

Functional Encryption: New Perspectives and Lower Bounds

Shweta Agrawal^{*1}, Sergey Gorbunov^{**2}, Vinod Vaikuntanathan^{***2}, and Hoeteck Wee^{†3}

¹ University of California, Los Angeles

² University of Toronto

³ George Washington University

Abstract. Functional encryption is an emerging paradigm for public-key encryption that enables fine-grained control of access to encrypted data. In this work, we present new lower bounds and impossibility results on functional encryption, as well as new perspectives on security definitions. Our main contributions are as follows:

- We show that functional encryption schemes that satisfy even a weak (non-adaptive) simulation-based security notion are impossible to construct in general. This is the *first* impossibility result that exploits *unbounded* collusions in an essential way. In particular, we show that there are no such functional encryption schemes for the class of weak pseudo-random functions (and more generally, for any class of incompressible functions). More quantitatively, our technique also gives us a lower bound for functional encryption schemes secure against *bounded* collusions. To be secure against q collusions, we show that the ciphertext in any such scheme must have size $\Omega(q)$.
- We put forth and discuss a simulation-based notion of security for functional encryption, with an unbounded simulator (called USIM). We show that this notion interpolates indistinguishability and simulation-based security notions, and is inspired by results and barriers in the zero-knowledge and multi-party computation literature.

* Research supported in part from a DARPA/ONR PROCEED award, NSF grants 1136174, 1118096, 1065276, 0916574 and 0830803, a Xerox Faculty Research Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant (Amit Sahai). This material is based upon work supported by the Defense Advanced Research Projects Agency through the U.S. Office of Naval Research under Contract N00014-11-1-0389. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

** Supported by Ontario Graduate Scholarship (OGS).

*** Supported by an NSERC Discovery Grant and by DARPA under Agreement number FA8750-11-2-0225. 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.

† Supported by NSF CAREER Award CNS-1237429.

1 Introduction

Functional encryption [SW05,SW08] is a new paradigm for public-key encryption that enables fine-grained control of access to encrypted data. It extends several previous notions, most notably identity-based encryption [Sha84,BF01,Coc01], and provides, for instance, the ability to generate and release secret keys associated with a keyword that can decrypt only those documents that contain the keyword. More generally, functional encryption allows the owner of a “master” secret key to release restricted secret keys that reveal a specific function of encrypted data. This stands in stark contrast to traditional encryption, where access to the encrypted data is all or nothing: namely, given the secret key, one can decrypt and read the entire plaintext, but without it, nothing about the plaintext is revealed at all (other than its length).

Functional Encryption. A functional encryption scheme for a circuit family [BSW11,O’N10] \mathcal{C} , associates secret keys SK_C with every circuit $C \in \mathcal{C}$ and ciphertext CT with input messages x .⁴

In broad terms, functional encryption requires that the owner of a secret key SK_C and a ciphertext CT (corresponding to an input message x) be able to compute $C(x)$, but learn nothing else about x itself. (Typically, and throughout this work, we assume that the circuit family \mathcal{C} as well as the circuit queries C are public, in the sense that they are not hidden from the key holders.)

Moreover, security should hold in the presence of collusions amongst “key holders”, that is, malicious users should not be able to combine their secret keys to learn unauthorized information. More formally, a collusion of users that hold secret keys $\text{SK}_{C_1}, \dots, \text{SK}_{C_q}$ and an encryption of x should learn nothing else about x apart from $C_1(x), \dots, C_q(x)$, for any polynomial q .

An important subclass of functional encryption is that of public-index predicate encryption. Here, the input x is a pair (ind, μ) where ind is an index and μ the payload message. Let P be a Boolean predicate defined on indices, the circuit family \mathcal{C} is given by:

$$C_P(\text{ind}, \mu) = \begin{cases} (\text{ind}, \mu) & \text{if } P(\text{ind}) = 1 \\ (\text{ind}, \perp) & \text{otherwise} \end{cases}$$

Even though public index predicate encryption seems like a weak object, it already captures identity-based encryption, and is also very useful in constructing protocols for verifiably delegating computation as shown recently by Parno, Raykova and Vaikuntanathan [PRV12].

Predicate encryption captures and generalizes a large number of previous constructions, including identity-based encryption (IBE) [Sha84,BF01,Coc01,BW06], fuzzy IBE [SW05,ABV⁺12], attribute-based encryption (ABE) [GPSW06,LOS⁺10],

⁴ An alternative approach is associate secret keys to inputs and ciphertexts to circuits. This is equivalent to our approach by taking a new “universal” family U_x that on input C outputs $C(x)$.

and inner product encryption [KSW08,LOS+10,AFV11]. Specifically, IBE corresponds to P encoding a point function. Moreover, essentially all known constructions are examples of public-index predicate encryption schemes or its variants, with a few exceptions – constructions in [BF01,BW06,KSW08] achieve a stronger private-index security notion in which the index ind also remains hidden from the adversary.

Security Notions. Boneh, Sahai and Waters [BSW11] and O’Neill [O’N10] were the first to put forth a general definitional framework for functional encryption. They considered two security notions for functional encryption, namely: *indistinguishability* (IND) based security and *simulation* (SIM) based security. The former stipulates that it is infeasible to distinguish encryptions of any two messages, without getting a secret key that decrypts the ciphertexts to distinct values; the latter stipulates the existence of an efficient simulator that given $C_1(x), \dots, C_q(x)$, outputs the view of the colluders that are given an encryption of x as well as secret keys $\text{SK}_{C_1}, \dots, \text{SK}_{C_q}$.

Both of these notions may be further refined in two ways:

- *adaptive* (AD) versus *non-adaptive* (NA) which capture whether the adversary’s queries to the key derivation oracle may or may not depend on the challenge ciphertext; and
- *one* versus *many*, referring to whether the adversary receives a single or multiple challenge ciphertexts.

Together, these give rise to eight security notions xx-yy-zzz , where $\text{xx} \in \{1, \text{many}\}$, $\text{yy} \in \{\text{NA}, \text{AD}\}$, and $\text{zzz} \in \{\text{IND}, \text{SIM}\}$.

Recent work. We briefly outline the known relationships amongst these eight notions. We note that in general, indistinguishability based security provides a weaker guarantee than simulation based security (that is, xx-yy-SIM implies xx-yy-IND and xx-yy-IND does not imply xx-yy-SIM in general); on the other hand, we have that 1-yy-IND implies many-yy-IND . Boneh, et al. [BSW11] pointed out that indistinguishability based security is vacuous and inadequate for certain circuit families, which indicate that we should opt for simulation-based security whenever possible.⁵ O’Neill [O’N10] showed that NA-IND and NA-SIM are equivalent for some subclass of circuit families that are roughly speaking, “easy to invert”.

