

Constrained Key-Homomorphic PRFs from Standard Lattice Assumptions

Or: How to Secretly Embed a Circuit in Your PRF

Zvika Brakerski*

Vinod Vaikuntanathan†

Abstract

Boneh et al. (Crypto 13) and Banerjee and Peikert (Crypto 14) constructed pseudorandom functions (PRFs) from the Learning with Errors (LWE) assumption by embedding combinatorial objects, a path and a tree respectively, in instances of the LWE problem. In this work, we show how to generalize this approach to embed *circuits*, inspired by recent progress in the study of Attribute Based Encryption.

Embedding a universal circuit for some class of functions allows us to produce *constrained keys* for functions in this class, which gives us the first standard-lattice-assumption-based constrained PRF (CPRF) for general bounded-description bounded-depth functions, for arbitrary polynomial bounds on the description size and the depth. (A constrained key w.r.t a circuit C enables one to evaluate the PRF on all x for which $C(x) = 1$, but reveals nothing on the PRF values at other points.) We rely on the LWE assumption and on the one-dimensional SIS (Short Integer Solution) assumption, which are both related to the worst case hardness of general lattice problems. Previous constructions for similar function classes relied on such exotic assumptions as the existence of multilinear maps or secure program obfuscation. The main drawback of our construction is that it does not allow collusion (i.e. to provide more than a single constrained key to an adversary). Similarly to the aforementioned previous works, our PRF family is also *key homomorphic*.

Interestingly, our constrained keys are very short. Their length does not depend directly either on the size of the constraint circuit or on the input length. We are not aware of any prior construction achieving this property, even relying on strong assumptions such as indistinguishability obfuscation.

1 Introduction

A pseudorandom function family (PRF) [GGM86] is a finite set of functions $\{F_s : D \rightarrow R\}_s$, indexed by a seed (or key) s , such that for a random s , F_s is efficiently computable given s , and is computationally indistinguishable from a random function from D to R , given oracle access. Since the introduction of this concept, PRFs have been one of the most fundamental building blocks in cryptography. Many variants of PRFs with additional properties have been introduced and have found a plethora of applications in cryptography. In this work, we will focus on *Constrained PRFs* and *Key-Homomorphic PRFs*.

*Weizmann Institute of Science, Rehovot, Israel. Email: zvika.brakerski@weizmann.ac.il. Supported by ISF grant 468/14, and by an Alon Young Faculty Fellowship.

†MIT, Cambridge, MA, USA. Email: vinodv@csail.mit.edu. Research supported by DARPA Grant number FA8750-11-2-0225, Alfred P. Sloan Research Fellowship, NSF CAREER Award CNS-1350619, NSF Frontier Grant CNS-1414119, Microsoft Faculty Fellowship, and a Steven and Renee Finn Career Development Chair from MIT.

Constrained PRFs. Constrained PRFs (CPRFs) have been introduced simultaneously by Boneh and Waters [BW13], Kiayias et al. [KPTZ13] (as “Delegatable PRFs”) and by Boyle, Goldwasser and Ivan [BGI14] (as “Functional PRFs”). Here an adversary is allowed to ask for a constrained key which should allow it to evaluate the PRF on a subset of the inputs, while revealing nothing about the values at other inputs. It has been shown [BW13, KPTZ13, BGI14] how to construct CPRFs for function classes of the form $x \in [i, j]$ (where the input is interpreted as an integer) based on any one-way function. This in particular allows for the “puncturing” technique of Sahai and Waters [SW14] that found many uses in the obfuscation literature. Further, [BW13] showed how to achieve more complicated function classes such as bit fixing functions and even arbitrary circuits, but those require use of cryptographic multilinear maps. They also introduce a number of applications for such CPRFs, including broadcast encryption schemes and identity based key exchange. Hofheinz et al. [HKKW14] show how to achieve adaptively secure CPRFs from indistinguishability obfuscation using a random oracle.

The original definition of CPRFs requires resilience to arbitrary collusion. Namely, a constrained key for C_1, C_2 should give no more information than a constrained key for $C_1 \vee C_2$ and must not reveal anything about values where $C_1(x) = C_2(x) = \text{false}$. Many of the applications of CPRFs (e.g. for broadcast encryption and identity based key exchange) rely on collusion resilience. Unfortunately, our construction in this work will not allow collusions, and therefore will not be useful for these applications. We hope that future works will be able to leverage our ideas into collusion resilient CPRFs.

Key-Homomorphic PRFs. In key-homomorphic PRFs, there is a group structure associated with the set of keys, and it is required that for any input x and keys s, t , $F_s(x) + F_t(x) = F_{s+t}(x)$. A construction in the random oracle model was given by Naor, Pinkas and Reingold [NPR99], and the first construction in the standard model was given by Boneh et al. [BLMR13] based on the Learning with Errors assumption (LWE), building on a (non key homomorphic) lattice-based PRF of Banerjee, Peikert and Rosen [BPR12]. This was followed by an improved construction by Banerjee and Peikert [BP14] based on quantitatively better lattice assumptions. The LWE based constructions achieved a slightly weaker notion, namely “almost” key-homomorphism, in which $\|(F_s(x) + F_t(x)) - F_{s+t}(x)\|$ is small, for an appropriately defined norm. This notion is sufficient for the known applications. Applications of key-homomorphic PRFs include distributed key-distribution, symmetric proxy re-encryption, updatable encryption and PRFs secure against related-key attacks [NPR99, BLMR13, LMR14].

Our Results. We view the main contribution of this work as showing how to impose *hidden semantics* into the evaluation process of LWE-based PRFs. Namely, we allow multiple computation paths for computing $F_s(x)$, such that we can selectively block some of these paths based on logic described by a circuit. This is done by extending ideas from the ABE literature, and in particular the ABE scheme of Boneh et al. [BGG⁺14] (see more about this connection below).

It is particularly interesting that previous constructions of PRFs [BLMR13, BP14] can be viewed as a special case of our framework, but ones that only allow a single computational path. Our work therefore highlights that the techniques used for constructing PRFs and for constructing ABE are special cases of the same grand schema. This could hopefully lead to new insights and constructions.

We employ our methods towards presenting a family of (single key secure) constrained key-

homomorphic PRFs based on worst-case general lattice assumptions. This is a first step in solving the open problem posed in [BW13] of achieving (collusion resilient) CPRFs from standard assumptions.

Our construction is selectively secure in the constraint query, namely the adversary needs to decide on the constraint before seeing the public parameters, but is adaptive with regards to PRF oracle queries. We achieve the latter without “complexity leveraging”, contrary to [BW13], and thus we do not require sub-exponential hardness assumptions as they do. This is done by employing our technique of embedding semantics into the evaluation process again. In particular, we embed the semantics of an *admissible hash function*, introduced by Boneh and Boyen [BB04] into the PRF, which allows us to handle adaptive queries.

Our proofs rely on two closely related hardness assumptions: The Learning with Errors (LWE) assumption, and the one-dimensional Short Integer Solution (1D-SIS) assumption. Both assumptions can be tied to the worst case hardness of general lattice problems such as GapSVP and SIVP, with similar parameters. LWE is sufficient for proving pseudorandomness in the absence of a constrained key. However, once the adversary is given a constrained key, the situation becomes more delicate. In particular, even showing *correctness* in this setting is not straightforward. (Correctness refers to the property that evaluation using the constrained key and using the actual seed result in the same output.) One can show unconditionally that the value computed using the constrained key is *close* (in norm) to the real value of the function but not that they are always equal. A similar issue comes up in the security proof (since the reduction “fabricates” oracle answers in a similar way to the constrained evaluation). Our solution is to use *computational* arguments. Namely to show that it is computationally intractable, under the 1D-SIS assumption, to come up with an input for which the constrained evaluation errs. Therefore even the correctness of our scheme relies on computational assumptions. We note that similar techniques can be used to strengthen the almost key-homomorphism property into computational key-homomorphism where it is computationally hard to find an input for which key homomorphism does not hold.

The following theorem presents the simplest application of our method, we explain how it can be extended below.

Theorem 1.1. *Let $\mathcal{C}_{\ell,d}$ be the class of size- ℓ depth- d circuits. Then for all polynomials ℓ, d , there exists a $\mathcal{C}_{\ell,d}$ -constrained (almost) key-homomorphic family of PRFs without collusion, based on the (appropriately parameterized) LWE and 1D-SIS assumptions (and hence on the worst-case hardness of appropriately parameterized GapSVP and SIVP problems).*

Interestingly, we can go beyond bounded size circuits. In fact, we can support any function family with bounded length description, so long as there is a universal evaluator of depth d that takes a function description and an input, and executes the function on the input. Namely, consider a sequence of universal circuits $\{\mathcal{U}_k\}_{k \in \mathbb{N}}$, where $\mathcal{U}_k : \{0, 1\}^\ell \times \{0, 1\}^k \rightarrow \{0, 1\}$. This sequence defines a class of functions $\{0, 1\}^* \rightarrow \{0, 1\}$, where each function F in the class is represented by a string $f \in \{0, 1\}^\ell$, and for $x \in \{0, 1\}^k$, it holds that $F(x) = \mathcal{U}_k(f, x)$. We call such a function class ℓ -uniform. We are only able to support \mathcal{U}_k whose depth is bounded by some a-priori polynomial in the security parameter d , however in some cases this is sufficient to support all k 's that are polynomial in the security parameter. The following theorem states our result with regards to such families.

Theorem 1.2. *Let $\mathcal{C}_{\ell,d}$ be a class of ℓ -uniform functions with depth- d evaluator. Then for all polynomials ℓ, d , there exists a $\mathcal{C}_{\ell,d}$ -constrained (almost) key-homomorphic family of PRFs without*

collusion, based on the (appropriately parameterized) LWE and 1D-SIS assumptions (and hence on the worst case hardness of appropriately parameterized GapSVP, SIVP).

Lastly, we show that the bit-length of the constrained keys in our scheme can be reduced to $\text{poly}(\lambda)$ for some *fixed* polynomial. Namely, completely independent of all of the parameters of the scheme. This is done by using an ABE scheme with short secret keys as a black box. In particular we resort to the same scheme, namely the ABE scheme of Boneh et al. [BGG⁺14], which inspired our constrained PRF construction. This is done by encrypting all of the “components” of the constrained key, and providing them in the public parameters of the construction. Then, the actual constrained key is an ABE secret key which only allows to decrypt the relevant components. We note that this short representation for constrained keys is not homomorphic (however the scheme is still almost key homomorphic with respect to the seed). A theorem statement follows.

Theorem 1.3. *There exists a constrained PRF scheme with the same properties as in Theorem 1.2, and under the same hardness assumptions, where the constrained keys are of asymptotic bit-length $\text{poly}(\lambda)$, for an a-priori fixed polynomial.*

See Section 2 for an extended overview of the construction.

Relation to the ABE Construction of Boneh et al. [BGG⁺14]. Our techniques are greatly influenced by the aforementioned LWE-based ABE construction of Boneh et al. [BGG⁺14]. Recall that in ABE, messages are encrypted relative to *attributes* and decryption keys are drawn relative to *functions*. Decryption is possible only if the function f of the decryption key accepts the attribute x of the ciphertext. In order to decrypt a ciphertext, [BGG⁺14] first applies a public procedure that depends on f, x on the ciphertext and then applies the decryption key on the resulting value. Their construction makes sure that for any f , encryptions with regards to all accepting x 's will derive a decryptable ciphertext (and all non-accepting x 's cannot be decrypted).

Our constrained key for a circuit C is almost identical to an encryption of 0 with attribute C in [BGG⁺14]. The randomness in the encryption roughly corresponds to the seed of the PRF. An application of the PRF on the constrained key includes applying the public procedure of the ABE on the ciphertext, with respect to the function $f = \mathcal{U}$, the universal circuit for the function class to which C belongs. However, there is the question of how to represent the input: We need to be able to evaluate C on any possible input while preserving security. One of our main technical ideas is in showing that this is possible, and in fact can be achieved regardless of the input length. Combined with the framework from [BGG⁺14], we can guarantee that for all x , regardless which C was used to generate the “ciphertext”, the output of the public procedure will only depend on x and not on C . The basic idea is therefore to use this value as the PRF value. This does not work as is (for example, it does not imply pseudorandomness for non-accepting x 's) and additional ideas are required.

As mentioned above, the PRFs of [BLMR13, BP14] that seem to stem from different ideas and have quite different proofs than [BGG⁺14] can be shown to be special cases of the above paradigm, except f is taken to be an arbitrary formula (a multiplication tree). For details see Section 2.

The novelty in our approach is to show the extra power that is obtained from generalizing these two approaches. We use the universal circuit as a way to embed an undisclosed computation into an LWE instance, and show how to achieve pseudorandomness using tools such as admissible hash functions (which are also embedded into an LWE instance).

Relation with the Constrained PRF of Hofheinz et al. [HKKW14]. The work of [HKKW14] constructs adaptively secure collusion-resistant CPRFs, namely ones where the challenge x^* needs not be provided ahead of time. Their building blocks are “universal parameters” and *adaptively secure* ABE, which are used as black-box. Note that we achieve adaptive security w.r.t the challenge (but not with respect to the constraint) while relying on techniques which are only known to imply *selectively secure* ABE. Further, whereas [HKKW14] use ABE only to implement access control and therefore need to rely on strong assumptions to implement the PRF so as to interface with the ABE, we use ABE techniques to achieve both pseudorandomness and access control. On the flip side, our construction is not collusion resistant, contrary to [HKKW14].

Open Problems. The main drawback of our CPRF is its vulnerability to collusion, which severely limits its applicability as a building block. It is an open problem to achieve bounded collusion resilience, even for two constrained keys instead of one and even at the cost of increasing the parameters. Any improvement on this front should be very interesting. Another avenue for research is trying to extend the construction so that there is no restriction on the constraint circuit size, similarly to the multilinear map based construction of [BW13]. Finally, it would also be interesting to apply this methodology of imposing semantics on a cryptographic computation to other primitives in order to allow more fine-grained access control.

2 Overview of Our Construction

We recall that the LWE assumption asserts that for a uniform vector \mathbf{s} and a matrix \mathbf{A} of appropriate dimensions (over \mathbb{Z}_q for an appropriate q), it holds that $(\mathbf{A}, \mathbf{s}^T \mathbf{A} + \mathbf{e}^T)$, is indistinguishable from uniform, where \mathbf{e} is taken from an appropriate distribution over *low norm* vectors and referred to as the *noise vector*. In this outline we will ignore the generation of \mathbf{e}^T and its evolution during computation process, and just denote it by *noise* (but of course care will need to be taken in the formal arguments).

