Communication Complexity in Algebraic Two-Party Protocols

Rafail Ostrovsky* William E. Skeith III†

Appeared in CRYPTO 2008: 379-396

Abstract

In cryptography, there has been tremendous success in building various two-party protocols with small communication complexity out of homomorphic semantically-secure encryption schemes, using their homomorphic properties in a black-box way. A few notable examples of such primitives include items like single database Private Information Retrieval (PIR) schemes (introduced in [15]) and private database update with small communication (introduced in [5]). In this paper, we illustrate a general methodology for determining what types of protocols can and cannot be implemented with small communication by using homomorphic encryption in a black-box way.

We hope that this work will provide a simple "litmus test" of feasibility for black-box use of known homomorphic encryption schemes by other cryptographic researchers attempting to develop new protocols with low communication. Additionally, a precise mathematical language for reasoning about such problems is developed in this work, which may be of independent interest. We stress that the class of algebraic structures for which we prove communication complexity lower bounds is large, and covers practically all known semantically-secure homomorphic cryptosystems (including those based upon bilinear maps).

Finally, we show the following equivalence which relates group homomorphic encryption and a major open question of designing a so-called fully-homomorphic cryptosystem: a fully homomorphic encryption scheme (over a non-zero ring) exists if and only if there exists homomorphic encryption over any finite non-abelian simple group. This result somewhat generalizes results of Barrington [1] (to any group containing a finite non-abelian simple subgroup) and of Maurer and Rhodes [18], and in fact gives a *constructive* proof of the 1974 result Werner [28]. (This also answers an open question posed by Rappe in [23], who in 2004 proved a special case of this result.

Keywords: homomorphic encryption, fully homomorphic encryption, private information retrieval, PIR writing, keyword search, communication complexity, algebraic lower bounds.

1 Introduction

One of the central problems in cryptography is that of finding a public key encryption scheme that would allow "computation on encrypted data". In its full generality the problem could be simply stated as follows: to find a public key encryption scheme such that given encryptions of arbitrary plaintexts $\mathcal{E}(x_1), \ldots, \mathcal{E}(x_n)$ it is possible *without the decryption key* to compute

*Supported in part by IBM Faculty Award, Xerox Innovation Group Award, NSF grants 0430254, 0716835, 0716389 and U.C. MICRO grant.

†Supported in part by U.C. Chancellor's Presidential Dissertation Fellowship 2006-2007 and by NSF grant 0430254.

$\mathcal{E}(f(x_1, \ldots, x_n))$ for any polynomial-time computable function f. Naturally, if one can find a public-key cryptosystem that is "fully homomorphic", i.e. allows operations on ciphertext that preserve the structure of a ring, and hence allows computation of the ubiquitous *"NAND"* operation on the underlying plaintext, it would give a general solution to the above problem. Indeed, the reason this is such a central problem is that it would create a powerful mechanism to arbitrarily manipulate encrypted data without sacrificing privacy. This problem was posed nearly 30 years ago by Rivest, Adelman and Dertouzos [24]. We do not know if such an encryption scheme exists in its full generality, though various partial answers are known: One partial answer is abelian group-homomorphic encryption: given $\mathcal{E}(x)$ and $\mathcal{E}(y)$, where x and y come from some abelian group, there exist cryptosystems that can compute $\mathcal{E}(x * y)$, where $*$ is the group operation. Examples include ElGamal [9], where the group operation is multiplication, Goldwasser and Micali [10] where the operation is addition modulo 2, and Pallier [22] where the group operation is addition modulo a large composite. Recent progress by Boneh, Goh and Nissim [3] has shown that more is possible: they designed a cryptosystem that allows an arbitrary number of additions and a single multiplication (of the underlying plaintext) by manipulating ciphertexts only. I.e., polynomials of total degree 2 can be computed on ciphertext. Another approach at building fully-homomorphic encryption was considered by Sander, Young, and Yung [26], but only applied to Boolean operations that doubled the ciphertext size at every step. As a result, one could only perform a few Boolean operations before the ciphertext size became impractical. A partial negative result was given by Boneh and Lipton [4].

Many useful protocols and primitives have been derived from such homomorphic schemes in a "black box" way, essentially just manipulating the homomorphic properties to construct various systems. Prominent examples include single-database private information retrieval (PIR), originally introduced by Kushilevitz and Ostrovsky [15] and collision-resistant hashing as shown by Ishai, Kushilevitz, and Ostrovsky [14]. (For more details regarding this approach to PIR, see the survey of Ostrovsky and Skeith [20].) In this work, we show a variety of communication complexity lower bounds for natural tasks when constructed in a similar, but somewhat more restricted manner (to further improve communication complexity, the aforementioned protocols often use repeated encryption, destroying the algebraic value of the resulting ciphertext). More accurately we'll illustrate a single basic task that cannot be algebraically accomplished (with small communication) in various structures (e.g., that of *any* abelian group). This result will give us a simple criterion or "litmus test" for determining the feasibility of constructing communication-efficient protocols in general, and a rule out the possibility for constructing many communication-efficient protocols based on the black box use of homomorphic encryption. Along the way, we'll also develop a mathematical language and technique for reasoning about such questions, which may be of independent interest. A lot of effort has been put into designing new cryptosystems that allow the structure to be as rich as possible, but our lower bounds capture an even larger class of algebraic structures than what current homomorphic encryption schemes provide.

1.1 Our Results

A few of the main results in this work are as follows, where n represents the database size in a PIR-writing scheme:

Theorem 1 (informal) We prove $\Omega(n)$ bound for algebraic **PIR-writing** based on **any** abelian group homomorphic encryption.

Theorem 2 (informal) We prove $\Omega(\sqrt{n})$ bound for algebraic PIR-writing based on the cryptosystem of Boneh, Goh and Nissim [3]. We note that the work of Boneh, Kushilevitz, Ostrovsky and Skeith [5] shows a matching upper bound for PIR-writing using [3] in a black-box way. Thus, we prove a *matching* black-box lower bound for [5].

