

Quantum Cryptanalysis of Hidden Linear Functions

(Extended Abstract)

Dan Boneh
dabo@cs.princeton.edu

Richard J. Lipton*
rjl@cs.princeton.edu

Department of Computer Science
Princeton University
Princeton, NJ 08544

Abstract. Recently there has been a great deal of interest in the power of “Quantum Computers” [4, 15, 18]. The driving force is the recent beautiful result of Shor that shows that discrete log and factoring are solvable in random quantum polynomial time [15]. We use a method similar to Shor’s to obtain a general theorem about quantum polynomial time. We show that any cryptosystem based on what we refer to as a ‘hidden linear form’ can be broken in quantum polynomial time. Our results imply that the discrete log problem is doable in quantum polynomial time over any group including Galois fields and elliptic curves. Finally, we introduce the notion of ‘junk bits’ which are helpful when performing classical computations that are not injective.

1 Introduction

The general discrete log problem can be phrased as follows: Let G be a finite group for which the group operation can be computed efficiently (given $x, y \in G$ we can find $x + y$). Let $h : \mathbb{Z} \rightarrow G$ be a homomorphism from the integers to G which can also be computed efficiently. Given $\beta = h(\alpha)$ the general discrete log problem is to find the smallest positive integer x such that $h(x) = \beta$. For example, in the standard discrete log problem over \mathbb{Z}_p^* the homomorphism h is defined by $h(\alpha) = g^\alpha \pmod{p}$ for some generator g of \mathbb{Z}_p^* . Here \mathbb{Z}_p^* is the multiplicative group of residues modulo a prime p .

A large variety of cryptosystems are based on the discrete log problem for various groups G . Specific groups that are being used are the multiplicative groups of large Galois fields [6], the multiplicative group of residues modulo a composite number [9, 10], elliptic curves over finite fields [11, 7] and the class group of imaginary quadratic fields [17].

Recently Shor [15] showed that the discrete log problem where $G = \mathbb{Z}_p^*$ can be solved in polynomial time on a quantum machine. We generalize this result to show that any type of cryptosystem which is based on what we refer to as

* Supported in part by NSF CCR-9304718.

a “hidden linear form” can be broken in quantum polynomial time(QP). An immediate application of this result shows that the general discrete log problem for any finite group G can be solved in QP. Thus, QP can break any of the cryptosystems discussed above.

Simon [14] observed that in QP it is possible to find a period of a function defined over \mathbb{Z}_2^n . We show that it is possible to detect the period of any function defined over \mathbb{Z} , even when the function is not one to one in its fundamental domain. Our method is similar to Shor’s factoring algorithm and is crucial for solving the general discrete log problem.

These results raise a natural question of trying to detect periods over arbitrary groups G . The problem can be stated as follows: given a function $f : G \rightarrow D$ for some range D , find an element $g \in G$ such that $f(x+g) = f(x)$ for all $x \in G$. For instance, the problem of detecting periods of functions over S_n is of significant importance since the problem of graph isomorphism can be reduced to it. Fourier analysis is a natural tool to use when trying to detect a period of a function. It is well known that one can define a Fourier transform over any group G ([13]). Now, suppose that for a given group G , the Fourier transform of G can be computed in QP (in time polynomial in $\log |G|$). Does this imply that a period of the function $f : G \rightarrow D$ can be found in QP? We have so far been unable to resolve this general problem. However, our results can be generalized to solve this problem for any finite Abelian group.

We assume that the reader is familiar with the general model of quantum computations. See [4, 15, 18] for further details.

2 Main Results

In this section we will state our main results. We begin by introducing some terminology. A function $h : \mathbb{Z} \rightarrow S$ has *period* q if for any integer x we have $h(x+q) = h(x)$. Such a function h can be regarded as a function from \mathbb{Z}_q to S . Here \mathbb{Z}_q is the group of residues modulo q . We say that the function h has *order at most* m provided that h does not map more than m elements of \mathbb{Z}_q to one, i.e. all $z \in S$ satisfy $|h^{-1}(z) \pmod{q}| \leq m$.

Let $f(x_1, \dots, x_k)$ be a function from the integers \mathbb{Z}^k to some arbitrary range S . Say that f has *hidden linear structure* over q provided there are integers $\alpha_2, \dots, \alpha_k$ and some function h with period q so that

$$f(x_1, \dots, x_k) = h(x_1 + \alpha_2 x_2 + \dots + \alpha_k x_k)$$

for all integers x_1, \dots, x_k . We say that f has order at most m if h has order at most m .

