# 第３部

# 特集3 分散型量子計算のネットワーク応用技術

Rodney Van Meter

## 第１章　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 2013, AQUA members published four papers in top-tier journals on a new means of executing logical gates on top of the favored surface code error correction mechanism; quantum computer architecture; and quantum repeater networks.

## 第２章　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 on general quantum computation. This is followed by a summary of 2013's major publications [35, 39, 42, 43]. 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.

## 第３章　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 [45]. 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 [15]. 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 [22, 24, 38].)

**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 3.1.[12] 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

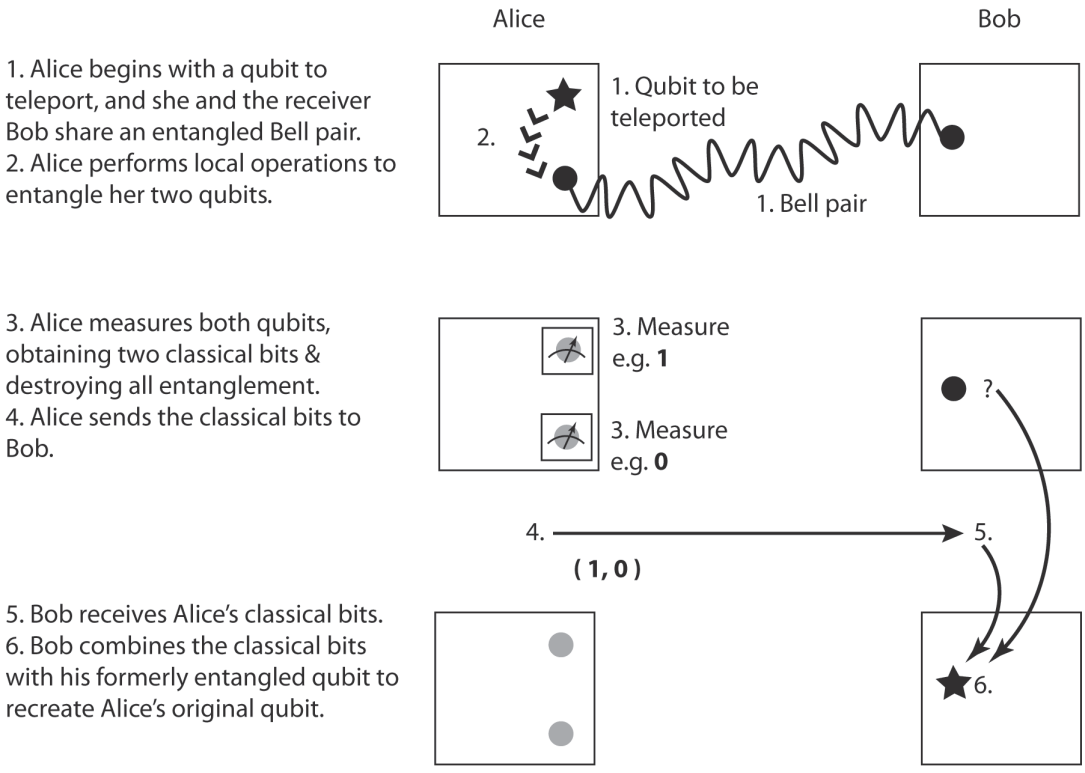


Figure 3.1 Operations in teleporting a qubit from Alice to Bob.

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.

## 第４章　Quantum Networks

In 2013, AQUA members published two papers on quantum networks: an overview of the field in the widely-read magazine *IEEE Communications*[43] (readership 45,000), and a paper establishing a form of Dijkstra's algorithm as a feasible approach to routing in quantum repeater networks in *Networking Science* [42]. Here, we first explain the basic idea of purify-and-swap repeaters, then turn to a more complete description of a quantum network architecture.

### 4.1　Purify and Swap - Early Quantum Repeaters

In the late 1990s, researchers recognized that teleportation can be used to extend entanglement, and that purification could be used to detect errors introduced in the process [15]. This first architecture we will call *purify and swap*, though the originators called it *nested purification*.

The process of *entanglement swapping* uses teleportation to splice two Bell pairs shared spanning adjacent short distances into one pair over the corresponding longer distance. If node *A* shares a Bell pair with node *B*, and node *B* shares another Bell pair with node *C*, then node *B* can teleport its member of the *A* ↔ *B* pair to node *C* using the *B* ↔ *C* Bell pair. In the process, the *B* ↔ *C* pair is consumed, and at the end we have a single *A* ↔ *C* Bell pair.

Enganglement swapping is independent of the distances between *A* and *B*, and between *B* and *C*. Only local quantum operations are required, supported by classical communication. We combine one-hop Bell pairs into two-hop Bell pairs, then combine two-hop pairs into four-hop pairs and so on, doubling the length of the remaining pair at each step, as shown in Figure 4.1.

To compensate for the errors introduced, purification is used, as shown in Figure 4.2: local quantum operations are performed at both nodes on two Bell pairs, then one of the Bell pairs is measured. The measurement results are exchanged and compared. If they agree, the pair's fidelity has improved, and it is kept for reuse. If the measurement results disagree, the pair

