

Finding Significant Fourier Coefficients: Clarifications, Simplifications, Applications and Limitations

Steven D. Galbraith Joel Laity Barak Shani

Received July 25, 2016; Revised March 5, 2018, and in final form December 13, 2018; Published December 20, 2018

Abstract: Ideas from Fourier analysis have been used in cryptography for the last three decades. Akavia, Goldwasser and Safra unified some of these ideas to give a complete algorithm that finds significant Fourier coefficients of functions on any finite abelian group. Their algorithm stimulated a lot of interest in the cryptography community, especially in the context of “bit security”. This manuscript attempts to be a friendly and comprehensive guide to the tools and results in this field. The intended readership is cryptographers who have heard about these tools and seek an understanding of their mechanics and their usefulness and limitations. A compact overview of the algorithm is presented with emphasis on the ideas behind it. We show how these ideas can be extended to a “modulus-switching” variant of the algorithm. We survey some applications of this algorithm, and explain that several results should be taken in the right context. In particular, we point out that some of the most important bit security problems are still open. Our original contributions include: a discussion of the limitations on the usefulness of these tools; an answer to an open question about the modular inversion hidden number problem.

1 Introduction

Let G be a finite abelian group. Fourier analysis provides a convenient basis for the space of functions $G \rightarrow \mathbb{C}$, namely the characters $\chi : G \rightarrow \mathbb{C}$. It follows that any function $f : G \rightarrow \mathbb{C}$ can be represented as a

Key words and phrases: significant Fourier transform, Goldreich–Levin algorithm, Kushilevitz–Mansour algorithm, bit security of Diffie–Hellman

linear combination $f(x) = \sum_{\alpha \in G} \widehat{f}(\alpha) \chi_{\alpha}(x)$, where \widehat{f} is the discrete Fourier transform of f . A standard problem is to approximate a function, up to any error term, using a linear combination of a small number of characters. This is not always possible, but for certain functions (which are called *concentrated*) it is possible. The coefficients in such an approximation are called *significant* Fourier coefficients, as their size is large relative to the function's norm. The simplest example of a concentrated function is a character itself.

A natural computational problem is to compute such an approximation. When doing this one might have a complete description of the function or, as will be the case in this paper, just a small set of values $f(x_i)$. The ability to choose specific x_i 's plays a crucial role in the ability to approximate f . Indeed, the main result in this subject is an algorithm that, given the ability to select the values x_i , efficiently computes a *sparse* approximation for any concentrated function on any abelian group G , by computing all its significant coefficients. On the other hand, when the x_i 's cannot be selected, such an algorithm is not known to exist in general. Furthermore it is conjectured that an efficient algorithm does not exist in the general case.

We use the general term *significant Fourier transform* (SFT) to refer to algorithms that compute a function's significant coefficients. SFT algorithms first appear explicitly in the work of Kushilevitz and Mansour [26], though some of the main ideas already appear in earlier works. Subsequently new algorithms were presented, in various special cases of groups or functions, until the work of Akavia, Goldwasser and Safra [3] who presented a generic algorithm for all finite abelian groups and all complex-valued functions. The algorithms in the literature are often presented very differently, and some of them are designed to fulfill a very particular task, but they are all based on the same mathematical principles.

The main aim of this paper is to present a complete study of the SFT algorithms. Our work unifies these algorithms by clarifying the core mathematics underlying them. Thus, our focus is on a broad mathematical overview using Fourier analysis on finite groups and elementary group theory. We remark that our work is not necessarily the best presentation of a specific SFT algorithm, but we believe that a reader who is interested in understanding the rules and framework of these algorithms would benefit from this work. Our study also leads to a new approach for some of the more complicated cases. Furthermore, this paper surveys applications of the SFT algorithm in the field of cryptography and also gives limitations for such applications.

The SFT algorithm and variants have received great attention in the literature outside the regime of cryptography. Researchers in engineering, concerned with practical applications in signal processing, have developed algorithms with greater efficiency (with respect to various metrics); for a recent survey on these algorithms see Gilbert, Indyk, Iwen and Schmidt [17]. Our work does not cover these developments.

Roadmap

Section 2 summarises the basic definitions. Section 3 presents the key ideas behind the SFT algorithm, and deals with some related issues. Specifically, with few examples we explain why being able to choose the inputs to the functions is essential and why one does not expect to have a similar tool when the inputs to the functions are chosen at random; In cases where the function values are given by an oracle, we analyze the case of working with unreliable oracles.

Section 3.1 reviews the development of ideas and highlights the contributions of Goldreich and Levin [19], Kushilevitz and Mansour [26], Mansour [33], Bleichenbacher [8] and Akavia, Goldwasser and Safra [3].

In Section 4 we outline our recent work [28] on applying *modulus switching* to this subject (namely to re-cast a function on \mathbb{Z}_p to a function on \mathbb{Z}_{2^n} for the nearest power of 2 to p). These ideas are very similar to the approach taken in Shor's (period-finding) algorithm [42]. The benefit of this new approach is twofold. Firstly, its analysis gives insights into the AGS algorithm. Secondly it provides a new approach for implementations and for proving concentration of functions. In particular we provide a new proof of a result by Morillo and Ràfols [38] (described in Section 4.1).

The SFT algorithm is a useful tool in the research area of bit security. Section 5 surveys bit security applications using the language of the *hidden number problem*: given f and oracle access to $f_s := f \circ \varphi_s$, for some function φ parameterized by an unknown value s , recover the value s . The main application is in the group $G = \mathbb{Z}_p$ for the particular function $\varphi_s(x) = sx \pmod{p}$, i.e. $f_s := f(sx)$. In this particular case the *scaling property* gives $\widehat{f}_s(\alpha) = \widehat{f}(\alpha s^{-1})$ for every $\alpha \in G$. It follows that f and f_s share the same coefficients in different order. If α is a significant Fourier coefficient of f and β is a significant Fourier coefficient of f_s then $\alpha\beta^{-1}$ is a candidate value for s .

Using this observation, Akavia, Goldwasser and Safra [3] showed that a number of bit security results (for RSA, Rabin, and discrete logs) can be re-proved using these tools. A classic result of this type, from Alexi, Chor, Goldreich and Schnorr (ACGS) [4], is that if one has an oracle that on input $x^e \pmod{N}$ (where (N, e) is an RSA public key) returns the least significant bit of x with probability noticeably better than $\frac{1}{2}$, then one can compute e -th roots modulo N . Håstad and Näslund [24] generalized this result for an oracle that returns any single bit of x (see also [20, Section 4.1]), but their method is very complex and requires complicated and adaptive manipulations of the bits. On the other hand, the algorithm given by AGS, which applies to functions with significant Fourier coefficients, is much clearer and is not adaptive.¹ Similar to Håstad and Näslund, Morillo and Ràfols [38] extended the AGS results to all single bit functions, by showing that each single bit function is concentrated and so has a significant Fourier coefficient (in particular, one can obtain the ACGS result for any bit). The SFT algorithm has also been used to show *search-to-decision* reductions for the *learning with errors* and *learning with rounding* problems [36, 9].

Subsequently, a number of papers [14, 15, 16, 49] have proved (or re-proved) various results on bit security in the context of Diffie–Hellman keys on elliptic curves and finite fields \mathbb{F}_{p^n} with $n > 1$, but these results consider an unconventional model that allows changing the curve or field representation. We emphasize that the requirement of chosen inputs for the functions restricts these applications. Indeed, the question of main interest, whether single bits of Diffie–Hellman shared keys are hardcore in a fixed representation, is still open. We elaborate on these applications in Section 5.

Section 6 explains a fundamental limitation to the approach described above: we prove that one can only solve the (chosen-multiplier) hidden number problem with these tools when the function φ_s is linear or affine. Therefore, these tools cannot be directly used to address the *elliptic curve hidden number problem* or the *modular inversion hidden number problem*. Our work therefore answers a question in [32].

¹We describe the notion of adaptiveness in Section 5.

2 Preliminaries

The following gives mathematical background needed to understand the paper and definitions that will be used throughout the paper. The main definitions and notation appear in the table in Section 2.4.

2.1 Fourier analysis on finite groups

We review basic background on Fourier analysis on discrete domains. Proofs and further details can be found in Terras [47].

Let $(R, +, \cdot)$ be a finite ring and denote by $G := (R, +)$ the corresponding additive abelian group. We are interested in the set of functions $L^2(R) := \{f : R \rightarrow \mathbb{C}\}$. The set $L^2(R)$ is a vector space over \mathbb{C} of dimension $|R|$, with the usual pointwise addition and scalar multiplication of functions. Convolution of two functions $f, g \in L^2(R)$ is defined by $(f * g)(x) = \frac{1}{|R|} \sum_{y \in R} f(x - y)g(y)$. The *expectation* of a function $f \in L^2(R)$ is defined to be $\mathbb{E}[f] = \frac{1}{|R|} \sum_{x \in R} f(x)$. The space $L^2(R)$ is equipped with an inner product $\langle f, g \rangle := \mathbb{E} \left[f(x) \overline{g(x)} \right] = \frac{1}{|R|} \sum_{x \in R} f(x) \overline{g(x)}$, where \bar{z} denotes the complex conjugate of $z \in \mathbb{C}$. The inner product induces a norm $\|f\|_2 = \sqrt{\langle f, f \rangle}$. We also define $\|f\|_\infty = \max_{x \in R} |f(x)|$.

One basis for this vector space is the set of *Kronecker delta functions* $\{\delta_i\}_{i \in R}$ ($\delta_i(j) = 1$ if $j = i$, otherwise $\delta_i(j) = 0$). This is an orthogonal basis with respect to the inner product. However, this basis is not as useful as the Fourier basis, as we will explain later in this section.

A *character* of an additive group G is a group homomorphism taking values in the non-zero complex numbers, namely $\chi : G \rightarrow \mathbb{C}^*$ such that $\chi(x + y) = \chi(x)\chi(y)$. Since $\chi(x)^{|G|} = \chi(|G|x) = \chi(0_G) = 1$, we see that the characters take values in the complex $|G|$ -th roots of unity. The set of characters of G forms a group (with respect to pointwise multiplication), isomorphic to G , which is often denoted \widehat{G} .

In general, we fix a choice of isomorphism $G \rightarrow \widehat{G}$ and denote it by $\alpha \mapsto \chi_\alpha$. In particular, for $G = \mathbb{Z}_N$ the characters are defined by $\chi_\alpha(x) := e^{\frac{2\pi i}{N} \alpha x}$ where $\alpha \in G$. For $G = \mathbb{Z}_{N_1} \times \dots \times \mathbb{Z}_{N_m}$, let $\alpha = (\alpha_1, \dots, \alpha_m)$ and $x = (x_1, \dots, x_m)$; the character χ_α is given by $\chi_\alpha(x) := \chi_{\alpha_1}(x_1) \cdot \dots \cdot \chi_{\alpha_m}(x_m) = e^{\frac{2\pi i}{N_1} \alpha_1 x_1} \cdot \dots \cdot e^{\frac{2\pi i}{N_m} \alpha_m x_m}$ and the map $\alpha \mapsto \chi_\alpha$ from G to \widehat{G} is an isomorphism. We sometimes write $\omega_N := e^{\frac{2\pi i}{N}}$ so that $\chi_\alpha(x) = \omega_N^{\alpha x}$.

The following relations are standard and can be used to show that the characters are orthonormal

$$\sum_{x \in G} \chi(x) = \begin{cases} |G| & \text{if } \chi \text{ is the identity in } \widehat{G}, \\ 0 & \text{otherwise,} \end{cases} \quad \sum_{\chi \in \widehat{G}} \chi(x) = \begin{cases} |G| & \text{if } x = 0, \\ 0 & \text{otherwise.} \end{cases}$$

If $G = \mathbb{Z}_{N_1} \times \dots \times \mathbb{Z}_{N_m}$ then for any subgroup $H \leq G$ we define the orthogonal set

$$H^\perp := \{a \in G \mid \chi_a(h) = 1 \text{ for all } h \in H\}. \tag{2.1}$$

This set is fundamental for the understanding of the SFT algorithm and appears frequently in Section 3.2. Using the relations above it can be shown that

$$\sum_{h \in H} \chi_h(x) = \begin{cases} |H|, & \text{if } x \in H^\perp, \\ 0, & \text{otherwise.} \end{cases} \tag{2.2}$$

The *Fourier basis* for $L^2(R)$ is the set \widehat{G} consisting of all the characters χ . It is an orthonormal basis. Therefore, we can represent each function $f : R \rightarrow \mathbb{C}$ uniquely as a linear combination $f(x) = \sum_{\alpha \in G} \widehat{f}(\alpha) \chi_\alpha(x)$ of the characters χ_α . The function $\widehat{f} : G \rightarrow \mathbb{C}$ given by $\widehat{f}(\alpha) = \langle f, \chi_\alpha \rangle$ is called the *discrete Fourier transform*. The map $f \mapsto \widehat{f}(\alpha)$ is \mathbb{C} -linear. Notice that a single Fourier coefficient encapsulates information about the function on the whole domain, unlike the representation in terms of Kronecker delta functions where one coefficient only holds information about the function at a single point.

Parseval's identity is the following relationship between the norms of f and \widehat{f} :

$$\|f\|_2^2 = \frac{1}{|G|} \sum_{x \in G} |f(x)|^2 = \langle f, f \rangle = \sum_{\alpha \in G} |\widehat{f}(\alpha)|^2 = |G| \cdot \|\widehat{f}\|_2^2.$$

Adopting signal-processing terminology, when we work with the values $f(x)$ for $x \in G$ we say that x is in the *time domain*. When we use the values $\widehat{f}(\alpha)$ we say $\alpha \in G$ is in the *frequency domain*. There does not seem to be a rigorous formulation of this terminology and we do not use it much, but the reader will find it very common in the engineering literature. We signal to the reader whether we are working in the time domain or frequency domain by using Latin letters x, y for elements in the former (elements of G), and Greek letters α, β for the latter (corresponding to elements of \widehat{G} , e.g. χ_α).

Let $R = \mathbb{Z}_{N_1} \times \dots \times \mathbb{Z}_{N_m}$ with componentwise addition and multiplication, and let $f, g \in L^2(R)$. Basic properties of the Fourier transform include the following (note that the basis of Kronecker delta functions does not satisfy these properties, which is one reason why it is less useful than the Fourier basis):

- (time) scaling: if $g(x) := f(cx)$ for $c \in R^*$, then $\widehat{g}(\alpha) = \widehat{f}(c^{-1}\alpha)$;
- (time) shifting: if $g(x) := f(c+x)$ for $c \in R$, then $\widehat{g}(\alpha) = \widehat{f}(\alpha)\chi_\alpha(c)$;
- (frequency) shifting: if $g(x) := f(x)\chi_c(x)$ for $c \in R$, then $\widehat{g}(\alpha) = \widehat{f}(\alpha - c)$;
- convolution-multiplication duality: $\widehat{f * g}(\alpha) = \widehat{f}(\alpha)\widehat{g}(\alpha)$.

We now recall some definitions from [3, 14, 38]. The same definitions can be made for functions over rings R where G is their additive group.

Definition 2.1 (Restriction). Given a function $f : G \rightarrow \mathbb{C}$ and a set of characters $\Gamma \subseteq \widehat{G}$, the *restriction of f to Γ* is the function $f|_\Gamma : G \rightarrow \mathbb{C}$ defined by $f|_\Gamma := \sum_{\chi_\alpha \in \Gamma} \widehat{f}(\alpha) \chi_\alpha$.

Definition 2.2 (ε -Concentration). Let $\varepsilon > 0$ be a real number. A family of functions $\{f_i : G_i \rightarrow \mathbb{C}\}_{i \in \mathbb{N}}$ is *Fourier ε -concentrated* if there exists a polynomial P and sets of characters $\Gamma_i \subseteq \widehat{G}_i$ such that $|\Gamma_i| \leq P(\log |G_i|)$ and $\|f_i - f_i|_{\Gamma_i}\|_2^2 \leq \varepsilon$ for all $i \in \mathbb{N}$.

Definition 2.3 (Concentration). A family of functions $\{f_i : G_i \rightarrow \mathbb{C}\}_{i \in \mathbb{N}}$ is *Fourier concentrated* if there exists a polynomial P and sets of characters $\Gamma_i \subseteq \widehat{G}_i$ such that $|\Gamma_i| \leq P(\log |G_i|/\varepsilon)$ and $\|f_i - f_i|_{\Gamma_i}\|_2^2 \leq \varepsilon$ for all $i \in \mathbb{N}$ and for all $\varepsilon > 0$.

Most applications are concerned with a single function that implicitly defines the entire family. In this case we informally say that the function, instead of the family, is concentrated. Examples of concentrated functions, and of this terminology, are given in Example 2.5.

Definition 2.4 (Heavy coefficient). For a function $f : G \rightarrow \mathbb{C}$ and a threshold $\tau > 0$, we say that a coefficient $\widehat{f}(\alpha)$ (of the character χ_α) is τ -heavy if $|\widehat{f}(\alpha)|^2 > \tau$.