Theorem 1. *Suppose that $f(x_1, \dots, x_k)$ is a function which has a hidden linear structure over q of order at most m . We impose two technical conditions:*

1. *Let $n = \log q$ then m and k are at most $n^{O(1)}$.*
2. *Let p be the smallest prime divisor of q ; then $m < p$.*

For such a function f , in random quantum polynomial time in n we can recover the values of all the $\alpha_2, \dots, \alpha_n \pmod{q}$ from an oracle for f .

The point of this theorem is that random quantum polynomial time is able to solve a kind of cryptanalysis problem. With just the ability to evaluate the function f we can find the “secret” linear structure of f . The two restrictions on the function f are critical. The first one restricts m , the order of h . This is crucial since for example, if h is a constant function then trivially it is impossible to recover the values of the α 's.

The second restriction on m ensures that the $\alpha_2, \dots, \alpha_n$ are unique modulo q . In fact, as we shall see in Section 6, this condition enables us to test if a proposed solution $\alpha'_2, \dots, \alpha'_n$ is the correct one. Note that when q has no small factors the second restriction is subsumed by the first.

Another important problem which can be solved in quantum polynomial time is that of determining the period of a function.

Theorem 2. *Suppose the function $h : \mathbb{Z} \rightarrow S$ is periodic. Let q be the smallest positive period of h and assume h has order at most m . We impose two conditions:*

1. *Let $n = \log q$ then m is at most $n^{O(1)}$.*
2. *Let p be the smallest prime divisor of q ; then $m < p$.*

For such a function h , in random quantum polynomial time in n it is possible to recover the period q of h .

The two technical conditions are required so that we will be able to test that the output of the algorithm is correct. Theorem 2 shows that the value of q need not be known for Theorem 1 to hold. Indeed, as we shall see, in many important applications the value of q is not known.

3 Applications

There are several applications of these theorems. First, we generalize the original results of Shor [15] to show how to compute discrete log over an arbitrary group. To achieve this we show how to phrase the general discrete log problem as a hidden linear form.

Let $h : \mathbb{Z} \rightarrow G$ be a homomorphism and let $\beta = h(\alpha)$. Given β we wish to find the smallest positive integer x such that $\beta = h(x)$. Let d be the order of $h(1)$ in the group G . Clearly, the homomorphism h has period d . Note that in general d is unknown, e.g. when $G = \mathbb{Z}_n^*$ for some composite n or when G is the class group of a quadratic field.

Define the function $f : \mathbb{Z}^2 \rightarrow G$ as $f(x, y) = h(x + \alpha y)$. By the remarks above, the function f has a hidden linear form over d of order 1. An important observation is that the function f can be efficiently evaluated as follows:

$$f(x, y) = h(x)h(\alpha y) = h(x)h(\alpha)^y = h(x)\beta^y \quad .$$

To solve the general discrete log problem we apply the following two steps: first use Theorem 2 to find d , the period of the homomorphism h . The theorem can be applied since the function h has order 1, i.e. $m=1$. Then apply Theorem 1 to find an integer $\alpha' < d$ such that $\alpha' \equiv \alpha \pmod{d}$. Since α' is the smallest positive integer such that $h(\alpha) = h(\alpha')$, it is the required solution to the general discrete log problem. We have proved the following corollary to Theorems 1 and 2.

Corollary 3. *The general Discrete Log problem can be solved in random quantum polynomial time.*

This shows that we can find Discrete Log over composite modulus, Galois fields, and elliptic curves. An immediate corollary of Theorem 2 is the following.

Corollary 4. *Factoring can be solved in random quantum polynomial time.*

Proof. Suppose we wish to factor an n bit odd integer q . For an element $g \in \mathbb{Z}_q^*$, define the function $h : \mathbb{Z} \rightarrow \mathbb{Z}_q^*$ by $h(x) = g^x \pmod{q}$. Let d be the order of g in \mathbb{Z}_q^* then the function h has period d and order 1, i.e. $m=1$. Theorem 2 can be used to find the period of h and hence the order of g . The ability to find the order of an element in \mathbb{Z}_q^* enables us to factor as is described in [15]. \square

Another application of Theorem 1 concerns what are sometimes called “garbled” linear equations. Consider the following family of linear equations over \mathbb{Z}_q :

$$\begin{aligned} \alpha_1 x_{11} + \dots + \alpha_n x_{1n} &= y_1 + e_1 \\ &\vdots \\ \alpha_1 x_{m1} + \dots + \alpha_n x_{mn} &= y_m + e_m \end{aligned}$$

where e_1, \dots, e_m are unknown “errors” and the x ’s are known values. The general garbled linear equation problem is to find the value of the α ’s given $m \gg n$ large enough and given that most of the errors are equal to 0. This is a known difficult problem. However, suppose that the errors are determined by some polynomial time rule, i.e. some polynomial time function $e(\cdot)$ satisfies $e(y_i) = e_i$. Then the function

$$f(x_1, \dots, x_n) = h(\alpha_1 x_1 + \dots + \alpha_n x_n) \quad \text{where} \quad h(y) = y + e(y)$$

has a hidden linear structure. By Theorem 1 we can, in random quantum polynomial time, find the α ’s provided h does not collapse too much. Note, that we assume that we have an oracle that given x_1, \dots, x_n supplies us with the value of $y + e(y)$. Of course we do not assume we know when $e(y) = 0$ or not.

