Laboratoire de l'Informatique du Parallélisme

École Normale Supérieure de Lyon
Unité Mixte de Recherche CNRS-INRIA-ENS LYON-UCBL n^o 5668

The Quantum Query Complexity of the Abelian Hidden Subgroup Problem

Pascal Koiran
Vincent Nesme
Natacha Portier

Mai 2005

Rapport de recherche N^o RR2005–17

École Normale Supérieure de Lyon
46 Allée d'Italie, 69364 Lyon Cedex 07, France
Téléphone : +33(0)4.72.72.80.37
Télécopieur : +33(0)4.72.72.80.80
Adresse électronique : `lip@ens-lyon.fr`

The Quantum Query Complexity of the Abelian Hidden Subgroup Problem

Pascal Koiran
Vincent Nesme
Natacha Portier

Mai 2005

Abstract

Simon in his FOCS'94 paper was the first to show an exponential gap between classical and quantum computation. The problem he dealt with is now part of a well-studied class of problems, the hidden subgroup problems. We stu<dy Simon's problem from the point of view of quantum query complexity and give here a first nontrivial lower bound on the query complexity of a hidden subgroup problem, namely Simon's problem. More generally, we give a lower bound which is optimal up to a constant factor for any Abelian group. At last we expose some elementary facts about the query complexity of hidden subgroup problems in weaker query models.

Keywords: quantum computation, query complexity, hidden subgroup, Simon's problem, lower bound.

Résumé

Dans son article de FOCS'94, Simon fut le premier à montrer un cas où le calcul quantique permet une accélération exponentielle par rapport au calcul classique. Il s'agissait d'un problème qui fait partie d'une classe de problèmes aujourd'hui très étudiés, les problèmes de sous-groupes cachés. Nous étudions le problème de Simon du point de vue de la complexité en requêtes quantiques. Nous donnons une première borne inférieure non triviale sur cette complexité et montrons comment on obtient en conséquence la complexité en requêtes quantiques du problème du sous-groupe caché Abélien, à un facteur constant près. Nous présentons enfin quelques résultats élémentaires de complexité pour des modèles de requêtes plus faibles.

Mots-clés: calcul quantique, complexité en requêtes, sous-groupe caché, problème de Simon, borne inférieure.

1 Introduction

Given an Abelian group G and a subgroup $H \leq G$, a function $f : G \to X$ is said to be hiding H if f can be defined in a one-to-one way on G/H. More precisely, f hides H if and only if

$$\forall g, g' \in G \ \big(f(g) = f(g') \iff \exists h \in H \ g = g' + h\big)$$

Suppose G is a fixed group and f is computed by an oracle : a quantum black-box. We are interested here in algorithms that find the hidden subgroup H. A large amount of documentation about the hidden subgroup problem can be found in the book of Nielsen and Chuang [16][1]. Among all work already done about such algorithms one can cite Shor's famous factoring algorithm [20] : it uses a period-finding algorithm, which is a special case of a hidden subgroup problem. In recent years, attention has shifted to non-Abelian hidden subgroup problems but we will restrict our attention here to Abelian groups, and in particular to groups of the form $(\mathbb{Z}/p\mathbb{Z})^n$.

In general, two kinds of complexity measures for black-box problems can be distinguished : query complexity, i.e., the number of times the function f is evaluated using the black-box, and computational or time complexity, i.e., the number of elementary operations needed to solve the problem. Typically, a hidden subgroup algorithm is considered efficient if its complexity (in query or in time, depending on the interest) is polynomial in the logarithm of the cardinality of G. For example, Kuperberg's algorithm [13] for the (non-Abelian) dihedral hidden subgroup problem is subexponential (but superpolynomial) in both time and query complexities.

Our main result is that the query complexity of finding a subgroup hidden in G is of order $r(G)$ for any Abelian group G, where $r(G)$ denotes the *rank* of G, that is, the minimal cardinality of a generating set of G (for instance, $r((\mathbb{Z}/p\mathbb{Z})^n) = n$ if $p \geq 2$ is an arbitrary integer). The proof of this result is naturally divided into an upper bound and a lower bound proof. The upper bound is achieved through a tight analysis of the standard Fourier sampling algorithm. It is a folklore theorem in quantum computation that this algorithm solves the hidden subgroup problem in Abelian groups with polynomial query complexity (see for instance [10], [8], [5] or [11]), but strangely enough no precise analysis seems to be available in the literature.

The greatest part of this paper is devoted to the lower bound proof. Here all the important ideas already appear in the analysis of Simon's problem, to which our preprint [12] is devoted. It is therefore fitting to recall the history of this problem, which is defined as follows. We are given a function f from $G = (\mathbb{Z}/2\mathbb{Z})^n$ to a known set X of size 2^n, and we are guaranteed that the function fulfills Simon's promises, that is either :

(1) f is one-to-one, or

(2) $\exists s \neq 0 \ \forall w, w' \ f(w) = f(w') \iff (w = w' \vee w = w' + s)$.

The problem is to decide whether (1) or (2) holds. Note that (1) is equivalent to "f hides the trivial subgroup $H = \{(0, \dots, 0)\}$" and (2) is equivalent to "f hides a subgroup $H = \{(0, \dots, 0), s\}$ of order 2". The original problem [21] was to compute s and the problem considered here is the associated decision problem. Of course, any lower bound on this problem will imply the same one on Simon's original problem. In his article, Simon shows that his problem can be solved by a quantum algorithm which makes $O(n)$ queries in the worst case and has a bounded probability of error. The time complexity of his algorithm is linear in the time required to solve an $n \times n$ system of linear equations over $(\mathbb{Z}/2\mathbb{Z})^n$. He also shows that any classical (probabilistic) algorithm for his problem must

[1]History of the problem on page 246 and expression of many problems (order-finding, dicrete logarithm...) in terms of hidden subgroup problems on page 241.

have exponential query complexity. In this paper we shall give a $\Omega(n)$ lower bound on the query complexity of Simon's problem, thus showing that Simon's algorithm is optimal in this respect. Our lower bound applies in fact to groups of the form $(\mathbb{Z}/p\mathbb{Z})^n$ where p is a prime number. The only difference with the special case $p = 2$ treated in our preprint [12] is that the formulas get more complicated. As a side remark, note that Simon also gives a Las Vegas version of his algorithm with expected query complexity $O(n)$. Even better, Brassard and Høyer [7] have given an "exact polynomial time" quantum algorithm for Simon's problem (i.e., their algorithm has a polynomial worst case running time and zero probability of error).

