Experimental requirements for Grover's algorithm in optical quantum computation

Jennifer L. Dodd,[1,2,3,*] Timothy C. Ralph,[1,2] and G. J. Milburn[1,2]
[1]*Centre for Quantum Computer Technology, The University of Queensland, Queensland 4072, Australia*
[2]*School of Physical Sciences, The University of Queensland, Queensland 4072, Australia*
[3]*Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125, USA*
(Received 11 June 2003; published 24 October 2003)

The field of linear optical quantum computation (LOQC) will soon need a repertoire of experimental milestones. We make progress in this direction by describing several experiments based on Grover's algorithm. These experiments range from a relatively simple implementation using only a single nonscalable controlled-NOT (CNOT) gate to the most complex, requiring two concatenated scalable CNOT gates, and thus form a useful set of early milestones for LOQC. We also give a complete description of basic LOQC using polarization-encoded qubits, making use of many simplifications to the original scheme of Knill, Laflamme, and Milburn [E. Knill, R. Laflamme, and G. J. Milburn, Nature (London) **409**, 46 (2001)].

DOI: 10.1103/PhysRevA.68.042328 PACS number(s): 03.67.Lx

I. INTRODUCTION

In the next few years, we can expect to see demonstrations of basic quantum gates in several implementations of quantum computation. With this in sight, it is natural to look ahead to what interesting quantum circuits can be built out of a small number of one- and two-qubit gates acting on a few qubits, as these circuits will provide milestones on the way to full-scale quantum computation [1].

Grover's search algorithm [2,3] is a good candidate for such a milestone. It is a quantum algorithm identifying one of N elements, marked by an oracle, with order $\sqrt{N}$ uses of the oracle. When the search space consists of four elements, the algorithm is guaranteed to produce the marked element after one use of the oracle, compared to the 2.25 uses expected in a classical search. We will see that it can be implemented using only seven one-qubit gates and two two-qubit gates, which makes it an excellent target once one- and two-qubit gates have been mastered. Not surprisingly, it was one of the first algorithms to be experimentally implemented in nuclear magnetic-resonance quantum computing (Chuang, Gershenfeld, and Kubinec [4] and Jones, Mosca, and Hansen [5]).

A promising quantum computing technology is the *linear optical quantum computation* (LOQC) scheme of Knill, Laflamme, and Milburn (KLM) [6] (see Gottesman, Kitaev, and Preskill [7] for an alternative approach). In this scheme, one-qubit gates are relatively straightforward. While implementing a scalable universal two-qubit gate such as a CNOT gate remains a challenge, such a gate is likely to be demonstrated in the next couple of years. Already, a non-scalable CNOT gate has been approximately implemented by Pittman *et al.* [8]. For these reasons, it is important to establish some specific LOQC milestones on the path toward building a large quantum computer, in the form of some simple algorithms on a few qubits.

This pursuit is dogged by conceptual difficulties associated with quantum algorithms on a very small number of qubits, summed up in the question: What is the criterion for ''quantumness''? A reasonable criterion, particularly in the context of Grover's algorithm, is to require a ''speedup'' over the best classical algorithm. However, this notion can be hard to make sense of when the number of steps is on the order of ten, rather than hundreds of thousands, and the problem can easily be done by hand (not to mention by a GHz classical processor). Furthermore, sometimes the reduction in the number of steps can be achieved in an implementation whose physical requirements grow exponentially with the number of qubits, trading off time for space. The question of whether or not this counts as ''quantum'' has received much attention (see, for example, Kwiat *et al.* [9], Bhattacharya, van Linden van den Heuvell, and Spreeuw [10]).

Perhaps the best solution to this problem is a pragmatic one. In the quest to build a quantum computer large enough to provide a genuine advantage over classical computers, two things must be achieved. First, a fine level of quantum control must be demonstrated for both single qubits and pairs of qubits. Second, it will be necessary to show that the number of components (qubits and gates) in a circuit can be increased without insurmountable increases in difficulty. In particular, we must avoid exponential increases in the amount of resource usage (either time or space)—the implementation must be *scalable*.[1]

Therefore, the importance of an experimental achievement of an early milestone (such as the four-element Grover's algorithm) should be measured primarily on these criteria. A demonstration that Grover's algorithm finds the marked item in fewer steps than is possible with a classical computer is an important goal, but it is less important than the fine level of quantum control that it implies. At this early stage of development of quantum computers, any such demonstration is a significant achievement, while a demonstration of such control in a scalable manner is likely to be significantly more difficult and consequently more impressive.

*Electronic address: www.physics.uq.edu.au/people/jdodd

[1]Blume-Kohout, Caves, and Deutsch [11] give a general characterization of the requirements for scalability.

This is illustrated by the experiment of Kwiat *et al.* [9], which demonstrated the ability to implement the search algorithm in a quantum optical system, but using an encoding that is not scalable—as they point out, the number of optical elements that they require grows exponentially in the number of qubits in their system. Thus, although their techniques might be successfully extended to a few qubits, they are not practical as the basis for an approach to building a quantum computer.

In contrast, we are explicitly concerned with developing experimental milestones on the path toward full-scale quantum computation in optical systems. We show that Grover's algorithm on four elements provides several experiments that gradually increase in complexity. The simplest version requires little more than a single, coincidence-basis CNOT gate together with a source of entangled photon pairs, while the most complex version requires two scalable CNOT gates and six photons.

Before describing these experiments and their requirements, we give a brief description of Grover's algorithm (Sec. II) and LOQC (Sec. III). Since the original proposal of LOQC, there have been many simplifications and improvements to the scheme. We give a concise description of the basics of LOQC making full use of these simplifications, focusing on a variant of the original scheme that uses polarization-encoded qubits. In Secs. IV and V, we describe and compare several optical circuits, all implementing Grover's algorithm on four elements. In Sec. VI we briefly discuss appropriate figures of merit for Grover's algorithm, and we conclude in Sec. VII.

