Approximate Quantum Counting on an NMR Ensemble Quantum Computer

J. A. Jones*

Centre for Quantum Computation, Clarendon Laboratory, Parks Road, Oxford OX1 3PU, United Kingdom
and Oxford Centre for Molecular Sciences, New Chemistry Laboratory, South Parks Road, Oxford OX1 3QT, United Kingdom

M. Mosca

Centre for Quantum Computation and Mathematical Institute, 24-29 St Giles', Oxford OX1 3LB, United Kingdom
(Received 26 August 1998)

We demonstrate the implementation of a quantum algorithm for estimating the number of matching items in a search operation using a two qubit nuclear magnetic resonance quantum computer.

PACS numbers: 03.67.Lx, 33.25.+k, 76.60.–k

Quantum computers [1,2] offer the tantalizing prospect of solving computational problems which are intractable for classical computers. A variety of algorithms have been developed, most notably Shor's algorithm for factorizing composite numbers in polynomial time [3,4], and Grover's quantum search algorithm [5,6]. Until recently these algorithms were only of theoretical interest, as it proved extremely difficult to build a quantum computer. In the last few years, however, there has been substantial progress [7–9] in the construction of small quantum computers based on nuclear magnetic resonance (NMR) studies [10] of the nuclei of small molecules in solution. NMR quantum computers have been used to implement a variety of simple quantum algorithms, including Deutsch's algorithm [11,12] and Grover's algorithm [13–15].

NMR quantum computers differ from other implementations in one important way: there is not one single quantum computer, but rather a statistical ensemble of them. For this reason NMR quantum computers should be described using density matrices rather than the more usual ket notation. In some cases this ensemble nature is irrelevant: it is possible to prepare the system with an initial density matrix indistinguishable from that of a pure eigenstate (a pseudopure eigenstate), and as long as the result is another pseudopure eigenstate the behavior of an ensemble quantum computer is identical to that of a conventional quantum computer. Some algorithms, however, produce a superposition of states (relative to the natural NMR computational basis) as their final result, and in such cases the behavior of an ensemble quantum computer will be quite different.

An important example is Grover's algorithm when there is more than one matching item to be found [16]. Suppose a search is made over N items among which there are k matching items. After $O(\sqrt{N/k})$ evaluations of Grover's search function the quantum search algorithm will produce an equally weighted superposition of the k matching items. With a conventional quantum computer this state allows one of the k matching items to be determined at random, as a measurement will result in one of the states contributing to the superposition. With an ensemble quantum computer, however, different members of the ensemble result in different states, and the final signal will be an average over the k matching values. In general it will be difficult or impossible to deduce anything about individual matching items from this ensemble average, and so NMR quantum computers will not be capable of carrying out conventional Grover searches when more than one item matches the search criteria.

An alternative approach to searching is to determine whether any matching items are found in some desired portion of the search space. A binary search will then permit the first matching item, for example, to be located in approximately $\log_2(N)$ attempts. This is a sensible strategy if some efficient algorithm for counting matches can be found. Fortunately this can be achieved by a simple modification of Grover's quantum search, approximate quantum counting [16–18].

Suppose we have a function $f(x)$ which maps n-bit binary strings to a single output bit, so that $f(x) = 0$ or 1. In general there will be $N = 2^n$ possible input values, with k values for which $f(x) = 1$. Grover's quantum search [5,6,16] allows one of these k items to be found, while quantum counting [16–18] allows the value of k to be estimated. The counting algorithm can be considered as a method for estimating an eigenvalue of the Grover iterate $G = HU_0H^{-1}U_{\overline{f}}$, which forms the basis of the searching algorithm [the operator H corresponds to the n-bit Hadamard transform, U_0 maps $|000\cdots0\rangle$ to $-|000\cdots0\rangle$ and leaves the remaining basis states alone, and $U_{\overline{f}}$ maps $|x\rangle$ to $(-1)^{f(x)+1}|x\rangle$].

