
Quantum Computing

H. Kimble, Chair
C. Callan, Jr.
K. Case
A. Despain
N. Fortson
J. Goodman
S. Koonin
H. Levine
N. Lewis
W. Press
O. Rothaus
P. Weinberger
R. Westervelt

July 1996

JSR-95-115

Approved for Public Release. Distribution Unlimited.

19960827 119

JASON
The MITRE Corporation
1820 Dolley Madison Boulevard
McLean, Virginia 22102-3481
(703) 883-6997

RECEIVED 19960827 119

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information estimated to average 1 hour per response, including the time for review instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE July 18, 1996	3. REPORT TYPE AND DATES COVERED	
4. TITLE AND SUBTITLE QUANTUM COMPUTING			5. FUNDING NUMBERS 07-95-8534-A4	
6. AUTHOR(S) C. Callan, K. Case, A. Despain, N. Fortson, J. Goodman, J. Kimble (Chair), S. Koonin, H. Levine, N. Lewis, W. Press, O. Rothaus, P. Weinberger, R. Westervelt				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The MITRE Corporation JASON Program Office, Z561 1820 Dolley Madison Blvd McLean, Virginia 22102			8. PERFORMING ORGANIZATION REPORT NUMBER JSR-95-115	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) DARPA/TIO 3701 North Fairfax Drive Arlington, Va 22030-1714			10. SPONSORING/MONITORING AGENCY REPORT NUMBER JSR-95-115	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release -Distribution Unlimited			12b. DISTRIBUTION CODE Limiter Statement A	
13. ABSTRACT (Maximum 200 words) An overview and assessment of the rapidly developing field of quantum computing is presented as a result of the 1996 JASON Summer Study. Interest in this field is fueled by the recent discovery by P. Shor of an efficient quantum algorithm for finding the prime factors of large numbers. Because factoring is a task of considerable importance within the domain of cryptography, the physical implementation of Shor's algorithm would have profound impact. In this report, some of the first designs for explicit quantum circuits are presented from which the scaling behavior in terms of space and time can be deduced. From these results, assessments of several physical systems are made together with estimates for the requirements for coherent to dissipative time scales. Beyond the factoring problem, preliminary investigations of new research directions to broaden the purview of quantum computation are presented.				
14. SUBJECT TERMS			15. NUMBER OF PAGES	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT SAR	

Contents

EXECUTIVE SUMMARY	1
1 INTRODUCTION	7
1.1 Study Charge	8
1.2 What is a Quantum Computer?	11
1.3 Quantum Factoring	13
1.4 Quantifying the Essential Character of Quantum Computation	14
2 POSSIBILITIES BEYOND SHOR'S ALGORITHMS	17
2.1 Level Sets; Solutions of Homogeneous Linear Equations	17
2.2 Simulating Quantum Systems — The Hubbard Model	22
2.3 Quantum Computers as Many-body Systems	26
2.4 Simulating quantum physics on a quantum computer	37
2.5 Cold Baths and Optimization	45
3 QUANTUM NETWORKS	49
3.1 Circuits for Quantum Computation	49
3.1.1 Qubit and State Representation	50
3.2 Operations	53
3.2.1 Measurements	53
3.2.2 Clear Operation	53
3.3 Gates	54
3.3.1 Controlled-Not	54
3.3.2 Controlled-Controlled-Not Operator	55
3.3.3 Rotate Operation	56
3.3.4 Conditional-Complex Rotation	58
3.3.5 Measure	59
3.3.6 Set of Operators	60
3.4 Basic Circuits	60
3.4.1 Full Adder	60
3.4.2 Multi-bit Full Adder	61
3.4.3 Multiplexor and Related Circuits	62
3.4.4 Multiplier	65
3.4.5 Subtractor	66
3.4.6 Adder for Recovering Scratch Space	67

3.4.7	Modulo Adder	68
3.4.8	Modulo Multiplier	70
3.5	A Circuit for Shor's Algorithm	71
3.5.1	Exponentiation	72
3.5.2	Modulo Multiplication	73
3.5.3	Calculation of $F(A)$	78
3.5.4	Quantum FFT	79
3.5.5	Conditioning	79
3.5.6	Overall Operations	80
3.6	Complexity	80
3.7	General Scaling for $F(a)$ Circuit	81
3.7.1	Shor's Algorithm	81
4	IMPACT OF DISSIPATION	85
5	ERROR CORRECTION	91
6	ASSESSMENT OF POSSIBLE PHYSICAL REALIZATIONS	95
6.1	Condensed Matter Systems	99
6.1.1	Introduction	99
6.1.2	Solid State Quantum Device Examples	105
6.2	Optical Systems	111
6.3	Trapped Ions	116
6.3.1	Ions in a linear trap	117
6.3.2	Laser cooling and interrogation of trapped ions	121
6.3.3	Requirements for quantum computing	123
7	RECOMMENDATIONS	129

EXECUTIVE SUMMARY

Overview

As part of the 1995 Summer Study, the JASONS undertook an assessment of the field of quantum computing. The current flurry of activity in this field is fueled by the discovery in 1994 by P. Shor of an efficient quantum algorithm for finding the prime factors of large numbers. Because no classical algorithm of comparable efficiency is known, the supposed fundamental difficulty of the factoring problem has ensured the security of the RSA public key cryptosystem, which Shor's algorithm now brings into question. However, beyond the domain of cryptology, the marriage of quantum mechanics and information science represents a *potentially* profound development with *possibly* far reaching implications.

Unfortunately any definite assessment of this field without qualifiers is currently impossible because of the dearth of significant results beyond the brilliant quantum algorithms that Shor presented 1.5 years ago. While proponents of the field believe in a prodigious future, this belief currently rests on a very small set of truly significant results. If quantum computation is to change the future in a major way, then there must emerge a much broader class of applications than are known today. Without a more wide ranging set of possibilities, the impetus for conquering the daunting issues associated with physical implementation is to large measure lost.

Apart from the actual implementation of quantum computation, the potential impact of this new paradigm on classical computation should not be

overlooked. It may well be that new perspectives offered by the investigation of quantum algorithms may well lead to more powerful classical algorithms. Unlike the physical implementation of quantum computation which is at best a long-term endeavor, such insight could more or less immediately advance the state of the art of classical computation.

JASON Activities

To attempt to address these and other issues related to quantum computation, the JASON Summer Study pursued a broad investigation principally along the following avenues.

1. A preliminary exploration of several new research directions to attempt to broaden the purview of quantum computation. Included are the simulation of general quantum systems, function minimization by simulated annealing, and the application of quantum many-body techniques.
2. Designs for the first explicit quantum circuits for Shor's factoring algorithm. From this work emerged estimates of the scaling behavior of certain quantum circuits and of the nontrivial nature of the "scratch-pad" problem. The inevitable tradeoff between intermediate storage and computation length was quantified to some extent.
3. Assessment of several physical systems for possible implementation of quantum computation, including condensed matter, optical, and ion-trap systems. In combination with the explicit JASON quantum circuits, estimates have been made of the requirements for the ratio of

coherent to dissipative time scales.

Recommendations

As a result of these activities over the course of the 1995 Summer Study, our rank-ordered recommendations for possible ARPA support of research in the area of quantum computation are as follows.

1. Establish a research program to investigate possibilities for quantum computation beyond Shor's algorithms. Here we have in mind the fostering of a fairly intense effort over the coming years to understand the types of problems for which quantum computation is well suited and whether or not new insights do indeed arise for developing more powerful classical algorithms. The two principal communities involved would probably be those of theoretical physics and computer science (but not to the exclusion of other groups). Clearly, as new quantum algorithms are developed, it will be essential to address the issue of error correction as well.
2. Seed research in various communities for quantitative minimization of algorithmic complexity and optimum circuit design. Given the extreme value of qubits of information in terms of the degree of difficulty of physical realization, it is quite important to have explicit quantum circuits with quantitative measures of resource requirements (beyond simply "a polynomial of order k ") in order to bridge the gulf between abstract quantum algorithms and actual physical implementations. The JASON quantum circuits provide an important step in this direction, even if

they prove to be less than optimum in the ultimate conservation of qubits and ops.

3. Supplement ongoing experimental research related to the isolation and control of discrete quantum systems suitable for quantum logic. Here the research objectives may not be specifically quantum computation, but may be instead fundamental components thereof such as the investigation of quantum dynamics in nontrivial Hilbert spaces (e.g., the generation of quantum-state entanglement for more than two qubits and the role of dissipation). Given the tremendous disparity between current laboratory capability and the requirements for a nontrivial implementation of Shor's algorithm, we would specifically advise against a program of "prototype development", but would rather supplement research on diverse fronts with modest goals rather disconnected from a "Holy Grail" pursuit of quantum computation.

Overall, we feel that the most pressing need with also the greatest potential, is for a broad theoretical exploration for opportunities beyond Shor's algorithms. If the "well" proves to be "dry", so be it. On the experimental front, we do not believe that there is a similarly pressing need for ARPA involvement. Although there are a variety of promising systems, the most optimistic near-term hope would be only extremely modest "proof-of-principle" demonstrations from which we doubt any profound new insights would emerge. Therefore, we recommend that ARPA should not shoulder the principal burden for funding of experimental efforts, which should not in any case be justified solely for their relevance to quantum computation.

However, ARPA could play an important role in ensuring that the experimental and theoretical communities remain engaged. Should an explosion of possibilities ensue from the theoretical investigations, then it may well be worthwhile to consider increasing the investment on the experimental front, bearing in mind that the time horizon for the experimental realization of quantum computation will still be distant.

Finally, we would urge the adoption of a broad-minded view for opportunities other than those related to large-scale computation, such as might arise in quantum cryptography [38] or coherent nanoscale electronics. This is a long-term endeavor of potentially profound significance where surprises are likely to emerge on diverse fronts.

1 INTRODUCTION

Although the field of quantum computing has a history that dates to the early 1980's, [1] there has recently been an explosion of activity driven by the discovery of P. Shor in 1994 of an efficient quantum algorithm for factoring a large number into its prime factors.[2, 3] Factoring is an important problem because its supposed fundamental difficulty provides the basis for the security of the widely employed RSA public key cryptosystem. Stated more quantitatively, if $L = \text{Log}_2 N$ is the number of binary bits in the number N that is to be factored, then Shor's quantum algorithm requires a number of computational steps G that scales as a polynomial function of L , $G \sim L^k$. By contrast, no such polynomial time algorithm is known for factoring on a classical (conventional) computer, where instead the degree of difficulty of the factoring problem is roughly (sub)exponential in L .

Beyond the example of factoring, quantum computation has the potential for a profound impact on the whole of information sciences. Expressed most succinctly in terms of computational complexity, the outstanding question is "Can quantum computers solve efficiently 'difficult' problems that are otherwise intractable on classical computers?" To answer this currently unresolved question, progress is required on several fronts, including the expansion of the library of quantum algorithms beyond those of Shor and the investigation of the feasibility of the actual physical implementation of a quantum computer.

1.1 Study Charge

Driven by these developments with respect to quantum computing and other recent results in the area of molecular computing with DNA, ARPA tasked a JASON study of the “Boundaries of Computing” as follows:

Advances in computing speed and power have, to date, mostly relied on packing transistors and resistors ever more densely on a silicon substrate. There is a consensus that this line of advance will eventually peter out (on a one- or two-decade time scale). Ideas for a totally different style of computing, dubbed quantum computing, have recently surfaced. It relies on the non-Boolean logic implicit in measurements on pure quantum states and promises such miracles as the prime factorization of numbers in polynomial time. Practical realization of such a scheme would have dramatic implications for cryptography and, no doubt, for other military applications. Most discussions of this subject have been at the purely mathematical level and JASON proposes to do a preliminary study of the practical issues and opportunities which will arise when material realizations of these schemes are attempted. Other issues, such as “biological computing” (ala Edelman) and new paradigms of networking will be discussed as appropriate.

As a result of this charge, two principal activities were undertaken dur-

ing the 1995 Summer Study. The results of the investigation of quantum computation are reported here; a separate JASON report (JSR-95-116) gives our findings relative to DNA computing.

To initiate our study of quantum computing, we held two days of briefings that included the following speakers.

- C. Bennett (IBM)-Quantum Communication and Cryptography
- D. DiVincenzo (IBM)-Physical Systems Overview
- A. Ekert (Oxford)-General Introduction to Quantum Computing
- R. Hughes (LANL)-A Proposed Prototype Experiment with Trapped Ions
- S. Lloyd (MIT)-General Physical Requirements and Quantum Networks
- D. Wineland (NIST)-Experiments with Trapped Ions for Quantum Logic

A. Garg (Northwestern) also participated as an expert (non-speaker) on dissipation. Later in the summer, we were briefed by D. Dubin (UCSD) on various aspects of binding and instabilities for trapped ions and had discussions with M. Roukes (Caltech) on coherence in nanoscale condensed matter systems.

Based upon these briefings and our own ensuing research, our study has addressed three principal topics. The first is an attempt to expand the purview of quantum computing beyond the singular examples of factoring and discrete log presented by Shor. Several promising research directions have been identified which are worthy of future investigation and are described in Section 2 of this report.

The second area of activity concerns the design of quantum circuits. At the outset of our study, there seemed to be only fairly vague statements about the quantitative resources that would be required for the implemen-

tation of Shor's factoring algorithm. There apparently existed no explicit layouts of quantum circuits that gave the number of primitive logic operations and quantum bits needed to implement the algorithm. In response to this situation, explicit JASON designs of appropriate quantum circuits have been developed, as described in Section 3. The importance of this work is that it not only enables us to make definite quantitative statements about quantum "resources", but it as well permits an understanding of the scaling behavior of the network and of the critical need for the efficient management of "scratch" or "garbage" qubits. Furthermore, from our designs we can glean some improved insight into the impact of dissipation (Section 4) and the consequent requirements for candidate physical systems for quantum computation, which is the third principal area of our study (Section 6).

In an assessment of possible physical realizations, we should emphasize at the outset that no system has yet been shown to be sufficient to the task of quantum logic with entangled states even at the level of a single two-bit gate, much less at the scale of a large quantum computation involving thousands of qubits. Nonetheless, in recent years there has been rapid and important progress in the isolation and controlled interaction of discrete degrees of freedom for simple quantum systems. Indeed, this is an area whose scientific growth has been fueled on diverse fronts quite independent of the recent arrival of possible applications to quantum computing. Still there is a daunting (some would say impossible) gap between theoretical requirements to implement Shor's quantum factoring algorithm on a nontrivial scale (e.g., for a 200 digit number) and current or near-term experimental capabilities. In our (incomplete) survey, we consider condensed matter and optical systems,

as well as an explicit scheme involving trapped ions.

1.2 What is a Quantum Computer?

We present here only the briefest sketch of the basic ideas related to quantum computing. An excellent overview of the field can be found in the review article by Ekert and Joza.[2]

To begin with, note that bits of information in a classical computer which take on values $[0,1]$ are replaced by quantum bits (or qubits) in a quantum computer. The bit values $[0,1]$ in the quantum case are associated with the states of some simple quantum system, such as an elementary spin [down, up] or internal atomic states [ground, excited]. A classical register such as $\{1, 0, 1\}$ which stores one of 2^L numbers ($L = 3$ bits here) is replaced with a quantum wavefunction to describe the state of the register, where for example $\{1, 0, 1\}$ represents the state of three particles with (spin up, spin down, spin up), respectively. An extremely important distinction between classical and quantum registers is that a quantum register can operate with superposition states such as

$$\begin{aligned}
 |\psi\rangle = \sum_A C_A |A\rangle &= C_{000}|0\rangle|0\rangle|0\rangle|0\rangle + C_{001}|0\rangle|0\rangle|0\rangle|1\rangle & (1-1) \\
 &+ C_{010}|0\rangle|1\rangle|0\rangle \\
 &+ C_{011}|0\rangle|1\rangle|1\rangle + C_{100}|1\rangle|0\rangle|0\rangle + C_{101}|1\rangle|0\rangle|1\rangle \\
 &+ C_{110}|1\rangle|1\rangle|0\rangle + C_{111}|1\rangle|1\rangle|1\rangle,
 \end{aligned}$$

which represent all 2^L numbers (here $L = 3$ again) and hence in effect store all

possible numbers $[0, 1, 2, \dots, 2^{L-1}]$ at once. Superposition states as in Equation (1-1) are the source of what has been termed “quantum parallelism”.

From a formal perspective, the evaluation of a function F for an input A on a quantum computer involves the unitary transformation

$$|0\rangle|0\rangle \rightarrow |A\rangle|0\rangle \rightarrow |A\rangle|F(A)\rangle, \quad (1-2)$$

with A an integer $[0, 1, 2, \dots, 2^{L-1}]$. Here, the bit values for A are held in a “data” register, while those for F are contained in a “function” register, with each “ \rightarrow ” representing a suitable unitary transformation on either the A or F registers or both. In less formal terms, we will see in Section 3 that the transformation of Equation (1-2) can be accomplished by a network or circuit of suitably arranged quantum gates operating on the bits of A and $F(A)$.

Because the A and F registers are quantum in nature, we can compute F via “quantum parallelism” for all possible A values with the same effort that it took to compute F for only one specific value of A . Instead of preparing a single input A , we extend Equation (1-2) to generate a uniform superposition of all A values (Equation (1-1) with all C_A set to $2^{-L/2}$, which is a state that can be generated efficiently). In this case, Equation (1-2) becomes

$$|0\rangle|0\rangle \rightarrow \sum_A |A\rangle|0\rangle \rightarrow \sum_A |A\rangle|F(A)\rangle, \quad (1-3)$$

where the sum is over all integers $[0, 1, 2, \dots, 2^{L-1}]$. Unfortunately, although the output register now contains all possible values of $F(A)$, it is not possible to “read” this register in the conventional sense of a classical register. A quantum measurement will project only one value for $F(A)$, with an exponen-

tially large number of reported calculations ($\sim 2^L$) required to tabulate F for all A . Hence, there would seem to be no gain associated with the “quantum parallelism” of the calculation.

This situation is quite different if however we wish to extract not the entire table of values for F , but instead some global property of F such as periodicity for a periodic function. In this case we proceed as above, where now the projection $F(A) \rightarrow F_0$ (with F_0 as the recorded value as a result of the measurement on the function register only) leaves the data register in a uniform superposition of states for which $F(A_0)=F_0$ and which exhibit the periodicity of $F(A)$:

$$\left\{ \sum_{A'_0} |A_0\rangle \right\} |F_0\rangle = \{|1\rangle + |6\rangle + |11\rangle + \dots\} |F_0\rangle, \quad (1-4)$$

for the case with $F(1) = F(6) = F(11)=F_0$ (period $r=5$). The (unknown) period r can be extracted efficiently by way of a quantum Fourier transform on the data register A . Hence, by preparing a uniform superposition for the input register A , we can with a “single” calculation of $F(A)$ extract a global property of F that would otherwise have required an extensive tabulation by repeated calculation of $F(A)$ over the $q=2^L$ values of A . We therefore apparently gain an exponential speedup in the calculation.

1.3 Quantum Factoring

Although we will not here go into the details of Shor’s factoring algorithm, suffice it to say that in the simplest possible terms, it is based on

the efficient determination of the period r of the function $F(A) = X^A \text{ Mod } N$, where N is the number to be factored and X is randomly chosen (and is coprime with N). Here A ranges over all integers from 0 to q , with $N^2 \leq q \leq 2N^2$, with $L = \text{Log}_2 q$. Since in practice N is a large number (≥ 100 digits), $F(A)$ must be computed over a gigantic set of numbers ($> 10^{100}$), leading to the difficulty (and hence the supposed security) of the factoring problem on a classical computer. However, because of the exponential speedup associated with “quantum parallelism” as illustrated in the preceding section, Shor’s quantum algorithm is in principle able to find the prime factors of N by finding the period r of $F(A)$ in a number of steps that is polynomial in L . No such polynomial time algorithm is known for factoring on a classical computer.

1.4 Quantifying the Essential Character of Quantum Computation

The JASON investigation attempted to address opportunities and difficulties associated with quantum computation. Apart from any issue associated with physical implementation, if we again consider a register with L qubits, then the “good news” of quantum computation is that L qubits can be used to describe states in a Hilbert space of dimension 2^L . Hence 2^L inputs can be accessed via the “quantum parallelism” of an L -bit register and correspondingly many function evaluations can be carried out in one fell swoop [2^L vs. L].

The “bad news” is multifold and begins with quantum measurement which restricts our ability to access completely all 2^L aspects of our output state with only a polynomial number of measurements L^k . Furthermore, although L qubits describe states in a 2^L dimensional Hilbert space, we are not able to make arbitrary transformations in this space, since a general unitary transformation requires “resources” of order 2^L (in terms of the number of parameters to describe the transformation or the number of gates to implement it). Hence, although “quantum parallelism” gives us access to a gigantic space of possibilities, we still have to live within an exponentially smaller budget of order L^k in our attempts to assess this space [L vs. 2^L]. Stated somewhat differently, because only an infinitesimal number of all possible unitary transformations in a space of dimension 2^L can be generated from L^k parameters [for $L \gg 1$], there is a demand for extreme ingenuity in finding “useful” quantum algorithms.

Finally, there is an important aspect of quantum computation related to the issue of the management of “garbage” qubits. Since a quantum calculation must proceed reversibly, any function evaluation which is not one-to-one (as in the evaluation of $F(A)$ for Shor’s quantum factoring algorithm) must have an accompanying register which contains information that could be used to reverse the calculation. While this situation is well known within the context of classical reversible computation,[4] quantum computation adds a new twist in that the garbage or scratch qubits become entangled with the input and function qubits. The consequence of this entanglement can be understood with reference to the example of the evaluation of a periodic function given in Section 1.2. The requirement for extra garbage qubits means that

the step of function evaluation in Equations (1-2) through (1-3) actually involves a third “scratch” register initially filled with 0’s and finally filled with the appropriate information for each $F(A)$ so as to allow the calculation to be reversed. Hence, the state in Equation (1-4) is not really simply a sum only over values A_0 corresponding to F_0 , but each term contains as well a different appended state specifying the state of the garbage qubits. The net result is that a quantum Fourier transform of the state will not yield the desired period r unless steps are taken to first clear the garbage. To circumvent this difficulty, the procedure to follow is that first suggested by Bennett within the context of classical reversible computation and essentially involves copying the desired result and then reversing the entire computation.[4] The cost of this procedure in the quantum context is high. There is a need for more qubits (which as we will see are a very precious resource) and more operations. State vectors are thereby forced to live longer in a higher dimensional Hilbert space than would be the case in the absence of garbage qubits, which leads to an increased sensitivity to dissipation.

