

HIDDEN SYMMETRY SUBGROUP PROBLEMS*

THOMAS DECKER[†], GÁBOR IVANYOS[‡], MIKLOS SANTHA[§], AND PAWEŁ WOCJAN[¶]

Abstract. We advocate a new approach for addressing hidden structure problems and finding efficient quantum algorithms. We introduce and investigate the hidden symmetry subgroup problem (HSSP), which is a generalization of the well-studied hidden subgroup problem (HSP). Given a group acting on a set and an oracle whose level sets define a partition of the set, the task is to recover the subgroup of symmetries of this partition inside the group. The HSSP provides a unifying framework that, besides the HSP, encompasses a wide range of algebraic oracle problems, including quadratic hidden polynomial problems. While the HSSP can have provably exponential quantum query complexity, we obtain efficient quantum algorithms for various interesting cases. To achieve this, we present a general method for reducing the HSSP to the HSP, which works efficiently in several cases related to symmetries of polynomials. The HSSP therefore connects in a rather surprising way certain hidden polynomial problems with the HSP. Using this connection, we obtain the first efficient quantum algorithm for the hidden polynomial problem for multivariate quadratic polynomials over fields of constant characteristic. We also apply the new methods to polynomial function graph problems and present an efficient quantum procedure for constant degree multivariate polynomials over any field. This result improves in several ways the currently known algorithms.

Key words. quantum algorithms, hidden subgroup problem, hidden polynomial problem

AMS subject classifications. 68Q12, 81P68

DOI. 10.1137/120864416

1. Introduction. The main goal of quantum computing is to identify suitable classes of problems and to find efficient quantum algorithms for them that provide a significant speed-up over their classical counterparts. The vast majority of such examples consists of group-theoretical problems that can be formulated within the framework of the hidden subgroup problem (HSP). This problem can be cast in the following terms: We are given a finite group G and a black-box function from G to some finite set. The level sets of the function correspond to the right cosets of some subgroup H . We say that f hides H , and the task is to determine this hidden subgroup. One query of the function counts as one step in the computation, and an algorithm is efficient if its running time is polynomial in the logarithm of the size of the group. While no classical algorithm is known to solve this problem with polynomial

*Received by the editors February 1, 2012; accepted for publication (in revised form) July 12, 2013; published electronically October 16, 2013. Most of this work was conducted while the first, second, and third authors were at the Centre for Quantum Technologies (CQT) in Singapore, partially funded by the Singapore Ministry of Education and the National Research Foundation, and also through the Tier 3 grant “Random Numbers from Quantum Processes.” This research was partially supported by the European Commission IST STREP projects Quantum Computer Science (QCS) 255961 and Quantum Algorithms (QALGO) 600700, by the French ANR Blanc program under contract ANR-12-BS02-005 (RDAM project), and by the Hungarian Scientific Research Fund (OTKA, grants NK105645 and K77476).

<http://www.siam.org/journals/sicomp/42-5/86441.html>

[†]Centre for Quantum Technologies, National University of Singapore, Singapore 117543 (t.d3ck3r@gmail.com).

[‡]Institute for Computer Science and Control, Hungarian Academy of Sciences, Budapest, Hungary (Gabor.Ivanyos@sztaki.mta.hu).

[§]LIAFA, University Paris 7, CNRS, 75205 Paris, France, and Centre for Quantum Technologies, National University of Singapore, Singapore 117543 (miklos.santha@liafa.univ-paris-diderot.fr).

[¶]Department of Electrical Engineering and Computer Science, University of Central Florida, Orlando, FL 32816-2362 (wocjan@eecs.ucf.edu). The work of this author was supported by NSF grant CCF-0726771 and NSF CAREER award CCF-0746600.

query complexity, the problem is computationally solvable in quantum polynomial time for every abelian group [1, 2, 3].

Several attempts were made to extend the quantum solution of the abelian HSP. Most of the research focused on the HSP in nonabelian groups since these include several algorithmically important problems. For example, it is known that efficient solutions for the dihedral and the symmetric group would imply efficient solutions for some lattice problems [4] and for graph isomorphism, respectively. While some progress has been made in this direction [5, 6, 7, 8, 9, 10, 11], the HSP for the dihedral and symmetric groups remains unsolved. It is already known that the methods for solving the abelian case fail for several nonabelian groups [12, 13]. The goal of obtaining efficient quantum algorithms for larger classes of nonabelian groups turned out to be rather elusive.

Another idea for generalizing the problem was proposed by Childs, Schulman, and Vazirani [14], who considered properties of algebraic sets hidden by black-box functions. One of these problems is the hidden polynomial problem (HPP), where the hidden object is a polynomial. To recover it we have at our disposal an oracle whose level sets coincide with the level sets of the polynomial. Childs et al. [14] showed that the quantum query complexity of this problem is polynomial in the logarithm of the field size, provided that the degree and the number of variables are held constant, leaving open the question of the time complexity. The authors also formulated computationally efficient quantum procedures for some related problems, such as the hidden radius and the hidden flat of centers. Nonetheless, to the best of our knowledge, no efficient quantum polynomial time algorithm has been proposed for the general HPP, not even for the simplest problem of hidden quadratic polynomials in one variable (HQPP).

In [15], Decker, Draisma, and Wocjan defined a related problem that we refer to as the hidden polynomial graph problem (HPGP) to distinguish it from the HPP. Here, similarly to the HPP, the hidden object is a polynomial, but the oracle is more powerful. They obtained a polynomial time quantum algorithm that correctly identifies the hidden polynomial when the degree and the number of variables are considered to be constant. Their proof applies to all finite fields whose characteristic is not in a finite set of exceptional characteristics that depend on the degree of the polynomials.

In this paper, we advocate a third possible approach for finding hidden structures. We consider a group G acting on some finite set M , and we suppose that we have at our disposal a black-box function whose level sets define a partition of M . The object we would like to recover is the group of symmetries of this partition inside G , i.e., the largest subgroup whose orbits under the action coincide with the classes of the partition. We call this problem the hidden symmetry subgroup problem (HSSP). It is easy to see that the HSP is a special case of the HSSP when the group acts on itself and the action corresponds to the group operation. But, for some actions, the HSSP is provably harder than any HSP. We show that Grover's search can be cast as an HSSP, establishing that certain cases of the HSSP have exponential quantum query complexity. This is in contrast to the HSP, which has polynomial quantum query complexity for all groups [16].

The potential of the HSSP lies mainly in the possibility of extending the HSP techniques to more general group actions that still admit efficient quantum procedures. We demonstrate the power of this new approach by designing and improving quantum algorithms for several algebraic problems. To achieve this we reduce both the HQPP and the univariate HPGP to appropriate HSSPs for which we can give efficient

quantum solutions in some interesting cases. Besides the construction of efficient algorithms, the formulation of problems as HSSP can also shed new light on their structure. For example, the apparent difficulty of the HQPP over prime fields might be explained by the equivalence of this problem to the HSP in the dihedral group, a connection discovered via their relations to the HSSP. It is also worth noting that the hidden shifted multiplicative character problem of van Dam, Hallgren, and Ip [17] is a version of the HSSP with an additional promise on the input.

To establish our algorithmic results, we first concentrate on the question of whether the HSSP can be reduced in some cases to the related HSP that we obtain by forgetting about the action. We design a reduction scheme, which involves the generalization of bases known from the theory of permutation groups. We are able to show that when the action has an efficiently computable generalized base then the HSSP is indeed efficiently reducible to the related HSP (Proposition 3.4). Then we describe a probabilistic construction of such bases for a large class of Frobenius groups. Therefore, the above reduction applies to these groups (Theorem 4.4). These groups include among others a large variety of affine groups, and the HSSP is efficiently solvable for these groups by a quantum algorithm. We remark that in [18] it is proved that the HSSP (in a slightly different formulation) can be solved efficiently for some of these affine groups. The proof uses essentially the same reduction technique.

We then establish several surprising connections between hidden polynomial problems and the HSSP. In fact, the HQPP turns out to be equivalent in a very strong sense to the HSSP over a related affine group. Combined with the above reduction to the related HSP, we are able to give the first ever quantum polynomial time solution for the HQPP over fields of constant characteristic (Theorem 4.7). We then give a quantum reduction of the multivariate quadratic HPP to the HQPP, which implies that over fields of constant characteristic this multivariate problem is also solvable in quantum polynomial time (Theorem 4.9).

Finally, for dealing with the HPGP, we define a class of semidirect product groups which we call function graph groups. We show that the HPGP for univariate polynomials of degree at most d coincides with the HSSP over a corresponding function graph group. These groups turn out to have a base of size d , and therefore our general reduction to the related HSP applies (Theorem 5.5). Based on this reduction, we improve the results of [15] by showing that there is a quantum polynomial time algorithm for the HPGP over every field when the degree of the polynomials is constant (Theorem 5.7).

2. Preliminaries. We first fix some useful notation: n denotes a positive integer, p a prime number, q a prime power, \mathbb{Z}_n the additive group of integers modulo n , \mathbb{F}_q the finite field of size q , and $\mathbb{F}_q^{(d)}[x]$ the set of univariate polynomials of degree at most d over \mathbb{F}_q .

2.1. Level sets and problem classes. Simply speaking, we study the general problem of determining hidden objects related to a given algebraic structure. The algebraic structure is specified by parameters of the problem, which are finite groups, families of subgroups of a given group, group actions, finite fields, and integers in the present case. We assume that we have access to an unknown member of a family of black-box functions $f : A \rightarrow S$, where A is part of the structure and S is some finite set. We consider this function f as the oracle input. We are restricted to identifying the hidden object solely from the information we obtain by querying the oracle f . In fact, the only useful information we can obtain is the structure of the *level sets* $f^{-1}(s) = \{a \in A : f(a) = s\}$, $s \in S$; that is, we can only determine whether two

elements in A are mapped to the same value or not. All nonempty level sets together constitute a partition of A which we denote by π_f .

DEFINITION 2.1. *The hidden subgroup problem HSP is parametrized by a finite group G and a family \mathcal{H} of subgroups of G .*

HSP(G, \mathcal{H})

Oracle input: A function f from G to some finite set S such that for some subgroup $H \in \mathcal{H}$, we have $f(x) = f(y) \iff Hx = Hy$.

Output: H .

The hidden polynomial problem HPP is parametrized by a finite field \mathbb{F}_q and two positive integers n and d .

HPP(\mathbb{F}_q, n, d).