By Parseval's identity it is evident the number of τ -heavy coefficients for a function $f : G \rightarrow \mathbb{C}$ is at most $\|f\|_2^2/\tau$ (see [26, Lemma 3.4] or [34, Lemma 4.8]). Thus, the cases of interest are where the latter value is polynomial in $\log(|G|)$, so there are at most polynomially many τ -heavy coefficients. This forces τ to be relatively large to $\|f\|_2$, e.g. $\tau = \|f\|_2^2/\text{poly}(\log(|G|))$. We remark that it might have been better to define a τ -heavy coefficient to satisfy $|\widehat{f}(\alpha)|^2 > \tau\|f\|_2^2$, however we keep the notion that is mostly used in the literature (as we show below most applications consider the specific case $\|f\|_2^2 = 1$).

The phrases *significant coefficient* and *heavy coefficient* are often used interchangeably to mean any coefficient $\widehat{f}(\alpha)$ which is large relative to the norm of the function, but without reference to any specific value of τ . In this paper our convention is to use “heavy” in a formal sense and “significant” in an informal sense.

The relationship between concentrated functions and functions with significant coefficients is subtle. If a function has a τ -heavy coefficient, then it is $(1 - \tau)$ -concentrated (with $|\Gamma| = 1$). But such a function is not necessarily ε -concentrated for all ε . The literature has tended to focus on concentrated functions, but for many of the bit security applications it is sufficient that the function has one or more significant coefficients. The distinction is important since it is harder to prove that a function is concentrated than to prove it has a significant coefficient.

Example 2.5. Here are some examples of functions with significant coefficients, most of which are concentrated:

- A single character is concentrated; that is, the family $\{\chi_\alpha : \mathbb{Z}_n \rightarrow \mathbb{C}\}_{n>\alpha}$ for some $\alpha \in \mathbb{N}$ is concentrated. The case $\alpha = 0$ corresponds to constant functions, which are concentrated but will be un-interesting in our applications.
- For the least-significant-bit function $\text{LSB}(x)$ on \mathbb{Z}_{2^n} , which gives the parity of x , the functions $f : \mathbb{Z}_{2^n} \rightarrow \mathbb{C}$ given by $f(x) := (-1)^{\text{LSB}(x)}$ are concentrated. Indeed, these functions correspond to the characters $f(x) = (-1)^x = \omega_{2^n}^{2^{n-1}x} = \chi_{2^{n-1}}(x)$.
- The functions $\text{half} : \mathbb{Z}_N \rightarrow \{-1, 1\}$, for which $\text{half}(x) = 1$ if $0 \leq x < \frac{N}{2}$ and $\text{half}(x) = -1$ otherwise, are concentrated; one has $\widehat{\text{half}}(\alpha) = \frac{1}{N}[\sum_{0 \leq x < \frac{N}{2}} \overline{\chi_\alpha}(x) - \sum_{\frac{N}{2} \leq x < N} \overline{\chi_\alpha}(x)]$. Elementary arguments (see Claim 4.1 below) show that

$$\left| \frac{1}{N} \sum_{0 \leq x < \frac{N}{2}} \overline{\chi_\alpha}(x) \right| = \left| \frac{1}{N} \sum_{0 \leq x < \frac{N}{2}} \omega_N^{-\alpha x} \right| < \frac{1}{\|\alpha\|_N}$$

where $\|\alpha\|_N$ denotes the unique integer in $(-N/2, N/2]$ that is congruent to α modulo N . Similarly $\left| \frac{1}{N} \sum_{\frac{N}{2} \leq x < N} \overline{\chi_\alpha}(x) \right| < \frac{1}{\|\alpha\|_N}$. These results can be used to show that half is concentrated on a set of characters α with small $\|\alpha\|_N$; See [3, Claim 4.1]. Similar arguments hold for the most-significant-bit function $f(x) := (-1)^{\text{MSB}(x)}$, thus it is also concentrated.

- For primes p , the functions $f : \mathbb{Z}_p \rightarrow \mathbb{C}$ given by $f(x) := (-1)^{\text{LSB}(x)}$ are concentrated. This follows from $f(x) = \text{half}(2^{-1}x)$ and the scaling property.

- The function $\text{LPN}_s : \{0, 1\}^n \rightarrow \{0, 1\}$, given by $\text{LPN}_s(x) = (-1)^{\langle x, s \rangle + e(x)}$ for e which is mostly 0 (and otherwise 1), has a significant coefficient and therefore is ε -concentrated (for some large ε). Let I be the set for which $e(x) = 1$, then $\widehat{\text{LPN}_s}(s) = \frac{1}{2^n} \sum_{x \notin I} 1 + \frac{1}{2^n} \sum_{x \in I} (-1) = 1 - \frac{2|I|}{2^n}$. Since the size $|I|$ is relatively small, the coefficient $\widehat{\text{LPN}_s}(s)$ is large, that is, the function LPN_s “behaves” like the character χ_s in $\{0, 1\}^n$. If $|I|$ is very small, for example $|I| = \text{poly}(\log |G|)$, then LPN_s is also concentrated. Moreover, one can show that $|\widehat{\text{LPN}_s}(v)| \leq \frac{|I|}{2^n}$, and on average is expected to be proportional to $\sqrt{2|I|/2^n(2^n - 1)} \approx \sqrt{2|I|/2^n}$.
- ‘Noisy characters’ given by $f(x) := \omega_p^{\alpha x + e(x)}$ for some suitable random functions e have a significant coefficient $\widehat{f}(\alpha)$ as we show in Section 6.1. An example of such a noisy character is the function $\text{LWE}_s : \mathbb{Z}_p^n \rightarrow \mathbb{Z}_p$, given by $\text{LWE}_s(x) = \omega_p^{\langle x, s \rangle + e(x)}$ for $e(x)$ drawn from a Gaussian distribution.

Another example of concentrated functions are the i -th bit functions, see Section 4.1 for details.

2.2 Learning model

Let $f : R \rightarrow \mathbb{C}$ be a function for which one wants to learn its significant coefficients. The learner gets *access* to samples of the form $(x, f(x))$. In the *random access* model the learner receives polynomially many samples for inputs $x \in R$ drawn independently and uniformly at random. As opposed to this model, in the *query access* model the learner can query the function on any chosen input $x \in R$ to receive the corresponding sample.

A *learning algorithm* for a function $f : G \rightarrow \mathbb{C}$ outputs a set containing all the significant Fourier coefficients of f . Formally, given a function f and $\varepsilon, \delta > 0$, the algorithm outputs a set Γ of size polynomial in $\log(|G|)$ and ε^{-1} , such that $\|f - f|_\Gamma\|_2^2 \leq \varepsilon$ with probability at least $1 - \delta$.

The main result of this subject (see Theorem 3.1 below) is that there is a randomised polynomial-time algorithm to compute a sparse approximation $f|_\Gamma$ to a concentrated function in the query access model. In other words, concentrated functions admit a polynomial-time learning algorithm in the query access model.

2.3 Probability

The Chernoff bound gives an upper bound on the probability that a sum of independent random variables deviates from its expected value. One can therefore derive a lower bound for the number of samples needed to estimate the sum of independent random variables, with any required probability and error term. For a random variable X on a set $A \subseteq \mathbb{C}$ we denote by $\mathbb{E}_{x \in A} X(x)$ the expected value $\sum_{x \in A} X(x) \Pr(x)$.

Theorem 2.6 (Chernoff). *Let A be a set of complex numbers such that $|x| \leq M$ for all $x \in A$. Let $x_i \in A$ be chosen independently and uniformly at randomly from A . Then*

$$\Pr \left[\left| \mathbb{E}_{x \in A} [x] - \frac{1}{m} \sum_{i=1}^m x_i \right| > \lambda \right] \leq 2e^{-\lambda^2 m / 2M^2}.$$

2.4 Table of notations

We summarize the main notation and definitions in the following table.

Notation/Definition	Meaning
ω_n	The complex n -th root of unity $e^{2\pi i/n}$.
χ	A character of G .
H^\perp	The orthogonal set $\{\alpha \in G \mid \chi_\alpha(h) = 1 \text{ for all } h \in H\}$.
\widehat{f}	The Fourier transform of f .
Scaling property	$\widehat{g}(\alpha) = \widehat{f}(c^{-1}\alpha)$ for $g(x) := f(cx)$ and $c \in R^*$.
τ -heavy coefficient	A coefficient satisfying $ \widehat{f}(\alpha) ^2 > \tau$.
Significant coefficient	A τ -heavy coefficient, for some $\tau^{-1} = \text{poly}(\log G , \ f\ _\infty)$.
Query access	The ability to ask for $f(x)$ for any input x .

3 Clarifications: Principles Underlying SFT Algorithms

In the last few decades several significant Fourier transform (SFT) algorithms were proposed in the literature in several scientific areas. The early algorithms treat specific functions, while the later algorithms apply to classes of functions. The principles underlying these algorithms come from elementary group theory. The aim of this section is to clarify the rules that govern these algorithms. Our analysis gives a unified presentation for all of these algorithms, which we believe brings clarity to the literature and will be more accessible to non-experts.

A precise statement of what an SFT algorithm does is given in Theorem 3.1. Section 3.1 gives an overview of the earlier algorithms. Section 3.2 presents the unified SFT algorithm in the query access model. The section starts with a high-level presentation of the SFT algorithm. We then describe the algorithm with a focus on the required algebraic relations between the queries, thus explaining the need for query access. For these relations to arise, the function’s domain needs to be “highly composite”, i.e. to contain many subgroups. We give examples of the requirement on the queries in some specific domains. We then turn to an analysis of the algorithm on domains of prime order. Moreover, the original approach that we present in Section 4 gives further insights on the connections between the different domains. We finish this section with two short descriptions. Section 3.3 discusses cases where some of the function’s outputs (i.e. the algorithm’s inputs) are “noisy”, that is where the actual values are replaced with some other values. Section 3.4 explains why an SFT algorithm in the random access model is unlikely to exist.

Theorem 3.1 ([1, SFT algorithm][3, Theorem 5]). *Let G be an abelian group represented by a set of generators of known orders. There is a learning algorithm that, given query access to a function $f : G \rightarrow \mathbb{C}$, a threshold $\tau > 0$ and $\delta > 0$, outputs a list L of size at most $2\|f\|_2^2/\tau$ such that*

- L contains all the τ -heavy Fourier coefficients of f with probability at least $1 - \delta$;
- L does not contain coefficients that are not $(\tau/2)$ -heavy with probability at least $1 - \delta$.

The algorithm runs in polynomial time in $\log(|G|)$, $\|f\|_\infty^2/\tau$ and $\log(\frac{1}{\delta})$.

3.1 History and special cases

Key ideas behind the SFT algorithm first arose in other settings, and the aim of this section is to put some of this early work in context. This section is not needed in order to understand the SFT algorithm. Readers who are mainly interested in understanding the general SFT algorithm should feel free to skip this section and go straight to Section 3.2.

3.1.1 Goldreich–Levin

Consider a ‘noisy’ inner product function $f_s : \{0, 1\}^n \rightarrow \{0, 1\}$ given by $f_s(x) = \langle x, s \rangle + \delta(x)$ (addition takes place mod 2) where $\delta(x) = 1$ with some small probability (noticeably smaller than $1/2$) and otherwise $\delta(x) = 0$. This is the same function as in the well-known learning parity with noise (LPN) problem. The task is to learn s given samples $f_s(x_i)$.

The connection to the Fourier basis can be seen by reformulating the problem as follows. Define $g : \{0, 1\}^n \rightarrow \{-1, 1\}$ by $g(x) = (-1)^{f_s(x)} = (-1)^{\langle x, s \rangle + \delta(x)}$. Notice that when $\delta \equiv 0$ then g is in fact the character $\chi_s(x) = (-1)^{\langle x, s \rangle}$. The fact that $\delta(x) = 0$ on most inputs guarantees that $\widehat{g}(s)$ is a significant Fourier coefficient for g , as shown in Example 2.5.

In the random access model, where one gets arbitrary samples, LPN is considered to be a hard computational problem (unless $\delta(x) = 0$ for all, or almost all, x ; then reconstructing f_s is an easy linear algebra problem). Goldreich and Levin [19] (GL) considered this problem in the query access model, and gave an efficient algorithm to solve it as we briefly explain. In the simplest setting there is a single τ -heavy coefficient for $\tau > 1/2$.

If one can choose the queries for f_s then an elementary approach is to query on the unit vectors $e_1 := (1, 0, \dots, 0), \dots, e_n := (0, \dots, 0, 1)$ to learn s bit-by-bit. However, since the query on e_i may return the answer $\langle e_i, s \rangle + 1$, one would like to generate a small set of independent values of the form $\langle e_i, s \rangle + \delta$, and determine s_i by majority rule, as $\delta = 0$ with probability noticeably greater than $1/2$. This can simply be achieved by querying on correlated values x and $x + e_i$ to get the results $\langle x, s \rangle + \delta(x)$ and $\langle x, s \rangle + \langle e_i, s \rangle + \delta(x + e_i)$. If both answers are not noisy (or if both are noisy) then by subtracting one from the other we get $\langle e_i, s \rangle$, which is the i -th coordinate of s . (For the interested reader: if the noise rate is at least $1/4$, then there may not be a unique solution (see Section 3.3); Rackoff (see [18, Section C.2]) suggested to use a trick due to Alexi et al. [4] to deal with this case.)

The original Goldreich–Levin paper [19] does not give a clear description of the learning algorithm. A description in the language of Fourier analysis was given in [26] by Kushilevitz and Mansour.

3.1.2 Bleichenbacher

Bleichenbacher [8] seems to have been the first to consider these problems in the case of functions on \mathbb{Z}_N where N is not a power of 2. He considers a ‘noisy’ product function $f_s : \mathbb{Z}_N \rightarrow \mathbb{Z}_N$ given by $f_s(x) = sx + \delta(x)$ where $|\delta(x)| < \frac{N}{2^\lambda}$, for some real number λ , with probability (noticeably) greater than $1/2$. This is usually viewed as outputting about λ most significant bits of the product $sx \in \mathbb{Z}_N$, as sx and f_s differ by a small number. The task, as before, is to learn s given samples $f_s(x_i)$. The connection to the Fourier basis can be seen by reformulating the problem as done in the previous GL case – see the ‘noisy character’ case in Example 2.5.

This problem is in fact the hidden number problem that was considered in [11] and which we further discuss in Section 5. Notice that if one can obtain any query, then this problem can be solved by successively multiplying by 2 to read the bits of s . Since some samples may be erroneous, majority rule is used, similar to the approach taken in the GL case. Moreover if $\delta(x)$ is very small, finding s and reconstructing f_s is easy (by ranging over all possible values for δ).

Bleichenbacher's original setting takes place in the random access model, so he gives a method (not efficient for large domains) to obtain samples $f_s(x)$ for which x lie in short intervals, and then gives a method to solve the original problem. We explain the latter method. Here however, one is not assumed to have any chosen query, but only that the queries lie in some (designated) intervals.

The main idea to solve this problem comes from the fact that one can use small (but gradually increasing) multipliers, not necessary powers of 2, to learn the bits of s . This comes from the following observation: if $s < \frac{N}{2^\eta}$, for some $\eta \geq 0$, then $sy < N$ for every $0 \leq y \leq 2^\eta$. In other words, the product sy does not 'wrap-around' the modulus N .

The latter observation can be used to determine upper bits of s : given y and $f_s(y) = sy + \delta(y)$, take $\lfloor f_s(y)/y \rfloor = \lfloor s + \delta(y)/y \rfloor$; assuming there is no wrap-around over N in $f_s(y)$, we get some upper bits of s . For example, if $2^{\eta-1} \leq y \leq 2^\eta$, then $|\delta(y)/y| < \frac{N}{2^{\eta+\lambda-1}}$ so we roughly learn $\lambda - 1$ of the upper bits of s that were not already known.

Now suppose one knows $\text{MSB}_\rho(s)$, the ρ most significant bits of s , then by subtracting it from s we have $s' := s - \text{MSB}_\rho(s) < \frac{N}{2^\rho}$. The goal now is to learn further (upper) bits of s' . One can define $f_{s'}(y)$ to be $f_s(y) - \text{MSB}_\rho(s)y = sy - \text{MSB}_\rho(s)y + \delta(y) = s'y + \delta(y)$. Thus, for appropriate multiplier y , say $2^{\rho-1} \leq y \leq 2^\rho$, we can determine more upper bits of s' as above. Repeating this procedure, one eventually learns all bits of s .

Notice that this approach requires having multipliers drawn from some interval $\{0, 1, \dots, 2^i - 1\}$ (specifically small multipliers in the first stages, which are the 'hardest' to get). Moreover, since it is not always the case that $|\delta(x)| < \frac{N}{2^x}$, we need to generate independent multipliers from these intervals. Similar to the approach in the GL case, this is done by fixing some z and querying on $z + r$ for r chosen uniformly in $\{0, 1, \dots, 2^i - 1\}$, then subtracting. Thus the queries have to be correlated such that their difference lies in the required interval.

