Elementary gates for quantum computation

Adriano Barenco
Oxford University [*]

Charles H. Bennett
IBM Research [†]

Richard Cleve
University of Calgary [‡]

David P. DiVincenzo
IBM Research [†]

Norman Margolus
MIT [§]

Peter Shor
AT&T Bell Labs [¶]

Tycho Sleator
New York Univ. [‖]

John Smolin
UCLA [**]

Harald Weinfurter
Univ. of Innsbruck [††]

Version 09.1; March 17, 1995

Abstract

We show that a set of gates that consists of all one-bit quantum gates (U(2)) and the two-bit exclusive-or gate (that maps Boolean values (x, y) to $(x, x \oplus y)$) is universal in the sense that all unitary operations on arbitrarily many bits n (U(2^n)) can be expressed as compositions of these gates. We investigate the number of the above gates required to implement other gates, such as generalized Deutsch-Toffoli gates, that apply a specific U(2) transformation to one input bit if and only if the logical AND of all remaining input bits is satisfied. These gates play a central role in many proposed constructions of quantum computational networks. We derive upper and lower bounds on the exact number of elementary gates required to build up a variety of two- and three-bit quantum gates, the asymptotic number required for n-bit Deutsch-Toffoli gates, and make some observations about the number required for arbitrary n-bit unitary operations.

PACS numbers: 03.65.Ca, 07.05.Bx, 02.70.Rw, 89.80.+h

[*] Clarendon Laboratory, Oxford OX1 3PU, UK
[†] Yorktown Heights, New York, NY 10598, USA.
[‡] Department of Computer Science, Calgary, Alberta, Canada T2N 1N4.
[§] Laboratory for Computer Science, Cambridge MA 02139 USA.
[¶] Murray Hill, NJ 07974 USA.
[‖] Physics Dept., New York, NY 10003 USA.
[**] Physics Department, Los Angeles, CA 90024, USA; also IBM Research.
[††] Inst. for Exptl. Physics, A-6020 Innsbruck, Austria.

1 Background

It has recently been recognized, after fifty years of using the paradigms of classical physics (as embodied in the Turing machine) to build a theory of computation, that quantum physics provides another paradigm with clearly different and possibly much more powerful features than established computational theory. In quantum computation, the state of the computer is described by a state vector Ψ, which is a complex linear superposition of all binary states of the bits $x_m \in \{0,1\}$:

$$\Psi(t) = \sum_{x \in \{0,1\}^m} \alpha_x |x_1, \ldots, x_m\rangle, \quad \sum_x |\alpha_x|^2 = 1.$$

The state's evolution in the course of time t is described by a *unitary* operator U on this vector space, i.e., a linear transformation which is bijective and length-preserving. This unitary evolution on a normalized state vector is known to be the correct physical description of an isolated system evolving in time according to the laws of quantum mechanics[1].

Historically, the idea that the quantum mechanics of isolated systems should be studied as a new formal system for computation arose from the recognition twenty years ago that computation could be made reversible within the paradigm of classical physics. It is possible to perform any computation in a way that is reversible both *logically*—i.e., the computation is a sequence of bijective transformations—and *thermodynamically*—the computation could in principle be performed by a physical apparatus dissipating arbitrarily little energy [2]. A formalism for constructing reversible Turing machines and reversible gate arrays (i.e., reversible combinational logic) was developed. Fredkin and Toffoli[3] showed that there exists a 3-bit "universal gate" for reversible computation, that is, a gate which, when applied in succession to different triplets of bits in a gate array, could be used to simulate any arbitrary reversible computation. (Two-bit gates like NAND which are universal for ordinary computation are not reversible.) Toffoli's version[4] of the universal reversible gate will figure prominently in the body of this paper.

Quantum physics is also reversible, because the reverse-time evolution specified by the unitary operator $U^{-1} = U^\dagger$ always exists; as a consequence, several workers recognized that reversible computation could be executed within a quantum-mechanical system. Quantum-mechanical Turing machines [5, 6], gate arrays [7], and cellular automata [8] have been discussed, and physical realizations of Toffoli's[9, 10, 11] and Fredkin's[12, 13, 14] universal three-bit gates within various quantum-mechanical physical systems have been proposed.

While reversible computation is contained within quantum mechanics, it is a small subset: the time evolution of a classical reversible computer is described by unitary operators whose matrix elements are only zero or one — arbitrary complex numbers are not allowed. Unitary time evolution can of course be *simulated* by a classical computer (e.g., an analog optical computer governed by Maxwell's equations)[15], but the dimension of the unitary operator thus attainable is bounded by the number of classical degrees of freedom—i.e., roughly proportional to the size of the apparatus.

By contrast a quantum computer with m physical bits (see definition of the state above) can perform unitary operations in a space of 2^m dimensions, exponentially larger than its physical size.

Deutsch[16] introduced a quantum Turing machine intended to generate and operate on arbitrary superpositions of states, and proposed that, aside from simulating the evolution of quantum systems more economically than known classical methods, it might also be able to solve certain *classical* problems—i.e., problems with a classical input and output—faster than on any classical Turing machine. In a series of artificial settings, with appropriately chosen oracles, quantum computers were shown to be qualitatively stronger than classical ones [17, 18, 19, 20], culminating in Shor's [21, 22] discovery of quantum polynomial time algorithms for two important natural problems, viz. factoring and discrete logarithm, for which no polynomial-time classical algorithm was known. The search for other such problems, and the physical question of the feasibility of building a quantum computer, are major topics of investigation today.

The formalism we use for quantum computation, which we call a quantum "gate array" was introduced by Deutsch [23], who showed that a simple generalization of the Toffoli gate (the three-bit gate $\wedge_2(R_x)$, in the language introduced later in this paper) suffices as a universal gate for quantum computing. The quantum gate array is the natural quantum generalization of acylic combinational logic "circuits" studied in conventional computational complexity theory. It consists of quantum "gates", interconnected without fanout or feedback by quantum "wires". The gates have the same number of inputs as outputs, and a gate of n inputs represents a unitary operation of the group $\mathrm{U}(2^n)$, i.e., a generalized rotation in a Hilbert space of dimension 2^n. Each wire represents a quantum bit, or *qubit* [25, 26], i.e., a quantum system with a 2-dimensional Hilbert space, capable of existing in a superposition of Boolean states and of being entangled with the states of other qubits. Where there is no danger of confusion, we will use the term "bit" in either the classical or quantum sense. Just as classical bit strings can represent the discrete states of arbitrary finite dimensionality, so a string of n qubits can be used to represent quantum states in any Hilbert space of dimensionality up to 2^n. The analysis of quantum Turing machines [20] is complicated by the fact that not only the data but also the control variables, e.g., head position, can exist in a superposition of classical states. Fortunately, Yao has shown [27] that acyclic quantum gate arrays can simulate quantum Turing machines. Gate arrays are easier to think about, since the control variables, i.e., the wiring diagram itself and the number of steps of computation executed so far, can be thought of as classical, with only the data in the wires being quantum.

Here we derive a series of results which provide new tools for the building-up of unitary transformations from simple gates. We build on other recent results which simplify and extend Deutsch's original discovery[23] of a three-bit universal quantum logic gate. As a consequence of the greater power of quantum computing as a formal system, there are many more choices for the universal gate than in classical reversible computing. In particular, DiVincenzo[28] showed that two-bit universal quantum gates are also possible; Barenco[29] extended this to show than almost any two-bit

gate (within a certain restricted class) is universal, and Lloyd[33] and Deutsch *et al.*[34] have shown that in fact almost any two-bit or n-bit ($n \geq 2$) gate is also universal. A closely related construction for the Fredkin gate has been given[35]. In the present paper we take a somewhat different tack, showing that a non-universal, classical two-bit gate, in conjunction with quantum one-bit gates, is also universal; we believe that the present work along with the preceding ones cover the full range of possible repertoires for quantum gate array construction.

With our universal-gate repertoire, we also exhibit a number of efficient schemes for building up certain classes of n-bit operations with these gates. A variety of strategies for constructing gate arrays efficiently will surely be very important for understanding the full power of quantum mechanics for computation; construction of such efficient schemes have already proved very useful for understanding the scaling of Shor's prime factorization[41]. In the present work we in part build upon the strategy introduced by Sleator and Weinfurter[9], who exhibited a scheme for obtaining the Toffoli gate with a sequence of exactly five two-bit gates. We find that their approach can be generalized and extended in a number of ways to obtain more general efficient gate constructions. Some of the results presented here have no obvious connection with previous gate-assembly schemes.

We will not touch at all on the great difficulties attendant on the actual physical realization of a quantum computer — the problems of error correction[36] and quantum coherence[37, 38] are very serious ones. We refer the reader to [39] for a comprehensive discussion of these difficulties.

