

Synthesis of Quantum Logic Circuits

Vivek V. Shende¹

vshende@eecs.umich.edu

Stephen S. Bullock²

stephen.bullock@nist.gov

Igor L. Markov¹

imarkov@eecs.umich.edu

¹Dept. of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor, MI 48109-2212, USA

²Mathematical and Computational Sciences Division, Natl. Inst. of Standards and Technology, Gaithersburg, MD 20899-8910, USA

Abstract — The pressure of fundamental limits on classical computation and the promise of exponential speedups from quantum effects have recently brought quantum circuits to the attention of the EDA community [10, 17, 4, 16, 9]. We discuss efficient circuits to initialize quantum registers and implement generic quantum computations. Our techniques yield circuits that are twice as small as the best previously published technique. Moreover, a theoretical lower bound shows that our new circuits can be improved by at most a factor of two. Further, the circuits grow by at most a factor of nine under severe architectural restrictions.

I. INTRODUCTION

As the ever-shrinking transistor approaches atomic proportions, Moore's law must confront the small scale granularity of the world: we cannot build wires thinner than atoms. Worse still, at atomic dimensions we must contend with the laws of quantum mechanics. For example, suppose one bit is encoded as the presence or the absence of an electron in a small region.¹ Since we know very precisely where the electron is located, the Heisenberg uncertainty principle dictates that we cannot know its momentum with high accuracy. Since its speed might be large, a large potential is needed to keep it in place. A quantitative analysis of the situation leads experts from NCSU, SRC and Intel [24] to derive fundamental limitations on the scalability of any computing device which moves electrons.

Yet these same quantum effects also facilitate a radically different form of computation [7]. Theoretically, *quantum* computers could outperform their classical counterparts when solving certain discrete problems [8]. A successful large-scale implementation of Shor's integer factorization [18] would compromise the RSA cryptosystem used in electronic commerce. On the other hand, quantum effects also allow perfectly secure public-key cryptography [3]. Indeed, such cryptography systems, based on single-photon communication, are already

¹Most current computing technologies use electron charges to store information; the exception is spintronics-based techniques, e.g. magnetic RAM.

commercially available from MagiQ Technologies in the U.S. and IdQuantique in Europe.

Quantum bit data states differ from classical states in two important ways. First, a single quantum bit may take on a continuum of values $z_1|0\rangle + z_2|1\rangle$ for z_1, z_2 complex numbers. Readings of the quantum bit return 0 or 1 with probability $|z_j|^2 / \sqrt{|z_1|^2 + |z_2|^2}$, so that quantum computers inherently allow for probabilistic computation. Second and far more significant, n qubits collectively may store more information than that stored by n isolated (local) one-qubit states. Meaning, the axioms of quantum mechanics demand an n -qubit quantum state be a sum of terms $z_{\bar{b}}|\bar{b}\rangle$ for each of the 2^n bit strings \bar{b} . Thus n quantum bits in particular store 2^n probabilities of observing each bit string. *Entanglement* is the physical effect allowing this. An example is the two-qubit register $(|00\rangle + |11\rangle)/\sqrt{2}$, where the strings 00 and 11 are observed with equal probability but 01, 10 are never observed.

Physically, qubits are stored in quantum-mechanical systems, such as the nuclear spins of atoms or ions, or the current in a superconductor. Quantum logic gates can be applied to selected qubits in an n -qubit register and modify the value of the register. The gate might be applied by an RF pulse or a laser beam. Usually, gates that act on three or more qubits are prohibitively difficult to implement directly and must be decomposed into a sequence of two-qubit gates [6]. Two-qubit gates may in turn be decomposed into circuits containing one-qubit gates, and a canonical two-qubit gate called the *controlled-not* (CNOT). The CNOT can be thought of as an XOR gate that prevents loss of information by preserving one of the input values.

