

From optimal measurement to efficient quantum algorithms for the hidden subgroup problem over semidirect product groups

Dave Bacon*
dabacon@santafe.edu

Andrew M. Childs†
amchilds@caltech.edu

Wim van Dam‡
vandam@cs.ucsb.edu

Abstract

We approach the hidden subgroup problem by performing the so-called pretty good measurement on hidden subgroup states. For various groups that can be expressed as the semidirect product of an abelian group and a cyclic group, we show that the pretty good measurement is optimal and that its probability of success and unitary implementation are closely related to an average-case algebraic problem. By solving this problem, we find efficient quantum algorithms for a number of nonabelian hidden subgroup problems, including some for which no efficient algorithm was previously known: certain metacyclic groups as well as all groups of the form $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$ for fixed r (including the Heisenberg group, $r = 2$). In particular, our results show that entangled measurements across multiple copies of hidden subgroup states can be useful for efficiently solving the nonabelian HSP.

1 Introduction

The hidden subgroup problem (HSP) stands as one of the major challenges for quantum computation. Shor’s discovery of an efficient quantum algorithm for factoring and calculating discrete logarithms [30], which essentially solves the abelian HSP [4, 19], focused attention on the question of what computational problems might be solved asymptotically faster by quantum computers than by classical ones. In particular, we would like to understand when the *nonabelian* hidden subgroup problem admits an efficient quantum algorithm.

Considerable progress on this question has been made since Shor’s discovery. Efficient quantum algorithms have been found for the case where the hidden subgroup is promised to be normal and there is an efficient quantum Fourier transform over the group [13], or where the group is “almost abelian” [12] or, more generally, “near-Hamiltonian” [11]. In addition, efficient algorithms have been found for several groups that can be written as semidirect products of abelian groups: the wreath product $\mathbb{Z}_2^n \wr \mathbb{Z}_2$ [28], certain groups of the form $\mathbb{Z}_{p^k}^n \rtimes \mathbb{Z}_2$ for a fixed prime power p^k [10], q -hedral groups with q sufficiently large [21], and particular groups of the form $\mathbb{Z}_{p^k} \rtimes \mathbb{Z}_p$ with p an odd prime [17].

Unfortunately, efficient algorithms have been elusive for two cases with particularly significant applications: the dihedral group and the symmetric group. An efficient algorithm for the HSP over the symmetric group would lead to an efficient algorithm for graph isomorphism [4, 7, 3, 16], and an efficient algorithm for the dihedral HSP (based on the standard approach described in Section 2.1) would lead to an efficient algorithm for certain lattice problems [26]. While no polynomial-time quantum algorithms for these problems are known, Kuperberg recently gave a subexponential (but superpolynomial) time and space algorithm for the dihedral HSP [20], and Regev improved the space requirement to be only polynomial [27].

*Santa Fe Institute, Santa Fe, NM 87501, USA

†Institute for Quantum Information, California Institute of Technology, Pasadena, CA 91125, USA

‡Department of Computer Science, University of California, Santa Barbara, Santa Barbara, CA 93106, USA

Recently, we have advocated an approach to solving hidden subgroup problems based on a state estimation technique known as the *pretty good measurement* (PGM). The standard approach to solving the HSP with a quantum computer (described in detail in Section 2.1) reduces the problem to the quantum mechanical task of distinguishing the members of an ensemble of *hidden subgroup states*. For the dihedral HSP, we showed that the PGM is in fact the optimal measurement for distinguishing any number of copies of the hidden subgroup states, that it successfully identifies the hidden subgroup with a polynomial number of copies of the states (for which we gave a tight lower bound), and that its implementation is closely related to an average-case subset sum problem [2]. Unfortunately, this subset sum problem appears to be difficult, so the approach did not yield an efficient algorithm for the dihedral HSP.

In this paper, we continue our study of the PGM as a tool for solving hidden subgroup problems. We apply the method to all groups that can be written as a semidirect product $A \rtimes \mathbb{Z}_p$ of an abelian group A and a cyclic group of prime order p . As in the case of the dihedral group, the PGM for these groups is closely related to a certain kind of average-case algebraic problem, which we call the *matrix sum problem*. For some groups, the matrix sum problem can be solved efficiently, leading to an efficient algorithm for the HSP. We demonstrate this for two classes of groups in Sections 5 and 6. In Section 5, we give an efficient algorithm for metacyclic groups $\mathbb{Z}_N \rtimes \mathbb{Z}_p$ with $N/p = \text{poly}(\log N)$, generalizing the result of [21] for p -hedral groups, which requires N prime. In Section 6, we give an efficient algorithm for any group of the form $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$ for fixed r , including the Heisenberg group, the unique nontrivial semidirect product $\mathbb{Z}_p^2 \rtimes \mathbb{Z}_p$. For the groups $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$, the matrix sum problem is a system of polynomial equations over a finite field, which can be solved efficiently using Buchberger’s (classical) algorithm for computing Gröbner bases [5]. For the metacyclic groups, the matrix sum problem requires the calculation of discrete logarithms, which can be done efficiently using Shor’s algorithm [30]. In both cases, the algorithm uses abelian Fourier transforms, but does not explicitly use a nonabelian Fourier transform.

To simplify the PGM approach to the HSP, we find it useful to focus on a specific set of subgroups rather than allowing the hidden subgroup to be arbitrary. For the dihedral group, Ettinger and Høyer showed that it is sufficient to consider the case where the hidden subgroup is either trivial or has order 2. In Section 3, we give an analogous reduction showing that it is sufficient to consider subgroups that are either trivial or cyclic and of order p . In fact, this reduction alone is sufficient to solve the HSP over some nonabelian groups, such as $\mathbb{Z}_2^n \wr \mathbb{Z}_2$ (cf. [28]) and the groups $P_{p,r}$ of [17].

It is well known that the query complexity of the hidden subgroup problem is polynomial, and in particular, that only polynomially many copies of the hidden subgroup states are sufficient to solve the problem in general [9]. However, measurements that operate on a single hidden subgroup state at a time are in general *not* sufficient—in particular, they are insufficient for the symmetric group [22]. Kuperberg’s algorithm for the dihedral HSP depends essentially on using entangled measurements on multiple copies of the hidden subgroup states [20], but unfortunately does not run in polynomial time. As far as we know, our algorithm for the HSP over groups of the form $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$ is the first efficient quantum algorithm to use entangled measurements across multiple copies (specifically, r copies) of the hidden subgroup states. This result provides hope that the distinguishability of polynomially many copies of hidden subgroup states may lead to further efficient algorithms through the implementation of entangled measurements.

The remainder of the article is organized as follows. In Section 2.1 we review the hidden subgroup problem and the standard approach to solving it with a quantum computer, and in Section 2.2 we review some relevant facts about semidirect product groups. Then, in Section 3, we give the reduction to cyclic subgroups. In Section 4, we present the pretty good measurement approach in detail, describing the relationship between hidden subgroup states and the matrix sum problem, computing the success probability of the PGM, proving its optimality, and explaining how

to implement it on a quantum computer. We apply the approach to specific groups in Sections 5 and 6, and we conclude in Section 7 with a discussion of the results and some open problems.

2 Definitions

2.1 Hidden subgroup problem

Let G be a finite group. We say that a function $f : G \rightarrow S$ (where S is a finite set) hides the subgroup $H \leq G$ if f is constant and distinct on left cosets of H in G . The hidden subgroup problem is the following: given the ability to query the function f , find a generating set for H .

To approach the HSP with a quantum computer, we must have the ability to query f in superposition. In particular, we are provided with a quantum oracle U_f acting as $U_f : |g, y\rangle \mapsto |g, y \oplus f(g)\rangle$ for all $g \in G$ and $y \in S$, where \oplus denotes the bitwise exclusive or operation, and the elements of S are represented by bit strings of length $\text{poly}(\log |G|)$. An efficient quantum algorithm for the HSP is an algorithm using this black box that finds a generating set for H in time $\text{poly}(\log |G|)$.

The standard approach to solving the HSP on a quantum computer is as follows. Create a superposition over all elements of the group and then query the function in superposition, giving

$$U_f : \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g, 0\rangle \mapsto \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g, f(g)\rangle. \quad (1)$$

Next, discard the second register, leaving the first register in a mixed state whose form depends on the hidden subgroup H ,

$$\rho_H := \frac{|H|}{|G|} \sum_{g \in K} |gH\rangle\langle gH| \quad \text{with the coset states} \quad |gH\rangle := \frac{1}{\sqrt{|H|}} \sum_{h \in H} |gh\rangle, \quad (2)$$

where $K \subset G$ is a complete set of left coset representatives of H in G . We call ρ_H the *hidden subgroup state* hiding the subgroup H . In the standard approach, one attempts to determine H using samples of hidden subgroup state ρ_H .