2 Introduction

We begin by introducing some basic ideas and notation. For any unitary

$$U = \begin{pmatrix} u_{00} & u_{01} \\ u_{10} & u_{11} \end{pmatrix},$$

and $m \in \{0, 1, 2, \ldots\}$, define the $(m+1)$-bit ($2^{(m+1)}$-dimensional) operator $\wedge_m(U)$ as

$$\wedge_m(U)(|x_1, \ldots, x_m, y\rangle) = \begin{cases} u_{y0}|x_1, \ldots, x_m, 0\rangle + u_{y1}|x_1, \ldots, x_n, 1\rangle & \text{if } \bigwedge_{k=1}^m x_k = 1 \\ |x_1, \ldots, x_m, y\rangle & \text{if } \bigwedge_{k=1}^m x_k = 0, \end{cases}$$

for all $x_1, \ldots, x_m, y \in \{0, 1\}$. (In more ordinary language, $\bigwedge_{k=1}^m x_k$ denotes the AND of the boolean variables $\{x_k\}$.) Note that $\wedge_0(U)$ is equated with U. The $2^{(m+1)} \times 2^{(m+1)}$ matrix corresponding to $\wedge_m(U)$ is

$$\begin{pmatrix} 1 & & & & & \\ & 1 & & & & \\ & & \ddots & & & \\ & & & 1 & & \\ & & & & u_{00} & u_{01} \\ & & & & u_{10} & u_{11} \end{pmatrix}$$

(where the basis states are lexicographically ordered, i.e., $|000\rangle, |001\rangle, \ldots, |111\rangle$).

When

$$U = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

$\wedge_m(U)$ is the so-called Toffoli gate[4] with $m+1$ input bits, which maps $|x_1, \ldots, x_m, y\rangle$ to $|x_1, \ldots, x_m, (\bigwedge_{k=1}^{m} x_k) \oplus y\rangle$. For a general U, $\wedge_m(U)$ can be regarded as a generalization of the Toffoli gate, which, on input $|x_1, \ldots, x_m, y\rangle$, applies U to y if and only if $\bigwedge_{k=1}^{m} x_k = 1$.

As shown by one of us [29], "almost any" single $\wedge_1(U)$ gate is universal in the sense that: by successive application of this gate to pairs of bits in an n-bit network, any unitary transformation may be approximated with arbitrary accuracy. (It suffices for U to be specified by Euler angles which are not a rational multiple of π.)

We show that in some sense this result can be made even simpler, in that any unitary transformation in a network can always be constructed out of only the "classical" two-bit gate $\wedge_1\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ along with a set of one-bit operations (of the form $\wedge_0(U)$). This is a remarkable result from the perspective of classical reversible computation because it is well known that the classical analogue of this assertion—which is that all invertible boolean functions can be implemented with $\wedge_1\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ and $\wedge_0\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ gates[42] — is false. In fact, it is well known that only a tiny fraction of Boolean functions (those which are linear with respect to modulo 2 arithmetic) can be generated with these gates[30].

We will also exhibit a number of explicit constructions of $\wedge_m(U)$ using $\wedge_1(U)$, which can all be made polynomial in m. It is well known[4] that the analogous family of constructions in classical reversible logic which involve building $\wedge_m\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ from the three-bit Toffoli gate $\wedge_2\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$, is also polynomial in m. We will exhibit one important difference between the classical and the quantum constructions, however; Toffoli showed[4] that the classical $\wedge_m$'s could not be built without the presence of some "work bits" to store intermediate results of the calculation. By contrast, we show that the quantum logic gates can always be constructed with the use of *no* workspace whatsoever. Similar computations in the classical setting (that use very few or no work bits) appeared in the work of Cleve[31] and Ben-Or and Cleve[32]. Still, the presence of a workspace plays an important role in the quantum gate constructions — we find that to implement a family of $\wedge_m$ gates exactly, the time required for our implementation can be reduced from $\Theta(m^2)$ to $\Theta(m)$ merely by the introduction of *one* bit for workspace.

3 Notation

We adopt a version of Feynman's[7] notation to denote $\wedge_m(U)$ gates and Toffoli gates in quantum networks as follows.

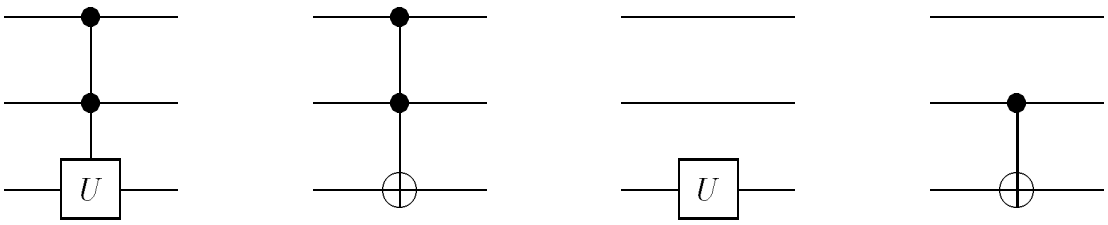

The first network contains a $\wedge_2(U)$ gate and the second one contains a 3-bit Toffoli gate[43]. The third and fourth networks contain a $\wedge_0(U)$ and a 2-bit reversible exclusive-or (simply called XOR henceforth) gate, respectively. The XOR gate is introduced as the "measurement gate" in [23], and will play a very prominent role in many of the constructions we describe below. Throughout this paper, when we refer to a *basic* operation, we mean either a $\wedge_0(U)$ gate or this 2-bit XOR gate.

4 Matrix Properties

Lemma 4.1: *Every unitary 2×2 matrix can be expressed as*

$$\begin{pmatrix} e^{i\delta} & 0 \\ 0 & e^{i\delta} \end{pmatrix} \cdot \begin{pmatrix} e^{-i\alpha/2} & 0 \\ 0 & e^{i\alpha/2} \end{pmatrix} \cdot \begin{pmatrix} \cos\theta/2 & \sin\theta/2 \\ -\sin\theta/2 & \cos\theta/2 \end{pmatrix} \cdot \begin{pmatrix} e^{-i\beta/2} & 0 \\ 0 & e^{i\beta/2} \end{pmatrix},$$

I have gone to the half-angle convention more usual in physics.– DDV

where δ, α, θ, and β are real-valued. Moreover, any special unitary 2×2 matrix (i.e., with unity determinant) can be expressed as

$$\begin{pmatrix} e^{-i\alpha/2} & 0 \\ 0 & e^{i\alpha/2} \end{pmatrix} \cdot \begin{pmatrix} \cos\theta/2 & \sin\theta/2 \\ -\sin\theta/2 & \cos\theta/2 \end{pmatrix} \cdot \begin{pmatrix} e^{-i\beta/2} & 0 \\ 0 & e^{i\beta/2} \end{pmatrix}.$$

Proof: Since a matrix is unitary if and only if its row vectors and column vectors are orthonormal, every 2×2 unitary matrix is of the form

$$\begin{pmatrix} e^{i(\delta-\alpha/2-\beta/2)}\cos\theta/2 & e^{i(\delta-\alpha/2+\beta/2)}\sin\theta/2 \\ -e^{i(\delta+\alpha/2-\beta/2)}\sin\theta/2 & e^{i(\delta+\alpha/2+\beta/2)}\cos\theta/2 \end{pmatrix},$$

where δ, α, θ, and β are real-valued. The first factorization above now follows immediately. In the case of special unitary matrices, the determinant of the first matrix must be 1, which implies $e^{i\delta} = \pm 1$, so the first matrix in the product can be absorbed into the second one.□

Definition: In view of the above lemma, we define the following.

- $\mathrm{Rot}(\theta) = \begin{pmatrix} \cos\theta/2 & \sin\theta/2 \\ -\sin\theta/2 & \cos\theta/2 \end{pmatrix}$ (a "rotation" by θ[40]).

- $\mathrm{Ph}(\alpha) = \begin{pmatrix} e^{-i\alpha/2} & 0 \\ 0 & e^{i\alpha/2} \end{pmatrix}$ (a "phase shift" with respect to α).
- $\mathrm{Scal}(\delta) = \begin{pmatrix} e^{i\delta} & 0 \\ 0 & e^{i\delta} \end{pmatrix}$ (a "scale" with respect to δ).
- $\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ (a "negation", or Pauli matrix).
- $I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ (the identity matrix).

Lemma 4.2: *The following properties hold:*

1. $\mathrm{Rot}(\theta_1) \cdot \mathrm{Rot}(\theta_2) = \mathrm{Rot}(\theta_1 + \theta_2)$
2. $\mathrm{Ph}(\alpha_1) \cdot \mathrm{Ph}(\alpha_2) = \mathrm{Ph}(\alpha_1 + \alpha_2)$
3. $\mathrm{Scal}(\delta_1) \cdot \mathrm{Scal}(\delta_2) = \mathrm{Scal}(\delta_1 + \delta_2)$
4. $\sigma_x \cdot \sigma_x = I$
5. $\sigma_x \cdot \mathrm{Rot}(\theta) \cdot \sigma_x = \mathrm{Rot}(-\theta)$
6. $\sigma_x \cdot \mathrm{Ph}(\alpha) \cdot \sigma_x = \mathrm{Ph}(-\alpha)$

Lemma 4.3: *For any special unitary matrix W ($W \in SU(2)$), there exist matrices A, B, and $C \in SU(2)$ such that $A \cdot B \cdot C = I$ and $A \cdot \sigma_x \cdot B \cdot \sigma_x \cdot C = W$.*