4 Basic Lemmas

Before we can prove Theorem 1 we need several lemmas. The following lemma is the main lemma which enables us to handle the fact that h may not be one-to-one in Theorems 1 and 2.

Lemma 5. *Let W be some integer and let $R < W$. Then for any integers b_1, \dots, b_m there are at least R/m^2 integers $0 \leq x \leq R$ satisfying*

$$\left| \sum_{k=1}^m \exp\left(\frac{2\pi i x b_k}{W}\right) \right| > \frac{1}{2} .$$

Lemma 5 relies on the following lemma.

Lemma 6. *Let $\lambda_1, \dots, \lambda_m$ be m complex numbers each of norm 1. Let $S_k = \sum_{j=1}^m \lambda_j^k$ then there exists a $1 \leq k \leq m$ such that $|S_k| > \frac{1}{2}$.*

Proof. Assume that for all $k = 1, \dots, m-1$ we have $|S_k| \leq \frac{1}{2}$. We show that this implies that $|S_m| > m/2$ proving the lemma. Let C_k be the m 'th symmetric polynomial in $\lambda_1, \dots, \lambda_m$, i.e.

$$C_k = \sum_{1 \leq j_1 < j_2 < \dots < j_k \leq m} \lambda_{j_1} \cdot \lambda_{j_2} \cdot \dots \cdot \lambda_{j_k} .$$

First we prove by induction on k that $|C_k| \leq \frac{1}{2}$ for $k = 1, \dots, m-1$. For $k = 1$ this is clear since $|C_1| = |S_1| \leq \frac{1}{2}$. Now, assume that $|C_j| \leq \frac{1}{2}$ for $j = 1, \dots, k-1 < m-1$. We show that $|C_k| \leq \frac{1}{2}$. For $k > 1$ define

$$A_k = C_1 S_{k-1} - C_2 S_{k-2} + \dots + (-1)^k C_{k-1} S_1 .$$

The Newton relations (see [8]) state that $S_k - A_k + (-1)^k k C_k = 0$ for $k \leq m$. The induction hypothesis implies that $|A_k| < \frac{k-1}{2}$ since the norm of each term in the sum is less than $1/2$. Hence,

$$|C_k| = \frac{1}{k} |S_k - A_k| \leq \frac{1}{k} (|S_k| + |A_k|) \leq \frac{1}{2} .$$

To conclude the proof of the lemma we show that $|S_m| > m/2$. The fact that for $k = 1, \dots, m-1$ we have $|S_k| \leq \frac{1}{2}$ and $|C_k| \leq \frac{1}{2}$ implies that $|A_m| \leq m/2$. Furthermore, Since $C_m = \prod_{k=1}^m \lambda_k$ we know that $|C_m| = 1$. Hence, by Newton's relations

$$|S_m| = |A_m - (-1)^m m C_m| \geq |m C_m| - |A_m| \geq m - \frac{m}{2} = m/2 .$$

□

Proof of Lemma 5. Define

$$\beta(x) = \left| \sum_{k=1}^m \exp\left(\frac{2\pi i x b_k}{W}\right) \right| .$$

By Lemma 6, for any x , one of $\beta(x), \beta(2x), \dots, \beta(mx)$ must be bigger than $\frac{1}{2}$. Observe that the integers $\{0, \dots, R\}$ can be partitioned into R/m^2 distinct sequences of the form $\{x, 2x, \dots, mx\}$. Hence, the lemma follows. \square

The following lemma provides a lower bound on the sum of roots of unity which are close to 1.

Lemma 7. *Suppose that $|\theta_k| \leq \epsilon < 1$ are real numbers for $k = 1, \dots, m$. Then,*

$$\left| \sum_{k=1}^m \exp(i \theta_k) \right| \geq (1 - \epsilon^2)m .$$

Proof. This follows directly from the fact that the real part of $\exp(i\theta_i)$ is at least $\cos(\epsilon) > 1 - \epsilon^2$. \square

5 An Overview of the Proofs

Before we present the proofs of Theorems 1 and 2 we will outline a general paradigm for proving that a problem of size n can be solved in quantum polynomial time. We will describe a certain quantum experiment \mathcal{E} . Each time we perform this experiment we will get some observable value. Let \mathcal{V} be some subset of all the possible observable values. We will arrange things so that the following are true:

1. Given *any* value from \mathcal{V} we can in polynomial time (on a conventional computer) solve the given problem.
2. The probability of observing a specific element of \mathcal{V} is at least $1/Wn^c$ for some integer W and constant c .
3. The cardinality of the set \mathcal{V} is at least $W/n^{c'}$ for some constant c' .