The first published algorithm to carry out a two-qubit gate decomposition implemented an n -qubit quantum gate by a circuit containing $O(n^3 4^n)$ CNOT gates [2]. Further improvements use clever circuit transformations and/or Gray codes [5, 1, 19]. Finally, a different technique [11] has led to circuits with a CNOT-count of $4^n - 2^{n+1}$. These numbers compare to the theoretical, dimension-based lower bound of $\lceil \frac{1}{4}(4^n - 3n - 1) \rceil$ [16]. Yet prior algorithms remain a factor of four away and fare poorly for small n . Each requires at least 8 CNOT gates for $n = 2$, while the lower bound is three. In contrast, hand-optimized two-qubit operators [23, 4] obtain three CNOTs [16, 22, 21]. Even special cases may be treated [16], using a simple procedure for finding a CNOT-optimal two-qubit circuits [14]. In contrast, in three qubits the lower bound is 14

while the generic decomposition of [11] achieves 48 CNOTs and a specialty circuit [20] achieves 40.

In our work, we implement an arbitrary n -qubit operator using $(1/2) \times 4^n - 3 \times 2^{n-1} + 1$ CNOT gates. This represents an improvement by a factor of two over the best known results for both the 3-qubit and n -qubit case. The 3-qubit count is 21 CNOTs, while the n -qubit count is a factor of two away from the lower bound of $(4^n - 3n - 1)/4$. We also discuss efficient circuits for initializing quantum registers and consider how architectural considerations can affect circuit size. As this paper reports exploratory work on a revolutionary computing technology, we do not necessarily seek working prototypes. Instead, we emphasize fundamental results and attempt to gain a better understanding of the structure of quantum circuits.

II. GATES FOR QUANTUM LOGIC

Let n be the number of qubits, $N = 2^n$. The *qubit* is the simplest possible quantum mechanical system, with only a two-dimensional state space. To bring out the analogy with a classical bit, we pick basis vectors $|0\rangle$ and $|1\rangle$. Note, however, that in general the state of a qubit is described by a complex vector $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. An n -qubit vector is a similar sum over bit strings $|\phi\rangle = \sum_{\bar{b}} \alpha_{\bar{b}} |\bar{b}\rangle$, i.e. vectors in \mathbb{C}^N . Computations are in particular unitary operators, i.e. maps $|\psi\rangle \mapsto u|\phi\rangle$, with $u \in \mathbb{C}^{N \times N}$ and $u\bar{u}^T = I_N$ for I_N an identity matrix.

We recall notation for Pauli matrices, commonly encountered in the quantum mechanics literature.

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

Appropriate physical operations may evolve a one-qubit state $|\phi\rangle \mapsto e^{i\theta\sigma_j} |\phi\rangle$, where σ_j is a Pauli matrix as above or a linear combination. An *elementary gate* is such an $R_j(\theta) = e^{-i\theta\sigma_j/2}$. One-qubit states may be seen as as vectors in space, and in this picture the R_j are spatial rotations[12, §4.2]. More explicitly:

- The x -axis rotation: $R_x(\theta) = \begin{pmatrix} \cos \theta/2 & i \sin \theta/2 \\ i \sin \theta/2 & \cos \theta/2 \end{pmatrix}$
- The y -axis rotation: $R_y(\theta) = \begin{pmatrix} \cos \theta/2 & \sin \theta/2 \\ -\sin \theta/2 & \cos \theta/2 \end{pmatrix}$
- The z -axis rotation: $R_z(\alpha) = \begin{pmatrix} e^{-i\alpha/2} & 0 \\ 0 & e^{i\alpha/2} \end{pmatrix}$

These operators suffice to implement an arbitrary one-qubit computation. In fact, an arbitrary 2×2 unitary matrix U has $U = e^{i\Phi} R_z(\alpha) R_y(\theta) R_z(\beta)$. We discard the leading scalar $e^{i\Phi}$ as it is physically unmeasurable. To derive this fact, we recall the Cosine-Sine decomposition [13] of matrix analysis.² It factors an even-dimensional unitary matrix u using smaller unitaries a, a', b, b' and real diagonal matrices c, s such that $s^2 + c^2 = 1$:

$$u = \begin{pmatrix} a & \\ & b \end{pmatrix} \begin{pmatrix} c & -s \\ s & c \end{pmatrix} \begin{pmatrix} a' & \\ & b' \end{pmatrix}$$

²Source code for computing the Cosine-Sine decomposition can be obtained from Matlab by typing `whic_h gsvd` at the Matlab prompt.

