

The quantum query complexity of the hidden subgroup problem is polynomial

Mark Ettinger

*Los Alamos National Laboratory**

Peter Høyer

University of Calgary†

Emanuel Knill

*Los Alamos National Laboratory**

January 12, 2004

Abstract

We present a quantum algorithm which identifies with certainty a hidden subgroup of an arbitrary finite group G in only a polynomial (in $\log |G|$) number of calls to the oracle. This is exponentially better than the best classical algorithm. However our quantum algorithm requires exponential time, as in the classical case. Our algorithm utilizes a new technique for constructing error-free algorithms for non-decision problems on quantum computers.

1 Introduction

Let G be a finite group, written multiplicatively with identity 1_G . A function f on G (with arbitrary range) is said to be H -periodic if f is constant on the left cosets of a subgroup H of G . If f also takes distinct values on distinct cosets we say f is *strictly* H -periodic and we call H the *hidden subgroup* of f . The *hidden subgroup problem* (HSP) is stated as follows: Given a description of G and a function f on G that is promised to be strictly H -periodic for some subgroup $H \leq G$, find a generating set for H .

Let r denote the number of distinct subgroups of G . Fix any ordering of the r subgroups (K_1, K_2, \dots, K_r) satisfying that $|K_\mu| \geq |K_{\mu+1}|$ for all $1 \leq \mu < r$. In the HSP, we are searching for a generating set for one out of r candidate subgroups. Let $N = |G|$ denote the order of G . We consider $n = \log |G|$ to be the input size. Since any subgroup of G is generated by a set of at most n elements of G , the number r of distinct subgroups of G is $2^{O(n^2)}$.

We assume the function f is given as an oracle so that the only way we can gain knowledge about f is by asking for its value on elements of G . Formally, on a quantum

* Email: {ettinger, knill}@lanl.gov.

† Email: hoyer@cpsc.ucalgary.ca.

computer, the oracle is a unitary operator O_f , that maps $|g\rangle|0\rangle$ to $|g\rangle|f(g)\rangle$ for all $g \in G$. We assume without loss of generality that algorithms for the HSP always output a subset of H . Suppose instead that an algorithm outputs $X \not\subseteq H$. Then we can find the intersection of X with H by evaluating f on each element $x \in X$ and only keeping x if $f(x) = f(1_G)$. This requires at most $|X| + 1$ evaluations of f .

If the group G is *Abelian*, then it is possible to solve the HSP in polynomial time with bounded error on a quantum computer. That is, we can efficiently find a subset $X \subseteq H$ that generates H with probability at least $\frac{2}{3}$. This result follows from the work of Simon [1], Shor [2] and Kitaev [3]. It is possible to improve the success probability to one for Abelian groups of smooth order [4] (a group is of *c-smooth order* if all prime factors of $|G|$ are at most $(\log |G|)^c$ for some constant c). For *non-Abelian* groups, our knowledge is much more limited [5, 6, 7, 8, 9, 10].

The efficient HSP algorithm for Abelian groups of smooth order implies that only a polynomial (in $\log |G|$) number of calls to the oracle are necessary to identify H with certainty. The main result of this paper is that this more limited result holds for all groups of finite order. That is, there exists a quantum algorithm that determines H using a polynomial number of calls to the oracle.

Theorem 1 *There exists a quantum algorithm that, given a finite group G and an oracle f on G promised to be strictly H -periodic for some subgroup $H \leq G$, calls the oracle $O(\log^4 |G|)$ times and outputs a generating set for H . The algorithm fails with probability exponentially small in $\log |G|$. The algorithm can be made exact in any model allowing arbitrary one-qubit gates.*

An important consequence of this result is that it rules out most known methods for proving super-polynomial lower bounds on the total complexity of bounded-error quantum algorithms for the HSP. Most of these methods bound the query complexity, including the recent ones by Aaronson [11] and Shi [12]. This works well for problems where the query complexity is at most poly-logarithmically smaller than the time complexity. Because of our result, one cannot obtain super-polynomial lower bounds on the total complexity of algorithms for the HSP by bounding the query complexity.

