On the Black-Box Complexity of Private-Key Inner-Product Functional Encryption

Mohammad Hajiabadi [1]
Roman Langrehr [2]
Adam O'Neill [3]
Mingyuan Wang [4]

[1] University of Waterloo
mdhajiabadi@uwaterloo.ca
[2] ETH Zurich
roman.langrehr@inf.ethz.ch
[3] Manning CICS, UMass Amherst
adamo@cs.umass.edu
[4] NYU Shanghai
mingyuan.wang@nyu.edu

Abstract

We initiate the study of the black-box complexity of *private-key* functional encryption (FE). Of central importance in the private-key setting is the *inner-product* functionality, which is currently only known from assumptions that imply public-key encryption, such as Decisional Diffie-Hellman or Learning-with-Errors. As our main result, we rule out black-box constructions of private-key inner-product FE from random oracles. This implies a black-box separation between private-key inner-product FE from all symmetric-key primitives implied by random oracles (e.g., symmetric-key encryption and collision-resistant hash functions).

Proving lower bounds for private-key functional encryption schemes introduces challenges that were absent in prior works. In particular, the combinatorial techniques developed by prior works for proving black-box lower bounds are only useful in the public-key setting and predicate encryption settings, which all fail for the private-key FE case. Our work develops novel combinatorial techniques based on Fourier analysis to overcome these barriers. We expect these techniques to be widely useful in future research in this area.

Keywords: Black-box impossibility, Functional encryption

1 Introduction

1.1 Background and Main Question

A major goal in cryptography is to identify minimal assumptions sufficient for realizing cryptographic primitives. Functional encryption (FE) [SW05, BSW11, O'N10] is a vast generalization of standard encryption whereby secret key holders

can decrypt a given ciphertext to various corresponding functions of its underlying plaintext. In particular, secret keys are associated with functions, and a secret key holder for a function f can learn $f(x)$ from an encryption of x. Security notions for FE capture the intuitive idea that secret keys for functions $f_1, \ldots, f_w$ should reveal only what can already be learned from the outputs of these functions on the underlying plaintexts.

For functional encryption, if the number of corruptions is *a priori* bounded such that the size of the system parameters can depend on this bound, the minimal complexity required is well-understood. In particular, one can build public-key (resp., private-key) bounded-collusion FE for any function family from the minimal assumption that CPA-secure public-key (resp., private-key) encryption schemes exist [SS10, GVW12, AV19].

However, things are unclear in the *unbounded collusion* case, where an adversary can obtain (a.k.a., corrupt) secret keys for many functions of her choosing. It is known that functional encryption with unbounded collusions for arbitrary polynomial-sized circuits implies Indistinguishability Obfuscation (iO) [BV15, AJ15]. Therefore, it is unlikely that we can build unbounded functional encryptions from plain CPA-secure encryption schemes. However, the question remains if, for some less expressive families of functions, unbounded functional encryptions can be built from minimal assumptions. Thus, much research has been devoted to realizing and improving the efficiency of unbounded FE for specific restrictive functionalities such as identity-based encryption [BF01], attribute-based encryption [SW05], predicate encryption [KSW13], inner-product FE [ABDP15], quadratic FE [BCFG17], and attribute-weighted sums [AGW20].

Private key vs public key: For most advanced encryption systems (e.g., key-dependent message (KDM) security [BRS02, CL01], homomorphic encryption [Gen09]) building private-key schemes appears to be as much challenging as their public-key counterparts, and sometimes even in a provable way [Rot11]. For FE, the situation seems to be different. For example, consider identity-based encryption (IBE) [BF01], which corresponds to point functions defined as $F_{\mathsf{id}}(\mathsf{id}', m) = (\mathsf{id}', m)$ if $\mathsf{id} = \mathsf{id}'$, and $F_{\mathsf{id}}(\mathsf{id}', m) = (\mathsf{id}', \perp)$, otherwise.[1] IBE is so far only possible in a black-box way from pairings/LWE, and is known to be black-box impossible from trapdoor permutations (TDPs) [BPR+08] or generic groups [PRV12, SS21, Zha22]. On the other hand, FE for point functions in the private-key setting can be trivially built from pseudorandom functions (PRFs) [GGM84] as follows. Let the master secret key msk be a PRF key, and let $k[\mathsf{id}'] := \mathsf{PRF}(\mathsf{msk}, \mathsf{id}')$ be a secret key for a point id'. Define $\mathrm{Enc}(\mathsf{msk}, (\mathsf{id}', m))$ as $\mathrm{Enc}_{\mathsf{priv}}(k[\mathsf{id}'], m)$, where $k[\mathsf{id}'] = \mathsf{PRF}(\mathsf{msk}, \mathsf{id}')$, and where $\mathrm{Enc}_{\mathsf{priv}}$ is the encryption function of a CPA-secure private-key encryption scheme.

What makes the above private-key construction possible are two points: (a) that a master secret key can implicitly generate exponentially many private keys and (b) each ciphertext can be decrypted by *exactly one* secret key, the same identity. In particular, the above observation readily generalizes to building private-key FE for any function family F under which for any plaintext x, for all but a polynomial number (in the security parameter κ) of keys $f \in F$, $f(x) = \perp$. For encrypting x, if $f_1, \ldots, f_{\mathsf{poly}(\kappa)}$ are the only functions for which $f_i(x) \neq \perp$

[1] The standard security notion for IBE allows the ciphertext is allowed to leak id, that is why we included it in the function output. IBEs that also hide the identity are called *anonymous*.

for $i \in [\mathsf{poly}(\kappa)]$, then encrypt $f_i(x)$ for every i under $\mathsf{PRF}(\mathsf{msk}, f_i)$, the secret key for f_i. The size of the final ciphertext remains polynomial because at most a polynomial number of $f_i(x)$ values are encrypted.

Thus, unlike in the public-key setting, for which we have lower bounds on FE for certain function families (e.g., IBE) [BPR+08, KY09] and a better understanding of the hierarchy between different functionalities (e.g., separations between IBE and more expressive functionalities such as attribute-based and predicate encryption [GKLM12]), our understanding of FE in the private-key setting is lacking. In particular, the lower bounds in the public key setting fail in the private setting, due to the positive result above and due to the more technical reason that in the private-key setting, encryptions are made relative to master secret keys, capable of generating exponentially many private keys. But in the public-key setting, encryptions are made relative to master public keys, which are bound to encode at most polynomially many public keys. We will elaborate more on these later.

Motivated by the above discussion, a seemingly basic characterization of what FE families can be built from OWFs in the private-key setting is missing. Therefore, the research direction our work aims to make progress on is:

For what function families F is private-key FE for F (im)possible from one-way functions?

In this work, we take the first step in this research direction. In particular, we consider the inner product functionality [ABDP15], where for a modulus q and dimension n, a function from the family is associated with $\mathbf{y} \in \mathbb{Z}_q^n$ and is defined as $f_{\mathbf{y}}(\mathbf{x}) = \langle \mathbf{x}, \mathbf{y} \rangle$ for all $\mathbf{x} \in \mathbb{Z}_q^n$. Inner-product is a simple and fundamental operation in both theory and practice, and there is a large body of work on this functionality and variants/extensions, *e.g.* [ABDP15, BJK15, ALS16, ALS16, AGRW17, BCFG17, ACF+18, CLT18, ABG19, Gay20, ACF+20, ACGU20, ALMT20]. This functionality is especially attractive in the *private-key setting* because it can be *function-hiding* [BJK15], namely, the key for $f_{\mathbf{y}}$ can hide $\mathbf{y}$ (to the extent possible given the inner products).[2] However, the only known constructions of private-key IPFE, even without function-hiding, are based on *algebraic* assumptions (*e.g.*, DDH/LWE), starting with the work of [ABDP15].[3]

While most positive results focus on IPFE for *restricted* inner-products (*i.e.*, where there is a polynomial number of possible inner-products recovered by decryption), we focus on the *unrestricted setting*. In particular, such a scheme is known from class groups [CLT18]. The unrestricted setting is arguably more natural, as the restricted setting came about as a result of limitations of the proposed constructions, not any external desire for achieving it. This brings us to the main question of this work:

Is private-key inner-product FE black-box possible from one-way functions?

We answer this question negatively in this work. The challenge we overcome in answering the above question is that all existing 'combinatorial' techniques

[2] Note that the function-hiding property is *unattainable* in the public-key setting since one can always recover $\mathbf{y}$ from $\mathsf{sk}_{\mathbf{y}}$ by first invoking the publicly available encryption algorithm on different $\mathbf{x}$ and then decrypting using $\mathsf{sk}_{\mathbf{y}}$ to learn $\langle \mathbf{x}, \mathbf{y} \rangle$.

[3] Note [ABDP15] construct public-key schemes, which trivially imply private-key ones. Later works starting with [BJK15] explicitly address the private-key setting.

employed in all the black-box impossibility results in the public-key setting (*e.g.*, for IBE [BPR+08, KY09] and ABE [GKLM12]) completely fail in the private-key setting (see Sections 1.3 and 2 for further discussions.) Thus, our work departs significantly from prior works and develops new combinatorial and other proof techniques to answer the above question.

1.2 Our Results

We prove that building private-key FE for inner-product functions (IPFE) is black-box impossible from OWFs, or more generally from any assumptions that hold relative to a random oracle (RO) — *e.g.*, collision-resistant hash functions (CRHFs), KDM-secure private-key encryption. We stress that our result does *not* require function-hiding for IPFE and holds relative to seemingly minimal formulations of its security.

Technically, we prove our result by showing that private-key inner-product FE (IPFE) cannot be constructed in the information-theoretic random oracle model [IR89]. Central to our impossibility proofs is a combinatorial lemma such that if proved for a function family, then we will have a black-box impossibility for that function family from ROs. We show that the combinatorial lemma holds for the inner-product functionality using techniques from Fourier analysis and covering problems in linear subspaces. The characterization of this combinatorial lemma and the proof of the lemma for inner products are the main novelties of our work, and will hopefully pave the way for characterizing FE functionalities provably impossible from ROs.

The main motivation behind the above question is to initiate the study of the black-box complexity of private-key functional encryption. Moreover, our work will facilitate future efforts to understand the black-box complexity of variants of IPFE. For example, function-hiding IPFE is so far only known from pairings [BJK15], and we have limited impossibility results for it from lattice assumptions [Ünа20, TÜ24]. However, we do not have any impossibility results for it in the generic-group model (GGM) without pairings. Some of the challenges that appear in ruling out function-hiding IPFE in the GGM also emerge in our setting, so we are hopeful that our work will also be useful for proving such an impossibility result.

1.3 Novelty, Comparison to Prior Work and Open Problems

As mentioned earlier, there are impossibility results for predicate encryption (PE) schemes in the public-key setting [BPR+08, KY09, GKLM12]. But all these results crucially make use of the public-key setting, and fail in the private-key setting (*c.f.*, the positive construction of private-key IBE from OWFs). Second, our results concern FE schemes that are of fine-grained access (*i.e.*, that each decryption reveals some partial information about the plaintext), whereas previous impossibility results concern only PE schemes, which are of the all-or-nothing nature.[4] Ruling out fine-grained FE schemes present additional challenges, as explained below. For example, for all we know IPFE might be possible from CPA-secure encryption schemes because IPFE is not know to imply

[4]That is, decryption reveals either the entire plaintext or nothing about it.

any PE schemes that are ruled out from CPA-secure schemes. For instance, we know how to build IPFE from black-box DDH [ABDP15], but we have black-box impossibility results for IBE from DDH [PRV12, SS21, Zha22]. This suggests that building IPFE might be 'easier' than IBE, or than other related PE primitives.

