

Succinct Functional Encryption and Applications: Reusable Garbled Circuits and Beyond

Shafi Goldwasser^{*} Yael Kalai[†] Raluca Ada Popa^{*}
Vinod Vaikuntanathan[✉] Nickolai Zeldovich^{*}

^{*} MIT CSAIL [†] Microsoft Research [✉] University of Toronto

December 31, 2012

Abstract

Functional encryption is a powerful primitive: given an encryption $\text{Enc}(x)$ of a value x and a secret key sk_f corresponding to a circuit f , it enables efficient computation of $f(x)$ without revealing any additional information about x . Constructing functional encryption schemes with succinct ciphertexts that guarantee security for even a single secret key (for a general function f) is an important open problem with far reaching applications, which this paper addresses.

Our main result is a functional encryption scheme for any general function f of depth d , with succinct ciphertexts whose size grows with the depth d rather than the size of the circuit for f . We prove the security of our construction based on the intractability of the learning with error (LWE) problem. More generally, we show how to construct a functional encryption scheme from any public-index predicate encryption scheme and fully homomorphic encryption scheme.

Previously, the only known constructions of functional encryption were either for specific inner product predicates, or for a weak form of functional encryption where the ciphertext size grows with the size of the circuit for f .

We demonstrate the power of this result, by using it to construct a reusable circuit garbling scheme with input and circuit privacy: an open problem that was studied extensively by the cryptographic community during the past 30 years since Yao's introduction of a one-time circuit garbling method in the mid 80's. Our scheme also leads to a new paradigm for general function obfuscation which we call token-based obfuscation. Furthermore, we show applications of our scheme to homomorphic encryption for Turing machines where the evaluation runs in input-specific time rather than worst case time, and to publicly verifiable and secret delegation.

Contents

1	Introduction	3
1.1	Our Results	4
1.1.1	Main Application: Reusable Garbled Circuits	6
1.1.2	Token-Based Obfuscation: a New Way to Circumvent Obfuscation Impossibility Results	6
1.1.3	Computing on Encrypted Data in Input-Specific Time.	7
1.1.4	Publicly Verifiable Delegation with Secrecy.	8
1.2	Technique Outline	8
2	Preliminaries	11
2.1	Notation	11
2.2	Background on Learning With Errors (LWE)	12
2.3	Background on Fully Homomorphic Encryption	12
2.4	Background on Garbled Circuits	13
2.5	Background on Functional Encryption	15
2.5.1	Security of Functional Encryption	15
2.5.2	Public-Index Functional Encryption	16
2.5.3	Two-Outcome Public-Index Functional Encryption	18
3	Our Functional Encryption Scheme	19
3.1	Construction	22
3.2	Proof	24
4	Reusable Garbled Circuits	29
4.1	Construction	30
4.2	Proof	31
4.3	Impossibility of Public-Key Reusable Garbled Circuits	34
5	Token-Based Obfuscation	35
5.1	Definition	35
5.2	Scheme	36
6	Computing on Encrypted Data in Input-Specific Time	37
6.1	Construction	39
6.2	Results	40
6.3	Input-Dependent Output Size	41
A	Detailed Background on Learning With Errors (LWE)	45
B	Construction of Two-Outcome Public-Index Functional Encryption	46
C	Homomorphic Encryption for Turing Machines: Definitions and Proofs	48
C.1	Proof	49

1 Introduction

Breaches of confidential data are commonplace: personal information of millions of people, such as financial, medical, customer, and employee data, is disclosed every year [Pri12, Ver12]. These disclosures often happen because untrustworthy systems handle confidential data. As applications move to cloud computing platforms, ensuring data confidentiality on third-party servers that may be untrustworthy becomes a top concern [Dav12].

A powerful technique for preventing data disclosures without having to ensure the server is trustworthy is to encrypt the data provided to the server and then compute on the encrypted data; thus, if the server does not have access to the plaintext or to the decryption key, it will be unable to disclose confidential data. The challenge in making this approach work lies in coming up with an encryption scheme that can simultaneously offer data privacy, the ability to perform (private) computation on encrypted data, and fine-grained access control to the computation results.

The big leap of the last decade towards addressing this challenge has been the invention of fully homomorphic encryption (FHE) [Gen09, DGHV10, SS10, BV11a, BV11b, Vai11, BGV12, GHS12a, GHS12b, LTV12, Bra12] allowing the remarkable ability of computing on encrypted data without access to the data itself. These exciting constructions open the doors to new applications. At the same time, though, they raise questions as to their ultimate usefulness.

Perhaps the most fundamental question is: *who can decrypt the results of computations on encrypted data?* Although computation on encrypted data using FHE can be performed by anyone (with the knowledge of the public key alone), only the holder of the secret key can decrypt the result of a computation; the secret key, however, allows decryption of the entire data. This model suffices for certain applications: for example, a client can save local computation time by delegating computation to a powerful cloud which returns encrypted computation results. However, it rules out a large class of applications in which the party computing on the encrypted data needs to determine the computation result on its own. For example, spam filters should be able to determine if an encrypted email is spam and discard it, without learning anything else about the email's content. With FHE, the spam filter can run the spam detection algorithm homomorphically on an encrypted email and obtain an encrypted result; but it cannot tell if the algorithm deems the email spam or not. Having the data owner provide the decryption key to the spam filter is not a solution: the spam filter can now decrypt all the emails as well!

A promising approach to this problem is *functional encryption* [SW05, GPSW06, KSW08, LOS⁺10, OT10, O'N10, BSW11]. In functional encryption, anyone can encrypt data with a master public key mpk and the holder of the master secret key can provide keys for functions, for example, sk_f for function f . Anyone with access to a key sk_f and a ciphertext c for x can obtain the result of the computation in plaintext form: $f(x)$. The security of FE requires that the adversary does not learn anything about x , other than the computation result $f(x)$. It is easy to see, for example, how to solve the above spam filter problem with a functional encryption scheme. A user Alice publishes her public key online and gives the spam filter a key for the filtering function. Users sending email to Alice will encrypt the email with her public key. The spam filter can now determine by itself, for each email, whether to pass it along to Alice's mailbox or to deem it as spam, but without ever learning anything about Alice's email (other than the fact that it was deemed a spam message or not).

Indeed, a general functional encryption scheme that supports all functions would be extremely useful for our needs. The recent impossibility result of Agrawal, Gorbunov, Vaikuntanathan and Wee [AGVW12] says that functional encryption schemes that can securely provide an arbitrary number of keys for general functions are impossible; stated differently, any functional encryption scheme that can securely provide q keys for general functions must have ciphertexts growing linearly in q . Since any scheme that can securely

provide a single key yields a scheme that can securely provide q keys by repetition, the question becomes if one can construct a *functional encryption scheme that can securely provide a single key for a general function*.

Up until our work, the known constructions of functional encryption were quite limited: (1) The works of Katz, Sahai and Waters [KSW08], Agrawal, Freeman and Vaikuntanathan [AFV11], and Shen, Shi and Waters [SSW09] show functional encryption schemes (based on different assumptions) for very simple functionalities: the inner product functionality f_y , that on input x outputs 1 if and only if $\langle x, y \rangle = 0$ [KSW08], but do not shed light on how to extend beyond inner products; (2) The more recent work of Gorbunov, Vaikuntanathan and Wee [GVW12a] address general functions but allow a *non-succinct* form of functional encryption where the size of the ciphertexts is at least the size of a universal circuit computing the functions allowed by the scheme.¹ They prove security when the adversary sees the secret key of a *single* function of his choice and then extend this construction to allow the adversary to see secret keys for q functions of his choice, by increasing the size of the ciphertexts linearly with q where q is known in advance.²

This leaves open the problem of constructing a single-key functional encryption scheme for general computations with *ciphertext size independent of the circuit size*. This is the subject of our work. We call such schemes *succinct* functional encryption schemes. We emphasize that the dependence on the circuit size (as in [GVW12a]) is particularly undesirable and it precludes many useful applications of functional encryption (e.g., delegation, reusable garbled circuits, FHE for Turing machines). For example, in the setting of delegation, a data owner wants to delegate her computation to a cloud, but the mere effort of encrypting the data is greater than computing the circuit directly, so the owner is better off doing the computation by herself.

We remark that functional encryption (FE) arises from, and generalizes, a beautiful sequence of papers on attribute-based encryption (including [SW05, GPSW06, BSW07, GJPS08, LOS⁺10, Wat11, Wat12, LW12]), and more generally predicate encryption (including [BW07, KSW08, OT09]). In this paper, we adopt the terminology of Boneh, Sahai and Waters [BSW11] as follows. We denote by a public-index functional encryption scheme, an encryption scheme where each ciphertext c of an underlying plaintext message m is tagged with a public index x . Each secret key sk_f is associated with a predicate f . Given a key sk_f and a ciphertext $c = \text{Enc}(x, m)$, the message m can be recovered if and only if $f(x)$ is true. Whether the message gets recovered or not, the index x is always public, so public-index functional encryption offers qualitatively weaker security than functional encryption. Such schemes were also called public-index predicate encryption schemes in the literature. (The name “attribute-based encryption” has also referred to public-index functional encryption for functions such as boolean formulas [BSW11].)

Very recently, the landscape of public-index functional encryption has significantly improved with the works of Gorbunov, Vaikuntanathan and Wee [GVW12b], and Sahai and Waters [SW12], who construct public-index functional encryption schemes for general functions, and are a building block for our results.

1.1 Our Results

Our main result is the construction of a *succinct* functional encryption scheme for *general functions*. We demonstrate the power of this result by showing that it can be used to address the long standing open problem in cryptography of reusing garbled circuits, as well as making progress on other open problems.

We can state our main result as a reduction from *any* public-index functional encryption (also called public-index predicate encryption scheme) and *any* fully homomorphic encryption scheme. In particular, we show how to construct a (single-key and succinct) functional encryption scheme for any class of functions

¹A universal circuit \mathcal{F} is a circuit that takes as input a description of a circuit f and an input string x , runs f on x and outputs $f(x)$.

²Namely, parameter q (the maximum number of keys allowed) is fixed during setup, and the ciphertexts size grows linearly with q .

\mathcal{F} , given an encryption scheme which can do homomorphic evaluations for any function in \mathcal{F} , and given a *public-index* functional encryption scheme for a “slightly larger” class of functions \mathcal{F}' ; \mathcal{F}' is the class of functions such that for any function $f \in \mathcal{F}$, the class \mathcal{F}' contains the function computing the i -th bit of the FHE evaluation of f .

Theorem 1.1 (Informal). *There is a single-key functional encryption scheme with succinct ciphertexts (independent of circuit size) for the class of functions \mathcal{F} assuming the existence of*

- *a fully homomorphic encryption scheme for the class of functions \mathcal{F} ; and*
- *a (single-key) public-index functional encryption scheme for a class of predicates \mathcal{F}' (as above).*

We consider both selective and full security definitions (see Sec. 2.5), and show that if the underlying public-index functional encryption is selectively (or fully) secure then our resulting functional encryption scheme is selectively (or fully) secure.

The leap from public-index functional encryption to general functional encryption provides a conceptual jump: with public-index, the input that the computation runs on (the index) is not hidden, whereas with general functional encryption, nothing leaks about the input of the computation (other than the computation result).

Two very recent results achieve public-index functional encryption for general functions. Gorbunov, Vaikuntanathan and Wee [GVW12b] achieve public-index FE for general circuits of bounded depth based on the subexponential Learning With Errors (LWE) intractability assumption. Sahai and Waters [SW12] achieve public-index FE for general circuits under the less standard k -Multilinear Decisional Diffie-Hellman (see [SW12] for more details); however, when instantiated with the only construction of multilinear maps currently known [GGH12], they also achieve public-index FE for general circuits of bounded depth.

When coupling our theorem with the public-index FE result of [GVW12a] and the FHE scheme of [BV11b, BGV12], we obtain:

Corollary 1.2 (Informal). *Under the subexponential LWE assumption, for any depth d , there is a single-key functional encryption scheme for general functions computable by circuits of depth d . The scheme has succinct ciphertexts: the ciphertext size is polynomial in the depth d (and does not depend on the circuit size).*

This corollary holds for both selective and full security definitions, since [GVW12a] constructs both selectively secure and fully secure public-index functional encryption schemes. However the parameters of the LWE assumptions are different in the two cases (Sec. 2.3).

Another corollary of our theorem is that given a universal public-index FE scheme (where the scheme is for all classes of circuits, independent of depth) and any fully homomorphic encryption, there is a universal (private-index) functional encryption scheme that is also independent of the depth.

As mentioned, extending our scheme to be secure against an adversary who receives q keys is straightforward. The basic idea is simply to repeat the scheme q times in parallel. This strategy results in the ciphertext size growing linearly with q , but this is unavoidable because of the discussed impossibility result [AGVW12]. Stated in these terms, our scheme is also a q -collusion-resistant functional encryption scheme like [GVW12a], but our scheme’s ciphertexts are succinct, whereas [GVW12a]’s are proportional to the circuit size.

From now on, we restrict our attention to the single-key case, which is the essence of the new scheme. In the body of the paper we often omit the single-key or succinct adjectives and whenever we refer to a functional encryption scheme, we mean a single-key succinct functional encryption scheme.

We next show how to use our main theorems to make significant progress on some of the most intriguing open questions in cryptography today: the *reusability* of garbled circuits, a new paradigm for *general function*

obfuscation, and applications to fully homomorphic encryption with evaluation running in *input-specific time* rather than in worst case time, and to publicly verifiable delegation.

1.1.1 Main Application: Reusable Garbled Circuits

Garbled circuits, which has been one of the most useful primitives in modern cryptography, is a construction originally suggested by Yao in the 80s in the context of secure two-party computation [Yao82]. This construction relies on the existence of a one-way function to encode an arbitrary circuit C (“garbling” the circuit) and then encode any input x to the circuit (where the size of the encoding is short, namely, it does not grow with the size of circuit C); a party given the garbling of C and the encoding of x can run the garbled circuit on the encoded x and obtain $C(x)$. The most basic properties of garbled circuits are circuit and input privacy: an adversary learns nothing about the circuit C or the input x other than the result $C(x)$.

Over the years, garbled circuits and variants thereof have found many applications: two party secure protocols [Yao86], multi-party secure protocols [GMW87], one-time programs [GKR08], KDM-security [BHHI10], verifiable computation [GGP10], homomorphic computations [GHV10] and others. However, a basic limitation of the original construction remains: it offers only *one-time* usage: providing an encoding of more than one input compromises the secrecy of the circuit. Thus, evaluating C on any new input requires an entirely new garbling of the circuit.

The problem of reusing garbled circuits has been open for 30 years. Using our newly constructed succinct functional encryption scheme we are now able to build *reusable garbled circuits* that achieve *circuit and input privacy*: a garbled circuit for any computation of depth d (where the parameters of the scheme depend on d), which can be run on *any polynomial number* of inputs without compromising the privacy of the circuit or the input. More generally, we prove the following:

Theorem 1.3 (Informal). *There exists a polynomial p , such that for any depth function d , there is a reusable garbling scheme for the class of all boolean circuits of depth d , assuming there is a single-key functional encryption scheme for all boolean circuits of depth $p(d)$.*³

Corollary 1.4 (Informal). *Under the subexponential LWE assumption, for any depth function d , there exists a reusable circuit garbling scheme with circuit and input privacy for all the boolean circuits of depth d .*

Reusability of garbled circuits (for depth bounded computations) implies a multitude of applications as evidenced by the research on garbled circuits over the last 30 years. We note that for many of these applications, depth bounded computation suffices. One of the more intriguing applications pertains to a new model for program obfuscation, token-based obfuscation, which we discuss next.

We remark that [GVW12b] gives a restricted form of reusable circuit garbling: it provides authenticity of the circuit output, but does not provide circuit privacy or input privacy as we do here. Informally, authenticity means that an adversary cannot obtain a different yet legitimate result from a garbled circuit. We note that most of the original garbling circuit applications (e.g., two party secure protocols [Yao86], multi-party secure protocols [GMW87]) rely on the privacy of the input or of the circuit.

1.1.2 Token-Based Obfuscation: a New Way to Circumvent Obfuscation Impossibility Results

Program obfuscation is the process of taking a program as input, and producing a functionally equivalent but modified program, so that the modified program reveals no information to a computationally bounded

³For this application we need to assume that the underlying functional encryption scheme is fully secure (as opposed to only selectively secure).

adversary about the original program, beyond what “black box access” to the program reveals. Whereas ad-hoc program obfuscators are built routinely, and are used in practice as the main software-based technique to fight reverse engineering of programs, in 2000 Barak et al. [BGI⁺01], followed by Goldwasser and Kalai [GK05], proved that program obfuscation for general functionalities is impossible using software alone, with respect to several strong but natural definitions of obfuscation.

The results of [BGI⁺01, GK05] mean that there exist functions which cannot be obfuscated. Still, the need to obfuscate or “garble” programs remains. A long array of works attempts to circumvent the impossibility results in various ways, including adding secure hardware components [GKR08, GIS⁺10, BCG⁺11], relaxing the definition of security [GR07], or considering only specific functionalities [Wee05, CKVW10].

The problem of obfuscation seems intimately related to the “garbled circuit” problem where given a garbling of a circuit C and an encoding for an input x , one can learn the result of $C(x)$ but nothing else. One cannot help but wonder whether the new reusable garbling scheme would *immediately* imply a solution for the obfuscation problem (which we know is impossible). Consider an example illustrating this intuition: a vendor obfuscates her program (circuit) by garbling it and then gives the garbled circuit to a customer. In order to run the program on (multiple) inputs x_i , the customer simply encodes the inputs according to the garbling scheme and thus is able to compute $C(x_i)$. Unfortunately, although close, this scenario does not work with reusable garbled circuits. The key observation is that encoding x requires knowledge of a secret key! Thus, an adversary cannot produce encoded inputs on its own, and needs to obtain “tokens” in the form of encrypted inputs from the data owner.

Instead, we propose a new *token-based* model for obfuscation. The idea is for a vendor to obfuscate an arbitrary program as well as provide tokens representing rights to run this program on specific inputs. For example, consider that some researchers want to obtain statistics out of an obfuscated database containing sensitive information (the obfuscated program is the program running queries with the secret database hardcoded in it). Whenever the researchers want to input a query x to this program, they need to obtain a token for x from the program owner. To produce each token, the program owner does little work. The researchers perform the bulk of the computation by themselves using the token and obtain the computation result without further interaction with the owner.

Claim 1.5. *Assuming a reusable garbling scheme for a class of circuits, there is a token-based obfuscation scheme for the same class of circuits.*

Corollary 1.6 (Informal). *Under the subexponential LWE assumption, for any depth function d , there exists a token-based obfuscation scheme for all boolean circuits of depth d .*

It is worthwhile to compare the token-based obfuscation model with previous work addressing obfuscation using trusted-hardware components such as [GIS⁺10, BCG⁺11]. In these schemes, after a user finishes executing the obfuscated program on an input, the user needs to interact with the trusted hardware to obtain the decryption of the result; in comparison, in our scheme, the user only needs to obtain a token before the computation begins, and can then run the computation and obtain the decrypted result by herself.

1.1.3 Computing on Encrypted Data in Input-Specific Time.

All current FHE constructions work according to the following template. For a fixed input size, a program is transformed into an arithmetic (or Boolean) circuit; homomorphic evaluation happens gate by gate on this circuit. The size of the circuit reflects the *worst-case* running time of the program: for example, every loop is unfolded into the maximum number of steps corresponding to the worst-case input, and each function is

called the maximum number of times possible. Such a circuit can be potentially very large, despite the fact that there could be many inputs on which the execution is short.

A fascinating open question has been whether it is possible to perform FHE following a “Turing machine”-like template: *the computation time is input-specific* and can terminate earlier depending on the input at hand. Of course, to compute in input-specific time, the running time must unavoidably leak to the evaluator; therefore such a scheme provides weaker security than homomorphic encryption (namely, nothing other than the running time leaks about the input), at the increase of efficiency.

Using our functional encryption scheme, we show how to achieve this goal. The idea is to use the scheme to test when an encrypted circuit computation has terminated, so the computation can stop earlier on certain inputs. We overview our technique in Sec. 1.2.

Because the ciphertexts in our functional encryption scheme grow with the depth of the circuits, such a scheme is useful only for Turing machines that can be expressed as circuits with small depth d . We refer to such Turing machines as *d-depth-bounded* and define them in Sec. 6.

Theorem 1.7. *There is a scheme for evaluating Turing machines on encrypted data in input-specific time for any class of d-depth-bounded Turing machines, assuming the existence of a single-key functional encryption scheme for circuits of depth d ,⁴ and a fully homomorphic encryption scheme for circuits of depth d .*

Corollary 1.8. *Under the subexponential LWE assumption, for any depth d , there is a scheme for evaluating Turing machines on encrypted data in input-specific time for any class of d-depth-bounded Turing machines.*

1.1.4 Publicly Verifiable Delegation with Secrecy.

Recently, Parno, Raykova and Vaikuntanathan [PRV12] showed how to construct a 2-message delegation scheme (for boolean functions) that is *publicly verifiable*, in the preprocessing model, from any public-index functional encryption scheme. This reduction can be combined with [GVW12a]’s public-index FE scheme to achieve such a delegation scheme.

However, this scheme does not provide *secrecy* of the inputs: the evaluator can learn the inputs. By replacing the public-index FE scheme in the construction of [PRV12] with our new functional encryption scheme, we add secrecy to the scheme; namely, we obtain a delegation scheme which is both *publicly verifiable* and *secret*, in the sense that the prover does not learn anything about the input or output of the function being delegated.⁵

More specifically, our delegation scheme is in the preprocessing model, and is for general depth bounded circuits, where the verifier works in time that depends on the *depth* of the circuit being delegated, but is independent of the size of the circuit, and the prover works in time proportional to the *size* of the circuit being delegated.

1.2 Technique Outline

Our functional encryption scheme. We first describe the ideas behind our main technical result: a reduction from public-index functional encryption (public-index FE) and fully homomorphic encryption (FHE) to (private-index/general) functional encryption (FE).

Compute on encrypted data with FHE. A natural starting point is FHE because it enables computation on encrypted data, which is needed with functional encryption. Using FHE, the FE encoding of an input x

⁴As in previous applications, we need to assume that the underlying functional encryption scheme is fully secure (as opposed to only selectively secure).