Oracle input: A function f from \mathbb{F}_q^n to some finite set S such that for some n -variate polynomial \mathcal{P} of degree d over \mathbb{F}_q , we have $f(x) = f(y) \iff \mathcal{P}(x) = \mathcal{P}(y)$.

Output: \mathcal{P} .

For every $u \in \mathbb{F}_q$ we define a monic quadratic polynomial over \mathbb{F}_q by $\mathcal{P}_u(x) = x^2 - 2ux$. The hidden quadratic polynomial problem HQPP is parametrized by some finite field \mathbb{F}_q .

HQPP(\mathbb{F}_q).

Oracle input: A function f from \mathbb{F}_q to some finite set S such that we have $f(x) = f(y) \iff \mathcal{P}_u(x) = \mathcal{P}_u(y)$.

Output: \mathcal{P}_u (or just u).

The hidden polynomial graph problem HPGP is parametrized by a finite field \mathbb{F}_q and two positive integers n and d .

HPGP(\mathbb{F}_q, n, d).

Oracle input: A function f from $\mathbb{F}_q^n \times \mathbb{F}_q$ to a finite set S such that for some n -variate polynomial Q of degree d over \mathbb{F}_q we have $f(x_1, y_1) = f(x_2, y_2) \iff y_1 - Q(x_1) = y_2 - Q(x_2)$.

Output: Q .

In all these problems we say that the input f hides the output of the problem.

In the definition of the HQPP we restrict our attention to monic polynomials with zero constant term, because adding a constant to a polynomial or multiplying all coefficients with the same nonzero constant does not change the partition π_f . Furthermore, observe that the HPGP is a special case of the HPP in $n+1$ variables, where the dependence on the $(n+1)$ st variable is linear. Also, the fact that an HPGP oracle $f(x, y)$ restricted to $y = 0$ is equivalent to an HPP oracle can be interpreted as the HPGP being a version of HPP with a more powerful oracle.

In all these problems the task is to determine the output hidden by the oracle input. We measure the time complexity of an algorithm by the overall running time when a query counts as one computational step. An algorithm is *efficient* if its time complexity is polynomial in the logarithm of the size of the group or field, and in the size of the integers in unary in the parametrization of the problem.

2.2. Semidirect product groups. Let K and H be finite groups and let $\phi : h \mapsto \phi_h$ be a homomorphism from H to the group of automorphisms of K . Then the semidirect product $K \rtimes_{\phi} H$ is the cartesian product of K and H equipped with the multiplication defined as $(k, h) \cdot (k', h') = (k \cdot \phi_h(k'), h \cdot h')$. We use the notation $K \rtimes H$ for $K \rtimes_{\phi} H$ whenever ϕ is clear from the context.

2.3. Group actions and partitions. A *left permutation action* of a group G on a set M is a binary function $\circ : G \times M \rightarrow M$, where we denote $\circ(g, m)$ by $g \circ m$, which for all $g, h \in G$ and $m \in M$ satisfies $g \circ (h \circ m) = (gh) \circ m$ and $e \circ m = m$ for the identity element e of G . For a subset $L \subseteq M$ we set $g \circ L = \{g \circ m : m \in L\}$. The *stabilizer* subgroup G_m of m is defined as $\{g \in G : g \circ m = m\}$, the set of elements in G which fix m . The action \circ is *faithful* if $\bigcap_{m \in M} G_m = \{e\}$. Throughout the paper we assume faithfulness. If G acts on M , then every subgroup H of G acts also naturally on M . The *H -orbit* of $m \in M$ is the set of elements of M to which m can be moved by elements of H , formally $H \circ m = \{h \circ m : h \in H\}$.

For each subgroup H , the H -orbits form a partition $H^* = \{S : \exists m \in M \text{ such that } S = H \circ m\}$ of M . For a partition $\pi = \{\pi_1, \dots, \pi_\ell\}$ of the set M , we define the subgroup $\pi^* = \{g \in G : (\forall i) g \circ \pi_i = \pi_i\}$. We call $\pi^* \leq G$ the group of *symmetries* of π within G . This is the subgroup of elements that stabilize every class of the partition π under the given action. Let $(\mathcal{S}(G), \subseteq)$ be the lattice of subgroups of G under the inclusion relation, and let $(\Pi(M), \leq)$ be the lattice of partitions of M , where by definition $\pi \leq \pi'$ if π' is finer than π . The maps $H \mapsto H^*$ and $\pi \mapsto \pi^*$ define an order-reversing Galois connection between $(\mathcal{S}(G), \subseteq)$ and $(\Pi(M), \leq)$, that is, $H \subseteq \pi^*$ if and only if $\pi \leq H^*$. The subgroup H^{**} is the *closure* of H [19]; it consists of the elements in G which stabilize every H -orbit. The *closure* of a partition π is π^{**} ; it consists of the orbits of its group of symmetries. It is always true that $H^{**} \supseteq H$ and $\pi^{**} \leq \pi$. The subgroup H is *\circ -closed* (or just closed if \circ is clear from the context) if $H = H^{**}$, or, equivalently, there exists a partition π such that $H = \pi^*$. Similarly, π is *closed* if $\pi = \pi^{**}$. We denote by $\mathcal{C}(G)$ the family of all closed subgroups in G .

2.4. The hidden symmetry subgroup problem. We now have all prerequisites to define the new problem class.

DEFINITION 2.2. *The hidden symmetry subgroup problem HSSP is parametrized by a finite group G , a finite set M , an action $\circ : G \times M \rightarrow M$ of G on M , and a family \mathcal{H} of closed subgroups of G .*

HSSP(G, M, \circ, \mathcal{H}).

Oracle input: A function f from M to some finite set S such that for some subgroup $H \in \mathcal{H}$, we have $f(x) = f(y) \iff H \circ x = H \circ y$.

Output: H .

In general, there can be several subgroups whose orbits coincide with the level sets of f , but the closures of these subgroups are the same. The unique closed subgroup that satisfies the promise is π_f^* , and this is exactly the output of the problem. We will say that f *hides H by symmetries*. In fact, it would be natural to extend HSSP to the more general setting where f is an arbitrary function on M and the task is to determine the (closed) subgroup π_f^* . The restriction we use in this paper is that π_f is a closed partition with $\pi_f^* \in \mathcal{H}$. We define an algorithm for solving the HSSP as *efficient* if it is polylogarithmic in $|G|$.

It is easy to see that the HSP is a special case of the HSSP when we set $M = G$ and choose the group action \circ to be the group operation, that is, $g \circ h = gh$. For this action every subgroup H of G is closed. Indeed, the stabilizer of any left coset of H is H , and hence H belongs to the partition of G into the left cosets of H . Furthermore, a function f hides a subgroup H if and only if f hides H by symmetries.

Given HSSP(G, M, \circ, \mathcal{H}), by forgetting about the action we obtain HSP(G, \mathcal{H}). We call this problem the *related HSP*.

2.5. Related results. While the HSP is generally hard in nonabelian groups, its query complexity is always small, due to a classical result of Ettinger, Høyer, and Knill [16].

FACT 1. *For every finite G , the $\text{HSP}(G, \mathcal{C}(G))$ has polynomial query complexity.*

Here $\mathcal{C}(G)$ is just the set of all the subgroups of G , since for the action corresponding to the HSP, every subgroup is closed. Among groups where the HSP is solvable in quantum polynomial time, some affine groups will be of importance for us. For a subgroup H of \mathbb{F}_q^* , let $\text{Aff}_q(H)$ denote the semidirect product $\mathbb{F}_q \rtimes H$, and let \mathcal{FC} be the family of conjugates of H by an element of \mathbb{F}_q (for a detailed discussion of these groups see section 4.2). The following positive results on the solvability of the HSP were obtained, respectively, by Moore et al. [11] and Friedl et al. [7].

FACT 2. *The following cases of the HSP can be solved in polynomial time:*

- (a) $\text{HSP}(\text{Aff}_q(H), \mathcal{FC})$, where q is a prime and $H \leq \mathbb{F}_q^*$ such that $1 < |H| \leq q-1$ and $|H| = \Omega(q/\text{polylog}(q))$.
- (b) $\text{HSP}(G, \mathcal{C}(G))$, where G is a finite group such that G' is commutative and every element of G' has an order bounded by a constant.

The query complexity of the HPP was investigated by Childs, Schulman, and Vazirani [14]. They showed the following.

FACT 3. *If $n \geq 2$ and d are constants, then for an $1 - o(1)$ fraction of the hidden polynomials, $\text{HPP}(\mathbb{F}_q, n, d)$ has polylogarithmic query complexity.*

Here, like in the case of the HQPP, polynomials are determined up to constant terms and scalar factors. We are not aware of any results regarding the quantum computational complexity even in the univariate quadratic case. Rötteler showed that multivariate quadratic (not hidden) Boolean functions can be identified using a linear number of quantum queries [20]. For the HPGP, Decker, Draisma, and Wocjan [15] showed the following.

FACT 4.

- (a) $\text{HPGP}(\mathbb{F}_q, n, d)$ can be reduced in polynomial time to $\text{HPGP}(\mathbb{F}_q, 1, d)$ for every constant n .
- (b) For every d there exists a finite set E_d of primes such that if d is constant and the characteristic of \mathbb{F}_q is not in E_d , then $\text{HPGP}(\mathbb{F}_q, 1, d)$ can be solved in quantum polynomial time.

3. A general reduction of the HSSP to the HSP. How much greater is the complexity of an HSSP compared to the complexity of the related HSP? To analyze this, we first give a simple example, which shows that the query complexity of the HSSP can be exponentially higher than the query complexity of the related HSP. Then, more interestingly, we will establish a general condition on the group action under which the HSSP can be reduced in polynomial time to the related HSP.

3.1. HSSP with exponential query complexity. While the quantum query complexity of the HSP is polylogarithmic in the size of the group, we show in this section that the query complexity of an HSSP can be in the order of $|G|^{1/4}$. More precisely, we show that Grover's search problem can be reduced to some specific HSSP.

For a prime power q , the *general affine group* Aff_q of invertible affine transformations over \mathbb{F}_q is defined as the semidirect product $\mathbb{F}_q \rtimes \mathbb{F}_q^*$, where \mathbb{F}_q^* denotes the multiplicative group of \mathbb{F}_q . The natural action of Aff_q on \mathbb{F}_q is defined as $(b, a) \circ x = ax + b$. For every $c \in \mathbb{F}_q$, the stabilizer of c is the subgroup $H_c = \{((1-a)c, a) : a \in \mathbb{F}_q^*\}$, which has two orbits, $\{c\}$ and $\{d \in \mathbb{F}_q : d \neq c\}$, and thus Aff_q is 2-transitive. Clearly, H_c is a closed subgroup. We set $\mathcal{H} = \{H_c : c \in \mathbb{F}_q\}$.