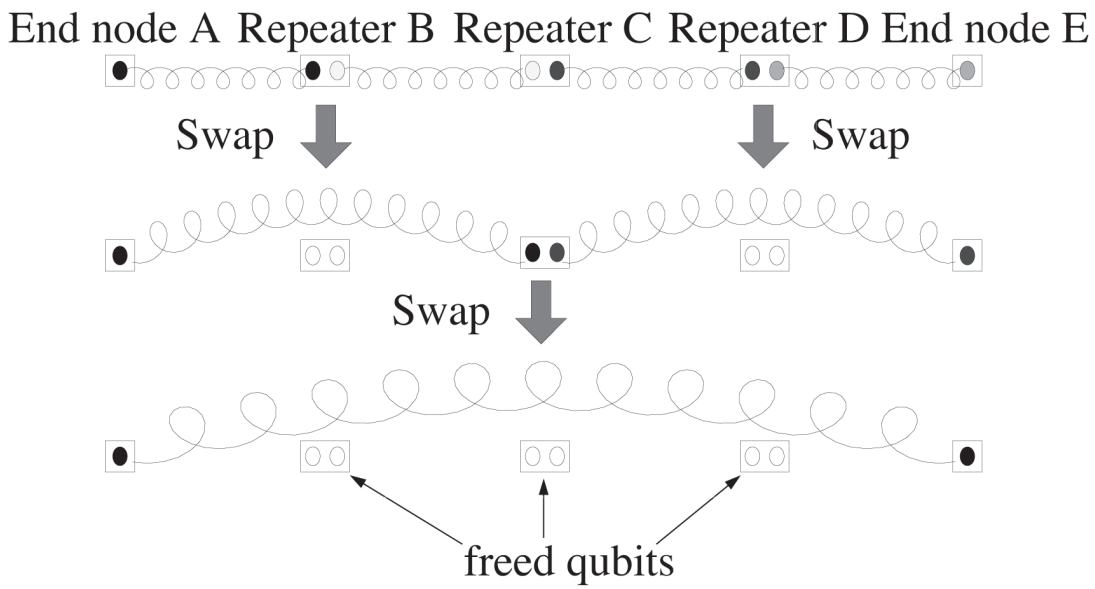


Figure 4.1　Entanglement swapping leverages teleportation to lengthen entanglement distances. Four one-hop Bell pairs become two two-hop pairs, then one four-hop Bell pair.

is discarded.

*Purify-and-swap* is the combination of these two concepts, interleaving purification with entanglement swapping. When purification is performed over one hop, then two hops, then four, resulting in a recursive, interleaved power-of-two approach called simply *nested purification*. The principles have proved to be flexible, so we refer to the entire group of specific designs as the *purify and swap session architectures*.

## 4.2 A Quantum Network Architecture

The design philosophy of our quantum network architecture is inspired by the Internet architecture, leveraging it as much as possible, except where modifications are absolutely necessary to distribute quantum state. As we noted earlier, there are two dimensions of an architecture: vertical layered communication and horizontal distributed group communication. Here, we describe layering in terms of the model we have developed [7], and group communication in terms of our Quantum Recursive Network Architecture (QRNA) [44]; these are presented in the following two subsections.

### 4.2.1 Vertical: Layered Quantum Communication

Layered communication describes how protocol functions are vertically composed within a communications node to provide increasingly complex capabilities. Layered quantum communication relies on five key vertical layer functions that are uniquely quantum.

**Physical layer.** We rely on a quantum physical layer using light to encode quantum state. Many technologies for this layer are under development.

**Link-level entanglement.** We rely on existing techniques to support entanglement across a link. Because most physical entanglement mechanisms are probabilistic, the link layer will include an acknowledgment to the sender indicating which attempts succeeded.

**Remote state composition.** In the Internet, links are composed by copying packets from one link to the next. In a quantum network, links are less readily composed due to the no-cloning theorem. Quantum paths thus either establish end-to-end entanglement from entangled links, or use that entanglement to

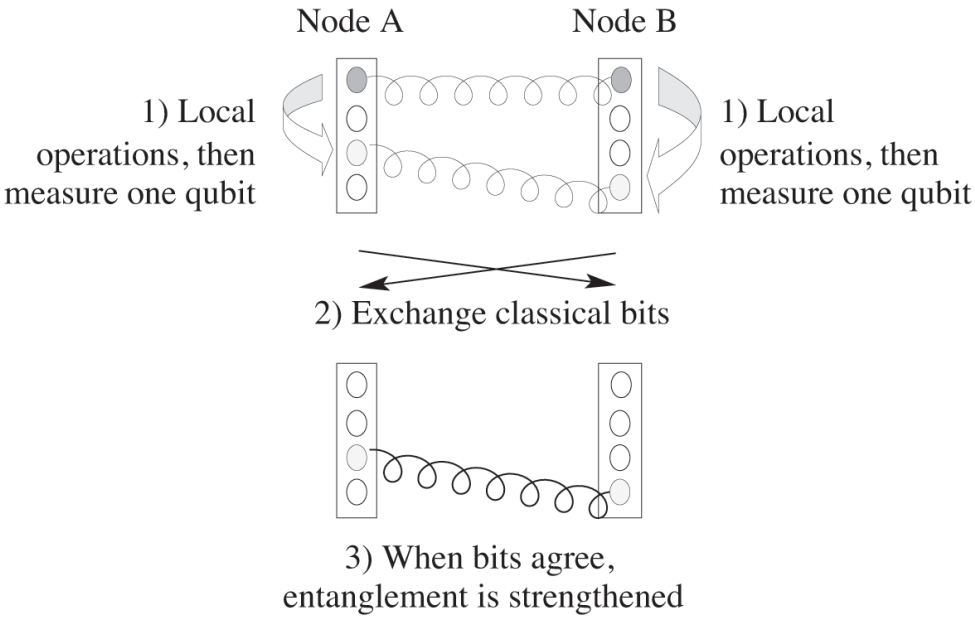


Figure 4.2 Two-node purification, converting two lower-fidelity Bell pairs into one higher-fidelity one. Note that this works independent of the Bell pairs' length.

teleport quantum state from one end to the other. This layer is very sensitive to the link-layer capabilities, as well as the error management mechanism.