Theorem 3 (informal) We prove $\Omega(\sqrt[t]{n})$ bound for algebraic PIR-writing based on homomorphic encryption that allows evaluation of total degree t multivariate polynomials on ciphertext. (We stress that cryptosystems for such structures are not known today beyond polynomials of total degree 2.)

Theorem 4 (informal) We show a constructive proof of a 1974 theorem of Werner [28] demonstrating the existence of a fully homomorphic encryption scheme (over a non-zero ring) if and only if there exists homomorphic encryption over any finite non-abelian simple group. (In the full version [21], we also show an explicit construction to implement a composable "NAND" gate from a group formula in any non-abelian simple group.) This also generalizes the result of Barrington [1] to all groups containing a finite non-abelian simple subgroup, as well as generalizing a result from Rappe [23].

A central element of this paper, from which we will derive a number of results, is an algebraic lower bound for a certain task- that of specifying "characteristic vectors" over a group. For a group G, we call a vector $(v_1, ..., v_n) \in G^n$ "characteristic" for a set $S \subset [n]$ if $v_i \neq \mathbf{0}_G$ if and only if $i \in S$, where $\mathbf{0}_G$ is the identity of G. We'll show that

> **Theorem (informal):** For *any abelian group*, communication complexity $\Omega(n)$ is required to "algebraically" specify characteristic vectors of arbitrary singleton subsets of $[n]$.

A formal statement of this idea appears as Theorem 2.7.

We stress that this statement holds for *all* abelian groups. For intuition, one may consider the case of linear algebra, in which the group G is of prime order, and has a field structure which could be put in place. It is a relatively simple exercise to prove this special case of our theorem, just arguing about the degree of vector spaces. However, note that this technique does not get very far. As the reader will see from Example 2.6 below, these ideas don't apply to general abelian groups G, even in the special case of cyclic groups. (Note that there is not even a well-defined notion of degree in this setting.) A "degree-based" argument could likely be carried out via free-module analysis, but it will substantially complicate and obfuscate matters, and furthermore it will yield a weaker version of the theorem. The abstract approach taken here will yield a strong algebraic result which will be of great utility later on, when we generalize to other structures.

Additionally, we prove a smooth trade-off in communication complexity as the number of non-identity elements in the characteristic vectors increases, and as mentioned, we also generalize to other algebraic structures, which contain virtually all structures that are preserved by known homomorphic encryption schemes. In particular, we prove results for *any abelian group* as well as results for arbitrary rings, in a setting restricted to polynomials of total degree t. (For an example of the case $t = 2$, see the cryptosystem of Boneh, Goh, and Nissim [3].) Finally, we'll show a number of natural cryptographic protocols that would imply the functionality of generating characteristic vectors, and hence derive algebraic lower bounds for the communication involved in these protocols as well.

As one will see after an examination of our algebraic results, they are in fact quite general. Since the results for abelian groups apply to all affine maps, this rules out many possibilities which do not necessarily come from group formulas. (For example, arbitrary endomorphisms may now be included in the class of "formulas" even though there is generally no way to compute all endomorphisms via an abelian group formula.) In particular, even if one changes their representation of data to be not just one group element, but many, and furthermore manipulates each of these elements independently, our results still apply (this is a simple consequence of Corollary 2.15).

Finally, regarding the equivalence of ring and group homomorphic encryption, we demonstrate that with *any* simple non-abelian group structure one can (constructively) compute all finite functions via group formulas and thus, the existence of any cryptosystem homomorphic over a simple non-abelian group implies a fully-homomorphic encryption scheme. This work can be found in the later sections, and somewhat generalizes results of Barrington [1] and Rhodes [18], however it is essentially a new and constructive version of the results of Werner [28] and may be of independent interest. This also answers an open question posed by Rappe in [23], who in 2004 proved a special case of this result.

1.2 Related Work

The lower bounds that we consider are most closely related to computational lower bounds on number theoretic problems when algorithms are restricted only to underlying group operations. For example, Boneh and Lipton [4] examine the computational difficulty breaking any algebraically homomorphic (over a field) cryptosystem. In contrast, our lower bounds are on communication complexity and apply to a wide variety of algebraic structures. Other related works are that of Shoup [27] and Maurer and Wolf [17], which consider computational difficulty of the discrete logarithm problem, and other number-theoretic problems in cyclic groups, provided that the algorithms do not exploit any specific properties of the representation of group elements.

Again, our lower bounds are geared towards communication complexity and program size, rather than computational complexity, but similar to these works, we focus only on algorithms that utilize nothing other than the underlying algebraic structures. However, we consider a greater variety of structures in our work (including arbitrary abelian groups and bounded degree multivariate polynomials over rings).

1.3 Overview, Motivation and Intuition

Often times, novel cryptographic protocols are developed using homomorphic encryption as building block (and often it is the only necessary ingredient). Many basic protocols can be constructed in this way, for example, private information retrieval, oblivious transfer, and collision-resistant hashing, to name a few. Indeed, such methods have accomplished much in the past, and continue to prove themselves as fruitful techniques. However, the types of algebraic structures available in homomorphic encryption schemes are limited. Not much beyond the structure of an abelian group can be preserved under an encryption scheme. Quite clearly, abelian groups have limited computing power. If one simply examines the number of distinct m-variable "formulas" in a finite abelian group G of order k in comparison to the number of G-valued functions (as set maps) that depend on m variables, one can't help but notice a great discrepancy in cardinality, so indeed, there is much that cannot be computed using only abelian group formulas. But what are these functions? Furthermore, in what sense can they not be computed or represented?

Using a black box model for homomorphic encryption, one is limited to only computing such formulas. However, there are a vast number of other types of "algebraic" maps which cannot necessarily be derived from any such formula that we study as well. We'll show that these maps also do not suffice for our tasks. As a somewhat trivial example, consider the endomorphism on $G = \mathbb{Z}_p \times \mathbb{Z}_p$ obtained by switching the coordinates. If one has only black box access to the group operation, then this endomorphism is not computable, however, if elements are represented as coordinate pairs, then computing this map is trivial. This is by no means the most complex example, but it does illustrate the benefits of an abstract approach (which will naturally cover all endomorphisms).