If u was a 2×2 matrix, then the left and right matrices are – up to scalars – R_z matrices. The center matrix is an R_y matrix. Collecting the scalars, we obtain the advertised decomposition.

We next describe a very useful two-qubit gate that can be implemented in practice. The controlled-not (CNOT) gate quantizes the classical reversible two-input two-output logic gate which inverts the second bit if the first is 1. Several CNOTs are depicted in Fig. 3. We write C_j^i for a CNOT gate that flips the i -th bit if the j -th is 1. The 4×4 unitaries for C_1^2 and C_2^1 are:

$$C_1^2 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad C_2^1 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}$$

If a $2^p \times 2^p$ unitary U_1 acts on a p -qubit register and a $2^q \times 2^q$ unitary U_2 acts on a q -qubit register, then the joint action on a combined $n = p + q$ register is *not* a block matrix but rather the Kronecker (tensor) product $U_1 \otimes U_2$ (e.g. [4]). As a matrix, $U_1 \otimes U_2$ is blockwise a matrix of multiples of U_2 , where the multiples are the entries of U_1 . Since tensors are cheap in the circuit language, we hope to recognize such factors.

III. THE QUANTUM MUX AND ITS IMPLEMENTATION

A *quantum multiplexor* is a gate acting on $k + 1$ qubits of which one is designated as the *control* qubit. Depending on whether the control bit carries 0 or 1, the gate performs either u_0 or u_1 on the remaining k bits. If the control bit is the highest order bit, the MUX matrix is block diagonal: $\begin{pmatrix} u_0 & 0 \\ 0 & u_1 \end{pmatrix}$.

The CNOT is a good example of a quantum multiplexor. Another variant is the *uniformly k -controlled R_z gate* [19]. Such a gate operates on $k + 1$ qubits, of which k are *controls* and one is the *target*. A different R_z is applied to the target for each control bit-string. If the target is the lowest order bit, then the matrix is block diagonal, with the i -th 2×2 block a $R_z(\theta_i)$ gate given control-string i . Encoding the parameters θ_i into a diagonal matrix δ , the uniformly controlled rotation is given by $e^{-i\sigma_z \otimes \delta/2}$. A uniformly controlled R_z gate with k controls requires only 2^k CNOT gates and 2^k gates R_z , per Fig. 1.

A $(k + 1)$ -qubit quantum multiplexor can be implemented using two k -qubit gates and a uniformly k -controlled R_z gate. To see this, we formulate an equation for the required gates and solve it. We want unitary v, w and unitary diagonal d satisfying

$$\begin{pmatrix} a & \\ & b \end{pmatrix} = \begin{pmatrix} v & \\ & v \end{pmatrix} \begin{pmatrix} d & \\ & \bar{d}^T \end{pmatrix} \begin{pmatrix} w & \\ & w \end{pmatrix}.$$

To find them, define d and v by diagonalizing $a\bar{b}^T = vd^2\bar{v}^T$. Then $w = d\bar{v}^T b$. As $d \oplus \bar{d}^T = e^{i\sigma_z \otimes \log d}$, it is a uniformly controlled R_z gate controlled on the low order bits.




Fig. 1. A uniformly controlled rotation $\exp(i\delta \otimes \sigma_z)$ is denoted by an R_z gate on σ_z 's wire and square controls on δ 's wires. We assert that such gates can be recursively decomposed as shown above. For, write $\delta = I_2 \otimes \delta_1 + \sigma_z \otimes \delta_2$, and use the well-known two-qubit circuit identity $C_1^2(\sigma_z \otimes \sigma_z)C_1^2 = I_2 \otimes \sigma_z$ to rewrite $\delta \otimes \sigma_z = I_2 \otimes \delta_1 \otimes \sigma_z + C_n^1(I_2 \otimes \delta_2 \otimes \sigma_z)C_n^1$. Exponentiating produces the circuit in the center. Recursive cancellations are shown at right.