**Error management.** In the classical Internet, errors are managed using redundancy (e.g., forward error correction) or error detection and retransmission. As noted earlier, the no-cloning theorem prevents straightforward use of either of these mechanisms. The fidelity of quantum states is critical in reducing the need for error management.

**Application.** The application may be a sensor network or a numeric computation or decision algorithm based on shared state [37]. The application will determine if end-to-end entanglement is required, or if our quantum states can be measured on a pay-as-you-go basis. Some applications may also desire quantum states other than Bell pairs, including any of several common forms of three-party or larger states. Of course, the application is driven by a classical program, presumably using a socket-like data structure.

**Composing Quantum Links -- Purify-and-Swap** To compose links, we use recursive (nested) purify-and-swap along a path. In classical networks, composition is a matter of copying and separately applying error correction and control. In purify-and-swap, composition and error management are sometimes viewed as an integrated operation, but in our layering they are natively distinct operations.

With this architectural background, let us return to the canonical purify-and-swap approach. Figure4.3 shows a five-node example. The physical and link layers are the two layers at the bottom, labeled Physical Entanglement (PE) and Entanglement Control (EC), respectively. The key feature in the communication session architecture is the recursive nature of the error management and remote state composition layers, which in purify-and-swap we call Purification Control (PC) and Entanglement Swapping Control (ESC), respectively. In this example, purification is run over individual links, two-hop entanglement, and finally four-hop entanglement. Entanglement swapping is run at all intermediate nodes, first at *B* and *D* to create two-hop entanglement, then after purification at *C* to create four-hop entanglement. One characteristic of this nested approach is that the end nodes of an *n*-node path must communicate with $log_2\ n$ other nodes along the path, which has implications for the path selection and composition mechanisms.

We can observe that entanglement swapping can be thought of as the middle node taking a Bell pair qubit from its left and teleporting it to the right using another Bell pair. As an example, node *B* in the figure in theory need only communicate with *C*. *A*'s role in the process is entirely passive. However, as the goal is to create end-to-end entanglement, *A* must participate; after entanglement swapping, the next operation may be purification, another swapping operation, or transfer of control to the application, any of which requires *A*'s involvement.

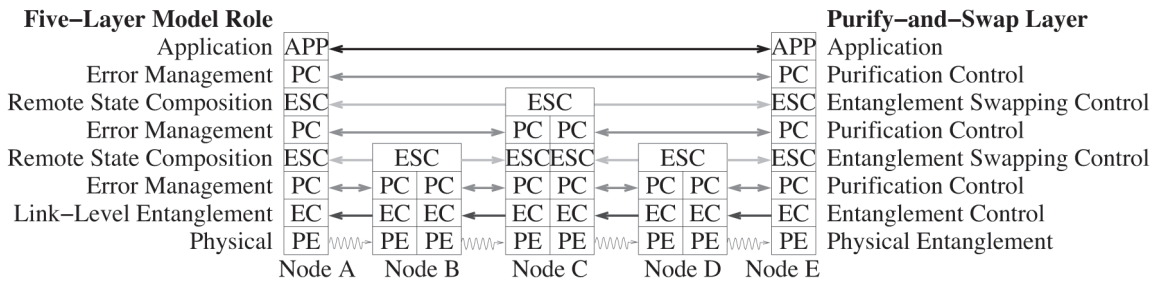


Figure 4.3 Protocol layers and their interaction in purify-and-swap repeaters, in a five-node, four-hop chain. The left labels indicate the model layer represented, and the boxed labels and right labels indicate the protocol name for purify-and-swap repeaters. Double-headed arrows indicate bidirectional classical communication is required. Only the physical layer is quantum, shown propagating left to right.

**Other aspects of quantum layering** Organizing the layering is the first step in developing the key purify-and-swap insight described in Sec.4.1 into a functional, robust, distributed implementation. To be practically implementable, details of the management of requests must be defined. Our approach to doing so is to use protocol state machines to govern the memories themselves [7]. One important facet of this problem is management of the Bell pairs to maximize the end-to-end success rate of purification and swapping, which in turn affects the overall system throughput. We call this the *purification scheduling* problem. Minor extensions are also required when a path is not a power of two hops long.

### 4.2.2 Horizontal: Distributed Quantum Communication

Distributed group communication describes how protocol functions are horizontally composed across different communication nodes. Distributed quantum communication extends this principle to quantum communication by explicitly managing distributed state through the use of recursive composition.

As with the Internet, our architecture composes links into paths, manages state and errors, and supports applications. This section introduces the key differences that result in quantum networks having different architectures from classical: type of information, state management, path composition, and identifiers.

Our Quantum Recursive Network Architecture (QRNA) [44] provides a general-purpose request mechanism abstracted from underlying layers, to accommodate any of the models presented above. Rather than explicit state transfer, it supports requests for creation of distributed states (including both two-party and multi-party states) and operations on those states. Requests may be recursively decomposed and distributed throughout the network in order to build the end-to-end state requested by an application. A link in QRNA may be a physical link or a recursively organized network. QRNA uses globally unique identifiers that represent the locations where the shared quantum state will be established; the structure of these identifiers affects how paths are determined, but is outside the scope of the architecture. Paths are constructed by classical means prior to communication. Table4.1 compares QRNA with the Internet architecture.