2.2 Semidirect product groups

The semidirect product of two groups A, B is defined in terms of a homomorphism $\varphi : B \rightarrow \text{Aut}(A)$, where $\text{Aut}(A)$ denotes the automorphism group of A . The semidirect product group $A \rtimes_{\varphi} B$ consists of the elements (a, b) with $a \in A$ and $b \in B$. With group operations in A and B written additively, the group operation in $A \rtimes_{\varphi} B$ is defined as $(a, b)(a', b') = (a + \varphi(b)(a'), b + b')$. It is not hard to show that group inversion satisfies $(a, b)^{-1} = (\varphi(-b)(-a), -b)$.

We consider the hidden subgroup problem for semidirect product groups $G = A \rtimes_{\varphi} \mathbb{Z}_p$, where A is an abelian group and p is prime. In this case, since φ is a homomorphism of a cyclic group, it is determined entirely by $\varphi(1)$ (in particular, $\varphi(p)$ is the identity map); hence, with a slight abuse of notation, we let $\varphi := \varphi(1)$ (an automorphism of A) henceforth, and define $\varphi^b := (\varphi \circ \dots \circ \varphi)$ (b times), giving $\varphi(b) = \varphi^b$.

We will be especially interested in cyclic subgroups of such a semidirect product group. For any $a \in A$, we have $\langle (a, 1) \rangle = \{(0, 0), (a, 1), (a + \varphi(a), 2), (a + \varphi(a) + \varphi^2(a), 3), \dots\}$. For convenience, we introduce the function $\Phi^{(b)} : A \rightarrow A$ defined by

$$\Phi^{(b)}(a) := \sum_{i=0}^{b-1} \varphi^i(a) \quad (3)$$

such that the elements of $\langle (a, 1) \rangle$ are $(a, 1)^b = (\Phi^{(b)}(a), b \bmod p)$ for $b \in \mathbb{N}$.

3 Reduction to cyclic subgroups

To simplify the hidden subgroup problem over semidirect product groups, we can reduce the problem to that of finding either the trivial hidden subgroup or a hidden cyclic subgroup of order p . This reduction generalizes a result of Ettinger and Høyer, who reduced the dihedral HSP to the problem of finding a trivial hidden subgroup or a hidden reflection (an order 2 subgroup).

Lemma 1 (cf. Theorem 2.3 of [8]). *To find an efficient algorithm for the HSP over $A \rtimes \mathbb{Z}_p$ with p prime, it suffices to find an efficient algorithm for the HSP over $A_2 \rtimes \mathbb{Z}_p$ for any $A_2 \leq A$ with the promise that either H is trivial or $H = \langle (d, 1) \rangle$ for some $d \in A_2$ with $|H| = p$.*

Proof. The proof proceeds along the lines of the Ettinger-Høyer reduction. Their idea is to factor out the part of the hidden subgroup that lies within A . For generalized dihedral groups ($p = 2$, $\varphi(a) = -a$), their reduction goes through unchanged, as noted in [10]. However, in some cases there is an additional complication coming from the fact that the A -part of the hidden subgroup may not be normal in G . Fortunately, we can deal with this case separately.

Let $f : G \rightarrow S$ hide a subgroup H . Let $G_1 := A \times \{0\}$, and let $H_1 := H \cap G_1 =: A_1 \times \{0\}$. Since f restricted to the abelian group G_1 hides the subgroup H_1 , we can efficiently find generators for H_1 by solving an abelian hidden subgroup problem.

We would like to factor out H_1 , so we check whether it is normal in G . For $g = (a, b) \in G$ and $h_1 = (h, 0) \in H_1$, we have $gh_1g^{-1} = (a, b)(h, 0)(\varphi^{-b}(-a), -b) = (\varphi^b(h), 0)$. Therefore $H_1 \trianglelefteq G$ if and only if $\varphi(H_1) = H_1$, where $\varphi(H_1) := \{(\varphi(h), 0) : h \in A_1\}$. If $H_1 \neq H$, then $\varphi(H_1) = H_1$: since p is prime, there must be some $(d, 1) \in H$, and for any $h \in A_1$, $(d, 1)(h, 0)(d, 1)^{-1} = (\varphi(h), 0) \in H$, hence $\varphi(h) \in A_1$.

If $H_1 = H$, it could be that $\varphi(H_1) \neq H_1$. However, given a generating set for H_1 , we can check whether $H_1 \trianglelefteq G$ (for example, we can use the results of [31]; note that G is solvable since its commutator subgroup is abelian¹). If we find $H_1 \not\trianglelefteq G$, then we know that $H = H_1$, and we are done. Otherwise, we learn that $H_1 \trianglelefteq G$, and we proceed to factor out H_1 .

The (left) cosets of H_1 in G can be represented as follows: for $g = (a, b) \in G$,

$$gH_1 = \{(a, b)(h, 0) : h \in A_1\} = \{(a + \varphi^b(h), b) : h \in A_1\} = \{(a + h, b) : h \in A_1\}. \quad (4)$$

Thus the cosets can be labeled by $b \in \mathbb{Z}_p$ and $a \in A_2 := A/A_1$. Now we work in the group $G_2 := G/H_1 \cong A_2 \rtimes_{\varphi_2} \mathbb{Z}_p$ whose elements are these cosets. Note that G_2 inherits its defining automorphism φ_2 from the original automorphism φ .

Since f is constant on cosets of $H_1 \leq H$, we can consider it as a function $f : G_2 \rightarrow S$ with the hidden subgroup $H_2 := H/H_1$. If $H_1 = H$ then H_2 is trivial. If $H_1 \neq H$, then we must have $H = \langle H_1, (a, 1) \rangle$ for some $a \in A$, which can be seen as follows. Since p is prime, there must be some $(a, 1) \in H$; but for any additional $(a', 1) \in H$, we have $(a', 1)(a, 1)^{-1} = (a' - a, 0) \in H_1$, and hence $(a', 1) \in \langle H_1, (a, 1) \rangle$. Also, note that $\langle (a, 1) \rangle \cap H_1 = \langle (\Phi^{(p)}(a), 0) \rangle$. Thus, by the second isomorphism theorem, $H_2 = H/H_1 \cong \langle (a, 1) \rangle / \langle (\Phi^{(p)}(a), 0) \rangle$, which is a cyclic group of order p generated by $(d, 1)$ for some $d \in A_2$. \square

From now on we assume that the hidden subgroup is $H = \langle (d, 1) \rangle = \{(\Phi^{(b)}(d), b) : b \in \mathbb{Z}_p\}$ for some $d \in A$. Note in particular that since $|H| = p$, we have $\Phi^{(p)}(d) = 0$.

This reduction alone is enough to solve the HSP over some semidirect product groups. For example, the HSP over the wreath product groups $\mathbb{Z}_2^n \wr \mathbb{Z}_2$ (for which an efficient quantum algorithm

¹In some cases, the results of [31] may not be required; for example, in the case $A = \mathbb{Z}_p^r$ considered in Section 6, H_1 can be viewed as a subspace, and it is sufficient to check whether this subspace is invariant under the action of the matrix μ that defines the automorphism φ .

was given in [28]) and $\mathbb{Z}_{p^r} \rtimes_{\varphi} \mathbb{Z}_p$ with $\varphi(a) = (p^{r-1} + 1)a$ (for which an efficient quantum algorithm was given in [17]) are both reduced to instances of the abelian HSP. However, in general, we are left with a nonabelian HSP, which we attempt to solve using the pretty good measurement.

4 The pretty good measurement approach

In this section, we present the pretty good measurement approach to the hidden subgroup problem over $G = A \rtimes \mathbb{Z}_p$. We begin in Section 4.1 by describing the hidden subgroup states and expressing them in terms of a certain algebraic problem called the matrix sum problem. Then, in Section 4.2, we describe the pretty good measurement for distinguishing these states. In Section 4.3, we give an expression for the success probability of the measurement (as well as general upper and lower bounds), and in Section 4.4, we prove that the measurement is optimal. Finally, in Section 4.5, we explain how the measurement can be implemented by solving the matrix sum problem.