II. GROVER'S ALGORITHM ON FOUR ELEMENTS

Grover's algorithm [2,3] (see also Nielsen and Chuang [12] for an elementary treatment on which most of this section is based) is a quantum algorithm that can speed up the solution to certain types of oracle-based computations. We will say more about oracles and their implementation after describing Grover's algorithm.

A. Grover's algorithm

Suppose our search space consists of $N \equiv 2^n$ elements, of which one is a solution to a given problem. Grover's algorithm identifies the solution (with high probability) using $n+1$ qubits according to the following algorithm.

(1) Prepare the state $|0\rangle^{\otimes n}|1\rangle$.

(2) Apply $R^{\otimes n+1}$, where $R = 1/\sqrt{2}\left[\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right]$ is the one-qubit Hadamard gate. (We use the symbol R instead of the usual H to avoid confusion with the horizontal polarization state.)

(3) Apply the oracle, which flips the ancilla qubit conditional on the other qubits being in the state corresponding to the solution.

(4) Apply $R^{\otimes n}$.

(5) Apply a phase shift to the data qubits conditional on not being in the state $|0\rangle^{\otimes n}$, described by the unitary operator $2|0\rangle\langle 0|^{\otimes n} - I_n$ where I_n is the identity operation on the data qubits.

(6) Apply $R^{\otimes n}$.

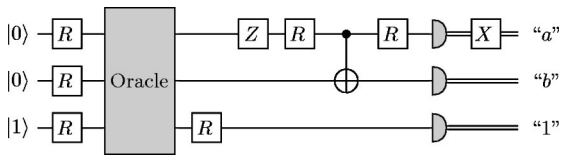

FIG. 1. A circuit diagram for the four-element Grover algorithm, based on the figure in Box 6.1 of Ref. [12]. The top two qubits are the data qubits, initialized in state $|0\rangle|0\rangle$, while the bottom qubit is the ancilla qubit, initialized in state $|1\rangle$. The boxes labeled R and Z represent the one-qubit Hadamard and Pauli σ_Z gates, respectively. The CNOT gate is denoted by the usual symbol, while the gray half circles represent one-qubit measurements in the computational basis, whose output appears on the classical output wires (double lines). The final X gate represents the classical NOT gate required to put the output into the correct form. The measurement always gives "1" on the ancilla qubit, while the data qubits give "a" and "b." It is straightforward to show that, in principle, ab is the state marked by the oracle.

(7) Repeat steps (3)–(6) a specified number of times, then measure the qubits in the computational basis.

The number of repetitions (which is also the number of uses of the oracle) that maximizes the probability of obtaining the correct answer is the nearest integer to

$$\frac{\arccos\sqrt{1/N}}{2\arccos\sqrt{(N-1)/N}} \tag{1}$$

(Boyer *et al.* [13], see also Ref. [12]). This number is bounded above by $\lceil \pi\sqrt{N}/4 \rceil$, hence the claim that Grover's algorithm uses $O(\sqrt{N})$ oracle calls, compared to the $O(N)$ oracle calls required in the classical case.

For the remainder of this paper, we restrict our attention to the case where the number of elements in the search space is $N=4$. In that case, the number of repetitions specified by Eq. (1) is exactly one. A simplified circuit based on the algorithm described above is shown in Fig. 1. It can be verified directly that this circuit, using only one oracle call, gives the correct answer with probability 1, compared to the average of 2.25 oracle calls that must be made with a classical circuit. For example, if the solution is 10, then the output of the circuit is $a=1$ and $b=0$.

B. Implementing the oracle

An oracle is a quantum circuit that *recognizes* solutions to a given problem. For example, suppose we wish to solve a version of the traveling salesman problem, where the goal is to find a route visiting a given collection of cities that is shorter than some specified length L. Although it is in general hard to find such a route, it is easy to recognize whether a proposed route solves the problem: simply add up the total distance the salesman would travel on the proposed route, and compare it to L.

Specifically, an oracle is a circuit that, given an input consisting of a potential solution to a problem, flips the sign of an ancilla qubit if and only if the input is a solution to the problem. Since the only action of the oracle is to recognize solutions, its internal structure is unimportant in a test of the

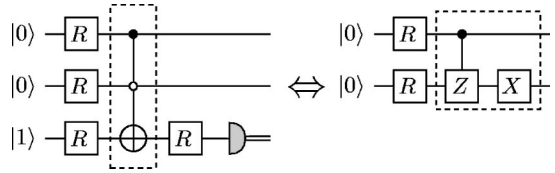

FIG. 2. The circuit on the left shows the beginning of the Grover circuit with an example oracle (inside the dashed box) marking the item 10. We have implemented the oracle using a variant of the Toffoli gate, where the state of the third qubit is flipped when the first two qubits are in the state $|10\rangle$, as indicated by the closed and open circles on the control qubits. We have moved the measurement on the third qubit forward since it plays no further role in the algorithm. In the text, we show that this circuit is in fact equivalent to the simplification on the right, where the Toffoli gate has been replaced by a controlled-Z (CSIGN) operation followed by an X on the appropriate qubit.

algorithm itself. Thus, for our purposes, the choice of oracle is arbitrary, and may be chosen to be as simple as possible.

Although the internal workings of the oracle are unimportant for the purposes of testing the algorithm, the complexity of implementing *some* oracle must be included to characterize the difficulty of performing the experiment. A simple implementation of an oracle marking one of the four states is a Toffoli gate, with the control qubits negated where necessary to specify any of the states 00, 01, 10, or 11 (see the left-hand side of Fig. 2 for the example where the marked state is 10).