All prior positive results achieve many-AD-IND security or relaxations thereof.⁶ The only known impossibility result we have for general functional encryption is that of Boneh et al. [BSW11] for realizing the IBE functionality under many-AD-SIM security. In particular, in light of known results, it is entirely

⁵ [BSW11, Section 5.3] presents an “equivalence” between many-AD-IND and many-AD-SIM in the *programmable* random oracle model for public-index predicate encryption. For this work, we consider only the standard model.

⁶ A commonly used relaxation of AD-IND security for predicate encryption is that of “selective security” [CHK03].

	realizable for public-index	realizable for all circuits
xx-yy-IND	[GVW13,GGH ⁺ 13] ⁷	open
xx-yy-SIM (xx = 1 OR yy = NA)	open	no (Section 4)
many-AD-SIM	no [BSW11]	no ←
xx-yy-USIM (xx = 1 OR yy = NA)	open	open
many-AD-USIM	no [BSW11] ‡	no ←

Fig. 1. Summary of results and open problems. Results from this work are marked with boldface. Results implicit in previous works are marked with ‡. Results that are trivially implied by results in a previous column are marked with ←. The second and third columns indicate whether the definition is realizable for all public-index predicate encryption schemes (e.g. IBE) and for all circuits respectively. USIM refers to the notion of unbounded simulation discussed in Section 1.2.

conceivable that we can realize functional encryption for all poly-size circuits under either 1-AD-SIM security (thus 1-AD-IND and many-AD-IND security) or many-NA-SIM security.

In this work, we narrow the gap between existing security definitions for functional encryption, as well as that between existing constructions and impossibility results. Our results are as follows.

1.1 New Lower Bound: Impossibility for Simulation-based Definitions

Our main result rules out general functional encryption under the one message secure, non-adaptive simulation definition (1-NA-SIM). In particular, this rules out both of the scenarios presented at the end of the preceding section (i.e. 1-AD-SIM or many-NA-SIM for all circuits) in a strong sense. This is the *first* lower bound that exploits *unbounded* collusions in an essential way. We compare the impossibility result from [BSW11] with ours in the full version.

Theorem 1 (Informal). *There exists a circuit family \mathcal{C} for which there is no 1-NA-SIM-secure functional encryption scheme.*

Specifically, assuming the existence of a family of weak pseudo-random function $w\text{PRF}(\cdot, \cdot)$ (See Definition 3) that outputs one bit, we show that there does not exist a functional encryption scheme for the family:

$$C_d(x) = w\text{PRF}(x, d), \text{ where the input message } x \text{ is the PRF seed}$$

We show that the ciphertext size in a 1-NA-SIM-secure scheme realizing this circuit family must grow with the size of the collusion; this yields a contradiction,

since the scheme must handle unbounded collusions. In fact, the result is unconditional since any non-trivial functional encryption scheme gives rise to a one-way function and thus pseudo-random functions.

The key observation is as follows. Suppose the adversary requests for q secret keys corresponding to random inputs C_{d_1}, \dots, C_{d_q} and then requests for an encryption of a random x . Then, the simulated ciphertext together with the q simulated secret keys constitute a description of the values $w\text{PRF}(x, d_1), \dots, w\text{PRF}(x, d_q)$, which is computationally indistinguishable from a sequence of q truly random bits via pseudo-randomness. By a standard information-theoretic argument, this means that the length of the ciphertext plus the secret keys must grow with q . To obtain a lower bound on the ciphertext size, we carefully exploit the fact that the simulator has to generate the secret keys before it sees the output of $w\text{PRF}(x, \cdot)$. Then, the simulator has to generate a small ciphertext that “explains” all these pseudorandom values which is impossible using a compressibility argument. More generally, we show that (1) weak pseudo-random family is “incompressible”, and (2) NA-SIM-secure functional encryption only exists for “compressible” circuit families. (In particular, the circuit family for all *public-index* predicate encryption is compressible.)

This idea is reminiscent of the obfuscation impossibility result of Goldwasser and Kalai [GK05], although the precise settings are quite different (in particular, functional encryption and program obfuscation seem incomparable, although related, objects).

Implications. The basic idea described above can be extended to a lower bound for even weaker forms of the simulation-based definition, including (a non-adaptive variant of) the definition of Boneh, Sahai and Waters [BSW11]. Here, we mention yet another implication of this idea.

Gorbunov, Vaikuntanathan and Wee [GVW12] recently presented a 1-AD-SIM-secure functional encryption scheme for all circuits, assuming that the adversary can only corrupt an a-priori bounded number of users (and thus, get the corresponding secret keys). One of the shortcomings of their bounded-collusion security notion as well as their construction is that the parameters of the system, and especially the size of the ciphertext depends on the collusion bound q . A natural question is whether their ciphertexts can be made to have size independent of q (or, at the very least, $o(q)$).⁸ Indeed, in light of the results of Dodis, Katz, Xu and Yung [DKXY02] and most recently, Goldwasser, Lewko and Wilson [GLW12] in the context of bounded-collusion IBE, one might expect that achieving “short” ciphertexts can actually be possible in general.

Unfortunately, our techniques result in a strong negative answer to this question.

⁸ The previous lower bound for many-AD-SIM IBE in [BSW11] (which says that the secret key size must grow with the number of challenge ciphertexts) is not applicable here as the [GVW12] construction considers only a single challenge ciphertext.

Corollary 1. *There exists a family of circuits \mathcal{C} such that for every $q = q(\kappa)$, there are no q -collusion resistant 1-NA-SIM-secure functional encryption schemes with ciphertexts of size $o(q)$.*

1.2 New Perspectives: Unbounded Simulation

The preceding lower bound together with those of Boneh, Sahai and Waters [BSW11] show that even fairly weak simulation-based definitions of functional encryption are unachievable for a large and natural class of circuits. This state of affairs begs the question:

What is a meaningful and generally realizable security notion for functional encryption?

While we do not provide a definitive answer to this question in our work, we believe that the quest for the right definition should incorporate insights from secure computation and zero knowledge. Indeed, Sahai and Seyalioglu [SS10] used Yao’s garbled circuits to construct a one-query secure functional encryption scheme for all circuits. Subsequently, Gorbunov et al. [GVW12] exploited more techniques and insights from secure computation [Yao86, BGW88, BMR90] to derive general feasibility results for functional encryption with bounded collusions.

