AQUA: Advancing Quantum Architecture
Annual Report 2011

January 19, 2012

Abstract

In 2011, 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 particular, this year AQUA completed and published the designs of basic protocol state machines for quantum repeaters, proposed a high-level recursive network architecture for quantum repeater networks, and made several advances in quantum error correction systems, compilation techniques and tools for quantum programs, and improved efficiency for some quantum programs. These results are contained in five journal papers and two peer-reviewed conference papers, one of which was awarded Best Student Paper. AQUA also conducted a 1.5-day tutorial and workshop and created animations explaining some of the key principles in quantum computation, as part of an educational outreach effort.

1 Introduction

2011 was a watershed year for quantum information. In May, the startup company D-Wave (Vancouver, B.C.) reported the sale of the first commercial "quantum computer", which was purchased by Lockheed Martin and installed at USC's Information Sciences Institute (USC/ISI). Although controversy exists about whether this particular machine is truly a quantum computer, and no one believes that the machine as currently configured is capable of outperforming classical computers on any important problems, the sale represents a milestone, and may spur further interest in the research and even business communities.

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 begins by introducing the AQUA group and work areas, then details results from calendar 2011. A brief introduction to the field of quantum information is included as Appendix A.

2 What is AQUA?

2.1 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.

2.2 Work Areas

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

- Devices: In conjunction with researchers at Stanford University, we are designing semiconductor-based chips using optically-controlled *quantum dots*.

- 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.

In particular, AQUA is currently focusing on projects in several areas, detailed below: quantum repeater networks, surface code error correction, and compilation of quantum programs and efficient use of resources.

3 Quantum Networks

The work on networks in AQUA this year covered several areas: development of protocol state machines for repeaters and analysis of their behavior, including multiplexing for networks under contention; the new Quantum Recursive Network Architecture, intended to support organizational autonomy and technological heterogeneity in an internetwork; and application-layer optimizations, notably quantum network coding.

3.1 Foundations of Quantum Repeater Networks

Swap-and-purify quantum repeaters use two mechanisms, entanglement swapping (Fig. 1) and purification (Fig. 2), to create high-fidelity, long-distance entanglement. Although these mechanisms have been studied by physicists, no formal protocol design exists. A layered architecture has been proposed [30], and WIDE members are now in the process of creating the protocol state machines and defining the contents and sequence of operations. WIDE members published work on the state machines, and on the uses of multiplexing in repeater networks, in 2011 [3, 2].

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 [10, 23], 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.

3.2 Quantum Recursive Network Architecture

While the issues of cross-technology transfer of quantum data are under investigation by experimental physicists, no prior work has looked at the issues of heterogeneity in network management. Truly global-scale quantum networks require solutions that allow networks to operate autonomously, and that do not require global information about the state of the network in order to e.g. select a path through the network or choose where to perform entanglement swapping. We have proposed a Quantum Recursive Network Architecture (QRNA) to address these problems [31].

In QRNA, a network can be abstracted as a single node, simplifying the topology of the network that must be considered when making decisions. The network protocols managing entanglement swapping and purification also "stack" nicely, creating a truly recursive protocol stack.

Rather than dealing only with the transmission of data, as in classical networks, quantum networks are also used to create specific distributed, entangled quantum states. Because the requests necessary to build these states are exactly the same as requests for distributed computation, QRNA in fact is more than just a network, it is an

architecture for a complete distributed system.

3.3 Session Layer: Quantum Network Coding

This research considers quantum network coding, which is a recent technique that enables quantum information to be sent on complex networks at higher rates than by using straightforward routing strategies. Kobayashi (NII) et al. have recently showed the potential of this technique by demonstrating how any classical network coding protocol gives rise to a quantum network coding protocol. They nevertheless primarily focused on an abstract model, in which quantum resource such as quantum registers can be freely introduced at each node.

In this work, we present a protocol for quantum network coding under weaker (and more practical) assumptions: our new protocol works even for quantum networks where adjacent nodes initially share one EPR-pair but cannot add any quantum registers or send any quantum information. A typically example of networks satisfying this assumption is quantum repeater networks, which are promising candidates for the implementation of large scale quantum networks. Our results thus show, for the first time, that quantum network coding techniques can increase the transmission rate in such quantum networks as well.

Quantum network coding, like classical network coding, is focused on enhancing the performance of the network. It is perhaps best viewed as a session layer, above the immediate issues of transport but below the actual application uses of the network itself.