We refer to the observables in \mathcal{V} as the “good” observables. By 2 and 3 above, The probability of sampling an observable from \mathcal{V} is at least $1/n^{O(1)}$. Once such an observable is found it will be used to solve the given problem. Hence, in expected polynomial time the problem will be solved.

An important point is that we do not know which observables lie in the set \mathcal{V} . When an observable is observed, we try to use it to solve the hidden linear problem as if it is in \mathcal{V} . Then, we check that the computed result works correctly. If it does we are done; otherwise, we try again.

6 The Proof of Theorem 1

We now turn to the proof of Theorem 1. We will prove the theorem for a hidden linear form with two variables $f(x, y) = h(x + \alpha y)$. This is enough to prove the general theorem, since we can find all the α 's one by one by setting all the irrelevant variables to zero.

Let $f(x, y) = h(x + \alpha y)$ be a hidden linear form over q , an n -bit number. The assumptions of Theorem 1 state that h has order at most $m = n^d$ for some constant d and if p is the smallest prime divisor of q , then $m < p$. Our objective is to find α .

We first show that given an α' it is easy to test if $\alpha \equiv \alpha' \pmod{q}$. This is the only place where we use the fact that $m < p$. Let $A_{\alpha'}$ be the set of pairs $\{(-k\alpha', k)\}$ for $k = 0, \dots, m$.

Lemma 8. *If for all $(x, y) \in A_{\alpha'}$ we have $f(x, y) = h(x + \alpha'y)$ then $\alpha \equiv \alpha' \pmod{q}$.*

Proof. Observe that for all $(x, y) \in A_{\alpha'}$ we have $x + \alpha'y = 0$. Hence, all $(x, y) \in A_{\alpha'}$ satisfy $h(x + \alpha y) = f(x, y) = h(0)$. Now, suppose $\alpha \not\equiv \alpha' \pmod{q}$. For two distinct pairs (x, y) and (x', y') in $A_{\alpha'}$ we have that $x + \alpha y \not\equiv x' + \alpha y'$. This follows from the fact that

$$\alpha \not\equiv \alpha' \equiv \frac{x - x'}{y' - y} \pmod{q} .$$

The division by $y - y'$ is valid since $|y - y'| \leq m < p$ where p is the smallest prime divisor of q . Hence, $y - y'$ is relatively prime to q and hence invertible. This shows that h maps the $m+1$ pairs in $A_{\alpha'}$ to the same value, $h(0)$. However, by assumption h had order at most m . This contradiction proves the lemma. \square

6.1 The Quantum Experiment

Let $W_1 < W_2 < \dots$ be the first primes that are relatively prime to q . Define $W = \prod_{i=1}^k W_i$ as the first product that exceeds $\max\{2q, mq\}$. Note that W and q are relatively prime. Since $m < n^{O(1)}$ we have $W < qn^{O(1)}$.

Let F_W be the Fourier transform unitary matrix:

$$(F_W)_{x,y} = \frac{1}{\sqrt{W}} e^{2\pi i xy/W} .$$

Shor shows that for the W constructed above the transformation F_W can be carried out by a quantum machine in polynomial time. In general this holds whenever W is smooth, i.e. contains no large prime factors,

The quantum experiment \mathcal{E} is as follows: First, the quantum machine writes two random numbers r_1, r_2 from \mathbb{Z}_q on its tape. So the state after this first step is

$$\frac{1}{q} \sum_{r_1, r_2} |r_1, r_2\rangle .$$

The algorithm next computes the function f in a reversible manner so that the machine is in state

$$\frac{1}{q} \sum_{r_1, r_2} |r_1, r_2, f(r_1, r_2) \rangle .$$

We now use the mapping $(F_W)_{x,y} = e^{2\pi i xy/W}$ to send each r_i to s_i for $i = 1, 2$ with amplitude $\frac{1}{\sqrt{W}} \exp(2\pi i r_i s_i/W)$. This places the machine in the state

$$\frac{1}{qW} \sum \exp(2\pi i(r_1 s_1 + r_2 s_2)/W) |s_1, s_2, f(r_1, r_2) \rangle$$

where the sum is over all r_1, r_2 and s_1, s_2 . Thus, the machine will end up in state $|s_1, s_2, b \rangle$ with probability

$$\left| \frac{1}{qW} \sum \exp(2\pi i(r_1 s_1 + r_2 s_2)/W) \right|^2$$

where the sum is over all r_1, r_2 such that $f(r_1, r_2) = b$.