If the marked state is 10, the action of the oracle on the three qubits is to take the state $(|00\rangle+|01\rangle+|10\rangle+|11\rangle)(|0\rangle-|1\rangle)$ to

$$(|00\rangle+|01\rangle+|11\rangle)(|0\rangle-|1\rangle)+|10\rangle(|1\rangle-|0\rangle)$$
$$=(|00\rangle+|01\rangle-|10\rangle+|11\rangle)(|0\rangle-|1\rangle) \quad (2)$$

(omitting the normalization). Thus the oracle simply has the effect of flipping the sign of the marked state. The ancilla is not used again, so it can be discarded at this point.

Toffoli gates are difficult to implement in LOQC because there is no known way to implement one without using several CNOT gates. However, for our purposes, a full Toffoli gate is not required because the ancilla qubit plays such a limited role. The two-qubit circuit on the right-hand side of Fig. 2 illustrates this for the case where the marked state is 10. A single controlled-Z (CSIGN) gate that flips the sign of the $|11\rangle$ state, followed by X gates to move the minus sign to the appropriate state, has the same action as the original oracle.

A simplified circuit to implement the four-element Grover algorithm is given in Fig. 3. This is the circuit that we will work with for the remainder of this paper.

III. LOQC WITH POLARIZATION ENCODING

In LOQC, qubits are encoded in *dual rail logic* [6]: Two modes A and B are used, and logical $|0\rangle$ and $|1\rangle$ are encoded as $|1\rangle_A|0\rangle_B$ and $|0\rangle_A|1\rangle_B$, respectively. The modes may represent two different *spatial* modes, or two different *polarization* modes of a single spatial mode [30].

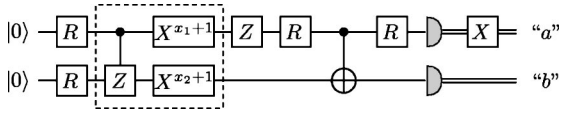

FIG. 3. Inserting the simplified oracle of Fig. 2 into the circuit of Fig. 1 gives this circuit. Note that the marked state is specified inside the oracle (the dashed box) by the values of x_1 and x_2 used to determine whether or not the X gates are applied. (Note that addition in the exponent of the X gates is modulo 2.) Under ideal circumstances, the output of the circuit is $a=x_1$ and $b=x_2$.

In practice, it is likely that polarization-encoded qubits will be used, so that logical $|0\rangle$ and $|1\rangle$ are encoded as $|H\rangle$ and $|V\rangle$, respectively, where H and V refer to horizontal and vertical polarization one-photon states of the same spatial mode. The main reasons for this are (1) it significantly simplifies the implementation of the CNOT gate (see below), (2) it allows one-qubit gates to be implemented using only wave plates and phase delays rather than beam splitters and interferometers, and (3) it reduces the effects of noise by ensuring that, unlike with spatial encoding, both states follow the same path on the quantum wires between gates. In this section we describe in some detail the construction of one-qubit gates and CNOT gates in polarization-encoded LOQC.

A. One-qubit gates

To our knowledge, no complete description of how to implement basic quantum gates with polarization encoding has been given in the literature, so we provide one here. For one-qubit gates, wave plates and phase delays are sufficient. A wave plate with slow axis $|H'\rangle$ and fast axis $|V'\rangle$ has action

$$|H'\rangle\rightarrow e^{i\phi}|H'\rangle,$$
$$|V'\rangle\rightarrow|V'\rangle, \quad (3)$$

where ϕ is the resulting relative phase difference. Special cases in common use are the half- and quarter-wave plates, with ϕ equal to half and a quarter of a wavelength, respectively. Now suppose $|H'\rangle$ is rotated counterclockwise (with respect to the direction of travel of the light) by an angle α from $|H\rangle$. If the input state is $\left[\begin{smallmatrix}h\\v\end{smallmatrix}\right]\equiv h|H\rangle+v|V\rangle$, then the output is given by

$$\begin{bmatrix} e^{i\phi}\cos^2\alpha+\sin^2\alpha & (e^{i\phi}-1)\cos\alpha\sin\alpha \\ (e^{i\phi}-1)\cos\alpha\sin\alpha & e^{i\phi}\sin^2\alpha+\cos^2\alpha \end{bmatrix}\begin{bmatrix} h \\ v \end{bmatrix}. \quad (4)$$

Special cases of this transformation for common one-qubit gates are set out in Table I. The Hadamard and $\pi/8$ gates, labeled R and T in the table, are a universal set for one-qubit quantum computation (Boykin *et al.* [14]), and so any one-qubit gate can be obtained by a sequence of wave plates, although it is convenient to allow phase delays as well. In the Grover circuit, the only one-qubit gates used are the R, X, and Z gates, and thus we only require half-wave plates.

TABLE I. Various one-qubit gates and their implementation in polarization-encoded LOQC. α and ϕ refer to the angle of the slow axis to the horizontal and the relative phase added to light parallel to the slow axis, respectively. Note that T requires a wave plate with a relative delay of one-eighth of a wavelength.