⁵We note that secrecy can be easily obtained by using an FHE scheme, however, this destroys public-verifiability.

consists of an FHE encryption of x , denoted \hat{x} , and the secret key for a function f is simply f itself. The semantic security of FHE provides the desired security (and more) because nothing leaks about x ; however, using FHE evaluation, the evaluator obtains an encrypted computation result, $\widehat{f(x)}$, instead of the decrypted value $f(x)$. An idea is to give the FHE decryption key to the evaluator, but this is obviously wrong because the evaluator can use it to decrypt x as well.

Attempt to decrypt using a Yao garbled circuit. We would like the evaluator to decrypt the FHE ciphertext $\widehat{f(x)}$ and only that ciphertext. An idea is for the owner to give the evaluator a Yao garbled circuit for the FHE decryption function with the FHE secret key hardcoded in it, $\text{HE.Dec}_{\text{hsk}}$. The owner also needs to give the evaluator the input labels corresponding to $\widehat{f(x)}$: namely, if garbling $\text{HE.Dec}_{\text{hsk}}$ produces labels $\{L_0^i, L_1^i\}_i$, the owner needs to give the evaluator the labels $\{L_{b_i}^i\}_i$ where b_i is the i -th bit of $\widehat{f(x)}$. But this is not possible because the owner does not know a priori $\widehat{f(x)}$ which is only determined after the FHE evaluation; furthermore, after providing more than one set of labels (which happens when encrypting a different input x'), the security of the garbled circuit (and hence of the FHE secret key) is compromised. One idea is to have the owner and the evaluator interact, but the syntax of functional encryption does not allow interaction. Therefore, the evaluator needs to determine the set of labels corresponding to $\widehat{f(x)}$ in a *non-interactive way*, and *should not obtain any other labels*.

Constraining decryption using public-index FE. It turns out that what we need here is very close to what public-index FE provides. Consider the following variant of public-index FE (called pFE_2) that can be constructed easily from a standard public-index FE scheme. One encrypts a value y together with two messages m_0, m_1 and obtains a ciphertext $c \leftarrow \text{pFE}_2.\text{Enc}(y, m_0, m_1)$. Then, one generates a key for a predicate g : $\text{sk}_g \leftarrow \text{pFE}_2.\text{KeyGen}(g)$. The decryption algorithm on input c and sk_g outputs m_0 if $g(y) = 0$ or outputs m_1 if $g(y) = 1$.

Now consider using pFE_2 multiple times, once for every $i \in \{1, \dots, \text{size of } \widehat{f(x)}\}$. For the i -th invocation of $\text{pFE}_2.\text{Enc}$, let m_0, m_1 be the garbled labels L_0^i, L_1^i , and let y be \hat{x} : $\text{pFE}_2.\text{Enc}(\hat{x}, L_0^i, L_1^i)$. Next, for the i -th invocation of $\text{pFE}_2.\text{KeyGen}$, let g be HE.Eval_f^i (the predicate returning the i -th bit of the evaluation of f on an input ciphertext): $\text{pFE}_2.\text{KeyGen}(\text{HE.Eval}_f^i)$. Then, the evaluator can use $\text{pFE}_2.\text{Dec}$ to obtain the needed label: $L_{b_i}^i$ where b_i is the i -th bit of $\widehat{f(x)}$. Armed with these labels and the garbled circuit, the evaluator decrypts $\widehat{f(x)}$. The security of the public-index FE scheme ensures the evaluator cannot decrypt any other labels, so the evaluator cannot learn more than $f(x)$. Finally, note that the one-time aspect of garbled circuits does not restrict the number of encryptions with our FE scheme because the encryption algorithm generates a new garbled circuit every time; since the garbled circuit is for the FHE decryption algorithm (which is a fixed algorithm), the size of the ciphertexts remains independent of the size of f .

We now explain how to use this result to obtain the aforementioned applications.

From FE to reusable garbled circuits. The goal of garbled circuits is to hide the circuit C . However, our FE scheme only hides the inputs to the circuit and not the circuit itself. To obtain circuit privacy, the idea is to somehow leverage the secrecy of the inputs to hide the circuit. The first idea that comes to mind is to generate a key for the universal circuit instead of C , and include C in the ciphertext when encrypting an input. However, this approach will yield large ciphertexts, as large as the circuit size.

Instead, the insight is to garble C by using our FE scheme, together with a semantically secure encryption scheme E.Enc : The garbling of C will be an FE secret key for the function $\text{E.Enc}_{\text{sk}}(C)$, which on input (sk, x) uses sk to decrypt C and then runs C on the input x . The token for an input x will be an FE encryption of (sk, x) . Now, even if the FE scheme does not hide $\text{E.Enc}_{\text{sk}}(C)$, the security of the encryption scheme E hides C .

Computing on encrypted data in input-specific time. We overview our approach to evaluate a Turing machine (TM) M homomorphically over encrypted data, while not running in worst-case time over all inputs to M .

Our idea is to use our functional encryption scheme to enable the evaluator to determine at various intermediary steps in the evaluation whether the computation finished or not. For each intermediary step, the client provides a secret key for a function that returns a bit indicating whether the computation finished or not. However, if the client provides a key for every computation step, then the amount of keys corresponds to the worst-case running time. Thus, instead, we choose intermediary points spaced at exponentially increasing intervals. In this way, the client only generates a logarithmic number of keys, namely for functions indicating if the computation finishes in $1, 2, 4, \dots, 2^i, \dots, 2^{\lceil \log t_{\max} \rceil}$ steps, where t_{\max} is the worst-case running time of M on all inputs of a certain size.

Because of the single-key requirements posed by the impossibility result of [AGVW12], the client cannot provide keys for an arbitrary number of TMs to the evaluator. However, this does not mean that the evaluator can only run an a priori fixed number of TMs on the encrypted data. The reason is that the client can provide keys for the universal TMs $U_0, \dots, U_{\lceil \log t_{\max} \rceil}$, where TM U_i is the TM that on input a TM M and a value x , runs M on x for 2^i steps and outputs whether M finished.

Therefore, in an offline preprocessing phase, the client provides $1 + \lceil \log t_{\max} \rceil$ keys where the i -th key is for a circuit corresponding to U_i , each key being generated with a different master secret key. The work of the client in this phase is at least t_{\max} which is costly, but this work only happens once and is amortized over all subsequent inputs in the online phase.

In an online phase, the client receives an input x and wants the evaluator to compute $M(x)$ for him. The client provides FE encryptions of (M, x) to the evaluator together with an FHE ciphertext (\hat{M}, \hat{x}) for (M, x) to be used for a separate FHE evaluation. The evaluator tries each key from the preprocessing phase and learns the smallest i for which the computation of M on x stops in 2^i steps. The evaluator then computes a universal circuit of size $\tilde{O}(2^i)$ and evaluates it homomorphically over (\hat{M}, \hat{x}) , obtaining the FHE encryption of $M(x)$. Thus, we can see that if t is the running time of M on x , the evaluator runs in roughly time t .

Publicly Verifiable Delegation with Secrecy. Delegation schemes aim to enable a weak verifier to delegate computation of a function f on an input x to a prover who can then prove to the verifier that he computed the function correctly. We now show that our single-key functional encryption scheme provides an improvement to publicly verifiable delegation by adding secrecy. We only present this improvement informally, because we prefer to focus on the other applications.

We now briefly recall the scheme of [PRV12] and then discuss how to modify it; we refer the reader to Section 2.5 for formal definitions of (public-index) functional encryption (FE). There are two phases in the delegation scheme: the preprocessing phase when the verifier prepares the computation f , and an online phase repeating many times, in which the verifier gives x to the prover who computes $f(x)$ and proves the computation was correct.

In the preprocessing phase, the verifier generates two pairs of master secret and public keys $(\text{msk}_1, \text{mpk}_1)$ and $(\text{msk}_2, \text{mpk}_2)$ for the underlying public-index functional encryption scheme. If f is the function to delegate, the verifier uses msk_1 to generate a key for f denoted sk_f , and msk_2 to generate a key for the negation of f , $\bar{f}(x) := 1 - f(x)$, denoted $\text{sk}_{\bar{f}}$. The verifier then sends both $(\text{mpk}_1, \text{mpk}_2)$ and $(\text{sk}_f, \text{sk}_{\bar{f}})$ to the prover. Generating sk_f and $\text{sk}_{\bar{f}}$ takes time that is proportional to the size of the circuit computing f , and thus is a costly operation. However, this is done only once in the preprocessing phase.

Whenever the verifier wants the prover to compute f on an input x , he chooses two random messages m_1, m_2 and sends the prover the encryptions of (x, m_*) under the two keys: $(\text{Enc}(\text{mpk}_1, x, m_1)$ and $\text{Enc}(\text{mpk}_2, x, m_2))$. The properties of the public-index functional encryption scheme guarantees that, if

$f(x) = 1$, the prover obtains m_1 using sk_f and \perp using $\text{sk}_{\bar{f}}$ so no information about m_0 , and vice versa if $f(x) = 0$. Therefore, the fact that the prover provides m_1 to the verifier is a proof that $f(x)$ was 1.

Importantly, this delegation scheme can be made to have the desired property of being *publicly verifiable*, meaning that the verifier can produce a “verification key” with which anyone can check the prover’s work. This is done by having the verifier also send two point function obfuscations, one of the point m_1 and the other of the point m_2 .

This reduction from public-index FE to publicly verifiable delegation can be combined with the recent result of [GVW12b] providing public-index FE schemes for any depth circuit: the result is a publicly verifiable 2-message delegation scheme in the preprocessing model for any depth d circuit with verifier’s work being proportional to the depth d and the prover’s work proportional to the circuit size.

Note however, that this scheme is not secret because the public-index FE does not hide the input x from the prover. It is well known that x can be made secret by encrypting everything using a fully homomorphic encryption scheme. However, this comes at the cost of losing the public verifiability property. Our idea is to replace the public-index FE scheme with our (private-index) functional encryption scheme in the protocol above; now the ciphertexts $\text{Enc}(\text{mpk}_1, x, m_1)$ and $\text{Enc}(\text{mpk}_2, x, m_2)$ hide x and the scheme provides secrecy because the prover learns nothing about x other than $f(x)$. The public verifiability of the scheme remains the same.

We remark that we could provide a stronger version of secrecy by also hiding the result $f(x)$ from the prover; such stronger secrecy is non-standard for delegation, so we do not delve on it. (The idea is for the client to concatenate a random bit to each input x and have the function f output the opposite result when the bit is set. In this way, the prover does not learn anything from seeing which ciphertext decrypts to non- \perp .)

2 Preliminaries

2.1 Notation

We let κ denote the security parameter throughout this paper. For a distribution \mathcal{D} , we say $x \leftarrow \mathcal{D}$ when x is sampled from the distribution \mathcal{D} . If S is a finite set, by $x \leftarrow S$, we mean x is sampled from the uniform distribution over the set S . We use $p(\cdot)$ to denote that p is a function that takes one input. Similarly, $p(\cdot, \cdot)$ denotes a function p that takes two inputs.

We say that a function f is negligible in an input parameter κ , if for all $d > 0$, there exists K such that for all $\kappa > K$, $f(\kappa) < \kappa^{-d}$. For brevity, we write: for all sufficiently large κ , $f(\kappa) = \text{negl}(\kappa)$. We say that a function f is polynomial in an input parameter κ , if there exists a polynomial p such that for all κ , $f(\kappa) \leq p(\kappa)$. We write $f(\kappa) = \text{poly}(\kappa)$. A similar definition holds for $\text{polylog}(\kappa)$.

Let $[n]$ denote the set $\{1, \dots, n\}$ for $n \in \mathbb{N}$. When saying that a Turing machine A is p.p.t. we mean that A is a non-uniform probabilistic polynomial time machine.

Two ensembles, $X = \{X_\kappa\}_{\kappa \in \mathbb{N}}$ and $Y = \{Y_\kappa\}_{\kappa \in \mathbb{N}}$, are said to be *computationally indistinguishable* (and denoted $\{X_\kappa\}_{\kappa \in \mathbb{N}} \stackrel{c}{\approx} \{Y_\kappa\}_{\kappa \in \mathbb{N}}$) if for every probabilistic polynomial-time algorithm D ,

$$|\Pr[D(X_\kappa, 1^\kappa) = 1] - \Pr[D(Y_\kappa, 1^\kappa) = 1]| = \text{negl}(\kappa).$$

In our security definitions, we will define probabilistic experiments and denote by random variables their outputs. For example, $\text{Exp}_{E,A}^{\text{real}}(1^\kappa)$ denotes the random variable representing the output of the real experiment for scheme E with adversary A on security parameter κ . Moreover, $\{\text{Exp}_{E,A}^{\text{real}}(1^\kappa)\}_{\kappa \in \mathbb{N}}$ denotes the ensemble of such random variables indexed by $\kappa \in \mathbb{N}$.

2.2 Background on Learning With Errors (LWE)

The security of our results will be based on the Learning with Errors (LWE) assumption, first introduced by Regev [Reg05]. Regev [Reg05] showed that solving the LWE problem *on the average* is (quantumly) as hard as solving the approximate version of several standard lattice problems, such as *gapSVP in the worst case*. Peikert [Pei09] later removed the quantum assumption from a variant of this reduction. Given this connection, we state all our results under worst-case lattice assumptions, and in particular, under (a variant of) the *gapSVP* assumption. We refer the reader to [Reg05, Pei09] for details about the worst-case/average-case connection.

The best known algorithms to solve these lattice problems with an approximation factor 2^{ℓ^ϵ} in ℓ -dimensional lattices run in time $2^{\tilde{O}(\ell^{1-\epsilon})}$ [AKS01, MV10]. Specifically, given the current state of the art on lattice algorithms, it is quite plausible that achieving approximation factors 2^{ℓ^ϵ} for these lattice problems for any constant $0 < \epsilon < 1$ is hard for polynomial time algorithms.

Appendix A provides more detailed background information on LWE.

2.3 Background on Fully Homomorphic Encryption

The notion of fully homomorphic encryption was first proposed by Rivest, Adleman and Dertouzos [RAD78] in 1978. The first fully homomorphic encryption scheme was proposed in a breakthrough work by Gentry in 2009 [Gen09]. A history and recent developments on fully homomorphic encryption is surveyed in [Vai11]. We recall the definitions and semantic security of fully homomorphic encryption; the definitions below are based on [Vai11] with some adaptations.

Definition 2.1. *A homomorphic (public-key) encryption scheme HE is a quadruple of polynomial time algorithms (HE.KeyGen, HE.Enc, HE.Dec, HE.Eval) as follows:*

- HE.KeyGen(1^κ) is a probabilistic algorithm that takes as input the security parameter 1^κ and outputs a public key pk and a secret key sk .
- HE.Enc($\text{pk}, x \in \{0, 1\}$) is a probabilistic algorithm that takes as input the public key pk and an input bit x and outputs a ciphertext ψ .
- HE.Dec(sk, ψ) is a deterministic algorithm that takes as input the secret key sk and a ciphertext ψ and outputs a message $x^* \in \{0, 1\}$.
- HE.Eval($\text{pk}, C, \psi_1, \psi_2, \dots, \psi_n$) is a deterministic algorithm that takes as input the public key pk , some circuit C that takes n bits as input and outputs one bit, as well as n ciphertexts ψ_1, \dots, ψ_n . It outputs a ciphertext ψ_C .

Compactness: For all security parameters κ , there exists a polynomial $p(\cdot)$ such that for all input sizes n , for all $x_1 \dots x_n$, for all C , the output length of HE.Eval is at most $p(n)$ bits long.

Definition 2.2 (C-homomorphism). Let $\mathcal{C} = \{C_n\}_{n \in \mathbb{N}}$ be a class of boolean circuits, where C_n is a set of boolean circuits taking n bits as input. A scheme HE is \mathcal{C} -homomorphic if for every polynomial $n(\cdot)$, for every sufficiently large security parameter κ , for every circuit $C \in \mathcal{C}_n$, and for every input bit sequence x_1, \dots, x_n ,

where $n = n(\kappa)$,

$$\begin{aligned} & \Pr[(pk, sk) \leftarrow \text{HE.KeyGen}(1^\kappa); \\ & \quad \psi_i \leftarrow \text{HE.Enc}(pk, x_i) \text{ for } i = 1 \dots n; \\ & \quad \psi \leftarrow \text{HE.Eval}(pk, C, \psi_1, \dots, \psi_n) : \\ & \quad \text{HE.Dec}(sk, \psi) \neq C(x_1, \dots, x_n)] = \text{negl}(\kappa). \end{aligned}$$

where the probability is over the coin tosses of HE.KeyGen and HE.Enc .

Definition 2.3 (Fully homomorphic encryption). *A scheme HE is fully homomorphic if it is homomorphic for the class of all arithmetic circuits over $\text{GF}(2)$.*

Definition 2.4 (Leveled fully homomorphic encryption). *A leveled fully homomorphic encryption scheme is a homomorphic scheme where HE.KeyGen receives an additional input 1^d and the resulting scheme is homomorphic for all depth- d arithmetic circuits over $\text{GF}(2)$.*

Definition 2.5 (IND-CPA security). *A scheme HE is IND-CPA secure if for any p.p.t. adversary \mathcal{A} ,*

$$\left| \Pr[(pk, sk) \leftarrow \text{HE.KeyGen}(1^\kappa) : \mathcal{A}(pk, \text{HE.Enc}(pk, 0)) = 1] - \Pr[(pk, sk) \leftarrow \text{HE.KeyGen}(1^\kappa) : \mathcal{A}(pk, \text{HE.Enc}(pk, 1)) = 1] \right| = \text{negl}(\kappa).$$

We now state the result of Brakerski, Gentry and Vaikuntanathan [BGV12] that shows a leveled fully homomorphic encryption scheme based on the LWE assumption:

Theorem 2.6 ([BV11b, BGV12]). *Assume that there is a constant $0 < \epsilon < 1$ such that for every sufficiently large ℓ , the approximate shortest vector problem gapSVP in ℓ dimensions is hard to approximate to within a $2^{O(\ell^\epsilon)}$ factor in the worst case. Then, for every n and every polynomial $d = d(n)$, there is a d -leveled fully homomorphic encryption scheme where encrypting n bits produces ciphertexts of length $\text{poly}(n, \kappa, d^{1/\epsilon})$, the size of the circuit for homomorphic evaluation of a function f is $\text{size}(C_f) \cdot \text{poly}(n, \kappa, d^{1/\epsilon})$ and its depth is $\text{depth}(C_f) \cdot \text{poly}(\log n, \log d)$.*

Note that while all known fully homomorphic encryption scheme (as opposed to merely a leveled one) require an additional assumption related to circular security of the associated encryption schemes, we do not need to make any such assumptions in this work as we will only use a leveled homomorphic encryption scheme in all our constructions.

2.4 Background on Garbled Circuits

We will now define garbled circuits. Initially, garbled circuits were presented by Yao [Yao82] in the context of secure two-party computation. They were then proven secure by Lindell and Pinkas [LP09]. Very recently, the notion has been formalized by Bellare et al. [BHR12]. For simplicity, we present more concise definitions of garbled circuits than in [BHR12].

Definition 2.7 (Garbling scheme). *A garbling scheme for a family of circuits $\mathcal{C} = \{C_n\}_{n \in \mathbb{N}}$ with C_n a set of boolean circuits taking as input n bits, is a tuple of p.p.t. algorithms $\text{Gb} = (\text{Gb.Garble}, \text{Gb.Enc}, \text{Gb.Eval})$ such that*

- $\text{Gb.Garble}(1^\kappa, C)$ takes as input the security parameter κ and a circuit $C \in \mathcal{C}_n$ for some n , and outputs the garbled circuit Γ and a secret key sk .

- $\text{Gb.Enc}(\text{sk}, x)$ takes as input $x \in \{0, 1\}^*$ and outputs an encoding c .
- $\text{Gb.Eval}(\Gamma, c)$ takes as input a garbled circuit Γ , an encoding c and outputs a value y which should be $C(x)$.

Correctness. For any polynomial n , for all sufficiently large security parameters κ , for all circuits $C \in \mathcal{C}_n$ and all $x \in \{0, 1\}^n$,

$$\Pr[(\Gamma, \text{sk}) \leftarrow \text{Gb.Garble}(1^\kappa, C); c \leftarrow \text{Gb.Enc}(\text{sk}, x); y \leftarrow \text{Gb.Eval}(\Gamma, c) : C(x) = y] = 1 - \text{negl}(\kappa).$$

Efficiency. There exists a polynomial $p = p(\kappa, n)$ independent of the size of circuits in \mathcal{C} such that the running time of Gb.Enc is at most $p(\kappa, n)$ where n is the size of the input x .

Intuitively, the efficiency property says that the encoding algorithm should do little work, independent of the size of C . Ideally, p should be independent of the class \mathcal{C} (and not merely of the size of circuits in \mathcal{C}), but in some cases p can depend on the depth of circuits in \mathcal{C} , but not on their sizes. Since Gb.Enc is a p.p.t., the size of the encoding c will also be compact: a polynomial in (κ, n) .

Yao garbled circuits. The garbled circuits presented by Yao have a specific property of the encoding scheme that is useful in various secure function evaluation protocols and in our construction as well. The secret key is of the form $\text{sk} = \{L_i^0, L_i^1\}_{i=1}^n$ and the encoding of an input x of n bits is of the form $c = (L_1^{x_1}, \dots, L_n^{x_n})$, where x_i is the i -th bit of x .

Two security guarantees are of interest: input privacy (the input to the garbled circuit does not leak to the adversary), and circuit privacy (the circuit does not leak to the adversary). All these properties hold only for one-time evaluation of the circuit: the adversary can receive at most one encoding of an input to use with a garbled circuit; obtaining more than one encoding breaks these security guarantees.

Bellare et al. [BHR12] also present a third property which they call authenticity; informally, this requires that an adversary should not be able to come up with a different result of the garbled circuit that could be “de-garbled” into a valid value. We do not present this property here because it is straightforward to show that a garbling scheme with input and circuit privacy as we define them below implies a different garbling scheme with the authenticity property and we would need to provide a slightly more complicated syntax for the definition of garbled circuits (with an additional “de-garbling” algorithm).