We now describe the special set of observables \mathcal{V} . We denote the residue of x modulo W by $\{x\}_W$. The observable (s_1, s_2, b) is in \mathcal{V} provided the following properties are satisfied:

1. $s_1 q \geq W$;
2. $\{s_1 q\}_W \leq W/m$;
3. Let $C = s_2 - s_1 \alpha + \frac{\alpha}{q} \{s_1 q\}_W$. Then $C = tW + \delta$ for some integer t and $|\delta| < 1$.
4. $\left| \sum_{k=1}^m \exp(2\pi i b_k s_1/W) \right| \geq 1/2$ where b_1, \dots, b_m are distinct elements so that $h(b_k) = b$ for $k = 1, \dots, m$. Recall that m is the order of the function h .

In what follows we will refer to these conditions as (1),(2),(3) and (4). It remains to prove that the set \mathcal{V} satisfies the three properties specified in Section 5.

6.2 Using a ‘‘Good’’ Observable

Let (s_1, s_2, b) be an observable from the set \mathcal{V} . We show how this observable can be used to find α . Condition (3) implies that

$$s_2 - \frac{\alpha}{q} (s_1 q - \{s_1 q\}_W) = tW + \delta .$$

Write $s_1 q = vW + u$ with $0 \leq u < W$. Observe that $v = \frac{s_1 q - \{s_1 q\}_W}{W}$. Since t is an integer, and $|\delta| < 1$, dividing the above equality by W leads to

$$\left\| \frac{s_2}{W} - \alpha \frac{v}{q} \right\| < \frac{1}{W}$$

where $\|x\|$ is the fractional part of x , i.e. $\min|x + i|$ over all integers i .

Let s be the integer which makes the values of $\frac{s}{q}$ the closest to $\frac{s_2}{W}$. That is, $\frac{s}{q}$ is the fraction we get when we round $\frac{s_2}{W}$ to the closest rational with denominator q . Since $W > 2q$ it is not difficult to see that for the above inequality to hold we must have

$$\left\| \frac{s}{q} - \alpha \frac{v}{q} \right\| = 0 .$$

This means that $s - \alpha v \equiv 0 \pmod{q}$. By condition (1) we know that $v \geq 1$. Hence, when q and v are relatively prime we can easily recover α .

When q and v are not relatively prime we proceed as follows: let $z = q/\gcd(q, v)$. Observe that v is invertible modulo z and let $\alpha' = s/v \pmod{z}$. Clearly $\alpha' \equiv \alpha \pmod{z}$. For $0 \leq \alpha' < z$ we have that $\alpha' \equiv \alpha \pmod{z}$ if and only if $\alpha' \frac{q}{z} \equiv \alpha \frac{q}{z} \pmod{q}$. Hence, it is easy to check that the resulting α' satisfies $\alpha' \equiv \alpha \pmod{z}$ by using Lemma 8 on the function $f'(x, y) = f(x, \frac{q}{z}y)$.

Once a pair α', z satisfying $\alpha' \equiv \alpha \pmod{z}$ is found, write $\alpha = \alpha' + zk$. Define a new function $f''(x, y) = f(zx - \alpha'y, y)$. Then

$$f''(x, y) = h(zx - \alpha'y + \alpha y) = h(z(x + ky)) .$$

Hence, $f''(x, y)$ has a hidden linear structure over q/z . We can now recursively apply the algorithm to f'' to find k and thus find $\alpha \pmod{q}$.

6.3 The Amplitude of a “Good” Observable

For an observable (s_1, s_2, b) , we denote by $\sigma(s_1, s_2, b)$ the probability of observing (s_1, s_2, b) at the end of the quantum experiment. To simplify the exposition in this section we assume that the order of the function f satisfies $m \geq 10$. This is not a restriction since a function which has order less than 10 may be regarded as a function with order 10.

Let (s_1, s_2, b) be an observable from the set \mathcal{V} . Recall that the probability of this observation is

$$\sigma(s_1, s_2, b) = \left| \frac{1}{qW} \sum \exp\left(\frac{2\pi i}{W}(r_1 s_1 + r_2 s_2)\right) \right|^2$$

where the sum is over all r_1, r_2 such that $f(r_1, r_2) = b$. The key is that f has a hidden linear structure, i.e. $f(r_1, r_2) = b$ if and only if $h(r_1 + \alpha r_2) = b$. Since h need not be one to one there are *distinct* $b_1, \dots, b_{m'}$ so that $h(b_k) = b$ for $k = 1, \dots, m'$ and $m' \leq m$. WLOG we assume $m = m'$. Thus, $\sigma(s_1, s_2, b)$ is equal to