Gate	Optical element
$e^{i\theta}I=e^{i\theta}\begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix}$	Phase delay of $-\theta$
$R=\frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}$	Wave plate with $\phi=180°,\ \alpha=-67.5°$
$T=\begin{bmatrix}1 & 0\\ 0 & e^{i\pi/4}\end{bmatrix}$	Wave plate with $\phi=45°,\ \alpha=90°$
$X=\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix}$	Wave plate with $\phi=180°,\ \alpha=-45°$
$Z=\begin{bmatrix}1 & 0\\ 0 & -1\end{bmatrix}$	Wave plate with $\phi=180°,\ \alpha=90°$
$Y=\begin{bmatrix}0 & -i\\ i & 0\end{bmatrix}$	Two wave plates and a phase delay $[Y=(e^{i\pi/4}I)XZ]$

B. Two-qubit gates

Since the publication of the original LOQC scheme of KLM [6], many simplifications of their CSIGN gate have been developed, with varying tradeoffs between simplicity and functionality. The different types may be divided into two classes, those that are scalable and those that are not. In this section, we describe both types. (Note that the CNOT and CSIGN gates are related by conjugation by Hadamard gates on the target bit, i.e., $\text{CNOT}=(I\otimes R)\text{CSIGN}(I\otimes R)$. Therefore, in the context of LOQC where one-qubit gates are relatively straightforward, these two gates are practically equivalent, and we will use the two almost interchangeably.)

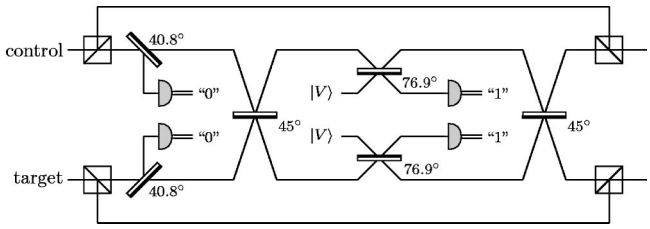

FIG. 4. The simplified KLM CSIGN gate of Ref. [15]. The top rail contains the control qubit and the bottom rail contains the target qubit, both encoded in the polarization of a single photon. A square with a diagonal line across it represents a polarizing beam splitter. By convention, we always assume that the horizontal polarization is 100% reflected while the vertical polarization is 100% transmitted. So, for example, after the first polarizing beam splitters, the topmost rail contains the horizontally polarized component of the control qubit. A thin rectangle represents an ordinary beam splitter, with a sign change for the mode reflected from the thick black side and reflectivity given by the cosine of the angle written next to it. (If the input modes to a beam splitter are $|a\rangle_{\text{in}}$ and $|b\rangle_{\text{in}}$, with the b mode receiving the sign change and with reflectivity given by $\cos x$, then the outputs are $\cos x|a\rangle_{\text{out}}+\sin x|b\rangle_{\text{out}}$ and $\sin x|a\rangle_{\text{out}}-\cos x|b\rangle_{\text{out}}$.) The circuit uses two vertically polarized ancilla photons. It succeeds if the first two measurements both count 0 photons and the second two measurements both count 1 photon.

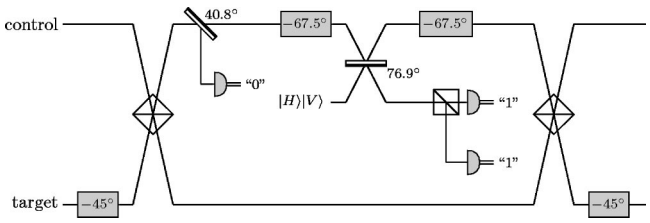

FIG. 5. A further simplified (but still scalable) polarization-encoded KLM CSIGN gate. A gray rectangle containing "$x°$" represents a half-wave plate with slow axis at an angle of $x°$ to the horizontal polarization. See Table I for the corresponding one-qubit gates. This circuit works similarly to the previous one (Fig. 4), but it takes fuller advantage of the orthogonality of the polarization states. For a full description, see the text.

1. Scalable two-qubit gates

The KLM scheme [6] has two properties that at first appear contradictory: the LOQC CSIGN gate is nondeterministic, but it is used to do computations in a scalable manner. The nondeterministic nature of the KLM CSIGN gate is essential to engineer a two-photon interaction without using highly nonlinear materials, but it poses a problem: if its success probability is $\epsilon<1$, then the success probability of a circuit with n CSIGN gates is ϵ^n, i.e., it decreases exponentially with n. A solution to this problem is the technique of *gate teleportation* described by Nielsen and Chuang [16] and Gottesman and Chuang [17]. This technique allows the gates to be prepared as an offline resource, and then "teleported in" whenever required for a computation. KLM showed that the teleportation step can be made near-deterministic using a sufficiently large number of repetitions. This technique is unlikely to be used in early experiments, however, because the extra difficulty involved in teleporting gates will more than cancel out the advantages of increasing the success probability when the number of CSIGN gates is small.

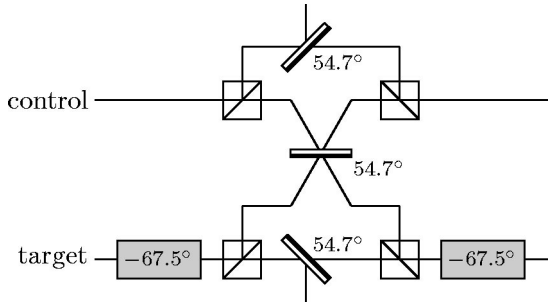

FIG. 6. The coincidence-basis CNOT gate of Refs. [21,22]. All three beam splitters have the same reflectivity $1/3\approx\cos 54.7°$. It can be turned into a CSIGN gate by removing the two half-wave plates. Note that it is not necessary to have detectors on the reflected modes of the topmost and bottommost beam splitters (even though measuring a photon in either of these modes would signal a failure), since other failures of this gate are undetectable until the end of the computation. The gate has worked if exactly one photon is found in each rail.