As mentioned before, there have been many protocols of great utility derived from homomorphic encryption over abelian groups (e.g. [15, 6, 14]). However, as the authors believe, for every such useful protocol in the literature, there are many dead ends, lying at the bottom of stacks of paper upon researchers' desks. But until now, there has not been much formal proof that these dead ends are actually just that. This work provides some basic proofs of lower bounds for a few simple protocols, based on these algebraic assumptions. More importantly, any task that can be reduced to our basic task is also immediately impossible to accomplish in an algebraic way with small communication complexity. We illustrate the power of these reductions on a number of examples below.

1.4 Implications of the Results

As we have mentioned, the applications of these results as lower bounds for cryptographic protocols are limited to an algebraic setting and are not absolutes. However, in many situations the bounds are quite practical. We'd like to take a moment to better illustrate and clarify where these results apply and where they do not. Additionally, we'll demonstrate the algebraic strength of the results, which are quite complete in the algebraic context.

The practical cryptographic significance of the results primarily deals with building protocols for computing on encrypted data. Let us consider single-database PIR, introduced by Kushilevitz and Ostrovsky [15]. PIR schemes are often based upon homomorphic encryption (e.g., [15],[16]). In the most efficient versions of these schemes, note that the answers to queries can be viewed as encryptions of the appropriate database elements- however, due to the repeated segmentation and application used to achieve better efficiency, these encryptions have *no algebraic value* after the second iteration (i.e. recursive calls in [15] scheme.) Roughly what is meant is that there is no way to combine two or more of these results (without the decryption key) to obtain an encryption of some other meaningful combination of the original elements. Looking at PIR alone, in its own context, this is not much of a problem. However, if PIR were to be used as a subroutine for some larger computation on encrypted data, this lack of algebraic value of the PIR "answer" could be very inconvenient. For example, the keyword search of [19] could be improved greatly if an efficient *algebraic* PIR protocol existed (see Section 3 for more details). To summarize very briefly, these results apply to situations where it is necessary to preserve (in the ciphertext) algebraic value of the results of underlying computation on encrypted data. In situations where algebraic value can safely be destroyed, there are often much more communication-efficient solutions.

We obtain results that hold *for all abelian groups*, as well as several other structures. We believe that this level of generality is a necessity. Details can be found in [21] (the full version of this work), but we summarize here. First of all, just using the simple structure of abelian groups, and the general abundance of cryptosystems that are homomorphic over cyclic groups, it is not hard to imagine constructing a homomorphic cryptosystem over virtually any abelian group. (For a more formal approach to this idea, see [11].) So, to make the results have any significance at all, a study at this level of generality is necessary, even though it requires additional machinery. In the case say of prime order cyclic groups, linear algebra suffices to solve the problem. However, this already breaks down for a general cyclic group (see Section 2). In addition to being insufficient from the start, a less general approach to the problem will also interfere the generalizations to other structures (see 2.16), as well as weakening the basic applications. For example, in addition to all algebraic formulas, the algebraic results as stated here cover the entire ring of endomorphisms (which in general may have no algebraic formula at all, much less a linear one). This would greatly complicate the set of functions to consider, and makes an elementary approach difficult. However, the abstract approach eliminates these issues[1].

2 Preliminaries and Basic Results

Most notations used are standard, and the algebraic notation used is typically consistent with [13]. However, a more comprehensive list of the notation used in this work can be found in [21], the full version.

2.1 Equivalence of Homomorphic Encryption over Non-abelian Simple Groups and Rings

We'll begin by stating a positive result regarding the equivalence of homomorphic encryption over non-abelian simple groups and rings. For more thorough formalizations, please see [21], the full version of this work.

Theorem 2.1 *Let G be a finite non-abelian simple group. Then any function $f : \{0,1\}^m \longrightarrow \{0,1\}^n$ can be represented solely in terms of the group operation of G.*

First, we'll prove a few elementary lemmas, and then the theorem (which again, uses only basic techniques from algebra). To begin, recall that from the Feit-Thompson theorem and Cauchy's theorem, we have that every non-abelian simple group of finite order has an element of order 2.

Lemma 2.2 *Let G be a finite group and suppose that $S \subset G$ is conjugation invariant (i.e., $\forall s \in S, g \in G$ we have $gsg^{-1} \in S$). Then $\langle S \rangle \lhd G$.*

The proof is straightforward, but can be found in its entirety in the full version [21].

Consider for a moment, the conjugacy classes. For an element $x \in G$, we will denote the conjugacy class by $\mathrm{Cl}_G(x)$. I.e.,

$$\mathrm{Cl}_G(x) = \{y \in G \mid y = gxg^{-1} \text{ some } g \in G\}$$

Recall that we can define a natural action of G on $\mathrm{Cl}_G(x)$ for any $x \in G$: for all $s \in \mathrm{Cl}_G(x)$, simply define $g \cdot s = gsg^{-1}$. Now, let G be a non-abelian simple group of finite order. From Cauchy's theorem, we know that there exists $x \in G$ such that x has order 2. Consider $\mathrm{Cl}_G(x)$. Let $|\mathrm{Cl}_G(x)| = k$. It must be the case that $k > 1$. If not, then every element of G conjugates x to itself, and hence we have $x \in \mathbf{Z}(G)$, the center of G. But of course this is impossible since the center of a group is always normal and we assumed that G is simple. So, the conjugacy class of x has at least

[1]Again, consider the simple example of $G = \mathbb{Z}_p \times \mathbb{Z}_p$ as a black box and as a direct product coordinate representation. The endomorphism $\varphi \in \mathrm{Hom}_{\mathbb{Z}}(G, G)$ by $(a, b) \mapsto (b, a)$ is not computable as a formula, but clearly is computable if given a coordinate representation.

two elements. Recall next, that whenever a group acts on a set S of size k, there is an induced homomorphism,

$$\varphi : G \longrightarrow S_k$$

Since the action of G on $\mathrm{Cl}_G(x)$ is obviously transitive, and since the size k of the class of x is greater than 1, we see that φ cannot be the trivial homomorphism which sends all elements to the identity, and hence $\ker(\varphi) \neq G$. But, since G is simple, we in fact know that $\ker(\varphi)$ must be the trivial subgroup $\{e\}$, since the kernel is always normal. *Therefore, every element of G acts non-trivially on the set $\mathrm{Cl}_G(x)$.*

We will extract the useful information into the following lemma which we have just now proved.