The two main methods for proving query complexity lower bounds in quantum computing are the adversary method of Ambainis and the polynomial method (for an excellent review of these methods in french, read [19]). We shall use the polynomial method, which was introduced in quantum complexity theory in [6]. There are recent interesting applications of this method to the collision and element distinctness problem [1, 15]. All previous applications of the polynomial method ultimately rely on approximation theory lemmas of Paturi [18] or Nisan and Szegedy [17].

Besides the application to a new type of problems (namely, the hidden subgroup problems) we also contribute to the development of the method by applying it in a situation where these lemmas are not applicable. Instead, we use an apparently new (and elementary) approximation theory result : Lemma 3 from section 3.

The remainder of this paper is organized follows. After some preliminaries in section 2 we give in section 3 an $\Omega(n)$ lower bound for groups of the form $(\mathbb{Z}/p\mathbb{Z})^n$, where p is a prime number. The general case of arbitrary Abelian groups (lower and upper bound) is treated in section 4. We then proceed to expose elementary lower bounds for other query models than the standard model presented in section 2. Obtaining tight bounds for non-Abelian groups is of course a natural open problem.

2 Preliminaries

From now on, p denotes a prime number and the problem of distinguishing the trivial subgroup from a group of order p in $(\mathbb{Z}/p\mathbb{Z})^n$ will be called "Simon's problem in $(\mathbb{Z}/p\mathbb{Z})^n$" (or sometimes just "Simon's problem"). More precisely, we are given a function f from $G = (\mathbb{Z}/p\mathbb{Z})^n$ to a known set X of size p^n, and we are guaranteed that the function fulfills Simon's promises, that is, either :

(1) f is one-to-one, or

(2) $\exists s \neq 0\ \forall w, w'\ [f(w) = f(w') \iff w - w' \in \langle s \rangle]$, where $\langle s \rangle$ is the group generated by s.

Again, the problem is to decide whether (1) or (2) holds. As pointed out in the introduction, Simon considered only the case $p = 2$.

We assume here that the reader is familiar with the basic notions of quantum computing [16, 9] and we now present the polynomial method. Let A be a quantum algorithm solving Simon's decision problem. Without loss of generality, we can suppose that for every n the algorithm A acts like a succession of operations

$$U_0, O, U_1, O, \ldots, O, U_{T(n)}, M$$

on a m-qubit, for some $m \geq 2n$, starting from state $|0\rangle^{\otimes m}$. The U_i are unitary operations independent of f and O is the call to the black-box function : if x and y are elements of $\{0,1\}^n$ then $O\,|x, y, z\rangle = |x, y \oplus f(x), z\rangle$. The operation M is the measure of the last qubit. There are some states of $(m-1)$-qubits $|\phi_0(f,n)\rangle$ and $|\phi_1(f,n)\rangle$ (of norm possibly less than 1) such that

$$U_{T(n)}OU_{T(n)-1}O\ldots OU_0\,|0\rangle^{\otimes m} = |\phi_0(n,f)\rangle \otimes |0\rangle + |\phi_1(n,f)\rangle \otimes |1\rangle\,.$$

After the measure M, the result is 0 (reject) with probability $||\phi_0(n,f)||^2$ and 1 (accept) with probability $||\phi_1(n,f)||^2$. The algorithm A is said to solve Simon's problem with bounded error probability ϵ if it accepts any bijection with probability at least $1-\epsilon$ and rejects every other function fullfilling Simon's promise with probability at least $1-\epsilon$. By definition, the query complexity of A is the function T. In section 3 we will prove the following lower bound.

Theorem 1 *If A is an algorithm which solves Simon's problem in $(\mathbb{Z}/p\mathbb{Z})^n$ with bounded error probability ϵ and query complexity T, then for every large enough integer n we have :*

$$T(n) \geq \frac{\log_2\left((2-4\epsilon)\frac{p^{n+3}}{p-1}\right)-1}{2\log_2\left(\frac{p^3}{p-1}\right)+2}.$$

Altghough it might not be self-evident that $T(n)=\Omega(n)$, this bound is indeed in the expected range. Indeed, it can be checked easily that the right-hand side is equivalent, for large values of n, to $A(p).n$, where $A(p)$ is positive and $\lim_{p\to+\infty} A(p)=\frac{1}{4}$. For $p=2$ we obtain the result presented in our preprint [12] : $T(n) \geq \frac{n+2+\log_2(2-4\epsilon)}{8}$.

As explained in the introduction, our proof of this theorem is based on the polynomial method. Lemma 1 below is the key observation on which this method relies. We state it using the formalism of [1] : if s is a partial function from $(\mathbb{Z}/p\mathbb{Z})^n$ to X and f a function from $(\mathbb{Z}/p\mathbb{Z})^n$ to X, $|\mathrm{dom}(s)|$ denotes the size of the domain of s. Moreover, we define :

$$I_s(f) = \begin{cases} 1 & \text{if } f \text{ extends } s \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 1 *[6] If A is an algorithm of query complexity T, there is a set S of partial functions from $(\mathbb{Z}/p\mathbb{Z})^n \to E$ such that for all functions $f:(\mathbb{Z}/p\mathbb{Z})^n \to E$, A accepts f with probability*

$$P_n(f) = \sum_{s\in S} \alpha_s I_s(f)$$

where for every $s \in S$ we have $|dom(s)| \leq 2T(n)$ and α_s is a real number.

The goal is now to transform $P_n(f)$ into a low-degree polynomial of a single real variable. This is achieved in Proposition 1. We can then prove and apply our lower bound result on real polynomials (Lemma 3).

3 Lower Bound Proof

An algorithm for Simon's problem is only supposed to distinguish between the trivial subgroup and a hidden subgroup of cardinality p (we recall that p is a prime number). To establish our lower bound, we will nonetheless need to examine its behavior on a black-box hiding a subgroup of arbitrary order (a similar trick is used in [1] and [15]). Note that this "generalized Simon problem" (finding an arbitrary hidden subgroup of $(\mathbb{Z}/p\mathbb{Z})^n$) can still be solved in $O(n)$ queries and bounded probability of error by essentially the same algorithm, see for instance [9].