Our result extends to exact quantum algorithms (algorithms that determine the answer with certainty) in any model that allows arbitrary one-qubit gates. If allowing only a restricted set of one-qubit gates, our work leaves a hope that one may be able to prove a super-polynomial lower bound on the query complexity for the *exact* case.

A proof of the upper bound on the query complexity only requires establishing the existence of a sufficiently short sequence of unitary operations and oracle calls on appropriately chosen quantum systems. The sequence depends on the group. Our proof explicitly constructs the sequence and makes it apparent how to realize the unitary operations using quantum gates from a universal set. In fact, the sequence can be obtained by means of a (classical) preprocessing algorithm with input a specification of G and whose output is the required sequence of gates and oracle calls. For solving the HSP exactly, the classical preprocessing algorithm requires exact real number arithmetic and access to trigonometric functions of rational angles. The preprocessing algorithms and the quantum networks they compute are inefficient.

2 The algorithm

Our proof of Theorem 1 consists of two stages. In subsection 2.1, we give a quantum algorithm that identifies the correct subgroup with exponentially small error probability, and in subsection 2.2, we then show how to reduce the error probability to zero. We begin with an overview.

We use $2 + 2s$ registers, where s is a positive integer that will be chosen to achieve sufficiently low error probability. The first register is the output register and contains an integer ν (a subgroup index) between 0 and r . The second register is used as a counter and contains an integer ℓ between 0 and r . The remaining $2s$ registers are grouped in s blocks, each consisting of 2 consecutive registers (a “couplet”) called the “subgroup” and the “function” register. Within each couplet, the first register contains an element of G and the second a value in the range of f .

We start by creating the initial state

$$|\Psi_{\text{init}}\rangle = |0\rangle|0\rangle \otimes \left(\frac{1}{\sqrt{N}} \sum_{g \in G} |g\rangle |f(g)\rangle \right)^{\otimes s}. \quad (1)$$

This superposition can be created efficiently using s applications of operator O_f . We then apply the unitary operator Test , to be defined in subsection 2.1, producing the superposition $|\Psi_{\text{final}}\rangle = \text{Test}|\Psi_{\text{init}}\rangle$. We measure the first register of $|\Psi_{\text{final}}\rangle$, yielding some subgroup-index ν as outcome. If $1 \leq \nu \leq r$, we output a generating set for K_ν , otherwise we output $\{1_G\}$, which may be the wrong answer. Our algorithm has exponentially small error probability.

Theorem 2 *Let $\text{Prob}[K_\nu|H]$ denote the probability that the outcome of the measurement of the first register of $|\Psi_{\text{final}}\rangle$ is ν , conditioned on the hidden subgroup being H . Then $\text{Prob}[H|H] \geq 1 - 4r/2^{s/2}$ for all subgroups $H \leq G$, where r is the number of subgroups of G and s is the number of queries. In particular, for $s \in \Theta(n^2 + \log(1/\epsilon))$, the algorithm outputs the correct subgroup with probability at least $1 - 1/\epsilon$.*

The theorem is proved in subsection 2.1, and in subsection 2.2, we make this algorithm exact by precomputing $\text{Prob}[K|H]$ for each subgroup pair (K, H) , adjusting the conditional probabilities to make them more uniform and applying amplitude amplification [13].

2.1 An algorithm with exponentially small error probability

A (left) *translation* for a subgroup K of G is a subset $T \subseteq G$ so that any element $g \in G$ can be written uniquely in the form $g = tk$ for some $t \in T$ and $k \in K$. Fix a translation T_μ for each of the r subgroups K_μ of G .