We now present the one-time security of garbling circuits. The security definition for reusable garbled will be presented later, in Sec. 4.

Definition 2.8 (Input and circuit privacy). *A garbling scheme Gb for a family of circuits $\{\mathcal{C}_n\}_{n \in \mathbb{N}}$ is input and circuit private if there exists a p.p.t. simulator $\text{Sim}_{\text{Garble}}$, such that for every p.p.t. adversaries A and D , for all sufficiently large security parameters κ ,*

$$\left| \Pr[(x, C, \alpha) \leftarrow A(1^\kappa); (\Gamma, \text{sk}) \leftarrow \text{Gb.Garble}(1^\kappa, C); c \leftarrow \text{Gb.Enc}(\text{sk}, x) : D(\alpha, x, C, \Gamma, c) = 1] - \Pr[(x, C, \alpha) \leftarrow A(1^\kappa); (\tilde{\Gamma}, \tilde{c}) \leftarrow \text{Sim}_{\text{Garble}}(1^\kappa, C(x), 1^{|C|}, 1^{|x|}) : D(\alpha, x, C, \tilde{\Gamma}, \tilde{c}) = 1] \right| = \text{negl}(\kappa),$$

where we only consider A such that for some n , $x \in \{0, 1\}^n$ and $C \in \mathcal{C}_n$.

Intuitively, this definition says that, for any circuit or input chosen adversarially, one can simulate in polynomial time the garbled circuit and the encoding solely based on the computation result (and relevant sizes). The variable α represents any state that A may want to convey to D .

A few variants of Yao garbling schemes exist (for example, [BHR12]) that provide both input and circuit privacy under the basic one-way function assumption. Any such construction is suitable for our scheme.

Theorem 2.9 ([Yao82, LP09]). *Assuming one-way functions exist, there exists a Yao (one-time) garbling scheme that is input- and circuit-private for all circuits over $GF(2)$.*

2.5 Background on Functional Encryption

We recall the functional encryption definition from the literature [KSW08, BSW11, GVW12a] with some notational changes.

Definition 2.10 (Functional Encryption (FE)). *A functional encryption scheme FE for a class of functions $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$, represented as boolean circuits with an n -bit input, is a tuple of four p.p.t. algorithms (FE.Setup, FE.KeyGen, FE.Enc, FE.Dec) such that:*

- FE.Setup(1^κ) takes as input the security parameter 1^κ and outputs a master public key fmpk and a master secret key fmsk.
- FE.KeyGen(fmsk, f) takes as input the master secret key fmsk and a function $f \in \mathcal{F}$ and outputs a key fsk_f .
- FE.Enc(fmpk, x) takes as input the master public key fmpk and an input $x \in \{0, 1\}^*$ and outputs a ciphertext c .
- FE.Dec(fsk_f, c) takes as input a key fsk_f and a ciphertext c and outputs a value y .

Correctness. *For any polynomial $n(\cdot)$, for every sufficiently large security parameter κ , for $n = n(\kappa)$, for all $f \in \mathcal{F}_n$, and all $x \in \{0, 1\}^n$,*

$$\Pr[(\text{fmpk}, \text{fmsk}) \leftarrow \text{FE.Setup}(1^\kappa); \text{fsk}_f \leftarrow \text{FE.KeyGen}(\text{fmsk}, f); c \leftarrow \text{FE.Enc}(\text{fmpk}, x) : \text{FE.Dec}(\text{fsk}_f, c) = f(x)] = 1 - \text{negl}(\kappa).$$

2.5.1 Security of Functional Encryption

Intuitively, the security of functional encryption requires that an adversary should not learn anything about the input x other than the computation result $C(x)$, for some circuit C for which a key was issued. Two notions of security have been used in previous work: full and selective security. Full security allows the adversary to provide the challenge ciphertext after seeing the public key, whereas in selective security, the adversary must provide the challenge ciphertext before seeing the public key. We present simulation-based definitions for functional encryption both in the full (FULL-SIM-security) and selective (SEL-SIM-security) cases, two standard security notions in the functional encryption literature. The reason we present both definitions is that we will achieve them with different parameters of the gapSVP assumption. The security definition states that whatever information an adversary is able to learn from the ciphertext and the function keys, can be simulated given only the function keys and the output of the function on the inputs.

Definition 2.11 (FULL-SIM-Security). *Let FE be a functional encryption scheme for the family of functions $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$. For every p.p.t. adversary $A = (A_1, A_2)$ and p.p.t. simulator S , consider the following two experiments:*

 $\text{Exp}_{\text{FE},A}^{\text{real}}(1^\kappa):$

- 1: $(\text{fmpk}, \text{fmsk}) \leftarrow \text{FE.Setup}(1^\kappa)$
- 2: $(f, \text{state}_A) \leftarrow A_1(\text{fmpk})$
- 3: $\text{fsk}_f \leftarrow \text{FE.KeyGen}(\text{fmsk}, f)$
- 4: $(x, \text{state}'_A) \leftarrow A_2(\text{state}_A, \text{fsk}_f)$
- 5: $c \leftarrow \text{FE.Enc}(\text{fmpk}, x)$
- 6: **Output** (state'_A, c)

 $\text{Exp}_{\text{FE},A,S}^{\text{ideal}}(1^\kappa):$

- 5: $\tilde{c} \leftarrow S(\text{fmpk}, \text{fsk}_f, f, f(x), 1^{|x|})$
 - 6: **Output** $(\text{state}'_A, \tilde{c})$
-

The scheme is said to be (single-key) FULL-SIM-secure if there exists a p.p.t. simulator S such that for all pairs of p.p.t. adversaries (A_1, A_2) , the outcomes of the two experiments are computationally indistinguishable:

$$\left\{ \text{Exp}_{\text{FE},A}^{\text{real}}(1^\kappa) \right\}_{\kappa \in \mathbb{N}} \stackrel{c}{\approx} \left\{ \text{Exp}_{\text{FE},A,S}^{\text{ideal}}(1^\kappa) \right\}_{\kappa \in \mathbb{N}}.$$

We now define selective security, which is a weakening of full security, by requiring the adversary to provide the challenge input x before seeing the public key or any other information besides the security parameter. We simply specify the difference from full security.

Definition 2.12 (SEL-SIM-Security). *The same as Def. 2.11, but modify the game so that the first step consists of A specifying the challenge input x given only the security parameter.*

It is easy to see that the full simulation definition (FULL-SIM-security) implies the selective definition (SEL-SIM-security).

The literature [BSW11, AGVW12] has considered another classification for simulation-based definitions: adaptive versus non-adaptive security. In the adaptive case, the adversary is allowed to ask for a function f after seeing the ciphertext c for an input x . In the non-adaptive case, the adversary must first provide f and only then ask for encryptions of inputs x . Our definition falls in the non-adaptive category. Boneh et al. [BSW11] have shown that adaptive simulation-based security is unachievable for even one single-key functional encryption scheme for even a simple functionality such as identity-based encryption. As such, the adaptive definition appears too strong and is unachievable for general functionalities, so we use non-adaptive security.

2.5.2 Public-Index Functional Encryption

We now define public-index functional encryption which is an important step in our main construction. Public-index FE leaks more information than a functional encryption scheme, so one way to think of public-index FE is to consider a specific class of functionalities for the general functional encryption definition where the information leaked is part of the output to the function; we can then use the same simulation-based security definitions as for FE.

Namely, consider functions whose plaintext space consists of pairs of values from $\{0, 1\}^n \times \mathcal{M}$, where $\{0, 1\}^n$ is referred to as the *index space* (with an index size of n) and \mathcal{M} is referred to as the *message space*. The function is more specific: there exists an associated predicate class $\mathcal{P} = \{\mathcal{P}_n\}_{n \in \mathbb{N}}$ to \mathcal{F} such that for every n , for every $f \in \mathcal{F}_n$, there is an associated predicate $P \in \mathcal{P}_n$ to f such that

$$f(x, M) = \begin{cases} (x, M), & \text{if } P(x) = 1, \\ (x, \perp), & \text{otherwise.} \end{cases}$$

The index x is in the output of the function no matter what P is, meaning that x leaks from the scheme no matter what. This justifies the name “public-index”. On the other hand, M is only revealed if $P(x) = 1$. Therefore, this functionality leads to weaker security guarantees than a general functional encryption scheme not revealing x . Yet, we will use this as a central building block in our construction.

Remark 2.13 (Public-index FE (pFE) versus private-index FE (sFE)). *Some works call the general form of functional encryption (from Defs. 2.10, 2.11, 2.12) “private-index functional encryption” to emphasize the difference from public-index functional encryption; with private-index functional encryption, one can compute on the index x without revealing it. In public-index FE, the value on which the predicate computes (the index x) gets revealed, and only the “message” M is kept secret. The security guarantee of a public-index FE scheme is thus significantly weaker and qualitatively different than the one of a private-index FE scheme. As mentioned, throughout this paper, we use the term functional encryption to denote private-index functional encryption. Moreover, from now on, when we refer to a (private-index) functional encryption scheme, we will denote it with sFE where “s” comes from secret; when we refer to a public-index FE scheme, we will use the notation pFE where “p” comes from public.*

Public index functional encryption schemes have been constructed for the class of Boolean formulas [GPSW06, LOS⁺10] and most recently for the class of all polynomial-size circuits [GVW12b, SW12] of bounded depth.

Many public-index FE schemes have been proven secure under indistinguishability-based definitions. Despite being weaker than simulation-based definitions, such definitions suffice for the security of our construction, so we present them here. As before, we only provide the definition for the case when the adversary can only ask for a *single key* because this is all we need for our results. As with functional encryption, we define both full and selective security.

Definition 2.14 (Public-index functional encryption security). *Let pFE be a functional encryption scheme for a class of predicates $\mathcal{P} = \{P_n\}_{n \in \mathbb{N}}$, and an associated message space \mathcal{M} , and let $A = (A_1, A_2, A_3)$ be a triple of p.p.t. adversaries. Consider the following experiment.*

$\text{Exp}_{\text{pFE}}(1^\kappa)$:

- 1: $(\text{fmpk}, \text{fmsk}) \leftarrow \text{pFE.Setup}(1^\kappa)$
 - 2: $(P, \text{state}_1) \leftarrow A_1(\text{fmpk})$
 - 3: $\text{fsk}_P \leftarrow \text{pFE.KeyGen}(\text{fmsk}, P)$
 - 4: $(M_0, M_1, x, \text{state}_2) \leftarrow A_2(\text{state}_1, \text{fsk}_P)$
 - 5: Choose a bit b at random and let $c \leftarrow \text{pFE.Enc}(\text{fmpk}, (x, M_b))$.
 - 6: $b' \leftarrow A_3(\text{state}_2, c)$. If $|M_0| = |M_1|$, $P(x) = 0$, and $b = b'$, output 1, else output 0.
-

We say that the scheme is a single-key fully secure public-index functional encryption if for all p.p.t. adversaries A , and for all sufficiently large κ :

$$\Pr[\text{Exp}_{\text{pFE}, A}(1^\kappa) = 1] = 1/2 + \text{negl}(\kappa).$$

We say that the scheme is single-key selectively secure if the same statements hold for a slightly modified game in which A provides x before receiving fmpk.

Before we state the results of Gorbunov, Vaikuntanathan and Wee [GVW12b], we will set up some notation. Let d and p be polynomial functions. Define $\mathcal{C}_{n, d(n), p(n)}$ to be the class of all Boolean circuits on n

inputs of depth at most $d(n)$ and size at most $p(n)$. Let $\mathcal{C}_{n,d(n)} := \bigcup_{\text{polynomial } p} \mathcal{C}_{n,d(n),p(n)}$. A (public-index) functional encryption scheme that supports circuits in $\mathcal{C}_{n,d(n)}$ is called a d -leveled (public-index) functional encryption scheme. We are now ready to state the theorem of [GVW12b].

Theorem 2.15 ([GVW12b]). *Assume that there is a constant $0 < \epsilon < 1$ such that for every sufficiently large ℓ , the approximate shortest vector problem gapSVP in ℓ dimensions is hard to approximate by a polynomial algorithm to within a $2^{O(\ell^\epsilon)}$ factor in the worst case. Then, for every n and every polynomial $d = d(n)$, there is a selectively secure d -leveled public index functional encryption scheme where encrypting n bits produces ciphertexts of length $\text{poly}(n, \kappa, d^{1/\epsilon})$.*

Furthermore, assuming that gapSVP in ℓ dimensions is hard to approximate to within a $2^{O(\ell^\epsilon)}$ factor in time $2^{O(\ell^\epsilon)}$, the scheme is fully secure with ciphertexts of length $\text{poly}(n, \kappa, d^{1/\epsilon})$.

In either case, the scheme is secure with polynomially many secret-key queries.

2.5.3 Two-Outcome Public-Index Functional Encryption

We make use of a public-index functional encryption scheme with a slightly modified definition. The setup and key generation algorithms are the same as in previous schemes. The difference is in the encryption and decryption algorithms: instead of encrypting one message M in one ciphertext, we encrypt two messages M_0 and M_1 in the same ciphertext such that M_0 is revealed if the predicate evaluates to zero on the index, and M_1 is revealed if the predicate evaluates to one. Since there are two possible outcomes of the decryption algorithm, we call the modified scheme a *two-outcome functional encryption*.

Definition 2.16 (Two-Outcome Public-Index Functional Encryption). *A two-outcome public-index functional encryption scheme (pFE₂) for a class of predicates $\mathcal{P} = \{\mathcal{P}_n\}_{n \in \mathbb{N}}$ represented as boolean circuits with n input bits and an associated message space \mathcal{M} is a tuple of algorithms (pFE₂.Setup, pFE₂.KeyGen, pFE₂.Enc, pFE₂.Dec) as follows:*

- pFE₂.Setup(1^κ): *Takes as input a security parameter 1^κ and outputs a public master key fmpk and a master secret key fmsk.*
- pFE₂.KeyGen(fmsk, P): *Given a master secret key fmsk and a predicate $P \in \mathcal{P}_n$, for some n , outputs a key fsk_P corresponding to P .*
- pFE₂.Enc(fmpk, (x, M_0, M_1)): *Takes as input the public key fmpk, an index $x \in \{0, 1\}^n$, for some n , and two messages $M_0, M_1 \in \mathcal{M}$ and outputs a ciphertext c .*
- pFE₂.Dec(fsk_P, c): *Takes as input a secret key for a predicate and a ciphertext and outputs $M^* \in \mathcal{M}$.*

Correctness. *For any polynomial $n(\cdot)$, for every sufficiently large security parameter κ , if $n = n(\kappa)$, for all predicates $P \in \mathcal{P}_n$, indexes $x \in \{0, 1\}^n$, messages $M_0, M_1 \in \mathcal{M}$:*

$$\begin{aligned} & \Pr[(\text{fmpk}, \text{fmsk}) \leftarrow \text{pFE}_2.\text{Setup}(1^\kappa); \\ & \quad \text{sk}_P \leftarrow \text{pFE}_2.\text{KeyGen}(\text{fmsk}, P); \\ & \quad c \leftarrow \text{pFE}_2.\text{Enc}(\text{fmpk}, (x, M_0, M_1)); \\ & \quad M^* \leftarrow \text{pFE}_2.\text{Dec}(\text{sk}_P, c) : \\ & \quad M^* = M_{P(x)}] = 1 - \text{negl}(\kappa) \end{aligned}$$

We now define the security for single-key two-outcome public-index functional encryption. Intuitively, the security definition requires that, using a token for a predicate P , an adversary can decrypt one of the two messages encrypted in C , but does not learn anything about the other message.

Definition 2.17 (Two-outcome public-index functional encryption security). *Let pFE_2 be a two-outcome public-index functional encryption scheme for the class of predicates $\mathcal{P} = \{\mathcal{P}_n\}_{n \in \mathbb{N}}$ and associated message space \mathcal{M} and let $A = (A_1, A_2, A_3)$ be a triple of p.p.t. adversaries. Consider the following experiment.*

$\text{Exp}_{\text{pFE}_2}(1^\kappa)$:

- 1: $(\text{fmpk}, \text{fmsk}) \leftarrow \text{pFE}_2.\text{Setup}(1^\kappa)$
- 2: $(P, \text{state}_1) \leftarrow A_1(\text{fmpk})$
- 3: $\text{sk}_P \leftarrow \text{pFE}_2.\text{KeyGen}(\text{fmsk}, P)$
- 4: $(M, M_0, M_1, x, \text{state}_2) \leftarrow A_2(\text{state}_1, \text{sk}_P)$
- 5: Choose a bit b at random. Then, let

$$c = \begin{cases} \text{pFE}_2.\text{Enc}(\text{fmpk}, (x, M, M_b)), & \text{if } P(x) = 0, \\ \text{pFE}_2.\text{Enc}(\text{fmpk}, (x, M_b, M)), & \text{otherwise.} \end{cases}$$

- 6: $b' \leftarrow A_3(\text{state}_2, c)$. If $b = b'$, $\exists n$ such that, for all $P \in \mathcal{P}_n$, messages $M, M_0, M_1 \in \mathcal{M}$, $|M_0| = |M_1|$, $x \in \{0, 1\}^n$, output 1, else output 0.
-

We say that the scheme is a secure single-key two-outcome public-index FE if for all p.p.t. adversaries A , and for all κ :

$$\Pr[\text{Exp}_{\text{pFE}_2, A}(1^\kappa) = 1] = 1/2 + \text{negl}(\kappa).$$

The scheme is single-key selectively secure if A needs to provide x before receiving fmpk .

As before, we only need a single-key public-index FE scheme for our construction.

A class of predicates $\{\mathcal{P}_n\}_n$ is closed under negation if for all input sizes n and for all predicates $p \in \mathcal{P}_n$, we have $\bar{p} \in \mathcal{P}_n$; \bar{p} is the negation of p , namely $\bar{p}(y) = 1 - p(y)$ for all y .

Claim 2.18. *Assuming there is a public-index FE scheme for a class of predicates closed under negation, there exists a two-outcome public-index FE scheme for the same class of predicates.*

The proof of this claim is straightforward and we present it in Appendix B.

3 Our Functional Encryption Scheme

In this section, we present our main result: the construction of a (private-index) functional encryption scheme sFE. We refer the reader to the introduction (Sec. 1.2) for an overview of our approach, and we proceed with the construction here.

We use three building blocks in our construction: a (leveled) fully homomorphic encryption scheme HE, a (leveled) public-index functional encryption scheme pFE_2 , and a Yao garbling scheme Gb .

We let $\text{HE.Eval}_f(\text{hpk}, \bar{\psi})$ denote the circuit that performs homomorphic evaluation of the function f on the vector of ciphertexts $\bar{\psi} := (\psi_1, \dots, \psi_n)$ using the public key hpk , and we will let $\text{HE.Eval}_f^i(\text{hpk}, \psi)$

denote the predicate that computes the i -th output bit of $\text{HE.Eval}_f(\text{hpk}, \bar{\psi})$. Namely,

$$\text{HE.Eval}_f(\text{hpk}, \bar{\psi}) = \left(\text{HE.Eval}_f^1(\text{hpk}, \bar{\psi}), \dots, \text{HE.Eval}_f^\lambda(\text{hpk}, \bar{\psi}) \right),$$

where $\lambda = \lambda(\kappa) = |\text{HE.Eval}_f(\text{hpk}, \bar{\psi})|$. Our main theorem then says:

Theorem 3.1. *There is a (fully/selectively secure) single-key (private-index) functional encryption scheme $\text{sFE} = (\text{sFE.Setup}, \text{sFE.KeyGen}, \text{sFE.Enc}, \text{sFE.Dec})$ for any class of circuits \mathcal{C} that take n bits of input and produce a one-bit output, assuming the existence of the following primitives:*

- a CPA-secure \mathcal{C} -homomorphic encryption scheme $\text{HE} = (\text{HE.KeyGen}, \text{HE.Enc}, \text{HE.Eval}, \text{HE.Dec})$;
- a public index (fully/selectively secure) single-key functional encryption scheme $\text{pFE} = (\text{pFE.Setup}, \text{pFE.KeyGen}, \text{pFE.Enc}, \text{pFE.Dec})$ for the class of predicates $\mathcal{P} = \mathcal{P}_{\mathcal{C}, \text{HE}}$ where

$$\mathcal{P} = \{ \text{HE.Eval}_{\mathcal{C}}^i : \mathcal{C} \in \mathcal{C} \text{ and } i \in \mathbb{N} \} \quad \text{and}$$

- a Yao garbling scheme $\text{Gb} = (\text{Gb.Garble}, \text{Gb.Enc}, \text{Gb.Eval})$ that is input and circuit private.

The succinctness property of the functional encryption scheme is summarized as follows: the size of the ciphertexts $\text{ctsize}_{\text{sFE}}(n)$ in the resulting scheme for n bits of input is

$$2 \cdot \text{ctsize}_{\text{HE}} \cdot [\text{ctsize}_{\text{pFE}}(n \cdot \text{ctsize}_{\text{HE}} + \text{pksize}_{\text{HE}})] + \text{poly}(\kappa, \text{ctsize}_{\text{HE}}, \text{sksize}_{\text{HE}}).$$

where $\text{ctsize}_{\text{pFE}}(k)$ denotes the size of the ciphertexts in the public index scheme for a k -bit index and a $\text{poly}(\kappa)$ -bit message, $\text{ctsize}_{\text{HE}}$ denotes the size of the ciphertexts in the homomorphic encryption scheme for a single bit message and $\text{pksize}_{\text{HE}}$ (resp. $\text{sksize}_{\text{HE}}$) denotes the size of the public key (resp. secret key) in the homomorphic encryption scheme.

Since garbling schemes can be constructed from one-way functions, our theorem says that we can move from *public-index* functional encryption, in which the part of the input that the function computes on leaks, to a (private-index) functional encryption scheme, in which no part of the input leaks, by leveraging the power of fully homomorphic encryption.

We can see that if the ciphertext size in the public-index FE scheme and the homomorphic encryption scheme does not depend on the circuit size (and thus, those schemes are by themselves succinct), then neither will the resulting ciphertexts of the private-index FE scheme depend on the circuit size; namely, the reduction does not blow up the ciphertexts and is “succinctness-preserving”. We know of both a leveled FHE schemes and a public-index scheme ([GVW12b]) with ciphertext lengths independent of the size of the circuits to evaluate; the ciphertext size in these schemes just depends on the depth of the circuits.