PROPOSITION 3.1. *The query complexity of $\text{HSSP}(\text{Aff}_q, \mathbb{F}_q, \circ, \mathcal{H})$ is $\Omega(q^{1/2})$.*

Proof. Grover's search over \mathbb{F}_q can be trivially reduced to this HSSP. Indeed, if the oracle input is f_c , defined by $f_c(x) = \delta_{c,x}$, where $\delta_{c,x}$ is the Kronecker delta, then f_c hides H_c as symmetry subgroup. From any generator (b, a) of H_c one recovers c simply by computing $(1-a)^{-1}b$. Hence, the query complexity of the HSSP is at least the query complexity $\Omega(q^{1/2})$ of Grover's search [21]. \square

3.2. A reduction scheme of the HSSP to the HSP. In this section, we describe a rather natural framework for reducing the HSSP to the related HSP. Essentially, the same idea was used in [18] for reducing certain hidden shift problems to the HSP in the affine group over prime fields. We assume that we are given a black-box function f over M , which hides some subgroup H of G by symmetries. With the help of f , we would like to construct a suitable function f_{HSP} over G , which hides H . A first approach could be to define $f_{\text{HSP}}(g) = f(g \circ m)$, where m is a fixed element of M . Unfortunately, this works only in very exceptional cases because f_{HSP} takes constant values on the left cosets of the stabilizer H_m of m . Therefore, even in the simple case when f hides the trivial subgroup, the function f_{HSP} will not work unless the stabilizer of m is trivial. As a straightforward refinement of this idea, we can pick several elements $m_1, \dots, m_t \in M$ and define

$$f_{\text{HSP}}(g) = (f(g \circ m_1), \dots, f(g \circ m_t)).$$

For the trivial hidden subgroup, this idea works when the common stabilizer of m_1, \dots, m_t is trivial, that is, when $\bigcap_{i=1}^t H_{m_i} = \{e\}$. In the theory of permutation groups such a system of elements is called a *base* [22]. Of course, bases exist only if the action of G is faithful. The following definition includes further conditions on m_1, \dots, m_t in order to make the above construction work in general.

DEFINITION 3.2. *Let G be a finite group and let $\circ : G \times M \rightarrow M$ be an action of G on the finite set M . Let $H \leq G$ be a subgroup of G , and let \mathcal{H} be a family of subgroups of G . A set $B \subseteq M$ is an H -strong base if for every $g \in G$, we have*

$$\bigcap_{m \in B} HG_{g \circ m} = H.$$

We call B an \mathcal{H} -strong base when it is H -strong for every subgroup $H \in \mathcal{H}$.

Observe that a strong $\{e\}$ -base is just a base in the conventional sense, and also that $\bigcap_{m \in M} HG_m = H^{**}$. Hence, M itself is always a $\mathcal{C}(G)$ -strong base. If B is an H -strong base, then B is also an (xHx^{-1}) -strong base for every $x \in G$, because $G_{xg \circ m} = xG_{g \circ m}x^{-1}$. Therefore, if \mathcal{H} consists of conjugated subgroups, then B is an \mathcal{H} -strong base if it is an H -strong base for some $H \in \mathcal{H}$. Also, if \mathcal{H} is closed under conjugation by elements of G , B is an \mathcal{H} -strong base if and only if $\bigcap_{m \in B} HG_m = H$ for every $H \in \mathcal{H}$.

The following lemma states that the HSSP is indeed reducible to the HSP via an \mathcal{H} -strong base.

LEMMA 3.3 (reduction of HSSP to HSP). *Let G be a finite group, and let \circ be an action of G on M . Suppose that the function $f : G \rightarrow S$ hides some $H \in \mathcal{H}$ by symmetries. Let $B = \{m_1, \dots, m_t\}$ be an \mathcal{H} -strong base. Then H is hidden by the function $f_{\text{HSP}}(g) = (f(g \circ m_1), \dots, f(g \circ m_t))$.*

Proof. We will show that for every $x, y \in G$, we have $f_{\text{HSP}}(x) = f_{\text{HSP}}(y)$ if and only if $y \in Hx$. To see the "only if" part, suppose that $f_{\text{HSP}}(x) = f_{\text{HSP}}(y)$. Then by definition $f(x \circ m) = f(y \circ m)$ for every $m \in B$. Therefore, for every $m \in B$ there

exists an element $h_m \in H$ such that $x \circ m = h_m \circ (y \circ m)$. This equality implies that $m = (x^{-1}h_my) \circ m$, that is, $x^{-1}h_my \in G_m$. Thus $y \in h_m^{-1}xG_m$ for every $m \in B$, from which we can deduce $y \in \bigcap_{m \in B} HxG_m$. Now observe that $xG_mx^{-1} = G_{x \circ m}$, and therefore $y \in \bigcap_{m \in B} HG_{x \circ m}x$. From this we can conclude $y \in Hx$ because B is an \mathcal{H} -strong base.

To show the reverse implication, suppose that $y = hx$ for some $h \in H$. This implies $y \circ m = h \circ (x \circ m)$ for all $m \in B$. Since f hides H as symmetry subgroup, we have $f(y \circ m) = f(x \circ m)$, again for all $m \in B$, implying $f_{\text{HSP}}(y) = f_{\text{HSP}}(x)$ by the definition of f_{HSP} . \square

The following statement is immediate from Lemma 3.3.

PROPOSITION 3.4. *Let G be a finite group, M a finite set, \circ a polynomial time computable action of G on M , and \mathcal{H} a family of subgroups of G . If there exists an efficiently computable \mathcal{H} -strong base in M , then $\text{HSSP}(G, M, \circ, \mathcal{H})$ is polynomial time reducible to $\text{HSP}(G, \mathcal{H})$.*

Together with Proposition 3.1, this result demonstrates that, in contrast to ordinary bases, finding (and even understanding the existence of) strong bases can be quite difficult. The results in the rest of the paper rely on constructing strong bases in two different contexts.

4. The HSSP for Frobenius complements and the HQPP. In view of Proposition 3.4, we are interested in group actions for which there exist easily computable (and therefore also small) bases for some interesting families of subgroups. If in addition the related HSP is easy to solve, then we have efficiently solvable HSSPs. It turns out not only that Frobenius groups under some conditions have these properties, but also that the HQPP can be cast as one of these HSSPs.

4.1. Strong bases in Frobenius groups. A *Frobenius group* is a transitive permutation group acting on a finite set such that only the identity element has more than one fixed point and some nontrivial element fixes a point (see, for example, [23]). Let us recall here some notions and facts about these groups. Let G be a Frobenius group with action \circ_M on M . The identity element together with the elements of G that have no fixed points form a normal subgroup K , the *Frobenius kernel*, for which we also have $|K| = |M|$. This latter fact and that K is closed under conjugation are easy to prove. Surprisingly, all the known proofs for the statement that K is a subgroup require representation theory. A subgroup H of G is a *Frobenius complement* if it is the stabilizer H_m of some element $m \in M$. It is a subgroup complementary to K , that is, $K \cap H = \{1\}$ and $G = KH$. Hence, the group G is a semidirect product $K \rtimes H$ of K and H . We define the binary operation $\circ_K : G \times K \rightarrow K$ by

$$g \circ_K x = yhxh^{-1},$$

when $x \in K$ and $g = yh$ with $y \in K$ and $h \in H$. It is a straightforward computation to check that \circ_K is an action of G on K . Observe that K acts on itself by multiplication from the left, while H acts on K by conjugation. Furthermore, we can identify the action \circ_M with the action \circ_K via the map $\phi : M \rightarrow K$ defined as follows. For any $n \in M$, there exists $g_n \in G$ such that $g_n \circ_M m = n$ since G is transitive. If $g_n = y_n h_n$ with $y_n \in K$ and $h_n \in H$, by definition we set $\phi(n) = y_n$. Note that ϕ depends on the choice of m , or, equivalently, on the choice of the complement H . Observe also that $\phi(n)$ can be characterized as the unique element y_n of K with $y_n \circ_M m = n$, and therefore ϕ is a bijection. Then indeed for every $g \in G$ and $n \in M$, we have $g \circ_K \phi(n) = \phi(g \circ_M n)$. From now on we will suppose without loss of generality that the action is \circ_K , which we denote for simplicity by \circ .

Observe then that with respect to \circ , the Frobenius complement H is the stabilizer of e , the identity element of K . The orbits of H are $\{e\}$ and some other subsets of K , each consisting of $|H|$ elements. The other Frobenius complements are $H_x = xHx^{-1}$ for $x \in K$. They are closed subgroups and their orbits form closed partitions $\{\{x\}, K \setminus \{x\}\}$. We denote by \mathcal{FC} the set of Frobenius complements in G . Since the Frobenius complements are all conjugates of H , being an \mathcal{FC} -strong base is equivalent to being an H -strong base.

To characterize H -strong bases it will be convenient to use the following notion. For $u, v \in K$ with $u \neq v$, we say that $z \in K$ separates u and v if $v \circ z \notin H \circ (u \circ z)$, or, in other words, vz is not a conjugate of uz by an element of H . We have the following characterization.

LEMMA 4.1. *Let $B \subseteq K$. Then B is an H -strong base if and only if for all $u \neq v$ in K there exists $z \in B$ which separates u and v .*

Proof. To see the “if” part of the statement, suppose that $g' \in \bigcap_{z \in B} HG_{g \circ z}$ for some $g', g \in G$. We will prove that $g' \in H$. Let $g = yh$ and $g' = y'h'$, where $y, y' \in K$ and $h, h' \in H$. Then for every $z \in B$, there exists $h_z \in H$ such that $g' \circ (g \circ z) = h_z \circ (g \circ z)$. Using the definition of \circ , this equality can be rewritten as $y'h'yzh^{-1}h'^{-1} = h_z y h z h^{-1} h_z^{-1}$. Multiplying both sides by $h^{-1}h'^{-1}$ from the left and by $h'h$ from the right gives $h^{-1}h'^{-1}y'h'yz = h^{-1}h'^{-1}h_z h h^{-1}y h z h^{-1}h_z^{-1}h'h$. Put $u = h^{-1}yh$, $v = h^{-1}h'^{-1}y'h'yh$, and $h''_z = h^{-1}h'^{-1}h_z h$. Then we can rewrite the equality as $vz = h''_z u z h''_z^{-1}$. We have $u, v \in K$ and $h''_z \in H$. Using that K acts on itself by multiplication, and that the action of H is conjugation, we obtain $v \circ z = h''_z \circ u \circ z$. This means that for every $z \in B$, we have $v \circ z \in H \circ (u \circ z)$; that is, no element in B separates u and v . Therefore, by the assumption we get $u = v$, which is equivalent to $y' = e$. Thus $g' = h'$ is indeed an element of H .