We put forth USIM security, where the simulator has unbounded computational power. In particular, this would allow us to circumvent our lower bound in the previous section, since the lower bound crucially relies on the existence of an efficient simulator in order to break the weak pseudo-random function. Similar notions have been considered for zero knowledge and secure computation [Pas03, PS04, BS05].⁹ In the more basic setting of public-key encryption, we know that IND and SIM are equivalent [GM82], and it follows readily that all of IND, USIM, and SIM are also equivalent.

We begin an intuitive interpretation of what USIM security buys us, via the real/ideal paradigm. Consider an efficient adversary \mathcal{A} holding a secret key sk_C . Then, an encryption of x leaks no more information about x apart from what a computationally *unbounded* adversary can learn from $C(x)$. Specifically, in the case of public-index predicate encryption where the predicate is false, $C(x)$ hides the payload message μ completely, even against unbounded adversaries. Thus, USIM security for public-index predicate encryption offers very meaningful simulation-based security.¹⁰ On the other hand, for circuits that only hide

⁹ The works on zero knowledge and secure computation focus on quasi-polynomial-time simulators. We observe that our lower bound also rules out quasi-polynomial-time simulators assuming the existence of one-way functions with sub-exponential hardness.

¹⁰ Prior work of O’Neill [O’N10, Section 4] implies that NA-IND, NA-USIM and NA-SIM are equivalent for public-index predicate encryption. This does not subsume the point we are making because our argument applies also to the adaptive setting, where AD-IND and AD-SIM are provably *not* equivalent for public-index predicate encryption.

information about x computationally, USIM security would be inadequate and SIM security remains the desirable notion.

We observe that USIM security is “sandwiched” between IND and SIM security, that is, for $yy \in \{\text{NA}, \text{AD}\}$:

$$yy\text{-IND} \Leftarrow yy\text{-USIM} \Leftarrow yy\text{-SIM}$$

This result holds for both single and many message definitions. Then, we build upon the results in [BSW11] to obtain separations and impossibility results for USIM security:

- We present a counter-example separating SIM and USIM security. In fact, the example (which encodes a one-way permutation into the circuit family) is exactly that in [BSW11, Section 4.2] for separating SIM and IND security.
- We show that it is impossible to achieve many-AD-USIM security for the IBE functionality. This strengthens the many-AD-SIM lower bound for IBE in [BSW11, Section 5.2]. That is, the latter is fundamentally about the limitations of simulation-based security notion, and not about efficiency.
- A discussion in [BSW11] pointed that IND security is inadequate whenever “the output of the functionality is supposed to have some computational hiding properties”; however, there was no precise formalization of the latter. USIM security provides a way to make this statement precise. Recall that USIM security implies IND security, and therefore, if USIM security is inadequate for some functionality, then IND security must be inadequate for the same functionality. Thanks to the real/ideal paradigm, we have a simple “litmus test” for checking whether USIM security is adequate or not. Specifically, USIM security is inadequate if $C(x)$ reveals more information about x to an unbounded adversary than to an efficient adversary. (Indeed, this is trivially the case for the separation for USIM and SIM security since an unbounded adversary can invert the one-way permutation.)

We leave as an intriguing open problem the question of establishing either a separation or an equivalence between USIM and IND security. As a first step, we establish an equivalence between USIM and IND security in the “fully non-adaptive” setting, where all queries and messages are generated by the adversary before it sees the public parameters (See Remark ?? for details).

Organization. We refer the reader to Figure 1 for a survey of our results and open problems, and to Appendix ?? for results on the unbounded simulator definition.

1.3 Discussion

FunctoMania. Let’s be wishful thinkers for a minute – suppose we can have whatever we hope for in functional encryption, call this world “Functomania”.

What does Functomania look like? In light of the existing (im)possibilities, there will be two incomparable “dream results”¹¹:

- 1-AD-SIM secure public index predicate encryption for all efficient predicates; such schemes also satisfy 1-AD-IND, 1-AD-USIM, and many-AD-IND security.
- 1-AD-USIM secure functional encryption for all poly-size circuits; such schemes also satisfy 1-AD-IND and many-AD-IND security.

The IND-(U)SIM Conundrum. From a definitional stand-point, SIM/USIM-based security notions are preferable to IND-based security notion, as they offer a stronger security guarantee that has a natural, intuitive and aesthetically pleasing interpretation via the real/ideal paradigm. On the other hand, IND-based security notion allows us to bypass the impossibility results given in [BSW11] and in this work; in addition, they guarantee *message composability* in that security with a single ciphertext implies security for multiple ciphertexts (and so does NA-SIM considered in [GVW12] and those considered in an independent work [BF13]). We do not offer a complete answer to this conundrum; instead, we point out that 1-AD-SIM and 1-AD-USIM appear to be an adequate compromise for predicate encryption and general functional encryption respectively. We also note that such a conundrum is not unique to functional encryption, and has indeed previously surfaced and widely studied in the context of zero knowledge [FS90,Pas03] and secure multi-party computation [PS04,BS05,MPR06]. One notable difference is that in zero knowledge and secure computation, super-polynomial time simulation offers concurrency; this is not the case for functional encryption. (The lower bound for many-AD-USIM-secure IBE indicates that even unbounded-time simulation does not help with message composability.)

Concurrent and independent work. In an independent work, Bellare and O’Neill [BO12] put forth simulation-based definitions for functional encryption with non-black-box simulators. In addition, they extended the [BSW11] lower bound for IBE to the setting of efficient, non-black-box simulators, assuming the existence of collision-resistant hash functions. At a high level, the work is similar in spirit to our results on USIM security in that both consider larger classes of simulators than that in [BSW11]. The independent work of Barbosa and Farshim [BF13] takes a orthogonal approach, namely to restrict the adversary’s key queries via some “potential leakage relation”.

As with [BSW11], the definitions we study in this work are “inherently black-box” since the simulator must explicitly provide the adversary with secret keys and ciphertexts. Moreover, our NA-SIM lower bound relies crucially on black-box simulation as the compression comes from the simulated ciphertext. This leaves as an open problem the question of realizing (or ruling out) many-AD-USIM IBE with a non-black-box simulator.

¹¹ Substantial progress were made recently on both of these problems in [SW12,GVW13,GKP⁺13].