Lemma 2.3 *Let G be a finite, non-abelian simple group, and let $x \in G$ be an element of order 2. Then there exists an element $y \in \mathrm{Cl}_G(x)$ such that $yxy^{-1} \neq x$, and hence, such that $[x, y] \neq e$.*

Using these facts, we can now prove Theorem 2.1.

Proof: We will simply show that the function $\mathrm{NAND}(a, b)$ is computable in this way, which suffices to prove the theorem since any such function $f : \{0,1\}^m \longrightarrow \{0,1\}^n$ can be written in terms of compositions of NAND alone. More precisely, we will show that for an element x of order 2, the set $\{e, x\}$ can be identified with $\{0, 1\}$ respectively, and the operation NAND can be computed solely in terms of the group operation of G.

So, to begin, let $x \in G$ be of order 2, which as we discussed exists by Cauchy's theorem. Define $C = \mathrm{Cl}_G(x)$. As discussed, $|C| > 1$. Consider $S = [C, C]$, the set of *commutators in C*. Note that the subset S is conjugation invariant since it is generated by $C = \mathrm{Cl}_G(x)$, which is quite clearly conjugation-invariant. Hence by Lemma 2.2, the subgroup generated by *these specific commutators*, is a normal subgroup: $\langle S \rangle = \langle [C, C] \rangle \lhd G$ However, by Lemma 2.3, we know that $|S| > 1$, as there are at least 2 non-commuting elements. But, we assumed that G was simple. Therefore, we have in fact that $\langle S \rangle = G$. So, in particular, there exists some product, $s_1 s_2 \cdots s_k$ of commutators in C such that $s_1 s_2 \cdots s_k = x$ So, each $s_i = [r_i, t_i]$ where r_i and t_i are both conjugate to x. Therefore we have sequences of group elements, $\{g_i\}_{i=1}^k$ and $\{h_i\}_{i=1}^k$ such that $[g_i x g_i^{-1}, h_i x h_i^{-1}] = s_i$ We are now ready to define our $\mathrm{NAND}(a, b)$. First, define the function $\mathrm{AND}(a, b)$ as follows:

$$\mathrm{AND}(a, b) = \prod_{i=1}^{k} [g_i a g_i^{-1}, h_i b h_i^{-1}]$$

It is now easy to observe that it performs the appropriate function on our inputs from $\{e, x\}^2$. Whenever a or b is set to the identity, every commutator will of course be the identity since all elements commute with e. However, if both a and b are set to the group element x, the by our design, we will have $\mathrm{AND}(x, x) = x$, exactly as desired. Now, since x has order 2, we can simply define $\mathrm{NAND}(a, b) = \mathrm{AND}(a, b)x$. This completes the proof. ■

Corollary 2.4 *Constructing a fully homomorphic encryption scheme over a ring with identity is equivalent to constructing a group homomorphic encryption over* any *finite non-abelian simple group. In particular, it is equivalent to constructing a homomorphic encryption scheme over A_5, the smallest such group.*

Proof: This is almost immediate, but see the full version [21], or [23] for more detail.

As mentioned, detailed examples and formalizations can be found in the full version [21] of this work.

2.2 Generating Encryptions of Characteristic Vectors: Motivation

This example provides a simple description of a protocol that can't be non-trivially implemented with abelian group algebra. Later, we'll show a variety of problems (usually related to PIR or PIR-writing) which would imply a protocol like this. Hence, these too cannot be implemented with abelian group algebra.

We could, at this point, formalize a cryptographic protocol about generating characteristic-type vectors over a group, but it may be convenient to postpone such a definition and instead get right to the main algebraic point. So for the moment, we'll just explain in simple terms the algebraic task we are trying to accomplish.

Consider the following problem: Let $n, m \in \mathbb{Z}^+$, and let G be an abelian group. Define the following elements $v_i \in G^n$:

$$v_i = (\mathbf{0}_G, ..., \mathbf{0}_G, x_i, \mathbf{0}_G, ..., \mathbf{0}_G)$$

where $x_i \neq \mathbf{0}_G$ appears in the i-th position.[2] Let $\{\mathfrak{m}_i\}_{i=1}^n \subset G^m$ and let f be an arbitrary affine group map in m variables from $G^m \longrightarrow G^n$, i.e., $f = f_m + c$ where $f_m : G^m \longrightarrow G^n$ is linear and $c \in G^n$. Note that these affine maps can express all possible abelian group formulas on a set of variables (see the full version [21] for complete formalization and definitions). The question is

Question 2.5 (Informal) *If $f(\mathfrak{m}_i) = v_i$ for all $i \in [n]$, what can be said about $|G^m|$? In particular, how small can it be?*

We will soon answer this question in a variety of contexts, but first we'll give an example to help motivate the question and our lower bound. The phrasing used regarding the size estimation was deliberate: we don't isolate or bound m alone, because we cannot bound m in a non-trivial way. It is in fact possible to accomplish the above result with $m = 1$, even for a cyclic group. However, as we'll show in our lower bound, this comes at the cost of increasing the size of G.

Example 2.6 *Let $n \in \mathbb{Z}^+$, and let $N = \prod_{i=1}^n p_i$, where p_i is the i-th prime number. Define $G = \mathbb{Z}_N$. Define integers $\{z_i\}_{i=1}^n$ as follows:*

$$z_i = \prod_{j \neq i} p_j$$

Then, since all the primes were distinct, it is easy to verify that

$$(z_i z_j \neq 0 \bmod N) \iff (i = j)$$

So, we could define a linear function $f = (f_1, ..., f_n)$ from $G \longrightarrow G^n$ by $f_i(x) = z_i \cdot x$, and we would have $f(z_i) = v_i$, for some elements $v_i \in G^n$ which fit the above description of a complete set of characteristic vectors.