The PRF of Banerjee and Peikert [BP14]. A high-level methodology for constructing PRFs, taken by [BLMR13, BP14] and also in this work, is to take \mathbf{s} as the seed, and to generate for each PRF input x , an LWE matrix \mathbf{A}_x such that the values $\mathbf{s}^T \mathbf{A}_x + \text{noise}$ for the different inputs x are jointly indistinguishable from uniform. Note that almost key homomorphism follows naturally for any implementation of this template, up to the accumulation of noise. The noise issue is handled by taking the PRF value to be a properly scaled down and rounded version of the above, so that the effect of the noise is minimal (and its norm can be bounded below 1). This property is also inherited by our scheme.

As a starting point for deriving our construction, let us revisit the key-homomorphic PRF construction of [BP14]. Their PRF family was associated with a combinatorial object – a binary tree. Each node v of the tree was associated with an LWE matrix \mathbf{A}_v , where the PRF input x determined the matrices for the leaves, and matrices for internal nodes are derived as follows. Given a node v whose children are associated with $\mathbf{A}_l, \mathbf{A}_r$, they define $\mathbf{A}_v = \mathbf{A}_l \cdot \mathbf{G}^{-1}(\mathbf{A}_r)$. In this notation, $\mathbf{G}^{-1}(\cdot)$ is the binary decomposition operator, which breaks each entry in the matrix into the bit vector of length $\log(q)$ of its binary representation. Note that $\mathbf{G}^{-1}(\cdot)$ will always have small norm, and that the inverse operator \mathbf{G} , representing binary composition, is linear so it can be represented by a matrix. Thus for all \mathbf{A} it holds that $\mathbf{G} \cdot \mathbf{G}^{-1}(\mathbf{A}) = \mathbf{A}$.

Going back to the PRF of [BP14], the derivation procedure described above allows to associate a matrix with the root of the tree, which depends only on the input x (and on the topology of the tree which is fixed). We will use the root's matrix as our \mathbf{A}_x . The proof hinges on the invariant that LWE instances will be multiplied on the right only by low-norm matrices (of the form $\mathbf{G}^{-1}(\cdot)$), and therefore $\mathbf{s}^T \mathbf{A}_l \mathbf{G}^{-1}(\mathbf{A}_r) + \text{noise} \approx (\mathbf{s}^T \mathbf{A}_l + \text{noise}) \mathbf{G}^{-1}(\mathbf{A}_r)$, which allows to replace $(\mathbf{s}^T \mathbf{A}_l + \text{noise})$ with a new uniform vector and propagate to the right.

From Embedded Trees to Embedded Circuits. We show that the operation $\mathbf{A}_v = \mathbf{A}_l \cdot \mathbf{G}^{-1}(\mathbf{A}_r)$ is in fact a special case of a more general operation, inspired by the recent Attribute Based Encryption (ABE) construction of Boneh et al. [BGG⁺14]. We will associate a matrix \mathbf{A}_v as well as a binary value x_v with each node, and pay special attention to the matrix $(\mathbf{A}_v - x_v \mathbf{G})$. In particular, considering a node v with children l, r , it holds that

$$(\mathbf{A}_l - x_l \mathbf{G}) \cdot \mathbf{G}^{-1}(\mathbf{A}_r) + (\mathbf{A}_r - x_r \mathbf{G}) \cdot x_l = \mathbf{A}_l \mathbf{G}^{-1}(\mathbf{A}_r) - x_r x_l \mathbf{G} .$$

This generalization associates the semantics of the multiplication operation with the syntactic definition $\mathbf{A}_v = \mathbf{A}_l \mathbf{G}^{-1}(\mathbf{A}_r)$, and it also maintains the invariant that the matrices $(\mathbf{A}_l - x_l \mathbf{G})$ and $(\mathbf{A}_r - x_r \mathbf{G})$ are only multiplied on the right by low norm elements, so that

$$\begin{aligned} \mathbf{s}^T \left((\mathbf{A}_l - x_l \mathbf{G}) \cdot \mathbf{G}^{-1}(\mathbf{A}_r) + (\mathbf{A}_r - x_r \mathbf{G}) \cdot x_l \right) + \text{noise} \approx \\ \left(\mathbf{s}^T (\mathbf{A}_l - x_l \mathbf{G}) + \text{noise} \right) \cdot \mathbf{G}^{-1}(\mathbf{A}_r) + \left(\mathbf{s}^T (\mathbf{A}_r - x_r \mathbf{G}) + \text{noise} \right) \cdot x_l , \end{aligned}$$

which will play an important role in the security proof. Put explicitly, if the evaluator holds $\mathbf{s}^T (\mathbf{A}_l - x_l \mathbf{G}) + \text{noise}$ and $\mathbf{s}^T (\mathbf{A}_r - x_r \mathbf{G}) + \text{noise}$, then it can compute $\mathbf{s}^T (\mathbf{A}_v - x_l \cdot x_r \mathbf{G}) + \text{noise}$ (and we will obviously define $x_v = x_l \cdot x_r$).

This semantic relation can be extended beyond multiplication gates, and in particular NAND gates can be supported in a fairly similar manner. Furthermore, there is no need to stick to tree structure and one can support arbitrary DAGs, which naturally correspond to circuits. Extending the above postulate, if our DAG corresponds to a circuit C , then having $\mathbf{s}^T (\mathbf{A}_i - x_i \mathbf{G}) + \text{noise}$, for all leaves (= inputs), allows to compute $\mathbf{s}^T (\mathbf{A}_x - C(x) \mathbf{G}) + \text{noise}$. Recalling that the value of the PRF on input x is $\mathbf{s}^T \mathbf{A}_x + \text{noise}$, the aforementioned information allows us to evaluate the PRF at points where $C(x) = 0$. It can also be shown that it is computationally hard to compute the value at points where $C(x) = 1$. We note that this process is practically identical to the public part of the decryption procedure in the [BGG⁺14] ABE (as we explained in Section 1). We also note that since [BP14] were trying to minimize the complexity of evaluating their PRF, it made no sense in their construction to consider DAGs which only increase the complexity. However, as we show here, there are benefits to embedding a computational process in the PRF evaluation.

Utilizing the Universal Circuit. The tools we describe so far indeed seem to get us closer to our goal of producing constrained keys, but we are still not quite there. What we showed is that for any circuit C , we can devise a PRF with a constrained key for C . Note that we use the negated definition to the one we used before, and allow to evaluate when $C(x) = 0$ and not when $C(x) = 1$. This will be our convention throughout this overview.

In order to reverse the order of quantifiers, we take C to be the universal circuit $\mathcal{U}(F, x)$, and the constrained keys will be of the form $\mathbf{s}^T (\mathbf{A}_i - f_i \mathbf{G}) + \text{noise}$, where the f_i is the i th bit of the description

of the constraint F , as well as values for the x wires, which will be of the form $\mathbf{s}^T(\hat{\mathbf{A}}_b - b\mathbf{G}) + \text{noise}$, for both $b \in \{0, 1\}$. These values will allow us to execute F on *any* input x . Note that we can use the same matrices $\hat{\mathbf{A}}_0, \hat{\mathbf{A}}_1$ for all input wires, hence we don't need to commit to the input size when we provide the constrained key.¹ From this description it is obvious why our construction is not collusion resistant: Given two constrained keys for two non identical functions, there exists an i such that the adversary gets both $\mathbf{s}^T\mathbf{A} + \text{noise}$ and $\mathbf{s}^T(\mathbf{A}_i - \mathbf{G}) + \text{noise}$. Recovering \mathbf{s}^T from these values is straightforward and hence all security is lost. Note that for the input values, unlike the function description, we use two different matrices for 0 and 1: $\hat{\mathbf{A}}_0, \hat{\mathbf{A}}_1$, so a similar problem does not occur.

The Problem with Correctness, and a Computational Solution. We introduced two ways to compute the value of the PRF at x : One is to compute \mathbf{A}_x and use the seed \mathbf{s}^T to compute $\mathbf{s}^T\mathbf{A}_x + \text{noise}$, and the other is to use the constrained key to obtain $\mathbf{s}^T(\mathbf{A}_x - F(x)\mathbf{G}) + \text{noise}$, which for $F(x) = 0$ gives $\mathbf{s}^T\mathbf{A}_x + \text{noise}$. The problem is that the *noise value* in these two methods could differ. It is possible to make the difference small by scaling down and rounding, but this is not going to suffice for our purposes (mostly because a similar problem comes up in the security proof). We solve this issue using the 1D-SIS assumption as follows. We first note that the evaluation using the constrained key is essentially evaluation of a linear function with small coefficients on the vectors constituting the constrained key (essentially they get multiplied by bits and by low norm matrices $\mathbf{G}^{-1}(\cdot)$). Secondly, the only way for the two computation paths to not agree is if the value $\mathbf{s}^T\mathbf{A}_x$ is very close to an integer multiple of a number p (which is part of the PRF description). Finally, we notice that by LWE, the vectors in the constrained key are indistinguishable from uniform and independent. Thus, if we encounter such x for which correctness does not work, we can also find a short linear combination of random elements whose scaled down rounded value is close to an integer. In other words, given a uniform vector \mathbf{v} in \mathbb{Z}_q , we can find \mathbf{z} such that $\lfloor \langle \mathbf{v}, \mathbf{z} \rangle / p \rfloor$ is “close” to an integer. This is similar to solving a one-dimensional instance of the SIS problem, i.e. $\langle \mathbf{v}, \mathbf{z} \rangle = 0 \pmod{p}$. Indeed, one can show that the 1D-SIS problem is as hard as standard worst-case hard lattice problems via a reduction from [Reg04].

Pseudorandomness and Adaptive Security. Given a constrained key for F , one can compute $\mathbf{s}^T(\mathbf{A}_x - F(x)\mathbf{G}) + \text{noise}$, and indeed if $F(x) = 1$ it is hard to compute $\text{PRF}_{\mathbf{s}}(x) = \mathbf{s}^T\mathbf{A}_x + \text{noise}$. However, we want to argue that this value is *pseudorandom* and furthermore that it remains pseudorandom after adaptive queries to the PRF. Namely, after the adversary sees as many values of the form $\text{PRF}_{\mathbf{s}}(x) = \mathbf{s}^T\mathbf{A}_x + \text{noise}$ as it wishes.

To achieve these goals, we add another feature to the PRF. We consider a new independent LWE matrix \mathbf{D} , and define $\text{PRF}_{\mathbf{s}}(x) = \mathbf{s}^T\mathbf{A}_x \cdot \mathbf{G}^{-1}(\mathbf{D}) + \text{noise}$. First of all, we note that given the constrained key, we can still compute the PRF for values where $C(x) = 0$, by first computing $(\mathbf{s}^T\mathbf{A}_x + \text{noise})$ as before, and then multiplying by $\mathbf{G}^{-1}(\mathbf{D})$, which has low norm. However, in general we have

$$\text{PRF}_{\mathbf{s}}(x) \approx \left(\mathbf{s}^T(\mathbf{A}_x - F(x)\mathbf{G}) + \text{noise} \right) \cdot \mathbf{G}^{-1}(\mathbf{D}) + F(x) \left(\mathbf{s}^T\mathbf{D} + \text{noise} \right),$$

and it can be shown that for $F(x) = 1$, the second term randomizes the expression, by the LWE assumption.

¹Recall that in [BLMR13, BP14] there are only two matrices altogether. This is sufficient here for the input wires for the same reason, but we need additional matrices to encode the constraint description.

This handles pseudorandomness for a single query, but not for the case of adaptive queries (since we can only use the pseudorandomness of $(\mathbf{s}^T \mathbf{D} + \text{noise})$ once). To handle adaptive queries we embed semantics into the matrix \mathbf{D} itself. Namely, $\mathbf{D} = \mathbf{D}_x$ will be derived by an application of the universal circuit to the input x and an *admissible hash function* h . Admissible hash functions, introduced by Boneh and Boyen [BB04], allow (at a very high level) to partition the input space such that with noticeable probability all of the adaptive queries have value $h(x) = 0$, but the challenge query will have $h(x) = 1$. This means that in the proof of security, we can hold a constrained key for h , which will allow us to compute $(\mathbf{s}^T \mathbf{D}_x + \text{noise})$, for all the queries of the adversary, but leave the challenge query unpredictable (to make it pseudorandom, we will multiply in the end by another final \mathbf{D}'). This concludes the security argument for adaptive queries.

Key-Homomorphism. As we mention above, key-homomorphism follows since we use the template $\text{PRF}_{\mathbf{s}}(x) = \mathbf{s}^T \mathbf{A}_x + \text{noise}$. We note that the existence of noise means that homomorphism may not be accurate and with some low probability $(\text{PRF}_{\mathbf{s}}(x) + \text{PRF}_{\mathbf{s}'}(x))$ will only be close to $\text{PRF}_{\mathbf{s}+\mathbf{s}'}(x)$ and not identical. However this property is sufficient for many applications.

We point out that our constrained keys are a collection elements of the form $(\mathbf{s}^T \mathbf{A}_i + \text{noise})$, and therefore the scheme is also homomorphic with respect to constrained keys, i.e. constrained keys for the same F w.r.t different keys \mathbf{s}, \mathbf{s}' can be added to obtain a constrained key w.r.t $\mathbf{s} + \mathbf{s}'$.

Reducing the Constrained Key Size. From the above, it follows that the constrained key contains $\ell + 2$ vectors, where ℓ is the bit length of a description of F relative to the universal circuit for the function class. Note that this does not depend directly on the input size to the function. However, indirectly the depth of the universal circuit affects the modulus q that needs to be used.

We show that we can remove the dependence on ℓ altogether using an ABE scheme with short secret keys, such as that of [BGG⁺14]. To do this, we notice that for each constraint function F , the adversary gets either $\mathbf{s}^T \mathbf{A}_i + \text{noise}$ or $\mathbf{s}^T (\mathbf{A}_i - \mathbf{G}) + \text{noise}$, according to the value of the bit f_i . We can prepare for both options by encrypting both vectors using the ABE, each with its own attribute $(i, 0)$ and $(i, 1)$ respectively. All of these encryptions, for all i , will be placed in the public parameters. Then in order to provide a constrained key, we will provide an ABE secret key for the function that takes (i, b) and returns 0 if and only if $f_i = b$. Given this key, the user can decrypt exactly those vectors that constitute its constrained key. Note that this function can be computed by a depth $O(\log(\ell)) = O(\log(\lambda))$ circuit, and thus the size of the secret key can be made asymptotically independent of all parameters except λ , e.g. by setting the parameters to support depth $\log^2(\lambda)$ circuits.

