use a single photon, this loss is fatal for that particular operation. Since losses in optical channels tend to be high, any communication system must be designed to manage this loss. Quantum optical states cannot be simply amplified without destroying the entanglement and superposition, so other techniques must be used. Losses in the channel generally force a return message to be used acknowledging success or failure.

## 4 Network Technologies

Quantum networks, like classical networks, will involve nodes and links, and a layered communication architecture with individual protocol modules communicating vertically up and down a protocol stack and horizontally with peers. This section focuses on the physical components, including a brief look at experimental progress on implementation of the lower levels, before the discussion of higher-level networking concepts in the next section.

Quantum communication channels are implemented by sending states of light down a physical channel. These states may be single photons, or other quantum optical states with either large or small numbers of photons. A channel may be a waveguide such as an optical fiber, or free space. It may involve a single transmitter and receiver, or multiple receivers that can individually be enabled or disabled in a shared bus configuration. A link uses a quantum channel and associated classical channel to connect two or more nodes.

A node may have quantum memory that can be used to store a qubit that is entangled with the pulse as it is sent out. When receiving a pulse, a node may either directly measure the pulse using, for example, an avalanche photodiode (APD), or may transfer its quantum state to a memory for later use or analysis. The pulses may come from weak lasers, flourescing atoms, or emission of single photons from a quantum dot, a structure created to exhibit some of the behavior of an atom.

One of the most promising hardware approaches for entangled networks uses microscopic pieces of diamond. When a carbon atom in the diamond lattice is replaced with a nitrogen atom, a positive electrical potential in the lattice capable of trapping a single electron is created. This approach, called *nitrogen vacancy (NV) diamond*, may work at room temperatures, in contrast to most other solid-state quantum systems, which require cryogenic temperatures. Other promising experimental approaches include various forms of quantum dots. Ion traps that hold individual atoms in a vacuum are perhaps the most experimentally advanced approach. Entanglement of up to fourteen qubits in a single trap has been accomplished.

All of these experimental approaches have drawbacks; most do not operate at telecom wavelengths, which will dramatically shorten feasible link distances, though wavelength conversion schemes are also under development. They suffer from short memory lifetimes to differing extents, and the probability of correctly transferring the optical state to the static qubit remains inadequate for reasons ranging from low optical coupling efficiency to basic physics. None of these approaches is ready for mass production, and currently all require hand-tuning and complex experimental setups.

The necessary classical messages include heralding at the physical layer to coordinate timing of the quantum pulses, and many messages for coordinating the higher-level error management, data movement, distributed state creation, and application functionality. Researchers often assume that the classical messages follow the same path through a network as the quantum messages, though except for the physical herald this not strictly necessary. When the classical messaging uses a different network topology, analysis of the communication efficiency must be done with care.

# 5 Quantum Repeaters

The purpose of the technologies just described is to create link-level entanglement. Interesting communication requires extending that entanglement across multiple hops while maintaining adequate fidelity. The *quantum repeater*, building on basic entanglement functionality with purification and teleportation, lays the foundation for quantum networks.

The most obvious method of moving quantum data from place to place is direct hop-by-hop transmission by transferring the qubit state onto a photon and firing that photon down the link. However, as noted above, this places the valuable quantum data at unacceptable risk of loss or corruption. Disappointingly, hop-by-hop teleportation is only marginally better because each hop degrades the fidelity of the data qubit. Alternatively, what about creating a Bell pair at the source, and performing hop-by-hop teleportation on one of the two Bell qubits? This will extend the length of the Bell pair, as shown in Fig. 1. If the qubit being sent is lost before reaching the destination, the Bell pair can be discarded and restarted. Once the qubit reaches the destination, the Bell pair can be used to teleport the important data qubit, without fear of loss. However, the same problem arises: the fidelity of the Bell pair degrades with each teleportation operation, as well as over time if the system keeps the qubit in memory.

[Figure 1 about here.]