Private-Key vs Public-Key. Let us illustrate why at a technical level the PE impossibility results of [BPR+08, KY09, GKLM12] fail in the private-key setting. In the public-key version for IBE, encryptions are made under mpk, which can encode at most polynomially-many base public keys (call this Property (*) below), while in the private-key version, encryptions are made under msk, which can potentially encode exponentially-many secret keys. We mentioned this point before. The above public-key impossibility results crucially rely on (*), implying that at most polynomially-many public keys $\mathsf{pk}_1, \ldots, \mathsf{pk}_t$ can be embedded into mpk. Specifically, if one corrupts $q \gg t$ identities and learns their secret keys, one has learned enough trapdoors associated with pk_i to be able to decrypt for an uncorrupted identity. But this intuition fails in the secret key setting as exemplified earlier.

The results of [GKLM12] shows that threshold predicate encryption is black-box impossible from IBE. This work gives an impossibility from IBE, a primitive for which the master secret key can generate exponentially-many trapdoors. But the main difference between [GKLM12] and ours is that in [GKLM12], too, they crucially rely on the public-key setting, used to argue that a master public key for the threshold FE can encode at most polynomially-many public-keys.

In light of the above, it is an open problem if, and to what extent, the PE impossibility results of [BPR+08, KY09, GKLM12] generalize to the secret-key setting. This tension between private-key and public-key concerns both the predicate encryption and our impossibilities for FE.

Functional Encryption vs Predicate Encryption. As mentioned earlier, the key difference between FE and PE is the partial decryption vs. full decryption natures of these primitives. Consequently, the security game of FE puts *more restrictions* on the set of keys that the adversary is allowed to corrupt during a successful attack. Hence, the 'combinatorial' techniques employed in the black-box impossibility result must be proven in a more stringent setting obeying such restrictions. We explain the difference between our combinatorial lemma and those of prior works in detail in Section 2.2.

Open Problems. The main problem left open by our work is to prove black-box impossibilities in the private-key setting for other FE functionalities. One concrete functionality that we were not able to handle is *fuzzy* FE [SW05]. Here plaintext and secret key vectors are all in $\mathbb{Z}_2^n$ and $F(\mathbf{x}, \mathbf{v}) = 1$ if the Hamming distance between $\mathbf{v}$ and $\mathbf{x}$ is at least ηn (for some fixed $\eta < 1$), and $F(\mathbf{x}, \mathbf{v}) = 0$, otherwise. Note that each decryption reveals some information about the plaintext $\mathbf{x}$ (because the decryption either outputs zero or one). Just like for our impossibility result, coming up with appropriate combinatorial lemmas and proving them will be the main challenges here.

It would also be interesting to see if and how the black-box impossibility results of [GKLM12] for PE extend to the private-key setting. In general,

understanding the black-box complexity of PE primitives in the private-key setting is a worthwhile goal.

Finally, as we mentioned, it is an intriguing open problem if we can extend our results to rule out the existence of function-hiding private-key IPFE in the GGM without pairings. Similarly, it would be interesting to extend our results to rule out black-box constructions of *restricted* IPFE from OWFs.

2 Technical Overview

A private-key IPFE relative to a random oracle O for vectors in $\mathbb{Z}_q^n$ is given by $\mathcal{E}^O := (\mathrm{KGen}^O, \mathrm{Enc}^O, \mathrm{Dec}^O)$, satisfying the following properties.[5] Let $w = w(\kappa)$ be the length of a master-secret key, where κ is the security parameter. The algorithm $\mathrm{KGen}^O(\mathsf{msk}, \mathbf{v})$ outputs a vector secret key $\mathsf{sk}[\mathbf{v}]$. One can use msk to encrypt a plaintext vector $\mathbf{x}$ as $C \leftarrow \mathrm{Enc}^O(\mathsf{msk}, \mathbf{x})$. Finally, decrypting C using $\mathsf{sk}[\mathbf{v}]$ as $\mathrm{Dec}^O(\mathsf{sk}[\mathbf{v}], C)$ returns $\langle \mathbf{v}, \mathbf{x} \rangle$.

We require the following weak notion of indistinguishability security (See Definition 3.3 for more details.) An adversary submits (non-adaptively, at once) t vectors $\mathbf{v}_1, \ldots, \mathbf{v}_t$, where $t = t(\kappa)$ can be an arbitrarily-large polynomial, as well as a *challenge secret-key vector* $\mathbf{v}^*$. It is required that $\mathbf{v}^* \notin \mathsf{Span}(\mathbf{v}_1, \ldots, \mathbf{v}_t)$, where $\mathsf{Span}()$ denotes the linear span of the corresponding vectors. In response, the adversary receives the secret keys $\{\mathsf{sk}[\mathbf{v}_i]\}_{i \in [t]}$ for these t vectors, as well as t ciphertexts $C_1, \ldots, C_t$ of t random plaintexts $\mathbf{x}_1, \ldots, \mathbf{x}_t$ sampled by the challenger, and $m_1, \ldots, m_t$, formed as follows:

- If the challenge bit $b = 0$, $m_i = \langle \mathbf{v}^*, \mathbf{x}_i \rangle$ for $i \in [t]$.
- If $b = 1$, $m_i \leftarrow \mathbb{Z}_q$ for all $i \in [t]$.

The adversary should be able to guess the value of b only with a probability negligibly greater than $1/2$. Note that the adversary is not given the underlying plaintexts $\mathbf{x}_1, \ldots, \mathbf{x}_t$.

Breaking $\mathcal{E}^O$ relative to ROs. Let $\mathcal{E}^O := (\mathrm{KGen}^O, \mathrm{Enc}^O, \mathrm{Dec}^O)$ be a candidate IPFE. Here we describe an adversary Brk^O that makes a polynomial number of queries to O, and then analyze its advantage. The attack is based on a polynomial $t(\kappa)$, instantiated later.

The adversary Brk^O chooses the challenge secret key vector $\mathbf{v}^*$ uniformly at random from $\mathbb{Z}_q^n$, chooses a random $(n-1)$-dimensional subspace S subject to $\mathbf{v}^* \notin S$ and chooses $\mathbf{v}_1, \ldots, \mathbf{v}_t$ uniformly at random from S. Let $(\{\mathsf{sk}[\mathbf{v}_i]\}_{i \in [t]}, \{C_i\}_{i \in [t]}, \{m_i\}_{i \in [t]})$ be the variables returned to Brk^O, as per the description of the game above. Also, suppose $\mathbf{x}_i$ is the underlying plaintext vector for C_i. The adversary should determine if m_i's are totally random values, or are the inner products of $\mathbf{x}_i$ with $\mathbf{v}^*$.

Simple case: Enc makes no O queries. To lay out our main techniques and to point out the challenges, let us make an overly simplifying assumption that Enc makes no queries. This is not a reasonable assumption, but we start

[5] There is also an additional Setup^O algorithm, that generates a master secret key msk. But we can remove this algorithm because an msk can be a uniformly random string from an appropriate space $\{0,1\}^w$, for some $w = w(\kappa)$.

with this assumption to describe our main techniques — later, we show how to remove this assumption using our combinatorial lemmas. Equipped with this assumption, we show how to design Brk to break the IPFE scheme by making a polynomial number of queries, and by making some other computation that takes exponential time but involves no queries. This will be sufficient for a black-box impossibility proof because ROs are OWFs against any adversary that can run in exponential time but which can make at most a polynomial number of queries. In particular, it shows that the adversary Brk, which breaks the IPFE scheme, cannot be used as a black-box to break the one-wayness of the oracle O.

Let S_g be the set of the query-answer (Q-A) pairs made to generate the challenge secret key $\mathsf{sk}[\mathbf{v}^*] \leftarrow \mathrm{KGen}(\mathsf{msk}, \mathbf{v}^*)$ and ρ the bit-length of the description of such a set for any vector secret key. Moreover, let μ be the length of a vector secret key. The adversary Brk^O attacks the scheme as follows.

(i) If there exists a secret key $\mathsf{sk} \in \{0,1\}^\mu$ and $\mathsf{S}_g \in \{0,1\}^\rho$ such that the following condition holds, return 0; otherwise, return 1:

   (a) $\mathrm{Dec}^{O'}(\mathsf{sk}, C_h) = m_h$ for all $h \in [t]$, where O' is defined as follows. If the query appears in the set S_g, respond to it accordingly, else, respond with a random value.

Note that Brk makes no queries to the oracle O at all — simulating the query responses based on S_g and random values, without invoking O itself. (The fact that Brk makes no queries is because of the above two simplifying assumptions.) Let us analyze the advantage of Brk.

Challenge bit $b = 0$: Suppose $\mathcal{E}$ is $(1-\alpha)$-correct, meaning that for any oracle O, and for any secret-key vector $\mathbf{v}$, the following holds. If we generate msk, $\mathsf{sk}[\mathbf{v}]$ (a secret key for $\mathbf{v}$) and C (a ciphertext for a random plaintext vector $\mathbf{x}$) all at random, the probability that $\langle \mathbf{v}, \mathbf{x} \rangle = \mathrm{Dec}^O(\mathsf{sk}[\mathbf{v}], C)$ is at least $1-\alpha$. (See Definition 3.2.) We show when $b = 0$, Brk outputs 0 with probability at least $(1-\alpha)$. To see this, we argue that Condition (i)a will hold with probability at least $(1-\alpha)$ when sk is set to $\mathsf{sk}[\mathbf{v}^*]$. The non-triviality of this lies in the fact that decryption is performed relative to O', and not relative to O itself. But since Enc makes no queries at all, and since we try all possible S_g — the set of Q-A pairs made to build $\mathsf{sk}[\mathbf{v}^*] \leftarrow \mathrm{KGen}^O(\mathsf{msk}; r)$ — had we run everything relative to O' instead (under the same randomness), we would have gotten the same $\mathsf{sk}[\mathbf{v}^*]$ and C, and hence $\mathrm{Dec}^{O'}(\mathsf{sk}, C_i)$ should output m_i with the same probability.

Challenge bit $b = 1$: We argue that in this case Brk^O outputs 1 with all but negligible probability. Fix a secret key sk and a set S_g. Since $b = 1$, all m_i's are chosen at random, and so the probability that $\mathrm{Dec}^{O'}(\mathsf{sk}, C_h) = m_h$ for all $h \in [t]$ is at most $\frac{1}{q^t}$. By the union bound over all sk, the probability that Brk mistakenly outputs 0 is at most $\frac{2^{\mu+\rho}}{q^t}$. By choosing t large enough, this probability will become negligible.