Table 4.1 Comparison of the Internet with the Quantum Recursive Network Architecture.

| | Internet | QRNA |
|---|---|---|
| Comm type | message transfer | distributed state |
| Info | byte sequences of varying length | entangled state |
| Link | simplex 2-party or broadcast N-party | 2-party with entangled state supporting quantum teleportation |
| Identifier | global fixed-length | global, any that supports the path composition algorithm |
| Path When selected | at node traversal | before communication |
| Composing hops into paths | intra-network routing chooses hops to transit; inter-network routing chooses networks to transit | same as for classical networks, but with more complex metrics |
| Relay | transmission of copies of classical bits, sent to neighboring node selected via longest-pre.x address match | recursive composition of operations on entangled quantum state distributed along path |

**Information Type and Group Communication** A quantum network architecture can be organized to present either the generation of distributed entangled state or the relocation of quantum state as the fundamental communication semantics. Relocation, using either direct transmission or simple teleportation, may seem easier and more natural, but distributed state generation natively supports a broader range of applications.

State relocation across a network would be sufficient for some applications. One-way teleportation from a client to a server is sufficient for *universal blind quantum computation*, in which the server is oblivious to the computation it performs for the client. State relocation also appears to extend smoothly from unentangled networks. Applications that need simultaneous, long-distance entangled states must build them, because state relocation does not provide entangled states. State relocation doesn't demand long-lived memory unless the session architecture itself does, but it also cannot easily take advantage of resources in the middle of the network to operate more efficiently.

Distributed state generation supports a more general distributed computation model. It works well with both two-party and multi-party entangled states. However, in the basic form it requires long-lived memory.

*Asynchronous* distributed state generation is actually the most general model, subsuming both of the above. This model, which QRNA adopts, provides the most direct match to applications such as entanglement-based quantum key distribution, in which long-distance Bell pairs are measured at each end soon after creation.

**Links, Nodes, and State** Classical network architectures are typically composed of three fundamental elements: nodes, links, and state. Nodes represent the communicating parties, or relays that assist those parties. Links represent one-hop communication paths, and state represents the information being communicated.

Nodes in a quantum network are much like their classical counterparts, except that they include memory that can encode qubits. Some architectures support nodes that interact with quantum state but avoid needing direct quantum memory. Links in a quantum network transmit both quantum state as well as classical information. Both types of information are required to support teleportation.

**Paths** Multihop networks require a means of selecting a path through the network [17]. One approach is to adapt Dijkstra's shortest path first algorithm to repeater networks.

The layered communication approach impacts whether paths are established before communication or on the fly. As shown in Figure 4.3, purify-and-swap requires continuous actions distributed among the nodes along the path, so it assumes that communications will follow the same path for the entire session. Pre-establishment of a path simplifies naming for mid-session operations and simplifies predictable resource allocation by assigning in-process quantum states to specific sessions. On-the-fly path construction is more flexible but could result in communications being interrupted if available memory or quantum states are exhausted, *e.g.,* by competing connections.

**Identifiers** Networks naturally require names for the nodes or communication end points. Unlike the Internet, purify-and-swap end nodes communicate directly with nodes along the path.

On the Internet, a packet is directed to transit a particular subnet (Internet Autonomous System), rather than given a complete, hop-by-hop source route. QRNA's recursive naming allows an operation, such as Bell pair creation or entanglement swapping, to be similarly directed to a subnetwork rather than to a specific node. Paths then can be transparently relocated within the subnetwork. This partially relaxes the path constraint, simplifying end node knowledge of network components and returning local operation decisions to the local neighborhood.

The entangled states built within the network also must be named, to facilitate their management and delivery to applications. On the Internet, packets are mapped to a connection using a tuple consisting of node addresses, a connection identifier (port numbers), and possibly an application-level identifier. In quantum networks, such a tuple may not yet exist because a distributed state, such as a Bell pair in the middle of the network, might not yet be assigned to serve a particular end-to-end session. QRNA is designed to accommodate this delayed association (a type of *late binding*) and to reassign state identifiers when necessary.

### 4.3 Quantum Communication Approaches

Purify-and-swap was developed because a perfect physical quantum link cannot exist. Purify-and-swap's demand for round-trip, end-to-end communication limits throughput and demands long memory lifetimes. The quest to better match available technological capabilities and improve performance have driven the development of several new approaches to the vertical layering and horizontal distributed communication interaction, as summarized in Table4.2.

A near-ideal technology would give long quantum memory lifetime, high-fidelity local operations, a high probability of entanglement success, and high-fidelity coupling, but no such technology exists today.

We can compensate for low memory lifetime by using quantum error correction in the repeater nodes, or reengineering the protocol stack to avoid round-trip delays. The encoded link [24] and quasi-asynchronous [31] approaches each require an individual memory lifetime longer than the link round-trip time (RTT), but for $n$ hops require this for $n$ separate memories, in which the total time a state is stored in memory to be proportional to the end-to-end latency. The surface code [22] and memoryless [32] approaches can tolerate short memory lifetimes, but at the expense of needing a high probability of entanglement success.

The availability of sufficient buffer memory is also a problem. The earliest purify-and-swap proposals required a few tens of qubits per node, proportional to the log of the number of repeater hops in a network's longest path. Although this suggests a scalable solution, it exceeds current experimental capabilities. An adapted version uses only two qubits per node [16]. Encoded link and surface code, which depend on QEC, require orders of magnitude more memory than purify and swap. The memoryless approach takes advantage of a clever encoding to avoid storing qubits in memory.

Table 4.2 Comparison of several quantum repeater communication session architectures. RTT is round trip time, E2E is end to end.

| Approach | Requirements | | | |
|---|---|---|---|---|
| | Memory Life.time | Local Operation Fidelity | Entanglement Success Probability | Entanglement Fidelity |
| Hop-by-hop teleportation | E2E RTT | Very high | Low | Very high |
| Purify & Swap [15] | multiple E2E RTTs | High | Low | Low |
| Encoded link [24] | E2E RTT | High enough for QEC | Low | Fairly high |
| Surface code [22] | local QEC cy.cle time | High enough for QEC | High | High |
| Quasiasynchronous [31] | E2E RTT / 2 | Very high | Low | Fairly high |
| Memoryless [32] | Very low | Very high | High | High |
| Measurementbased [46] | multiple E2E RTTs | Fairly high | N/A | Low |

If the generated entanglement is already of high fidelity, all of these schemes will work well. Purify-and-swap and measurement-based operate with low fidelity entanglement, but can reduce the round-trip purification delays when entanglement fidelity is high.