One solution is executing purification in a distributed fashion, as in Fig. 2. When purifying Bell pairs, node A holds one half of each of two pairs, and node B holds the other half of each. Using local quantum operations, including measurement of one of the Bell pairs, A and B can probabilistically improve the fidelity of the other. Note that, because purification does not require direct quantum communication, it can operate over any distance, provided the requisite Bell pairs and a classical communication channel are available. The biggest drawback to distributed purification is that it requires that each end convey the results of its local measurement to the far end. Assuming that both nodes can independently identify the set of operations to perform, the minimum time for completion of purification is the one-way classical message latency between the two nodes.

[Figure 2 about here.]

The one-hop-at-a-time extend-purify-extend approach will work, but fails to take full advantage of the fact that *the distributed states being created are generic*, which allows the network to effectively build the Bell pair in parallel. The network can choose to build from both ends of the needed connection, or from the middle; note that the teleportation operation shown in Fig. 1 operates independent of the length of the two Bell pairs. In the late 1990s, Dür, Briegel and their collaborators proposed “nesting” Bell pair purification and teleportation so that the length of entanglement *doubles* in each round, allowing a logarithmic-depth number of rounds to create end-to-end entanglement over a large number of short hops: eight one-hop Bell pairs become four two-hop Bell pairs, then two four-hop Bell pairs, and finally one eight-hop Bell pair. This has become the benchmark approach to repeaters, with much research assuming a power of two number of hops.

This purify-and-teleport architecture is not the only known approach; quantum error correction (QEC) can replace purification. Managed properly, QEC protects the data more completely, reducing the need for multi-hop purification and its associated need for round-trip delays. These advantages come at the expense of substantially more memory and computation resources at each node. Comparison of these basic choices continues in the research literature.

# 6 Quantum Network Architectures

These physical entanglement, end-to-end state building/transfer and error management technologies will get us to laboratory-scale demonstrations, but do not form a complete, deployable network architecture. While in simple arrangements the decisions that a repeater must make for each data request are straightforward, real-world networks must accommodate heterogeneous links arranged in complex topologies, competing traffic sources, ongoing network events such as node/link up/down, and autonomous management. An architecture for a large-scale network must support independent decision making by the nodes in a manner that will result in robust, efficient operation of the network as a whole.

A complete architecture must specify:

- a model for the requests themselves;
- a state propagation mechanism for fulfilling those requests;
- an error management technique or set of techniques;
- an approach to managing dynamic consumption of resources, for both individual requests and the network as a whole; and
- a means of managing the network itself.

The request model depends first on what the network is designed to *do*. In classical networks, traditionally the network layer only sends data, via unicast, multicast or broadcast, with other functions delegated to higher-layer protocols. To support the applications mentioned in the introduction, a quantum network can operate in one of three modes: it can teleport data from place to place, it can execute certain computational operations over a distance (a technique known as *teleporting gates*), or it can create distributed quantum states. Each of these options results in a different form of contract between the requesting end node and the network, some more akin to Active Networks or Named Data Networking than to IP. Put differently, what should the semantics of a network message be, and what does a quantum socket look like?

Perhaps the most fundamental operation, at the network layer, would be creation of high-fidelity distributed Bell pairs, which alone are adequate for building more complex distributed states, teleporting data, or executing remote operations. Conservative engineering practice would suggest that, as with operating system APIs and IP packet semantics, simple is best, favoring this as the lone network-supplied operation. However, data movement as a primitive may improve performance by operating more asynchronously. Providing remote computation requests, or a richer set of state-creation services in the network, may reduce application complexity or improve overall system performance by reducing the total number of operations that must be performed. With appropriate network protocols, all three modes can be mixed in the same network.

The *quasi-asynchronous* scheme proposed by Munro *et al.* [12] illustrates the impact of this design decision. One proposed use of this model begins with the creation of a Bell pair at a node halfway between the source and destination, then propagates the member qubits toward the endpoints in a hop-by-hop teleported fashion. For this to be realizable, an application must have a way to request that the node in the middle create the Bell pair and propagate it outward; equally critically, some means of identifying that middle node must exist, a decidedly non-trivial problem in an Internet-scale network. A more generalized form of distributed state creation also requires complex decisions. This may, of course, be done at the application level among a set of cooperating end nodes, but may be better done by the network itself.

