# AQUA: Advancing Quantum Architecture Annual Report 2012

January 15, 2013


## 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 2012, AQUA members published seven 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; workloads for quantum computers; and quantum repeater networks.


## 1 Introduction

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

This report first discusses recent work in WIDE on quantum networks, then on general quantum computation. This is followed by a summary of 2012's major publications [9, 10, 17, 18, 19, 29, 31, 32]. AQUA outreach efforts, centering on a hand-built demonstration of quantum key distribution, are discussed, followed by our operational efforts. An introduction to the AQUA group and work areas is included as Appendix A. A brief introduction to the field of quantum information is included as Appendix B.

## 2 Quantum Networks

In 2012, AQUA members two papers on quantum networks: an overview of the field in the widely-read magazine *IEEE Network* [32], and a paper integrating the concepts of practical quantum repeaters with the abstract concept of quantum network coding, improving the practicality of network coding [29].

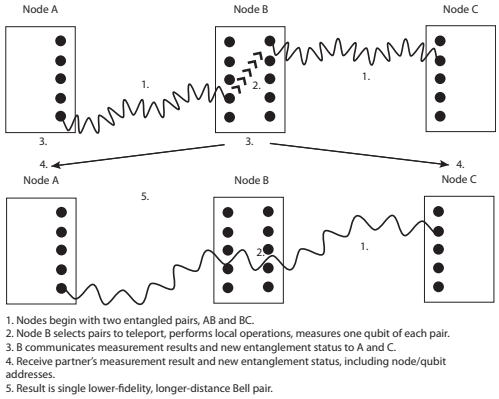


Figure 1: Teleportation can lengthen one Bell pair using another, in a process known as entanglement swapping. Image from [32].

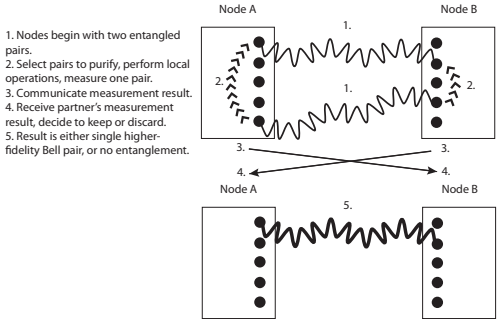


Figure 2: Steps involved in purification. Image from [32].

### 2.1 Foundations of Quantum Repeater Networks

As Van Meter noted [32],

> 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 architecture, of course, must fit the landscape: it is constrained underneath by the capabilities we have, and above by the needs of the applications. In AQUA, our work concerns most aspects of such an architecture, especially the inter-networking issues, but with only minimal attention to the physical layer. Indeed, our approach is to apply a layered abstraction in order to isolate functionality, and allow the physical layer to evolve independently of the error management and other layers.

The "standard" model of quantum repeater uses an approach known as *swap and purify*. 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 [35], and WIDE members are now in the process of creating the protocol state machines and defining the contents and sequence of operations, building on work from prior years [2, 3].

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 [11, 25], 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 is currently under review [36].

### 2.2 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 [29].

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.

## 3 Quantum Computation

### 3.1 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 [28, 34]. In 2011, we developed *lattice surgery*, which allows smaller numbers of qubits to be used for the surface code, and published a description in *New Journal of Physics* in 2012 [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.

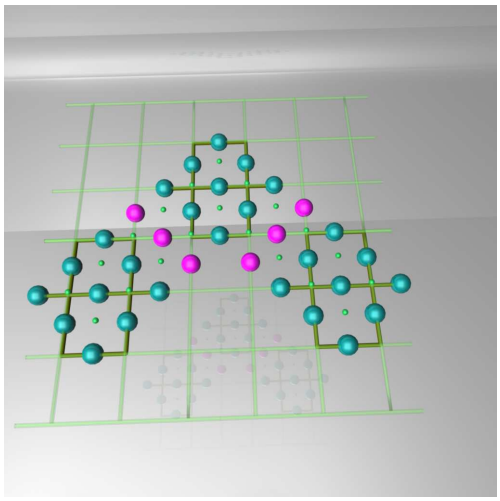

Figure 3: Depiction of the 53-qubit controlled-NOT gate using the lattice surgery method on the planar surface code. Image from [17].