4.1 Hidden subgroup states and the matrix sum problem

According to (2), the hidden subgroup states are uniform mixtures of uniform superpositions over the left cosets of H in G . A complete set of left coset representatives of $H = \langle (d, 1) \rangle$ in $G = A \rtimes \mathbb{Z}_p$ is given by $L = \{(\ell, 0) : \ell \in A\}$, and we have the coset states

$$|\psi_{\ell,d}\rangle := \frac{1}{\sqrt{p}} \sum_{b \in \mathbb{Z}_p} |(\ell, 0)(\Phi^{(b)}(d), b)\rangle = \frac{1}{\sqrt{p}} \sum_{b \in \mathbb{Z}_p} |\ell + \Phi^{(b)}(d), b\rangle. \quad (5)$$

The hidden subgroup state is

$$\rho_d = \frac{1}{|A|} \sum_{\ell \in A} |\psi_{\ell,d}\rangle \langle \psi_{\ell,d}|. \quad (6)$$

Fourier transforming the first register (over A) gives

$$\tilde{\rho}_d = \frac{1}{|A|} \sum_{x \in A} |\tilde{\psi}_{x,d}\rangle \langle \tilde{\psi}_{x,d}| \quad \text{where} \quad |\tilde{\psi}_{x,d}\rangle := \frac{1}{\sqrt{p}} \sum_{b \in \mathbb{Z}_p} \chi_x(\Phi^{(b)}(d)) |x, b\rangle \quad (7)$$

and where $\chi_x : A \rightarrow \mathbb{C}$ for $x \in A$ denotes the x th group character of A , satisfying $\chi_x \cdot \chi_{x'} = \chi_{x+x'}$ and $\chi_x(y) = \chi_y(x)$. (For example, for $A = \mathbb{Z}_N$, $\chi_x(y) = \exp(2\pi i xy/N)$; for $A = \mathbb{Z}_p^r$, $\chi_x(y) = \exp(2\pi i(x \cdot y)/p)$.)

By Lemma 7 in Appendix A, there exists a function $\hat{\Phi}^{(b)} : A \rightarrow A$ such that $\chi_x(\Phi^{(b)}(d)) = \chi_{\hat{\Phi}^{(b)}(x)}(d)$ for all $d, x \in A$; hence

$$\tilde{\rho}_d = \frac{1}{|G|} \sum_{x \in A} \sum_{b, c \in \mathbb{Z}_p} \chi_{\hat{\Phi}^{(b)}(x)}(d) \bar{\chi}_{\hat{\Phi}^{(c)}(x)}(d) |x, b\rangle \langle x, c|. \quad (8)$$

For k copies, we have simply

$$\tilde{\rho}_d^{\otimes k} = \frac{1}{|G|^k} \sum_{x \in A^k} \sum_{b, c \in \mathbb{Z}_p^k} \chi_{\hat{\Phi}^{(b)}(x)}(d) \bar{\chi}_{\hat{\Phi}^{(c)}(x)}(d) |x, b\rangle \langle x, c| \quad (9)$$

where for $b \in \mathbb{Z}_p^k$ and $x \in A^k$,

$$\hat{\Phi}^{(b)}(x) := \sum_{j=1}^k \hat{\Phi}^{(b_j)}(x_j). \quad (10)$$

To simplify this expression further, for any $x \in A^k$ and $w \in A$, let $S_w^x := \{b \in \mathbb{Z}_p^k : \hat{\Phi}^{(b)}(x) = w\}$ denote the set of solutions to the equation $\hat{\Phi}^{(b)}(x) = w$, let $\eta_w^x := |S_w^x|$ denote the number of such solutions, and let

$$|S_w^x\rangle := \frac{1}{\sqrt{\eta_w^x}} \sum_{b \in S_w^x} |b\rangle \quad (11)$$

denote the (normalized) uniform superposition over all solutions. (If $\eta_w^x = 0$, then no such state can be defined, and we use the convention $|S_w^x\rangle = 0$.) Then the hidden subgroup state can be written

$$\tilde{\rho}_d^{\otimes k} = \frac{1}{|G|^k} \sum_{x \in A^k} \sum_{w, v \in A} \chi_w(d) \bar{\chi}_v(d) \sqrt{\eta_w^x \eta_v^x} |x, S_w^x\rangle \langle x, S_v^x|. \quad (12)$$

Clearly, the hidden subgroup states are closely connected to the problem of finding solutions $b \in \mathbb{Z}_p^k$ to the equation

$$\hat{\Phi}^{(b)}(x) = w, \quad (13)$$

where $x \in A^k$ and $w \in A$ are chosen uniformly at random. We refer to this problem as the *matrix sum problem* because we can represent φ as a matrix μ , and hence $\hat{\Phi}^{(b)}$ as a sum of matrices $M^{(b)}$ (see Appendix A). In the case $A = \mathbb{Z}_N$ of Section 5, the matrix sum is a single scalar $M^{(b)} \in \mathbb{Z}_N^\times$, while in the case $A = \mathbb{Z}_p^r$ of Section 6, $M^{(b)} \in \mathbb{Z}_p^{r \times r}$ is the transpose of the matrix representing $\Phi^{(b)}$.

In general, the matrix sum problem will have many solutions when k is large, and few solutions when k is small. Since $\sum_{w \in A} \eta_w^x = p^k$, the expected number of solutions is

$$\mathbb{E}_{x \in A^k, w \in A} [\eta_w^x] = \frac{p^k}{|A|}. \quad (14)$$

Thus, we typically expect the matrix sum problem to have many solutions for $k \gg \log_p |A|$ and few solutions for $k \ll \log_p |A|$. In fact, we often find a sharp transition at $k \sim \log_p |A|$.

Note that for the dihedral group of order $2N$, $A = \mathbb{Z}_N$ and $p = 2$ with $\varphi(x) = -x$, the matrix sum problem is the average-case subset sum problem [2], while for the abelian case with $\varphi(x) = x$ (so that $\Phi^{(b)}(x) = bx$), the matrix sum problem is simply $bx = w$.

4.2 The measurement

The *pretty good measurement* (PGM, also known as the *square root measurement* or *least squares measurement*) is a positive operator valued measure (POVM) that often does a pretty good job of distinguishing members of an ensemble of quantum states [14]. For the ensemble of states σ_j with equal a priori probabilities, the PGM $\{E_j\}$ is given by

$$E_j := \Sigma^{-1/2} \sigma_j \Sigma^{-1/2} \quad \text{where} \quad \Sigma := \sum_j \sigma_j, \quad (15)$$

and the inverse is taken over the support of Σ . Clearly, the PGM is a POVM over the support of the ensemble.

For our ensemble of hidden subgroup states, using k copies of the state, we have

$$\Sigma := \sum_{j \in A} \rho_j^{\otimes k} = \frac{|A|}{|G|^k} \sum_{x \in A^k} \sum_{w \in A} \eta_w^x |x, S_w^x\rangle \langle x, S_w^x| \quad (16)$$

where we have assumed for simplicity that the subgroup $\langle(j, 1)\rangle$ has order p for every $j \in A$. Since Σ is diagonal, the square root of its inverse (over its support) is particularly easy to calculate. Inserting (12) and (16) into (15), we find the measurement operators

$$E_j = \frac{1}{|A|} \sum_{x \in A^k} \sum_{w, v \in A} \chi_w(j) \bar{\chi}_v(j) |x, S_w^x\rangle \langle x, S_v^x|. \quad (17)$$

This defines the pretty good measurement for distinguishing order p hidden subgroup states of G .

To solve the HSP using Lemma 1, we must also identify the trivial subgroup. However, if the PGM correctly identifies an order p subgroup when one exists, we can simply look for an order p subgroup and check whether the function is constant on the identity and some generator, which it will not be if the hidden subgroup is in fact trivial. Therefore, from now on we focus on finding an order p hidden subgroup.

In general, G will have some subgroups $\langle(j, 1)\rangle$ of order p and some such subgroups whose orders are larger integer multiples of p . In this case, (17) is not, strictly speaking, the PGM for distinguishing the order p hidden subgroups alone. Furthermore, the state (12) for d corresponding to a non-order p subgroup is not even a hidden subgroup state. However, since (12) is always a valid quantum state, the resulting PGM is well defined. It is convenient to work with this PGM even when not all subgroups $\langle(j, 1)\rangle$ have order p . If the measurement identifies the order p hidden subgroups with reasonable probability, then by Lemma 1, this is sufficient to solve the HSP.