The operator Test tests the hidden subgroup for each of the r candidate subgroups, one by one. It is defined by

$$\text{Test} = \text{Test}_r \cdot \dots \cdot \text{Test}_2 \cdot \text{Test}_1, \quad (2)$$

where each Test_μ is a unitary operator that tests whether f is K_μ -periodic. If a function is K -periodic, it is also K' -periodic for any proper subgroup K' of K , so we test for bigger subgroups first by requiring that $|K_\mu| \geq |K_{\mu+1}|$ for all $1 \leq \mu < r$. When we find that f is K_μ -periodic for some subgroup K_μ , we record this in the first register, and we begin counting in the second register. For every subgroup $K_\mu \leq G$, let Q_μ be any unitary operator acting on the first two registers that satisfies

$$Q_\mu : \begin{cases} |0\rangle|0\rangle & \mapsto |\mu\rangle|1\rangle \\ |\nu\rangle|\ell\rangle & \mapsto |\nu\rangle|\ell+1\rangle, \end{cases} \quad \text{if } \ell > 0.$$

Once the count ℓ in the second register is increased from its initial value of 0 to 1, the contents of the first register are never changed. The purpose of the counter is to ensure unitarity and that once some test succeeds, no future test affects the contents of the first register.

We test for K_μ -periodicity by acting on the s couplets. If function f is K_μ -periodic then the s subgroup registers are in a superposition of the coset states $|tK_\mu\rangle = \frac{1}{\sqrt{|K_\mu|}} \sum_{k \in K_\mu} |tk\rangle$. Let $P_{s,\mu}$ be the projector of the s couplets defined by

$$P_{s,\mu} = \left(\sum_{t \in T_\mu} |tK_\mu\rangle \langle tK_\mu| \otimes I \right)^{\otimes s},$$

where I denotes the identity operator, and let $P_{s,\mu}^\perp$ denote its complement. Define operator Test_μ by

$$\text{Test}_\mu = Q_\mu \otimes P_{s,\mu} + I \otimes P_{s,\mu}^\perp, \quad (3)$$

which is unitary by construction. Its effect is an application of Q_μ on the first two registers, conditioned on having the s subgroup registers in coset states of K_μ . The condition can be implemented with the help of any pair of unitary operators U_μ and V_μ , where U_μ maps $|1_G\rangle$ to $|K_\mu\rangle$ and V_μ maps $|t\rangle|k\rangle$ to $|1_G\rangle|tk\rangle$ for all $t \in T_\mu$ and $k \in K_\mu$. The procedure is as follows: Adjoin an ancilla register to each subgroup register and apply V_μ^\dagger to these s register pairs. Then apply U_μ^\dagger to the subgroup registers. Next, coherently apply Q_μ if all subgroup registers are in $|1_G\rangle$ and finally reverse the previous steps. It is possible to realize each of these steps with a network of gates of complexity polynomial in N and s .

Lemma 3 *If f is K_μ -periodic, then*

$$\text{Test}_\mu |\Psi_{\text{init}}\rangle = |\mu\rangle|1\rangle \otimes \left(\frac{1}{\sqrt{N}} \sum_{g \in G} |g\rangle |f(g)\rangle \right)^{\otimes s}.$$

Proof We assume in the lemma that f is K_μ -periodic, that is, $f(t) = f(tk)$ for all $t \in T_\mu$ and $k \in K_\mu$, and hence the state $\frac{1}{\sqrt{N}} \sum_{g \in G} |g\rangle |f(g)\rangle = \frac{1}{\sqrt{N}} \sum_{t \in K_\mu} |tK_\mu\rangle |f(t)\rangle$ is in the $+1$ -eigenspace of $P_{1,\mu}$. It follows that $P_{s,\mu}$ acts as the identity on the s couplets, and thus applying operator Test_μ as defined in Eq. 3 on the initial state $|\Psi_{\text{init}}\rangle$ yields the state given on the right hand side in the equation of the lemma. \square

Since we iterate through r tests, we require that if f is not K_μ -periodic, then the state is so marginally altered that it is safe to continue to test for $K_{\mu+1}$ -periodicity.