4 Quantum Computation

4.1 Surface Code Error Correction

The surface code is considered to be one of the most viable forms of quantum error correction, but the resource demands for it are high [29, 26]. In 2011, we developed *lattice surgery*, which allows smaller numbers of qubits to be used for the surface code [17]. Fig. 3 depicts the arrangement of qubits for a code distance 3 execution of a controlled-NOT gate, using 53 qubits. This approach is a good compromise between practical execution requirements and strength of error correction, and will make experimental tests of the surface code feasible in the coming years.

A critical issue in device design is dealing with hard faults (non-functional qubits) in the surface code error correction mechanism. Our ability to allow the system to work with hard faults will likely determine the success or failure of a promising hardware/software approach to large-scale system architecture, and therefore is high-priority work.

We are simulating complete systems (classically, rather than full quantum simulations), including the code to actually perform error correction, in conjunction with Austin Fowler (Melbourne). Preliminary results indicate that, contrary to some earlier analytic estimates, surface code systems on systems with hardware faults of more than a few percent cannot operate well. This has profound implications for the effort to build quantum computing systems; either the fabrication process must be nearly perfect, or the system must be organized such that the error correction process is unaware of physical defects in the system. More detailed results will appear in early 2012.

4.2 Compilation and Resource Management

Quantum picturalism is a graphical mathematical formalism for describing the processes that create quantum states, especially those created using the surface code. It allows for straightforward comparison of two quantum circuits to see if they calculate the same state, and includes a simple set of rules that help optimization of quantum circuits [16].

Graph embedding is another tool, focused on effective management of the motion of quantum variables throughout the quantum computer. The structure of the physical system can be described as a graph, and the program or circuit can be described as a directed graph connecting nodes (qubits) based on dependency chains in the computation (edges) [9].

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.

5 Educational Outreach and Animations

WIDE members conducted a 1.5-day event, the FIRST/Quantum Cybernetics/CREST Joint 1.5-day Surface Code Quantum Error Correction Tutorial/Workshop, held at Osaka University, Feb. 23-24, 2011. Most of the sessions were conducted in Japanese, for about 25 attendees. Eight hours of video were collected and posted on the web [1]. Attendees included both device-level theorists and experimentalists, and may lead to experimental demonstrations of basic operations related to the surface code in the next couple of years.

For Keio University's Shonan Fujisawa Campus's annual Open Research Forum, SFC students (most of whom are WIDE/AQUA members) created a series of animations explaining key concepts in quantum information [2]. Fig. 4 shows an example screen shot from one of the animations, depicting the process of quantum key distribution. Other animations cover the basic concepts of interference for both classical and quantum optics, entanglement, and quantum measurement.

6 Publications

AQUA members had four journal papers published in 2011 and a fifth accepted for publication.

- Rodney Van Meter, Joe Touch and Clare Horsman, "Recursive Quantum Repeater Networks," *Progress in Informatics*, 8, pp. 65–79, Mar. 2011.
  Abstract Internet-scale quantum repeater networks will be heterogeneous in physical technology, repeater functionality, and management. The classical control necessary to use the network will therefore face similar issues as Internet data transmission. Many scalability and management problems that arose during the development of the Internet might have been solved in a more uniform fashion, improving flexibility and reducing redundant engineering effort. Quantum repeater network development is currently at the stage where we risk similar duplication when separate systems are combined. We propose a unifying framework that can be used with all existing repeater designs. We introduce the notion of a Quantum Recursive Network Architecture, developed from the emerging classical concept of *recursive networks*, extending recursive mechanisms from a focus on data forwarding to a more general distributed computing request framework. Recursion abstracts independent transit networks as single relay nodes, unifies software layering, and virtualizes the addresses of resources to improve information hiding and resource management. Our architecture is useful for building arbitrary distributed states, including fundamental distributed states such Bell pairs and GHZ, W, and cluster states.

