Algorithms for Quantum Computation: Discrete Log and Factoring

Extended Abstract

Peter W. Shor
AT&T Bell Labs
Room 2D-149
600 Mountain Ave.
Murray Hill, NJ 07974 USA

email: shor@research.att.com

Abstract

This paper gives algorithms for the discrete log and the factoring problems that take random polynomial time on a quantum computer (thus giving the first examples of quantum cryptanalysis).

1 Introduction

Since the discovery of quantum mechanics, people have found the behavior of the laws of probability in quantum mechanics counterintuitive. Because of this behavior, quantum mechanical phenomena act quite differently than the phenomena of classical physics that we are used to. Feynman seems to have been the first to ask what effect this has on computation [Fey]. He gave arguments as to why this behavior might make it intrinsically computationally expensive to simulate quantum mechanics on a classical (Von Neumann) computer. He also mentioned the possibility of using a computer based on quantum mechanical principles to avoid this problem, thus implicitly asking the converse question: by using quantum mechanics in a computer can you compute more efficiently than on a classical computer. Other early work in the field of quantum mechanics and computing was done by Benioff [Beni]. Although he did not ask whether quantum mechanics conferred extra power to computation, he did show that a Turing machine could be simulated by the unitary evolution of a quantum process, which is a necessary prerequisite for quantum computation. Deutsch [Deu1, Deu2] was the first to give an explicit model of quantum computation. He defined both quantum Turing machines and quantum circuits and investigated some of their properties.

This question of whether using quantum mechanics in a computer allows one to obtain more computational power has not yet been satisfactorily answered. The question was addressed in [DJ, BB1, BB2], but it was not shown how to solve any problem in quantum polynomial time that was not solvable in BPP (the class of problems which can be solved in polynomial time with a bounded probability of error). Recent work on this problem was stimulated by Bernstein and Vazirani's seminal paper [BV], which contained two results. One of these results was an oracle problem which can be done in polynomial time on a quantum Turing machine and requires super-polynomial time on a classical computer; this was the first indication (other than the fact that nobody knew how to simulate a quantum computer on a classical computer without an exponential slowdown) that quantum computation might obtain a greater than polynomial speedup over classical computation augmented by randomization. This result was improved by Simon [Sim], who gave a much simpler construction of an oracle problem which takes polynomial time on a quantum computer and requires exponential time on a classical computer. Indeed, by viewing Simon's oracle as a subroutine, this result becomes a promise problem which takes polynomial time on a quantum computer and looks as if it might be very difficult on a classical computer. The algorithm for the "easy case" of discrete log given in this paper is directly analogous to Simon's algorithm, except with the group $Z_2{}^n$ replaced by the group Z_{p-1}; I was only able to discover this algorithm after seeing Simon's paper.

In the other part of Bernstein and Vazirani's paper, a particular class of quantum Turing machine was rigorously defined and a universal quantum Turing machine was given which could simulate any other quantum Turing machine of this class. Unfortunately, it was not clear whether these quantum Turing machines could simulate other classes of quantum Turing machines, so this result was not entirely satisfactory. Yao [Yao] has remedied the situation by showing that quantum Turing machines can simulate, and be simulated by, uniform families of polynomial size quantum circuits, with at most polynomial slowdown. He has further defined quantum Turing machines with k heads and showed that these machines can be simulated with slowdown of a factor of 2^k. This seems to show that the class of problems which can be solved in polynomial time on one of these machines, possibly with a bounded probability of error, is reasonably robust. This class is called BQP in analogy to the classical complexity class BPP, and could be considered the class of problems that are efficiently solvable on a quantum Turing machine.

Since Quantum P is contained in $P^{\#P}$ [BV], any non-relativized proof that Quantum P is strictly larger than P would imply a structural complexity result that is not yet proven. In view of this difficulty, two approaches come to mind; one is showing that Quantum P $\subseteq$ P would lead to a collapse of classical complexity classes which are believed to be different: Showing that NP $\subseteq$ Quantum P, for example, would show that if Quantum P $\subseteq$ P, then the polynomial hierarchy would collapse. A second approach, which we take, is solve in Quantum P some well-studied problem for which no polynomial time algorithm is known. This would show that the extra power conferred by quantum interference is at least hard to achieve using classical computation.

The discrete log and factoring problems are two number-theory problems which have

been studied extensively but for which no polynomial-time algorithms are known. In fact, these problems are so widely believed to be hard that cryptosystems based on their hardness have been proposed, and the RSA public key cryptosystem [RSA], based on the hardness of factoring, is in use. We show that these problems can be solved in random quantum polynomial time, or RQP, which is the class of problems which run in polynomial time on a quantum computer, but are allowed a small probability of (one-sided) error.