The request model and state propagation scheme, whether data teleportation, remote computation request, or distributed state creation, are intimately connected to the choice of error management mechanism. Purification forces high-latency delays as the required bidirectional messages are transferred. The quasi-async model depends on reduction in bidirectional messaging, achieved by application of QEC. While the trend in theoretical work is away from purification due to its long memory lifetime requirements and high-latency operations, the QEC-based schemes require amounts of memory and numbers of quantum operations that remain prohibitively large. Recently, researchers have been expending considerable effort on circumventing these restrictions, using readily-available resources such as light pulses and classical messages to optimize use of scarcer resources, especially long-term quantum memory. There is as yet no clear winner, and the various alternatives may all be implemented as the technology evolves.

Ultimately, we have identified five types of protocol layers directly involved in each data request: physical entanglement, link entanglement control, error management, quantum state propagation, and application. Defining and optimizing the messages and semantics for these mechanisms forms a core topic of research in quantum networking.

In addition to the layered protocols involved in each data request, the architecture must specify resource management for data requests. Circuit switching is the most obvious approach, especially as fragile quantum memories impose fairly stringent real-time constraints. However, recent work suggests that time division or statistical multiplexing may raise the aggregate throughput of the network if implemented carefully.

Because the states being created are generic, an interesting dynamic resource assignment problem arises: does an entangled Bell pair between two nodes "belong" to a specific end-to-end request, or are all Bell pairs "up for grabs"? How are the operations to fulfill a request organized? Again, this issue is affected by the choice of message semantics.

Finally, the architecture must provide some guidance concerning management of the network itself, most obviously routing protocols. Quantum networks differ from classical enough that classical approaches to such issues require adaptation and reevaluation of fitness for the tasks at hand.

The need to understand the network topology manifests itself both in the choice of path through the static, physical topology, and in the dynamic operation of the network, as in Fig. 3. Each node needs enough information to make decisions that will be consistent with those made by other nodes. Fig. 4 shows one possible case that protocol designs must be careful to avoid.

[Figure 3 about here.]

[Figure 4 about here.]

Beyond a single, technologically homogeneous network architecture lies internetworking, with autonomy, heterogeneity, and sheer scale as key driving forces. A recursive network architecture can be viewed as a natural fit for quantum repeater internetworks, especially when coupled with the distributed computation model, which requires only modest extension above the distributed state model to support fully universal distributed computation [13]. The services each network must provide to its peers, guidance on the use of management information such as routing tables, and the decision-making for in-network modification of requests all flow smoothly from a recursive architecture. Direct solution of the problem in Fig. 3 in an non-recursive internetwork is often infeasible; recursion can reduce the topological complexity and retain network autonomy and privacy, making such problems tractable.

In an internetwork, the issue of naming is especially critical. The intermediate states necessary for fulfilling an application request will cross the boundaries between end and transit networks. An internetwork must virtualize the names of both physical resources and in-process states, so that a message can uniquely identify a state and request that a node or network perform certain actions, preserving operational autonomy and privacy for individual network operators and minimizing the knowledge each node must have about the internetwork as a whole.

The SECOQC QKD network in Vienna, which does not yet build multi-hop Bell pairs, uses a network architecture centering around three *planes*, with the quantum plane at the bottom, the secrets plane in the middle, and the data plane at the top, with each plane consisting of several protocols [14]. Each higher plane provides an overlay topology, perhaps distinct from the physical topology. The planes include transport and routing protocols as necessary. The data plane uses keys generated below to encrypt classical data, providing encrypted virtual links over which protocols such as IP can be run. Thus, the architecture is focused on the specific task of key generation to support secure classical communication, rather than a broader class of distributed quantum tasks.

## 7 Conclusions

Implementations of QKD are well beyond the experimental phase [15, 6]. A few commercial products are available, and metropolitan-area testbed networks exist in Boston, Vienna, Geneva, Barcelona, Durban, Tokyo, several sites in China, and elsewhere throughout the world; several of these networks incorporated technologically disparate link types from the beginning, although the coupling at intermediate nodes is strictly classical. In contrast, entangled networks remain rudimentary, existing only at small scales in laboratories, and have yet to demonstrate all necessary functionality in a single experiment [8].