- Byung-Soo Choi and Rodney Van Meter, "On the Effect of Quantum Interaction Distance on Quantum Addition Circuits," *ACM Journal of Emerging Technologies in Computing Systems*, 7(3), Aug. 2011.
  Abstract We investigate the theoretical limits of the effect of the quantum interaction distance on the speed of exact quantum addition circuits. For this study, we exploit graph embedding for quantum circuit analysis. We study a logical mapping of qubits and gates of any $\Omega(\log n)$-depth quantum adder circuit for two n-qubit registers onto a practical architecture, which limits interaction distance to the nearest neighbors only and supports only one- and two-qubit logical gates. Unfortunately, on the chosen k-dimensional practical architecture, we prove that the depth lower bound of any exact quantum addition circuits is no longer $\Omega(\log n)$, but $\Omega(\sqrt[k]{n})$. This result, the first application of graph embedding to quantum circuits and devices, provides a new tool for compiler development, emphasizes the impact of quantum computer architecture on performance, and acts as a cautionary note when evaluating the time performance of quantum algorithms.

- Clare Horsman, "Quantum picturalism for topological cluster-state computing," *New Journal of Physics*, 13, 095011, Sept. 2011.
  Abstract Topological quantum computing is a way of allowing precise quantum computations to run on noisy and imperfect hardware. One implementation uses *surface codes* created by forming defects in a highly-entangled cluster state. Such a method of comput-

[1] Videos are available at `http://www.soi.wide.ad.jp/class/20110030/`.

[2] These animations are available on the WIDE Annual Report CD, as well as online at `http://aqua.sfc.wide.ad.jp/ORF2011/ORF2011-teaching-videos.html`.

ing is a leading candidate for large-scale quantum computing. However, there has been a lack of sufficiently powerful high-level languages to describe computing in this form without resorting to single-qubit operations, which quickly become prohibitively complex as the system size increases. In this paper we apply the category-theoretic work of Abramsky and Coecke to the topological cluster-state model of quantum computing to give a high-level graphical language that enables direct translation between quantum processes and physical patterns of measurement in a computer – a "compiler language". We give the equivalence between the graphical and topological information flows, and show the applicable rewrite algebra for this computing model. We show that this gives us a native graphical language for the design and analysis of topological quantum algorithms, and finish by discussing the possibilities for automating this process on a large scale.

- Byung-Soo Choi and Rodney Van Meter, "A $\Theta(\sqrt{n})$-depth Quantum Adder on a 2D NTC Quantum Computer Architecture," *ACM Journal on Emerging Technologies in Computing Systems*, to appear; available from the arXiv as `quant-ph:1008.5093`.
  Abstract In this work, we propose an adder for the 2D NTC architecture, designed to match the architectural constraints of many quantum computing technologies. The chosen architecture allows the layout of logical qubits in two dimensions and the concurrent execution of one- and two-qubit gates with nearest-neighbor interaction only. The proposed adder works in three phases. In the first phase, the first column generates the summation output and the other columns do the carry-lookahead operations. In the second phase, these intermediate values are propagated from column to column, preparing for computation of the final carry for each register position. In the last phase, each column, except the first one, generates the summation output using this column-level carry. The depth and the number of qubits of the proposed adder are $\Theta(\sqrt{n})$ and $O(n)$, respectively. The proposed adder executes faster than the adders designed for the 1D NTC architecture when the length of the input registers n is larger than 58.

- Clare Horsman, Katherine L. Brown, William J. Munro, and Vivien M. Kendon, "Reduce, reuse, recycle for robust cluster-state generation," *Phys. Rev. A* 83, 042327 (2011), Apr. 2011
  Abstract Efficient generation of cluster states is crucial for engineering large-scale measurement-based quantum computers. Hybrid matter-optical systems offer a robust, scalable path to this goal. Such systems have an ancilla which acts as a bus connecting the qubits. We show that by generating smaller cluster "Lego bricks", reusing one ancilla per brick, the cluster can be produced with maximal efficiency, requiring fewer than half the operations compared with no bus reuse. By reducing the time required to prepare sections of the cluster, bus reuse more than doubles the size of the computational workspace that can be used before decoherence effects dominate. A row of buses in parallel provides fully scalable cluster state generation requiring only 20 controlled-PHASE gates per bus use.

AQUA had two peer-reviewed papers accepted to international conferences this year, one of which was awarded Best Student Paper.

- Luciano Aparicio, Rodney Van Meter, and Hiroshi Esaki, "Protocol design for quantum repeater networks," *Proc. AINTEC 2011*, Nov. 2011. *Awarded Best Student Paper.*
  Abstract When built, quantum repeater networks will require classical network protocols to control the quantum operations. However, existing work on repeaters has focused on the quantum operations themselves, with less attention paid to the contents, semantics, ordering and reliability of the classical control messages. In this paper we define and describe our implementation of the classical control protocols. The state machines and packet sequences for the three protocol layers are presented, and operation confirmed by running the protocols over simulations of the physical network. We also show that proper management of the resources in a bottleneck link allows the aggregate throughput of two end-to-end flows to substantially exceed that of a single flow. Our layered architectural framework will support independent evolution of the separate protocol layers.