IV. SYNTHESIS OF UNITARY OPERATORS

Recall from Section II the *Cosine-Sine* decomposition:

$$u = \begin{pmatrix} a & \\ & b \end{pmatrix} \begin{pmatrix} c & -s \\ s & c \end{pmatrix} \begin{pmatrix} d' & \\ & b' \end{pmatrix}$$

The left and right factors $a \oplus b$ and $d' \oplus b'$ are quantum multiplexors. The central factor may be written $e^{i\sigma_y \otimes \log(c-is)/2}$. In analogy with uniformly controlled R_z gates, we call this a uniformly controlled R_y gate. It may also be implemented using Fig. 1. Simplifying multiplexors per Section III, we obtain:

NQ Decomposition: For $u \in \mathbb{C}^{N \times N}$, $u\bar{u}^T = I_N$, there exist $v_1, v_2, v_3, v_4, \delta_1, \delta_2, \delta_3 \in \mathbb{C}^{N/2 \times N/2}$, $v_j \bar{v}_j^T = I_{N/2}$, δ_j diagonal, real, such that

$$u = (I_2 \otimes v_1) e^{i\sigma_z \otimes \delta_1} (I_2 \otimes v_2) e^{i\sigma_y \otimes \delta_2} (I_2 \otimes v_3) e^{i\sigma_z \otimes \delta_3} (I_2 \otimes v_4)$$


Hence an arbitrary n -qubit operator can be implemented by a circuit containing three uniformly controlled rotations and four $(n-1)$ -qubit operators, as illustrated in Figure 2. We next count gates for the resulting recursive unitary synthesis algorithm. Let c_j be the number of CNOT gates needed to implement a j -qubit operator. Then $c_j \leq 4c_{j-1} + 3 \times 2^{j-1}$. In particular, if ℓ -qubit operators may be implemented using $\leq c_\ell$ CNOT gates, then the following inequality for c_n results.

$$c_n \leq 4^{n-\ell} (c_\ell + 3 \times 2^{\ell-1}) - 3 \times 2^{n-1}$$

Apply the decomposition until only one-qubit operators remain, $c_n \leq (3/4) \times 4^n - 3 \times 2^{n-1}$ CNOT gates. (Cf. [11].) If we rather terminate the recursion with two-qubit operators, hand-optimized 3-CNOT circuit decompositions [16, 22, 21] lower the CNOT-count to $(9/16) \times 4^n - 3 \times 2^{n-1}$.

For $(4^n - 3 \times 2^n + 2)/2$ CNOTs, a final optimization is needed. End the recursion once 4^{n-2} two-qubit operators remain. These two-qubit operators are all on the same lines and are separated by the controls of uniformly controlled rotations.

Fig. 2. A circuit diagram illustrating the NQ matrix decomposition.



Diagonal operators pass through controls. Also, for any two-qubit operator v , there is a diagonal two-qubit operator d so that vd and dv may be implemented using two CNOTs ([14], Prop. III.3.) Passing through the $I_{n-2} \otimes d$'s, the remaining two-qubit operators cost two CNOT gates. Since we save one CNOT in the implementation of every two-qubit gate but the first, the count above results. Note that for $n=3$, 21 CNOTs are needed. This is the best known circuit at present (Cf. [20].)


V. PRACTICAL CONSIDERATIONS

Certain common primitives in classical computing are not available for quantum computation. For example, initializing an n -quantum bit register requires $\Omega(2^n)$ gates rather than n , since an amplitude α_b is set for each bit string. Moreover, some architectures only allow gate operations on neighboring qubits.