Enc^O making O queries. We now show how to lift the assumption, that Enc^O makes no O queries. Let us re-run the previous description of Brk (Step (i)a)

to see where it fails when Enc makes O queries. Our analysis for $b = 0$ fails: because a query made during $\mathrm{Dec}^{O'}(\mathsf{sk}[\mathbf{v}^*], C_h)$ might be one that was asked before when generating $C_h \leftarrow \mathrm{Enc}^{O}(\mathsf{msk}, \mathbf{x}_h)$, and if O' replies to it randomly, we will have an inconsistency (i.e., we cannot argue the simulated decryption $\mathrm{Dec}^{O'}(\mathsf{sk}[\mathbf{v}^*], C_h)$ outputs m_h when $b = 0$). Letting Q_h be the set Q-A pairs made during the generation of $C_h \leftarrow \mathrm{Enc}^{O}(\mathsf{msk}, \mathbf{x}_i)$, we should somehow learn all those Q-A pairs in Q_h that also appear during $\mathrm{Dec}^{O}(\mathsf{sk}[\mathbf{v}^*], C_h)$. Fix the index h below. Consider the following strategy for decrypting the hth ciphertext.

1. For all $i \in [t]$, run $\mathrm{Dec}^{O}(\mathsf{sk}[\mathbf{v}_i], C_h)$, and record all Q-A pairs in the set S_g.
2. Perform Step (i) from before.

The intuition is that if t is large enough and if a query is to appear during $\mathrm{Dec}^{O}(\mathsf{sk}[\mathbf{v}^*], C_h)$, then it should have also appeared during one of the $\mathrm{Dec}^{O}(\mathsf{sk}[\mathbf{v}_i], C_h)$ executions in Step 1, except with small probability. However, proving this leads to the following challenge: the vector $\mathbf{v}^*$ is sampled from the entire space $\mathbb{Z}_q^n$, while $\mathbf{v}_i$ vectors are sampled from an $(n-1)$-dimensional subspace, leading to $\mathsf{sk}[\mathbf{v}^*]$ having a different distribution from $\mathsf{sk}[\mathbf{v}_i]$. The above statement would have been easy to prove if all of $\mathbf{v}_i$'s were picked from the entire space $\mathbb{Z}_q^n$, but that is not the case here.

The above challenge is about a *covering* problem. Suppose ℓ queries are made during $C \leftarrow \mathrm{Enc}^{O}(\mathsf{msk}, \mathbf{x})$. (We replace C_h with C for better readability.) Let us number these queries as $1, \ldots, \ell$. Consider a function $F\colon \mathbb{Z}_q^n \to 2^{[\ell]}$, where $j \in F(\mathbf{y})$ if Query j appears during the decryption of $\mathrm{Dec}^{O}(\mathsf{sk}[\mathbf{y}], C)$, where $\mathsf{sk}[\mathbf{y}] \leftarrow \mathrm{KGen}^{O}(\mathsf{msk}, \mathbf{y})$. Here $2^{[\ell]}$ denotes the set of subsets of $[\ell]$. We would like to prove that with high probability $F(\mathbf{v}^*) \subseteq \cup_i F(\mathbf{v}_i)$, where $\mathbf{v}^*$ is the challenge secret key vector, and $\mathbf{v}_i$'s are the vectors from the $(n-1)$-dimensional subspace, whose secret keys are given to Brk.

We abstract out the above problem as a combinatorial lemma.

Lemma 2.1 (Combinatorial Lemma). *Let $n = n(\kappa) \geqslant 3$ be such that $\frac{1}{q^n}$ is negligible. Let $\ell = \ell(\kappa)$ be an arbitrary polynomial. Let $F : \mathbb{Z}_q^n \to 2^{[\ell]}$ be an arbitrary function, assigning a subset of $[\ell]$ to every vector. Then, for all large enough polynomial values of $t = t(\kappa)$, with overwhelming probability*

$$F(y^*) \subseteq \bigcup_{i=1}^{t} F(y_i), \tag{1}$$

where $y^* \leftarrow \mathbb{Z}_q^n$, and $y_1, \ldots, y_t$ are sampled as follow: sample an $(n-1)$-dimensional subspace $V \subseteq \mathbb{Z}_q^n$ uniformly at random conditioned on $y^* \notin V$ and then sample $y_1, \ldots, y_t$ uniformly at random from V.

2.1 Proof of the Combinatorial Lemma

In this overview we will focus on the case $q = 2$. In the main part of the paper, we will prove the same result for any prime q with slightly looser bounds.

A simpler problem. Before describing how to prove this result, we focus on a related but simpler problem: We show that there is no polynomial ℓ that satisfies the above condition is violated in a *worst-case sense*. In other words, there exists no polynomial ℓ and $F : \mathbb{Z}_2^n \to 2^{[\ell]}$ such that for every $\mathbf{y}^*$ and every $(n-1)$-dimensional subspace V with $\mathbf{y}^* \notin V$ we have

$$F(\mathbf{y}^*) \not\subseteq \bigcup_{\mathbf{y} \in V} F(\mathbf{y}). \tag{2}$$

We will refer in the following to the set $[\ell]$ as colors and to F as a coloring of the vectors in $\mathbb{Z}_q^n$. A simple coloring strategy that satisfies Condition (2) is to set $\ell = 2^n - 1$ and assign each non-zero vector one individual color. Another simple strategy is to also use $\ell = 2^n - 1$ and pick for every $(n-1)$-dimensional subspace one new individual color and add it to the colorings of each vector *not* contained in the subspace.

We show that using $\ell \geqslant 2^n - 1$ colors is necessary to satisfy Condition (2) by a double-counting argument on the vectors $\mathbf{y}^*$ and $(n-1)$-dimensional subspaces V with $\mathbf{y}^* \notin V$.

For a vector $\mathbf{y}^*$ and an $(n-1)$-dimensional subspace V we say the color cl is *useful* for $(\mathbf{y}^*, V)$ if $\mathbf{y}^* \notin V$, $\mathsf{cl} \in F(\mathbf{y}^*)$ and $\mathsf{cl} \notin \bigcup_{\mathbf{y} \in V} F(\mathbf{y})$. Condition (2) says for every combination of $\mathbf{y}^*$ and V we need at least one useful color.

In $\mathbb{Z}_2^n$, each $(n-1)$-dimensional subspace is uniquely defined by a non-zero vector; *i.e.*, the subspace that is orthogonal to the underlying non-zero vector. Thus, we have $2^n - 1$ different $(n-1)$-dimensional subspaces. Each subpsace V has 2^{n-1} different vectors $\mathbf{y}^*$ not contained in V. Thus, the number of triples $(V, \mathbf{y}^*, \mathsf{cl})$ where cl is useful for $(\mathbf{y}^*, V)$ is at least $(2^n - 1)2^{n-1}$. We now use another way of counting to argue the number of such triples is at most $\ell 2^{n-1}$, implying $\ell \geqslant 2^n - 1$.

Fix a color cl. We argue that the color cl is useful for at most 2^{n-1} different combinations $(\mathbf{y}^*, V)$, implying the number of such triples as above is at most $\ell 2^{n-1}$. Let $\mathcal{S}$ be the set of all the $(n-1)$-dimensional subspaces V such that there exists at least one $\mathbf{y}^\star$ with cl being useful for $(\mathbf{y}^*, V)$. Then, no vector in $U := \bigcup_{V \in \mathcal{S}} V$ has the color cl and, in the worst case, every vector in $\mathbb{Z}_2^n \setminus U$ has the color cl. Each of the subspaces $V \in \mathcal{S}$ can be described by one linear equation, i.e. for each $V \in \mathcal{S}$ we can pick a vector $\mathbf{x}_i$ such that $V = \{\mathbf{v} \mid \langle \mathbf{x}_i, \mathbf{v} \rangle = 0\}$, because V is $(n-1)$-dimensional. When we have $t := |\mathcal{S}|$ subspaces, we get at least $d \geqslant \lceil \log_2(t) \rceil$ linear independent equations. Let $\mathbf{x}_1, \ldots, \mathbf{x}_d$ be the vectors representing these linear independent equations. In order for a vector $\mathbf{y}^*$ to be in $\mathbb{Z}_2^n \setminus U$, it is necessary that

$$\langle \mathbf{y}^*, \mathbf{x}_1 \rangle \neq 0 \wedge \cdots \wedge \langle \mathbf{y}^*, \mathbf{x}_d \rangle \neq 0.$$

Since we are working over $\mathbb{Z}_2$, this is equivalent to

$$\langle \mathbf{y}^*, \mathbf{x}_1 \rangle = 1 \wedge \cdots \wedge \langle \mathbf{y}^*, \mathbf{x}_d \rangle = 1.$$

This version shows us that $\mathbb{Z}_2^n \setminus U$ is contained in an $(n-d)$-dimensional affine subspace of $\mathbb{Z}_2^n$ and thus there can be at most 2^{n-d} vectors in $\mathbb{Z}_2^n \setminus U$. Thus the number of combinations for which a color can be useful is at most $t \cdot 2^{n-d} \leqslant 2^d \cdot 2^{n-d} = 2^n$. Since we need for each of the combination at least one useful color, we need at least

$$\ell \geqslant \frac{(2^n - 1)2^{n-1}}{2^n} = \frac{(2^n - 1)}{2} \text{ colors.}$$

In the main body we enhance this argument and observe that not every combination of $\mathbf{x}_1, \ldots, \mathbf{x}_t$ is allowed. Concretely, we show that they must be a sum-free set. This can be used to improve the bound on the dimension to $d \geqslant \lceil \log_2(t) \rceil + 1$, which then implies $\ell \geqslant 2^n - 1$.

We present a formal proof for the case $\mathbb{Z}_2$ in Section A.1. A formal proof for $\mathbb{Z}_q$ is given in Section 4.

The Combinatorial Lemma over $\mathbb{Z}_2$. Next, we sketch how to generalize the worst-case arguments to the average case, as required by Lemma 2.1. In this setting we are restricted to use $\ell = \mathsf{poly}(\kappa)$ different colors and we have to argue that we will hit with noticeable probability a vector $\mathbf{y}^*$ and an $(n-1)$-dimensional subspace $V \not\ni \mathbf{y}^*$ such that every color that appears in $\mathbf{y}^*$ also appears in many vectors in V. The latter condition will ensure that if we sample $\mathbf{y}_1, \ldots, \mathbf{y}_t$ from V (for a large enough polynomial t), they will satisfy Condition (1) with high probability.

We assume here that every color in $F(\mathbf{y}^*)$ is used often in the whole space $\mathbb{Z}_2$ (concretely, in at least $2^{n-1}/\ell$ vectors). With noticeable probability (here: at least $1/2$), this is satisfied when picking $\mathbf{y}^*$ uniformly at random.

We then prove that a color that appears that often in the whole space must also appear often (concretely in $2^{n-3}/\ell$ vectors) in all but negligible many $(n-1)$-dimensional subspaces. The argument for this is a generalization of the simpler problem we described above and uses the Fourier transform for hypercubes.

We present a formal proof for the Combinatorial Lemma over $\mathbb{Z}_2$ in Section A.2.