- Luciano Aparicio and Rodney Van Meter, "Multiplexing schemes for quantum

repeater networks," Quantum Communications and Quantum Imaging IX, Aug. 2011. *Proc. SPIE* 8163, 816308.

Abstract When built, quantum repeaters will allow the distribution of entangled quantum states across large distances, playing a vital part in many proposed quantum systems. Enabling multiple users to connect through the same network will be key to their real-world deployment. Previous work on repeater technologies has focused only on simple entanglement production, without considering the issues of resource scarcity and competition that necessarily arise in a network setting. We simulated a thirteen-node network with up to five flows sharing different parts of the network, measuring the total throughput and fairness for each case. Our results suggest that the Internet-like approach of statistical multiplexing use of a contested link gives the highest aggregate throughput. Time division multiplexing and buffer space multiplexing were slightly less effective, but all three schemes allow the sum of multiple flows to substantially exceed that of any one flow, improving over circuit switching by taking advantage of resources that are forced to remain idle in circuit switching. All three schemes proved to have excellent fairness. The high performance, fairness and simplicity of implementation support a recommendation of statistical multiplexing for shared quantum repeater networks.

One student, Luciano Aparicio, completed his master's thesis on network protocols for quantum repeaters:

- Luciano Aparicio, "Design and Evaluation of Communication Protocols for Quantum Repeater Networks," master's thesis, University of Tokyo, 2011.

7 AQUA's Plans For the Coming Year

Indications are that 2012 and 2013 are likely to be years of significant flux for quantum information technology and the R&D community. A number of experimental groups throughout the world stand on the verge of being able to perform experiment demonstrations at small scales (5-15 qubits) of the viability of quantum algorithms and especially topological surface code computation, some perhaps using simplified versions of the lattice surgery method described above. Two years from now, sources of funding (government, corporate internal, venture capital), which laboratories and organizational approaches are key (e.g., university laboratory, national laboratory, corporate laboratory or startup company), and which implementation technologies are considered viable (nanophotonic, nitrogen-vacancy diamond, superconducting flux qubit, quantum dot) are all likely to change.

In 2012, we expect to lay the groundwork for a testbed for large-scale emulation of quantum repeater networks. Such a testbed will allow implementation of the classical protocols, evaluating their performance and robustness, and in particular carefully evaluating the information that each node in the network has about the state of entanglement throughout the network. Such a testbed will support testing many different protocol variants for both propagation of entanglement through the network, and suppression of errors.

The I-D for IPsec with QKD is currently under revision, and is expected to be resubmitted as an "individual submission" to the IETF Security Area Directors during 2012. ITU is currently in the process of standarding low-level aspects of QKD, and would like to have this portion of the technology documented and standardized.

Papers on qDijkstra, surface code on a defective lattice, and implementation of quantum simulation algorithms are expected to be submitted early in 2012. Technical work will continue on these topics, as well as on QRNA and quantum repeaters in general.

Issues of optimizing circuits and matching circuits to architectures continue to be important, and work on graph embedding (both tools and principles) will continue in 2011.

AQUA WG meetings are expected to continue during 2012 at WIDE Camps and Kenkyuukai, as well as informal weekly cross-institutional research meetings.

References

[1] F. Adler. Minimum energy cost of an observation. *Information Theory, IEEE Transactions on*, 1(3):28–32, 1955. IRE Transactions on Information Theory.

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

[3] L. Aparicio, R. Van Meter, and H. Esaki. Protocol design for quantum

repeater networks. In *Proc. Asian Internet Engineering Conference*, Nov. 2011.

[4] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon. Simulated quantum computation of molecular energies. *Science*, 309:1704–1707, 2005.

[5] D. Bacon and W. van Dam. Recent progress in quantum algorithms. *Commun. ACM*, 53(2):84–93, Feb. 2010.

[6] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proc. IEEE International Conference on Computers, Systems, and Signal Processing*, pages 175–179. IEEE, Dec. 1984.

[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] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM J. Computing*, 26(5):1411–1473, 1997.