4.3 Success probability

The probability of successfully identifying an order p hidden subgroup $\langle(d, 1)\rangle$ is independent of d , and is given by

$$\Pr(\text{success}) := \text{tr } E_d \rho_d^{\otimes k} = \frac{p}{|G|^{k+1}} \sum_{x \in A^k} \left(\sum_{w \in A} \sqrt{\eta_w^x} \right)^2. \quad (18)$$

For the PGM to successfully solve the HSP, this probability must not be too small. Whether the success probability is appreciable essentially depends on whether the corresponding matrix sum problem has many solutions. Specifically, we have the following:

Lemma 2 (Cf. Theorem 2 of [2]). *If $\Pr(\eta_w^x \geq \alpha) \geq \beta$ for uniformly random $x \in A^k$ and $w \in A$ (i.e., if most instances of the matrix sum problem have many solutions), then $\alpha\beta^2|A|/p^k \leq \Pr(\text{success}) \leq p^k/|A|$.*

Proof. For the upper bound, we have

$$\Pr(\text{success}) \leq \frac{p}{|G|^{k+1}} \sum_{x \in A^k} \left(\sum_{w \in A} \eta_w^x \right)^2 = \frac{p^k}{|A|} \quad (19)$$

since the η 's are integers and $\sum_{w \in A} \eta_w^x = p^k$ for any x . For the lower bound, we have

$$\Pr(\text{success}) \geq \frac{|A|}{p^k} \left(\frac{1}{|A|^{k+1}} \sum_{x \in A^k} \sum_{w \in A} \sqrt{\eta_w^x} \right)^2 \quad (20)$$

by Cauchy's inequality applied to (18). Now

$$\frac{1}{|A|^{k+1}} \sum_{x \in A^k} \sum_{w \in A} \sqrt{\eta_w^x} \geq \sqrt{\alpha} \Pr(\eta_w^x \geq \alpha), \quad (21)$$

so by the hypothesis, $\Pr(\text{success}) \geq \alpha\beta^2|A|/p^k$ as claimed. \square

4.4 Optimality

In fact, the pretty good measurement is the optimal measurement for distinguishing the states (12), in the sense that it maximizes the success probability (18). This can be seen using the following result of Holevo and Yuen, Kennedy, and Lax:

Theorem 3 ([15, 32]). *Given an ensemble of states σ_j with equal a priori probabilities, the measurement with POVM elements E_j maximizes the probability of successfully identifying the state if and only if $\sum_i \sigma_i E_i = \sum_i E_i \sigma_i$ and $\sum_i \sigma_i E_i \geq \sigma_j$ for all j .*

The optimality of the PGM (17) can be proved by directly verifying these conditions. Note that this optimality does not necessarily follow from [23], as the hidden subgroups may not be conjugates. However, the same basic principle is at work.

Identifying the optimal measurement can be useful for proving lower bounds on the number of hidden subgroup states required to solve the HSP [2]. Of course, for the purpose of finding an efficient algorithm, it is not necessary for the measurement to be optimal; rather, it is sufficient for it to identify the hidden subgroup with reasonable probability. In fact, it may be that a suboptimal measurement which nevertheless is sufficient to solve the HSP is significantly easier to implement than the optimal one. However, it is encouraging to find that a particularly straightforward measurement is in fact optimal, and we will see that this measurement does lead to efficient algorithms in some cases. Together with the result of Ip showing that Shor's algorithm implements the optimal measurement for the abelian hidden subgroup problem [18], this suggests that identifying optimal measurements may be a useful guiding principle for discovering quantum algorithms.

4.5 Implementation

To find an efficient quantum algorithm based on the PGM, we must show how to implement the measurement efficiently on a universal quantum computer. Unsurprisingly, this implementation is also closely related to the matrix sum problem.

According to Neumark's theorem [24], any POVM can be implemented by a unitary transformation on the system together with an ancilla, followed by a measurement in the standard basis. In particular, for a POVM consisting of N rank one operators $E_j = |e_j\rangle\langle e_j|$ in a D -dimensional Hilbert space, U has the block form

$$U = \begin{pmatrix} V & X \\ Y & Z \end{pmatrix} \quad (22)$$

where the columns of the $N \times D$ matrix V are the D -vectors $|e_j\rangle$, i.e., $V = \sum_{j=1}^N |j\rangle\langle e_j|$.

Recall from (17) that the POVM operators for the PGM on hidden subgroup states can be written

$$E_j = \sum_{x \in A^k} |x\rangle\langle x| \otimes E_j^x \quad \text{where} \quad E_j^x := |e_j^x\rangle\langle e_j^x| \quad \text{with} \quad |e_j^x\rangle := \frac{1}{\sqrt{|A|}} \sum_{w \in A} \chi_w(j) |S_w^x\rangle. \quad (23)$$

In other words, each E_j is block diagonal, with blocks labeled by $x \in A^k$, and where each block is rank one. Thus, the measurement can be implemented in a straightforward way by first measuring the block label x and then performing the POVM $\{E_j^x\}_{j \in A}$ conditional on the first measurement result.

To implement the POVM $\{E_j^x\}_{j \in A}$ using Neumark's theorem, we would like to implement the unitary transformation U^x with the upper left submatrix

$$V^x = \frac{1}{\sqrt{|A|}} \sum_{j, w \in A} \bar{\chi}_w(j) |j\rangle\langle S_w^x|. \quad (24)$$

It is convenient to perform a Fourier transform (over A) on the left (i.e., on the index j), giving a unitary operator \tilde{U}^x with upper left submatrix

$$\tilde{V}^x = \frac{1}{|A|} \sum_{j,w,v \in A} \chi_w(j) \bar{\chi}_v(j) |w\rangle \langle S_v^x| = \sum_{w \in A} |w\rangle \langle S_w^x|. \quad (25)$$

Therefore, the PGM can be implemented efficiently if we can efficiently perform the transformation

$$|x, w\rangle \mapsto \begin{cases} |x, S_w^x\rangle & \eta_w^x > 0 \\ |\xi_w^x\rangle & \eta_w^x = 0 \end{cases} \quad (26)$$

where $|\xi_w^x\rangle$ is any state allowed by the unitarity of \tilde{U}^x . We refer to (26) as *quantum sampling* of solutions to the matrix sum problem. If we can efficiently quantum sample from matrix sum solutions, then by running the circuit in reverse, we can efficiently implement \tilde{U}^x , and hence the desired measurement.

5 Metacyclic groups

In this section, we present our first application of the PGM approach to a particular class of groups, those of the form $\mathbb{Z}_N \rtimes \mathbb{Z}_p$ with p prime. All such groups are *metacyclic*. The possible automorphisms φ correspond to multiplication by some $\mu \in \mathbb{Z}_N^\times$, so that $\varphi(a) = \mu a$. For φ^p to be the identity map, we must have $\mu^p = 1 \pmod{N}$. The function $\Phi^{(b)}$ can be represented by the sum

$$M^{(b)} := \sum_{i=0}^{b-1} \mu^i, \quad (27)$$

so that $\Phi^{(b)}(a) = \hat{\Phi}^{(b)}(a) = M^{(b)}a$.

Note that because $\langle \mu \rangle \leq \mathbb{Z}_N^\times$, Lagrange's theorem implies that p divides $|\mathbb{Z}_N^\times| = \phi(N)$, where $\phi(N)$ denotes the Euler totient function of N , i.e., the number of elements of \mathbb{Z}_N that are relatively prime to N . To give an efficient algorithm, we require that N and p be fairly close; in particular, we require $N/p = \text{poly}(\log N)$.

5.1 Solution of the matrix sum problem

To apply the PGM to the HSP over groups of the form $\mathbb{Z}_N \rtimes \mathbb{Z}_p$, we must understand the matrix sum problem for such groups. In general, the problem is the following: given uniformly random $x \in \mathbb{Z}_N^k$ and $w \in \mathbb{Z}_N$, find $b \in \mathbb{Z}_p^k$ such that $\sum_{j=1}^k M^{(b_j)} x_j = w$. We note in passing that if $\mu - 1 \in \mathbb{Z}_N^\times$ (which is always the case if, for example, N is prime), then we can sum the geometric series (27) to obtain $M^{(b)} = (\mu^b - 1)/(\mu - 1)$, and in particular, $M^{(p)} = 0$, so that $|\langle (d, 1) \rangle| = p$ for all $d \in \mathbb{Z}_N$. However, for the following, we do not need to assume that N is prime, that $\mu - 1 \in \mathbb{Z}_N^\times$, or even that $M^{(p)} = 0$.