2 Functional Encryption

Let $\mathcal{X} = \{\mathcal{X}_\kappa\}_{\kappa \in \mathbb{N}}$ and $\mathcal{Y} = \{\mathcal{Y}_\kappa\}_{\kappa \in \mathbb{N}}$ denote ensembles where each \mathcal{X}_κ and \mathcal{Y}_κ is a finite set. Let $\mathcal{C} = \{\mathcal{C}_\kappa\}_{\kappa \in \mathbb{N}}$ denote an ensemble where each \mathcal{C}_κ is a finite collection of circuits, and each circuit $C \in \mathcal{C}_\kappa$ takes as input a string $x \in \mathcal{X}_\kappa$ and outputs $C(x) \in \mathcal{Y}_\kappa$.

A functional encryption scheme \mathcal{FE} for \mathcal{C} consists of four algorithms $\mathcal{FE} = (\text{FE.Setup}, \text{FE.Keygen}, \text{FE.Enc}, \text{FE.Dec})$ defined as follows¹².

- **Setup** $\text{FE.Setup}(1^\kappa)$ is a p.p.t. algorithm that takes as input the unary representation of the security parameter and outputs the master public and secret keys (MPK, MSK).
- **Key Generation** $\text{FE.Keygen}(\text{MSK}, C)$ is a p.p.t. algorithm that takes as input the master secret key MSK and a circuit $C \in \mathcal{C}_\kappa$ and outputs a corresponding secret key SK_C .
- **Encryption** $\text{FE.Enc}(\text{MPK}, x)$ is a p.p.t. algorithm that takes as input the master public key MPK and an input message $x \in \mathcal{X}_\kappa$ and outputs a ciphertext CT.
- **Decryption** $\text{FE.Dec}(\text{SK}_C, \text{CT})$ is a deterministic algorithm that takes as input the secret key SK_C and a ciphertext CT and outputs $C(x)$.

Definition 1 (Correctness). A functional encryption scheme \mathcal{FE} is correct if for all $C \in \mathcal{C}_\kappa$ and all $x \in \mathcal{X}_\kappa$,

$$\Pr \left[\begin{array}{l} (\text{MPK}, \text{MSK}) \leftarrow \text{FE.Setup}(1^\kappa); \\ \text{FE.Dec}(\text{FE.Keygen}(\text{MSK}, C), \text{FE.Enc}(\text{MPK}, x)) \neq C(x) \end{array} \right] = \text{negl}(\kappa)$$

where the probability is taken over the coins of FE.Setup , FE.Keygen , and FE.Enc .

2.1 A Simulation-based Definition of Security

In this section, we present a simulation-based definition of functional encryption, similar in spirit to the way one defines security for secure computation via the ideal/real paradigm. We define the security game for a single message since our lower bounds apply to this weaker setting. However, this definition can be easily extended to many messages setting (see Appendix ??).

Definition 2 (1-NA-SIM- and 1-AD-SIM- Security). Let \mathcal{FE} be a functional encryption scheme for a circuit family \mathcal{C} . Consider a p.p.t. adversary $A = (A_1, A_2)$ and a stateful p.p.t. simulator Sim .¹³ Let $U_x(\cdot)$ denote a universal oracle, such that $U_x(C) = C(x)$. Consider the following two experiments:

¹² Unlike in [BSW11], we do not consider the “empty key”.

¹³ One can replace a stateful simulator can be replaced by a regular (stateless) simulator that outputs a state st_s upon each invocation which is carried over to its next invocation.

$\text{Exp}_{\mathcal{FE}, A}^{\text{real}}(1^\kappa)$:	$\text{Exp}_{\mathcal{FE}, \text{Sim}}^{\text{ideal}}(1^\kappa)$:
1: $(\text{MPK}, \text{MSK}) \leftarrow \text{FE.Setup}(1^\kappa)$ 2: $(x, st) \leftarrow A_1^{\text{FE.Keygen}(\text{MSK}, \cdot)}(\text{MPK})$ 3: $\text{CT} \leftarrow \text{FE.Enc}(\text{MPK}, x)$ 4: $\alpha \leftarrow A_2^{\mathcal{O}(\text{MSK}, \cdot)}(\text{MPK}, \text{CT}, st)$ 5: <i>Output</i> (x, α)	1: $\text{MPK} \leftarrow \text{Sim}(1^\kappa)$ 2: $(x, st) \leftarrow A_1^{\text{Sim}(\cdot)}(\text{MPK})$ 3: $\text{CT} \leftarrow \text{Sim}^{U_x(\cdot)}(1^\kappa, 1^{ x })$ 4: $\alpha \leftarrow A_2^{\mathcal{O}'(\cdot)}(\text{MPK}, \text{CT}, st)$ 5: <i>Output</i> (x, α)

We distinguish between two cases of the above experiment:

1. The adaptive experiment, where:
 - the oracle $\mathcal{O}(\text{MSK}, \cdot) = \text{FE.Keygen}(\text{MSK}, \cdot)$ and
 - the oracle $\mathcal{O}'(\cdot)$ is the simulator, namely $\text{Sim}^{U_x(\cdot)}(\cdot)$.
 We call a stateful simulator algorithm Sim admissible if, on each input C , Sim makes just a single query to its oracle $U_x(\cdot)$ on C itself. The functional encryption scheme \mathcal{FE} is then said to be simulation-secure for one message against adaptive adversaries (1-AD-SIM-secure, for short) if there is an admissible stateful p.p.t. simulator Sim such that for every p.p.t. adversary $A = (A_1, A_2)$, the following two distributions are computationally indistinguishable:

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

2. The non-adaptive experiment, where the oracles $\mathcal{O}(\text{MSK}, \cdot)$ and $\mathcal{O}'(\cdot)$ are both the “empty oracles” that return nothing. The functional encryption scheme \mathcal{FE} is then said to be simulation-secure for one message against non-adaptive adversaries (1-NA-SIM-secure, for short) if there is an admissible stateful p.p.t. simulator Sim such that for every p.p.t. adversary $A = (A_1, A_2)$, the two distributions above are computationally indistinguishable.