Apart from this “bad news” arising from issues of principle, there are a variety of challenges associated with the physical implementation of quantum computation. We stress again that no system has displayed sufficient capabilities for the realization of even simple quantum logic, much less quantum computation in a large dimensional Hilbert space. Motivated by this observation, we will discuss in Sections 3 through 6 a variety of topics relevant to bridging the gulf between abstract quantum algorithms and actual physical implementations.

2 POSSIBILITIES BEYOND SHOR'S ALGORITHMS

In this section, we review the progress made by various JASONS in attempting to find new quantum algorithms and applications for quantum computing beyond Shor's algorithms for discrete log and factoring. Our objective has been to gain further insight into the operation of quantum computers. In addition to quantum algorithms per se, it is also a worthwhile endeavor to explore the possibility of "quantum inspired" algorithms for classical computers.

The initial condition for our investigation is a paucity of substantial problems for which the quantum computer has a natural applicability. This situation exists not for want of effort; there has been an intense effort by many groups around the world to rush forward into the same stream where Shor discovered his golden nuggets. However, the absence of significant output from this endeavor is troubling, since if quantum computing is to have a profound impact on the futures of the physical or information sciences, then necessarily the range of topics encompassed must be greatly expanded.

2.1 Level Sets; Solutions of Homogeneous Linear Equations

One of the more problematic aspects of quantum computing is the

paucity of substantial problems for which it has a natural applicability. Indeed, the only problems where quantum computing apparently succeeds and conventional computing fails (at least at this time) are those of factoring a very large number and the discrete log problem. Curiously, these two problems appear to be the drivers behind much of the current enthusiasm for quantum computing, but if the development of this technology is to go forward rationally, other large problem areas to which it is naturally and advantageously suited must be found.

Now the investigators in this area are well aware of this shortcoming, and have tried to find such problems, but so far without success. The principal thrust of research to date appears to be the construction of potentially useful quantum gates and theoretical research on foundational problems.

In order to examine the situation further, we decided to have a hard look at the factoring and discrete log algorithms as discovered by Shor, and try to broaden their purview.

From a fundamental point of view, the algorithms appear to succeed because of modular periodicity, coupled with the power of the Fourier Transform for exposing periodicity, and all carried out at the quantum computing level. The Fourier Transform is natural in this setting because it is a unitary transformation, and so can represent the evolution of the state vector of a physical system, but also because it can be built out of simpler unitary transformations, the parts of the Fast Fourier Transform.

If one could take directly at the quantum level the Fourier Transform of

a function, periodicity would show up immediately. But the direct Fourier Transform is not a natural operation in quantum computing. In Shor's algorithms the Fourier Transform is introduced by a subterfuge, and some further analysis is needed to extract the desired answer.

What we are going to try to do here is extract the underlying principle of some of Shor's procedure, and show incidentally that it is not exactly periodicity which makes it work.

We suppose we have a quantum mechanical system in a uniform superposition of states labelled $(0, 0)$ to $(q - 1, q - 1)$, $q = 2^L - 1$. The states might correspond, for example, to two sequences each of L particles, with spin labelled 0 or 1. We are interested in deriving properties of a function f (f may be vector valued), defined on the states (i, j) , $0 \leq i, j \leq q - 1$, and by some means, quantum mechanical in nature, we have got our quantum computer into the uniform superposition of states $|a, b, f(a, b)\rangle$. Physicists like to write $|a\rangle|b\rangle|f(a, b)\rangle$. We would really like to compute the discrete Fourier Transform of f , and resort to the following alternate device.

Using a sequence of unitary transforms arising from steps of the Fast Fourier Transform, we bring our computer, instantaneously, into the state

$$\frac{1}{q} \sum_{a,b} e^{2\pi i \frac{ac+bd}{q}} |c, d, f(a, b)\rangle .$$

A measurement which yields $|c, d, e_0\rangle$ has probability

$$\left| \frac{1}{q} \sum_{f(a,b)=e_0} e^{2\pi i \frac{ac+bd}{q}} \right|^2$$

where the above sum is over all pairs (a, b) for which $f(a, b) = e_0$. What we are seeing, in fact, is just the magnitude of a component of the Fourier Transform of the characteristic function of $f^{-1}(e_0)$. We never see more than this, just Fourier transforms of characteristic functions of subsets.

The probability of $|c, d, e_0\rangle$ is clearly greatest, c and d variable, when $c = d = 0$, and unless we get coherent sums for some reason, we will never see $|c, d, e_0\rangle$ for other (c, d) .

Moreover, if the particular level set $f^{-1}(e_0)$ is substantially larger than many others, we will never see Fourier information about these others.

So let us take a level set $f^{-1}(e_0)$ which is as large or larger than any other. What does it take for $\frac{1}{q} \sum_{f(a,b)=e_0} e^{2\pi i \frac{ac+bd}{q}}$ to be roughly as large in magnitude as the same for $c = d = 0$. A relatively easy criterion, which is roughly the one that Shor uses, requires that the set $\{ac + bd\}$, as the pair (a, b) runs through $f^{-1}(e_0)$, shall, when taken with least residue mod q , satisfy

$$-\frac{q}{4} < \{ac + bd + K\}_q < \frac{q}{4}$$

for some K . We call the pair (c, d) a multiplier for the level set.

Geometrically, this simply means the least residues lie in a narrowish stripe. The phases of the corresponding summands now all lie on a half circle, and they will add semi-coherently. The expected Fourier component has a magnitude $\frac{4}{\pi^2}$ times the $(0, 0)$ component, on the assumption of uniform distribution in the half circle.

The condition described above is sufficient, but by no means necessary. It is the criterion exploited by Shor in a setting which carried substantially more structure. Particularly, all his level sets were of the same cardinality, and all had the same set of multipliers. The desired information on periodicity was extracted from opportune multipliers.

For a general function such as we have been describing, the high probability measurements suggest that, and nothing more, a certain level set has “strong” linear regression mod q .

We are not certain how to use this structure algorithmically.

It is interesting, however, to look at the special case where the function being analyzed is (vector valued) linear and thus, for example,

$$f(a, b) = (L_1(a, b), L_2(a, b), L_3(a, b)) .$$

A level set is either empty, or an affine “plane” of some dimension.

Let N be the subspace of solutions (mod q) to the homogeneous equations $L_1 = 0, L_2 = 0, L_3 = 0$.

If N consists of the zero vector alone, the only states with substantial probability after Fourier transforming are $|0, 0, e, f, g\rangle$, where $L_1 = e, L_2 = f, L_3 = g$ has a solution (mod q).

The situation with N non-trivial is similar. The only states with substantial probability are $|c, d, e, f, g\rangle$ with (c, d) perpendicular to N and the equation $L_1 = e, L_2 = f, L_3 = g$ having a solution (mod q).

In more detail, if (c, d) is not perpendicular to N , the state $|c, d, e, f, g\rangle$ has zero probability for any e, f, g . This statement is not entirely trivial, depending as it does on the fundamental theorem of finitely generated abelian groups. For our particular case, it says that the subspace N of solutions to the homogeneous system is freely generated by some appropriate solutions u_1, u_2, \dots, u_p . That is to say, every $n \in N$ is of the form, and uniquely:

$$n = t_1 u_1 + t_2 u_2 + \dots + t_p u_p$$

where each t_i goes from 0 to $2^{h_i} - 1$, and $2^{h_i} u_i = 0 \pmod{q}$, no smaller power of 2 annihilates u_i . With this result in hand, the more detailed statement about (c, d) not perpendicular to N follows readily.

To assist the reader in understanding this result, we remark that N is not a vector subspace in the usual sense, so does not have a basis in the usual sense. Fortunately, it has a “basis” in the sense just described.

Thus by quantum computing we can decide whether a homogeneous linear system has non-trivial solutions. This facility could possibly be incorporated into useful numerical algorithms.

2.2 Simulating Quantum Systems — The Hubbard Model

One possible use for a quantum computer involves the efficient simulation of other quantum mechanical systems. This application is different in spirit than the classical application such as factoring. There, the basic idea is that the quantum coherence accomplishes many calculations in parallel,

with the hard problem being the collection of the information desired from the final state. Here, the desired answer is some expectation value in the original quantum system which maps onto some (“few-body”) operator in the qubit quantum dynamics — this part is in some cases straightforward. The harder part is finding an efficient encoding¹ of the true Hilbert space into the qubit space which allows for simple time evolution steps. Here we discuss how one might do this for one particular case, the Hubbard model.

The Hubbard model is a simple Hamiltonian meant to describe the physics of electrons in solids made out of transition metals such as copper. The idea is that there are two allowed states per site, a specific orbital with a choice of spin $\pm 1/2$. Electrons can hop from one atom to the next with some amplitude w which depends ultimately on the overlap of atomic orbital wave functions; the size of this then determines the width of the electronic band. In addition, electrons on the same site suffer a large Coulomb repulsion U . This model and variants thereof have been studied extensively in recent years in the context of high T_c copper-oxide superconductors.

Because of the interest in this and other quantum systems, much effort has gone into devising simulation strategies. For example, the equivalence between quantum mechanics in imaginary time and classical statistical mechanics allow the use of the Monte Carlo technique for computing a variety of interesting objects. However, simulations which offer information about the real time correlation functions (and these are the ones most relevant to experimental measurements) have to date been ad-hoc. Basically, one could

¹Hard for Shor-type computation as well.

study either very small systems or alternatively employ ad-hoc truncations of the Hilbert space so as to get around the exponential number of coefficients in the real space time evolution. Here, we argue that one can do at least some quantum simulations in polynomial time on a quantum computer.

For the Hubbard model

$$H = -W \sum_{\vec{x}} \sum_{\vec{a}, \sigma} \psi_{\vec{x}, \sigma}^\dagger \psi_{\vec{x}+\vec{a}, \sigma} + U \sum_{\vec{x}} \psi_{x\uparrow}^\dagger \psi_{x\uparrow} \psi_{x\downarrow}^\dagger \psi_{x\downarrow}$$

where $\sigma = \uparrow$ or \downarrow represents a spin, \vec{x} labels points on some lattice and \vec{a} is a lattice vector. For $U = 0$, we get a free electron band with energies $E_K = -w \cos Ka$, $0 \leq K \leq \pi/a$. The basic idea here is to use two qubits to represent one site, with the identification

$$\begin{aligned} |00\rangle &\equiv |0\rangle \\ |10\rangle &\equiv |\uparrow\rangle \\ |01\rangle &\equiv |\downarrow\rangle \\ |11\rangle &\equiv |\downarrow\uparrow\rangle \end{aligned}$$

We then write the overall time evolution operator as the product of a large number of time evolution operators for a small time interval Δt

$$\begin{aligned} T \exp \frac{-i}{\hbar} \int_0^{t_f} H dt' &= \prod_n \exp -\frac{i\Delta t}{\hbar} H \\ &= \prod_n \left[\prod_{\vec{x}} \left(\exp -\frac{i\Delta t}{\hbar} U \psi_{x\uparrow}^\dagger \psi_{x\uparrow} \psi_{x\downarrow}^\dagger \psi_{x\downarrow} \right) \cdot \right. \\ &\quad \left. \prod_{\vec{a}, \sigma} \exp -\frac{i\Delta t}{\hbar} (-W) (\psi_{\vec{x}, \sigma}^\dagger \psi_{\vec{x}+\vec{a}, \sigma} + \text{h.c.}) \right] \end{aligned}$$

up to terms of order $(\Delta t)^2$. Each of the individual unitary transformations acts on a two qubit Hilbert space. The U term, for example, corresponds to

a controlled phase rotation

$$\begin{aligned}
|00\rangle_{\vec{x}} &\rightarrow |00\rangle_{\vec{x}} \\
|01\rangle_{\vec{x}} &\rightarrow |01\rangle_{\vec{x}} \\
|10\rangle_{\vec{x}} &\rightarrow |10\rangle_{\vec{x}} \\
|11\rangle_{\vec{x}} &\rightarrow e^{i\alpha}|11\rangle_{\vec{x}}
\end{aligned}$$

The hopping term couples qubits corresponding to the same spin at different spatial sites. This term gives the unitary transformation (for $\sigma = \uparrow$, say)

$$\begin{aligned}
|0\cdot\rangle_{\vec{x}} |0\cdot\rangle_{\vec{x}+\vec{a}} &\quad \text{unchanged} \\
|1\cdot\rangle_{\vec{x}} |1\cdot\rangle_{\vec{x}+\vec{a}} &\quad \text{unchanged} \\
|0\cdot\rangle_{\vec{x}} |1\cdot\rangle_{\vec{x}+\vec{a}} &\rightarrow i \sin \beta |1\cdot\rangle_{\vec{x}} |0\cdot\rangle_{\vec{x}+\vec{a}} + \cos \beta |0\cdot\rangle_{\vec{x}} |1\cdot\rangle_{\vec{x}+\vec{a}} \\
|1\cdot\rangle_{\vec{x}} |0\cdot\rangle_{\vec{x}+\vec{a}} &\rightarrow i \sin \beta |0\cdot\rangle_{\vec{x}} |1\cdot\rangle_{\vec{x}+\vec{a}} + \cos \beta |1\cdot\rangle_{\vec{x}} |0\cdot\rangle_{\vec{x}+\vec{a}}
\end{aligned}$$

where α and β are simple numbers, and the \cdot means that we don't care about the state of the ion representing the down spin at either site.

Both of these operations are easy to implement. Also, the objects which are of interest for physics purposes are matrix elements of few body Heisenberg operators between states; for example, the density-density correlation function is

$$\langle \bar{\Psi}_0 | \hat{\psi}_{x',\sigma}^\dagger(t) \psi_{x',\sigma}(t) \psi_{x,\sigma}^\dagger(0) \psi_{x,\sigma}(0) | \bar{\Psi}_0 \rangle$$

where the “hat” above the operator denotes the Heisenberg picture. Going over to the Schroedinger picture which we have been using, this becomes

$$\langle \psi_0(t) | \psi_{x',\sigma}^\dagger \psi_{x',\sigma} e^{-iHt} \psi_{x,\sigma}^\dagger \psi_{x,\sigma} | \psi_0(0) \rangle .$$

Since we are mostly interested in poles of this structure thought of as a function of space and time, we can usually take any reasonable choice for the initial state — if we actually need to have the ground state expectation, we would need to use this formulation to adiabatically turn on the U interaction starting from the non-interacting Fermi sea filled up to the Fermi energy. Once we have the initial state, acting on it with the Hermitian operator $\psi^\dagger\psi$ is just a projector onto all states with an electron at site X . Once this is done, the time evolution carries the system to time t , another projection is done and then the inverse time evolution (from t to 0) leaves one with the final step of taking the overlap with the initial wavefunction.

One remaining point to consider is the size of the time step. A reasonably conservative choice might be related to the inverse of the maximum energy in the problem — this is of order $N_{\text{electrons}}U$ (for $U > W$), and hence the time step probably scale as $1/N_{\text{electrons}} \sim 1/N_{\text{sites}}$ for a fixed filling fraction. Since this is a polynomial slow-down in N for a fixed total time of simulation, it does not significantly detract from the efficiency of the simulation.

2.3 Quantum Computers as Many-body Systems

Quantum computers can be thought of as peculiar quantum many-body systems with non-local and time-dependent interactions between the component quantum bits (qubits). The qubits are prepared in some initial state, carried through a prescribed series of unitary transformations (which we refer

to loosely as “time evolution”) and then various observables are measured to give “the answer”. The “program” of a quantum computer then resides in both the sequence of unitary transformations and the particular measurements made in the final state.

Quantum many-body systems have long been a part of physics. Numerous theoretical methods have been developed for describing such systems either approximately or exactly (mean-field or semiclassical, RPA, variational, stochastic simulation) and these correctly answered many physically interesting questions *before* experiments became possible. It is therefore intriguing to ask whether the application of these same methods to quantum systems “that compute” might be similarly successful. Modest success would be defined as some insight into how quantum computers work, their error rates, etc. A more substantial achievement would be application of these methods to Shor’s quantum factoring scheme to produce a “quantum inspired” factoring algorithm that could be implemented on an ordinary (deterministic) computer, thus circumventing the need for explicit physical realization of the quantum computer.

The purpose of this section is to begin the application of classic many-body methods to quantum computers by considering the semiclassical approximation (mean field plus corrections). The approach is to reduce the very high-dimensional quantum dynamics (Hilbert space of dimension 2^L for an L -qubit system) to the evaluation of a high-dimensional (of order L or L^2) integral, which is then approximated by the stationary phase method.

Quantum computers as many-body systems

A quantum computer consists of L input qubits, a mechanism for effecting a unitary transformation upon the high-dimensional Hilbert space that they define, and a mechanism for measuring one or more of the qubits after transformation. Conventionally, the qubits are spin-1/2 (two-state) quantum systems (so that the Hilbert space has dimension 2^L) although it may be convenient later to consider generalizations where each is a spin- j ($2j + 1$ -state) quantum system, in which case the Hilbert space has dimension $(2j + 1)^L$.

Computation is effected by specifying some initial state of the input, $|i\rangle$, applying the unitary transformation to obtain a final state $|f\rangle = U|i\rangle$, and then observing one or more qubits in the final state. The input state is most commonly (but need not be) taken to be of the direct product form,

$$|i\rangle = |i\rangle_1 |i\rangle_2 \dots |i\rangle_L$$

where $|i\rangle_l$ is the state of the l 'th qubit. For example, in Shor's factoring algorithm, the initial state is the coherent sum of all integers,

$$\frac{1}{\sqrt{2^L}} \sum_{a=0}^{2^L-1} |a\rangle = \frac{1}{\sqrt{2}} (|0\rangle_1 + |1\rangle_1) \frac{1}{\sqrt{2}} (|0\rangle_2 + |1\rangle_2) \dots \frac{1}{\sqrt{2}} (|0\rangle_L + |1\rangle_L).$$

The unitary transformation produced by a network of G quantum gates can be written as an ordered product of G one- and two-bit unitary transformations of the forms $\exp(i\alpha A)$ and $\exp(i\alpha AB)$, where A, B are operators affecting bits a and b , respectively. Thus,

$$U = U_G U_{G-1} \dots U_2 U_1$$

where each of the U_i is a one- or two-bit unitary transform. It has been shown [22] that one- and two-bit unitary transformations are universal, in

the sense that all quantum computations can be realized by a sequence of such gates; an efficient quantum algorithm will have G polynomial in L .

In the measurement process, the L output bits are divided into three classes: P bits that are prescribed to values (b_1, \dots, b_P) , M bits that are measured to have values (m_1, \dots, m_M) and T bits that are traced over (ignored), such that $P + M + T = L$. Thus, the probability of measuring outcome (m_1, \dots, m_M) for the M measured bits is

$$P(m) = \frac{\sum_t |\langle pmt|U|i\rangle|^2}{\sum_{mt} |\langle pmt|U|i\rangle|^2},$$

where the sums are over a complete set of states for the traced and/or measured bits.

If we define a "final state" operator F as

$$F = \left(\prod_{p=1}^P |b_p\rangle\langle b_p| \right) O_1 O_2 \dots O_M \equiv PO$$

where P is an hermitian projector for the prescribed bits and the O_m are the observables for the measured bits, we can write the mean value of a measurement as

$$\langle O \rangle = \frac{\langle i|U^+POU|i\rangle}{\langle i|U^+PU|i\rangle}.$$

It is quantities such as this that we will attempt to approximate via semi-classical methods.

Hubbard-Stratonovich representation of quantum networks

The difficulty in constructing or simulating a quantum network arises from the two-bit gates, as one-bit transformations can be represented by

2×2 matrices. The strategy of the mean-field method is to represent the two-bit gates as an infinite sum over a product of two one-bit gates, and then to use only the “most important” terms in that sum. The more general semiclassical method considers also terms “near” the most important ones.

A two-bit quantum gate effects a unitary transformation, U , on two qubits that is generated by a two-bit operator. The form of the gates required is

$$U = e^{-i\alpha AB}$$

where α is a c -number and A, B are commuting hermitian operators referring to the two bits. [More generally, there may be further one-body transformations, which pose no additional problem.]

A Hubbard-Stratonovich representation of U is

$$U = \frac{\alpha}{2\pi} \int_{-\infty}^{+\infty} d\sigma d\tau e^{i\alpha\sigma\tau} e^{-i\alpha\tau A} e^{-i\alpha\sigma B} ,$$

where α and τ are two real fields. Thus, U is represented as an infinite superposition of one-body transformations.

As an example, consider the Controlled Not gate acting on two qubits a (the “control”) and b (the “target”):

$$C = (-i)^a e^{i\pi a J_{xb}} = e^{-i\pi a/2} e^{i\pi a J_{xb}}$$

where $a = (J_{za} + 1/2)$ has eigenvalues 0 and 1. Thus, C leaves the state of b unchanged if $a = 0$, while it inverts b (rotates by π about the x -axis) if $a = 1$. The Hubbard-Stratonovich representation of C is

$$C = -\frac{1}{2} \int d\sigma d\tau e^{-i\pi\sigma\tau} e^{i\pi a(\tau-1/2)} e^{i\pi\sigma J_{xb}} .$$

Similarly, the Controlled Phase operator appearing in the Quantum Fourier Transform, $Q = \exp(i\omega ab)$, which applies the phase ω if both a and b are 1, can be written as

$$Q = -\frac{\omega}{2\pi} \int d\sigma d\tau e^{-i\omega\sigma\tau} e^{i\omega\tau a} e^{i\omega\tau b}.$$

By invoking the HS representation for each quantum gate in a network, we can represent the total unitary transformation involved as an integral over $2G$ fields:

$$U = \int D\sigma e^{i\sum_g \alpha_g \sigma_g \tau_g} U_\sigma$$

where $D\sigma$ is the measure over all auxiliary fields, the sum in the exponent is over all two-bit gates, and U_σ is an ordered product of one-body evolution operators stemming from both the Hubbard-Stratonovich decompositions and the one-bit gates.