For $k = 1$, the matrix sum problem $M^{(b)}x = w$ is particularly easy to solve. Assume that $x \in \mathbb{Z}_N^\times$, which occurs with probability $\phi(N)/N \in \Omega(1/\log \log N)$, so that we can rewrite the equation as $M^{(b)} = w/x$.

If such a b exists, it must be unique, since $M^{(b)} \neq M^{(b')}$ for $b \neq b' \in \mathbb{Z}_p$, which can be seen as follows. Supposing the claim is false, with $b > b'$ without loss of generality, $M^{(b)} - M^{(b')} = \mu^{b'} M^{(b-b')} = 0$, so $M^{(b-b')} = 0$. For $\mu \neq 1$, $\langle \mu \rangle \leq \mathbb{Z}_N^\times$ has $|\langle \mu \rangle| = p$ since $\mu^p = 1$. Using the identity

$$(\mu - 1)M^{(b)} = \mu^b - 1, \quad (28)$$

it is clear that $M^{(b-b')} = 0$ implies $\mu^{b-b'} = 1$, so that $b - b'$ is a multiple of p . But since $b, b' \in \mathbb{Z}_p$ with $b \neq b'$, this is a contradiction.

Now from (28), we see that by solving the discrete logarithm problem $\mu^b = 1 + (\mu - 1)w/x$ for b , we solve the matrix sum problem. The discrete logarithm can be calculated efficiently using Shor's algorithm [30], which implies that the unique solution to the matrix sum problem can be found efficiently when it exists.

For the PGM to be successful with reasonably high probability, the equation $M^{(b)} = w/x$ must be likely to have a solution. Since the various values of $b \in \mathbb{Z}_p$ yield p distinct possibilities for $M^{(b)}$, and since w/x is uniformly random in \mathbb{Z}_N under the assumption $x \in \mathbb{Z}_N^\times$, a solution exists with probability p/N , which is not too small given $N/p = \text{poly}(\log N)$. Taking into account the probability that $x \in \mathbb{Z}_N^\times$, a solution exists with probability at least $\phi(N)p/N^2$. Therefore, by Lemma 2, the PGM succeeds with probability $1/\text{poly}(\log Np)$, and we have

Theorem 4. *For p prime and N arbitrary with $N/p = \text{poly}(\log N)$, the hidden subgroup problem over $\mathbb{Z}_N \rtimes \mathbb{Z}_p$ can be solved in time $\text{poly}(\log Np)$.*

This generalizes a result of [21], which proves Theorem 4 in the case where N is prime.

If N/p were superpolynomial, then we could imagine solving the problem by implementing the PGM with $k > 1$ copies of the hidden subgroup state. However, in this case, the matrix sum problem is not simply a discrete logarithm, so it is not clear whether it can be solved efficiently.

5.2 Stripped down algorithm

Since the PGM approach described above requires only $k = 1$ copy of the hidden subgroup state, it yields a very simple quantum algorithm. Here we present this result without reference to the general framework, giving an algorithm that is quite straightforward, especially in comparison to the earlier algorithm of [21], which requires a nonabelian Fourier transform.

Consider the order p hidden subgroup $H = \{(M^{(b)}d, b) : b \in \mathbb{Z}_p\}$ for $d \in \mathbb{Z}_N$ with $M^{(p)}d = 0$. The left cosets are of the form $(\ell, 0)H$ for $\ell \in \mathbb{Z}_N$, and the corresponding coset states (5) are

$$|\psi_{\ell,d}\rangle = \frac{1}{\sqrt{p}} \sum_{b \in \mathbb{Z}_p} |\ell + M^{(b)}d, b\rangle. \quad (29)$$

Applying the Fourier transform over \mathbb{Z}_N to the first register, we obtain

$$|\tilde{\psi}_{\ell,d}\rangle = \frac{1}{\sqrt{Np}} \sum_{x \in \mathbb{Z}_N} \sum_{b \in \mathbb{Z}_p} \omega^{x(\ell + M^{(b)}d)} |x, b\rangle \quad (30)$$

where $\omega := \exp(2\pi i/N)$. Now measure the first register and assume that the result is some $x \in \mathbb{Z}_N^\times$, which happens with probability $\phi(N)/N$. Then append an ancilla register to this state and perform the calculation $|b, 0\rangle \mapsto |b, xM^{(b)}\rangle$, giving the state

$$\frac{1}{\sqrt{p}} \sum_{b \in \mathbb{Z}_p} \omega^{xM^{(b)}d} |b, xM^{(b)}\rangle. \quad (31)$$

Note that for any fixed $b \in \mathbb{Z}_p$, $M^{(b)}$ can be calculated efficiently by repeated squaring, using the fact that $M^{(2b)} = (1 + \mu^b)M^{(b)}$.

Now we perform a classical computation to erase the value of b . With the expression (28) and knowledge of $x \in \mathbb{Z}_N^\times$ and μ , we can efficiently perform the computation $|xM^{(b)}, 0\rangle \mapsto |xM^{(b)}, \mu^b\rangle$.

Then, using Shor's discrete logarithm algorithm, we efficiently compute $|xM^{(b)}, \mu^b\rangle \mapsto |xM^{(b)}, b\rangle$. Using this procedure to erase the value of b , we see that we can produce the state

$$\frac{1}{\sqrt{p}} \sum_{b \in \mathbb{Z}_p} \omega^{xM^{(b)}d} |xM^{(b)}\rangle. \quad (32)$$

Finally, we perform an inverse Fourier transform over \mathbb{Z}_N and observe the register in the hope of obtaining d . To calculate the probability of this happening, consider the perfect state

$$|\tilde{d}\rangle := \frac{1}{\sqrt{N}} \sum_{j \in \mathbb{Z}_N} \omega^{jd} |j\rangle, \quad (33)$$

which would be guaranteed to yield the answer d . The overlap between the perfect state (33) and the actual state (32) is $\sqrt{p/N}$, so the probability of observing d is at least p/N . Thus, overall, we find a success probability of at least $\phi(N)p/N^2$. Because $\phi(N)/N \in \Omega(1/\log \log N)$, and $N/p = \text{poly}(\log N)$ by assumption, simply repeating the above protocol $N^2/\phi(N)p = \text{poly}(\log Np)$ times gives an efficient quantum algorithm for the HSP over $\mathbb{Z}_N \times \mathbb{Z}_p$.

6 Groups of the form $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$

We now turn to another class of semidirect product groups, those of the form $\mathbb{Z}_p^r \rtimes_{\varphi} \mathbb{Z}_p$. We find an efficient quantum algorithm for such groups for any φ , so long as r is constant. In other words, the running time of the algorithm is $\text{poly}(\log p)$. The algorithm is particularly simple in the case $r = 2$, where the only nontrivial semidirect product is known as the Heisenberg group (for which it was recently shown that there is an efficient quantum algorithm whose output information theoretically determines the solution of the HSP [25]). We present the algorithm for $r = 2$ in Section 6.1, and then proceed to the general case in Section 6.2.

In general, when A is the elementary abelian p -group $A = \mathbb{Z}_p^r$, its automorphism group is $\text{Aut}(\mathbb{Z}_p^r) \cong \text{GL}_r(\mathbb{Z}_p)$. Therefore, φ can be identified with a nonsingular matrix $\mu \in \text{GL}_r(\mathbb{Z}_p)$ such that $\mu^p = I$. As before, we define $M^{(b)} := \sum_{i=0}^{b-1} \mu^i$ so that $\hat{\Phi}^{(b)}(a) = M^{(b)}a$ (see Appendix A).