From now on we suppose that A is an algorithm solving Simon's problem with probability of error bounded by $\epsilon < \frac{1}{2}$ and query complexity T. Moreover, $P_n(f) = \sum_{s \in S} \alpha_s I_s(f)$ as given by Lemma 1.

For $0 \leq d \leq n$ and $D = p^d$, let $Q_n(D)$ be the probability that A accepts f when f is chosen uniformly at random among the functions from $(\mathbb{Z}/p\mathbb{Z})^n$ to X hiding a subgroup of $(\mathbb{Z}/p\mathbb{Z})^n$ of order D. Of course, $Q_n(D)$ is only defined for some integer values of D and it can be extended in many different ways. By abuse of language we will say that Q_n is a polynomial of degree δ if it can be interpolated by a polynomial of degree δ.

The point of this definition is that we have a bound on some values of Q_n, and a gap between two of them. Namely, we have :

1. for any integer $d \in [0; n]$, $0 \leq Q_n(p^d) \leq 1$ (this number is a probability), and
2. $Q_n(1) \geq 1 - \epsilon$ and $Q_n(p) \leq \epsilon$, hence $|Q_n'(x_0)| \geq \frac{1-2\epsilon}{p-1} > 0$ for some $x_0 \in [1; 2]$.

If we denote by X_D the set of functions hiding a subgroup of order D, by Lemma 1 we have $Q_n(D) = \sum_{s \in S} \left(\frac{\alpha_s}{|X_D|} \sum_{f \in X_D} I_s(f) \right)$. Hence

$$Q_n(D) = \sum_{s \in S} \alpha_s Q_n^s(D), \tag{1}$$

where $Q_n^s(D)$ is the probability that a random function f hiding a subgroup of order D extends s. We now prove that Q_n is a low-degree polynomial. By (1), it suffices to bound the degree of Q_n^s. Let us start by counting subgroups :

Lemma 2 *Let n and k be nonnegative integers.*
The group $(\mathbb{Z}/p\mathbb{Z})^n$ has exactly $\beta_p(n,k) = \prod_{0 \leq i < k} \frac{p^{n-i}-1}{p^{k-i}-1}$ distinct subgroups of order p^k.

Proof: We look at $(\mathbb{Z}/p\mathbb{Z})^n$ as a vector space over the field $\mathbb{Z}/p\mathbb{Z}$: from this point of view the subgroups are the subspaces. We start by counting the number of free k-tuples of vectors. For the first v_0, we can choose anything but 0, so there are $p^n - 1$ choices. For the second vector v_1 we can choose any element not in the subspace generated by v_0 ; $p^n - p$ possibilities remain. For the third vector, any linear combinaison of v_0 and v_1 is forbidden : there are p^2 of them. In general, the number of free k-tuples of vectors is $\alpha_p(n,k) = \prod_{0 \leq i < k} (p^n - p^i)$. Each subspace of dimension k can be generated by $\alpha_p(k,k)$ different k-tuples, so the total number of subspaces of dimension k is $\frac{\alpha_p(n,k)}{\alpha_p(k,k)} = \prod_{0 \leq i < k} \frac{p^{n-i}-1}{p^{k-i}-1}$. Note that this formula is correct even if $k > n$, in which case $\alpha_p(n,k) = 0$. □

Proposition 1 *The polynomial Q_n is of degree at most $2T(n)$.*

Proof: By (1), it suffices to show that for all partial functions $s : (\mathbb{Z}/p\mathbb{Z})^n \to E$ such that $|\mathrm{dom}(s)| \leq 2T(n)$, the probability $Q_n^s(D)$ that a random function f hiding a subgroup of order D extends s is a polynomial in D of degree at most $2T(n)$. So, let s be such a partial function. We will proceed in three steps : we first examine the case where s is a constant function, then the case where s is injective and finally the general case.

Let us therefore suppose that s is constant and note $\mathrm{dom}(s) = \{a_i / i = 1 \ldots k\}$, with $k \leq 2T(n)$, the a_i's being of course all different. A function f hiding a subgroup H extends

s if and only if $\{a_i - a_1 / i = 1 \dots k\} \subseteq H$ and $f(a_1) = s(a_1)$. So $Q_n^s(D) = Q_n^{s'}(D)$ where $s'(x) = s(x - a_1)$. We will thus suppose without loss of generality that $a_1 = 0$. Since E, the possible range for f, is of size p^n, we have $Q_n^s(D) = \frac{\lambda}{p^n}$, where λ is the proportion, among the subgroups of order D, of those containing $\mathrm{dom}(s)$. Let H' be the subgroup generated by $\mathrm{dom}(s)$, and $D' = p^{d'}$ its order, d' being the dimension of H' as a vector space. The number of subgroups of order D containing H' is equal to the number of subgroups of order $\frac{D}{D'}$ of $(\mathbb{Z}/p\mathbb{Z})^n / H'$, which is isomorphic to $(\mathbb{Z}/p\mathbb{Z})^{n-d'}$; so there are $\beta(n-d', d-d')$ of them. We then have $Q_n^s(D) = \frac{1}{p^n} \frac{\beta(n-d', d-d')}{\beta(n,d)} = \frac{1}{p^n} \prod\limits_{0 \le i < d'} \frac{p^{d-i}-1}{p^{n-i}-1}$, which is a polynomial in D of degree $d' < |\mathrm{dom}(s)| \le 2T(n)$.

Let us now suppose that s is injective. We still note in the same way $\mathrm{dom}(s) = \{a_i / i = 1 \dots k\}$. A function f hiding a subgroup H extends s if and only if the a_i's lie in distinct cosets of H and f takes appropriate values on these cosets ; so $Q_n^s(D) = \nu\lambda$, where λ is the probability for a subgroup H of order D to contain none of the $a_i - a_j (i \neq j)$ and ν is the probability to extend s for a function h hiding a subgroup H of order D that does not contain any of the $a_i - a_j (i \neq j)$. First we compute ν. For each subgroup H of order D that does not contain any of the $a_i - a_j (i \neq j)$ there are $(p^n)(p^n - 1) \dots (p^n - p^n/D + 1)$ possible functions f : choose a different value for each coset of H. Among these functions, the number of them extending s is $(p^n - k)(p^n - k - 1) \dots (p^n - p^n/D + 1)$: choose a value for each coset not containing any a_i. So $\nu = \frac{(p^n-k)!}{(p^n)!}$. The probability λ is equal to $1 - \mu$, where μ is the probability for a subgroup H of order D to contain some $a_i - a_j$ for some $i \neq j$.