On generalizing to any prime modulus q. Generalizing our results to arbitrary prime moduli q is non-trivial. This can be best seen when focusing on the simpler problem we described in the beginning. Recall that there we needed to count the number of vectors that are non-orthogonal to every vector in a fixed set of vectors. We used there that

$$\langle \mathbf{y}^*, \mathbf{x}_1 \rangle \neq 0 \wedge \cdots \wedge \langle \mathbf{y}^*, \mathbf{x}_d \rangle \neq 0$$

is equivalent to

$$\langle \mathbf{y}^*, \mathbf{x}_1 \rangle = 1 \wedge \cdots \wedge \langle \mathbf{y}^*, \mathbf{x}_d \rangle = 1$$

and thus the solutions are an $(n-d)$-dimensional subspace. Over $\mathbb{Z}_q$, it seems that this problem does not have a nice algebraic structure. Of course, we could use

$$\langle \mathbf{y}^*, \mathbf{x}_1 \rangle \in \mathbb{Z}_q^* \wedge \cdots \wedge \langle \mathbf{y}^*, \mathbf{x}_d \rangle \in \mathbb{Z}_q^* \tag{3}$$

instead. However, this makes the resulting bound worse by a factor of $(q-1)^d$ and the result is no longer useful. The reason the (almost) tight bound from before becomes here very loose is that many of the vectors satisfying Eq. (3) will in fact be orthogonal to one of the linear combinations of $\mathbf{x}_1, \ldots, \mathbf{x}_d$.[6] We prove results for $\mathbb{Z}_q$ that are similar to $\mathbb{Z}_2$, but use different techniques. The main technique is to use the Cauchy-Schwartz inequality to get a lower bound on the "overlap" of many $(n-1)$-dimensional subspaces, which then again allows us to

[6] This also happens with $q = 2$ for sums of $\mathbf{x}_1, \ldots, \mathbf{x}_d$ with an even number of summands. However, this costs us there only a factor of 2 and we manage to get ride of this factor by using that $\mathbf{x}_1, \ldots, \mathbf{x}_t$ has to be sum-free.

argue that if many vectors are colored in the whole subspace $\mathbb{Z}_q^n$, in almost all $(n-1)$-dimensional subspaces many vectors are colored.

We present a formal proof for the Combinatorial Lemma over $\mathbb{Z}_q$ in Section 4.

2.2 Comparison with Prior Combinatorial Lemmas

Katz and Yerukhimovich [KY09] generalize the results of Boneh et al. [BPR+08] to rule out a broader class of PE primitives from trapdoor permutations. Here is a simplified version of the combinatorial lemma of [KY09, Lemma 1].

If there exists predicates $f_1, \ldots, f_q$ together with attributes $A_1, \ldots, A_q$ such that for all i: $f_i(A_i) = 1$ but $f_{i+1}(A_i) = \ldots = f_q(A_i) = 0$, then they show an impossibility. The idea is to corrupt all they keys for $(i+1, \ldots, q)$ and use the info to decrypt for A_i. And [KY09] shows that certain predicates (e.g., IBE, broadcast encryption) satisfy this property, using the Pigeonhole principle. The above property can be established also for zero inner-product encryption (where the predicate is satisfied iff the inner product is zero). Let $f_i = (id_i, 1)$ and let $A_i = (-1, id_i)$, allowing one to rule out public-key inner-product predicate encryption from PKE.

However, for FE, the combinatorial lemma becomes much more complicated, both in terms of its description and also establishing it for a functionality. The main reason is: under PE, only certain decryptions reveal information about the plaintext (an all-or-nothing property), whereas under FE, all decryptions do. For instance, extending the above combinatorics, established for zero inner-product encryption as above, to IPFE faces the following challenge: the vector A_i will be the whole plaintext vector, and we must make sure for all i, the vector f_i is not in the span of $(f_{i+1}, \ldots, f_q)$. This will limit how large q can become, while being able to make q arbitrarily large was crucial in the arguments of [KY09]. Thus, we need to come up with a more specialized combinatorial lemma. And we cannot establish the resulting combinatorial lemma using simple combinatorial techniques anymore (*e.g.*, the pigeonhole principle as in [BPR+08, KY09]) and need more advanced tools.

3 Preliminaries

Notation. We use κ to denote the security parameter. We use $[n] := \{1, \ldots, n\}$ for $n \in \mathbb{N}$ and 2^S to denote the power set of S.

Lemma 3.1 *Let $X_1, \ldots, X_{t+1}$ be independent, Bernoulli random variables, where $\Pr[X_i = 1] = p$, for all $i \leqslant t+1$. Then*

$$\Pr[X_1 = 0 \wedge \cdots \wedge X_t = 0 \wedge X_{t+1} = 1] \leqslant \frac{1}{t} \ .$$

We give the definitions for private-key IPFE schemes relative to an oracle. As we mentioned, we consider *unrestricted* IPFE where there is no restriction on the number of possible inner-products recovered by decryption. We do not explicitly mention this point in the remainder of the paper.

A private-key IPFE scheme $\mathcal{E}^O = (\mathrm{KGen}^O, \mathrm{Enc}^O, \mathrm{Dec}^O)$ is given by three algorithms. We assume without loss of generality that a master secret key is chosen uniformly at random from $\{0,1\}^w$, for some $w := w(\kappa)$. Moreover, we assume $\mathbb{Z}_q$ is the underlying field.

- $\mathrm{KGen}^O(\mathsf{msk}, \mathbf{v})$: On input a master secret key msk and a vector $\mathbf{v}$, the key generation algorithm outputs a secret key $\mathsf{sk}[\mathbf{v}]$.
- $\mathrm{Enc}^O(\mathsf{msk}, \mathbf{x})$: On input a master secret key msk and a vector $\mathbf{x}$, the encryption algorithm outputs a ciphertext C.
- $\mathrm{Dec}^O(\mathsf{sk}[\mathbf{v}], C)$: On input a secret key $\mathsf{sk}[\mathbf{v}]$ and a ciphertext C, the decryption algorithm outputs $y \in \mathbb{Z}_q$.

We require the following properties.

Definition 3.2 (IPFE Correctness). Fix an oracle O. We say an oracle-aided IPFE $(\mathrm{KGen}^O, \mathrm{Enc}^O, \mathrm{Dec}^O)$ is ν-correct for dimension n relative to O, if for any key vector $\mathbf{v} \in \mathbb{Z}_q^n$, the following experiment outputs one with probability at least ν. Sample msk uniformly at random, $\mathsf{sk}[\mathbf{v}] \leftarrow \mathrm{KGen}^O(\mathsf{msk}, \mathbf{v})$, $\mathbf{x} \leftarrow \mathbb{Z}_q^n$ and $C \leftarrow \mathrm{Enc}^O(\mathsf{msk}, \mathbf{x})$. The experiment outputs one if $\langle \mathbf{v}, \mathbf{x} \rangle = \mathrm{Dec}^O(\mathsf{sk}[\mathbf{v}], C)$.

Next, we give a definition for selective-security for a private-key IPFE below.

Definition 3.3 (IPFE Security). We work with a selective-security definition for IPFE. Fix a dimension n. The adversary designates a challenge secret-key vector $\mathbf{v}^*$ (sent to the challenger) and the adversary can make two types of queries, as follows, but all the adversary's queries should be made non-adaptively at once. The challenger starts by sampling a master secret key msk and a challenge bit $b \leftarrow \{0, 1\}$ uniformly at random.

- Key Queries: The adversary submits a vector $\mathbf{v}$ and receives a secret key for $v \leftarrow \mathrm{KGen}(\mathsf{msk}, \mathbf{v})$.
- Inner-Product Queries: Upon calling this oracle, the challenger samples $\mathbf{w} \leftarrow \mathbb{Z}_q^n$ and returns $(\mathrm{Enc}(\mathsf{msk}, \mathbf{w}), m^*)$ to the adversary, where $m^* = \langle \mathbf{w}, \mathbf{v}^* \rangle$ if $b = 0$, and m^* is chosen freshly for each query if $b = 1$.

We say an adversary $\mathcal{A}$ is *admissible* if $\mathcal{A}$ makes all its queries non-adaptively at once, and if $\mathbf{v}^* \notin \mathsf{Span}(\mathbf{v}_1, \ldots, \mathbf{v}_t)$, where $\mathbf{v}_1, \ldots, \mathbf{v}_t$ are all the adversary's key queries. We say a private-key IPFE is selectively secure if any admissible PPT adversary $\mathcal{A}$ has at most $1/2 + \mathrm{negl}(\kappa)$ advantage in guessing the value of b.

Our IPFE definition is strictly weaker than standard ones [ABDP15], making our impossibility result stronger. Specifically, ours requires security only for random plaintext vectors, while the standard ones concern *all* vectors. Our definition is implied by the standard definitions via a simple hybrid argument. To see why it is strictly weaker, change a given scheme $\mathcal{E} = (\mathrm{KGen}, \mathrm{Enc}, \mathrm{Dec})$ meeting the standard definitions into a scheme $\mathcal{E}' = (\mathrm{KGen}', \mathrm{Enc}', \mathrm{Dec}')$ so that $\mathrm{Enc}'(\mathsf{msk}, \mathbf{e}_1)$ outputs $\mathrm{Enc}(\mathsf{msk}, \mathbf{e}_1) \| \mathbf{e}_1$, where $\mathbf{e}_1$ is the first unit vector, and Dec' is defined accordingly. The rest of the scheme remains the same. The new scheme satisfies our notion but not the standard ones. Intuitively, it satisfies ours because we only need security with respect to random plaintext vectors.

4 The Combinatorial Problem over $\mathbb{Z}_q$

Lemma 4.1 *Fix $n = n(\kappa)$ and suppose $q^{-n} \in \mathrm{negl}(\kappa)$ and $n \geqslant 3$. Let $\ell = \mathsf{poly}(\kappa)$, q be a prime number and $F : \mathbb{Z}_q^n \to 2^{[\ell]}$. Fix a constant c. Then, there exists a polynomial $t = t(\kappa)$ such that with probability at least $1 - \kappa^{-c}$*

$$F(\mathbf{y}^*) \subseteq \bigcup_{i=1}^{t} F(\mathbf{y}_i),$$

where $\mathbf{y}^* \leftarrow \mathbb{Z}_q^n$, and we sample a random $(n-1)$-dimensional subspace V subject to $\mathbf{y}^* \notin V$ and we sample $\mathbf{y}_1, \ldots, \mathbf{y}_t$ all uniformly at random from V.

The following theorem is the key tool in proving the above Lemma.

Theorem 4.2 *For any subset $S \subseteq \mathbb{Z}_q^n$ with $|S| = p \cdot q^n$, there exists at most $4q/p^2$ $(n-1)$-dimensional subspaces h of $\mathbb{Z}_q^n$ with*

$$|S \cap h| \leqslant \frac{|S|}{2q} = \frac{p}{2} q^{n-1}.$$

of Lemma 4.1 using Theorem 4.2. Fix a mapping $F : \mathbb{Z}_q^n \to 2^{[\ell]}$. We will refer to the set $[\ell]$ as the set of colors and say that F colors each vector of $\mathbb{Z}_q^n$.

For a color cl, let $F^{-1}(\mathsf{cl}) = \{\mathbf{y} \mid \mathsf{cl} \in F(\mathbf{y})\}$. We say that a color cl is p-heavy if $\left|F^{-1}(\mathsf{cl})\right| \geqslant p \cdot q^n$. For $V \subseteq \mathbb{Z}_q^n$, we say that cl is p-heavy in V if $\left|F^{-1}(\mathsf{cl}) \cap V\right| \geqslant p \cdot |V|$. We use the heaviness threshold

$$p = \frac{1}{2\ell\kappa^c}.$$

For uniformly random $\mathbf{y}^*$ the probability that $F(\mathbf{y}^*)$ contains a non-p-heavy color is less than

$$\ell \cdot \frac{1}{2\ell\kappa^c} = \frac{1}{2\kappa^c}.$$

The rest of the proof assumes all colors in $F(\mathbf{y}^*)$ are p-heavy.