To evaluate quantities of the form (as shown above), it is convenient to introduce an additional field χ and consider the quantity

$$Z(\chi) = \langle i|U^+ P e^{i\chi O} P U|i\rangle,$$

so that $\langle O \rangle = -i\partial \ln Z / \partial \chi|_{\chi=0}$. Upon introducing Hubbard-Stratonovich representations for both U and U^+ (auxiliary fields σ and σ' , respectively), we can write

$$Z = \int D\sigma e^{iS}$$

where the action is

$$S(\sigma) = \sum_g \alpha_g (\sigma_g \tau_g - \sigma'_g \tau'_g) - i \ln \langle i|U_{\sigma'}^+ P e^{i\chi O} P U_\sigma|i\rangle.$$

The stationary phase approximation

The simplest approximation to $\ln Z$ is that of stationary phase, where we seek those field configurations at which the action is an extremum; i.e., $\partial S/\partial\sigma = 0$. Let σ_0 be such a mean field configuration. Then the naive stationary phase approximation (SPA) is $\ln Z = S(\sigma_0)$, from which it follows that

$$\langle O \rangle = \frac{\langle i|U_0^+ POPU_0|i \rangle}{\langle i|U_0^+ PU_0|i \rangle},$$

where $U_0 = U_{0G} \dots U_{01}$ is the one-body evolution operator under the stationary fields.

The condition that S be extremal implies a set of $4G$ equations of the form

$$\sigma_g = \text{Re} \frac{\langle i|U_0^+ PU_{0G} \dots U_{0g+1} AU_{0g} \dots U_{01}|i \rangle}{\langle i|U_0^+ PU_0|i \rangle},$$


where A is an operator for one of the qubits involved in gate g ; these determine the stationary fields self-consistently.

Several remarks should be made at this point. First, there may be more than one stationary solution, in which case their individual contributions must be added coherently and interferences can occur. Second, it is easy to show that time-symmetric solutions ($\sigma'\tau' = \sigma\tau$) always exist, and that solutions breaking this symmetry occur in time-conjugate pairs. Third, when $P = 1$, so that no qubits are prescribed in the final state, the self-consistency conditions reduce to the familiar Time Dependent Hartree Fock equations, in which each qubit of a gate evolves in a field generated by the instantaneous wavefunction of the other; the mean fields are then independent of the final

state. However, when P is different from 1, the evolution is non-local in time and the mean fields depend upon the measurements made. Fourth, quadratic corrections to the naive SPA can be readily evaluated; these involve taking determinants of matrices of dimension $4G$. Finally, it is important to note that while the mean field method is most easily carried out when the initial state is a direct product, correlated states can be handled and any correlations present will influence the mean field results.

Example: The Controlled Not gate

As an example of the application of the mean-field technique, consider a single Controlled Not gate as defined by the operator G . Suppose that for an arbitrary initial state $|I\rangle$, the final state of the control qubit is prescribed to be $|p\rangle = p_0|0\rangle + p_1|1\rangle$ and the average value of the control bit, $\langle b \rangle$, is measured. We can represent this computation graphically as



Since it follows that $C^+bC = b + a(b_- - b)$, where the operator $b_- = 1 - b$, the exact result is

$$\langle b \rangle = \frac{\langle i|C^+P_\psi bC|i\rangle}{\langle i|C^+P_\psi C|i\rangle} = \frac{\langle i|C^+P_\psi C(b + a(\bar{b} - b))|i\rangle}{\langle i|C^+P_\psi C|i\rangle}$$

where $P_\psi = |\psi\rangle\langle\psi|$ is the projector for the final state of the control qubit.

the exact result if

$$\langle b \rangle = \frac{\langle i|C^+P_\psi bC|i\rangle}{\langle i|C^+P_\psi C|i\rangle} = \frac{\langle i|C^+P_\psi C(b+a(\bar{b}-b))|i\rangle}{\langle i|C^+P_\psi C|i\rangle}$$

where $P_\psi = |\psi\rangle\langle\psi|$ is the projector for the final state of the control qubit.

In the mean field approximation, the stationary configuration is

$$\begin{aligned}\sigma' = \sigma &= \text{Re} \frac{\langle i|U_0^+ P_\psi a U_0|i\rangle}{\langle i|U_0^+ P_\psi U_0|i\rangle} \\ \tau' = \tau &= \text{Re} \frac{\langle i|U_0^+ P_\psi J_{xb} U_0|i\rangle}{\langle i|U_0^+ P_\psi U_0|i\rangle}\end{aligned}$$

Here,

$$U_0 = e^{i\pi a(\tau-1/2)} e^{i\pi\sigma J_{xb}} \equiv U_a U_b$$

is the mean field evolution operator. If we define $\tilde{P}_\psi = U_a^+ P_\psi U_a$ the mean fields can be written as

$$\sigma = \text{Re} \frac{\langle i|\tilde{P}_\psi a|i\rangle}{\langle i|\tilde{P}_\psi|i\rangle}; \quad \tau = \text{Re} \frac{\langle i|\tilde{P}_\psi J_{xb}|i\rangle}{\langle i|\tilde{P}_\psi|i\rangle}$$


and the required observable is

$$\langle b \rangle = \frac{\langle i|\tilde{P}_\psi U_b^+ b U_b|i\rangle}{\langle i|\tilde{P}_\psi|i\rangle} = \cos^2\left(\frac{\pi\sigma}{2}\right) \frac{\langle i|\tilde{P}_\psi b|i\rangle}{\langle i|\tilde{P}_\psi|i\rangle} + \sin^2\left(\frac{\pi\sigma}{2}\right) \frac{\langle i|\tilde{P}_\psi \bar{b}|i\rangle}{\langle i|\tilde{P}_\psi|i\rangle} - \sin(\pi\sigma) \frac{\langle i|\tilde{P}_\psi J_{yb}|i\rangle}{\langle i|\tilde{P}_\psi|i\rangle}.$$

To understand the validity and limits of the mean field approximation, one can consider several limiting cases. When $|\psi\rangle$ is a ‘‘pure’’ state (i.e., $|\psi\rangle = |0\rangle$ or $|1\rangle$), we find $\sigma = \psi$ and mean field and exact results agree. This is significant observation in that it demonstrates conditional dynamics: results of observing the target bit depend upon the prescribed final state of the control bit. This is not an obvious property of an approximation involving only one-body evolution, which might be supposed to make all qubits independent of

one another. However, it arises at the mean field level through the final-state dependence of the mean fields.

When $|a\rangle$ is not a pure state, insight is more difficult. However, there is some simplification when $|i\rangle$ is a product state, such as $|i\rangle = |i_a\rangle|i_b\rangle$. Consider for example the case for which the control bit is both nput and measured in a state $|a\rangle = a_0|0\rangle + a_1|1\rangle$, the target bit is input in the state $|0\rangle$, and then its mean value is measured in the output. Thus,



The exact result is easily found to be

$$\langle b \rangle_{ex} = \frac{|a_1|^4}{|a_0|^4 + |a_1|^4} = \frac{\langle a \rangle^2}{(1 - \langle a \rangle)^2 + \langle a \rangle^2},$$

where $\langle a \rangle = |a_l|^2$ is the average value of the control bit (in either the initial or final state).

To evaluate $\langle b \rangle$ in the SPA, we find the time-symmetric mean fields to be $\tau = 0$ and

$$\sigma = \text{Re} \frac{\langle a | a e^{-i\pi a/2} | a \rangle}{\langle a | e^{-i\pi a/2} | a \rangle} = \frac{|a_1|^4}{|a_0|^4 + |a_1|^4} = \langle b \rangle_{ex}$$

so that the required observable is

$$\langle b \rangle = \langle 0 | e^{-i\pi \sigma J_x} b e^{i\pi \sigma J_x} | 0 \rangle = \sin^2 \left(\frac{\pi \sigma}{2} \right) = \sin^2 \left(\frac{\pi \langle b \rangle_{ex}}{2} \right).$$

Note that the mean-field and exact results agree when $\langle b \rangle = 0, 0.5$, and 1 , and differ by at most 0.11 when $\langle b \rangle_{ex} = 0.22$ and 0.78 .

Other examples like those above can be worked out, with the effort growing as the number of two-bit gates involved. For example, the inclusive observation of one of the output bits of the two-bit Discrete Quantum Fourier Transform [6] yields a trigonometric function of the input one-body expectation values that is exact in the limits of an “aligned” input state (all input qubits either 0 or 1) and interpolates between them, much as in our results above.

Summary and future considerations

We have considered the application of one of the standard techniques of many-body dynamics (the mean field approximation) to quantum computation. A Hubbard-Stratonovich decomposition of the action of each two-bit gate represents the exponentially difficult many-body evolution as a superposition of many one-body evolutions, each in a different one-body field. The stationary phase (mean-field) approximation to the resulting integral reduces the problem to that of one-qubit quantum dynamics, which is only polynomially difficult.

The formalism and methods were demonstrated by application to a single Controlled Not gate. For an arbitrary initial state, the mean field approximation to the average value of the target qubit in the final state is exact when the control qubit is measured to be in either “aligned” state, and interpolates between these two limits when the control qubit is in a rotated state.

However, this apparently fails when a specification of the target bit is used to determine a measurement of the control bit, the prototype of quantum function inversion.

We have not considered explicitly the quadratic corrections to the mean field results. These would involve the correlations among the fields, ultimately coming down to the manipulation (determinant, inverse, etc.) of $2G \times 2G$ matrices. Consistent with experience in other quantum systems, it is likely that these would improve the naive mean field observables, although the extent to which this is so (and the degree of dependence upon the specific quantum computation considered) remains to be demonstrated.

The failure of the mean-field in solving the simplest function inversion problem is perhaps the most troubling result. It is clearly associated with the fact that although the control qubits are unaffected by the unitary transformation, measurements of them are affected by the projection associated with measurements of the target bits. However, a deeper understanding of the phenomenon is desirable. It is possible that the defect could be remedied, since mean field results often depend upon the particular way the Hubbard-Stratonovich decomposition is implemented.

2.4 Simulating quantum physics on a quantum computer

Feynman[5] has pointed out that simulation of even simple nonrelativistic quantum-mechanical systems on a classical digital computer can involve

exponential effort because of the very large dimension of the Hilbert spaces involved. A system of L interacting spins at fixed locations, for example, occupies a Hilbert space of $(2s + 1)^L$ dimensions (for spin s). This has led to the suggestion that some exponentially difficult computational problems, even if they bear no relation to physics, might be solvable with polynomial effort on a quantum computer. In this section, we return to Feynmann's original observation and ask what class of physical quantum systems might one simulate in polynomial time on a quantum computer? The answer appears to be that it is (or includes) a large and interesting class.

Simple closed nondissipative systems

The dynamics of a single spinless point particle of mass m moving in a potential $V(x)$ is described by the Hamiltonian

$$H = \frac{p^2}{2m} + V(x), \quad (2-5)$$

where the momentum p is Fourier conjugate to x . The unitary operator that evolves this system from time t_1 to t_2 is an exponential in H :

$$|\psi(t_2)\rangle = U(t_2, t_1)|\psi(t_1)\rangle = \exp[-2\pi i(t_2 - t_1)H/h]|\psi(t_1)\rangle. \quad (2-6)$$

Henceforth we adopt units in which Planck's constant $h = 1$. The position x is a single real number for a particle moving in one dimension, or a set of d real numbers for a particle in d dimensions. For notational convenience, $d = 1$ until further notice.

Since x is a real number, the position eigenstates $|x\rangle$ of the physical system live in an infinite-dimensional Hilbert space. To simulate this system

on a quantum computer, we are forced to reduce $|x\rangle$ to a finite number (L) of qubits, whose collective eigenstates we may regard as fixed-point binary approximations to x :

$$x \approx (q_0 + 2q_1 + \dots + 2^{L-1}q_{L-1})\Delta x, \quad (2-7)$$

where Δx is the smallest-resolvable difference in position, and each $q_k \in \{0, 1\}$ labels the eigenstates $\{|0\rangle, |1\rangle\}$ of a two-state system (qubit). For convenience, we write simply $|x\rangle$ for a simultaneous eigenstate of the L qubits approximating x . Since we cannot represent numbers larger than $x_{\max} \equiv (2^L - 1)\Delta x$ in the form Equation (2-7), we regard the system as periodic in x :

$$|x + x_{\max}\rangle = |x\rangle, \quad V(x + x_{\max}) = V(x). \quad (2-8)$$

As always, the physical system can be represented by momentum eigenstates $|p\rangle$ instead of position eigenstates $|x\rangle$. In the simulation, these two bases are related by a unitary transformation

$$U_{FFT}|x\rangle = 2^{-L/2} \sum_p e^{2\pi i p x} |p\rangle, \quad (2-9)$$

which can be implemented by Coppersmith's [6] quantum mechanical Fast Fourier Transform (QFFT). The representable values of p are integral multiples of $2^{-L}\Delta x^{-1}$, and the simulation is periodic in momentum as well as position space:

$$|p + \Delta x^{-1}\rangle = |p\rangle. \quad (2-10)$$

The simulation will proceed iteratively in time steps of size Δt , where to avoid aliasing one should choose

$$\Delta t \leq (2E_{\max})^{-1}, \quad (2-11)$$

if E_{\max} is the largest energy representable in the simulation. In addition to the L “coordinate” qubits needed for x or p , the simulation will require N_f “function” qubits to represent $V(x)\Delta t$ or $(p^2/2m)\Delta t$, plus some number N_g of “garbage” qubits needed during the computations of $V(x)\Delta t$ and $(p^2/2m)\Delta t$. The garbage bits include tables of fixed constants such as the mass m and Δt . Thus a typical state of computer would be

$$\sum_x \alpha(x) |x\rangle_c |f(x)\rangle_f |g(x)\rangle_g \quad (2-12)$$

where the subscripts on the ket vectors distinguish the coordinate, function, and garbage sectors, and each $\alpha(x)$ is a complex amplitude.

It is important that the function and garbage qubits return their initial state ($\equiv |f_0\rangle_f |g_0\rangle_g$) at the end of each iteration; fortunately this can always be arranged by a method due to Feynmann, as explained below.

Each time step is divided into the following stages:

1. Compute $V(x)\Delta t$, by applying a unitary transformation U_{pot} that is diagonal in the position representation:

$$U_{\text{pot}} \left(\sum_x \alpha(x) |x\rangle_c |f_0\rangle_f |g_0\rangle_g \right) = \sum_x \alpha(x) |x\rangle_c |V(x)\Delta t\rangle_f |g(x)\rangle_g. \quad (2-13)$$

We assume that U_{pot} can be implemented in a polynomial (in L) number of elementary quantum gates when $V(x)$ is a discrete approximation to a mathematical function that can be written in closed form.

2. Apply a phase to each position eigenstate according to the value represented by the function qubits:

$$U_{\text{ph}} \left(\sum_x \alpha(x) |x\rangle_c |V(x)\Delta t\rangle_f |g(x)\rangle_g \right)$$

$$= \sum_x \exp[-2\pi i V(x)\Delta t] \alpha(x) |x\rangle_c |V(x)\Delta t\rangle_f |g(x)\rangle_g. \quad (2-14)$$

This will require at least N_f distinct gates. The n^{th} gate takes the n^{th} function qubit as input and applies a phase factor $\exp[-2\pi i 2^{n-N_f}]$ if the qubit is in the $|1\rangle$ eigenstate, but leaves the $|0\rangle$ eigenstates unchanged. Here $n \in \{0, \dots, N_f-1\}$. Gates like these are needed by the QFFT also. Since a position eigenstate of the entire computer is a direct product of qubit eigenstates, applying a phase to any qubit is equivalent to applying the same phase to the computer as a whole. Clearly this stage requires N_f single-qubit operations.

3. Next apply U_{pot}^{-1} to the output of stage 2, obtaining

$$\sum_x \exp[-2\pi i V(x)\Delta t] \alpha(x) |x\rangle_c |f_0\rangle_f |g_0\rangle_g. \quad (2-15)$$

This is a version of Feynmann's trick for "resetting" garbage bits. In our case, the trick relies not only on the reversibility of U_{pot} , but also on the fact that U_{ph} is diagonal in any basis where all of the function qubits have definite values, so that the application of the phase Equation (2-14) does not prevent U_{pot}^{-1} from resetting the f and g qubits.

4. Now apply the FFT Equation (2-9) to obtain

$$U_{FFT} \left(\sum_x \beta(x) |x\rangle_c |f_0\rangle_f |g_0\rangle_g \right) = \sum_p \hat{\beta}(p) |p\rangle_c |f_0\rangle_f |g_0\rangle_g, \quad (2-16)$$

where

$$\beta(x) \equiv \exp[-2\pi i V(x)\Delta t] \alpha(x), \quad (2-17)$$

$$\hat{\beta}(p) \equiv 2^{-L/2} \sum_x e^{2\pi i p x} \beta(x). \quad (2-18)$$

5,6,7. By a sequence of unitary transformations $U_{\text{kin}}^{-1}U_{\text{ph}}U_{\text{kin}}$ analogous to the sequence $U_{\text{pot}}^{-1}U_{\text{ph}}U_{\text{pot}}$ of stages 1, 2, 3 above, apply a phase $\exp[-iF(p)\Delta t]$ to each momentum eigenstate $|p\rangle_c|f_0\rangle|g_0\rangle$ in the sum Equation (2-15). Ideally we would take $F(p) = p^2/2m$, but we must be contented with some more complicated approximation to this, for two reasons: First, $p^2/2m$ doesn't satisfy the periodicity condition Equation (2-10). We could cure this by using $F_1(p) = [1 - \cos(2\pi p\Delta x)]/(\pi^2\Delta x^2 m)$, which is equivalent to a finite-difference approximation; that is,

$$F_1(p)|x\rangle = -\frac{|x + \Delta x\rangle - 2|x\rangle + |x - \Delta x\rangle}{2m\Delta x^2}. \quad (2-19)$$

$F_1(p)$ has the periodicity property Equation (2-10). Second and more importantly, however, $p^2/2m$ and $F_1(p)$ are too "large". The largest eigenvalue of $F_1(p)$ is $2/m\Delta x^2$, so that to satisfy the anti-aliasing condition Equation (2-11), we would be forced to a timestep $\Delta t \leq \pi m\Delta x^2/2$, which is exponentially small in L for fixed x_{max} [cf. Equation (2-7)]. So instead, we should use something like

$$F(p) \equiv \frac{V_0 F_1(p)}{\sqrt{V_0^2 + F_1(p)^2}} \quad (2-20)$$

for the kinetic-energy operator, where V_0 is comparable to the maximum value of the potential $V(x)$. The operator Equation (2-20) is $\approx p^2/2m$ for small p , is monotonic in $F_1(p)$, and has the required periodicity.

8. The timestep is completed by applying the inverse Fourier Transform U_{FFT}^{-1} to return to the position representation.

Steps 1 – 8 apply the unitary operator

$$\exp(-2\pi i\Delta t H_{\text{kin}}) \exp(-2\pi i\Delta t H_{\text{pot}}), \quad (2-21)$$

$$H_{\text{kin}} \equiv F(\mathbf{p}), \quad (2-22)$$

$$H_{\text{pot}} \equiv V(x), \quad (2-23)$$

to the state of the computer. The momentum \mathbf{p} has been put in boldface to signify that it is a differential operator—not a number—in the position representation. What we really wanted, however, was not the operator (2-21) but rather $\exp[-2\pi i\Delta t(H_{\text{kin}} + H_{\text{pot}})]$, and this is not the same as the operator ((2-21)) because H_{kin} and H_{pot} do not commute, since one is a function of momentum and the other of position. Nevertheless, (2-21) is unitary, and therefore it can be written as $\exp(-2\pi i\hat{H}\Delta t)$ for some hermitian operator \hat{H} , which we can regard as the effective Hamiltonian of our simulation. It follows from the Baker-Hausdorf formula that

$$\hat{H} = H_{\text{kin}} + H_{\text{pot}} - i\Delta t[H_{\text{kin}}, H_{\text{pot}}] + O(\Delta t^2). \quad (2-24)$$

In fact one can approximate $H_{\text{kin}} + H_{\text{pot}}$ to second order with no extra effort simply by alternating the order in which one applies the kinetic and potential operators at successive iterations:

$$\begin{aligned} & \exp(-2\pi i\Delta t H_{\text{pot}}) \exp(-2\pi i\Delta t H_{\text{kin}}) \exp(-2\pi i\Delta t H_{\text{kin}}) \exp(-2\pi i\Delta t H_{\text{pot}}) = \\ & \exp\left\{-2\pi i(2\Delta t)(H_{\text{kin}} + H_{\text{pot}} + O[\Delta t^2])\right\} \end{aligned} \quad (2-25)$$

Have we avoided the exponential slowdown?

If one wants to simulate a one-dimensional quantum system with a Hamiltonian of the type Equation (2-5) on a classical computer, one does

not keep track of a complex amplitude $\alpha(x, t)$ for every one of the 2^L possible x 's that can be represented on a machine with an L -bit word; to do so would be exponentially difficult. Instead, one assumes that the wavefunction $\alpha(x, t)$ has some smoothness with respect to x and represents it on a grid with spacing δx much larger than the machine precision Δx . This is equivalent to restricting the simulation to a polynomially large number of states with momenta $\leq \hbar/\Delta x$. The scheme of the previous subsection, which requires a quantum computer, uses a wavefunction with 2^L distinct spatial arguments. Thus it can represent states with momentum as large as $\hbar/\Delta x$. This much larger state space is somewhat illusory, however, because in order to avoid an exponentially small time step, we were forced to replace the finite-difference kinetic energy operator Equation (2-19), which dices the wavefunction on the scale of the machine precision, with a regularized version such as Equation (2-21) that has much less resolution.