3 Preliminaries

We first recall some background. For an integer modulus q , let $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ denote the ring of integers modulo q . For an integer $p \leq q$, we define the modular “rounding” function

$$\lfloor \cdot \rfloor_p : \mathbb{Z}_q \rightarrow \mathbb{Z}_p \text{ that maps } x \rightarrow \lfloor (p/q) \cdot x \rfloor$$

and extend it coordinate-wise to matrices and vectors over \mathbb{Z}_q . We denote the elements of the standard basis by $\mathbf{u}_1, \mathbf{u}_2, \dots$, where the dimension will be clear from the context.

We denote distributions (or random variables) that are computationally indistinguishable by $X \stackrel{c}{\approx} Y$. This refers to the standard notion of negligible distinguishing gap for any polynomial

time distinguisher. Our reductions preserve the uniformity of the adversary so by assuming the hardness of our assumption for uniform adversary we get security for our construction against uniform adversaries, and likewise for non-uniform assumptions and adversaries.

The Gadget Matrix. Let $\ell = \lceil \log q \rceil$ and define the “gadget matrix” $\mathbf{G} = \mathbf{g} \otimes \mathbf{I}_n \in \mathbb{Z}_q^{n \times n\ell}$ where

$$\mathbf{g} = (1, 2, 4, \dots, 2^{\ell-1}) \in \mathbb{Z}_q^\ell$$

We will also refer to this gadget matrix as the “powers-of-two” matrix. We define the inverse function $\mathbf{G}^{-1} : \mathbb{Z}_q^{n\ell \times m} \rightarrow \{0, 1\}^{n\ell \times m}$ which expands each entry $a \in \mathbb{Z}_q$ of the input matrix into a column of size ℓ consisting of the bit decomposition of a . We have the property that for any matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$,

$$\mathbf{G} \cdot \mathbf{G}^{-1}(\mathbf{A}) = \mathbf{A}$$

Norms for Vectors and Matrices. We will always use the infinity norm for vectors and matrices. Namely for a vector \mathbf{x} , the norm $\|\mathbf{x}\|$ is the maximal absolute value of an element in \mathbf{x} . Similarly, for a matrix \mathbf{A} , $\|\mathbf{A}\|$ is the maximal absolute value of any of its entries. If \mathbf{x} is n -dimensional and \mathbf{A} is $n \times m$, then $\|\mathbf{x}^T \mathbf{A}\| \leq n \cdot \|\mathbf{x}\| \cdot \|\mathbf{A}\|$. We remark that L_1 or L_2 norms can also be used and even achieve somewhat tighter parameters, but the proofs become more complicated.

3.1 Constrained Pseudorandom Function: Definition

In a constrained PRF family [BW13, BGI14, KPTZ13], one can compute a constrained PRF key K_C corresponding to any Boolean circuit C . Given K_C , anyone can compute the PRF on inputs x such that $C(x) = 0$. Furthermore, K_C does not reveal any information about the PRF values at the other locations. Below we recall their definition, as given by [BW13].

Syntax A *constrained* pseudo-random function (PRF) family is defined by a tuple of algorithms (KeyGen, Eval, Constrain, ConstrainEval) where:

- **Key Generation** $\text{KeyGen}(1^\lambda, 1^{k_{\text{in}}}, 1^{k_{\text{out}}})$ is a PPT algorithm that takes as input the security parameter λ , an input length k_{in} and an output length k_{out} , and outputs a PRF key K ;
- **Evaluation** $\text{Eval}(K, x)$ is a deterministic algorithm that takes as input a key K , a string $x \in \{0, 1\}^{k_{\text{in}}}$ and outputs $y \in \{0, 1\}^{k_{\text{out}}}$;
- **Constrained Key Generation** $\text{Constrain}(K, C)$ is a PPT algorithm that takes as input a PRF key K , a circuit $C : \{0, 1\}^{k_{\text{in}}} \rightarrow \{0, 1\}$ and outputs a constrained key K_C ;
- **Constrained Evaluation** $\text{ConstrainEval}(K_C, x)$ is a deterministic algorithm that takes as input a constrained key K_C and a string $x \in \{0, 1\}^{k_{\text{in}}}$ and outputs either a string $y \in \{0, 1\}^{k_{\text{out}}}$ or \perp .

We define the notion of (*single key*) *selective-function* security for constrained PRFs.

Definition 3.1. A family of PRFs (KeyGen, Eval, Constrain, ConstrainEval) is a single-key selective-function constrained PRF (henceforth, referred to simply as *constrained PRF*) if it satisfies the following properties:

- **Functionality computationally preserved under constraining.** For every PPT adversary (A_0, A_1) , consider an experiment where we choose $K \leftarrow \text{KeyGen}(1^\lambda, 1^{k_{\text{in}}}, 1^{k_{\text{out}}})$, $(C, \sigma_0) \leftarrow A_0(1^\lambda)$, and $K_C \leftarrow \text{Constrain}(K, C)$. Then:

$$\Pr \left[x^* \leftarrow A_1^{\text{Eval}(K, \cdot)}(1^\lambda, K_C, \sigma_0); \quad : \quad C(x^*) = 0 \wedge \text{Eval}(K, x^*) \neq \text{ConstrainEval}(K_C, x^*) \right]$$

is negligible in the security parameter, where C, K, K_C are selected as described above.

In words, it is computationally hard to find an x^* such that $C(x^*) = 0$, and yet the result of the constrained evaluation differs from the actual PRF evaluation.

- **Pseudorandom at constrained points.** For every PPT adversary (A_0, A_1, A_2) , consider an experiment where $K \leftarrow \text{KeyGen}(1^\lambda, 1^{k_{\text{in}}}, 1^{k_{\text{out}}})$, $(C, \sigma_0) \leftarrow A_0(1^\lambda)$, and $K_C \leftarrow \text{Constrain}(K, C)$. Then:

$$\Pr \left[\begin{array}{l} b \leftarrow \{0, 1\}; \\ (x^*, \sigma_1) \leftarrow A_1^{\text{Eval}(K, \cdot)}(1^\lambda, K_C, \sigma_0); \\ \text{If } b = 0, y^* = \text{Eval}(K, x^*), \\ \text{Else } y^* \leftarrow \{0, 1\}^{k_{\text{out}}} \end{array} ; \quad \begin{array}{l} C(x^*) = 1 \wedge \\ A_2(1^\lambda, y^*, \sigma_1) = b \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

The correctness and security properties could potentially be combined into one game, but we choose to present them as two distinct properties for the sake of clarity.

3.2 Learning with Errors

The Learning with Errors (LWE) problem was introduced by Regev [Reg05] as a generalization of “learning parity with noise” [BFKL93, Ale03]. We now define the decisional version of LWE. (Unless otherwise stated, we will treat all vectors as column vectors in this paper).

Definition 3.2 (Decisional LWE (DLWE) [Reg05]). Let λ be the security parameter, $n = n(\lambda)$, $m = m(\lambda)$, and $q = q(\lambda)$ be integers and $\chi = \chi(\lambda)$ be a probability distribution over \mathbb{Z} . The DLWE $_{n,q,\chi}$ problem states that for all $m = \text{poly}(n)$, letting $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow \chi^m$, and $\mathbf{u} \leftarrow \mathbb{Z}_q^m$, the following distributions are computationally indistinguishable:

$$(\mathbf{A}, \mathbf{s}^T \mathbf{A} + \mathbf{e}^T) \stackrel{c}{\approx} (\mathbf{A}, \mathbf{u}^T)$$

There are known quantum (Regev [Reg05]) and classical (Peikert [Pei09]) reductions between DLWE $_{n,q,\chi}$ and approximating short vector problems in lattices. Specifically, these reductions take χ to be a discrete Gaussian distribution $D_{\mathbb{Z},\alpha,q}$ for some $\alpha < 1$. We write DLWE $_{n,q,\alpha}$ to indicate this instantiation. We now state a corollary of the results of [Reg05, Pei09, MM11, MP12]. These results also extend to additional forms of q (see [MM11, MP12]).

Corollary 3.1 ([Reg05, Pei09, MM11, MP12]). Let $q = q(n) \in \mathbb{N}$ be either a prime power $q = p^r$, or a product of co-prime numbers $q = \prod q_i$ such that for all i , $q_i = \text{poly}(n)$, and let $\alpha \geq \sqrt{n}/q$. If there is an efficient algorithm that solves the (average-case) DLWE $_{n,q,\alpha}$ problem, then:

- There is an efficient quantum algorithm that solves $\text{GapSVP}_{\tilde{O}(n/\alpha)}$ (and $\text{SIVP}_{\tilde{O}(n/\alpha)}$) on any n -dimensional lattice.

- If in addition $q \geq \tilde{O}(2^{n/2})$, there is an efficient classical algorithm for $\text{GapSVP}_{\tilde{O}(n/\alpha)}$ on any n -dimensional lattice.

Recall that GapSVP_γ is the (promise) problem of distinguishing, given a basis for a lattice and a parameter d , between the case where the lattice has a vector shorter than d , and the case where the lattice doesn't have any vector shorter than $\gamma \cdot d$. SIVP is the search problem of finding a set of “short” vectors. The best known algorithms for GapSVP_γ ([Sch87]) require at least $2^{\tilde{\Omega}(n/\log \gamma)}$ time. We refer the reader to [Reg05, Pei09] for more information.

In this work, we will only consider the case where $q \leq 2^n$. Furthermore, the underlying security parameter λ is assumed to be polynomially related to the dimension n .

3.3 One-Dimensional Short Integer Solution (SIS) and Variants

We present a special case of the well known Short Integer Solution (SIS) problem [Ajt96].

Definition 3.3. *The One-Dimensional Short Integer Solution problem, denoted $1\text{D-SIS}_{q,m,t}$, is the following problem. Given a uniformly distributed vector $\mathbf{v} \xleftarrow{\$} \mathbb{Z}_q^m$, find $\mathbf{z} \in \mathbb{Z}^m$ such that $\|\mathbf{z}\| \leq t$ and also $\langle \mathbf{v}, \mathbf{z} \rangle \in [-t, t] + q\mathbb{Z}$.*

For appropriately chosen moduli q , the $1\text{D-SIS}_{q,m,t}$ problem is as hard as worst-case lattice problems. This follows from the techniques in the classical worst-case to average-case reduction of Ajtai [Ajt96]. We state below the version due to Regev [Reg04].

Corollary 3.2 (Section 4 in [Reg04] and Proposition 4.7 in [GPV07]). *Let $n \in \mathbb{N}$ and $q = \prod_{i \in \mathbb{N}} p_i$, where all $p_1 < p_2 < \dots < p_n$ are co-prime. Let $m \geq c \cdot n \log q$ (for some universal constant c). Assuming that $p_1 \geq t \cdot \omega(\sqrt{mn \log n})$, the one-dimensional SIS problem $1\text{D-SIS}_{q,m,t}$ is at least as hard as $\text{SIVP}_{t \cdot \tilde{O}(\sqrt{mn})}$ and $\text{GapSVP}_{t \cdot \tilde{O}(\sqrt{mn})}$.*

Proof. The hardness of a closely related problem is established by combining the techniques in [Reg04, Section 4] and [GPV07, Proposition 4.7]: Given $\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^{m+1}$, find \mathbf{y} with $\|\mathbf{y}\| \leq t$ such that $\langle \mathbf{a}, \mathbf{y} \rangle = 0 \pmod{q}$.

We now show how to convert an instance for this problem into an instance of 1D-SIS . Given an instance $\mathbf{a} \in \mathbb{Z}_q^{m+1}$, we consider the first component a_1 . If this element is not a unit (i.e. invertible) in \mathbb{Z}_q , then the reduction aborts. Otherwise it defines $\mathbf{v} = a_1^{-1} \cdot [a_2, \dots, a_{m+1}]$. Given a solution \mathbf{z} for 1D-SIS on input \mathbf{v} , we define \mathbf{y} by letting $\mathbf{y} = [-\langle \mathbf{v}, \mathbf{z} \rangle, x_1, \dots, x_m]$. It is easy to verify that $\langle \mathbf{a}, \mathbf{y} \rangle = a_1 \cdot (-\langle \mathbf{v}, \mathbf{z} \rangle + \langle \mathbf{v}, \mathbf{z} \rangle) = 0 \pmod{q}$. Further, by definition, $\|\mathbf{y}\| \leq t$. \square

Next, we define a related problem which will be useful for our reductions.

Definition 3.4. *Let $q = p \cdot \prod_{i \in \mathbb{N}} p_i$, where all $p_1 < p_2 < \dots < p_n$ are all co-prime and co-prime with p as well. Further let $m \in \mathbb{N}$. The $1\text{D-SIS-R}_{q,p,t,m}$ problem is the following: Given $\mathbf{v} \xleftarrow{\$} \mathbb{Z}_q^m$, find $\mathbf{z} \in \mathbb{Z}^m$ with $\|\mathbf{z}\| \leq t$ such that $\langle \mathbf{v}, \mathbf{z} \rangle \in [-t, t] + (q/p)\mathbb{Z}$.*

The following corollary establishes the hardness of 1D-SIS-R based on 1D-SIS .

Corollary 3.3. *Let q, p, t, m be as in Definition 3.4. Then $1\text{D-SIS-R}_{q,p,t,m}$ is at least as hard as $1\text{D-SIS}_{q/p,t,m}$.*

Proof. The reduction works in the obvious way: Given an input $\mathbf{v} \in \mathbb{Z}_{q/p}^m$ for 1D-SIS $_{q/p,t,m}$, we embed \mathbf{v} in $\mathbf{v}' \in \mathbb{Z}_q^m$, using CRT representation. Namely $\mathbf{v}' = \mathbf{v} \pmod{q/p}$ and $\mathbf{v}' = \mathbf{r} \pmod{p}$, where $\mathbf{r} \stackrel{\$}{\leftarrow} \mathbb{Z}_p^m$. Then given a solution \mathbf{z} for 1D-SIS- $R_{q,p,t,m}$ with input \mathbf{v}' , we claim that \mathbf{z} is also a solution for 1D-SIS $_{q/p,t,m}$ with input \mathbf{v} . This follows since by definition $\|\mathbf{z}\| \leq t$, and since $\langle \mathbf{v}, \mathbf{z} \rangle \equiv \langle \mathbf{v}', \mathbf{z} \rangle \pmod{q/p}$. \square

3.4 Admissible Hash Functions

The concept of admissible hash functions was defined by Boneh and Boyen [BB04] to convert selectively secure identity based encryption (IBE) schemes into fully secure ones. In this paper, we use admissible hash functions for our PRF construction. Our definition of admissible hash functions below will follow that of Cash, Hofheinz, Kiltz and Peikert [CHKP12] with minor changes (in particular, note that we do not require that the bad set is efficiently recognizable).