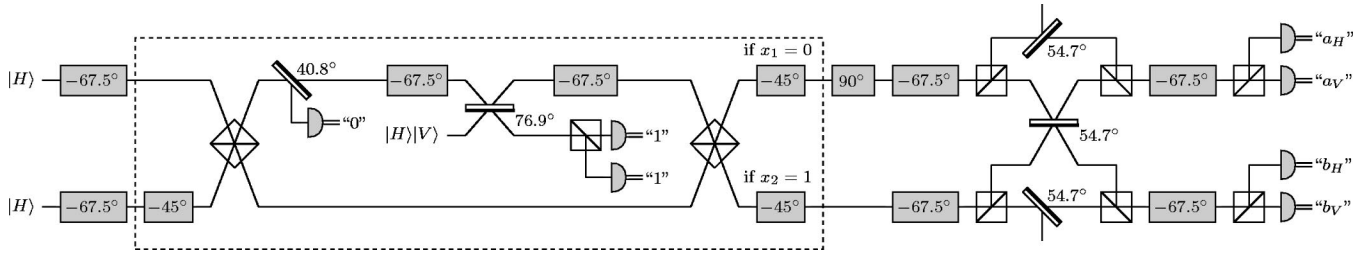

FIG. 7. An optical implementation of Grover's algorithm on four elements based on the circuit in Fig. 3. The oracle part is contained in the dashed box. This circuit is essentially the concatenation of the circuits for the scalable CSIGN gate (Fig. 5) and the coincidence-basis CNOT gate (Fig. 6), together with a few extra wave plates. The output of the circuit is discarded unless the first measurement counts 0 photons, the second two measurements both count 1 photon, and one photon is found in each pair of detectors at the end, i.e., $a_H+a_V=b_H+b_V=1$. Note that we have omitted the final correcting NOT gate on the classical output in this diagram, but it should still be done. For example, if the oracle marks state 10, then the algorithm has successfully identified the marked state if measurements return $a_H=1$, $a_V=0$, $b_H=1$, $b_V=0$.

An essential feature required to make this work is that it must be possible to determine when the gate has succeeded. The KLM CNOT gate has this property—although it only succeeds once in 16 attempts, whether or not it has succeeded is determined by the outcomes of measurements on ancilla photons. We use the term *scalable* to describe a CSIGN (or CNOT) gate that has the property that it is known when it succeeds.

In this paper, we will not work directly with the KLM CSIGN gate since there are simpler alternatives, such as the closely related simplification proposed by Ralph *et al.* [15] and the substantial modification proposed by Knill [18]. There is also a promising alternative approach using entangled ancillas discovered by Pittman, Jacobs, and Franson [19] that we will not consider further here. We focus on the CSIGN gate of Ralph *et al.*, shown in Fig. 4.[2]

In fact, there is a further, substantial simplification to this circuit that is achieved by making fuller use of the polarization encoding, resulting in the circuit in Fig. 5. This gate still requires two ancilla photons. However, it uses fewer detectors, beam splitters, and polarizing beam splitters, and eliminates two interferometers. Its effect on qubit states is unchanged, up to an unimportant overall phase of -1. If we denote the beam splitter reflectivities as $\eta_1 \equiv 5-3\sqrt{2}$ and $\eta_2 \equiv (3-\sqrt{2})/7$ (which are approximated as $\cos 40.8°$ and $\cos 76.9°$ in the diagram), then the action of the gate is the following:

$$|00\rangle \rightarrow \sqrt{\eta_1 \eta_2}(2\eta_2 - 1)|00\rangle = -\sqrt{p}|00\rangle,$$

$$|01\rangle \rightarrow \eta_1(3\eta_2^2 - 2\eta_2)|01\rangle = -\sqrt{p}|01\rangle,$$

$$|10\rangle \rightarrow \sqrt{\eta_1 \eta_2}(2\eta_2 - 1)|00\rangle = -\sqrt{p}|10\rangle, \quad (5)$$

$$|11\rangle \rightarrow \eta_2|11\rangle = \sqrt{p}|11\rangle,$$

where the success probability p is given by $p \equiv \eta_2^2 = (11 - 6\sqrt{2})/49 \approx 0.05$. Thus the gate works approximately once out of every 20 attempts. For the remainder of this paper, we will refer to this gate simply as a ''scalable CSIGN gate.''

2. Coincidence-basis two-qubit gates

An even simpler, but nonscalable CNOT gate was discovered by Hofmann and Takeuchi [21] and Ralph *et al.* [22]. It succeeds once in 9 attempts, but it only works in the *coincidence basis*, i.e., when the results of the whole computation are selected to contain an allowed distribution of photons among detectors. We call this a ''coincidence-basis CNOT gate.'' See Fig. 6. This circuit has been designed so that if exactly one photon is measured in the top rail (in either polarization) and one in the bottom rail, it has worked with certainty. Otherwise, the result is discarded and the experiment is repeated. It cannot, in general, be followed by further two-qubit gates, as it is possible for a later gate to mask a failure. Thus it cannot be used to do scalable quantum computation.

The useful purpose served by this gate (as well as the coincidence-basis gate of Ref. [8]) is as a simpler intermediate step before the full complexity of a scalable CNOT gate. In a general circuit, it may be possible to replace one or more scalable CNOT gate with a coincidence-basis CNOT gate, thereby significantly reducing the complexity of circuits containing a few CNOT gates. In the following sections on constructing optical circuits to perform the four-element Grover algorithm, we will see some of these ideas in action.

IV. THE TWO-QUBIT GROVER IN LOQC

A simplified circuit for the four-element Grover algorithm was given in Fig. 3. In Fig. 7, this circuit is translated directly into an optical circuit, using the prescriptions and circuits of the preceding section.