We note that fully homomorphic encryption schemes with succinct ciphertexts independent of depth as well are known, albeit under the stronger assumption of circular security of the underlying schemes. Thus, if the result of [GVW12b] can be improved to remove the depth dependency of the ciphertexts in the public-index FE scheme, one automatically obtains a corresponding result for (private-index) FE using our reduction.

Our theorem only needs the public-index scheme to be secure with a single-key, even though the recent constructions [GVW12b] and [SW12] can tolerate an arbitrary number of keys.

Our main theorem is thus a reduction, which has a number of useful corollaries. The first and perhaps the most important one shows how to combine the leveled fully homomorphic encryption scheme from

[BV11b, BGV12] with the recent construction of a leveled public-index functional encryption scheme from [GVW12b] to obtain a (leveled, private index) functional encryption scheme based solely on the hardness of LWE. In other words, the corollary says that for every depth d , there is a (private index) functional encryption scheme for the class of all Boolean circuits of (arbitrary) polynomial size and depth at most d . The size of the ciphertexts in the scheme grows with d , and is of course independent of the size of the circuits it supports.

Let d and p be polynomial functions. Define $\mathcal{C}_{n,d(n),p(n)}$ to be the class of all Boolean circuits on n inputs of depth at most $d(n)$ and size at most $p(n)$. Let $\mathcal{C}_{n,d(n)} := \bigcup_{\text{polynomial } p} \mathcal{C}_{n,d(n),p(n)}$.

Corollary 3.2 (The LWE Instantiation). *We have the following two constructions of functional encryption based on the worst-case hardness of lattice problems:*

- Assume that there is a constant $0 < \epsilon < 1$ such that for every sufficiently large ℓ , the approximate shortest vector problem gapSVP in ℓ dimensions is hard to approximate to within a $2^{O(\ell^\epsilon)}$ factor (in polynomial time) in the worst case. Then, for every n and every polynomial $d = d(n)$, there is a selectively secure functional encryption scheme for the class $\mathcal{C}_{n,d(n)}$ where encrypting n bits produces ciphertexts of length $\text{poly}(n, \kappa, d^{1/\epsilon})$.
- Assume that there is a constant $0 < \epsilon < 1$ such that for every sufficiently large ℓ , the approximate shortest vector problem gapSVP in ℓ dimensions is hard to approximate to within a $2^{O(\ell^\epsilon)}$ factor in time $2^{O(\ell^\epsilon)}$ in the worst case. Then, for every n and every polynomial $d = d(n)$, there is a fully secure functional encryption scheme for the class $\mathcal{C}_{n,d(n)}$ where encrypting n bits produces ciphertexts of length $\text{poly}(n^{1/\epsilon}, \kappa, d^{1/\epsilon^2})$.

The corollary follows directly from Theorem 3.1, by invoking the leveled fully homomorphic encryption scheme of [BV11b] (see Theorem 2.6) and the leveled public-index functional encryption scheme of [GVW12a] (see Theorem 2.15). The concrete constructions and proofs in fact go through the learning with errors (LWE) problem; we refer to [BV11b, GVW12a] for the concrete setting of parameters.

Letting *universal* (public/private-index) functional encryption denote a single functional encryption scheme that supports the class of all polynomial-size circuits, we have the following corollary:

Corollary 3.3 (Universal Functional Encryption). *Assuming that fully homomorphic encryption schemes exist and universal single-key public-index functional encryption schemes exist, there is a universal single-key (private-index) functional encryption scheme.*

Of the two prerequisites mentioned above, we know that fully homomorphic encryption schemes exist (albeit under stronger assumptions than merely LWE). Thus, the corollary provides a way to immediately translate any universal public-index scheme into a private-index scheme. We point out that universal functional encryption schemes, by definition, have succinct ciphertexts.

A recent result of Gorbunov, Vaikuntanathan and Wee [GVW12a] shows how to generically convert single-key functional encryption schemes into q -keys functional encryption schemes for any bounded q , where the latter provide security against an attacker that can obtain secret keys of up to q functions of her choice. The size of the ciphertexts in the q -keys scheme grows polynomially with q .

Corollary 3.4 (Many queries, using [GVW12a]). *For every $q = q(\kappa)$, there is a (fully/selectively secure) q -keys functional encryption scheme for any class of circuits \mathcal{C} that take n bits of input and produce a one-bit output, assuming the existence of primitives as in Theorem 3.1. The size of the ciphertexts $\text{ctsize}_{\text{SE}}(n)$ in the resulting scheme is q times as large as in Theorem 3.1.*

Finally, a functional encryption scheme for circuits that output multiple bits can be constructed by thinking of the circuit as many circuits each with one-bit output, and modifying the key generation procedure to produce keys for each of them. This gives us the following corollary although we remark that more efficient methods of achieving this directly are possible using homomorphic encryption schemes that pack multiple bits into a single ciphertext [SV11, BGV12, GHS12a].

Corollary 3.5 (Many queries, many output bits). *For every $q = q(\kappa)$ and $k = k(n)$, there is a (fully/selectively secure) q -keys functional encryption scheme for any class of circuits \mathcal{C} that take n bits of input and produce k bits of output, assuming the existence of primitives as in Theorem 3.1. The size of the ciphertexts $\text{ctsize}_{\text{sFE}}(n)$ in the resulting scheme is qk times as large as in Theorem 3.1.*

Remark 3.6 (On the necessity of single-key security). *We note that even though the work of [GVW12b] provides a public-index functional encryption scheme that is secure even if the adversary obtains secret keys for polynomially many functions, our theorem only gives us a single-key secure scheme. Indeed, this is inherent by the impossibility result of [AGVW12] if we ask for (even a very weak notion of) simulation security. Corollary 3.4 gives us a way to get (simulation-)security with q queries for any a-priori bounded q , albeit at the expense of the ciphertext growing as a function of q .*

Remark 3.7 (On composing our functional encryption scheme). *One might wonder if chaining is possible with our FE scheme. Namely, one could try to generate keys for a function f that computes another function f_1 on an input x and then outputs $f_1(x)$ together with a new encryption of x under a different public key for the FE scheme. The new encryption of x could be used to compute a second function $f_2(x)$ and an encryption of x under yet another public key. This chain could potentially repeat. However, this approach only allows a very small number of iterations because, in order to produce one bit of output from sFE.Dec , the ciphertexts output by sFE.Enc are polynomial in κ . To obtain a sFE ciphertext as result of sFE.Dec , one needs to have started with ciphertexts of size quadratic in the first polynomial. If we want to chain the scheme q times, the original ciphertext must have been exponential in q .*

3.1 Construction

For simplicity, we construct sFE for functions outputting one bit; functions with larger outputs can be handled by repeating our scheme below for every output bit. The construction of $\text{sFE} = (\text{sFE.Setup}, \text{sFE.KeyGen}, \text{sFE.Enc}, \text{sFE.Dec})$ proceeds as follows.

From Claim 2.18, the existence of a secure single-key public-index FE scheme implies the existence of a two-outcome single-key public-index FE scheme, which we denote pFE_2 .

Let $\lambda = \lambda(\kappa)$ be the length of the ciphertexts in the HE scheme; this refers to both the ciphertexts produced by the HE encryption algorithm and to the ciphertexts produced by the HE evaluation algorithm.

Setup $\text{sFE.Setup}(1^\kappa)$: Run the setup algorithm for the two-outcome public-index FE scheme λ times:

$$(\text{fmpk}_i, \text{fmsk}_i) \leftarrow \text{pFE}_2.\text{Setup}(1^\kappa) \text{ for } i \in [\lambda].$$

Output:

$$\text{MPK} = (\text{fmpk}_1, \dots, \text{fmpk}_\lambda) \text{ and } \text{MSK} = (\text{fmsk}_1, \dots, \text{fmsk}_\lambda)$$

Key Generation $\text{sFE.KeyGen}(\text{MSK}, f)$: Let n be the number of bits in the input to the circuit f . As before, let

$$\text{HE.Eval}_f(\text{hpk}, \psi_1, \dots, \psi_n) = \text{HE.Eval}(\text{hpk}, f, \psi_1, \dots, \psi_n)$$

be the homomorphic evaluation algorithm for f . HE.Eval_f takes a public key hpk and n ciphertexts of λ bits each and produces a λ -bit ciphertext as the output. Let $\text{HE.Eval}_f^i : \{0, 1\}^{|\text{hpk}|} \times \{0, 1\}^{n\lambda} \rightarrow \{0, 1\}$ be the function that produces the i -th output bit of $\text{HE.Eval}_f(\text{hpk}, \psi_1, \dots, \psi_n)$ for a given f where $i \in [\lambda]$.

$\text{sFE.KeyGen}(\text{MSK}, f)$ does the following:

1. Run the key generation algorithm of pFE_2 for the functions HE.Eval_f^i (under the different master secret keys) to construct secret keys:

$$\text{fsk}_i \leftarrow \text{pFE}_2.\text{KeyGen}(\text{fmsk}_i, \text{HE.Eval}_f^i) \text{ for } i \in [\lambda].$$

2. Output the tuple $\text{fsk}_f := (\text{fsk}_1, \dots, \text{fsk}_\lambda)$ as the secret key for the function f .

Encryption $\text{sFE.Enc}(\text{MPK}, x)$: Let n be the number of bits of x , namely $x = x_1 \dots x_n$. Encryption proceeds in three steps.

1. Generate a fresh key pair $(\text{hpk}, \text{hsk}) \leftarrow \text{HE.KeyGen}(1^\kappa)$ for the (leveled) fully homomorphic encryption scheme. Encrypt each bit of x homomorphically: $\psi_i \leftarrow \text{HE.Enc}(\text{hpk}, x_i)$. Let $\psi := (\psi_1, \dots, \psi_n)$ be the encryption of the input x .
2. Run the Yao garbled circuit generation algorithm to produce a garbled circuit for the HE decryption algorithm $\text{HE.Dec}(\text{hsk}, \cdot) : \{0, 1\}^\lambda \rightarrow \{0, 1\}$ together with 2λ labels L_i^b for $i \in [\lambda]$ and $b \in \{0, 1\}$. Namely,

$$\left(\Gamma, \{L_i^0, L_i^1\}_{i=1}^\lambda \right) \leftarrow \text{Gb.Garble}(1^\kappa, \text{HE.Dec}(\text{hsk}, \cdot)),$$

where Γ is the garbled circuit and the L_i^b are the input labels.

3. Produce encryptions c_1, \dots, c_λ using the pFE_2 scheme in the following way. Let

$$c_i \leftarrow \text{pFE}_2.\text{Enc}(\text{fmpk}_i, ((\text{hpk}, \psi), L_i^0, L_i^1))$$

where (hpk, ψ) comes from the first step, and the labels (L_i^0, L_i^1) come from the second step.

The output of the encryption algorithm is $c = (c_1, \dots, c_\lambda, \Gamma)$.

Decryption $\text{sFE.Dec}(\text{fsk}_f, c)$:

1. First, run the pFE_2 decryption algorithm on the ciphertexts c_1, \dots, c_λ to recover the labels for the garbled circuit. In particular, let

$$L_i^{d_i} \leftarrow \text{pFE}_2.\text{Dec}(\text{fsk}_i, c_i) \text{ for } i \in [\lambda],$$

where d_i is a bit representing $\text{HE.Eval}_f^i(\text{hpk}, \psi)$.

2. Now, armed with the garbled circuit Γ and the labels $L_i^{d_i}$, run the garbled circuit evaluation algorithm to compute

$$\text{Gb.Eval}(\Gamma, L_1^{d_1}, \dots, L_\lambda^{d_\lambda}) = \text{HE.Dec}(\text{hsk}, d_1 d_2 \dots d_\lambda) = f(x)$$

3.2 Proof

We now proceed to prove Theorem 3.1 by proving that the theorem holds for our construction above.

Proof of Theorem 3.1. We first argue correctness.

Claim 3.8. *The above scheme is a correct functional encryption scheme (Def. 2.10).*

Proof. Let us examine the values we obtain in $\text{sFE.Dec}(\text{fsk}_f, c_1, \dots, c_\lambda, \Gamma)$. In Step (1), by the correctness of the pFE_2 scheme used, d_i is the i -th bit of $\text{HE.Eval}_f(\text{hpk}, \psi)$.

Therefore, the inputs to the garbled circuit Γ in Step (2) are the labels corresponding to $\text{HE.Eval}_f(\text{hpk}, \psi)$. By the correctness of the HE scheme, decrypting $\text{HE.Eval}_f(\text{hpk}, \psi)$ results in $f(x)$. Finally, by the correctness of the garbling scheme, the HE ciphertext gets decrypted correctly, yielding $f(x)$ as the output of sFE.Dec . \square

We now prove the succinctness property which follows directly from our construction. The output of sFE.Enc consists of λ pFE_2 ciphertexts and a garbled circuit. First, λ equals $\text{ctsize}_{\text{HE}}$. Second, each pFE_2 ciphertext consists of two pFE ciphertexts generated by pFE.Enc on input $n \cdot \text{ctsize}_{\text{HE}} + \text{psize}_{\text{HE}}$ bits. The labels of the garbled circuit are $\text{poly}(\kappa)$ in size. Third, the garbled circuit is the output of Gb.Garble so its size is polynomial in the size of the input circuit, which in turn is polynomial in $\text{sksize}_{\text{HE}}$ and $\text{ctsize}_{\text{HE}}$. Therefore, overall, we obtain $2 \cdot \text{ctsize}_{\text{HE}} \cdot \text{ctsize}_{\text{pFE}}(n \cdot \text{ctsize}_{\text{HE}} + \text{psize}_{\text{HE}}) + \text{poly}(\kappa, \text{sksize}_{\text{HE}}, \text{ctsize}_{\text{HE}})$. We can thus see that if HE and pFE produce ciphertexts independent of the circuit size, then so will our functional encryption scheme.

We focus on the full security case: namely, assuming pFE_2 is fully secure, we show that the resulting FE scheme is fully secure. We then discuss the proof for the selective case.

For full security, we construct a p.p.t. simulator S that achieves Def. 2.11. S receives as input $(\text{MPK}, \text{fsk}_f, f, f(x), 1^n)$ and must output \tilde{c} such that the real and ideal experiments in Def. 2.11 are computationally indistinguishable. Intuitively, S runs a modified version of sFE.Enc to mask the fact that it does not know x .

Simulator S on input $(\text{MPK}, \text{fsk}_f, f, f(x), 1^n)$:

1. Choose a key pair $(\text{hpk}, \text{hsk}) \leftarrow \text{HE.KeyGen}(1^\kappa)$ for the homomorphic encryption scheme (where S can derive the security parameter κ from the sizes of the inputs it gets). Encrypt 0^n (n zero bits) with HE by encrypting each bit individually and denote the ciphertext $\hat{0} := (\hat{0}_1 \leftarrow \text{HE.Enc}(\text{hpk}, 0), \dots, \hat{0}_n \leftarrow \text{HE.Enc}(\text{hpk}, 0))$.
2. Let $\text{Sim}_{\text{Garble}}$ be the simulator for the Yao garbling scheme (described in Def. 2.8) for the class of circuits corresponding to $\text{HE.Dec}(\text{hsk}, \cdot)$. Run $\text{Sim}_{\text{Garble}}$ to produce a simulated garbled circuit $\tilde{\Gamma}$ for the HE decryption algorithm $\text{HE.Dec}(\text{hsk}, \cdot) : \{0, 1\}^\lambda \rightarrow \{0, 1\}$ together with the simulated encoding consisting of one set of λ labels \tilde{L}_i for $i = 1 \dots \lambda$. Namely,

$$\left(\tilde{\Gamma}, \{\tilde{L}_i\}_{i=1}^\lambda \right) \leftarrow \text{Sim}_{\text{Garble}}(1^\kappa, f(x), 1^{|\text{HE.Dec}(\text{hsk}, \cdot)|}, 1^\lambda).$$

The simulator S can invoke $\text{Sim}_{\text{Garble}}$ because it knows $f(x)$, and can compute the size of the $\text{HE.Dec}(\text{hsk}, \cdot)$ circuit, and λ from the sizes of the input parameters.

3. Produce encryptions $\tilde{c}_1, \dots, \tilde{c}_\lambda$ under the pFE₂ scheme in the following way. Let

$$\tilde{c}_i \leftarrow \text{pFE}_2.\text{Enc} \left(\text{fmpk}_i, ((\text{hpk}, \hat{0}), \tilde{L}_i, \tilde{L}_i) \right),$$

where S reuses the simulated labels \tilde{L}_i .

4. Output $\tilde{c} = (\tilde{c}_1, \dots, \tilde{c}_\lambda, \tilde{\Gamma})$.

To prove indistinguishability of the real and ideal experiments (Def. 2.11), we define a sequence of hybrid experiments, and then invoke the security definitions of the underlying schemes (HE, garbled circuit, and pFE₂ respectively) to show that the hybrids are indistinguishable.

Hybrid 0 is the ideal experiment from Def. 2.11 for our sFE construction with simulator S . We denote it $\text{Exp}_{\text{sFE}, A}^{H_0}$ ($= \text{Exp}_{\text{sFE}, A, S}^{\text{ideal}}$).

Hybrid 1 ($\text{Exp}_{\text{sFE}, A}^{H_1}$) is the same as Hybrid 0, except that the ciphertext provided by the simulator, denoted $\tilde{c}^{(1)}$ for Hybrid 1, changes. Let $\tilde{c}^{(1)}$ be the ciphertext obtained by running the algorithm of S , except that in Step (1), we encrypt x instead of 0, namely:

$$\tilde{c}_i^{(1)} \leftarrow \text{pFE}_2.\text{Enc} \left(\text{fmpk}_i, ((\text{hpk}, \psi), \tilde{L}_i, \tilde{L}_i) \right),$$

where $\psi \leftarrow (\text{HE}.\text{Enc}(\text{hpk}, x_1), \dots, \text{HE}.\text{Enc}(\text{hpk}, x_n))$ and

$$\tilde{c}^{(1)} = (\tilde{c}_1^{(1)}, \dots, \tilde{c}_\lambda^{(1)}, \tilde{\Gamma}).$$

Hybrid 2 ($\text{Exp}_{\text{sFE}, A}^{H_2}$) is the same as Hybrid 1, except that in Step (2), the ciphertext contains a real garbled circuit

$$\left(\Gamma, \{L_i^0, L_i^1\}_{i=1}^\lambda \right) \leftarrow \text{Gb.Garble}(\text{HE}.\text{Dec}(\text{hsk}, \cdot)).$$

Let $d_i = \text{HE}.\text{Eval}_f^i(\text{hpk}, \psi)$. In Step (3), include L^{d_i} twice; that is

$$\tilde{c}_i^{(2)} \leftarrow \text{pFE}_2.\text{Enc} \left(\text{fmpk}_i, ((\text{hpk}, \psi), L_i^{d_i}, L_i^{d_i}) \right),$$

and

$$\tilde{c}^{(2)} = (\tilde{c}_1^{(2)}, \dots, \tilde{c}_\lambda^{(2)}, \Gamma).$$

Hybrid 3 ($\text{Exp}_{\text{sFE}, A}^{H_3}$) is the output of the real experiment from Def. 2.11 for our sFE construction.

We prove each pair of consecutive hybrids to be computationally indistinguishable in the following three lemmas, Lemmas 3.9, 3.10, and 3.11.

Lemma 3.9. *Assuming HE is CPA-secure, Hybrid 0 and Hybrid 1 are computationally indistinguishable.*

Proof. We proceed by contradiction. We assume that there exist p.p.t. adversaries $A = (A_1, A_2)$ and a p.p.t. distinguisher D such that D (with A) can distinguish between Hybrid 0 and Hybrid 1 above. Namely, there exists a polynomial $p(\cdot)$ such that, for infinitely many κ ,

$$|\Pr[D(\text{Exp}_{\text{sFE}, A}^{H_0}(1^\kappa)) = 1] - \Pr[D(\text{Exp}_{\text{sFE}, A}^{H_1}(1^\kappa)) = 1]| \geq 1/p(\kappa). \quad (1)$$

We construct a p.p.t. adversary $R = (R_1, R_2)$ that can break the semantic security of HE. Adversary R_1 outputs an n -bit value x for some n , and adversary R_2 receives as input either homomorphic encryption of x or of 0^n , and it will distinguish between these two. Distinguishing successfully implies that there is an adversary that can distinguish successfully in Def. 2.5, by a standard hybrid argument.

To determine x , adversary R_1 works as follows:

1. Run $\text{Exp}_{\text{sFE},A,S}^{\text{ideal}}(1^\kappa)$ (Def. 2.11) from Step (1) to Step (4) and let x be the output of A_2 in Step (4).
2. Output x .

To distinguish between encryption of x or 0^n , adversary R_2 receives input hpk^* , the HE public key, and an encryption E^* of x or 0^n and works as follows:

1. Run a modified algorithm of S by using hpk^* instead of generating fresh HE keys and using E^* instead of encrypting 0^n . Namely:
 - (a) Generate $(\tilde{\Gamma}, \{\tilde{L}_i\}_{i=1}^\lambda)$ as in Step (2) of S .
 - (b) Output $c^* = (c_1^*, \dots, c_\lambda^*)$ for $c_i^* = \text{pFE}_2.\text{Enc}(\text{fmpk}_i, ((\text{hpk}^*, E^*), \tilde{L}_i, \tilde{L}_i))$.
2. Feed $(c^*, \tilde{\Gamma})$ to D and output the decision of D .