We next describe how to optimally initialize a quantum register from $|0\rangle$ to a given n -qubit state $|\phi\rangle$. Suppose first that the vector describing ϕ has only real entries. Partition the vector representing $|\phi\rangle$ into 2-element blocks, and consider each as a vector in \mathbb{R}^2 . Let the j -th such vector have length λ_j and form an angle of θ_j with the x -axis. Taking $|\phi'\rangle = \sum \lambda_j |j\rangle$ and $\delta = \sum \theta_j |j\rangle \langle j|$, we see that $\exp(i\delta \otimes \sigma_y) |\phi'\rangle |0\rangle = |\phi\rangle$. The recursive technique suggested by this equation yields a circuit with $2^n - 2$ CNOTs. The real $|\phi\rangle$ differs from the general case by a diagonal unitary d , which differ from a uniformly controlled R_z by an $n-1$ qubit diagonal d' . Hence a circuit for d . Cancellations between these two circuits reduce the count to $2^{n+1} - 2n - 2$ CNOTs [15].

Second, we note that our circuits adapt well to qubit-chain libraries, where qubits are ordered in a sequence and only CNOTs between adjacent qubits are allowed. Most CNOT gates used in our decomposition already act on nearest neighbors, e.g. those gates implementing the two-qubit operators. Moreover, Fig. 1 shows that only 2^{n-k} CNOT gates of length k (where the length of a local CNOT is 1) will appear in the cir-

Fig. 3. Implementing a long-range CNOT gate with nearest-neighbor CNOTs.



	Number of Qubits						
	1	2	3	4	5	6	7
VMS [19]	0	4	64	536	4156	22618	108760
MVBS [11]	0	8	48	224	960	3968	16128
NQ	0	3	21	105	465	1953	8001

TABLE I

A COMPARISON OF CNOT-COUNTS VERSUS NUMBER OF QUBITS FOR SEVERAL RECENT SYNTHESIS ALGORITHMS.

cuit implementing a uniformly-controlled rotation with n control bits. Fig. 3 decomposes a length k CNOT into $4k - 4$ length 1 CNOTs. It follows that $9 \times 2^{n-1} - 8$ nearest-neighbor CNOTs suffice to implement the uniformly controlled rotation. Therefore restricting CNOT gates to nearest-neighbor interactions increases CNOT count by at most a factor of nine.

VI. CONCLUSIONS AND FUTURE WORK

Our approach to quantum circuit synthesis emphasizes simplicity, a well-pronounced top-down structure, and practical computation via the Cosine-Sine decomposition. By introducing the quantum multiplexor, we have reinterpreted the Cosine-Sine decomposition to allow recursive implementation of quantum gates. While applying our methods to the problem of 3-qubit circuit synthesis is presently the best approach, future specialty techniques developed to solve this problem can be used as terminal cases of our recursion. We have also discussed various problems specific to quantum computation, specifically initialization of quantum registers and mapping to the nearest-neighbor gate library.

As seen in Table I, the universal circuit reported in this work achieves the best known controlled-not counts, both for small numbers of qubits and asymptotically. However, ultimately this just means that our exponentially large circuits are a constant factor smaller than the next best exponentially large circuits. More telling is the fact that our technique performs well in finding small circuits when this is possible.

Acknowledgements. This work is funded by the DARPA QuIST program and an NSF grant. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing official policies or endorsements of employers and funding agencies. We are grateful to Professors Dianne O’Leary from the Univ. of Maryland and Joseph Shinnerl from UCLA for their help with computing the CS decomposition in Matlab; to Gavin Brennen at NIST and Jun Zhang at UC Berkeley for their helpful comments, and the authors of `quant-ph/0406003`, whose package `Qcircuit.tex` produced all figures.

NIST disclaimer. Certain commercial equipment or instruments may be identified in this paper to specify experimental procedures. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology.