Starting from the state $|000\cdots0\rangle\langle000\cdots0|$ apply the Hadamard operator H to obtain an equally weighted superposition of all basis states. For $0 < k < N$

$$H|000\cdots0\rangle = (|\Psi_+\rangle + |\Psi_-\rangle)/\sqrt{2}, \qquad (1)$$

where $|\Psi_+\rangle$ and $|\Psi_-\rangle$ are two of the eigenvectors of G with eigenvalues $e^{\pm i\phi_k}$, where $\sin(\phi_k/2) = \sqrt{k/N}$. These eigenvalues show why the probability of success in a Grover search is a periodic function of the number of iterations with period $1/\phi_k \approx 2\sqrt{N/k}$ [16]. A more detailed description is given in [18,19]. For the two extreme cases,

 0031-9007/99/83(5)/1050(4)$15.00

$k = 0$ and $k = N$, $H|000\cdots0\rangle$ is itself an eigenvector with eigenvalue given by the formulas above.

Eigenvalue estimation is most easily described by considering a register which begins the calculation in an eigenvector of G, say $|\Psi_+\rangle$. An additional *control* qubit is needed which begins in the state $(|0\rangle + |1\rangle)/\sqrt{2}$; this may be obtained from $|0\rangle$ by a Hadamard transform. The operator G is then applied to the target register when the control bit is in state $|1\rangle$, that is, a controlled G. The controlled G produces the result

$$\frac{1}{\sqrt{2}}\left(|0\rangle + e^{i\phi_k}|1\rangle\right)|\Psi_+\rangle, \tag{2}$$

or after r repetitions of the controlled G

$$\frac{1}{\sqrt{2}}\left(|0\rangle + e^{ir\phi_k}|1\rangle\right)|\Psi_+\rangle. \tag{3}$$

A second Hadamard transform on the control qubit gives

$$\left(\frac{1 + e^{ir\phi_k}}{2}|0\rangle + \frac{1 - e^{ir\phi_k}}{2}|1\rangle\right)|\Psi_+\rangle; \tag{4}$$

tracing out the target register and expanding the exponential terms gives for the final state of the control qubit

$$\rho = \frac{1}{2}\begin{pmatrix} 1 + \cos(r\phi_k) & i\sin(r\phi_k) \\ -i\sin(r\phi_k) & 1 - \cos(r\phi_k) \end{pmatrix}. \tag{5}$$

The same result is obtained if we replace $|\Psi_+\rangle$ with $|\Psi_-\rangle$, except that the two off-diagonal elements are negated. Thus the same diagonal elements are also obtained from any superposition or statistical mixture of the two, such as $H|000\cdots0\rangle$ [Eq. (1)]. Starting with $H|000\cdots0\rangle$ in the target register will entangle that register to the control register (except when $k = 0$ or N) and tracing out the target register gives the state

$$\rho = \frac{1}{2}\begin{pmatrix} 1 + \cos(r\phi_k) & 0 \\ 0 & 1 - \cos(r\phi_k) \end{pmatrix} \tag{6}$$

in the control register. This will also be the state when $k = 0$ or N.

A variety of different ensemble measurements can be performed to characterize the final state of the control qubit, but the simplest approach is to measure the expectation value of σ_z. This corresponds to determining the population difference between the $|0\rangle\langle0|$ and $|1\rangle\langle1|$ states and is proportional to $\cos(r\phi_k)$. Note that in this case ensemble quantum computers have an advantage: with a single quantum computer it would be necessary to repeat the calculation several times in order to obtain a statistical estimate of $\cos(r\phi_k)$.