However, in the preceding example, notice that n different primes had to divide the order of G. Hence, $|G| > 2^n$ is of exponential size in n. We will show that even using affine maps, this is always the case: to generate n orthogonal-type characteristic vectors with m group elements always requires a group G such that G^m has exponential size in n, although the statement we prove has a more abstract setting.

[2]We give x an index i simply to show that it need not be uniform across all vectors.

2.3 A Basic Algebraic Result

Here, we will make precise the relationship regarding n and the size of an abelian group that can algebraically generate a complete set of n characteristic vectors over an abelian group G. Again, to conserve space, we direct the reader to the full version [21] for most of the proofs.

Theorem 2.7 *Let $n \in \mathbb{Z}^+$ and let G, A be abelian groups. Let $V = \{v_i\}_{i=1}^n \subset G^n$ be any collection of elements so that the j-th position of v_i is $\mathbf{0}_G$ if and only if $i \neq j$. Then if $F = f + c$ is an affine map from $A \longrightarrow G^n$ such that $V \subset F(A)$ then we have $\log(|A|) \in \Omega(n)$. More specifically, if $A \subset G^m$, we have that*

$$\log(|G|) \geq \frac{n}{m+1}$$

The proof of this theorem is given in the full version [21]. The pieces used are outlined below, and their proofs can also be found in [21] as well. To begin, we'll prove the following lemma which will help us analyze affine maps and translated characteristic vectors.

Lemma 2.8 *Let R be a finite ring with identity, and let M be a (unitary) R-module. Let $\Omega = \{\omega_i\}_{i=1}^k \subset M$ be a finite collection of elements. Let $\Omega' = \{(\omega_i + c)\}_{i=1}^k$ for some fixed element $c \in M$. Then $\langle \Omega' \rangle$, the module generated by Ω', increases in size by at most a factor of $|R|$ over the size of $\langle \Omega \rangle$. I.e.,*

$$\frac{|\langle \Omega' \rangle|}{|\langle \Omega \rangle|} \leq |R|$$

Proof: See the full version [21].

In light of Lemma 2.8, we need only to analyze "un-translated" characteristic-type vectors. If they generate a large module, then so will the translated vectors. It is quite clear any such module generated by elements like those in V will be exponential in size, however to be complete, we provide a formal proof.

Observation 2.9 *Let G be a finite abelian group. Let $n \in \mathbb{Z}^+$. Define elements $v_i \in G^n$ by $v_{ij} = \delta_{ij} \cdot \alpha_i$ for some $\alpha_i \neq 0 \in G$, and $\delta_{ij} \in \mathbb{Z}$ with $\delta_{ii} = 1$ for all i and $\delta_{ij} = 0$ for $i \neq j$. Let $H = \langle \{v_i\}_{i=1}^n \rangle$, the subgroup of G^n generated by the v_i. Then $|H| \geq 2^n$.*

Proof: See the full version [21].

We'll also make use of a few very elementary observations from group theory. As elementary as they may be, proofs can none the less be found in the full version [21].

Observation 2.10 *Let G be an abelian group and let $a, b \in G$ with $x = \mathrm{ord}(a), y = \mathrm{ord}(b)$. Then* $\mathrm{ord}(ab) \mid \mathrm{lcm}(x, y)$.

Observation 2.11 *Let G, H be groups, and let $f : G \longrightarrow H$ be a homomorphism. Then for all $g \in G$, we have that* $\mathrm{ord}(f(g)) \mid \mathrm{ord}(g)$.

Observation 2.12 *Let G be a group, and let $(a, b) \in G \times G$. Then* $\mathrm{ord}((a, b)) = \mathrm{lcm}(\mathrm{ord}(a), \mathrm{ord}(b))$.

Observation 2.13 *Let G be an abelian group, and suppose that there exists $N \in \mathbb{Z}^+$ such that $N \cdot g = \mathbf{0}_G$ for all $g \in G$, where $\cdot$ denotes $\mathbb{Z}$-module action. Then, G is a $\mathbb{Z}_N$-module, where the action is inherited from that of $\mathbb{Z}$.*

These basic observations and lemmas are enough for the proof of Theorem 2.7 (which again, can be found in the full version [21]).

2.4 Functions that Change Multiple Values

We can also generalize this algebraic result to include other types of vectors, where $F(\mathfrak{m}_i)$ has the i-th component non-identity, but possibly some other number of positions are non-identity elements as well. If the function F has the ability to change arbitrary subsets of c elements for a constant c, then our original results clearly apply, as you could re-organize G^n as a product $G^c \times \cdots \times G^c$ with n/c components. (Without loss of generality, we assume $c|n$.) However, the bounds still apply for less powerful classes of functions. We will show that *any* function that produces vectors with $c(n)$ or fewer non-identity positions at a time has communication complexity $\Omega(n/c(n))$, provided only that it is complete- i.e., for every position, it has the ability to produce a vector that is non-identity in that position. Here, $c(n)$ is any positive function of n, and note also that the number of non-identity positions per $\mathfrak{m}_i$ need not be uniform- we only ask that it is bounded by $c(n)$. We'll prove this by showing that we can always re-organize G^n into a product of larger components (of size $c(n)$) so that the original function F produces orthogonal characteristic-type vectors in the original sense, only over $(G^{c(n)})^{n/c(n)}$. Then, the proof follows immediately from the original result. Consider the following lemma.

Lemma 2.14 *Let $c \in \mathbb{Z}^+$. Let $\{S_k\}_{k\in\Gamma}$ be a collection of sets such that $S_k \subseteq [n]$, $|S_k| \leq c$ for all $k \in [n]$ and such that the $\{S_k\}$ form a cover of $[n]$, i.e., $\bigcup_{k\in\Gamma} S_k = [n]$. Then there exists $X \subseteq [n]$ and a sub-collection of sets $\{S_{k_j}\}_{k_j\in\Lambda\subseteq\Gamma}$ such that $S_{k_j} \cap S_{k_{j'}} \cap X = \varnothing$ whenever $j \neq j'$ yet $S_{k_j} \cap X \neq \varnothing$ for at least $\lceil n/c \rceil$ of the sets S_{k_j}.*

Proof: See the full version [21].