Measurement-based can be considered a new implementation of purify-and-swap, and a carefully-engineered protocol stack would allow it as a drop-in replacement for individual nodes. Conversely, memoryless is a new link architecture whose benefits are realized only when the entire protocol stack is optimized. Encoded link, surface code and quasi-asynchronous are not specific to a particular link layer, and may as a group be able to support the same upper layer protocols, including ESC.

### 4.4 Routing for Quantum Repeater Networks

WIDE members are the first researchers to explore the issue of path selection in realistic, heterogeneous quantum networks. As in classical networks, the selection of a path between two nodes must be done efficiently in a distributed fashion, and perhaps with imperfect information about the state of the network. The path selection algorithm impacts the stability and performance of the entire network, as well as the single communication being requested.

This problem demonstrates perfectly the operational methodology of AQUA: many classical networks use Dijkstra's shortest path first (SPF) algorithm [18, 30], but it cannot be used as-is in quantum networks. Rather than deriving a new, untested approach to path selection, we chose to adapt Dijkstra. By properly defining the link cost, we have discovered that SPF can indeed be used to select a high-bandwidth path through a network of quantum repeaters. A paper on this topic was published at the end of 2013 [42].

We present the results of three sets of simulations of various paths using four different qualities of links. The first set of forty-six paths vary in length from one to nine hops, while the second set covers 256 link combinations in four-hop paths, both using a target fidelity of $F \geq 0.98$. The third data

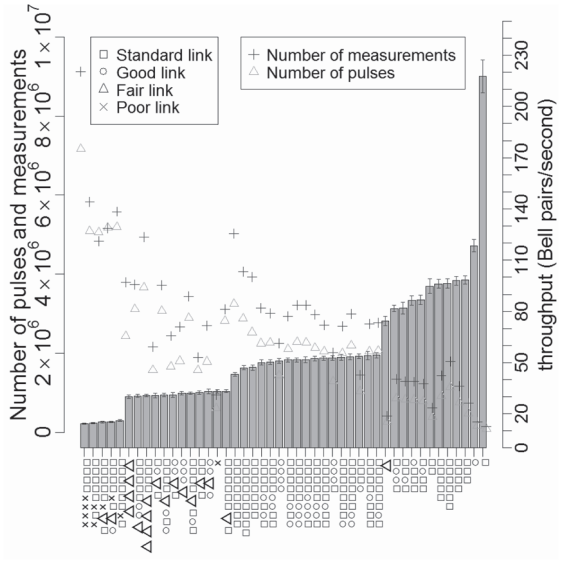


Figure 4.4 Total number of pulses (△) and measurements (+) for forty-six of our candidate paths (right scale), for end-to-end fidelity $F \geq 0.98$. The paths vary in length from one to nine hops. They are ordered left to right according to ascending throughput, plotted using bars (left scale). The legend below the graph shows the individual path configurations; □, ○, △ and × represent our standard, good, fair, and poor links, respectively. The stair-step behavior reflects increasing numbers of rounds of purification.

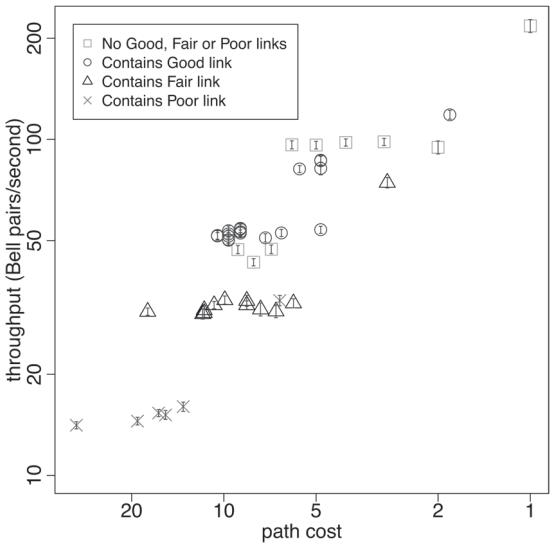


Figure 4.5 Throughput $F \geq 0.98$ versus BellGenT path cost for forty-six of our candidate paths. Each path is represented by the symbol for the weakest type of link in the path. The clustering of each type of data point shows clearly that throughput is limited by the bottleneck link. The length of the vertical bar (mostly contained within each symbol) shows the std. dev. of the throughput.

set replicates the first forty-six paths, but for a target fidelity of $F \geq 0.90$, which is too low for some distributed quantum computations but high enough for successful quantum key distribution [14]. The Open Shortest Path First (OSPF) protocol built on a distributed form of Dijkstra's algorithm is typically deployed in networks of up to a thousand nodes or so, with a diameter typically less than twenty (often less than five in modern practice) and an average path length of four to seven even in the largest networks [10, 23, 30]. Thus, we believe that examining cases of up to nine hops provides adequate coverage of the likely usage scenarios for our approach.