Notice that if E^* is encryption of 0^n , R_2 simulates Hybrid 0 perfectly; when E^* is encryption of x , R_2 simulates Hybrid 1 perfectly. Therefore, D must have a probability of distinguishing between the two cases of at least $1/p(\kappa)$ (Eq. (1)); moreover, whenever D distinguishes correctly, R also outputs the correct decision. Therefore:

$$\begin{aligned} & |\Pr[x \leftarrow R_1(1^\kappa); (\text{hsk}^*, \text{hpk}^*) \leftarrow \text{HE.KeyGen}(1^\kappa) : R_2(\text{hpk}^*, \text{HE.Enc}(\text{hpk}^*, x)) = 1] - \\ & \Pr[(\text{hsk}^*, \text{hpk}^*) \leftarrow \text{HE.KeyGen}(1^\kappa) : R_2(\text{hpk}^*, \text{HE.Enc}(\text{hpk}^*, 0^n)) = 1]| = \\ & |\Pr[D(\text{Exp}_{\text{sFE},A}^{H_0}(1^\kappa)) = 1] - \Pr[D(\text{Exp}_{\text{sFE},A}^{H_1}(1^\kappa)) = 1]| \geq 1/p(\kappa), \end{aligned}$$

which contradicts the CPA security of the HE scheme. \square

Lemma 3.10. *Assuming the garbled circuit is circuit- and input-private (Def. 2.8), Hybrid 1 and Hybrid 2 are computationally indistinguishable.*

Proof. We proceed by contradiction. Assume there exist p.p.t. adversaries $A = (A_1, A_2)$ and a p.p.t. distinguisher D such that D (with A) can distinguish Hybrid 1 and Hybrid 2 above. Namely, there exists a polynomial p such that, for infinitely many κ ,

$$|\Pr[D(\text{Exp}_{\text{sFE},A}^{H_1}(1^\kappa)) = 1] - \Pr[D(\text{Exp}_{\text{sFE},A}^{H_2}(1^\kappa)) = 1]| \geq 1/p(\kappa). \quad (2)$$

We construct a stateful p.p.t. adversary $R = (R.A, R.D)$ that can break the security of the garbling scheme from Def. 2.8. The adversary $R.A$ has to provide a circuit G and an input I and then $R.D$ needs to distinguish between the simulated and the real garbled circuits and input encodings.

The adversary $R.A$ computes I and G as follows.

1. Run Steps (1)–(4) from Def. 2.11, which are the same in Hybrid 1 and Hybrid 2 and obtain f from A_1 and x from A_2 .
2. Generate $(\text{hsk}, \text{hpk}) \leftarrow \text{HE.KeyGen}(1^\kappa)$ and let $\psi \leftarrow \text{HE.Enc}(\text{hpk}, x)$.
3. Output $G(\cdot) := \text{HE.Dec}(\text{hsk}, \cdot)$ and $I := \text{HE.Eval}_f(\text{hpk}, \psi)$ and the following state for $R.D$: $\alpha = (\psi, \text{fmpk}_i, \text{hpk})$.

The adversary $R.D$ receives as input a garbled circuit Γ^* and a set of labels, one for each i : $\{L_i^*\}_{i=1}^\lambda$. These could be outputs of either $\text{Sim}_{\text{Garble}}$ or $\text{Gb.Garble}/\text{Gb.Enc}$ and $R.D$ decides which is an output of as follows:

1. Compute $c^* = (\{\text{pFE}_2.\text{Enc}(\text{fmpk}_i, ((\text{hpk}, \psi), L_i^*, L_i^*))\}_{i=1}^\lambda, \Gamma^*)$.
2. Run D on c^* and output what D outputs.

Notice that if $(\Gamma^*, \{L_i^*\}_{i=1}^\lambda)$ are outputs of $\text{Sim}_{\text{Garble}}$, R simulates Hybrid 1 perfectly; when $(\Gamma^*, \{L_i^*\}_{i=1}^\lambda)$ are outputs of the real garbling scheme, R simulates Hybrid 2 perfectly. Therefore, the probability that D distinguishes between the two cases at least is $1/p(\kappa)$ (Eq. (2)); moreover, whenever D distinguishes correctly, R also outputs the correct decision. Therefore:

$$\begin{aligned} & |\Pr[(G, I) \leftarrow R.A(1^\kappa) : R.D(\tilde{\Gamma}, \{\tilde{L}_i\}_{i=1}^\lambda) = 1] - \Pr[(G, I) \leftarrow R.A(1^\kappa) : R.D(\Gamma, \{L_i\}_{i=1}^\lambda) = 1]| = \\ & |\Pr[D(\text{Exp}_{\text{sFE}, A}^{H_1}(1^\kappa)) = 1] - \Pr[D(\text{Exp}_{\text{sFE}, A}^{H_2}(1^\kappa)) = 1]| \geq 1/p(\kappa), \end{aligned}$$

where, $(\tilde{\Gamma}, \{\tilde{L}_i\}_{i=1}^\lambda)$ are outputs of $\text{Sim}_{\text{Garble}}$ and $(\Gamma, \{L_i\}_{i=1}^\lambda)$ are outputs of $\text{Gb.Garble}/\text{Gb.Enc}$. This relation contradicts the security of the garbling scheme Def. 2.8. \square

Lemma 3.11. *Assuming the underlying pFE₂ scheme is fully secure, Hybrid 2 and Hybrid 3 in the fully secure setting above are computationally indistinguishable.*

Proof. In Hybrid 2 and Hybrid 3, there are λ pFE₂ encryptions, each with a pair of independent pFE₂ keys. First, we would like to prove that if Hybrid 2 and Hybrid 3 are computationally indistinguishable with only one of these encryptions, then they are computationally indistinguishable with λ encryptions. This would enable us to focus on only one pFE₂ ciphertext for the proof.

The argument proceeds in a standard way with a set of sub-hybrids, one for each index $i = 0 \dots \lambda$. The argument is straightforward because \tilde{c}_i and \tilde{c}_j (for $i \neq j$) use independently generated keys and the values encrypted with these keys are known to R . Hence, we only present the hybrid argument briefly. Sub-hybrid 0 corresponds to Hybrid 2 and sub-hybrid λ corresponds to Hybrid 3. Sub-hybrid i has the first i ciphertexts as in Hybrid 2 and the rest $\lambda - i$ as in Hybrid 3.

If an adversary A can distinguish between sub-hybrids $i - 1$ and i , for some i , then he can distinguish Hybrid 2 and Hybrid 3 for only one pair of ciphertexts (c_i^2, c_i^3) ; the reason is that we can build an adversary B : B places the challenge ciphertext in slot i of the challenge to A and produces the ciphertexts for all other slots $j \neq i$ with the correct distribution; B can do so because these ciphertexts are encrypted with fresh pFE₂ keys and B has all the information it needs to generate them correctly.

Now we are left to prove that Hybrid 2 and Hybrid 3 are indistinguishable when there is only one ciphertext, say the ℓ -th ciphertext. Namely, we need to prove that:

$$\left\{ (\text{state}'_A, \tilde{c}_\ell^{(2)}) \leftarrow \text{Exp}_{\text{sFE}, A}^{H_2}(1^\kappa) \right\} \stackrel{c}{\approx} \left\{ (\text{state}'_A, \tilde{c}_\ell^{(3)}) \leftarrow \text{Exp}_{\text{sFE}, A}^{H_3}(1^\kappa) \right\}. \quad (3)$$

We prove this statement by contradiction. Assume there exist p.p.t. adversaries $A = (A_1, A_2)$ and distinguisher D that can distinguish the distributions in (3); namely, there exists a polynomial $p(\cdot)$ such that, for infinitely many κ ,

$$|\Pr[D(\text{Exp}_{\text{sFE}, A}^{H_2}(1^\kappa)) = 1] - \Pr[D(\text{Exp}_{\text{sFE}, A}^{H_3}(1^\kappa)) = 1]| \geq 1/p(\kappa). \quad (4)$$

We construct a p.p.t. adversary $R = (R_1, R_2, R_3)$ that breaks the security of pFE₂ from Def. 2.17. R_1 , R_2 and R_3 send state to each other as in Def. 2.17, but for simplicity we will not denote this explicitly. R_3 aims to guess b in this definition.

Intuition. A and D can distinguish between Hybrid 2 and Hybrid 3. The only difference between these hybrids is that \tilde{c}_ℓ contains encryption of $(L_\ell^{d_\ell}, L_\ell^{d_\ell})$ versus $(L_\ell^{d_\ell}, L_\ell^{1-d_\ell})$. However, the pFE_2 scheme does not decrypt $L_\ell^{1-d_\ell}$ by the definition of d_ℓ , so its security hides the value of $L_\ell^{1-d_\ell}$. Since A and D can distinguish between these hybrids, they must be breaking the security of pFE_2 . Therefore, R will use L_ℓ^ℓ and $L_\ell^{1-\ell}$ as part of its answers to C and then use D to distinguish its challenge.

Specifically, the adversary R_1 receives as input fmpk^* in Step 2 of Def. 2.17 and computes P as follows:

1. Interact with adversary A_1 by running Steps (1)–(2) from Defs. 2.11 as follows.
 - (a) Let $\text{fmpk}_\ell := \text{fmpk}^*$. Generate the rest of pFE_2 keys using the $\text{pFE}_2.\text{Setup}$ algorithm: $(\text{fmpk}_i, \text{fmsk}_i) \leftarrow \text{pFE}_2.\text{Setup}(1^\kappa)$ for $i \neq \ell$.
 - (b) Receive f from A_1 and output $P := \text{HE.Eval}_f^\ell$.

Adversary R_2 receives sk_P^* in Step 4 of Def. 2.17 and computes M, M_0, M_1, x_c as follows:

1. Continue interaction with A_2 . To provide fsk_f to A_2 , compute $\text{fsk}_i \leftarrow \text{pFE}_2.\text{KeyGen}(\text{fmsk}_i, \text{HE.Eval}_f^i)$ for $i \neq \ell$, and let $\text{fsk}_\ell := \text{sk}_P^*$.
2. Receive x from A_2 .
3. Run the real garbled circuit generation as in Hybrid 2 and 3. Let $L_\ell^{d_\ell}$ be defined as in Hybrid 2. Provide $M := L_\ell^{d_\ell}, M_0 := L_\ell^{d_\ell}$ and $M_1 := L_\ell^{1-d_\ell}$.
4. Let $x_c := (\text{hpk}, \psi)$ where $\psi \leftarrow (\text{HE.Enc}(\text{hpk}, x_1), \dots, \text{HE.Enc}(\text{hpk}, x_n))$, the bitwise FHE encryption of x .
5. Output (M, M_0, M_1, x_c) .

Adversary R_3 receives as input a challenge ciphertext c^* and decides if it corresponds to M_0 or to M_1 as follows:

1. Let $\tilde{c}_\ell := c^*$ and provide $(\text{state}'_A, \tilde{c}_\ell)$ to D .
2. Output D 's guess.

In order for D to distinguish (as in Eq. (4)), the input distribution to A must be the one from Hybrid 2 or 3. We can see that this is the case: if $b = 0$, R simulates Hybrid 2 perfectly, and if $b = 1$, R simulates Hybrid 3 perfectly. Moreover, whenever D distinguishes correctly, R also outputs the correct decision. Therefore, by a simple calculation, we can see that

$$\Pr[\text{Exp}_{\text{pFE}_2, R}(1^\kappa) = 1] \geq 1/2 + 1/2p(\kappa),$$

which contradicts the security of the pFE_2 scheme, Def. 2.17. □

Returning to the proof of our theorem, by transitivity of computational indistinguishability, we showed that Hybrid 0 (the ideal experiment) is equivalent to Hybrid 3 (the real experiment), thus concluding our proof.

Selective security. The proof for the selective case follows similarly. The simulator S and the four hybrids are the same. Lemmas 3.9 and 3.10 proceed similarly, except that R now interacts with A as in the selective FE definition Def. 2.12 rather than Def. 2.11. The argument of Lemma 3.11 is the same, except that the order of some operations changes. This lemma makes the resulting FE scheme selective if one starts from a selective pFE_2 scheme. □

4 Reusable Garbled Circuits

In this section, we show how to construct garbled circuits that can be reused; namely, a garbled circuit that can run on an arbitrary number of encoded inputs without compromising the privacy of the circuit or of the input. For this goal, we build on top of our functional encryption scheme.

The syntax and correctness of the reusable garbling schemes remains the same as the one for one-time garbling schemes (Def. 2.7). In Sec. 2.4, we provided the one-time security definition for circuit and input privacy, Def. 2.8. We begin by defining security for more than one-time usage.

Definition 4.1 (Input and circuit privacy with reusability). *Let RGb be a garbling scheme for a family of circuits $\mathcal{C} = \{C_n\}_{n \in \mathbb{N}}$. For a pair of p.p.t. algorithms $A = (A_1, A_2)$ and a p.p.t. simulator $S = (S_1, S_2)$, consider the following two experiments:*

$\text{Exp}_{\text{RGb}, A}^{\text{real}}(1^\kappa)$:	$\text{Exp}_{\text{RGb}, A, S}^{\text{ideal}}(1^\kappa)$:
<ol style="list-style-type: none"> 1: $(C, \text{state}_A) \leftarrow A_1(1^\kappa)$ 2: $(\text{gsk}, \Gamma) \leftarrow \text{RGb.Garble}(1^\kappa, C)$ 3: $\alpha \leftarrow A_2^{\text{RGb.Enc}(\text{gsk}, \cdot)}(C, \Gamma, \text{state}_A)$ 4: Output α 	<ol style="list-style-type: none"> 1: $(C, \text{state}_A) \leftarrow A_1(1^\kappa)$ 2: $(\tilde{\Gamma}, \text{state}_S) \leftarrow S_1(1^\kappa, 1^{ C })$ 3: $\alpha \leftarrow A_2^{\text{O}(\cdot, C)[[\text{state}_S]]}(C, \tilde{\Gamma}, \text{state}_A)$ 4: Output α

In the above, $\text{O}(\cdot, C)[[\text{state}_S]]$ is an oracle that on input x from A_2 , runs S_2 with inputs $C(x)$, $1^{|x|}$, and the latest state of S ; it returns the output of S_2 (storing the new simulator state for the next invocation).

We say that the garbling scheme RGb is input- and circuit-private with reusability if there exists a p.p.t. simulator S such that for all pairs of p.p.t. adversaries $A = (A_1, A_2)$, the following two distributions are computationally indistinguishable:

$$\left\{ \text{Exp}_{\text{RGb}, A}^{\text{real}}(1^\kappa) \right\}_{\kappa \in \mathbb{N}} \stackrel{c}{\approx} \left\{ \text{Exp}_{\text{RGb}, A, S}^{\text{ideal}}(1^\kappa) \right\}_{\kappa \in \mathbb{N}}.$$

We can see that this security definition enables reusability of the garbled circuit: A_2 is allowed to make as many queries for input encodings as it wants.

From now on, by *reusable garbling scheme*, we will implicitly refer to a garbling scheme that has input and circuit privacy with reusability as in the definition above, Def. 4.1.

Remark 4.2. *We can provide an alternate syntax for a reusable garbling scheme, and we can also construct a scheme with this syntax (and a similar security definition) from our functional encryption scheme. This syntax has an additional setup algorithm (separate from the garble algorithm) that produces the secret key necessary for encoding and for circuit garbling; such a syntax would allow the garbled circuit to be generated after the encodings.*

Remark 4.3. *We do not provide a definition of authenticity because it is a straightforward extension of our scheme and is already achieved by [GVW12b]. We focus on circuit and input privacy, which have not been achieved by previous work.*

Recall the class of circuits $\mathcal{C}_{n, d(n)}$ defined for Corollary 3.2.

Theorem 4.4. *There exists a polynomial p , such that for every depth $d = d(n)$ function of the input size n , there is a reusable garbling scheme for any class of boolean circuits $\{\mathcal{C}_{n, d}\}_{n \in \mathbb{N}}$, assuming there is a fully secure single-key functional encryption scheme for any class of boolean circuits $\{\mathcal{C}_{n, p(d)}\}_{n \in \mathbb{N}}$.*

Corollary 4.5 (The LWE Instantiation). *For every integer $n \in \mathbb{N}$, polynomial function $d = d(n)$, there is a reusable garbling scheme for the class $\mathcal{C}_{n,d(n)}$, under the following assumption: there is a constant $0 < \epsilon < 1$ such that for every sufficiently large ℓ , the approximate shortest vector problem gapSVP in ℓ dimensions is hard to approximate to within a $2^{O(\ell^\epsilon)}$ factor in time $2^{O(\ell^\epsilon)}$ in the worst case.*

The proof of this corollary follows from Theorem 4.4 when instantiating the functional encryption scheme with the one from Corollary 3.2.

Denote by *universal reusable garbling scheme*, a reusable garbling scheme for the class of all polynomial-sized circuits. Then, the following corollary follows directly from Theorem 4.4:

Corollary 4.6 (Universal reusable garbled circuits). *If there is a universal single-key fully secure functional encryption scheme, there is a universal reusable garbling scheme.*

Notice that our functional encryption tool (sFE) almost gives reusable garbled circuits: the garbling of C is $\text{sFE.KeyGen}(\text{fmsk}, C)$, whereas the encoding of the input x is $\text{sFE.Enc}(\text{fmpk}, x)$. However, the problem is that sFE does not hide the circuit C , which is a required property of garbling schemes.

The insight in achieving circuit privacy is to use the input-hiding property of the sFE scheme to hide the circuit as well. The first idea that comes to mind is to hide C by including it in the ciphertext together with the input x . Specifically, instead of providing a key for circuit C , the encryptor runs sFE.KeyGen on a universal circuit U that on input (C, x) computes $C(x)$. Notice that U can be public because it carries no information about C other than its size. Now the encryption of x consists of an encryption of (C, x) using sFE.Enc . In this way, we can see that the resulting garbled circuit satisfies the correctness property. Moreover, for security, sFE hides the input (C, x) so it would hide the circuit C as well.

However, this approach is not useful because the encoding is as large as the circuit C (in particular, RGb.Enc no longer satisfies the efficiency property in Def. 2.7). Moreover, in this case, the standard one-time garbling schemes would be enough because one could produce a fresh garbled circuit with each ciphertext.

To overcome this problem, the idea is to provide, together with the ciphertext of x , *the ability to decrypt C* rather than the entire description of C . Specifically, let E be the encryption of the circuit C with a semantically secure symmetric encryption scheme under a secret key sk . The garbling of C consists of running the key generation sFE.KeyGen on a circuit U_E that includes E and works as follows. On input (x, sk) the circuit U_E decrypts E to obtain C , and outputs the result of running C on x . Even though $\text{sFE.KeyGen}(\text{fmsk}, U_E)$ does not hide U_E , the description of U_E does not leak C because C is encrypted. An encoding by RGb.Enc of x thus consists of running the encryption algorithm sFE.Enc on (x, sk) .

4.1 Construction

We construct a reusable garbling scheme $\text{RGb} = (\text{RGb.Garble}, \text{RGb.Enc}, \text{RGb.Eval})$ as follows. Let $\text{E} = (\text{E.KeyGen}, \text{E.Enc}, \text{E.Dec})$ be a semantically secure symmetric encryption scheme.

Garbling $\text{RGb.Garble}(1^\kappa, C)$:

1. Generate sFE keys $(\text{fmpk}, \text{fmsk}) \leftarrow \text{sFE.Setup}(1^\kappa)$ and a secret key $\text{sk} \leftarrow \text{E.KeyGen}(1^\kappa)$.
2. Let $E := \text{E.Enc}(\text{sk}, C)$.
3. Define U_E to be the following universal circuit:

U_E takes as input a secret key sk and a value x :

- (a) Compute $C := \text{E.Dec}(\text{sk}, E)$.
- (b) Run C on x .

4. Let $\Gamma \leftarrow \text{sFE.KeyGen}(\text{fmsk}, U_E)$ be the garbled circuit.
5. Output $\text{gsk} := (\text{fmpk}, \text{fmsk}, \text{sk})$ as the secret key and Γ as the garbling of C .

Encoding $\text{Rb.Enc}(\text{gsk}, x)$: Compute $c_x \leftarrow \text{sFE.Enc}(\text{fmpk}, (\text{sk}, x))$ and output c_x .

Evaluation $\text{Rb.Eval}(\Gamma, c_x)$: Compute and output $\text{sFE.Dec}(\Gamma, c_x)$.

The existence of a semantically secure encryption scheme does not introduce new assumptions because the sFE scheme itself is a semantically secure encryption scheme if no key (computed by sFE.KeyGen) is ever provided to an adversary.

Tightness of the scheme. The astute reader may have observed that the resulting scheme requires that the encodings be generated in the secret key setting because the encoding of x includes sk . It turns out that generating encodings privately is in fact necessary; if the encodings were publicly generated, the power of the adversary would be the same as in traditional obfuscation which was shown impossible [BGI⁺01, GK05] (see discussion in Sec. 1.1.2).

One might wonder though, whether a reusable garbling scheme exists where the encoding generation is secret key, but Rb.Garble is public key. We prove in Sec. 4.3 that such a scheme is also impossible; hence, with regard to the public versus private key settings, our reusable garbling result is tight.

4.2 Proof

Proof of Theorem 4.4. We first argue the scheme satisfies the correctness and efficiency properties in Def. 2.7.

Claim 4.7. *The above scheme Rb is a correct and efficient garbling scheme.*

Proof. We can easily see correctness of Rb.Eval :

$$\begin{aligned} \text{Rb.Eval}(\Gamma, c_x) &= \text{sFE.Dec}(\Gamma, c_x) \\ &= U_E(\text{sk}, x) \quad (\text{by the correctness of sFE}) \\ &= C(x). \end{aligned}$$

The efficiency of Rb depends on the efficiency of the sFE.Enc algorithm. If the runtime of sFE.Enc does not depend on the class of circuits to be computed at all, the same holds for Rb.Enc 's efficiency. If sFE.Enc depends on the depth of the circuits to be computed, as is the case in our LWE instantiation, Rb 's runtime also depends on the depth of the circuits, but still remains independent of the size of the circuits, which could potentially be much larger. \square

We can see that to obtain a Rb scheme for circuits of depth d , we need a sFE scheme for polynomially deeper circuits: the overhead comes from the fact that U is universal and it also needs to perform decryption of E to obtain C .

To prove security, we need to construct a simulator $S = (S_1, S_2)$ satisfying Def. 4.1, assuming there is a simulator Sim_{sFE} that satisfies Def. 2.11.

To produce a simulated garbled circuit $\tilde{\Gamma}$, S_1 on input $(1^\kappa, 1^{|C|})$ runs:

1. Generate fresh fmpk , fmsk , and sk as in Rb.Garble .
2. Compute $\tilde{E} := E.\text{Enc}(\text{sk}, 0^{|C|})$. (The reason for encrypting $0^{|C|}$ is that S_1 does not know C).
3. Compute and output $\tilde{\Gamma} \leftarrow \text{sFE.KeyGen}(\text{fmsk}, U_{\tilde{E}})$.