Remarks on the Definition. Our definition is stronger than that in [BSW11] but weaker than that in [GVW12]; our lower bound in Section 4 holds for all three definitions. Amongst the three, the one in [GVW12] is the only for which we know a composition theorem where security for one message implies security for many messages, in the non-adaptive setting. Note that composition in the non-adaptive setting is the “best” we can hope for; composition in the adaptive setting is essentially impossible by many-AD-SIM lower bound for IBE [BSW11]. In more detail:

- In [BSW11], the simulator is given oracle access to A_2 , which it can call on any ciphertext. Therefore, it can “rewind” the adversary A_2 and adaptively reconstruct the view, which is problematic for composition [PRS02, Lin04, BMQU07]. We call this a “rewinding” definition. In our “straight-line” definition, the simulator must commit to a ciphertext once and for all, which makes it stronger.

- Unlike our definition, the [GVW12] definition does not allow the simulator to fake or “program” the setup parameters and the secret keys. The difficulty in proving a composition theorem for our definition lies in that the simulator may use “trapdoor” information from faking the setup parameters and secret keys while simulating the ciphertext.

We note that in the equivalence of NA-IND and NA-SIM under pre-image sampleability in [O’N10, Section 4], the NA-SIM-simulator actually satisfies the stronger definition in [GVW12].

The Indistinguishability-based Definition of Security. We refer the reader to the full version for the non-adaptive NA-IND and the adaptive AD-IND notions of security.

3 Preliminaries

Notations. Let \mathcal{D} denote a distribution over some finite set S . Then, $x \leftarrow \mathcal{D}$ is used to denote the fact that x is chosen from the distribution \mathcal{D} . When we say $x \leftarrow S$, we simply mean that x is chosen from the uniform distribution over S . Let κ denote the security parameter.

Definition 3 (wPRF). Let $\text{wPRF} = \{\text{wPRF}_\kappa\}_{\kappa \in \mathbb{N}}$ denote a family of efficiently computable functions where $\text{wPRF}_\kappa : \{0, 1\}^{n(\kappa)} \times \{0, 1\}^{m(\kappa)} \rightarrow \{0, 1\}^{k(\kappa)}$, the first argument of which is called the seed to the wPRF and the second argument is the input.

For every probabilistic polynomial time oracle distinguisher Dist , consider the following two experiments:

- $\text{Real}_{\text{Dist}}(1^\kappa)$: Choose $x \xleftarrow{\$} \{0, 1\}^{n(\kappa)}$ and run Dist with access to a probabilistic oracle $\mathcal{O}_{\text{real}}(x)$ which, when invoked, chooses a uniformly random $d \leftarrow \{0, 1\}^{m(\kappa)}$ and returns the pair $(d, \text{wPRF}_\kappa(x, d))$. This experiment outputs whatever Dist outputs.
- $\text{Rand}_{\text{Dist}}(1^\kappa)$: Choose a uniformly random function $R : \{0, 1\}^{m(\kappa)} \rightarrow \{0, 1\}^{k(\kappa)}$ and run Dist with access to a probabilistic oracle $\mathcal{O}_{\text{rand}}(R)$ which, when invoked, chooses a uniformly random $d \leftarrow \{0, 1\}^{m(\kappa)}$ and returns the pair $(d, R(d))$. This experiment outputs whatever Dist outputs.

We say wPRF is a weak pseudo-random function if for all p.p.t. distinguishers Dist ,

$$|\Pr[\text{Real}_{\text{Dist}}(1^\kappa) = 1] - \Pr[\text{Rand}_{\text{Dist}}(1^\kappa) = 1]| = \text{negl}(\kappa)$$

where the probabilities are over the choice of x and R , as well as the coin-tosses of Dist and the oracles $\mathcal{O}_{\text{real}}$ and $\mathcal{O}_{\text{rand}}$.

This is in contrast to the stronger notion of (regular) pseudo-random functions where the distinguisher Dist gets query access to the function, namely it can query the function on inputs x of its choice and get either the output of the function (in the real world) or independent random bits (in the ideal world).

In our impossibility result, we will use a weak pseudo-random function with seed length $n(\kappa) = \kappa$ and output length $k(\kappa) = 1$.

4 Impossibility Results for Functional Encryption

In this section, we present our main lower bound for 1-NA-SIM-secure functional encryption. We begin with a notion of “incompressible” circuits. Then, we show that (1) weak pseudo-random functions are “incompressible”, and (2) 1-NA-SIM-secure functional encryption only exists for “compressible” circuits. Putting the two together yields our lower bound.

4.1 Incompressible Circuits

We first define a family of compressible circuits. Informally, we say that a family of circuits $\{\mathcal{G}_\kappa\}$ is (ℓ, t) -compressible if for a list of uniformly random circuit descriptions $G_1, \dots, G_\ell \in \mathcal{G}_\kappa$ and a uniformly chosen input x , there is some efficiently computable description of $G_1(x), \dots, G_\ell(x)$ of size t .

Definition 4 (Incompressible Circuits). *Let $\ell = \ell(\kappa)$ and $t = t(\kappa)$ be functions of the security parameter κ . A family of circuits $\mathcal{G} = \{\mathcal{G}_\kappa\}_{\kappa \in \mathbb{N}}$ is (ℓ, t) -compressible if there exists a family of (deterministic) compressor circuits $\{\mathbf{C}_\kappa\}_{\kappa \in \mathbb{N}}$ and a family of decompressor circuits $\{\mathbf{D}_\kappa\}_{\kappa \in \mathbb{N}}$ such that:*

- (polynomial size) the circuits \mathbf{C}_κ and \mathbf{D}_κ have size $\text{poly}(\kappa, \ell)$.
- (mild compression) for sufficiently large κ and all x ,

$$|\mathbf{C}_\kappa(G_1, \dots, G_\ell, y_1, \dots, y_\ell)| = t$$

where $y_i = G_i(x)$.

- (correctness) there is a polynomial $p = p(\kappa)$ such that

$$\Pr[x \xleftarrow{\$} \{0, 1\}^\kappa, G_1, \dots, G_\ell \xleftarrow{\$} \mathcal{G}_\kappa, y_i = G_i(x) : \mathbf{D}_\kappa(G_1, \dots, G_\ell, \mathbf{C}_\kappa(G_1, \dots, G_\ell, y_1, \dots, y_\ell)) = (y_1, \dots, y_\ell)] \geq 1/p(\kappa)$$

where the probability is taken over the choice of x as well as the circuits G_1, \dots, G_ℓ .

The family \mathcal{G} is (ℓ, t) -incompressible if it is not (ℓ, t) -compressible.