The value ϕ_k can be estimated by varying r (the number of repetitions of the controlled G) in a manner based on the technique of Kitaev [20]. For small k, ϕ_k is $O(\sqrt{k/N})$, so when r is on the order of $\sqrt{N/k}$ we will observe a significant σ_z. To count exactly, we need to distinguish ϕ_k from all other ϕ_j, the most difficult cases being ϕ_{k+1} and ϕ_{k-1}. Since $|\phi_k - \phi_{k-1}|$ and $|\phi_{k+1} - \phi_k|$ are $O[1/\sqrt{k(N-k)}]$, then when r is on the order of $\sqrt{k(N-k)}$ we can distinguish ϕ_k from ϕ_{k+1}, ϕ_{k-1}, and the other possibilities. In fact, determining ϕ_k with sufficient accuracy to determine k *requires* roughly $\sqrt{k(N-k)}$ applications of G [21], while a classical algorithm would require roughly N evaluations of f. It is also possible to estimate k to some desired accuracy: to obtain an estimate $\tilde{k}$ with accuracy ϵ, that is,

$$|\tilde{k} - k| \le \epsilon k, \tag{7}$$

uses on the order of $(1/\epsilon)\sqrt{N/k}$ applications of G [17,18,22], while a classical algorithm usually requires about $(1/\epsilon^2)N/k$ evaluations of f [25].

A quantum circuit for implementing this algorithm on a two qubit NMR quantum computer is shown in Fig. 1. This differs from the conventional circuit in three ways. First, pairs of Hadamard gates are replaced by an NMR pseudo-Hadamard gate (a 90°_y rotation) and its inverse [11]. Second, the controlled-Hadamard gates inside the controlled-G propagator have been replaced by uncontrolled gates; this is permitted as the intervening U_0 gate has no effect when the control spin is in state $|0\rangle$. Finally we implement the function f not using a complex network of gates and auxillary qubits, as would be needed for some real function whose properties were unknown, but by simply implementing the corresponding phase shift gates directly. This simplified implementation is necessary, as the small size of current quantum computers does not permit any other approach. The function implemented can, however, be selected by a third party, and so its properties remain unknown to the operator.

This algorithm was implemented using our two-qubit NMR quantum computer [11], which uses two ^{1}H nuclei in a solution of cytosine in D_2O; pseudopure states were generated using the approach of Cory *et al.* [7,8]. All NMR experiments were carried out on a homebuilt spectrometer

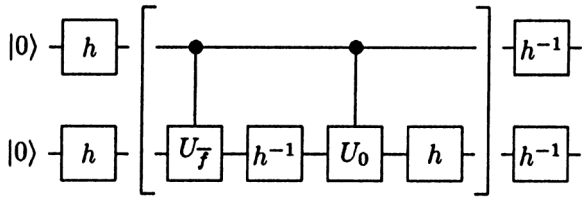

FIG. 1. A quantum circuit for implementing quantum counting on a two qubit NMR quantum computer; the central sequence of gates, surrounded by brackets, is applied r times. The upper line corresponds to the control bit, while the lower line corresponds to the target bit. A similar circuit can be constructed for a larger search space by replacing the target bit by a register and replacing gates applied to the target by multibit versions. Gates marked h implement the NMR pseudo-Hadamard operation, while those marked h^{-1} implement the inverse operation. Controlled gates are marked by a circle and a vertical "control line."

at the Oxford Centre for Molecular Sciences, with a ^{1}H operating frequency of 500 MHz. The two spin states of the ^{1}H nuclei act as qubits, and it is necessary to address each spin individually. Previous experiments on this system [11,14] have used soft pulses to achieve selective excitation, and errors in these pulses have resulted in significant distortions in observed spectra. For these experiments a different approach was adopted, using nonselective hard pulses whenever possible.