To see the reverse implication, assume that there exist $u, v \in K, u \neq v$, such that none of the elements $z \in B$ separate u and v . This means that for every $z \in B$ there exists an element $h_z \in H$ such that $v \circ z = h_z \circ (u \circ z)$. Using $v \circ z = (vu^{-1}) \circ (u \circ z)$, this equality implies $vu^{-1}(u \circ z) = h_z \circ (u \circ z)$, whence $h_z^{-1}vu^{-1}(u \circ z) = u \circ z$, that is, $h_z^{-1}vu^{-1} \in G_{u \circ z}$. This gives $vu^{-1} \in h_z G_{u \circ z} \subseteq HG_{u \circ z}$ for every $z \in B$, that is, $vu^{-1} \in \bigcap_{z \in B} HG_{u \circ z}$. As $vu^{-1} \notin H$, this contradicts the definition of a strong base. \square

Our next lemma gives a lower bound on the number of elements in K that separate u and v .

LEMMA 4.2. *Let $|H| \neq |K| - 1$. Then for any two distinct elements u and v of K we have*

$$|\{z \in K : z \text{ separates } u \text{ and } v\}| > |K|/2.$$

Proof. If z does not separate u and v , then there exists an element $h \in H$ such that $vz = huzh^{-1}$ which can also be written as $hu^{-1}h^{-1}v = hzh^{-1}z^{-1}$. We say that such an element h belongs to z . The identity element $h = e$ does not belong to any element $z \in K$ since $u \neq v$. We claim that $h \neq e$ cannot belong to two distinct elements of K . Indeed, if $hzh^{-1}z^{-1} = hz'h^{-1}z'^{-1}$, then $hz'^{-1}zh^{-1} = z'^{-1}z$, which in turn implies that $z'^{-1}z = e$ as e is the only element of K stabilized by the elements of H . Therefore, there are at most $|H| - 1$ elements in K which do not separate u and v . In other words, at least $|K| - |H| + 1$ of the elements of K separate u and v . Note that H has $(|K| - 1)/|H|$ orbits of length $|H|$ on the nontrivial elements of K , and thus $|H|$ divides but is not equal to $|K| - 1$, which implies $|H| \leq (|K| - 1)/2$. From this we can indeed conclude, since then $|K| - |H| + 1 > |K|/2$. \square

We have the following result regarding the existence of small strong bases for \mathcal{FC} .

PROPOSITION 4.3. *Let G be a Frobenius group with kernel K such that the cardinality of the Frobenius complements is different from $|K| - 1$. Let $B \subseteq K$ be a uniformly random set of size ℓ , where $\ell = \Theta(\log |K| \log 1/\epsilon)$. Then B is an \mathcal{FC} -strong base with probability of at least $1 - \epsilon$.*

Proof. Let B be a uniformly random subset of K of size ℓ . By Lemma 4.1 it is sufficient to prove that with a probability of at least $1 - \epsilon$, for every $u \neq v$, there exists an element in B which separates u and v . We will in fact upper bound the probability of the opposite event. For a fixed pair $u \neq v$, by Lemma 4.2, the probability that a random z does not separate u and v is at most $1/2$. Therefore, the probability that none of the elements in B separates u and v is less than $2^{-\ell}$. Thus, the probability that for some pair $u \neq v$ none of the elements in B separates u and v is less than $\binom{|K|}{2} 2^{-\ell}$, which is at most ϵ by the choice of ℓ . \square

If G is a Frobenius group that satisfies the condition of Proposition 4.3, then we can compute efficiently a small base for the Frobenius complements, because there are efficient algorithms for random sampling nearly uniformly in black-box groups [24]. Therefore, by Proposition 3.4 we can efficiently reduce the HSSP to the related HSP, and we obtain the following result.

THEOREM 4.4. *Let $G = K \rtimes H$ be a Frobenius group with action \circ such that $|H| < |K| - 1$. Then $\text{HSSP}(G, K, \circ, \mathcal{FC})$ is reducible in probabilistic polynomial time to $\text{HSP}(G, \mathcal{FC})$.*

We remark that the reduction of Grover's search to a specific HSSP in Proposition 3.1 can be extended to arbitrary Frobenius groups when $|H| = |K| - 1$, that is, sharply 2-transitive groups. Therefore, for such groups it not only follows that small H -bases fail to exist, but it also follows that even the quantum query complexity of the HSSP is $\Omega(|G|^{1/4})$. Also, the only strong base in a sharply 2-transitive group is the whole K .

4.2. Affine groups. As any affine group, the general affine group $\text{Aff}_q = \mathbb{F}_q \rtimes \mathbb{F}_q^*$ defined in section 3.1 is a Frobenius group. Its kernel is \mathbb{F}_q . In the terminology of Frobenius groups, we have proved in Proposition 3.1 that for Aff_q the HSSP for the complements is difficult. Let H be a proper subgroup of \mathbb{F}_q^* which is not the trivial group. We define the group $\text{Aff}_q(H)$ as $\mathbb{F}_q \rtimes H$. With the restriction of the natural action, denoted here by \circ , $\text{Aff}_q(H)$ is also a Frobenius group. In contrast to the difficulty in the full affine group, we obtain the following positive results for the smaller Frobenius groups. They are consequences of the analogous results for the related HSP stated in Facts 1 and 2, via the reduction of Theorem 4.4. Statements (a) and (b) are not new; they are proved in a slightly different formulation in [18], using implicitly the randomized construction for a strong base. For (c) note that the derived subgroup in $\text{Aff}_q(H)$ is indeed commutative.

COROLLARY 4.5. *Let q be a prime power and let $H \leq \mathbb{F}_q^*$ such that $1 < |H| < q - 1$. The following results hold for $\text{HSSP}(\text{Aff}_q(H), \mathbb{F}_q, \circ, \mathcal{FC})$:*

- (a) *It has polynomial query complexity.*
- (b) *It can be solved in quantum polynomial time when q is prime and $|H| = \Omega(q/\text{polylog}(q))$.*
- (c) *It can be solved in quantum polynomial time when q is the power of a fixed prime.*

The case of $\text{Aff}_q(\{\pm 1\})$, when q is an odd prime power, is particularly interesting. If q is itself an odd prime, then $\text{Aff}_q(\{\pm 1\})$ is just the familiar dihedral group D_q . It turns out that the HSSP over $\text{Aff}_q(\{\pm 1\})$ for the Frobenius complements is essentially the same problem as the HQPP over \mathbb{F}_q .

PROPOSITION 4.6. *The following problems are polynomially equivalent:*

1. HQPP(\mathbb{F}_q).
2. HSSP($\text{Aff}_q(\{\pm 1\}), \mathbb{F}_q, \circ, \mathcal{FC}$).
3. HSP($\text{Aff}_q(\{\pm 1\}), \mathcal{FC}$).

Proof. The first two problems are equivalent as we claim that every $f : \mathbb{F}_q \rightarrow S$, as oracle input for HQPP(\mathbb{F}_q) hides the polynomial \mathcal{P}_u if and only if as oracle input for HSSP($\text{Aff}_q(\{\pm 1\}), \mathbb{F}_q, \circ, \mathcal{FC}$) it hides the Frobenius complement H_u . To see this, observe that the level sets of \mathcal{P}_u are of the form $\{x + u, -x + u\}$, which are exactly the orbits of H_u . Therefore, we have the following equivalences:

$$\begin{aligned} f \text{ hides } p_u &\iff \pi_f = \{\{x + u, -x + u\} : x \in \mathbb{F}_q\} \\ &\iff \pi_f^* = H_u \\ &\iff f \text{ hides } H_u \text{ by symmetries.} \end{aligned}$$

The reduction from the second problem to the third one is provided by Theorem 4.4. Note that we can construct a base deterministically by choosing two different elements of order two. For a reduction in the reverse direction, consider a function f on $\text{Aff}_q(\{\pm 1\})$ which hides the subgroup $H_u = \{(0, 1), (2u, -1)\}$. Then all the collisions taken by f on elements of $\text{Aff}_q(\{\pm 1\})$ are $f(2u - b, -1) = f(b, 1)$ for $b \in \mathbb{F}_q$. We define a new function f° on \mathbb{F}_q as $f^\circ(b) = \min(f(b, 1), f(b, -1))$. Examining the possible collisions gives that for $b \neq b' \in \mathbb{F}_q$, we have $f^\circ(b) = f^\circ(b')$ if and only if $b' = 2u - b = (2u, -1) \circ b$. \square

Together with Corollary 4.5 (c) the statements of this proposition imply the following result.

THEOREM 4.7. *HQPP(\mathbb{F}_q) is solvable in quantum polynomial time over constant characteristic fields.*

We observe that in contrast to the constant characteristic case, the HQPP appears to be difficult over prime fields \mathbb{F}_p , as it is equivalent to the HSP in the dihedral group $D_p \cong \text{Aff}_p(\{\pm 1\})$.

Note that in [17] van Dam, Hallgren, and Ip gave a polynomial time solution to a problem which can be considered as a version of HSSP($\text{Aff}_q(H), \mathbb{F}_q, \circ, \mathcal{FC}$), where the function hiding the complement is promised to be a shifted multiplicative character $\chi : \mathbb{F}_q^* \rightarrow \mathbb{C}^*$. This strong promise (in our oracle model we can only check for equality of the output values) makes the problem efficiently solvable even in the case $H = \{\pm 1\}$, where the HSSP with general hiding function appears to be difficult.

We also remark that strong bases in the Frobenius group $\text{Aff}_q(H)$ with $|H| = (q - 1)/2$ play an important role (under the name *factoring sets*) in certain algorithms for factoring univariate polynomials over \mathbb{F}_q ; see [25]. This is because a set B which separates two (unknown) field elements u and v can be used to find a proper decomposition of a polynomial having both u and v as roots. In fact, an efficient deterministic construction of strong bases for such affine groups over prime fields would imply an efficient deterministic algorithm for factoring polynomials over finite fields.