Now, for a uniformly random subspace V^* conditioned on $\mathbf{y}^* \notin V^*$ we can claim that, with overwhelming probability, all colors $\mathsf{cl} \in F(\mathbf{y}^*)$ are also $(p/2)$-heavy in V^*. This is because applying Theorem 4.2 with $S := F^{-1}(\mathsf{cl})$ shows that if cl is p-heavy in the entire space, there exist at most $\frac{4q}{p^2}$ $(n-1)$-dim subspaces where cl is not $p/2$-heavy in those subspaces. Here we are using the fact that an $(n-1)$-dim subspace has q^{n-1} elements. Applying the union bound for all colors in $F(\mathbf{y}^*)$ shows that the number of $(n-1)$-dimensional subspaces S where at least one color of $F(\mathbf{y}^*)$ is not $(p/2)$-heavy in S ("bad subspaces") is at most

$$|F(\mathbf{y}^*)| \frac{4q}{p^2} \leqslant \frac{4\ell q}{p^2} = 16\ell^3 q \kappa^{2c}.$$

The number of $(n-1)$-dimensional subspaces containing $\mathbf{y}^*$ is $\frac{q^{n-1}-1}{q-1}$, because each such subspace can be identified with an $n-2$-dimensional subspace in the $n-1$-dimensional quotient space $(\mathbb{Z}_q^n)_{/\langle \mathbf{y}^* \rangle}$. Thus, the total number of $(n-1)$-dimensional subspaces not containing $\mathbf{y}^*$ is $\frac{q^n - q^{n-1}}{q-1}$. Hence, by sampling one of these subspaces uniformly at random, we hit a bad subspace with probability

$$\frac{16\ell^3 q \kappa^{2c}(q-1)}{q^n - q^{n-1}} = \frac{16\ell^3 \kappa^{2c}(q-1)}{q^{n-1} - q^{n-2}},$$

which is negligible in κ since $n \geqslant 3$ and q^{-n} is negligible.

We now analyze the probability of $F(\mathbf{y}^*) \subseteq \bigcup_{i=1}^{t} F(\mathbf{y}_i)$ when setting $t = \lceil 2\kappa/p \rceil$. In the following, the probability is taken over the choice of $\mathbf{y}^*, \mathbf{y}_1, \ldots, \mathbf{y}_t$.

$$\begin{aligned}\Pr[\forall \mathsf{cl} \in F(\mathbf{y}^*)\ \exists i \in [t] : \mathsf{cl} \in F(y_i)] &= 1 - \Pr[\exists \mathsf{cl} \in F(\mathbf{y}^*) \forall i \in [t] : \mathsf{cl} \notin F(y_i)] \\ &\geqslant 1 - \ell \max_{\mathsf{cl} \in F(\mathbf{y}^*)} \Pr[\forall i \in [t] : \mathsf{cl} \notin F(y_i)] \\ &\geqslant 1 - \ell \left(1 - \frac{p}{2}\right)^t \geqslant 1 - \ell e^{-t \cdot \frac{p}{2}} = 1 - \ell e^{-\kappa}.\end{aligned}$$

The last inequality follows from $1 - \frac{p}{2} \leqslant e^{-\frac{p}{2}}$ (the Bernoulli inequality) and taking both sides to the t-th power. □

of Theorem 4.2. Bennett proved an existence-only version of Theorem 4.2 [Ben18, Lemma 4.1]. The following proof follows Bennett's ideas.

Let H be the set of $(n-1)$-dimensional subspaces h of $\mathbb{Z}_q^n$ with

$$|S \cap h| \leqslant \frac{|S|}{2q}.$$

Fix a bijection $\phi : [q^n] \to \mathbb{Z}_q^n$ and define the vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^{(q^n)}$ via

$$\mathbf{u}_i = \begin{cases} 1 & \text{if } \phi(i) \notin S \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \mathbf{v}_i = |\{h \in H \mid \phi(i) \in h\}|.$$

For a set $Y \subseteq \mathbb{Z}_q^n$ we use χ_Y to denote the characteristic function of Y.

We get

$$\langle \mathbf{u}, \mathbf{u} \rangle = \sum_{\mathbf{x} \notin S} 1^2 = q^n - |S| = q^n(1-p)$$

$$\langle \mathbf{v}, \mathbf{v} \rangle = \sum_{\mathbf{x} \in \mathbb{Z}_q^n} \left(\sum_{h \in H} \chi_h(\mathbf{x}) \right)^2 = \sum_{h_1, h_2 \in H} \sum_{\mathbf{x} \in \mathbb{Z}_q^n} \chi_{h_1}(\mathbf{x}) \chi_{h_2}(\mathbf{x}) \leqslant q^{n-1}|H| + q^{n-2}|H|^2.$$

In the last inequality we use that $\sum_{\mathbf{x} \in \mathbb{Z}_q^n} \chi_h(\mathbf{x}) \chi_h(\mathbf{x}) = q^{n-1}$ and for $h_1 \neq h_2$ we have $\sum_{\mathbf{x} \in \mathbb{Z}_q^n} \chi_{h_1}(\mathbf{x}) \chi_{h_2}(\mathbf{x}) \leqslant q^{n-2}$.

Each $h \in H$ has by definition $|S \cap h| \leqslant \frac{|S|}{2q}$ and thus $|(\mathbb{Z}_q^n \setminus S) \cap h| = q^{n-1} - |S \cap h| \geqslant q^{n-1} - \frac{|S|}{2q}$ which gives us

$$\langle \mathbf{u}, \mathbf{v} \rangle = \sum_{\mathbf{x} \notin S, h \in H} \chi_h(\mathbf{x}) \geqslant |H| \left(q^{n-1} - \frac{|S|}{2q} \right) = |H| q^{n-1} \left(1 - \frac{p}{2} \right).$$

Applying the Cauchy-Schwartz inequality to the vectors $\mathbf{u}$ and $\mathbf{v}$ gives us $\langle \mathbf{u}, \mathbf{v} \rangle^2 \leqslant \langle \mathbf{u}, \mathbf{u} \rangle \cdot \langle \mathbf{v}, \mathbf{v} \rangle$. Plugging in the bounds from above leads to the following

inequality

$$
\begin{aligned}
& |H|^2 q^{2n-2} \left(1-\frac{p}{2}\right)^2 \leqslant q^n(1-p)(q^{n-1}|H| + q^{n-2}|H|^2) \\
\iff \quad & \frac{(1-p)}{q^{n-2}\left(1-\frac{p}{2}\right)^2} \geqslant \frac{|H|^2}{(q^{n-1}|H| + q^{n-2}|H|^2)} = \frac{|H|}{q^{n-2}(q+|H|)} \\
\iff \quad & \frac{(1-p)}{\left(1-\frac{p}{2}\right)^2} \geqslant \frac{|H|}{(q+|H|)} \\
\iff \quad & 1 - \frac{(1-p)}{\left(1-\frac{p}{2}\right)^2} \leqslant 1 - \frac{|H|}{(q+|H|)} = \frac{q}{(q+|H|)} \leqslant \frac{q}{|H|} \\
\iff \quad & |H| \leqslant \frac{q}{1-\frac{(1-p)}{\left(1-\frac{p}{2}\right)^2}} = \frac{q\left(1-\frac{p}{2}\right)^2}{\left(1-\frac{p}{2}\right)^2 - (1-p)} = \frac{4q\left(1-\frac{p}{2}\right)^2}{p^2} \\
\iff \quad & |H| \leqslant \frac{4q}{p^2} \,.
\end{aligned}
$$

□

5 Separating IPFE from OWFs

The definition below gives a procedure that allows one to sample a random vector $\mathbf{v}^*$ together with t random vectors from a random $(n-1)$-dimensional subspace of $\mathbb{Z}_q^n$ which does not span $\mathbf{v}^*$.

Definition 5.1 (Sampling Spanning Vectors). The procedure $(\mathbf{v}_1, \ldots, \mathbf{v}_t, \mathbf{v}^*) \leftarrow \text{SubSpcSamp}(\mathbb{Z}_q^n, t)$ works as follows. Sample a random vector $\mathbf{v}^* \leftarrow \mathbb{Z}_q^n$, and sample a random $(n-1)$-dimensional subspace S of $\mathbb{Z}_q^n$ subject to $\mathbf{v}^* \notin S$. Sample $\mathbf{v}_1, \ldots, \mathbf{v}_t$ uniformly at random from S. The sampling procedure of SubSpcSamp can be performed in $\mathsf{poly}(n, t, \log q)$ time.

Description of the Attack. Let $\mathcal{E}^O := (\text{KGen}^O, \text{Enc}^O, \text{Dec}^O)$ be a candidate IPFE construction for vectors in $\mathbb{Z}_q^n$. Assume without loss of generality that $\mathcal{E}^O$ is $(1-\frac{1}{2^\kappa})$-correct.[7] Our goal is to remove O queries from the decryption algorithm Dec^O, while impacting correctness and security only minimally. In order to do this, if one knows the set Q_s of all the Q-A pairs asked during the generation of a secret key $\mathsf{sk}[\mathbf{v}]$ and the set Q_e of all the Q-A pairs formed to generate a ciphertext C, then one can remove O queries from $\text{Dec}(\mathsf{sk}[\mathbf{v}], C)$ as follows: if an issued query appears in $Q_s \cup Q_e$, reply to it accordingly; else, reply with a random response, without calling O. In fact, this argument still holds if we just know the subset $Q_e \cap Q_d$, where Q_d is the set of Q-A pairs that appear during $\text{Dec}(\mathsf{sk}[\mathbf{v}], C)$.

The adversary will then decrypt all the ciphertexts obtained via the inner-product queries with all secret keys it received and store all Q-A pairs that appeared during this process in a set L. The combinatorial Lemma from the previous section guarantees that with high probability $Q_e \cap Q_d \subseteq \mathsf{L}$, if the number of keys and ciphertexts is large enough.

[7] If the scheme is $1/2 + \frac{1}{\mathsf{poly}(\kappa)}$ correct, we can boost its correctness all the way up to $(1-\frac{1}{2^\kappa})$ by encrypting the vector many times and taking the majority during decryption.

Finally, the adversary makes a brute-force search over $\mathsf{sk}[\mathbf{v}^*]$ and the set Q_s and check each candidate by decrypting all challenge ciphertexts without asking any queries.

The attack in detail.

Attack 5.2 Let $\mathcal{E}^O := (\text{Setup}, \text{KGen}^O, \text{Enc}^O, \text{Dec}^O)$ be an IPFE. We give a polynomial query adversary Brk^O which breaks the security of $\mathcal{E}^O$.

Parameters. The adversary's algorithm is based on integers t and η, instantiated later. Also, let $\mu := \mu(\kappa, n)$ be the bit size of a secret key, generated by $\text{KGen}^O(\mathsf{msk}, *)$ and $\rho := \rho(\kappa, n)$ the bit size of a the Q-A pairs of all queries made during secret key generation.[8]

Phase 1: Corrupting Keys and Setting Up the Challenge.

1. Sample $(\mathbf{v}_1, \ldots, \mathbf{v}_t, \mathbf{v}^*) \leftarrow \text{SubSpcSamp}(\mathbb{Z}_q^n, t)$. The vector $\mathbf{v}^*$ will be the challenge secret-key vector.
2. For all $i \in [t]$, make a key query $\mathbf{v}_i$ to receive $\mathsf{sk}[\mathbf{v}_i]$.
3. For all $i \in [\eta]$ make an inner-product query to receive (C_i, m_i), where recall that $m_i \in \mathbb{Z}_q$.