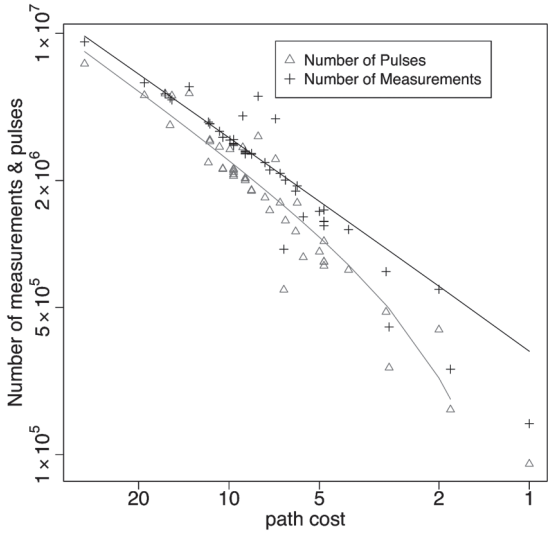


Figure 4.6 Total work ($F \geq 0.98$), measured in Pulses ( △ ) and Measurements (+), versus BellGenT path cost for forty-six of our candidate paths. The coefficient of determination of each linear fit is 0.88, showing that our path cost is a strong predictor of total work.

Across the first two data sets, the coefficient of determination is 0.88 or better between the path cost and the total work performed (counted as the number of *quantum measurements* performed along the whole path), supporting our choice of link cost and the effectiveness of Dijkstra for this type of quantum network. Comparing the results of pairs of simulations, the path with the lower cost also has higher throughput in more than 80% of all tested cases. We demonstrate that, in direct analogy to classical networks, the performance of a quantum path will be limited by the throughput of the *bottleneck link*, while total work is a function of both the path length and the quality of all the links.

Figures 4.4, 4.5 and 4.6 show our simulation results for the paths. In Fig 4.4, the throughput for each specific path can

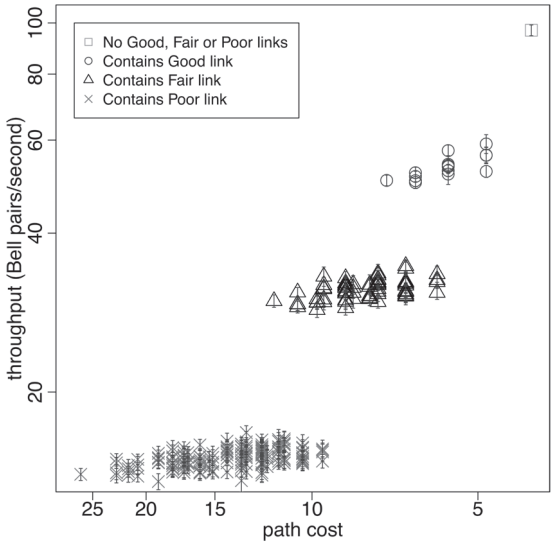


Figure 4.7 Throughput ($F \geq 0.98$) versus BellGenT path cost for all 256 four-hop candidate paths. Each path is represented by the symbol for the weakest type of link in the path. The clustering of each type of data point shows clearly that throughput is limited by the bottleneck link.

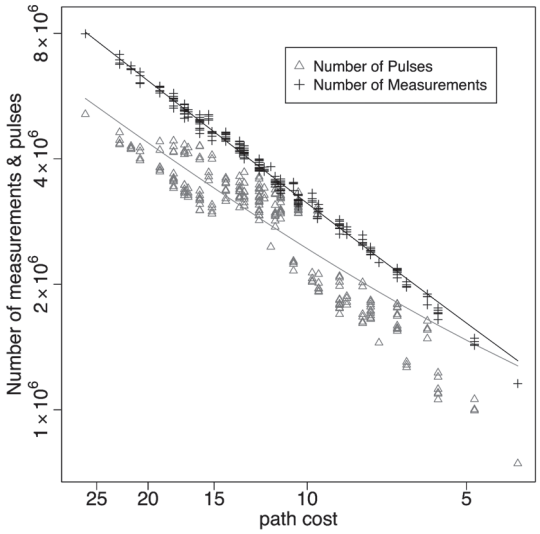


Figure 4.8 Total work to achieve output fidelity $F \geq 0.98$, measured in pulses ( △ ) and measurements (+), versus BellGenT path cost for all 256 four-hop candidate paths, with linear fits. The coefficient of determination for the number of pulses is 0.81, and for the number of measurements is 0.99.

be seen, as well as the two measures of total work, Pulse and Meas. In Fig 4.5, throughput is plotted against the Dijkstra-calculated path cost using BellGenT as our cost metric, and in Fig 4.6 the total work measures are plotted against calculated path cost. Figures 4.7 and fig 4.8 plot the results for all 256 four-hop paths we simulated.

# 第５章　Quantum Computation

## 5.1　Quantum Architecture

In order to encourage more research into quantum computer architectures, we have published a paper titled “A Blueprint for Building a Quantum Computer,” in *Communications of the ACM* [39], with a readership of 100,000. To create a stronger understanding of the various subfields and their contribution to a complete ecosystem including applications, programming tools, and architectures, the subfields and their relationship are shown in Fig 5.1.

## 5.2　Compilation and Resource Management

Finally, we are developing new optimizations for specific quantum gates. Quantum operations, are specified along a continuum, but often must be implemented using a small set of discrete gates. The standard approach is known as Solovay-Kitaev decomposition. Ongoing research is centered around improvements in the search mechanism for finding good decompositions. Preliminary results indicate a factor of three improvement in run time on the quantum computer, while producing higher accuracy. Fig 5.2 shows the length of the output sequence length on the vertical axis, as a function of the accuracy of the resulting sequence on the horizontal axis. The Recursive SSE technique consistently outperforms the standard algorithm by a large margin.

