

Hierarchical Functional Encryption^{*†}

Zvika Brakerski¹, Nishanth Chandran², Vipul Goyal³,
Aayush Jain⁴, Amit Sahai⁵, and Gil Segev⁶

- 1 Weizmann Institute of Science, Rehovot, IL
zvika.brakerski@weizmann.ac.il
- 2 Microsoft Research India, Bangalore, IN
nichandr@microsoft.com
- 3 CMU, Pittsburgh, USA
goyal@cs.cmu.edu
- 4 UCLA, Los Angeles, USA
aayushjain1728@gmail.com
- 5 UCLA, Los Angeles, USA
sahai@cs.ucla.edu
- 6 Hebrew University of Jerusalem, Jerusalem, IL
segev@cs.huji.ac.il

Abstract

Functional encryption provides fine-grained access control for encrypted data, allowing each user to learn only specific functions of the encrypted data. We study the notion of *hierarchical* functional encryption, which augments functional encryption with *delegation* capabilities, offering significantly more expressive access control.

We present a *generic transformation* that converts any general-purpose public-key functional encryption scheme into a hierarchical one without relying on any additional assumptions. This significantly refines our understanding of the power of functional encryption, showing that the existence of functional encryption is equivalent to that of its hierarchical generalization.

Instantiating our transformation with the existing functional encryption schemes yields a variety of hierarchical schemes offering various trade-offs between their delegation capabilities (i.e., the depth and width of their hierarchical structures) and underlying assumptions. When starting with a scheme secure against an unbounded number of collusions, we can support *arbitrary* hierarchical structures. In addition, even when starting with schemes that are secure against a bounded number of collusions (which are known to exist under rather minimal assumptions such as the existence of public-key encryption and shallow pseudorandom generators), we can support hierarchical structures of bounded depth and width.

1998 ACM Subject Classification E.3 Data Encryption

Keywords and phrases Functional encryption

* In this extended abstract we present results from [23] and [25].

† Z. Brakerski is supported by the Israel Science Foundation (Grant No. 468/14), the Alon Young Faculty Fellowship, Binational Science Foundation (Grant No. 712307) and Google Faculty Research Award.

A. Sahai and A. Jain are supported in part from a DARPA/ARL SAFEWARE award, NSF Frontier Award 1413955, NSF grants 1619348, 1228984, 1136174, and 1065276, a Xerox Faculty Research Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant. This material is based upon work supported by the Defense Advanced Research Projects Agency through the ARL under Contract W911NF-15-C-0205. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

G. Segev is supported by the European Union's 7th Framework Program (FP7) via a Marie Curie Career Integration Grant, by the Israel Science Foundation (Grant No. 483/13), by the Israeli Centers of Research Excellence (I-CORE) Program (Center No. 4/11), by the US-Israel Binational Science Foundation (Grant No. 2014632), and by a Google Faculty Research Award.



1 Introduction

The rapidly evolving vision of functional encryption [49, 15, 48] offers tremendous flexibility when accessing encrypted data. Specifically, functional encryption schemes support restricted decryption keys that allow users to learn specific functions of the encrypted data and nothing else. Motivated by the early examples of functional encryption schemes for specific functionalities (such as identity-based encryption [51, 13, 27]), extensive research has recently been devoted to the study of functional encryption (see, for example, [49, 15, 48, 38, 3, 9, 19, 30, 37, 5, 53, 32, 7, 22, 42, 21] and the references therein).

In a functional encryption scheme, a trusted authority holds a master secret key known only to the authority. When the authority is given the description of some function f as input, it uses its master secret key to generate a functional key sk_f associated with the function f . Now, anyone holding the functional key sk_f and an encryption of any message x , can compute $f(x)$ but cannot learn any additional information about the message x . Such fine-grained access to encrypted data is extremely useful for a wide variety of applications, including expressive access control, spam filtering, mining encrypted databases, and more (we refer the reader to the survey by Boneh, Sahai and Waters [16] for an in-depth discussion of these applications).

Hierarchical functional encryption. Motivated by the applicability of functional encryption to expressive access-control systems, in this paper we study the notion of *hierarchical* functional encryption, which was introduced by Ananth, Boneh, Garg, Sahai and Zhandry [4]. The hierarchical notion augments standard functional encryption with delegation capabilities, enabling significantly more expressive access control.

Specifically, recall that in a functional encryption scheme, the holder of the master secret key msk can generate a functional key sk_f corresponding to any given function f . In a hierarchical functional encryption scheme, the holder of any such functional key sk_f can in turn generate a functional key $\text{sk}_{g \circ f}$ corresponding to the function $g \circ f$ for any given function g . Now, anyone holding the delegated functional key $\text{sk}_{g \circ f}$ and an encryption of any message x , can compute $(g \circ f)(x) = g(f(x))$ but cannot learn any additional information about the message x . Such expressive delegation capabilities give rise to *hierarchical* access control, which is a sought-after ingredient in modern access control systems. In particular, the notion of hierarchical functional encryption generalizes those of hierarchical attribute-based encryption, hierarchical identity-based encryption and more (see the discussion at the end of Section 3 on the delegation capabilities of functional encryption).

Ananth et al. formalized a notion of security for hierarchical functional encryption schemes, and sketched how the functional encryption scheme of Garg et al. [30] can be transformed into a hierarchical one by using a general-purpose indistinguishability obfuscator [8, 30].¹ Their approach, however, is both tailored to the specific functional encryption scheme of Garg et al. [30], and can only support hierarchical structures of *constant* depth (i.e., can only support a constant number of successive delegations).²

¹ It was recently shown by Ananth and Jain [6] and by Bitansky and Vaikuntanathan [10] that indistinguishability obfuscation can be constructed from some flavors of functional encryption. Specifically, from *succinct* functional encryption with *sub-exponential security*. Our approach is both more direct and requires no such properties.

² In the hierarchical scheme of Ananth et al. [4], a delegated functional key $\text{sk}_{g \circ f}$ for the function $g \circ f$ is

1.1 Our Contributions

We present a *generic transformation* that converts any general-purpose public-key functional encryption scheme into a hierarchical one without relying on any additional assumptions (and, in particular, without relying on indistinguishability obfuscation). Our transformation allows instantiations based both on unbounded-collusion functional encryption schemes and on bounded-collusion ones. This level of generality yields a variety of hierarchical schemes based on various assumptions in the standard model. These include strong assumptions such as indistinguishability obfuscation [30, 53] or multilinear maps [32], and much milder assumptions such as learning with errors [37], and even the existence of any public-key encryption scheme and low-depth pseudorandom generator [38].

We stress that our result stands even if it turns out that indistinguishability obfuscation is impossible to achieve or requires strong computational assumptions. One could view it as evidence that the hierarchical properties do not stem from the “obfuscation-like” features of functional encryption, but rather from the more rudimentary properties that are achievable under minimal assumptions.

Security and efficiency. In terms of security, the schemes resulting from our transformation guarantee semi-adaptive security. Moreover, by assuming that the underlying functional encryption scheme is sub-exponentially secure, we obtain adaptive security via a standard complexity leveraging argument. In terms of efficiency, our approach results in keys and ciphertexts with a rather low overhead compared to the efficiency of the underlying functional encryption scheme. A ciphertext in our scheme is essentially a ciphertext of the underlying scheme. A delegated functional key for $g \circ f$ contains two functional keys for functions that essentially compute f and g , respectively, in addition to some basic cryptographic computation (an evaluation of a PRF, and an encryption of a ciphertext of the underlying scheme). We refer the reader to the overview below for more details on the security and efficiency of the schemes resulting from our approach.

Instantiations. The variety of schemes resulting from our transformation offer different trade-offs between their underlying assumptions and their delegation capabilities (i.e., the depth and width of the hierarchical structures that they support). For example, instantiating our transformation with the schemes of Garg et al. [30] and Waters [53] results in hierarchical schemes that support hierarchical structures of any polynomial depth and any polynomial width (where these polynomials do not have to be specified in advance). In addition, instantiating our transformation with the schemes of Goldwasser et al. [37] or Gorbunov et al. [38] results in hierarchical schemes that support hierarchical structures of any constant depth and any pre-determined polynomial width. This should be compared to the hierarchical scheme of Ananth et al. [4] that is constructed from the specific functional encryption scheme of Garg et al. [30], and can only support hierarchical structures of *constant* depth, and to the alternative hierarchical scheme in this work (see Section 1.3) that is constructed based on even stronger assumptions and still requires an a-priori bound on the depth of the support hierarchical structures. We refer the reader to Table 1 for a comparison of the assumptions underlying the known hierarchical functional encryption schemes and of their supported hierarchical structures.

computed from sk_f by applying the obfuscator to a program that contains sk_f as part of its description. Thus, since sk_f itself consists of such an obfuscated program, this allows for only a constant number of successive delegations.

■ **Table 1** A comparison of the assumptions underlying the known hierarchical functional encryption schemes and of their supported hierarchical structures. We note that indistinguishability obfuscation ($i\mathcal{O}$) is known to imply unbounded-collusion functional encryption [53], which in turn clearly implies bounded-collusion functional encryption. In addition, bounded-collusion functional encryption is implied by much milder assumptions such as learning with errors [37], and even the existence of any public-key encryption scheme and low-depth pseudorandom generator [38].

| Assumption | Hierarchical Structure | |
|--|------------------------|----------------------|
| | Depth | Width |
| $i\mathcal{O}$ [4] | Constant | Unbounded |
| Sub-exponentially-secure $i\mathcal{O}$ and sub-exponentially-secure HIBE (this work [25]) | Any fixed polynomial | Unbounded |
| Unbounded-collusion FE (this work [23]) | Unbounded | Unbounded |
| Bounded-collusion FE (this work [23]) | Constant | Any fixed polynomial |

1.2 Overview of Our Approach

Formally, a hierarchical FE scheme contains the standard (Setup, KG, Enc, Dec) algorithms, in addition to a new *delegation* algorithm `Delegate`. The delegation algorithm `Delegate(hskf, g)` is identical in syntax to the KG algorithm, except that it takes a functional key hsk_f (which can itself be the output of a previous delegation) instead of the master secret key msk . Its output is a key $\text{hsk}_{g \circ f}$ corresponding to the composed function $g \circ f$.

Indeed, the way we implement this functionality is by associating a unique master secret key with any delegable functional key. Namely, a delegable key hsk_f (with respect to the master key pair (msk, mpk)) will contain a fresh master secret key msk' in addition to a “standard” functional key for a re-encryption function $\text{sk}_{\text{ReEnc}_{f, \text{mpk}'}}$ (the key pair $(\text{msk}', \text{mpk}')$ is generated using the standard setup procedure). The function $\text{ReEnc}_{f, \text{mpk}'}(x)$, intuitively, takes an input x and outputs a *functional encryption* of $f(x)$ under the new key mpk' . It is obvious that since msk' is a part of hsk_f , then the owner of $\text{hsk}_f = (\text{sk}_{\text{ReEnc}_{f, \text{mpk}'}, \text{msk}'})$ can derive $f(x)$ itself if it so desires. The beauty of this procedure is that it can then be repeated. If hsk_f needs to be delegated via `Delegate(hskf, g)`, then one only needs to generate a new pair $(\text{msk}'', \text{mpk}'')$, use msk' to obtain $\text{sk}'_{\text{ReEnc}_{g, \text{mpk}''}}$ and output $\text{hsk}_{g \circ f} = (\text{sk}_{\text{ReEnc}_{f, \text{mpk}'}, \text{sk}'_{\text{ReEnc}_{g, \text{mpk}''}, \text{msk}''})$. In the decryption process, we start with some $\text{ct} = \text{FE.Enc}_{\text{mpk}}(x)$, use the first component of the key to obtain $\text{ct}' = \text{FE.Enc}_{\text{mpk}'}(f(x))$, and then using the second component to obtain $\text{ct}'' = \text{FE.Enc}_{\text{mpk}''}(g(f(x)))$. Finally, msk'' is used to decrypt ct'' and thus learn the value $g(f(x))$.

Care needs to be taken in order to securely realize the above intuition. In particular, one has to come up with a source of randomness for the re-encryption process. This is done by slightly modifying the encryption algorithm of the hierarchical scheme such that $\text{Enc}(\text{mpk}, x) = \text{FE.Enc}(\text{mpk}, (x, K, \perp))$, where K is a key to a puncturable pseudorandom function PRF, and \perp is a placeholder that will only be used in the proof. Similarly, we will generate functional keys of the form $\text{sk}_{\text{ReEnc}_{f, t, \text{mpk}', c}}$, where t is a random tag and c is a random string that will be used in the proof. The function $\text{ReEnc}_{f, t, \text{mpk}', c}(x, K, \perp)$ will compute $f(x)$ and encrypt the tuple $(f(x), K', \perp)$ under msk' using randomness r' . The randomness for the generation of K' and r' is produced by evaluating $\text{PRF}_K(t)$.