By the inclusion-exclusion formula, we can expand λ as follows :

$$\lambda = 1 - \left[\begin{array}{ll} & \sum\limits_{i \neq j} \Pr(a_i - a_j \in H) \\ - & \sum\limits_{\substack{i_1 \neq j_1 \\ i_2 \neq j_2 \\ \{i_1; j_1\} \neq \{i_2; j_2\}}} \Pr(a_{i_1} - a_{j_1} \in H \wedge a_{i_2} - a_{j_2} \in H) \\ + & \cdots \\ - & \cdots \\ \vdots & \\ + & \Pr(\forall i \neq j \; a_i - a_j \in H) \end{array} \right]$$

Our study of the first case above shows that each term in this sum is a polynomial in D of degree less than d', where the order of the subgroup generated by the $a_i - a_j$'s is $p^{d'}$. Since $a_i - a_j$ is always in the subgroup generated by $\mathrm{dom}(s)$, $d' \le |\mathrm{dom}(s)| \le 2T(n)$.

Finally, in the general case the partial function s is defined by conditions of the form

$$\left\{ \begin{array}{c} s(a_1^1) = s(a_2^1) = \cdots = s(a_{k_1}^1) = b_1 \\ s(a_1^2) = s(a_2^2) = \cdots = s(a_{k_2}^2) = b_2 \\ \vdots \\ s(a_1^l) = s(a_2^l) = \cdots = s(a_{k_l}^l) = b_l \end{array} \right.$$

with $b_1, \dots, b_l$ all different. In the same way as before, we will suppose without loss of generality that $a_1^1 = 0$. Furthermore, since $f(a_i^j) = f(a_1^j)$ is equivalent to $f(a_i^j - a_1^j) = f(0)$ (i.e. a_i^j and a_1^j are in the same coset of H) we can remove each a_i^j, for $i, j > 1$ from $\mathrm{dom}(s)$ and replace them by adding the point $a_i^j - a_1^j$ to $\mathrm{dom}(s)$ associated to the value b_1. The size of $\mathrm{dom}(s)$ does not increase. It may happen that s was already defined on one of these entries and that our new definition is contradictory. In that case there is simply no subgroup-hiding function f extending s, so Q_n^s is simply the null polynomial and we are done. We will therefore consider only conditions of the form :

$$\begin{cases} s(0) = s(a_2^1) = \cdots = s(a_{k_1}^1) = b_1 \\ s(a^2) = b_2 \\ \vdots \\ s(a^l) = b_l \end{cases}$$

The probability $Q_n^s(D)$ that a function f hiding a subgroup of order D extends s is the probability Q_1 that f satisfies $f(0) = f(a_2^1) = \cdots = f(a_{k_1}^1) = b_1$ times the probability Q_2 that f extends s given that $f(0) = f(a_2^1) = \cdots = f(a_{k_1}^1) = b_1$. We have already computed the first probability : this is the case where s is constant. Let H' be the subgroup generated by the a_i^1's and $D' = p^{d'}$ its order ; then $Q_1 = \frac{1}{p^n} \prod\limits_{0 \le i < d'} \frac{p^{d-i}-1}{p^{n-i}-1}$. Let us define s' on G/H' as the quotient of s if it exists (if not, this means again that Q_n^s is the null polynomial, and we are done). If f satisfies $f(0) = f(a_2^1) = \cdots = f(a_{k_1}^1) = b_1$ then we can define f' on G/H' as the quotient of f ; the condition "f extends s and hides a subgroup of order D" is equivalent to "f' extends s' and hides a subgroup of order D/D'". Since s' is defined by the condition $s'(H') = b_1, s'(a^2 + H') = b_2, \ldots, s'(a^l + H') = b_l$ and is injective, our study of the second case shows that $Q_2 = Q_n^{s'}(D/D')$ is a polynomial in D of degree less than $|\mathrm{dom}(s')|$. Hence, $Q_n^s(D)$ is a polynomial in D of degree at most $d' + |\mathrm{dom}(s')| \le |\mathrm{dom}(s)| \le 2T$. $\square$

Now that we have an upper bound on the degree of Q, let us find a lower bound. The following analogue of the lemmas of Paturi [18] and Nisan-Szegedy [17] will help.

Lemma 3 *Let $c > 0$ and $\xi > 1$ be constants and P a polynomial with the following properties :*

1. *For any integer $0 \le i \le n$ we have $\left|P(\xi^i)\right| \le 1$.*
2. *For some real number $1 \le x_0 \le \xi$ we have $|P'(x_0)| \ge c$.*

Then $\deg(P) = \Omega(n)$, and more precisely : $\deg(P) \ge \min\left(\frac{n}{2}, \frac{\log_2(\xi^{n+3}c)-1}{\log_2\left(\frac{\xi^3}{\xi-1}\right)+1}\right)$.

Proof: Let d be the degree of P, and let us write $P'(X) = \lambda \prod\limits_{i=1}^{d-1}(X - \alpha_i)$, where the α_i's are real or complex numbers. The polynomials P' and P'' are respectively of degree $d-1$ and $d-2$, so there exists an integer $a \in [n-2d+2; n-1]$ such that P'' has no real root in $\left(\xi^a; \xi^{a+1}\right)$, and P' has no root whose real part is in this same interval. If $d \ge n/2$ there is nothing to prove, so we may and we will assume that $d \le \frac{n}{2}$. This implies in particular that $\xi^a \ge \xi^2$.