Proof: By Lemma 4.1, there exist α, θ, and β such that $W = \mathrm{Ph}(\alpha) \cdot \mathrm{Rot}(\theta) \cdot \mathrm{Ph}(\beta)$. Set $A = \mathrm{Ph}(\alpha) \cdot \mathrm{Rot}(\frac{\theta}{2})$, $B = \mathrm{Rot}(-\frac{\theta}{2}) \cdot \mathrm{Ph}(-\frac{\alpha+\beta}{2})$, and $C = \mathrm{Ph}(\frac{\beta-\alpha}{2})$. Then

$$\begin{aligned} A \cdot B \cdot C &= \mathrm{Ph}(\alpha) \cdot \mathrm{Rot}(\tfrac{\theta}{2}) \cdot \mathrm{Rot}(-\tfrac{\theta}{2}) \cdot \mathrm{Ph}(-\tfrac{\alpha+\beta}{2}) \cdot \mathrm{Ph}(\tfrac{\beta-\alpha}{2}) \\ &= \mathrm{Ph}(\alpha) \cdot \mathrm{Ph}(-\alpha) \\ &= I, \end{aligned}$$

and

$$\begin{aligned} A \cdot \sigma_x \cdot B \cdot \sigma_x \cdot C &= \mathrm{Ph}(\alpha) \cdot \mathrm{Rot}(\tfrac{\theta}{2}) \cdot \sigma_x \cdot \mathrm{Rot}(-\tfrac{\theta}{2}) \cdot \mathrm{Ph}(-\tfrac{\alpha+\beta}{2}) \cdot \sigma_x \cdot \mathrm{Ph}(\tfrac{\beta-\alpha}{2}) \\ &= \mathrm{Ph}(\alpha) \cdot \mathrm{Rot}(\tfrac{\theta}{2}) \cdot \sigma_x \cdot \mathrm{Rot}(-\tfrac{\theta}{2}) \cdot \sigma_x \cdot \sigma_x \cdot \mathrm{Ph}(-\tfrac{\alpha+\beta}{2}) \cdot \sigma_x \cdot \mathrm{Ph}(\tfrac{\beta-\alpha}{2}) \\ &= \mathrm{Ph}(\alpha) \cdot \mathrm{Rot}(\tfrac{\theta}{2}) \cdot \mathrm{Rot}(\tfrac{\theta}{2}) \cdot \mathrm{Ph}(\tfrac{\alpha+\beta}{2}) \cdot \mathrm{Ph}(\tfrac{\beta-\alpha}{2}) \\ &= \mathrm{Ph}(\alpha) \cdot \mathrm{Rot}(\theta) \cdot \mathrm{Ph}(\beta) \\ &= W. \end{aligned}$$

□

5 Two-Bit Networks

5.1 Simulation of General $\wedge_1(U)$ Gates

Lemma 5.1: *For a unitary 2×2 matrix W, a $\wedge_1(W)$ gate can be simulated by a network of the form*

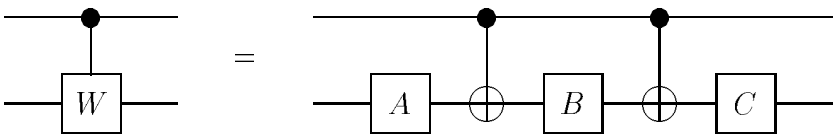

where A, B, and $C \in SU(2)$, if and only if $W \in SU(2)$.

Proof: For the "if" part, let A, B, and C be as in Lemma 4.3. If the value of the first (top) bit is 0 then $A \cdot B \cdot C = I$ is applied to the second bit. If the value of the first bit is 1 then $A \cdot \sigma_x \cdot B \cdot \sigma_x \cdot C = W$ is applied to the second bit.

For the "only if" part, note that $A \cdot B \cdot C = I$ must hold if the simulation is correct when the first bit is 0. Also, if the network simulates a $\wedge_1(W)$ gate then $A \cdot \sigma_x \cdot B \cdot \sigma_x \cdot C = W$. Therefore, since $\det(A \cdot \sigma_x \cdot B \cdot \sigma_x \cdot C) = 1$, W must also be special unitary.□

Lemma 5.2: *For any δ and $S = Scal(\delta)$, a $\wedge_1(S)$ gate can be simulated by a network of the form*

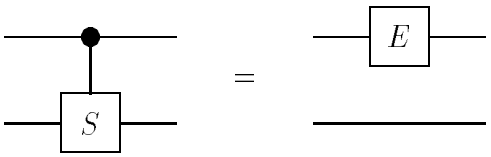

where E is unitary.

Proof: Let

$$E = \mathrm{Ph}(\delta) \cdot \mathrm{Scal}(\tfrac{\delta}{2}) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\delta} \end{pmatrix}.$$

Then the observation is that the 4×4 unitary matrix corresponding to each of the above networks is

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & e^{i\delta} & 0 \\ 0 & 0 & 0 & e^{i\delta} \end{pmatrix}.$$

□

Clearly, $\wedge_1(S)$ composed with $\wedge_1(W)$ yields $\wedge_1(S \cdot W)$. Thus, by noting that any unitary matrix U is of the form $U = S \cdot W$, where $S = \mathrm{Scal}(\delta)$ (for some δ) and W is $\in SU(2)$, we obtain the following.

Corollary 5.3: *For any unitary 2×2 matrix U, a $\wedge_1(U)$ gate can be simulated by at most six basic gates: four 1-bit gates ($\wedge_0$), and two XOR gates ($\wedge_1(\sigma_x)$).*

5.2 Special Cases

In Section 5.1, we have established a general simulation of a $\wedge_1(U)$ gate for an arbitrary unitary U. For special cases of U that may be of interest, a more efficient construction than that of Corollary 5.3 is possible. Clearly, Lemma 5.1 immediately yields a more efficient simulation for all special unitary matrices. For example, the "x-axis rotation matrix" (to use the language suggested by the mapping between SU(2) and SO(3), the group of rigid-body rotations[44])

$$R_x = \begin{pmatrix} \cos\theta/2 & i\sin\theta/2 \\ i\sin\theta/2 & \cos\theta/2 \end{pmatrix} = \mathrm{Ph}(-\tfrac{\pi}{2}) \cdot \mathrm{Rot}(\theta) \cdot \mathrm{Ph}(\tfrac{\pi}{2})$$

is special unitary. (R_x is of special interest because $\wedge_2(R_x)$ is the "Deutsch gate"[23], which was shown to be universal for quantum logic.) For other specific SU(2) matrices an even more efficient simulation is possible.

Lemma 5.4: *A $\wedge_1(W)$ gate can be simulated by a network of the form*

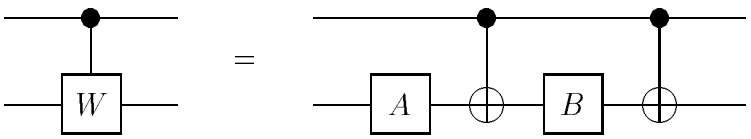

where A and $B \in SU(2)$ if and only if W is of the form

$$W = \mathrm{Ph}(\alpha) \cdot \mathrm{Rot}(\theta) \cdot \mathrm{Ph}(\alpha) = \begin{pmatrix} e^{-i\alpha}\cos\theta/2 & \sin\theta/2 \\ -\sin\theta/2 & e^{i\alpha}\cos\theta/2 \end{pmatrix},$$

where α and θ are real-valued.

Proof: For the "if" part, consider the simulation of $\wedge_1(W)$ that arises in Lemma 5.1 when $W = \mathrm{Ph}(\alpha) \cdot \mathrm{Rot}(\theta) \cdot \mathrm{Ph}(\alpha)$. In this case, $A = \mathrm{Ph}(\alpha) \cdot \mathrm{Rot}(\frac{\theta}{2})$, $B = \mathrm{Rot}(-\frac{\theta}{2}) \cdot \mathrm{Ph}(-\alpha)$ and $C = I$. Thus, $B = A^\dagger$ and C can be omitted.

For the "only if" part, note that $B = A^\dagger$ must hold for the simulation to be valid when when the first bit is 0. Therefore, if the first bit is 1 then $A \cdot \sigma_x \cdot A^\dagger \cdot \sigma_x$ is applied to the second bit. Now, the matrix $A \cdot \sigma_x \cdot A^\dagger$ has determinant -1 and is traceless (since its trace is the same as that of σ_x). By specializing the characterization of unitary matrices in Lemma 4.1 to traceless matrices with determinant -1, we conclude that $A \cdot \sigma_x \cdot A^\dagger$ must be of the form

$$A \cdot \sigma_x \cdot A^\dagger = \begin{pmatrix} \sin\theta/2 & e^{-i\alpha}\cos\theta/2 \\ e^{i\alpha}\cos\theta/2 & -\sin\theta/2 \end{pmatrix}.$$

Therefore,

$$A \cdot \sigma_x \cdot A^\dagger \cdot \sigma_x = \begin{pmatrix} e^{-i\alpha}\cos\theta/2 & \sin\theta/2 \\ -\sin\theta/2 & e^{i\alpha}\cos\theta/2 \end{pmatrix},$$

as required.□

Examples of matrices of the form of Lemma 5.4 are the y- and z-axis rotation matrices $R_y = \mathrm{Rot}(\theta) = \mathrm{Ph}(0) \cdot \mathrm{Rot}(\theta) \cdot \mathrm{Ph}(0)$ and $R_z = \mathrm{Ph}(\alpha) = \mathrm{Ph}(\frac{\alpha}{2}) \cdot \mathrm{Rot}(0) \cdot \mathrm{Ph}(\frac{\alpha}{2})$ (but R_x is not of this form).