For the sake of our security proof, one last addition is made to the description of $\text{ReEnc}_{f,t,\text{mpk}',c}$. If its input is of the form (\cdot, \cdot, k) , where k is a key for a symmetric encryption scheme, then the first two arguments are ignored and $\text{SKE.Dec}_k(c)$ is output. Thus we implement a “trapdoor circuit” (or a “Trojan”) as per [29, 5].

Security notion. The notions of security that we consider in this work are those formalized by Ananth et al. [4]. Specifically, we consider adversaries that obtain functional keys for various functions of their choice by issuing key-generation queries and delegation queries. We require that such adversaries have only a negligible advantage in distinguishing the encryptions of two challenge messages, x_0^* and x_1^* , of their choice as long as for any function f for which they obtain a functional key it holds that $f(x_0^*) = f(x_1^*)$, where such a key may be produced either as a result of a key-generation query or a delegation query (we refer the reader to Section 3 for more details).

We prove that if the underlying scheme \mathcal{FE} is selectively secure then the resulting hierarchical scheme is selectively secure, and if \mathcal{FE} is semi-adaptively secure then the resulting hierarchical scheme is semi-adaptively secure.³ We leave it as an intriguing open problem to design a hierarchical functional encryption scheme that guarantees adaptive security. We note that already in the less-expressive setting of identity-based encryption, designing adaptively-secure hierarchical schemes is extremely challenging. In particular, Lewko and Waters [46] recently showed why known proof methods fall short of proving adaptive security even for adaptively-secure hierarchical identity-based encryption (which is a special case of hierarchical FE) without degrading exponentially with the number of delegation levels.

Proof overview. Let us focus on the case of selective security, semi-adaptive security follows by a practically identical argument. In the selective security game, the adversary first specifies challenge messages x_0^* and x_1^* , receives mpk , and then makes a sequence of key-generation and delegation queries. One could visualize the structure that is generated by these queries as a tree, whose root is (msk, mpk) and whose nodes are the key pairs that are generated upon each call to KG or Delegate . Each such call generates a new child for one of the nodes in the tree, as per the adversary’s choice. Each node is associated with a function f which was input to $\text{KG}/\text{Delegate}$ in its creation, and also with a function \tilde{f} , which is the composition of all functions from the root to that node. If $\tilde{f}(x_0^*) = \tilde{f}(x_1^*)$ then we say that the node is *observable*, since the adversary is allowed to see the functional key $\text{hsk}_{\tilde{f}}$ associated with that node. We can assume w.l.o.g that all the leaves of the tree are observable.

The high-level intuition of the proof is the following. Let us pretend that ReEnc is actually capable of outputting a re-encrypted ciphertext which is encrypted with true randomness, rather than with pseudorandomness. Now, consider an unobservable node i (i.e., a node corresponding to f_i and \tilde{f}_i for which $\tilde{f}_i(x_0^*) \neq \tilde{f}_i(x_1^*)$) that only has observable children. This means that all functions $\text{ReEnc}_{f,t,\text{mpk}',c}$ that are generated relative to this node’s msk_i will output the same value whether the challenge ciphertext is an encryption of x_0^* or of x_1^* . The security of the underlying scheme will guarantee that the re-encrypted ciphertext cannot be used to distinguish x_0^* from x_1^* . Let us take another leap of faith and pretend that

³ We briefly remind the reader the differences between selective, semi-adaptive, and adaptive security. *Selective* security considers adversaries that specify their challenge messages before seeing the public parameters or making any key queries. *Semi-adaptive* security, as recently defined by Chen and Wee [26], considers adversaries that specify their challenge messages after seeing the public parameters but before making any key queries. Finally, *adaptive* security considers adversaries that specify their challenge messages even after seeing the public parameters and making key queries.

not only the re-encrypted ciphertext cannot distinguish x_0^* from x_1^* but it is in fact identical in both cases. Then the above process can propagate towards the root of the tree, where at every step we increase the number of nodes whose output is the same regardless of whether the challenge ciphertext encrypts x_0^* or x_1^* . Once this process gets all the way to the root and applied to the challenge ciphertext itself, the proof is complete.

This intuition is implemented using the mechanisms of punctured programming [50] and “trapdoor circuits” [29] (or “Trojans” [5]). We will replace the c values in $\text{ReEnc}_{f,t,\text{mpk}',c}$ with symmetric encryptions of our “fantasy ciphertexts” (ones that are encrypted with true randomness), and append the challenge ciphertext with the appropriate symmetric decryption key (in fact, multiple symmetric keys will be needed, one for every level of the hierarchy, and one has to carefully control their propagation along the tree). Puncturable PRFs will be used to argue that the use of fantasy ciphertexts is indistinguishable from the actual output of $\text{ReEnc}_{f,t,\text{mpk}',c}$, which will allow the proof idea from above to go through. This requires a careful and delicate argument since we can only puncture a PRF key that had been generated with fresh randomness, hence one has to also consider fantasy PRF keys and propagate them along the tree as well together with the fantasy ciphertexts. The formal proof thus contains many fine points and a large number of steps, and is provided in Section 4.

1.3 The Multi-Authority Setting and an Alternative Hierarchical Scheme

While above, we have focused on the question of delegating functional keys to other parties, a closely related problem is that of achieving functional encryption in the context of multiple key-issuing authorities. In Multi-Authority Functional Encryption (MAFE), we allow n authorities to “independently” generate their private and public keys. An encryptor should be able to encrypt a message m along with a policy F over the various authorities. Any authority i , should be able to generate a token for a user with identity UID and property U_i . A user with identity UID with tokens for U_i from authority $i \in [n]$, should be able to decrypt the ciphertext to recover $F(U_1, \dots, U_n, m)$. We require that colluding users, say UID_1 and UID_2 , (possibly, in collusion with some corrupt authorities) should not jointly learn anything more from the ciphertext than what they are authorized to.

Our main construction idea is natural – we view a ciphertext as an obfuscated program and functional keys as signatures generated with respect to each authority’s unique public key. To elucidate this main idea, let us first analyze the problem in a single authority scenario. The starting point of our construction is a construction of a functional encryption scheme due to [19] based on *differing-inputs obfuscation* (diO). In their construction, a ciphertext is an obfuscated circuit that checks the signature for a function f and computes f if the signature is valid. The security proof relied on the fact that if the adversary distinguishes the ciphertext (or rather an obfuscated circuit) encrypting m_0 from that of m_1 , then one must be able to extract a signature on a function g such that $g(m_0) \neq g(m_1)$ thereby producing a forgery. This construction is easily scalable to the multi-authority setting. However, given the implausibility of differing-inputs obfuscation [31], our objective is a construction based on indistinguishability obfuscation (iO). While it was observed in [19] that iO behaves like diO for the setting where the circuits in question differ only on a polynomially bounded number of points, unfortunately, there are simply too many points on which these functions would differ for this strategy to yield a construction from iO.

Constructing MAFE based on indistinguishability obfuscation. Our first idea is to use signatures that are *unique*, in the sense, that for every message there exists only one signature that verifies. In order to instantiate unique signatures, we use puncturable PRF's, indistinguishability obfuscation and an injective one way function following the punctured programming technique from [50]. The main idea of the security proof is that we can build exponentially many hybrids corresponding to the function space. We index each hybrid with a function x . In hybrid x , the ciphertext takes as input f with a signature σ and checks if $f < x$. If that is the case, it outputs $f(m_0)$ otherwise $f(m_1)$. We argue indistinguishability between hybrids x and $x + 1$ by intermediate hybrids where we puncture the PRF at x and replace the PRF evaluations with a random sample. Then we further note that if the adversary now distinguishes these hybrids, then it can be used to invert an injective one way function in sub-exponential time. In order to handle key corruptions, we now let the signing key be an obfuscation of a circuit that evaluates a PRF on the input. This allows us to "program" the secret keys using puncturing techniques [50] and the techniques described above so that the proof can be made to go through.

Constructing MAFE based on LWE. Since MAFE for arbitrary policies imply obfuscation it is unlikely to construct them from simpler assumptions. We also study a very natural policy referred to as the $n - \text{out} - \text{of} - n$ threshold policy. In this policy for any function f the decryptor must get tokens from all the authorities to learn $f(m)$. We describe our construction at a very high level below. Our starting point is the recent construction of threshold homomorphic encryption (TFHE) by [47]. In a TFHE scheme the secret key can be shared among n parties and any (evaluated) ciphertext can be partially decrypted using each of the n key shares. The partial decryptions can be finally added to yield the output. The security guarantee is two fold: First, given at most $n - 1$ key shares the semantic security of the encryption holds. Second, partial decryption of any ciphertext using any key share can be statistically simulated using the remaining $n - 1$ shares and the output of the decryption. The (simplified) scheme can be described as follows. Each authority runs an FE setup. To encrypt a message m , the encryptor runs a fresh setup of a TFHE system. He encrypts m as ct and outputs n FE ciphertexts, one corresponding to each authority. For authority i he encrypts $(\text{ct}, \text{sk}_i, *)$ where $*$ represents some additional strings (for example, PRF keys, e.t.c used for programming the proof) and sk_i is i^{th} partial TFHE decryption key. The key for a function f is now simply an FE key for a function that takes as input a TFHE ciphertext ct and evaluates it for function f and then computes a partial decryption using sk_i .

The solution described above suffers from one key flaw. The problem is that the scheme is not *user-collusion-resistant*. One user uid_1 can query tokens for f from say k authorities, while another user uid_2 can query them from other $n - k$ authorities and later collude to recover $f(m)$. We prevent this by relying on Pseudo-Random Zero-Sharing (PRSS) [28]. Using this primitive one can secret share 0 pseudorandomly and deterministically using initially distributed shares $(\beta_1, \dots, \beta_n)$ and any common input. Concretely there exists a publicly known function g such that for any index x , $\{g(\beta_i, x)\}_{i \in [n]}$ forms a pseudorandom secret sharing of 0. We now modify the existing scheme as follows: FE ciphertext for each authority i now also consists β_i and the FE key now produces TFHE partial decryptions masked by $g(\beta_i, \text{uid}, f)$. This masking prevents the adversary to collude in the manner described above.

Applications. We also show that MAFE immediately implies multi-authority attribute based encryption and this yields the first decentralised ABE scheme without setup and without the use of random oracles.

Alternate construction of Hierarchical FE. An application of the unique signature idea described above is that it can be extended to allow for delegation of function keys. In what we described above, if the signature scheme is delegatable (i.e. if it allows us to compute a signature on $f||g$ for any g given a signature on f) then it allows us to compute keys for $g \cdot f$ given a key for f , where $g \cdot f$ denotes the composition of functions f and g and on input x computes $g(f(x))$. In order to construct this primitive, we make use of the exponential hybrid approach and make interesting use of hierarchical identity based encryption (HIBE). The main idea behind this construction is as follows: to encrypt a message m , we let the ciphertext be an obfuscated program that takes as input a “function” f . On input f , the obfuscated program will compute a HIBE encryption of $f(m)$ on identity f using randomness generated from a PRF applied on f . Any decryptor can then decrypt this ciphertext using a HIBE key for identity f or its prefix. Here too, the security proof goes input by input over the space of all circuits. Our construction supports delegation up to any (a-priori bounded) polynomial number of times, in contrast to prior schemes which allowed delegation only up to $O(1)$ number of times [4].

A detailed description of the MAFE construction as well as the alternate construction for HFE can be found in [25].

1.4 Related Work

Hierarchical encryption schemes. Encryption schemes supporting a hierarchical structure have been extensively studied in the setting of identity-based encryption, and have been recently studied in the more general setting of attribute-based encryption and functional encryption.

The line of research on hierarchical identity-based encryption has been extremely successful, starting with schemes in the random oracle model, evolving through selectively-secure schemes in the standard model and graduating to adaptively secure schemes for polynomially many levels. It is far beyond the scope of this paper to provide an extensive overview of this line of research, and we refer the reader to [34, 40, 11, 12, 18, 33, 52, 1, 2, 43, 44, 45, 24, 46] and the references therein.

Recently, Boneh et al. [14] constructed an attribute-based encryption scheme that supports delegation of keys. This scheme enables anyone holding a key sk_P corresponding to a predicate P to generate a key $\text{sk}_{P \wedge Q}$ corresponding to the predicate $P \wedge Q$ for any given predicate Q . Now, given the key $\text{sk}_{P \wedge Q}$ and an encryption of any pair (x, m) , one can recover m if and only if $(P \wedge Q)(x) = 1$. Although the setting of attribute-based encryption is significantly more expressive than the identity-based one, it does not seem to come close to the general setting of functional encryption that we consider in this paper.