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 2013.

## 3.2 Quantum Architecture

In order to encourage more research into quantum computer architectures, and to create a stronger understanding of the various subfields and their contribution to a complete ecosystem including applications, programming tools, and architectures, we have developed a concept titled "A Blueprint for Building a Quantum Computer," and a paper on this topic has been accepted to *Communications of the ACM*. The subfields and their relationship are shown in Fig. 4.

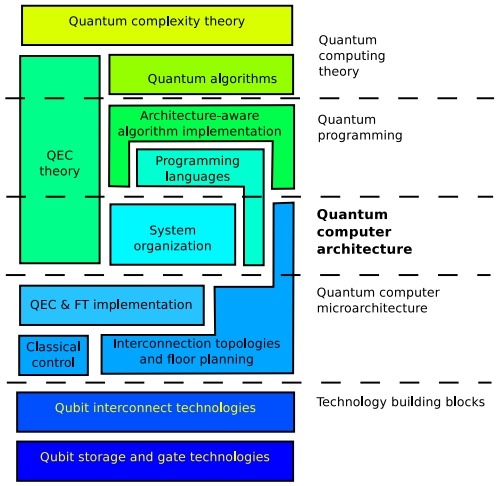


Figure 4: Subfields that all contribute to a complete quantum computing architecture and ecosystem. Image from a forthcoming paper in *Communications of the ACM*.

In 2012, in conjunction with a group from Stanford University, we published a description of a specific architecture using quantum dots [19]. More importantly, this paper contains a decomposition of an architecture into layers, with specific responsibilities, much as a classical architecture has the device, micro- and macro-architectures, as shown in Fig. 5. Our hope is that other researchers will be able to use this framework in developing their own architectures, which will provide consistent division of responsibilities, allow greater reuse of engineering effort, and make apples-to-apples comparisons of tradeoffs straightforward.

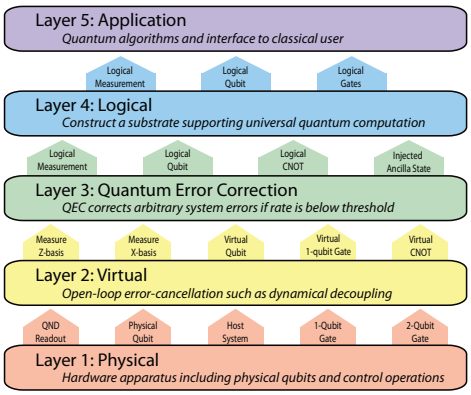


Figure 5: Layers of a quantum architecture. Image from [19].

### 3.3 Workloads and Architecture-Aware Implementations

Few quantum algorithms have been described in enough detail to determine their actual runtime on realistic architectures. An important step in this process is architecture-aware development of key building blocks, such as arithmetic subroutines. In many architectures, the connectivity of qubits at both the physical and logical layer is limited to neighbors in a plane. We have developed an addition subroutine that is optimized for use in such an environment [10].

### 3.4 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. 6 shows

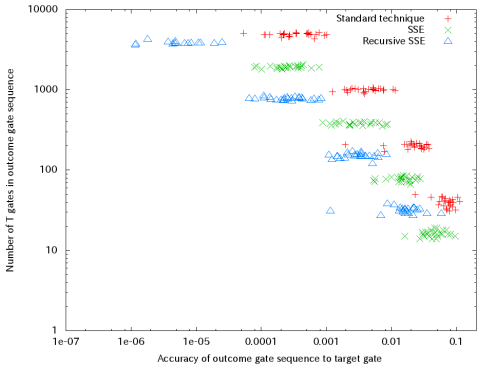


Figure 6: Compilation using geometric near-neighbor trees and search space expansion (SSE) is more computationally efficient, allowing improved accuracy of gate sequences used to approximate difficult-to-execute-directly arbitrary single-qubit rotations needed for quantum algorithms.

## 4 Publications

AQUA members had seven journal papers published in 2012 and another accepted for publication, one invited conference paper, and several international conference poster presentations.