This description presents the core ideas behind Bleichenbacher's algorithm in a manner similar to the description of the GL algorithm above. Bleichenbacher's description, which involves terminology from Fourier analysis, resembles the Kushilevitz–Mansour modification to the GL algorithm (see below) and the ideas described in Section 4. For the full details we refer to Bleichenbacher [8] (see also Section 6.1 below). This method does not seem to have been used for cryptographic applications until the recent works [13, 5].

3.1.3 Following work

The early work did not explicitly mention Fourier coefficients, but it was realised that one can re-phrase the problems as finding significant Fourier coefficients of related functions, as we show above. The Goldreich–Levin case was generalized by Kushilevitz and Mansour [26] (KM) to any real-valued function over $\{0, 1\}^n$ and this work was the first to explicitly treat functions with more than one significant Fourier coefficient.

Subsequently, Mansour [33] gave an algorithm for functions $f : \mathbb{Z}_{2^n} \rightarrow \mathbb{C}$. Unlike other works, Mansour’s algorithms computes the significant coefficients from the least significant bit to the most significant bit (a link between these works [26, 33] is explained in Remark 3.3 below). The approach of Mansour was extended, thereby giving a generalisation of Bleichenbacher’s result, by Akavia, Goldwasser and Safra [3] (AGS).

Notice that combining the KM and AGS ideas gives an algorithm for all groups $\mathbb{Z}_{N_1} \times \cdots \times \mathbb{Z}_{N_r}$, since one can easily collapse from the latter to \mathbb{Z}_{N_j} (by choosing appropriate queries, for example queries of the form $r \cdot e_j$ for desired values $r \in \mathbb{Z}_{N_j}$). Therefore, the case of most interest is $G = \mathbb{Z}_p$ which we present below. As further evidence for the unity of all these ideas we remark that the KM and AGS algorithms query on exactly the same set of queries as GL and Bleichenbacher (and subsequently reveal the significant coefficients bit-by-bit from MSB to the LSB).

3.2 The SFT algorithm

Let $f : G \rightarrow \mathbb{C}$. Given a threshold $\tau \in \mathbb{R}$, the algorithm outputs all τ -heavy Fourier coefficients of f (and potentially some other $\tau/2$ -heavy coefficients) with overwhelming probability.

We first give a high-level view of how the algorithm works. The method is a form of binary search: the algorithm divides the set of Fourier coefficients into two (disjoint) sets, say A and B , and checks each set separately to determine whether it potentially contains a τ -heavy coefficient. To do this the algorithm defines two new functions, one for each set of coefficients. A clever use of Parseval’s identity allows the algorithm to check the size of all coefficients in each set simultaneously, given the norm of each function. Hence, the task is to determine the norms of the two new functions, which requires a method to compute the function outputs. The structure of the sets A, B is important: for some sets we have useful formulas to compute the functions at required values. Instead of precisely calculating these values, it is sufficient to have approximations of the outputs of the functions and to approximate the norm of each function. The Chernoff bound is then used to bound the error term in the approximations.

Schematically, the algorithm operates as follows, where we initially take $D = G$:

- Partition $D = A \cup B$, and define $f_A(x) := \sum_{\alpha \in A} \widehat{f}(\alpha) \chi_\alpha(x)$ and $f_B(x) := \sum_{\beta \in B} \widehat{f}(\beta) \chi_\beta(x)$.
- Approximate the values $f_A(x_i)$ and $f_B(y_j)$ for polynomially many samples x_i, y_j , chosen uniformly at random. This is done using the fundamental relation in (3.1) below.
- Using the values from the previous step, approximate the norms $\|f_A\|_2^2$ and $\|f_B\|_2^2$. See (3.3).
- Using Parseval’s identity $\|f_A\|_2^2 = \sum_{\alpha \in A} |\widehat{f}(\alpha)|^2$, if the approximation of the norm is smaller than $\frac{3}{4}\tau$ then with overwhelming probability f does not have a τ -heavy coefficient in A . Hence, dismiss A . Act similarly for f_B .
- Run the algorithm recursively on the remaining sets and stop when it reaches singletons.

Remark 3.2. We emphasize that the algorithm can work with any function f and with any threshold τ . Specifically, if f does not have any τ -heavy coefficients, then the algorithm will output an empty list.

²A lower threshold $\frac{3}{4}\tau$ is needed since the algorithm only approximates the norm. As a consequence, the final list may contain coefficients that are $\frac{\tau}{2}$ -heavy but not τ -heavy.

However, the running time is polynomial in $\|f\|_\infty^2/\tau$ so the algorithm will not be efficient if the threshold is chosen to be too low.

3.2.1 Domains of size 2^n

We now sketch an algorithm that unifies the KM and Mansour algorithms. Our presentation is more group-theoretic than the original works. We refer to [26] and [34] for exact details and proofs.

Let $f : G \rightarrow \mathbb{C}$ and $\tau \in \mathbb{R}$. At each iteration the algorithm takes a set D (starting with $D = G$) and proceeds as follows.

Partial functions. Partition $D = A \cup B$ into two sets that are defined below. Define the function $f_A : G \rightarrow \mathbb{C}$ by $f_A(x) = \sum_{\alpha \in A} \widehat{f}(\alpha) \chi_\alpha(x)$. If f has a τ -heavy coefficient α and $\alpha \in A$, then f_A has a τ -heavy coefficient. All arguments hold similarly for the set B .

Estimating $f_A(x)$. We need a method to estimate values of the function f_A using values of the original function f . We define a *filter function* $h_A : G \rightarrow \mathbb{C}$ by $h(x) = \sum_{\alpha \in A} \chi_\alpha(x)$, and then use the property $\widehat{f * h_A} = \widehat{f} \cdot \widehat{h_A}$. Since

$$\widehat{h_A}(\alpha) = \begin{cases} 1 & \alpha \in A, \\ 0 & \text{otherwise,} \end{cases}$$

we have

$$\widehat{f * h_A}(\alpha) = \begin{cases} \widehat{f}(\alpha) & \alpha \in A, \\ 0 & \text{otherwise.} \end{cases}$$

In other words,

$$f * h_A = f_A. \quad (3.1)$$

Convolution is not a task we have an efficient method to calculate in general, let alone efficiently calculating $h_A(x) = \sum_{\alpha \in A} \chi_\alpha(x)$. Therefore, the structure of the sets is important and plays a key role in the ability to apply the algorithm. Notice that if A is an arithmetic progression, then $\sum_{\alpha \in A} \chi_\alpha(x) = \sum_j \chi_{qj+r}(x) = \chi_r(x) \sum_j \chi_q(jx)$, and so it can be evaluated by the formula for geometric series. More generally, assume $D \leq G$ is a subgroup and let $H \leq D$ be a subgroup (of index 2). We take A to be a coset $A = z + H$ for some $z \in G$ (then B is taken to be the other coset). Then,

$$h_{z+H}(x) = \sum_{h \in H} \chi_{z+h}(x) = \sum_{h \in H} \chi_z(x) \chi_h(x) = \chi_z(x) \sum_{h \in H} \chi_h(x),$$

and the latter is zero unless $x \in H^\perp$ (H^\perp is defined in (2.1) above). Thus the function h_A is given by

$$h_A(x) = h_{z+H}(x) = \begin{cases} \chi_z(x) \cdot |H|, & \text{if } x \in H^\perp, \\ 0, & \text{otherwise.} \end{cases} \quad (3.2)$$

We therefore get, since $|H||H^\perp| = |G|$,

$$\begin{aligned} f_A(x) &= f * h_A(x) = \mathbb{E}_{y \in G} [f(x-y)h_A(y)] = \frac{1}{|G|} \sum_{y \in G} f(x-y)h_A(y) \\ &= \frac{1}{|G|} |H| \sum_{y \in H^\perp} f(x-y)\chi_z(y) = \mathbb{E}_{y \in H^\perp} [f(x-y)\chi_z(y)]. \end{aligned}$$

Estimating $\|f_A\|_2$. We can now write $\|f_A\|_2^2$ as

$$\|f_A\|_2^2 = \mathbb{E}_{x \in G} |(f * h_A)(x)|^2 = \mathbb{E}_{x \in G} \left| \mathbb{E}_{y \in G} [f(x-y)h_A(y)] \right|^2 = \mathbb{E}_{x \in G} \left| \mathbb{E}_{y \in H^\perp} [f(x-y)\chi_z(y)] \right|^2.$$

Again, an approximation of the norm is sufficient (a consequence of the approximation is that we have to lower the threshold τ a little bit).

We can therefore approximate $\|f_A\|_2^2$ by choosing m_1, m_2 sufficiently large (given by the Chernoff bound), randomly choosing³ $x_i \in G$ where $1 \leq i \leq m_1$, randomly choosing $y_{ij} \in H^\perp$ for each i where $1 \leq j \leq m_2$ and calculating

$$\frac{1}{m_1} \sum_{i=1}^{m_1} \left| \frac{1}{m_2} \sum_{j=1}^{m_2} f(x_i - y_{ij}) \chi_z(y_{ij}) \right|^2 \approx \|f_A\|_2^2 = \sum_{\alpha \in A} |\widehat{f}(\alpha)|^2. \quad (3.3)$$

One then checks if this value is smaller than $3\tau/4$. If so then with overwhelming probability there is no $\alpha \in A$ such that $\widehat{f}(\alpha)$ is τ -heavy, and so the set A can be dismissed. Notice that if this value is greater than $3\tau/4$ it does not necessarily mean that A contains a significant coefficient. In this case the algorithm sets $D = A$ and repeats until all sets are singletons or dismissed.

We give the pseudocode of the algorithm. At start, set $z = 0$ and $k = n$, so $H_k = G$.

Algorithm 1: MainProcedure

Input: A coset $z + H_k$.

if $|H_k| = 1$ **then**

if $|\text{Est} \widehat{f}(z)|^2 \geq 3\tau/4$ **then**

return $\{z\}$

else

return \emptyset

else

 Let W be a set of coset representatives for H_{k-1} in H_k

 Let $W' = \{w \in W \mid \text{EstNormSq}(f_{(z+w)+H_{k-1}}) \geq 3\tau/4\}$

return $\cup_{w \in W'} \text{MainProcedure}((z+w) + H_{k-1})$

Algorithm 2: EstNormSq

Input: $f_{z+H} : G \rightarrow \mathbb{C}$.

Choose $x_i \in G$ where $1 \leq i \leq m_1$

For each i , choose $y_{ij} \in H^\perp$ where $1 \leq j \leq m_2$

return $\frac{1}{m_1} \sum_{i=1}^{m_1} \left| \frac{1}{m_2} \sum_{j=1}^{m_2} f(x_i - y_{ij}) \chi_z(y_{ij}) \right|^2$

³Note that as in [26, 33] one can define the function f_A over H (and not G), and therefore choose the values x_i from H .

Algorithm 3: $\widehat{\text{Est}f}$

Input: $z \in G$.

 Choose $x_i \in G$ where $1 \leq i \leq m_1$
return $\frac{1}{m_1} \sum_{i=1}^{m_1} f(x_i) \chi_z(-x_i)$

3.2.2 Examples

Notice that in (3.3) above for each x_i one needs the samples $f(x_i - y_{ij})$. This explains the importance of having query access to the function. To illustrate this point, we give some concrete examples.

Kushilevitz and Mansour [26] consider a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$. Write $x = x_1 \dots x_n$. At the first iteration define A to contain all n -bit strings that start with 0 and B to contain all the n -bit strings that start with 1. Then we have

$$h_A(x) = \begin{cases} 2^{n-1}, & \text{if } x = 0 \dots 0 \text{ or } x = 10 \dots 0, \\ 0, & \text{otherwise,} \end{cases} \quad (3.4)$$

and indeed

$$\widehat{h}_A(\alpha) = \frac{1}{2^n} \sum_x h_A(x) (-1)^{\langle \alpha, x \rangle} = \frac{1}{2} \left((-1)^0 + (-1)^{\alpha_1} \right) = \begin{cases} 1 & \alpha \in A; \\ 0 & \text{otherwise.} \end{cases}$$

One can only evaluate $f * h_A(x)$ if one has the values $f(x)$ and $f(x + e_1)$. This shows that the KM approach requires (in the first iteration) queries on pairs of vectors that differ by a unit vector, exactly as in the elementary approach to the GL theorem as sketched in Section 3.1.1.

Mansour [33] considers a function $f : \mathbb{Z}_{2^n} \rightarrow \mathbb{C}$. At the first iteration define A to contain all the even numbers in \mathbb{Z}_{2^n} and B to contain all the odd numbers. Then, we have

$$h_A(x) = \begin{cases} 2^{n-1}, & \text{if } x = 0 \text{ or } x = 2^{n-1}, \\ 0, & \text{otherwise,} \end{cases} \quad (3.5)$$

and indeed

$$\widehat{h}_A(\alpha) = \frac{1}{2^n} \sum_x h_A(x) \omega_{2^n}^{\alpha x} = \frac{1}{2} (1 + (-1)^\alpha) = \begin{cases} 1 & \alpha \in A; \\ 0 & \text{otherwise.} \end{cases}$$

One can only evaluate $f * h_A(x)$ if one has $f(x)$ and $f(x + 2^{n-1})$.

The analysis of this algorithm is useful for the prime case below, and so we present its later stages. In stage l of this algorithm, one defines the subgroup H to contain all multiples of 2^l in \mathbb{Z}_{2^n} . Hence the cosets used to partition the solution space contain all numbers that agree on their remainder modulo 2^l , and $H^\perp = \{x \in \mathbb{Z}_{2^n} \mid x2^l \equiv 0 \pmod{2^n}\} = \{0, 2^{n-l}, 2 \cdot 2^{n-l}, 3 \cdot 2^{n-l}, \dots, (2^l - 1)2^{n-l}\}$. Define $A = A_r = \{x \in \mathbb{Z}_{2^n} \mid x \equiv r \pmod{2^l}\} = H + r$. Then, the filter function h_A satisfies

$$h_A(x) = \begin{cases} \chi_r(x) \cdot 2^{n-l}, & \text{if } x \in H^\perp, \\ 0, & \text{otherwise.} \end{cases} \quad (3.6)$$

Again, to approximate $f * h_A(x)$, one needs enough samples $f(x_i)$ for $x_i \in H^\perp$.

Remark 3.3. Readers familiar with lattice cryptography may be interested to know that the idea that underlies the modulus-dimension tradeoff [29] already appears in the relationship between the KM [26] algorithm on $\{0, 1\}^n$ and the Mansour [33] algorithm on \mathbb{Z}_{2^n} . We briefly sketch this idea. Let $\mathbf{a} = (a_0, \dots, a_{n-1}) \in \mathbb{Z}_p^n$, $\mathbf{s} = (s_0, \dots, s_{n-1}) \in \{0, 1\}^n$, and suppose

$$b \equiv \mathbf{a} \cdot \mathbf{s} + e \equiv \sum_{i=0}^{n-1} a_i s_i + e \pmod{p}.$$

Writing $a = a_0 p^{n-1} + a_1 p^{n-2} + \dots + a_{n-2} p + a_{n-1}$ and $s = s_0 + s_1 p + \dots + s_{n-1} p^{n-1}$ we have

$$as \equiv (a_0 s_0 + \dots + a_{n-1} s_{n-1}) p^{n-1} + \text{lower term} \pmod{p^n}$$

and some of its MSBs agree with the MSBs of bp^{n-1} , when p is large.

As shown in equation (3.4) above, at the first iteration over $\{0, 1\}^n$ the filter function is nonzero on the inputs $\mathbf{0}$ and $\mathbf{a} = (1, 0, \dots, 0)$ in \mathbb{Z}_2^n . These vectors correspond to the values $a = 0$ and $a = 2^{n-1}$ in \mathbb{Z}_{2^n} , which are exactly the values appearing in equation (3.5). Since the lower terms of $a \cdot s$ are zero, when $a = 0, p^{n-1}$, the MSB of as and bp^{n-1} agree even for $p = 2$. In both domains, we use these values to recover s_0 . The generalization to all inputs \mathbf{a} arising in the algorithms is straightforward.

3.2.3 Domains of prime order

The ideas behind the algorithm presented above make use of the fact that the domain's order can be factored as a product of small primes (especially for powers of 2, as been shown for $\{0, 1\}^n$ in [26] and for \mathbb{Z}_{2^n} in [33]). A case of interest, from the theoretical and practical sides, is domains of (large) prime order. Notice that each additive group \mathbb{Z}_N can be decomposed into a direct product of prime subgroups $\mathbb{Z}_{p_1} \times \dots \times \mathbb{Z}_{p_n}$. The query access allows us to work over each subgroup separately, to recover the coefficients prime-by-prime, similar to the bit-by-bit approach in the GL case above (Section 3.1.1). Indeed one can query on $\mathbf{x}_j = (0, \dots, 0, x_j, 0, \dots, 0)$ to work over the group \mathbb{Z}_{p_j} .⁴ Thus being able to find heavy coefficients for functions over a prime group \mathbb{Z}_p will allow us to find heavy coefficients for functions over any \mathbb{Z}_N .