4.3. Multivariate quadratic hidden polynomials. In this part, we reduce the HPP for multivariate polynomials of degree at most two to the univariate HQPP. As already noted, adding a constant term does not change the level sets; therefore we consider polynomials with zero constant term. Thus, we assume that the hidden polynomial is of the form

$$(4.1) \quad \mathcal{P}(x_1, \dots, x_n) = \sum_{1 \leq i < j \leq n} a_{ij} x_i x_j + \sum_{1 \leq k \leq n} b_k x_k.$$

Also, as the partition $\pi_{\mathcal{P}}$ remains the same when we multiply all coefficients with the same nonzero element from \mathbb{F}_q , we consider that the HPP has been solved if we determine the ratios between all the pairs of the $n(n+1)/2$ coefficients a_{ij} and b_k .

PROPOSITION 4.8. *The problem HPP($\mathbb{F}_q, n, 2$) can be reduced on a quantum computer to $O(n^2)$ instances of HQPP(\mathbb{F}_q) in time $(n + \log q)^{O(1)}$.*

Proof. In order to simplify the following discussions we define a_{ji} to be a_{ij} for $j > i$. Additionally, if $q = 2$, then we also assume $a_{ii} = 0$ because $x^2 = x$ holds over \mathbb{F}_2 . We assume that we have a procedure \mathcal{R} for determining the coefficients of a univariate quadratic polynomial up to a common factor. Its oracle input is a function on \mathbb{F}_q that has the same level set structure as a polynomial of the form $ax^2 + bx$. We assume that \mathcal{R} decides whether a is zero, and if $a \neq 0$, then \mathcal{R} returns the quotient b/a .

We start with the case $n = 2$. We have an oracle with the same level sets as the polynomial

$$\mathcal{P}(x_1, x_2) = a_{11}x_1^2 + a_{22}x_2^2 + a_{12}x_1x_2 + b_1x_1 + b_2x_2.$$

We use the oracle with the inputs $(x_1, x_2) := (x, 0)$. This way, we obtain an instance of HQPP for the univariate polynomial $a_{11}x^2 + b_1x$. We use \mathcal{R} to decide whether a_{11} is zero or not, and if $a_{11} \neq 0$, then we compute the quotient b_1/a_{11} . Furthermore, we set $(x_1, x_2) := (x, 1)$ for the inputs of the oracle to compute $(a_{12} + b_1)/a_{11}$ in the second step. From this result we can easily compute the quotient a_{12}/a_{11} . Similarly, using the substitutions $(x_1, x_2) := (0, x)$ and $(x_1, x_2) := (1, x)$, we decide whether a_{22} is zero or not. If $a_{22} \neq 0$, then we obtain the quotients a_{12}/a_{22} and b_2/a_{22} . We now consider the following different cases.

- $a_{11}, a_{22} \neq 0$: If $a_{12} \neq 0$, then we have determined all coefficients of \mathcal{P} up to a common factor. If $a_{12} = 0$, then we use the inputs $(x_1, x_2) := (x, x)$ and obtain HQPP for $(a_{11} + a_{22})x^2 + (b_1 + b_2)x$. With \mathcal{R} we can determine whether $a_{11} + a_{22}$ is zero or not. If it is nonzero, then we find an element $r \in \mathbb{F}_q$ such that $b_1 + b_2 = r(a_{11} + a_{22})$. When we write $b_i/a_{ii} = c_i$, then the equation $(r - c_1)a_{11} = (c_2 - r)a_{22}$ follows. Since $a_{ii} \neq 0$, we can compute easily all coefficients of \mathcal{P} up to a common factor. If $a_{11} + a_{22} = 0$, then we also can compute all coefficients easily.
- $a_{11} \neq 0, a_{22} = 0$: If $a_{12} = 0$, then we use the inputs $(x_1, x_2) := (x, x)$ and obtain HQPP for the polynomial $a_{11}x^2 + (b_1 + b_2)x$. With \mathcal{R} we can determine the quotient $(b_1 + b_2)/a_{11}$ and together with the already known value b_1/a_{11} we obtain the missing b_2/a_{11} . If $a_{12} \neq 0$, then we pick $\alpha \in \mathbb{F}_q \setminus \{0\}$ such that $1 + \alpha a_{12}/a_{11} \neq 0$ and we use the inputs $(x_1, x_2) := (x, \alpha x)$. We obtain HQPP for $(a_{11} + \alpha a_{12})x^2 + (b_1 + \alpha b_2)x$, which can be used to find $r \in \mathbb{F}_q$ such that $(b_1 + \alpha b_2) = r(a_{11} + \alpha a_{12})$. This gives us the missing fraction b_2/a_{11} . The case $a_{22} \neq 0$ and $a_{11} = 0$ can be treated in a similar way.
- $a_{11} = a_{22} = 0, q \neq 2$: We use the inputs $(x_1, x_2) := (x, x)$ and obtain HQPP for the polynomial $a_{12}x^2 + (b_1 + b_2)x$ that can be used to decide whether $a_{12} = 0$ or not. If it is nonzero, then we compute $(b_1 + b_2)/a_{12}$. Furthermore, we can choose $\alpha \in \mathbb{F}_q^\times, \alpha \neq 1$, and we use the inputs $(x_1, x_2) := (x, \alpha x)$ to compute the fraction $(b_1 + \alpha b_2)/(\alpha a_{12})$. From these two fractions we can determine b_1/a_{12} and b_2/a_{12} . If $a_{12} = 0$, then we have the polynomial $b_1x_1 + b_2x_2$ and can determine the ratio between b_1 and b_2 by the algorithm for the abelian HSP over the additive group of \mathbb{F}_q^2 . Note that we use a quantum computer for an efficient implementation of this step of the reduction.

- $a_{11}, a_{22} = 0, q = 2$: We use the inputs $(x_1, x_2) := (x, 0)$ and obtain HPP for the polynomial b_1x . We can easily test whether it is constant, i.e., $b_1 = 0$, or not. The coefficient b_2 can be computed in a similar way. The input $(x_1, x_2) := (x, 1)$ give us $a_{12} + b_1$.

This shows that we can find all coefficients of a bivariate polynomial up to a common factor when we use \mathcal{R} only a constant number of times and some additional operations, which can be performed efficiently on a quantum computer.

Next we consider the case $n = 3$. Substituting zero in x_3 , we can use the algorithm for the bivariate case to test whether $a_{11} = 0$. If $a_{11} \neq 0$, we can determine the quotient of the remaining coefficients (except for a_{23}) and a_{11} by substituting zero in x_2 or x_3 and using the algorithm for the bivariate case. For a_{23} we can substitute $(x_1, x_2, x_3) = (x, y, y)$ and obtain the polynomial

$$a_{11}x^2 + (a_{12} + a_{13})xy + (a_{22} + a_{23} + a_{33})y^2 + b_1x + (b_2 + b_3)y.$$

Then the algorithm for the bivariate case gives us $(a_{22} + a_{23} + a_{33})/a_{11}$ from which we can compute a_{23}/a_{11} . The cases where any of the coefficients b_1, a_{22}, b_2, a_{33} , or b_3 is nonzero can be treated in a similar way. It remains to handle the case of a polynomial of the form $a_{12}x_1x_2 + a_{13}x_1x_3 + a_{23}x_2x_3$. Then substituting 1 in x_3 gives the polynomial $a_{12}x_1x_2 + a_{13}x_1 + a_{23}x_2$, and the ratio between the three coefficients can be found by the bivariate algorithm.

The case $n = 4$ can be handled as follows. We apply the algorithm of the preceding paragraph to the four polynomials obtained by substituting zero in x_1, x_2, x_3 , and x_4 , respectively. Observe that these steps determine the ratio between pairs of coefficients that have indices that fit in a three-element subset of $\{1, 2, 3, 4\}$. By transitivity, we are done unless our polynomial is of the form $a_{12}x_1x_2 + a_{34}x_3x_4, a_{13}x_1x_3 + a_{24}x_2x_4$, or $a_{14}x_1x_4 + a_{23}x_2x_3$. If it is of the form $a_{12}x_1x_2 + a_{34}x_3x_4$, then we can determine the ratio between the coefficients by using the bivariate algorithm by substituting x_1 in x_2 and x_3 in x_4 . The two remaining polynomials can be treated in a similar way.

Finally we consider the case $n > 4$. Using $O(n^2)$ applications of the bivariate algorithm, we find indices $i \neq j$ such that at least one of a_{ii}, b_i , and a_{ij} is nonzero. The ratio between this coefficient and any other can be computed using the algorithm for two, three, or four variables. The cost of these steps amounts to $O(n^2)$ applications of the procedure \mathcal{R} and a polynomial number of other operations. \square

THEOREM 4.9. *HPP($\mathbb{F}_q, n, 2$) can be solved by a polynomial time quantum algorithm over fields of constant characteristic.*

5. Function graph groups and the HPGP. For dealing with the HPGP we define a family of semidirect product groups that we call function graph groups. We show that each instance of the HPGP($\mathbb{F}_q, 1, d$) can be reduced to the HSP for an appropriate function graph group corresponding to univariate polynomials of degree at most d . These special function graph groups are semidirect products of groups of q -power order. Therefore, they cannot be Frobenius groups.

5.1. The HPGP as HSSP over function graph groups. It will be convenient to work in a more general setting.

DEFINITION 5.1. *Let A and B be two abelian groups. The family of functions mapping A to B forms an abelian group \mathcal{F} with the addition defined as $(Q_1 + Q_2)(x) = Q_1(x) + Q_2(x)$. For every $t \in A$, the shift map a_t defined as $(a_tQ)(x) = Q(x - t)$ is an automorphism of this group. A function group from A to B is a subgroup K of \mathcal{F} which is closed under the shift maps. We denote the restriction of a_t to K also with a_t .*

Then the map $t \mapsto a_t$ is a homomorphism from A to the automorphism group of K . The function graph group $\text{Fg}(K)$ is defined as the semidirect product $K \rtimes_{t \mapsto a_t} A$.