The circuit, which succeeds once in approximately $180 = 20 \times 9$ (the product of the number of attempts per success

[2]Recent numerical work by Lund, Bell, and Ralph [20] shows that the simplified KLM CSIGN gate of Ref. [15] is more resilient to detector and ancilla inefficiencies than the other two, perhaps because it acts symmetrically on the two qubits. For example, the fidelity of this gate (calculated as the fidelity of the actual output with the ideal output, minimized over input states) is larger than the fidelities of the other two gates for detector efficiencies up to approximately 95%. However, it remains to be seen what effects other sources of error, such as mode-matching errors, and imperfect beam splitter reflectivities, will have on the relative merits of each gate.

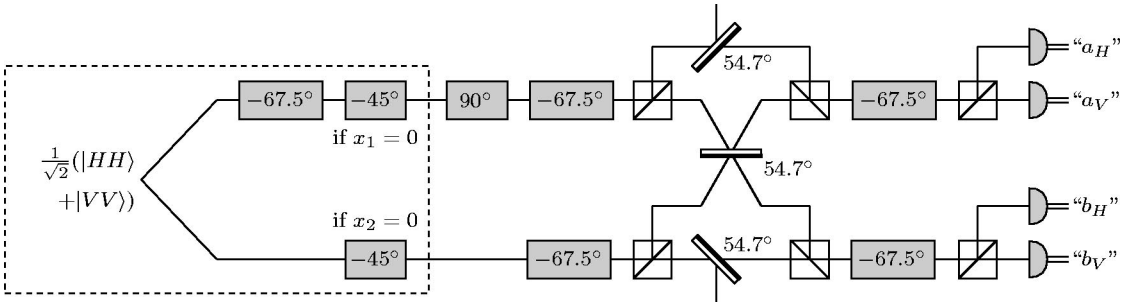

FIG. 8. Grover's algorithm using a parametric down-conversion input. This circuit works similarly to the previous one, but the oracle is no longer demarcated from the initial part of the circuit. The dashed box in this figure now contains both the oracle and the initialization to the state $|HH\rangle+|VV\rangle$. The advantage of this circuit is that it makes use of a natural source of optical entanglement (parametric down-conversion) to replace the very difficult scalable CSIGN gate. The outputs from this circuit are accepted under the same conditions as the previous circuit ($a_H+a_V=b_H+b_V=1$), and the final classical NOT gate has again been omitted.

for each CNOT gate) attempts, uses 10–12 half-wave plates,[3] five beam splitters (two of which must be mode matched), nine polarizing beam splitters (four of which must be mode matched), four photons that must be simultaneously produced in desired polarization states, and seven single-photon detectors. The second CNOT gate can be done in the coincidence basis since there are no interactions between the two qubits following it. Therefore, if the final measurement contains an allowed distribution of photons (exactly one in the top two detectors and one in the bottom two detectors), we know that the second CNOT gate worked, which is sufficient for our purposes here.[4]

However, it is important to note that the output of this circuit (before the measurement) could not be used to do further calculations because of the uncertainty in the outcome of the second CNOT gate. If, for example, there were two photons in the top rail after the second CNOT, the system's state would no longer be in the ''qubit space.'' A third CNOT gate might bring the system back into the qubit space, but it is unlikely to have performed the transformation we expected. In this case, the overall circuit fails, but we have no way of detecting the failure (except to compare with the answer that we can calculate by hand for this simple case).

To ensure reliability for further calculations, the second CNOT gate should be replaced by a scalable CNOT gate. The optical circuit for this case would work once in 400 attempts, and would contain of the order of 14–16 wave plates, eight polarizing beam splitters (six of which would be mode matched), four ordinary beamsplitters (two of which would be mode matched), six photons produced in desired polarization states simultaneously, and ten single-photon detectors. This would be considerably more difficult to achieve experimentally. Since we are (in principle) guaranteed to be in the qubit space at the end of this circuit, the output of each pair of detectors should contain exactly one photon. Therefore, it is possible to simplify the final detection process by simply blocking out one of the polarizations (horizontal, say), and then looking to see if a photon is detected. This would reduce the number of polarizing beam splitters to six and the number of detectors to eight, at the cost of introducing two polarization filters. However, in practice the number of photons at the output will sometimes be incorrect. Thus, the increase in simplicity would have to be weighed against the failures that would go undetected.

V. SIMPLIFICATIONS

By far the most difficult aspect of the experiments just described is implementing the scalable CSIGN gate. However, the CSIGN gate in the oracle is only used in a very restricted way, and it turns out that we can replace it with a much simpler circuit. Since only one input state is ever used, namely, $(|H\rangle+|V\rangle)(|H\rangle+|V\rangle)$, only one state is ever output from the CSIGN gate, namely, $|HH\rangle+|VH\rangle+|HV\rangle-|VV\rangle$. (We will continue to neglect normalization constants.) If a source of entangled input states were available, then the CSIGN gate could be replaced. In optics, such a source is in fact readily available: a parametric down-conversion source can be used to produce the state $|HH\rangle+|VV\rangle$, which can be converted into our desired state by a Hadamard gate on the first qubit, $|H\rangle\rightarrow|H\rangle+|V\rangle$, $|V\rangle\rightarrow|H\rangle-|V\rangle$. Using this fact, a much simplified version of Grover's algorithm is presented in Fig. 8.

The simplicity of this circuit compared with the previous one is emphasized by comparing the number of components. This circuit works once in every nine attempts, and requires 6–8 wave plates, six polarizing beam splitters (of which two must be mode matched), three ordinary beam splitters (one of which must be mode matched), two photons which are

[3]Note that the 90° and 67.5° half-wave plates cannot be combined into a single wave plate: their product $1/\sqrt{2}\left[\begin{smallmatrix}1 & -1\\ 1 & 1\end{smallmatrix}\right]$ has terms of opposite sign in the off-diagonal terms, while the wave plate equation [Eq. (4)] has these entries equal.