For prime groups the analysis we presented for the algorithm above does not apply as \mathbb{Z}_p does not have any proper subgroups, specifically not those of small index. The importance of the subgroups is in the evaluation of exponential sums (such as equation (2.2) above), which subsequently allows us to have useful formulas for the filter functions (such as equation (3.2)). We now show that one can still follow the steps in the algorithm above. Natural candidates for the partitioning sets are intervals (of similar size) of consecutive numbers or classes of numbers with the same remainder modulo 2^l (where l represents the stage we work at), which is similar to the approach taken over \mathbb{Z}_{2^n} (see Section 3.2.2).⁵ In fact, using the frequency-shifting and scaling properties of the Fourier transform, one can show that these two partitions are equivalent (where there is a correspondence between the size of the intervals and the size of the classes), in the sense that one can transform the coefficients in an interval to coefficients of the same class modulo 2^l and vice versa. We show this equivalence below.

⁴Since deterministic queries are not desirable, additional randomization is used in practice.

⁵Note that both are arithmetic progressions, which allow evaluating h_A .

The algorithm over \mathbb{Z}_p [3] works in the same steps as explained in Section 3.2. The main obstacle is to show how to efficiently calculate the function f_A , for some appropriate set A . We therefore focus on this step. The other steps are similar to the algorithm for domains of size 2^n .

Working in the ‘frequency domain’. In order to show the difficulty working in a domain of prime size, we start with a naive imitation of the approach taken in the algorithm for domains of size 2^n . Let A be an arithmetic progression in \mathbb{Z}_p , and define $f_A = \sum_{\alpha \in A} \widehat{f}(\alpha) \chi_\alpha(x)$ and $h_A(x) = \sum_{\alpha \in A} \chi_\alpha(x)$ as above. Then $f_A(x) = f * h_A(x) = \mathbb{E}_{y \in G} [f(x-y)h_A(y)] = \mathbb{E}_{y \in G} [f(x-y) \sum_{\alpha \in A} \chi_\alpha(y)]$. Since A is an arithmetic progression, $\sum_{\alpha \in A} \chi_\alpha(x)$ is a geometric progression for which we have a formula. We get that $f_A(x)$ is an expectation over values each of which we can calculate exactly. Moreover, unlike in the algorithm above, the filter function here is nonzero over a very large set, and therefore one can hope that specific queries are not needed in this case (as shown in Section 3.2.2 the previous filter functions are zero almost everywhere, so in order to get a good approximation of $f_A(x)$ we need the specific inputs where the filter function is not zero). This turns out to be a disadvantage. Indeed, in order to determine $f_A(x)$ in polynomial time, we can only approximate this expectation, but as the values of this geometric progression can be as large as $|A|$, one derives from the Chernoff bound that the number of samples needed to have a good approximation of $f_A(x)$ is roughly $|A|$, which is exponential in $\log(p)$ in the first stages of the algorithm. Hence this approach is not practical.

Working in the ‘time domain’. Instead of working in the ‘frequency domain’, we can work in the ‘time domain’. In this case we define A to be a class of numbers with the same remainder mod 2^l . We adapt the filter function in (3.6) to the \mathbb{Z}_p case. As in Section 3.2.2, let H be the set containing all multiples of 2^l in \mathbb{Z}_p . Define $H^\perp := \{0, 2^{-l}, 2 \cdot 2^{-l}, \dots, (2^l - 1)2^{-l}\}$. Notice that while H^\perp is not orthogonal to H , it contains all numbers that give small remainder (mod p) when multiplied by 2^l . Let $z \in \mathbb{Z}_p$ such that $z \equiv r \pmod{2^l}$ and define $A = A_r = \{x \in \mathbb{Z}_p \mid x \equiv r \pmod{2^l}\}$ to be the class in \mathbb{Z}_p for which the remainder mod 2^l is r . We define

$$h_A(x) = h_{z+H}(x) = \begin{cases} \frac{p}{2^l} \chi_z(x), & \text{if } x \in H^\perp, \\ 0, & \text{otherwise.} \end{cases}$$

It turns out that this function, which is a simple adaptation of (3.6) to \mathbb{Z}_p , is a ‘noisy’ version of a ‘pure’ filter function: the size of the coefficients $|\widehat{h}_A(\alpha)|$ is close to 1 for $\alpha \in A$ and close to 0 for $\alpha \notin A$. Indeed,

$$\widehat{h}_A(\alpha) = \frac{1}{p} \sum_{x \in \mathbb{Z}_p} h_A(x) \overline{\chi_\alpha(x)} = \frac{1}{2^l} \sum_{x \in H^\perp} \chi_{z-\alpha}(x).$$

Write $\alpha = 2^l k + j$, $z = 2^l q + r$ and $x = d2^{-l}$ for $0 \leq j, r < 2^l$ and $0 \leq d \leq \lfloor \frac{p}{2^l} \rfloor$. Then,

$$\widehat{h}_A(\alpha) = \frac{1}{2^l} \sum_{0 \leq d \leq \lfloor \frac{p}{2^l} \rfloor} \chi_{2^l q + r - 2^l k - j}(d2^{-l}) = \frac{1}{2^l} \sum_{0 \leq d \leq \lfloor \frac{p}{2^l} \rfloor} \chi_{q-k}(d) \chi_{r-j}(2^{-l} d).$$

One can show that the last sum is large if and only if $j = r$ as $\chi_{r-j} \equiv 1$, that is if and only if $\alpha \in A$, and so that $|\widehat{h}_A(\alpha)| \approx 1$, and otherwise it is close to 0. More precisely, for $\alpha = z$ we have $|\widehat{h}_A(\alpha)| = 1$ and as k gets further away from q , the size of $\widehat{h}_A(2^l k + r)$ slowly decays (follows from Claim 4.1 below). The function h_A is said to be “centered around” z . The results in Section 4 below give further insights for the reasons why this adaptation of the filter function from \mathbb{Z}_{2^n} to \mathbb{Z}_p in the time domain, only slightly affects its frequency domain.

The work of AGS. The approach taken in [3, 1] is to work over intervals. We show how, using the scaling and frequency-shifting properties, one can transform from the set A to an interval I of the same size. Define $h_I(x) := h_A(2^{-l}x)$, then $\widehat{h}_I(\alpha) = \widehat{h}_A(2^l\alpha)$. This is a permutation of the coefficients of h_A . If $A = \{r, 2^l + r, \dots, t2^l + r\}$, then $I = \{r2^{-l}, r2^{-l} + 1, \dots, r2^{-l} + t\}$, and the coefficients which were large on A and small outside A are now large over I and small outside it. Moreover, if we define $h_I(x) := h_A(2^{-l}x)\chi_c(x)$ then by the shifting property the previous interval I shifts to $I - c$.

AGS consider an interval $[a, b]$ of size $\lfloor \frac{p}{2^l} \rfloor$, for which $c = \lfloor \frac{a+b}{2} \rfloor$ is a middle point. They then define

$$h_{a,b}(x) = \begin{cases} \frac{p}{2^l}\chi_c(x), & \text{if } 0 \leq x < 2^l, \\ 0, & \text{otherwise.} \end{cases}$$

A direct calculation using the definition of $\widehat{h}_{a,b}(\alpha)$ shows that

$$\widehat{h}_{a,b}(\alpha) = \mathbb{E}_{0 \leq x < 2^l} [\chi_c(x)\overline{\chi_\alpha(x)}] = \mathbb{E}_{0 \leq x < 2^l} [\chi_{c-\alpha}(x)].$$

Again, one can show that $|\widehat{h}_{a,b}(\alpha)| \approx 1$ if $a \leq \alpha \leq b$ and $|\widehat{h}_{a,b}(\alpha)| \approx 0$ for α outside this interval (see, for example, Claim 4.1). For further details see [3, 1]. This function is “centered around” c , that is, for $\alpha = c$ we have $|\widehat{h}_A(\alpha)| = 1$ and while α gets further away from c , the size of $\widehat{h}_A(\alpha)$ slowly decays.

Remark 3.4. There is a technical issue which we ignore in this description. As the size of $\widehat{h}_A(\alpha)$ slowly decays while α moves away from c , when α reaches the end of the interval $[a, b]$ the value $|\widehat{h}_A(\alpha)|$ is close to the value $|\widehat{h}_A(\beta)|$ for β just outside this interval. This imposes some complexities in the filtering process; specifically one should take overlapping intervals, so the sets A, B are not distinct as in the case of domains of size 2^n . Moreover, the choice of the point c (therefore the choice of the interval) also affects the filtering process. We refer to Sections 7.2.3 and 7.2.4 in [3] and to [1, Section 3] for the technical details.

With this filter function (either h_A or $h_{a,b}$) f_A can be approximated efficiently, as shown in the previous section. The algorithm now proceeds as the algorithm for domains of size 2^n .

3.3 Working with unreliable oracles

It is sometimes desirable to describe access to the function f as querying an oracle. The oracle can be perfect – always provides the correct value $f(x)$ – or imperfect. Working with unreliable oracles is of importance in several applications. This section is dedicated to analyzing these cases.

Sometimes the samples $f(x_i)$ are given by an unreliable oracle O . By this we mean the oracle satisfies $O(x) = f(x)$ only with high probability. One can think of O as a ‘noisy version’ of f . A common approach to this situation is to generate several independent values, each of which gives the value $f(x)$ with good probability; then, by applying majority rule, one can obtain the correct value $f(x)$ with overwhelming probability. Examples of this approach are presented in Section 3.1.

We show how the language of Fourier analysis gives a very general approach to analyze situations for working with unreliable oracles. The main idea is that if a function f has a significant Fourier coefficient,

then its noisy version also has a significant coefficient. Note however that if f is concentrated, then its ‘noisy’ version is not necessarily concentrated.

To be precise, let $f : G \rightarrow \mathbb{C}$. We describe the oracle as a function $O : G \rightarrow \mathbb{C}$ such that $O(x) = f(x)$ on the majority of $x \in G$. We assume that $\|O\|_\infty \leq \|f\|_\infty$. Define $R : G \rightarrow \mathbb{C}$ by $R(x) = O(x) - f(x)$ and let $I = \{x \in G : R(x) \neq 0\}$. We want to show that if $\widehat{f}(\alpha)$ is τ -heavy, then $\widehat{O}(\alpha)$ is τ' -heavy, for some τ' relatively large (its precise size depends on the success rate of the oracle).

Since $O = f + R$, then $\widehat{O}(\alpha) = \widehat{f}(\alpha) + \widehat{R}(\alpha)$. Note that $\|R\|_\infty \leq 2\|f\|_\infty$. Hence

$$\left| \widehat{O}(\alpha) \right| \geq \left| \widehat{f}(\alpha) \right| - \left| \frac{1}{|G|} \sum_{x \in I} R(x) \overline{\chi_\alpha(x)} \right| \geq \left| \widehat{f}(\alpha) \right| - \frac{2|I|}{|G|} \|f\|_\infty.$$

As I is small, if $\widehat{f}(\alpha)$ is significant then so is $\widehat{O}(\alpha)$. Note that as the reliability rate of the oracle decreases, so does the size of $\widehat{O}(\alpha)$, while other coefficients increase in size. One can see that, similarly to majority rule, more samples are needed when the reliability rate of the oracle decreases. Indeed, the number of samples is proportional to τ^{-1} and as the size of the threshold τ decreases, τ^{-1} increases.

It is well-known that the GL theorem finds the unique function in case of low noise rate, namely if the the noise rate is smaller than $\frac{1}{4} - \varepsilon$. One immediately sees this from our analysis: the original function satisfies $|\widehat{f}(s)| = 1$, for the secret vector s , and so only one Fourier coefficient of O is larger than $\frac{1}{2}$.

3.4 Hardness of finding significant coefficients in the random access model

The SFT algorithm requires chosen queries. The aim of this section is to explain that one does not expect a general learning algorithm for problems where the function values cannot be chosen. Indeed, we will show that if such a learning algorithm existed then the learning parity with noise (LPN) and learning with errors (LWE) problems would be easy.

Recall the LPN problem: an instance is a list of samples $(a, b = \langle a, s \rangle + e(a)) \in \mathbb{Z}_2^n \times \mathbb{Z}_2$ for some secret value s and a function $e : \{0, 1\}^n \rightarrow \{0, 1\}$ which determines the noise. Define LPN : $\{0, 1\}^n \rightarrow \{0, 1\}$ by $\text{LPN}(a) := (-1)^b$. This is a ‘noisy version’ of the function $f(x) := (-1)^{\langle a, s \rangle}$ for which $\widehat{f}(s)$ is the only non-zero Fourier coefficient. For a small noise rate (as in LPN), as shown in Section 3.3, the coefficient $\widehat{\text{LPN}}(s)$ is a significant coefficient for this function. Hence, if one could find significant coefficients in $\{0, 1\}^n$ on random samples, then one could solve LPN given the samples (a, b) . Since LPN is believed to be hard, one does not expect such a variant of the SFT algorithm to exist. Further evidence for the hardness of this problem in the random access model is that it is related to the problem of decoding a random binary linear code.

The same argument holds for LWE in \mathbb{Z}_p^n . In LWE one has samples $a \in \mathbb{Z}_p^n$ and $b = \langle a, s \rangle + e(a) \pmod{p}$ where $e(a)$ is “small” relative to p . Defining $\text{LWE}(a) := \omega_p^b$ one can show that the coefficient of the character $\chi_s(x) = \omega_p^{\langle x, s \rangle}$ is significant. Hence, if one could find the significant coefficients when given random samples, then one could solve LWE given the samples (a, b) . Since we have good evidence that LWE is a hard problem, this shows that we do not expect to be able to learn significant Fourier coefficients in the random access model.

The modulus-dimension tradeoff for LWE [29] shows how to transform LWE in \mathbb{Z}_p^n to LWE in $\mathbb{Z}_{p^d}^{n/d}$ (albeit with a different error distribution), and so one can conclude that finding significant coefficients in

\mathbb{Z}_p^n on random samples is at least as hard as solving LWE in \mathbb{Z}_p^n with binary secrets. This is an example of the connection between \mathbb{Z}_2^n and \mathbb{Z}_{2^n} as explained in Remark 3.3.

4 Simplifications: Modulus Switching

The SFT algorithm is considerably simpler to understand and implement for \mathbb{Z}_2^n or \mathbb{Z}_{2^n} than for \mathbb{Z}_p . Furthermore, for domains of size 2^n , considerable effort has been invested by researchers in the engineering community into making this algorithm more efficient with respect to various measures [17] (see also Mansour and Sahar [35]). Hence, it is natural to try to work with functions over \mathbb{Z}_{2^n} instead of functions over \mathbb{Z}_p . We now sketch an approach that shows how one can transform functions on \mathbb{Z}_p into functions on \mathbb{Z}_{2^n} where $2^n \approx p$, while maintaining a relation between their significant coefficients. In analogy to similar ideas in lattice cryptography we call this “modulus switching”.

These ideas are implicit in the work of Shor [42] on factoring with quantum computers. Shor extends a periodic function to a larger domain. The core idea is that if a function is periodic, then the period, which is a feature of the time domain, is preserved over any (large enough) domain. This fact is exploited by Shor, where his further ideas take place in the frequency domain. Shor’s analysis provides a clear interaction between the representation of a (periodic) function in the time and frequency domains.

We extend these ideas to show that a much larger class of functions keeps the properties of their frequency domain representation, when extending their time domain. Specifically, significant coefficients are “preserved” even when the time domain representation of the function is extended (by “preserved” we mean that there is a clear relation between the significant coefficients of both functions). We refer to Laity and Shani [28] for the technical details.

Let $N = 2^n > p$ be the smallest power of two greater than p . For a function $f : \mathbb{Z}_p \rightarrow \mathbb{C}$, we define

$$\tilde{f}(x) := \begin{cases} f(x) & \text{when } 0 \leq x < p, \\ 0 & \text{when } p \leq x < N. \end{cases}$$

Note that the operation $f \mapsto \tilde{f}$ is \mathbb{C} -linear. The basic observation (see Figure 1) is that for a character χ_α on \mathbb{Z}_p , $\tilde{\chi}_\alpha(x)$ is a function on \mathbb{Z}_N that is also concentrated.

To explain this observation we state the following basic fact and sketch a proof of it. It is straightforward to turn this result into a rigorous upper bound.

