Architecture of a Quantum Multicomputer

Rodney Van Meter[1 *] W. J. Munro[2 3 †] Kae Nemoto[3 ‡]

[1] *Faculty of Environment and Information Studies, Keio University*
5322 Endo, Fujisawa, Kanagawa, 252-8520, Japan
[2] *Hewlett-Packard Laboratories*
Filton Road, Stoke Gifford, Bristol BS34 8QZ, United Kingdom
[3] *National Institute of Informatics*
2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan

Abstract. We have created the architecture of a quantum multicomputer and analyzed its performance for running Shor's algorithm for factoring large numbers. In this paper, we combine fault-tolerance techniques with performance goals for our architecture, which uses a linear interconnect and six logical qubits per node. Our performance target of factoring a 1,024-bit number in one month requires teleporting 6.2 logical qubits per second on each link in the system, which translates to 3,300 physical teleportations per second on each link. A qubus-based interconnect generates Bell pairs of intermediate fidelity that must be purified. Starting from a Bell state with fidelity $F = 0.638$, as a qubus-based cavity QED interconnect might generate with a qubit-to-qubit loss of 3.4dB, about 1.5 million physical entanglement attempts per second are enough to reach this level of performance.

Keywords: distributed quantum computation, quantum computer architecture

1 Introduction

Researchers have begun designing systems for distributed quantum computation (DQC). The basic principle of distributed quantum computation has been known for a decade [1, 2]. A few uses for geographically distributed entanglement have been developed, but more recently, interest has grown in the use of distributed quantum computation within a single laboratory. We refer to such distributed-memory systems as quantum multicomputers.

The interest in DQC stems from the difficulty of scaling up any individual quantum computer to hold the millions of physical qubits that may be necessary for some applications. Any quantum computing technology will have a limit to the number of physical qubits that can be supported in a single device. These limits are not yet well understood, but may range into the low thousands for solid-state systems. Quantum error correction reduces the number available to applications to only a handful of logical qubits. Thus, researchers are designing devices to be connected into multicomputers [3, 4, 5].

In previous work, we have determined that a linear network of nodes will perform well on the addition subroutine, and that the most efficient form of the algorithm on a quantum multicomputer teleports data qubits, rather than gates [6]. The mechanics of the original Vedral-Barenco-Ekert (VBE) carry-ripple adder [7] work best for nodes that hold at least four logical qubits. The links between nodes may be serial, and two layers of the [[23,1,7]] Steane quantum error correction code will allow teleportation error rates of around one percent. We have also established a performance goal of factoring a 1,024-bit number in one month of wall-clock time on our system, which should exceed the performance of the best classical systems available [8].

In this paper, we establish performance requirements for the interconnect links, increase the level of detail on the node requirements and on QEC, and provide a more complete analysis of the application algorithm. We show that physical qubits must be teleported between nodes at about 3,300 teleportations per second on each link in the system, and individual nodes must contain about 10,000 physical qubits. Using the qubus scheme for the interconnect, initial fidelities are low, and purification must be used. For a cavity QED system as a candidate technology, the resulting requirement is about 1.5 million entanglement attempts per second on each link.

2 Multicomputer Architecture

Our quantum multicomputer (QMC) architecture consists of a group of semi-autonomous nodes, connected by a quantum network and a real-time classical network, all controlled from a classical front-end computer that determines the program to be run and the role to be played by each node. Each node contains four logical qubits for the algorithmic data and two logical buffer qubits for communications, which are used to send and receive data through a qubus entangling channel, giving a total of six logical qubits per node. All quantum error correction is performed locally, within a single node. Our complete system will consist of 1,024 nodes for the arithmetic unit, plus a few more for the control variables for the algorithm. An overview of the quantum components of the architecture is shown in Figure 1.

A computer system cannot be designed without an understanding of its target workload and expected performance. Above, we suggested a goal of factoring a 1,024-bit number in one month of wall-clock execution time. Next we show that achieving this performance level will require 6.2 logical qubit teleportations per second per node.

The computational core of Shor's algorithm is the quantum modular exponentiation. The modular exponentiation can be performed in roughly $4n^2$ calls to the addition subroutine, or about four million calls for a 1,024-bit number. The execution time of the adder can be optimized down to the time to teleport two logical qubits between a pair of nodes, when the

*rdv@sfc.wide.ad.jp
†bill.munro@hp.com
‡nemoto@nii.ac.jp

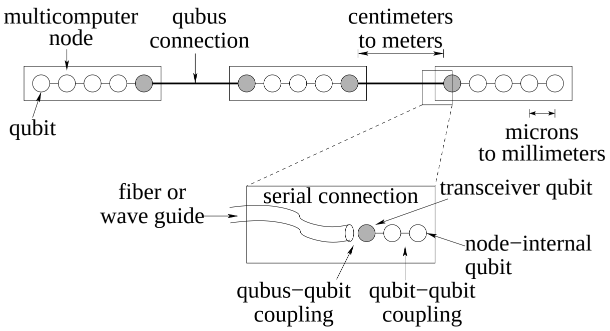

Figure 1: Quantum portions of our quantum multicomputer

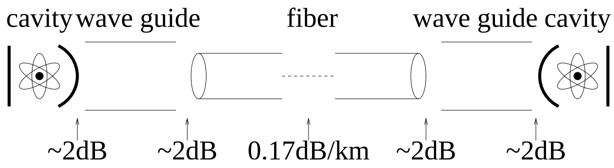

Figure 2: Qubit-to-qubit losses in a cavity QED system

time to generate a high-fidelity Bell pair is long compared to the local gate time [6]. To compose a series of additions into a multiplication, a control variable must be used to set one of the operands. We must teleport the control variable into the node before each addition, and again to clear the operand after the addition. Thus, we use a total of four teleportation operations per addition, half for algorithm control and half for the arithmetic itself. Sixteen million logical qubit teleportations can be accomplished in one month if the rate of logical teleportations is 6.2 per second.