Finally, for certain U, we obtain an even greater simplification of the simulation of $\wedge_1(U)$ gates.

Lemma 5.5: *A $\wedge_1(V)$ gate can be simulated by a construction of the form*

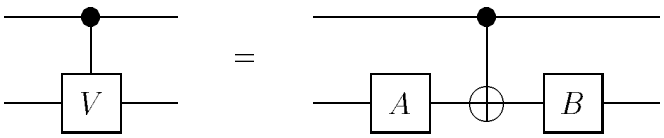

where A and B are unitary if and only if V is of the form

$$V = \mathrm{Ph}(\alpha) \cdot \mathrm{Rot}(\theta) \cdot \mathrm{Ph}(\alpha) \cdot \sigma_x = \begin{pmatrix} \sin\theta/2 & e^{-i\alpha}\cos\theta/2 \\ e^{i\alpha}\cos\theta/2 & -\sin\theta/2 \end{pmatrix},$$

where α and θ are real-valued.

Proof: If an additional $\wedge_1(\sigma_x)$ is appended to the end of the network in Lemma 5.4 then, the network is equivalent to that above (since $\wedge_1(\sigma_x)$ is an involution), and also simulates a $\wedge_1(W \cdot \sigma_x)$ gate (since $\wedge_1(W)$ composed with $\wedge_1(\sigma_x)$ is $\wedge_1(W \cdot \sigma_x)$).□

Examples of matrices of the form of Lemma 5.5 are the Pauli matrices

$$\sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} = \mathrm{Ph}(\tfrac{\pi}{2}) \cdot \mathrm{Rot}(0) \cdot \mathrm{Ph}(\tfrac{\pi}{2}) \cdot \sigma_x$$

and

$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} = \mathrm{Ph}(0) \cdot \mathrm{Rot}(\pi) \cdot \mathrm{Ph}(0) \cdot \sigma_x$$

(as well as σ_x itself).

Lemma 5.5 permits an immediate generalization of Corollary 5.3:

Corollary 5.6: *For any unitary 2×2 matrix U, a $\wedge_1(U)$ gate can be simulated by at most six basic gates: four 1-bit gates ($\wedge_0$), and two gates ($\wedge_1(V)$), where V is of the form $V = Ph(\alpha) \cdot Rot(\theta) \cdot Ph(\alpha) \cdot \sigma_x$.*

A particular feature of the $\wedge_1(\sigma_z)$ gates is that they are symmetric with respect to their input bits. In view of this, as well as for future reference, we introduce the following special notation for $\wedge_1(\sigma_z)$ gates.

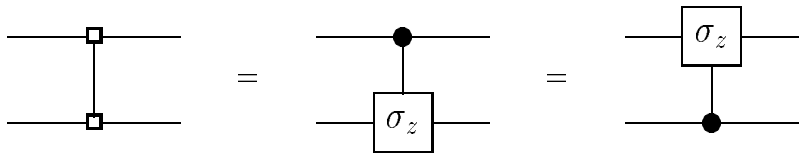

6 Three-Bit Networks

6.1 Simulation of General $\wedge_2(U)$ Gates

Lemma 6.1: *For any unitary 2×2 matrix U, a $\wedge_2(U)$ gate can be simulated by a network of the form*

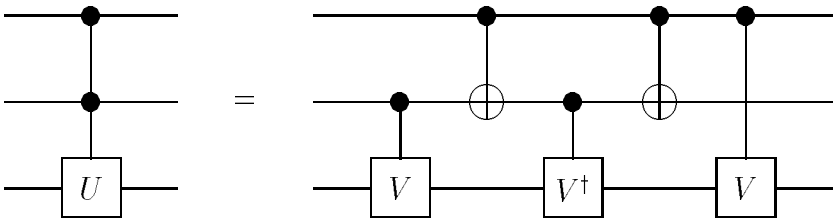

where V is unitary.

Proof: Let V be such that $V^2 = U$. If the first bit or the second bit are 0 then the transformation applied to the third bit is either I or $V \cdot V^\dagger = I$. If the first two bits are both 1 then the transformation applied to the third is $V \cdot V = U$.□

Some of the intuition behind the construction in the above Lemma is that, when the first two input bits are x_1 and x_2, the sequence of operations performed on the third bit is: V iff $x_1 = 1$, V iff $x_2 = 1$, and $V^\dagger$ iff $x_1 \oplus x_2 = 1$. Since

$$x_1 + x_2 - (x_1 \oplus x_2) = 2 \cdot (x_1 \wedge x_2)$$

(where "+", "−", and "·" are the ordinary arithmetic operations), the above sequence of operations is equivalent to performing V^2 on the third bit iff $x_1 \wedge x_2 = 1$, which is the $\wedge_2(V^2)$ gate. (This approach generalizes to produce a simulation of $\wedge_m(V^{2^{m-1}})$, for $m > 2$, which is considered in Section 7.)

We can now combine Lemma 6.1 with Corollary 5.3 to obtain a simulation of $\wedge_2(U)$ using only basic gates ($\wedge_1(\sigma_x)$ and $\wedge_0$). The number of these gates is reduced when it is recognized that a number of the one-bit gates can be *merged* and eliminated. In particular, the $\wedge_0(C)$ from the end of the simulation of the first $\wedge_1(V)$ gate, and the $\wedge_0(C^\dagger)$ from the $\wedge_1(V^\dagger)$ gate combine to form the identity and are eliminated entirely. This same sort of merging occurs to elimiate a $\wedge_0(A)$ gate and a $\wedge_0(A^\dagger)$ gate. We arrive at the following count:

Corollary 6.2: *For any unitary 2×2 matrix U, a $\wedge_2(U)$ gate can be simulated by at most sixteen basic gates: eight 1-bit gates ($\wedge_0$) and eight XOR gates ($\wedge_1(\sigma_x)$).*

A noteworthy case is when $U = \sigma_x$, where we obtain a simulation of the 3-bit Toffoli gate $\wedge_2(\sigma_x)$, which is the primitive gate for classical reversible logic [4]. Later we will use the fact that because $\wedge_2(\sigma_x)$ is its own inverse, either the simulation of Lemma 6.1 or the time-reversed simulation (in which the order of the gates is reversed, and each unitary operator is replaced by its Hermitian conjugate) may be used.

6.2 Three-bit gates congruent to $\wedge_2(U)$

We now show that more efficient simulations of three-bit gates are possible if phase shifts of the quantum states other than zero are permitted. If we define the matix W as

$$W = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} = \mathrm{Scal}(\tfrac{\pi}{2}) \cdot \sigma_y,$$

then the gates $\wedge_2(W)$ and $\wedge_2(\sigma_x)$ can be regarded as being "congruent modulo phase shifts", because the latter gate differs only in that it maps $|111\rangle$ to $-|110\rangle$ (instead of $|110\rangle$). This is perfectly acceptable if the gate is part of an operation which merely mimics classical reversible computation, or if the gate is paired with another similar one to cancel out the extra phase, as is sometimes the case in reversible gate arrangements (see Corollary 7.4); however, this phase difference is dangerous in general if non-classical unitary operations appear in the computation. Gates congruent to $\wedge_2(\sigma_x)$ modulo phase shifts have been previously investigated in [45].

The following is a more efficient simulation of a gate congruent to $\wedge_2(\sigma_x)$ modulo phase shifts:

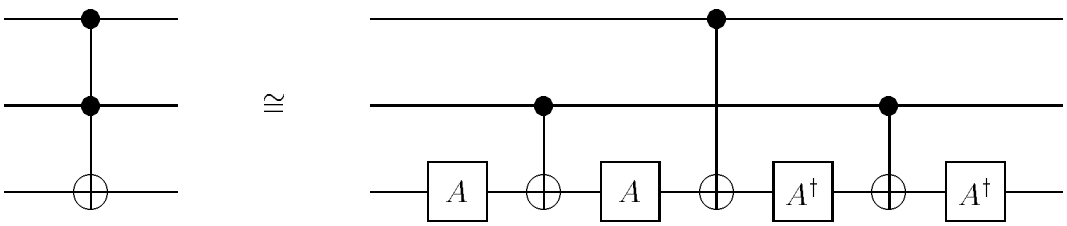

where $A = \mathrm{Rot}(\frac{\pi}{4})$. In the above, the "$\cong$" indicates that the networks are not identical, but differ at most in the phases of their amplitudes, which are all ± 1 (the phase of the $|101\rangle$ state is reversed in this case).

An alternative simulation of a gate congruent to $\wedge_2(\sigma_x)$ modulo phase shifts (whose phase shifts are identical to the previous one) is given by

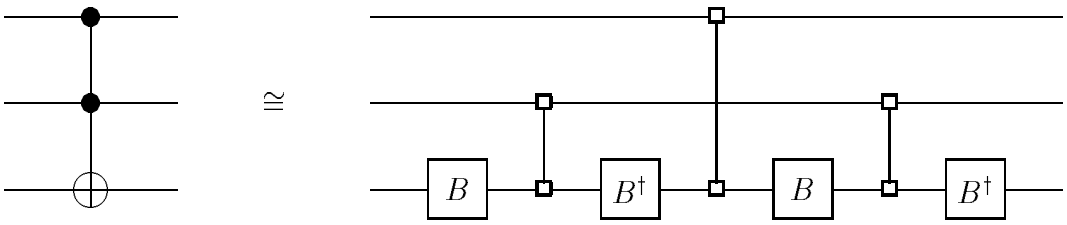

where $B = \mathrm{Rot}(\frac{3\pi}{4})$.