Claim 4.1. *Let $N > 1$, $\omega_N = e^{\frac{2\pi i}{N}}$ and let $\alpha \in \mathbb{R}$, $\alpha \neq 0$, $|\alpha| < N/2$ and $K \in \mathbb{N}$. Define*

$$S_{\alpha,K} = \sum_{x=0}^{K-1} \omega_N^{\alpha x}.$$

Then

$$|S_{\alpha,K}| \approx N \frac{|1 - \omega_N^{\alpha K}|}{2\pi|\alpha|}.$$

To see this note that the geometric series sums to $(1 - \omega_N^{\alpha K}) / (1 - \omega_N^\alpha)$ and the denominator is $(1 - \cos(2\pi\alpha/N)) - i \sin(2\pi\alpha/N)$ which has norm squared equal to $2(1 - \cos(2\pi\alpha/N))$. Finally, since $(1 - \cos(x)) \approx x^2/2$ (indeed $\frac{x^2}{2}(1 - \frac{x^2}{12}) \leq 1 - \cos(x) \leq \frac{x^2}{2}$), the result follows.




Figure 1: The magnitude of the Fourier coefficients $\widehat{\chi}_\alpha(z)$. Here $p = 37$, $N = 64$ and $\alpha = 5$.

We now compute the Fourier transform of $\widetilde{\chi}_\alpha$ as a function on \mathbb{Z}_N where $N = 2^n$. We have

$$\widehat{\widetilde{\chi}}_\alpha(\beta) = \langle \widetilde{\chi}_\alpha, \chi_\beta \rangle = \frac{1}{N} \sum_{x=0}^{p-1} \exp\left(2\pi i \left(\frac{\alpha}{p} - \frac{\beta}{N}\right)x\right).$$

If $\frac{\alpha}{p} - \frac{\beta}{N} \neq 0$, which will be satisfied in general since $\alpha, \beta \in \mathbb{Z}$ while $\gcd(p, N) = 1$, then applying Claim 4.1 gives the approximation

$$\left| \widehat{\widetilde{\chi}}_\alpha(\beta) \right| \approx \frac{|1 - \exp(2\pi i(\alpha/p - \beta/N))|}{2\pi|\alpha/p - \beta/N|}.$$

If $\beta \approx N\alpha/p$ then this coefficient is large and so the function $\widetilde{\chi}_\alpha$ has a significant Fourier coefficient at $\lfloor N\alpha/p \rfloor$. Moreover, the size of $\widehat{\widetilde{\chi}}_\alpha(\lfloor N\alpha/p + k \rfloor)$, for $0 < |k| < N/2$, is bounded by $O(1/k)$, and so $\widetilde{\chi}_\alpha$ is concentrated in a small set $\Gamma \subseteq \mathbb{Z}_N$ of characters represented by values around $N\alpha/p$.

Since the maps $f \mapsto \widehat{f}$ and $g \mapsto \widehat{g}$ are \mathbb{C} -linear, for any $f(x) = \sum_{\alpha \in G} \widehat{f}(\alpha) \chi_\alpha(x)$ we have

$$\widehat{\widehat{f}}(\beta) = \sum_{\alpha=0}^{p-1} \widehat{f}(\alpha) \widehat{\widetilde{\chi}}_\alpha(\beta).$$

Thus, if $\widehat{f}(\alpha)$ is a significant coefficient for f , then one expects that for $\beta = \lfloor N\alpha/p \rfloor$, the coefficient $\widehat{\widehat{f}}(\beta)$ is significant for \widehat{f} . The work of Laity and Shani [28] made these arguments to a precise theorem.

Theorem 4.2 ([28, Theorem 1.1]). *Let $\{n_k\}_{k \in \mathbb{N}}, \{m_k\}_{k \in \mathbb{N}}$ two sequences of positive integers with $m_k \geq n_k/2$ for every $k \in \mathbb{N}$. Let $Q \in \mathbb{R}[x]$ be a polynomial. Let $\{f_k : \mathbb{Z}_{n_k} \rightarrow \mathbb{C}\}_{k \in \mathbb{N}}$ be a concentrated family of functions such that $\|f_k\|_2^2 \leq Q(\log(n_k))$ for all $k \in \mathbb{N}$. Then $\{f_k : \mathbb{Z}_{m_k} \rightarrow \mathbb{C}\}_{k \in \mathbb{N}}$ is a concentrated family of functions.*

Specifically, if $f(x)$ is a concentrated function on \mathbb{Z}_p then $\tilde{f}(x)$ is a concentrated function on \mathbb{Z}_{2^n} . A similar result holds where f is ε -concentrated. We refer to [28] for the technical details.

As a consequence, one sees that it is not necessary to develop a variant of the SFT algorithm for the group \mathbb{Z}_p . Instead one can simply modulus-switch to a power of two and apply the SFT algorithm for the group \mathbb{Z}_{2^n} . This is addressed in [28, Section 6.1]. Since the algorithms for \mathbb{Z}_{2^n} have been optimised significantly (see [17, 35]) we believe that the resulting algorithms will be no less efficient than applying the AGS algorithm directly. Moreover, unlike the complexities working directly over \mathbb{Z}_p as explained in Remark 3.4, this technique (although it might introduce new “noise”) overcomes the need to take overlapping intervals and is not subject to the choice of the interval.

4.1 The i -bit function is concentrated

We now explain that modulus switching provides an alternative proof of the Morillo–Ràfols result that every single-bit functions is concentrated [38].

The above discussion assumed the function \tilde{f} extends f from \mathbb{Z}_p to \mathbb{Z}_{2^n} where 2^n is slightly larger than p . As Theorem 4.2 shows, one can consider modulus switching for domains of any size, including switching to a smaller domain. The results about concentration hold in this greater generality, and this provides a new technique to prove concentration of (some) families of functions, by showing that a subfamily of functions, defined on domains of specific forms, is concentrated.

Theorem 4.3 ([28, Theorem 6.1]). *Consider a family of functions $\mathcal{J} = \{f_{2^k} : \mathbb{Z}_{2^k} \rightarrow \mathbb{C}\}_{k \in \mathbb{N}}$ and define the family $\mathcal{J}' = \{f_n : \mathbb{Z}_n \rightarrow \mathbb{C}\}_{n \in \mathbb{N}}$, where for each $2^{k-1} < n \leq 2^k$ we let $f_n(x) := f_{2^k}(x)$ for every $x \in \mathbb{Z}_n$. If \mathcal{J} is concentrated then \mathcal{J}' is concentrated.*

As an application, one can prove that the i -th bit function is concentrated by showing that the family of the i -th bit function on domains \mathbb{Z}_{2^k} is concentrated, that is, that $\{\text{bit}_i : \mathbb{Z}_{2^k} \rightarrow \{-1, 1\}\}_{i < k \in \mathbb{N}}$ is concentrated. Here i can be a function of k , so for example the most-significant-bit function is given by $i = k - 1$. The latter can be easily proven using the structure of these functions under these domains. This is summarized in the following lemma, where we define $|x|_N := \min\{x, N - x\}$.

Lemma 4.4 ([28, Lemma 6.2]). *Let $k \in \mathbb{N}$ and $0 \leq i < k$. Define $\text{bit}_i : \mathbb{Z}_{2^k} \rightarrow \{-1, 1\}$ by $\text{bit}_i(x) = (-1)^{x_i}$ where $x = \sum_{j=0}^{k-1} x_j 2^j$ and $x_j \in \{0, 1\}$. Let $\alpha \in \mathbb{Z}_{2^k}$. Then $\widehat{\text{bit}_i}(\alpha) = 0$ unless α is an odd multiple of 2^{k-i-1} in which case $|\widehat{\text{bit}_i}(\alpha)| = O(2^{k-i}/|\alpha|_{2^k})$.*

The lemma shows that, when i is small there are a few non-zero coefficients (especially for $i = 0$, there is only one non-zero coefficient at $\alpha = 2^{k-1}$). When i is “medium” then there are non-zero coefficients at all multiples $\alpha = j2^{k-i-1}$, j odd, and they decrease in size with $1/|j|_{2^k}$. When i is large (e.g., $i = k - 1$) then the significant coefficients are all close to 0 and are spaced at distance $2 \cdot 2^{k-1-i}$ (i.e., when $i = k - 1$ they are 2 apart; for the second most significant bit they are spaced 4 apart, and so on).

A corollary is that the i -th bit function on \mathbb{Z}_{2^k} is concentrated. See arguments on the function half in Example 2.5 and [3, Claim 4.1]. For clarification, we state again that i can be a fixed constant ($i = 0$ corresponds to the least significant bit) or be dependent on k ($i = k - 1$ corresponds to the most significant bit).

Having established that the i -th bit function is concentrated on \mathbb{Z}_{2^k} , our modulus switching approach shows that the i -th bit function (on any domains \mathbb{Z}_N) is concentrated by Theorem 4.3. This general approach gives a new and simpler proof of the result in [38] (the proof in [38] is very technical; they decompose $N = k2^i \pm m$ and consider different cases of m).

5 Applications: Cryptography

The SFT algorithm has been used to reprove known results on the hardness of recovering bits of the secret values in the discrete logarithm problem (DLP) and RSA problem. It has been used to give reductions for the learning with errors (LWE) [41] and learning with rounding (LWR) [6] problems, that prove that the ‘search’ and ‘decision’ problems are equivalently hard even when the number of samples is fixed. It has also been used to prove results about the hardness of recovering bits of Diffie–Hellman shared secrets keys in both (non-prime) finite fields and elliptic curves. This section surveys how the SFT algorithm is used in these applications. In addition, we explain the specific model for which the Diffie–Hellman results hold, and clarify that the question whether single bits of Diffie–Hellman shared keys are hardcore (in the usual model) is still open.

5.1 Background and motivation

A one-way function h , if it exists, assures that while given x it is easy to compute $h(x)$, retrieving x from $h(x)$ is hard. This hardness does not necessarily mean that given $h(x)$ one cannot find some partial information of x . Naturally, the main interest is in trying to learn some bits of x , but other sorts of partial information have also been considered. Bits of x that cannot be learnt from $h(x)$, or more generally cannot be predicted noticeably better than a guess, are called *hardcore bits*. In other words, a hardcore bit is a bit which is as hard to compute (or to predict) as the entire secret value. For a historical overview see [20]. To show that a bit (or a set of bits) is hardcore, one usually tries to construct an algorithm that inverts h , given a target value $h(x)$ and an oracle that takes $h(t)$ and outputs a bit of t . In order to do so, one first needs to establish a way to query the oracle on values $h(t)$ such that there is some known relation between t and x , for example $t = \alpha x$ for known α 's.

A useful language to describe these ideas is the *hidden number problem*, which was introduced by Boneh and Venkatesan [11] in order to study bit security of secrets keys arising from Diffie–Hellman key exchange. This problem turned out to be general enough to be applied to other cryptographic problems like DLP and RSA. In fact, the generality of the problem allows it to be used also outside of the scope of bit security (see [40, Section 4.4] and references within, also [13, 5]). Therefore, the hidden number problem is of theoretical interest and is studied today in its own right. It has many extensions and different variants; see [44] for a comprehensive survey.

Definition 5.1 (Hidden number problem). Let (G, \cdot) be a group, let $s \neq 0$ be a secret (unknown) element of G and let f be a function defined over G . Find s using oracle access to the function $f_s(x) := f(s \cdot x)$.

We use the term *oracle access* as a general term for either of the following oracle models: in the *random access* model the solver receives polynomial many samples $(x, f_s(x))$ where the values x are drawn independently and uniformly at random from G ; in the *query access* model the solver can query

the oracle on any input $x \subseteq G$ and receive the answer $(x, f_s(x))$. To emphasize the difference between these models, we refer to the hidden number problem in the latter model as chosen-multiplier hidden number problem (CM-HNP). This problem can also be divided into two models, namely *adaptive access* where the solver has a continuous access to the oracle and can query it at any time of the recovery process, and *non-adaptive access* where the solver is not allowed to query the oracle once the recovery process has started. Other types of access models could be also considered. For example, the original work on the hidden number problem [11] considers an oracle for which on the query $x \in \mathbb{Z}_p$ replies with $(x, f(sg^x))$.

An interesting case is when the oracle is unreliable. That is, the oracle does not give a correct answer all the time, but with some probability. It is common to call an oracle that always provides a correct answer a *perfect* oracle. An oracle that is correct only with some noticeable advantage is called an *unreliable* or *imperfect* oracle.

The following table summarizes some of the known results on the hidden number problem in different models. Here p is a prime number and ‘imperfect’ under the ‘Oracle’ column refers to an oracle with any non-negligible advantage over trivial guessing. The starting point of this work is the Boneh–Venkatesan result [11] which requires a perfect oracle and uses lattice methods rather than Fourier learning methods; this work was adapted to unreliable oracles by [21], but there is a complex tradeoff with the number of bits and so we do not include it in our table.

Problem	Access	Group	Bits	Oracle	Remarks
HNP	random	\mathbb{Z}_p^*	$\sqrt{\log p} + \log \log p$ MSB ⁶	perfect	Given by [11]
CM-HNP	adaptive	\mathbb{Z}_p^*	LSB	imperfect	Given by [4]
CM-HNP	adaptive	\mathbb{Z}_p^*	any single bit	imperfect	Given by [24]
CM-HNP	non-adaptive	\mathbb{Z}_N^*	MSB & LSB	imperfect	Given by [8]
CM-HNP	non-adaptive	\mathbb{Z}_N^*	each single bit for the outer $\log \log p$ bits	imperfect	Given by [3]
CM-HNP	non-adaptive	\mathbb{Z}_N^*	any single bit	imperfect	Given by [38]

Most early works such as [4, 8, 24] require complicated algebraic manipulations such as tweaking and untweaking bits. Using the SFT algorithm [3] gives a uniform and clear approach. We present this solution to CM-HNP, using different terminology than the original one, for functions of norm 1, as the subsequent applications involve single bit functions (with the convention that $\text{bit}_i(x) = (-1)^{x_i}$ where x_i is the i -th bit of x).

Theorem 5.2 ([3]). *Let $f : \mathbb{Z}_N \rightarrow \{-1, 1\}$ be a function with a τ -heavy Fourier coefficient $\alpha \in \mathbb{Z}_N^*$ for $\tau^{-1} = \text{poly}(\log |G|)$. Then, the chosen-multiplier hidden number problem in \mathbb{Z}_N with $s \in \mathbb{Z}_N^*$ and the function f can be solved in polynomial time.*

In particular, the theorem holds for every concentrated function.

Remark 5.3 (Coding Theory terminology). Theorem 5.2 rephrases Theorem 2 of [3]. The latter work gives a polynomial time list-decoding algorithm for *concentrated codes* with corrupted code words

⁶Since one can easily transform HNP with the LSB function to HNP with the MSB function, HNP can also be solved given $\sqrt{\log p} + \log \log p$ LSB. A generalization of this technique [39, Section 5.1] allows to transform HNP with $2d$ consecutive inner bits to HNP with d MSB, hence HNP can also be solved given $2(\sqrt{\log p} + \log \log p)$ consecutive inner bits.

(Theorem 1) and subsequently a general list-decoding methodology for proving hardcore functions (Theorem 2). Most subsequent works on hardcore bits adopt this coding-theoretic language. Thus, in order to apply Theorem 2 of [3], these works use Theorem 1 of [3], which applies to concentrated codes. This caused the authors of these works to put effort into proving that a particular code is concentrated. However, we emphasize that to apply the CM-HNP approach of [3] there is no need for the function to be concentrated. Instead it suffices that the function has a significant Fourier coefficient, and this is usually much easier to prove. We make this clear in our formulation of Theorem 5.2. In other words, while concentration is sufficient for a code to be *recoverable* it is not a necessary condition. For these reasons (and others) we find the coding-theoretic language unhelpful and do not use it in this paper.

We now sketch the proof of Theorem 5.2: run the SFT algorithm on f and f_s to get short lists L, L_s of τ -heavy coefficients for each function, respectively. By the scaling property $\widehat{f}_s(\alpha) = \widehat{f}(\alpha s^{-1})$ for every α . Therefore, for every $\alpha \in L_s$ for which $\widehat{f}_s(\alpha)$ is τ -heavy there exists $\beta \in L$ such that $\beta = \alpha s^{-1}$. The secret s can be recovered efficiently. Notice that while the hidden number problem takes place in a multiplicative group, this solution involves Fourier analysis over an additive group.

A template for algorithms for CM-HNP is the following: show that (i) the “partial information” function f has a significant coefficient, (ii) the function f_s has a significant coefficient, and (iii) some (recoverable) relation between the coefficients of f and f_s exists. If one succeeds in showing these conditions, then using the SFT algorithm one can solve this instance of CM-HNP. This template allows bit security researchers to look for settings where a solution to CM-HNP is already known (namely, cases where these three conditions are already known to hold, like single-bit functions over \mathbb{Z}_N) and try to convert their problem of interest to this setting.

5.1.1 The multivariate hidden number problem

Another case of interest is the *multivariate hidden number problem* (MVHNP), which we define as follows.