Another way to view the situation is as follows. Let $N(E)$ be the number of states of the system with energy $\leq E$. In a simulation with time step Δt , energy eigenstates with energies differing by integer multiples of $1/\Delta t$ suffer the same phase change at each time step. Therefore, the number of states that can be simulated without aliasing is $N(1/2\Delta t)$. Since the number of operations required to evolve the system over a fixed interval T scales with the time step as $T/\Delta t$, the algorithm we have described here is efficient if $N(E)$ is a rapidly increasing function of its argument.

Quantum simulation is much more advantageous when the number of degrees of freedom (d) is large because $N(E)$ tends to increase exponentially

with d . Consider for example a particle moving in a spherically symmetric power-law potential $V(\vec{x}) = V_0|\vec{x}|^\alpha$ in d dimensions, with $\alpha > 0$. Then $N(E) \propto E^{\nu d}$, where $\nu = (\alpha+2)/(2\alpha)$. Note that the number of qubits needed to represent the state rises only linearly with d , and the number of operations required to compute the kinetic and potential energies normally rises only polynomially with d ; this depends upon $V(\vec{x})$, but it is true for many-body problems with two-body interactions (in this case d is proportional to the number of bodies). For the power-law potential above, the number of states that can be followed with a fixed time step increases as $(T/\Delta t)^{\nu d}$.

Therefore, if one measures the size of the problem by the number of degrees of freedom, then the quantum computer does offer an exponential advantage over the classical one.

2.5 Cold Baths and Optimization

A broad class of important computational problems require minimization of a continuous function f of several variables $(x_1, x_2, \dots, x_d) \equiv \vec{x}$. When d is large, the absolute minimum of f can be exponentially hard to find because the number of local minima is proportional to $(R/\ell)^d$ if R is the typical size of the domain in any one dimension, and ℓ is the scale over which f varies significantly. But if one is content with a suboptimal solution, a choice of \vec{x} 's for which f is much smaller than usual but not absolutely minimal, then the problem may be only polynomially hard.

For large d , such problems are often attacked by the method of simulated annealing. This algorithm executes a weighted random walk in the \vec{x} domain, with weight $\exp(-\beta f)$, and is inspired by a physical analog: it imitates a classical particle in a potential $f(\vec{x})$ and in contact with a heat bath at temperature β^{-1} . The “temperature” is gradually lowered in hopes that the “particle” will settle into a deep potential well.

We now consider whether something analogous to simulated annealing can be carried out on a quantum computer. Let the simulated system consist of two parts: a particle moving in d dimensions with potential $V(\vec{x})$ equal to $f(\vec{x})$, or possibly some monotonic function of f ; and a set of M spin-1/2 spins. The two parts will be weakly coupled. The hamiltonian for the total system is

$$H = H_d + H_s + \epsilon H_{\text{int}}, \quad (2-26)$$

where

$$H_d = F(\vec{p}) + V(\vec{x}) \quad (2-27)$$

depends only on the position and momentum of the particle, and

$$H_s = G(\sigma_z^1, \dots, \sigma_z^M) \quad (2-28)$$

depends only on the z components of the spins. Here σ_z^m measures the component of the m^{th} spin along the z direction and has eigenvalues $\pm 1/2$. The function G should probably be chosen so that H_s has as many as possible distinct eigenvalues, 2^M . For example one could use a pseudorandom number generator of period 2^M that maps every M -bit integer (representing a spin configuration) to another M -bit integer (proportional to the energy of that configuration). Since H_s depends only on the σ_z^m , which commute with

one another, the energy of the spin system can be measured precisely by measuring all of its spins.

The interaction hamiltonian H_{int} must be chosen not to commute with H_d or H_g . For the sake of argument we will take

$$H_{\text{int}} = C(\vec{\mathbf{p}}) \sum_m \sigma_x^m, \quad (2-29)$$

but there may be better choices. Notice that since $C(\vec{\mathbf{p}})$ depends only on momentum, the interaction does not “measure” the position of the particle at all. The parameter ϵ , imagined to be small, controls the strength of the interaction.

The scheme for minimizing $V(\vec{x})$ is as follows.

1. Prepare the particle in a uniform superposition of all positions, and the spins in the configuration of lowest energy (Σ_0).
2. Evolve the system through some number of time steps $N_\epsilon \propto \epsilon^{-1}$.
3. Measure the spin configuration and reset it to Σ_0 .
4. Repeat steps 2 and 3 as many times as necessary.

The total energy of the system, the expectation value of the total hamiltonian H , is not changed by the evolution in step 2. (Actually the conserved hamiltonian \hat{H} will differ slightly from H by $O(\Delta t)$ or $O(\Delta t^2)$, as explained above, but this is not important.) Because of the interaction, the spin configuration will generally change and will increase its energy. To the extent that

the interaction is weak, the increased energy of the spins (i.e. $\langle H_s \rangle$) must come at the expense of the energy of the particle ($\langle H_d \rangle$). Thus by iterating steps 2 and 3, we gradually lower the energy of the total system and hence the energy of the particle. Eventually, the amplitude for the particle to be in its ground state will be large, so that we can measure \vec{x} and have a good chance of finding the particle in the deepest potential well. At that point classical Newton-Raphson iteration will locate the absolute minimum.

As yet, however, we have not determined whether this procedure offers any speedup (i.e., saves operations) compared to classical simulated annealing, even in principle.

3 QUANTUM NETWORKS

3.1 Circuits for Quantum Computation

It is traditional to refer to the data flow graph of a program of reversible computation as a circuit, as could be directly implemented in hardware if desired. Quantum computation must always be reversible and thus quantum programs are also called circuits. The operations are specified by instructions (fetched and issued by an ordinary computer) and are called gates in the circuit descriptions. Bits are known as qubits, and their quantum state is considerably more complex than that of ordinary bits.

All fundamental operations (gates) and all series of operations (circuits) for quantum computation must be fully reversible. In general no measurements are allowed *within* the circuits. This means that intermediate nodes (qubits) must be returned through reversed operations to constant values, before they can be reused. The physical implementation of quantum computing apparently can maintain system coherence only over a limited product of the number of qubits and the number of operations. As a result, of this coherence limit, the first important resource to conserve in quantum circuits is the permanent consumption of qubits. Second, the number of operation steps (gates) must also be conserved. Third, the complexity of primitive operations (gates) should be minimized by choosing the least expensive

operation (in terms of physical operation complexity) when operator choice is possible.

It can be shown that only a single bit Rotate operator and a two-bit Controlled-Not operator can realize any quantum calculation. However, we are more concerned with what can be efficiently implemented. Thus we will employ two more operations, the three-bit Controlled-Controlled-Not, and a two-bit Rotate. Thus these four gates will be our choice for a set of gates. In addition we need operations to set a qubit to zero (Clear) and operations to read the values of qubits (Measure). Cirac and Zoller [9] have shown how to realize all these operations by simple physical steps within the context of a linear array of trapped ions.

3.1.1 Qubit and State Representation

The above set of four primitive operations, referred to as ‘gates’, are used to perform quantum calculations. These operations must be performed by a physical system such as a laser shining on a suitable array of trapped ions. Each ion, if it were independent of the others, would have two spin states, down represented for our purposes by zero and up represented by one. The coupled array is represented by a superposition of all possible combinations of all up/down states. Thus for L ions, we have 2^L possible combinations of spins, to represent each possible system state. Each state of the system will have an associated complex amplitude α_k whose squared magnitude represents the probability that the system of ion spins is in that

state. This is denoted as

$$\alpha_k |b_{n-1} b_{n-2} \dots b_0 \rangle$$

where the b's are 'qubits' (quantum bits) set to represent k expressed in binary. For example

$$\alpha_3 |011 \rangle .$$

For economy of representation we will represent a system state as a list of possibly complex numbers, α with α_0 appearing on the right hand side (or bottom) of the list. For example for a two spin system

$$S = \alpha_3 |11 \rangle + \alpha_2 |10 \rangle + \alpha_1 |01 \rangle + \alpha_0 |00 \rangle .$$

We will usually only show

$$S = [\alpha_3, \alpha_2, \alpha_1, \alpha_0],$$

so that individual spin state identifiers are only represented by their position in the list. Complex numbers, when used, will be denoted either by the notation

$$(\alpha_{nr} + i\alpha_{ni})$$

or

$$\alpha_n e^{i\phi_n}$$

or more frequently only by the 2-element list

$$[\alpha_r, \alpha_i]$$

where the real part of the complex number appears to be the left and the imaginary part to the right. It will generally be embedded in a larger state list as:

$$S = [\alpha_{nr}, \alpha_{ni}], \dots, [\alpha_{0r}, \alpha_{0i}].$$

For example the system ground state would be:

$$S_0 = [[0, 0][0, 0], \dots [1, 0]].$$

Operations are applied to a state so to produce a new state s_1 as

$$\text{opr}(x)|S_0 \rangle = |S_1 \rangle .$$

Operations must be, by the laws of quantum mechanics, unitary transforms. Thus in general $\text{opr}(x)$ will always have a unitary matrix representation. In practice, this matrix may be too large to express explicitly. The primitive operations all manipulate the values and positions of the α 's, and each represents one or more physical manipulations applied to one or more of the ions.

Operations compose. That is, a string of operators can be composed to form a new operator using a higher level compose operator \circ . The new operator can be applied to a state to produce a new state. Its components are evaluated right to left. For example

$$|S_1 \rangle = \text{opr}(x_2) \circ \text{opr}(x_1) \circ \text{opr}(x_0) |S_0 \rangle = (\text{opr}(x_2)(\text{opr}(x_1)(\text{opr}(x_0)|S_0 \rangle))).$$

In general we can employ a compose operator and express the above formula for example as

$$|S_0 \rangle = \prod_{j=0}^2 \text{opr}(x_j) |S_0 \rangle$$

3.2 Operations

3.2.1 Measurements

After a series of operations on qubits, as represented by a circuit (or program), the 'answer' is *measured* by the physical system. The measurement determines an actual state of the quantum computer that will exist after the measurement is made. Any given state, that is possible, might be observed. It will occur with a probability factor $|\alpha|^2$ associated with the observed state. A measurement constitutes the only output of the quantum computer.

3.2.2 Clear Operation

The clear operation 'cools' a qubit to zero and is designated clear (k). By successive application clear (k) can 'cool' the system to zero such that $\alpha_0 = 1$ and all other α 's are zero. This means all ions have a down spin after all the clear operations occur. This is the operation 'clear' and is the first step in any calculation.

3.3 Gates

3.3.1 Controlled-Not

The Controlled-Not gate (operation) $x(I,J)$ was devised and named by Feynman [7]. $x(I,J)$ performs a type of logical exclusive-or operation on two qubits, I and J where qubit I remains unchanged but qubit J becomes the exclusive-or of I and J. It performs a unitary transform specified by the matrix:

$$\begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{vmatrix}.$$


In our list representation, this means the α 's with qubit I = 1 are swapped by pairs identified by qubit J. For example

$$\begin{aligned} & x(2,0)[\alpha_7|111, \alpha_6|110, \alpha_5|101, \alpha_4|100, \alpha_3|011, \alpha_2|010, \alpha_1|001, \alpha_0|000] \\ &= [\alpha_6|111, \alpha_7|110, \alpha_4|101, \alpha_5|100, \alpha_3|011, \alpha_2|010, \alpha_1|001, \alpha_0|000] \end{aligned}$$

Since we generally only show the values of the alphas, this example would be

$$S_1 = [\alpha_6, \alpha_7, \alpha_4, \alpha_5, \alpha_3, \alpha_2, \alpha_1, \alpha_0].$$

The circuit diagram for this operation (gate) is:



with truth table:

INPUT		OUTPUT	
I	J	I	J'
0	0	0	0
0	1	0	1
1	0	1	1
1	1	1	0

In physical terms, a Controlled-Not can be implemented on an array of ions as a laser pulse that only causes a state change if the system has pairs of energy levels that match the laser frequency (energy) and pairs that do not.

3.3.2 Controlled-Controlled-Not Operator


The Controlled-Controlled-Not $x(K,I,J)$ is equivalent to a logic circuit composed of an AND operation an XOR operation. It is a Controlled-Not additionally conditioned by bit K. It performs a unitary transform specified by the matrix:

1	0	0	0	0	0	0	0	0
0	1	0	0	0	0	0	0	0
0	0	1	0	0	0	0	0	0
0	0	0	1	0	0	0	0	0
0	0	0	0	1	0	0	0	0
0	0	0	0	0	1	0	0	0
0	0	0	0	0	0	0	1	0
0	0	0	0	0	0	1	0	0

It has a truth table:

K	I	J	K	I	J'
0	0	0	0	0	0
0	0	1	0	0	1
0	1	0	0	1	0
0	1	1	0	1	1
1	0	0	1	0	0
1	0	1	1	0	1
1	1	0	1	1	1
1	1	1	1	1	0

and symbol



3.3.3 Rotate Operation

The rotate operation $r(\phi, J)$ is applied to a single qubit (ion) and shifts the probability amplitudes α_0 and α_1 of the individual qubit. In the full state representation of the system, this operation, chooses by pairs, states that are identical except for bit J, and rotates the pair by ϕ .

For example

$$r(\phi, 1)[\alpha_7|111 \rangle, \alpha_6|110 \rangle, \alpha_5|101 \rangle, \alpha_4|100 \rangle, \alpha_3|011 \rangle, \alpha_2|010 \rangle, \alpha_1|001 \rangle, \alpha_0|000 \rangle] =$$

$$\begin{vmatrix} \alpha_7 \cos \phi/2 - \alpha_5 \sin \phi/2 \\ \alpha_6 \cos \phi/2 - \alpha_4 \sin \phi/2 \\ \alpha_5 \cos \phi/2 + \alpha_7 \sin \phi/2 \\ \alpha_4 \cos \phi/2 + \alpha_6 \sin \phi/2 \\ \alpha_3 \cos \phi/2 - \alpha_1 \sin \phi/2 \\ \alpha_2 \cos \phi/2 - \alpha_0 \sin \phi/2 \\ \alpha_1 \cos \phi/2 + \alpha_3 \sin \phi/2 \\ \alpha_0 \cos \phi/2 + \alpha_2 \sin \phi/2 \end{vmatrix}$$

In general α is a complex number. The unitary transform matrix for $r(\phi, J)$ is

$$\begin{vmatrix} \cos \phi/2 & \sin \phi/2 \\ -\sin \phi/2 & \cos \phi/2 \end{vmatrix}$$

The gate symbol for rotate $r(\phi, J)$ is

$$J - \boxed{\phi} - .$$

One specialization of the Rotate operator is $b(J) = r(\pi/2, J)$. Because of its role in the QFFT we will name it “Butterfly” in analogy to the butterfly operation of the FFT. Coppersmith [8] calls this P_J . Its symbol is:

$$J - \boxed{\pi/2} - .$$

It has the unitary transform matrix:

$$\begin{vmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ -1/\sqrt{2} & 1/\sqrt{2} \end{vmatrix}$$

3.3.4 Conditional-Complex Rotation

All ordinary logical operations can be composed from Rotate and Controlled-Not operators. However Coppersmith's Fast Fourier Transform [8] QFFT on qubits requires a special operation $t(J,K)$ that is applied to a set of L qubits. Qubits J and K are caused to interact in such a manner to rotate the amplitudes that correspond to a state with a one in both positions J and K . The rotation factor is a power of a primitive root of unity $\omega = \exp(2\pi i/N)$ where $N = 2^L$. Let $m = 2^{L-1-K+J}$. The rotation factor is:

$$\omega^{m \bmod 2\pi i}.$$

For a 3-ion example, let $L = 3$, $J = 0$ and $K = 1$. Then

$$\omega^2 = e^{i\pi/2}.$$


Therefore starting from our usual state, the resulting state is:

$$\begin{pmatrix} \alpha_7 \cdot e^{i\pi/2} |111\rangle \\ \alpha_6 |110\rangle \\ \alpha_5 |101\rangle \\ \alpha_4 |100\rangle \\ \alpha_3 \cdot e^{i\pi/2} |011\rangle \\ \alpha_2 |010\rangle \\ \alpha_1 |001\rangle \\ \alpha_0 |000\rangle \end{pmatrix}$$

We will call this two-bit rotation operator "Twiddle" in honor of its application in the QFFT. This operator is the Q_{JK} transformation of Coppersmith [8]. It has the unitary transformation matrix

$$\begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \omega^m & 0 \\ 0 & 0 & 1 & \omega^m \end{vmatrix}$$

and has the symbol



3.3.5 Measure

The operation measure(k) measures the state of qubit (k) of the system. This operation generally happens after all other operations are complete and will provide a read-out of one likely answer if several (or many) are possible. While we simulate the operations with a list of possible states

$$S = [\alpha_{n-1}, \dots, \alpha_0],$$

when a measurement of S is made, the result is a state $|\alpha(k)\rangle$ that will appear with probability $|\alpha_k|^2$. Such a measurement of the k qubit will in general change the chances of a given observation on any of the remaining qubits with which it is entangled.

3.3.6 Set of Operators

We list our selection of primitive operators as discussed above:

- 1) measure (K) : Measure
- 2) clear (K) : Clear
- 3) x(I,J) : Controlled-Not
- 4) x(K,I,J) : Controlled-Controlled-Not
- 5) r(ϕ ,J) : Rotate
- 6) b(J) : Butterfly
- 7) t(J,K) : Twiddle


3.4 Basic Circuits

3.4.1 Full Adder

To perform arithmetic, we need a simple addition operator. Ideally we would like to not consume any scratch space in summing numbers. Thus the truth table for an ideal, full adder would be:

INPUT			OUTPUT		
A	B	C_i	A	S	C_o
0	0	0	0	0	0
0	0	1	0	1	0
0	1	0	0	1	0
0	1	1	0	0	1
1	0	0	1	1	0
1	0	1	1	0	1
1	1	0	1	0	1
1	1	1	1	1	1

This truth table is *not* reversible, that is, given any output state there is not necessarily a unique input state. Therefore a full adder must consume scratch space. In order to find the most efficient full adder, a computer design program was written to search for it, given one bit of scratch space. It discovered an improvement over the full adder given by Feynman [7]. The improved full adder is shown below.



This circuit is fully reversible but does, as it must, consume one scratch (zero) bit. It is represented as


$$fa(A, B, C, Z) = x(B, C) \circ x(B, C, Z) \circ x(A, C) \circ x(A, C, Z).$$

3.4.2 Multi-bit Full Adder

Multi-bit full adders for M bits are composed of M full adders in the classical manner (ripple-carry.)

$$\text{add}(A, B, S) = \bigcirc_{j=0}^{M-1} fa(a_j, b_j, s_j, s_{j+1})$$

The symbol for this adder is:




3.4.3 Multiplexor and Related Circuits

The design program also discovered a particularly good implementation of a multiplexor (mx). Previously the best multiplexor was the use of the more costly Fredkin gate (see below). The function is to select A or B depending upon C. The truth table for the multiplexor is:

INPUT			OUTPUT		
A	B	C	A'	B'	C
0	0	0	0	0	0
0	0	1	0	0	1
0	1	0	1	1	0
0	1	1	1	0	1
1	0	0	1	0	0
1	0	1	1	1	1
1	1	0	0	1	0
1	1	1	0	1	1

The circuit is




A multibit multiplexor is:


$$mx(A, B, c) = \bigoplus_{j=0}^{L-1} mx(a_j, b_j, c).$$

Note that it does *not* consume any scratch space.


Sometimes qubit values must be moved. This can be accomplished with a switch module:




with the symbol



A useful specialization of the multiplexor occurs when both of the inputs are classical (non-quantum) but the control is a quantum variable.



Because C, C₂ can be precomputed, this becomes:



This will be very useful later.


A Fredkin gate can interchange A & B depending upon C:

$$f(A, B, C) = x(B, A) \circ x(A, C, B) \circ x(B, A).$$

It has truth table

INPUT			OUTPUT		
A	B	C_i	A	B	C_0
0	0	0	0	0	0
0	0	1	0	0	1
0	1	0	0	1	0
0	1	1	1	0	1
1	0	0	1	0	0
1	0	1	0	1	1
1	1	0	1	1	0
1	1	1	1	1	1


and the circuit:



A multi-bit Fredkin gate is


$$f(A, B, c,) = \bigcirc_{j=0}^{L-1} f(a_j, b_j, c).$$

We will represent it with the symbol:



3.4.4 Multiplier


A one-bit wide multiplier circuit is simply an AND gate which is just the Controlled-Controlled Not with a zero input on P: $\text{mult}(A,B,P) = x(A,B,P)$. To produce larger bit-width multipliers we proceed in several directions. We will first show a 1 by L bit multiplier:



This has the formula

$$\text{mult}(A,b,P) = \sum_{j=0}^{L-1} x(a_j, b, p_j).$$

A two bit by two bit multiplier, designed by hand is:



The truth table for this multiplier is

A_o	A_1	B_o	B_1	P_o	P_1	P_2	P_3
0	0	0	0	0	0	0	0
0	0	0	1	0	0	0	0
0	0	1	0	0	0	0	0
0	0	1	1	0	0	0	0
0	1	0	0	0	0	0	0
0	1	0	1	0	0	0	1
0	1	1	0	0	0	1	0
0	1	1	1	0	0	1	1
1	0	0	0	0	0	0	0
1	0	0	1	0	0	1	0
1	0	1	0	0	1	0	0
1	0	1	1	0	1	1	0
1	1	0	0	0	0	0	0
1	1	0	1	0	0	1	1
1	1	1	0	0	1	1	0
1	1	1	1	1	0	0	1

3.4.5 Subtractor


A Subtraction consists of one's complementing the subtrahend and adding it, one and the addend. The one is conveniently added by asserting one on the carry-in signal s_0 . First we define $x(a_j)$, as shorthand for $x(1, a_j)$, the 1 being a classical bit. Then we define

$$X(A) = \bigcirc_{j=0}^{L-1} x(a_j)$$

where L is the size of A . Subtraction is now:


$$\text{sub}(A, B, D) = X(A) \circ \text{add}(A, B, D) \circ x(d_o) \circ X(A)$$

where it is assumed D is initially set to zero. A symbol for subtraction is




3.4.6 Adder for Recovering Scratch Space

Qubits are precious. A very powerful and general technique for recovering qubits is given in a paper by Beckman, et al. [10]. It employs the reverse of the inverse of the function to recover qubits. For adders the inverse is subtraction. The symbol for a reversed subtractor is:



Now placing this after an adder realizes the desired circuit in which B disappears



The formula for this adder is then

$$\text{plus}(A, B, S) = \text{subr}(A, S, B) \circ \text{add}(A, B, S)$$


where subr is just the reverse of sub:

$$\text{subr}(A, B, D) = x(A) \circ x(d_o) \circ \text{add}(A, B, D) \circ x(A).$$


Also we note that the reverse of plus is just minus, a subtractor that recovers scratch space.