S_2 receives queries for values $x_1, \dots, x_t \in \{0, 1\}^*$ for some t and needs to output a simulated encoding for each of these. To produce a simulated encoding for x_i , S_2 receives inputs $(C(x_i), 1^{|x_i|}$ and the latest simulator's state) and invokes the simulator Sim_{sFE} of the sFE scheme and outputs

$$\tilde{c}_x := \text{Sim}_{\text{sFE}}(\text{fmpk}, \text{fsk}_{U_{\tilde{E}}}, U_{\tilde{E}}, C(x), 1^{|\text{sk}|+|x_i|}).$$

A potentially alarming aspect of this simulation is that S generates a key for the circuit $0^{|C|}$. Whatever circuit $0^{|C|}$ may represent, it may happen that there is no input x to $0^{|C|}$ that results in the value $C(x)$. The concern may then be that Sim_{sFE} may not simulate correctly. However, this is not a problem because, by semantic security, E and \tilde{E} are computationally indistinguishable so Sim_{sFE} must work correctly, otherwise it breaks semantic security of the encryption scheme E .

We now prove formally that the simulation satisfies Def. 4.1 for any adversary $A = (A_1, A_2)$. Let us assume that the α output of A_2 is its view, namely, all the information A_2 receives in the protocol, $(C, \text{state}_A, \Gamma, \{x_i, c_{x_i}\}_{i=1}^t)$. If the outcome of the real and ideal experiments are computationally indistinguishable in this case, then they are computationally indistinguishable for any other output strategy of A_2 because D can always run A_2 on its view since A_2 is p.p.t.. Therefore, we would like to show that:

$$\left\{ (C, \text{state}_A, \Gamma, \{x_i, c_{x_i}\}_{i=1}^t) \leftarrow \text{Exp}_{\text{RGb}, A}^{\text{real}}(1^\kappa) \right\}_\kappa \stackrel{c}{\approx} \left\{ (C, \text{state}_A, \tilde{\Gamma}, \{x_i, \tilde{c}_{x_i}\}_{i=1}^t) \leftarrow \text{Exp}_{\text{RGb}, A, S}^{\text{ideal}}(1^\kappa) \right\}_\kappa.$$

We show this relation with three hybrids. Consider the following games:

Game 0: The ideal game of Def. 4.1 with simulator S ; we recall that the output distribution in this case is

$$(C, \text{state}_A, \text{sFE.KeyGen}(\text{fmsk}, U_{\tilde{E}}), \{x_i, \text{Sim}_{\text{sFE}}(\text{fmpk}, \text{fsk}_{U_{\tilde{E}}}, U_{\tilde{E}}, C(x_i), 1^{|x_i|+|\text{sk}|})\}_{i=1}^t).$$

Game 1: The same as Game 0, but \tilde{E} is replaced with $E = E.\text{Enc}(\text{sk}, C)$. That is, the output distribution is

$$(C, \text{state}_A, \Gamma, \{x_i, \text{Sim}_{\text{sFE}}(\text{fmpk}, \text{fsk}_{U_E}, U_E, C(x_i), 1^{|x_i|+|\text{sk}|})\}_{i=1}^t).$$

Game 2: The real game with our construction for RGb. It consists of the output distribution

$$(C, \text{state}_A, \Gamma, \{x_i, c_{x_i}\}_{i=1}^t).$$

First, let us argue that the distributions output by Game 0 and Game 1 are computationally indistinguishable. Note that these two distributions only differ in E and \tilde{E} . Since these distributions do not contain sk or any other function of sk other than E/\tilde{E} , by semantic security of the encryption scheme, these two distributions are computationally indistinguishable. Finally, Lemma 4.8 proves that the outputs of Game 1 and Game 2 are also computationally indistinguishable, which concludes our proof.

Lemma 4.8. *Assuming sFE is FULL-SIM-secure, the outputs of Game 1 and Game 2 are computationally indistinguishable.*

Proof. The proof of the lemma is by contradiction. We assume there exist p.p.t. adversaries $A = (A_1, A_2)$ and p.p.t. distinguisher D such that D with A can distinguish Game 1 and Game 2. Namely, there exists a polynomial $p(\cdot)$ such that, for infinitely many κ ,

$$|\Pr[D(\text{Exp}_{\text{sFE}, A}^{\text{Game1}}(1^\kappa)) = 1] - \Pr[D(\text{Exp}_{\text{sFE}, A}^{\text{Game2}}(1^\kappa)) = 1]| \geq 1/p(\kappa). \quad (5)$$

We construct adversaries that break the full security of the functional encryption scheme Def. 2.11. We call these adversaries $A^{\text{sFE}} = (A_1^{\text{sFE}}, A_2^{\text{sFE}})$ and D^{sFE} using the “sFE” superscript to differentiate them from the adversaries distinguishing Game 1 and 2. In fact, we construct adversaries A^{sFE} and D^{sFE} that break a modified version of Def. 2.11: the modification is that A^{sFE} can repeat Steps (4–5) as many times as it wishes and adaptively; more precisely, for the i -th repetition of Steps (4–5), A_2^{sFE} can ask for an encryption of an input x_i where x_i could be determined based on the previous values and encryptions of x_1, \dots, x_{i-1} ; A_2^{sFE} receives either a real encryption or a simulated encryption as in Step (5), but either all encryptions are real or all are simulated. We can see that if A^{sFE} and D^{sFE} break this modified definition, then they must break the original definition (with a polynomially smaller advantage): this implication follows from a standard hybrid argument possible because the encryption of x_i is public key.

On input fmpk , adversary A_1^{sFE} works as follows:

1. Run A_1 on input 1^κ and obtain C and state_A .
2. Choose $\text{sk} \leftarrow \text{E.KeyGen}(1^\kappa)$, encrypt $E \leftarrow \text{E.Enc}(\text{sk}, C)$, and let U_E be the circuit described above.
3. Output function U_E and $\text{state}_A^{\text{sFE}} := (\text{sk}, U_E, \text{state}_A)$.

On input $(\text{fsk}_{U_E}, \text{state}_A^{\text{sFE}})$, adversary A_2^{sFE} works as follows:

1. Let $\Gamma := \text{fsk}_{U_E}$.
2. Run A_2 on U_E, Γ and state_A by answering to its oracle queries as follows.
 - (a) Consider the i -th oracle query (x_i, state_A) . Output (x_i, sk) .
 - (b) Receive as input CT_i which is either the real ciphertext $c_i \leftarrow \text{sFE.Enc}(\text{fmpk}, (x_i, \text{sk}))$ or the simulated ciphertext $\tilde{c}_i \leftarrow \text{Sim}_{\text{sFE}}(\text{fmpk}, \Gamma, U_E, C(x_i), 1^{|x_i|+|\text{sk}|})$. Respond to A_2 with $(\text{CT}_i, \text{state}_A)$.
 - (c) Repeat these steps until A_2 finishes querying for encodings, and outputs α .
3. Output α .

Adversary D^{sFE} is the same as D .

When the encodings CT_i are the ideal ciphertexts, we can see that $(A_1^{\text{sFE}}, A_2^{\text{sFE}})$ simulate perfectly Game 1; hence

$$\Pr[D^{\text{sFE}}(\text{Exp}_{\text{sFE}, A^{\text{sFE}}}^{\text{ideal}}(1^\kappa)) = 1] = \Pr[D(\text{Exp}_{\text{sFE}, A}^{\text{Game 1}}(1^\kappa)) = 1].$$

When the encodings CT_i are the real ciphertexts, $(A_1^{\text{sFE}}, A_2^{\text{sFE}})$ simulate perfectly Game 2 and thus

$$\Pr[D^{\text{sFE}}(\text{Exp}_{\text{sFE}, A^{\text{sFE}}}^{\text{real}}(1^\kappa)) = 1] = \Pr[D(\text{Exp}_{\text{sFE}, A}^{\text{Game 2}}(1^\kappa)) = 1].$$

By Eq. (5), we have

$$|\Pr[D^{\text{sFE}}(\text{Exp}_{\text{sFE}, A^{\text{sFE}}}^{\text{ideal}}(1^\kappa)) = 1] - \Pr[D^{\text{sFE}}(\text{Exp}_{\text{sFE}, A^{\text{sFE}}}^{\text{real}}(1^\kappa)) = 1]| \geq 1/p(\kappa),$$

which contradicts FULL-SIM-security of sFE. \square

Having proved that Game 0 and Game 1 are computationally indistinguishable, and that Game 1 and Game 2 are computationally indistinguishable, we conclude that Game 0 and Game 2 are computationally indistinguishable, and therefore that garbling scheme RGB is input- and circuit-private with reusability. \square

4.3 Impossibility of Public-Key Reusable Garbled Circuits

In this section, we show that a public-key reusable garbling scheme is impossible. Our argument is at a high level because it follows from existing results straightforwardly.

A public-key reusable garbling scheme would have the following syntax:

Definition 4.9 (Public-key garbling scheme). *A public-key garbling scheme PubGb for the class of circuits $\{\mathcal{C}_n\}_{n \in \mathbb{N}}$, with \mathcal{C}_n a set of boolean circuits taking n bits as input, is a tuple of p.p.t. algorithms (PubGb.Setup, PubGb.Garble, PubGb.Enc, PubGb.Eval) such that*

- PubGb.Setup(1^κ): Takes as input the security parameter 1^κ and outputs a secret key gsk and a public key gpk.
- PubGb.Garble(gpk, C): Takes as input a public key gpk and a circuit C , and outputs the garbled circuit Γ of the circuit C .
- PubGb.Enc(gsk, x): Takes as input the secret key gsk and an input x , and outputs an encoding c_x .
- PubGb.Eval(Γ, c_x): Takes as input a garbled circuit Γ and an encoding c_x and outputs a value y .

Correctness. For all polynomials $n(\cdot)$, for all sufficiently large security parameters κ , for $n = n(\kappa)$, for all circuits $C \in \mathcal{C}_n$, and for all $x \in \{0, 1\}^n$,

$$\Pr[(\text{gsk}, \text{gpk}) \leftarrow \text{PubGb.Setup}(1^\kappa); \Gamma \leftarrow \text{PubGb.Garble}(\text{gpk}, C); c_x \leftarrow \text{PubGb.Enc}(\text{gsk}, x) : \text{PubGb.Eval}(\Gamma, c_x) = C(x)] = 1 - \text{negl}(\kappa).$$

The natural security definition of circuit-private definition of this new scheme is similar in flavor to Def. 2.11, but we do not elaborate. (In fact, this definition can be relaxed to not require input privacy for the impossibility result to still hold.)

The first step in the impossibility argument is to note that the syntax and correctness of a public-key garbling scheme is the same as the syntax of a functional encryption scheme (Def. 2.10) with the following correspondence of algorithms: PubGb.Setup corresponds to FE.Setup, PubGb.Garble corresponds to the encryption algorithm FE.Enc, PubGb.Enc corresponds to FE.KeyGen and PubGb.Eval corresponds to FE.Dec. Note that PubGb.Enc does not correspond to FE.Enc but to FE.KeyGen because PubGb.Enc is a secret key algorithm and FE.Enc is a public-key algorithm. Therefore, an encoding of an input x in the reusable garbling scheme corresponds to a secret key for a function f_x in the functional encryption scheme.

Moreover, considering this mapping, it is straightforward to show that a circuit-private public-key garbling scheme implies a secure functional encryption scheme. Since the reusable garbling scheme allows an arbitrary number of inputs being encoded, it implies that the functional encryption scheme can generate an arbitrary number of secret function keys sk_{f_x} ; furthermore, in this functional encryption scheme, the size of the ciphertexts does not depend on the number of keys generated (because this number is nowhere provided as input in the syntax of the scheme). This conclusion directly contradicts the recent impossibility result of Agrawal et al. [AGVW12]: they show that any functional encryption scheme that can securely provide q keys must have the size of the ciphertexts grow in q ; therefore, a reusable circuit-private public-key garbling scheme is unachievable.

5 Token-Based Obfuscation

Following the discussion of obfuscation in Sec. 1.1.2, the purpose of this section is to cast reusable garbled circuits in the form of obfuscation and to show that this provides a new model for obfuscation, namely *token-based obfuscation*.

Reusable garbled circuits come close to obfuscation: a reusable garbled circuit hides the circuit while permitting circuit evaluation on an arbitrary number of inputs. While they come close, reusable garbled circuits do not provide obfuscation, because the encoding of each input requires knowledge of the secret key: namely, to run an obfuscated program on an input, one needs to obtain a token for the input from the obfuscator. This requirement of our scheme is in fact necessary: as argued in the tightness discussion in Sec. 4, a scheme in which one can publicly encode inputs is impossible because it falls directly onto known impossibility results for obfuscation.

Therefore, we propose a new token-based model for obfuscation. The idea is for a program vendor to obfuscate his program and provide tokens representing rights to run this program on specific inputs. For example, consider the case when some researchers want to compute statistics on a database with sensitive information. The program to be obfuscated consists of the database service program with the secret database hardcoded in it, U_{DB} . When researchers want to compute statistics x , they request a token for x from the database owner. Using the obfuscated program and the token, the researchers can compute $U_{\text{DB}}(x)$, the statistics result by themselves without having to contact the owner again. It is crucial that the time to compute the token for x is much smaller than the time to compute U_{DB} on x , so that the owner does not have to do a lot of work. We also note, that in certain cases, one has to anyways request such a token from the owner for other reasons: for example, the database owner can check that the statistics the researchers want to compute is not too revealing and only grant a token if this is the case.

5.1 Definition

We now provide the definition for token-based obfuscation and the desired simulation security. These definitions are very similar to the definitions for reusable garbled circuits (Def. 2.7 and Def. 4.1): the syntax, correctness and efficiency are the same except that garbling schemes have an additional Eval algorithm.

Definition 5.1 (Token-based Obfuscation). *A token-based obfuscation scheme for the class of circuits $\{\mathcal{C}_n\}_{n \in \mathbb{N}}$ with $\mathcal{C}_n : \{0, 1\}^n \rightarrow \{0, 1\}$ is a pair of p.p.t. algorithms (tOB.Obfuscate, tOB.Token) such that*

- tOB.Obfuscate($1^\kappa, C$): Takes as input the security parameter 1^κ , and a circuit $C \in \mathcal{C}_n$, and outputs a secret key osk and the obfuscation O of the circuit C .
- tOB.Token(osk, x): Takes as input the secret key osk and some input $x \in \{0, 1\}^n$, and outputs tk_x .

Efficiency. *The running time of tOB.Token is independent of the size of C .*

Correctness. *For all polynomials $n(\cdot)$, for all sufficiently large security parameters κ , if $n = n(\kappa)$, for all circuits $C \in \mathcal{C}_n$, and for all $x \in \{0, 1\}^n$,*

$$\Pr[(\text{osk}, O) \leftarrow \text{tOB.Obfuscate}(1^\kappa, C); \text{tk}_x \leftarrow \text{tOB.Token}(\text{osk}, x) : O(\text{tk}_x) = C(x)] = 1 - \text{negl}(\kappa).$$

Remark 5.2. *We could use an alternative definition of token-based obfuscation that separates the generation of osk (in an additional tOB.Setup algorithm with input the security parameter) from the tOB.Obfuscate algorithm. Such a formulation would force osk and thus the token computation tOB.Token(osk, x) to be*

independent of the circuit obfuscated; moreover, C could be chosen later, even after all inputs x have been encrypted with tOB.Token .

Our construction satisfies this definition as well because it generates the secret key osk independent of C .

However, we did not choose such a formulation because we wanted to be consistent with the definition of obfuscation, which does not have a separate setup phase.

Intuitively, in a secure token-based obfuscation scheme, an adversary does not learn anything about the circuit C other than $C(x)$ and the size of C .

Definition 5.3 (Secure token-based obfuscation). *Let tOB be a token-based obfuscation scheme for a family of circuits $\mathcal{C} = \{C_n\}_{n \in \mathbb{N}}$. For $A = (A_1, A_2)$ and $S = (S_1, S_2)$, pairs of p.p.t. algorithms, consider the following two experiments:*

$\text{Exp}_{\text{tOB}, A}^{\text{real}}(1^\kappa)$:	$\text{Exp}_{\text{tOB}, A, S}^{\text{ideal}}(1^\kappa)$:
<ol style="list-style-type: none"> 1: $(C, \text{state}_A) \leftarrow A_1(1^\kappa)$ 2: $(\text{osk}, O) \leftarrow \text{tOB.Obfuscate}(1^\kappa, C)$ 3: $\alpha \leftarrow A_2^{\text{tOB.Token}(\text{osk}, \cdot)}(C, O, \text{state}_A)$ 4: Output α 	<ol style="list-style-type: none"> 1: $(C, \text{state}_A) \leftarrow A_1(1^\kappa)$ 2: $(\tilde{O}, \text{state}_S) \leftarrow S_1(1^\kappa, 1^{ C })$ 3: $\alpha \leftarrow A_2^{\text{OS}(\cdot, C)[[\text{state}_S]]}(C, \tilde{O}, \text{state}_A)$ 4: Output α

In the above, $\text{OS}(\cdot, C)[[\text{state}_S]]$ is an oracle that on input x from A_2 , runs S_2 with inputs $C(x)$, $1^{|x|}$, and the current state of S , state_S . S_2 responds with tk_x and a new state state'_S which OS will feed to S_2 on the next call. OS returns tk_x to A_2 .

We say that the token-based obfuscation tOB is secure if there exists a pair of p.p.t. simulators $S = (S_1, S_2)$ such that for all pairs of p.p.t. adversaries $A = (A_1, A_2)$, the following two distributions are computationally indistinguishable:

$$\left\{ \text{Exp}_{\text{tOB}, A}^{\text{real}}(1^\kappa) \right\}_{\kappa \in \mathbb{N}} \stackrel{c}{\approx} \left\{ \text{Exp}_{\text{tOB}, A, S}^{\text{ideal}}(1^\kappa) \right\}_{\kappa \in \mathbb{N}}.$$

Note that, in this security definition, a token tk_x hides x as well because S_2 never receives x . This is usually not required of obfuscation, but we achieve this property for free.

5.2 Scheme

The construction of a token-based obfuscation scheme is very similar to the construction of reusable garbled circuits, the technical difference being minor: we need to specify how to construct the algorithm tOB.Obfuscate from RgB.Garble and RgB.Eval . We construct a token-based obfuscation $\text{tOB} = (\text{tOB.Obfuscate}, \text{tOB.Token})$ as follows based on a reusable garbled scheme $\text{RgB} = (\text{RgB.Garble}, \text{RgB.Enc}, \text{RgB.Eval})$. The token algorithm tOB.Token is the same as RgB.Enc .

Obfuscation $\text{tOB.Obfuscate}(1^\kappa, C \in \mathcal{C}_n)$:

1. Let $(\Gamma, \text{sk}) \leftarrow \text{RgB.Garble}(1^\kappa, C)$.
2. Construct the circuit O (the obfuscation of C) as follows. The circuit O has Γ hardcoded. It takes as input a token tk_x , computes $\text{RgB.Eval}(\Gamma, \text{tk}_x)$, and outputs the result.
3. Output sk as the secret key, and the description of O as the obfuscation of C .

Since the construction is essentially the same as the one of reusable garbled circuits and the security is the same, the same claims and proofs as for reusable garbled circuits hold here, based on Theorem 4.4 and Corollary 4.5. We state them here for completeness.

Claim 5.4. *Assuming a reusable garbling scheme for the class of circuits \mathcal{C} , there is a token-based obfuscation scheme for \mathcal{C} .*

Recall the class of circuits $\mathcal{C}_{n,d(n)}$ defined for Corollary 3.2.

Corollary 5.5 (The LWE Instantiation). *For every integer $n \in \mathbb{N}$, polynomial function $d = d(n)$, there is a token-based obfuscation scheme for the class $\mathcal{C}_{n,d(n)}$, under the following assumption: there is a constant $0 < \epsilon < 1$ such that for every sufficiently large ℓ , the approximate shortest vector problem gapSVP in ℓ dimensions is hard to approximate to within a $2^{O(\ell^\epsilon)}$ factor in time $2^{O(\ell^\epsilon)}$ in the worst case.*

Denote by *universal token-based obfuscation scheme*, a token-based obfuscation scheme for the class of all polynomial-sized circuits. Then,

Corollary 5.6 (Universal token-based obfuscation). *If there is a universal fully secure single-key functional encryption scheme, there is a universal token-based obfuscation scheme.*

6 Computing on Encrypted Data in Input-Specific Time

We initiate the study of fully homomorphic encryption where the runtime of the homomorphic evaluation is input-specific rather than worst-case time. We show how to use our functional encryption scheme to evaluate Turing machines on encrypted data in input-specific time.

Let us recall the setting of computation on encrypted data. A client gives various encrypted inputs and a function f to an evaluator. The evaluator should compute f on the encrypted inputs and return the encrypted result, while learning nothing about the inputs.

Fully homomorphic encryption has been the main tool used in this setting. It was first constructed in a breakthrough work by Gentry [Gen09] and refined in subsequent work [DGHV10, SS10, BV11b, Vai11, BGV12, GHS12a, GHS12b]. Since then, FHE has found many great applications to various problems.

However, one of the main drawbacks of FHE is that when evaluating a Turing machine (TM) over encrypted data, the running time is at least the worst-case running time of the Turing machine over all inputs. The reason is that, one needs to transform the TM into a circuit. If t_{\max} is the maximum running time of the TM on inputs of a certain size—namely, the running time on the worst-case input—then the size of the resulting circuit is at least t_{\max} . Thus, even if the TM runs in a short time on most of the inputs, but for a very long time (t_{\max}) on only one input, *homomorphic evaluation will still run in t_{\max} for all inputs*. This property often results in inefficiency in practice; for example, consider a TM having a loop that depends on the input. For specific inputs, it can loop for a very long time, but for most inputs it does not loop at all.