The ^{1}H transmitter frequency was set in the center of the spectrum, so that the two spins have angular frequencies in the rotating frame of $\pm\omega/2$. The Hamiltonian can then be written in product operator notation [26] as

$$\mathcal{H} = \frac{\omega}{2} I_z - \frac{\omega}{2} S_z + \pi J_{IS} 2 I_z S_z , \qquad (8)$$

where J_{IS} is the spin-spin coupling constant, and weak coupling has been assumed (i.e., $\omega \gg J_{IS}$). Using a combination of nonselective pulses and carefully chosen periods of free evolution under $\mathcal{H}$ it is possible to implement many of the necessary gates without the use of selective pulses. For example, the controlled-$U_{\overline{f_{01}}}$ gate, which implements the function when $f(0) = 0$ and $f(1) = 1$, can be constructed using the pulse sequence shown in Fig. 2.

Some gates, however, cannot be implemented without using selective pulses; for example, the pseudo-Hadamard gates within the controlled G should be applied only to the target spin. Fortunately it is possible to create selective pulses using only hard pulses and delays, and this process is particularly simple when only two spins are involved. For short periods of evolution under $\mathcal{H}$ the small spin-spin coupling term can be neglected, and $\mathcal{H} \approx (\omega/2)(I_z - S_z)$. Thus after a time $\epsilon_{45} = \pi/2\omega$ the two spins will have undergone rotations of $\pm 45°$ about their respective z axes. This $\pm z$ rotation can be converted to a $\pm y$ rotation by sandwiching the τ period between 90°_x and 90°_{-x} pulses (a variant of the more composite z pulse [10]). Combining this with a 45° pulse along the y axis gives an overall 90°_y rotation for the first spin (I), but no nett rotation for the second spin (S).

With minor variations this approach can be used to generate selective pulses along any axis, and which excite either I or S as desired. These selective pulses can then be used to implement the remaining gates: for example, a controlled $U_{\overline{f_{10}}}$ can be implemented using the circuit for controlled $U_{\overline{f_{01}}}$ with a selective 180° pulse applied immediately before and after the other pulses.

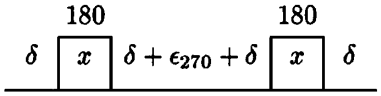

FIG. 2. A pulse sequence implementing a controlled-$U_{\overline{f_{01}}}$ gate using only hard pulses and periods of free precession. Pulse rotation angles (in degrees) are marked above each pulse, while pulse phases are marked within a pulse. Other periods correspond to free precession under the Hamiltonian $\mathcal{H}$ for the time indicated. These times are chosen such that $4\delta + \epsilon_{270} = 1/(2J_{IS})$ and $\epsilon_{270} = 3\pi/\omega$.

The circuit shown in Fig. 1 encodes the result of the calculation in the state of the control qubit, which could be characterized by measuring the expectation value of σ_z for the spin. This cannot be achieved directly, as z magnetization is not a direct NMR observable, but an equivalent measurement can be made by exciting the spin with a 90°_y pulse and observing the NMR spectrum. After phase correction the integrated intensity of the corresponding signal gives the desired result. The phase correction step requires a reference spectrum [11,14], which can be obtained by acquiring a spectrum with $r = 0$.

Prior to acquisition a magnetic field gradient pulse was applied to destroy the homogeneity of the main field. This has the effect of dephasing (thus rendering undetectable) all off-diagonal terms in the final density matrix [14], with the exception of those corresponding to zero quantum coherence [10]. The zero quantum terms can also be removed as they evolve at frequencies of $\pm\omega$ under the Hamiltonian $\mathcal{H}$. This zero quantum filter is easily combined with a standard four-step CYCLOPS phase cycle [10], to reduce instrumental imperfections.