- Clare Horsman, Austin G. Fowler, Simon Devitt, and Rodney Van Meter, "Surface code quantum computing by lattice surgery," *New Journal of Physics*, 14, 123011, Dec. 2012.
  **Abstract** In recent years, surface codes have become a leading method for quantum error correction in theoretical large-scale computational and communications architecture designs. Their comparatively high fault-tolerant thresholds and their natural two-dimensional nearest-neighbour (2DNN) structure make them an obvious choice for large scale designs in experimentally realistic systems. While fundamentally based on the toric code of Kitaev, there are many variants, two of which are the planar- and defect-based codes. Planar codes require fewer qubits to implement (for the same strength of error correction), but are restricted to encoding a single qubit of information. Interactions between encoded qubits are achieved via transversal operations, thus destroying the inherent 2DNN nature of the code. In this paper we introduce a new technique enabling the coupling of two planar codes without transversal operations, maintaining the 2DNN of the encoded computer. Our lattice surgery technique comprises splitting and merging planar code surfaces, and enables us to perform universal quantum computation (including magic state injection) while removing the need for braided logic in a strictly 2DNN design, and hence reduces the overall qubit resources for logic operations. Those resources are further reduced by the use of a rotated lattice for the planar encoding. We show how lattice surgery allows us to distribute encoded GHZ states in a more direct (and overhead friendly) manner, and how a demonstration of an encoded CNOT between two distance-3 logical states is possible with 53 physical qubits, half of that required in any other known construction in 2D.

- N. Cody Jones, James D. Whitfield, Peter L. McMahon, Man-Hong Yung, Rodney Van Meter, Alan Aspuru-Guzik, and Yoshihisa Yamamoto, "Faster quantum chemistry simulation on fault-tolerant quantum computers," *New Journal of Physics*, 14, 115023, Nov. 2012.
  **Abstract** Quantum computers can in

principle simulate quantum physics exponentially faster than their classical counterparts, but some technical hurdles remain. We propose methods which substantially improve the performance of a particular form of simulation, ab initio quantum chemistry, on faulttolerant quantum computers; these methods generalize readily to other quantum simulation problems. Quantum teleportation plays a key role in these improvements and is used extensively as a computing resource. To improve execution time, we examine techniques for constructing arbitrary gates which perform substantially faster than circuits based on the conventional Solovay-Kitaev algorithm (Dawson and Nielsen 2006 Quantum Inform. Comput. 6 81). For a given approximation error $\epsilon$ arbitrary single-qubit gates can be produced fault-tolerantly and using a restricted set of gates in time which is $O(\log \epsilon)$ or $O(\log \log \epsilon)$; with sufficient parallel preparation of ancillas, constant average depth is possible using a method we call programmable ancilla rotations. Moreover, we construct and analyze efficient implementations of first- and second-quantized simulation algorithms using the fault-tolerant arbitrary gates and other techniques, such as implementing various subroutines in constant time. A specific example we analyze is the ground-state energy calculation for lithium hydride.

- Satoh, Takahiko and Le Gall, François and Imai, Hiroshi, *Physical Review A*, 86(3), 032331, Sept. 2012.
  **Abstract** This paper considers quantum network coding, which is a recent technique that enables quantum information to be sent on complex networks at higher rates than straightforward routing strategies. Kobayashi 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 resources such as additional 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 Einstein-Podolsky-Rosen pair but cannot add any additional quantum registers or send any quantum information. A typical 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 that quantum network coding techniques can increase the transmission rate in such quantum networks as well.

- 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 (JETC)*, 8(3), 22, Aug. 2012.
  **Abstract** In this work, we propose an adder for the 2-Dimensional Nearest-Neighbor, Two-Qubit gate, Concurrent (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 with $\sqrt{n}$ columns where each column has $\sqrt{n}$ qubits 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 51.