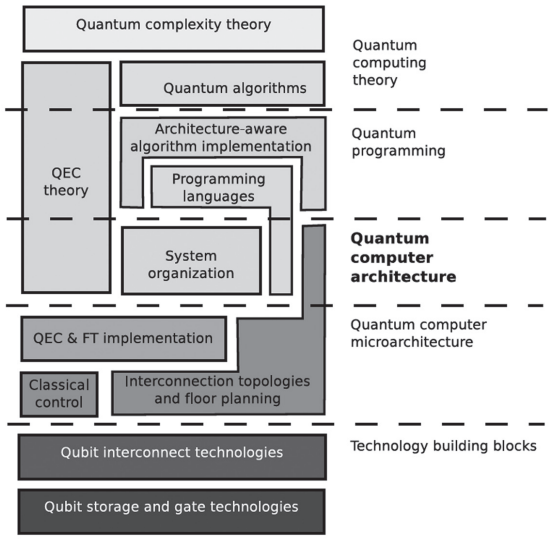


Figure 5.1　Subfields that all contribute to a complete quantum computing architecture and ecosystem. Image from a forthcoming paper in *Communications of the ACM*.

# 第６章　Publication

AQUA members had four journal papers published in 2013 and several international conference poster presentations.

- Rodney Van Meter, Takahiko Satoh, Thaddeus D. Ladd, William J. Munro, and Kae Nemoto, “Path Selection for Quantum Repeater Networks,” *Networking Science*, Dec. 2013.

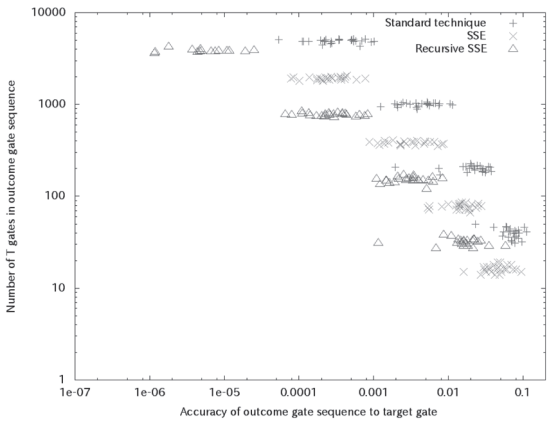


Figure 5.2　Compilation using geometric near-neighbor trees and search space expansion (SSE) is more computationally efficient, allowing improved accuracy of gate sequences used to approximate difficult-to-execute-directly arbitrary single-qubit rotations needed for quantum algorithms.

**Abstract** Quantum networks will support long-distance quantum key distribution (QKD) and distributed quantum computation, and are an active area of both experimental and theoretical research. Here, we present an analysis of topologically complex networks of *quantum repeaters* composed of heterogeneous links. Quantum networks have fundamental behavioral differences from classical networks; the delicacy of quantum states makes a practical path selection algorithm imperative, but classical notions of resource utilization are not directly applicable, rendering known path selection mechanisms inadequate. To adapt Dijkstra's algorithm for quantum repeater networks that generate entangled Bell pairs, we quantify the key differences and define a link cost metric, *seconds per Bell pair* of a particular fidelity, where a single Bell pair is the resource consumed to perform one quantum teleportation. Simulations that include both the physical interactions and the extensive classical messaging confirm that Dijkstra's algorithm works well in a quantum context. Simulating about three hundred heterogeneous paths, comparing our path cost and the total work along the path gives a coefficient of determination of 0.88 or better.

- Rodney Van Meter and Clare Horsman, "A Blueprint for Building a Quantum Computer," *Communications of the ACM*, 56(10), 84--93.

**Abstract** Small-scale quantum computing devices built on a variety of underlying physical implementations exist in the laboratory, where they have been evolving for over a decade, and have demonstrated the fundamental characteristics necessary for building systems. The challenge lies in extending these systems to be large enough, fast enough, and accurate enough to solve problems that are intractable for classical systems, such as the factoring of large numbers and the exact simulation of other quantum mechanical systems. The architecture of such a computer will be key to its performance. Structurally, when built, a "quantum computer" will in fact be a hybrid device, with quantum computing units serving as coprocessors to classical systems. The program, much control circuitry, and substantial pre and post-processing functions will reside on the classical side of the system. The organization of the quantum system itself, the algorithmic workloads for which it is designed, its speed and capabilities in meeting those goals, its interfaces to the classical control logic, and the design of the classical control systems are all the responsibility of quantum computer architects. In this article we review the progress that has been made in developing architectures for full-scale quantum computers. We highlight the process of integrating the basic elements that have already been developed, and introduce the challenges that remain in delivering on the promise of quantum computing.

- Rodney Van Meter and Joe Touch, "Designing Quantum Repeater Networks," *IEEE Communications*, 51(8), 64-71.

**Abstract** Quantum networks generate distributed entangled state or relocate quantum state, uniquely ensuring eavesdropper detection or reaching agreement more quickly than their classical counterparts. These capabilities rely on the composition of link and multihop mechanisms into a coherent system, with particular attention to managing errors in and loss of delicate quantum states. This document explores quantum networking in terms of fundamental network architecture principles and explains where and how it diverges from its classical counterparts. It discusses engineering principles that ensure robust and interoperable communication by introducing new protocol layers to support quantum sessions, and considers how these layers interact with quantum link mechanisms to support user-level quantum-enabled applications.

- Tieng Trung Pham, Rodney Van Meter and Clare Horsman, "Optimizing the Solovay-Kitaev Algorithm," *Physical Review A*, 87, 052332.

**Abstract** The Solovay-Kitaev algorithm is the standard method used for approximating arbitrary single-qubit gates for fault-tolerant quantum computation. In this paper we introduce a technique called search space expansion, which modifies the initial stage of the Solovay-Kitaev algorithm,

increasing the length of the possible approximating sequences but without requiring an exhaustive search over all possible sequences. This technique is combined with an efficient space search method called geometric nearest-neighbor access trees, modified for the unitary matrix lookup problem, in order to reduce significantly the algorithm run time. We show that, with low time cost, our techniques output gate sequences that are almost an order of magnitude smaller for the same level of accuracy. This therefore reduces the error correction requirements for quantum algorithms on encoded fault-tolerant hardware.