Lemma 4 *If f is not K_μ -periodic, then the distance $|\langle \text{Test}_\mu | \Psi_{\text{init}} \rangle - |\Psi_{\text{init}} \rangle|$ is at most $\frac{2}{2^{s/2}}$.*

Proof Let H denote the hidden subgroup of f . Consider the case $s = 1$. Then

$$\begin{aligned} |\langle P_{s,\mu} | H \rangle |f(H)\rangle|^2 &= \sum_{t \in T_\mu} |\langle tK_\mu | H \rangle|^2 = \sum_{t \in T_\mu: tK_\mu \cap H \neq \emptyset} |\langle tK_\mu | H \rangle|^2 \\ &= (|H|/|K_\mu \cap H|) |K_\mu \cap H|^2 / (|K_\mu| |H|) = |K_\mu \cap H| / |K_\mu| \leq \frac{1}{2}. \end{aligned}$$

It follows that for arbitrary s , the amplitude squared of $(Q_\mu \otimes P_{s,\mu}) |\Psi_{\text{init}}\rangle$ is upper bounded by $(\frac{1}{2})^s$. Since Test_μ acts trivially on the orthogonal component $(I \otimes P_{s,\mu}^\perp) |\Psi_{\text{init}}\rangle$, the result follows. \square

For each $1 \leq j \leq r$, let $|\Psi_j\rangle = \text{Test}_j \cdots \text{Test}_1 |\Psi_{\text{init}}\rangle$ denote the state of the system after j tests. By the above lemma, it is safe to iterate through all tests, since distances can add up only linearly.

Lemma 5 *If f is not K_μ -periodic for any $1 \leq \mu \leq j$, then the distance $|\langle \Psi_j | \Psi_{\text{init}} \rangle|$ is at most $\frac{2j}{2^{s/2}}$.*

Suppose that the input function f is strictly K_ν -periodic. Then, by Lemma 5, the state $|\Psi_{\nu-1}\rangle$ just prior the test Test_ν is at most at a distance $\epsilon = \frac{2r}{2^{s/2}}$ away from the initial state $|\Psi_{\text{init}}\rangle$. Thus the probability that test Test_ν fails in producing the correct answer $|\nu\rangle$ in the first register is at most $2\epsilon = \frac{4r}{2^{s/2}}$ by Lemma 3.

We note that operator Test never acts on the s function registers. One can therefore measure these prior to the application of Test without affecting the error probability of the bounded error algorithm. However, our exact algorithm requires unitarity and assumes that the function registers are not measured.

The probability of measuring the outcome μ depends on which subgroup H is the hidden subgroup, but it is independent of the values f takes on the different cosets of H . That is, for any two functions f and f' having the same hidden subgroup H , the probabilities of measuring μ are the same. We may therefore let $\text{Prob}[K_\mu | H]$ denote the probability that μ is the outcome of measuring the first register of $|\Psi_{\text{final}}\rangle$, conditioned on the hidden subgroup being H . Theorem 2 follows.

2.2 An exact algorithm

We next use amplitude amplification to make our algorithm exact. This requires the ability to compute exactly the conditional probabilities $\text{Prob}[K_\mu | H]$ without using the oracle. One method for computing $\text{Prob}[K_\mu | H]$ is to pick an arbitrary function f that is strictly H -periodic, and simulate the quantum computation of Test on oracle f with a classical computer. Note that the classical computer implements arithmetic on exact

real numbers. However, neither this nor the high complexity of the algorithm is relevant to our proof of low query complexity. For this purpose we only need to know that the appropriate unitary transformations between queries exist. Thus, our quantum algorithm runs in exponential time, but uses only polynomially many queries in any model allowing arbitrary one-qubit and two-qubit gates, where each gate is given (implicitly) via the result of a classical computation. This model is of course not realistic, but it suffices to rule out easy query complexity lower bounds, as discussed in the Introduction.

Let $Y_{\frac{1}{4}} \cup Y_{\frac{3}{4}}$ be any partitioning of the set of subgroups $\{K_1, \dots, K_r\}$. The algorithm `Test` of the previous subsection succeeds in identifying the hidden subgroup with high probability. We now describe a new algorithm `ExactTest` that merely distinguishes between the two above sets of subgroups, but does so with known and desirable probabilities.