We now give examples of (in)compressible circuits. First, consider the notion of pre-image samplable family of circuits introduced by O’Neill [O’N10] which requires that given $G_1(x), \dots, G_\ell(x)$, there is a polynomial-time algorithm that returns an arbitrary x' such that $G_i(x') = G_i(x)$ for all i . In our language, this says that the family \mathcal{G} is $(\ell, |x'|)$ -compressible; the compression algorithm simply outputs x' .

Next, consider an arbitrary public-index circuit family parametrized by predicates P and given by:

$$G_P(\text{ind}, \mu) = \begin{cases} (\text{ind}, \mu) & \text{if } P(\text{ind}) = 1 \\ (\text{ind}, \perp) & \text{otherwise} \end{cases}$$

It is easy to see that this circuit family is $(\ell, |(\text{ind}, \mu)|)$ -compressible. On input

$$G_{P_1}(\text{ind}, \mu), \dots, G_{P_\ell}(\text{ind}, \mu)$$

the compression algorithm always learn ind . In addition, if $P_i(\text{ind}) = 1$ for some i , then the compressor also learns μ and hence it outputs (ind, μ) . If $P_i(\text{ind}) = 0$ for all i , then the compressor outputs (ind, \perp) . Given $(G_{P_1}, \dots, G_{P_\ell}, (\text{ind}, \mu))$ the decoding algorithm outputs $y_i = (\text{ind}, \mu)$ if $G_{P_i}(\text{ind}) = 1$ and $y_i = (\text{ind}, \perp)$ otherwise. Given $(G_{P_1}, \dots, G_{P_\ell}, (\text{ind}, \perp))$ the decoder simply outputs $y_i = (\text{ind}, \perp)$ for all i .

On the other hand, as we show below (see Lemma 1), any family of (weak) pseudo-random functions is incompressible in a strong sense. More precisely, consider a family of circuits $\mathcal{G} = \{G_{d_i}(\cdot) = \text{wPRF}(\cdot, d_i)\}$ where d_i serves as the input to the pseudo-random function. Informally, the incompressibility is due to the fact that a sequence $(G_{d_1}(x), \dots, G_{d_\ell}(x)) = (\text{wPRF}(x, d_1), \dots, \text{wPRF}(x, d_\ell))$ is indistinguishable from a sequence of uniformly random bits, which are clearly incompressible.

Lemma 1 (weak PRFs are $(\ell, \ell - \kappa)$ -incompressible). *Let $\text{wPRF} = \{\text{wPRF}_\kappa : \{0, 1\}^\kappa \times \{0, 1\}^{m(\kappa)} \rightarrow \{0, 1\}\}_{\kappa \in \mathbb{N}}$ be a family of weak pseudo-random functions, where $m(\kappa) = \omega(\log \kappa)$. Define $G_d(x) = \text{wPRF}(x, d)$. Consider a family $\mathcal{G} = \{\mathcal{G}_\kappa\}_{\kappa \in \mathbb{N}}$ defined as*

$$\mathcal{G}_\kappa = \{G_d(\cdot) : |d| = m(\kappa)\}$$

Then, \mathcal{G} is $(\ell, \ell - \kappa)$ -incompressible.

We refer the reader to the full version for the formal proof.

4.2 The Impossibility Result

We are now ready to state and prove our main theorem.

Theorem 2. *There exists a family of circuits \mathcal{G} for which there are no 1-NA-SIM-secure functional encryption schemes.*

Proof. We consider two cases.

Case 1: Assume there exists a circuit family of weak pseudo-random functions

$$\text{wPRF} = \{\text{wPRF}_\kappa : \{0, 1\}^\kappa \times \{0, 1\}^{m(\kappa)} \rightarrow \{0, 1\}\}_{\kappa \in \mathbb{N}}$$

where $m(\kappa) = \omega(\log \kappa)$. Let $G_d(x) = \text{wPRF}(x, d)$ and consider a family $\mathcal{G} = \{\mathcal{G}_\kappa\}_{\kappa \in \mathbb{N}}$ defined as

$$\mathcal{G}_\kappa = \{G_d(\cdot) : |d| = m(\kappa)\}$$

Assume, for the sake of contradiction, there exist a 1-NA-SIM-secure function encryption scheme \mathcal{FE} for \mathcal{G} , and let $|\text{CT}|$ denote the length of a ciphertext in the scheme. Let $\ell = \ell(\kappa) = |\text{CT}| + \kappa$.

From Lemma 1, we know that \mathcal{G} is $(|\text{CT}| + \kappa, |\text{CT}|)$ -incompressible. However, Lemma 2 below tells us that since there is a 1-NA-SIM secure scheme for \mathcal{G} , the family \mathcal{G} is $(|\text{CT}| + \kappa, |\text{CT}|)$ -compressible. This gives us the desired contradiction, and therefore, there cannot exist a 1-NA-SIM-secure functional encryption scheme for \mathcal{G} .

Case 2: Assume there does not exist a family of weak pseudo-random functions. Also, for the sake of contradiction, assume there exists a 1-NA-SIM-secure function encryption scheme for all families of circuits \mathcal{G} .

In particular, this means that there is a functional encryption scheme for the empty circuit family (namely, a family \mathcal{G} that does not contain any circuits at all). A 1-NA-SIM-secure scheme \mathcal{FE} for \mathcal{G} is also a secure public-key encryption scheme. Since public-key encryption implies one-way functions, which in turn imply pseudo-random functions [GGM86,HILL99], we obtain the desired contradiction.

Lemma 2 (1-NA-SIM \Rightarrow $(\ell, |\text{CT}|)$ -compressibility). *Let $\mathcal{G} = \{\mathcal{G}_\kappa\}_{\kappa \in \mathbb{N}}$ be a family of circuits. Suppose there exists a 1-NA-SIM-secure functional encryption scheme for the \mathcal{G} . Then, the family \mathcal{G} is $(\ell, |\text{CT}|)$ -compressible for any polynomially bounded $\ell = \ell(\kappa)$, where $|\text{CT}|$ denotes size of the encryption of input x .*

Informally, the compression algorithm works as follows: on input G_1, \dots, G_ℓ and $G_1(x), \dots, G_\ell(x)$, the output is the simulated ciphertext corresponding to an encryption of x . The decompression algorithm then evaluates the decryption algorithm, which is guaranteed to produce $G_1(x), \dots, G_\ell(x)$.