3.4.7 Modulo Adder

To add two numbers A and B modulo N first add the numbers, then determine if the sum equals or exceeds N by subtracting N from the sum and checking the sign bit. This sign-bit (select) is then used to choose the proper value $A+B$ or $A+B-N$. Next scratch space is recovered using the reversed-inverse technique described above. The circuit to do this is:



The symbol for madd is:



The formula to describe the modulo adder is:


$$\text{madd}(A, B, b_L, Z, z_L) = \text{subr}(A, B, Z|b_L) \circ x(b_L, z_L) \circ \text{sub}(A, B, Z|b_L) \\ \circ a/sr(z_L, N, B|B_L, Z) \circ x(z_L) \circ f(B, Z, z_L) \circ x(z_L) \circ \text{sub}(N, B|b_L, Z|z_L) \circ \text{plus}(A, B, Z|b_L).$$

The symbol “|” represents catenation of qubits. Thus $Z|z_L = z_0 z_1 \dots z_L$. Note that the most significant bit (msb) originally associated with $A+B$ is arbitrarily associated with the value opposite to $A+B \bmod N$. This is permissible because addition is symmetric.

A very clever method given by Beckman, et. al. [10] was used to perform the final return to zero of the select line z_L . This was to compare $A+B \bmod N$ to A , both readily available, and employ the result to clear z_L .

The module a/sr was introduced here. It is just the reverse of add or subtract of the classical constant N in order to clear the unused result of the selection. z_L (sel) is the select line that controls this. The adder is special only in that either $+N$ or $-N$ is added to B , as selected by sel, and the N bits are classical. While the circuit shows a three input gate, in reality, since it is connected to a classical bit (N), if this bit is one, the corresponding two input operation $x(\text{sel}, Z_k, Z_{k+1})$ will be issued. If N , the classical bit, is zero,

no operation at all will be issued. The bit slice is:



The formula for this function is:

$$a/s(sel, +n_k | -n_k, b_k, z_k, z_{k+1}) = x(b_k, z_k) \circ x(b_k, z_k, z_{k+1}) \circ x(1, sel) \circ x(-n_k, sel, z_k) \circ x(-n_k, sel, z_k, z_{k+1}) \circ x(1, sel) \circ x(+n_k, sel, z_k) \circ x(+n_k, sel, z_k, z_{k+1}).$$

Now the a/sr module for L-bits is just the straight reverse of this bit slice arrayed in reverse order:

$$a/sr(sel, N, B, Z) = \prod_{k=L}^0 a/sr(sel, +n_k | -n_k, b_k, z_k, z_{k-1}).$$

The symbol for this circuit is shown above.

3.4.8 Modulo Multiplier

The general method of multiplying two L bit numbers A and B modulo N is to select each bit of B, and only if it is one, multiply A by the corresponding power of two, then reduce it modulo N and modulo add it to a growing sum. This essentially as it is usually done with pencil and paper. For our special purpose that follows, we need a simpler, but special modulo multiplier that

multiplies a quantum variable P_k by a conditioned classical constant X to produce $P_{k+1}=XP_k$. We will also need to recover scratch bits by use of the reverse-inverse technique.

Our purpose for developing all these circuits is to implement the Shor quantum factoring algorithm. This will require a number of variations of the general circuits that are specialized for the algorithm, primarily to reduce to a minimum the number of required qubits. We now will develop a circuit to implement Shors algorithm.

3.5 A Circuit for Shor's Algorithm

Ever since Feynman [7] suggested the possibility of quantum computing, there has been a search for an important task for it to solve. Shor [11] has found such a task, the only one to date, the factoring of large numbers. We propose to show a complete circuit of Shor's algorithm.

The task is given N , find p such that p divides N . The method of solution is to find the period r of a function $F(A)$ where

$$F(A) = X^A \text{ mod } N$$

and X is relatively prime to N . The algorithm consists of a quantum calculation of $F(A)$, followed by a FFT of A and a search for a "frequency peak" in the spectrum of A via repeated trials. Two circuits are therefore required: one for $F(A)$ and one for $\text{FFT}(A)$. We will derive, from the circuit modules developed above, circuits for each.

3.5.1 Exponentiation

We want a circuit to calculate $F(A) = X^A \bmod N$, where X is to be chosen to be relatively prime to N . The first step is to write a binary expansion of A

$$A = a_0 2^0 + a_1 2^1 + a_2 2^2 \dots + a_k 2^k \dots$$

Then

$$\begin{aligned} F(A) &= X^{a_0 2^0 + a_1 2^1 + a_2 2^2 + \dots + a_k 2^k} \\ F(A) &= (X)^{a_0} \circ (X^2 \bmod N)^{a_1} \circ (X^4 \bmod N)^{a_2} \\ &\quad \circ \dots (X^{2^k} \bmod N)^{a_k} \dots \bmod N. \end{aligned}$$

Because all the powers of $X \bmod N$ can be precomputed this reduces to

$$F(A) = X_0^{a_0} \circ X_1^{a_1} \circ \dots \circ X_k^{a_k} \dots$$

The term X_k is easily calculated as a_k is either zero or one. Thus the term is either one or X_k . However we will not actually perform exponentiation this way as we will combine it with multiplication (below.)

3.5.2 Modulo Multiplication


We desire to multiply a value P by $X^{a_k} \bmod N$:

$$\begin{aligned} X_k P_k &= \sum_{j=0}^{L-1} p_j 2^j X_k^{a_k} \bmod N \\ &= \sum_{j=0}^{L-1} p_j (a_k 2^j X_k \bmod N + \bar{a}_k 2^j \bmod N) \bmod N \\ &= \sum_{j=0}^{L-1} p_j (a_k X_{kj} + \bar{a}_k Y_{kj}) \end{aligned}$$

where $X_{kj} = 2^j X_k \bmod N$ and $Y_{kj} = 2^j \bmod N$. To calculate this we will employ a variable S initially set to zero to which we successively add the X and Y terms. Each X and Y are precomputed classical variables.

$$S' = S + p \wedge (a \wedge X + \bar{a} \wedge Y) .$$

Except for the first time, each additional X and Y addition will require modular reduction so we adopt the madd module developed above. The plus module inside madd must be slightly modified to use X+Y as input but since one of either X or Y must be zero, this is not a problem. A bit slice of the module ema that corresponds to that of add in plus is:



It has a formula:

$$\begin{aligned} \text{ema}(x, y, a, p, s, z_j, z_{j+1}) &= x(a) \circ x(y, a, p, z_j) \circ x(y, a, p, z_j, z_{j+1}) \circ x(a) \circ \\ &x(x, a, p, z_j) \circ x(x, a, p, z_j, z_{j+1}) \circ x(s, z_j) \circ x(s, z_j, z_{j+1}). \end{aligned}$$

Therefore the formula for emod is

$$\text{emad}(X, Y, a, p, S, Z|z_L) = \bigcirc_{j=0}^{L-1} \text{ema}(x_j, y_j, a, p, s_j, z_j, z_{j+1}).$$

This circuit “emad” replaces the add module inside the plus module. The second half of a plus module is the subr module. We need to subtract X and Y as conditioned by a_k just as we did above for addition. The result will be new modules ‘emsb” and its reverse ‘emsr” to replace these modules in the plus and madd modules.

The function of emsb is to calculate

$$S' = S + p(a \wedge \bar{X} + \bar{a} \wedge \bar{Y} + 1)$$

The bit-slice to do this is the same as for ema except for the negation of X and Y. Therefore

$$\text{emsb}(\bar{X}, \bar{Y}, a, p, S, Z|z_L) = \bigcirc_{j=0}^{L-1} \text{ema}(x_j, y_j, a, p, s_j, z_j, z_{j+1}) \circ x(z_0).$$

and

$$\text{emsr}(X, Y, a, p, S, Z|z_L) = x(z_0) \circ \bigcirc_{j=L-1}^0 \text{emar}(x_j, y_j, a, p, s_j, z_j, z_{j+1}).$$

Now we can define mplus:

$$\text{mplus}(X, Y, \bar{X}, \bar{Y}, a, p, B|b_L, Z) =$$


$$\text{emsr}(\bar{X}, \bar{Y}, a, p, Z|b_L, B) \circ \text{emad}(X, Y, a, p, B, Z|b_L).$$

At this point we drop $X, Y, \bar{X}, \bar{Y}, N$ from our notation as they are classical variables that do not further interact with the calculation. Thus we just write $\text{mplus}(a, p, B|b_L, Z)$, etc.

We now have all the modules we need to implement the equivalent of madd discussed above. We will call the new addition module 'maddp'. It will include all of the multiplication and exponentiation circuits discussed above. Any classical input to a gate such as X will not be issued when X is zero and the smaller operation (with X deleted) will be issued when X is one. Finally it can be seen that three-input gates are employed. These can be either direct implementations, or by the use of one extra qubit, be realized by three two-input gates in the obvious manner.

$$\begin{aligned} \text{maddp}(a, p, B, z_L) = & \text{emsr}(a, p, S, Z|z_L) \circ x(b_L, z_L) \circ \text{emsb}(a, p, S, Z|z_L) \circ \\ & a/\text{sr}(z_L, N, B|b_L, Z) \circ x(z_L) \circ f(B, Z, z_L) \circ x(z_L) \circ \\ & \text{sub}(N, B|b_L, Z|z_L) \circ \text{mplus}(a, p, B|b_L, Z). \end{aligned}$$

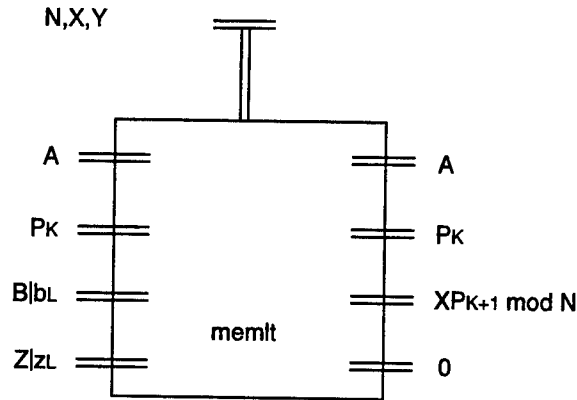
A full multiplier can now be formed from L maddp modules:



The formula is:


$$\text{memlt}(A, P, B|b_L, Z|z_L) = \prod_{j=0}^{L-1} \text{maddp}(a_j, p_j, B, b_L, Z, z_L).$$

A symbol for memlt is:




We must now clear P_k . Again we employ the ‘reverse-inverse’ idea. For the inverse we need division so we can divide XP by X to yield P . Then we reverse this so that P is cleared just as we did for addition as described above.


For Shor’s algorithm, X is relatively prime to N . This means only a selected subset of X ’s will ever be used and each of these X ’s has an inverse mod N . For $N=15$ valid X ’s are $X = \{1,2,4,7,8,11,13,14\}$. $X=1$ would not be chosen however as it would not produce a factorization. All X ’s except the inverse pairs $\{2,8\}$ and $\{7,13\}$ are self inverses. This means that the reverse division module is just a reversed multiplication by X^{-1} . Consider forward multiplication by X^{-1} (X^{-1} is precalculated).



This is reversed and combined with the multiplier




The following is the symbol for this multiplier circuit that clears P, but employs memlt as the basic multiplier.




3.5.3 Calculation of F(A)

The times modules can now be combined to calculate F(A).



This is the circuit to calculate F(A) and has the symbol



and has the formula

$$f(A, P, B|b_L, Z|z_L) = \left(\begin{array}{c} L-1 \\ \mathbf{O} \text{ times}(A_k, P_k, B|b_L, Z|z_L) \\ k=1 \end{array} \right) \circ \text{pmult}(A_0, P_0)$$

where

$$\text{pmult}(A, P) = \begin{array}{c} L-1 \\ \mathbf{O} \text{ pmt}(x_j, a_j, p_j) \\ k=1 \end{array}$$


and

$$\text{pmt}(x, a, p) = x(a) \circ x(a, p) \circ x(a) \circ x(x, a, p).$$

Note that because classical variables are always available we do not show them as arguments except when a specific classical value is defined.

3.5.4 Quantum FFT

The Quantum FFT (QFFT) is a simple calculation in comparison to calculating $F(A)$. Coppersmith [2] devised a very clever algorithm for the QFFT. It uses the Twiddle and Butterfly modules defined above. Here is the QFFT for four qubits.



It is easy to see that for L qubits, the QFFT will require only L Butterfly modules and $L(L-1)/2$ Twiddle modules.

3.5.5 Conditioning

Prior to starting, all qubits must be cleared. This is just $5L+3$ clear operations. Next, all the A bits are conditioned by rotating each of them with a Butterfly operation. This will require $2L+1$ operations.

3.5.6 Overall Operations

The steps to calculate Shors algorithm are then:

- 1) Clear all qubits
- 2) Condition the A qubits with Butterfly operations
- 3) Calculate $F(A)$
- 4) Calculate $QFFT(A)$
- 5) Measure A

3.6 Complexity

Both the space (qubits) and time (gates) complexity of the circuit are easy to determine. The $F(A)$ module requires the most space at $5L+4$ qubits. By following through all the modules, gates can be counted. If only up to three bit gates are allowed, then the total gate count is about $14L + 16L^2 + 292L^3$ gates and operations.

For factoring 15, $L=4$ so this implies about 24 qubits and 18,940 gates. If four bit gates were to be allowed, these numbers would drop to about 23 qubits and 12,900 gates. It may also be possible to reduce qubits allocated to A down to say six bits. This could mean that 20 qubits would suffice. Gate count would drop to roughly 10,000.

The gate counts given here are an upper bound because whenever a classical constant bit is found to be zero, the classical computer will simply not issue the corresponding operation. Also it may be possible to discover more efficient implementations. The table below details the numbers of the different types of gates for $F(A)$ for this implementation.

Gate Degree	Number Required for $F(A)$
0	$2L + 12L^2 + 48L^3$
1	$L + 4L^2 + 44L^3$
2	$L + 56L^3$
3	$48L^3$
TOTAL	$4L + 16L^2 + 196L^3$

Gate Degree	Number Required for $F(A)$
0	$2L + 12L^2 + 48L^3$
1	$L + 4L^2 + 44L^2$
2	$L + 200L^3$
TOTAL	$4L + 16L^2 + 292L^3$

An interesting factoring would be for numbers of about 800 bits This would require about 4000 qubits and 10^{11} gates!

3.7 General Scaling for $F(a)$ Circuit

3.7.1 Shor's Algorithm

In considering quantum computing, one encounters some delicate and essential questions concerning trade-offs, in practice, between scratch space

and ops (operations). The whole situation is a little vague and the terms are not fully defined now — basically there is very little experience in programming a quantum computer.

One should bear in mind however, that programming a general purpose computer in any significant way depends on the existence of a large battery of well defined, easily executed, moderately complex subroutines.

In the theoretical discussions of quantum computing, a great deal of emphasis is laid on the back of the fact that every $K \times K$ unitary transformation is the product of at most $\frac{K(K-1)}{2}$ unitary transformations which are just rotations in two dimensional coordinate planes. As a mathematical observation, this is just a little more than gauss diagonalization of a symmetric matrix, and an elementary proof is easily contrived.

One should remember that in quantum computing, the number K is going to be extraordinarily large (it is as large as 2^{800} in Shor's algorithm), and the typical or generic unitary matrix is going to take at least $K/2$ 2-dimensional rotations as factors in aforementioned decomposition. This is quite easy to see, since the product of L 2-D coordinate rotations will have at least $K - 2L$ ones on the diagonal; i.e. the product of L coordinate rotations is the direct sum of a $2L$ dimensional unitary matrix with a $K - 2L$ dimensional identity matrix.

Clearly, this is no way then to program the generic unitary transformation; the number of ops is daunting.

Nor is this the decomposition which actually occurs in some of the computational procedures described in the literature, the Fourier Transform being the prime example.

Let us note some other examples. A common task in quantum computing will probably be the generation of the uniform superposition of the states $|a, f(a)\rangle$, $a = 0$ to $2^L - 1$, from the uniform superposition of the states $|a\rangle$.

Shor's algorithm depends on the detailed case $f(a) = x^a \pmod n$.

Even the simple case of copying a bit, say $f(a) = a_0$, the first bit "a", is not readily executed as a product of 2-D rotation. In the Cirac-Zoller scheme it is executed with one pair of laser pulses, even though it would appear to require about 2^{L-1} 2-D rotations.

Let us look in a little more detail at what the computation of $\frac{1}{\sqrt{q}} \sum |a, f(a)\rangle$ requires. One might imagine starting with two registers, one for "a", the other, initially filled with all zeros, is for ultimate load by $f(a)$; along the way of the computation it is filled with $g(a)$, a function computed on the way to computing our final $f(a)$. At some point along the computational way, we may have to multiply two large numbers, because we need a tensor product of state vectors, or because of the way we have arranged the computations.

We cannot carry out the computation in the two registers already described, we must make work space for the product elsewhere, and we must do it reversibly because we are "quantum computing". If we generate trash in our computation on the side we can eat it up by the Bennett/Feynman

scheme, but we do pay the price of at least doubling ops and making many essential scratch lines. If somewhat later we carry out an additional computation which uses the product routine just devised, it will fold in the doubled ops already created into the doubling of ops of its own by the outer computation.

All of this discussion simply points out that without very careful control of subroutines, the number of ops in a subroutine is going to repeatedly double, and the side workspace (trash) required will also grow to unpleasant magnitude.

We saw a clear illustration of the above in the development of the JASON program for the computation of Shor's function $x^a \pmod n$ [Section 3.5]. More generally, the following are outstanding questions that need to be more fully explained and quantified. What are the building blocks for efficient quantum computing subroutines, and what are the unitary transformations which can be synthesized from said subroutines? In the synthesis, what is the trade-off between scratch space and ops?

4 IMPACT OF DISSIPATION

Thus far we have assumed an ideal setting in which the computational basis states and superpositions thereof are unperturbed by their surroundings. Of course in practice, no physical system exists in absolute isolation from its environment. The interaction of the qubits of a quantum register with the degrees of freedom in the external environment causes decay of the state of the qubits, with both energy loss and decay of coherence. For example, if a qubit is an elementary spin, then interactions with stochastic external fields (arising from perhaps thermally fluctuating spins in the host medium) can cause a spin in an external field to flip (energy decay), but can as well cause the angular orientation of the spin around the axis of the field to become indeterminate (phase decay).

As should be apparent, for our application the time scale of principal concern is that for the decay of quantum coherence for the overall wavefunction of the state of the quantum computer.[29] Unfortunately, independent of a specific model system, it is not possible to make quantitative statements which characterize the diverse impacts of dissipation on a complex quantum system. Moreover, various possible quantum states will in general be affected quite differently by dissipation. However, we can draw from work in the quantum theory of measurement over the past 15 years to make some simple estimates about the likely role of dissipation and the consequent requirements for the physical implementation of quantum computation.[30]

Towards this end, we first assume independent decay by the various qubits into individual uncorrelated reservoirs (which we take to be in their vacuum states), with the time for the decay of phase coherence for a single qubit denoted by τ . Now, a superposition state formed from J qubits will not decay uniformly at this same rate $1/\tau$. Instead the off-diagonal elements of the density matrix for the previously perfectly coherent quantum computer will decay with time constant set by the global separation of its components in the computational eigenbasis.[31, 32] If we take as the characteristic “distance” between components in a complex superposition state to be the Hamming distance, then an order of magnitude estimate for this distance is J itself, with the time scale for decay then given roughly by τ/J . Stated more physically, superposition states which exhibit “macroscopic” characteristics (such as a coherent superposition of 1,000 spins up plus 1,000 spins down) are incredibly sensitive to dissipation. In the case of a quantum computer, an estimate for the time τ_o for decay (the decoherence time) of an entangled state of J qubits is $\tau_o \sim \tau/J \ll \tau$ for $J \gg 1$.

Now clearly for the operation of the quantum computer we will require that this decoherence time τ_o is small as compared to the total time for the quantum computation. Given a basic clock time T_o for an elementary gate involving the coherent (reversible) interaction of two qubits as discussed in Section 6, then the time for the computation is of order $M T_o$ with M as the total number of operations (or “ops” as estimated here by the number of primitive gates). The requirement that $\tau_o \gg M T_o$ then leads to the following inequality for the ratio of decay time τ for an individual qubit to

the clock time per step T_o

$$\tau/T_o \gg JM. \quad (4-1)$$

To understand the scaling that this inequality implies, note that $\{\tau, T_o\}$ and $\{J, M\}$ will in general be functions of the size of the problem L , where L is the number qubits in the input data register. Without reference to a specific physical system for implementation,[33] we take $\{\tau, T_o\}$ to be constant with respect to L and consider only the implications of the complexity of the quantum network as described in the preceding Section 3.