[9] B.-S. Choi and R. Van Meter. On the effect of quantum interaction distance on quantum addition circuits. *J. Emerg. Technol. Comput. Syst.*, 7:11:1–11:17, August 2011.

[10] E. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959.

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

[12] 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.

[13] C. Elliott, D. Pearson, and G. Troxel. Quantum cryptography in practice. In *Proc. SIGCOMM 2003*. ACM, ACM, Aug. 2003.

[14] ESIA, JEITIA, KSIA, TSIA, and SIA. International technology roadmap for semiconductors. Technical report, ESIA and JEITIA and KSIA and TSIA and SIA, 2009. http://public.itrs.net/.

[15] R. P. Feynman. Simulating physics with computers. In A. J. G. Hey, editor, *Feynman and Computation*. Westview Press, 2002.

[16] C. Horsman. Quantum picturalism for topological cluster-state computing. *New Journal of Physics*, 13:095011, 2011.

[17] C. Horsman, A. Fowler, S. Devitt, and R. Van Meter. Surface code quantum computing by lattice surgery. *Arxiv preprint arXiv:1111.4022*, 2011.

[18] R. Landauer. Irreversibility and heat generation in the computing process. *IBM J. of Research and Development*, 5(3):183–191, 1961. reprinted in IBM J. R.&D. Vol. 44 No. 1/2, Jan./Mar. 2000, pp. 261–269.

[19] R. Landauer. Energy needed to send a bit. *Proceedings: Mathematical, Physical and Engineering Sciences*, 454(1969):305–311, 1998.

[20] H.-K. Lo and Y. Zhao. Quantum cryptography. In *Encyclopedia of Complexity and System Science*. Springer, 2008. arXiv:0803.2507v4 [quant-ph].

[21] G. E. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8), Apr. 1965.

[22] M. Mosca. Quantum algorithms. *Arxiv preprint arXiv:0808.0369*, 2008.

[23] J. Moy. OSPF version 2. RFC 2178, July 1997.

[24] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[25] M. Peev, C. Pacher, R. Alleaume, C. Barreiro, J. Bouda, W. Boxleitner, T. Debuisschert, E. Diamanti, M. Dianati, J. F. Dynes, S. Fasel, S. Fossier, M. Furst, J.-D. Gautier, O. Gay, N. Gisin, P. Grangier, A. Happe, Y. Hasani, M. Hentschel, H. Hubel, G. Humer, T. Langer, M. Legre, R. Lieger, J. Lodewyck, T. Lorunser, N. Lutkenhaus, A. Marhold, T. Matyus, O. Maurhart, L. Monat, S. Nauerth, J.-B. Page, A. Poppe, E. Querasser, G. Ribordy, S. Robyr, L. Salvail, A. W. Sharpe, A. J. Shields, D. Stucki, M. Suda, C. Tamas, T. Themel, R. T. Thew, Y. Thoma, A. Treiber, P. Trinkler, R. Tualle-Brouri, F. Vannel, N. Walenta, H. Weier, H. Weinfurter, I. Wimberger, Z. L. Yuan, H. Zbinden, and A. Zeilinger. The SECOQC quantum key distribution network in Vienna. *New Journal of Physics*, 11(7):075001 (37pp), 2009.

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

[27] P. W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Proc. 35th Symposium on Foundations of Computer Science*, pages 124–134, Los Alamitos, CA, 1994. IEEE Computer Society Press.

[28] R. Van Meter, K. M. Itoh, and T. D. Ladd. Architecture-dependent execution time of Shor's algorithm. In *Proc. Int. Symp. on Mesoscopic Superconductivity and Spintronics (MS+S2006)*, Feb. 2006.

[29] R. Van Meter, T. D. Ladd, A. G. Fowler, and Y. Yamamoto. Distributed quantum computation architecture using semiconductor nanophotonics. *International Journal of Quantum Information*, 8:295–323, 2010. preprint available as arXiv:0906.2686v2 [quant-ph].

[30] R. Van Meter, T. D. Ladd, W. J. Munro, and K. Nemoto. System design for a long-line quantum repeater. *IEEE/ACM Transactions on Networking*, 17(3):1002–1013, June 2009.

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

A Background: FAQ on Quantum Computing

A.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".

A.2 Why is Quantum Computing Valuable?