Definition 5.4 (Multivariate hidden number problem). Let R be a ring, let $\mathbf{s} = (s_1, \dots, s_m) \neq (0, \dots, 0)$ be a secret (unknown) element in R^m and let f be a function defined over R . Find \mathbf{s} using oracle access to the function $f_{\mathbf{s}}(\mathbf{x}) := f(\mathbf{s} \cdot \mathbf{x}) = f(s_1 x_1 + \dots + s_m x_m)$.

Specific instances of this problem are LWE and LWR, and it is related to *trace-HNP* [31] and *polynomial HNP* [43]. Similar to the solution to HNP in \mathbb{Z}_p , one can give a solution in \mathbb{Z}_{p^m} in the random access model for a function f that outputs $\sqrt{\log(p^m)} = \sqrt{m \log(p)}$ MSB’s of its input (derived from [43], for example).

One can also define CM-MVHNP, the chosen-multiplier version of the multivariate hidden number problem, similar to CM-HNP. To solve this variant we need an analogue of the Fourier scaling property in higher dimensions. Such an analogue, which we call the multivariate scaling property, is given in [16, Lemma 13] and we sketch it now.

Multivariate scaling property. Let $f : \mathbb{Z}_p \rightarrow \mathbb{C}$, let $\mathbf{s} = (s_1, \dots, s_m) \in \mathbb{Z}_p^m$ such that not all $s_i = 0$, and

define $f_{\mathbf{s}} : \mathbb{Z}_p^m \rightarrow \mathbb{C}$ by $f_{\mathbf{s}}(\mathbf{x}) := f(\mathbf{s} \cdot \mathbf{x})$. For any $s_k \neq 0$, the Fourier transform of $f_{\mathbf{s}}$ satisfies

$$\widehat{f}_{\mathbf{s}}(\mathbf{z}) = \widehat{f}_{\mathbf{s}}(z_1, \dots, z_m) = \begin{cases} \widehat{f}(c) & \text{if } (z_1, \dots, z_m) = (cs_1, \dots, cs_m), c \in \mathbb{Z}_p; \\ 0 & \text{otherwise.} \end{cases}$$

This allows generalizing Theorem 5.2 to CM-MVHNP. The proof, which we omit, follows from the proof to Theorem 5.2 given above.

Theorem 5.5 ([16]). *Let $f : \mathbb{Z}_p \rightarrow \{-1, 1\}$ be a function with a τ -heavy Fourier coefficient $\alpha \in \mathbb{Z}_p^*$ for $\tau^{-1} = \text{poly}(\log |G|)$. Then, the chosen-multiplier multivariate hidden number problem in \mathbb{Z}_p^m with the function f can be solved in polynomial time.*

5.2 Applications

We present some of the applications in cryptography of the SFT algorithm. They are all based on reducing some problems to the CM-HNP or CM-MVHNP. In the following we assume to have an oracle O that solves some problem, and show how to use this oracle to solve a harder problem, thus establishing the hardness equivalence between the two problems.

5.2.1 Proving known results: bit security of RSA and DLP

The first application of the algorithm was given in [3], where it is shown that the most significant bit and least significant bit are hardcore for the RSA function $RSA_{N,e}(x) := x^e \pmod{N}$ and for exponentiation $EXP_g(x) := g^x$, where g is an element of prime order ℓ in some group. The results hold for imperfect oracles that have noticeable advantage over guessing. These results were already known, as [4] first shows that the LSB is hardcore for the RSA function and [24] shows that every bit is hardcore for both functions. Nevertheless, the approach based on SFT is more general (holds for every function with significant coefficients) and simpler. We explain how to derive these results.

Claim 5.6. *Each single bit is hardcore for the RSA function. That is, predicting any bit of x , given $RSA_{N,e}(x)$, is as hard as inverting the RSA function.*

We sketch the proof: One direction is trivial. In the other direction, given an instance $RSA_{N,e}(x) = x^e \pmod{N}$, we want to recover x . Suppose the (imperfect) oracle O takes $RSA_{N,e}(t) := t^e \pmod{N}$ and outputs $\text{bit}_i(t)$, the i -th bit of t . Since the values e, N are public in the RSA setting, for every number r one can compute $RSA_{N,e}(rx \pmod{N})$ by $(r^e \pmod{N})(x^e \pmod{N}) = (rx)^e \pmod{N}$. Hence, given $RSA_{N,e}(x)$ one can query the oracle on $RSA_{N,e}(rx)$ to get the i -th bit of rx for every chosen r . The problem therefore becomes the CM-HNP in \mathbb{Z}_N^* , and this can be solved using the SFT algorithm over the additive group $(\mathbb{Z}_N, +)$, which has known order. Indeed, bit_i is concentrated (see Section 4.1), thus has a significant coefficient. The oracle function O also has a significant coefficient (see Section 3.3). The rest follows from Theorem 5.2.

Claim 5.7. *Each single bit is hardcore for the exponentiation function EXP_g for prime-order element g . That is, predicting any bit of x , given $EXP_g(x)$ is as hard as inverting the function EXP_g , i.e. solving DLP in the corresponding group.*

The proof, which we leave as an exercise, is similar to the previous case, using the fact that $(g^x)^r = g^{rx}$. This proves bit security results for the DLP in finite fields and elliptic curves. Similar results also hold for other functions (problems), as *Rabin* (see [1, Chapter 7]) and the Paillier trapdoor permutation (see [38, Section 7]).

5.2.2 Bit security of the Diffie–Hellman protocol and related schemes

An open question is to prove that single bits of Diffie–Hellman keys are hardcore. Here we consider an oracle O that on g, g^a, g^b returns a single bit of the Diffie–Hellman key $s = g^{ab}$. To interact with the oracle, notice that given g^b one can compute $g^{b+r} = g^b g^r$ for any r . One can then query the oracle O with g, g^a, g^{b+r} and receive a bit of $g^{a(b+r)} = g^{ab} g^{ar} = st$. This is how the hidden number problem was originally identified. This interaction does not correspond to the CM-HNP, since choosing the multiplier $t = (g^a)^r$ (for the secret s) is equivalent to finding discrete logarithms for the base g^a in \mathbb{Z}_p^* .

Advice bits. For related schemes where the exponent a is fixed (unlike schemes using ephemeral exponents, as in Diffie–Hellman key exchange), Akavia [2] followed Boneh–Venkatesen [12] to get around this problem by assuming an “advice” that provides the discrete logarithms of the chosen multipliers t to the base g^a , but this is not realistic in actual applications (see also our remark in Section 6.2). There is currently no method known to prove the hardness of single bits of Diffie–Hellman keys in the usual model.

New Diffie–Hellman model. To overcome this problem, Boneh and Shparlinski [10] suggested (in the context of elliptic curves) a different model where the oracle O takes as input, in addition to the values g, g^a, g^b , a group homomorphism $\phi : G \rightarrow G'$, and then outputs partial information (e.g. a single bit) of $\phi(g^{ab})$. The approach is then to keep the inputs g, g^a, g^b fixed and to use ϕ as the way to choose multipliers for $s = g^{ab}$ in the hidden number problem. This model corresponds to a variant of Diffie–Hellman key exchange, where a representation of the group is not fixed. We call this the *representation changing model*.

This is an example of our discussion at the end of Section 5.1 above on converting a given problem to a setting (a new model, in this case) that allows to apply the solution to CM-HNP. We now explain how in this model one can reduce the original problem to variants of CM-HNP.

In this model one can think of the bit security problem for any secret element s (not necessarily a Diffie–Hellman key as the interaction with the oracle does not come from the key exchange setting). Let $\mathbf{s} = (s_1, \dots, s_n) \in G$ and write $\phi(\mathbf{s}) = (\phi_1(\mathbf{s}), \dots, \phi_n(\mathbf{s})) \in G'$, and suppose that the oracle O returns a bit of some component $\phi_i(\mathbf{s})$. Write also $\mathbf{r} = (r_1, \dots, r_n)$. Suppose there exists a family of homomorphisms $\phi^{\mathbf{r}}$ for every⁷ \mathbf{r} such that for some $1 \leq i \leq n$ the i -th component of $\phi^{\mathbf{r}}$ satisfies $\phi_i^{\mathbf{r}}(\mathbf{x}) := \sum_{j=1}^n r_j x_j$. Then, getting a single bit of $\phi_i^{\mathbf{r}}(\mathbf{s}) = \sum_{j=1}^n r_j s_j$ for chosen \mathbf{r} , gives rise to CM-MVHNP for a single-bit function and the secret \mathbf{s} . A special case is where \mathbf{r} is of the form $r_j \cdot \mathbf{e}_j = (0, \dots, 0, r_j, 0, \dots, 0)$. Then $\phi_i^{\mathbf{r}}(\mathbf{s}) = r_j s_j$, which gives rise to CM-HNP for a single-bit function and secret s_j .

Therefore, if one can find a group for which the condition on the homomorphisms $\phi^{\mathbf{r}}$ holds, then proving the hardness of single bits in this model reduces to either CM-MVHNP or CM-HNP (note that in the latter case one only recovers a component of s , and therefore needs other methods for recovering the entire value s ; for the case in which $s = g^{ab}$ is a Diffie–Hellman key in \mathbb{F}_{p^m} that we describe below,

⁷It is sufficient that there is a ‘large enough’ subfamily of homomorphisms.

one can use the results involving “summing functions” from [48] and recover the entire secret \mathbf{s} from the algorithm that recovers a single (fixed) component s_i ; for the case of elliptic curves it is sufficient to know one coordinate, as there are at most 3 values for the other coordinate). We give a brief overview of the known results in the literature.

As mentioned above, this idea was introduced by Boneh and Shparlinski [10] for the LSB of (both the x, y coordinates of) Diffie–Hellman keys in elliptic curve groups over prime fields. It is shown there that changing the Weierstrass equation is an isomorphism that gives rise to the desired multipliers. Indeed, it is well known that twists of the curve give $\phi(x, y) = (u^2x, u^3y)$. Therefore, given a request for desired multiplier r (for example by the SFT algorithm), one can obtain it if there is a solution to $u^d = r$ (where $d = 2$ or $d = 3$, depends on the coordinate) and flip a coin to guess the bit if a solution does not exist. The work [10] uses the same technique as in [4] to prove hardness of LSB. This approach was then applied by [14] (see also [25]) to every single bit of a larger class of elliptic curve secrets, that also includes Diffie–Hellman keys in elliptic curves, using the SFT algorithm (that is, using the solution to CM-HNP for single-bit functions, as in Theorem 5.2).

The idea of changing group representations can also be used for finite fields. The works [15, 49] consider the computational Diffie–Hellman (CDH) problem in groups $\mathbb{F}_{p^m}^*$ for $m > 1$. They show that some polynomial representations of \mathbb{F}_{p^m} give rise to the desired homomorphisms $\phi^{\mathbf{r}}$ for $\mathbf{r} = r_j \cdot \mathbf{e}_j$, and therefore reduce to CM-HNP.

For a detailed overview of these techniques we refer the reader to the exposition of Sections 5, 5.1, 5.2 and subsections within of [16]. This latter work gives applications of the solution for CM-MVHNP to show bit security of the computational Diffie–Hellman problem in groups of higher dimension in models similar to those mentioned above; specifically, for elliptic curves over extension fields, and for $\mathbb{F}_{p^m}^*$ with different (non-polynomial) representations of the field \mathbb{F}_{p^m} .

We stress that these models do not tell a lot about the hardness of specific bits in real-life implementations of Diffie–Hellman key exchange, where the representation of the group is fixed. One should interpret results in the representation changing model as follows: assuming hardness of CDH in a group G (where G can be the multiplicative group of a finite extension field or an elliptic curve over a finite field), there is no algorithm that takes $g, g^a, g^b \in G$ and outputs the i -th bit of g^{ab} for many representations of G (more precisely, for representations corresponding to the specific isomorphisms used in the reduction). Nevertheless, given an instance g^a, g^b in a specific representation of G , this result does not tell us whether it is hard to compute a specific bit of the secret g^{ab} . Indeed, this problem is still open.

5.2.3 Sample-preserving search-to-decision reductions for LWE and LWR

We assume the reader is familiar with the *search* and *decision* variants of the LWE and LWR problems [6, 41]. The problem at hand is to reduce the search problem to the decision problem. That is, to show that the decision problem is at least as hard as the search problem. This is done in a similar fashion to the bit security reductions above: one assumes an oracle to the decision problem is available, and uses it to solve the search problem. We explain the reduction and show how the SFT algorithm is used to get a reduction in the stronger “sample preserving” model. This is done, as above, by reducing the problem to CM-MVHNP.

We only focus on the part of the reduction which involves the SFT algorithm; the entire reduction is more involved. By a “hybrid” argument (see [20, Theorem 1] or [9, Lemma 3]), one can reduce the

decision problem to distinguishing a specific LWE sample, among the set of all samples.⁸ We therefore consider a single LWE sample.

The standard method to show that the decision problem is as hard as the search problem is as follows. Suppose one has a perfect decision oracle. Given an LWE sample $b = \langle \mathbf{a}, \mathbf{s} \rangle + e = a_1 s_1 + \dots + a_n s_n + e \pmod{p}$ one makes a guess s' for s_1 and re-randomises the sample as $\mathbf{a}' = (a_1 + r, a_2, \dots, a_n)$, $b' = b + rs' \pmod{p}$. If the guess is correct (i.e., if $s' = s_1$) then (\mathbf{a}', b') is a valid LWE sample whereas if the guess is incorrect then b' is uniform and independent of the other samples. Hence the decision oracle determines whether the guess s' of the secret value s_1 is correct. After at most pn queries to the decision oracle one can compute the secret \mathbf{s} .

When the oracle is not perfect one will have to repeat this procedure with different inputs (\mathbf{a}, b) and follow majority rule. When the success rate of the oracle is low, one may not have enough initial inputs (\mathbf{a}, b) to satisfactorily apply the majority rule, and therefore would need to draw more samples. A *sample-preserving* reduction is a reduction that uses only the initial given samples, and does not ask for more samples during the procedure. Micciancio and Mol [36] used the SFT algorithm to give a sample-preserving search-to-decision reduction for the learning with errors problem. We now explain this reduction.

The standard method above involves choosing a unit vector \mathbf{e}_j and guessing $\langle \mathbf{e}_j, \mathbf{s} \rangle$. Micciancio and Mol observe that one can choose any vector \mathbf{v} and guess $\langle \mathbf{v}, \mathbf{s} \rangle$, then let the decision oracle to advise whether this guess is correct or incorrect. Again, if the oracle is perfect then one determines the correct guesses, denotes them by $b_{\mathbf{v}}$, and eventually obtains n linear equations in \mathbf{s} and hence can solve the problem. However if the oracle is not perfect (but has a noticeable advantage over a random guess), then in the case where the oracle says that the guess for $\langle \mathbf{v}, \mathbf{s} \rangle$ is incorrect (more precisely, that the distribution is uniform), one sets $b_{\mathbf{v}}$ to be some value from the remaining $p - 1$ possibilities, chosen uniformly. Then for a selection of chosen vectors \mathbf{v} we have the values $b_{\mathbf{v}}$, for which $b_{\mathbf{v}} = \langle \mathbf{v}, \mathbf{s} \rangle$ with some noticeable bias from $\frac{1}{p}$. In other words, we have query access to a noisy version of the function $f(\mathbf{v}) = \langle \mathbf{v}, \mathbf{s} \rangle \pmod{p}$.

This is an instance of CM-MVHNP with an unreliable oracle. The function $\omega_p^{b_{\mathbf{v}}}$, which is a noisy version of $\omega_p^{\langle \mathbf{v}, \mathbf{s} \rangle}$, has a significant coefficient for the character $\chi_{\mathbf{s}}$ (see Section 6.1). Thus, one can run the SFT algorithm on the function $\omega_p^{b_{\mathbf{v}}}$, to find this significant coefficient, hence the character, and thus solve this problem.

A very similar approach is taken in [9] for the learning with rounding problem. We remark that in the case of a non-prime p the reduction is more subtle, and requires some restrictions (see [36, 9] for more details). We also remark that the reduction is an average-case reduction, and does not hold for the worst case (more precisely, there may be a set of initial samples $\{(\mathbf{a}_i, b_i)\}$ for which the reduction fails). A sample-preserving reduction for the latter is still an open problem.

6 Limitations: Non-Linear Problems

This section presents limitations on natural generalisation of the approaches taken above to a larger class of applications. We show that the linearity in the hidden number problem, induced from the operation

⁸The reduction given in [36] uses the duality of the LWE and knapsack functions.

$s \cdot x$, is essential for the SFT to be useful. In particular, we give an answer (in the negative) to an open question in [32].