The multiplication of $\text{Fg}(K)$ is given by the rule

$$(Q_1, t_1)(Q_2, t_2) = (Q_1 + a_{t_1}Q_2, t_1 + t_2).$$

The *shifting action* \circ of $\text{Fg}(K)$ on $A \times B$ is defined as

$$(Q, t) \circ (x, y) = (x + t, y + Q(x + t)).$$

For $t \in A$ and $Q \in K$, we set $a_{Q,t} = (Q - a_t Q, t)$, the conjugate of the element $(0, t)$ by $(Q, 0)$. Furthermore, let $A_Q = \{a_{Q,t} : t \in A\}$ be the conjugate of the subgroup $\{(0, t) : t \in A\}$ by $(Q, 0)$. Then every A_Q is a subgroup of $\text{Fg}(K)$ that is complementary to the normal subgroup $\{(Q, 0) : Q \in K\}$. We call them *standard complements*, and we denote by \mathcal{SC} the family $\{A_Q : Q \in K\}$ of the standard complements.

We are now ready to show a connection between function graph problems and the orbits of the standard complements in function graph groups.

PROPOSITION 5.2. *Let $\text{Fg}(K)$ be a function graph group, let \circ be its shifting action on $A \times B$, and let A_Q be a standard complement. Then A_Q is closed and the orbits of A_Q are the level sets of the function $f : (x, y) \mapsto y - Q(x)$ on $A \times B$.*

Proof. Assume that A_Q is not closed. Then, as A_Q is a complement of $\{(Q', 0) : Q' \in K\}$, there exists $Q' \in K \setminus \{0\}$ such that $(x, y + Q'(x)) = (Q', 0) \circ (x, y) \in A_Q \circ (x, y)$ for every pair $(x, y) \in A \times B$. This is a contradiction since $a_{Q,t}(x, y) = (x, y')$ is only possible if $t = 0$ and $y' = y$.

To see the second part of the statement, observe that $f(x, y) = f(x', y')$ if and only if $\exists t \in A : (x', y') = (x + t, y - Q(x) + Q(x + t))$ if and only if $\exists t \in A : (x', y') = a_{Q,t} \circ (x, y)$ if and only if $(x', y') \in A_Q \circ (x, y)$. \square

We now specialize function graph groups to polynomials which relate them to the HPGP. Let A and B be the additive group of \mathbb{F}_q , and let K be $\mathbb{F}_q^{(d)}[x]$, the set of polynomials of degree at most d . Observe that we include also polynomials with nonzero constant terms in order to be closed under the shifts. Then Proposition 5.2 translates to the following statement.

PROPOSITION 5.3. *Let $f : \mathbb{F}_q \times \mathbb{F}_q \rightarrow S$ be a function. Then f hides for HPGP($\mathbb{F}_q, 1, d$) the polynomial Q if and only if for $\text{HSSP}(\text{Fg}(\mathbb{F}_q^{(d)}[x]), \mathbb{F}_q \times \mathbb{F}_q, \circ, \mathcal{SC})$ it hides the standard complement A_Q by symmetries.*

5.2. Small bases for standard complements. In this section, we construct strong bases for the standard complements in function graph groups. The next lemma gives a simple characterization of such bases.

LEMMA 5.4. *Let $\text{Fg}(K) = K \rtimes A$ be a function graph group with action \circ on $A \times B$. Let $D = \{(x_1, y_1), \dots, (x_\ell, y_\ell)\}$ be a subset of $A \times B$. Then D is an \mathcal{SC} -strong base if and only if for all $Q \in K$ the equation $Q(x_1) = \dots = Q(x_\ell) = 0$ implies $Q = 0$.*

Proof. As \mathcal{SC} is closed under conjugation, by the remarks following Definition 3.2, D is an \mathcal{SC} -strong base if and only if $\bigcap_{i=1}^{\ell} A_Q \text{Fg}(K)_{(x_i, y_i)} = A_Q$ for every $Q \in K$. The statement $(Q', t') \in A_Q \text{Fg}(K)_{(x_i, y_i)}$ is true if and only if there is a $t_i \in A$ such that $(Q', t') \circ (x_i, y_i) = a_{Q,t_i} \circ (x_i, y_i)$. This can be rewritten as $(x_i + t', y_i + Q'(x_i + t')) = (x_i + t_i, y_i - Q(x_i) + Q(x_i + t_i))$. The equality holds if and only if $t_i = t'$ and $(a_{-t'}Q' - a_{-t'}Q + Q)(x_i) = 0$. Hence, an element (Q', t') is in the intersection $\bigcap_{i=1}^{\ell} A_Q \text{Fg}(K)_{(x_i, y_i)}$ if and only if $t_i = t'$ and $(a_{-t'}Q' - a_{-t'}Q + Q)(x_i) = 0$ holds for all i .

We first prove the “only if” part of the lemma. To this end, let D be an A_Q -strong base, and let $R \in \mathcal{C}$ be a function such that $R(x_i) = 0$ for all i . Let

$$(Q', t') \in \bigcap_{i=1}^{\ell} A_Q \text{Fg}(K)_{(x_i, y_i)}$$

be any element. Since D is a base, we know that the intersection is equal to A_Q , and from this $(Q', t') = (Q - a_{t'}Q, t')$ follows. We also know that $(a_{-t'}Q' - a_{-t'}Q + Q)(x_i) = 0$ for all i , and the same is true for $a_{-t'}Q' + R - a_{-t'}Q + Q$. Hence, we also have $(Q' + a_{t'}R, t')$ in the intersection and $(Q' + a_{t'}R, t') = (Q - a_{t'}Q, t')$ follows. We have $(Q', t') = (Q' + a_{t'}R, t')$, and this directly implies $a_{t'}R = R = 0$.

To see the “if” part of the lemma, observe that the second statement of the lemma, applied to the function $a_{-t'}Q' - a_{-t'}Q + Q$, implies that $a_{-t'}Q' - a_{-t'}Q + Q = 0$. We apply the shift map $a_{t'}$ to this equality and obtain $Q' = Q - a_{t'}Q$. Hence, we have $(Q', t') = a_{Q, t'} \in A_Q$, and this shows that D is an \mathcal{SC} -strong base. \square

Combining the statements of this section, we obtain the following result.

THEOREM 5.5. *HPGP($\mathbb{F}_q, 1, d$) can be reduced to HSP($\text{Fg}(\mathbb{F}_q^{(d)}[x]), \mathcal{SC}$) in polynomial time in d and $\log q$.*

Proof. Univariate polynomials of degree d have at most d roots over a field. Therefore, by Proposition 5.3 and Lemmas 5.4 and 3.3, we can associate in polynomial time an instance of HPGP($\mathbb{F}_q, 1, d$) that hides a polynomial Q with symmetries to an instance of HSP($\text{Fg}(\mathbb{F}_q^{(d)}[x]), \mathcal{SC}$) that hides the subgroup A_Q . Then the polynomial Q (up to a constant term) can be recovered from generators for A_Q as follows. The elements $(Q - a_{t_1}Q, t_1), \dots, (Q - a_{t_\ell}Q, t_\ell)$ generate A_Q if and only if t_1, \dots, t_ℓ generate the additive group of \mathbb{F}_q . It follows that for arbitrary $s \in \mathbb{F}_q$, we can efficiently compute $Q - a_sQ$ using the group operation in $A_Q \leq \text{Fg}(\mathbb{F}_q^{(d)}[x])$. Substituting s into $Q - a_sQ$ gives $Q(s) - Q(0)$. We do this for d different values $s \in \mathbb{F}$ and compute $Q - Q(0)$ using Lagrange interpolation. \square

We remark that the group $\text{Fg}(\mathbb{F}_q^{(d)}[x])$ is of nilpotency class $d + 1$. However, we can actually give a reduction to the HSP in a group of class d . To this end, observe that the hidden subgroup is a conjugate of the complement \mathbb{F}_q . Therefore, it can be found in the subgroup generated by the commutator of \mathbb{F}_q with $\mathbb{F}_q^{(d)}[x]$ (this is an abelian normal subgroup) and the complement \mathbb{F}_q . This semidirect product group has nilpotency class d .

Note that semidirect product groups $\text{Fg}(\mathbb{F}_q^{(d)}[x])$ have commutative commutator subgroups and that their exponent is the characteristic of \mathbb{F}_q . Therefore, for fixed characteristic, we can apply Fact 2 (b) to obtain the following result.

COROLLARY 5.6. *Assume that q is a power of a fixed prime p . Then we can solve HPGP($\mathbb{F}_q, 1, d$) by a quantum algorithm in time polynomial in d and $\log q$.*

This corollary allows us to complete Fact 4 (b), because it can be applied to fields of characteristic in the set E_d that were left open. Since E_d is finite, it follows that for fixed d we can solve HPGP($\mathbb{F}_q, 1, d$) in quantum polynomial time for all finite fields. Together with Fact 4 (a) this improves the overall result of [15]: the HPGP(\mathbb{F}_q, n, d) can be solved efficiently for all finite fields when n and d are constant. We further improve this result in the next section, where we present a more powerful reduction of the multivariate problem to the univariate case.

We conclude this section by showing that the HSP for semidirect product groups of the form $\mathbb{Z}_p^m \rtimes \mathbb{Z}_p$ can be reduced to a multidimensional analogue of the HPGP. This HSP is discussed in [5], where it is shown that the HSP for all possible subgroups

can be reduced to the HSP in which the hidden subgroups are complements of \mathbb{Z}_p^m . Let H be such a subgroup. Following arguments of [5], we show that the cosets of H can be considered as level sets of a polynomial map from \mathbb{Z}_p^{m+1} to \mathbb{Z}_p^m of the form $y - Q(x)$, where $y = (y_1, \dots, y_m)$ and $Q(x) = (Q_1(x), \dots, Q_m(x))$, each $Q_i(x)$ being a univariate polynomial over \mathbb{Z}_p of degree at most d . Here $d \leq \min(m, p)$ depends on the structure of G (actually its nilpotency class).

To this end, notice that the semidirect product structure is given by a linear transformation A on \mathbb{Z}_p^m . (This is the action of the generator 1 of \mathbb{Z}_p on \mathbb{Z}_p^m .) We have $A^p = I$, whence $B = A - I$ satisfies $B^p = (A - I)^p = A^p - I^p = 0$. Therefore, there exists a smallest positive integer $d \leq \min(m, p)$ such that $B^d = 0$.