Currently, nobody knows how to build a quantum computer, although it seems as though it should be possible within the laws of quantum mechanics. It is hoped that this paper will stimulate research on whether it is feasible to actually construct one.

Even if no quantum computer is ever built, this research does illuminate the problem of simulating quantum mechanics on a classical computer. Any method of doing this for an arbitrary Hamiltonian would necessarily be able to simulate a quantum computer. Thus, any general method for simulating quantum mechanics with at most a polynomial slowdown would lead to a polynomial algorithm for factoring.

2 Quantum Computation

In this section we will give a brief introduction to quantum computation, emphasizing the properties that we will use. For a more complete overview I refer the reader to [BV, Yao].

In quantum physics, an experiment behaves as if it proceeds down all possible paths simultaneously. Each of the paths has a complex amplitude. The probability of any particular outcome of the experiment is then proportional to the square of the absolute value of the sum of the amplitude of all the paths leading to that outcome. A quantum computer behaves in much the same way. The computation proceeds down all possible paths at once, and each path has associated with it a complex amplitude. To determine the probability of any final state of the machine, we add the amplitudes of all the paths which reach that final state, and then square the absolute value of this sum.

An equivalent way of looking at this process is to imagine that the machine is in some *superposition of states* at every step of the computation. We will represent this superposition of states as

$$\sum_i a_i \left|S_i\right\rangle$$

where the amplitudes a_i are complex numbers such that $\sum_i |a_i|^2 = 1$ and each $|S_i\rangle$ is a state of the machine; in a quantum Turing machine, a state is defined by what is written on the tape at that step and by the position and state of the head. In a quantum circuit a state is the values of the signals on all the wires at some level of the circuit. If the machine is examined at a particular step, the probability of seeing state $|S_i\rangle$ is $|a_i|^2$; however, by the Heisenberg uncertainty principle, looking at the machine interferes with the rest of the computation. The laws of quantum mechanics only permit unitary transformations of the state. A unitary matrix is one whose conjugate transpose is equal to its inverse, and requiring unitary matrices ensures that the probabilities of obtaining each possible result will add up to one. Further, the definitions of quantum Turing machine and quantum circuit

only allow local unitary transformations, that is, unitary transformations on a fixed number of bits.

Perhaps an example will be informative at this point. Suppose our machine is in the superposition of states

$$\tfrac{i}{\sqrt{2}}\,|000\rangle + \tfrac{1}{2}\,|100\rangle - \tfrac{1}{2}\,|110\rangle$$

and we apply the unitary transformation

	00	01	10	11
00	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
01	$\frac{1}{2}$	$\frac{i}{2}$	$-\frac{1}{2}$	$-\frac{i}{2}$
10	$\frac{1}{2}$	$-\frac{1}{2}$	$\frac{1}{2}$	$-\frac{1}{2}$
11	$\frac{1}{2}$	$-\frac{i}{2}$	$-\frac{1}{2}$	$\frac{i}{2}$

to the last two bits of our state. Then the machine will then go to the superposition of states

$$\tfrac{i}{2\sqrt{2}}\,|000\rangle + \tfrac{i}{2\sqrt{2}}\,|001\rangle + \tfrac{i}{2\sqrt{2}}\,|010\rangle + \tfrac{i}{2\sqrt{2}}\,|011\rangle + \tfrac{1}{2}\,|101\rangle + \tfrac{1}{2}\,|111\rangle\,.$$

Notice that the result would have been different if we had started with the superposition of states

$$\tfrac{i}{\sqrt{2}}\,|000\rangle + \tfrac{1}{2}\,|100\rangle + \tfrac{1}{2}\,|110\rangle$$

which has the same probabilities of being in any particular configuration if we observe it.

We now give certain properties of quantum computation that will be useful. These facts are not apparent from the definition of quantum Turing machine or quantum circuit, and they are very useful for constructing algorithms for quantum machines.

Fact 1: A deterministic computation is performable on a quantum computer if and only if it is reversible [BV]. From results on reversible computation [Benn, BV], this means that we can compute any polynomial time function $f(a)$ as long as we keep the input, a, on the machine. To erase a and replace it with $f(a)$ we need in addition that f is one-to-one and that a is computable in polynomial time from $f(a)$; i.e., that both f and f^{-1} are polynomial.

Fact 2: Any polynomial size unitary matrix can be approximated using a polynomial number of elementary unitary transformations [Deu2, BV, Yao] and thus can be approximated in polynomial time on a quantum computer. Further, this approximation is good enough so as to introduce at most a bounded probability of error into the results of the computation.

3 Building Unitary Transformations

Since quantum computation deals with unitary transformations, it is helpful to be able to build certain useful unitary transformations. In this section we give some techniques

for constructing unitary transformations on quantum machines, which will result in our showing how to construct one particular unitary transformation in polynomial time. These transformations will generally be given as matrices, with both rows and columns indexed by states. These states will correspond to representations of integers on the computer; in particular, the rows and columns will be indexed beginning with 0 unless otherwise specified.