Definition 3.5 ([BB04, CHKP12]). *Let $\mathcal{H} = \{\mathcal{H}_\lambda\}_\lambda$ be a family of hash functions such that $\mathcal{H}_\lambda \subseteq (\{0,1\}^* \rightarrow \{0,1\}^\ell)$ for some $\ell = \ell(\lambda)$. We say that \mathcal{H} is a family of admissible hash functions if for every $H \in \mathcal{H}$ there exists a set \mathbf{bad}_H of “bad string-tuples” such that the following two properties hold:*

1. *For every PPT algorithm \mathcal{A} , there is a negligible function ν such that*

$$\Pr[(x^{(0)}, \dots, x^{(t)}) \in \mathbf{bad}_H \mid H \leftarrow \mathcal{H}_\lambda, (x^{(0)}, \dots, x^{(t)}) \leftarrow \mathcal{A}(1^\lambda, H)] \leq \nu(\lambda)$$

where the probability is over the choice of $H \leftarrow \mathcal{H}_\lambda$ and the coins of \mathcal{A} .

2. *Let $\mathcal{L} = \{0,1\}^{2\ell}$, and for all $L \in \mathcal{L}$ define $\Pi_L : \{0,1\}^\ell \rightarrow \{0,1\}$ to be the string comparison with wildcards function. Namely, write L as a pair of strings $(\alpha, \beta) \in \{0,1\}^\ell$, and define*

$$\Pi_{L=(\alpha,\beta)}(w) = 1 \Leftrightarrow \forall i \in [\ell] \left((\alpha_i = 0) \vee (\beta_i = w_i) \right) .$$

Intuitively, Π is a string comparison function with wildcards. It compares w and β only at those points where $\alpha_i = 1$. Note that this representation is somewhat redundant but it will be useful for our application.

Then, we require that for every polynomial $t = t(\lambda)$ there exists a noticeable function $\Delta_t(\lambda)$ and an efficiently sampleable distribution \mathcal{L}_t over \mathcal{L} such that for every $H \in \mathcal{H}_\lambda$ and sequences $(x^{(0)}, \dots, x^{(t)}) \notin \mathbf{bad}_H$ with $x^{(0)} \notin \{x^{(1)}, \dots, x^{(t)}\}$, we have:

$$\Pr_{L \leftarrow \mathcal{L}_t} [\Pi_L(H(x^{(0)})) \wedge \overline{\Pi_L(H(x^{(1)}))} \wedge \dots \wedge \overline{\Pi_L(H(x^{(t)}))}] \geq \Delta_t(\lambda)$$

It has been shown by [BB04] that a family of admissible hash functions can be constructed based on any collision resistant hash function. In particular one can instantiate it based on the SIS problem (for virtually any parameter setting for which the problem is hard), which is at least as hard as LWE. Therefore throughout this manuscript we assume the existence of an LWE-based family of admissible hash functions, which will not add an additional assumption to our construction.

3.5 Attribute-Based Encryption

We define (leveled) attribute-based encryption, following [GPSW06, GVW13]. An attribute-based encryption scheme for a class of predicate circuits \mathcal{C} (namely, circuits with a single bit output) consists of four algorithms ($\mathcal{ABE}.\text{Setup}$, $\mathcal{ABE}.\text{KeyGen}$, $\mathcal{ABE}.\text{Enc}$, $\mathcal{ABE}.\text{Dec}$).

$\mathcal{ABE}.\text{Setup}(1^\lambda, 1^\ell, 1^d) \rightarrow (\text{pp}, \text{msk})$: The setup algorithm gets as input the security parameter λ , the length ℓ of the attributes and the maximum depth of the predicate circuits d , and outputs the public parameter (pp, mpk) , and the master key msk . All the other algorithms get pp as part of their input.

$\mathcal{ABE}.\text{KeyGen}(\text{msk}, C) \rightarrow \text{sk}_C$: The key generation algorithm gets as input msk and a predicate specified by $C \in \mathcal{C}$ (of depth at most d). It outputs a secret key (C, sk_C) .

$\mathcal{ABE}.\text{Enc}(\text{pp}, \mathbf{x}, m) \rightarrow \text{ct}$: The encryption algorithm gets as input mpk , attributes $\mathbf{x} \in \{0, 1\}^\ell$ and a message $m \in \mathcal{M}$. It outputs a ciphertext (\mathbf{x}, ct) .

$\mathcal{ABE}.\text{Dec}((C, \text{sk}_C), (\mathbf{x}, \text{ct})) \rightarrow m$: The decryption algorithm gets as input a circuit C and the associated secret key sk_C , attributes \mathbf{x} and an associated ciphertext ct , and outputs either \perp or a message $m \in \mathcal{M}$.

Correctness. We require that for all ℓ, d , all (\mathbf{x}, C) such that $\mathbf{x} \in \{0, 1\}^\ell$, C has depth at most d and $C(\mathbf{x}) = 1$, for all $(\text{pp}, \text{msk}) \leftarrow \mathcal{ABE}.\text{Setup}(1^\lambda, 1^\ell, 1^d)$, all $\text{sk}_C \leftarrow \mathcal{ABE}.\text{KeyGen}(\text{msk}, C)$, all $\text{ct} \leftarrow \mathcal{ABE}.\text{Enc}(\text{pp}, \mathbf{x}, m)$, and all $m \in \mathcal{M}$,

$$\text{Dec}((C, \text{sk}_C), (\mathbf{x}, \text{ct})) = m .$$

Security Definition. We define selective security of ABE, which is sufficient for our purposes. We allow the adversary to make multiple challenge message queries, which is equivalent to the single query case but will be easier for us to work with.

Definition 3.6. For a stateful adversary \mathcal{A} , we define the advantage function $\text{Adv}_{\mathcal{A}}^{\text{ABE}}$ to be

$$\Pr \left[\begin{array}{l} b \xleftarrow{\$} \{0, 1\}; \\ \mathbf{x}_1, \dots, \mathbf{x}_Q \leftarrow \mathcal{A}(1^\lambda, 1^\ell, 1^d); \\ (\text{pp}, \text{msk}) \leftarrow \mathcal{ABE}.\text{Setup}(1^\lambda, 1^\ell, 1^d); \\ \{(m_{0,i}, m_{1,i})\}_{i \in [Q]} \leftarrow \mathcal{A}^{\mathcal{ABE}.\text{KeyGen}(\text{msk}, \cdot)}(\text{pp}), \forall i. |m_{0,i}| = |m_{1,i}|; \\ \text{ct}_i \leftarrow \mathcal{ABE}.\text{Enc}(\text{pp}, \mathbf{x}_i, m_{b,i}); \\ b' \leftarrow \mathcal{A}^{\mathcal{ABE}.\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}_1, \dots, \text{ct}_Q) \end{array} \right] - \frac{1}{2}$$

with the restriction that all queries C that \mathcal{A} makes to $\mathcal{ABE}.\text{KeyGen}(\text{msk}, \cdot)$ satisfies $C(\mathbf{x}_i) = 0$ for all i (that is, sk_C does not decrypt the ciphertext corresponding to any of the \mathbf{x}_i). An attribute-based encryption scheme is selectively secure if for all PPT adversaries \mathcal{A} , the advantage $\text{Adv}_{\mathcal{A}}^{\text{ABE}}$ is a negligible function in λ .

We will use a special type of attribute-based encryption scheme with succinct keys, namely one where $|\text{sk}_C|$ does not grow with the size of the circuit C , but rather only its depth.

Theorem 3.4 ([BGG⁺14]). *Let λ be the security parameter, and $d \in \mathbb{N}$. Let $n = n(\lambda, d)$, $q = q(\lambda, d) = n^{O(d)}$, and let χ be a $\text{poly}(n)$ -bounded error distribution. Then, there is a selectively secure ABE scheme for the class of depth- d -bounded circuits, based on the hardness of $\text{DLWE}_{n,q,\chi}$. Furthermore, the secret key sk_C for a circuit C has size $\text{poly}(\lambda, n, d)$.*

4 Embedding Circuits into Matrices

In this section, we present the core techniques that we use in our construction. In essence, we use a method, developed in a recent work by Boneh et al. [BGG⁺14] to “embed” bits x_1, \dots, x_k into matrices $\mathbf{A}_1, \dots, \mathbf{A}_k$ and compute a circuit F on these matrices. This is done through a pair of algorithms (`ComputeA`, `ComputeC`) satisfying the following properties:

1. The deterministic algorithm `ComputeA` takes as input a circuit $F : \{0, 1\}^k \rightarrow \{0, 1\}$ and k matrices $\mathbf{A}_1, \dots, \mathbf{A}_k$, and outputs a matrix \mathbf{A}_F ; and
2. The deterministic algorithm `ComputeC` takes as input a bit string $\mathbf{x} = (x_1, \dots, x_k) \in \{0, 1\}^k$, and k LWE samples $\mathbf{s}^T(\mathbf{A}_i + x_i \mathbf{G}) + \mathbf{e}_i$, and outputs an LWE sample $\mathbf{s}^T(\mathbf{A}_F + F(\mathbf{x}) \cdot \mathbf{G}) + \mathbf{e}_F$ associated to the output matrix \mathbf{A}_F and the output bit $F(\mathbf{x})$.

These algorithms are closely modeled on the work of Boneh et al. [BGG⁺14]. We now describe how these algorithms work, and what their properties are.

The Algorithm `ComputeA`. Given a circuit F , input matrices $\mathbf{A}_1, \dots, \mathbf{A}_k$ (corresponding to the k input wires) and an auxiliary matrix \mathbf{A}_0 , the `ComputeA` procedure works inductively, going through the gates of the circuit F from the input to the output. Assume without loss of generality that the circuit F is composed of NOT and AND gates. For every AND gate $g = (u, v; w)$, assume inductively that we have computed matrices \mathbf{A}_u and \mathbf{A}_v for the input wires u and v . Define

$$\mathbf{A}_w = -\mathbf{A}_u \cdot \mathbf{G}^{-1}(\mathbf{A}_v)$$

For every NOT gate $g = (u; w)$, define

$$\mathbf{A}_w = \mathbf{A}_0 - \mathbf{A}_u$$

The Algorithm `ComputeC`. Given a circuit F , an input $x \in \{0, 1\}^k$ and LWE samples $(\mathbf{A}_i, \mathbf{y}_i)$, the `ComputeC` algorithm works as follows. For each AND gate $g = (u, v; w)$, assume that we have computed LWE samples $(\mathbf{A}_u, \mathbf{y}_u)$ and $(\mathbf{A}_v, \mathbf{y}_v)$ for the input wires u and v . Define

$$\mathbf{y}_w = x_u \cdot \mathbf{y}_v - \mathbf{y}_u \cdot \mathbf{G}^{-1}(\mathbf{A}_v)$$

where x_u and x_v are the bits on wires u and v when evaluating the circuit F on input x . For every NOT gate $g = (u; w)$, define

$$\mathbf{y}_w = \mathbf{y}_0 - \mathbf{y}_u$$

We will need the following lemma about the behavior of `ComputeA` and `ComputeC`. (We remind the reader that we use $\|\cdot\|$ to denote the ℓ_∞ norm).

Lemma 4.1. *Let F be a depth- d Boolean circuit on k input bits, and let $x \in \{0, 1\}^k$ be an input. Let $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_k \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{y}_0, \dots, \mathbf{y}_k \in \mathbb{Z}_q^m$ be such that*

$$\|\mathbf{y}_i - \mathbf{s}^T(\mathbf{A}_i + x_i \mathbf{G})\| \leq B \quad \text{for } i = 0, 1, \dots, k.$$

for some $\mathbf{s} \in \mathbb{Z}_q^n$ and $B = B(\lambda)$. Let $\mathbf{A}_F \leftarrow \text{ComputeA}(F, \mathbf{A}_0, \dots, \mathbf{A}_k)$ and $\mathbf{y}_F \leftarrow \text{ComputeC}(F, x, \mathbf{A}_0, \dots, \mathbf{A}_k, \mathbf{y}_0, \dots, \mathbf{y}_k)$. Then, $\|\mathbf{y}_F - \mathbf{s}^T(\mathbf{A}_F + F(\mathbf{x}) \cdot \mathbf{G})\| \leq m^{O(d)} \cdot B$.

Furthermore, \mathbf{y}_F is a “low-norm” linear function of $\mathbf{y}_0, \dots, \mathbf{y}_k$. That is, there are matrices $\mathbf{Z}_0, \dots, \mathbf{Z}_k$ (which depend on the function F , the input \mathbf{x} , and the input matrices $\mathbf{A}_0, \dots, \mathbf{A}_k$) such that $\mathbf{y}_F = \sum_{i=0}^k \mathbf{y}_i \mathbf{Z}_i$ and $\|\mathbf{Z}_i\| \leq m^{O(d)} \cdot B$.

Proof. We show this by induction on the levels of the circuit F , starting from the input. Consider two cases.