7 n-Bit Networks

The technique for simulating $\wedge_2(U)$ gates in Lemma 6.1 generalizes to $\wedge_m(U)$ gates for $m > 2$. For example, to simulate a $\wedge_3(U)$ gate for any unitary U, set V so that $V^4 = U$ and then construct a network as follows.

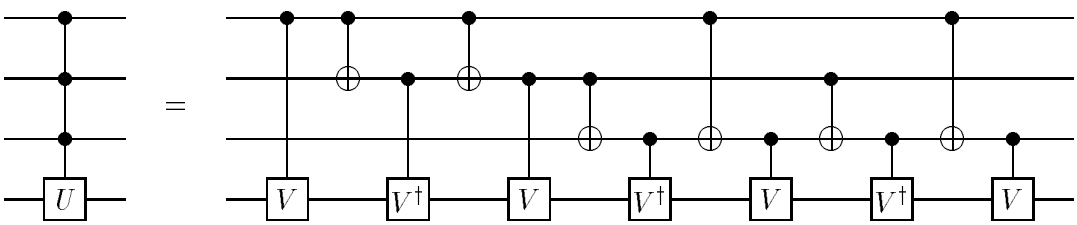

The intuition behind this construction is similar to that behind the construction of Lemma 6.1. If the first three input bits are x_1, x_2, and x_3 then the sequence of operations performed on the fourth bit is:

V	iff $x_1 = 1$	(100)
$V^\dagger$	iff $x_1 \oplus x_2 = 1$	(110)
V	iff $x_2 = 1$	(010)
$V^\dagger$	iff $x_2 \oplus x_3 = 1$	(011)
V	iff $x_1 \oplus x_2 \oplus x_3 = 1$	(111)
$V^\dagger$	iff $x_1 \oplus x_3 = 1$	(101)
V	iff $x_3 = 1$	(001).

The strings on the right encode the condition for the operation V or $V^\dagger$ at each step—the "1"'s indicate which input bits are involved in the condition. For an efficient implementation of $\wedge_3(U)$, these strings form a grey code sequence. Note also that the parity of each bit string determines whether to apply V or $V^\dagger$. By comparing this sequence of operations with the terms in the equation

$$x_1 + x_2 + x_3 - (x_1 \oplus x_2) - (x_1 \oplus x_3) - (x_2 \oplus x_3) + (x_1 \oplus x_2 \oplus x_3) = 4 \cdot (x_1 \wedge x_2 \wedge x_3),$$

it can be verified that the above sequence of operations is equivalent to performing V^4 on the fourth bit iff $x_1 \wedge x_2 \wedge x_3 = 1$, which is the $\wedge_3(V^4)$ gate.

The foregoing can be generalized to simulate $\wedge_m(U)$ for larger values of m.

Lemma 7.1: *For any $n \geq 3$ and any unitary 2×2 matrix U, a $\wedge_{n-1}(U)$ gate can be simulated by an n-bit network consisting of $2^{n-1}-1$ $\wedge_1(V)$ and $\wedge_1(V^\dagger)$ gates and $2^{n-1}-2$ $\wedge_1(\sigma_x)$ gates, where V is unitary.*

We omit the proof of Lemma 7.1, but point out that it is a generalization of the $n = 4$ case above and based on setting V so that $V^{2^{n-2}} = U$ and "implementing" the identity

$$\sum_{k_1} x_{k_1} - \sum_{k_1<k_2} (x_{k_1} \oplus x_{k_2}) + \sum_{k_1<k_2<k_3} (x_{k_1} \oplus x_{k_2} \oplus x_{k_3}) - \cdots + (-1)^{m-1}(x_1 \oplus x_2 \oplus \cdots \oplus x_m)$$
$$= 2^{m-1} \cdot (x_1 \wedge x_2 \wedge \cdots \wedge x_m)$$

with a grey-code sequence of operations.

For some specific small values of n (for $n = 3$, 4, 5, 6, 7, and 8), this is the most efficient technique that we are aware of for simulating arbitrary $\wedge_{n-1}(U)$ gates as well as $\wedge_{n-1}(\sigma_x)$ gates; taking account of mergers (see Corollary 6.2), the simulation requires $3 \cdot 2^{n-1} - 4$ $\wedge_1(\sigma_x)$'s and $2 \cdot 2^{n-1}$ $\wedge_0$'s. However, since this number is $\Theta(2^n)$, the simulation is very inefficient for large values of n. For the remainder of this section, we focus on the asymptotic growth rate of the simulations with respect to n, and show that this can be quadratic in the general case and linear in many cases of interest.

7.1 Linear Simulation of $\wedge_{n-2}(\sigma_x)$ Gates on n-Bit Networks

Lemma 7.2: *If $n \geq 5$ and $m \in \{3, \ldots, \lceil \frac{n}{2} \rceil\}$ then a $\wedge_m(\sigma_x)$ gate can be simulated by a network consisting of $4(m-2)$ $\wedge_2(\sigma_x)$ gates that is of the form*

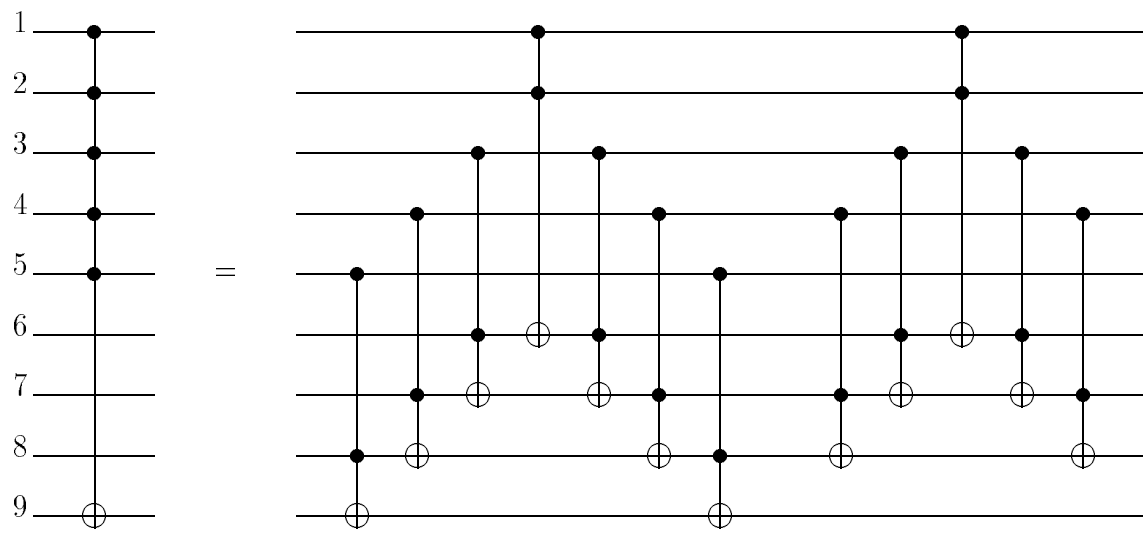

(illustrated for $n = 9$ and $m = 5$).

Proof: Consider the group of the first 7 gates in the above network. The sixth bit (from the top) is negated iff the first two bits are 1, the seventh bit is negated iff the

first three bits are 1, the eighth bit is negated iff the first four bits are 1, and the ninth bit is negated iff the first five bits are 1. Thus, the last bit is correctly set, but the three preceding bits are altered. The last 5 gates in the network reset the values of these three preceding bits.□

Note that in this construction and in the ones following, although many of the bits not involved in the gate are operated upon, the gate operation is performed correctly independent of the initial state of the bits (i.e., they do not have to be "cleared" to 0 first), and they are reset to their initial values after the operations of the gate (as in the computations which occur in [32] and [31]). This fact makes constructions like the following possible.

Lemma 7.3: *For any $n \geq 5$, and $m \in \{2, \ldots, n-3\}$ a $\wedge_{n-2}(\sigma_x)$ gate can be simulated by a network consisting of two $\wedge_m(\sigma_x)$ gates and two $\wedge_{n-m-1}(\sigma_x)$ gates which is of the form*

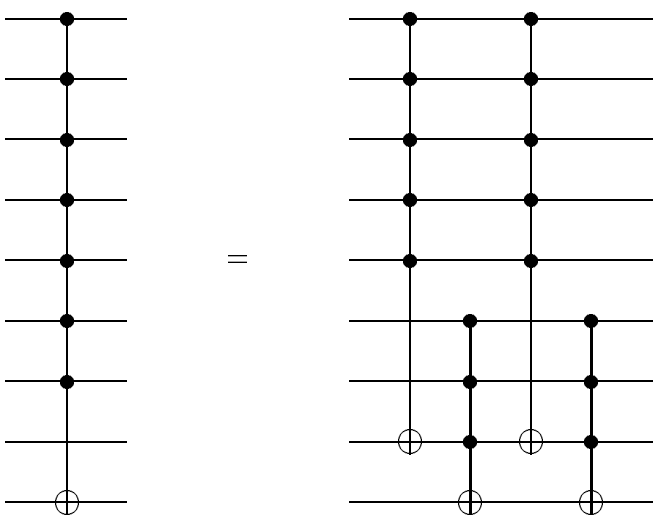

(illustrated for $n = 9$ and $m = 5$).