As a result, researchers have tried to find input-specific schemes. A first observation is that this goal is impossible: input-specific evaluation implies that the evaluator learns the runtime of the TM on each input, which violates CPA-security of the homomorphic scheme (Def. 2.5). Hence, we must relax the security definition and allow the evaluator to learn the runtime for each input, but require that *the evaluator learns nothing else besides the running time*. This goal is not possible with FHE because the evaluator cannot decrypt any bit of information, so it cannot tell whether the computation finished or not; thus, we must look for new solutions.

A second observation is that the evaluator must no longer be able to evaluate TMs of his choice on the client’s data: if he could, the evaluator would run TMs whose running times convey the value of the input x (for example, the evaluator could run $|x|$ TMs, where the i -th TM stops early if the i -th bit of x is zero, and otherwise, it stops later; in this way, the evaluator learns the exact value of x).

Based on these observations, we can see that functional encryption is the natural solution: it hides the inputs to the computation, enables the evaluator to decrypt the running time, and requires the evaluator to obtain a secret key from the client to evaluate each TM.

Due to the impossibility result for functional encryption [AGVW12] discussed in Sec. 1, the client cannot give keys for an arbitrary number of Turing machines to the evaluator. The best we can hope to achieve is for the client to provide a single key for a function to the evaluator (or equivalently, for a constant number q of keys if the client runs the scheme q times). Fortunately, the single-key restriction does not mean that the client can evaluate only one Turing machine. In fact, the client can give a key to the evaluator for a universal Turing machine U that takes as input a TM M and a value x , and outputs $M(x)$. Then, the client must specify together with each input x the TM M he wants to run on x . Such a strategy is even desirable in certain cases: the client may not want the evaluator to compute a TM on every input the client has provided and learn the running time on that input; the client may prefer to specify what inputs to run each Turing machine on.

Using our functional encryption scheme, we achieve a construction that enables computation in input-specific time. We call such a scheme *Turing machine homomorphic encryption*, or shortly TMHE.

As discussed (Corollary 3.2), our functional encryption scheme is succinct in that the ciphertexts grow with the depth of the circuit rather than the size of the circuit. Therefore, our input-specific computation is useful only for Turing machines that can be represented in circuits whose depths are smaller than the running time – because otherwise the client would have to do a lot of work and could instead just run the Turing machine on its own. Moreover, for these machines, we cannot use the Pippenger-Fischer [PF79] transformation because the resulting circuits have depth roughly equal to the running time of the transformed machines. Specifically, our input-specific scheme makes sense for the following class of circuits, with a bound on their depth.

Definition 6.1 (*d -depth-bounded class of Turing machines*). *A finite class of Turing machines \mathcal{M} is d -depth-bounded for a function d , if there exists a class of efficiently computable transformations $\{\mathcal{T}_n\}_{n \in \mathbb{N}}$ with $\mathcal{T}_n : \mathbb{N} \rightarrow \{\text{all circuits}\}$ such that $\mathcal{T}_n(t) = C_{n,t}$ where $C_{n,t}$ is a circuit as follows.*

- *On input a Turing machine $M \in \mathcal{M}$ and a value $x \in \{0, 1\}^n$, $C_{n,t}$ outputs $M(x)$ if M on input x stops in t steps, or \perp otherwise.*
- *The depth of $C_{n,t}$ is at most $d(n)$ and the size of $C_{n,t}$ is $\tilde{O}(t)$.*

Remark 6.2. *Notice that, if we remove the depth constraint (but still keep the circuit size constraint), any finite class of Turing machines satisfies the definition because of the Pippenger-Fischer transformation applied to the universal circuit of this class of Turing machines. Specifically, let U_t be a universal Turing machine that runs any given machine $M \in \mathcal{M}$ for t steps. This machine has $O(t)$ running time and when applying the Pippenger-Fischer transformation [PF79] to it, we get a circuit of size $O(t \log t)$.*

We next present our construction. For completeness, we provide formal definitions and proofs of our theorems and claims in Appendix C. Our security notion (Def. C.2 in the appendix) is called *runtime-CPA security*, which straightforwardly captures the fact that the evaluator should learn nothing about the computation besides the running time.

6.1 Construction

A TMHE scheme consists of four algorithms: $\text{TMHE} = (\text{TMHE.KeyGen}, \text{TMHE.Enc}, \text{TMHE.Eval}, \text{TMHE.Dec})$. The client runs TMHE.KeyGen once in an offline preprocessing stage. Later, in the online phase, the client sends a potentially large number of encrypted inputs to the evaluator. For every input (x, M) consisting of a value x and a Turing machine M , the client runs TMHE.Enc to encrypt the input and then TMHE.Dec to decrypt the result from the evaluator. The evaluator runs TMHE.Eval to evaluate M on x homomorphically in input-specific running time. The work of the client in the offline phase is proportional to t_{\max} , the worst-case input running time. However, for each input in the online phase, the client does little work (independent of the running time of M) and thus the cost is amortized.

We first provide intuition for our construction. As mentioned, we use our functional encryption scheme sFE to enable the evaluator to determine at various intermediary points whether the computation finished or not. For each intermediary step, the client has to provide the evaluator with a function secret key fsk (using the sFE scheme) for a function that returns a bit indicating whether the computation has finished. However, if the client provides a key for every computation step, the offline work of the client becomes quadratic in t_{\max} , which can be very large in certain cases. The idea is to choose intermediary points spaced at exponentially increasing intervals. In this way, the client only generates a logarithmic number of keys, while the evaluator runs in roughly twice the time of M on an input.

As part of TMHE.Enc , besides providing the FE encryptions for a pair (M, x) , the client also provides a homomorphic encryption for x and the machine M , so that once the evaluator learns the running time of M on x , it can then perform the homomorphic computation on x in that running time.

We present our construction for a class of d -depth-bounded Turing machines. By Def. 6.1, such a class has a transformation \mathcal{T}_n that enables transforming a universal TM into a circuit. Let HE be any homomorphic encryption scheme (as defined in Sec. 2.3) for circuits of depth d and let sFE be any functional encryption scheme for circuits of depth d . For simplicity, we present our scheme for Turing machines that output only one bit; we discuss in Sec. 6.3 multiple output bits and how to avoid having the output size be worst-case.

Key generation $\text{TMHE.KeyGen}(1^\kappa, 1^n, 1^{t_{\max}})$ takes as input the security parameter κ , an input size n , and a maximum time bound t_{\max} .

1. Let $\tau = \lceil \log t_{\max} \rceil$. For each $i \in [\tau]$, let $D_i = \mathcal{T}_n(2^i)$ be the circuit that outputs $M(x)$ if M finishes in 2^i steps on input x or \perp otherwise. Construct circuit C_i based on D_i : the circuit C_i , on input a TM M and a value x , outputs 1 if M finished in 2^i steps when running on input x or 0 otherwise; C_i is the same as circuit D_i but it just outputs whether the first output bit of C_i is non- \perp or \perp , respectively.
2. Generate functional encryption secret keys for C_1, \dots, C_τ by running:

$$(\text{fmpk}_i, \text{fmsk}_i) \leftarrow \text{sFE.Setup}(1^\kappa) \quad \text{and} \quad \text{fsk}_i \leftarrow \text{sFE.KeyGen}(\text{fmsk}_i, C_i) \quad \text{for } i \in [\tau].$$

3. Generate HE keys $(\text{hsk}, \text{hpk}) \leftarrow \text{HE.KeyGen}(1^\kappa)$.
4. Output the tuple $\text{PK} := (\text{fmpk}_1, \dots, \text{fmpk}_\tau, \text{hpk})$ as the public key, $\text{EVK} := (\text{fsk}_1, \dots, \text{fsk}_\tau, \text{hpk})$ as the evaluation key, and $\text{SK} := \text{hsk}$ as the secret key.

Encryption $\text{TMHE.Enc}(\text{PK}, (M, x))$: takes as input the public key PK of the form $(\{\text{fmpk}_i\}_i, \text{hpk})$, a TM M and a value x of n bits long.

1. Let $\hat{x} \leftarrow (\text{HE.Enc}(\text{hpk}, x_1), \dots, \text{HE.Enc}(\text{hpk}, x_n))$, where x_i is the i -th bit of x . Similarly, let $\hat{M} \leftarrow (\text{HE.Enc}(\text{hpk}, M_1), \dots, \text{HE.Enc}(\text{hpk}, M_n))$, which is the homomorphic encryption of M (the string description of TM M) bit by bit.
2. Compute $c_i \leftarrow \text{sFE.Enc}(\text{fmpk}_i, (M, x))$ for $i \in [\tau]$.
3. Output the ciphertext $c = (\text{"enc"}, \hat{x}, \hat{M}, c_1, \dots, c_\tau)$.

Evaluation $\text{TMHE.Eval}(\text{EVK}, c)$: takes as input an evaluation key EVK of the form $(\{\text{fsk}_i\}_i, \text{hpk})$ and a ciphertext c of the form $(\text{"enc"}, \hat{x}, \hat{M}, c_1, \dots, c_\tau)$.

1. Start with $i = 1$. Repeat the following:
 - (a) $b \leftarrow \text{sFE.Dec}(\text{fsk}_i, c_i)$.
 - (b) If $b = 1$, (computation finished and we can now evaluate homomorphically on \hat{x})
 - i. Compute D_i , the circuit that evaluates a Turing machine in \mathcal{M} for 2^i steps, using $\mathcal{T}_n(2^i)$.
 - ii. Evaluate and output $(\text{"eval"}, \text{HE.Eval}(\text{hpk}, D_i, (\hat{M}, \hat{x})))$.
 - (c) Else ($b = 0$), proceed to the next i .

Decryption $\text{TMHE.Dec}(\text{SK}, c)$: takes as input a secret key $\text{SK} = \text{hsk}$ and a ciphertext c of the form $(\text{"enc"}, \hat{x}, \hat{M}, c_1, \dots, c_\tau)$ or $(\text{"eval"}, c)$.

1. If the ciphertext is of type "enc" , compute and output $\text{HE.Dec}(\text{hsk}, \hat{x})$.
2. Else (the ciphertext is of type "eval"), compute and output $\text{HE.Dec}(\text{hsk}, c)$.

6.2 Results

We now state our results.

Theorem 6.3. *For any class of d -depth-bounded Turing machines that take n bits of input and produce one bit of output, there is a Turing machine homomorphic encryption scheme, assuming the existence of a fully secure functional encryption scheme sFE for any class of circuits of depth d , and an d -leveled homomorphic encryption scheme HE , where:*

- *The online work of the client is*

$$(n + \log t_{\max}) \cdot \text{poly}(\kappa, d(n))$$

- *The online work of the server in evaluating M on an encryption of x is*

$$\text{poly}(n, d(n), \text{time}(M, x)),$$

where $\text{time}(M, x)$ is the runtime of M on x .

This theorem shows that our TMHE scheme comes as a reduction from any functional encryption scheme. The proof of this theorem is in Appendix C. We can see that the work of the client is indeed smaller than computing the circuit especially if the polynomial d is smaller than the running time. Moreover, we can also see that the server runs in input-specific time: the evaluation time depends on the actual running time and the depth of the circuit.

When instantiating our TMHE construction with our functional encryption sFE construction from Sec. 3 and using Corollary 3.2, we obtain a scheme under an LWE assumption.

Corollary 6.4 (LWE Instantiation). *For every integer $n \in \mathbb{N}$ and polynomial function $d = d(n)$, there is a Turing machine homomorphic encryption scheme for any class of d -depth bounded Turing machines, under the following assumption: there is a constant $0 < \epsilon < 1$ such that for every sufficiently large ℓ , the approximate shortest vector problem gapSVP in ℓ dimensions is hard to approximate to within a $2^{O(\ell^\epsilon)}$ factor in time $2^{O(\ell^\epsilon)}$ in the worst case.*

Remark 6.5. *If the underlying sFE scheme is selectively secure (Def. 2.12), one can still obtain an input-specific homomorphic encryption scheme, but with selective security; namely, the scheme would achieve a modified version of Def. C.2 in Appendix C (the adversary A must choose x before seeing EVK and PK). The scheme would then be secure under the following assumption: there is a constant $0 < \epsilon < 1$ such that for every sufficiently large ℓ , the approximate shortest vector problem gapSVP in ℓ dimensions is hard to approximate to within a $2^{O(\ell^\epsilon)}$ factor in the worst case by polynomial-time adversaries.*

Let us discuss what kind of Turing machines classes are d -depth-bounded.

Fact 6.6. *The class of Turing machines running in log-space is \log^2 -depth-bounded.*

This fact follows directly from the known relation that the LOGSPACE complexity class is in NC2.

In general, the following pattern of computation would fit in d -depth-boundedness and would benefit from input-specific evaluation. Consider a computation that on different types of inputs, it performs different kinds of computation; all these computations are of the same (shallow) depth, but the computation can be much larger in one case.

A few remarks are in order:

Remark 6.7. *Denote by universal TMHE scheme to be a scheme for any finite class of Turing machines. Based on Remark 6.2, we can see that if there is a universal succinct functional encryption scheme and a fully homomorphic scheme, there is a universal TMHE scheme with online client and server work independent of depth:*

- *The online work of the client becomes*

$$(n + \log t_{\max}) \cdot \text{poly}(\kappa)$$

- *The online work of the server in evaluating M on an encryption of x becomes*

$$\text{poly}(n, \text{time}(M, x)),$$

where $\text{time}(M, x)$ is the runtime of M on x .

6.3 Input-Dependent Output Size

The construction above considered Turing machines that output only one bit. To allow TMs that output more than one bit, one can simply use the standard procedure of running one instance of the protocol for each bit of the output. However, as with running time, this would result in repeating the protocol as many times as the worst-case output size for every input. Certain inputs can result in small outputs while others can result in large outputs, so it is desirable to evaluate in *input-specific output size*.

We can use the same approach as above to obtain input-specific output size: The client provides keys to the evaluator to decrypt the size of the output. Then, the evaluator can simply use homomorphic evaluation on a circuit whose output size is the determined one.

Acknowledgments

This work was supported by an NSERC Discovery Grant, by DARPA awards FA8750-11-2-0225 and N66001-10-2-4089, by NSF awards CNS-1053143 and IIS-1065219, and by Google.

The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

References

- [AFV11] Shweta Agrawal, David Mandell Freeman, and Vinod Vaikuntanathan. Functional encryption for inner product predicates from learning with errors. In *ASIACRYPT*, pages 21–40, 2011.
- [AGVW12] Shweta Agrawal, Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption: New perspectives and lower bounds. *Cryptology ePrint Archive, Report 2012/468*, 2012.
- [AKS01] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, pages 601–610, 2001.
- [Ale03] Michael Alekhnovich. More on average case vs approximation complexity. In *FOCS*, pages 298–307, 2003.
- [BCG⁺11] Nir Bitansky, Ran Canetti, Shafi Goldwasser, Shai Halevi, Yael Tauman Kalai, and Guy N. Rothblum. Program obfuscation with leaky hardware. In *ASIACRYPT*, pages 722–739, 2011.
- [BFKL93] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In *CRYPTO*, pages 278–291, 1993.
- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In *CRYPTO*, pages 1–18, 2001.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In *ITCS*, pages 309–325, 2012.
- [BHHI10] Boaz Barak, Iftach Haitner, Dennis Hofheinz, and Yuval Ishai. Bounded key-dependent message security. In *EUROCRYPT*, pages 423–444, 2010.
- [BHR12] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Garbling schemes. *Cryptology ePrint Archive, Report 2012/265*, 2012.
- [BKW03] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, 2003.
- [Bra12] Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In *CRYPTO*, pages 868–886, 2012.
- [BSW07] John Bethencourt, Amit Sahai, and Brent Waters. Ciphertext-policy attribute-based encryption. In *Proceedings of the 28th IEEE Symposium on Security and Privacy*, pages 321–334, 2007.
- [BSW11] Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In *TCC*, pages 253–273, 2011.
- [BV11a] Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ring-LWE and security for key dependent messages. In *CRYPTO*, pages 505–524, 2011.
- [BV11b] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *FOCS*, pages 97–106, 2011.

- [BW07] Dan Boneh and Brent Waters. Conjunctive, subset, and range queries on encrypted data. In *TCC*, pages 535–554, 2007.
- [CKVW10] Ran Canetti, Yael Tauman Kalai, Mayank Varia, and Daniel Wichs. On symmetric encryption and point obfuscation. In *TCC*, pages 52–71, 2010.
- [Dav12] Michael A. Davis. Cloud security: Verify, don’t trust. *Information Week*, August 2012. <http://reports.informationweek.com/abstract/5/8978/>.
- [DGHV10] Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In *EUROCRYPT*, pages 24–43, 2010.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
- [GGH12] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices and applications. Cryptology ePrint Archive, Report 2012/610, 2012.
- [GGP10] Rosario Gennaro, Craig Gentry, and Bryan Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *CRYPTO*, pages 465–482, 2010.
- [GHS12a] Craig Gentry, Shai Halevi, and Nigel P. Smart. Fully homomorphic encryption with polylog overhead. In *EUROCRYPT*, 2012.
- [GHS12b] Craig Gentry, Shai Halevi, and Nigel P. Smart. Homomorphic evaluation of the AES circuit. Cryptology ePrint Archive, Report 2012/099, 2012.
- [GHV10] Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. A simple BGN-type cryptosystem from LWE. In *EUROCRYPT*, pages 506–522, 2010.
- [GIS⁺10] Vipul Goyal, Yuval Ishai, Amit Sahai, Ramarathnam Venkatesan, and Akshay Wadia. Founding cryptography on tamper-proof hardware tokens. In *TCC*, pages 308–326, 2010.
- [GJPS08] Vipul Goyal, Abhishek Jain, Omkant Pandey, and Amit Sahai. Bounded ciphertext policy attribute based encryption. In *ICALP*, pages 579–591, 2008.
- [GK05] Shafi Goldwasser and Yael Tauman Kalai. On the impossibility of obfuscation with auxiliary input. In *FOCS*, pages 553–562, 2005.
- [GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. One-time programs. In *CRYPTO*, pages 39–56, 2008.
- [GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *STOC*, pages 218–229, 1987.
- [GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *Proceedings of the 13th ACM CCS*, pages 89–98, 2006.
- [GR07] Shafi Goldwasser and Guy N. Rothblum. On best-possible obfuscation. In *TCC*, pages 194–213, 2007.
- [GVW12a] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In *CRYPTO*, pages 162–179, August 2012.
- [GVW12b] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Predicate encryption for circuits: How to filter computation. In submission to *STOC 2013*, 2012.
- [KSW08] Jonathan Katz, Amit Sahai, and Brent Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *EUROCRYPT*, pages 146–162, 2008.
- [LOS⁺10] 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 *EUROCRYPT*, pages 62–91, 2010.

- [LP09] Yehuda Lindell and Benny Pinkas. A proof of security of Yao’s protocol for two-party computation. *J. Cryptol.*, 22:161–188, April 2009.
- [LTV12] Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *STOC*, pages 1219–1234, 2012.
- [LW12] Allison B. Lewko and Brent Waters. New proof methods for attribute-based encryption: Achieving full security through selective techniques. In *CRYPTO*, 2012.
- [MV10] Daniele Micciancio and Panagiotis Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. In *STOC*, pages 351–358, 2010.
- [O’N10] Adam O’Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010.
- [OT09] Tatsuaki Okamoto and Katsuyuki Takashima. Hierarchical predicate encryption for inner-products. In *ASIACRYPT*, pages 214–231, 2009.
- [OT10] Tatsuaki Okamoto and Katsuyuki Takashima. Fully secure functional encryption with general relations from the decisional linear assumption. In *CRYPTO*, pages 191–208, 2010.
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342, 2009.
- [PF79] Nicholas Pippenger and Michael J. Fischer. Relations among complexity measures. *J. ACM*, 26(2):361–381, 1979.
- [Pri12] Privacy Rights Clearinghouse. Chronology of data breaches. <http://www.privacyrights.org/data-breach>, 2012.
- [PRV12] Bryan Parno, Mariana Raykova, and Vinod Vaikuntanathan. How to delegate and verify in public: Verifiable computation from attribute-based encryption. In *TCC*, pages 422–439, 2012.
- [RAD78] R. Rivest, L. Adleman, and M. Dertouzos. On data banks and privacy homomorphisms. In *Foundations of Secure Computation*, pages 169–177. Academic Press, 1978.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005.
- [SS10] Damien Stehlé and Ron Steinfeld. Faster fully homomorphic encryption. In *ASIACRYPT*, pages 377–394, 2010.
- [SSW09] Emily Shen, Elaine Shi, and Brent Waters. Predicate privacy in encryption systems. In *TCC*, pages 457–473, 2009.
- [SV11] Nigel P. Smart and Frederik Vercauteren. Fully homomorphic SIMD operations. Cryptology ePrint Archive, Report 2011/133, 2011.
- [SW05] Amit Sahai and Brent Waters. Fuzzy identity-based encryption. In *EUROCRYPT*, pages 457–473, 2005.
- [SW12] Amit Sahai and Brent Waters. Attribute-based encryption for circuits from multilinear maps. Cryptology ePrint Archive, Report 2012/592, 2012.
- [Vai11] Vinod Vaikuntanathan. Computing blindfolded: New developments in fully homomorphic encryption. In *FOCS*, pages 5–16, 2011.
- [Ver12] Verizon RISK Team. 2012 data breach investigations report. http://www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf, 2012.
- [Wat11] Brent Waters. Ciphertext-policy attribute-based encryption: an expressive, efficient, and provably secure realization. In *PKC*, pages 53–70, 2011.

- [Wat12] Brent Waters. Functional encryption for regular languages. In *CRYPTO*, pages 218–235, 2012.
- [Wee05] Hoeteck Wee. On obfuscating point functions. In *STOC*, pages 523–532, 2005.
- [Yao82] Andrew C. Yao. Protocols for secure computations. In *FOCS*, pages 160–164, 1982.
- [Yao86] Andrew C. Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167, 1986.

A Detailed Background on Learning With Errors (LWE)

The LWE problem was introduced by Regev [Reg05] as a generalization of “learning parity with noise” [BFKL93, BKW03, Ale03]. Regev showed that solving the LWE problem *on the average* is as hard as (quantumly) solving several standard lattice problems *in the worst case*. This result bolstered our confidence in the LWE assumption and generated a large body of work building cryptographic schemes under the assumption, culminating in the construction of a fully homomorphic encryption scheme [BV11b].