AND gate. Consider an AND gate $g = (u, v; w)$ where the input wires are at level L , and assume that $\mathbf{y}_u = \mathbf{s}^T(\mathbf{A}_u + x_u \mathbf{G}) + \mathbf{e}_u$ and $\mathbf{y}_v = \mathbf{s}^T(\mathbf{A}_v + x_v \mathbf{G}) + \mathbf{e}_v$, with $\|\mathbf{e}_u\|, \|\mathbf{e}_v\| \leq (m+1)^L \cdot B$. Now,

$$\begin{aligned} \mathbf{y}_w &= x_u \cdot \mathbf{y}_v - \mathbf{y}_u \cdot \mathbf{G}^{-1}(\mathbf{A}_v) \\ &= x_u \cdot (\mathbf{s}^T(\mathbf{A}_v + x_v \mathbf{G}) + \mathbf{e}_v) - \left(\mathbf{s}^T(\mathbf{A}_u + x_u \mathbf{G}) + \mathbf{e}_u \right) \cdot \mathbf{G}^{-1}(\mathbf{A}_v) \\ &= \mathbf{s}^T \left(x_u \mathbf{A}_v + x_u x_v \mathbf{G} - \mathbf{A}_u \mathbf{G}^{-1}(\mathbf{A}_v) - x_u \mathbf{A}_v \right) + \left(-\mathbf{e}_u \mathbf{G}^{-1}(\mathbf{A}_v) + x_u \mathbf{e}_v \right) \\ &= \mathbf{s}^T(\mathbf{A}_w + x_w \mathbf{G}) + \mathbf{e}_w \end{aligned}$$

where $\mathbf{A}_w = -\mathbf{A}_u \cdot \mathbf{G}^{-1}(\mathbf{A}_v)$, $x_w = x_u x_v$, and

$$\|\mathbf{e}_w\| \leq m \cdot \|\mathbf{e}_u\| + \|\mathbf{e}_v\| \leq (m+1) \cdot (m+1)^L \cdot B \leq (m+1)^{L+1} \cdot B$$

NOT gate. In a similar vein, for a NOT gate $g = (u; w)$, assume that $\mathbf{y}_u = \mathbf{s}^T(\mathbf{A}_u + x_u \mathbf{G}) + \mathbf{e}_u$, with $\|\mathbf{e}_u\| \leq (m+1)^L \cdot B$. Then,

$$\begin{aligned} \mathbf{y}_w &= \mathbf{y}_0 - \mathbf{y}_u = \mathbf{s}^T(\mathbf{A}_0 + \mathbf{G} - \mathbf{A}_u - x_u \mathbf{G}) + (\mathbf{e}_0 - \mathbf{e}_u) \\ &= \mathbf{s}^T(\mathbf{A}_w + (1 - x_u) \mathbf{G}) + \mathbf{e}_w \end{aligned}$$

where $\mathbf{A}_w = \mathbf{A}_0 - \mathbf{A}_u$, $x_w = 1 - x_u$, and

$$\|\mathbf{e}_w\| \leq \|\mathbf{e}_0\| + \|\mathbf{e}_u\| \leq B + (m+1)^L \cdot B \leq (m+1)^{L+1} \cdot B$$

Thus, $\mathbf{y}_F = \mathbf{s}^T \mathbf{A}_F + \mathbf{e}_F$ where $\|\mathbf{e}_F\| \leq m^{O(d)} \cdot B$. Furthermore, both transformations are linear functions on \mathbf{y}_u and \mathbf{y}_v , as required. \square

5 Constrained PRF

5.1 Construction

A family of functions $\mathcal{F} \subseteq (\{0, 1\}^* \rightarrow \{0, 1\})$ is z -uniform if each function $F \in \mathcal{F}$ can be described by a string in $\{0, 1\}^z$ (we associate F with its description), and there exists a uniform circuit

family $\{\mathcal{U}_k\}_{k \in \mathbb{N}}$ such that $\mathcal{U}_k : \{0, 1\}^z \times \{0, 1\}^k \rightarrow \{0, 1\}$ such that for all $x \in \{0, 1\}^k$ it holds that $\mathcal{U}_k(F, x) = F(x)$. We assume for the sake of simplicity that the depth of \mathcal{U}_k grows monotonically with k and for all d we let k_d to be the maximal input size for which \mathcal{U}_k has depth at most d . We define \mathcal{F}_d to be such that $F \in \mathcal{F}$ is undefined for inputs of length $k > k_d$. We call such a family d -depth-bounded.

Our constrained PRF for a z -uniform d -depth-bounded family \mathcal{F} works as follows.

- **KeyGen**($1^\lambda, 1^z, 1^d$): The key generation algorithm takes as input the maximum size z and depth d of the constraining circuits. Let \mathcal{H} be a family of admissible hash functions (see Section 3.4) and let $\ell = \ell(\lambda)$ be the output length of hash functions in the family.

Let $n = n(\lambda, d)$, $q = q(\lambda, d)$, $p = p(\lambda, d)$ be parameters chosen as described in Section 5.2 below, let $m = n \lceil \log q \rceil$.

Generate $z+2\ell+3$ matrices as follows: let \mathbf{A}_0 and \mathbf{A}_1 be the “input matrices”, let $\mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_z$ be the “function matrices”, let $\mathbf{C}_1, \dots, \mathbf{C}_{2\ell}$ be the “partitioning matrices”, and let \mathbf{D} be an “auxiliary matrix”. All of these matrices are uniform in $\mathbb{Z}_q^{n \times m}$ (note that the “gadget matrix” \mathbf{G} has the same dimensions). In addition sample an admissible hash function $H \xleftarrow{\$} \mathcal{H}_\lambda$.

The public parameters consist of

$$\mathcal{PP} = (H, \mathbf{A}_0, \mathbf{A}_1, \mathbf{B}_1, \dots, \mathbf{B}_z, \mathbf{C}_1, \dots, \mathbf{C}_{2\ell}, \mathbf{D})$$

The seed of the PRF is a uniformly random vector $\mathbf{s} \in \mathbb{Z}_q^n$.

- **Eval**($\mathbf{s}, \mathcal{PP}, \mathbf{x}$) takes as input the PRF seed \mathbf{s} , the public parameters \mathcal{PP} , and an input $x \in \{0, 1\}^k$ such that $k \leq k_d$ (i.e. \mathcal{U}_k is of depth $\leq d$), and works as follows.

Recall that $\mathcal{U}_k : \{0, 1\}^z \times \{0, 1\}^k \rightarrow \{0, 1\}$ is the universal circuit that takes a description of a function F and an input x and outputs $\mathcal{U}_k(F, x) = F(x)$. Let $\Pi : \{0, 1\}^{2\ell} \times \{0, 1\}^\ell \rightarrow \{0, 1\}$ denote the circuit that computes $\Pi(L, w) = \Pi_L(w)$ from Definition 3.5. Note that Π can be implemented by a binary circuit of depth $\log(\ell) + O(1)$.

Let (x_1, \dots, x_k) denote the bits of x . Let $w = H(x)$, and let w_1, \dots, w_ℓ be its bits. Compute

$$\mathbf{B}_\mathcal{U} \leftarrow \text{ComputeA}(\mathcal{U}_k, \mathbf{B}_1, \dots, \mathbf{B}_z, \mathbf{A}_{x_1}, \mathbf{A}_{x_2}, \dots, \mathbf{A}_{x_k}) \quad (1)$$

$$\mathbf{C}_\Pi \leftarrow \text{ComputeA}(\Pi, \mathbf{C}_1, \dots, \mathbf{C}_{2\ell}, \mathbf{A}_{w_1}, \mathbf{A}_{w_2}, \dots, \mathbf{A}_{w_\ell}) \quad (2)$$

and output

$$\text{PRF}_\mathbf{s}(\mathbf{x}) = \lfloor \mathbf{s}^T \mathbf{B}_\mathcal{U} \cdot \mathbf{G}^{-1}(\mathbf{C}_\Pi) \cdot \mathbf{G}^{-1}(\mathbf{D}) \rfloor_p$$

- **Constrain**($\mathbf{s}, \mathcal{PP}, F$) takes as input the PRF key \mathbf{s} and a circuit F (of size at most z) and does the following. Compute

$$\mathbf{a}_b = \mathbf{s}^T (\mathbf{A}_b + b \cdot \mathbf{G}) + \mathbf{e}_{1,b}^T \in \mathbb{Z}_q^m \text{ for } b \in \{0, 1\}$$

$$\mathbf{b}_i = \mathbf{s}^T (\mathbf{B}_i + f_i \cdot \mathbf{G}) + \mathbf{e}_{2,i}^T \in \mathbb{Z}_q^m \text{ for all } i \in [z]$$

where the vectors \mathbf{e} are drawn from an error distribution χ to be specified later (in Section 5.2).

The constrained seed K_F is the tuple $(\mathbf{a}_0, \mathbf{a}_1, \mathbf{b}_1, \dots, \mathbf{b}_z) \in (\mathbb{Z}_q^m)^{z+2}$.

- $\text{ConstrainEval}(K_F, \mathcal{PP}, \mathbf{x})$ takes as input the constrained key K_F and an input \mathbf{x} . It computes

$$\mathbf{b}_{\mathcal{U}, \mathbf{x}} \leftarrow \text{ComputeC}\left(\mathcal{U}, (\mathbf{b}_1, \dots, \mathbf{b}_z, \mathbf{a}_{x_1}, \dots, \mathbf{a}_{x_k}), (f_1, \dots, f_z, x_1, \dots, x_k)\right)$$

and outputs $\lfloor \mathbf{b}_{\mathcal{U}, \mathbf{x}} \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi}) \cdot \mathbf{G}^{-1}(\mathbf{D}) \rfloor_p$, where \mathbf{C}_{Π} is defined as above.

5.2 Setting the Parameters

Let us start by providing a typical parameter setting, and then explain how parameters can be modified and the effect on security.

Consider setting $n(\lambda, d) = (\lambda \cdot d)^c$, for a constant c that will be discussed shortly. We will set χ to be a discrete Gaussian distribution $D_{\mathbb{Z}, \alpha q}$ s.t. $\alpha q = \Theta(\sqrt{n})$. We define $n' = \lambda$ and let $p_1, \dots, p_{n'} = m^{O(d+\log \ell)}$ be all primes, and $p = \text{poly}(\lambda)$ (in fact, there is a lot of freedom in the choice of p , and it can be as large as $m^{O(d+\log \ell)}$ under the same asymptotic hardness). Finally, let $q = p \cdot (\alpha q) \cdot \prod_{i \in [n']} p_i = m^{n' \cdot O(d+\log \ell)} = 2^{\tilde{O}(\lambda \cdot d)} = 2^{\tilde{O}(n^{1/c})}$ (recall that $\ell = \text{poly}(\lambda)$).

This parameter setting translates into a PRF with $m = n \lceil \log q \rceil \cdot \Theta(\log \lambda)$ output bits per input, whose security is based (as we show in the next section) on the hardness of approximating lattice problems to within a factor of $2^{\tilde{O}(n^{1/c})}$.

Taking larger values of c will increase the hardness of the underlying lattice problem, but at the cost of considerably increasing the element sizes.

5.3 Security

Throughout this section, we let \mathcal{F} be a family of z -uniform functions and let d be a depth bound (both can depend on λ). We let $n = n(\lambda, d)$, $m = m(\lambda, d)$, $q = q(\lambda, d)$, $p = p(\lambda, d)$ and the noise distributions $\chi = \chi(\lambda, d)$ be as defined in Section 5.2. We let \mathcal{H} be the family of admissible hash functions as described in Section 3.4, with range $\{0, 1\}^\ell$.

Theorem 5.1. *Let \mathcal{F} be a family of z -uniform functions and let d be a depth bound (both can depend on λ). Let $n = n(\lambda, d)$, $m = m(\lambda, d)$, $q = q(\lambda, d)$, $p = p(\lambda, d)$ and the noise distributions $\chi = \chi(\lambda, d)$ be as defined in Section 5.2. Further let $m' = m \cdot (z + 2\ell + 3)$, and $\gamma = \omega(\sqrt{n \log \lambda}) \cdot p \cdot m^{O(d+\log \ell)}$. Assuming the hardness of $\text{DLWE}_{n, q, \chi}$, $\text{1D-SIS-R}_{q, p, \gamma, m'}$ and the admissible hash function family \mathcal{H} , the scheme $\mathcal{CPRF} = (\text{KeyGen}, \text{Eval}, \text{Constrain}, \text{ConstrainEval})$ is a single-key secure selective-function secure constrained PRF for \mathcal{F} .*

We note that the hardness of all three assumptions translates to the worst case hardness of approximating lattice problems such as GapSVP and SIVP to within sub-exponential factors.

Proof. Let \mathcal{A} be a PPT selective-constraint adaptive-input adversary against $\mathcal{CPRF}_{z, d}$. Let $t = \text{poly}(\lambda)$ be the (polynomial) number of input queries made by \mathcal{A} (w.l.o.g). Let ϵ be the advantage of \mathcal{A} in the constrained PRF game. We let $B = \alpha q \cdot \omega(\sqrt{\log \lambda})$. It holds that with all but negligible probabilities, all samples that we take from χ will have absolute value at most B . For the duration of the proof we assume that this is indeed the case.

The proof will proceed by a sequence of hybrids (or experiments) where the challenger samples a bit $b \in \{0, 1\}$ and interacts with \mathcal{A} . We let $\text{Adv}_{\mathcal{H}}(\mathcal{A})$ denote the probability that \mathcal{A} outputs b in hybrid \mathcal{H} .

Hybrid H₀. This hybrid is the legitimate constrained PRF security game. The challenger generates $(\mathbf{s}, \mathcal{PP}) \leftarrow \text{KeyGen}(1^\lambda, 1^z, 1^d)$. It gets $F \in \{0, 1\}^z$ from \mathcal{A} and produces a constrained key $K_F \leftarrow \text{Constrain}(\mathbf{s}, \mathcal{PP}, F)$. It then sends \mathcal{PP}, K_F to \mathcal{A} . At this point \mathcal{A} adaptively makes queries $x^{(i)} \in \{0, 1\}^*$, and the challenger computes $y^{(i)} \leftarrow \text{Eval}(\mathbf{s}, \mathcal{PP}, x^{(i)})$ and returns it to \mathcal{A} . Finally, \mathcal{A} outputs $x^* \in \{0, 1\}^*$. If $b = 0$ then the challenger returns $y^* \leftarrow \text{Eval}(\mathbf{s}, \mathcal{PP}, x^*)$, and if $b = 1$ it returns a random y^* . Therefore, we have

$$\text{Adv}_{\text{H}_0}(\mathcal{A}) \geq 1/2 + \epsilon .$$

Hybrid H₁. This is the notorious “artificial abort” phase. Let $\Delta_t = \Delta_t(\lambda)$ be the noticeable function from Definition 3.5. This hybrid is identical to the previous one, except in the last step the challenger flips a coin and with probability $1 - \Delta_t/2$ aborts the experiment (hence giving the adversary no information on b).

The adversary’s advantage thus degrades appropriately:

$$\text{Adv}_{\text{H}_1}(\mathcal{A}) \geq (\Delta_t/2) \cdot (1/2 + \epsilon) + (1 - \Delta_t/2) \cdot (1/2) = 1/2 + \epsilon \cdot \Delta_t/2 .$$

Hybrid H₂. In this hybrid, we associate some meaning with the artificial abort. Intuitively, the abort will be associated with a failure of the admissible hash function to partition the queries correctly. We are guaranteed that correct partitioning happens with probability $\geq \Delta_t$ (except for sequences that are hard to generate), but we would like to make it (*almost*) *exactly* $\Delta_t/2$ so as to not correlate the adversary’s success probability with the string L (the loss of the 2 factor is due to probability estimation).