E-prints. Most recent references in quantum computing are available only as e-prints at <http://arxiv.org> For example, “e-print, `quant-ph/0406176`” is located at <http://arxiv.org/abs/quant-ph/0406176>

REFERENCES

- [1] A. V. Aho and K. M. Svore. Compiling quantum circuits using the palindrome transform. e-print, `quant-ph/0311008`.
- [2] A. Barenco et al. Elementary gates for quantum computation. *Physical Review A*, 52:3457, 1995.
- [3] C. H. Bennett and G. Brassard. Quantum cryptography: Public-key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, page 175179, Bangalore, India, 1984. IEEE Press.
- [4] S. S. Bullock and I. L. Markov. An elementary two-qubit quantum computation in twenty-three elementary gates. In *Proceedings of the 40th ACM/IEEE Design Automation Conference*, pages 324–329, Anaheim, CA, June 2003. Journal: *Physical Review A* 68, p. 012318 (2003).
- [5] G. Cybenko. Reducing quantum computations to elementary unitary operations. *Computing in Sci. and Engineering*, 3:27–32, March 2001.
- [6] D. P. DiVincenzo. Two-bit gates are universal for quantum computation. *Physical Review A*, 15:1015, 1995.
- [7] R. P. Feynman. Quantum mechanical computers. *Found. Phys.*, 16:507–531, 1986.
- [8] L. K. Grover. Quantum mechanics helps with searching for a needle in a haystack. *Physical Review Letters*, 79:325, 1997.
- [9] W. N. N. Hung, X. Song, G. Yang, J. Yang, and M. Perkowski. Quantum logic synthesis by symbolic reachability analysis. In *Proceedings of the 41st Design Automation Conference*, San Diego, CA, June 2004.
- [10] K. Iwama, Y. Kambayashi, and S. Yamashita. Transformation rules for designing cnot-based quantum circuits. In *Proceedings of the 39th Design Automation Conference*, pages 419–425, 2002.
- [11] M. Mottonen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa. Quantum circuits for general multiqubit gates. *Physical Review Letters*, 93:130502, September 2004. e-print, `quant-ph/0404089`.
- [12] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [13] C. C. Paige and M. Wei. History and generality of the cs decomposition. *Linear Algebra and Applications*, 208:303, 1994.
- [14] V. V. Shende, S. S. Bullock, and I. L. Markov. Recognizing small-circuit structure in two-qubit operators. *Physical Review A*, 70:012310, 2004.
- [15] V. V. Shende and I. L. Markov. Quantum circuits for incompletely specified operators. e-print, `quant-ph/0401162`.
- [16] V. V. Shende, I. L. Markov, and S. S. Bullock. Smaller two-qubit circuits for quantum communication and computation. In *Design, Automation, and Test in Europe*, pages 980–985, Paris, France, February 2004. Journal version in *Physical Review A* 69, p. 062321 (2004).
- [17] V. V. Shende, A. K. Prasad, I. L. Markov, and J. P. Hayes. Synthesis of reversible logic circuits. *IEEE Transactions on Computer Aided Design*, 22:710, 2003.
- [18] P. Shor. Polynomial-time algorithms for prime factorization and discrete logarithm on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
- [19] J. J. Vartiainen, M. Mottonen, and M. M. Salomaa. Efficient decomposition of quantum gates. *Physical Review Letters*, 92:177902, 2004.
- [20] F. Vatan and C. Williams. Realization of a general three-qubit quantum gate. e-print, `quant-ph/0401178`.
- [21] F. Vatan and C. Williams. Optimal quantum circuits for general two-qubit gates. *Physical Review A*, 69:032315, 2004.
- [22] G. Vidal and C. M. Dawson. A universal quantum circuit for two-qubit transformations with three cnot gates. *Phys. Rev. A*, 69:010301, 2004.
- [23] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley. Exact two-qubit universal quantum circuit. *Physical Review Letters*, 91:027903, 2003.
- [24] V. V. Zhirnov, R. K. Cavin, J. A. Hutchby, and G. I. Bourianoff. Limits to binary logic switch scaling — a gedanken model. *Proceedings of the IEEE*, 91(11):1934–1939, November 2003.