A tool we will use repeatedly in this paper is the following unitary transformation. Consider a number a with $0 \leq a < q$ for some q where the number of bits of q is polynomial. We will transform it into a number b, $0 \leq b < q$, with amplitude $\frac{1}{q^{1/2}} \exp(2\pi iab/q)$. This transformation is at the heart of our algorithms, and we will call the associated unitary matrix A_q. Since we use it for q of exponential size, we must show how this transformation can be done in polynomial time. In fact, we will only be able to do this for *smooth* numbers q, that is, ones with small prime factors. For the purposes of this paper, a *smooth* number q will be one that contains no prime power factor that is larger than $(\log q)^c$ for some fixed c.

If we know a factorization $q = q_1 q_2 q_3 \cdots q_k$ where $\gcd(q_i, q_j) = 1$ and where k and all of the q_i are of polynomial size we will show how to build the transformation A_q in polynomial time. For this, we first need a lemma on quantum computation.

Lemma 3.1 *Suppose the matrix B is a block-diagonal $mn \times mn$ unitary matrix composed of n identical unitary $m \times m$ matrices B' along the diagonal and 0's everywhere else. Suppose further that the state transformation B' can be done in time $T(B')$ on a quantum Turing machine. Then the matrix B can be done in $T(B') + (\log mn)^c$ time on a quantum Turing machine, where c is a constant.*

We will call this matrix B the direct sum of n copies of B' and use the notation $B = \bigoplus_n B'$.

Proof: Suppose that we have a number a on our tape. We can reversibly compute α_1 and α_2 from a where $a = m\alpha_1 + \alpha_2$. This computation erases a from our tape and replaces it with α_1 and α_2. Now α_1 tells in which block the row a is contained, and α_2 tells which row of the matrix within that block is the row a. We can then apply B' to α_2 to obtain β_2 (erasing α_2 in the process). Now, combining α_1 and β_2 to obtain $b = m\alpha_1 + \beta_2$ gives the result of B applied to A (with the right amplitudes). The computation of B' takes $T(B')$ time, and the rest of the computation is polynomial in $\log m + \log n$. ∎

We now show how to obtain A_q for smooth q. We will decompose A_q into a product of a polynomial number of unitary transformations, all of which are performable in polynomial time; this enables us to construct A_q in polynomial time. Suppose that we have $q = q_1 q_2$ with $\gcd(q_1, q_2) = 1$. What we will do is represent $A_q = CD$, where by rearranging the rows and columns of D we obtain $\bigoplus_{q_1} A_{q_2}$ and rearranging the rows and columns of C we obtain $\bigoplus_{q_2} A_{q_1}$. As long as these rearrangements of the rows and columns of C and D are performable in polynomial time (i.e., given row r, we can find in polynomial time row r' to which it is taken) and the inverse operations are also performable in polynomial time, then by using the lemma above and recursion we can obtain a polynomial-time way to perform A_q on a quantum computer.

We now need to define C and D and check that $A_q = CD$. To define C and D we need some preliminary definitions. Recall that $q = q_1 q_2$ with q_1 and q_2 relatively prime. Let $\omega = \exp(2\pi i/q)$. Let u be the number (mod q) such that $u \equiv 0 \pmod{q_1}$ and $u \equiv -1 \pmod{q_2}$. Such a number exists by the Chinese remainder theorem, and can be computed in polynomial time. We will decompose row and column indices a, b and c as follows: $a = \alpha_1 q_2 + \alpha_2$, $b = \beta_1 q_1 + \beta_2$, and $c = \gamma_1 q_1 + \gamma_2$. Note the asymmetry in the definitions of a, b and c.

We can now define C and D:

$$C(a,b) = \begin{cases} 0 & \text{if } \alpha_2 \neq \beta_1 \\ \frac{1}{q_1^{1/2}}\, \omega^{\alpha_1\beta_2 q_2 + \beta_1\beta_2(u+1)} & \text{otherwise} \end{cases}$$

$$D(b,c) = \begin{cases} 0 & \text{if } \beta_2 \neq \gamma_2 \\ \frac{1}{q_2^{1/2}}\, \omega^{\beta_1\gamma_1 q_1 - \beta_1\beta_2 u} & \text{otherwise} \end{cases}$$