6.1 The Heisenberg group

For $r = 2$, there are only two nonisomorphic semidirect product groups $\mathbb{Z}_p^2 \rtimes \mathbb{Z}_p$: the abelian group \mathbb{Z}_p^3 and the *Heisenberg group*, for which

$$\mu = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}. \quad (34)$$

The matrix sum problem for the Heisenberg group is the following: given uniformly random $x, y \in \mathbb{Z}_p^k$ and $w, v \in \mathbb{Z}_p^k$, find $b \in \mathbb{Z}_p^k$ such that

$$\sum_{j=1}^k \sum_{i=0}^{b_j-1} \mu^i \begin{pmatrix} x_j \\ y_j \end{pmatrix} = \begin{pmatrix} w \\ v \end{pmatrix}. \quad (35)$$

Clearly,

$$\mu^i = \begin{pmatrix} 1 & i \\ 0 & 1 \end{pmatrix} \quad \text{so} \quad M^{(b)} := \sum_{i=0}^{b-1} \mu^i = \begin{pmatrix} b & sb(1-b) \\ 0 & b \end{pmatrix} \quad (36)$$

where s is the multiplicative inverse of -2 in the finite field \mathbb{Z}_p , defined by $p = 2s + 1$. Therefore the matrix sum problem can be written

$$\sum_{j=1}^k \begin{pmatrix} b_j & sb_j(1-b_j) \\ 0 & b_j \end{pmatrix} \begin{pmatrix} x_j \\ y_j \end{pmatrix} = \begin{pmatrix} w \\ v \end{pmatrix}. \quad (37)$$

If $k = 1$, then it is not hard to see that the probability of having a solution is $O(1/p)$, i.e., exponentially small in $\log p$. However, if we take $k = 2$, then there are as many variables as (scalar) equations, and we find the matrix sum problem

$$b_1x_1 + sb_1(1-b_1)y_1 + b_2x_2 + sb_2(1-b_2)y_2 = w \quad (38)$$

$$b_1y_1 + b_2y_2 = v \quad (39)$$

with uniformly random $x_1, x_2, y_1, y_2, w, v \in \mathbb{Z}_p$.

The equations (38–39) can be solved as follows. Define

$$\Delta := (2wy_1 + vy_1 - v^2 - 2vx_1)(y_1 + y_2)y_2 + (vy_2 + x_1y_2 - x_2y_1)^2. \quad (40)$$

If Δ is a nonzero square in \mathbb{Z}_p and $y_1, y_2, y_1 + y_2 \neq 0$, then there are two solutions for (b_1, b_2) :

$$b_1 = \frac{vy_1 + x_2y_1 - x_1y_2 \pm \sqrt{\Delta}}{y_1(y_1 + y_2)} \quad \text{and} \quad b_2 = \frac{vy_2 + x_1y_2 - x_2y_1 \mp \sqrt{\Delta}}{y_2(y_1 + y_2)}. \quad (41)$$

It is straightforward to check that (b_1, b_2) is indeed a solution. If $\Delta = 0$ then these two solutions are the same, and there is only one solution. Finally, if Δ is nonsquare, then there are no solutions.

To calculate the success probability of the PGM, we must determine the probability that $y_1, y_2, y_1 + y_2 \neq 0$ and Δ is a square in \mathbb{Z}_p . First, note that since y_1, y_2 are uniformly random, we have $y_1, y_2, (y_1 + y_2) \neq 0$ with probability $(p-1)(p-2)/p^2 = 1 - O(1/p)$. Now rewrite Δ as

$$\Delta = 2wy_1y_2(y_1 + y_2) + [v(y_1 - 2x_1 - v)(y_1 + y_2)y_2 + (vy_2 + x_1y_2 - x_2y_1)^2]. \quad (42)$$

Assuming $2y_1y_2(y_1 + y_2) \neq 0$, Δ depends linearly on w . Hence, if we fix x_1, x_2, y_1, y_2, v and choose w uniformly at random from \mathbb{Z}_p , then Δ will also be uniformly random in \mathbb{Z}_p . Of the possible values of $\Delta \in \mathbb{Z}_p$, we have $(p-1)/2$ cases with $\sqrt{\Delta} \neq 0$, $(p-1)/2$ cases with Δ not a square, and one case with $\Delta = 0$. Therefore, under the assumption $y_1, y_2, y_1 + y_2 \neq 0$, we have

$$\Pr(\eta_{w,v}^{x,y} = 0) = \frac{1}{2} - \frac{1}{2p} \quad (43)$$

$$\Pr(\eta_{w,v}^{x,y} = 1) = \frac{1}{p} \quad (44)$$

$$\Pr(\eta_{w,v}^{x,y} = 2) = \frac{1}{2} - \frac{1}{2p}. \quad (45)$$

In particular, we see that $\Pr(\eta_{w,v}^{x,y} = 2) = \frac{1}{2} - O(1/p)$. Therefore, by Lemma 2, $\Pr(\text{success}) \geq 1 - O(1/p)$.

Since the above discussion gives an explicit solution of the matrix sum problem, and since arithmetic in the finite field \mathbb{Z}_p can be performed in time $\text{poly}(\log p)$, it is straightforward to efficiently implement the quantum sampling transformation (26).² Combined with the fact that the success probability of the PGM is large, this shows that the hidden subgroup problem in the Heisenberg group can be solved efficiently.

²For example, given the ability to compute a list of solutions b_1, \dots, b_η with η small, the following simple trick can be used to efficiently create the uniform superposition. Create the labeled superposition $\sum_{j \in \mathbb{Z}_\eta} |j, b_j\rangle / \sqrt{\eta}$, Fourier transform the first register (over \mathbb{Z}_η), and measure the first register. When the outcome is 0, which occurs with probability $1/\eta$, the desired state is obtained. Finally, use $O(\eta)$ repetitions to boost the success probability close to 1, and implement the measurement unitarily.

6.2 The general case

More generally, consider the group $\mathbb{Z}_p^r \rtimes_{\varphi} \mathbb{Z}_p$ for any constant r and for any automorphism φ , i.e., for any matrix μ satisfying $\mu^p = I$. It is easy to see that matrices related by a similarity transformation correspond to isomorphic groups. Thus, without loss of generality, we can assume that μ is in Jordan canonical form, or in other words, that it is zero everywhere except the diagonal and first superdiagonal, and that the elements on the first superdiagonal are either 0 or 1. Since $\mu^p = I$, the diagonal elements of this μ must all be equal to 1. Thus the various nonisomorphic groups of the form $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$ correspond to the partitions of r , where the partitions describe the sizes of the Jordan blocks.

For simplicity, we consider the case of a single Jordan block of size r ; the extension to other cases will be clear. In other words, we consider the $r \times r$ matrix

$$\mu = \begin{pmatrix} 1 & 1 & 0 & \cdots & 0 \\ 0 & 1 & 1 & \ddots & \vdots \\ 0 & 0 & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & 1 & 1 \\ 0 & \cdots & 0 & 0 & 1 \end{pmatrix}. \quad (46)$$

Then the b th matrix sum $M^{(b)} := I + \mu + \mu^2 + \cdots + \mu^{b-1}$ is given by

$$M^{(b)} = \begin{pmatrix} \binom{b}{1} & \binom{b}{2} & \binom{b}{3} & \cdots & \binom{b}{r} \\ 0 & \binom{b}{1} & \binom{b}{2} & \ddots & \vdots \\ 0 & 0 & \ddots & \ddots & \binom{b}{3} \\ \vdots & \ddots & \ddots & \binom{b}{1} & \binom{b}{2} \\ 0 & \cdots & 0 & 0 & \binom{b}{1} \end{pmatrix} \pmod{p}. \quad (47)$$

Note that $M^{(0)} = M^{(p)} = 0$. If μ consists of several Jordan blocks, then (47) can be applied in each block.

The matrix sum problem is thus the following: given uniformly random $x \in (\mathbb{Z}_p^r)^k$ and $w \in \mathbb{Z}_p^r$, find $b \in \mathbb{Z}_p^k$ such that $\sum_{j=1}^k M^{(b_j)} x_j = w$. This is a set of r polynomial equations over \mathbb{Z}_p in k variables. For example, in the case of a single Jordan block, we have

$$\sum_{j=1}^k \left[\binom{b_j}{1} x_{i,j} + \cdots + \binom{b_j}{i} x_{1,j} \right] = w_i \quad (48)$$

for all $1 \leq i \leq r$, where the i th equation is of degree i .

If $k < r$, so that there are fewer unknowns than equations, we expect that the matrix sum problem will typically have no solutions. On the other hand, if $k > r$, we expect that the ideal generated by (48) will typically have dimension $k - r$, and there will be $O(p^{k-r})$ solutions. By choosing $k = r$, we can ensure that there are typically $O(1)$ solutions, so that the PGM succeeds, yet the solutions of the matrix sum problem are few in number and thus relatively easy to find (and to quantum sample).