$$\frac{1}{q^2 W^2} \left| \sum_{k=1}^m \sum \exp(2\pi i(r_1 s_1 + r_2 s_2)/W) \right|^2$$

where the inner sum is over all r_1, r_2 so that $r_1 \equiv b_k - \alpha r_2 \pmod{q}$. Since $1 \leq r_1 < q$, given an r_2 the value of r_1 is equal to $b_k - \alpha r_2 - q \lfloor (b_k - \alpha r_2)/q \rfloor$. Thus, the key is to bound the absolute value of the following double summation,

$$\sum_{k=1}^m \exp(2\pi i b_k s_1 / W) \sum_{r_2=0}^{q-1} \exp\left[\frac{2\pi i}{W}(r_2 s_2 - \alpha s_1 r_2 - s_1 q \lfloor (b_k - \alpha r_2)/q \rfloor)\right] .$$

First we bound the inner sums. For a given k , rewrite the inner sum as

$$\sum_{r_2=0}^{q-1} \exp \left[\frac{2\pi i}{W} r_2 (s_2 - \alpha s_1 + \frac{\alpha}{q} \{s_1 q\}_W) \right] \exp \left[-\frac{2\pi i}{W} \left(\frac{\alpha r_2}{q} + \left\lfloor \frac{b_k - \alpha r_2}{q} \right\rfloor \right) \{s_1 q\}_W \right] .$$

By condition (3), and the fact that $r_2/W < q/W < 1/m$, the argument of the first exponent is always less than $2\pi i/m$. For the second exponent we know $b_k < q$. The fact that all reals $A, B > 0$ satisfy $|B + \lfloor A - B \rfloor| \leq \lfloor A \rfloor + 1$ implies that

$$\left| \frac{\alpha r_2}{q} + \left\lfloor \frac{b_k - \alpha r_2}{q} \right\rfloor \right| \leq \left\lfloor \frac{b_k}{q} \right\rfloor + 1 \leq 1 .$$

Combining this with condition (2) we see that the argument of the second exponent is always less, in absolute value, than $2\pi i/m$. Hence, the total exponent is less than $4\pi i/m$. Using Lemma 7, we get that the inner sum is always bigger than $\lfloor 1 - O(\frac{1}{m^2}) \rfloor q$. On the other hand the inner sum is clearly less than q . It follows that $\sigma(s_1, s_2, b)$ is equal to

$$\sigma(s_1, s_2, b) = \frac{1}{W^2} \left| \sum_{k=1}^m (1 - \epsilon_k) \exp(2\pi i b_k s_1 / W) \right|^2$$

where $0 \leq \epsilon_k \leq O(\frac{1}{m^2})$ for all $k = 1, \dots, m$. Now, since the ϵ_k are small it is not difficult to see that condition (4) implies that $\sigma(s_1, s_2, b) > \Omega(\frac{1}{W^2})$. Hence, a “good” observable (s_1, s_2, b) has the required probability.

6.4 Cardinality of Set of “Good” Observables

The last step is to show that \mathcal{V} has the required cardinality. First, observe that for any s_1 there exists an s_2 satisfying condition (3). This follows by setting s_2 to the integer closest to $\alpha s_1 + \frac{\alpha}{q} \{s_1 q\}_W$. We only need to lower bound the number of s_1 satisfying

1. $s_1 q \geq W$;
2. $\{s_1 q\}_W \leq W/m$;
3. $\left| \sum_{k=1}^m \exp(2\pi i b_k s_1 / W) \right| \geq 1/2$

We will show that the number of s_1 satisfying conditions (2) and (3) is at least W/m^3 . The number of s_1 violating condition (1) is at most W/q which is negligible in comparison. Hence, throwing away the s_1 that violate condition (1) will make no difference.

Let $x = q s_1 \pmod{W}$ and $c_k = b_k q^{-1} \pmod{W}$. Since q and W are relatively prime by construction, q^{-1} exists modulo W . Conditions (2) and (3) can now be rewritten as

1. $0 \leq x \leq W/m$
2. $\left| \sum_{k=1}^m \exp(2\pi i c_k x / W) \right| \geq 1/2$

By Lemma 5, the number of x that satisfy these two conditions is at least W/m^3 . Since $m < n^{O(1)}$, the number of such x is at least $W/n^{O(1)}$.

Hence, the total number of pairs s_1, s_2 satisfying conditions (1),(2),(3) and (4) in the definition of \mathcal{V} is $W/n^{O(1)}$. Putting this together with the fact that there are q possible value for b , we get that the number of triplets (s_1, s_2, b) in \mathcal{V} is $qW/n^{O(1)}$. By definition of W we know that $W = qn^{O(1)}$. Hence, $|\mathcal{V}| > W^2/n^{O(1)}$, which is what we had to show.