- N. Cody Jones, Rodney Van Meter, Austin G. Fowler, Peter L. McMahon, Jungsang Kim, Thaddeus D. Ladd, Yoshihisa Yamamoto, "A Layered Architecture for Quantum Computing Using Quantum Dots," *Physical Review X (PRX)*, 2, 031007, July 2012.
  **Abstract** We develop a layered quantum-computer architecture, which is a systematic framework for tackling the individual challenges of developing a quantum computer while constructing a cohesive device design. We discuss many of the prominent techniques for

implementing circuit-model quantum computing and introduce several new methods, with an emphasis on employing surface-code quantum error correction. In doing so, we propose a new quantum-computer architecture based on optical control of quantum dots. The time scales of physical-hardware operations and logical, error-corrected quantum gates differ by several orders of magnitude. By dividing functionality into layers, we can design and analyze subsystems independently, demonstrating the value of our layered architectural approach. Using this concrete hardware platform, we provide resource analysis for executing fault-tolerant quantum algorithms for integer factoring and quantum simulation, finding that the quantum-dot architecture we study could solve such problems on the time scale of days.

- Rodney Van Meter, "Quantum Networking and Internetworking," *IEEE Network*, July/August 2012, pp. 59-64.
  **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.

- Katherine L. Brown, Clare Horsman, Viv Kendon, William J. Munro, "Layer by layer generation of cluster states," *Physical Review A (PRA)* 85, 052305, May 2012.
  **Abstract** Cluster states can be used to perform measurement-based quantum computation. The cluster state is a useful resource, because once it has been generated only local operations and measurements are needed to perform universal quantum computation. In this paper, we explore techniques for quickly and deterministically building a cluster state. In particular, we consider generating cluster states on a qubus quantum computer, a computational architecture which uses a continuous variable ancilla to generate interactions between qubits. We explore several techniques for building the cluster, with the number of operations required depending on whether we allow the ability to destroy previously created controlled-phase links between qubits. In the case where we cannot destroy these links, we show how to create an $n \times m$ cluster using just $3nm - 2n - \lceil 3m/2 \rceil + 3$ operations. This gives more than a factor of 2 savings over a naive method. Further savings can be obtained if we include the ability to destroy links, in which case we only need $\lceil 1/3(8nm - 4n - 4m - 8) \rceil$ operations. Unfortunately, the latter scheme is more complicated so choosing the correct order to interact the qubits is considerably more difficult. A halfway scheme, that keeps a modular generation but saves additional operations over never destroying links requires only $3nm - 2n - 2m + 4$ operations. The first scheme and the last scheme are the most practical for building a cluster state because they split up the generation into the repetition of simple sections.

One invited paper appeared in the proceedings of the 21st Asian Test Symposium:

- Rodney Van Meter, "Counting Gates, Moving Qubits: Evaluating the Execution Cost of Quantum Circuits," 21st Asian Test Symposium, Niigata, Nov. 2012.
  **Abstract** Quantum algorithms can be written down in several forms; one of the most common is the quantum circuit representation using discrete gates. The challenge in assessing the computational cost then becomes counting those gates, with realistic costs assigned to each gate. Moreover, interacting pairs of qubits inside most quantum computers will require moving qubits. In many architectures, this will involve cellular automaton-like swapping of qubits. In general, the depth will be described

in number of quantum error correction (QEC) cycles, while the total cost will be space-time "volume" consisting of the number of qubits involved over that set of QEC cycles. This implies that accurate estimates can be made only in the context of a particular architecture and error correction mechanism.

One student, Shota Nagayama, completed his master's thesis and two students, Kaori Ishizaki and Pham Tien Trung, completed their bachelor's theses on topics related to quantum compilation:

- Shota Nagayama, "Surface Code Quantum Computation on a Defective Physical Lattice," master's thesis, Keio University, Graduate School of Media and Governance, Mar. 2012.
- Kaori Ishizaki, "An algorithm for optimizing movement of quantum variables on arbitrary physical qubit structures," bachelor's thesis, Keio University, Faculty of Environment and Information Studies, Mar. 2012 (in Japanese).
- Pham Tien Trung, "Constructing a software framework for synthesizing high accuracy quantum circuits," bachelor's thesis, Keio University, Keio University Faculty of Environment and Information Studies, Mar. 2012.