Phase 2: Learning Important Decryption Queries.

1. Let $\mathsf{L} := \emptyset$. For all $i \in [t]$ and all $j \in [\eta]$, execute $\text{Dec}^O(\mathsf{sk}[\mathbf{v}_i], C_j)$ and add all the Q-A pairs to the set L.

Phase 3: Leveraging the Set of Learned Q-A Pairs to Decrypt. In this phase, Brk uses the set L to break the security game. In this phase, Brk does not make any queries to O.

1. If there exists $\mathsf{sk} \in \{0,1\}^\mu$ and a set of Q-A pairs $Q_s \in \{0,1\}^\rho$ such that for all $h \in [\eta]$, $\text{Dec}^{O'}(\mathsf{sk}, C_i) = m_i$, where O' is a random oracle sampled uniformly at random on-the-fly subject to being consistent with L and Q_s, return 0; else, return 1.

We now show how to set the parameters, and then discuss the effectiveness of the attack.

Parameters 5.3 (Setting the parameters.). Set the parameters as follows.

- Set η such that $\eta \geqslant \kappa + \mu + \rho$.
- Let ℓ be the number of queries made by $\text{Enc}^O(\mathsf{msk}, \cdot)$; i.e., the number of queries made to generate each of $C_1, \ldots, C_\eta$ in the attack above. Choose a constant c such that $\kappa^c \geqslant 2\eta\kappa$. Choose t based on ℓ and κ^{-c} as per Lemma 4.1.

[8] The assumption that the length of a secret key and and the Q-A pairs is a fixed function of κ and n is without loss of generality.

Lemma 5.4 *Let* Brk^O *be as in Attack 5.2, and let* b' *be the output of* Brk^O. *Suppose* $n \geqslant 3$ *and* $q^n \in \omega(\mathsf{poly}(\kappa))$. *We have* $\Pr[b' = 0 \mid b = 0] \geqslant 1 - \frac{1}{\kappa}$.

Proof. First, recall that $\mathcal{E}^O$ has correctness $1 - \frac{1}{2^\kappa}$ (c.f. the footnote of Page 15). Let $\mathsf{sk}[\mathbf{v}^*]$ be the secret key for $\mathbf{v}^*$ relative to the real oracle O, namely $\mathsf{sk}[\mathbf{v}^*] \leftarrow \mathrm{KGen}^O(\mathsf{msk}, \mathbf{v}^*)$. By correctness of $\mathcal{E}^O$, for each $h \in [\eta]$, with probability at least $1 - \frac{1}{2^\kappa}$, $\mathrm{Dec}^O(\mathsf{sk}[\mathbf{v}^*], C_h) = m_h$. We claim that performing the decryption relative to O' (as opposed to the real oracle O) does not impact the decryption result much. In particular, for any $h \in [\eta]$,

$$\Pr[\mathrm{Dec}^{O'}(\mathsf{sk}[\mathbf{v}^*], C_h) = m_h] \geqslant 1 - \frac{1}{2^\kappa} - \kappa^{-c}. \tag{4}$$

Thus, the probability that Brk mistakenly outputs one when $b = 0$ is at most $\eta(\frac{1}{2^\kappa} + \kappa^{-c})$. This is because Brk goes through all choices of vector secret keys and Q-A pairs, hitting $\mathsf{sk}[\mathbf{v}^*]$ and Q_s at some point. The reason that the multiplicative factor η appears is that we require for all $h \in [\eta]$, the decryption result be m_h. (Line 1 of Phase 3 of Brk's procedure.) Since by Parameters 5.3, $\kappa^c \geqslant 2\eta\kappa$

$$\Pr[b' = 1 \mid b = 0] \leqslant \eta\left(\frac{1}{2^\kappa} + \kappa^{-c}\right) \leqslant \eta\left(\frac{1}{2^\kappa} + \frac{1}{2\eta\kappa}\right) \leqslant \frac{1}{\kappa}, \tag{5}$$

for all large enough κ.

To argue about Equation 4, fix $h \in [\eta]$. Let S_0 and S_1 be the set of Q-A pairs made during the generation of $\mathsf{sk}[\mathbf{v}^*]$ and during the generation of $C_h \leftarrow \mathrm{Enc}^O(\mathsf{msk}, \mathbf{x}_h)$. Define the event Bad as follows.

- Bad: the event that a query in $\mathsf{S}_1 \setminus \mathsf{L}$ is asked during the decryption of $\mathrm{Dec}^O(\mathsf{sk}[\mathbf{v}^*], C_h)$, where recall that the set L is defined in Step 1 of Phase 2 of Brk's procedure.

If $Q_s = \mathsf{S}_0$ and $\overline{\mathsf{Bad}}$ holds, then the decryption execution of $\mathrm{Dec}^{O'}(\mathsf{sk}[\mathbf{v}^*], C_h)$ proceeds identically to that of $\mathrm{Dec}^O(\mathsf{sk}[\mathbf{v}^*], C_h)$. Thus, we show $\Pr[\mathsf{Bad}] \leqslant \kappa^{-c}$, which implies

$$\begin{aligned}\Pr[\mathrm{Dec}^{O'}(\mathsf{sk}[\mathbf{v}^*], C_h) = m_h] &\geqslant \Pr[\mathrm{Dec}^O(\mathsf{sk}[\mathbf{v}^*], C_h) = m_h] - \kappa^{-c} \\ &\geqslant 1 - 2^{-\kappa} - \kappa^{-c},\end{aligned}$$

as desired.

By Lemma 4.1 we obtain $\Pr[\mathsf{Bad}] \leqslant \kappa^{-c}$. To see this, set ℓ to be the number of queries in S_1. (Parameters 5.3.) Name these queries as $1, \ldots, \ell$. Let $F\colon \mathbb{Z}_q^n \to 2^{[\ell]}$ be a function, where $F(\mathbf{y})$ contains those queries in $[\ell]$ that appear during the decryption of $\mathrm{Dec}^O(\mathsf{sk}[\mathbf{y}], C_h)$, where $\mathsf{sk}[\mathbf{y}] \leftarrow \mathrm{KGen}^O(\mathsf{msk}, \mathbf{y})$. Now by Lemma 4.1 the set L contains all the queries in $\cup_{i \in [t]} F(\mathbf{v}_i)$, except with probability at most $\frac{1}{\kappa^c}$. The proof is now complete. □

Lemma 5.5 *Let* Brk^O *be as in Attack 5.2, and let* b' *be the output of* Brk^O. *Then,* $\Pr[b' = 1 \mid b = 1] \geqslant 1 - 2^{\mu+\rho-\eta}$. *By setting the parameters as in Parameters 5.3 (specifically that* $\eta \geqslant \kappa + \mu + \rho$*),* $\Pr[b' = 1 \mid b = 1] \geqslant 1 - 2^\kappa$.

Proof. Fix a vector secret key $\mathsf{sk} \in \{0,1\}^\mu$ and a set of Q-A pairs $Q_s \in \{0,1\}^\rho$. For any $i \in [\eta]$, the probability that $\mathrm{Dec}^{O'}(\mathsf{sk}, C_i) = m_i$ is at most $1/2$ because the righthand side is completely independent of the lefthand side (because $b = 1$; see Definition 3.3). Thus, the probability that for all $i \in [\eta]$, $\mathrm{Dec}^{O'}(\mathsf{sk}, C_i) = m_i$ is at most $\frac{1}{2^\eta}$. Doing a union bound over all vector secret keys $\mathsf{sk} \in \{0,1\}^\mu$ and all Q-A pairs $Q_s \in \{0,1\}^\rho$, the probability that Brk outputs zero is at most $2^{\mu+\rho-\eta}$. The proof is now complete. □

Putting together the above lemmas, we achieve the final impossibility result.

Theorem 5.6 *Suppose $n \geqslant 3$ and q^n is super-polynomial in the security parameter. There exists no black-box construction of IPFE for dimension n and modulus q from any primitive that exists relative to a random oracle.*

The condition of q^n being super-polynomial in κ in the above theorem is necessary. For example, there exists trivial black-box IPFE constructions for $\mathbb{Z}_2^n$ from OWFs, where $n = \log \kappa$, as follows. Let the master secret key be a PRF key, and let the secret key for a vector be the PRF output for that vector. When encrypting a plaintext vector $\mathbf{x}$ under msk, encrypt the result of $\langle \mathbf{x}, \mathbf{v} \rangle$ under the corresponding secret key for $\mathbf{v}$, for all $\mathbf{v}$. The size of the ciphertext remains polynomial.

Acknowledgments. We thank the anonymous reviewers for their helpful comments, in particular for pointing out a technical issue in an earlier draft of this work and suggesting how to resolve it.

Mohammad Hajiabadi was supported in part by an NSERC Discovery Grant 03270, and a Meta Research Award. Part of the work was done while Mingyuan Wang was at UC Berkeley. Mingyuan Wang is supported in part by an NYU startup fund.

References

[ABDP15] Michel Abdalla, Florian Bourse, Angelo De Caro, and David Pointcheval. Simple functional encryption schemes for inner products. In Jonathan Katz, editor, *PKC 2015: 18th International Conference on Theory and Practice of Public Key Cryptography*, volume 9020 of *Lecture Notes in Computer Science*, pages 733–751, Gaithersburg, MD, USA, March 30 – April 1, 2015. Springer, Berlin, Heidelberg, Germany. doi:10.1007/978-3-662-46447-2_33. 2, 3, 5, 12

[ABG19] Michel Abdalla, Fabrice Benhamouda, and Romain Gay. From single-input to multi-client inner-product functional encryption. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019, Part III*, volume 11923 of *Lecture Notes in Computer Science*, pages 552–582, Kobe, Japan, December 8–12, 2019. Springer, Cham, Switzerland. doi:10.1007/978-3-030-34618-8_19. 3

[ACF+18] Michel Abdalla, Dario Catalano, Dario Fiore, Romain Gay, and Bogdan Ursu. Multi-input functional encryption for inner products:

Function-hiding realizations and constructions without pairings. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 597–627, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Cham, Switzerland. doi:10.1007/978-3-319-96884-1_20. 3

[ACF+20] Shweta Agrawal, Michael Clear, Ophir Frieder, Sanjam Garg, Adam O'Neill, and Justin Thaler. Ad hoc multi-input functional encryption. In Thomas Vidick, editor, *ITCS 2020: 11th Innovations in Theoretical Computer Science Conference*, volume 151, pages 40:1–40:41, Seattle, WA, USA, January 12–14, 2020. LIPIcs. doi:10.4230/LIPIcs.ITCS.2020.40. 3

[ACGU20] Michel Abdalla, Dario Catalano, Romain Gay, and Bogdan Ursu. Inner-product functional encryption with fine-grained access control. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2020, Part III*, volume 12493 of *Lecture Notes in Computer Science*, pages 467–497, Daejeon, South Korea, December 7–11, 2020. Springer, Cham, Switzerland. doi:10.1007/978-3-030-64840-4_16. 3

[AGRW17] Michel Abdalla, Romain Gay, Mariana Raykova, and Hoeteck Wee. Multi-input inner-product functional encryption from pairings. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017, Part I*, volume 10210 of *Lecture Notes in Computer Science*, pages 601–626, Paris, France, April 30 – May 4, 2017. Springer, Cham, Switzerland. doi:10.1007/978-3-319-56620-7_21. 3