Specifically, in this hybrid, rather than flipping a coin at the end of the experiment, the challenger does the following. For all $\vec{x} = (x^{(1)}, \dots, x^{(t)}, x^*)$, we define the event $\text{GoodPartition}_{L, \vec{x}}$ to be the event in which $\Pi_L(H(x^{(1)})) = \dots = \Pi_L(H(x^{(t)})) = 0$ and $\Pi_L(H(x^*)) = 1$, and define $\delta_{\vec{x}} = \Pr_{L \leftarrow \mathcal{L}_t} [\text{GoodPartition}_{\vec{x}, L}]$. The challenger will first compute an estimate $\tilde{\delta}_{\vec{x}}$ of $\delta_{\vec{x}}$ by sampling multiple values of L from \mathcal{L}_t and using Chernoff (both additive and multiplicative). Using $\text{poly}(\lambda)$ -many samples we can compute $\tilde{\delta}_{\vec{x}}$ such that

$$\Pr \left[\left| \delta_{\vec{x}} - \tilde{\delta}_{\vec{x}} \right| > \Delta_t/4 \right] \leq 2^{-\lambda} .$$

and in addition if $\delta_{\vec{x}} \geq \Delta_t/2$ then

$$\Pr \left[\left| \frac{\tilde{\delta}_{\vec{x}}}{\delta_{\vec{x}}} - 1 \right| > \epsilon/2 \right] \leq 2^{-\lambda} .$$

The challenger will then perform as follows: (i) It first verifies that $\tilde{\delta}_{\vec{x}} \geq \frac{3}{4}\Delta_t$, and aborts if this is not the case. (ii) It then samples $L \leftarrow \mathcal{L}_t$ and aborts if $\text{GoodPartition}_{\vec{x}, L}$ did not occur (note that by our definitions above, this happens with probability $1 - \delta_{\vec{x}}$ over the choice of L). (iii) Then it flips a coin with probability $\frac{\tilde{\delta}_{\vec{x}} - \Delta_t/2}{\tilde{\delta}_{\vec{x}}}$ and aborts if the outcome is 1. Otherwise it carries out the experiment towards completion.

To analyze the effect on the success probability, we first notice that the probability that $\tilde{\delta}_{\vec{x}} < \frac{3}{4}\Delta_t$ (abortion is step (i)) is negligible. This is since, except with $2^{-\lambda}$ probability, this indicates that $\delta_{\vec{x}} < \Delta_t$, which implies that $\vec{x} \in \text{bad}_H$. Definition 3.5 guarantees that this happens with probability at most $\nu(\lambda) = \text{negl}(\lambda)$.

If the above abort did not occur, we know that $\delta_{\bar{x}} \geq \Delta_t/2$ (except with probability $2^{-\lambda}$), we first notice that the total probability of abort in steps (ii) + (iii)

$$1 - \delta_{\bar{x}} + \delta_{\bar{x}} \cdot \frac{\tilde{\delta}_{\bar{x}} - \Delta_t/2}{\tilde{\delta}_{\bar{x}}} = 1 - \frac{\delta_{\bar{x}}}{\tilde{\delta}_{\bar{x}}} \Delta_t/2 \in [(1 - \Delta_t/2) - \epsilon \Delta_t/4, (1 - \Delta_t/2) + \epsilon \Delta_t/4]$$

It therefore follows that if there was no abort in step (i), then the adversary's view in \mathbf{H}_2 is within statistical distance $2^{-\lambda} + \epsilon \Delta_t/4$ from its view in \mathbf{H}_1 .

Putting all steps together, we get that

$$\text{Adv}_{\mathbf{H}_2}(\mathcal{A}) \geq 1/2 + \epsilon \cdot \Delta_t/2 - \nu(\lambda) - O(2^{-\lambda}) - \epsilon \Delta_t/4 = 1/2 + \epsilon \cdot \Delta_t/4 - \text{negl}(\lambda) .$$

Hybrid \mathbf{H}_3 . In this hybrid, the challenger first samples $L \stackrel{\$}{\leftarrow} \mathcal{L}_t$, and then, for each $x^{(i)}$ in turn, it checks whether $\Pi_L(H(x^{(i)})) = 0$, and immediately aborts if not. Similarly, upon receiving x^* , it checks whether $\Pi_L(H(x^*)) = 1$ and immediately aborts if not. Otherwise it continues the same as \mathbf{H}_2 .

It is rather straightforward to see that the \mathcal{A} 's advantage does not change. The cases in which we abort are exactly the same as the ones in the previous hybrid (since it is sufficient that a single $x^{(i)}$ does not give the required value in order to abort). Further, the sampling of L has been completely independent of all the other randomness in the experiment so it might as well happen in the beginning. We conclude that

$$\text{Adv}_{\mathbf{H}_3}(\mathcal{A}) = \text{Adv}_{\mathbf{H}_2}(\mathcal{A}) \geq 1/2 + \epsilon \cdot \Delta_t/4 - \text{negl}(\lambda) .$$

Hybrid \mathbf{H}_4 . In this hybrid, the challenger changes the way the matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}$ are generated. Recall that our security game is constraint-selective, namely \mathcal{A} produces the constraint F before seeing the public parameters.

Therefore, here, the challenger waits until receiving F from \mathcal{A} and only generates the public parameters at that point (note that by then L has also been specified). To generate the public parameters, the matrix \mathbf{D} is produced identically to before. In addition, the challenger samples matrices $\{\hat{\mathbf{A}}_\beta\}_{\beta \in \{0,1\}}$, $\{\hat{\mathbf{B}}_i\}_{i \in [z]}$, $\{\hat{\mathbf{C}}_i\}_{i \in [2\ell]}$. It then sets

$$\begin{aligned} \mathbf{A}_\beta &= \hat{\mathbf{A}}_\beta - \beta \mathbf{G} \\ \mathbf{B}_i &= \hat{\mathbf{B}}_i - f_i \mathbf{G} \\ \mathbf{C}_i &= \hat{\mathbf{C}}_i - L_i \mathbf{G} \end{aligned}$$

The remainder of the experiment remains unchanged.

Since the distributions of the $\mathbf{A}, \mathbf{B}, \mathbf{C}$ matrices is identical to their original uniform distributions, it follows that

$$\text{Adv}_{\mathbf{H}_4}(\mathcal{A}) = \text{Adv}_{\mathbf{H}_3}(\mathcal{A}) .$$

Hybrid \mathbf{H}_5 . In this hybrid, the adversary changes the way it computes the outputs $y^{(i)}$. Recall that $K_F = (\mathbf{a}_0, \mathbf{a}_1, \mathbf{b}_1, \dots, \mathbf{b}_z)$ is the constrained key given to \mathcal{A} . Let us denote

$$\begin{aligned} \mathbf{c}_i &= \mathbf{s}^T (\mathbf{C}_i + L_i \mathbf{G}) + \mathbf{e}_{3,i}^T \quad \text{for all } i \in [z] \\ \mathbf{d} &= \mathbf{s}^T \mathbf{D} + \mathbf{e}_4^T \end{aligned}$$

where $\mathbf{e}_{3,i}$ are sampled coordinate-wise from χ , and \mathbf{e}_4 is sampled coordinate-wise from χ' .

In this hybrid, in order to answer input queries, the challenger first computes

$$\mathbf{b}_{\mathcal{U},x^{(i)}} \leftarrow \text{ComputeC}\left(\mathcal{U}, (\mathbf{b}_1, \dots, \mathbf{b}_z, \mathbf{a}_{x_1}, \dots, \mathbf{a}_{x_k}), (f_1, \dots, f_z, x_1^{(i)}, \dots, x_k^{(i)})\right)$$

and then, letting $w^{(i)} = H(x^{(i)})$

$$\mathbf{c}_{\Pi,w^{(i)}} \leftarrow \text{ComputeC}\left(\Pi, (\mathbf{c}_1, \dots, \mathbf{c}_{2\ell}, \mathbf{a}_{w_1}, \dots, \mathbf{a}_{w_\ell}), (L_1, \dots, L_{2\ell}, w_1^{(i)}, \dots, w_\ell^{(i)})\right)$$

We recall that by Lemma 4.1 it holds that

$$\begin{aligned} \mathbf{b}_{\mathcal{U},x^{(i)}}^T &= \mathbf{s}^T (\mathbf{B}_{\mathcal{U},x^{(i)}} + F(x^{(i)}) \cdot \mathbf{G}) + \mathbf{e}_{\mathcal{U}}^T \\ \mathbf{c}_{\Pi,w^{(i)}}^T &= \mathbf{s}^T (\mathbf{C}_{\Pi,w^{(i)}} + \Pi_L(w^{(i)}) \cdot \mathbf{G}) + \mathbf{e}_{\Pi}^T, \end{aligned}$$

for some $\mathbf{e}_{\mathcal{U}}, \mathbf{e}_{\Pi}$ for which $\|\mathbf{e}_{\mathcal{U}}\| \leq B \cdot m^{O(d)}$, $\|\mathbf{e}_{\Pi}\| \leq B \cdot m^{O(\log \ell)}$.

We recall that by definition

$$\begin{aligned} \text{PRF}_s(x^{(i)}) &= \left[\mathbf{s}^T \mathbf{B}_{\mathcal{U},x^{(i)}} \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi,w^{(i)}}) \mathbf{G}^{-1}(\mathbf{D}) \right]_p \\ &= \left[\mathbf{s}^T (\mathbf{B}_{\mathcal{U},x^{(i)}} + F(x^{(i)}) \mathbf{G}) \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi,w^{(i)}}) \mathbf{G}^{-1}(\mathbf{D}) \right. \\ &\quad \left. - F(x^{(i)}) \mathbf{s}^T \mathbf{C}_{\Pi,w^{(i)}} \mathbf{G}^{-1}(\mathbf{D}) \right]_p \\ &= \left[\mathbf{s}^T (\mathbf{B}_{\mathcal{U},x^{(i)}} + F(x^{(i)}) \mathbf{G}) \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi,w^{(i)}}) \mathbf{G}^{-1}(\mathbf{D}) \right. \\ &\quad \left. - F(x^{(i)}) \mathbf{s}^T (\mathbf{C}_{\Pi,w^{(i)}} + \Pi_L(w^{(i)}) \mathbf{G}) \mathbf{G}^{-1}(\mathbf{D}) \right. \\ &\quad \left. + F(x^{(i)}) \Pi_L(w^{(i)}) \mathbf{s}^T \mathbf{D} \right]_p \\ &= \left[\mathbf{b}_{\mathcal{U},x^{(i)}}^T \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi,w^{(i)}}) \mathbf{G}^{-1}(\mathbf{D}) - F(x^{(i)}) \mathbf{c}_{\Pi,w^{(i)}}^T \mathbf{G}^{-1}(\mathbf{D}) \right. \\ &\quad \left. + F(x^{(i)}) \Pi_L(w^{(i)}) \mathbf{d}^T + \mathbf{e}'^T \right]_p, \end{aligned} \tag{3}$$

where

$$\mathbf{e}'^T = -\mathbf{e}_{\mathcal{U}}^T \mathbf{G}^{-1}(\mathbf{C}_{\Pi,w^{(i)}}) \mathbf{G}^{-1}(\mathbf{D}) + F(x^{(i)}) \mathbf{e}_{\Pi}^T \mathbf{G}^{-1}(\mathbf{D}) - F(x^{(i)}) \Pi_L(w^{(i)}) \mathbf{e}_4^T \tag{4}$$

which implies that $\|\mathbf{e}'\| \leq E$ for some $E = (m^{O(d)} + m^{O(\log \ell)}) \cdot B$.

To analyze the distinguishing probability between these hybrids, for any input x (and $w = H(x)$) we define the event Borderline_x as the event where there exists $j \in [m]$ such that:

$$\begin{aligned} &(\mathbf{b}_{\mathcal{U},x}^T \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi,w}) \cdot \mathbf{G}^{-1}(\mathbf{D}) - F(x) \cdot \mathbf{c}_{\Pi,w}^T \cdot \mathbf{G}^{-1}(\mathbf{D}) \\ &\quad + F(x) \cdot \Pi_L(w) \cdot \mathbf{d}^T) \cdot \mathbf{u}_j \in [-E, E] + (q/p)\mathbb{Z}, \end{aligned}$$

where we recall that \mathbf{u}_j is the j th indicator vector. Namely, this is the probability that one of the coordinates of the vector $\mathbf{b}_{\mathcal{U},x}^T \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi,w}) \mathbf{G}^{-1}(\mathbf{D}) - F(x) \mathbf{c}_{\Pi,w}^T \mathbf{G}^{-1}(\mathbf{D}) + F(x) \Pi_L(w) \mathbf{d}^T$ is “dangerously close” to being rounded in the wrong direction.

By definition of rounding, if $\neg\text{Borderline}_{x^{(i)}}$, then

$$\begin{aligned} \text{PRF}_{\mathbf{s}}(x^{(i)}) &= \lfloor \mathbf{b}_{\mathcal{U},x^{(i)}}^T \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi,w^{(i)}}) \mathbf{G}^{-1}(\mathbf{D}) - F(x^{(i)}) \mathbf{c}_{\Pi,w^{(i)}}^T \mathbf{G}^{-1}(\mathbf{D}) \\ &\quad + F(x^{(i)}) \Pi_L(w^{(i)}) \mathbf{d}^T \rfloor_p . \end{aligned}$$

The challenger in this hybrid, given a query $x^{(i)}$, will first check whether $\text{Borderline}_{x^{(i)}}$. If the event happens, the challenger aborts. Otherwise it returns $\text{PRF}_{\mathbf{s}}(x^{(i)})$ as defined above. Note that the challenger only needs to respond to queries $x^{(i)}$ for which $\Pi_L(w^{(i)}) = \Pi_L(H(x^{(i)})) = 0$, which do not depend on \mathbf{d} , a fact that will be important later on.

Finally, on the challenge query x^* , unless abort is needed, it holds that $F(x^*) = 1$ and $\Pi_L(w^*) = 1$ (where $w^* = H(x^*)$) and therefore, unless the event Borderline_{x^*} happens, it holds that

$$\text{PRF}_{\mathbf{s}}(x^*) = \left\lfloor \mathbf{b}_{\mathcal{U},x^*}^T \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi,w^{(i)}}) \mathbf{G}^{-1}(\mathbf{D}) - \mathbf{c}_{\Pi,w^*}^T \mathbf{G}^{-1}(\mathbf{D}) + \mathbf{d}^T \right\rfloor_p .$$

The challenger will therefore abort if Borderline_{x^*} and return the aforementioned value otherwise (that is if the bit b is 0; if $b = 1$ then of course a uniform value is returned).