Lemma 6 *The probability that the outcome of a measurement of the ancilla qubit of the state $\text{ExactTest}(|\Psi_{\text{init}}\rangle \otimes |0\rangle)$ is 1 is $\frac{3}{4}$ if the hidden subgroup H is in $Y_{\frac{3}{4}}$, and it is $\frac{1}{4}$ if H is in $Y_{\frac{1}{4}}$.*

Before we describe the algorithm `ExactTest`, let M be an $r \times r$ matrix over $[0, 1]$ with each row and column indexed by a subgroup. Let entry (H, K_μ) of M be the conditional probability $\text{Prob}[K_\mu|H]$ that a measurement of the first register of $|\Psi_{\text{final}}\rangle$ yields the outcome μ conditional on f being strictly H -periodic.

Let $s = \lceil 2 \log(4r^3) \rceil \in O(\log^2 N)$ so that by Theorem 2, any diagonal entry of M is at least $1 - \frac{1}{r^2}$, and since the entries of any row of M sum to 1, any off-diagonal entry of M is between 0 and $\frac{1}{r^2}$. Thus, we can express M as $M = I - \Delta$, where each entry of Δ has absolute value bounded by $\frac{1}{r^2}$. It follows that $M^{-1} = I + \Delta + \Delta^2 + \Delta^3 + \dots$, subject to the convergence of $\Gamma = \Delta + \Delta^2 + \Delta^3 + \dots$, which we now show. By induction on i , each entry of Δ^i has absolute value bounded by $\frac{1}{r^{i+1}}$. Therefore, each entry of Γ has absolute value bounded by $\sum_{i=1}^{\infty} \frac{1}{r^{i+1}} = \frac{1}{r(r-1)}$.

Let y be any $r \times 1$ column vector with entries from $\{\frac{1}{4}, \frac{3}{4}\}$ and with each row indexed by a subgroup. Set $x = M^{-1}y$. Then, since $M^{-1}y = y + \Gamma y$, every entry of x is within $\frac{3}{4(r-1)}$ of the corresponding entry of y , and thus every entry of x is in $[0, 1]$ for $r \geq 4$.

Algorithm `ExactTest` acts on the initial state $|\Psi_{\text{init}}\rangle \otimes |0\rangle$, where the last register holds an ancilla qubit in state $|0\rangle$, and is defined as

$$\text{ExactTest} = R \cdot (\text{Test} \otimes I). \quad (4)$$

First, it applies `Test` on the first part of the system. It then applies R , which, conditionally on the output register holding the subgroup index μ , rotates the ancilla qubit from $|0\rangle$ to $\sqrt{1-x_\mu}|0\rangle + \sqrt{x_\mu}|1\rangle$. Because $R = \sum_\mu P_\mu \otimes R_\mu$ for projectors $P_\mu = |\mu\rangle\langle\mu|$ and certain qubit rotations R_μ , it can be implemented unitarily. The probability that a measurement of the ancilla qubit of the resulting state $\text{ExactTest}(|\Psi_{\text{init}}\rangle \otimes |0\rangle)$ yields a 1 is thus

$$\sum_\mu x_\mu \text{Prob}[K_\mu|H_\nu],$$

which, by definition of the column vector x , is equal to y_ν , where H_ν is the hidden subgroup. In other words, the probability of measuring a 1 depends only on the index of the hidden subgroup. Set $y_\nu = \frac{3}{4}$ if $K_\nu \in Y_{\frac{3}{4}}$, and set $y_\nu = \frac{1}{4}$ if $K_\nu \in Y_{\frac{1}{4}}$. Lemma 6 follows.

Lemma 6 provides us with a method for distinguishing between two complementary subsets of subgroups with probabilities $\frac{3}{4}$ and $\frac{1}{4}$. Using amplitude amplification [13], we can alter those probabilities into being equal to 0 and 1, and hence distinguish between the two sets $Y_{\frac{3}{4}}$ and $Y_{\frac{1}{4}}$ with certainty. Applying binary search on the set of subgroups with various choices of $Y_{\frac{3}{4}}$ and $Y_{\frac{1}{4}}$ then yields the second half of Theorem 1.