For entangled networks, by far the most important ongoing research is on the physical layer – if quantum memories and local operations do not reach sufficiently high fidelities, quantum networks will remain a laboratory exercise. Applications for distributed quantum states, whether numerical computation or sensor networks, will drive the need for quantum networks; without them, no one will buy and deploy quantum networking equipment. Both of these areas are being addressed in depth, the first by experimental physicists, the second by theorists in both computer science and physics. It bears pointing out that the performance required for some of these applications remains several orders of magnitude beyond even optimistic hardware predictions for the next several years.

To bridge the considerable gap between theoretical large-scale, wide-area applications and small-scale experiments, an overarching network architecture and matching protocols must be developed. These protocols must emphasize optimized use of quantum memory, both spatially, by reducing the number of qubits that must be stored, and temporally, by reducing the length of time a qubit must be held in an intermediate state, e.g. by eliminating round-trip messaging where possible. The real-time factors in physical memory decay, or the high resource requirements of error correction-protected memories, must be managed properly. Moreover, the form of requests within the network is critical to efficiency. Improvements in the request model can alter the demands on the size, quality and capabilities of the physical system. Ultimately, the problems exemplified in Figs. 3 and 4 are matters of giving each node enough information to make high-quality, autonomous decisions. The design of robust, efficient classical protocols usable in multi-user, multi-technology quantum internetworks will demand the technical skills of the data networking community.

The Quantum Internet, once realized, will allow us to exploit entanglement over long distances for new computational capabilities and for new physical capabilities such as eavesdropping detection. More speculatively, we can imagine uses such as sensor networks for quantum-enhanced telescopes and tests of the fundamental correctness of quantum mechanics. Within a few years, quantum networking and teleportation will move out of the physics laboratory and into the network engineering domain, offering some of the most exciting and intellectually challenging research and development topics of the coming decade.

## ACKNOWLEDGMENTS

This work was supported by JSPS KAKENHI 21500020.

## Author Biography

Rodney Van Meter holds degrees from the California Institute of Technology, the University of Southern California, and Keio University. His research interests include storage systems, networking, and post-Moore's Law computer architecture. He has held positions in both industry and academia in the U.S. and Japan. He is now an Associate Professor of Environment and Information Studies at Keio University's Shonan Fujisawa Campus.

## References

[1] Nicolas Gisin and Rob Thew. Quantum communication. *Nature Photonics*, 1:165–171, March 2007.

[2] H. J. Kimble. The quantum Internet. *Nature*, 453:1023–1030, June 2008.

[3] D. Gottesman, T. Jennewein, and S. Croke. Longer-baseline telescopes using quantum repeaters. *Arxiv preprint arXiv:1107.2939 [quant-ph]*, July 2011.

[4] 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.

[5] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, January 1983.

[6] D. Dodson, M. Fujiwara, P. Grangier, M. Hayashi, K. Imafuku, K. Kitayama, P. Kumar, C. Kurtsiefer, G. Lenhart, N. Luetkenhaus, et al. Updating Quantum Cryptography Report ver. 1. *Arxiv preprint arXiv:0905.4325*, 2009.

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

[8] N. Sangouard, C. Simon, H. De Riedmatten, and N. Gisin. Quantum repeaters based on atomic ensembles and linear optics. *Reviews of Modern Physics*, 83(1):33, 2011.

[9] T.D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J.L. O'Brien. Quantum computers. *Nature*, 464:45–53, March 2010.

[10] Dave Bacon and Wim van Dam. Recent progress in quantum algorithms. *Communications of the ACM*, 53(2):84–93, February 2010.

[11] Simon J. Devitt, Kae Nemoto, and William J. Munro. Quantum error correction for beginners. arXiv:0905.2794v3 [quant-ph], September 2011.

[12] WJ Munro, KA Harrison, AM Stephens, SJ Devitt, and K. Nemoto. From quantum multiplexing to high-performance quantum networking. *Nature Photonics*, 2010.