Corollary 2.15 *Let $n \in \mathbb{Z}^+$ and let G, A be abelian groups. Let $w(x)$ be a positive valued function and let $V = \{v_i\}_{i=1}^n \subset G^n$ be any collection of elements so that the i-th position of v_i is not equal to $\mathbf{0}_G$, and at most $w(n)$ total positions of v_i are non-identity for all $i \in [n]$. Then if $F = f + c$ is an affine map from $A \longrightarrow G^n$ such that $V \subset F(A)$ then we have* $\log(|A|) \in \Omega(n/w(n))$.

Proof: See the full version [21].

2.5 Polynomials of Bounded Total Degree

Recently, new cryptosystems have been developed with additional homomorphic properties (see [3]), which provide the ability to compute on ciphertext, polynomials of total degree at most 2. Here, we will generalize our original algebraic result to apply to algebraic functions of the form of any polynomial of total degree t, over a ring R. Although the following result will apply to the ring of polynomials over any ring R (it need not have an identity or be commutative), this result has the most meaning in the case of commutative rings with identity, since in this case the ring of multivariate polynomials coincides precisely with our notion of "algebraic formula", which is formalized in [21], the full version of this work. (For a non-commutative ring, there's a more general structure that serves as the set of all formulas.)

Corollary 2.16 *Let $n \in \mathbb{Z}^+$ and let R be any ring. Let $V = \{v_i\}_{i=1}^n \subset R^n$ be any collection of elements so that the j-th position of v_i is not equal to $\mathbf{0}_R$ precisely when $j = i$, for all $i, j \in [n]$. Then if $F : R^m \longrightarrow R^n$ is such that $F = (F_1, ..., F_n)$ with each $F_i \in R[X_1, ..., X_m]$ of total degree less than or equal to t (a constant) and has $V \subset F(R^m)$ then we have $\left(\sqrt[t]{\log(|R|)}\right) m \in \Omega(\sqrt[t]{n})$. In particular, if $|R|$ is independent of n, then $m \in \Omega(\sqrt[t]{n})$.*

Proof: See the full version [21].

3 Applications of Algebraic Results

We will discuss here a number of protocols which are both easy to state, and would provide desirable functionalities, yet under algebraic assumptions, they cannot be very well implemented with existing technology.

3.1 Private Database Modification (PIR Writing)

As seen in [5], the ability to privately modify an encrypted database in a communication efficient way could provide a valuable tool for private computation. One very natural approach to such a problem, is to proceed in a manner analogous to many PIR protocols, and use homomorphic encryption as a building block (as was done in [5]).

The protocol would then communicate encrypted values which encode the modification to take place, and then the database owner would execute some algebraic operations on the encrypted database and the description given by the user to update the database contents. Since all of the communication consisted only of encrypted values, CPA-type security comes easily from a hybrid type argument. Unfortunately, we have very limited structures available to homomorphic encryption schemes. Almost always, what is preserved is the operation of an abelian group. At best, the ability to evaluate polynomials of total degree 2 is provided (see [3]). It will follow from our preliminary algebraic results, that these types of algebraic protocols cannot be very well implemented with existing encryption schemes. We'll often speak of "algebraic" maps, which will usually mean functions that are obtainable from some type of formula involving only the operations of the algebraic structure. A precise, formal, and detailed exposition of this idea is given in [21], the full version of this work. Also, we've omitted some of the formal protocol-type definitions to improve readability, since there isn't much surprising about them, and most readers of this paper could likely re-invent them in a few minutes. We'll instead give an informal description of the protocol here. For precise statements, again we direct the reader to the full version [21].

Let **U** be a user that wishes to update the database, and denote by **DB** the database owner. We'll summarize a protocol for algebraic database modification between **U** and **DB** via the following steps, in which we assume that G is an abelian group. Below, we'll just describe the algebra involved. In an actual privacy preserving protocol, everything will of course be computed on ciphertext in some homomorphic encryption scheme over G.

1. **U** selects $\mathfrak{m}_i = (g_1, ..., g_m)$ to modify position i and sends $\mathfrak{m}_i$ to **DB**.
2. **DB** computes an algebraic function $F(X, \mathfrak{m}_i, H)$ of the database $X \in G^n$, the modification description $\mathfrak{m}_i$, and other inputs of his own, $H \in G^\epsilon$.
3. **DB** replaces X by $X' = F(X, \mathfrak{m}_i, H)$

Clearly the algebra involved in this protocol implies the ability to algebraically generate complete sets of characteristic vectors:

Claim 3.1 *An algebraic protocol for database modification over an abelian group implies an algebraic function (affine map) with a complete set of characteristic vectors in the image.*

Proof sketch: Define a database $X = \{\mathbf{0}_G\}_{i=1}^n$, which is the identity in all positions. Apply **DB**'s function to obtain $X' = F(X, \mathfrak{m}_i, H)$ where $\mathfrak{m}_i$ describes a modification for position i. Then clearly $X' = v_i$, a characteristic vector in G^n, non-identity at position i. ■

Therefore, by Theorem 2.7, if we build such a protocol based on a homomorphic cryptosystem over *any* abelian group, it will necessarily have linear communication complexity. Note the strong sense in which this is true: abelian group formulas always correspond to affine maps, but certainly not every affine map comes from such a formula.[3] So, we've shown that even if we allow **DB** to somehow compute arbitrary affine maps on the ciphertext values, it still does not suffice to accomplish this task. Furthermore, Theorem 2.7 did not even assume that the groups were the same. So, even if the database elements are encrypted in some other cryptosystem which is homomorphic over a group different than that of the descriptions, and even if we were provided the ability to compute all algebraic maps from one to the other on ciphertext, we still couldn't produce a non-trivial protocol over abelian groups. We'll summarize these ideas as

Corollary 3.2 *There are no non-trivial Algebraic Oblivious Database Modifiers over an abelian group. I.e., any oblivious database modifier based on the operations of an abelian group has communication complexity* $\Omega(n)$.