It is easy to see that $CD(a,c) = C(a,b)D(b,c)$ where $b = \alpha_2 q_1 + \gamma_2$ since we need $\alpha_2 = \beta_1$ and $\beta_2 = \gamma_2$ for non-zero entries in $C(a,b)$ and $D(b,c)$. Now,

$$\begin{aligned} CD(a,c) &= \tfrac{1}{q_1^{1/2} q_2^{1/2}}\, \omega^{\alpha_1\beta_2 q_2 + \beta_1\beta_2(u+1) + \beta_1\gamma_1 q_1 - \beta_1\beta_2 u} \\ &= \tfrac{1}{q^{1/2}}\, \omega^{\alpha_1\gamma_2 q_2 + \alpha_2\gamma_1 q_1 + \alpha_2\gamma_2} \\ &= \tfrac{1}{q^{1/2}}\, \omega^{(\alpha_1 q_2 + \alpha_2)(\gamma_1 q_1 + \gamma_2)} \\ &= \tfrac{1}{q^{1/2}}\, \omega^{ac} \end{aligned}$$

so $CD(a,c) = A_q(a,c)$.

We will now sketch how to rearrange the rows and columns of C to get the matrix $\bigoplus_{q_2} A_{q_1}$. The matrix C can be put in block-diagonal form where the blocks are indexed by $\alpha_2 = \beta_1$ (since all entries with $\alpha_2 \neq \beta_1$ are 0). Let $u + 1 \equiv tq_2 \pmod{q}$. Within a given block $\alpha_2 = \beta_1$, the entries look like

$$\begin{aligned} \sqrt{q_1}\, C(a,b) &= \omega^{\alpha_1\beta_2 q_2 + \beta_1\beta_2(u+1)} \\ &= \exp(2\pi i(\alpha_1\beta_2 + \beta_1\beta_2 t)q_2/q) \\ &= \exp(2\pi i(\alpha_1 + \alpha_2 t)\beta_2/q_1). \end{aligned}$$

Thus, if we rearrange the rows within this block so that they are indexed by $\alpha' \equiv \alpha_1 + \alpha_2 t \pmod{q_1}$, we obtain the transformation $\alpha' \to \beta_2$ with amplitude $\frac{1}{q_1^{1/2}} \exp(2\pi i \alpha'\beta_2/q_1)$, which is A_{q_1}. The matrix D can similarly be rearranged to obtain the matrix $\bigoplus_{q_1} A_{q_2}$.

We also need to show how to find a smooth q that lies between n and $2n$ in polynomial time. There are actually smooth q much closer to n than this, but this is all we need. It is also not known how to find smooth numbers very close to n in polynomial time.

Lemma 3.2 *Given n, there is a polynomial-time algorithm to find a number q with $n \leq q < 2n$ such that no prime power larger than $c \log q$ divides q, for some constant c independent of n.*

Proof: To find such a q, multiply the primes $2 \cdot 3 \cdot 5 \cdot 7 \cdot 11 \cdots p_k$ until the product is larger than n. Now, if this product is larger than $2n$, divide by the largest prime that keeps the number larger than n. This produces the desired q. There is always a prime between m and $2m$ ([HW], Theorem 418), so $n \leq q < 2n$. The prime number theorem ([HW], Theorem 6) and some calculation show that the largest prime dividing q is of size $O(\log n)$. ∎

4 Discrete Log: The Easy Case

The discrete log problem is: given a prime p, a generator g of the multiplicative group (mod p) and an x (mod p), find an r such that $g^r \equiv x \pmod{p}$. We will start by giving a polynomial-time algorithm for discrete log on a quantum computer in the case that $p-1$ is smooth. This algorithm is analogous to the algorithm in Simon's paper [Sim], with the group $\mathrm{Z}_2{}^n$ replaced by Z_{p-1}. The smooth case is not in itself an interesting accomplishment, since there are already polynomial time algorithms for classical computers in this case; however, explaining this case is easier than explaining the general case, and as the two algorithms are similar, the easy case will illuminate how the general case works.

We will start our algorithm with x, g and p on the tape. We are trying to compute r such that $g^r \equiv x \pmod{p}$. Since we will never delete them, x, g, and p are constants, and we will specify a state of our machine by the other contents of the tape.

The algorithm starts out by picking random numbers a and b (mod $p-1$), so the state of the machine after this step is

$$\frac{1}{p-1} \sum_{a=0}^{p-2} \sum_{b=0}^{p-2} |a, b\rangle .$$

The algorithm next computes $g^a x^{-b}$ reversibly, so we must keep the values a and b on the tape. The state of the machine is now

$$\frac{1}{p-1} \sum_{a=0}^{p-2} \sum_{b=0}^{p-2} |a, b, g^a x^{-b}\rangle .$$

What we do now is map $a \to c$ with amplitude $\frac{1}{(p-1)^{1/2}} \exp(2\pi i ac/(p-1))$ and $b \to d$ with amplitude $\frac{1}{(p-1)^{1/2}} \exp(2\pi i bd/(p-1))$. As was discussed in the previous section, this is a unitary transformation, and since $p-1$ is smooth it can be accomplished in polynomial time on a quantum machine. This leaves the machine in state

$$\frac{1}{(p-1)^2} \sum_{a=0}^{p-2} \sum_{b=0}^{p-2} \sum_{c=0}^{p-2} \sum_{d=0}^{p-2} \exp(2\pi i(ac+bd)/(p-1)) \; |c, d, g^a x^{-b}\rangle .$$