A subgroup H_v complementary to \mathbb{Z}_p^m in $\mathbb{Z}_p^m \rtimes \mathbb{Z}_p$ consists of the powers of an element of the form $(v, 1)$ for some $v \in \mathbb{Z}_p^m$. With the map

$$(5.1) \quad Q_v : \begin{cases} \mathbb{Z}_p \rightarrow \mathbb{Z}_p^m, \\ t \mapsto \sum_{j=0}^{t-1} A^j v \end{cases}$$

these powers are the pairs $(Q_v(t), t)$ for $t \in \mathbb{Z}_p$, and the right cosets of H_v are the sets of the pairs $(Q_v(t) + y, t)$ for $t \in \mathbb{Z}_p$, where $y \in \mathbb{Z}_p^m$. It turns out that the entries of the matrix of $\sum_{j=0}^{t-1} A^j$, as functions in t , are polynomials of degree at most d with zero constant term (see [5]). Therefore, the same holds for the coordinates $Q_v^{(i)}$ of the vector $Q_v(t)$. In other words, the map $t \mapsto Q_v(t)$ is a polynomial map from \mathbb{Z}_p to \mathbb{Z}_p^m of degree d with zero constant term. Hence, the cosets of H_v are exactly the level sets for the polynomial map

$$(5.2) \quad (y_1, \dots, y_m, x) \mapsto (y_1 - Q_v^{(1)}(x), \dots, y_m - Q_v^{(m)}(x))$$

from \mathbb{Z}_p^{m+1} to \mathbb{Z}_p^m .

It follows that any function on $\mathbb{Z}_p^m \rtimes \mathbb{Z}_p$ that hides the subgroup $H_v = \langle (v, 1) \rangle$ directly defines an instance of the m -dimensional analogue of the HPGP for Q_v as defined in (5.1). If we solve the m -dimensional HPGP for these instances, i.e., if we determine Q_v , then we obtain v by calculating $v = Q_v(1)$.

This shows that the HSP of $\mathbb{Z}_p^m \rtimes \mathbb{Z}_p$ can be indeed efficiently reduced to the m -dimensional analogue of the HPGP. Plugging $A = \mathbb{Z}_p$, $B = \mathbb{Z}_p^m$, and $K = (\mathbb{Z}_p^{(d)}[x])^m$ into Proposition 5.2, we obtain that this problem can be viewed as an instance of the HSSP over a semidirect product K with \mathbb{Z}_p . Here the functions are vectors of univariate polynomials. Therefore, by Lemma 5.4, small bases exist and can be found easily and the reduction to a HSP works. Note, however, that the new group is in general much bigger than the original one.

These reductions explain why it was possible to construct the algorithm of [15] in close analogy with the pretty good measurement framework of [5] for semidirect product groups.

5.3. Reduction of multivariate HPGP to univariate case. The scheme of [15] for reducing the multivariate HPGP to the univariate case can be improved with the help of a generalized Vandermonde matrix.

THEOREM 5.7. *An instance of HPGP(\mathbb{F}_q, n, d) can be reduced to $O\binom{d+n}{n}$ instances of HPGP($\mathbb{F}_q, 1, d$) by a classical algorithm with running time polynomial in $\binom{d+n}{n}$. If d is constant, then HPGP(\mathbb{F}_q, n, d) can be solved by a polynomial time quantum algorithm.*

We prove the theorem in the remainder of the subsection. For this, we consider n -variate polynomials of the special form

$$(5.3) \quad y - Q(x_1, \dots, x_n) \quad \text{with} \quad Q \in \mathbb{F}_q[x_1, \dots, x_n].$$

Note that we changed the notation by replacing the polynomials $Q(x_1)$ in Definition 2.1 by polynomials $Q(x_1, \dots, x_n) \in \mathbb{F}_q[x_1, \dots, x_n]$. This makes the following discussion easier. Recall that the constant term of the polynomials $Q(x_1, \dots, x_n)$ is assumed to be zero since it cannot be determined. Furthermore, the identity $x^q = x$ in \mathbb{F}_q implies that we can only distinguish polynomials that are reduced modulo $x_i^q - x_i$ for all variables x_i . Hence, for a maximum total degree d we consider only local degrees of at most $\min\{d, q-1\}$; i.e., the power of each x_i in all monomials occurring in $Q(x_1, \dots, x_n)$ is less or equal to this minimum.

For each j with $1 \leq j \leq n$ let

$$\mathcal{I}^{(j)} := \left\{ \alpha \in \mathbb{N}^j : \sum_{i=1}^j \alpha_i \leq d, \alpha_i \leq \min\{d, q-1\} \text{ for } i = 1, \dots, j \right\} \setminus \{(0, \dots, 0)\}$$

be the set of all exponent vectors for the monomials of total degree at most d when the variables are restricted to x_1, \dots, x_j . For each $\alpha \in \mathcal{I}^{(j)}$ let

$$m_\alpha := x_1^{\alpha_1} \cdot \dots \cdot x_j^{\alpha_j}$$

denote the corresponding monomial. For j and j' with $1 \leq j < j' \leq k$, a monomial m_α with $\alpha = (\alpha_1, \dots, \alpha_j) \in \mathcal{I}^{(j)}$ is also defined by $\tilde{\alpha} = (\alpha_1, \dots, \alpha_j, 0, \dots, 0) \in \mathcal{I}^{(j')}$. Finally, for $v = (v_1, \dots, v_j) \in \mathbb{F}_q^j$ let

$$m_\alpha(v) := v_1^{\alpha_1} \cdot \dots \cdot v_j^{\alpha_j}$$

denote the evaluation of the monomial m_α at the point v . For $q > \ell$, the number of such monomials is given by the simple expression

$$|\mathcal{I}^{(j)}| = \binom{d+j}{j} - 1.$$

For $q \leq \ell$, the number of such monomials is determined with the inclusion-exclusion principle, which leads to the expression

$$|\mathcal{I}^{(j)}| = \sum_{i=0}^j (-1)^i \binom{j}{i} \binom{d-iq+j}{j} - 1.$$

We use the convention that the binomial coefficient is zero if the number at the top is negative. With the help of $\mathcal{I}^{(j)}$ we can define the generalized Vandermonde matrix and describe an efficient construction.

LEMMA 5.8. *Let d be the maximum total degree of the monomials in $\mathcal{I}^{(j)}$ over the field \mathbb{F}_q . Then there is a classical algorithm for constructing a set $\mathcal{V}^{(j)} \subset \mathbb{F}_q^j$ of cardinality $|\mathcal{I}^{(j)}|$ such that the square matrix*

$$(5.4) \quad M^{(j)} := \left[m_\alpha(v) \right]_{v \in \mathcal{V}^{(j)}, \alpha \in \mathcal{I}^{(j)}}$$

has full rank. This matrix is called the generalized Vandermonde matrix. The running time of the algorithm is polynomial in $|\mathcal{I}^{(j)}|$.

Proof. This statement is proved in [26]. For the sake of completeness we present here another proof which is also much simpler than the original one. The condition that $M^{(j)}$ has full rank is equivalent to the following condition: for every (nonzero) polynomial

$$(5.5) \quad F(x_1, \dots, x_j) = \sum_{\alpha \in \mathcal{I}^{(j)}} c_\alpha m_\alpha$$

there is at least one $v \in \mathcal{V}^{(j)}$ such that $F(v) \neq 0$.

Let $b := \min\{d, q-1\}$ denote the upper bound on the local degrees. For $j = 1$, we have $\mathcal{I}^{(1)} = \{(1), \dots, (b)\}$ and the corresponding set of monomials is $\{x_1, x_1^2, \dots, x_1^b\}$. We can choose $\mathcal{V}^{(1)} := \{v_1, v_2, \dots, v_b\}$ to be a set containing b different nonzero elements of \mathbb{F}_q . Then the matrix

$$M^{(1)} = \begin{pmatrix} v_1^1 & v_1^2 & \cdots & v_1^b \\ v_2^1 & v_2^2 & \cdots & v_2^b \\ \vdots & \vdots & \ddots & \vdots \\ v_b^1 & v_b^2 & \cdots & v_b^b \end{pmatrix}$$

has full rank $|\mathcal{I}^{(1)}| = b$. Observe that we obtain a (square) Vandermonde matrix by multiplying $M^{(1)}$ with $\text{diag}(v_1^{-1}, v_2^{-1}, \dots, v_b^{-1})$ from the left. We choose v_1 to be equal to 1.

Assume that we have already determined a suitable $\mathcal{V}^{(j-1)}$ for some $j \geq 2$. We show how to obtain $\mathcal{V}^{(j)}$ using $\mathcal{V}^{(j-1)}$.

1. Set $\mathcal{V}^{(j)} \leftarrow \{(1, \dots, 1)\} \subseteq \mathbb{F}_q^j$
2. Set

$$(5.6) \quad L^{(j)} \leftarrow \left[m_\alpha(v) \right]_{v \in \mathcal{V}^{(j)}, \alpha \in \mathcal{I}^{(j)}}$$

3. REPEAT

4. Determine a (nontrivial) vector $c = (c_\alpha) \in \mathbb{F}_q^{|\mathcal{I}^{(j)}|}$ in the kernel of $L^{(j)}$
5. Set

$$G(x_1, \dots, x_j) \leftarrow \sum_{\alpha \in \mathcal{I}^{(j)}} c_\alpha m_\alpha$$

6. Determine a vector $u \in \mathbb{F}_q^j$ such that $G(u) \neq 0$
7. Set $\mathcal{V}^{(j)} \leftarrow \mathcal{V}^{(j)} \cup \{u\}$
8. Add the row vector $(m_\alpha(u))_{\alpha \in \mathcal{I}^{(j)}}$ at the bottom of $L^{(j)}$
9. UNTIL the rank of $L^{(j)}$ is maximal

We now explain how the different steps can be implemented efficiently and why the algorithm produces a valid $\mathcal{V}^{(j)}$.