The results of our NMR experiments are shown in Fig. 3. Measurements were made for each of the four possible functions: f_{00} ($k = 0$), f_{01} ($k = 1$), f_{10} ($k = 1$), and f_{11} ($k = 2$). In each case the predicted signal is a cosinusoidal modulation of the signal intensity as a function of r, the number of times the controlled G is applied, where the frequency of the modulation, ϕ_k, depends on $\sqrt{k/N}$. For the two-qubit case, where $N = 2$, the behavior is particularly simple, with modulation frequencies of 0 ($k = 0$), $\pi/2$ ($k = 1$), and π ($k = 2$). In this case it

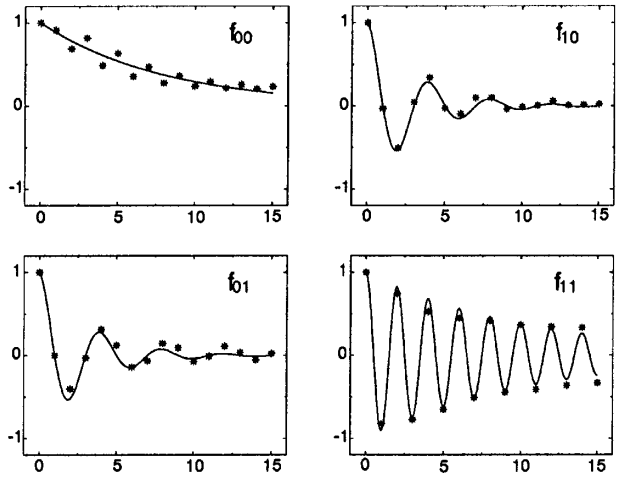

FIG. 3. Experimental results from our NMR quantum computer for each of the four possible functions, f. The observed signal intensity is plotted as a function of r, the number of times the controlled-G operator is applied, and all intensities are normalized relative to the case of $r = 0$. The solid lines are exponentially damped cosinusoids with the theoretically predicted frequencies and are plotted merely to guide the eye.

is possible to determine k using just *one* experiment, with $r = 1$, but spectra were also acquired with larger values of r, both to demonstrate the principle involved and to explore the buildup of errors in the calculation.

The experimental results do indeed show a cosinusoidal modulation as expected, but they deviate from these simple predictions in several ways. First, all the signals show a clear decay in signal intensity as r is increased; this decay is most rapid for f_{01} and f_{10} (where $k = 1$), and least rapid for f_{11} (where $k = 2$). One possible explanation is decoherence, but the observed decay rates are too rapid for this to be the sole explanation: for the case of f_{11} with $n = 15$ the entire pulse sequence lasts around 1 s, compared to a T_2 value of about 3 s, and so a signal loss of only about 30% might be expected. Another likely cause is imperfections in the pulses applied, in particular, those arising from variations in the strength of the resonant rf field across the sample (B_1 inhomogeneity). Both effects are expected to be most severe when $k = 1$, as these cases have complex $U_{\overline{f}}$ gates which take a long time to implement, and least severe when $k = 2$, in which case $U_{\overline{f}}$ is just the identity operation.

In addition to the main exponential decay other deviations from the simple behavior predicted by theory can be seen. These effects are clearest for f_{00} ($k = 0$), where signal intensities are seen to lie alternately above and below the main curve. Such effects could in principle arise from many different causes, but numerical simulations indicate that the major cause is off-resonance effects. These occur because the applied rf field is not perfectly resonant with the NMR transitions, but instead is applied a small distance ($\pm\omega/2$) away. Thus the effect of the field (in the rotating frame) is not simply to cause a rotation around itself, but rather to cause a rotation around a tilted axis [10]. We are currently seeking ways to reduce the size of such effects.

Despite these small errors the results are remarkably good, especially for f_{11}. In this case the experiments have been repeated with much larger values of r, and the cosinusoidal variation remains clearly visible after 60 or more iterations (data not shown). Thus our NMR quantum computer is capable of demonstrating quantum algorithms involving several hundred quantum gates.

We thank S. C. Wimperis and A. Nayak for helpful discussions and C. M. Dobson and A. Ekert for their encouragement. J. A. J. thanks the Royal Society of London for financial support. OCMS is supported by the U.K. EPSRC, BBSRC, and MRC. M. M. thanks CESG (U.K.) for support.