We now compute the probability that the computation ends with the machine in state $|c, d, y\rangle$ with $y = g^k$. This probability is the absolute square of the sum over all ways the

machine could produce this state, or

$$\left| \frac{1}{(p-1)^2} \sum_{\substack{a,b \\ a-rb \equiv k}} \exp(2\pi i(ac+bd)/(p-1)) \right|^2 .$$

where the sum is over all a, b satisfying $a - rb \equiv k \pmod{p-1}$. This condition arises from the fact that computational paths can only interfere when they give the same $y \equiv g^{a-rb} \equiv g^k \pmod{p}$. We now substitute the equation $a \equiv k + rb \pmod{p-1}$ in the above exponential. The above sum then reduces to

$$\left| \frac{1}{(p-1)^2} \sum_{b=0}^{p-2} \exp(2\pi i(kc+b(d+rc))/(p-1)) \right|^2 .$$

However, if $d + rc \not\equiv 0 \pmod{p-1}$ the above sum is over a set of $(p-1)$st roots of unity evenly spaced around the unit circle, and thus the probability is 0. If $d \equiv -rc$ the above sum is $(p-1)e^{2\pi ikc/(p-1)}$, so the probability is $1/(p-1)^2$. We can check that these probabilities add up to one by counting that there are $(p-1)^2$ states $|c, -rc, y\rangle$ since there are $p-1$ choices of $c \pmod{p-1}$ and $p-1$ choices of $y \not\equiv 0 \pmod{p}$.

Our computation thus produces a random $c \pmod{p-1}$ and $d \equiv -rc \pmod{p-1}$. If c and $p-1$ are relatively prime, we can find r by division. Because we are choosing among all possible c's with equal probability, the chance that c and $p-1$ are relatively prime is $\phi(p-1)/(p-1)$, where ϕ is the Euler ϕ-function. It is easy to check that $\phi(p-1)/(p-1) > 1/\log(p)$. (Actually, from [HW] Theorem 328, $\liminf \phi(p-1)/(p-1) \approx e^{-\gamma}/\log\log p$.) Thus we only need a number of experiments polynomial in $\log p$ to obtain r with high probability. In fact, we can find a set of c's such that at least one is relatively prime to every prime divisor of $p-1$ by repeating the experiment only an expected constant number of times. This also gives us enough information to obtain r.

5 A Note on Precision

The number of bits of precision needed in the amplitude of quantum mechanical computers could be a barrier to practicality. The generally accepted theoretical dividing line between feasible and infeasible is that polynomial precision (i.e., logarithmic number of bits) is feasible and that more is infeasible. This is because on a quantum computer the phase angle would need to be obtained through some physical device, and constructing such devices with better than polynomial precision is clearly impractical. In fact, even polynomial precision may prove to be impractical; however, using it as the dividing line results in nice theoretical properties.

We thus need to show that the computations in the previous section need to use only polynomial precision in the amplitudes. The very act of writing down the expression $\exp(2\pi iac/(p-1))$ seems to imply that we need exponential precision, as this phase angle is exponentially precise. Fortunately, this is not the case. Consider the same matrix A_{p-1}

with every term $\exp(2\pi iac/(p-1))$ replaced by $\exp(2\pi iac/(p-1) \pm \pi i/20)$. Each positive case, i.e., one resulting in $d \equiv -rc$, will still occur with nearly as large probability as before; instead of adding $p-1$ amplitudes which have exactly the same phase angle, we add $p-1$ amplitudes which have nearly the same phase angle, and thus the size of the sum will only be reduced by a constant factor. The algorithm will thus give a (c,d) with $d \equiv -rc$ with constant probability (instead of probability 1).

Recall that we obtain the matrix A_{p-1} by multiplying at most $\log p$ matrices A_{q_i}. Further, each entry in A_{p-1} is the product of at most $\log p$ terms. Suppose each phase angle were off by at most $\epsilon/\log p$ in the A_{q_i}'s. Then in the product, each phase angle would be off by at most ϵ, which is enough to perform the computation with constant probability of success. A similar argument shows that the magnitude of the amplitudes in the A_{q_i} can be off by a polynomial fraction. Almost exactly the same arguments hold for the general case of discrete log and for factoring to show that we need only polynomial precision for the amplitudes in these cases as well.