7 The Proof of Theorem 2

Say we are given a function $h : \mathbb{Z} \rightarrow S$ which is periodic. We wish to find the smallest period q of h . Let $n = \log q$. We assume that h is of order at most m where $m = n^{O(1)}$.

Without loss of generality we can assume that we are given an upper bound q' on q such that $q' < 2q$. This upper bound can be found by guessing some initial q' and running the algorithm. If the algorithm fails to find the period, double q' and rerun the algorithm. After at most n steps q' will be the required upper bound.

Let p be the smallest prime factor of q . As in the previous section, the assumption of Theorem 2 that $m < p$ implies that when the algorithm outputs q' as the period, we can test that $q = q'$.

7.1 The Quantum Experiment

Let W be a smooth number constructed as in the previous section such that $W > \max\{q'^2, mq'^2\}$ and $W < q'^2 n^{O(1)}$. The quantum experiment \mathcal{E} is as follows: First, the quantum machine writes a random numbers r from \mathbb{Z}_W on its tape. So the state after this first step is

$$\frac{1}{\sqrt{W}} \sum_r |r\rangle .$$

The algorithm next computes the function h in a reversible manner so that the state of the machine is now

$$\frac{1}{\sqrt{W}} \sum_r |r, h(r)\rangle .$$

We now use the Fourier unitary transformation F_W to send r to s with amplitude $\frac{1}{\sqrt{W}} \exp(2\pi i r s / W)$. It places the machine in the state

$$\frac{1}{W} \sum_{r,s} \exp(2\pi i r s / W) |s, h(r)\rangle .$$

The probability that the machine ends in the state $|s, b\rangle$ is

$$\left| \frac{1}{W} \sum \exp(2\pi i r s / W) \right|^2$$

where the sum is over all r such that $h(r) = b$.

As before, we now describe the special set of observables \mathcal{V} . An observable (s, b) is in \mathcal{V} provided the following properties are satisfied:

1. $\{sq\}_W < q/m$;
2. $\left| \sum_{k=1}^m \exp(2\pi i b_k s/W) \right| \geq 1/2$ where b_1, \dots, b_m are distinct elements so that $h(b_k) = b$ for $k = 1, \dots, m$. Recall that m is the order of the function h .

It remains to prove that the set \mathcal{V} satisfies the three properties specified in Section 5:

1. Given an observable (s, b) in \mathcal{V} Condition (1) implies that we can find a non trivial factor z of q using a method similar to Shor's [15]. We can then define a new function $h'(x) = h(zx)$ which will have period q/z . The algorithm can be applied recursively on h' to recover q/z . This shows that given a "good" observable we can find the period q .
2. Using condition (2) and an argument similar to the one in the previous section we can show that the amplitude of a "good" observable is $\Omega(\frac{1}{q^2})$.
3. Using Lemma 5 we can show that the cardinality of \mathcal{V} is at least $q^2/n^{O(1)}$.

8 Junk Bits

In both algorithms described in the previous sections the first step was to pick a random number between 1 and $q - 1$ for some integer q . This means that the machine should be in state

$$\frac{1}{\sqrt{q}} \sum_{r=0}^{q-1} |r\rangle .$$

However, when q is a large prime, this state can not be easily constructed using a quantum circuit.

An easy method for generating a random number between 0 and $q - 1$ is to pick an integer W which is the closest power of 2 to q . Then generate a random number $x \pmod{W}$. If $x < q$ then use x , otherwise generate a new x and repeat this until a number in the required range is generated. This will clearly generate a number uniformly distributed on $0, \dots, q-1$. The problem is that this procedure can not be carried out on a quantum machine since all the "bad" samples (the ones larger than q) can not be erased from the tape. Erasure is not a reversible operation. Clearly the bad samples can not be left on the tape since they would prevent the interference effects which are so useful in quantum computing.

Another approach is to pick some large integer $W > q^2$ which is a power of 2. Then generate a random number $x \pmod{W}$ and compute $x \pmod{q}$. The resulting value will be exponentially close to being uniformly distributed between 0 and $q - 1$ which is good enough. However, as before, we run into the problem that the map sending x to $x \pmod{q}$ is not reversible. As before keeping extra information on the tape to make this map reversible is risky since it may prevent interference effects.

The solution is to keep just enough extra information on the tape so that the computation is reversible, however the extra information on the tape should be independent of the computation taking place. We call this extra information *Junk bits*.

Definition 9. Let $f : \{0, 1\}^n \rightarrow Y$ be some polynomial time computable function which is not one to one. A function $J : \{0, 1\}^n \rightarrow Y'$ will be called a “junk” function for f if the following are satisfied:

1. The map $x \rightarrow (f(x), J(x))$ is one to one and polynomial time computable. Furthermore, the inverse map is in QP;
2. $|\Pr[f(x) = y \mid J(x) = j] - \Pr[f(x) = y]| < 2^{-\Omega(n)}$.