It follows that if we define $\text{Borderline} = (\bigvee_i \text{Borderline}_{x^{(i)}}) \vee \text{Borderline}_{x^*}$, then

$$|\text{Adv}_{\text{H}_5}(\mathcal{A}) - \text{Adv}_{\text{H}_4}(\mathcal{A})| \leq \Pr_{\text{H}_5}[\text{Borderline}] .$$

We will bound $\Pr_{\text{H}_5}[\text{Borderline}]$ as a part of our analysis in the next hybrid.

As a final remark on this hybrid, we note that in order to execute this hybrid, the challenger does not need to access \mathbf{s} itself, but rather only the $\mathbf{a}_\beta, \mathbf{b}_i, \mathbf{c}_i, \mathbf{d}$ vectors. This will be useful in the next hybrid.

Hybrid H₆. In this hybrid, all $\mathbf{a}_\beta, \mathbf{b}_i, \mathbf{c}_i, \mathbf{d}$ are sampled from the uniform distribution. Everything else remains the same. We note that by definition, in hybrid H₅:

$$\begin{aligned} \mathbf{a}_\beta^T &= \mathbf{s}^T \hat{\mathbf{A}}_\beta + \mathbf{e}_{1,\beta}^T \\ \mathbf{b}_i^T &= \mathbf{s}^T \hat{\mathbf{B}}_i + \mathbf{e}_{2,i}^T \\ \mathbf{c}_i^T &= \mathbf{s}^T \hat{\mathbf{C}}_i + \mathbf{e}_{3,i}^T \\ \mathbf{d}^T &= \mathbf{s}^T \mathbf{D} + \mathbf{e}_4^T , \end{aligned}$$

where all $\hat{\mathbf{A}}_\beta, \hat{\mathbf{B}}_i, \hat{\mathbf{C}}_i, \mathbf{D}$ are uniformly distributed, and all $\mathbf{e}_{1,\beta}^T, \mathbf{e}_{2,i}^T, \mathbf{e}_{3,i}^T, \mathbf{e}_4^T$ are sampled coordinate-wise from χ . The DLWE $_{n,q,\chi}$ assumption therefore asserts that:

$$|\text{Adv}_{\text{H}_6}(\mathcal{A}) - \text{Adv}_{\text{H}_5}(\mathcal{A})| \leq \text{negl}(\lambda) .$$

Furthermore, since Borderline is an efficiently recognizable event, it also holds that

$$\left| \Pr_{\text{H}_6}[\text{Borderline}] - \Pr_{\text{H}_5}[\text{Borderline}] \right| = \text{negl}(\lambda) . \quad (5)$$

In H₆, the probability of Borderline can be bounded under the 1D-SIS-R assumption.

Claim 5.1.1. *Under the 1D-SIS-R $_{q,p,\gamma,m'}$ assumption, it holds that $\Pr_{\text{H}_6}[\text{Borderline}] = \text{negl}(\lambda)$, where $m' = m \cdot (2 + z + 2\ell + 1)$, and $\gamma = p \cdot B \cdot m^{O(d+\log \ell)}$.*

Proof. Let $\mathbf{v} \in \mathbb{Z}_q^{(2+z+2\ell+1)m}$ be an input to 1D-SIS- $\mathbf{R}_{q,p,\gamma,m'}$. Then define $\mathbf{a}_\beta, \mathbf{b}_i, \mathbf{c}_i, \mathbf{d}$ be so that their concatenation is \mathbf{v} .

The reduction executes \mathbf{H}_6 as the challenger, using the vectors defined above. We claim that if **Borderline** occurs, then we solve 1D-SIS-R. This follows since if **Borderline** occurs then we found x, j such that

$$\begin{aligned} (\mathbf{b}_{\mathcal{U},x}^T \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi,w})\mathbf{G}^{-1}(\mathbf{D}) - F(x^{(i)})\mathbf{c}_{\Pi,w}^T \mathbf{G}^{-1}(\mathbf{D}) + F(x)\Pi_L(w)\mathbf{d}^T)\mathbf{u}_j \\ \in [-E, E] + (q/p)\mathbb{Z} . \end{aligned}$$

However, by Lemma 4.1, it follows that

$$\begin{aligned} \mathbf{b}_{\mathcal{U},x}^T &= \sum_{\beta \in \{0,1\}} \mathbf{a}_\beta^T \mathbf{R}'_{1,\beta} + \sum_{i \in [z]} \mathbf{b}_i^T \mathbf{R}'_{2,i} \\ \mathbf{c}_{\Pi,x}^T &= \sum_{\beta \in \{0,1\}} \mathbf{a}_\beta^T \mathbf{R}''_{1,\beta} + \sum_{i \in [2\ell]} \mathbf{c}_i^T \mathbf{R}''_{3,i} \end{aligned}$$

where $\|\mathbf{R}'_{1,\beta}\|, \|\mathbf{R}'_{2,i}\| \leq m^{O(d)}$ and $\|\mathbf{R}''_{1,\beta}\|, \|\mathbf{R}''_{3,i}\| \leq m^{O(\log \ell)}$. It follows that there exists an (efficiently derivable) matrix \mathbf{R}_0 such that

$$\mathbf{b}_{\mathcal{U},x}^T \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi,w})\mathbf{G}^{-1}(\mathbf{D}) - F(x^{(i)})\mathbf{c}_{\Pi,w}^T \mathbf{G}^{-1}(\mathbf{D}) + F(x)\Pi_L(w)\mathbf{d}^T = \mathbf{v}^T \mathbf{R}_0 ,$$

and $\|\mathbf{R}_0\| \leq m^{O(d+\log \ell)}$.

Finally,

$$\langle \mathbf{v}, \mathbf{R}_0 \cdot \mathbf{u}_j \rangle \in [-E, E] + (q/p)\mathbb{Z} ,$$

with $\|\mathbf{R}_0 \cdot \mathbf{u}_j\| \leq \|\mathbf{R}_0\| \leq m^{O(d+\log \ell)}$ and $E = B \cdot m^{O(d+\log \ell)} = m^{O(d+\log \ell)}$. Thus $\mathbf{R}_0 \cdot \mathbf{u}_j$ is a valid solution for 1D-SIS- $\mathbf{R}_{q,p,\gamma,m'}$. The claim thus follows. \blacksquare

Putting together Claim 5.1.1 and Eq. (5), we get that

$$\Pr_{\mathbf{H}_5}[\mathbf{Borderline}] \leq \Pr_{\mathbf{H}_6}[\mathbf{Borderline}] + \text{negl}(\lambda) \leq \text{negl}(\lambda) .$$

and thus, finally

$$|\text{Adv}_{\mathbf{H}_5}(\mathcal{A}) - \text{Adv}_{\mathbf{H}_6}(\mathcal{A})| \leq \text{negl}(\lambda) .$$

Finally, we notice that the vector \mathbf{d} is only used when answering the challenge query in the case of $b = 0$. This means that in the adversary's view, the answer it gets when $b = 0$ is uniform and independent of its view so far, exactly the same as the case $b = 1$ where an actual random vector is returned. It follows that

$$\text{Adv}_{\mathbf{H}_6}(\mathcal{A}) = 1/2 .$$

On the other hand

$$\text{Adv}_{\mathbf{H}_6}(\mathcal{A}) \geq 1/2 + \epsilon \Delta_t / 4 - \text{negl}(\lambda) ,$$

and thus

$$\epsilon \leq \frac{\text{negl}(\lambda)}{\Delta_t / 4} = \text{negl}(\lambda) .$$

It follows that \mathcal{A} cannot achieve a noticeable advantage in the constrained PRF experiment under the DLWE $_{q,n,\chi}$ assumption. \square

5.4 Computational Functionality Preserving

We now prove the computational functionality preservation of our scheme, as per Definition 3.1. Throughout this section, we let \mathcal{F} be a family of z -uniform functions and let d be a depth bound (both can depend on λ). We let $n = n(\lambda, d)$, $m = m(\lambda, d)$, $q = q(\lambda, d)$, $p = p(\lambda, d)$ and the noise distributions $\chi = \chi(\lambda, d)$ be as defined in Section 5.2. We let \mathcal{H} be the family of admissible hash functions as described in Section 3.4, with range $\{0, 1\}^\ell$.

Theorem 5.2. *Let \mathcal{F} be a family of z -uniform functions and let d be a depth bound (both can depend on λ). Let $n = n(\lambda, d)$, $m = m(\lambda, d)$, $q = q(\lambda, d)$, $p = p(\lambda, d)$ and the noise distributions $\chi = \chi(\lambda, d)$ be as defined in Section 5.2. Further let $m' = m \cdot (z + 2\ell + 3)$, and $\gamma = \omega(\sqrt{n \log \lambda}) \cdot p \cdot m^{O(d + \log \ell)}$.*

Assuming the hardness of $\text{DLWE}_{n,q,\chi}$ and $\text{1D-SIS-R}_{q,p,\gamma,m'}$, the scheme CPRF is computationally functionality preserving.

We note that the hardness of both assumptions translates to the worst case hardness of approximating lattice problems such as GapSVP and SIVP to within sub-exponential factors.

outline. The theorem follows from an argument practically identical to that made in Hybrids H_5, H_6 of the proof of Theorem 5.1.

Recall that we showed that **Borderline** events only happen with negligible probability, and therefore with all but negligible probability, it holds that the PRF value at point $x^{(i)}$ is exactly equal to

$$\left[\mathbf{b}_{\mathcal{U}, x^{(i)}}^T \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi, w^{(i)}}) \mathbf{G}^{-1}(\mathbf{D}) - F(x^{(i)}) \mathbf{c}_{\Pi, w^{(i)}}^T \mathbf{G}^{-1}(\mathbf{D}) + F(x^{(i)}) \Pi_L(w^{(i)}) \mathbf{d}^T \right]_p.$$

However, when $F(x^{(i)}) = 0$, this term simplifies to

$$\left[\mathbf{b}_{\mathcal{U}, x^{(i)}}^T \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi, w^{(i)}}) \mathbf{G}^{-1}(\mathbf{D}) \right]_p$$

which is exactly $\text{ConstrainEval}(K_F, \mathcal{PP}, x^{(i)})$ by definition. Functionality is thus preserved with all but negligible probability. \square

5.5 Other Properties

We describe several other properties that our construction satisfies.

Unconditional Almost-Correctness. We have shown that our constrained PRF satisfies a computational correctness property, namely that it is hard to find an input \mathbf{x} such that $\text{PRF}_K(\mathbf{x}) \neq \text{ConstrainEval}(K_F, \mathcal{PP}, \mathbf{x})$. We are also able to show unconditionally that the constrained evaluation and the actual PRF evaluation do not differ by much, for any input \mathbf{x} . Indeed, by Equation 3 and 4, we have

$$\|\text{PRF}_K(\mathbf{x}) - \text{ConstrainEval}(K_F, \mathcal{PP}, \mathbf{x})\|_\infty \leq m^{O(d)} \cdot B$$

Key Homomorphism. Our PRF is also “almost key homomorphic” in the sense that $\text{PRF}_s(\mathbf{x}) + \text{PRF}_{s'}(\mathbf{x})$ is close to $\text{PRF}_{s+s'}(\mathbf{x})$ for any keys s and s' and any input \mathbf{x} . Recall that our PRF is

$$\text{PRF}_s(\mathbf{x}) = \left[\mathbf{s}^T \mathbf{B}_{\mathcal{U}} \cdot \mathbf{G}^{-1}(\mathbf{C}_{\Pi}) \cdot \mathbf{G}^{-1}(\mathbf{D}) \right]_p$$

For any keys \mathbf{s}_i and input \mathbf{x} , denoting $\mathbf{s}_i^T \mathbf{B}_U \cdot \mathbf{G}^{-1}(\mathbf{C}_\Pi) \cdot \mathbf{G}^{-1}(\mathbf{D})$ as \mathbf{h}_i , we have

$$\|\text{PRF}_{\sum \mathbf{s}_i}(\mathbf{x}) - \sum \text{PRF}_{\mathbf{s}_i}(\mathbf{x})\|_\infty = \left\| \left[\sum_i \mathbf{h}_i \right]_p - \sum_i [\mathbf{h}_i]_p \right\|_\infty \leq k + 1$$

Constrained-Key Homomorphism. Our constrained keys are “almost homomorphic” as well, in the same sense as above. That is, if K_F and K'_F are constrained versions of PRF keys K and K' for the same function F , the summation $K_F + K'_F$ is a constrained version of $K + K'$ for the function F . For any input \mathbf{x} , we then have that $\text{ConstrainEval}(K_F + K'_F, \mathcal{PP}, \mathbf{x})$ is close to $\text{PRF}_{K+K'}(\mathbf{x})$.

We remark that techniques similar to what we used in showing computational correctness can be used to strengthen the almost key-homomorphism property into computational key-homomorphism where it is computationally hard to find an input for which key homomorphism does not hold.

6 Succinct Constrained Keys

In this section we show how to reduce the size of the constrained key so that asymptotically it depends only on the security parameter and independent of the function class. The construction builds upon the scheme \mathcal{CPRF} from Section 5 but reduces the key size by utilizing an attribute based encryption scheme (ABE). In particular, the constrained keys in our new system have size $\text{poly}(\lambda)$, independent of the parameters of the constraining circuit (namely, its size or depth).

Our *succinct* constrained PRF \mathcal{SCPRF} for a z -uniform d -depth-bounded family \mathcal{F} works as follows.

- $\text{KeyGen}(1^\lambda, 1^z, 1^d)$: The key generation algorithm takes as input the maximum size z and depth d of the constraining circuits. Let $t = O(\log z)$ to be specified later.

It starts by calling $\mathcal{CPRF}.\text{KeyGen}(1^\lambda, 1^z, 1^d)$ to obtain the seed \mathbf{s} , and public parameters $\mathcal{PP} = (H, \mathbf{A}_0, \mathbf{A}_1, \{\mathbf{B}_i\}_{i \in [z]})$.

It then generates: $\mathbf{a}_\beta = \mathbf{s}^T(\mathbf{A}_\beta + \beta \mathbf{G}) + \mathbf{e}_{1,\beta}^T$ and $\mathbf{b}_{i,\beta} = \mathbf{s}^T(\mathbf{B}_i + \beta \mathbf{G}) + \mathbf{e}_{2,i,\beta}^T$. Note that any possible constrained key of \mathcal{CPRF} consists of \mathbf{a}_0 and \mathbf{a}_1 , together with a subset of $\{\mathbf{b}_{i,\beta}\}_{i \in [z], \beta \in \{0,1\}}$.