For some problems, quantum computers are believed to be much faster than classical computers [22, 5]. The most famous result to date is Peter Shor's algorithm for factoring large numbers [27], 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 [29, 28].

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 [15]. Quantum computers with as few as 40 high-quality qubits may prove to be useful for solving problems in quantum chemistry [4]. 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 [24, 8, 18, 19, 1].

A.3 Why is Quantum Computing Necessary?

The economic imperative of Moore's Law [21] 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 [14].

A.4 What is Quantum Key Distribution?

Quantum key distribution (QKD) uses quantum effects to detect the presence of an eavesdropper on a communications channel [6, 20]. 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 [13], Vienna [25], and Tokyo.

A.5 What is a Quantum Repeater?

Loss of photons in a fiber is exponential in the length of the fiber, and the fidelity (quality) of the quantum state also declines, limiting practical direct quantum connections to perhaps 150km. Quantum repeaters [12, 11] connect a series of shorter hops (perhaps as little as 10km, depending on technology), creating entangled states over a long distance and potentially allowing the creation of a global quantum network.

Quantum repeaters use *purification* (a quantum-specific type of error correction) and *entanglement swapping* (based on *teleportation* [7]), and must have high-quality quantum memory.

A.6 What is a Quantum Network?

Quantum networks come in two flavors: those that use long-lived entanglement, and those that do not. The latter kind are primarily useful for QKD, whereas the former are expected to be used for various distributed applications beyond QKD, such as the quantum metrology mentioned above.

Except for the physical mechanism of entangling qubits using an optical fiber (or even through free space), the problems of quantum networks are the same as for classical networks: how to choose an efficient route through a network with imperfect information, how to reliably transmit information, and how to manage the resources of the network in a distributed fashion.

Beyond the simple transfer of quantum data from one location to another, quantum networks actual act as fully distributed quantum computing systems [31]. Thus, the classical requests that support quantum communication effectively become requests for the execution of quantum algorithms. This feature of quantum networks remains to be explored.

A.7 Where is World-Leading Quantum Information Research Being Done?

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. 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 [3]. 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.

[3] `http://first-quantum.net/e/index.html`

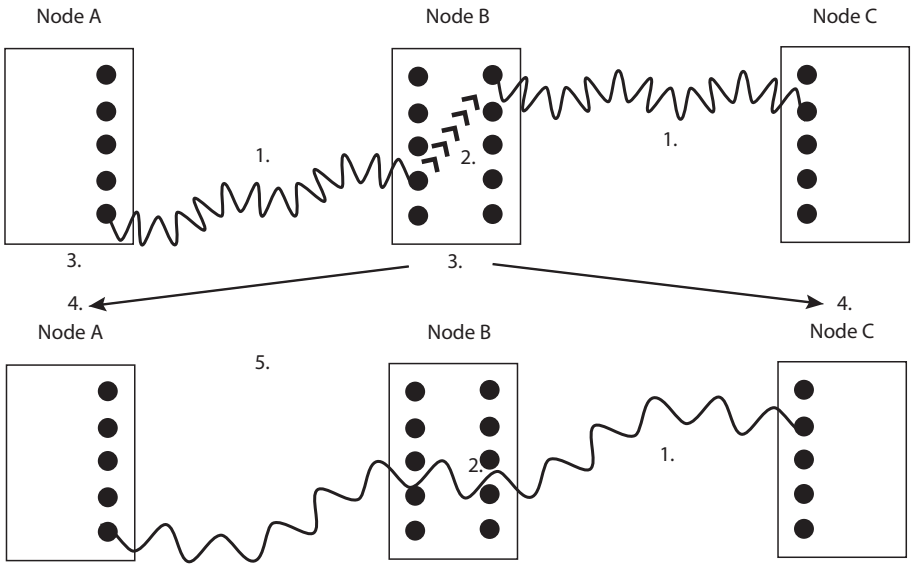

Figure 1: Teleportation can lengthen one Bell pair using another, in a process known as entanglement swapping.