We can compute a nontrivial vector c in the kernel of $L^{(j)}$ in step 4 with Gaussian elimination. To find a u with $G(u) \neq 0$ in step 5, we write G in the form

$$G(x_1, \dots, x_j) = \sum_{i=1}^b F_i(x_1, \dots, x_{j-1}) \cdot x_j^i \in \mathbb{F}_q[x_1, \dots, x_{j-1}][x_j].$$

At least one of the polynomials F_i is nonzero because G is nonzero. Set F to be the nonzero F_i with the smallest i . We can write F as

$$F(x_1, \dots, x_{j-1}) = \sum_{\beta \in \mathcal{I}^{(j-1)}} d_\beta m_\beta \in \mathbb{F}_q[x_1, \dots, x_{j-1}]$$

with appropriate coefficients $d_\beta \in \mathbb{F}_q$. There exists a vector $v = (v_1, \dots, v_{j-1}) \in \mathcal{V}^{(j-1)}$ with $F(v) \neq 0$. This is because otherwise we would have a nontrivial linear dependency of the rows of $L^{(j-1)}$ corresponding to the elements in $\mathcal{V}^{(j-1)}$. Hence, the polynomial $P(x) := G(v_1, \dots, v_{j-1}, x)$ is a nonzero univariate polynomial that can be written as a linear combination of monomials m_γ with $\gamma \in \mathcal{I}^{(1)}$. The element $w \in \mathbb{F}_q$ with $P(w) \neq 0$ can be found among the elements of $\mathcal{V}^{(1)}$. We obtain the desired vector u by setting it equal to (v_1, \dots, v_{j-1}, w) .

By adding the new row vector to $L^{(j)}$ in step 8, we achieve that the vector c is no longer in the kernel of the new augmented matrix. Hence, we reduced the dimension of the kernel of the linear map defined by this matrix. In other words, we have increased its rank by exactly 1. This shows that the algorithm terminates. \square

We are now ready to describe the improved reduction.

LEMMA 5.9. *Let $\mathcal{V}^{(n)}$ be as in Lemma 5.8. Then the coefficients of the hidden polynomial of (5.3) can be determined by solving the univariate HPGP for the polynomials $Q(v_1x, v_2x, \dots, v_nx)$ for all $v \in \mathcal{V}^{(n)}$.*

Proof. The unknown polynomials can be expressed as

$$Q(x_1, \dots, x_n) = \sum_{\ell=1}^d Q_\ell(x_1, \dots, x_n),$$

where Q_ℓ denotes the homogeneous part of total degree ℓ .

For each $v = (v_1, \dots, v_n) \in \mathbb{F}_q^n$, the substitution $x_i \mapsto v_ix$ in the hidden multivariate polynomial Q leads to the univariate polynomial

$$P_v(x) := Q(v_1x, \dots, v_nx) = \sum_{\ell=1}^d Q_\ell(v)x^\ell.$$

We determine the coefficients $Q_\ell(v)$ of $P_v(x)$ by using the quantum algorithm for the univariate case. Let $z = [q_\alpha]_{\alpha \in \mathcal{I}^{(n)}}^T$ be the column vector whose entries are the unknown coefficients we seek to learn. Let $y = [Q_1(v) + \dots + Q_d(v)]_{v \in \mathcal{V}^{(n)}}$ be the column vector whose entries are the sum of evaluations of the homogeneous part Q_ℓ at the points $v \in \mathcal{V}^{(n)}$. We have $M^{(n)}z = y$. Hence, we can recover y since the generalized Vandermonde matrix $M^{(n)}$ has full rank. \square

Theorem 5.7 follows directly from Lemma 5.9. Note that in the course of the above reduction, we learn $d|\mathcal{I}^{(n)}|\log_2(q)$ bits by solving $|\mathcal{I}^{(n)}|$ instances of the univariate case (each instance yielding exactly d coefficients in \mathbb{F}_q). The absolute lower bound is given by $|\mathcal{I}^{(n)}|\log_2(q)$, which corresponds to the number of bits necessary to specify all coefficients of the hidden polynomial Q . This discussion shows that our method is optimal up to the factor d .

Theorem 5.7 also gives an instance of the HSSP which is solvable in quantum polynomial time, although no small strong bases exist (and therefore the reduction scheme of section 3 does not work directly). Let A and B be the additive group of \mathbb{F}_q^n and \mathbb{F}_q , respectively, and let K be $\mathbb{F}_q^{(d)}[x_1, \dots, x_n]$, the set of polynomials in n variables of total degree at most d . Then Proposition 5.2 translates to the following statement.

PROPOSITION 5.10. *Let $f : \mathbb{F}_q^n \times \mathbb{F}_q \rightarrow S$ be a function. Then f hides for $\text{HPGP}(\mathbb{F}_q, n, d)$ the polynomial Q if and only if for $\text{HSSP}(\text{Fg}(\mathbb{F}_q^{(d)}[x_1, \dots, x_n]), \mathbb{F}_q^n \times \mathbb{F}_q, \circ, \mathcal{SC})$ it hides the standard complement A_Q by symmetries.*

Hence, by Theorem 5.7, for constant d , $\text{HSSP}(\text{Fg}(\mathbb{F}_q^{(d)}[x_1, \dots, x_n]), \mathbb{F}_q^n \times \mathbb{F}_q, \circ, \mathcal{SC})$ can be solved in quantum polynomial time. On the other hand, as multivariate

polynomials have many zeros, there are no bases of polynomial size for the action \circ for $n \geq 2$; see Lemma 5.4.

Acknowledgment. The authors are grateful to the anonymous referees for their helpful remarks and suggestions on a previous version of the paper.

REFERENCES

- [1] P. W. SHOR, *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*, SIAM J. Comput., 26 (1997), pp. 1484–1509.
- [2] D. BONEH AND R. LIPTON, *Quantum cryptanalysis of hidden linear functions*, in Proceedings of Crypto'95, Lecture Notes in Comput. Sci. 963, Springer, Berlin, 1995, pp. 424–437.
- [3] A. KITAEV, *Quantum Measurements and the Abelian Stabilizer Problem*, preprint, arXiv:quant-ph/9511026v1, 1995.
- [4] O. REGEV, *Quantum computation and lattice problems*, SIAM J. Comput., 33 (2004), pp. 738–760.
- [5] D. BACON, A. CHILDS, AND W. VAN DAM, *From optimal measurement to efficient quantum algorithms for the hidden subgroup problem over semidirect product groups*, in Proceedings of the 46th IEEE Symposium on Foundations of Computer Science (FOCS), 2005, pp. 469–478.
- [6] A. DENNEY, C. MOORE, AND A. RUSSELL, *Finding conjugate stabilizer subgroups in $PSL(2; q)$ and related groups*, Quantum Inf. Comput., 10 (2010), pp. 282–291.
- [7] K. FRIEDL, G. IVANYOS, F. MAGNIEZ, M. SANTHA, AND P. SEN, *Hidden translation and orbit coset in quantum computing*, in Proceedings of the 35th ACM Symposium on Theory of Computing (STOC), 2003, pp. 1–9.
- [8] M. GRIGNI, L. SCHULMAN, M. VAZIRANI, AND U. VAZIRANI, *Quantum mechanical algorithms for the nonabelian hidden subgroup problem*, in Proceedings of the 33rd ACM Symposium on Theory of Computing (STOC), 2001, pp. 68–74.
- [9] S. HALLGREN, A. RUSSELL, AND A. TA-SHMA, *The hidden subgroup problem and quantum computation using group representations*, SIAM J. Comput., 32 (2003), pp. 916–934.
- [10] G. KUPERBERG, *A subexponential-time quantum algorithm for the dihedral hidden subgroup problem*, SIAM J. Comput., 35 (2005), pp. 170–188.
- [11] C. MOORE, D. ROCKMORE, A. RUSSELL, AND L. J. SCHULMAN, *The power of basis selection in Fourier sampling: Hidden subgroup problems in affine groups*, in Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2004, pp. 1106–1115.
- [12] C. MOORE, A. RUSSELL, AND L. SCHULMAN, *The symmetric group defies strong Fourier sampling*, in Proceedings of the 46th IEEE Symposium on Foundations of Computer Science (FOCS), 2005, pp. 479–488.
- [13] S. HALLGREN, C. MOORE, M. RÖTTELER, A. RUSSELL, AND P. SEN, *Limitations of quantum coset states for graph isomorphism*, in Proceedings of the 38th ACM Symposium on Theory of Computing (STOC), 2006, pp. 604–617.
- [14] A. CHILDS, L. SCHULMAN, AND U. VAZIRANI, *Quantum algorithms for hidden nonlinear structures*, in Proceedings of the 48th IEEE Symposium on Foundations of Computer Science (FOCS), 2007, pp. 395–404.
- [15] T. DECKER, J. DRAISMA, AND P. WOCJAN, *Quantum algorithm for identifying hidden polynomial function graphs*, Quantum Inf. Comput., 9 (2009), pp. 215–230.
- [16] M. ETTINGER, P. HØYER, AND E. KNILL, *The quantum query complexity of the hidden subgroup problem is polynomial*, Inform. Process. Lett., 91 (2004), pp. 43–48.
- [17] W. VAN DAM, S. HALLGREN, AND L. IP, *Quantum algorithms for some hidden shift problems*, in Proceedings of the 14th Annual ACM-SIAM Symposium On Discrete Algorithms (SODA), 2003, pp. 489–498.
- [18] C. MOORE, D. ROCKMORE, A. RUSSELL, AND L. J. SCHULMAN, *The power of strong Fourier sampling: Quantum algorithms for affine groups and hidden shifts*, SIAM J. Comput., 37 (2007), pp. 938–958.
- [19] T. BLYTH, *Lattices and Ordered Algebraic Structures*, Springer, London, 2005.
- [20] M. RÖTTELER, *Quantum algorithms to solve the hidden shift problem for quadratics and for functions of large Gowers norm*, in Proceedings of the 34th International Symposium on Mathematical Foundations of Computer Science (MFCS), Lecture Notes in Comput. Sci. 5734, Springer, Berlin, 2009, pp. 663–674.
- [21] C. H. BENNETT, E. BERNSTEIN, G. BRASSARD, AND U. VAZIRANI, *Strengths and weaknesses of quantum computing*, SIAM J. Comput., 26 (1997), pp. 1510–1523.
- [22] Á. SERESS, *Permutation Group Algorithms*, Cambridge University Press, Cambridge, UK, 2003.

- [23] B. HUPPERT, *Endliche Gruppen I*, Springer, Berlin, 1983.
- [24] L. BABAI, *Local expansion of vertex-transitive graphs and random generation in finite groups*, in Proceedings of the 23rd ACM Symposium on Theory of Computing (STOC), 1991, pp. 164–174.
- [25] P. CAMION, *Improving an algorithm for factoring polynomials over finite fields and constructing large irreducible polynomials*, IEEE Trans. Inform. Theory, 29 (1983), pp. 378–385.
- [26] Z. ZILIC AND Z. VRANESIC, *A deterministic multivariate polynomial interpolation algorithm for small finite fields*, IEEE Trans. Comput., 37 (2002), pp. 1100–1105.