*To whom correspondence should be addressed at the Clarendon Laboratory.

[1] R. P. Feynman, Int. J. Theor. Phys. **21**, 467 (1982).
[2] D. Deutsch, Proc. R. Soc. London A **400**, 97 (1985).
[3] P. W. Shor, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science, Santa Fe, NM, 1994* (IEEE Computer Society Press, Los Alamitos, CA, 1994).
[4] A. Ekert and R. Jozsa, Rev. Mod. Phys. **68**, 733 (1996).
[5] L. K. Grover, in *Proceedings of the 28th Annual ACM Symposium on the Theory of Computing* (ACM Press, Philadelphia, PA, 1996).
[6] L. K. Grover, Phys. Rev. Lett. **79**, 325 (1997).
[7] D. G. Cory, A. F. Fahmy, and T. F. Havel, in *PhysComp'96: Proceedings of the 4th Workshop on Physics and Computation, Boston, 1996* (New England Complex Systems Institute, Cambridge, MA, 1996).
[8] D. G. Cory, A. F. Fahmy, and T. F. Havel, Proc. Natl. Acad. Sci. U.S.A. **94**, 1634 (1997).
[9] N. Gershenfeld and I. L. Chuang, Science **275**, 350 (1997).
[10] R. R. Ernst, G. Bodenhausen, and A. Wokaun, *Principles of Nuclear Magnetic Resonance in One and Two Dimensions* (Clarendon Press, Oxford, 1987).
[11] J. A. Jones and M. Mosca, J. Chem. Phys. **109**, 1648 (1998).
[12] I. L. Chuang, L. M. K. Vandersypen, X. Zhou, D. W. Leung, and S. Lloyd, Nature (London) **393**, 143 (1998).
[13] I. L. Chuang, N. Gershenfeld, and M. Kubinec, Phys. Rev. Lett. **80**, 3408 (1998).
[14] J. A. Jones, M. Mosca, and R. H. Hansen, Nature (London) **393**, 344 (1998).
[15] J. A. Jones, Science **280**, 229 (1998).
[16] M. Boyer, G. Brassard, P. Høyer, and A. Tapp, Fortschr. Phys. **46**, 493 (1998).
[17] G. Brassard, P. Høyer, and A. Tapp, in *Automata, Languages, and Programming: Proceedings of the 25th International Colloquium, ICALP'98, Aalborg, Denmark, 1998* (Springer, Berlin, 1997); see also LANL e-print quant-ph/9805082.
[18] M. Mosca, Theor. Comput. Sci. (to be published).
[19] G. Brassard, P. Høyer, M. Mosca, and A. Tapp (to be published).
[20] A. Y. Kitaev, LANL e-print quant-ph/9511026.
[21] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf, in *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science* (IEEE Computer Society Press, Palo Alto, CA, 1998). In this paper it is shown that that many queries are *necessary* to successfully count with high probability. Sufficiency is shown in [17,18].
[22] Note that exact quantum counting gives a roughly square-root speedup only for small values of k, while approximate quantum counting gives a square-root speedup in all cases. Similar results hold for related problems, as discussed in [23], such as the threshold problem [21] and parity determination [24]. A. Nayak and F. Wu (in LANL e-print quant-ph/9804066) show matching lower bounds for $k \leq N/2$ (up to a constant factor).
[23] L. K. Grover, Science **281**, 792 (1998).
[24] E. Farhi, J. Goldstone, S. Gutmann, and M. Sisper, LANL e-print quant-ph/9802045.
[25] R. Canetti, G. Even, and O. Goldreich, Inf. Process. Lett. **53**, 17–25 (1995).
[26] O. W. Sørensen, G. W. Eich, M. H. Levitt, G. Bodenhausen, and R. R. Ernst, Prog. Nucl. Magn. Reson. Spectrosc. **16**, 163 (1983).