Finally, as discussed above, Ananth et al. [4] sketched how the functional encryption scheme of Garg et al. [30] can be transformed into a hierarchical one by using a general-purpose indistinguishability obfuscator. When compared to their scheme our approach offers two main advantages. First, whereas Ananth et al. rely on a specific scheme and on indistinguishability obfuscation, we can rely on any underlying general-purpose scheme. This enables us to rely on a variety of underlying assumptions, including learning with errors and even the existence of any public-key encryption scheme and low-depth pseudorandom

generators, as discussed in Section 1.1. Furthermore, as new functional encryption schemes are presented, they can immediately be plugged in to our construction. Second, the schemes resulting from our transformation guarantee semi-adaptive security, whereas the scheme of Ananth et al. guarantees only the somewhat less realistic notion of selective security.

Encapsulation techniques in functional encryption. Key encapsulation is a very useful approach for improving both the efficiency and the security of encryption schemes. Specifically, key encapsulation typically means that instead of encrypting a message x under a fixed key sk , one can instead sample a fresh key k , encrypt x under k , and then encrypt k under sk . Recently, Ananth et al. [5], followed by Brakerski et al. [21], showed that key encapsulation is useful also for functional encryption, and can be used for generically enhancing the functionality and the security of such schemes. Their approaches suggest that encapsulation techniques may in fact be a general tool that is useful in the design of functional encryption schemes. As discussed in Section 1.2, our construction relies on encapsulation techniques as a key ingredient, significantly extending the initial ideas of Ananth et al. and Brakerski et al. from encapsulating keys to realizing a re-encryption mechanism that generates a hierarchical structure.

1.5 Paper Organization

The remainder of this paper is organized as follows. In Section 2 we provide an overview of the notation, definitions, and tools underlying our constructions. In Section 3 we present the notion of a hierarchical functional encryption scheme and define its security. In Section 4 we present our generic construction of a hierarchical functional encryption scheme.

The details of the results described in Section 1.3 are deferred from this extended abstract and can be found in [25].

2 Preliminaries

In this section we present the notation and basic definitions that are used in this work. For a distribution X we denote by $x \leftarrow X$ the process of sampling a value x from the distribution X . Similarly, for a set \mathcal{X} we denote by $x \leftarrow \mathcal{X}$ the process of sampling a value x from the uniform distribution over \mathcal{X} . For an integer $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, \dots, n\}$. Throughout the paper, we denote by λ the security parameter. A function $\text{neg} : \mathbb{N} \rightarrow \mathbb{R}$ is *negligible* if for every constant $c > 0$ there exists an integer N_c such that $\text{neg}(\lambda) < \lambda^{-c}$ for all $\lambda > N_c$. Two sequences of random variables $X = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $Y = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are *computationally indistinguishable* if for any probabilistic polynomial-time algorithm \mathcal{A} there exists a negligible function $\text{neg}(\cdot)$ such that $|\Pr[\mathcal{A}(1^\lambda, X_\lambda) = 1] - \Pr[\mathcal{A}(1^\lambda, Y_\lambda) = 1]| \leq \text{neg}(\lambda)$ for all sufficiently large $\lambda \in \mathbb{N}$.

2.1 Pseudorandom Functions

Let $\{\mathcal{K}_\lambda, \mathcal{X}_\lambda, \mathcal{Y}_\lambda\}_{\lambda \in \mathbb{N}}$ be a sequence of sets and let $\mathcal{PRF} = (\text{PRF.Gen}, \text{PRF.Eval})$ be a function family with the following syntax:

- PRF.Gen is a probabilistic polynomial-time algorithm that takes as input the unary representation of the security parameter λ , and outputs a key $K \in \mathcal{K}_\lambda$.
- PRF.Eval is a deterministic polynomial-time algorithm that takes as input a key $K \in \mathcal{K}_\lambda$ and a value $x \in \mathcal{X}_\lambda$, and outputs a value $y \in \mathcal{Y}_\lambda$.

8:10 Hierarchical Functional Encryption

The sets \mathcal{K}_λ , \mathcal{X}_λ , and \mathcal{Y}_λ are referred to as the *key space*, *domain*, and *range* of the function family, respectively. For $K \in \mathcal{K}_\lambda$, we use the notation $\text{PRF.Eval}(K, \cdot)$, $\text{PRF.Eval}_K(\cdot)$ and $\text{PRF}_K(\cdot)$ interchangeably.

The following is the standard definition of a pseudorandom function family.

► **Definition 2.1** (Pseudorandomness). A function family $\mathcal{PRF} = (\text{PRF.Gen}, \text{PRF.Eval})$ is *pseudorandom* if for every probabilistic polynomial-time algorithm \mathcal{A} there exists a negligible function $\text{neg}(\cdot)$ such that

$$\text{Adv}_{\mathcal{PRF}, \mathcal{A}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr_{K \leftarrow \text{PRF.Gen}(1^\lambda)} [\mathcal{A}^{\text{PRF.Eval}_K(\cdot)}(1^\lambda) = 1] - \Pr_{f \leftarrow F_\lambda} [\mathcal{A}^{f(\cdot)}(1^\lambda) = 1] \right| \leq \text{neg}(\lambda),$$

for all sufficiently large $\lambda \in \mathbb{N}$, where F_λ is the set of all functions that map \mathcal{X}_λ into \mathcal{Y}_λ .

In addition to the standard notion of a pseudorandom function family, we rely on the seemingly stronger (yet existentially equivalent) notion of a *puncturable* pseudorandom function family [41, 17, 50, 20]. In terms of syntax, this notion asks for an additional probabilistic polynomial-time algorithm, PRF.Punc , that takes as input a key $K \in \mathcal{K}_\lambda$ and a set $S \subseteq \mathcal{X}_\lambda$ and outputs a “punctured” key K_S . The properties required by such a puncturing algorithm are captured by the following definition.

► **Definition 2.2** (Puncturable PRF). A pseudorandom function family $\mathcal{PRF} = (\text{PRF.Gen}, \text{PRF.Eval}, \text{PRF.Punc})$ is *puncturable* if the following properties are satisfied:

1. **Functionality:** For all sufficiently large $\lambda \in \mathbb{N}$, for every set $S \subseteq \mathcal{X}_\lambda$, and for every $x \in \mathcal{X}_\lambda \setminus S$ it holds that

$$\Pr_{\substack{K \leftarrow \text{PRF.Gen}(1^\lambda); \\ K_S \leftarrow \text{PRF.Punc}(K, S)}} [\text{PRF.Eval}_K(x) = \text{PRF.Eval}_{K_S}(x)] = 1.$$

2. **Pseudorandomness at punctured points:** Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be any probabilistic polynomial-time algorithm such that $\mathcal{A}_1(1^\lambda)$ outputs a set $S \subseteq \mathcal{X}_\lambda$, a value $x \in S$, and state information *state*. Then, for any such \mathcal{A} there exists a negligible function $\text{neg}(\cdot)$ such that

$$\text{Adv}_{\mathcal{PRF}, \mathcal{A}}(\lambda) \stackrel{\text{def}}{=} |\Pr[\mathcal{A}_2(K_S, \text{PRF.Eval}_K(x), \text{state}) = 1] - \Pr[\mathcal{A}_2(K_S, y, \text{state}) = 1]| \leq \text{neg}(\lambda)$$

for all sufficiently large $\lambda \in \mathbb{N}$, where $K \leftarrow \text{PRF.Gen}(1^\lambda)$, $(S, x, \text{state}) \leftarrow \mathcal{A}_1(1^\lambda)$, $K_S = \text{PRF.Punc}(K, S)$, and $y \leftarrow \mathcal{Y}_\lambda$.

For our constructions we rely on pseudorandom functions that need to be punctured only at a single point (i.e., in both parts of Definition 2.2 it holds that $S = \{x^*\}$ for some $x^* \in \mathcal{X}_\lambda$). As observed by [41, 17, 50, 20] the GGM construction [36] of PRFs from any one-way function can be easily altered to yield such a puncturable pseudorandom function family.

Augmented evaluation. When dealing with pseudorandom functions that need to be punctured only at a single point, we find it natural to consider an “augmented” evaluation algorithm that outputs a pre-determined value y^* at the punctured point. That is, we extend the functionality of PRF.Eval such that given an augmented key of the form $(K_{x^*}, (x^*, y^*))$, it holds that

$$\text{PRF.Eval}_{(K_{x^*}, (x^*, y^*))}(x) = \begin{cases} y^*, & \text{if } x = x^* \\ \text{PRF.Eval}_{K_{x^*}}(x), & \text{if } x \neq x^* \end{cases}.$$

2.2 Private-Key Encryption with Pseudorandom Ciphertexts

A private-key encryption scheme over a message space $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$ is a triplet $\Pi = (\text{KG}, \text{Enc}, \text{Dec})$ of probabilistic polynomial-time algorithms. The key-generation algorithm KG takes as input the unary representation 1^λ of the security parameter $\lambda \in \mathbb{N}$ and outputs a secret key k . The encryption algorithm Enc takes as input a secret key k and a message $x \in \mathcal{X}_\lambda$, and outputs a ciphertext c . The decryption algorithm Dec takes as input a secret key k and a ciphertext c , and outputs a message $x \in \mathcal{X}_\lambda$ or the dedicated symbol \perp . In terms of correctness we require that for any key k that is produced by $\text{KG}(1^\lambda)$ and for every message $x \in \mathcal{X}_\lambda$ it holds that $\text{Dec}(k, \text{Enc}(k, x)) = x$ with probability 1 over the internal randomness of the algorithms Enc and Dec . We also require that a uniformly distributed string does not decrypt to a valid message with overwhelming probability, i.e. $\text{Dec}(k, \rho) = \perp$ with probability $(1 - \text{neg}(\lambda))$ over the randomness of the key k and a uniformly distributed string ρ of the same length as the ciphertext⁴. In terms of security, we rely on the following standard notion of pseudorandom ciphertexts which can be based on any one-way function (see, for example, [35]).

► **Definition 2.3** (Pseudorandom ciphertexts). A private-key encryption scheme $\Pi = (\text{KG}, \text{Enc}, \text{Dec})$ has *pseudorandom ciphertexts* if for any probabilistic polynomial-time adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function $\text{neg}(\cdot)$ such that

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{PC}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr \left[\text{Exp}_{\Pi, \mathcal{A}}^{\text{PC}}(\lambda) = 1 \right] - \frac{1}{2} \right| \leq \text{neg}(\lambda),$$

for all sufficiently large $\lambda \in \mathbb{N}$, where the random variable $\text{Exp}_{\Pi, \mathcal{A}}^{\text{PC}}(\lambda)$ is defined via the following experiment:

1. $k \leftarrow \text{KG}(1^\lambda)$, $b \leftarrow \{0, 1\}$.
2. $(x^*, \text{state}) \leftarrow \mathcal{A}_1^{\text{Enc}(k, \cdot)}(1^\lambda)$, where $x^* \in \mathcal{X}_\lambda$.
3. $c_0^* \leftarrow \text{Enc}(k, x^*)$, $c_1^* \leftarrow \{0, 1\}^{|c_0^*|}$.
4. $b' \leftarrow \mathcal{A}_2^{\text{Enc}(k, \cdot)}(c_b^*, \text{state})$.
5. If $b' = b$ then output 1, and otherwise output 0.

2.3 Public-Key Functional Encryption

A public-key functional encryption scheme over a message space $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$ and a function space $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ is a quadruple $\Pi = (\text{Setup}, \text{KG}, \text{Enc}, \text{Dec})$ of probabilistic polynomial-time algorithms. The setup algorithm Setup takes as input the unary representation 1^λ of the security parameter $\lambda \in \mathbb{N}$ and outputs a master-secret key msk and a master-public key mpk . The key-generation algorithm KG takes as input a master-secret key msk and a function $f \in \mathcal{F}_\lambda$, and outputs a functional key sk_f . The encryption algorithm Enc takes as input a master-public key mpk and a message $x \in \mathcal{X}_\lambda$, and outputs a ciphertext ct . In terms of correctness we require that for all sufficiently large $\lambda \in \mathbb{N}$, for every function $f \in \mathcal{F}_\lambda$ and message $x \in \mathcal{X}_\lambda$ it holds that $\text{Dec}(\text{KG}(\text{msk}, f), \text{Enc}(\text{mpk}, x)) = f(x)$ with all but a negligible probability over the internal randomness of the algorithms Setup , KG , and Enc .