3 Concluding remarks

Let HSP denote the decision problem of determining if the hidden subgroup is non-trivial. Let $Q_E(\mathcal{P})$ denote the quantum *query* complexity of determining some decision problem \mathcal{P} with certainty, and let $Q_1(\mathcal{P})$ denote the quantum query complexity of determining \mathcal{P} with one-sided error. Then $Q_E(\text{HSP}) \in O(\log^2 |G|)$ since it suffices to use one round of amplitude amplification, because binary search among subgroups is not needed (let $Y_{\frac{3}{4}}$ be the singleton containing only the trivial subgroup). Also $Q_1(\text{HSP}) \in O(\log |G|)$ since we need to test only for the cyclic subgroups, of which there are at most $|G|$, and then use one round of amplitude amplification for the case where the subgroup is trivial. If the subgroup is non-trivial, the algorithm may output an incorrect answer and thus the algorithm has one-sided error.

The technique to construct exact quantum algorithms presented here relies on the property that we can compute the conditional probabilities with arbitrary precision. The technique seems to be applicable both to proving lower bounds as well as to designing efficient algorithms. It can rule out easy lower bounds for exact quantum computation, or it can be used to give simple and efficient exact quantum algorithms for problems for which the number of distinct success probabilities is polynomially bounded in the running time of the given bounded-error algorithm.

Acknowledgements

We are grateful to Richard Cleve for suggesting studying the exact query complexity of the HSP and for valuable comments and encouragement. We appreciate the constructive comments of the referees. M. E. and E. K. were supported by the DOE, contract W-7405-ENG-36, and by the NSA. P. H. received support from Alberta Ingenuity Fund and the Pacific Institute for the Mathematical Sciences.

References

- [1] D. R. Simon, On the power of quantum computation, *SIAM J. Comput.* 26 (1997) 1474–1483.

- [2] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM J. Comput.* 26 (1997) 1484–1509.
- [3] A. Y. Kitaev, Quantum computations: Algorithms and error correction, *Russian Math. Surveys* 52 (1997) 1191–1249.
- [4] G. Brassard, P. Høyer, An exact quantum polynomial-time algorithm for Simon’s problem, in: *Proc. of the 5th Israeli Symposium on Theory of Computing Systems, Israel, 1997*, pp. 12–23.
- [5] M. Ettinger, P. Høyer, On quantum algorithms for noncommutative hidden subgroups, *Advances in Applied Mathematics* 25 (2000) 239–251.
- [6] M. Grigni, L. Schulman, M. Vazirani, U. Vazirani, Quantum mechanical algorithms for the nonabelian hidden subgroup problem, in: *Proc. of the 33rd Annual ACM Symposium on the Theory of Computation*, ACM Press, 2001, pp. 68–74.
- [7] S. Hallgren, A. Russell, A. Ta-Shma, Normal subgroup reconstruction and quantum computation using group representations, in: *Proc. of the 32nd Annual ACM Symposium on the Theory of Computation*, ACM Press, 2000, pp. 627–635.
- [8] G. Ivanyos, F. Magniez, M. Santha, Efficient quantum algorithms for some instances of the non-Abelian hidden subgroup problem, in: *Proc. of the 13th ACM Symposium on Parallel Algorithms*, ACM Press, 2001, pp. 263–270.
- [9] M. Rötteler, T. Beth, Polynomial-time solution to the hidden subgroup problem for a class of non-Abelian groups, *quant-ph/0112086* (2001).
- [10] C. Zalka, On a particular non-Abelian hidden subgroup problem, <http://qso.lanl.gov/~zalka/QC/QC.html> (1999).
- [11] S. Aaronson, Quantum lower bound for the collision problem, in: *Proc. of the 34th Annual ACM Symposium on the Theory of Computation (STOC)*, ACM Press, 2002, pp. 635–642.
- [12] Y. Shi, Quantum lower bounds for the collision and the element distinctness problems, in: *Proc. of the 43rd Annual Symposium on the Foundations of Computer Science*, 2002, pp. 513–519.
- [13] G. Brassard, P. Høyer, M. Mosca, A. Tapp, Quantum amplitude amplification and estimation, in: J. S. J. Lomonaco, H. E. Brandt (Eds.), *Quantum Computation and Quantum Information: A Millennium Volume*, AMS Contemporary Mathematics Series, Am. Math. Soc. USA, 2002, pp. 53–74.