For positive integers ℓ and $q \geq 2$, a vector $\mathbf{s} \in \mathbb{Z}_q^\ell$, and a probability distribution χ on \mathbb{Z}_q , let $A_{\mathbf{s},\chi}$ be the distribution obtained by choosing a vector $\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^\ell$ uniformly at random and a noise term $e \xleftarrow{\$} \chi$, and outputting $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in \mathbb{Z}_q^\ell \times \mathbb{Z}_q$. A formal definition follows.

Definition A.1 (LWE). *For an integer $q = q(\ell)$ and an error distribution $\chi = \chi(\ell)$ over \mathbb{Z}_q , the learning with errors problem $\text{LWE}_{\ell,m,q,\chi}$ is defined as follows: Given m independent samples from $A_{\mathbf{s},\chi}$ (for some $\mathbf{s} \in \mathbb{Z}_q^\ell$), output \mathbf{s} with noticeable probability.*

The (average-case) decision variant of the LWE problem, denoted $\text{dLWE}_{\ell,m,q,\chi}$, is to distinguish (with non-negligible advantage) m samples chosen according to $A_{\mathbf{s},\chi}$ (for uniformly random $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^\ell$), from m samples chosen according to the uniform distribution over $\mathbb{Z}_q^\ell \times \mathbb{Z}_q$.

We denote by $\text{LWE}_{\ell,q,\chi}$ (resp. $\text{dLWE}_{\ell,q,\chi}$) the variant where the adversary gets oracle access to $A_{\mathbf{s},\chi}$, and is not a-priori bounded in the number of samples.

For cryptographic applications we are primarily interested in the average case decision problem dLWE , where $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^\ell$. We will also be interested in assumptions of the form: no t -time adversary can solve dLWE with non-negligible advantage, which we will call the t -hardness of dLWE .

There are known quantum [Reg05] and classical [Pei09] reductions between $\text{dLWE}_{\ell,m,q,\chi}$ and approximating short vector problems in lattices. Specifically, these reductions take χ to be (discretized versions of) the Gaussian distribution, which is B -bounded for an appropriate B . Since the exact distribution χ does not matter for our results, we state a corollary of the results of [Reg05, Pei09] in terms of the bound on the distribution.

Let $B = B(\ell) \in \mathbb{N}$. A family of distributions $\chi = \{\chi_\ell\}_{\ell \in \mathbb{N}}$ is called B -bounded if the support of χ_ℓ is (a subset of) $[-B(\ell), \dots, B(\ell)]$. Then:

Lemma A.2 ([Reg05, Pei09]). *Let $q = q(\ell) \in \mathbb{N}$ be a product of co-prime numbers $q = \prod q_i$ such that for all i , $q_i = \text{poly}(\ell)$, and let $B \geq \ell$. Then there exists an efficiently sampleable B -bounded distribution χ such that if there is an efficient algorithm that solves the (average-case) $\text{dLWE}_{\ell,q,\chi}$ problem. Then:*

- *There is a quantum algorithm that solves SIVP with approximation factor $\tilde{O}(\ell\sqrt{\ell} \cdot q/B)$ and gapSVP with approximation factor $\tilde{O}(\ell\sqrt{\ell} \cdot q/B)$ on any ℓ -dimensional lattice, and runs in time $\text{poly}(\ell)$.*
- *There is a classical algorithm that solves the ζ -to- γ decisional shortest vector problem $\text{gapSVP}_{\zeta,\gamma}$, where $\gamma = \tilde{O}(\ell\sqrt{\ell} \cdot q/B)$, and $\zeta = \tilde{O}(q\sqrt{\ell})$, on any ℓ -dimensional lattice, and runs in time $\text{poly}(\ell)$.*

We remark that this connection is time-preserving, in the sense that given an LWE algorithm that runs in time t , these reductions produce algorithms to solve lattice problems that run in time $\text{poly}(t)$.

We refer the reader to [Reg05, Pei09] for the formal definition of these lattice problems, as they have no direct connection to this work. We only note here that the best known algorithms for these problems run in time nearly exponential in the dimension ℓ [AKS01, MV10]. More generally, the best algorithms that approximate these problems to within a factor of 2^k run in time $2^{\tilde{O}(\ell/k)}$. Specifically, given the current state of the art on lattice algorithms, the $\text{LWE}_{\ell, q, \chi}$ assumption is quite plausible for a $\text{poly}(\ell)$ -bounded distribution χ and q as large as 2^{ℓ^ϵ} (for any constant $0 < \epsilon < 1$).

Given this state of affairs, we will abuse notation slightly and conflate the LWE dimension ℓ with the security parameter κ .

B Construction of Two-Outcome Public-Index Functional Encryption

Let us construct a two-outcome public-index functional encryption scheme, denoted pFE_2 , from a public-index FE scheme, pFE .

The idea is to use two pFE instantiations, one encrypting M_0 and the other M_1 . To make sure that exactly one of these messages gets revealed when a predicate is evaluated, we provide secret keys for the predicate and the negation of the predicate for the two instantiations.

Setup $\text{pFE}_2.\text{Setup}(1^\kappa)$:

1. Run $(\text{fmsk}_0, \text{fmpk}_0) \leftarrow \text{pFE}.\text{Setup}(1^\kappa)$ and $(\text{fmsk}_1, \text{fmpk}_1) \leftarrow \text{pFE}.\text{Setup}(1^\kappa)$.
2. Let $\text{fmsk} := (\text{fmsk}_0, \text{fmsk}_1)$ and $\text{fmpk} := (\text{fmpk}_0, \text{fmpk}_1)$. Output fmsk and fmpk .

Key generation $\text{pFE}_2.\text{KeyGen}(\text{fmsk}, P)$: Let $\text{fsk}_0 \leftarrow \text{pFE}.\text{KeyGen}(\text{fmsk}_0, \bar{P})$ and $\text{fsk}_1 \leftarrow \text{pFE}.\text{KeyGen}(\text{fmsk}_1, P)$, where \bar{P} is the negation of P , namely $\bar{P}(x) = 1 - P(x)$. Output $\text{fsk}_P = (\text{fsk}_0, \text{fsk}_1)$.

Encryption $\text{pFE}_2.\text{Enc}(\text{fmpk}, (x, M_0, M_1))$: Let $C_0 \leftarrow \text{pFE}.\text{Enc}(\text{fmpk}_0, (x, M_0))$ and $C_1 \leftarrow \text{pFE}.\text{Enc}(\text{fmpk}_1, (x, M_1))$. Output $C = (C_0, C_1)$.

Decryption $\text{pFE}_2.\text{Dec}(\text{fsk}_P, C)$:

1. Parse $\text{fsk}_P = (\text{fsk}_0, \text{fsk}_1)$ and $C = (C_0, C_1)$.
2. Run $M_0 \leftarrow \text{pFE}.\text{Dec}(\text{fsk}_0, C_0)$ and if $M_0 \neq \perp$, output M_0 .
3. Run $M_1 \leftarrow \text{pFE}.\text{Dec}(\text{fsk}_1, C_1)$ and if $M_1 \neq \perp$, output M_1 .

We next prove that this construction yields a secure two-outcome public-index FE scheme. Note that our construction requires a public-index FE scheme where the predicate class \mathcal{P}_n is closed under negation: for every predicate $P \in \mathcal{P}_n$, the predicate \bar{P} is also included in \mathcal{P}_n .

Proof of Claim 2.18. Correctness of pFE_2 is straightforward: If $P(x) = 0$, C_0 will decrypt to M_0 by the correctness of pFE , and mutatis mutandis for $P(x) = 1$.

We prove security by contradiction. Assume there exists p.p.t. $A = (A_1, A_2, A_3)$ that breaks the security of our pFE_2 construction: Def. 2.17; namely, there exists a polynomial p such that, for infinitely many κ ,

$$\Pr[\text{Exp}_{\text{pFE}_2, A}(1^\kappa) = 1] \geq 1/2 + 1/p(\kappa). \quad (6)$$

We construct a p.p.t. adversary $R = (R_1, R_2, R_3)$ that breaks the security of pFE, Def. 2.14.

The adversary R_1 receives as input fmpk^* and outputs a predicate P^* as follows. The adversary A_1 expects two public keys. R_1 uses fmpk^* as one of these public keys and generates the other public key freshly $(\text{fmsk}, \text{fmpk}) \leftarrow \text{pFE.KeyGen}(1^\kappa)$. The order in which R_1 provides these keys to A_1 depends on the value of $P(x)$ not known at this step. If $P(x)$ will be 0, R will have to give A the ability to decrypt a ciphertext encrypted with the first key. If that key is fmpk^* , R cannot accomplish this task because it does not have the corresponding secret key. Therefore, R will try to guess $P(x)$ by flipping a random coin. Concretely, R_1 runs:

1. Guess $P(x)$ at random, namely draw a random bit denoted guess . If guess is 0:
 - (a) Provide $(\text{fmpk}, \text{fmpk}^*)$ to A_1 .
 - (b) Receive P from A_1 and output $P^* := P$.
2. Else [guess is 1]:
 - (a) Provide $(\text{fmpk}^*, \text{fmpk})$ to A_1 .
 - (b) Receive P from A_1 and output $P^* := \bar{P}$.

Adversary R_2 receives as input fsk_{P^*} and generates M_0^*, M_1^* , and x^* as follows.

1. Generate $\text{fsk}_{\bar{P}^*} \leftarrow \text{pFE.KeyGen}(\text{fmsk}, \bar{P}^*)$.
2. If guess is 0, provide $(\text{fsk}_{\bar{P}^*}, \text{fsk}_{P^*})$ to A_2 , else (guess was 1) provide $(\text{fsk}_{P^*}, \text{fsk}_{\bar{P}^*})$ to A_2 .
3. Receive (M, M_0, M_1, x) from A_2 . Output $M_0^* := M_0$, $M_1^* := M_1$ and $x^* := x$.

Adversary R_3 receives as input c^* and outputs a guess bit as follows:

1. Check that $P(x)$ equals guess . If this is not the case, namely, R_1 had guessed incorrectly the value of $P(x)$, output a random bit and exit. Otherwise, continue.
2. Feed the following input to A_3 : if $\text{guess} = 0$, feed inputs $(\text{pFE.Enc}(\text{fmpk}, (x, M)), c^*)$, else ($\text{guess} = 1$), feed inputs $(c^*, \text{pFE.Enc}(\text{fmpk}, (x, M)))$. Output whatever A_3 outputs.

R_1 guesses $P(x)$ correctly with a chance of half. When R_1 does not guess $P(x)$ correctly, R_3 outputs a correct bit with chance $1/2$ (because it outputs a random guess). When R_1 guesses $P(x)$ correctly, we can see that R simulates the pFE_2 game with A correctly. Therefore, in this case, whenever A guesses correctly, R also guesses correctly. Using Eq. (6), we have

$$\Pr[\text{Exp}_{\text{pFE}, R}(1^\kappa) = 1] \geq 1/2 \cdot 1/2 + 1/2(1/2 + 1/2p(\kappa)) = 1/2 + 1/2p(\kappa), \quad (7)$$

which provides the desired contradiction. \square

C Homomorphic Encryption for Turing Machines: Definitions and Proofs

Let us first define the syntax of a Turing machine homomorphic encryption scheme.

Definition C.1. A Turing machine homomorphic encryption scheme TMHE for a class of Turing machines \mathcal{M} is a quadruple of p.p.t. algorithms $(\text{TMHE.KeyGen}, \text{TMHE.Enc}, \text{TMHE.Dec}, \text{TMHE.Eval})$ as follows:

- $\text{TMHE.KeyGen}(1^\kappa, 1^n, 1^{t_{\max}})$ takes as input a security parameter κ , an input size n , and a time bound t_{\max} , and outputs a public key PK , an evaluation key EVK , and a secret key SK .
- $\text{TMHE.Enc}(\text{PK}, M, x)$ takes as input the public key PK , a Turing machine M with one bit of output, and an input $x \in \{0, 1\}^n$, for some n , and outputs a ciphertext c .
- $\text{TMHE.Dec}(\text{SK}, c)$ takes as input the secret key SK and a ciphertext c , and outputs a message x .
- $\text{TMHE.Eval}(\text{EVK}, c)$ takes as input the evaluation key EVK , and a ciphertext c , and outputs a ciphertext c' .

Correctness: For every polynomial $n(\cdot)$, for every polynomial $t_{\max}(\cdot)$, for every sufficiently large security parameter κ , for $n = n(\kappa)$, for every Turing machine $M \in \mathcal{M}$ with upper bound on running time for inputs of size n of $t_{\max}(n)$, and for every input $x \in \{0, 1\}^n$,

$$\begin{aligned} \text{Pr}[(\text{PK}, \text{EVK}, \text{SK}) \leftarrow \text{TMHE.KeyGen}(1^\kappa, 1^n, 1^{t_{\max}(n)}); \\ c \leftarrow \text{TMHE.Enc}(\text{PK}, M, x); \\ c^* \leftarrow \text{TMHE.Eval}(\text{EVK}, M, c) : \\ \text{TMHE.Dec}(\text{SK}, c^*) \neq M(x)] = \text{negl}(\kappa). \end{aligned}$$

Note that the correctness property constraints t_{\max} to be a polynomial. However, t_{\max} can still be a very large polynomial and we would like the server to not have to run in that time for all inputs. (In fact, this constraint can be eliminated if we use a FHE scheme and a public-index scheme that have no correctness error).

Definition C.2 (Runtime-CPA Security). Let TMHE be an input-specific homomorphic encryption scheme for the class of Turing machines \mathcal{M} . For every p.p.t. adversary $A = (A_1, A_2)$ and p.p.t. simulator S , consider the following two experiments:

$\underline{\text{Exp}_{\text{TMHE}, A}^{\text{real}}(1^\kappa) :}$ <ol style="list-style-type: none"> 1: $(1^{t_{\max}}, 1^n, \text{state}_A) \leftarrow A_1(1^\kappa)$. 2: $(\text{PK}, \text{EVK}, \text{SK}) \leftarrow \text{TMHE.KeyGen}(1^\kappa, 1^n, 1^{t_{\max}})$ 3: $(M, x, \text{state}'_A) \leftarrow A_2(\text{state}_A, \text{PK}, \text{EVK})$ 4: $c \leftarrow \text{TMHE.Enc}(\text{PK}, M, x)$ 5: Output (state'_A, c) 	$\underline{\text{Exp}_{\text{TMHE}, A, S}^{\text{ideal}}(1^\kappa) :}$ <ol style="list-style-type: none"> 4: $\tilde{c} \leftarrow S(M, 1^n, 1^t, \text{EVK}, \text{PK})$ with $t = \text{time}(M, x)$ 5: Output $(\text{state}'_A, \tilde{c})$
---	---

The scheme is said to be runtime-CPA-secure if there exists a p.p.t. simulator S such that for all pairs of p.p.t. adversaries $A = (A_1, A_2)$ for which A_2 outputs $M \in \mathcal{M}$ and $x \in \{0, 1\}^n$, we have

$$\left\{ \text{Exp}_{\text{TMHE}, A}^{\text{real}}(1^\kappa) \right\}_{\kappa \in \mathbb{N}} \stackrel{c}{\approx} \left\{ \text{Exp}_{\text{TMHE}, A, S}^{\text{ideal}}(1^\kappa) \right\}_{\kappa \in \mathbb{N}} .$$

This definition essentially captures our security goal: one can simulate any information learned from the scheme by using only the Turing machine M and the running time of M on x , but without any other information about x .

In fact, we can achieve a scheme that hides M as well in a straightforward way: since our construction passes M and x as inputs to universal circuits, M could also be hidden in the same way as x is.

C.1 Proof

Proof of Theorem 6.3. We first prove the correctness and efficiency claims of the theorem and then we prove security.

If the underlying sFE scheme is correct, then TMHE is correct; whenever 2^i for some i is an upper bound on the running time of M on x , then $C_i(M, x)$'s output is 1. Based on the correctness of the FHE scheme HE, the evaluation of D_i on \hat{M}, \hat{x} will be correct, so HE.Dec will return $M(x)$.

Lemma C.3. *The online work of the client in the TMHE scheme is $(\log t_{\max} + n) \cdot \text{poly}(\kappa, d(n))$.*

Proof. The work of the client in the online phase consists of running TMHE.Enc(PK, x) and TMHE.Dec(SK, c). The work of the client for TMHE.Enc(PK, x) is $n \text{poly}(d(\kappa))$ to compute the HE ciphertexts and $(1 + \lceil \log t_{\max} \rceil) \cdot \text{poly}(d(n), \kappa)$ to compute the sFE ciphertexts. Since n depends polynomially in κ , we obtain that total cost is at most $(\log t_{\max} + n) \text{poly}(\kappa, d(n))$ (where we incorporated the constant values in the poly notation).

The runtime of TMHE.Dec(SK, c) is $\text{poly}(d(\kappa))$ because HE.Enc runs polynomial in κ and $d(\kappa)$. Therefore, the total online work of the client is $(\log t_{\max} + n) \text{poly}(d(\kappa), d(n), \kappa)$. \square

Lemma C.4. *The work of the evaluator in the TMHE scheme is $\text{poly}(n, d(n), \text{time}(M, x))$.*

Proof. The work of the evaluator consists of running TMHE.Eval(EVK, M, c). This depends on the number of times the loop in TMHE.Eval is repeated and the cost within each loop. Let us evaluate the cost at the i -th repetition of the loop. Let $t_i = 2^i$.

By the properties of the transformation \mathcal{T}_n , the size of C_i is at most $t_i \text{polylog } t_i$. The cost of evaluating sFE.Dec(fsk_i, c_i) is therefore $\text{poly}(n, d(n), t_i \text{polylog } t_i)$.

If t is the runtime of M on x , the index i at which the loop will halt (because the evaluator obtained a value the bit b being one) is at most $1 + \lceil \log t \rceil$. Therefore, the loop will repeat at most $1 + \lceil \log t \rceil$ times.

$$\begin{aligned} \text{Runtime of TMHE.Eval(EVK, } c) &= \sum_{i=1}^{1+\lceil \log t \rceil} \text{poly}(n, d(n), t_i \text{polylog } t_i) \\ &\leq (1 + \lceil \log t \rceil) \text{poly}(n, d(n), t \text{polylog } t) \\ &\leq \text{poly}(n, d(n), t \text{polylog } t) = \text{poly}(n, d(n), t), \end{aligned}$$

where the last equality comes from adjusting the implicit polynomial in poly. Note that even though EVK consists of $\log t_{\max}$ such fsk_i keys, TMHE.Eval does not have to read all of EVK. \square

Finally, we prove security of the scheme.

Lemma C.5. *The TMHE protocol is runtime-CPA-secure.*

Proof. To prove that our TMHE construction is secure, we provide a simulator S , as in Def. C.2. The simulator S invokes the simulator of the functional encryption scheme, as in Def. 2.11, which we denote Sim_{sFE} . The simulator S receives inputs $M, 1^n, 1^t, \text{EVK}$, and PK , and proceeds as follows:

1. Compute $\hat{0}^n \leftarrow (\text{HE.Enc}(\text{hpk}, 0), \dots, \text{HE.Enc}(\text{hpk}, 0))$ (n times).
2. For each $i \in [\tau]$, compute the circuits $D_i = T_n(2^i)$ and then C_i as before; we remind the reader that C_i , on input a TM M and a value x , outputs 1 if M finished in 2^i steps when running on input x or 0 otherwise.
3. For each i such that $2^i < t$:
 - (a) Call the simulator Sim_{sFE} to simulate a computation result of 0 because M could not have finished its computation at step i . Specifically, compute $\tilde{c}_i \leftarrow \text{Sim}_{\text{sFE}}(\text{fmpk}_i, \text{fsk}_i, C_i, 0, 1^{n+|M|})$.
4. For each i such that $2^i \geq t$:
 - (a) Call the simulator Sim_{sFE} to simulate an answer of 1 because M finished computation on the input (unknown to S). Thus, compute $\tilde{c}_i \leftarrow \text{Sim}_{\text{sFE}}(\text{fmpk}_i, \text{fsk}_i, C_i, 1, 1^{n+|M|})$.
5. Output $\tilde{c} = (\hat{0}, \tilde{c}_1, \dots, \tilde{c}_\tau)$.

To prove that S satisfies Def. C.2, we use three hybrids:

Hybrid 0: The ideal experiment with simulator S .

Hybrid 1: The same as Hybrid 0 but $\hat{0}^n$ gets replaced with $\hat{x} = (\text{HE.Enc}(\text{hpk}, x_1), \dots, \text{HE.Enc}(\text{hpk}, x_n))$.

Hybrid 2: The real experiment.

It is easy to see that the outcome of Hybrid 0 and the outcome of Hybrid 1 are computationally indistinguishable because HE is semantically secure: the encryptions of 0^n in Hybrid 0 and the encryption of x in Hybrid 1 are both generated with fresh randomness, and the secret key hsk (or any function of hsk other than a fresh encryption) is never released to any adversary.

Now let us look at Hybrid 1 and Hybrid 2. These are computationally indistinguishable based on a standard hybrid argument invoking the security of Sim_{sFE} as follows.

The simulator Sim_{sFE} is called τ times. Let $\tilde{c}_i^{(1)}$ be the ciphertext output by the simulator for the i -th invocation in Hybrid 1, and let c_i be the ciphertext output in Hybrid 2 on the i -th invocation. It is enough to prove that the outcome of these two experiments consisting of state $'_A$ and only one of the ciphertexts (e.g., $\tilde{c}_i^{(1)}$ or c_i) are computationally indistinguishable. The reason is that one can employ a standard hybrid argument consisting of $\tau + 1$ sub-hybrids, the 0-th sub-hybrid being Hybrid 1 and the τ -th sub-hybrid being Hybrid 2 and any intermediary sub-hybrid i has the first i ciphertexts as in Hybrid 2 and the rest as in Hybrid 1. Such an argument is possible because τ is polynomial in the security parameter and each ciphertext is encrypted with independently generated public keys.

Therefore, all we need to argue is that the outcome of Hybrid 1 and Hybrid 2 consisting of state $'_A$ and only $\tilde{c}_i^{(1)}$ (c_i respectively) are computationally indistinguishable. This follows directly because Sim_{sFE} satisfies the FULL-SIM-secure functional encryption definition, Def. 2.11. \square

The three lemmas above complete the proof of the theorem. \square