We rely on the standard indistinguishability-based notion of adaptive security for public-key functional encryption (see, for example, [15, 48, 9, 5]), asking that encryptions of any

⁴ More accurately, since the ciphertext length may (in general) not be a fixed function of the security parameter, the uniform string ρ is sampled as follows: Given the key k , encrypt a fixed message (say, the message 0) to obtain a ciphertext c , and then uniformly sample $\rho \leftarrow \{0, 1\}^{|c|}$, where $|c|$ denotes the bit-length of c .

two messages, x_0 and x_1 , are computationally indistinguishable given access to functional keys for any function f such that $f(x_0) = f(x_1)$.

► **Definition 2.4** (Adaptive security). A functional encryption scheme $\Pi = (\text{Setup}, \text{KG}, \text{Enc}, \text{Dec})$ over a message space $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$ and a function space $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ is *adaptively secure* if for any probabilistic polynomial-time adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ there exists a negligible function $\text{neg}(\cdot)$ such that

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{FE}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr \left[\text{Exp}_{\Pi, \mathcal{A}}^{\text{FE}}(\lambda) = 1 \right] - \frac{1}{2} \right| \leq \text{neg}(\lambda),$$

for all sufficiently large $\lambda \in \mathbb{N}$, where the random variable $\text{Exp}_{\Pi, \mathcal{A}}^{\text{FE}}(\lambda)$ is defined via the following experiment:

1. $(\text{msk}, \text{mpk}) \leftarrow \text{Setup}(1^\lambda)$, $b \leftarrow \{0, 1\}$.
2. $(x_0^*, x_1^*, \text{state}) \leftarrow \mathcal{A}_1^{\text{KG}(\text{msk}, \cdot)}(1^\lambda, \text{mpk})$, where $x_0^*, x_1^* \in \mathcal{X}_\lambda$, and for each function $f \in \mathcal{F}_\lambda$ with which \mathcal{A}_1 queries $\text{KG}(\text{msk}, \cdot)$ it holds that $f(x_0^*) = f(x_1^*)$.
3. $\text{ct}^* \leftarrow \text{Enc}(\text{mpk}, x_b^*)$.
4. $b' \leftarrow \mathcal{A}_2^{\text{KG}(\text{msk}, \cdot)}(\text{ct}^*, \text{state})$, where for each function $f \in \mathcal{F}_\lambda$ with which \mathcal{A}_2 queries $\text{KG}(\text{msk}, \cdot)$ it holds that $f(x_0^*) = f(x_1^*)$.
5. If $b' = b$ then output 1, and otherwise output 0.

In addition to the above notion of adaptive security we consider two natural relaxations: semi-adaptive security, and selective security. Semi-adaptive security is defined via an experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{semiFE}}(\lambda)$ that is obtained from the experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{FE}}(\lambda)$ by asking the adversary to determine the challenge messages before making any key-generation queries (but *after* receiving the public key). Selective security is defined via an experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{selfFE}}(\lambda)$ that is obtained from the experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{FE}}(\lambda)$ by asking the adversary to determine the challenge messages in advance (i.e., *before* receiving the public key).

Known constructions. General-purpose functional encryption schemes that satisfy the above notion of adaptive security are known to exist based on a variety of assumptions. Ananth et al. [5] gave a generic transformation from selective security to adaptive security, implying that schemes that are adaptively secure for any number of key-generation queries can be based on indistinguishability obfuscation [30, 53], differing-input obfuscation [19, 4], and multilinear maps [32]. In addition, schemes that are adaptively secure for a bounded number $B = B(\lambda)$ of key-generation queries can be based on the Learning with Errors (LWE) assumption (where the length of ciphertexts grows with B and with a bound on the depth of allowed functions) [37], based on any public-key encryption scheme and pseudorandom generators computable by small-depth circuits (where the length of ciphertexts grows with B and with an upper bound on the circuit size of the functions) [38], and even based on any public-key encryption scheme (for $B = 1$) [38].

3 Hierarchical Functional Encryption

In this section we define the notion of a hierarchical functional encryption scheme and formalize several notions of security for such schemes (based on [4]). A hierarchical functional encryption scheme is a functional encryption scheme that supports delegation of functional keys: Given a functional key sk_f corresponding to a function f , and given a function g , it is possible to efficiently compute a functional key $\text{sk}_{g \circ f}$ corresponding to the function $g \circ f$

(i.e., the function that first applies f and then applies g). This capability is provided via a delegation algorithm denote `Delegate`.

Formally, a hierarchical functional encryption scheme over a message space $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$ and a function space $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ is a tuple $\Pi = (\text{Setup}, \text{KG}, \text{Enc}, \text{Dec}, \text{Delegate})$ of probabilistic polynomial-time algorithms, where $(\text{Setup}, \text{KG}, \text{Enc}, \text{Dec})$ is a functional encryption scheme (see Section 2.3), and `Delegate` is a delegation algorithm that operates as follows: It takes as input a functional key sk_f (which had been produced either by the key-generation algorithm or by the delegation algorithm itself) corresponding to a function $f \in \mathcal{F}_\lambda$, and a function $g \in \mathcal{F}_\lambda$, and outputs a functional key $\text{sk}_{g \circ f}$.

Correctness. In terms of correctness we require that for every $\lambda \in \mathbb{N}$, for every polynomial $\ell = \ell(\lambda)$, for every sequence of functions $f_1, \dots, f_\ell \in \mathcal{F}_\lambda$, and for every message $x \in \mathcal{X}_\lambda$, it holds that

$$\text{Dec}(\text{sk}_{f_\ell \circ \dots \circ f_1}, \text{Enc}(\text{mpk}, x)) = (f_\ell \circ \dots \circ f_1)(x)$$

with all but a negligible probability over the internal randomness of the algorithms `Setup`, `KG`, `Enc` and `Delegate`, where $\text{sk}_{f_1} \leftarrow \text{KG}(\text{msk}, f_1)$ and $\text{sk}_{f_{i+1} \circ \dots \circ f_i} \leftarrow \text{Delegate}(\text{sk}_{f_i \circ \dots \circ f_1}, f_{i+1})$ for every $i \in [\ell - 1]$. One can also consider schemes that support ℓ delegation levels for some *fixed* polynomial $\ell = \ell(\lambda)$, although we note that our scheme in this paper supports *any* polynomial number of delegation levels.

Security. As in the work of Ananth et al. [4, Appendix E] we consider the natural extensions of the existing indistinguishability-based definitions of functional encryption [15, 48] to the hierarchical setting. Specifically, we consider adversaries that obtain functional keys for various functions of their choice by issuing key-generation queries and delegation queries. We require that such adversaries have only a negligible advantage in distinguishing the encryptions of two challenge messages, x_0^* and x_1^* , of their choice as long as for any function f for which they obtain a functional key it holds that $f(x_0^*) = f(x_1^*)$.

The experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{HFE}}(\lambda)$. Let $\Pi = (\text{Setup}, \text{KG}, \text{Enc}, \text{Dec}, \text{Delegate})$ be a hierarchical public-key functional encryption scheme over a message space $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$ and a function space $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$, and let \mathcal{A} be a probabilistic polynomial-time adversary. For each $\lambda \in \mathbb{N}$ we denote by $\text{Exp}_{\Pi, \mathcal{A}}^{\text{HFE}}(\lambda)$ the random variable that is defined via the following experiment involving the scheme Π , the adversary \mathcal{A} , and a challenger:

1. **Setup phase:** The challenger samples $(\text{msk}, \text{mpk}) \leftarrow \text{Setup}(1^\lambda)$ and $b \leftarrow \{0, 1\}$.
2. **Pre-challenge phase:** \mathcal{A} on input $(1^\lambda, \text{mpk})$ adaptively issues queries of the form $(f, \text{parent}, \text{mode})$, where $f \in \mathcal{F}_\lambda$, $\text{parent} \in \mathbb{N} \cup \{0\}$ and $\text{mode} \in \{\text{OutputKey}, \text{StoreKey}\}$. The i th query $(f_i, \text{parent}_i, \text{mode}_i)$ is answered by the challenger as follows:
 - a. If $\text{parent} = 0$ then the challenger generates $\text{hsk}_i \leftarrow \text{KG}(\text{msk}, f)$.
 - b. Else, if $\text{hsk}_{\text{parent}_i}$ had already been generated (and is not \perp), then the challenger generates $\text{hsk}_i \leftarrow \text{Delegate}(\text{hsk}_{\text{parent}_i}, f)$. Otherwise set $\text{hsk}_i = \perp$.
 - c. Finally, if $\text{mode}_i = \text{OutputKey}$ then the challenger outputs hsk_i , and if $\text{mode} = \text{StoreKey}$ then the challenger outputs \perp .
3. **Challenge phase:** \mathcal{A} outputs $(x_0^*, x_1^*) \in \mathcal{X}_\lambda \times \mathcal{X}_\lambda$, and then the challenger computes $\text{ct}^* \leftarrow \text{Enc}(\text{mpk}, x_b^*)$ and sends it to \mathcal{A} .
4. **Post-challenge phase:** \mathcal{A} adaptively issues queries as in the pre-challenge phase.
5. **Output phase:** \mathcal{A} outputs b' , and the output of the experiment is 1 if and only if $b' = b$.

Valid adversaries. As standard in functional encryption, we rule out adversaries that can easily distinguish between the two challenge messages, x_0^* and x_1^* , using their queries. Specifically, we say that an adversary is *valid* if for any query $(f_i, \text{parent}_i, \text{mode}_i)$ where $\text{mode}_i = \text{OutputKey}$, it holds that $\tilde{f}_i(x_0^*) = \tilde{f}_i(x_1^*)$, where \tilde{f} is defined recursively by $\tilde{f}_i = f_i \circ \tilde{f}_{\text{parent}_i}$ and $f_0(x) = x$ (if any of these values is not well defined then $\tilde{f}_i(x) \equiv \perp$ for all x).

Having defined the experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{HFE}}(\lambda)$ and the notion of a valid adversary, we are now ready to present our notion of adaptive security for hierarchical functional encryption schemes.

► **Definition 3.1.** A hierarchical functional encryption scheme $\Pi = (\text{Setup}, \text{KG}, \text{Enc}, \text{Dec}, \text{Delegate})$ over a message space $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$ and a function space $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ is *adaptively secure* if for any probabilistic polynomial-time valid adversary \mathcal{A} there exists a negligible function $\text{neg}(\cdot)$ such that

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{HFE}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr \left[\text{Exp}_{\Pi, \mathcal{A}}^{\text{HFE}}(\lambda) = 1 \right] - \frac{1}{2} \right| \leq \text{neg}(\lambda),$$

for all sufficiently large $\lambda \in \mathbb{N}$.

In addition to our notion of adaptive security we consider two natural relaxations: semi-adaptive security, and selective security. Semi-adaptive security is defined via an experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{semiHFE}}(\lambda)$ that is obtained from the experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{HFE}}(\lambda)$ by eliminating the pre-challenge query phase (note that the adversary determines the challenge messages *after* receiving the public key). Selective security is defined via an experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{selHFE}}(\lambda)$ that is obtained from the experiment $\text{Exp}_{\Pi, \mathcal{A}}^{\text{HFE}}(\lambda)$ by asking the adversary to determine the challenge messages in advance (i.e., *before* receiving the public key).

Discussion: The delegation capabilities of functional encryption. It is important to point out that given a functional key sk_f , one cannot hope to delegate anything beyond the set of functions $g \circ f$ while maintaining the security properties of functional encryption. To see this, assume towards contradiction that there exists a function h such that h cannot be expressed as $g \circ f$, but sk_h can be derived from sk_f . Since the value of $h(x)$ cannot be inferred just by examining the value of $f(x)$, there must exist two inputs, x_0 and x_1 such that $f(x_0) = f(x_1)$ but $h(x_0) \neq h(x_1)$. Given sk_f , therefore, one should not be able to distinguish encryptions of x_0 and x_1 , but by delegating to sk_h , this becomes possible, hence the contradiction.

The above optimality claim may seem a little confusing when we think about special cases such as attribute-based encryption (ABE) or even identity-based encryption (IBE). In ABE for example, each ciphertext contains an attribute x and a message m , and sk_f reveals m if and only if $f(x) = 1$. In hierarchical ABE (HABE) [39, 14], given sk_f , one should be able to derive $\text{sk}_{f \wedge f'}$ for all f' . At first glance, this seems to not be covered by our definition since $f \wedge f'$ cannot be expressed as $g \circ f$. However, we notice that in fact when thinking about HABE as a special case of functional encryption, it must be the case that what we call sk_f , is in fact a functional key for the function $f^+(x, m) = ((f(x) = 1)?(x, m) : \perp)$ (i.e., the function that takes (x, m) as input, and if $f(x) = 1$ it returns (x, m) and otherwise it returns \perp). This is because if $f(x) = 1$ then x can always be recovered by considering delegated keys that fix the value of each bit of x to 0 or 1, and check if decryption still works. It is clear from this viewpoint that $(f \wedge f')^+$ can be seen as $g \circ f^+$ for an appropriate g . Therefore, our definition and construction are fully compatible also with the more restricted settings of HABE and HIBE.