We still need to show how to construct A_{q_i} from constant size unitary matrices with limited precision. There is not room to do that in this extended abstract, but in fact, Bernstein and Vazirani [Vaz] have shown that it is sufficient to use polynomial precision for any computation on a quantum Turing machine to obtain the answer with high probability. An interesting open question is whether it is possible to do this computation with less than polynomial precision, or whether some precision–time tradeoff is possible.

6 Discrete Log: The General Case

For the general case, we first find a smooth number q such that q is close to p, i.e., with $p \le q \le 2p$ (see Lemma 3.2).

Now, we do the same thing as before, we choose a and b at random (mod $p-1$), and then compute $g^a x^{-b} \pmod p$. This leaves our machine in state

$$\frac{1}{p-1} \sum_{a=0}^{p-2} \sum_{b=0}^{p-2} |a, b, g^a x^{-b}\rangle .$$

As before, we next send $a \to c$ and $b \to d$ (mod q), with amplitude $\frac{1}{q}\exp(2\pi i(ac+bd)/q)$, giving us the state

$$\frac{1}{(p-1)q} \sum_{a=0}^{p-2} \sum_{b=0}^{p-2} \sum_{c=0}^{q-1} \sum_{d=0}^{q-1} \exp(2\pi i(ac+bd)/q)\ |c, d, g^a x^{-b}\rangle .$$

Note that we now have two moduli to deal with, $p-1$ and q. While this makes keeping track of things more confusing, we will still be able to obtain r using a similar algorithm to the easy case. The probability of observing a state $|c, d, y\rangle$ with $y \equiv g^k \pmod p$ is, almost as before,

$$\left| \frac{1}{(p-1)q} \sum_{\substack{a,b \\ a-rb \equiv k}} \exp(2\pi i(ac+bd)/q) \right|^2$$

where the sum is over all (a, b) such that $a - rb \equiv k \pmod{p-1}$. We now use the relation

$$a = br + k - (p-1)\left\lfloor \tfrac{br+k}{p-1} \right\rfloor$$

and substitute in the above expression to obtain the amplitude

$$\frac{1}{(p-1)q} \sum_{b=0}^{p-2} \exp\left(2\pi i \left(brc + kc + bd - c(p-1)\left\lfloor \tfrac{br+k}{p-1} \right\rfloor\right)/q\right).$$

We will now analyze this expression. First, a factor of $\exp(2\pi ikc/q)$ can be taken out of all the terms and ignored, because it does not change the probability. Next, we split the exponent into two parts and factor out b to obtain

$$\frac{1}{(p-1)q} \sum_{b=0}^{p-2} \exp\left(\frac{2\pi ib}{q}\Big(rc + d - \tfrac{r}{p-1}\{c(p-1)\}_q\Big)\right) \exp\left(\frac{2\pi i}{q}\Big(\tfrac{br}{p-1} - \left\lfloor \tfrac{br+k}{p-1} \right\rfloor\Big)\{c(p-1)\}_q\right) \tag{1}$$

where by $\{a\}_q$ we mean the least residue of $a \pmod{q}$. We will show that if we get a "good" output, then we still can deduce r, and that furthermore, the chance of getting a "good" output is constant. The idea is that if

$$\left| rc + d - \frac{r}{p-1}\{c(p-1)\}_q \right| \leq \frac{1}{2} \tag{2}$$

then as b varies between 0 and $p-1$, the first exponential term in Eq. (1) only varies over at most half of the unit circle. Further, if

$$\{c(p-1)\}_q \leq q/20, \tag{3}$$

the second exponential term never is farther than $\exp(\pi i/10)$ from 1. By combining these two observations, we show that if both conditions hold, then the contribution to the probability from the corresponding term is significant. Furthermore, both conditions will hold with constant probability, and a reasonable sample of c's for which condition (2) holds allows us to deduce r.

We will now figure out a lower bound on the probability of each good output. By the condition (2), we know that as b ranges from 0 to $p-2$, the phase ranges from 1 to

$$\frac{2\pi i(p-2)}{q}\Big(rc + d - \tfrac{r}{p-1}\{c(p-1)\}_q\Big)$$

and thus the component of the vector in the direction

$$\exp\left(\frac{\pi ip}{q}\Big(rc + d - \tfrac{r}{p-1}\{c(p-1)\}_q\Big)\right) \tag{4}$$

is at least $\sin \pi b/p$. Now, by condition (3), the phase can vary by at most $\pi i/10$ due to the second exponential. Applying this variation in the manner that minimizes the component

in direction (4), we get that the component in this direction is at least $\sin(\pi b/p - \pi/10)$ if $b < p/2$ and at least $\sin(\pi b/p + \pi/10)$ if $b > p/2$. Putting everything together, the probability of arriving at a state $|c, d, y\rangle$ that satisfies both conditions (2), (3) is at least