Next it generates parameters for the ABE scheme $(\mathcal{ABE}.\text{msk}, \mathcal{ABE}.\text{pp}) \leftarrow \mathcal{ABE}.\text{Setup}(1^\lambda, 1^t)$, and generates $\text{ct}_{i,\beta} \leftarrow \mathcal{ABE}.\text{Enc}(\mathcal{ABE}.\text{pp}, (i, \beta), \mathbf{b}_{i,\beta})$, encryptions with (i, β) as the “attributes” and $\mathbf{b}_{i,\beta}$ as the “message”.

The public parameters consist of

$$\mathcal{SCPRF}.\mathcal{PP} = (\mathcal{CPRF}.\mathcal{PP}, \mathcal{ABE}.\mathcal{PP}, \mathbf{a}_0, \mathbf{a}_1, \{\text{ct}_{i,\beta}\}_{i,\beta})$$

The seed for \mathcal{SCPRF} contains a seed for \mathcal{CPRF} , namely a uniformly random vector $\mathbf{s} \in \mathbb{Z}_q^n$, and in addition the ABE master secret key $\mathcal{ABE}.\text{msk}$. We note that in fact \mathbf{s} can be retrieved from the public parameters using $\mathcal{ABE}.\text{msk}$ and therefore it is not necessary to give it explicitly. However, it is more natural to think of \mathbf{s} as a part of the seed. In particular \mathbf{s} will be used to evaluate \mathcal{SCPRF} (see Eval below) and $\mathcal{ABE}.\text{msk}$ will be used to produce constrained keys (see Constrain below).

- $\text{Eval}(\mathbf{s}, \mathcal{PP}, \mathbf{x})$ takes as input the PRF seed \mathbf{s} , the public parameters \mathcal{PP} which contains $\mathcal{CPRF}.pp$, and an input $x \in \{0, 1\}^k$ such that $k \leq k_d$ (i.e. \mathcal{U}_k is of depth $\leq d$), and outputs the result of the CPRF evaluation, namely $\mathcal{CPRF}.Eval(\mathbf{s}, \mathcal{CPRF}.pp, \mathbf{x})$.
- $\text{Constrain}(\mathcal{ABE}.msk, F)$ takes as input the ABE master secret key $\mathcal{ABE}.msk$ and a circuit F (represented as a string in $\{0, 1\}^z$) and does the following. Consider the function:

$$\phi_F(i, \beta) = \begin{cases} 1, & \text{if } F_i = \beta \\ 0, & \text{otherwise} \end{cases}$$

Note that ϕ_F can be computed by a depth $O(\log z)$ circuit (whose depth is independent of the depth of F itself), the parameter t from above is set to be equal to this depth. We recall Section 3.5

The constrained key for F is the ABE token for ϕ_F , namely

$$K_F = \mathcal{ABE}.KeyGen(\mathcal{ABE}.msk, \phi_F)$$

- $\text{ConstrainEval}(K_F, \mathcal{PP}, \mathbf{x})$ takes as input the constrained key K_F , the public parameters \mathcal{PP} and an input \mathbf{x} .

Recalling that $\mathcal{PP} = (\mathcal{CPRF}.pp, \mathcal{ABE}.pp, \mathbf{a}_0, \mathbf{a}_1, \{\text{ct}_{i,\beta}\})$, and that K_F is the ABE decryption key for the function ϕ_F , it first decrypts to obtain $\mathbf{b}_i = \mathcal{ABE}.Dec(K_F, \text{ct}_{i,F_i})$, and then applies the constrained evaluation algorithm $\mathcal{CPRF}.ConstrainEval((\mathbf{a}_0, \mathbf{a}_1, \{\mathbf{b}_i\}), \mathcal{CPRF}.PP, \mathbf{x})$.

The correctness follows in a straightforward manner from the correctness of \mathcal{ABE} and \mathcal{CPRF} . The constrained key size of \mathcal{SCPRF} is derived from that of \mathcal{ABE} and is $\text{poly}(\lambda, t) = \text{poly}(\lambda, \log z)$. It follows that there exists a $\text{poly}(\lambda)$ asymptotic upper bound on the key sizes that applies for all polynomial values of z . Security is proven in the following theorem.

Theorem 6.1. *If \mathcal{CPRF} is a single-key secure constrained pseudorandom function for function class \mathcal{F} (Definition 3.1), which is built according to the template in Section 5, and if \mathcal{ABE} is a selectively secure ABE scheme (Definition 3.6), then the scheme \mathcal{SCPRF} described above is a secure single-key CPRF for \mathcal{F} .*

Proof. Let \mathcal{A} be a CPRF adversary against \mathcal{SCPRF} . The proof will proceed by a sequence of hybrids where in each hybrid the challenger will sample a random bit b and the adversary's success in inferring b will be \mathcal{A} 's advantage in the hybrid.

Hybrid H_0 . This hybrid is the constrained PRF security game for \mathcal{SCPRF} . The challenger generates $(\mathcal{SK}, \mathcal{PP}) \leftarrow \text{KeyGen}(1^\lambda, 1^z, 1^d)$. It gets $F \in \{0, 1\}^z$ from \mathcal{A} and produces a constrained key $K_F \leftarrow \text{Constrain}(\mathcal{SK}, \mathcal{PP}, F)$. It then sends \mathcal{PP}, K_F to \mathcal{A} . At this point \mathcal{A} adaptively makes queries $x^{(i)} \in \{0, 1\}^*$, and the challenger computes $y^{(i)} \leftarrow \text{Eval}(\mathbf{s}, \mathcal{PP}, x^{(i)})$ and returns it to \mathcal{A} . Finally, \mathcal{A} outputs $x^* \in \{0, 1\}^*$. If $b = 0$ then the challenger returns $y^* \leftarrow \text{Eval}(\mathbf{s}, \mathcal{PP}, x^*)$, and if $b = 1$ it returns a random y^* .

$$\text{Adv}_{H_0}(\mathcal{A}) \geq 1/2 + \epsilon .$$

Hybrid H_1 . In this hybrid, the challenger does the following. It first receives the function F and then generates $(\mathcal{SK}, \mathcal{PP})$ with one change compared to the previous hybrid. The ciphertexts $\text{ct}_{i,1-F_i}$ will now be generated as $\mathcal{ABE}.\text{Enc}(\mathcal{ABE}.\text{pp}, (i, \beta), 0)$.

Claim 6.1.1. *Under the selective security of \mathcal{ABE} , it holds that*

$$|\text{Adv}_{H_0}(\mathcal{A}) - \text{Adv}_{H_1}(\mathcal{A})| \leq \text{negl}(\lambda) ,$$

Proof. Let \mathcal{B} be the following adversary against multi-message selective security of \mathcal{ABE} (see Definition 3.6), which will work by simulating the interaction of \mathcal{A} in the CPRF security game against \mathcal{SCPRF} . First of all \mathcal{B} generates a key pair $(\mathbf{s}, \mathcal{CPRF}.\mathcal{PP}) \leftarrow \mathcal{CPRF}.\text{KeyGen}(1^\lambda, 1^z, 1^d)$, and produces the vectors $\mathbf{a}_\beta, \mathbf{b}_{i,\beta}$ as in the key generation process of \mathcal{SCPRF} .

Then \mathcal{B} runs \mathcal{A} to obtain the function description $F \in \{0, 1\}^z$. It sends to the \mathcal{ABE} challenger the attribute sequence $\{(i, 1 - F_i)\}_{i \in [z]}$. Then, for each i , it will send to the \mathcal{ABE} challenger the message pair $m_{0,i} = \mathbf{b}_{i,1-F_i}, m_{1,i} = 0$. It receives $\mathcal{ABE}.\mathcal{PP}$ and ciphertexts $\text{ct}_{i,1-F_i}$ which encrypt either $m_{0,i}$ or $m_{1,i}$. Using $\mathcal{ABE}.\mathcal{PP}$ it generates ct_{i,F_i} by itself as in $\mathcal{SCPRF}.\text{KeyGen}$. Further, \mathcal{B} generates $\mathcal{SCPRF}.\mathcal{PP}$ using $\mathcal{CPRF}.\mathcal{PP}, \mathcal{ABE}.\mathcal{PP}$ and $\text{ct}_{i,\beta}$, and forwards this value to \mathcal{A} . Note that this is distributed identically to an $\mathcal{SCPRF}.\mathcal{PP}$ in H_0 if $b = 0$ and identically to $\mathcal{SCPRF}.\mathcal{PP}$ in H_1 if $b = 1$.

Next, \mathcal{B} queries the \mathcal{ABE} challenger on the function token ϕ_F , noting that $\phi_F(i, 1 - F_i) = 0$ for all i . The challenger responds with the appropriate token, which will be forwarded to \mathcal{A} as K_F . Note that this value is correctly distributed.

The adversary \mathcal{B} continues to simulate \mathcal{A} , answering its oracle queries using the seed \mathbf{s} . Finally, when \mathcal{A} halts and outputs some b' , \mathcal{B} halts as well and outputs b' as its own output.

By definition, the advantage of \mathcal{B} against \mathcal{ABE} is exactly $\text{Adv}_{H_0}(\mathcal{A}) - \text{Adv}_{H_1}(\mathcal{A})$, and the claim follows from the selective security of \mathcal{ABE} . \blacksquare

Next, we notice that the adversary's advantage in this hybrid cannot be noticeable without breaking the security of \mathcal{CPRF} .

Claim 6.1.2. *If \mathcal{CPRF} is a secure single-key constrained PRF then $|\text{Adv}_{H_1}(\mathcal{A}) - 1/2| = \text{negl}(\lambda)$.*

Proof. We present an adversary \mathcal{B} against \mathcal{CPRF} whose advantage is $|\text{Adv}_{H_1}(\mathcal{A}) - 1/2|$ as follows. It will first get F from \mathcal{A} and forward it to the \mathcal{CPRF} challenger. Then, upon receiving $\mathcal{CPRF}.\mathcal{PP}, \mathbf{a}_\beta, \mathbf{b}_i$, it will generate $(\mathcal{ABE}.\text{msk}, \mathcal{ABE}.\mathcal{PP})$ by itself. Then it will encrypt \mathbf{b}_i as \mathbf{b}_{i,F_i} to obtain ct_{i,F_i} , and will encrypt zero to obtain $\text{ct}_{i,1-F_i}$. Finally it will generate K_F by running $\mathcal{ABE}.\text{KeyGen}$ on the function ϕ_F . This will allow generating $\mathcal{SCPRF}.\mathcal{PP}, K_F$ which are consistent with the distribution that \mathcal{A} receives in H_1 .

The values $\mathcal{SCPRF}.\mathcal{PP}, K_F$ will be sent to \mathcal{A} , and when \mathcal{A} makes PRF queries they will be forwarded to the \mathcal{CPRF} challenger, and the response forwarded back to \mathcal{A} . In addition, \mathcal{A} 's challenge will be forwarded, and the response forwarded back. When \mathcal{A} terminates and returns b' , the same b' will be returned by \mathcal{B} .

It is straightforward to see that whenever \mathcal{A} wins in H_1 , \mathcal{B} wins against \mathcal{CPRF} . The claim follows. \blacksquare

Putting the two claims together, it follows that

$$1/2 + \epsilon \leq \text{Adv}_{H_0}(\mathcal{A}) \leq \text{Adv}_{H_1}(\mathcal{A}) + \text{negl}(\lambda) \leq 1/2 + \text{negl}(\lambda) ,$$

which completes the proof of the theorem. \square

References

- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 99–108. ACM, 1996.
- [Ale03] Michael Alekhnovich. More on average case vs approximation complexity. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, pages 298–307. IEEE Computer Society, 2003.
- [BB04] Dan Boneh and Xavier Boyen. Secure identity based encryption without random oracles. In *Advances in Cryptology - CRYPTO 2004, 24th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings*, pages 443–459, 2004.
- [BFKL93] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, pages 278–291, 1993.
- [BGG⁺14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, pages 533–556, 2014.
- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In *Public-Key Cryptography - PKC 2014 - 17th International Conference on Practice and Theory in Public-Key Cryptography, Buenos Aires, Argentina, March 26-28, 2014. Proceedings*, pages 501–519, 2014.
- [BLMR13] Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic prfs and their applications. In *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, pages 410–428, 2013.
- [BP14] Abhishek Banerjee and Chris Peikert. New and improved key-homomorphic pseudorandom functions. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 353–370. Springer, 2014.
- [BPR12] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 719–737. Springer, 2012.

- [BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part II*, pages 280–300, 2013.
- [CHKP12] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. *J. Cryptology*, 25(4):601–639, 2012.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986. Extended abstract in FOCS 84.
- [GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *Proceedings of the 13th ACM Conference on Computer and Communications Security, CCS 2006, Alexandria, VA, USA, Ioctober 30 - November 3, 2006*, pages 89–98. ACM, 2006.
- [GPV07] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(133), 2007.
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 545–554. ACM, 2013.
- [HKKW14] Dennis Hofheinz, Akshay Kamath, Venkata Koppula, and Brent Waters. Adaptively secure constrained pseudorandom functions. Cryptology ePrint Archive, Report 2014/720, 2014. <http://eprint.iacr.org/>.
- [KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013*, pages 669–684. ACM, 2013.
- [LMR14] Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Improved constructions of prfs secure against related-key attacks. In Ioana Boureanu, Philippe Owesarski, and Serge Vaudenay, editors, *Applied Cryptography and Network Security - 12th International Conference, ACNS 2014, Lausanne, Switzerland, June 10-13, 2014. Proceedings*, volume 8479 of *Lecture Notes in Computer Science*, pages 44–61. Springer, 2014.
- [MM11] Daniele Micciancio and Petros Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, pages 465–484, 2011.

- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, pages 700–718, 2012.
- [NPR99] Moni Naor, Benny Pinkas, and Omer Reingold. Distributed pseudo-random functions and kdcs. In Jacques Stern, editor, *Advances in Cryptology - EUROCRYPT '99, International Conference on the Theory and Application of Cryptographic Techniques, Prague, Czech Republic, May 2-6, 1999, Proceeding*, volume 1592 of *Lecture Notes in Computer Science*, pages 327–346. Springer, 1999.
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 333–342, 2009.
- [Reg04] Oded Regev. Lattices in computer science - average case hardness. Lecture Notes for Class (scribe: Elad Verbin), 2004. http://www.cims.nyu.edu/~regev/teaching/lattices_fall_2004/ln/averagecase.pdf.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93, 2005.
- [Sch87] Claus-Peter Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 475–484. ACM, 2014.