This intuition can be formalized by calculating the mean and variance of the number of solutions. Using such an argument, we find the following:

Lemma 5. *For the matrix sum problem of any group $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$ with $k = r$, $\Pr(\eta_w^x = 1 \text{ or } 2) \geq \frac{1}{4}$.*

Proof. See Appendix B. □

By Lemma 2, this shows that $\Pr(\text{success}) \geq 1/16$, even if we only consider the cases in which the number of solutions is at most 2. Thus, the pretty good measurement succeeds in identifying the hidden subgroup.

To efficiently implement the PGM, we quantum sample from solutions of the matrix sum problem, which can be done by computing a list of all solutions. A list of solutions can be found by first using Buchberger’s algorithm to compute a Gröbner basis for the ideal and then using elimination theory [5]. In general, upper bounding the complexity of Buchberger’s algorithm is a difficult problem, and it is known that the algorithm may be very inefficient in terms of the number of variables, the number of equations, and the degree of those equations. However, since we consider r constant, such inefficiency is not an issue. The running time in terms of p , the size of the field, enters only as an overall $\text{poly}(\log p)$ factor, accounting for the cost of performing arithmetic operations in a finite field (see for example [6]). Thus, for our purposes, a Gröbner basis can be computed efficiently. Since we consider only the cases in which there are at most 2 solutions, elimination is also efficient, giving an overall $\text{poly}(\log p)$ time algorithm for computing a list of matrix sum solutions, and hence for quantum sampling from those solutions.

Collecting these results, we find

Theorem 6. *The hidden subgroup problem over any group of the form $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$ with r fixed can be solved in time $\text{poly}(\log p)$ on a quantum computer.*

7 Discussion

In this paper, we have studied the pretty good measurement for semidirect product groups of the form $A \rtimes \mathbb{Z}_p$ with A abelian and p prime. We found that the PGM is closely connected to the matrix sum problem, and we exploited this connection to find efficient quantum algorithms for certain metacyclic groups (Section 5) as well as all groups of the form $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$ with r fixed (Section 6). The latter algorithm demonstrates that entangled measurements may be useful for efficiently solving the nonabelian HSP.

Aside from the fact that these particular nonabelian groups admit efficient quantum algorithms, our results suggest two general directions for further investigation. First, in the standard approach to the nonabelian HSP, it would be helpful to have a better understanding of when entangled measurements are necessary and when they can be implemented to give efficient algorithms. Second, for the HSP or for other problems that can be viewed as quantum state distinguishability problems, identifying an optimal measurement (or considering a particularly nice measurement such as the PGM, which in general may or may not be optimal) can be used as a principle for discovering new quantum algorithms.

While the reduction of Lemma 1 combined with the PGM approach outlined in Section 4 appears to efficiently solve the HSP in most of the semidirect product groups where efficient algorithms are known, there is one notable exception. The groups $\mathbb{Z}_p^n \rtimes \mathbb{Z}_2$ (for which an efficient quantum algorithm is given in [10]) give rise to a subset sum problem over \mathbb{Z}_p^n , which appears to be essentially as difficult as the subset sum problem over \mathbb{Z}_N arising from the dihedral group. Thus, it would be interesting to understand what allows the algorithm in [10] to be efficient even though the matrix sum problem is (apparently) hard.

Of course, there are many nonabelian groups that are semidirect products of nonabelian groups, or that cannot be nontrivially decomposed into semidirect products at all. The PGM approach is well

defined for any group, so it would be interesting to explore the approach in such cases, regardless of whether the PGM is optimal.

Acknowledgments AMC received support from the National Science Foundation under Grant No. EIA-0086038.

References

- [1] N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley Interscience, New York, 2nd edition, 2000.
- [2] D. Bacon, A. M. Childs, and W. van Dam. Optimal measurements for the dihedral hidden subgroup problem. arXiv:quant-ph/0501044.
- [3] R. Beals. Quantum computation of Fourier transforms over symmetric groups. In *Proc. 29th Annual ACM Symposium on Theory of Computing*, pages 48–53, 1997.
- [4] D. Boneh and R. J. Lipton. Quantum cryptanalysis of hidden linear functions. In *Lecture Notes in Computer Science*, volume 963, pages 424–437. Springer-Verlag, Berlin, 1995.
- [5] D. Cox, J. Little, and D. O’Shea. *Ideals, Varieties, and Algorithms*. Springer-Verlag, Berlin, 2nd edition, 1997.
- [6] T. Dubé, B. Mishra, and C.-K. Yap. Complexity of Buchberger’s algorithm for Gröbner bases. 1995.
- [7] M. Ettinger and P. Høyer. A quantum observable for the graph isomorphism problem. arXiv:quant-ph/9901029.
- [8] M. Ettinger and P. Høyer. On quantum algorithms for noncommutative hidden subgroups. *Adv. in Appl. Math.*, 25:239–251, 2000. arXiv:quant-ph/9807029.
- [9] M. Ettinger, P. Høyer, and E. Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Inform. Process. Lett.*, 91:43–48, 2004. arXiv:quant-ph/0401083.
- [10] K. Friedl, G. Ivanyos, F. Magniez, M. Santha, and P. Sen. Hidden translation and orbit coset in quantum computing. In *Proc. 35th Annual ACM Symposium on Theory of Computing*, pages 1–9, 2003. arXiv:quant-ph/0211091.
- [11] D. Gavinsky. Quantum solution to the hidden subgroup problem for poly-near-hamiltonian groups. *Quantum Inf. Comput.*, 4:229–235, 2004.
- [12] M. Grigni, L. Schulman, M. Vazirani, and U. Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. In *Proc. 33rd Annual ACM Symposium on Theory of Computing*, pages 68–74, 2001.
- [13] S. Hallgren, A. Russell, and A. Ta-Shma. Normal subgroup reconstruction and quantum computing using group representations. In *Proc. 32nd Annual ACM Symposium on Theory of Computing*, pages 627–635, 2000.
- [14] P. Hausladen and W. K. Wootters. A ‘pretty good’ measurement for distinguishing quantum states. *J. Mod. Opt.*, 41:2385–2390, 1994.

- [15] A. S. Holevo. Statistical decision theory for quantum systems. *J. Multivar. Anal.*, 3:337–394, 1973.
- [16] P. Høyer. Efficient quantum transforms. arXiv:quant-ph/9702028.
- [17] Y. Inui and F. Le Gall. An efficient algorithm for the hidden subgroup problem over a class of semi-direct product groups. arXiv:quant-ph/0412033.
- [18] L. Ip. Shor’s algorithm is optimal. 2003.
- [19] A. Kitaev. Quantum measurements and the abelian stabilizer problem. arXiv:quant-ph/9511026.
- [20] G. Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. arXiv:quant-ph/0302112.
- [21] C. Moore, D. N. Rockmore, A. Russell, and L. J. Schulman. The power of basis selection in Fourier sampling: Hidden subgroup problems in affine groups. In *Proc. 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1113–1122, 2004. arXiv:quant-ph/0211124; extended version in arXiv:quant-ph/0503095.
- [22] C. Moore, A. Russell, and L. J. Schulman. The symmetric group defies strong Fourier sampling: Part I. arXiv:quant-ph/0501056.
- [23] C. Moore and A. Russell. For distinguishing conjugate hidden subgroups, the pretty good measurement is as good as it gets. arXiv:quant-ph/0501177.
- [24] M. A. Neumark. On a representation of additive operator set functions. *C. R. (Dokl.) Acad. Sci. URSS*, 41:359–361, 1943.
- [25] J. Radhakrishnan, M. Rötteler, and P. Sen. On the power of random bases in Fourier sampling: Hidden subgroup problem in the Heisenberg group. arXiv:quant-ph/0503114.
- [26] O. Regev. Quantum computation and lattice problems. In *Proc. 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 520–529, 2002. arXiv:cs.DS/0304005.
- [27] O. Regev. A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space. arXiv:quant-ph/0406151.
- [28] M. Rötteler and T. Beth. Polynomial-time solution to the hidden subgroup problem for a class of non-abelian groups. arXiv:quant-ph/9812070.
- [29] K. Shoda. Über die Automorphismen einer endlichen Abelschen Gruppe. *Math. Ann.*, 100:674–686, 1928.
- [30] P. W. Shor. Algorithms for quantum computation: Discrete log and factoring. In *Proc. 35th Annual IEEE Symposium on Foundations of Computer Science*, pages 124–134, 1994.
- [31] J. Watrous. Quantum algorithms for solvable groups. In *Proc. 33rd Annual ACM Symposium on Theory of Computing*, pages 60–67, 2001. arXiv:quant-ph/0011023.
- [32] H. P. Yuen, R. S. Kennedy, and M. Lax. Optimum testing of multiple hypotheses in quantum detection theory. *IEEE Trans. Inform. Theory*, 21:125–134, 1975.