$$\left(\frac{1}{q}\frac{1}{\pi}\int_{-\pi/10}^{9\pi/10} \sin b \; db\right)^2,$$

or at least $.366/q^2$. Now, the number of pairs (c, d) such that condition (2) holds is exactly the number of possible c's, since for every c there is exactly one d such that it holds (round off the fraction to the nearest integer to obtain this d). The number of c's for which condition (3) holds is approximately $q/20$. Thus, there are $q/20$ (c, d) pairs satisfying both conditions. Multiplying by $p-1$, which is the number of possible y's, gives roughly $pq/20$ states $|c, d, y\rangle$. Putting this calculation together with the lower bound on the probability of each good state gives us that the probability of obtaining any good state is at least $p/60q$, or at least $1/120$ (since $q < 2p$).

We must now recover r from a pair c, d such that

$$-\frac{1}{2q} < \frac{d}{q} + \frac{r}{q}\left(c - \frac{\{c(p-1)\}_q}{p-1}\right) < \frac{1}{2q} \qquad (\text{mod } 1).$$

The first thing to notice is that multiplier on r is a fraction with denominator $p-1$, since q divides $c(p-1) - \{c(p-1)\}_q$. Thus, we need only round d/q off to the nearest multiple of $1/(p-1)$ and divide (mod $p-1$) to find a candidate r. To show that this experiment need only be repeated a polynomial number of times to find the correct r requires only a few more details. The problem is again that we cannot divide by a number which is not relatively prime to $p-1$.

What we have is that each good (c, d) pair is generated with probability at least $1/6q$. These are mapped from c/q to

$$\frac{c}{q} - \frac{\{c(p-1)\}_q}{p-1} = \frac{\left\lfloor \frac{c(p-1)}{q} \right\rfloor}{p-1},$$

i.e., c/q is rounded down to the nearest multiple of $1/(p-1)$, say $c'/(p-1)$. Further, the good c's are exactly those in which c/q is close to $c'/(p-1)$. Thus, each good c corresponds with exactly one good c'. We would like to show that for any prime p_i dividing $p-1$, a random good c' is unlikely to contain p_i. If we are willing to accept a large constant for the algorithm, we can just ignore the primes under 40; if we know r modulo all primes over 40, we can try all possible residues for primes under 40 with a constant factor penalty in running time. For a prime p_i over 40, each good c' is divisible by p_i with probability at most $20/p_i$. Thus, if we have for t good c''s chosen uniformly from all good c', the probability of having a prime that divides all of them is at most

$$\sum_{\substack{p_i > 40 \\ p_i | p}} \left(\frac{20}{p_i}\right)^t.$$

This sum (over all primes > 40) converges for $t = 2$, and goes down by at least a factor of 2 for each further increase of t by 1; thus for some large constant t it is less than $1/2$. Now, each experiment gives every good (c, d) pair with probability at least $1/6q$. Since there are $q/20$ good (c, d) pairs, after $120t$ experiments, we are likely to obtain a sample of t good (c, d) pairs chosen equally likely from all good (c, d) pairs. Thus, we will be able to find a set of c's such that all primes $p_i > 40$ dividing $p - 1$ are relatively prime to at least one of these c's. This will enable us to deduce r.

If one were to actually program this algorithm (which must wait until a quantum computer is built) there are many ways in which the efficiency could be increased over the efficiency shown in this paper.

7 Factoring

This extended abstract will only have a brief section on factoring, which sketches the proof. Filling in the details of the proof can be done by looking at the corresponding steps in the proof of the general case of the discrete log.

We will give a quantum computational algorithm for finding the order of an element (mod n); that is, the least integer r such that $x^r \equiv 1 \pmod{n}$. To factor n, given a way to compute the order of an element, we choose a random x, find the order r_x of x, and compute $\gcd(x^{r_x/2} - 1, n)$. This fails to give a non-trivial divisor of n only if r_x is odd or if $x^{r_x/2} \equiv -1 \pmod{n}$. Using this criterion, it can be shown that the algorithm finds a factor of n with probability at least $1/2$.

Given x and n, to find r such that $x^r \equiv 1 \pmod{n}$, we do the following. First, we find a smooth q with $5n^2 < q \leq 10n^2$. Next, we choose a random number $a \pmod{q}$. This leaves our machine in state

$$\frac{1}{q^{1/2}} \sum_{a=0}^{q-1} |a\rangle .$$

As in the algorithm for discrete log, we will not write x and n in the state of our machine, because we never erase these values.