For the first case, consider the quantum FFT circuit, which we recall is the role model for efficiency. Here the number of ops M scales as L^2 , while the number of “scratch” qubits is 0, so that $J = L$. Hence, Equation (4-1) implies that

$$\tau/T_o \gg L^3 \sim 10^9 \quad (4-2)$$

for the quantum factoring problem with a 200 digit number (roughly 800 bits). Although this is a daunting number with respect to current laboratory demonstrations, one might well be bold enough to state “Nonetheless, let’s get on with it!” However, the second case of the quantum circuit for $F(A)$ provides a discouraging deflation of such enthusiasm. Recall that now the number of ops M scales as L^3 , while the number of “scratch” qubits $J \sim L$. Hence, Equation (4-1) implies that

$$\tau/T_o \gg L^4 \sim 10^{12}, \quad (4-3)$$

where the estimate is again for the quantum factoring problem with a 200 digit number. This number is beyond any reasonable assessment of projected physical capabilities well into the next century.

As emphasized in Section 3, there is no guarantee that the JASON circuit for computing $F(A)$ is optimum in minimizing either the number of elementary operations or the size of the scratch space. However, our circuit together with recent work at Caltech by Professor J. Preskill [10] and at Oxford by Barenco et al. [11] are the only explicit such networks of which we are aware, and it is fair to say that they are certainly not the most inefficient circuits that could be designed. Our quantum circuit together with the above estimates do emphatically make the point of the critical need both for better quantitative analyses of the role of dissipation and for excruciating care in the actual design of quantum circuits to minimize both ops and scratch space. Estimates of computational complexity which give the scaling (e.g., polynomial of order k) are an important first step, but cannot be taken as a substitute for explicit, quantitative results. For example, for our quantum circuit the numerical prefactors in the polynomial expressions for $M \sim L^3$ and $J \sim L$ increase the estimate of τ/T_o by almost 1,000 fold! The fact that quantum bits of information are such a fragile resource underscores the need for explicit constructions if reliable assessments of feasibility of implementation are to be made.

We should stress that although the above estimates represent a reasonable first attempt to assess the role of dissipation, they are based on simple models that describe the decay of coherent superpositions for states of an otherwise isolated “system” into an external “environment”. These results should be taken as qualitative indicators of the nontrivial difficulties presented by dissipation; however, they should not be viewed as setting universal constraints for quantum computation. For the most part the community

that has considered the role of dissipation (with results as sketched above) has relied on a Master Equation formalism for the density matrix of the system for which the degrees of freedom of the surrounding environment are eliminated and no information about system decay is recorded. By contrast, in some physical settings it is possible to observe certain decay channels of the system (e.g., photoelectric detection of escaping photons) and thereby to monitor continuously the system's evolution[25, 30, 34]. In this case, the results of the Master Equation formalism are not necessarily applicable to subensembles chosen relative to some criteria such as the absence of emitted quanta from the system. Furthermore, it seems reasonable that one should consider not only "freely" decaying systems, but should as well investigate the nature of decoherence for driven quantum systems ("quantum food for quantum dynamics").

Of particular note in this regard are recent developments in the area of Quantum Optics, with a premier example being the work of the group of P. Zoller and colleagues at the University of Innsbruck in Austria. As we will discuss in more detail in Section 6, Zoller's group has made two noteworthy proposals for the physical implementation of quantum computation, both of which involve linear strings of trapped ions (or atoms) which are "wired" together on the one hand by a phonon mode [9] and on the other hand by a photon mode.[25] While both of these suggestions may well be realizable in the laboratory (at least with modest numbers of qubits), perhaps their most important near-term impact is that they provide plausible model systems for realistic analyses of quantum computation, including the impact of dissipation and possibilities for various error correction schemes.

Indeed, we feel that it is very important to go beyond the situation of “oracle” pronouncements of the evils of decoherence and to focus instead on quantitative analyses of suitable model systems, where appropriate systems are those for which there exist reasonable microscopic understandings of the system-reservoir interaction in both experimental and theoretical terms. Such analyses can provide an essential bridge between formal quantum algorithms and nascent experimental capabilities.

5 ERROR CORRECTION

Diverse imperfections in the components of classical computers could have devastating impact were it not for the mitigating effects of various error correcting schemes. In broad brush, such schemes have in common a “majority” voting philosophy, where the effect of errors with probability $p < 1/2$ are exponentially suppressed with repeated trials (or parallel processing). If quantum computation is to be implemented successfully, it seems reasonable to assume that analogous error correcting protocols must be developed to deal with the adverse consequences of diverse quantum imperfections, including the fidelity with which quantum transformations might be accomplished as well as (and perhaps much more importantly) the loss of coherence due to interactions with the environment.

At first sight, it might seem a fool’s errand to suppose that a quantum calculation could be “corrected” in route, since any measurement of a quantum system necessarily perturbs the system. Although the business of quantum error correction is a nascent enterprise with no firm “software release dates”, there are some preliminary analyses that offer faint rays of hope. [36, 37] The common theme of the research thus far is to operate $P \gg 1$ quantum computers in “parallel”. For a “correct” calculation free from error, the overall wavefunction for the P computers evolves exclusively in a symmetric subspace of the total Hilbert space, while “errors” in any one of the computers causes the overall state to develop components outside of the symmetric subspace. Since the ratio of “volumes” of symmetric to

total spaces is exponentially small in P , any such error creates a new component of the overall wavefunction that is approximately orthogonal to the symmetric subspace. One then attempts to devise measurement strategies that squelch these orthogonal components by active intervention in the total Hilbert space, thus constraining the quantum evolution (i.e., the “calculation” with P computers) to remain in the symmetric subspace without actually having disturbed this realm. While current research has identified this approach as a promising avenue for error correction in the quantum domain, it remains to be seen whether or not robust algorithms can be developed that can also be physically implemented.

Apart from work to find general formalisms for quantum error correction, it is also important to look to specific model systems both to assess the actual (as opposed to “generic”) impact of imperfections as well as to test candidate correction protocols. In this regard, we would once again point to the work of P. Zoller and colleagues as exemplary. Particularly in their work with atoms interacting via photons in a cavity, [25] one has a detailed microscopic understanding of both the reversible, coherent evolution of the qubits (internal atomic states coupled via a quantized cavity field) as well as of dissipation by way of atomic spontaneous emission and cavity decay. Since the cavity dissipative channel can in principle be monitored with high efficiency, the information thus gleaned can be exploited for error correction, with some counterintuitive results emerging from Zoller’s analysis.

Quite recently two other very promising developments have occurred related to the possibility of actively correcting errors at the level of individual

qubits. Mabuchi and Zoller [34] have identified subspaces which are “invariant” with respect to quantum jumps and for which state decay can be restored by continuously monitoring decay channels. Their scheme involves a redundant encoding of information in multiple qubits. A more general scheme which does not require the detection of decay events in the external environment has been presented by Shor [39].

6 ASSESSMENT OF POSSIBLE PHYSICAL REALIZATIONS

In considering the physical implementation of primitive “gates” for quantum logic, there are two times scales of particular relevance. The first is set by the interaction energy $\Delta E = \hbar\chi$ between a pair of qubits, with $T_o = 1/\chi$ being the time required to effect a coherent change of state of the qubits due to their mutual interaction. Clearly the “clock” cycle for quantum computation can proceed no faster than T_o . The second time scale is that set by the grim reaper of dissipation, with Γ as the damping rate for the coherence of any one qubit and $\tau = 1/\Gamma$ as the single qubit decay time.

For reversible, coherent computation, we require that

$$\tau/T_o \gg 1 \tag{6-1}$$

and hence that

$$\chi/\Gamma \gg 1. \tag{6-2}$$

Since $m = (\Gamma/\chi)^2$ has the physical significance of a critical number of quanta in many settings, Equation (6-2) is equivalent to the requirement that

$$m \ll 1, \tag{6-3}$$

which is to say that a single quanta (e.g., a phonon, photon, elementary charge, ...) must be capable of significantly modifying the interaction between qubits. Condition (6-3) goes by a variety of pseudonyms in different fields of physics; here, we shall adopt the terminology of optical physics with Equation (6-3) being the condition for strong coupling for interactions between qubits.

Although these are quite general qualitative considerations, we might attempt to better understand the order of magnitude implied by the above inequalities by recalling our discussion of dissipation in Section 4. In combination with the explicit JASON quantum circuits relevant to Shor's factoring algorithm, we arrived at the estimate (Equation (4-1))

$$\tau/T_o \gg JM,$$

which in the current setting implies that

$$\chi/\Gamma \gg JM. \tag{6-4}$$

For the case of the quantum FFT circuit, the number of ops M scales as L^2 , while the number of "scratch" qubits is 0, so that $J \sim L$. Hence, Equation (6-4) implies that

$$\chi/\Gamma \gg L^3 \sim 10^9 \tag{6-5}$$

for the quantum factoring problem with a 200 digit number (roughly $L = 800$ bits). As we shall see in the sections that follow, this is a number well beyond current laboratory capabilities. The quantum circuit for $F(A)$ is even more challenging, since in this case the number of ops M scales as L^3 , while the number of "scratch" qubits $J \sim L^2$. Hence, Equation (6-5) implies that

$$\chi/\Gamma \gg L^4 \sim 10^{12}, \tag{6-6}$$

where the estimate is again for the quantum factoring problem with a 200 digit number ($L = 800$ bits). This number is so large as to be ridiculous into the foreseeable future and would seem to preclude the physical implementation of quantum computation for problems with the above stated scaling properties (which apparently includes the quantum algorithms described by

Shor, at least via the JASON quantum circuits). Clearly it is imperative to find more “efficient” algorithms which avoid the simple scaling that arises from a demand for global coherence throughout the entire Hilbert space for the total duration of the computation (as for example by suitable division of the computation into restricted domains of the Hilbert space or by error correction protocols).

In the alternative, one should perhaps broaden the perspective beyond the view of quantum computation as supplanting conventional large scale computation with a concomitant requirement for large numbers of qubits ($> 10^3$) and should consider the field of quantum information more globally. More specifically, there appear to be interesting problems in the domain of quantum communication which require a much more modest scale for the number of globally coherent qubits.[38] In addition, it might well be the case that there could be non-negligible benefits to be gleaned from the exploitation of quantum coherence in the setting of “conventional” computation. As the physical scale of components continues to progress in the realm of nanostructures, one should be alert to opportunities associated with quantum state entanglement on a modest scale of a few qubits.

Against this backdrop of the daunting physical requirements implicit in Shor’s algorithms and of considerable uncertainty as to the specific requirements for other as yet unnamed opportunities, we next turn to consider specific physical systems as candidates for the components of quantum logic gates.[8] Independent of the current excitement over the quantum computer, there has been a clear trend over more than a decade in several areas of

physics to move into the domain of strong coupling as expressed by Equation (6-3). This work represents an intellectual as well as a technical frontier, with potential import well beyond the realm of the quantum computer.

As will become clear in the following sections, no physical system has demonstrated performance sufficient to the task of the implementation on a nontrivial scale of a quantum algorithm such as Shor's. However, modest steps have been taken with two recent demonstrations of conditional logic at the single quantum level. In one experiment a quantum phase gate has been investigated within the setting of cavity quantum electrodynamics (COED).[23] Here the individual qubits are optical photons, with the internal state specified by polarization and with interactions between qubits proceeding via strong coupling to an atom in an optical cavity. In the other experiment, a Controlled-Not gate has been implemented for a single trapped ion.[40] Here, the qubits are the ion with two internal states and phonons associated with the quantized oscillation of the ion in a Paul RF trap. One should stress that in both cases no explicit demonstration of quantum-state entanglement has been made. Hence although these experiments explore conditional dynamics at the level of single quanta (strong coupling as in Equation (6-3)), they come with no warranty that the interactions are sufficient for the actual implementation of quantum logic with entangled states.

In the following three sections, we review some promising physical systems, including the two mentioned above as well as three possible condensed matter systems. This should not be viewed as an exhaustive listing (especially given the rate at which new possibilities have been theoretically

suggested in the past year), but rather as indicative of some of the issues to be considered when evaluating candidate systems.


6.1 Condensed Matter Systems

6.1.1 Introduction

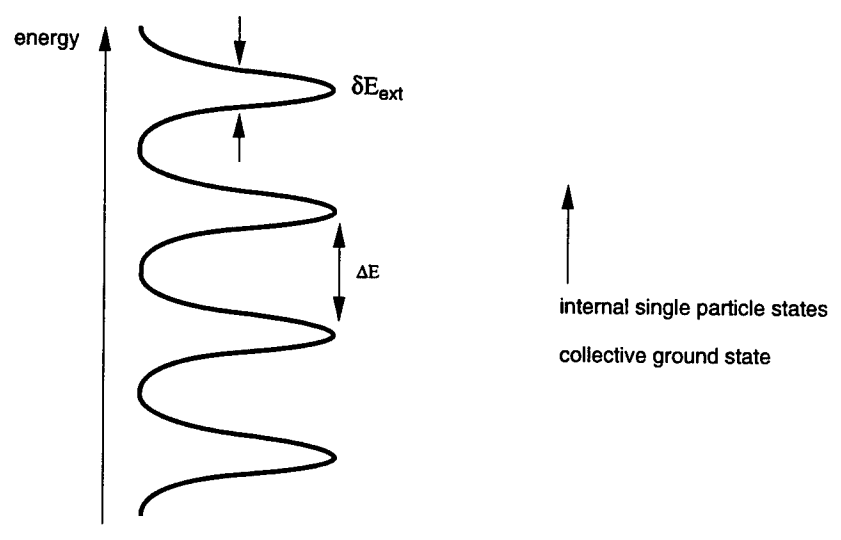
The construction of solid state quantum computers is appealing for many reasons – essentially the same reasons we use solid state computers today: they are compact, rugged, and reliable. Perhaps most importantly, solid state technology has the potential to fabricate complex quantum circuits with desired couplings using extensions of advanced processing techniques developed for the semiconductor industry. At present, atoms in traps are the leading contenders for the experimental realization of a quantum computer, because the high degree of isolation of these atoms from their environment preserves the coherence of quantum states for relatively long times. Q-bits of information in solid state quantum computers will likely be represented by collective states of many particles in close proximity with their environment, and the challenge will be to preserve the coherence of the q-bit against the internal collective modes and external perturbations.

Figure 6-1 illustrates a hypothetical spin system which serves to illustrate the challenges facing solid state quantum computers. The circle represents a very small spherical ferromagnetic particle containing N aligned spins. Small particles of this type are the basis of super-paramagnetic materials, so-called because the total spin of the particle can rotate freely in response to an applied magnetic field. Large ferromagnetic spheres are used as microwave resonators. Suppose that the quantum variable of interest is the z -component $m_z h$ of the total spin along an externally applied magnetic field. The total spin J represents a collective state of the N individual spins j aligned by their ferromagnetic interaction. For use in a quantum computer this collective state should be isolated both from internal excitations and from external interference, because either can destroy quantum coherence of the many body state.

Internal excitations exist for any quantum system composed of more than one particle. Ideally, one would like a spectrum having no excitations with energy below that of the q-bit collective state, so that break up into internal excitations is energetically forbidden. This is trivially true for nuclear excitations, for example. However, solid state systems composed of more than one atom possess excitations associated with the internal motion of the atoms, spin, and charge. The lowest energy internal excitations for the super-paramagnetic particle example illustrated in Figure 6-1(a) are spin waves consisting of continuous rotations of the local magnetization. For large systems the spin wave spectrum is essentially continuous and extends to very small energies. However, for small superparamagnetic particles the lowest energy spin wave excitation is finite, determined by the phase velocity



(a)



(b)

Figure 6-1.

and the size of the particle. Both the spectrum of internal excitations and the strength of their coupling to a collective quantum variable determines how fast coherence is destroyed. Although solid state systems generally possess less symmetry than isolated atoms, constraints such as energy and momentum conservation still apply and play a large role in determining the relaxation rate. For example, in a perfect ferromagnetic sphere, the total spin does not couple well to spin waves even when they are energetically allowed.

External perturbations of the quantum states in a quantum computer change their energy spectrum and cause a loss of coherence. For a discussion of these issues in mesoscopic quantum systems see Altshuler et al. (1991) [12] and Beenakker and van Houten (1991) [14]. Figure 6-1(b) schematically illustrates the energy level diagram for the super-paramagnetic sphere in a magnetic field, including perturbations due to external spins. For a fixed configuration of external spin states, the energy levels of the entire system consisting of sphere and external spins remains sharp, but each level of the superparamagnetic sphere is shifted by an amount typically $\sim \delta E_{\text{ext}}$. The pattern of energy level shifts is typically complex and depends sensitively on the configuration of external spins. Performing an ensemble average over external spin states leads to a statistical distribution of levels represented by the level width in Figure 6-1(b). In order to avoid a dependence of the q-bit on the state of external spins, the broadening due to the time of observation Δt_{obs} must be made larger than the interaction energy $\delta E_{\text{ext}} \ll \hbar/\Delta t_{\text{obs}}$. It is clear that external interactions must be minimized in order to provide an observation time long enough to perform useful computation. This is difficult at present in solid state systems, due to the low characteristic energies $< 1 \text{ eV}$

and due to the close proximity of the quantum device to particles in its environment.

Ideally the concept of temperature would not be relevant to quantum computers, because they are assumed to be in well defined quantum states, out of equilibrium with their environment. In reality of course, temperature is a major consideration. Often it is useful to separate the internal thermalization of quantum states of the device and external thermalization to the environment. The internal thermalization rate is determined by the relaxation rate of the quantum variable via inelastic processes into unwanted internal excitations as discussed above. For certain systems, the internal thermalization rate can be made quite slow, either by arranging for the excitations to be energetically forbidden, or by fully accounting for all accessible internal states as for atoms. The external thermalization rate is determined by inelastic processes due to coupling of the quantum system to its environment. A first step toward reducing this coupling is to make the energy level spacing much greater than the ambient temperature, as for atoms in traps, so thermal electromagnetic radiation rarely causes transitions. For solid state systems this typically means placing the device inside an enclosure cooled below room temperature. Other inelastic processes also contribute to thermalization and destroy coherence; for the hypothetical spin system above an additional inelastic process would be flipping an external spin. Enumerating the important inelastic processes for a given system and evaluating their relative strengths is a difficult problem subject to great uncertainty: some inelastic processes are fundamental in nature and intrinsic to the system, others such as impurities or structural imperfections differ with material and

the manufacturing process, all are generally temperature dependent (see Altshuler et al. 1991, Beenakker and van Houten 1991) [12, 14].

Solid state quantum devices typically differ from isolated atoms in that atoms are identical and possess a high degree of symmetry. The internal excitations of atoms are the same from one atom to the next and they are incorporated into the energy level spectrum, rather than being thought of as unwanted internal modes. As a consequence of the high degree of symmetry, the level spacing of atoms is not uniform and one can select transitions of the desired energy for use in a quantum computer, e.g. optical transitions with energies much greater than room temperature. These advantages are degraded to some extent as one couples many atoms together to form a quantum computer, as in the Cirac and Zoller scheme discussed elsewhere in this report.


By contrast, man-made solid state devices such as ferromagnetic particles, quantum dots, or superconducting quantum interference devices (SQUID'S) typically consist of many atoms; the nominally identical copies of a given device are neither precisely identical nor perfectly symmetric. The energy of quantum transitions in solid state devices typically lie below the optical range, and current solid state quantum devices generally require cooling. As a consequence of disorder the level spacing of states in solid state quantum devices is often approximately uniform, a defining characteristic of chaotic quantum systems (see Altshuler et al. 1991, Gutzwiller 1991)[12] [17]. Because copies of a quantum device are not precisely the same, their characteristics will differ. These disadvantages are offset by the great advantage that

solid state quantum devices and circuits can be designed and constructed in custom designed configurations not found in nature.


6.1.2 Solid State Quantum Device Examples

Single electron logic and single flux quantum logic are advanced approaches to solid state computation which represent each logical bit by one electron and by one magnetic flux quantum respectively (see Altshuler et al. 1991, Grabert and Devoret 1991, Tinkham 1995) [12] [16] [18]. Although based on the quantization of charge and magnetic flux, both are classical in the sense that bits of information are passed through the circuit as incoherent classical particles. Single electron and single flux quantum logic provide good starting points to evaluate possible extensions for use in a quantum computer, because they are the subject of a great deal of current experimental and theoretical research, and many properties are well understood.

Figure 6-2(a) illustrates a quantum device, the single electron transistor (SET) which is the basis of single electron logic (see Altshuler et al. 1991, Grabert and Devoret 1991, Tinkham 1995)[12] [16] [18]. The SET consists of a small quantum dot coupled to the circuit via a pair of tunnel contacts, as indicated, and an electrostatic gate. If the tunnel conductance of each contact is small, $G_c \ll 2e^2/h$, the number of electrons on the dot is a good quantum number. The charging energy necessary to place a single excess electron on the dot is $e^2/2C_\Sigma$, where C_Σ is the total capacitance of the dot. This energy acts as a barrier to tunneling if the charging energy is much



(a) Single electron logic



(b) Single flux quantum logic

Figure 6-2.

greater than the temperature, known as the Coulomb blockade. In the region of applied contact and gate voltages corresponding to the blockade, the number of electrons on the dot is fixed, and can be used to represent a bit of information. In single electron logic, transitions between charge states are generally assumed to take place incoherently, so that all phase information is lost and the computation is classical.

One can imagine forming a quantum computer from a collection of quantum dots connected together into a circuit via tunnel contacts, electron waveguides, and electrostatic gates. Under ideal conditions the phase coherence of the electron charge representing the logical bit would be preserved over the entire circuit, making quantum computation possible. To date, phase coherence has been observed for tunnel-coupled double and triple quantum dots in a GaAs/AlGaAs heterostructure (Waugh et al. 1995),[20] but not for larger circuits.