Proof: By inspection.□

Corollary 7.4: *On an n-bit network (where $n \geq 7$), a $\wedge_{n-2}(\sigma_x)$ gate can be simulated by $8(n-5)$ $\wedge_2(\sigma_x)$ gates (3-bit Toffoli gates), as well as by $48n-204$ basic operations.*

Proof: First apply Lemma 7.2 with $m_1 = \lceil \frac{n}{2} \rceil$ and $m_2 = n - m_1 - 1$ to simulate $\wedge_{m_1}(\sigma_x)$ and $\wedge_{m_2}(\sigma_x)$ gates. Then combine these by Lemma 7.3 to simulate the $\wedge_{n-2}(\sigma_x)$ gate. Then, each $\wedge_2(\sigma_x)$ gate in the above simulation may be simulated by a set of basic operations (as in Corollary 6.2). We find that almost all of these Toffoli gates need only to be simulated modulo phase factors as in Sec. 6.2; in particular, only 4 of the Toffoli gates, the ones which involve the last bit in the diagram above, need to be simulated exactly according to the construction of Corollary 6.2. Thus these 4 gates are simulated by 16 basic operations, while the other $8n - 36$ Toffoli gates are simulated in just 6 basic operations. A careful accounting of the mergers

of $\wedge_0$ gates which are then possible leads to the total count of basic operations given above. $\Box$

The above constructions, though asymptotically efficient, requires at least one "extra" bit, in that an n-bit network is required to simulate the $(n-1)$-bit gate $\wedge_{n-2}(\sigma_x)$. In the next subsection, we shall show how to construct $\wedge_{n-1}(U)$ for an arbitrary unitary U using a quadratic number of basic operations on an n-bit network, which includes the n-bit Toffoli gate $\wedge_{n-1}(\sigma_x)$ as a special case.

7.2 Quadratic Simulation of General $\wedge_{n-1}(U)$ Gates on n-Bit Networks

Lemma 7.5: *For any unitary 2×2 matrix U, a $\wedge_{n-1}(U)$ gate can be simulated by a network of the form*

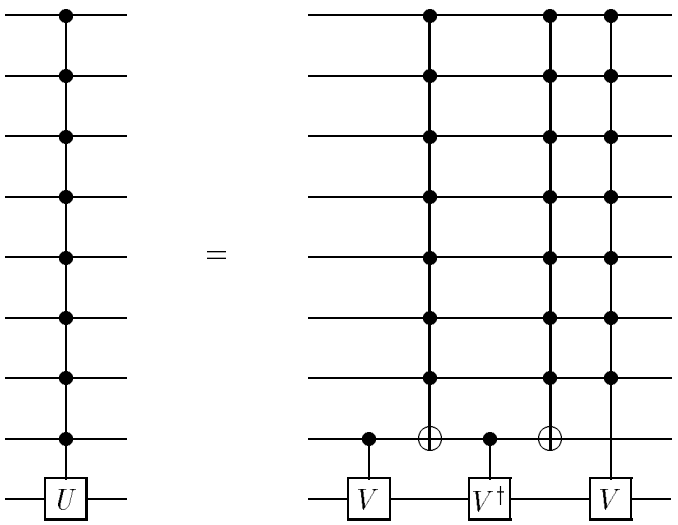

(illustrated for $n = 9$), where V is unitary.

Proof: The proof is very similar to that of Lemma 6.1, setting V so that $V^2 = U$.$\Box$

Corollary 7.6: *For any unitary U, a $\wedge_{n-1}(U)$ gate can be simulated in terms of $\Theta(n^2)$ basic operations.*

Proof: This is a recursive application of Lemma 7.5. Let C_{n-1} denote the cost of simulating a $\wedge_{n-1}(U)$ (for an arbitrary U). Consider the simulation in Lemma 7.5. The cost of simulating the $\wedge_1(V)$ and $\wedge_1(V^\dagger)$ gates is $\Theta(1)$ (by Corollary 5.3). The cost of simulating the two $\wedge_{n-2}(\sigma_x)$ gates is $\Theta(n)$ (by Corollary 7.4). The cost of simulating the $\wedge_{n-2}(V)$ gate (by a recursive application of Lemma 7.5) is C_{n-2}. Therefore, C_{n-1} satisfies a recurrence of the form

$$C_{n-1} = C_{n-2} + \Theta(n),$$

which implies that $C_{n-1} \in \Theta(n^2)$.$\Box$

In fact, we find that using the gate-counting mentioned in Corollary 7.4, the number of basic operations is $48n^2 + O(n)$.
Although Corollary 7.6 is significant in that it permits any $\wedge_{n-1}(U)$ to be simulated with "polynomial complexity", the question remains as to whether a subquadratic simulation is possible. The following is an $\Omega(n)$ lower bound on this complexity.

Lemma 7.7: *Any simulation of a nonscalar $\wedge_{n-1}(U)$ gate (i.e. where $U \neq Scal(\delta) \cdot I$) requires at least $n-1$ basic operations.*

Proof: Consider any n-bit network with arbitrarily many 1-bit gates and fewer than $n-1$ $\wedge_1(\sigma_x)$ gates. Call two bits *adjacent* if there there is a $\wedge_1(\sigma_x)$ gate between them, and *connected* if there is a sequence of consecutively adjacent bits between them. Since there are fewer than $n-1$ $\wedge_1(\sigma_x)$ gates, it must be possible to partition the bits into two nonempty sets $\mathcal{A}$ and $\mathcal{B}$ such that no bit in $\mathcal{A}$ is connected to any bit in $\mathcal{B}$. This implies that the unitary transformation associated with the network is of the form $A \otimes B$, where A is $2^{|\mathcal{A}|}$-dimensional and B is $2^{|\mathcal{B}|}$-dimensional. Since the transformation $\wedge_{n-1}(U)$ is not of this form, the network cannot compute $\wedge_{n-1}(U)$.□

It is conceivable that a linear size simulation of $\wedge_{n-1}(U)$ gates is possible. Although we cannot show this presently, in the remaining subsections, we show that something "similar" (in a number of different senses) to a linear size simulation of $\wedge_{n-1}(U)$ gates is possible.

7.3 Linear Approximate Simulation of General $\wedge_{n-1}(U)$ Gates on n-Bit Networks

Definition: We say that one network *approximates* another one *within* ε if the distance (induced by the Euclidean vector norm) between the unitary transformations associated with the two networks is at most ε.

This notion of approximation in the context of reducing the complexity of quantum computations was introduced by Coppersmith[41], and is useful for the following reason. Suppose that two networks that are approximately the same (in the above sense) are executed with identical inputs and their outputs are observed. Then the probability distributions of the two outcomes will be approximately the same in the sense that, for any event, its probability will differ by at most 2ε between the two networks.

Lemma 7.8: *For any unitary 2×2 matrix U and $\varepsilon > 0$, a $\wedge_{n-1}(U)$ gate can be approximated within ε by $\Theta(n \log(\frac{1}{\varepsilon}))$ basic operations.*

Proof: The idea is to apply Lemma 7.5 recursively as in Corollary 7.6, but to observe that, with suitable choices for V, the recurrence can be terminated after $\Theta(\log(\frac{1}{\varepsilon}))$ levels.

Since U is unitary, there exist unitary matrices P and D, such that $U = P^\dagger \cdot D \cdot P$ and

$$D = \begin{pmatrix} e^{id_1} & 0 \\ 0 & e^{id_2} \end{pmatrix}$$

where d_1 and d_2 are real. e^{id_1} and e^{id_2} are the eigenvalues of U. If V_k is the matrix used in the k^{th} recursive application of Lemma 7.3 ($k \in \{0, 1, 2, \ldots\}$) then it is sufficient that $V_{k+1}^2 = V_k$ for each $k \in \{0, 1, 2, \ldots\}$. Thus, it suffices to set $V_k = P^\dagger \cdot D_k \cdot P$, where

$$D_k = \begin{pmatrix} e^{id_1/2^k} & 0 \\ 0 & e^{id_2/2^k} \end{pmatrix},$$

for each $k \in \{0, 1, 2, \ldots\}$. Note that then

$$\begin{aligned} \|V_k - I\|_2 &= \|P^\dagger \cdot D_k \cdot P - I\|_2 \\ &= \|P^\dagger \cdot (D_k - I) \cdot P\|_2 \\ &\leq \|P^\dagger\|_2 \cdot \|D_k - I\|_2 \cdot \|P\|_2 \\ &= \|D_k - I\|_2 \\ &\leq \pi/2^k. \end{aligned}$$

Therefore, if the recursion is terminated after $k = \lceil \log_2(\frac{\pi}{\varepsilon}) \rceil$ steps then the discrepancy between what the resulting network computes and $\wedge_{n-1}(U)$ is an $(n-k)$-bit transformation of the form $\wedge_{n-k-1}(V_k)$. Since $\| \wedge_{n-k-1}(V_k) - \wedge_{n-k-1}(I)\|_2 = \|V_k - I\|_2 \leq \pi/2^{\lceil \log_2(\frac{\pi}{\varepsilon}) \rceil} \leq \varepsilon$, the network approximates $\wedge_{n-1}(U)$ within ε.□