The solution to the CM-HNP in \mathbb{Z}_N (Theorem 5.2) is based on Fourier analysis in the additive group $(\mathbb{Z}_N, +)$ and it exploits the scaling property of the Fourier transform for the function $f_s(x) := f(sx)$. In other words, the function f_s is the composition of f with a linear map on \mathbb{Z}_N . It is natural to consider whether this approach can be used for other algebraic groups (such as elliptic curves and algebraic tori). The hidden number problem in the case of elliptic curves is to determine a secret point $S \in E(\mathbb{F}_p)$ given samples $(P, f(S+P))$ where a typical choice for the function would be $f(Q) = \text{bit}_i(x(Q))$. The natural approach is to still use Fourier analysis in the additive group $(\mathbb{Z}_p, +)$ but instead of composing with a linear map, to compose with a rational function (e.g., coming from the translation map $t_S(P) = P + S$). Another generalisation would be Fourier analysis in other groups (G, \cdot) .

If such tools could be developed we might have an approach to the bit security of Diffie–Hellman key exchange in the group of elliptic curve points in certain models. There are also other interesting problems that could be approached with Fourier analysis on general groups. For example, the authors of [32] raise the question whether it is possible to apply these results to the *modular inversion hidden number problem*.

Unfortunately, there is a major obstacle to applying the SFT algorithm to these sorts of problems. Namely, if f is a concentrated function then the composition $f \circ \varphi$ is concentrated only when φ is affine. In fact, $f \circ \varphi$ has significant coefficients only when φ is affine. The aim of this section is to explain this obstacle. Since the translation map for the elliptic curve group law is a non-affine rational function, this explains why the method cannot be directly applied to the elliptic curve hidden number problem. Our argument also answers the question of [32] in the negative.

Let $f : G \rightarrow \mathbb{C}$ be a function and let $f_s(x) = f \circ \varphi_s(x)$, where $\varphi_s : G \rightarrow G$ is an efficiently computable function (that depends on some unknown value s). To generalise the proof of Theorem 5.2 one needs the following three conditions:

1. the function f has significant coefficients;
2. the function f_s has significant coefficients;
3. there exists a relation between the significant coefficients of f and f_s that allows to determine s (or at least a small set of candidates for s).

One special case is when f is a constant function. Then f_s is also a constant function and both conditions 1 and 2 are satisfied. The problem is that a constant function cannot tell us anything about the secret s , and so condition 3 does not hold. Hence, we need to focus on functions that are far from constant, which we formalise in our proof by requiring that $\widehat{f}(0) = 0$ (in other words, f is “balanced”).

Having dispensed with this special case we focus on the first two conditions. We first consider the case when f is concentrated. If $\varphi_s(x) = ax + b$ is affine then we already know from the scaling and time-shifting properties that all Fourier coefficients of f are preserved in f_s , and so if f is concentrated then f_s is also concentrated. Our aim is to show a converse to this fact: if φ_s is a rational function and if conditions 1 and 2 both hold then φ_s must be affine. This result is closely related to the Beurling–Helson Theorem [7] (see [27, 30] for related results in \mathbb{Z}_p) and the work of Green and Konyagin [23] on the Fourier transform of balanced functions.

For our result we need the following lemma [39, Lemma 7] (a proof, for general fields \mathbb{F}_{p^m} , can be found in [37, Theorem 2]).

Lemma 6.1. *Let q be prime. For any polynomials $f, g \in \mathbb{F}_q[x]$ such that the rational function $h = \frac{f}{g}$ is not constant in \mathbb{F}_q , the following bound holds*

$$\left| \sum_{\lambda \in \mathbb{F}_p}^* \omega_q^{h(\lambda)} \right| \leq (\max\{\deg(f), \deg(g)\} + u - 2)\sqrt{q} + \delta,$$

where \sum^* means that the summation is taken over all $\lambda \in \mathbb{F}_q$ which are not poles of h and

$$(u, \delta) = \begin{cases} (v, 1) & \text{if } \deg(f) \leq \deg(g), \\ (v + 1, 0) & \text{if } \deg(f) > \deg(g), \end{cases}$$

and v is the number of distinct zeros of g in the algebraic closure of \mathbb{F}_q .

We formulate the following result for functions on \mathbb{Z}_q for a prime q , but it can be generalised to finite fields \mathbb{F}_{p^m} with $m > 1$. Let $g, h \in \mathbb{Z}_q[x]$ be polynomials where h is not the constant zero. Let Z_h be the set of zeroes in \mathbb{Z}_q of h . We define $\varphi(x) = g(x)/h(x)$ for all $x \in \mathbb{Z}_q \setminus Z_h$ and $\varphi(x) = 0$ otherwise (since we will assume Z_h is small compared with q it does not matter how we define φ on Z_h).

Recall that the definition of concentration applies to families of functions. To keep the formulation of the following proposition clean, we call a single function concentrated as explained after the definition above.

Proposition 6.2. *Let q be a sufficiently large prime. Let f be a concentrated function on \mathbb{Z}_q such that $\|f\|_2 = 1$ and $\widehat{f}(0) = 0$. Let $g, h \in \mathbb{Z}_q[x]$ be polynomials of degree bounded by $\text{poly}(\log(q))$ and let Z_h be the set of zeroes of h . Define $\varphi(x)$ as above and suppose this function is non-constant. Let $\tau = 1/\text{poly}(\log(q))$. If $f \circ \varphi$ has any τ -heavy Fourier coefficients then $\varphi(x) = ax + b$ for some $a, b \in \mathbb{Z}_q$.*

Proof. Let $G = \mathbb{Z}_q$ and write $f = \sum_{\alpha \in G} \widehat{f}(\alpha) \chi_\alpha$. Let $d = \max\{\deg(g(x)), \deg(h(x))\}$. Let $\varepsilon = \frac{\tau}{32d^2}$. Since f is concentrated there is a set Γ of size $\text{poly}(\log(|G|))$ such that

$$\|f - f|_\Gamma\|_2^2 \leq \varepsilon = \frac{\tau}{32d^2}.$$

Since $\widehat{f}(0) = 0$ it follows that Γ does not contain zero.

Now consider $f_\varphi(x) = f(\varphi(x)) = \sum_{\alpha \in G} \widehat{f}(\alpha) \chi_\alpha(\varphi(x))$. Assume it has a τ -heavy coefficient; for contradiction we suppose $\varphi(x) \neq ax + b$ for any a, b . For every $\beta \in G$ we have

$$\begin{aligned} \widehat{f}_\varphi(\beta) &= \frac{1}{|G|} \sum_{x \in G} f_\varphi(x) \overline{\chi_\beta(x)} = \frac{1}{|G|} \sum_{x \in G} f(\varphi(x)) \overline{\chi_\beta(x)} = \\ &= \frac{1}{|G|} \sum_{x \in G} \sum_{\alpha \in G} \widehat{f}(\alpha) \chi_\alpha(\varphi(x)) \overline{\chi_\beta(x)} = \frac{1}{|G|} \sum_{\alpha \in G} \widehat{f}(\alpha) \sum_{x \in G} \chi_\alpha(\varphi(x)) \overline{\chi_\beta(x)} = \\ &= \frac{1}{|G|} \sum_{\alpha \in G} \widehat{f}(\alpha) \sum_{x \in G} \chi_1(\alpha\varphi(x) - \beta x) = \frac{1}{|G|} \sum_{\alpha \in G} \widehat{f}(\alpha) \sum_{x \in G} \chi_1(\psi_\alpha^\beta(x)), \end{aligned}$$

where we denote $\psi_\alpha^\beta(x) = \alpha\varphi(x) - \beta x$. Since $\widehat{f}(0) = 0$ we can ignore the case $\alpha = 0$ and by our supposition that $\varphi \neq ax + b$ we know that there are no α, β such that ψ_α^β is constant. Hence, the last sum is a character sum satisfying the conditions of Lemma 6.1. Furthermore, $\psi_\alpha^\beta = (\alpha g(x) - \beta xh(x))/h(x)$ and so the value u in Lemma 6.1 is bounded by $\max\{\deg(g), \deg(h)\} \leq d$. Applying Lemma 6.1, we get that for every $\alpha \neq 0$ and every β it holds that $|\sum_{x \in G \setminus Z_h} \chi(\psi_\alpha^\beta(x))| \leq C$ where $C = 2d\sqrt{q}$.

Now note that

$$\widehat{f}_\varphi(\beta) = \frac{1}{|G|} \sum_{\alpha \in G} \widehat{f}(\alpha) \sum_{x \in Z_h} \chi_1(\psi_\alpha^\beta(x)) + \frac{1}{|G|} \sum_{\alpha \in \Gamma} \widehat{f}(\alpha) \sum_{x \in G \setminus Z_h} \chi_1(\psi_\alpha^\beta(x)) + \frac{1}{|G|} \sum_{\alpha \notin \Gamma} \widehat{f}(\alpha) \sum_{x \in G \setminus Z_h} \chi_1(\psi_\alpha^\beta(x)).$$

For the first term we note that $|\sum_{x \in Z_h} \chi_1(\psi_\alpha^\beta(x))| \leq |Z_h| \leq d$ and that $\|f\|_2 = 1$ implies $\sum_{\alpha \in G} |\widehat{f}(\alpha)| \leq \sqrt{|G|} = \sqrt{q}$ and $|\widehat{f}(\alpha)| \leq 1$ for all α . Therefore

$$\left| \widehat{f}_\varphi(\beta) \right| \leq \frac{d}{\sqrt{q}} + \left| \frac{1}{|G|} \sum_{\alpha \in \Gamma} \widehat{f}(\alpha) \sum_{x \in G \setminus Z_h} \chi(\psi_\alpha^\beta(x)) \right| + \left| \frac{1}{|G|} \sum_{\alpha \notin \Gamma} \widehat{f}(\alpha) \sum_{x \in G \setminus Z_h} \chi(\psi_\alpha^\beta(x)) \right|.$$

We apply the triangle inequality on the first sum and the Cauchy–Schwarz inequality on the second. Let $k = |\Gamma|$ and write $\Gamma = \{\alpha_1, \dots, \alpha_k\}$. Then using Lemma 6.1 we get

$$\begin{aligned} \left| \frac{1}{|G|} \sum_{\alpha \in \Gamma} \widehat{f}(\alpha) \sum_{x \in G \setminus Z_h} \chi(\psi_\alpha^\beta(x)) \right| &= \left| \frac{1}{|G|} \sum_{j=1}^k \widehat{f}(\alpha_j) \sum_{x \in G \setminus Z_h} \chi(\psi_{\alpha_j}^\beta(x)) \right| \leq \left| \frac{1}{q} \sum_{j=1}^k \widehat{f}(\alpha_j) \cdot C \right| \\ &\leq \frac{1}{q} \sum_{j=1}^k |\widehat{f}(\alpha_j)| C = \frac{2kd}{\sqrt{q}}. \end{aligned}$$

Since $k = |\Gamma| = \text{poly}(\log(q))$ we have that this bound (similarly for the earlier bound d/\sqrt{q}) is negligible, so we have for example

$$\frac{d}{\sqrt{q}} + \frac{2kd}{\sqrt{q}} < 2d\sqrt{\varepsilon}.$$

From Parseval’s identity $\sum_{\alpha \notin \Gamma} |\widehat{f}(\alpha)|^2 = \|f - f|_\Gamma\|_2^2 \leq \varepsilon$. Therefore, by the Cauchy–Schwarz inequality we have

$$\left| \frac{1}{|G|} \sum_{\alpha \notin \Gamma} \widehat{f}(\alpha) \sum_{x \in G \setminus Z_h} \chi(\psi_\alpha^\beta(x)) \right| \leq \frac{1}{|G|} \left(\sum_{\alpha \notin \Gamma} |\widehat{f}(\alpha)|^2 \right)^{\frac{1}{2}} \left(\sum_{\alpha \notin \Gamma} \left| \sum_{x \in G \setminus Z_h} \chi(\psi_\alpha^\beta(x)) \right|^2 \right)^{\frac{1}{2}} \leq \frac{1}{|G|} \sqrt{\varepsilon} \left(\sum_{\alpha \notin \Gamma} C^2 \right)^{\frac{1}{2}}.$$

Then

$$\left| \frac{1}{|G|} \sum_{\alpha \notin \Gamma} \widehat{f}(\alpha) \sum_{x \in G \setminus Z_h} \chi(\psi_\alpha^\beta(x)) \right| \leq \frac{\sqrt{\varepsilon} \sqrt{q - k} 2d\sqrt{q}}{q} \leq 2d\sqrt{\varepsilon}.$$

Finally, combining the bounds we get

$$\left| \widehat{f}_\varphi(\beta) \right|^2 \leq \left(\frac{d}{\sqrt{q}} + \frac{2kd}{\sqrt{q}} + 2d\sqrt{\varepsilon} \right)^2 < \left(4d\sqrt{\varepsilon} \right)^2 = \left(4d \frac{\sqrt{\tau}}{4d\sqrt{2}} \right)^2 = \frac{\tau}{2}.$$

Therefore, for every β the coefficient $\widehat{f}_\varphi(\beta)$ is not τ -heavy for any noticeable τ . This gives the required contradiction and so we conclude that φ is affine. \square

6.1 ε -concentrated functions

Proposition 6.2 shows that if f is concentrated (and far from constant) and $f \circ \varphi$ has significant coefficients, then φ is affine. It is natural to wonder whether the condition that f is concentrated is necessary. In fact, the result cannot be weakened in general: if $\varphi(x) = g(x)/h(x)$ is non-affine and invertible almost everywhere (such as a Möbius function $\varphi(x) = (ax+b)/(cx+d)$ where $ad - bc = 1$) then $f(x) = \chi_\alpha(x) + \chi_\beta(\varphi^{-1}(x))$ is such that $f(x)$ has a significant coefficient at α and $f \circ \varphi$ has a significant coefficient at β .

However, a version of Proposition 6.2 is true for some non-concentrated functions of interest. Since Theorem 5.2 does not require the function to be concentrated, it is of interest to also show that composing with non-affine $\varphi(x)$ is an obstruction to the solution to CM-HNP for these functions as well. Hence, for the rest of this section we consider a ‘noisy character’, $f(x) := \omega_N^{\alpha x + e(x)}$. We first show that these functions have a significant coefficient, then we show that $f \circ \varphi$ does not have a significant coefficient when φ is not affine.

To formalise the problem we think of $e(x)$ as a random variable from some distribution (e.g., a discrete Gaussian or a uniform distribution on some small interval compared with N). We treat $e(x)$ as being independent of x , in which case we can write

$$\widehat{f}(\beta) = \mathbb{E} \left(\omega_N^{\alpha x - \beta x + e(x)} \right) = \mathbb{E} \left(\omega_N^{(\alpha - \beta)x} \omega_N^{e(x)} \right) = \mathbb{E} \left(\omega_N^{(\alpha - \beta)x} \right) \mathbb{E} \left(\omega_N^{e(x)} \right).$$

To show that $|\widehat{f}(\alpha)|$ is large it suffices to give a lower bound for $|\mathbb{E} \left(\omega_N^{e(x)} \right)|$. We do this by following an argument due to Bleichenbacher [8].

Bleichenbacher defines the *bias* of a random variable X on \mathbb{Z} as

$$B_N(X) = \mathbb{E} \left(\exp(2\pi i X / N) \right).$$

Assume X is the uniform distribution in some interval $[0, T - 1]$ for some $0 < T \leq N$. Then

$$B_N^U(X) := B_N(X) = \frac{1}{T} \sum_{0 \leq x < T} \exp(2\pi i x / N).$$

Some properties of $B_N^U(X)$ appear in Lemma 1 of [13]. Since the latter is a geometrical progression,

$$B_N^U(X) = \frac{1}{T} \frac{\sin(\pi T / N)}{\sin(\pi / N)}.$$

Suppose $e(x)$ follows the uniform distribution X . That is, for each $x \in \mathbb{Z}_N$ the value $e(x)$ is chosen uniformly and independently at random in $[0, T - 1]$. From linearity it is easy to see that

$$\mathbb{E} \left(\omega_N^{e(x)} \right) = \frac{1}{N} \sum_{x \in \mathbb{Z}_N} \left(\exp(2\pi i e(x)/N) \right) = \frac{N/T}{N} \sum_{0 \leq t < T} \exp(2\pi i t/N) = \frac{1}{T} \sum_{0 \leq t < T} \exp(2\pi i t/N) = B_N^U(X).$$

It is obvious that if $T = N$ then $B_N^U(X) = 0$. In applications $e(x)$ usually represents some given bits, and so it is natural to restrict $T \leq N/2$ as we do, though the following argument also holds given a “fraction of a bit”, i.e. for $T > N/2$. For $T \leq N/2$ one has⁹ $|B_N^U(X)| > 0.5$, and so $|\mathbb{E} \left(\omega_N^{e(x)} \right)|^2 = |B_N^U(X)|^2 > 0.25$. The desired lower bound is provided.

A similar approach holds when e follows a Gaussian distribution. In this case the size of the bias is even larger, as $e(x) = 0$ on a large set (and $e(x)$ is small on an even larger set) and so most of the “energy” is distributed around zero.

Hence, we have established that a noisy character has a significant coefficient. Finally, we address the result of Proposition 6.2 for such a function.