3.2 Algebraic and Homomorphic Protocols for PIR

As a second corollary, we consider "algebraic", or "homomorphic" protocols for private information retrieval. One may have observed, as the authors have, that the query results for PIR protocols usually fall into one of two categories: either (a) they have no (or very limited) algebraic value[4] or homomorphic properties, or (b) the server side communication is non-constant, i.e., the results of a query return many items, not just an encryption of one value in the database[5]. A protocol for private information retrieval that returns encryptions of single values which retain algebraic and homomorphic properties could be a very useful tool in private computation[6], especially in non-interactive settings. In what follows, we present evidence that the absence of such protocols is perhaps to be expected.

We'll try to establish a basic definition that captures the properties that we desire, and encapsulates most existing work possessing these properties. Suppose that the values in a database have some algebraic structure. For now, say that of an abelian group which we'll denote $(G, \cdot)$. We will denote the return value of a PIR query for the i-th position of a database by $\mathrm{PIR}(i)$, which consists of one or more encrypted database elements. Let $S_i = \{s_j\}_{j=1}^k$ denote the set of values from the database that are returned by a PIR query for position i.

Suppose for a moment that the domain from which PIR query returns reside has the algebraic structure of a group as well, say $(G', \star)$[7]. To name just a few examples, one can see that the PIR protocols of [15], [6], [7] all fit this description. We could then make the following definition:

[3]Again, consider $G = \mathbb{Z}_p \times \mathbb{Z}_p$ and $\varphi \in \mathrm{Hom}_{\mathbb{Z}}(G, G)$ by $(a, b) \mapsto (b, a)$.

[4]See the work of [7] for an example of such a PIR protocol having "limited" algebraic value.

[5]See [15] for such an example, but many PIR protocols based on homomorphic encryption (over an abelian group) have this property.

[6]For example, in the keyword search of [19], the dictionary size could be reduced.

[7]There is no assumption that the group representing the query returns are the same as the database elements, or even that they are encryptions of database elements, exactly. It could be the case that as a part of the encryption, the group that the database elements come from is first homomorphically transformed, and then transformed back as a part of decryption. The general way that we've stated our algebraic results will be useful for such a definition.

Definition 3.3 *Using the notation established above, we say that a PIR protocol is* **homomorphic** *if for a given database $X \in G^n$, we have that $\mathcal{D}(\mathrm{PIR}(i) \star \mathrm{PIR}(j)) = S_i \cdot S_j$ where $\mathcal{D}$ is the function from the PIR protocol that decrypts the query results.*

Note also that for such a PIR protocol to be of much utility as a subroutine in some non-interactive private computation, it is almost essential to have $|S_i| = 1$, or at least bounded by a small constant. If not, then the party which is to perform a computation on the return values of a homomorphic PIR query will likely not have any information about where the relevant element is in the query results. Hence, if such a party wishes to perform a computation on t variables obtained via homomorphic PIR queries, it would in general require repeatedly performing the computation on all $|S_i|^t$ possible sequences to ensure that the right variables were involved at least once. Furthermore, it may not be possible for any party to distinguish which of the resulting outputs in fact corresponds to the desired computation, even after decryption.

Finally note that from the definition of homomorphic PIR, we see that the results of queries are in fact encryptions of elements in some homomorphic cryptosystem. To create such a PIR protocol, a very natural approach is to manipulate the algebraic structure of some such homomorphic cryptosystem. This motivates the following definition.

Definition 3.4 *We say that a PIR protocol is* **algebraic** *if the following hold:*

1. *A query consists of an ordered sequence of ciphertexts in some cryptosystem where the plaintext set A has some algebraic structure.*
2. *To process a query, the database owner computes on ciphertext some algebraic function of the query's array, this function being determined by the contents of the database to obtain an array of ciphertext which will be the results of the query.*

For precise definitions of the term "algebraic function", we again direct the reader to [21], the full version of this work. In the case of abelian groups, these definitions yield affine maps as our model of algebraic functions.

Corollary 3.5 *Consider an abelian group algebraic PIR protocol with sender-side communication complexity $g(n)$ and server-side communication complexity $h(n)$. Then $g(n)h(n) = \Omega(n)$. More specifically, if $k(n)$ is any positive integer-valued function and if the server's response consists of $k(n)$ encrypted values, then the sender-side communication complexity is $\Omega(n/k(n))$.*

Proof: See the full version [21].

Using Corollary 2.16, we can generalize this result to cryptosystems that may have additional homomorphic properties (see [3]), showing $\Omega(\sqrt[t]{n})$ bounds if total degree t polynomials over a ring R can be computed on ciphertext.

For example, if given a cryptosystem that allows polynomials of fixed total degree t to be computed on ciphertext over some ring R, we can easily construct an algebraic PIR protocol with sender-side communication $\Theta(\sqrt[t]{n})$ and server-side complexity $\Theta(1)$ (see [3], or [21] for details of a simple example). However, this is in fact meets a lower bound: In general, if such a protocol has sender-side complexity $g(n)$ and server-side complexity $h(n)$, then we can show that $g(n)h(n) = \Omega(\sqrt[t]{n})$, which is a simple consequence of Corollary 2.16.

3.3 Private Keyword Searching [19]

As another relatively simple corollary, we resolve (under our algebraic assumptions) an open problem posed by Ostrovsky and Skeith [19] regarding extending the query semantics for private searching on streaming data. We show that without new homomorphic encryption schemes with additional properties, their methods cannot be extended to perform conjunctive queries.

Corollary 3.6 *The problem of private keyword search on streaming data as proposed in [19], has no non-trivial algebraic solution for a conjunctive query of two or more terms if the underlying cryptosystem is only group homomorphic over an abelian group.*

Remark: We will assume the same basic framework as developed in [19] for a solution and show that there is no such solution that performs conjunctive queries. Specifically, we assume that a dictionary with an associated array of ciphertexts is used to conditionally encrypt documents as in [19].