Proof. Let $(\text{FE.Setup}, \text{FE.Keygen}, \text{FE.Enc}, \text{FE.Dec})$ denote the encryption scheme for the family \mathcal{G} . Consider the adversary $A = (A_1, A_2)$ in the 1-NA-SIM security experiment that acts as follows:

- A_1 chooses $G_1, \dots, G_\ell \xleftarrow{\$} \mathcal{G}$ independently at random and requests for the corresponding secret keys $\text{SK}_1, \dots, \text{SK}_\ell$. In addition, it chooses $x \xleftarrow{\$} \{0, 1\}^{m(\kappa)}$ and outputs x as the challenge message and state

$$(G_1, \dots, G_\ell, \text{SK}_1, \dots, \text{SK}_\ell)$$

- A_2 outputs α composed of the challenge ciphertext and the state

$$(G_1, \dots, G_\ell, \text{SK}_1, \dots, \text{SK}_\ell)$$

Let Sim denote the (admissible) stateful p.p.t. simulator guaranteed by 1-NA-SIM security. We show how to use the simulator to construct a family of (deterministic) compressor and decompressor circuits \mathbf{C}_ρ and \mathbf{D}_ρ , indexed by a random string ρ corresponding to the random tape for the simulator:

- The compressor \mathbf{C}_ρ , on input G_1, \dots, G_ℓ and y_1, \dots, y_ℓ works as follows: first, compute $\text{MPK} \leftarrow \text{Sim}(1^\kappa; \rho)$ and secret keys $\{\text{SK}_i : \text{SK}_i \leftarrow \text{Sim}(G_i; \rho)\}_{i \in [\ell]}$.

Then compute and output CT as the compressed string, where queries $G_i(x)$ are answered with y_i :

$$\text{CT} \leftarrow \text{Sim}^{U_x(\cdot)}(1^{|\mathcal{m}(\kappa)|})$$

- The decompressor \mathbf{D}_ρ , on input G_1, \dots, G_ℓ and CT first reconstructs the master public key $\text{MPK} \leftarrow \text{Sim}(1^\kappa; \rho)$ and the set of secret keys:

$$\{\text{SK}_i : \text{SK}_i \leftarrow \text{Sim}(G_i; \rho)\}_{i \in [\ell]}$$

Note that \mathbf{D}_ρ has the same randomness ρ hard-wired, and so the secret keys SK_i are exactly the same as those used by \mathbf{C}_ρ . Finally, it computes and outputs:

$$\{y_i \leftarrow \text{FE.Dec}(\text{SK}_i, \text{CT})\}_{i \in [\ell]}$$

Formally, we output $(\mathbf{C}_\rho, \mathbf{D}_\rho)$ for a random ρ , which is a pair of polynomial-size circuits. Clearly, we achieve *mild compression* (where $|\text{CT}|$ is the compressor's output size), since the size of CT is determined by the functional encryption scheme and independent of ℓ . To establish correctness, it suffices to show that:

$$\Pr_{\rho, x, G_1, \dots, G_\ell} [\mathbf{D}_\rho(G_1, \dots, G_\ell, \mathbf{C}_\rho(G_1, \dots, G_\ell, G_1(x), \dots, G_\ell(x))) = (G_1(x), \dots, G_\ell(x))] \geq 1 - \text{negl}(\kappa)$$

Here, we will rely on the correctness of the functional encryption scheme as well as 1-NA-SIM-security. First, consider the distinguisher Dist that given the output $(x, \text{CT}, G_1, \dots, G_\ell, \text{SK}_1, \dots, \text{SK}_\ell)$ of the adversary A_2 proceeds as follows:

Output 1 iff for all $i \in [\ell]$, $\text{FE.Dec}(\text{SK}_i, \text{CT}) = G_i(x)$.

Observe that by correctness of the encryption scheme, Dist outputs 1 with probability $1 - \text{negl}(\kappa)$ given the output of the adversary A_2 in the 1-NA-SIM experiment. Therefore, by 1-NA-SIM-security, Dist also outputs 1 with probability $1 - \text{negl}(\kappa)$ given the output of the (admissible) simulator, where the randomness is taken over the coin tosses ρ of the simulator, along with the random choices of x, G_1, \dots, G_ℓ .

This shows that the pair of circuits $(\mathbf{C}_\rho, \mathbf{D}_\rho)$ for a uniformly random ρ is a correct compressor-decompressor pair. Therefore, we obtain a $(\ell, |\text{CT}|)$ -compressor and a decompressor, thus establishing the lemma.

We point out here that our lower bound extends to the setting where the simulator is not required to be admissible, by using a family of (standard) pseudo-random functions.

Finally, the argument here generalizes to showing that functional encryption secure against an a-priori bounded number $q = q(\kappa)$ of collisions is impossible if one insists on small ciphertexts (namely, ciphertexts with much fewer than q bits). This matches the recent result of [GVW12] who construct such functional encryption schemes with ciphertexts of size polynomial in q .

Corollary 2. *There exists a family of circuits \mathcal{G} such that for every $q = q(\kappa)$, there are no q -collision resistant 1-NA-SIM-secure functional encryption schemes with ciphertexts of size $o(q)$.*

4.3 Extensions: Impossibility of Weaker Simulation-based Definitions

The idea behind our impossibility result is robust enough to apply to various relaxations of the simulation-based security definition. In this section, we describe a number of such extensions of our result.

Impossibility for the selective and random-input definitions. In the selective model, the adversary is required to commit to the secret key queries G_1, \dots, G_q as well as the challenge input x before the setup phase. In particular, this means that the adversary will not be able to pick up the circuits or the challenge input depending on the system parameters. Variants of the selective security model are frequently considered in the literature as a relaxations of regular security notions (see, e.g., [BB11, GPSW06, AFV11]). Another relaxation one can consider is one where the adversary is not allowed to choose the circuits or the challenge, but instead, they are chosen uniformly at random.

Our lower bound easily extends to these weaker notions, simply because the adversary we consider in the proof of Lemma 2 chooses the circuits and the challenge uniformly at random, and independent of the system parameters.

Impossibility for the non-adaptive BSW Definition (the “Rewinding Definition”). The main difference between the definition proposed by [BSW11] and our definition in Section 2 is that whereas our definition restricts the simulator to be “straight-line”, the BSW definition allows the simulator to “rewind” the adversary and interact with it in order to generate the view. For more details, we direct the reader to the full version.

We also state the impossibility extension for secret-key functional encryption in the full version.

Acknowledgments

We thank Shafi Goldwasser, Yael Kalai, Raluca Ada Popa and Charles Rackoff for a number of insightful conversations that helped improve the presentation of the impossibility result. We would also like to thank Adam O’Neill and Amit Sahai for helpful pointers and discussions.