4 Our Generic Transformation

In this section we show how to transform any general-purpose public-key functional encryption scheme into a hierarchical one. Our construction relies on the following building blocks:

1. A general-purpose public-key functional encryption scheme $\mathcal{FE} = (\text{FE.Setup}, \text{FE.KG}, \text{FE.Enc}, \text{FE.Dec})$.
2. A private-key encryption scheme $\mathcal{SKE} = (\text{SKE.KG}, \text{SKE.Enc}, \text{SKE.Dec})$.
3. A puncturable pseudorandom function family $\mathcal{PRF} = (\text{PRF.Gen}, \text{PRF.Eval}, \text{PRF.Punc})$.

Our hierarchical scheme $\mathcal{HFE} = (\text{Setup}, \text{KG}, \text{Enc}, \text{Dec}, \text{Delegate})$ is defined as follows.

- **The setup algorithm.** On input the security parameter 1^λ the setup algorithm samples and outputs $(\text{msk}, \text{mpk}) \leftarrow \text{FE.Setup}(1^\lambda)$.
- **The encryption algorithm.** On input the public key mpk and a message x , the encryption algorithm first samples a PRF key $K \leftarrow \text{PRF.Gen}(1^\lambda)$. Then, it computes and outputs $\text{ct} \leftarrow \text{FE.Enc}(\text{mpk}, (x, K, \perp))$. (Note that the message space of the resulting scheme is thus smaller than that of the original scheme.)
- **The key-generation algorithm.** On input the master secret key msk and a function f , the key-generation algorithm first generates a fresh key pair $(\text{msk}', \text{mpk}') \leftarrow \text{FE.Setup}(1^\lambda)$ and uniformly samples a tag $t \leftarrow \{0, 1\}^\lambda$. Then, it uniformly samples $c \leftarrow \{0, 1\}^\ell$ where $\ell = \ell(\lambda)$ is the length of an \mathcal{SKE} encryption of an \mathcal{FE} ciphertext.⁵ Finally, it computes $\text{sk}_f \leftarrow \text{FE.KG}(\text{msk}, \text{ReEnc}_{f,t,\text{mpk}',c})$, where ReEnc is defined in Figure 1, and outputs $\text{hsk}_f = (\text{sk}_f, \text{msk}')$.
- **The delegation algorithm.** On input a (possibly delegated) functional key of the form $\text{hsk}_{f_i \circ \dots \circ f_1} = (\text{sk}_{f_1}, \dots, \text{sk}_{f_i}, \text{msk}')$ for some integer $i \geq 1$, and a function f_{i+1} , the delegation algorithm uses the key-generation algorithm described above to compute $(\text{sk}_{f_{i+1}}, \text{msk}'') \leftarrow \mathcal{HFE.KG}(\text{msk}', f_{i+1})$, and outputs $\text{hsk}_{f_{i+1} \circ \dots \circ f_1} = (\text{sk}_{f_1}, \dots, \text{sk}_{f_i}, \text{sk}_{f_{i+1}}, \text{msk}'')$.
- **The decryption algorithm.** On input a functional key $\text{hsk}_{f_i \circ \dots \circ f_1} = (\text{sk}_{f_1}, \dots, \text{sk}_{f_i}, \text{msk}')$ for some integer $i \geq 1$, and a ciphertext ct , the decryption algorithm first sets $\text{ct}_0 = \text{ct}$ and computes $\text{ct}_j \leftarrow \text{FE.Dec}(\text{sk}_{f_j}, \text{ct}_{j-1})$ for $j = 1, \dots, i$. Then, ct_i is decrypted by using msk' for generating a functional key for the identity function $\text{ID} \in \mathcal{F}$:

$$w \leftarrow \text{FE.Dec}(\text{FE.KG}(\text{msk}', \text{ID}), \text{ct}_i).$$

Finally, w is parsed as a triplet $w = (y, \cdot, \cdot)$, of which the first element y is returned as output.

In what follows we first discuss the correctness of our resulting scheme, then discuss its parameters and overhead, and then state and prove its security based on that of its underlying building blocks.

⁵ To be accurate, ℓ is also a function of the message space of the scheme and of the specific properties of the master secret key. We refrain from mentioning these implicit parameters to avoid cluttering of notation. We note however that this imposes an a-priori bound on the length of the ciphertext and thus also on the message space of our resulting scheme. Lifting this restriction is an interesting research direction.

ReEnc $_{f,t,\text{mpk}',c}(x, K, \mathbf{k})$

1. Compute $\text{ct} \leftarrow \text{SKE.Dec}(\mathbf{k}, c)$, $(s, r) = \text{PRF.Eval}(K, t)$, and $K' = \text{PRF.Gen}(1^\lambda; s)$.
2. If $\text{ct} \neq \perp$ then output ct , and otherwise output $\text{FE.Enc}(\text{mpk}', (f(x), K', \perp); r)$.

■ 1 The function $\text{ReEnc}_{f,t,\text{mpk}',c}$.

Correctness. The correctness of our scheme follows easily by induction on the delegation depth i . Let $(\text{msk}, \text{mpk}) \leftarrow \text{Setup}(1^\lambda)$, and fix a message $x \in \mathcal{X}_\lambda$ and a sequence of functions $f_1, \dots, f_i \in \mathcal{F}_\lambda$.

For $i = 1$ the correctness of decrypting a ciphertext $\text{ct}_0 \leftarrow \text{FE.Enc}(\text{mpk}, (x, K_0, \perp))$ using a key $\text{hsk}_{f_1} = (\text{sk}_{f_1}, \text{msk}_1) \leftarrow \text{KG}(\text{msk}, f_1)$ follows from that of the underlying scheme \mathcal{FE} . Specifically, the decryption algorithm first computes $\text{ct}_1 \leftarrow \text{FE.Dec}(\text{sk}_{f_1}, \text{ct}_0)$, and by the correctness of \mathcal{FE} with an overwhelming probability it holds that $\text{ct}_1 = \text{FE.Enc}(\text{mpk}_1, (f_1(x), K_1, \perp); r_1)$, where $(s_1, r_1) = \text{PRF.Eval}(K_0, t)$ for some t chosen during the key generation, $K_1 = \text{PRF.Gen}(1^\lambda; s_1)$, and mpk_1 is a master public key that is sampled together with msk_1 . Next, the decryption algorithm decrypts ct_1 with msk_1 , and noting that r_1 is pseudorandom given the triplet $(\text{mpk}_1, f_1(x), K_1)$ we can once again rely on the correctness of the underlying scheme \mathcal{FE} and argue that the decryption algorithm outputs $f_1(x)$ with an overwhelming probability.⁶

Assume that the scheme is correct for up to $i - 1$ levels of delegation, and consider decrypting a ciphertext $\text{ct}_0 \leftarrow \text{FE.Enc}(\text{mpk}, (x, K_0, \perp))$ using a key $\text{hsk}_{f_i \circ \dots \circ f_1} = (\text{sk}_{f_1}, \dots, \text{sk}_{f_i}, \text{msk}_i) \leftarrow \text{Delegate}(\text{hsk}_{f_{i-1}}, f_i)$ that is generated using i levels of delegation. Then, the correctness for up to $i - 1$ levels guarantees that by repeatedly applying the keys $\text{sk}_{f_1}, \dots, \text{sk}_{f_{i-1}}$ starting with the initial ciphertext ct_0 as prescribed by the decryption algorithm, we obtain with an overwhelming probability a ciphertext $\text{ct}_{i-1} = \text{FE.Enc}(\text{mpk}_{i-1}, ((f_{i-1} \circ \dots \circ f_1)(x), K_{i-1}, \perp); r_{i-1})$ for some mpk_{i-1} , K_{i-1} and r_{i-1} , where r_{i-1} is pseudorandom given the triplet $(\text{mpk}_{i-1}, (f_{i-1} \circ \dots \circ f_1)(x), K_{i-1})$. Next, the decryption algorithm computes $\text{ct}_i \leftarrow \text{FE.Dec}(\text{sk}_{f_i}, \text{ct}_{i-1})$, and by the correctness of \mathcal{FE} with an overwhelming probability it holds that $\text{ct}_i = \text{FE.Enc}(\text{mpk}_i, ((f_i \circ \dots \circ f_1)(x), K_i, \perp); r')$, where $(s_i, r_i) = \text{PRF.Eval}(K_{i-1}, t)$ for some t chosen during the key generation, $K_i = \text{PRF.Gen}(1^\lambda; s_{i-1})$, and mpk_i is a master public key that is sampled together with msk_i . Note that again r_i is pseudorandom given the triplet $(\text{mpk}_i, (f_i \circ \dots \circ f_1)(x), K_i)$. Therefore, when the decryption algorithm decrypts ct_i with msk_i , it outputs $(f_i \circ \dots \circ f_1)(x)$ with an overwhelming probability.

Parameters and overhead. We now discuss the parameters that govern the properties that are required of the underlying scheme and thus the overhead of our construction. We address two parameters of the hierarchy: The *width* which is the maximal number of delegated keys that are derived from each key at the previous level, and the *depth* which is the maximal number of successive derivations.⁷ The functionality and security of our scheme hold for arbitrary and a-priori unbounded width and depth. However, if the underlying scheme is

⁶ If we further assume that either the underlying functional encryption scheme is perfectly correct, or that the underlying pseudorandom function produces outputs whose marginal distribution is uniform, the argument significantly simplifies and there is no need to argue that r_1 is pseudorandom given the triplet $(\text{mpk}_1, f_1(x), K_1)$.

⁷ One could consider a more fine-grained view of the parameters, e.g. that the maximal width itself depends on the depth of the key. Such analyses follow the same principles presented here.

restricted in some way, then this restriction could propagate through our reduction. For example, if the underlying scheme only supports bounded collusion, then the maximal width will be restricted. Furthermore, since the `ReEnc` function produces a functional ciphertext with respect to the next level of the hierarchy, certain instantiations could produce a cascading effect that will increase the overhead. We analyze these restrictions below and show that in some cases they can be overcome completely and in others they can be managed.

Define the *compactness parameter* of a (standard) FE scheme, denoted $C(\lambda, S)$, as the computational complexity of encrypting a message of length λ (or some other fixed length which does not depend on S), while allowing to produce functional keys for size S functions. Note that C is also a bound on the length of the ciphertext, and in the currently-known schemes it also governs the complexity of key generation (see Section 2.3 for the currently-known schemes). Then in our construction, the ciphertext encryption complexity at depth i , which we denote by C_i is at most $C_i \leq C(\lambda, C_{i+1} \cdot \text{poly}(\lambda))$. This relation follows immediately from the description of the scheme.

For a scheme which only allows bounded collusion, the compactness is $C(\lambda, S, B)$, where B is the bound on the number of collusions. In this case, the width factors in as well such that for a scheme with width w it holds that $C_i \leq C(\lambda, C_{i+1} \cdot \text{poly}(\lambda), w)$.

In particular, in the known schemes with unbounded collusion [30, 53, 32], the encryption complexity is independent of S and therefore instantiating our construction with such a scheme will support arbitrary polynomial depth and width while keeping the encryption complexity polynomial. In fact, one can show, via a little modification of [5], that any scheme that supports unbounded collusions can be modified using randomized encodings to one where the compactness is independent of S .

For known schemes with bounded collusion, such as those based on public-key encryption [38] or on LWE [37], the compactness is $C(\lambda, S, B) \leq \text{poly}(\lambda) \cdot S \cdot B$, which implies that C_i is bounded by $C_{i+1} \cdot \text{poly}(\lambda) \cdot w$. If we intend to support a total depth d , then unfolding the reduction, the bound we have is $C_0 \leq w^d \cdot \lambda^{O(d)}$. This means that if we wish to keep the encryption complexity polynomial in λ , we can only allow $d = O(1)$ and $w = \text{poly}(\lambda)$. Furthermore, we must know w ahead of time in order to instantiate the parameters of the scheme.

Security. The following theorem captures the security of our resulting scheme. We note that the assumptions stated in the theorem are all known to be implied by the existence of any (selectively-secure) general-purpose public-key functional encryption scheme (see Section 2 for formal descriptions of our building blocks and their known instantiations).