As discussed above, both internal and external excitations act to destroy phase coherence. As indicated in Figure 6-2(a), the lowest energy internal excitations for quantum dots are typically single particle excitations with energy spacing E_F/N , where E_F is the Fermi energy, and N is the total number of electrons on the dot. The level spacing is approximately uniform as a consequence of disorder as for quantum chaotic systems (see Gutzwiller 1991) [17]. In order to avoid thermal excitation of single particle excitations and loss of coherence, the temperature must be made much less than the level spacing.

For two dimensional electron gases (2DEG) in high quality GaAs/AlGaAs

heterostructures, the dominant external mechanism which destroys coherence at low temperatures is electron-electron scattering (see Ando et al. 1982, Altshuler et al. 1991) [13] [12]. The electron-electron scattering rate $1/\tau_{ee}$ is temperature dependent, and has two parts, one due to particle-particle collisions, and one due to collective charge fluctuations which can be thought of as dephasing by Nyquist noise (see Altshuler et al. 1991) [12]. The electron-electron scattering rate is sensitive to the geometry of the device, and differs for two-dimensional sheets, one-dimensional wires, and zero-dimensional dots. For a two-dimensional electron gas in GaAs at the base temperature of a typical dilution refrigerator ~ 10 mK, the coherence time is of the order $\tau_{coh} \sim \tau_{ee} \sim 10$ nsec. For a quantum dot isolated by tunnel contacts, the coherence time could be considerably longer, because the scattering between electrons in the dot and external electrons is cut off by their physical separation; no data in this limit are currently available.

An important figure of merit for quantum computers is the ratio of coherence time to switching time. The switching time τ_{sw} for a quantum dot is limited by the charging time $\tau_{RC} = C_{\Sigma}/G_C$ necessary to charge the dot through the tunnel contact conductance $G_c < 2e^2/h \cong 1/(12k\Omega)$. For quantum dots of size $\sim 0.1\mu\text{m}$ made using current lithography the switching time is of the order $\tau_{sw} \sim 10$ psec. Thus, at present, the coherence time for single electron logic is a factor $\sim 10^3$ longer than the switching time, with sizable uncertainties associated with operating temperature, and device characteristics. This ratio could improve by orders of magnitude in the future as improved technology permits the fabrication of small quantum dots approaching the size of large molecules.

Single flux quantum logic is based on the trapping of a single flux quantum by a superconducting quantum interference device (SQUID), as indicated schematically in Figure 6-2(b) (see Grabert and Devoret 1991, Tinkham 1995)[16] [18]. The SQUID consists of a superconducting ring which traps magnetic flux via the Meissner effect, with a Josephson junction or weak link through which magnetic flux quanta can pass, either by activation or by tunneling. It is interesting to note that a SQUID is approximately the electromagnetic dual of a single electron transistor: the quantum of magnetic flux trapped in the ring replaces the quantum of electric charge trapped on the dot. As electronic devices, SQUID's are far more highly developed than SET'S, and they are available commercially for a number of applications. Macroscopic quantum tunneling of single flux quanta through a Josephson junction is a relatively recent discovery, and is not the basis of operation of conventional SQUID'S. Macroscopic quantum coherence of a single flux quantum tunneling coherently between two energetically equivalent states in a SQUID has been searched for experimentally for a number of years, but not yet observed.

The internal excitations of an ideal SQUID are quasi-particles produced by excitation across the superconducting energy gap. For submicron superconducting devices, it is possible experimentally to freeze out quasi-particle excitations entirely at a modest fraction of the superconducting transition temperature, so that every electron is paired (Tuominen et al. 1992) [19]. Much less is understood about the loss of coherence of flux quanta in SQUID's than for electrons in semiconductors, and it is difficult to make quantitative comparisons. At first one might think that the coherence of the supercon-

ducting state itself is the relevant issue, and that the coherence times are very long, but this is not the case. Each magnetic flux quantum which represent a bit of information in a single flux quantum or quantum computer is a collective excitation of the superconducting ground state. Magnetic coupling to nearby impurities or conductors can produce a loss of coherence. The reason why macroscopic quantum coherence has not yet been observed in SQUIDs at low temperatures is thought to be external dissipation due to eddy currents in nearby metallic conductors (see Tinkham 1995) [18].

Solid state quantum devices are attractive candidates for use in quantum computers of the future, because one can build on a large base of technological expertise. Incoherent quantum circuits – single electronics and single flux quantum logic – are at the edge of current technology, and the coherent quantum circuits needed for a quantum computer are just beginning to be studied in research laboratories. The ratio of coherence time to switching time in current quantum devices is adequate in some cases to permit measurements on small circuits, but quantum computers that perform useful tasks will require much larger ratios as discussed elsewhere in this report. The construction of a practical quantum computer by any means probably lies at least several decades in the future, so there is enormous uncertainty concerning its architecture and the technology available at that time. The capabilities of solid state quantum devices will likely continue to advance at a rapid pace, driven by applications in conventional computers and electronics, and they may well prove to be useful for quantum computation.

6.2 Optical Systems

Within the context of optical physics, the most promising candidates for quantum logic and computation are currently systems in cavity quantum electrodynamics (CQED), where the basic configuration is that of an atom situated inside a resonant cavity [21]. The atom and cavity field interact via a dipole coupling of the atomic transition moment to the electric field of the cavity. For the case of coincident atomic and cavity resonance frequencies (detuning $\delta = 0$), the interaction energy $\Delta E = \hbar g$ is expressed in terms of the rate g for oscillation between atomic ground and excited states for a single photon in the cavity (more specifically, g is half the Rabi nutation frequency induced by a single photon for $\delta = 0$). Dissipation proceeds by way of either atomic spontaneous emission at rate γ or by cavity decay at rate κ , with now two rates replacing our previous single decay rate Γ . The condition (6-3) for strong coupling translates into a statement that a single photon should drive coherent evolution over a time scale much shorter than the time scale for dissipation ($g \gg (\gamma, \kappa)$), which has been achieved in CQED. We should emphasize that although the condition for strong coupling is easy to express, the experimental realization of this regime has been a rare event in physics. Because of this, a number of theoretical schemes have been proposed for implementing quantum logic via CQED.[22] Indeed, the two demonstrations to date of quantum logic at the level of single quanta, one has been implemented via a nonlinear interaction between pairs of photons which is mediated by an atom in a cavity.[23]

More generally, in the setting of CQED the qubits can be either photons (with internal state specified by polarization) or atoms (with nondegenerate ground and excited states or with degenerate Zeeman levels). Conditional logic can be achieved by either “absorptive” or “dispersive” type processes, where the presence or absence of a photon (atom) leads to a different dynamic for the evolution of the atom (photon). As a somewhat more quantitative example, we consider low-loss dispersive interactions for which the atom-cavity resonance frequencies are detuned by an amount δ . The effective rate χ for coherent atomic evolution then becomes $\chi = g^2/\delta$ for one intracavity photon and $\chi = 0$ for no intracavity photons, so that internal atomic evolution can be controlled by a single photon in the cavity. The requirements for the dominance of coherent evolution over dissipation are then stated as $(\chi = g^2/\delta)T \geq 1$ (i.e., the coherent interaction must do something non-negligible during the interaction time T) and $1/T \gg (\gamma, \kappa)$ (i.e., the time taken for the interaction of atom and cavity field must be small compared to the dissipative time scale), where T is the interaction time of atom and cavity. For a sequence of operations, necessarily then the clock-cycle time $T_0 \geq T$.

Independent of the cloak of quantum computation, important progress has been made in the area of (CQED) over the past decade, with experiments having been carried out in both the optical and microwave domains.[21] In the microwave domain (20–50 GHz), Rydberg atoms are employed in superconducting cavities of extremely high quality factor ($Q \sim 10^{10}$). The coupling frequency $g/2\pi \sim 10\text{kHz}$, so that the time scale for coherent evolution of two qubits is of order 10^{-5} sec. For circular Rydberg states $\gamma/2\pi \sim 5\text{Hz}$,

hence the ratio of damping time to that for resonant coherent evolution (for $\delta = 0$) is $g/(\gamma, \kappa) \sim 10^3 - 10^4$ (τ/T_0 as in Equation (6-1)), which seems quite promising. The interaction time T is set by the transit time of a thermal atom through the resonant cavity and is such that $g T > 1$, with however $(\chi = g^2/\delta)T < 1$. In principle, the interaction time T could be increased by employing laser-cooled atoms, but at the expense of fewer “ops” per decay time. Note that a system with simultaneously all the best attributes of large coupling and high Q and circular Rydberg states has not yet been operated, but that one should be coming “on-line” soon.

On a somewhat less positive note, experiments in the microwave require temperatures in the milliKelvin domain to eliminate thermal photons and involve cavities of centimeter scale cooled with a dilution refrigerator (or ^3He evaporative cooling). Hence, although research in the microwave domain offers interesting avenues for explorations with small numbers of qubits, it seems doubtful whether systems involving large numbers (e.g., tens much less thousands) of qubits would be possible in the foreseeable future.

In the optical domain, considerably larger values of the coupling frequency g have been obtained ($g/2\pi \sim 2 \times 10^7$ Hz) for interactions in small optical cavities of very high finesse $F \sim 10^5 - 10^6$ ($Q \sim 10^8$). The time scale for coherent evolution of two qubits is then of order 10^{-8} sec. For allowed dipole transitions in the optical domain, $\gamma/2\pi \sim 5 \times 10^6$ Hz, so that the ratios of damping times to that for coherent evolution ($\delta = 0$) are $g/\gamma \sim 4$ and $g/\kappa \sim 10^2$, which are certainly smaller than the corresponding quantities in the microwave domain. On the other hand, the product gT is roughly ten

times larger in the optical domain. Further, there are proposals for achieving substantial increases both in the magnitude of the coupling rate g as well as in the cavity Q by employing new types of optical cavities, with projected values of $g/\gamma \sim 40$ and $g/\kappa \sim 10^4 - 10^5$. [24] Although the value of g/γ achievable in the optical domain will probably never rival that obtained in the microwave, there are theoretical schemes that greatly mitigate the role of excited-state spontaneous emission through the use of hyperfine ground-state levels and so-called “dark-state” resonances that transfer population between ground states via resonant excited-state interactions but (somewhat surprisingly) without excited-state population. The effective atomic damping rate can thereby be greatly reduced.

Relative to microwave schemes (or indeed to condensed matter systems), optical cavity QED has the attractive features of the absence of a need for a cryogenic environment. The much smaller wavelength seems to afford reasonable opportunities for extensions to larger number of qubits. However, a daunting prospect would be the simultaneous operation of thousands of high finesse cavities and the control of atom trajectories or photon paths to interconnect the whole array. Fortunately, a very promising new scheme has been proposed by P. Zoller and colleagues at the University of Innsbruck in Austria. [25] Here a large number of atoms are placed at distinct sites to interact with the field of a single cavity. Individual atoms serve as qubits with information encoded in internal atomic states. Quantum registers for the computer are formed from this array of atoms, with the overall quantum state of the computer existing as an entangled state of the atom array. Computation proceeds by a series of one and two bit interactions between

the atoms. These interactions are accomplished via single photons in the intracavity field (with which all atoms interact) and an array of “classical” laser beams (which address selectively distinct pairs of atoms). Apart from employing a single cavity, this scheme has the attractive features of exploiting the aforementioned dark states (which greatly reduce the impact of atomic dissipation) and of populating the cavity mode with a photon only briefly during a transient period in the transfer of state information between atoms. Zoller’s group has made a detailed theoretical investigation of this system (including the impact of dissipation and the prospects for error correction) for realistic experimental parameters, with promising results both for gate operation (C-NOT) and for a “calculation” (a quantum Fourier transform with 5 atoms).

In this section, we have stressed CQED systems in the context of atomic physics. In addition to having achieved strong coupling, a promising aspect of these systems is that they have been shown experimentally to be faithful realizations of theoretical model systems, with defects in the correspondence being reasonably well understood. Field states exhibiting manifestly quantum or nonclassical photon statistics have been predicted theoretically and observed experimentally.[26, 27]

Beyond the setting of CQED in atomic physics, there are other promising systems in CQED. In the context of condensed matter physics, we mention in particular optical excitation of electron-hole pairs and interactions of these excitons.[28] Although there has been burgeoning progress in recent years with such systems (which are certainly promising for a variety of ap-

plications, including quantum logic), strong coupling (in the sense defined by Equation (6-3) with, for example, an appreciable nonlinear response for a single intracavity photon) has not been achieved to date, nor have the interactions been explicitly demonstrated to be sufficient to the task of generating nonclassical field states, which would seem to be necessary conditions for quantum logic at the single quantum level. Still, this is a rapidly developing area which warrants attention for its potential.

6.3 Trapped Ions

Atomic ions can be held in vacuum in a radio-frequency electric trap and cooled by laser fields to mK temperatures or less. A single cooled ion under typical conditions moves in an orbit smaller than $0.1 \mu\text{m}$ at the center of the trap, and can be confined to the lowest zero-point vibrational mode along one or more directions. In a recent achievement, a beryllium ion has been cooled to zero-point motion in all 3 spatial dimensions and the basic elements of a quantum logic gate have been demonstrated.[40] Groups of ions in traps of suitable geometries can be cooled to form a single linear array, with the spacing between adjacent ions determined by their mutual Coulomb repulsion. Cooling of such arrays of ions to the quantum limit of zero point motion has not yet been achieved, but is believed to be technically feasible in the near future.

Proven techniques exist for selectively inducing transitions among the vibrational states of the ions as well as among internal energy states of each

individual ion, and for determining with nearly 100 % efficiency which vibrational and internal state is occupied. The ions tend to remain in a selected state because they are trapped in a high vacuum free of perturbing atomic collisions and can be well shielded from stray external electric and magnetic fields.

Therefore the means exist, or should exist shortly, for studying arrays of ions in definite quantum states, and for creating quantum logic gates along the lines proposed by Cirac and Zoller.[9] Below, we will describe enough of the techniques of ion confinement, laser cooling, and ion state selection and detection to allow an assessment of the main issues and limitations of ion trapping for quantum computing.

It appears to us that the number of ions, with one qubit per ion, that can be successfully utilized in quantum computing will steadily increase with further research, perhaps reaching as many as 1000 ions with currently foreseeable technology. If such a level is reached, identifiable issues of vibrational mode stability and isolation of the ions from external heating, as well as general technical complexity, would likely make further progress much more difficult.

6.3.1 Ions in a linear trap

Of the many trap configurations that have been tested, we concentrate on the Paul linear trap, which is useful for confining an array of ions as

needed in a quantum computer. In this kind of trap, radio-frequency fields bind ions in the radial direction (perpendicular to the trap axis) and static fields prevent escape along the trap axis. The linear trap is formed by a group of four parallel conducting rods arranged symmetrically about the trap axis, which is denoted here by z . As shown in Figure 6-3, diagonally opposite rods are connected to a common potential, and an rf voltage is applied between adjacent rods, creating a time-varying potential in the space between the rods that can be approximated near the trap axis by an oscillating quadrupole potential:

$$V_{\text{rf}} = \frac{V_0}{2} \left[1 + \frac{x^2 - y^2}{a^2} \right] \sin \Omega t, \quad (6-7)$$

in which the x and y axes are oriented as shown in the figure and a is the shortest distance from the z axis to the inner surface of each rod.

To see how radial confinement takes place, note that the magnitude of the oscillating electric field in the xy plane increases linearly with distance from the origin, leading to a corresponding increase in the driven amplitude of oscillation of an ion (the so-called micromotion) as the ion moves further from the trap axis. The energy of this micromotion creates an effective potential energy that can bind the ion radially to the trap axis. In the usual regime of operation, the size of the micromotion is much smaller than that of the bound orbit, or equivalently, the orbit has a frequency of oscillation in the radial direction, ω_r , that is much less than Ω . In this limit, the effective two-dimensional radio-frequency potential energy that binds an ion of mass m is:

$$U_{\text{rf}} = \frac{e^2 V_0^2}{4m\Omega^2 a^4} (x^2 + y^2) = \frac{m}{2} \omega_r^2 r^2, \quad (6-8)$$

where x, y locate the position of the ion (or more precisely the center of the




Figure 6-3. Cross section of a linear quadrupole trap. An alternating rf voltage $V_0 \cos \Omega t$ is applied to a pair of diagonally opposite rods. The other pair of rods is maintained at rf ground.

small micromotion that is superposed on the orbital motion of the ion), and

$$\omega_r = \frac{eV_0}{\sqrt{2m\Omega a^2}}.$$

As a numerical example, a $^{136}\text{Ba}^+$ ion bound in a trap of size $\cong a = 200\mu\text{m}$ will have a well depth in the xy plane of 11 eV, and corresponding oscillation frequency in this plane of $\omega_r/2\pi = 3$ MHz, when an rf voltage of $V_0 = 250$ volts and frequency $\Omega/2\pi = 25$ MHz is applied between adjacent rods. Note that $\Omega \gg \omega_r$ as required. In normal operation, $\Omega > 5\omega_r$; imposing this condition leads to a convenient expression for the required rf voltage amplitude:

$$V_0 > (75 \text{ volts}) \left(\frac{a}{100\mu\text{m}} \right)^2 \left(\frac{\omega_r/2\pi}{5\text{MHz}} \right)^2 \left(\frac{m_{\text{ion}}}{100\text{amu}} \right). \quad (6-9)$$

To complete the linear trap the ions must also be confined to a region of the z axis. It is possible to break up each of the 4 rods into segments (keeping the rf potential along a given rod constant), and apply DC voltages between segments to create an electrostatic potential that varies along z , with a minimum at some point in the middle of the trap which we call $z = 0$. Near this minimum the electrostatic potential has approximately a quadrupole shape, leading to an electrostatic potential energy of the ion near the origin,

$$U_s = \frac{m}{2}\omega_z^2[z^2 - \frac{1}{2}(x^2 + y^2)], \quad (6-10)$$

where ω_z is the frequency of axial oscillations given by

$$\omega_z = (2eV_s/mz_0^2)^{\frac{1}{2}},$$

with V_s the on-axis electrostatic potential at $z = \pm z_0$ relative to $z = 0$. If $2z_0$ is the length of the middle segment of the trap, then V_s is approximately equal to the applied DC potential difference between trap segments.

Combining Equations (6-8) and (6-10), the complete 3-dimensional potential energy of an ion in the trap becomes:

$$U = U_{\text{rf}} + U_s = \frac{m}{2} [(\omega'_r)^2(x^2 + y^2) + \omega_z^2 z^2] \quad (6-11)$$

with

$$(\omega'_r)^2 = \omega_r^2 - \frac{1}{2}\omega_z^2.$$

Note the weakening of the effective potential well in the radial direction by the addition of the static potential.

6.3.2 Laser cooling and interrogation of trapped ions

A trapped ion can be cooled by driving an optical transition between ground and excited electronic states of the ion using a laser tuned somewhat below resonance with the transition frequency. Through the Doppler effect, there is a net loss of translational energy in each absorption and emission cycle, and the ion motion in the trap is cooled. The ultimate Doppler cooling limit is

$$k_B T > \hbar \Gamma / 2, \quad (6-12)$$

where $\Gamma/2\pi$ is the frequency width of the laser cooling transition, often the natural width of the excited state. A typical allowed-transition decay rate, $\Gamma \cong 10^8 \text{ sec}^{-1}$, thus implies a limiting temperature of about 10 mK. Γ can be

greatly reduced by using E2 (or M1) excitation to metastable states or using stimulated Raman transitions between two near-by sublevels of the ground state; in such ways, much lower ion temperatures have been achieved.

Because the ion motion in the trap is periodic, the Doppler effect (to first order in v/c) actually produces an unshifted carrier laser frequency together with sidebands separated by the frequency of ion oscillations (ω_r' and/or ω_z) along the light propagation axis. When the ion orbital size b becomes small compared to the optical wavelength λ , the so-called Lamb-Dicke regime, most of the power is in the carrier; the nearest pair of sidebands has relative amplitude $\cong 2\pi b/\lambda$. Further laser cooling can still be accomplished by tuning the laser to the lower Doppler sideband.

In a few experiments, a single trapped ion has been cooled until at least one vibrational mode is in the lowest quantum state, and recently a ${}^9\text{Be}^+$ ion was cooled to the quantum limit in all 3 dimensions. In the latter case, cooling was accomplished by a Raman transition between the two hyperfine sub-levels of the electronic ground state.

Arrays of 30 ions or more have been trapped in a linear Paul trap, cooled, and observed to form a linear crystal chain along the trap axis. It appears to be a challenging but straightforward extension of single-ion techniques to cool such an array of ions to the zero point vibrational limit.

Individual ions are detected by their scattering of laser radiation. Usually what is observed is laser fluorescence on the same electronic transition as is used for cooling the ion. This fluorescence can also be used to determine

the internal state of the ion with nearly 100% efficiency, by employing the technique of 'shelving'. A shelving laser beam excites the ion to a metastable or 'dark' state with a probability depending upon which initial state the ion is in. The ion is then illuminated for fluorescence, and the presence or absence of fluorescence determines whether the ion was initially in the state that allowed it to be shelved to the metastable level.

6.3.3 Requirements for quantum computing

As discussed by Cirac and Zoller, a linear array of trapped ions offers some interesting possibilities related to logic gates and quantum computing. N ions would give us N qubits, a qubit being identified with a pair of long-lived internal atomic ion levels such as two hfs sub-levels (or Zeeman sub-levels) in the ground electronic state. (A suitable pair might also be the ground state and an excited metastable state, viz. the 6S and 5D states of Ba^+ , though this pair is subject to extra perturbations from the quadrupole trapping fields.) A laser beam could transfer an ion from one internal level to another, using stimulated Raman transitions if the qubits are a pair of ground state sub-levels. Thus the ions need to be separated by only a few optical wavelengths in order to address different qubits independently, and as mentioned later, even smaller separations might be permissible. The vibrational coupling among the ions serves to link the qubits of different ions, and to allow entangled quantum product states among the qubits to be created. Cirac and Zoller illustrate the possibilities using the lowest frequency (center of mass) longitudinal mode of vibration of a string of ions.