[4]A small but potentially useful simplification is to remove the 40.8° beam splitter, as described in Ref. [20]. They show that, until detector and source efficiencies of up to approximately 99.5% are reached, the fidelity of the gate can be substantially increased by removing this beam splitter and adjusting the reflectivity of the 76.9° beam splitter. Given that beam splitter reflectivities are imperfect, removing this beamsplitter is likely to decrease that source of error, while also decreasing the complexity of the circuit by removing a detector. There is a catch, however: the probability of success decreases by a factor of 4–5 for efficiencies of 80–95 %.

produced as the output of a parametric down-conversion source, and four single-photon detectors.

What have we traded for this enormous gain in simplicity? It turns out that we have compromised the versatility of the algorithm. Most significantly, the oracle is no longer easily replaceable. In principle, the oracle should be a ''plug-in'' component able to have many different forms corresponding to different potential problems. In this simplified scheme, however, we have obscured the line between the oracle and nonoracle parts of the circuit, making it difficult to see how to make the circuit solve a problem using a different oracle. In Fig. 8, a dashed box outlines the ''oracle'' part of the circuit for comparison with the previous diagrams, but there is in fact no clear line dividing the oracle from the earlier part of the circuit.

This change affects how the circuit could be used. One example is demonstrating the variation in the success probability of Grover's algorithm as a function of the number of repetitions of steps (3)–(6) described in Sec. II. In the circuit in Fig. 7, the oracle can be reused with some small changes.[5] On the other hand, in Fig. 8, this is not possible—the oracle can only be used once.[6]

VI. FIGURES OF MERIT

An important question that has so far not been addressed is what the appropriate figures of merit are for this experiment. There are two related but distinct notions of success here. The first is to what extent the actual goal of Grover's algorithm has been achieved, i.e., how successfully the experiment distinguishes between the four different oracles. The second is how similar the actual operation of circuit is to the ideal operation. This second notion is important for using these experiments as tests of the ability to combine the basic elements of quantum computation. It is clearly related to the first—if the experiment cannot reliably distinguish between the oracles, then the actual behavior of the circuit must be very far from the ideal operation.

Note that, since the two-qubit gates in these circuits fail the majority of the time, the *average* performance of the circuits will be very far from ideal. However, ultimately the probability of success of two-qubit gates in LOQC will be boosted arbitrarily close to one using gate teleportation, as discussed in Sec. III B 1, and so we restrict our attention to the performance of the circuits when the two-qubit gates succeed.

In order to be able to compare experiments (and also to optimize the performance of a particular experimental setup), we need to be more precise about how to measure the success of these experiments. We suggest calculating figures of merit reflecting each of the two notions of success described above. The first is to simply measure the distinguishability of the distribution of measurement results output by the circuit for different oracles. For example, suppose that for the oracle marking the state 00, the results 00, 01, 10, and 11 occur with probabilities $p_{00} \equiv \{0.9,0.04,0.02,0.04\}$, while the corresponding results when the oracle marks state 10 are $p_{10} \equiv \{0.01,0.08,0.8,0.11\}$. A simple indicator of the distinguishability of these two distributions is their fidelity

$$F(p_{00},p_{10}) \equiv \sum_x \sqrt{p_{00}(x)p_{10}(x)}, \tag{6}$$

where x ranges over the measurement outcomes $00,\ldots,11$ and $p_{ab}(x)$ is the probability of obtaining result x given that the oracle marked state ab. This quantity has the property that it is 1 precisely when the two distributions are identical and 0 precisely when the two distributions are nonoverlapping, that is, when the set of results for which the first distribution is nonzero has no elements in common with the set of results for which the second distribution is nonzero.

In the context of Grover's algorithm, it is desirable to make the fidelity between the distributions arising from each pair of oracles as small as possible. (For an introduction to the fidelity, see, for example, Refs. [12,24]. The relationship of the fidelity to distinguishability is explored by Wootters [25] and in Ref. [24].)

The second figure of merit is related to the similarity of the actual operation implemented ($\mathcal{E}$) to the desired unitary U. U is obtained by simply multiplying together the circuit elements in Fig. 3. $\mathcal{E}$, on the other hand, must be determined experimentally. Ideally, $\mathcal{E}$ should be determined precisely using a method such as quantum process tomography (Chuang and Nielsen [26] and Poyatos, Cirac, and Zoller [27]). Although process tomography can be done using only product-state inputs and one-qubit measurements, it requires an enormous number of runs of the experiment since the output states resulting from 16 different input density matrices must be determined via quantum state tomography.

A less stringent, but much more easily calculated, criterion is that the probability distributions for each oracle should be close to the ideal distributions. Thus, it is desirable to have the fidelity of the actual distribution to the ideal distribution for each oracle as close to 1 as possible.[7] This

[5]The oracle on the right-hand side of Fig. 2 is designed to work with inputs that are equal superpositions of computational basis states. If the oracle is used twice in the same circuit, then it is unlikely that the input state will always be the same. In order to make the oracle work for an arbitrary input state, it is necessary to simply duplicate the X gates following the CSIGN gate, before the CSIGN gate. For the example in Fig. 2, where the oracle marks the state $|01\rangle$, the oracle should consist of the following: an X gate acting on the bottom qubit, followed by the CSIGN gate, followed by the X acting on the bottom qubit.

[6]For a more speculative example, Grover's algorithm can be used to obtain upper bounds on an entanglement monotone called the *Groverian entanglement*, as described by Biham, Nielsen, and Osborne [23]. The basic idea is that if an n-qubit state ρ (possibly mixed) is used as input rather than $|0\rangle^{\otimes n}$, the square root of 1 minus the success probability gives a good measure of the entanglement of ρ. This application requires input states with varying degrees of entanglement, and thus is not possible in the simplified circuit.