7.4 Linear Simulation in Special Cases

Lemma 7.9: *For any SU(2) matrix W, a $\wedge_{n-1}(W)$ gate can be simulated by a network of the form*

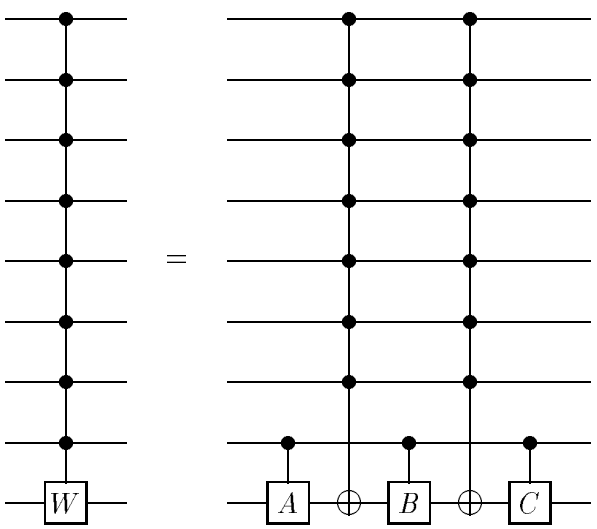

where A, B, and $C \in SU(2)$.

Proof: The proof is very similar to that of Lemma 5.1, referring to Lemma 4.3.□

Combining Lemma 7.9 with Corollary 7.4, we obtain the following.

Corollary 7.10: *For any $W \in SU(2)$, a $\wedge_{n-2}(W)$ gate can be simulated by $\Theta(n)$ basic operations.*

As in Section 5, a noteworthy example is when

$$W = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} = \mathrm{Scal}(\tfrac{\pi}{2}) \cdot \sigma_y.$$

In this case, we obtain a linear simulation of a transformation congruent modulo phase shifts to the n-bit Toffoli gate $\wedge_{n-1}(\sigma_x)$.

7.5 Linear Simulation of General $\wedge_{n-2}(U)$ Gates on n-Bit Networks With One Bit Fixed

Lemma 7.11: *For any unitary U, a $\wedge_{n-2}(U)$ gate can be simulated by an n-bit network of the form*

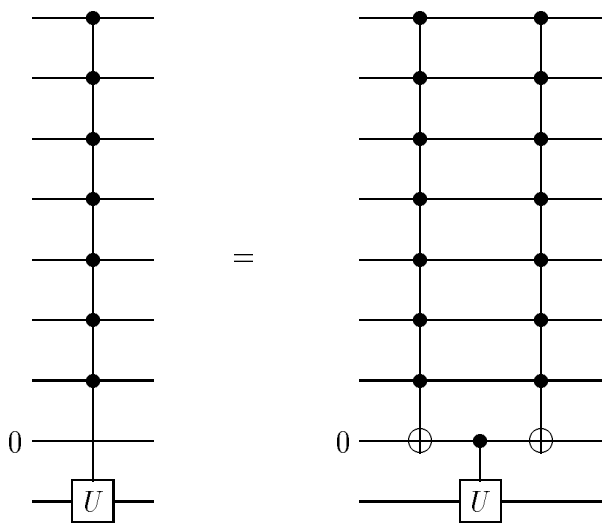

(illustrated for $n = 9$), where the initial value of one bit (the second to last) is fixed at 0 (and it incurs no net change).

Proof: By inspection.□

Combining Lemma 7.11 with Corollary 7.4, we obtain the following.

Corollary 7.12: *For any unitary U, a $\wedge_{n-2}(U)$ gate can be simulated by $\Theta(n)$ basic operations in n-bit network, where the initial value of one bit is fixed and incurs no net change.*

Note that the "extra" bit above may be reused in the course of several simulations of $\wedge_m(U)$ gates.

8 Efficient general gate constructions

In this final discussion we will change the ground rules slightly by considering the "basic operation" to be *any* two-bit operation. This may or may not be a physically reasonable choice in various particular implementations of quantum computing, but for the moment this should be considered as just a mathematical convenience which will permit us to address somewhat more general questions than the ones considered above. When the arbitrary two-bit gate is taken as the basic operation, then as we have seen, 5 operations suffice to produce the Toffoli gate (recall Lemma 6.1), 3 produce the Toffoli gate modulo phases (we permit a merging of the operations in the constuction of Sec. 6.2), and 13 can be used to produce the 4-bit Toffoli gate (see Lemma 7.1). In no case do we have a proof that this is the most economical method for producing each of these functions; however, for most of these examples we have compelling evidence from numerical study that these are in fact minimal[45].

In the course of doing these numerical investigations we discovered a number of interesting additional facts about two-bit gate constructions. It is natural to ask, how many two-bit gates are required to perform *any arbitrary* three-bit unitary operation, if the two-bit gates were permitted to implement any member of U(4)? The answer is six, as in the gate arrangement shown here.

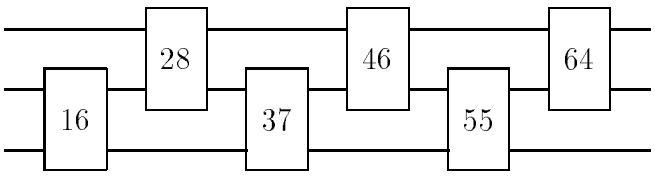

We find an interesting regularity in how the U(8) operation is built up by this sequence of gates, which is summarized by the "dimensionalities" shown in the diagram. The first U(4) operation has $4^2 = 16$ free angle parameters; this is the dimensionality of the space accessible with a single gate, as indicated. With the second gate, this dimensionality increases only by 12, to 28. It does not double to 32, for two reasons. First, there is a single global phase shared by the two gates. Second, there is a set of operations acting only on the bit shared by the two gates, which accounts for the additional reduction of 3. Formally, this is summarized by noting that 12 is the dimension of the *coset space* SU(4)/SU(2). The action of the third gate increases the dimensionality by another $9 = 16 - 1 - 3 - 3$. 9 is the dimension of the coset space SU(4)/SU(2)×SU(2). The further subtraction by 3 results from the duplication of one-bit operations on both bits of the added gate. At this point the dimensionality increases by nine for each succeeding gate, until the dimensionality reaches exactly 64, the dimension of U(8), at the sixth gate. In preliminary tests on four-bit operations, we found that the same rules for the increase of dimensionality applied. This permits us to make a conjecture, just based on dimension counting, of a lower bound on the number of two-bit gates required to produce an arbitrary n-bit unitary transformation: $\Omega(n) = \frac{1}{9}4^n - \frac{1}{3}n - \frac{1}{9}$. It is clear that "almost all" unitary transformations will be computationally uninteresting, since they will require exponentially many operations

to implement.

Finally, we mention that by combining the quantum gate constuctions introduced here with the decomposition formulas for unitary matrices as used by Reck *al.*[15], an explicit, exact simulation of *any* unitary operator on n bits can be constructed using a finite number ($\Theta(n^3 4^n)$) of two-bit gates, and using no work bits. In outline, the procedure is as follows: Reck *et al.*[15] note that a formula exists for the decomposition of any unitary matrix into matrices only involving a U(2) operation acting in the space of pairs of *states* (not *bits*):

$$U = (\prod_{x1,x2 \in \{0,1\}^m,\ x1<x2} T(x1,x2)) \cdot D.$$

$T(x1,x2)$ performs a U(2) rotation involving the two basis states $x1$ and $x2$, and leaves all other states unchanged; D is a diagonal matrix involving only phase factors, and thus can also be thought of as a product of 2^{n-1} matrices which perform rotations in two-dimensional subspaces. Using the methods introduced above, each $T(x1,x2)$ can be simulated in polynomial time, as follows: write a grey code connecting $x1$ and $x2$; for example, if $n = 8$, $x1 = 00111010$, and $x2 = 00100111$:

1	00111010	*x1*
2	00111011	
3	00111111	
4	00110111	
5	00100111	*x2*

Operations involving adjacent steps in this grey code require a simple modification of the $\wedge_{n-1}$ gates introduced earlier. The $(n-1)$ control bits which remain unchanged are not all 1 as in our earlier constructions, but they can be made so temporarily by the appropriate use of "NOT" gates ($\wedge_0(\sigma_x)$) before and after the application of the $\wedge_{n-1}$ operation. Now, the desired $T(x1,x2)$ operation is constructed as follows: first, permute states down through the grey code, performing the permutations (1,2), (2,3), (3,4), ... (m-2,m-1). These numbers refer to the grey code elements as in the table above, where m, the number of elements in the grey code, is 5 in the example. Each of these permutations is accomplished by a modified $\wedge_{n-1}(\sigma_x)$. Second, the desired U(2) rotation is performed by applying a modified $\wedge_{n-1}(U)$ involving the states $(m-1)$ and (m). Third, the permutations are undone in reverse order: (m-2,m-1), (m-3,m-2), ... (2,3), (1,2).