For two levels of the [[23,1,7]] Steane error correcting code, teleporting a logical qubit requires $23^2 = 529$ physical qubit teleportations. Because we are using serial links between the multicomputer nodes, each link must support $6.2 \times 529 = 3300$ physical teleportations per second.

Teleportation consumes Bell pairs. The interconnect subsystem is tasked with creating the high-fidelity Bell pairs that we need. The qubus system is one candidate for the interconnect [9]. Simulations of a cavity QED form of qubus for quantum repeaters reveal an upper bound of about 5dB loss qubit-to-qubit, beyond which entanglement fidelity is too low for purification to operate [10]. For a more practical limit, a loss of 3.4dB will give an initial fidelity $F = 0.638$. The entanglement attempts succeed about 40% of the time, and purification itself is a probabilistic process, so using purification to reach a final fidelity $F = 0.98$ or better requires an average of about 450 qubus entanglement attempts. To execute 6.2 logical teleportations per second requires $450 \times 3300 = 1.5 \times 10^6$ physical entanglement attempts per second.

Figure 2 shows approximate, somewhat optimistic signal strength losses from qubit to qubit for a cavity QED system. For a multicomputer configuration, with a fiber length of only a few meters, the loss is clearly dominated by the couplings in the system. To reach our preferred loss level of three to four dB, modest improvement in coupling efficiency is needed. Switching or demultiplexing of incoming signals to multiple transceiver qubits seems desirable, but switch losses are often several dB. This figure makes it clear that avoiding any unnecessary losses is an important goal when designing qubus-based interconnects.

Single-qubit rotations on logical qubits are difficult, and fault tolerance demands that direct data qubit-to-data qubit interactions be minimized. Fowler therefore recommends using three registers for each logical qubit [11], boosting the number of qubits per node to $6 \times 3 \times 23^2 = 9522$ physical qubits. Including transceiver qubits for the interconnect and purification buffering, we require about 10,000 physical qubits per node.

3 Conclusion

In this paper, we have described a quantum multicomputer built from about ten million physical qubits packaged in a thousand separate nodes, connected by a qubus-based interconnect. We believe such a multicomputer structure can be developed more easily than a large-scale monolithic system, allowing a variety of physical qubit technologies to scale quickly beyond currently-perceived limits. Thus, the multicomputer has the potential to dramatically accelerate the arrival of quantum computers that generate results beyond the reach of classical systems.

Acknowledgments. The authors thank MEXT for partial support for this research, and Kohei M. Itoh, Thaddeus D. Ladd, and Austin Fowler for technical guidance.

References

[1] Lov K. Grover. Quantum telecomputation. http://arXiv.org/quant-ph/9704012, April 1997.

[2] J.I. Cirac, A. Ekert, S.F. Huelga, and C. Macchiavello. *Physical Review A*, 59:4249, 1999.

[3] S.M. Clark, K.M.C. Fu, T.D. Ladd, and Y. Yamamoto. *Physical Review Letters*, 99:040501, 2007.

[4] L. Jiang, J. M. Taylor, A. S. Sorensen, and M. D. Lukin. Scalable quantum networks based on few-qubit registers. quant-ph/0703029, 2007.

[5] Daniel K. L. Oi, Simon J. Devitt, and Lloyd C. L. Hollenberg. *Physical Review A*, 74:052313, 2006.

[6] Rodney Van Meter, W. J. Munro, Kae Nemoto, and Kohei M. Itoh. *ACM Journal of Emerging Technologies in Computing Systems*, 2007. to appear; available as arXiv:quant-ph/0607160v2.

[7] Vlatko Vedral, Adriano Barenco, and Artur Ekert. *Phys. Rev. A*, 54:147–153, 1996.

[8] A. Lenstra, E. Tromer, A. Shamir, W. Kortsmit, B. Dodson, J. Hughes, and P. Leyland. In *AsiaCrypt 2003*, Lecture Notes in Computer Science, New York, 2003. Springer-Verlag.

[9] T. P. Spiller, Kae Nemoto, Samuel L. Braunstein, W. J. Munro, P. van Loock, and G. J. Milburn. *New Journal of Physics*, 8:30, February 2006.

[10] T.D. Ladd, P. van Loock, K. Nemoto, W. J. Munro, and Y. Yamamoto. *New Journal of Physics*, 8:184, 2006.

[11] Austin G. Fowler. Constructing arbitrary single-qubit fault-tolerant gates. quant-ph/0411206, December 2005.