In order for the Cirac-Zoller scheme to work, the ions must be placed in nearly pure electronic and vibrational states, and remain free of significant dissipative external perturbations over a complete quantum computation cycle. The situation regarding the internal atomic ion states is quite favorable. The coherence lifetimes can be several seconds or longer, while optical transitions between these states can be driven rapidly. The cycling (Rabi) frequency, even on forbidden or Raman transitions, can be $\Omega_{Rabi} \cong 2\pi \times 10^7 \text{ rad/sec}$ using only modest laser power focused on a single ion. Of course the laser power must be controlled quite accurately to induce precise $\pi/2$, 2π , etc. pulses between states. The Cirac-Zoller proposal also requires that the laser beam acting on any one of the ions drive a vibrational transition of the entire linear array of ions. This coupling varies as $N^{-\frac{1}{2}}$, but the reduction in Rabi frequency for a given laser power should not cause major problems.

Other than the sheer technical complexity of implementing the Cirac-Zoller scheme with a large number of ions, the most serious issues seem to be connected with the vibrational states of the ions, namely: 1. vibrational mode-stability, 2. cooling to the vibrational quantum limit, and 3. outside heating and dephasing of the vibrational modes; each of these issues we now discuss.

1. Vibrational mode instabilities must be prevented.

To remain in a region of linear stability and prevent the onset of zigzag modes, the ions in a linear trap must be confined much more strongly radially

than axially. The onset of instability has been studied in numerical simulations [see J. P. Schiffer, Phys. Rev. Lett. **70** 818 (1993)]. Analytical results in agreement with the numerical ones have been obtained by Dan Dubin of UCSD [D. Dubin, Phys. Rev. Lett. **71** 2753 (1993)] for a bound Coulomb chain similar to the chain of ions in a linear Paul trap. For a linear chain of N ions, $N \gg 1$, Prof. Dubin finds that Δz , the mean spacing between ions, is given by

$$(\Delta z)^3 = \left(\frac{6e^2}{\pi \epsilon_0 m \omega_z^2} \right) \left(\frac{\log_e N}{N^2} \right), \quad (6-13)$$

and that the region of stability against zigzag modes occurs for

$$\left(\frac{\omega'_r}{\omega_z} \right)^2 > 0.1 \frac{N^2}{\log_e N}. \quad (6-14)$$

Combining this last equation with Equation (6-9), we find the following requirement on the rf trapping voltage amplitude:

$$V_0 > (430 \text{ volts}) \left(\frac{a}{100 \mu\text{m}} \right)^2 \left(\frac{\omega_z/2\pi}{100 \text{kHz}} \right)^2 \left(\frac{m_{\text{ion}}}{100 \text{amu}} \right) \left(\frac{N}{1000 \text{ ions}} \right)^2 \left(\frac{\log_e 1000}{\log_e N} \right). \quad (6-15)$$

It thus appears possible using a trap of reasonable dimensions to confine up to 1000 ions, or perhaps a few thousand, in a stable linear chain having an axial center-of-mass vibrational frequency as high as 100 kHz. However, V_0 grows as N^2 and quickly becomes prohibitively large as N is increased. The spacing between ions is also an issue. It follows from Equation (6-13) that even light ions such as ${}^9\text{Be}^+$ would be spaced by only $\Delta z = 1.5 \mu\text{m}$ in a string of 1000 ions, and thus barely resolvable optically. Although it is possible to distinguish more closely spaced ions by using field gradients to shift resonance frequencies of adjacent ions, or to redesign the trap to spread the ions out, the ion spacing clearly becomes one more difficulty when increasing N beyond a few thousand.

2. The vibrational motion must be cooled to the quantum limit.

Thus the axial motion must be cooled below $k_B T_{ions} = \hbar\omega_z$, or below $T_{ions} = 10\mu\text{K}$ for an axial frequency of 100 kHz. Thus far, cooling to the quantum limit has been accomplished only at higher vibrational frequencies and temperatures, 2 MHz and 50 μK in the case of 2D cooling of a single ^{199}Hg ion, and $\cong 10$ MHz for 3D cooling of a ^9Be ion. Achieving the quantum limit at a frequency as low as 100 kHz, with an array of hundreds or thousands of ions, presents a truly formidable technical challenge, but one that may well be met by normal evolution of current techniques.

3. The vibrational modes must be free of dissipative coupling or outside heating for the duration of a complete quantum computation cycle.

Perhaps representative of the degree of isolation attained in current experiments was the observed heating rate, $\Gamma_{heat} \cong 10^3$ vibrational quanta per second, when ^9Be was cooled to the quantum limit. Improvement in isolation by several orders of magnitude should be possible before approaching any fundamental limitations posed by coupling to the trapping electrodes. To get an idea of the isolation needed in quantum computing, we return to Equation (4-1) and note that for J qubits and M ops, the internal ion (qubit) states must be switched at a frequency $\Omega_{Rabi} > JM\Gamma_{heat}$. But Ω_{Rabi} cannot be made arbitrarily large, because it is necessary that $\Omega_{Rabi} \ll \omega_z$ to prevent mixing of the wrong vibrational levels and also to prevent energy shifts comparable to the vibrational splitting. Thus we require that $\omega_z \gg JM\Gamma_{heat}$. As before, J and M will be functions of the algorithm employed. Assuming

$J \sim N$ and $M \sim N^3$, we find that to use 1000 ions with an axial frequency of 100 kHz means that Γ_{heat} would have to be much less than 10^{-7} quanta per second.

Perhaps the ultimate limit on the number of ions that could be used in this proposed method of quantum computing will depend upon the dissipation rate Γ_{heat} . If this rate can be made very slow, and lower ion temperatures can be reached, then ω_z can be reduced to ease the restriction on N in Equations (6-13) and (6-15). In that case Equation (6-15), with ω_z^2 replaced by $J^2 M^2 \Gamma_{heat}^2$, would still offer a useful way of estimating the required rf trapping voltage.

7 RECOMMENDATIONS

Quantum computation represents a *potentially* profound development with *possibly* far-reaching implications in both the physical and information sciences. Unfortunately any definite assessment without qualifiers is essentially impossible because of the dearth of significant results beyond the brilliant quantum algorithms that Shor presented 1.5 years ago. While proponents of the field believe that the marriage of quantum mechanics and information science is ripe with potential for prodigious progeny, this belief currently rests on a very small set of truly significant results. If quantum computation is to change the future of computation in a major way, then there must emerge a much broader class of applications than are known today. Without a more wide ranging set of possibilities, the impetus for conquering the daunting issues associated with physical implementation is to large measure lost.

Apart from the actual implementation of quantum computation, the potential impact of this new paradigm on classical computation should not be overlooked. It may be that new perspectives offered by the investigation of quantum algorithms might well lead to more powerful classical algorithms. Unlike the physical implementation of quantum computation which is at best a long-term endeavor, such insight could more or less immediately advance the state of the art of classical computation.

Within the context of the this outlook, our recommendations for possible

ARPA support of research in the area of quantum computation are as follows.

1. Establish a research program to investigate possibilities for quantum computation beyond Shor's algorithms. Here we have in mind the fostering of a fairly intense effort over the coming years to understand the types of problems for which quantum computation is well suited and whether or not new insights do indeed arise for developing more powerful classical algorithms. The two principal communities involved would probably be those of theoretical physics and computer science (but not to the exclusion of other groups). Clearly, as new quantum algorithms are developed, it will be essential to address the issue of error correction as well.
2. Seed research in various communities for quantitative minimization of algorithmic complexity and optimum circuit design. Given the extreme value of qubits of information in terms of the degree of difficulty of physical realization, it is quite important to have explicit quantum circuits with quantitative measures of resource requirements (beyond simply "a polynomial of order k ") in order to bridge the gulf between abstract quantum algorithms and actual physical implementations. The JASON quantum circuits provide an important step in this direction, even if they prove to be less than optimum in the ultimate conservation of qubits and ops.
3. Supplement ongoing experimental research related to the isolation and control of discrete quantum systems suitable for quantum logic. Here the research objectives might not be directly quantum computation,

but might be instead fundamental aspects thereof such as the investigation of quantum dynamics in nontrivial Hilbert spaces (e.g., the generation of quantum-state entanglement for more than two qubits and the role of dissipation). Given the tremendous gulf between laboratory capability and the requirements for a nontrivial implementation of Shor's algorithm, we would specifically advise against a program of "prototype development", but would rather supplement research on diverse fronts with modest goals rather disconnected from a "Holy Grail" pursuit of quantum computation.

Overall, we feel that the most pressing need with also the greatest potential is for a broad theoretical exploration for opportunities beyond Shor's algorithms. If the "well" proves to be "dry", so be it. On the experimental front, we do not believe that there is a similarly pressing need for ARPA involvement. Although there are a variety of promising systems, the most optimistic near-term hope would be only extremely modest "proof-of-principle" demonstrations from which we doubt any profound new insights would emerge. Therefore, we recommend that ARPA should not shoulder the principal burden for funding of experimental efforts, which should not in any case be justified solely for their relevance to quantum computation. However, ARPA could play an important role in ensuring that the experimental and theoretical communities remain engaged. Should an explosion of possibilities ensue from the theoretical investigations, then it may well be worthwhile to consider increasing the investment on the experimental front, bearing in mind that the time horizon for the experimental realization of quantum computation will still be distant.

Finally, we would urge the adoption of a broad-minded view for opportunities other than those related to large-scale computation, such as might arise in quantum cryptography or coherent nanoscale electronics. This is a long-term endeavor of potentially profound significance where surprises are likely to emerge on diverse fronts.

References

- [1] Benioff, P. Phys. Rev. Lett. **48**, 1581 (1982); Feynman, R. P., Found Phys. **16**, 507 (1986); D. Deutsch, Proc. Roy. Soc. A **400**, 97 (1985).
- [2] For an excellent review of quantum computing see A. Ekert and R. Jozsu, "Review of Modern Physics" (to be published, 1995).
- [3] Shor, P. W. in *Proceedings of the 35th Annual Symposium on FOCS*, edited by S. Goldwasser (IEEE Computer Society Press, New York, 1994).
- [4] Bennett, C. IBM J. Res. Develop. (Nov. 1973), 525.
- [5] Feynman, R. P. Int. J. Theor. Phys. **21**, 467 (1982)
- [6] Coppersmith, D. "An Approximate Fourier Transform Useful in Quantum Factoring," IBM Research Report No. RC19642 (1994)
- [7] Feynman, R. P. Int. J. Theor. Phys. **21**, 467(1982)
- [8] S. Lloyd, Science **261**, 1569 (1993).
- [9] Cirac, J. I. and P. Zoller, *Quantum Computations with Cold Trapped Ions*, Phy. Rev. Ltrs. **74**, 20, 15 May 1995
- [10] Beckman, David, A. N. Chari, S. Devabhaktuni, and J. Preskill *Machine Language for Modular Exponentiation on a Quantum Computer*, unpublished report, Div. Phys., CIT, 20 August 1995
- [11] Barenco, A., and A. Ekert, private communication and to be published (1995).

- [12] Altshuler, B. L., P. A. Lee, and R. A. Webb, eds., *Mesoscopic Phenomena in Solids* (North-Holland, New York, 1991).
- [13] Ando, T., A.B. Fowler, and F. Stem, *Rev. Mod. Phys.* 54, 437 (1982).
- [14] Beenakker, C.W.J., and H. van Houten, in *Solid State Physics* 44, Ed. H. Ehrenreich.
- [15] Turnbull, D., (Academic Press, San Diego, 1991).
- [16] Grabert, H., and M. H. Devoret, *Single Charge Tunneling - Coulomb Blockade Phenomena in Nanostructures* (Plenum Press, New York, 1992).
- [17] Gutzwiller, M. C., *Chaos in Classical and Quantum Mechanics* (Springer Verlag, New York, 1991).
- [18] Tinkham, M., *Introduction to Superconductivity* (McGraw Hill, New York, 1995).
- [19] Tuominen, M.T., J.M. Hergenrother, T.S. Tighe, and M. Tinkham, *Phys. Rev. Lett.* 69, 1997 (1992).
- [20] Waugh, F.R., R.M. Westervelt, K. Campman and A.C. Gossard, *Phys. Rev. Lett.* 75, 705 (1995).
- [21] For a recent review of CQED, see *Advances in Atomic, Molecular, and Optical Physics*, Suppl.2, Berman, P., ed. (Academic Press, New York, 1991), p.57
- [22] Barenco, A. et al., *Phys. Rev. Lett.*, **74**, 4083 (1995); T. Sleator and H. Weinfurter, *ibid.*, 4087; J. I. Cirac and P. Zoller, *ibid.*, 4091.

- [23] Turchette, Q. A., C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble, submitted to *Phys. Rev. Lett.* **75**, 4710 (1995).
- [24] Braginsky, V. B. Gorodetsky, M. L. and Ilchenko, V. S., *Phys. Lett. A* **137**, 393 (1989).
H. Mabuchi, and H. J. Kimble, *Opt. Lett.* **19**, 749 (1994).
V. B. Braginsky, M. L. Gorodetsky, and V. S. Ilchenko, in *Laser Optics '93-Proceedings of the S.P.I.E.* (1994)
L. Collot, V. Lefevre-Sequin, M. Brune, J. M. Raimond, and S. Haroche, *Europhys. Lett* **23** 327 (1993).
- [25] Pellizzari, T., S. A. Gardiner, J. I. Cirac, and P. Zoller, submitted to *Phys. Rev. Lett.* (1995).
- [26] Rempe, G., R. J. Thompson, R. J. Brecha, W. D. Lee and H. J. Kimble, *Phys. Rev. Lett.* **67** 1727 (1991).
- [27] Rempe, G., F. Schmidt-Kaler, and H. Walther, *Phys. Rev. Lett.* **64**, 2783 (1990).
- [28] Cao, H., Klomowitch, G. Bjork, and Y. Yamamoto, *Phys. Rev. Lett.* **75** 1146 (1995); I. Chang and Y. Yamamoto, *Phys. Rev. A* (1995).
- [29] Landauer, R., *Proc. Roy Soc. Lond* (1995).
- [30] For a recent review, see Howard Carmichael, *An Open Systems Approach to Quantum Optics*, [Lecture Notes in Physics, m18], Springer Verlag, Berlin (1993).
- [31] Calderia, A. O. and A. J. Leggett, *Phys. Rev. A* **31**, 1059 (1985).

DISTRIBUTION LIST

Director of Space and SDI Programs
SAF/AQSC
1060 Air Force Pentagon
Washington, DC 20330-1060

CMDR & Program Executive Officer
U S Army/CSSD-ZA
Strategic Defense Command
PO Box 15280
Arlington, VA 22215-0150

A R P A Library
3701 North Fairfax Drive
Arlington, VA 22209-2308

Dr Arthur E Bisson
Director
Technology Directorate
Office of Naval Research
Room 407
800 N. Quincy Street
Arlington, VA 20350-1000

Dr Albert Brandenstein
Chief Scientist
Office of Nat'l Drug Control Policy
Executive Office of the President
Washington, DC 20500

Mr. Edward Brown
Assistant Director
ARPA/SISTO
3701 North Fairfax Drive
Arlington, VA 22203

Dr H Lee Buchanan, I I I
Director
ARPA/DSO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr Curtis G Callan, Jr
Princeton University
P O Box 708
Princeton, NJ 08540

Dr Ashton B Carter
Nuclear Security & Counter Proliferation
Office of the Secretary of Defense
The Pentagon, Room 4E821
Washington, DC 20301-2600

Dr Kenneth M Case
1429 Calle Altura
La Jolla, CA 92037

Dr Collier
Chief Scientist
U S Army Strategic Defense Command
PO Box 15280
Arlington, VA 22215-0280

DTIC [2]
Cameron Station
Alexandria, VA 22314

Mr John Darrah
Senior Scientist and Technical Advisor
HQAF SPACOM/CN
Peterson AFB, CO 80914-5001

Dr Victor Demarines, Jr.
President and Chief Exec Officer
The MITRE Corporation
202 Burlington Road
A210
Bedford, MA 01730-1420

Dr Alvin M Despain
3273 Corinth Avenue
Los Angeles, CA 90066

Mr Dan Flynn [5]
OSWR
Washington, DC 20505

Dr Paris Genalis
Deputy Director
OUSD(A&T)/S&TS/NW
The Pentagon, Room 3D1048
Washington, DC 20301

DISTRIBUTION LIST

Dr Lawrence K. Gershwin
NIC/NIO/S&T
7E47, OHB
Washington, DC 20505

Dr Jeremy Goodman
Princeton University Observatory
Peyton Hall
Princeton, NJ 08544

Mr. Thomas H Handel
Office of Naval Intelligence
The Pentagon, Room 5D660
Washington, DC 20350-2000

Dr Helmut Hellwig
Deputy Asst. Secretary
Office SAF/AQR
Science, Tech & Engineering
1919 S. Eads Street, Suite 100
Arlington, VA 22202-3053

Dr Robert G Henderson
Director
JASON Program Office
The MITRE Corporation
7525 Colshire Drive
Mailstop Z561
McLean, VA 22102

Dr William E Howard III [2]
Director of Advanced Concepts &
Systems Design
The Pentagon Room 3E480
Washington, DC 20301-0103

Dr Gerald J Iafrate
U S Army Research Office
PO Box 12211
4330 South Miami Boulevard
Research Triangle NC 27709-2211

J A S O N Library [5]
The MITRE Corporation
Mail Stop W002
7525 Colshire Drive
McLean, VA 22102

Dr Anita Jones
Department of Defense
DOD, DDR&E
The Pentagon, Room 3E1014
Washington, DC 20301

Mr. O' Dean P. Judd
Los Alamos National Laboratory
Mailstop F650
Los Alamos, NM 87545

Dr Bobby R Junker
Office of Naval Research
Code 111
800 North Quincy Street
Arlington, VA 22217

Dr H Jeff Kimble
CA Institute of Technology
12-33
Norman Bridge Laboratory of Physics
Pasadena, CA 91125

Dr Steven E Koonin
California Institute of Technology
Vice President and Provost
206-31
Pasadena, CA 91125

Dr Ken Kress
Office of Research and Development
809 Ames
Washington, DC 20505

Lt Gen, Howard W. Leaf, (Retired)
Director, Test and Evaluation
HQ USAF/TE
1650 Air Force Pentagon
Washington, DC 20330-1650

DISTRIBUTION LIST

Dr Herbert Levine
University of California/San Diego
Department of Physics
Mayer Hall, B019
La Jolla, CA 92093

Dr Nathan Lewis
California Institute of Technology
Division of Chemistry and
Chemical Engineering: 127-72
Pasadena, CA 91125

Mr. Larry Lynn
Director
ARPA/DIRO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr. John Lyons
Director of Corporate Laboratory
US Army Laboratory Command
2800 Powder Mill Road
Adelphi, MD 20783-1145

Col Ed Mahen
ARPA/DIRO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr. Arthur Manfredi
OSWR
Washington, DC 20505

Dr George Mayer
Office of Director of Defense
Reserach and Engineering
Pentagon, Room 3D375
Washington, DC 20301-3030

Dr Bill Murphy
ORD
Washington, DC 20505

Mr Ronald Murphy
ARPA/ASTO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr Julian C Nall
Institute for Defense Analyses
1801 North Beauregard Street
Alexandria, VA 22311

Dr Ari Patrinos
Director
Environmental Sciences Division
ER74/GTN
US Department of Energy
Washington, DC 20585

Dr Bruce Pierce
USD(A)D S
The Pentagon, Room 3D136
Washington, DC 20301-3090

Dr William Press
Harvard College Observatory
Center for Astrophysics MS-51
60 Garden Street
Cambridge, MA 02138

Mr John Rausch [2]
Division Head 06 Department
NAVOPINTCEN
4301 Suitland Road
Washington, DC 20390

Records Resource
The MITRE Corporation
Mailstop W115
7525 Colshire Drive
McLean, VA 22102

Dr Victor H Reis
US Department of Energy
DP-1, Room 4A019
1000 Independence Ave, SW
Washington, DC 20585

DISTRIBUTION LIST

Dr Oscar Rothaus
Cornell University
Math Department
Ithaca, NY 12853

Dr Fred E Saalfeld
Director
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217-5000

Dr Dan Schuresko
O/DDS&T
Washington, DC 20505

Dr John Schuster
Technical Director of Submarine
and SSBN Security Program
Department of the Navy OP-02T
The Pentagon Room 4D534
Washington, DC 20350-2000

Dr Michael A Stroschio
US Army Research Office
P. O. Box 12211
Research Triangle NC 27709-2211

Superintendent
Code 1424
Attn Documents Librarian
Naval Postgraduate School
Monterey, CA 93943

Ambassador James Sweeney
Chief Science Advisor
USACDA
320 21st Street NW
Washington, DC 20451

Dr George W Ullrich [3]
Deputy Director
Defense Nuclear Agency
6801 Telegraph Road
Alexandria, VA 22310

Dr Walter N Warnick [25]
Deputy Director
Office of Planning & Analysis, ER-5.1
Office of Energy Research
U S Department of Energy
Germantown, MD 2074

Dr Peter J Weinberger
22 Clinton Avenue
Maplewood, NJ 07040

Dr Robert M Westervelt
Harvard University
Division of Applied Sciences
Boston, MA 02138

Dr Edward C Whitman
Dep Assistant Secretary of the Navy
C3I Electronic Warfare & Space
Department of the Navy
The Pentagon 4D745
Washington, DC 20350-5000

Capt H. A. Williams, U S N
Director Undersea Warfare Space
& Naval Warfare Sys Cmd
PD80
2451 Crystal Drive
Arlington, VA 22245-5200