The number of basic operations to perform all these steps may be easily estimated. Each $T(x1,x2)$ involves $2m-3$ (modified) $\wedge_{n-1}$ gates, each of which can be done in $\Theta(n^2)$ operations. Since m, the number of elements in the grey code sequence, cannot exceed $n+1$, the number of operations to simulate $T(x1,x2)$ is $\Theta(n^3)$. There are $O(4^n)$ T's in the product above, so the total number of basic operations to simulate any $U(2^n)$ matrix exactly is $\Theta(n^3 4^n)$. (The number of steps to simulate the D matrix is smaller and does not affect the count.) So, we see that this strict upper bound differs only by a polynomial factor (which likely can be made better than n^3) from the expected lower bound quoted earlier, so this Reck procedure is relatively "efficient" (if something which scales exponentially may be termed so). A serious problem with this

procedure is that it is extremely unlikely, so far as we can tell, to provide a polynomial-time simulation of those special $U(2^n)$ which permit it, which of course are exactly the ones which are of most interst in quantum computation. It still remains to find a truly efficient and useful design methodology for quantum gate construction.

Acknowledgments

We are very grateful to D. Coppersmith, D. Deutsch, A. Ekert and T. Toffoli for helpful discussions. We are also most happy to thank the Institute for Scientific Interchange in Torino, and its director Professor Mario Rasetti, for making possible the most recent workshop on quantum computation at which much of this work was performed. A. B. thanks the financial support of the Berrow's fund in Lincoln College (Oxford).

References

[1] P. A. M. Dirac, *The Principles of Quantum Mechanics*, (Oxford, 1958), Chap. 5; for a modern treatment, see A. Peres, *Quantum Theory: Concepts and Methods*, (Kluwer, 1993), Chap. 8.6.

[2] R. Landauer, "Irreversibility and heat generation in the computing process", IBM J. Res. Dev. **5**, 183 (1961); C. H. Bennett, "Logical reversibility of computation", IBM Journal of Research and Development, **17**, 525 (1973); Yves Lecerf, "Machines de Turing reversibles. Recursive insolubilite en $n\varepsilon N$ de l'equation $u = \theta^n$ ou θ est un "isomorphism de codes". *Comptes Rendus* **257**, 2597-2600 (1963); C. H. Bennett, "Time/space trade-offs for reversible computation", SIAM J. Comput. **18**, 766 (1989). For a review, see C. H. Bennett and R. Landauer, "Physical limits of computation", Scientific American, July, 1985, p. 48.

[3] E. Fredkin and T. Toffoli, "Conservative Logic", Internat. J. Theoret. Phys. **21**, 219 (1982).

[4] T. Toffoli "Reversible Computing", in *Automata, Languages and Programming*, eds. J. W. de Bakker and J. van Leeuwen (Springer, New York, 1980), p. 632; Technical Memo MIT/LCS/TM-151, MIT Lab. for Comp. Sci. (unpublished).

[5] P. Benioff, "Quantum mechanical Hamiltonian models of Turing machines", J. Stat. Phys. **29**, 515 (1982).

[6] A. Peres, "Reversible logic and quantum computers", Phys. Rev. A **32**, 3266 (1985).

[7] R. P. Feynman, "Quantum mechanical computers", Optics News, February 1985, **11**, p. 11.

[8] N. Margolus, "Parallel Quantum Computation", in Complexity, Entropy, and the Physics of Information, Santa Fe Institute studies in the Sciences of Complexity, vol. VIII, ed. W. H. Zurek, (Addison-Wesley, 1990), p. 273.

[9] Tycho Sleator and Harald Weinfurter, "Realizable universal quantum logic gates", preprint (1994).

[10] A. Barenco, D. Deutsch, A. Ekert, and R. Jozsa, "Logic gates for quantum circuits", preprint (1994).

[11] J. I. Cirac and P. Zoller, "Quantum computations with cold trapped ions", preprint (November 1994).

[12] I. Chuang and Y. Yamamoto, "A simple quantum computer", submitted to Phys. Rev. A (November 1994).

[13] Y. Yamamoto, M. Kitegawa, and K. Igeta, in *Proc. 3rd Asia-Pacific Phys. Conf.* (World Scientific, Singapore, 1988); G. J. Milburn, Phys. Rev. Lett. **62**, 2124 (1989).

[14] S. Lloyd, "A potentially realizable quantum computer", Science, **261**, 1569 (1993); "Envisioning a quantum supercomputer", Science, **263**, 695 (1994).

[15] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, "Experimental realization of any discrete unitary operator", Phys. Rev. Lett. **73**, 58 (1994).

[16] D. Deutsch, "Quantum theory, the Church-Turing principle and the universal quantum computer", Proc. R. Soc. Lond. A **400**, 97 (1985).

[17] Deutsch, D. and Jozsa, R., "Rapid solution of problems by quantum computation", *Proceedings of the Royal Society*, London, vol. A439, 1992, pp. 553−558.

[18] Berthiaume, A. and Brassard, G., "Oracle quantum computing", *Proceedings of the Workshop on Physics and Computation — PhysComp '92*, October 1992, IEEE Press, pp. 195−199.

[19] Simon, D., "On the power of quantum computation", Proceedings of the 35th Annual Symposium on the Foundations of Computer Science (IEEE Computer Society Press, Los Alamitos, CA, 1994), p. 116.

[20] Bernstein, E. and Vazirani, U., "Quantum complexity theory", *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, (ACM Press, New York, 1993), pp. 11−20.

[21] P. W. Shor, "Algorithms for quantum computation: discrete log and factoring", Proceedings of the 35th Annual Symposium on the Foundations of Computer Science (IEEE Computer Society Press, Los Alamitos, CA, 1994), p. 124.

[22] For a review see A. Ekert and R. Jozsa, "Shor's quantum algorithm for factorising numbers", in preparation for Reviews of Modern Physics (1995); see also J. Brown, New Scientist **133**, No. 1944, p. 21 (1994).

[23] D. Deutsch, "Quantum computational networks", Proc. Roy. Soc. Lond. A **425**, 73 (1989).

[24] The angle-halving of this definition conforms to the usual relation between operations in SO(3) and SU(2).

[25] B. Schumacher, "On Quantum Coding," *Phys. Rev. A* (in press to appear 1995).

[26] R. Jozsa and B. Schumacher, "A New Proof of the Quantum Noiseless Coding Theorem,"*J. Modern Optics* **41**, 2343-2349, (1994).

[27] A. C.-C. Yao, "Quantum Circuit Complexity", Proceedings of the 34th Annual Symposium on the Foundations of Computer Science (IEEE Computer Society Press, Los Alamitos, CA, 1993), p. 352.

[28] D. P. DiVincenzo, "Two-bit gates are universal for quantum computation", Phys. Rev. A **50**, 1015 (1995).

[29] A. Barenco, "A universal two-bit gate for quantum computation", preprint (1994).

[30] D. Coppersmith and E. Grossman, "Generators for Certain Alternating Groups with Applications to Cryptography", SIAM J. Applied Math. **29**, No. 4, pp. 624-627 (1975).

[31] Richard Cleve, "Reversible Programs and Simple Product Ciphers" in *Methodologies for Designing Block Ciphers and Cryptographic Protocols*, PhD Thesis, University of Toronto, pp. 2-56 (1989).

[32] Michael Ben-Or and Richard Cleve, "Computing algebraic formulas using a constant number of registers", SIAM J. Comput. **21**, 54 (1992).

[33] S. Lloyd, "Almost any quantum logic gate is universal", preprint (1994).

[34] D. Deutsch, A. Barenco, and A. Ekert, "Universality in quantum computation", submitted to Proc. R. Soc. Lond. (1995).

[35] H. F. Chau amd F. Wilczek, "Realization of the Fredkin gate using a series of one- and two-body operators", Report IASSNS-HEP-95/15, (1995) (quant-ph/9503005).

[36] A. Berthiaume, D. Deutsch, and R. Jozsa, "The stabilisation of quantum computations", Proceedings of the Workshop on Physics and Computation, PhysComp '94 (Los Alamitos: IEEE Comp. Soc. Press, 1994), p. 60.

[37] W. G. Unruh, "Maintaining coherence in quantum computers", Phys. Rev. A **51**, 992 (1995).

[38] I. L. Chuang, R. Laflamme, P. Shor, and W. H. Zurek, "Quantum Computers, Factoring and Decoherence", Report LA-UR-95-241 (1995) (quant-ph/9503007).

[39] R. Landauer, "Is quantum mechanics useful?", Proc. Roy. Soc. Lond., (to be published).

[40] The angle-halving of this definition conforms to the usual relation between operations in SO(3) and SU(2).

[41] D. Coppersmith, "An approximate Fourier transform useful in quantum factoring", IBM Research Report RC19642 (1994); R. Cleve, "A note on computing Fourier transforms by quantum programs", unpublished note (1994).

[42] The only non-trivial one-bit classical invertible operation is $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$.

[43] Corresponding with the definition in the previous section, the top two wires contain the input bits x_1 and x_2, and the third wire contains the input y).

[44] J. Mathews and R. L. Walker, Mathematical Methods of Physics, (2nd edition, W. A. Benjamin, Menlo Park, CA, 1970), p. 464.

[45] D. P. DiVincenzo and J. Smolin, "Results on two-bit gate design for quantum computers", Proceedings of the Workshop on Physics and Computation, PhysComp '94 (Los Alamitos: IEEE Comp. Soc. Press, 1994), p. 14.