[AGW20] Michel Abdalla, Junqing Gong, and Hoeteck Wee. Functional encryption for attribute-weighted sums from k-Lin. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020, Part I*, volume 12170 of *Lecture Notes in Computer Science*, pages 685–716, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Cham, Switzerland. doi:10.1007/978-3-030-56784-2_23. 2

[AJ15] Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *Advances in Cryptology – CRYPTO 2015, Part I*, volume 9215 of *Lecture Notes in Computer Science*, pages 308–326, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Berlin, Heidelberg, Germany. doi:10.1007/978-3-662-47989-6_15. 2

[ALMT20] Shweta Agrawal, Benoît Libert, Monosij Maitra, and Radu Titiu. Adaptive simulation security for inner product functional encryption. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020: 23rd International Conference on Theory and Practice of Public Key Cryptography, Part I*, volume 12110 of *Lecture Notes in Computer Science*, pages 34–64, Edinburgh,

UK, May 4–7, 2020. Springer, Cham, Switzerland. doi:10.1007/978-3-030-45374-9_2. 3

[ALS16] Shweta Agrawal, Benoît Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016, Part III*, volume 9816 of *Lecture Notes in Computer Science*, pages 333–362, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Berlin, Heidelberg, Germany. doi:10.1007/978-3-662-53015-3_12. 3

[AV19] Prabhanjan Ananth and Vinod Vaikuntanathan. Optimal bounded-collusion secure functional encryption. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019: 17th Theory of Cryptography Conference, Part I*, volume 11891 of *Lecture Notes in Computer Science*, pages 174–198, Nuremberg, Germany, December 1–5, 2019. Springer, Cham, Switzerland. doi:10.1007/978-3-030-36030-6_8. 2

[BCFG17] Carmen Elisabetta Zaira Baltico, Dario Catalano, Dario Fiore, and Romain Gay. Practical functional encryption for quadratic functions with applications to predicate encryption. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 67–98, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Cham, Switzerland. doi:10.1007/978-3-319-63688-7_3. 2, 3

[Ben18] Michael Bennett. Occurrence of right angles in vector spaces over finite fields. *European Journal of Combinatorics*, 70:155–163, 2018. doi:10.1016/j.ejc.2017.12.005. 14

[BF01] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. In Joe Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 213–229, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Berlin, Heidelberg, Germany. doi:10.1007/3-540-44647-8_13. 2

[BJK15] Allison Bishop, Abhishek Jain, and Lucas Kowalczyk. Function-hiding inner product encryption. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology – ASIACRYPT 2015, Part I*, volume 9452 of *Lecture Notes in Computer Science*, pages 470–491, Auckland, New Zealand, November 30 – December 3, 2015. Springer, Berlin, Heidelberg, Germany. doi:10.1007/978-3-662-48797-6_20. 3, 4

[BPR+08] Dan Boneh, Periklis A. Papakonstantinou, Charles Rackoff, Yevgeniy Vahlis, and Brent Waters. On the impossibility of basing identity based encryption on trapdoor permutations. In *49th Annual Symposium on Foundations of Computer Science*, pages 283–292, Philadelphia, PA, USA, October 25–28, 2008. IEEE Computer Society Press. doi:10.1109/FOCS.2008.67. 2, 3, 4, 5, 11

[BRS02] John Black, Phillip Rogaway, and Thomas Shrimpton. Black-box analysis of the block-cipher-based hash-function constructions from PGV. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 320–335, Santa Barbara, CA, USA, August 18–22, 2002. Springer, Berlin, Heidelberg, Germany. doi:10.1007/3-540-45708-9_21. 2

[BSW11] Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In Yuval Ishai, editor, *TCC 2011: 8th Theory of Cryptography Conference*, volume 6597 of *Lecture Notes in Computer Science*, pages 253–273, Providence, RI, USA, March 28–30, 2011. Springer, Berlin, Heidelberg, Germany. doi:10.1007/978-3-642-19571-6_16. 1

[BV15] Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. In Venkatesan Guruswami, editor, *56th Annual Symposium on Foundations of Computer Science*, pages 171–190, Berkeley, CA, USA, October 17–20, 2015. IEEE Computer Society Press. doi:10.1109/FOCS.2015.20. 2

[CL01] Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In Birgit Pfitzmann, editor, *Advances in Cryptology – EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 93–118, Innsbruck, Austria, May 6–10, 2001. Springer, Berlin, Heidelberg, Germany. doi:10.1007/3-540-44987-6_7. 2

[CLT18] Guilhem Castagnos, Fabien Laguillaumie, and Ida Tucker. Practical fully secure unrestricted inner product functional encryption modulo p. In Thomas Peyrin and Steven Galbraith, editors, *Advances in Cryptology – ASIACRYPT 2018, Part II*, volume 11273 of *Lecture Notes in Computer Science*, pages 733–764, Brisbane, Queensland, Australia, December 2–6, 2018. Springer, Cham, Switzerland. doi:10.1007/978-3-030-03329-3_25. 3

[Gay20] Romain Gay. A new paradigm for public-key functional encryption for degree-2 polynomials. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020: 23rd International Conference on Theory and Practice of Public Key Cryptography, Part I*, volume 12110 of *Lecture Notes in Computer Science*, pages 95–120, Edinburgh, UK, May 4–7, 2020. Springer, Cham, Switzerland. doi:10.1007/978-3-030-45374-9_4. 3

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st Annual ACM Symposium on Theory of Computing*, pages 169–178, Bethesda, MD, USA, May 31 – June 2, 2009. ACM Press. doi:10.1145/1536414.1536440. 2

[GGM84] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions (extended abstract). In *25th Annual Symposium on Foundations of Computer Science*, pages 464–479, Singer

Island, Florida, October 24–26, 1984. IEEE Computer Society Press. doi:10.1109/SFCS.1984.715949. 2

[GKLM12] Vipul Goyal, Virendra Kumar, Satyanarayana V. Lokam, and Mohammad Mahmoody. On black-box reductions between predicate encryption schemes. In Ronald Cramer, editor, *TCC 2012: 9th Theory of Cryptography Conference*, volume 7194 of *Lecture Notes in Computer Science*, pages 440–457, Taormina, Sicily, Italy, March 19–21, 2012. Springer, Berlin, Heidelberg, Germany. doi:10.1007/978-3-642-28914-9_25. 3, 4, 5

[GVW12] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 162–179, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Berlin, Heidelberg, Germany. doi:10.1007/978-3-642-32009-5_11. 2

[IR89] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *21st Annual ACM Symposium on Theory of Computing*, pages 44–61, Seattle, WA, USA, May 15–17, 1989. ACM Press. doi:10.1145/73007.73012. 4

[KSW13] Jonathan Katz, Amit Sahai, and Brent Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. *Journal of Cryptology*, 26(2):191–224, April 2013. doi:10.1007/s00145-012-9119-4. 2

[KY09] Jonathan Katz and Arkady Yerukhimovich. On black-box constructions of predicate encryption from trapdoor permutations. In Mitsuru Matsui, editor, *Advances in Cryptology – ASIACRYPT 2009*, volume 5912 of *Lecture Notes in Computer Science*, pages 197–213, Tokyo, Japan, December 6–10, 2009. Springer, Berlin, Heidelberg, Germany. doi:10.1007/978-3-642-10366-7_12. 3, 4, 5, 11

[O'N10] Adam O'Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010. URL: https://eprint.iacr.org/2010/556. 1

[PRV12] Periklis A. Papakonstantinou, Charles W. Rackoff, and Yevgeniy Vahlis. How powerful are the DDH hard groups? Cryptology ePrint Archive, Report 2012/653, 2012. URL: https://eprint.iacr.org/2012/653. 2, 5

[Rot11] Ron Rothblum. Homomorphic encryption: From private-key to public-key. In Yuval Ishai, editor, *TCC 2011: 8th Theory of Cryptography Conference*, volume 6597 of *Lecture Notes in Computer Science*, pages 219–234, Providence, RI, USA, March 28–30, 2011. Springer, Berlin, Heidelberg, Germany. doi:10.1007/978-3-642-19571-6_14. 2

[SS10] Amit Sahai and Hakan Seyalioglu. Worry-free encryption: functional encryption with public keys. In Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov, editors, *ACM CCS 2010: 17th Conference on Computer and Communications Security*, pages 463–472, Chicago, Illinois, USA, October 4–8, 2010. ACM Press. doi:10.1145/1866307.1866359. 2

[SS21] Gili Schul-Ganz and Gil Segev. Generic-group identity-based encryption: A tight impossibility result. In Stefano Tessaro, editor, *2nd Conference on Information-Theoretic Cryptography, ITC 2021, July 23-26, 2021, Virtual Conference*, volume 199 of *LIPIcs*, pages 26:1–26:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. 2, 5

[SW05] Amit Sahai and Brent R. Waters. Fuzzy identity-based encryption. In Ronald Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 457–473, Aarhus, Denmark, May 22–26, 2005. Springer, Berlin, Heidelberg, Germany. doi:10.1007/11426639_27. 1, 2, 5

[TÜ24] Erkan Tairi and Akin Ünal. Lower bounds for lattice-based compact functional encryption. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024, Part II*, volume 14652 of *Lecture Notes in Computer Science*, pages 249–279, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland. doi:10.1007/978-3-031-58723-8_9. 4

[Üna20] Akin Ünal. Impossibility results for lattice-based functional encryption schemes. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 169–199, Zagreb, Croatia, May 10–14, 2020. Springer, Cham, Switzerland. doi:10.1007/978-3-030-45721-1_7. 4

[Yap69] H. P. Yap. Maximal sum-free sets of group elements. *Journal of the London Mathematical Society*, s1-44(1):131–136, 1969. doi:10.1112/jlms/s1-44.1.131. 24

[Zha22] Mark Zhandry. Augmented random oracles. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part III*, volume 13509 of *Lecture Notes in Computer Science*, pages 35–65, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland. doi:10.1007/978-3-031-15982-4_2. 2, 5

A The Combinatorial Problem over $\mathbb{Z}_2$

Proving the combinatorial problem over $\mathbb{Z}_2$ is significantly easier than over $\mathbb{Z}_q$ for any prime q. Thus, we present it separately for $\mathbb{Z}_2$.

A.1 Existential version

As a warm-up, we prove a weaker statement that just establishes the existence of vectors $\mathbf{y}^*, \mathbf{y}_1, \dots \mathbf{y}_t$ with the desired properties.

Lemma A.1 *Fix $n = n(\kappa) \in \omega(\log \kappa)$. For $\ell = \mathsf{poly}(\kappa)$ and any map $F : \mathbb{Z}_2^n \to 2^{[\ell]}$ there exists $t = \mathsf{poly}'(\kappa)$, $\mathbf{y}_1, \dots, \mathbf{y}_t, \mathbf{y}^* \in \mathbb{Z}_2^n$ with*

$$\mathbf{y}^* \notin \mathsf{Span}\,(\mathbf{y}_1, \dots, \mathbf{y}_t) \quad \text{and} \quad F(\mathbf{y}^*) \subseteq \bigcup_{i=1}^{t} F(\mathbf{y}_i)$$

when κ is sufficiently large.

The proof uses the following result of [Yap69] on the maximal size of sum-free sets in finite groups. A subset S of a finite group is sum free, if there exist no $a, b, c \in S$ with $a + b = c$.