The polynomial P' is monotone on $\left(\xi^a; \xi^{a+1}\right)$, for P'' has no root in it. This means that P is either convex or concave on this interval, so that the graph of P is either over or under its tangent at the middle point of the interval, which is equal to $\frac{\xi^a + \xi^{a+1}}{2} = \frac{1+\xi}{2}\xi^a$. Suppose that $P'\left(\frac{1+\xi}{2}\xi^a\right)$ is nonnegative (the case when it is negative is similar). Then P is increasing on $\left(\xi^a; \xi^{a+1}\right)$, since P' has no root in this interval. Let $y = t(x)$ be the equation of the tangent of P at $\frac{1+\xi}{2}\xi^a$. If $t\left(\xi^{a+1}\right) > 1$, then $P\left(\xi^{a+1}\right) < t\left(\xi^{a+1}\right)$, so P is concave on $\left(\xi^a; \xi^{a+1}\right)$, hence $-1 \le P\left(\xi^a\right) \le t\left(\xi^a\right)$. But, since P is monotone on $\left(\xi^a; \xi^{a+1}\right)$, $t\left(\frac{1+\xi}{2}\xi^a\right) = P\left(\frac{1+\xi}{2}\xi^a\right) \le 1$. Since $t(\xi^{a+1}) - t\left(\frac{1+\xi}{2}\xi^a\right) = t\left(\frac{1+\xi}{2}\xi^a\right) - t(\xi^a)$, it follows that $t\left(\xi^{a+1}\right) \le 3$ and $t\left(\xi^{a+1}\right) - t\left(\xi^a\right) \le 4$. The same inequality can also be derived if we assume $t\left(\xi^a\right) < -1$, and it is of course still true if $t\left(\xi^a\right) \ge -1$ and $t\left(\xi^{a+1}\right) \le 1$.

We conclude that the inequality $t\left(\xi^{a+1}\right) - t\left(\xi^{a}\right) \leq 4$ always holds, which implies that $0 \leq P'\left(\frac{1+\xi}{2}\xi^a\right) \leq \frac{4}{\xi^a(\xi-1)}$. If we now include the case where P' is negative, we obtain the inequality

$$\left|P'\left(\frac{1+\xi}{2}\xi^a\right)\right| \leq \frac{4}{\xi^a(\xi-1)}.$$

We therefore have

$$\left|\frac{P'\left(\frac{1+\xi}{2}\xi^a\right)}{P'(x_0)}\right| \leq \frac{4}{c\xi^a(\xi-1)} \leq \frac{4}{c\xi^{n-2d+2}(\xi-1)}. \tag{2}$$

To conclude we need to state a simple geometric fact. Let MBC be a triangle, M' the orthogonal projection of M onto (BC), and (d) the perpendicular bissector of $[BC]$. Let us suppose that M is "at the right of (d)", i.e. $MC \leq MB$.

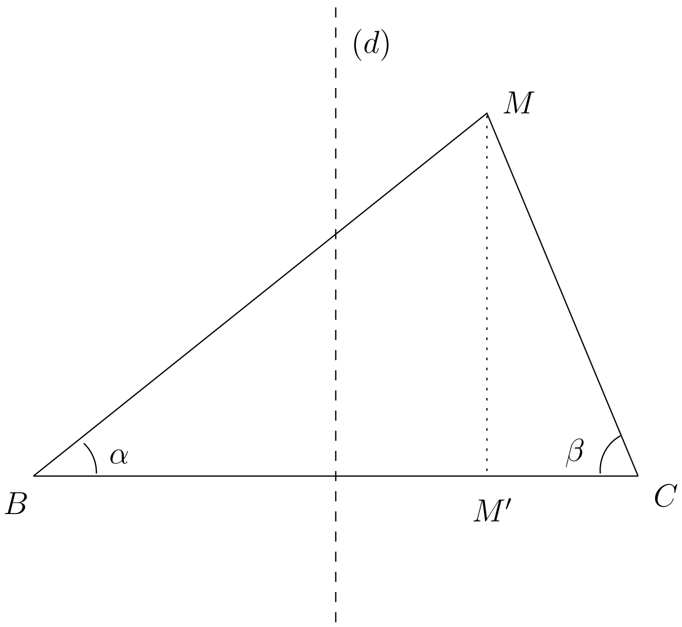

Since C is closer to the line (MM') than B, $\tan\alpha = MM'/BM' \leq \tan\beta = MM'/CM'$. Hence $\alpha \leq \beta$, and $\cos\alpha \geq \cos\beta$, i.e. :

$$\frac{MC}{MB} \geq \frac{M'C}{M'B}. \tag{3}$$

Let $f : \begin{pmatrix} \mathbb{R}\setminus\{x_0\} & \rightarrow & \mathbb{R} \\ x & \mapsto & \left|\frac{\frac{1+\xi}{2}\xi^a - x}{x_0 - x}\right| \end{pmatrix}$. Since $x_0 < \xi^a < \frac{1+\xi}{2}\xi^a < \xi^{a+1}$, a quick study of this function shows that for all $x \in \mathbb{R}\backslash\left(\{x_0\} \cup \left(\xi^a;\xi^{a+1}\right)\right)$, $f(x) \geq \min(1, f(\xi^a), f(\xi^{a+1})) \geq \frac{\xi-1}{2\xi}$.

We will distinguish two cases for each $i \in \{1;\ldots;d-1\}$.

1. If $\Re(\alpha_i) \leq \frac{1}{2}\left(\frac{1+\xi}{2}\xi^a + x_0\right)$, then $\left|\frac{\frac{1+\xi}{2}\xi^a - \alpha_i}{x_0 - \alpha_i}\right| \geq 1$.
2. If $\Re(\alpha_i) > \frac{1}{2}\left(\frac{1+\xi}{2}\xi^a + x_0\right)$, let us apply (3) to the points $M = \alpha_i$, $M' = \Re(\alpha_i)$, $B = x_0$ and $C = \frac{1+\xi}{2}\xi^a$. We obtain the inequality

$$\left|\frac{\frac{1+\xi}{2}\xi^a - \alpha_i}{x_0 - \alpha_i}\right| \geq \left|\frac{\frac{1+\xi}{2}\xi^a - \Re(\alpha_i)}{x_0 - \Re(\alpha_i)}\right|.$$

Remember though that no root of P' has its real part in $\left(\xi^a;\xi^{a+1}\right)$, so that $\left|\frac{\frac{1+\xi}{2}\xi^a-\alpha_i}{x_0-\alpha_i}\right| \geq \frac{\xi-1}{2\xi}$.