► **Theorem 4.1.** *Assuming that (1) \mathcal{FE} is semi-adaptively (resp., selectively) secure (2) \mathcal{SKE} has pseudorandom ciphertexts, and (3) \mathcal{PRF} is a puncturable pseudorandom function family, then \mathcal{HFE} is a semi-adaptively-secure (resp., selectively-secure) hierarchical functional encryption scheme.*

Proof. For ease of exposition we focus here on the case where the underlying scheme \mathcal{FE} is semi-adaptively secure. The proof for the case where \mathcal{FE} is only selectively secure is identical, except for requiring the adversary to provide the challenge messages prior to receiving the public parameters. Let \mathcal{A} be a valid probabilistic polynomial-time adversary (as defined in Section 3). We present a sequence of experiments and upper bound \mathcal{A} 's advantage in distinguishing each two consecutive experiments. The first experiment is the experiment $\text{Exp}_{\mathcal{HFE}, \mathcal{A}}^{\text{semiHFE}}(\lambda)$ and the last experiment is completely independent of the bit b . This enables

us to prove that there exists a negligible function $\text{neg}(\cdot)$ such that

$$\text{Adv}_{\mathcal{HFE}, \mathcal{A}}^{\text{semiHFE}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr \left[\text{Exp}_{\mathcal{HFE}, \mathcal{A}}^{\text{semiHFE}}(\lambda) = 1 \right] - \frac{1}{2} \right| \leq \text{neg}(\lambda)$$

for all sufficiently large $\lambda \in \mathbb{N}$.

How to read this proof. To read our proof, one starts from the first hybrid and proceeds in order to the next, each adjacent hybrid is shown to be computationally indistinguishable from its predecessor. When a loop is encountered, this means that a sequence of hybrids is now being defined, one hybrid for each “iteration” of the loop. The hybrid defined in the first iteration needs to be indistinguishable from the last hybrid before the loop, and all hybrids except the first need to be indistinguishable from the hybrid of the previous iteration. In a *nested loop*, each iteration of the external loop represents a generation of many hybrids, as many as the internal loop generates. In such case, in the first iteration of the external loop, and the first iteration of the internal loop, the hybrid being defined needs to be indistinguishable from the one preceding the loop. However, in the next execution of the external loop, the first iteration of the internal needs to be indistinguishable with the *last* iteration of the internal loop that have been carried out in the previous iteration of the external loop. For example, say that the external loop iterates for $i = 1, \dots, S$ and the internal loop iterates for $j = 1, \dots, T$. Then what we prove for $\mathcal{H}^{(i,j)}$ is that: $\mathcal{H}^{(1,1)}$ is indistinguishable from the last hybrid before the loop, $\mathcal{H}^{(i,1)}$ for $i > 1$ is indistinguishable from $\mathcal{H}^{(i-1,T)}$, and for $i, j > 1$ that $\mathcal{H}^{(i,j)}$ is indistinguishable from $\mathcal{H}^{(i,j-1)}$.

In order to explain the purpose of the different steps in the proof, we also include *invariants* which are properties of the distribution of the current experiment. The invariant holds *only* at that point in the proof where it appears and does not necessarily hold in following hybrids. An invariant inside a loop holds whenever the flow of the proof reaches that point in the loop. Namely, going back to our nested loop example from above, an invariant that appears after the “for $i = 1, \dots, S$ ” statement, holds for the experiment immediately preceding the loop, and for all hybrids $\mathcal{H}^{(i,T)}$, except $\mathcal{H}^{(S,T)}$. An invariant that appears after the “for $j = 1, \dots, T$ ” statement, should hold for the hybrid immediately preceding the loop, as well as for all $\mathcal{H}^{(i,T)}$, except $\mathcal{H}^{(S,T)}$.

We advise the reader to read our proof as if it was an execution of a computer program. We believe that while this proof writing method is still not very widely used, it is quite beneficial in writing complicated proofs, and will find additional uses. In what follows we first describe the notation used throughout the proof, and then describe the experiments.

Notation. Let $Q = Q(\lambda)$ denote a polynomial upper bound on the number of queries that are made by \mathcal{A} in the experiment $\text{Exp}_{\mathcal{HFE}, \mathcal{A}}^{\text{semiHFE}}(\lambda)$. We denote these queries by $\{(f_i, \text{parent}_i, \text{mode}_i)\}_{i \in [Q]}$ and we also consider an implicit “zeroth” query which generates the master key pair (msk, mpk) of the scheme. This allows us to define the *depth* of the i th query, denoted $d(i)$, s.t. $d(0) = 0$ and $d(i) = d(\text{p}(i)) + 1$ for $i > 0$, where we use $\text{p}(i)$ as shorthand for parent_i . Thus we can view \mathcal{A} ’s queries as a tree rooted by the zeroth query, where each query $(f_i, \text{parent}_i, \text{mode}_i)$ is the child of the query $\text{p}(i)$ and has depth $d(i)$ in the tree.

For any query $i \in \{0, \dots, Q\}$, we define a function \tilde{f}_i as follows: \tilde{f}_0 is the identity function, and for all $i > 0$ we define $\tilde{f}_i = f_i \circ \tilde{f}_{\text{p}(i)}$. In other words, the i th query $(f_i, \text{parent}_i, \text{mode}_i)$ generates a delegated key that allows to compute the function $\tilde{f}_i(x)$ given an encryption of x . We say that the i th query is *observable* if $\tilde{f}_i(x_0^*) = \tilde{f}_i(x_1^*)$, and *unobservable* otherwise. We note that if the i th query is unobservable then necessarily $\text{mode}_i = \text{StoreKey}$.

We let $(\text{msk}_i, \text{mpk}_i)$ denote the key pair generated by the challenger while answering the i th query, and let $(\text{msk}_0, \text{mpk}_0)$ be the master key pair (msk, mpk) that is generated by the setup algorithm. Similarly, we let t_i denote the tag that is sampled while answering the i th query.

We denote by x_0^* and x_1^* the challenge messages that are chosen by \mathcal{A} , and by K^* the PRF key that is used for computing the challenge ciphertext. We further define $K_0^* = K^*$, and for all $i > 0$ we define K_i^* , s_i^* , and r_i^* as follows: $(s_i^*, r_i^*) = \text{PRF.Eval}(K_{\text{p}(i)}^*, t_i)$, and $K_i^* = \text{PRF.Gen}(1^\lambda; s_i^*)$. Note that these are exactly the values that are computed by the ReEnc function produced in the i th query, when evaluated on the challenge ciphertext.

Finally, throughout the proof we find it convenient to denote by $\$$ a fresh value that is sampled uniformly and independently of all other existing values.

Experiment \mathcal{H}_0 . This is the experiment $\text{Exp}_{\mathcal{HFE}, \mathcal{A}}^{\text{semiHFE}}(\lambda)$ (see Section 3).

Experiment \mathcal{H}_1 . This experiment is obtained from the experiment \mathcal{H}_0 by having the challenger sample in advance the tags and the key pairs that are used for replying to \mathcal{A} 's queries. In fact, we will sample these values in a redundant manner so that we prepare several such triplets for each query, and the choice of which triplet to use is determined by the depth of the query. We thus have the following claim:

Specifically, at the beginning of the experiment, for all $i, d \in [Q]$ the challenger samples $t_{i,d}^{(o)}, t_{i,d}^{(u)} \leftarrow \{0, 1\}^\lambda$ and $(\text{msk}_{i,d}, \text{mpk}_{i,d}) \leftarrow \text{FE.Setup}(1^\lambda)$. Then, the experiment proceeds exactly as in \mathcal{H}_1 , and whenever the challenger needs to sample t_i and $(\text{msk}_i, \text{mpk}_i)$ for replying to the i th query, it will use $t_i = t_{i,d(i)}^{(o)}$ if i is an observable query, and $t_i = t_{i,d(i)}^{(u)}$ otherwise. It will further use $(\text{msk}_i, \text{mpk}_i) = (\text{msk}_{i,d(i)}, \text{mpk}_{i,d(i)})$.

Looking ahead, this experiment allows the challenger to know in advance, for every possible depth, a polynomial superset of the tags and key pairs that will be produced for replying to queries of this depth. The view of the adversary in this experiment is distributed identically to its view in the experiment \mathcal{H}_0 , yielding the following observation:

► **Observation 4.2.** For all $\lambda \in \mathbb{N}$ it holds that

$$\Pr[\mathcal{H}_0(\lambda) = 1] = \Pr[\mathcal{H}_1(\lambda) = 1].$$

Experiment \mathcal{H}_2 . This experiment is obtained from the experiment \mathcal{H}_1 as follows. After the generation of the tags $t_{i,d}^{(o)}$ and $t_{i,d}^{(u)}$, and before interacting with the adversary, the challenger checks if any of the values $t_{i,d}^{(o)}$ or $t_{i,d}^{(u)}$ for some $(i, d) \in [Q]^2$ appears more than once. In such case the output of the experiment is defined as \perp , and otherwise the experiment is identical to the experiment \mathcal{H}_1 . A standard union bound implies that the experiments \mathcal{H}_1 and \mathcal{H}_2 differ with probability at most $2(Q+1)^4 \cdot 2^{-\lambda} = \text{neg}(\lambda)$, yielding the following observation:

► **Observation 4.3.** For all $\lambda \in \mathbb{N}$ it holds that

$$|\Pr[\mathcal{H}_1(\lambda) = 1] - \Pr[\mathcal{H}_2(\lambda) = 1]| \leq \frac{2(Q+1)^4}{2^\lambda}.$$

Experiment \mathcal{H}_3 . This experiment is obtained from the experiment \mathcal{H}_2 by sampling a sequence $k_0, \dots, k_{Q-1} \leftarrow \text{SKE.KG}(1^\lambda)$ of symmetric keys (one for each possible depth – recall that Q is always an upper bound on the maximal depth), and modifying the symmetric ciphertext c that is generated by the key-generation algorithm when replying to each query as

follows: When replying to the i th query $(f_i, \text{parent}_i, \text{mode}_i)$, instead of sampling c uniformly, the key-generation algorithm computes

$$c_i = \text{SKE.Enc}(k_{d(i)-1}, \text{ct}_i; \$)$$

where $\text{ct}_i = \text{FE.Enc}(\text{mpk}_i, (\tilde{f}_i(x_b^*), K_i^*, \perp); r_i^*)$ (recall that throughout the proof we find it convenient to denote by $\$$ a fresh value that is sampled uniformly and independently of all other existing values).

Note that ct_i is exactly the same as the “ ct_i ” value that is computed in the process of decrypting the challenge ciphertext using the i th functional key (and is also computed as an intermediate value when decrypting the challenge ciphertext with any descendant of the i th key). See the decryption algorithm above.

It thus makes sense to extend our notation and denote the challenge ciphertext by ct_0 (as in the decryption algorithm). Note that while ct_0 is encrypted with true randomness and includes a properly generated PRF key, all other ct_i 's are encrypted using pseudorandomness and contain PRF keys that were generated pseudorandomly. We further say that ct_i is observable if the i th query is an observable query and unobservable otherwise.

To see why the adversary's view in \mathcal{H}_3 is indistinguishable from \mathcal{H}_2 , we note that in \mathcal{H}_3 , the symmetric keys k_0, \dots, k_{Q-1} are used only for generating the c_i 's. In other words, this experiment can be carried out given only access to an encryption oracle $\text{SKE.Enc}(k_d, \cdot)$ for each $d \in \{0, \dots, Q-1\}$ (instead of explicit access to the actual keys k_0, \dots, k_{Q-1}). This enables us to use the ciphertext pseudorandomness of \mathcal{SKE} to prove computational indistinguishability from \mathcal{H}_2 , yielding the following claim in a rather straightforward manner:

► **Claim 4.4.** *Assuming that \mathcal{SKE} has pseudorandom ciphertexts, there exists a negligible function $\text{neg}(\cdot)$ such that*

$$|\Pr[\mathcal{H}_2(\lambda) = 1] - \Pr[\mathcal{H}_3(\lambda) = 1]| \leq \text{neg}(\lambda)$$

for all sufficiently large $\lambda \in \mathbb{N}$.

For $d = 0, \dots, Q$:

► **Invariant 4.5.** *In the previous experiment, it should hold that all ciphertexts ct_i that correspond to unobservable queries (i.e., queries for which $\tilde{f}_i(x_0^*) \neq \tilde{f}_i(x_1^*)$) such that $d(i) < d$ are of the form $\text{FE.Enc}(\text{mpk}_i, (\perp, K_i^*, k_{d(i)}); \$)$, and all such ciphertext such that $d(i) = d$ are of the form $\text{ct}_i = \text{FE.Enc}(\text{mpk}_i, (\tilde{f}_i(x_b^*), K_i^*, \perp); \$)$. Further, if $d(i) \leq d$ then $K_i^* = \text{PRF.Gen}(1^\lambda; \$)$. More specifically:*

■ *If i is such that $d(i) < d$ and $\tilde{f}_i(x_0^*) \neq \tilde{f}_i(x_1^*)$, then it holds that*

$$\text{ct}_i = \text{FE.Enc}(\text{mpk}_i, (\perp, K_i^*, k_{d(i)}); \$)$$

and $K_i^* = \text{PRF.Gen}(1^\lambda; \$)$.