## 第7章　State of the Community

Last year, we wrote that 2013 would be a year of significant flux for the quantum R&D community, and that has proved to be true. IARPA killed a major program aimed at engineering of quantum systems, and many researchers around the world shuffled their locations. More of the same can be expected in 2014.

Venture capital has begun taking a serious interest in quantum technology. Google is apparently increasing its investment in the area. Microsoft remains cautious, but is building a strong group. HRL in Malibu continues to grow, while BBN and other participants in the IARPA program have had to retrench. The Singapore government has given very large grants to some mid-career researchers to build groups, and rumors abound that the Korean government is considering upping its investment.

The FIRST and Quantum Cybernetics programs that have provided funding for the last several years end in March 2014, so change will be coming to the Japanese community as well.

## 第8章　Appendix A: What is AQUA?

### 8.1　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.

### 8.2　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.

# 第９章　Appendix B: Background: FAQ on Quantum Computing

## 9.1 What is Quantum Computing?

Quantum computing brings new capabilities, including the ability to solve some problems efficiently for which no efficient classical solutions are known, such as factoring large numbers (which impacts encryption key exchange mechanisms), and new, secure means for sharing information based on the physics of quantum effects rather than the mathematical difficulty of certain problems.

Classically, a device that holds binary data can be in only one state at a time, either zero or one. However, when data is stored on systems controlled by quantum effects, the device (or *qubit*) can be in a *superposition* of states, partially in the zero state and partially in the one state. With some restrictions, this allows a *quantum computer* to operate on an exponentially large number of inputs at the same time, e.g., $n$ qubits can hold $2^n$ values at the same time. When multiple qubits are in a highly correlated state, they are *entangled*.

The difficult part, and the true art in designing algorithms for quantum computers, is extracting useful answers from the superposition state. *Interference* is used to cancel out incorrect answers and reinforce correct answers, so that *measuring* the quantum state has a high probability of giving the correct answer to a problem.

Quantum technologies initially will not be standalone: they need to integrate with classical systems and networks. In fact, they may be deployed as coprocessors for large-scale classical systems, improving precision and runtime for large computations through “quantum-assisted computing”.

## 9.2 Why is Quantum Computing Valuable?

For some problems, quantum computers are believed to be much faster than classical computers [9, 29]. The most famous result to date is Peter Shor's algorithm for factoring large numbers [36], which may potentially impact encryption technology, as mechanisms such as Diffie-Hellman key exchange and public-key cryptography (e.g., RSA) may be vulnerable to a practical solution to this problem. However, machines for running Shor's algorithm are known to be very large, far beyond currently-viable technology [40, 41].

Before Shor machines become viable, then, it is likely that quantum computers will be deployed for other uses. They were, in fact, originally conceived as a means for simulation other quantum systems [21]. Quantum computers with as few as 40 high-quality qubits may prove to be useful for solving problems in quantum chemistry [8]. This approach may lead to the custom design of new materials, and possibly an improved understanding of the quantum effects that result in superconductivity. Related quantum technologies are also expected to advance quantum metrology, improving our ability to measure gravitional fields and to create high-accuracy clocks capable of measuring time to an accuracy of $10^{-19}$.

Above all, quantum computation promises to be a completely new theory of information, based on recognizing that information is not abstract, but must be connected to its physical representation [6, 13, 25, 26, 33].

## 9.3 Why is Quantum Computing Necessary?

The economic imperative of Moore's Law [28] dictates that companies in the semiconductor industry increase the density of silicon chips every year, while reducing the per-transistor price correspondingly. In recent years, the pace of improvement has slowed somewhat to a doubling approximately every three years, but the net result remains an exponential growth in the

number of transistors in a chip, and therefore a reduction in the size of each transistor [20].

### 9.4 What is Quantum Key Distribution?

Quantum key distribution (QKD) uses quantum effects to detect the presence of an eavesdropper on a communications channel [11, 27]. QKD creates a stream of bits shared between two parties that are guaranteed by physics, rather than mathematics, to be secret (subject, of course, to the usual issues of correct and safe implementation). These secret bits are then useful as keys for standard, symmetric encryption, replacing keys generated using the Diffie-Hellman protocol. Experimental networks of QKD systems have been deployed in Boston [19], Vienna [34], and Tokyo.

### 9.5 Where is World-Leading Quantum Information Research Being

Outstanding experimental work on quantum technologies is being done in over thirty laboratories here in Japan, as well as in the United States (Caltech, Stanford, Harvard, Berkeley, Duke, MIT, Los Alamos National Lab, NIST, and many others), Canada (especially Waterloo and Calgary), the United Kingdom (Bristol, Oxford and others), Austria, Australia, France, and elsewhere. Within Japan, leading institutions include U. Tokyo, Osaka U., Tohoku U., NICT, NEC, RIKEN, NTT, Keio and others. Top-level theory work is also a broad international effort covering the same countries. IBM has had a long-standing, broad-based effort in this area, and recently companies such as Microsoft have begun contributing. In Japan, leading theorists work at NII, U. Tokyo, Keio, NTT, RIKEN, Osaka U., Tohoku U., and elsewhere.

Many of the researchers in Japan, including WIDE Board member Rodney Van Meter, are members of the FIRST Quantum Information Processing Project[*1]. This four-year project, begun in 2010, is supported with 3,000,000,000 yen from the Japanese government. Most of the money is expected to be used to support continuing leading-edge experimental work.

---

*1 http://first-quantum.net/e/index.html