A Matrix representation of $\Phi^{(b)}$

In this appendix, we show that $\Phi^{(b)}$ and $\hat{\Phi}^{(b)}$ can be represented by (simply related) matrices, and furthermore, that $\chi_x(\Phi^{(b)}(d)) = \chi_{\hat{\Phi}^{(b)}(x)}(d)$.

Lemma 7. *Let A be a finite abelian group, $\varphi \in \text{Aut}(A)$, $b \in \mathbb{N}$ with the corresponding $\Phi : A \rightarrow A$ defined by $\Phi := \sum_{i=0}^{b-1} \varphi^i$, and $\chi : A \rightarrow \mathbb{C}$ a character of A . Then there exists a function $\hat{\Phi} : A \rightarrow A$ such that $\chi_x(\Phi(d)) = \chi_{\hat{\Phi}(x)}(d)$ for all $d, x \in A$.*

Proof. Let $A \cong A_{p_1} \times \cdots \times A_{p_r}$ be the elementary divisor decomposition of A , i.e., the decomposition into p -groups where the p_i are distinct primes. Accordingly, let $d = (d_1, \dots, d_r)$, $x = (x_1, \dots, x_r)$, and $\chi_{x_i}^{(i)} : A_{p_i} \rightarrow \mathbb{C}$ such that $\chi_x(d) = \chi_{x_1}^{(1)}(d_1) \cdots \chi_{x_r}^{(r)}(d_r)$. It is known that any automorphism $\varphi \in \text{Aut}(A)$ can be decomposed as $\varphi = (\varphi_1, \dots, \varphi_r)$ with $\varphi_i \in \text{Aut}(A_{p_i})$ [29]. Similarly, we have the decomposition $\Phi(d) = (\Phi_1(d_1), \dots, \Phi_r(d_r))$. Hence if we prove the lemma for p -groups, then this proves it for all finite abelian groups.

Assume therefore that A is a p -group, $A \cong \mathbb{Z}_{p^{e_1}} \times \cdots \times \mathbb{Z}_{p^{e_k}}$ with $e_1 \geq \cdots \geq e_k$. By [29], we can represent $\varphi \in \text{Aut}(A)$ as a matrix transformation $(x_1, \dots, x_k) \mapsto (x_1, \dots, x_k)\mu$ where $\mu \in \mathbb{Z}^{k \times k}$ and $p^{e_j - e_i} | \mu_{ij}$ for all $i > j$. Consequently, we can also represent the transformation Φ by a matrix M with $p^{e_j - e_i} | M_{ij}$ for all $i > j$. Now define a conjugate matrix $\hat{M}_{ji} := p^{e_i - e_j} M_{ij}$ for all i, j (note that all entries of \hat{M} are integers). For any character $\chi_x : A \rightarrow \mathbb{C}$, we find

$$\chi_x(\Phi(d)) = \exp\left(2\pi i \sum_{ji} d_i M_{ij} x_j / p^{e_j}\right) = \exp\left(2\pi i \sum_{ij} d_i x_j \hat{M}_{ji} / p^{e_i}\right) = \chi_{\hat{\Phi}(x)}(d), \quad (49)$$

where $\hat{\Phi} : A \rightarrow A$ is the matrix transformation defined by $(x_1, \dots, x_k) \mapsto (x_1, \dots, x_k)\hat{M}$. \square

Note that if $A = \mathbb{Z}_N$ (as in Section 5), then Φ can be represented by a single scalar $M \in \mathbb{Z}_N$, and hence $\hat{\Phi} = \Phi$. Also, if $A = \mathbb{Z}_p^r$ (as in Section 6), then $M \in \mathbb{Z}_p^{r \times r}$ and $\hat{M} = M^T$. For notational simplicity, since we only need the matrix representation of $\hat{\Phi}$ (and not of Φ) in these two cases, we put $\mu \rightarrow \mu^T$ and $M \rightarrow M^T$ throughout the body of the paper.

B Proof of Lemma 5

Proof. For $x \in (\mathbb{Z}_p^n)^k$, $w \in \mathbb{Z}_p^n$, and $b \in \mathbb{Z}_p^k$, let $\Xi_w^x(b) = 0$ denote the system of polynomial equations (48). We want to understand the typical behavior of

$$\eta_w^x := |\{b : \Xi_w^x(b) = 0\}| = \sum_b \delta[\Xi_w^x(b) = 0] \quad (50)$$

for uniformly random x, w . Specifically, we want to show that it is typically close to its mean,

$$\mu := \mathbb{E}_{x \in A^k, w \in A} [\eta_w^x] = p^{k-r} \quad (51)$$

where we have used (14). To compute the variance, note that

$$\mathbb{E}_{x \in A^k, w \in A} [(\eta_w^x)^2] = \frac{1}{p^{rk+r}} \sum_{x,w} (\eta_w^x)^2 \quad (52)$$

$$= \frac{1}{p^{rk+r}} \sum_{x,w} \left(\sum_b \delta[\Xi_w^x(b) = 0] \right) \left(\sum_c \delta[\Xi_w^x(c) = 0] \right) \quad (53)$$

$$= \frac{1}{p^{rk+r}} \sum_{x,w} \left(\sum_b \delta[\Xi_w^x(b) = 0] + \sum_{b \neq c} \delta[\Xi_w^x(b) = \Xi_w^x(c) = 0] \right). \quad (54)$$

The first (diagonal) term is just the mean. For the second (off-diagonal) term, note that the condition $\Xi_w^x(b) = \Xi_w^x(c)$ actually does not depend on w , so we can write

$$\mathbb{E}_{x \in A^k, w \in A} [(\eta_w^x)^2] = \mu + \frac{1}{p^{rk+r}} \sum_{b \neq c} \sum_x \delta[\Xi_0^x(b) = \Xi_0^x(c)] \sum_w \delta[\Xi_w^x(b) = 0] \quad (55)$$

$$= \mu + \frac{1}{p^{rk+r}} \sum_{b \neq c} \sum_x \delta[\Xi_0^x(b) = \Xi_0^x(c)] \quad (56)$$

where we have used the fact that for any fixed x, b , there is exactly one w that satisfies the equation. Now for $b \neq c$, choose some j such that $b_j \neq c_j$. The system of equations $\Xi_0^x(b) = \Xi_0^x(c)$ is linear in x , and in the i th equation (48), the coefficient of x_{ij} is nonzero (it is simply $b_j - c_j$), so for any values of the $x_{ij'}$ for $j' \neq j$, we can solve the equations uniquely for the x_{ij} , giving

$$\mathbb{E}_{x \in A^k, w \in A} [(\eta_w^x)^2] = \mu + \frac{1}{p^{rk+r}} \sum_{b \neq c} \sum_x \delta[\forall i \ x_{ij} \text{ fixed}] \quad (57)$$

$$= \mu + \frac{1}{p^{rk+r}} \sum_{b \neq c} p^{rk-r} \quad (58)$$

$$= \mu + p^{2(k-r)} - p^{k-2r}. \quad (59)$$

Thus we find $\sigma^2 := \mathbb{E}_{x,w} [(\eta_w^x)^2] - \mu^2 = \mu - p^{k-2r} = \mu(1 - p^{-r})$.

Since the variance is small, Chebyshev's inequality shows that the probability of deviating far from the mean number of solutions is small:

$$\Pr(|\eta_w^x - \mu| \geq c) \leq \frac{\sigma^2}{c^2}. \quad (60)$$

For $k = r$, we have $\mu = 1$ and $\sigma^2 = 1 - p^{-r}$, so by putting $c = 2$ in (60), we find $\Pr(\eta_w^x \geq 3) \leq \frac{1}{4}$.

To see that we are unlikely to have no solutions, we need a slightly stronger bound than the Chebyshev inequality. Since η_w^x is a nonnegative, integer-valued random variable, we have [1, p. 58] $\Pr(\eta_w^x = 0) \leq \sigma^2 / (\mu^2 + \sigma^2) = (1 - p^{-r}) / (2 - p^{-r}) \leq 1/2$. Combining these results, we see that $\Pr(\eta_w^x = 1 \text{ or } 2) \geq \frac{1}{4}$ as claimed. \square