■ *If i is such that $d(i) = d$ and $\tilde{f}_i(x_0^*) \neq \tilde{f}_i(x_1^*)$, then it holds that*

$$\text{ct}_i = \text{FE.Enc}(\text{mpk}_i, (\tilde{f}_i(x_b^*), K_i^*, \perp); \$)$$

and $K_i^* = \text{PRF.Gen}(1^\lambda; \$)$.

■ *If i is such that $d(i) > d$ or $\tilde{f}_i(x_0^*) = \tilde{f}_i(x_1^*)$, then it holds that*

$$\text{ct}_i = \text{FE.Enc}(\text{mpk}_i, (\tilde{f}_i(x_b^*), K_i^*, \perp); r_i^*)$$

and $K_i^* = \text{PRF.Gen}(1^\lambda; s_i^*)$.

We note that this indeed holds for $d = 0$ in experiment \mathcal{H}_3 .

For $i = 0, \dots, Q$:

Experiment $\mathcal{H}_4^{(i,d)}$. In this experiment, the challenger changes the way ct_i is computed as follows. Before generating ct_i , the challenger checks if both $\mathbf{d}(i) = d$ and ct_i is unobservable ($\tilde{f}_i(x_0^*) \neq \tilde{f}_i(x_1^*)$). If both conditions hold then it sets

$$\text{ct}_i = \text{FE.Enc}(\text{mpk}_i, (\perp, K_i^*, \mathbf{k}_d); \mathfrak{s}) .$$

Otherwise ct_i is computed as in the previous experiment.

To see why the adversary's view in this experiment is indistinguishable from the previous hybrid, we note that for any child of i , i.e., j such that $\mathbf{p}(j) = i$,

$$\text{ReEnc}_{f_j, t_j, \text{mpk}_j, c_j}(\tilde{f}_i(x_b^*), K_i^*, \perp) = \text{ReEnc}_{f_j, t_j, \text{mpk}_j, c_j}(\perp, K_i^*, \mathbf{k}_d) = \text{ct}_j .$$

This is because necessarily $\mathbf{d}(j) = d + 1 > d$ and due to Invariant 4.5. Thus, the security of the $(\text{msk}_{i,d}, \text{mpk}_{i,d})$ key pair guarantees that this hybrid is indistinguishable from the previous one: Since ct_i is unobservable, then necessarily the adversary cannot access $\text{msk}_{i,d}$ directly, but rather only via further delegation (i.e., via functional keys to $\text{ReEnc}_{f_j, t_j, \text{mpk}_j, c_j}$). This yields the following claim in rather straightforward manner:

► **Claim 4.6.** *Assuming that \mathcal{FE} is semi-adaptively secure, there exists a negligible function $\text{neg}(\cdot)$ such that*

$$\left| \Pr[\mathcal{H}_4^{(0,0)}(\lambda) = 1] - \Pr[\mathcal{H}_3(\lambda) = 1] \right| \leq \text{neg}(\lambda)$$

and

$$\left| \Pr[\mathcal{H}_4^{(i,d)}(\lambda) = 1] - \Pr[\mathcal{H}_4^{(i-1,d)}(\lambda) = 1] \right| \leq \text{neg}(\lambda)$$

for all $d \in \{0, \dots, Q\}$ and $i \in \{1, \dots, Q\}$, and for all sufficiently large $\lambda \in \mathbb{N}$.

End For i .

► **Invariant 4.7.** *In the previous experiment, it should hold that all ciphertexts ct_i corresponding to unobservable queries such that $\mathbf{d}(i) \leq d$ are of the form $\text{FE.Enc}(\text{mpk}_i, (\perp, K_i^*, \mathbf{k}_{\mathbf{d}(i)}); \mathfrak{s})$ and further $K_i^* = \text{PRF.Gen}(1^\lambda, \mathfrak{s})$.*

Recall that our goal is to restore Invariant 4.5 for value $(d + 1)$. To this end, we next need to replace r_j^* and s_j^* for all j such that $\mathbf{d}(j) = d + 1$, with random values (rather than values that are generated from $K_{\mathbf{p}(j)}^*$).

For $j = 0, \dots, Q$:

► **Invariant 4.8.** *This is similar to Invariant 4.7, but for all ciphertexts $\text{ct}_{j'}$ corresponding to unobservable queries such that $j' < j$ and $\mathbf{d}(j') = d + 1$ it holds that $r_{j'}^*$ and $s_{j'}^*$ had already been replaced with random.*

For $i = 0, \dots, Q$:

Experiment $\mathcal{H}_5^{(i,j,d)}$. In this hybrid, we again change ct_i as follows. If ct_i is unobservable and $d(i) = d$, then define $K_i^* = \text{PRF.Gen}(1^\lambda, \$)$ (as before), $K_i^\otimes = \text{PRF.Punc}(K_i^*, t_{j,d+1}^{(u)})$, $y_{i,j,d+1} = \text{PRF.Eval}(K_i^*, t_{j,d+1}^{(u)})$. We now set:

$$\text{ct}_i = \text{FE.Enc} \left(\text{mpk}_i, \left(\perp, \left(K_i^\otimes, (t_{j,d+1}^{(u)}, y_{i,j,d+1}) \right) \right), \text{k}_{d(i)} \right); \$ \ .$$

Namely, we replace the PRF key with a punctured key at the point $t_{j,d+1}^{(u)}$, and supply the value at that point⁸. We note that the functionality of $\text{PRF.Eval}((K_i^\otimes, (t_{j,d+1}^{(u)}, y_{i,j,d+1})), \cdot)$ is identical to $\text{PRF.Eval}(K_i^*, \cdot)$. The security of the key pair $(\text{msk}_{i,d}, \text{mpk}_{i,d})$ guarantees the indistinguishability of this hybrid (again relying on ct_i being unobservable and thus $\text{msk}_{i,d}$ is not given to the adversary). This yields the following claim in rather straightforward manner:

► **Claim 4.9.** *Assuming that \mathcal{FE} is semi-adaptively secure, there exists a negligible function $\text{neg}(\cdot)$ such that*

$$\left| \Pr \left[\mathcal{H}_5^{(0,0,d)}(\lambda) = 1 \right] - \Pr \left[\mathcal{H}_4^{Q,d}(\lambda) = 1 \right] \right| \leq \text{neg}(\lambda)$$

for all $d \in \{0, \dots, Q\}$, and

$$\left| \Pr \left[\mathcal{H}_5^{(i,j,d)}(\lambda) = 1 \right] - \Pr \left[\mathcal{H}_5^{(i-1,j,d)}(\lambda) = 1 \right] \right| \leq \text{neg}(\lambda)$$

for all $d, j \in \{0, \dots, Q\}$ and $i \in \{1, \dots, Q\}$, and for all sufficiently large $\lambda \in \mathbb{N}$.

End For i .

► **Invariant 4.10.** *In the current experiment, it holds that the PRF key for all depth- d ciphertexts which are unobservable had been punctured at point $t_{j,d+1}^{(u)}$, namely at the point on which it will be evaluated if indeed ct_j is of level $d+1$.*

For $i = 0, \dots, Q$:

Experiment $\mathcal{H}_6^{(i,j,d)}$. In this hybrid, we again change ct_i in the case where ct_i is unobservable and $d(i) = d$. The change from the previous experiment is only that now $y_{i,j,d+1} \leftarrow \$$, namely sampled randomly. We notice that already in the previous hybrid we never use K^* for unobservable queries, only the respective K^\otimes and y values. Therefore swapping the y value to a completely random will be indistinguishable to the adversary by the punctured PRF property. This yields the following claim in rather straightforward manner:

► **Claim 4.11.** *Assuming that PRF is a puncturable pseudorandom function, there exists a negligible function $\text{neg}(\cdot)$ such that*

$$\left| \Pr \left[\mathcal{H}_6^{(0,j,d)}(\lambda) = 1 \right] - \Pr \left[\mathcal{H}_5^{Q,j,d}(\lambda) = 1 \right] \right| \leq \text{neg}(\lambda)$$

⁸ As discussed in Section 2.1, we find it natural to consider an “augmented” evaluation algorithm that outputs a pre-determined value at the punctured point. That is, the augmented evaluation algorithm is given an augmented key $(K_i^\otimes, (t_{j,d+1}^{(u)}, y_{i,j,d+1}))$, where K_i^\otimes is punctured at $t_{j,d+1}^{(u)}$, and given an input t it outputs $\text{PRF.Eval}_{K_i^\otimes}(t)$ if $t \neq t_{j,d+1}^{(u)}$, and it outputs $y_{i,j,d+1}$ if $t = t_{j,d+1}^{(u)}$.

and

$$\left| \Pr \left[\mathcal{H}_6^{(i,j,d)}(\lambda) = 1 \right] - \Pr \left[\mathcal{H}_6^{(i-1,j,d)}(\lambda) = 1 \right] \right| \leq \text{neg}(\lambda)$$

for all $d, j \in \{0, \dots, Q\}$ and $i \in \{1, \dots, Q\}$, and for all sufficiently large $\lambda \in \mathbb{N}$.

End For i .

► **Invariant 4.12.** *In the current experiment, it holds that the PRF key for all depth- d ciphertexts which are unobservable had been punctured at point $t_{j,d+1}$, and further that the punctured value had been substituted with random.*

Experiment $\mathcal{H}_7^{(j,d)}$. In this hybrid, we (finally) change the way ct_j is generated in the case where ct_j is unobservable and $\text{d}(j) = d + 1$ (if these conditions don't hold then we proceed as in the previous experiment). In particular, we change the way the randomness for ct_j and K_j^* is generated. Note that if ct_j is unobservable then it must be the case that $\text{ct}_{\text{p}(j)}$ is also unobservable (since $\tilde{f}_j = f_j \circ \tilde{f}_{\text{p}(j)}$). In the previous experiment, we had

$$(s_j^*, r_j^*) = \text{PRF.Eval} \left(\left(K_{\text{p}(j)}^{\otimes}, (t_{j,d+1}^{(u)}, y_{\text{p}(j),j,d+1}) \right), t_{j,d+1}^{(u)} \right) .$$

We now define instead $(s'_j, r'_j) = y_{\text{p}(j),j,d+1}$. We set $K_j^* = \text{PRF.Gen}(1^\lambda, s'_j)$ and

$$\text{ct}_j = \text{FE.Enc}(\text{mpk}_j, (\tilde{f}_j(x_b^*), K_j^*, \perp); r'_j) .$$

The view of the adversary here remains exactly the same, since $(s'_j, r'_j) = (s_j^*, r_j^*)$. However, conceptually this means that (s_j^*, r_j^*) are detached from the value that is embedded in $\text{ct}_{\text{p}(i)}$. As we will see in the next experiment, we will remove $y_{i,j,d+1}$ from ct_i , but (s'_j, r'_j) will still be well defined. This yields the following observation:

► **Observation 4.13.** *For all $\lambda \in \mathbb{N}$ it holds that*

$$\Pr \left[\mathcal{H}_7^{(j,d)}(\lambda) = 1 \right] = \Pr \left[\mathcal{H}_6^{Q,j,d}(\lambda) = 1 \right]$$

for all $d, j \in \{0, \dots, Q\}$.