We conclude that $\left|\frac{\frac{1+\xi}{2}\xi^a-\alpha_i}{x_0-\alpha_i}\right| \geq \frac{\xi-1}{2\xi}$ in both cases. Taking (2) into account, we finally obtain the inequality $\left(\frac{\xi-1}{2\xi}\right)^{d-1} \leq \frac{4}{c\xi^{n-2d+2}(\xi-1)}$, hence $d \geq \frac{\log_2\left(\xi^{n+3}c\right)-1}{\log_2\left(\frac{\xi^3}{\xi-1}\right)+1}$. $\square$

We can now complete the proof of Theorem 1. Let A be our algorithm solving Simon's problem with bounded error probability ϵ and query complexity T. As pointed out before Lemma 2, the associated polynomial Q_n satisfies $|Q_n'(x_0)| \geq 1-2\epsilon$ for some $x_0 \in [1,\xi]$ and $Q_n(\xi^i) \in [0,1]$ for any $i \in \{0,1,\ldots,n\}$. An application of Lemma 3 to the polynomial $P = 2Q_n - 1$ therefore yields the inequality $\deg(Q_n) \geq \min\left(\frac{n}{2}, \frac{\log_2\left((2-4\epsilon)\frac{p^{n+3}}{p-1}\right)-1}{\log_2\left(\frac{p^3}{p-1}\right)+1}\right)$. Theorem 1 follows since $\deg(Q_n) \leq 2T(n)$ by Proposition 1.

4 Abelian groups

In this section we give lower and upper bounds for the quantum query complexity of Abelian hidden subgroup problems. As explained in the introduction, our bounds are optimal up to constant factors.

Let G be a finite Abelian group, $\hat{G}$ its dual group, i.e. the group of its characters (see for example [9]). For each subgroup H of G, we note $H^\perp$ the orthogonal of H, which is a subgroup of $\hat{G}$ consisted of those characters χ such that $\chi(h) = 1$ for all $h \in H$. According to basic representation theory, $\hat{G}$ is isomorphic to G and, for all subgroup $H \leq G$, the index of $H^\perp$ in $\hat{G}$ is equal to the order of H.

The well-established method of Fourier sampling allows one, with one query to the black-box function, to pick a uniformly random element of the orthogonal of the hidden subgroup. In order to solve the hidden subgroup problem for G, this routine is run k times so as to generate k random elements $x_1,\ldots,x_k \in H^\perp$. The algorithm outputs the orthogonal of the group generated by $x_1,\ldots,x_k$. This output is correct if $x_1,\ldots,x_k$ generate all of $H^\perp$.

We will now show that this algorithm is optimal if we know when to stop, i.e., how many random elements should be picked in $H^\perp$. The following lemma implies that the query complexity of the cyclic subgroup problem is constant. Note that this fact is already pointed out (without proof) in [20]. We give the proof here for the sake of completeness.

Lemma 4 *For any integer $M \geq 1$, two random elements chosen uniformly and independently in $\mathbb{Z}/M\mathbb{Z}$ generate all of this group with probability at least $\frac{1}{2}$.*

Proof: Let us write $M = \prod_{i=1}^{n} p_i^{\alpha_i}$ where the p_i's are distinct primes. Let $x_1,\ldots,x_k$ be k elements of $\mathbb{Z}/M\mathbb{Z}$. These elements generate all of $\mathbb{Z}/M\mathbb{Z}$ iff for each $i \in \{1,\ldots,n\}$ there exists $j \in \{1,\ldots,k\}$ such that p_i does not divide x_j. Let X_i, for $i = 1,\ldots,n$, be the random variable which, to a random element x of $\mathbb{Z}/M\mathbb{Z}$, associates 0 if p_i divides x, and 1 otherwise. It is easily verified that the X_i's are independent random variables (for instance, $\mathbb{P}[X_i = 0 \wedge X_j = 0] = \mathbb{P}[X_i = 0]\ \mathbb{P}[X_j = 0] = \frac{1}{p_i}\ \frac{1}{p_j}$ for $i \neq j$). The probability $\mathcal{P}(M,k)$ that the x_j's generate $\mathbb{Z}/M\mathbb{Z}$ is therefore equal to the product over the p_i's of the probabilities that p_i does not divide all of the x_j's. Namely, $\mathcal{P}(M,k) = \prod_{i=1}^{n}\left(1-p_i^{-k}\right)$.

Note that $\log_2 \mathcal{P}(M,k) = \sum_{i=1}^{n} \log_2 \left(1 - p_i^{-k}\right) \geq -2 \sum_{i=1}^{n} p_i^{-k}$. Let $\mathbb{P} = \{2,3,5,\ldots\}$ be the set of prime numbers and let $k_1 \in \mathbb{N}$ be such that $\sum_{p\in\mathbb{P}} p^{-k_1} \leq -\frac{\log_2\left(1-\frac{1}{2}\right)}{2} = \frac{1}{2}$. Using the fact that $\sum_{n\in\mathbb{N}^*} n^{-2} = \frac{\pi^2}{6}$, it can be easily verified that $k_1 = 2$ is suitable. Then $\mathcal{P}(M,2) \geq \frac{1}{2}$ and we are done. □

We recall that (following for instance [14]) the rank $r(G)$ of a group G is the minimal cardinality of a generating set of G. According to the fundamental theorem of finite Abelian groups, G is isomorphic to $\mathbb{Z}/m_1\mathbb{Z} \times \mathbb{Z}/m_2\mathbb{Z} \times \cdots \times \mathbb{Z}/m_{r(G)}\mathbb{Z}$ where m_i divides m_{i-1} for every $i \in \{2,\ldots,r(G)\}$, and this decomposition is unique.

Proposition 2 *For any $\epsilon > 0$ there exists an integer k such that for any finite Abelian group G, $k.r(G)$ random elements chosen uniformly and independently in G generate all of this group with probability at least $1-\epsilon$.*

Proof: Let us denote by $\mathcal{E}_n$ the supremum of the expectations of the number of random elements of G needed to generate G, taken over the groups G such that $r(G) \leq n$. We can assume that $G = \mathbb{Z}/m_1\mathbb{Z} \times \cdots \times \mathbb{Z}/m_{r(G)}\mathbb{Z}$, where $m_{r(G)}|\ldots|m_1$. To generate G we can proceed with the two following steps.

First we pick enough random elements $\left(x_1^1,\ldots,x_1^{r(G)}\right),\ldots,\left(x_k^1,\ldots,x_k^{r(G)}\right)$ in G so that $x_1^1,\ldots,x_k^1$ generate $\mathbb{Z}/m_1\mathbb{Z}$; the expectation of k is at most $\mathcal{E}_1$. By Lemma 4, $\mathcal{E}_1$ is finite ; we can very roughly bound it in the following way.