Proof: First note that the protocol inherently gives rise to an algebraic method for generating complete sets of characteristic vectors: Suppose that the dictionary D has size m. Each word has its role in the query encoded via an encrypted group element, say in some group G. Look at the encoded dictionary (un-encrypted) as the set G^m. Suppose we have a protocol as described in [19] for some query that involves k variables. Running this protocol on m^k documents which run over all unique k-tuples from the dictionary gives us a set of characteristic vectors inside of $G^{(m^k)}$. So, we can think of this as an algebraic map from $G^m \longrightarrow G^{(m^k)}$, which (unless the query is somewhat trivial) contains a complete set of characteristic vectors in the image. But, now the question is how many positions are non-identity in each vector? This of course depends on the query. Suppose that the query is a disjunction of terms. Each vector in $G^{(m^k)}$ will have at least $m^{k-1}k$ positions that are non-identity, since $k-1$ entries could be arbitrary as long as one contains a keyword. So, the ratio of total positions to non-identity positions is less than m and our algebraic lower bounds give no contradiction (which of course should be the case since [19] gives such a construction). But now consider a conjunctive query, just of two terms. In the same way as described above, this gives rise to an algebraic function for characteristic vectors from $G^m \longrightarrow G^{(m^2)}$, however this time we have $\mathcal{O}(1)$ positions of each vector are non-identity. So, applying Corollary 2.15, we see that no such protocol can exist based on an abelian group. More generally, from Corollary 2.16, we see that if given the ability to compute total degree t polynomials, we can construct a protocol that executes a conjunction of *at most* t terms. ■

We believe that this example illustrates particularly well a situation in which the bounds proved in this work are especially useful. The entire method of [19] critically depends on the ability to generate these types of characteristic vectors so that the final representation is an encryption in a homomorphic scheme. This is the case since the functionality of characteristic vectors is used as a subroutine for the larger procedure, and so to continue the computation (i.e., writing to the buffer, etc.) it is necessary that the output have algebraic value. So, since we have proven that this subroutine is impossible to implement in the required manner, it seems that improving the work of [19] would require either a completely new approach, or new designs of homomorphic encryption schemes, such as fully-homomorphic encryption.

It is this type of information that we hope will save researchers time and effort in the future. Applying these bounds may not give an absolute impossibility, but it can quickly eliminate a very large space of what might otherwise seem to be feasible approaches to the problem.

References

[1] D. Barrington. Bounded-Width Polynomial-Size Branching Programs Recognize Exactly Those Languages in NC. STOC 1986: 1-5

[2] D. Boneh, G. Crescenzo, R. Ostrovsky, G. Persiano. Public Key Encryption with Keyword Search. EUROCRYPT 2004: 506-522

[3] D. Boneh, E. Goh, K. Nissim. Evaluating 2-DNF Formulas on Ciphertexts. TCC 2005: 325-341

[4] D. Boneh, R. Lipton. Searching for Elements in Black Box Fields and Applications. CRYPTO'96, LNCS1109, pp. 283-297.

[5] D. Boneh, E. Kushilevitz, R. Ostrovsky, W. Skeith. Public Key Encryption that Allows PIR Queries. CRYPTO-2007.

[6] Y. C. Chang. Single Database Private Information Retrieval with Logarithmic Communication. ACISP 2004

[7] C. Cachin, S. Micali, and M. Stadler. Computationally private information retrieval with polylogarithmic communication. EUROCRYPT '99 pp. 402–414.

[8] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. In *J. of the ACM*, 45:965–981, 1998.

[9] T. ElGamal. A Public-Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms. IEEE Transactions on Information Theory, v. IT-31, n. 4, 1985, pp469472.

[10] S. Goldwasser and S. Micali. Probabilistic encryption. In J. Comp. Sys. Sci, 28(1):270–299, 1984.

[11] D. Grigoriev and I. Ponomarenko. Homomorphic public-key cryptosystems over groups and rings. In *Quaderni di Mathematica*, 2004, vol. 13, p. 305–325

[12] I. N. Herstein. Abstract Algebra. Prentice-Hall, 1986, 1990, 1996.

[13] T. W. Hungerford. Algebra. Springer-Verlag, Berlin, 1984.

[14] Y. Ishai, E. Kushilevitz, R. Ostrovsky. Sufficient Conditions for Collision-Resistant Hashing. TCC-2005

[15] E. Kushilevitz and R. Ostrovsky. Replication is not needed: Single database, computationally-private information retrieval. FOCS-97, pp 364–373.

[16] H. Lipmaa. An Oblivious Transfer Protocol with Log-Squared Communication. IACR ePrint Cryptology Archive 2004/063

[17] U. Maurer and S. Wolf. Lower bounds on generic algorithms in groups. EUROCRYPT '98, pp. 72–84.

[18] W. Maurer and J. Rhodes. A property of finite non-Abelian simple groups. In proc. Am. Math. Soc., vol. 16, pages 522-554 (1965).

[19] R. Ostrovsky and W. Skeith. Private Searching on Streaming Data. CRYPTO 2005, and Journal of Cryptology Volume 20:4, pp. 397-430, October 2007.

[20] R. Ostrovsky and W. Skeith A Survey of Single Database PIR: Techniques and Applications. PKC 2007) LNCS vol. 4450/2007.

[21] R. Ostrovsky and W. Skeith Algebraic Lower Bounds for Computing on Encrypted Data. Electronic Colloquium on Computational Complexity report TR07-22.

[22] P. Paillier. Public Key Cryptosystems based on CompositeDegree Residue Classes. EUROCRYPT 99 pp. 223-238.

[23] D. K. Rappe. Homomorphic Cryptosystems and their Applications Ph.D. Thesis, 2004, under E. Becker and J. Patarin

[24] R. L. Rivest, L. Adleman and M. L. Dertouzos, On data banks and privacy homomorphisms, In Foundations of Secure Computation, eds. R. A. DeMillo et al., Academic Press, 1978, pp. 169-179.

[25] R. L. Rivest, A. Shamir, and L. Adleman A method for obtaining digital signatures and public key cryptosystems, Commun. ACM 21 (1978), 120126.

[26] T. Sander, A. Young, M.Yung. Non-Interactive CryptoComputing For NC1 FOCS 1999: 554-567

[27] V. Shoup. Lower Bounds for Discrete Logarithms and Related Problems. Eurocrypt '97 pp. 256–266.

[28] H. Werner. Finite simple non-Abelian groups are functionally complete. In Bull. Soc. Roy. Sci. Liège, vol. 43, pp. 400, (1974)