For $i = 0, \dots, Q$:

Experiment $\mathcal{H}_8^{(i,j,d)}$. In this hybrid, we again change ct_i in the case where ct_i is unobservable and $\text{d}(i) = d$. We will now undo the puncturing of the PRF keys.

$$\text{ct}_i = \text{FE.Enc}(\text{mpk}_i, (\perp, K_i^*, k_{\text{d}(i)}); \mathfrak{s}) .$$

Indistinguishability holds since in all positions except $t_{j,d+1}^{(u)}$ the new and old keys, K_i^* and $(K_i^{\otimes}, (t_{j,d+1}, y_{i,j,d+1}))$ are functionally equivalent. Furthermore, the function PRF.Eval is never evaluated at $t_{j,d+1}^{(u)}$ (since if ct_j is unobservable then (r'_j, s'_j) are used instead of (r_j^*, s_j^*)). The functional encryption security of $(\text{msk}_{i,d}, \text{mpk}_{i,d})$ therefore implies indistinguishability. This yields the following claim in rather straightforward manner:

► **Claim 4.14.** *Assuming that \mathcal{FE} is semi-adaptively secure, there exists a negligible function $\text{neg}(\cdot)$ such that*

$$\left| \Pr \left[\mathcal{H}_8^{(0,j,d)}(\lambda) = 1 \right] - \Pr \left[\mathcal{H}_7^{j,d}(\lambda) = 1 \right] \right| \leq \text{neg}(\lambda)$$

for all $d, j \in \{0, \dots, Q\}$, and

$$\left| \Pr \left[\mathcal{H}_8^{(i,j,d)}(\lambda) = 1 \right] - \Pr \left[\mathcal{H}_8^{(i-1,j,d)}(\lambda) = 1 \right] \right| \leq \text{neg}(\lambda)$$

for all $d, j \in \{0, \dots, Q\}$ and $i \in \{1, \dots, Q\}$, and for all sufficiently large $\lambda \in \mathbb{N}$.

End For i .

The proof of the following claim is almost identical to that of Claim 4.9 and is therefore omitted:

► **Claim 4.15.** *Assuming that \mathcal{FE} is semi-adaptively secure, there exists a negligible function $\text{neg}(\cdot)$ such that*

$$\left| \Pr \left[\mathcal{H}_8^{(Q,j,d)}(\lambda) = 1 \right] - \Pr \left[\mathcal{H}_5^{(0,j+1,d)}(\lambda) = 1 \right] \right| \leq \text{neg}(\lambda)$$

for all $d \in \{0, \dots, Q\}$ and $j \in \{0, \dots, Q-1\}$, and for all sufficiently large $\lambda \in \mathbb{N}$.

End For j .

The proof of the following claim is almost identical to that of Claim 4.6 and is therefore omitted:

► **Claim 4.16.** *Assuming that \mathcal{FE} is semi-adaptively secure, there exists a negligible function $\text{neg}(\cdot)$ such that*

$$\left| \Pr \left[\mathcal{H}_8^{(Q,Q,d)}(\lambda) = 1 \right] - \Pr \left[\mathcal{H}_4^{(0,d+1)}(\lambda) = 1 \right] \right| \leq \text{neg}(\lambda)$$

for all $d \in \{0, \dots, Q-1\}$, and for all sufficiently large $\lambda \in \mathbb{N}$.

End For d .

We now notice that the proof is practically finished, since the last hybrid $\mathcal{H}_8^{(Q,Q,Q)}$ is completely independent of the bit b . To see this, note that the only values that depend on b in the experiment are the values $\tilde{f}_i(x_b^*)$ that appear inside the ciphertexts ct_i (in particular inside the challenge ciphertext $\text{ct}^* = \text{ct}_0$). We first point out that the value $\tilde{f}_i(x_b^*)$ is in fact *independent* of b in observable ciphertexts, since by definition $\tilde{f}_i(x_0^*) = \tilde{f}_i(x_1^*)$. As for unobservable ciphertexts, in $\mathcal{H}_8^{(Q,Q,Q)}$ none of them contains $\tilde{f}_i(x_b^*)$ at all, as this value had been replaced by \perp . This yields the following observation:

► **Observation 4.17.** *For all $\lambda \in \mathbb{N}$ it holds that*

$$\Pr \left[\mathcal{H}_8^{(Q,Q,Q)}(\lambda) = 1 \right] = \frac{1}{2}.$$

We presented a sequence of polynomially-many experiments starting with the experiment $\mathcal{H}_0 = \text{Exp}_{\mathcal{HFE}, \mathcal{A}}^{\text{semiHFE}}$ and ending with the experiment $\mathcal{H}_8^{(Q,Q,Q)}$ which is completely independent of the bit b . We showed that the output distributions of each two consecutive experiments are negligibly close, which implies that there exists a negligible function $\text{neg}(\cdot)$ such that

$$\begin{aligned} \text{Adv}_{\mathcal{HFE}, \mathcal{A}}^{\text{semiHFE}}(\lambda) &\stackrel{\text{def}}{=} \left| \Pr \left[\text{Exp}_{\mathcal{HFE}, \mathcal{A}}^{\text{semiHFE}}(\lambda) = 1 \right] - \frac{1}{2} \right| \\ &= \left| \Pr \left[\mathcal{H}_0(\lambda) = 1 \right] - \Pr \left[\mathcal{H}_8^{(Q,Q,Q)}(\lambda) = 1 \right] \right| \\ &\leq \text{neg}(\lambda) \end{aligned}$$

for all sufficiently large $\lambda \in \mathbb{N}$, as required. ◀

References

- 1 Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In *Advances in Cryptology – EUROCRYPT’10*, pages 553–572, 2010.
- 2 Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *Advances in Cryptology – CRYPTO’10*, pages 98–115, 2010.
- 3 Shweta Agrawal, Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption: New perspectives and lower bounds. In *Advances in Cryptology – CRYPTO’13*, pages 500–518, 2013.
- 4 Prabhanjan Ananth, Dan Boneh, Sanjam Garg, Amit Sahai, and Mark Zhandry. Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689, 2013.
- 5 Prabhanjan Ananth, Zvika Brakerski, Gil Segev, and Vinod Vaikuntanathan. From selective to adaptive security in functional encryption. In *Advances in Cryptology – CRYPTO’15*, pages 657–677, 2015.
- 6 Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In *Advances in Cryptology – CRYPTO’15*, pages 308–326, 2015.
- 7 Gilad Asharov and Gil Segev. Limits on the power of indistinguishability obfuscation and functional encryption. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science*, pages 191–209, 2015.
- 8 Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *Journal of the ACM*, 59(2):6, 2012.
- 9 Mihir Bellare and Adam O’Neill. Semantically-secure functional encryption: Possibility results, impossibility results and the quest for a general definition. In *Proceedings of the 12th International Conference on Cryptology and Network Security*, pages 218–234, 2013.
- 10 Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science*, pages 171–190, 2015.
- 11 Dan Boneh and Xavier Boyen. Efficient selective-ID secure identity-based encryption without random oracles. In *Advances in Cryptology – EUROCRYPT’04*, pages 223–238, 2004.
- 12 Dan Boneh, Xavier Boyen, and Eu-Jin Goh. Hierarchical identity based encryption with constant size ciphertext. In *Advances in Cryptology – EUROCRYPT’05*, pages 440–456, 2005.
- 13 Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. *SIAM Journal on Computing*, 32(3):586–615, 2003. Preliminary version in *Advances in Cryptology – CRYPTO’01*, pages 213–229, 2001.
- 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’14*, pages 533–556, 2014.
- 15 Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In *Proceedings of the 8th Theory of Cryptography Conference*, pages 253–273, 2011.
- 16 Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: a new vision for public-key cryptography. *Communications of the ACM*, 55(11):56–64, 2012.
- 17 Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In *Advances in Cryptology – ASIACRYPT’13*, pages 280–300, 2013.
- 18 Xavier Boyen and Brent Waters. Anonymous hierarchical identity-based encryption (without random oracles). In *Advances in Cryptology – CRYPTO’06*, pages 290–307, 2006.

- 19 Elette Boyle, Kai-Min Chung, and Rafael Pass. On extractability obfuscation. In *Proceedings of the 11th Theory of Cryptography Conference*, pages 52–73, 2014.
- 20 Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In *Proceedings of the 17th International Conference on Practice and Theory in Public-Key Cryptography*, pages 501–519, 2014.
- 21 Zvika Brakerski, Ilan Komargodski, and Gil Segev. Multi-input functional encryption in the private-key setting: Stronger security from weaker assumptions. In *Advances in Cryptology – EUROCRYPT’16*, pages 852–880, 2016.
- 22 Zvika Brakerski and Gil Segev. Function-private functional encryption in the private-key setting. In *Proceedings of the 12th Theory of Cryptography Conference*, pages 306–324, 2015.
- 23 Zvika Brakerski and Gil Segev. Hierarchical functional encryption. Cryptology ePrint Archive, Report 2015/1011, 2015.
- 24 David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. *Journal of Cryptology*, 25(4):601–639, 2012.
- 25 Nishanth Chandran, Vipul Goyal, Aayush Jain, and Amit Sahai. Functional encryption: Decentralised and delegatable. Cryptology ePrint Archive, Report 2015/1017, 2015.
- 26 Jie Chen and Hoeteck Wee. Semi-adaptive attribute-based encryption and improved delegation for Boolean formula. In *Proceedings of the 9th International Conference on Security and Cryptography for Networks*, pages 277–297, 2014.
- 27 Clifford Cocks. An identity based encryption scheme based on quadratic residues. In *Proceedings of the 8th IMA International Conference on Cryptography and Coding*, pages 360–363, 2001.
- 28 Ronald Cramer, Ivan Damgård, and Yuval Ishai. Share conversion, pseudorandom secret-sharing and applications to secure computation. In *Proceedings of the 2nd Theory of Cryptography Conference*, pages 342–362, 2005.
- 29 Angelo De Caro, Vincenzo Iovino, Abhishek Jain, Adam O’Neill, Omer Paneth, and Giuseppe Persiano. On the achievability of simulation-based security for functional encryption. In *Advances in Cryptology – CRYPTO’13*, pages 519–535, 2013.
- 30 Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, pages 40–49, 2013.
- 31 Sanjam Garg, Craig Gentry, Shai Halevi, and Daniel Wichs. On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In *Advances in Cryptology – CRYPTO’14*, pages 518–535, 2014.
- 32 Sanjam Garg, Craig Gentry, Shai Halevi, and Mark Zhandry. Functional encryption without obfuscation. In *Proceedings of the 1th Theory of Cryptography Conference*, pages 480–511, 2016.
- 33 Craig Gentry and Shai Halevi. Hierarchical identity based encryption with polynomially many levels. In *Proceedings of the 6th Theory of Cryptography Conference*, pages 437–456, 2009.
- 34 Craig Gentry and Alice Silverberg. Hierarchical ID-based cryptography. In *Advances in Cryptology – ASIACRYPT’02*, pages 548–566, 2002.
- 35 Oded Goldreich. *Foundations of Cryptography – Volume 2: Basic Applications*. Cambridge University Press, 2004.
- 36 Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, 1986.

- 37 Shafi Goldwasser, Yael Kalai, Raluca Ada Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, pages 555–564, 2013.
- 38 Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In *Advances in Cryptology – CRYPTO’12*, pages 162–179, 2012.
- 39 Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, pages 545–554, 2013.
- 40 Jeremy Horwitz and Ben Lynn. Toward hierarchical identity-based encryption. In *Advances in Cryptology – EUROCRYPT’02*, pages 466–481, 2002.
- 41 Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In *Proceedings of the 20th Annual ACM Conference on Computer and Communications Security*, pages 669–684, 2013.
- 42 Ilan Komargodski, Gil Segev, and Eylon Yogev. Functional encryption for randomized functionalities in the private-key setting from minimal assumptions. In *Proceedings of the 12th Theory of Cryptography Conference*, pages 352–377, 2015.
- 43 Allison B. Lewko, Tatsuaki Okamoto, Amit Sahai, Katsuyuki Takashima, and Brent Waters. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In *Advances in Cryptology – EUROCRYPT’10*, pages 62–91, 2010.
- 44 Allison B. Lewko and Brent Waters. New techniques for dual system encryption and fully secure HIBE with short ciphertexts. In *Pocceedings of the 7th Theory of Cryptography Conference*, pages 455–479, 2010.
- 45 Allison B. Lewko and Brent Waters. Unbounded HIBE and attribute-based encryption. In *Advances in Cryptology – EUROCRYPT’11*, pages 547–567, 2011.
- 46 Allison B. Lewko and Brent Waters. Why proving HIBE systems secure is difficult. In *Advances in Cryptology – EUROCRYPT’14*, pages 58–76, 2014.
- 47 Pratyay Mukherjee and Daniel Wichs. Two round multiparty computation via multi-key FHE. In *Advances in Cryptology – EUROCRYPT’16*, pages 735–763, 2016.
- 48 Adam O’Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010.
- 49 Amit Sahai and Brent Waters. Slides on functional encryption. Available at <http://www.cs.utexas.edu/~bwaters/presentations/files/functional.ppt>, 2008.
- 50 Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 475–484, 2014.
- 51 Adi Shamir. Identity-based cryptosystems and signature schemes. In *Advances in Cryptology – CRYPTO’84*, pages 47–53, 1984.
- 52 Brent Waters. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In *Advances in Cryptology – CRYPTO’09*, pages 619–636, 2009.
- 53 Brent Waters. A punctured programming approach to adaptively secure functional encryption. In *Advances in Cryptology – CRYPTO’15*, pages 678–697, 2015.