[13] Rodney Van Meter, Joe Touch, and Clare Horsman. Recursive quantum repeater networks. *Progress in Informatics*, (8):65–79, March 2011.

[14] M Peev et al. The SECOQC quantum key distribution network in Vienna. *New Journal of Physics*, 11(7):075001 (37pp), 2009.

[15] Chip Elliott, David Pearson, and Gregory Troxel. Quantum cryptography in practice. In *Proc. SIGCOMM 2003*. ACM, ACM, August 2003.

# List of Figures

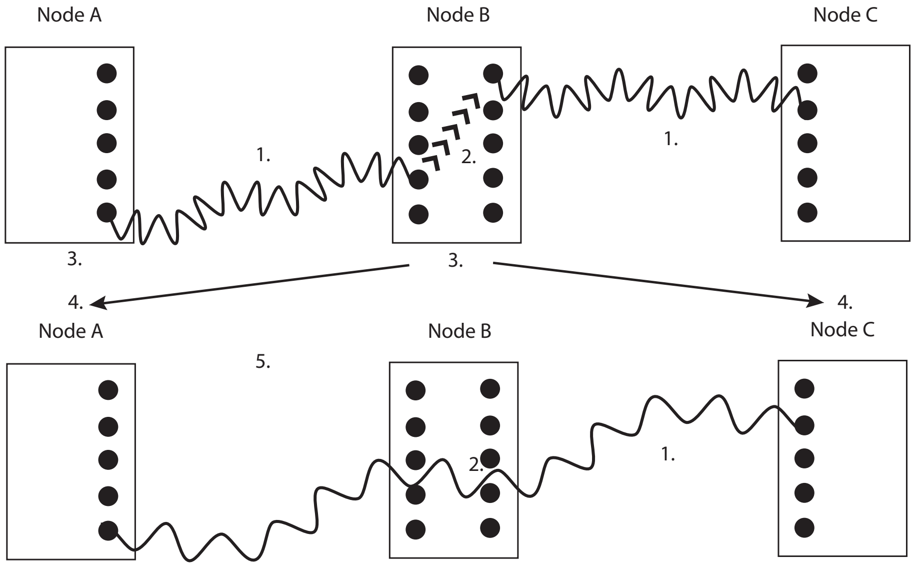

1. Nodes begin with two entangled pairs, AB and BC.
2. Node B selects pairs to teleport, performs local operations, measures one qubit of each pair.
3. B communicates measurement results and new entanglement status to A and C.
4. Receive partner's measurement result and new entanglement status, including node/qubit addresses.
5. Result is single lower-fidelity, longer-distance Bell pair.

Figure 1: Teleportation can lengthen one Bell pair using another.

1. Nodes begin with two entangled pairs.
2. Select pairs to purify, perform local operations, measure one pair.
3. Communicate measurement result.
4. Receive partner's measurement result, decide to keep or discard.
5. Result is either single higher-fidelity Bell pair, or no entanglement.

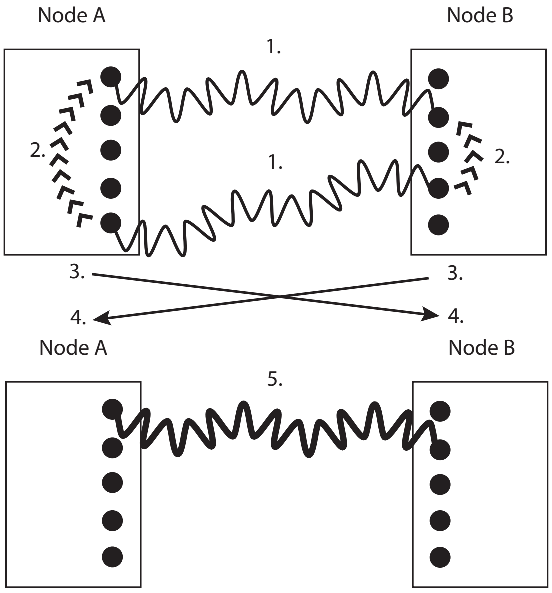

Figure 2: Steps involved in purification of Bell pairs.