First pick two random elements in $\mathbb{Z}/m_1\mathbb{Z}$. With probability $p_{\leq 2}$ they generate $\mathbb{Z}/m_1\mathbb{Z}$ and with probability $p_{>2}$ they do not ; when they fail to generate, just forget about them and renew the experiment with two new random elements. In the first case the expectation of the number of elements is 2, in the second case it is at most $2+\mathcal{E}_1$, so we have $\mathcal{E}_1 \leq 2p_{\leq 2} + (2+\mathcal{E}_1)\,p_{>2}$. Clearly $p_{\leq 2} + p_{>2} = 1$ and according to Lemma 4 we have $p_{\leq 2} \geq \frac{1}{2}$. This shows that $\mathcal{E}_1 \leq 4$.

Then the subgroup generated by these elements contains some element $y = \left(y^1,\ldots,y^{r(G)}\right)$ such that the of order of y^1 is m_1. The rank of $G/\langle y\rangle$ is equal to $r(G)-1$ since $G/\langle y\rangle$ is isomorphic to $\mathbb{Z}/m_2\mathbb{Z}\times\cdots\times\mathbb{Z}/m_{r(G)}\mathbb{Z}$. This isomorphism follows from the fact the classes of $e_2,\ldots,e_{r(G)}$ generate $G/\langle y\rangle$, where e_i denotes the element of G whose i^{th} coordinate is equal to 1 and all other coordinates equal to 0. We now pick enough random elements $x_{k+1},\ldots,x_{k+l} \in G$ so that their images in $G/\langle y\rangle$ generate all of it ; the expectation of l is of course at most $\mathcal{E}_{r(G)-1}$. Putting it together, we get $\mathcal{E}_{n+1} \leq \mathcal{E}_1 + \mathcal{E}_n$, so $\mathcal{E}_n \leq 4n$. By Markov's inequality, if we choose $\left\lfloor \frac{4}{\epsilon} \right\rfloor r(G)$ random elements in a group G, we generate all of this group with probability at least $1-\epsilon$. □

We can now prove our main result.

Theorem 2 *The quantum query complexity of the hidden subgroup problem in a finite Abelian group G is $\Theta(r(G))$.*

Proof: The upper bound is achieved with the standard method : one just applies Proposition 2 to the orthogonal of the hidden subgroup, which is isomorphic to a subgroup of G, using the fact that r is an nondecreasing function on finite Abelian groups.

The lower bound of course comes from Theorem 1. Since for every finite Abelian group G there is some prime p such that $(\mathbb{Z}/p\mathbb{Z})^{r(G)}$ is isomorphic to some subgroup of G, we need only to state that the hidden subgroup problem for a subgroup of G reduces correctly to the hidden subgroup problem for G. Indeed, let H be a subgroup of G and let $H+t_0,\ldots,H+t_k$

be the cosets of H in G, where $t_0 = 0$. If $\gamma : H \to X$ hides a subgroup of H, we can define a function $\gamma' : G \to X \times \{t_i/0 \leq i \leq k\}$ which hides the same subgroup. Namely, we define $\gamma'(x + t_i) = (\gamma(x), t_i)$ for $x \in H$. Moreover, a call to γ' uses just one call to γ, so we are done.

$\square$

5 Other query models

We will consider two other (weaker) query models, the *test model* and the *collision model*. The test model was introduced in the context of quantum computing in [19]. A *comparison model* similar in spirit to our collision model is studied in [2].

5.1 The collision model

In the standard query model, the black box outputs $F(x)$ on input x. This model is formally defined in section 2, and used in the first four sections of this paper. In the collision model, the black box can only test whether $F(x) = F(y)$ for two inputs x and y. This model would seem at first rather natural for hidden subgroup problems since the actual values taken by F do not matter. It is only the fact that F takes distinct values on distinct cosets that matters. Nevertheless, we shall see that the query complexity of hidden subgroup problems can be much higher in this model than in the standard model.

The collision model can be formally defined as follows. As in section 2, we describe an algorithm A as a succession of operations

$$U_0, O, U_1, O, \ldots, O, U_{T(n)}, M$$

on a m-qubit, starting from state $|0\rangle^{\otimes m}$.

The U_i are unitary operations independent of f and O is the call to the black-box function : if x and y are elements of G then $O\,|x, y, z, t\rangle = \left|x, y, z \oplus \delta_{F(x)F(y)}, t\right\rangle$. The operation M is the measure of the last qubit.

Proposition 3 *Let G be a group containing n subgroups $H_1, \ldots, H_n$ such that $H_i \cap H_j = \{0\}$ for $i \neq j$, and $H_i \neq \{0\}$ for all i. In the collision model, the query complexity of the hidden subgroup problem for G is $\Omega(\sqrt{n})$.*

Proof
We proceed by reduction from the search problem in an unordered list of n boolean items, which admits a well-known $\Omega(\sqrt{n})$ lower bound ([6, 4]). More precisely, let $f : \{1, \ldots, n\} \to \{0, 1\}$ be a function which is either identically zero, or takes the value 1 at a single point i_0. To such a function we associate a function $F : G \to \mathbb{N}$ which hides the trivial subgroup if f is identically zero, and hides H_{i_0} if $f(i_0) = 1$. If we have access to a black-box for f we can easily simulate the collision black-box for F since $F(x) = F(y) \Leftrightarrow x - y \in H_{i_0}$. To decide whether $x - y \in H_{i_0}$, we first determine whether $x - y$ belongs to one of the groups H_i. If not, we know that $F(x) \neq F(y)$. If $x - y$ does belong to one of these groups, i is unique by the hypothesis on G if $x - y \neq 0$. If $x - y = 0$ we give of course a positive answer to the collision query. If $x - y \neq 0$, we give a positive answer iff $f(i) = 1$. To answer one collision query we thus need to perform a single call to f. An algorithm which decides whether F hides the trivial subgroup in T collision queries can therefore be turned into an algorithm which determines in T queries whether f is identically zero. $\square$

Note that the proof does not use the hypothesis that G is Abelian.