## 5 Educational Outreach

For Keio University's Shonan Fujisawa Campus's annual Open Research Forum, two SFC students, Iori Mizutani and Koji Murata, created a demonstration of the principles behind *quantum key distribution* (QKD). In QKD, single photons are transmitted and used to create shared random string of classical bits that are guaranteed secret due to our ability to detect the changes to the state caused by an eavesdropper (see App. B.4 for an additional, brief description of QKD). The demonstration used strong laser pulses and simple polaroid filters rather than single photons, and therefore cannot be truly secure, but the strong pulses allow for a visible beam and a visual, easy-to-understand demo. The demo itself, as shown in Fig. 7, is otherwise a fairly complete implementation, including an eavesdropper and detection of her presence or absence. An elegant web interface (Fig. 10) controls the entire setup for both the sender Alice (Fig. 8) and the receiver Bob (with Eve, in Fig. 9). Fig. 11 shows the beam from the laser diode at Alice passing through Eve's filter.

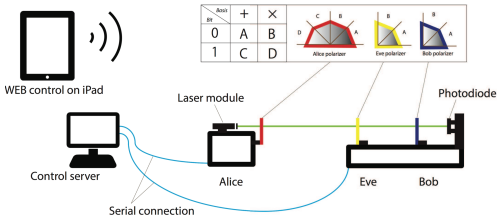


Figure 7: Setup of the QKD demo, using hand-built hardware and a custom web interface.

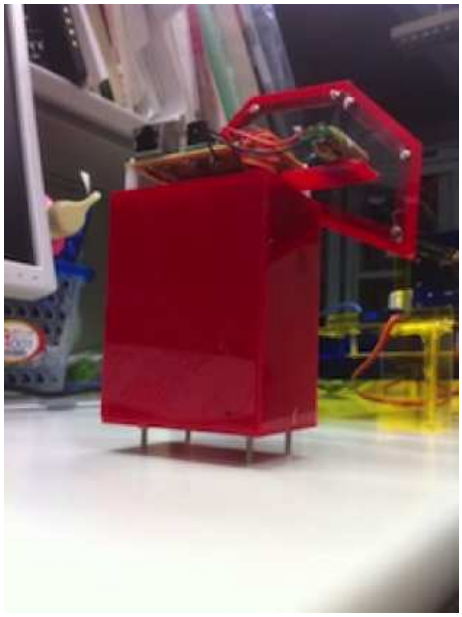

Figure 8: Photograph of the Alice hardware.

## 6 Operational Support for FIRST Summer School

In 2012, WIDE members Shota Nagayama, Kaori Ishizaki, Takahiko Satoh and Rodney Van Meter attended the FIRST Project's quantum computing summer school in Miyako-jima, Okinawa-ken. As in 2010, we supplied WLAN connectivity and a local server to host educational materials as well as a photography upload center for about 75 participants. This year's effort was handled by Nagayama and Ishizaki.

## 7 AQUA's Plans For the Coming Year

Indications are that 2013 is likely to be a year 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

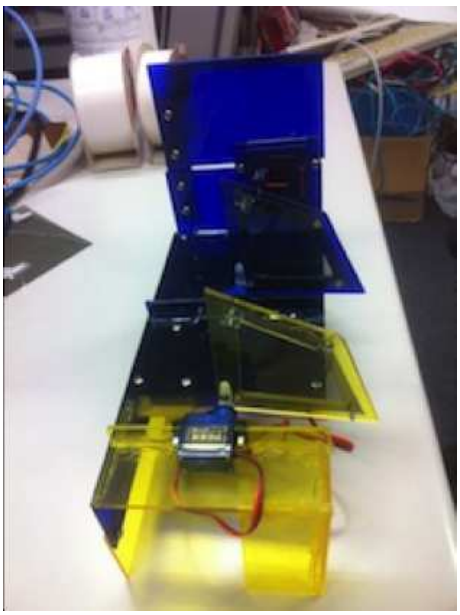

Figure 9: Photograph of the Bob (blue) and Eve (yellow) hardware.

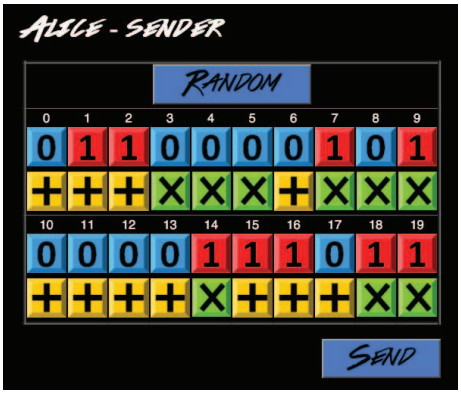


Figure 10: The web interface for the QKD demo. Alice chooses both a basis ($\times$ or $+$) and a value (0 or 1) to send. Bob chooses only a measurement basis.