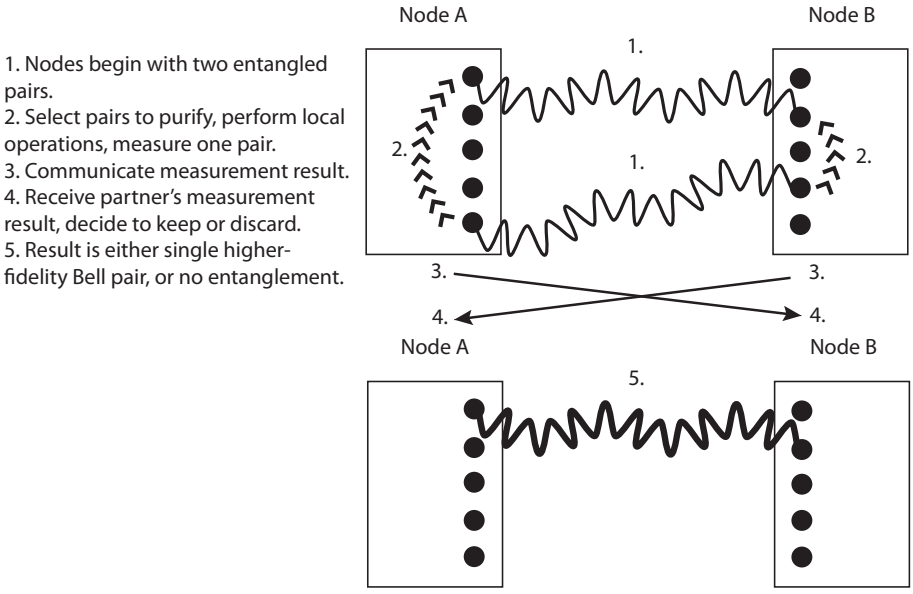

Figure 2: Steps involved in purification.

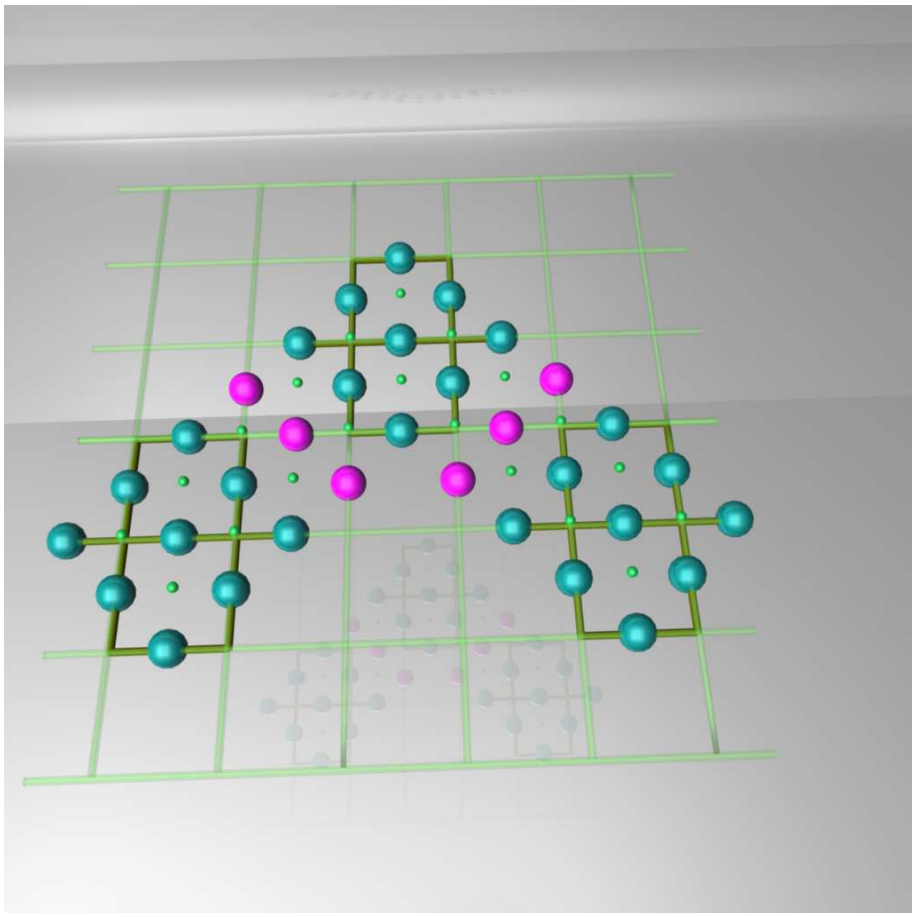

Figure 3: Depiction of the 53-qubit controlled-NOT gate using the lattice surgery method on the planar surface code.

Figure 4: Animation explaining the principles of quantum key distribution. The complete animation is available on the WIDE Annual Report CD.