Corollary 1 *In the collision model, the query complexity of the hidden subgroup problem for $(\mathbb{Z}/2\mathbb{Z})^n$ is $\Theta(\sqrt{2^n})$. For $\mathbb{Z}/N\mathbb{Z}$, the query complexity is $\Omega(\sqrt{n})$, where n is the number of prime factors of N.*

Proof Let $p_1, \ldots, p_n$ be the prime factors of N. In $\mathbb{Z}/N\mathbb{Z}$ there is exactly one subgroup of order p_i, and these n subgroups have pairwise trivial intersections. In $(\mathbb{Z}/2\mathbb{Z})^n$, there are $2^n - 1$ subgroups of order two. $\square$

5.2 The test model

In the test model, a black box decides whether $F(x) = y$ given two inputs x and y. The formal definition of this model is identical to that of the collision model, except that the gate O is now defined by $O\,|x, y, z, t\rangle = \big|x, y, z \oplus \delta_{F(x)y}, t\big\rangle$.

The following lower bound is probably far from optimal, but suffices to separate the test model from the standard query model (see [19] for other examples).

Proposition 4 *In the test model the query complexity of the hidden subgroup problem for $\mathbb{Z}/N\mathbb{Z}$ is $\Omega(\log n)$, where n is the number of prime factors of N.*

Proof
Let $p_1 < p_2 \ldots < p_n$ be the prime factors of N and H_i the subgroup of $\mathbb{Z}/N\mathbb{Z}$ generated by $\frac{N}{p_i}$. We proceed by reduction from the search problem in an ordered list of n elements, which admits a $\Omega(\log n)$ lower bound [3]. Let $f : \{1, \ldots, n\} \to \{0, 1\}$ be a function such that $f(i) = 1$ iff $i \geq i_0$, where $i_0 \in \{1, \ldots, n\}$. To such a function we associate the function

$$F : \begin{pmatrix} \mathbb{Z}/N\mathbb{Z} & \to & \mathbb{Z}/N\mathbb{Z} \\ k & \mapsto & k \mod \frac{N}{p_{i_0}} \end{pmatrix},$$

which hides the subgroup H_{i_0}. In order to answer a query of the form "$F(x) = y\,$?" using a bounded number of calls to f, we distinguish the following cases.

1. If $y > x$ we always answer "no" the query "$F(x) = y\,$?".
2. If $y < x$, there is at most one i such that $x - y \in H_i$. If there is no such i, we answer "no" to the collision query. If there is such an i, we can test with at most 2 calls to f whether $i = i_0$. If $i \neq i_0$, we answer "no". If $i = i_0$, we accept iff $y < N/p_i$.
3. If $y = x$, we should answer "yes" iff $x < N/p_{i_0}$. If there is no i such that $x < N/p_i$, we may therefore answer "no". Otherwise, let i_1 be the biggest such i. Then the answer is "yes" iff $i_0 \leq i_1$, that is iff $f(i_1) = 1$.

An algorithm which finds in T test queries the subgroup hidden by F can therefore be turned into an algorithm of query complexity $O(T)$ which finds $i_0 = \min\{i;\ f(i) = 1\}$. $\square$

We conjecture that there is in the test model an $\Omega(\sqrt{2^n})$ lower bound for the query complexity of the hidden subgroup problem in $(\mathbb{Z}/2\mathbb{Z})^n$. Unfortunately, there does not seem to be any straightforward way of adapting the techniques of this section to obtain such a lower bound. For instance, if one tries to mimic the proof of Proposition 3 it is natural to define $F(x) = \min\{x, x + i_0\}$ where $i_0 = f^{-1}(1)$, or $i_0 = 0$ if f is identically zero. It is however not clear how one could answer a query of the form "$F(x) = x\,$?" with a constant number of calls to f.

Acknowledgments

Thanks go to Xavier Caruso, Yves de Cornulier, Joël Riou and Frédéric Magniez for useful help and bibliographical hints.

Références

[1] Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4) :595–605, July 2004.

[2] Andris Ambainis. Quantum walk algorithm for element distinctness. http ://www.arxiv.org/pdf/quant-ph/0311001.

[3] Andris Ambainis. A better lower bound for quantum algorithms searching an ordered list. In *FOCS '99 : Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, page 352. IEEE Computer Society, 1999.

[4] Andris Ambainis. Quantum lower bounds by quantum arguments. *J. Comput. Syst. Sci.*, 64(4) :750–767, 2002.

[5] R. Beals. Quantum computation of Fourier transforms over symmetric groups. In *Proceedings of the 29th Annual ACM Symposium on the Theory of Computation (STOC)*, pages 48–53. ACM Press, 1997.

[6] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4) :778–797, 2001.

[7] Gilles Brassard and Peter Høyer. An exact quantum polynomial-time algorithm for Simon's problem. In *Israel Symposium on Theory of Computing Systems*, pages 12–23, 1997.

[8] Lisa R. Hales. *The Quantum Fourier Transform and Extensions of the Abelian Hidden Subgroup Problem*. PhD thesis, UC Berkeley, 2002.

[9] Mika Hirvensalo. *Quantum Computing (Natural Computing Series)*. SpringerVerlag, 2001.

[10] Peter Høyer. Conjugated operators in quantum algorithms. *Phys. Rev. A*, 59 :3280–3289, may 1999.

[11] R. Jozsa. Quantum algorithms and the Fourier transform. *Proc. R. Soc. of London A*, 454, 1998.

[12] Pascal Koiran, Vincent Nesme, and Natacha Portier. A quantum lower bound for the query complexity of Simon's problem. http ://www.arxiv.org/pdf/quant-ph/0501060.

[13] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. Quantum Physics e-Print Archive, 2003.

[14] Hans Kurzweil and Bernd Stellmacher. *The Theory of Finite Groups, An Introduction*. Universitext. Springer, 2004.

[15] Samuel Kutin. Quantum lower bound for the collision problem. *quant-ph/0304162*, 2003.

[16] Michael A. Nielsen and Isaac L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2000.

[17] Noam Nisan and Mario Szegedy. On the degree of boolean functions as real polynomials. *Comput. Complex.*, 4(4) :301–313, 1994.

[18] Ramamohan Paturi. On the degree of polynomials that approximate symmetric boolean functions (preliminary version). In *STOC '92 : Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*, pages 468–474, 1992.

[19] Pierre Philipps. Bornes inférieures en calcul quantique : Méthode par adversaire vs. méthode des polynômes. Rapport de stage de DEA, effectué au LRI sous la direction de Frédéric Magniez, http ://www.lri.fr/˜magniez/stages-dea.html, 2003.

[20] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5) :1484–1509, 1997.

[21] David R. Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5) :1474–1483, 1997.