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 addition to those papers already published, our paper on qDijkstra is under review at a journal, and our paper on Solovay-Kitaev decomposition is under revision for resubmission. Papers on surface code on a defective lattice, and implementation of quantum simulation algorithms are expected to be submitted early in 2013. Technical work will continue on these topics, as well as on our Quantum Recursive Network Architecture (QRNA) and quantum repeaters in general.

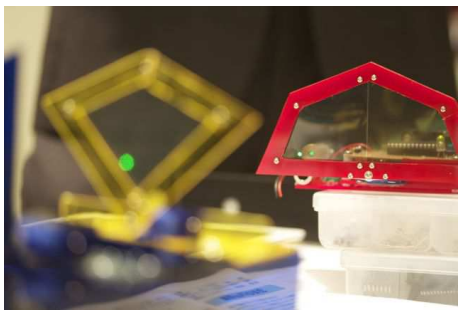

Figure 11: Photograph of Alice (background) firing the laser at Eve (foreground). If Eve is eavesdropping on the conversation, Alice and Bob will be able to detect her presence due to her impact on the quantum states being transmitted.

Finally, Prof. Van Meter has been appointed to the editorial board of ACM's *Journal of Emerging Technologies in Computing Systems (JETC)*, and will be contributing to the work of the journal this year.

## 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. *Communications of the 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] K. L. Brown, C. Horsman, V. Kendon, and W. J. Munro. Layer by layer generation of cluster states. *Physical Review A*, 85:052305, May 2012.

[10] B.-S. Choi and R. Van Meter. A $\theta(\sqrt{n})$-depth quantum adder on a 2d NTC quantum computer architecture. 8(3):24:1–24:22, Aug. 2012.

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

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

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

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

[15] 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/.

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

[17] C. Horsman, A. Fowler, S. Devitt, and R. Van Meter. Surface code quantum computing by lattice surgery. *New Journal of Physics*, 14:123011, 2012.

[18] N. Jones, J. Whitfield, P. McMahon, M. Yung, R. Van Meter, A. Aspuru-Guzik, and Y. Yamamoto. Faster quantum chemistry simulation on fault-tolerant quantum computers. *New Journal of Physics*, 14(11):115023, 2012.

[19] N. C. Jones, R. Van Meter, A. G. Fowler, P. L. McMahon, J. Kim, T. D. Ladd, and Y. Yamamoto. Layered architecture for quantum computing. *Phys. Rev. X*, 2:031007, Jul 2012.

[20] 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.

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

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

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

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

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

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

[27] 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.

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

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

[30] 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.

[31] R. Van Meter. Counting gates, moving qubits: Evaluating the execution cost of quantum circuits. In *Proc. 21st Asian Test Symposium*, Nov. 2012.

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

[33] 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.

[34] 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.

[35] 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.

[36] R. Van Meter, T. Satoh, T. D. Ladd, W. J. Munro, and K. Nemoto. Path selection for quantum repeater networks. *arXiv:1206.5655v1*, 2012.

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

# A What is AQUA?

## A.1 Goals

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

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

## A.2 Work Areas

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

- Devices: In conjunction with researchers at Stanford University, 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.

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

# B Background: FAQ on Quantum Computing

## B.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".

## B.2 Why is Quantum Computing Valuable?

For some problems, quantum computers are believed to be much faster than classical computers [24, 5]. The most famous result to date is Peter Shor's algorithm for factoring large numbers [30], 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 [34, 33].

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 [16]. 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 [26, 8, 20, 21, 1].

## B.3 Why is Quantum Computing Necessary?

The economic imperative of Moore's Law [23] 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 [15].

## B.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, 22]. 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 [14], Vienna [27], and Tokyo.

## B.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 [13, 12] 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.

## B.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 [37]. 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.

## B.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. 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`