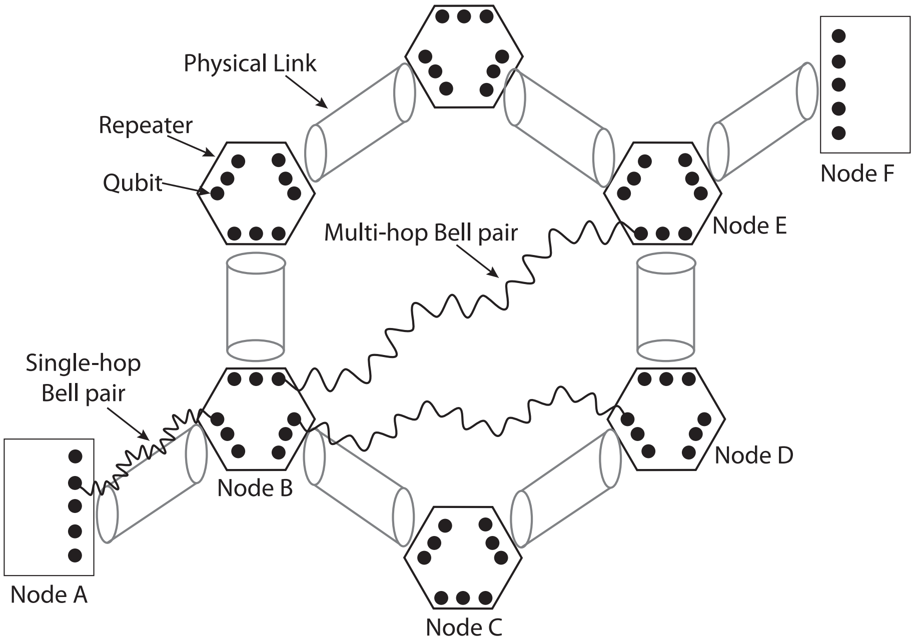

Figure 3: Even when Node B knows that A is trying to build a Bell pair with F, B may be uncertain whether its Bell pair connected to Node D or Node E is “closer” to the destination.

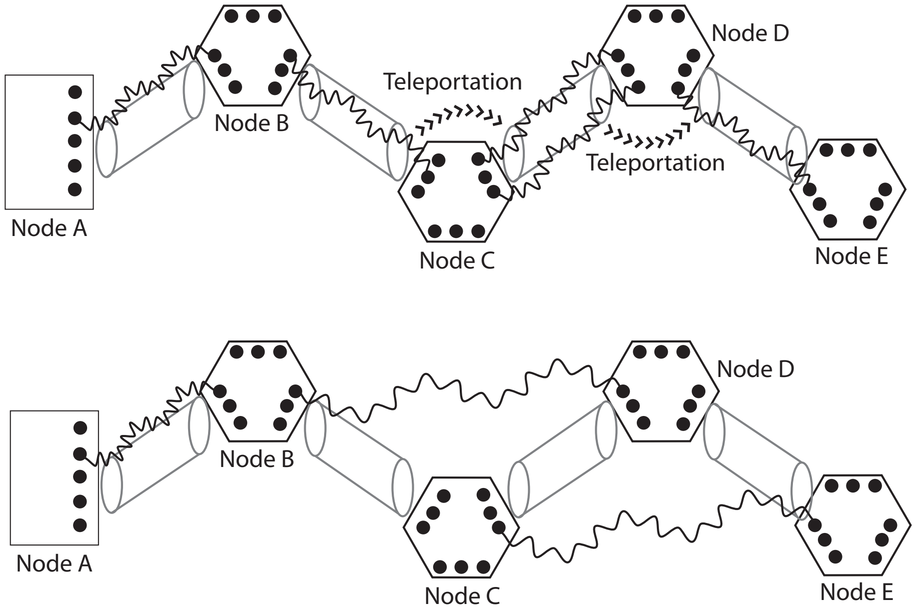

Figure 4: Conflicting teleportation choices by Nodes C and D in the top figure may result in Bell pairs "leapfrogging" each other, as in the bottom figure, leaving no easy path to connect A to E.