Thus, the value of $J(x)$ and $f(x)$ should be almost independent of one another. Condition (1) implies that the map sending x to $(f(x), J(x))$ can be computed in QP using a result due to Bennett [2]. It should be clear that once we have computed $(f(x), J(x))$, the computation can proceed to use the value of $f(x)$ as if $J(x)$ was not written on the tape. The independence property will guarantee that the interference effects will change by an exponentially small amount. The full details of this method will be given in the final version of the paper.

To generate a random number between 0 and $q - 1$ we follow the second method. Let $W > q^2$ be a large power of 2. Generate a random number between 0 and $W - 1$. We now wish to compute the function $f(x) = x \bmod q$. A possible junk function for f is $J(x) = \lfloor x/q \rfloor$. It is not difficult to see that $J(x)$ is indeed a junk function for $f(x)$. Using similar methods we can show that it is possible to generate random permutations and other random objects.

9 Conclusions and Open Problems

We have shown that QP can solve two types of problems: recovering the hidden linear structure of a function and detecting periods over \mathbb{Z} . Our results hold even when the function h used is not one to one. Using both theorems we were able to show that the discrete log problem can be solved in quantum polynomial time over any group.

The problem of recovering the hidden linear structure can be generalized to any ring. Similarly, the problem of detecting periods can be generalized to any group. As was mentioned in the introduction, graph isomorphism is reducible to the problem of detecting periods of functions defined over the symmetric group S_n . This example shows the importance of these generalizations. We hope that Fourier methods analogous to the ones used in this paper can be used to detect periods over S_n . This will show that the graph isomorphism problem can be solved in random quantum polynomial time. We mention that Beals [1] has shown that the Fourier transform over the group S_n can be carried out in quantum polynomial time.

We have also introduced the concept of Junk bits which enables quantum machine to carry out certain non invertible functions in a way that does not effect the interference patterns. A natural problem is to try and understand which deterministic computations can be done using junk bits.

Acknowledgments

We wish to thank Robert Beals and Merrick Furst for helpful discussions about this work.

References

1. R. Beals, *Computing Fourier Transform over S_n in QP*, unpublished manuscript.
2. C. Bennett, *Logical reversibility of computation*, IBM J. Res. Develop. vol. 17, 1973, pp. 525-532.
3. C. Bennett, E. Bernstein, G. Brassard, U. Vazirani, *Strengths and Weaknesses of Quantum Computing*, to appear.
4. E. Bernstein and U. Vazirani, *Quantum Complexity Theory*, Proc. 25th ACM Symp. on Theory of Computation, 1993.
5. D. Coppersmith, *An Approximate Fourier Transform Useful in Quantum Factoring*, IBM Research Report 19642, 1994.
6. W. Diffie and M. Hellman, *New Directions in Cryptography*, IEEE transactions on Information Theory, vol. 22, no. 6, pp. 644-654, 1976.
7. N. Koblitz, *Elliptic Curve Cryptosystems*, Mathematics of Computations 48, 1987, pp. 203-209.
8. S. Lang, *Algebra*.
9. U. Maurer and Y. Yacobi, *Non-interactive public-key cryptography*, EuroCrypt-91, pp.498-507, 1991.
10. K. McCurley, *A Key Distribution System Equivalent to Factoring*, Journal of Cryptology, vol. 1, no. 2, pp. 95-105.
11. V. Miller, *Uses of Elliptic Curves in Cryptography*, In Proceedings of Crypto 1985, pp. 417-426.
12. B. Preneel, R. Govaerts, J. Vandewalle, *Hash Functions Based on Block Ciphers: A Synthetic Approach*, in Proc. of Advances in Cryptology, CRYPTO '93.
13. J. P. Serre, *Linear Representations of Finite Groups*, Springer-Verlag, 1977.
14. D. Simon, *On the Power of Quantum Computation*, Proc. FOCS, 1994, pp. 116-123.
15. P. Shor, *Algorithms for Quantum Computation*, Proc. FOCS, 1994, pp. 124-134.
16. L. Washington, *Introduction to Cyclotomic Fields*, Springer-Verlag, 1982.
17. J. Buchmann and H. Williams, *A Key Exchange System Based on Imaginary Quadratic Fields*, Journal of Cryptology, vol. 1, no. 2, pp. 107-118, 1988.
18. A. Yao, *Quantum Circuit Complexity*, Proc. 34th IEEE Symp. on Foundations of Computer Science, 1993, pp. 352-360.

This article was processed using the \LaTeX macro package with LLNCS style