Lemma A.2 *[Yap69] The maximal size of a sum-free subset $S \subseteq \mathbb{Z}_2^n, n \geqslant 1$ is 2^{n-1}.*

Yap [Yap69] proves a bound for arbitrary finite abelian groups. We use this to prove the following theorem, which is the existence-only analog of Thm. A.5.

Theorem A.3 *For every non-empty subset $S \subseteq \mathbb{Z}_2^n, n \geqslant 1$, there are at most $2^{n-1}/|S|$ subspaces $V \subseteq \mathbb{Z}_2^n$ of dimension $n-1$ with $V \cap S = \emptyset$.*

Proof. Let t be the number of subspaces of dimension $(n-1)$ that do not contain any of the vectors in S. Each such subspace V can be described by a vector $\mathbf{v}$ as follows: $V = \{\mathbf{v} \mid \langle \mathbf{x}, \mathbf{v} \rangle = 0\}$. So let $\mathbf{x}_1, \dots, \mathbf{x}_t$ be the vectors describing all the subspaces V with $V \cap S = \emptyset$. No subset of 3 of these vectors can be colinear: Assume there exists indices $i, j, k \in [t]$ with $\mathbf{x}_i + \mathbf{x}_j = \mathbf{x}_k$. Then for every vector $\mathbf{v} \in \mathbb{Z}_2^n$ we have $\langle \mathbf{x}_i, \mathbf{v} \rangle = 0$ or $\langle \mathbf{x}_j, \mathbf{v} \rangle = 0$ or ($\langle \mathbf{x}_i, \mathbf{v} \rangle = 1$ and $\langle \mathbf{x}_j, \mathbf{v} \rangle = 1$). In the last case $\langle \mathbf{x}_k, \mathbf{v} \rangle = \langle (\mathbf{x}_i + \mathbf{x}_j), \mathbf{v} \rangle = 0$. This means every vector $\mathbf{v}$ is contained in the subspace associated to $\mathbf{x}_i$, $\mathbf{x}_j$ or $\mathbf{x}_k$ and thus S would have to be empty. This is a contradiction.

Lemma A.2 implies that the vectors $\mathbf{x}_1, \dots, \mathbf{x}_t$ must contain at least $d := \lceil \log(t+1) \rceil$ linear independent vectors. Let this be without loss of generality $\mathbf{x}_1, \dots, \mathbf{x}_d$. Clearly, every vector $\mathbf{v} \in S$ must satisfy $\langle \mathbf{x}_i, \mathbf{v} \rangle = 1$. At most $2^{n-d} \leqslant 2^{n-(\log(t)+1)} = 2^{n-1}/t$ can satisfy this equation. Thus $|S| \leqslant 2^{n-1}/t$ which can be rearranged to $t \leqslant 2^{n-1}/|S|$. □

The proof uses two facts that are specific to $\mathbb{Z}_2$:

1. The union of three subspaces, where each one is orthogonal to one of the three vectors $\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k$ with $\mathbf{x}_i + \mathbf{x}_j = x_k$, is all of $\mathbb{Z}_2$.

2. For fixed vectors $\mathbf{x}_1, \dots, \mathbf{x}_t$ the set

$$\{\mathbf{v} \mid \forall i \in [t] : \langle \mathbf{v}, \mathbf{x}_k \rangle \neq 0\}$$

   is an affine subspace of $\mathbb{Z}_2^n$.

In particular that the set is in the second point seems to be algebraically much more complicated over $\mathbb{Z}_q$.

We next give a proof of Thm. A.1 using Thm. A.3. This part can be easily generalized to $\mathbb{Z}_q$.

Proof of Thm. A.1. First of all, note that the requirement that t is polynomial is not necessary here: If there exists $\mathbf{y}^*$ and an $(n-1)$-dimensional subspace V^* with $\mathbf{y}^* \notin V^*$ and $F(\mathbf{y}^*) \subseteq \bigcup_{i=1}^{t} F(\mathbf{y}_i)$, we can set $t := |F(\mathbf{y}^*)| \leqslant m$ and pick for each color $\mathsf{cl} \in F(\mathbf{y}^*)$ a vector $\mathbf{y}_i$.

We prove the Lemma with the double counting technique.

We count the number triples $(\mathsf{cl}, \mathbf{y}^*, V) \in [\ell] \times \mathbb{Z}_q^n \times 2^{\mathbb{Z}_q^n}$ such that

1. $\mathbf{y}^* \notin V$,
2. V is an $(n-1)$-dimensional subspace,
3. $\mathsf{cl} \in F(\mathbf{y}^*)$, and $\mathsf{cl} \notin \bigcup_{\mathbf{y} \in V} F(\mathbf{y})$.

We say that a triple is valid, if it satisfies all of the above conditions. To satisfy the condition in the lemma, for every $\mathbf{y}^*$ and V that satisfy 1 and 2, there must be at least one valid triple. We have $2^n - 1$ choices for a non-zero vector $\mathbf{y}^*$ and then 2^{n-1} choices for V such that 1 and 2 are satisfied. Thus there must be at least $(2^n - 1) \cdot 2^{n-1}$ valid triples.

On the other hand, if a color is used for s vectors, by Thm. A.3 there are at most $2^{n-1}/s$ $(n-1)$-dimensional subspaces not containing any vector having this color. Thus, for a fixed color there can be at most 2^{n-1} contributions.

This leads to the following inequality

$$\ell 2^{n-1} \geqslant (2^n - 1) \cdot 2^{n-1} \quad \iff \quad \ell \geqslant (2^n - 1).$$

Since ℓ grows only polynomial in κ but 2^n grows exponential in κ, this inequality can not hold for sufficiently large κ. □

A.2 Probabilistic version

We now give the proof for the probabilistic version over $\mathbb{Z}_2^n$. The proof uses Fourier transforms, but the techniques do not seem to extend beyond $q > 2$.

To prove our impossibility result, we need to strengthen the theorem in two ways:

1. The vector $\mathbf{y}^*$ needs to be uniformly random.
2. The vectors $\mathbf{y}_1, \ldots, \mathbf{y}_t$ have to be efficiently sampleable *without* knowing F.

This is formalized by the following lemma:

Lemma A.4 *Fix $n = n(\kappa) \in \omega(\log \kappa)$. Let $\ell = \mathsf{poly}(\kappa)$. Fix a constant c. Then, there exists a polynomial $t = t(\kappa)$ such that with probability at least $1 - n^{-c}$*

$$F(\mathbf{y}^*) \subseteq \bigcup_{i=1}^{t} F(\mathbf{y}_i),$$

where $\mathbf{y}^* \leftarrow \mathbb{Z}_2^n$, and we sample a random $(n-1)$-dimensional subspace V subject to $\mathbf{v}^* \notin S$ and we sample $\mathbf{y}_1, \ldots, \mathbf{y}_t$ all uniformly at random from V.

This lemma can be proven as Thm. 4.1, but we can replace Thm. 4.2 (that is used for the proof) with the following version that has an easier proof and gives a slightly better bound at the cost of being specific for $\mathbb{Z}_2$.

Theorem A.5 *For any subset $S \subseteq \mathbb{Z}_2^n$, there exist at most $\frac{2^{n+2}}{|S|}$ linear subspaces V of dimension $n-1$ such that*

$$|S \cap V| \leqslant \frac{1}{2} \cdot \frac{|S|}{2}.$$

To prove Theorem A.5 we recall the Fourier transform for Boolean hypercubes.

Definition A.6 (Fourier Transform). For any function $f : \mathbb{Z}_2^n \to \mathbb{R}$, its Fourier coefficient $\widehat{f}(\mathbf{u})$ for any $\mathbf{u} \in \mathbb{Z}_2^n$ is defined as

$$\widehat{f}(\mathbf{u}) = \frac{1}{2^n} \cdot \sum_{\mathbf{x} \in \mathbb{Z}_2^n} (-1)^{\langle \mathbf{x}, \mathbf{u} \rangle} \cdot f(\mathbf{x}).$$

Theorem A.7 (Parseval's Identity).

$$\sum_{\mathbf{x} \in \mathbb{Z}_2^n} f(\mathbf{x})^2 = 2^n \cdot \sum_{\mathbf{u} \in \mathbb{Z}_2^n} \widehat{f}(\mathbf{u})^2.$$

of Theorem A.7.

$$\begin{aligned}
\sum_{\mathbf{u} \in \mathbb{Z}_2^n} \widehat{f}(\mathbf{u})^2 &= \sum_{\mathbf{u} \in \mathbb{Z}_2^n} \left(\frac{1}{2^n} \cdot \sum_{\mathbf{x} \in \mathbb{Z}_2^n} (-1)^{\langle x, \mathbf{u} \rangle} \cdot f(x) \right)^2 \\
&= \frac{1}{2^{2n}} \sum_{\mathbf{u}, \mathbf{x}, \mathbf{y} \in \mathbb{Z}_2^n} (-1)^{\langle \mathbf{x}, \mathbf{u} \rangle} \cdot (-1)^{\langle \mathbf{y}, \mathbf{u} \rangle} \cdot f(\mathbf{x}) \cdot f(\mathbf{y}) \\
&= \frac{1}{2^{2n}} \sum_{\mathbf{x}, \mathbf{y} \in \mathbb{Z}_2^n} f(\mathbf{x}) \cdot f(\mathbf{y}) \cdot \sum_{\mathbf{u} \in \mathbb{Z}_2^n} (-1)^{\langle \mathbf{x}+\mathbf{y}, \mathbf{u} \rangle} \\
&= \frac{1}{2^{2n}} \sum_{\mathbf{x} \in \mathbb{Z}_2^n} f(\mathbf{x})^2 \cdot 2^n = \frac{1}{2^n} \sum_{\mathbf{x} \in \mathbb{Z}_2^n} f(\mathbf{x})^2
\end{aligned}$$

□

We are now ready for the proof.

of Theorem A.5. Let f be the indicator function for S. We know

$$\sum_{\mathbf{x} \in \mathbb{Z}_2^n} f(x)^2 = |S|.$$

By Theorem A.7,

$$\sum_{\mathbf{u} \in \mathbb{Z}_2^n} \widehat{f}(\mathbf{u})^2 = \frac{|S|}{2^n}.$$

Observe that,

$$\sum_{\mathbf{u} \neq 0^n} \widehat{f}(\mathbf{u})^2 = \frac{|S|}{2^n} - \left(\frac{|S|}{2^n} \right)^2 \leqslant \frac{|S|}{2^n}.$$

For every non-zero $\mathbf{u}$, let $V_{\mathbf{u}}$ denote the subspace orthogonal to $\mathbf{u}$. Observe that,

$$\widehat{f}(\mathbf{u}) = \frac{|S \cap V_{\mathbf{u}}| - |S \setminus V_{\mathbf{u}}|}{2^n} = \frac{2 \cdot |S \cap V_{\mathbf{u}}| - |S|}{2^n}.$$

Hence,

$$|S \cap V_{\mathbf{u}}| \leqslant \frac{1}{2} \cdot \frac{|S|}{2} \implies \widehat{f}(\mathbf{u})^2 \geqslant \left(\frac{|S|}{2^{n+1}}\right)^2.$$

The number of such $\mathbf{u}$ can be upper bounded by

$$\frac{2^{n+2}}{|S|}.$$

□