References

- ABV⁺12. Shweta Agrawal, Xavier Boyen, Vinod Vaikuntanathan, Panagiotis Voulgaris, and Hoeteck Wee. Functional encryption for threshold functions (or fuzzy IBE) from lattices. In *PKC*, 2012.
- AFV11. Shweta Agrawal, David Mandell Freeman, and Vinod Vaikuntanathan. Functional encryption for inner product predicates from learning with errors. In *Asiacrypt*, 2011.
- BB11. Dan Boneh and Xavier Boyen. Efficient selective identity-based encryption without random oracles. *J. Cryptology*, 24(4):659–693, 2011.

- BF01. Dan Boneh and Matthew K. Franklin. Identity-based encryption from the weil pairing. In *CRYPTO*, pages 213–229, 2001.
- BF13. Manuel Barbosa and Pooya Farshim. On the semantic security of functional encryption schemes. In *Public Key Cryptography*, pages 143–161, 2013.
- BGW88. Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, STOC '88, pages 1–10, New York, NY, USA, 1988. ACM.
- BMQU07. Michael Backes, Jörn Müller-Quade, and Dominique Unruh. On the necessity of rewinding in secure multiparty computation. In *Proceedings of the 4th conference on Theory of cryptography*, TCC'07, pages 157–173, Berlin, Heidelberg, 2007. Springer-Verlag.
- BMR90. Donald Beaver, Silvio Micali, and Phillip Rogaway. The round complexity of secure protocols (extended abstract). In *STOC*, pages 503–513, 1990.
- BO12. Mihir Bellare and Adam O'Neill. Semantically-secure functional encryption: Possibility results, impossibility results and the quest for a general definition. Cryptology ePrint Archive, Report 2012/515, 2012.
- BS05. Boaz Barak and Amit Sahai. How to play almost any mental game over the net - concurrent composition via super-polynomial simulation. In *FOCS*, pages 543–552, 2005.
- BSW11. Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In *TCC*, pages 253–273, 2011.
- BW06. Xavier Boyen and Brent Waters. Anonymous hierarchical identity-based encryption (without random oracles). In *CRYPTO*, pages 290–307, 2006.
- CHK03. Ran Canetti, Shai Halevi, and Jonathan Katz. A forward-secure public-key encryption scheme. In *EUROCRYPT*, pages 255–271, 2003.
- Coc01. Clifford Cocks. An identity based encryption scheme based on quadratic residues. In *IMA Int. Conf.*, pages 360–363, 2001.
- DKXY02. Yevgeniy Dodis, Jonathan Katz, Shouhuai Xu, and Moti Yung. Key-insulated public key cryptosystems. In *In EUROCRYPT*, pages 65–82. Springer-Verlag, 2002.
- FS90. Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *STOC*, pages 416–426, 1990.
- GGH⁺13. Sanjam Garg, Craig Gentry, Shai Halevi, Amit Sahai, and Brent Waters. Attribute based encryption for circuits from multilinear maps. *CRYPTO '13*, 2013.
- GGM86. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
- GK05. Shafi Goldwasser and Yael Tauman Kalai. On the impossibility of obfuscation with auxiliary input. In *FOCS*, pages 553–562, 2005.
- GKP⁺13. Shafi Goldwasser, Yael Kalai, Raluca Ada Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. Succinct functional encryption and its power: Reusable garbled circuits and beyond. In *STOC*, 2013. To appear.
- GLW12. Shafi Goldwasser, Allison B. Lewko, and David A. Wilson. Bounded-collusion IBE from key homomorphism. In *TCC*, pages 564–581, 2012.
- GM82. Shafi Goldwasser and Silvio Micali. Probabilistic encryption and how to play mental poker keeping secret all partial information. In *STOC*, pages 365–377, 1982.
- GPSW06. Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *ACM Conference on Computer and Communications Security*, pages 89–98, 2006.

- GVW12. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In *CRYPTO*, pages 162–179, 2012. Also Cryptology ePrint Archive, Report 2012/521.
- GVW13. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *Proceedings of the 45th annual ACM symposium on Symposium on theory of computing*, STOC '13, pages 545–554, New York, NY, USA, 2013. ACM.
- HILL99. Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- KSW08. Jonathan Katz, Amit Sahai, and Brent Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *EUROCRYPT*, pages 146–162, 2008.
- Lin04. Yehuda Lindell. Lower bounds and impossibility results for concurrent self composition. *the Journal of Cryptology*, 2004.
- 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.
- MPR06. Silvio Micali, Rafael Pass, and Alon Rosen. Input-indistinguishable computation. In *FOCS*, pages 367–378, 2006.
- O’N10. Adam O’Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010. <http://eprint.iacr.org/>.
- Pas03. Rafael Pass. Simulation in quasi-polynomial time and its application to protocol composition. In *EUROCRYPT*, pages 160–176, 2003.
- PRS02. Manoj Prabhakaran, Alon Rosen, and Amit Sahai. Concurrent zero knowledge with logarithmic round-complexity. In *In 43rd FOCS*, pages 366–375, 2002.
- PRV12. Bryan Parno, Mariana Raykova, and Vinod Vaikuntanathan. How to delegate and verify in public: Verifiable computation from attribute-based encryption. In Ronald Cramer, editor, *TCC*, volume 7194 of *Lecture Notes in Computer Science*, pages 422–439. Springer, 2012.
- PS04. Manoj Prabhakaran and Amit Sahai. New notions of security: achieving universal composability without trusted setup. In *STOC*, pages 242–251, 2004.
- Sha84. Adi Shamir. Identity-based cryptosystems and signature schemes. In *CRYPTO*, pages 47–53, 1984.
- SS10. Amit Sahai and Hakan Seyalioglu. Worry-free encryption: functional encryption with public keys. In *ACM Conference on Computer and Communications Security*, pages 463–472, 2010.
- SW05. Amit Sahai and Brent Waters. Fuzzy identity-based encryption. In *EUROCRYPT*, pages 457–473, 2005.
- SW08. Amit Sahai and Brent Waters. Slides on functional encryption. power point presentation, 2008. <http://www.cs.utexas.edu/~bwaters/presentations/files/functional.ppt>.
- SW12. Amit Sahai and Brent Waters. Attribute-based encryption for circuits from multilinear maps. Cryptology ePrint Archive, Report 2012/592, 2012.
- Yao86. Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167, 1986.