Claim 6.3. *Let φ be as in Proposition 6.2, and let $e(x)$ given by the uniform distribution (over some interval in \mathbb{Z}_N) or by a Gaussian distribution. If $f_\varphi(x) := \omega_N^{\varphi(x)+e(x)}$ has a significant coefficient then $\varphi(x) = ax + b$ for some $a, b \in \mathbb{Z}_N$.*

Proof sketch. We observe that for every β

$$\widehat{f}_\varphi(\beta) = \mathbb{E} \left(\omega_N^{\varphi(x)-\beta x+e(x)} \right) = \mathbb{E} \left(\omega_N^{\psi_1^\beta(x)} \omega_N^{e(x)} \right) = \mathbb{E} \left(\omega_N^{\psi_1^\beta(x)} \right) \mathbb{E} \left(\omega_N^{e(x)} \right),$$

where $\psi_1^\beta(x) = \varphi(x) - \beta x$. Since $|\mathbb{E} \left(\omega_N^{e(x)} \right)| \leq 1$, it suffices to upper-bound $|\mathbb{E} \left(\omega_N^{\psi_1^\beta(x)} \right)|$. Such a bound follows from Lemma 6.1 in the same way as in the proof of Proposition 6.2. \square

6.2 Hidden number problem in subgroups

Another limitation on the applications of the SFT algorithm is the following. Suppose that the multipliers in the hidden number problem are drawn from some set $H \subseteq G$. One can consider the multipliers to be in a proper subgroup $H < G$, as done in [22, 45]. It is not clear how to apply the SFT algorithm to solve this variant of the (chosen-multiplier) hidden number problem. Specifically, the chosen queries in the algorithm have to be correlated, but it is not guaranteed that these correlated queries will all lie in the same subgroup. If the index $[G : H]$ is small (e.g., $[G : H] = 2$, as in the case of the set of squares in \mathbb{F}_p^*) then the issue can be managed, but if $[G : H]$ is large then no results are known. Therefore, for results (on Diffie–Hellman related schemes) that rely on advice of the form of discrete logarithms to some base g (as in [2, 12, 46]), if g generates a relatively small subgroup, it is not guaranteed that the desired correlated multipliers are indeed in the group generated by g . This restricts, for example, the result given in [2, Section 5]. This observation is similar to the one in [44, Section 2.5], and was handled in [10, Section 5] and [14, 4.1] since the set of squares in \mathbb{F}_p^* has index 2 in \mathbb{F}_p^* .

⁹See [13, Table 1] for some values $|B_N^U(X)|$ for different $T \leq N/2$.

Acknowledgements

We thank Ben Green for providing some insights and references. We also thank two anonymous referees for their helpful comments on an earlier version of the paper.

References

- [1] Akavia, A. (2008) “Learning Noisy Characters, Multiplication Codes and Hardcore Predicates.” Ph.D. Thesis. Massachusetts Institute of Technology. [8](#), [17](#), [26](#)
- [2] Akavia, A. (2009) “Solving Hidden Number Problem with One Bit Oracle and Advice,” in Halevi, S. (ed.) *Advances in Cryptology – CRYPTO 2009*. LNCS, vol. 5677, pp. 337–354. Springer, Heidelberg. [26](#), [33](#)
- [3] Akavia, A., Goldwasser, S., and Safra, S. (2003) “Proving Hard-Core Predicates Using List Decoding,” in *FOCS 2003*, pp. 146–157. IEEE Computer Society, Washington, DC. [2](#), [3](#), [5](#), [6](#), [8](#), [11](#), [16](#), [17](#), [21](#), [23](#), [24](#), [25](#)
- [4] Alexi, W., Chor, B., Goldreich, O., and Schnorr, C.P. (1988) “RSA and Rabin Functions: Certain Parts are as Hard as the Whole,” in *SIAM Journal on Computing*, **17**(2), 194–209. [3](#), [9](#), [23](#), [25](#), [27](#)
- [5] Aranha, D.F., Fouque, P.-A., Gérard B., Kammerer, J.-G., Tibouchi, M., and Zapalowicz, J.-C. (2014) “GLV/GLS Decomposition, Power Analysis, and Attacks on ECDSA Signatures with Single-Bit Nonce Bias,” in Sarkar, P., Iwata, T. (eds.) *Advances in Cryptology – ASIACRYPT 2014*. LNCS, vol. 8873, pp. 262–281. Springer, Heidelberg. [10](#), [22](#)
- [6] Banerjee, A., Peikert, C., and Rosen, A. (2012) “Pseudorandom Functions and Lattices,” in D. Pointcheval and T. Johansson (eds.), *Advances in Cryptology – EUROCRYPT 2012*, LNCS, vol. 7237, pp. 719–737. Springer, Heidelberg. [22](#), [27](#)
- [7] Beurling, A., and Helson, H. (1953) “Fourier–Stieltjes Transforms with Bounded Powers,” in *Math. Scand.*, **1**, 120–126. [29](#)
- [8] Bleichenbacher, D. (2000) “On the Generation of One-Time Keys in DL Signature Schemes,” Presentation at IEEE P1363 Working Group meeting. [2](#), [9](#), [10](#), [23](#), [32](#)
- [9] Bogdanov, A., Guo, S., Masny, D., Richelson, R., and Rosen, A. (2016) “On the Hardness of Learning with Rounding over Small Modulus,” in E. Kushilevitz and T. Malkin (eds.), *Theory of Cryptography – TCC 2016-A*, LNCS, vol. 9562, pp. 209–224. Springer, Heidelberg. [3](#), [27](#), [28](#)
- [10] Boneh, D., and Shparlinski, I.E. (2001) “On the Unpredictability of Bits of the Elliptic Curve Diffie–Hellman Scheme,” in Kilian, J. (ed.) *Advances in Cryptology – CRYPTO 2001*. LNCS, vol. 2139, pp. 201–212. Springer, Heidelberg. [26](#), [27](#), [33](#)
- [11] Boneh, D., and Venkatesan, R. (1996) “Hardness of Computing the Most Significant Bits of Secret Keys in Diffie–Hellman and Related Schemes,” in Koblitz, N. (ed.) *Advances in Cryptology – CRYPTO ’96*. LNCS, vol. 1109, pp. 129–142. Springer, Heidelberg. [10](#), [22](#), [23](#)

- [12] Boneh, D., and Venkatesan, R. (1997) “Rounding in Lattices and its Cryptographic Applications,” in Saks, M.E. (ed.) SODA 1997, pp. 675–681. ACM/SIAM, Philadelphia. [26](#), [33](#)
- [13] De Mulder, E., Hutter, M., Marson, M.E., and Pearson, P. (2013) “Using Bleichenbacher’s Solution to the Hidden Number Problem to Attack Nonce Leaks in 384-Bit ECDSA,” in G. Bertoni and J.S. Coron (eds.), CHES 2013, LNCS, vol. 8086, pp. 435–452. Springer, Heidelberg. [10](#), [22](#), [32](#), [33](#)
- [14] Duc, A., and Jetchev, D. (2012) “Hardness of Computing Individual Bits for One-Way Functions on Elliptic Curves,” in Safavi-Naini, R. and Canetti, R. (eds.) Advances in Cryptology – CRYPTO 2012. LNCS, vol. 7417, pp. 832–849. Springer, Heidelberg. [3](#), [5](#), [27](#), [33](#)
- [15] Fazio, N., Gennaro, R., Perera I.M., and Skeith, W.E. III (2013) “Hard-Core Predicates for a Diffie–Hellman Problem over Finite Fields,” in Canetti, R. and Garay, J.A. (eds.) Advances in Cryptology – CRYPTO 2013. LNCS, vol. 8043, pp. 148–165. Springer, Heidelberg. [3](#), [27](#)
- [16] Galbraith, S.D., and Shani, B. (2015) “The Multivariate Hidden Number Problem,” in Lehmann, A. and Wolf, S. (eds.) Information Theoretic Security – ICITS 2015. LNCS, vol. 9063, pp. 250–268. Springer, Heidelberg. [3](#), [24](#), [25](#), [27](#)
- [17] Gilbert, A.C., Indyk, P., Iwen, M., and Schmidt, L. (2014) “Recent Developments in the Sparse Fourier Transform,” in IEEE Signal Processing Magazine, **31**(5), 91–100. [2](#), [19](#), [21](#)
- [18] Goldreich, O. (1999) “Modern Cryptography, Probabilistic Proofs and Pseudorandomness.” Springer-Verlag Berlin Heidelberg. [9](#)
- [19] Goldreich, O., and Levin, L.A. (1989) “A Hard-Core Predicate for all One-Way Functions,” in STOC 1989, ACM, pp. 25–32. New York, NY [2](#), [9](#)
- [20] González Vasco, M.I., and Näslund, M. (2001) “A Survey of Hard Core Functions,” in Lam, K.-Y., Shparlinski, I., Wang, H. and Xing, C. (eds.) Proc. Workshop on Cryptography and Computational Number Theory 1999. Progress in Computer Science and Applied Logic, vol. 20, pp. 227–255. Birkhäuser, Basel. [3](#), [22](#), [27](#)
- [21] González Vasco, M.I., Näslund, M., and Shparlinski, I.E. (2004) “New Results on the Hardness of Diffie–Hellman Bits,” in Bao, F., Deng, R. and Zhou, J. (eds.) Public Key Cryptography – PKC 2004. LNCS, vol. 2947, pp. 159–172. Springer, Heidelberg. [23](#)
- [22] González Vasco, M.I., and Shparlinski, I.E. (2001) “On the Security of Diffie-Hellman Bits,” in Lam, K.-Y., Shparlinski, I., Wang, H. and Xing, C. (eds.) Proc. Workshop on Cryptography and Computational Number Theory 1999. Progress in Computer Science and Applied Logic, vol. 20, pp. 257–268. Birkhäuser, Basel. [33](#)
- [23] Green, B., and Konyagin, S. (2009) “On the Littlewood Problem Modulo a Prime,” in Canad. J. Math., vol. **61**(1), 141–164. [29](#)
- [24] Håstad, J., and Näslund, M. (2003) “The Security of all RSA and Discrete Log Bits,” in Journal of the ACM, **51**(2), 187–230. [3](#), [23](#), [25](#)

- [25] Kiltz, E. (2001) “A Primitive for Proving the Security of Every Bit and About Universal Hash Functions & Hard Core Predicates,” full version of paper published in the Proceedings of the 13th International Symposium on Fundamentals of Computation Theory, pp. 388–391. Springer-Verlag. http://homepage.ruhr-uni-bochum.de/Eike.Kiltz/papers/hash_full.pdf 27
- [26] Kushilevitz, E., and Mansour, Y. (1991) “Learning Decision Trees Using the Fourier Spectrum,” in Koutsougeras, C. and Vitter, J.S. (eds.) STOC 1991, ACM, 455–464. 2, 6, 9, 10, 11, 12, 13, 14, 15
- [27] Konyagin, S.V., and Shkredov, I.D. (2015) “A Quantitative Version of the Beurling–Helson Theorem,” in Functional Analysis and Its Applications, vol. 49(2), 110–121. 29
- [28] Laity, J., and Shani, B. (2017) “On Sets of Large Fourier Transform Under Changes in Domain,” in Applied and Computational Harmonic Analysis 45, 216–232. 3, 19, 20, 21
- [29] Langlois, A., Peikert, C., Regev, O., and Stehlé, D. (2013) “Classical Hardness of Learning with Errors,” in STOC 2013, ACM, pp. 575–584. New York, NY 15, 18
- [30] Lebedev, V. (2012) “Absolutely Convergent Fourier Series. An Improvement of the Beurling–Helson Theorem,” in Functional Analysis and Its Applications, vol. 46(2), 121–132. 29
- [31] Li, W.-C.W., Näslund, M., and Shparlinski, I.E. (2002) “Hidden Number Problem with the Trace and Bit Security of XTR and LUC,” in Yung, M. (ed.) Advances in Cryptology – CRYPTO 2002. LNCS, vol. 2442, pp. 433–448. Springer, Heidelberg. 24
- [32] Ling, S., Shparlinski, I.E., Steinfeld, R., and Wang, H. (2011) “On the Modular Inversion Hidden Number Problem,” in Journal of Symbolic Computation, 47(4), 358–367. 3, 29
- [33] Mansour, Y. (1992) “Randomized Interpolation and Approximation of Sparse Polynomials,” in Kuich, W. (ed.) Proceedings of the 19th International Colloquium on Automata, Languages and Programming, pp. 261–272. Springer, Heidelberg. 2, 11, 13, 14, 15
- [34] Mansour, Y. (1994) “Learning Boolean Functions via the Fourier transform,” in Roychowdhury, V., Kai-Yeung, S. and Orlitsky, A. (eds.) Theoretical Advances in Neural Computation and Learning, pp. 391–424. Springer US, Boston. 6, 12
- [35] Mansour, Y. and Sahar, S. (2000) “Implementation Issues in the Fourier Transform Algorithm,” in Machine Learning, vol. 40(1), 5–33. 19, 21
- [36] Micciancio, D., and Mol, P. (2011) “Pseudorandom Knapsacks and the Sample Complexity of LWE Search-to-Decision Reductions,” in P. Rogaway (ed.), Advances in Cryptology – CRYPTO 2011, LNCS, vol. 6841, 465–484. Springer, Heidelberg. 3, 28
- [37] Moreno, C.J., and Moreno, O. (1991) “Exponential sums and Goppa codes: I,” in Proc. Amer. Math. Soc., 111, 523–531. 30
- [38] Morillo, P., and Ràfols, C. (2009) “The Security of All Bits Using List Decoding,” in Jarecki, S. and Tsudik, G. (eds.) Public Key Cryptography – PKC 2009. LNCS, vol. 5443, pp. 15–33. Springer, Heidelberg. 3, 5, 21, 22, 23, 26

- [39] Nguyen, P.Q., and Shparlinski, I.E. (2002) “The Insecurity of the Digital Signature Algorithm with Partially Known Nonces,” in *Journal of Cryptology*, **15**(3), 151–176. [23](#), [30](#)
- [40] Nguyen, P.Q., and Stern, J. (2001) “The Two Faces of Lattices in Cryptology,” in J. H. Silverman (ed.), *Cryptography and Lattices 2001*, LNCS, vol. 2146, pp. 146–180. Springer, Heidelberg. [22](#)
- [41] Regev, O. (2005) “On Lattices, Learning with Errors, Random Linear Codes, and Cryptography,” in *STOC 2005*, ACM, pp. 84–93. New York, NY [22](#), [27](#)
- [42] Shor, P.W. (1994) “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer,” in *FOCS 1994*, pp. 124–134. IEEE Computer Society, Washington, DC. [3](#), [19](#)
- [43] Shparlinski, I.E. (2001) “Sparse Polynomial Approximation in Finite Fields,” in *Proc. 33rd ACM Symposium on Theory of Computing – STOC 2001*, pp. 209–215. ACM, New York. [24](#)
- [44] Shparlinski, I.E. (2005) “Playing “Hide-and-Seek” with Numbers: The Hidden Number Problem, Lattices and Exponential Sums,” in Garrett, P. and Lieman, D. (eds.) *Public-Key Cryptography; Proceedings of Symposia in Applied Mathematics*, vol. 62, AMS, pp. 153–177. [22](#), [33](#)
- [45] Shparlinski, I.E., and Winterhof, A. (2005) “A Hidden Number Problem in Small Subgroups,” in *Mathematics of Computation*, vol. 74, pp. 2073–2080. [33](#)
- [46] Shparlinski, I.E., and Winterhof, A. (2004) “A Nonuniform Algorithm for the Hidden Number Problem in Subgroups,” in Bao, F., Deng, R.H. and Zhou, J. (eds.) *Public Key Cryptography – PKC 2004*. LNCS, vol. 2947, pp. 416–424. Springer, Heidelberg. [33](#)
- [47] Terras, A. (1999) “Fourier Analysis on Finite Groups and Applications,” in *London Mathematical Society Student Texts* (No. 43), Cambridge University Press. Cambridge. [4](#)
- [48] Verheul, E.R. (2000) “Certificates of Recoverability with Scalable Recovery Agent Security,” in Imai, H., Zheng, Y. (eds.) *Public Key Cryptography – PKC 2000*. LNCS, vol. 1751, pp. 258–275. Springer, Heidelberg. [27](#)
- [49] Wang, M., Zhan, T., and Zhang, H. (2016) “Bit Security of the CDH Problems over Finite Fields,” in O. Dunkelman and L. Keliher (eds.), *SAC 2015*, LNCS, vol. 9566, 441–461. Springer, Heidelberg. [3](#), [27](#)

AUTHORS

Steven D. Galbraith
 Professor
 University of Auckland, New Zealand
s.galbraith@auckland.ac.nz
<https://www.math.auckland.ac.nz/~sgal018/>

Joel Laity
Work done as a postgraduate student
University of Auckland, New Zealand
jlai052@aucklanduni.ac.nz

Barak Shani
Work done as a postgraduate student
University of Auckland, New Zealand
barak.shani@auckland.ac.nz