[7]Knill *et al.* [28] have a useful discussion of these issues where they advocate the *entanglement fidelity* to measure the quality of an

approach certainly does not completely characterize the behavior of the circuit. For example, it does not determine whether the circuit behaves correctly for inputs other than $|H\rangle|H\rangle$. It is an open question to determine whether there exist methods characterizing how well a circuit implements a desired operation, which are simpler than full process tomography.

VII. A HIERARCHY OF EXPERIMENTS

This collection of different implementations of the same algorithm could be used as the basis for a series of experiments, each building on the last, each more complicated than the last, each demonstrating improved quantum control. For example, once a basic coincidence-basis CNOT gate is working, it would be relatively simple to add a small number of wave plates and a source of entangled photons to do the circuit in Fig. 8. Once a scalable CNOT gate is achieved, these two different CNOT gate circuits could be combined to do the more complicated implementation of Grover's algorithm in Fig. 7, demonstrating the ability to combine a scalable CNOT with further nontrivial quantum computations. Finally, in the more distant future, the implementation using two scalable CNOT gates would make a good testing ground for techniques for combining LOQC components.

experimental implementation of the five-qubit code. They describe a simple way of measuring the entanglement fidelity that could be easily generalized to the setting of Grover's algorithm.

ACKNOWLEDGMENTS

J.L.D. thanks Jeremy O'Brien, Geoff Pryde, and Andrew White for providing insight into the world of experiments, Alexei Gilchrist and Geoff Pryde for a careful reading of the manuscript, Andrew Doherty for help in understanding wave plates, Michael Nielsen for helpful discussions on Grover's algorithm, and Paul Cochrane and Alexei Gilchrist for help using their program for drawing quantum circuits [29], which produced all of the diagrams in this paper. J.L.D. and G.J.M. thank the Institute for Quantum Information for their hospitality. This work was supported in part by the National Science Foundation under Grant No. EIA–0086038, and by the Australian Research Council and ARDA.

[1] See URL http://qist.lanl.gov
[2] L.K. Grover, Phys. Rev. Lett. **79**, 325 (1997).
[3] L. Grover, in *Proceedings of the 28th Annual ACM Symposium on the Theory of Computation* (ACM Press, New York, 1996), pp. 212–219.
[4] I.L. Chuang, N. Gershenfeld, and M. Kubinec, Phys. Rev. Lett. **80**, 3408 (1998).
[5] J.A. Jones, M. Mosca, and R.H. Hansen, Nature (London) **393**, 344 (1998).
[6] E. Knill, R. Laflamme, and G.J. Milburn, Nature (London) **409**, 46 (2001).
[7] D. Gottesman, A. Kitaev, and J. Preskill, Phys. Rev. A **64**, 012310 (2001).
[8] T.B. Pittman, M.J. Fitch, B.C. Jacobs, and J.D. Franson, e-print quant-ph/0303095.
[9] P.G. Kwiat, J.R. Mitchell, P.D.D. Schwindt, and A.G. White, J. Mod. Opt. **47**, 257 (2000).
[10] N. Bhattacharya, H.B. van Linden van den Heuvell, and R.J.C. Spreeuw, Phys. Rev. Lett. **88**, 137901 (2002).
[11] R. Blume-Kohout, C.M. Caves, and I.H. Deutsch, Found. Phys. **32**, 1641 (2002).
[12] M.A. Nielsen and I.L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
[13] M. Boyer, G. Brassard, P. Høyer, and A. Tapp, Fortschr. Phys. **46**, 493 (1998).
[14] P.O. Boykin, T. Mor, M. Pulver, V. Roychowdhury, and F. Vatan, Inf. Process. Lett. **75**, 101 (2000).
[15] T.C. Ralph, A.G. White, W.J. Munro, and G.J. Milburn, Phys. Rev. A **65**, 012314 (2002).
[16] M.A. Nielsen and I.L. Chuang, Phys. Rev. Lett. **79**, 321 (1997).
[17] D. Gottesman and I.L. Chuang, Nature (London) **402**, 390 (1999).
[18] E. Knill, Phys. Rev. A **66**, 052306 (2002).
[19] T.B. Pittman, B.C. Jacobs, and J.D. Franson, Phys. Rev. A **64**, 062311 (2001).
[20] A. P. Lund, T. B. Bell, and T. C. Ralph, Phys. Rev. A **68**, 022313 (2003).
[21] H.F. Hofmann and S. Takeuchi, Phys. Rev. A **66**, 024308 (2002).
[22] T.C. Ralph, N.K. Langford, T.B. Bell, and A.G. White, Phys. Rev. A **65**, 062324 (2002).
[23] O. Biham, M.A. Nielsen, and T.J. Osborne, Phys. Rev. A **65**, 062312 (2002).
[24] C.A. Fuchs, Ph.D. thesis, The University of New Mexico, Albuquerque, NM, 1996 (unpublished).
[25] W.K. Wootters, Phys. Rev. D **23**, 357 (1981).
[26] I.L. Chuang and M.A. Nielsen, J. Mod. Opt. **44**, 2455 (1997).
[27] J.F. Poyatos, J.I. Cirac, and P. Zoller, Phys. Rev. Lett. **78**, 390 (1997).
[28] E. Knill, R. Laflamme, R. Martinez, and C. Negrevergne, Phys. Rev. Lett. **86**, 5811 (2001).
[29] See http://pyscript.sourceforge.net
[30] A good introduction to LOQC in spatial encoding is provided by a set of lectures by Knill, available online at http://online.itp.ucsb.edu/online/qinfo01/