Next, we compute x^a. We then map $a \to c$ with amplitude $\frac{1}{q^{1/2}} \exp(2\pi iac/q)$. This leaves our machine in state

$$\frac{1}{q} \sum_{a=0}^{q-1} \exp(2\pi iac/q) \, |c, x^a\rangle .$$

We now compute the probability that our machine ends in this state. Writing $a = br + k$, we obtain that this probability is

$$\left| \frac{1}{q} \sum_{b=0}^{\lfloor (q-k)/r \rfloor} \exp(2\pi i(br + k)c/q) \right|^2 .$$

Using the same argument as in the algorithm for discrete log, if $\{rc\}_q$ is small relative to q, all the amplitudes will point in nearly the same direction, giving a big probability. This

probability of seeing a state $|c, x^k\rangle$ will thus be at least $1/3r^2$ if

$$\frac{-r}{2} \leq \{rc\}_q \leq r/2,$$

i.e., if there is a c' such that

$$\frac{-r}{2} \leq rc - c'q \leq r/2.$$

Dividing by rq and rearranging the terms gives

$$\left|\frac{c}{q} - \frac{c'}{r}\right| \leq \frac{1}{2q}.$$

Now, if this holds and c' is relatively prime to r, we can obtain r by rounding c/q to the nearest fraction with a denominator of at most n. Because $q > 5n^2$, there is at most one such fraction. This fraction can be found in polynomial time by using a continued fraction expansion of c/q, which finds all the best approximations of c/q by fractions.

There are r^2 states $|c, x^k\rangle$ with $\left|\frac{c}{q} - \frac{c'}{r}\right| \leq \frac{1}{2q}$ because there are r possible c', each one giving one c, and there are r possible k. Using similar arguments as in the case of discrete log, we find that with probability at least $\phi(r)/3r$, we obtain c'/r with c' relatively prime to r. We can thus find r a polynomial fraction of the time, so by repeating this experiment only polynomially many times, we are assured of a high probability of success.

Note that in the algorithm for order, we did not use many of the properties of multiplication (mod n). In fact, if we have a permutation f mapping the set $\{0, 1, 2, \ldots, n-1\}$ into itself such that its kth iterate, $f^{(k)}(a)$, is computable in time polynomial in $\log n$ and $\log k$, the same algorithm will be able to find the order of an element a under f, i.e., the minimum r such that $f^{(r)}(a) = a$.

Acknowledgements

I would like to thank Jeff Lagarias for finding and fixing a critical bug in an early version of this paper and Andrew Odlyzko for showing me how to factor given an algorithm for the order of an element. I would also like to thank Charles Bennett, Gilles Brassard, Dan Simon, and Umesh Vazirani for productive discussions, for corrections and improvements of early drafts of this paper, and for pointers to the literature.

References

[Beni] P. Benioff, *Quantum mechanical Hamiltonian models of Turing machines*, J. Stat. Phys., Vol. 29, pp. 515–546 (1982).

[Benn] C. H. Bennett, *Logical Reversibility of Computation*, IBM J. Res. Develop., Vol 17, pp. 525–532 (1973).

[BV] E. Bernstein and U. Vazirani, *Quantum Complexity Theory*, Proc. 25th ACM Symp. on Theory of Computation, pp. 11–20 (1993).

[BB1] A. Berthiaume and G. Brassard, *The Quantum Challenge to Structural Complexity Theory*, Proc. 7th IEEE Conference on Structure in Complexity Theory (1992).

[BB2] A. Berthiaume and G. Brassard, *Oracle Quantum Computing*, Proc. Physics of Computation (1992).

[Deu1] D. Deutsch, *Quantum Theory, the Church–Turing Principle and the Universal Quantum Computer*, Proc. R. Soc. Lond., Vol. A400, pp. 96–117 (1985).

[Deu2] D. Deutsch, *Quantum Computational Networks,* Proc. R. Soc. Lond., Vol. A425, pp. 73–90 (1989).

[DJ] D. Deutsch and R. Jozsa, *Rapid Solution of Problems by Quantum Computation*, Proc. R. Soc. Lond., Vol. A439, pp. 553–558 (1992).

[Fey] R. Feynman, *Simulating Physics with Computers*, International Journal of Theoretical Physics, Vol. 21, nos. 6/7, pp. 467–488 (1982).

[HW] G. H. Hardy and E. M. Wright, *An Introduction to the Theory of Numbers*, Fifth Edition, Oxford University Press, New York, 1979.

[RSA] R. L. Rivest, A. Shamir, and L. Adelman *A Method of Obtaining Digital Signatures and Public-Key Cryptosystems*, Comm. ACM, Vol. 21, No. 2, pp. 120–126 (1978).

[Sim] D. Simon, *On the Power of Quantum Computation*, manuscript.

[Vaz] U. Vazirani, personal communication.

[Yao] A. Yao, *Quantum Circuit Complexity*, Proc. 34th IEEE Symp. on Foundations of Computer Science, 1993.