

Cliptography: Clipping the Power of Kleptographic Attacks

Alexander Russell* Qiang Tang† Moti Yung‡ Hong-Sheng Zhou§

July 11, 2015

Abstract

Kleptography, originally introduced by Young and Yung [Crypto '96], studies how to steal information securely and subliminally from cryptosystems. Secure cryptosystems can be broken if they are maliciously implemented since the adversary may have some backdoors embedded in the implementation. Although kleptographic attacks have been investigated about two decades ago, for too long the possibility of kleptographic attacks have been dismissed and been viewed only as a far-fetched theoretical concept. This is dramatically changed when real-world examples were recently revealed by Edward Snowden, demonstrating that such deliberate attacks (directly inspired by the original work) exist and probably have been used for massive surveillance. In light of such possible failures of basic protective technology, the security community started to seriously re-investigate this important issue: one notable example is the work of Bellare, Paterson, and Rogaway [Crypto '14], which initiated the formal studies of attacks on symmetric key encryption algorithms.

Motivated by the original examples of subverting key generation algorithms in the kleptography papers from Young and Yung [Crypto '96, Eurocrypt '97], we initiate the study of cryptography in the case that *all* algorithms are subject to kleptographic attacks—we call it **cliptography**. As a first step, we formally study the fundamental primitives of one-way function and trapdoor one-way function in this complete subversion model. And more interesting, we investigate the general immunization strategy to clip the power of kleptographic subversions; concretely, we propose a general framework for sanitizing the (trapdoor) one-way function generation algorithm by hashing the function index, and prove that such procedure indeed destroys the connection between a subverted function generation procedure and any possible backdoor. Along the way, we propose a split program model for practical deployment.

We then examine the applications of (trapdoor) one way function secure in the complete subversion model in two ways. First we consider to build “higher level” primitives via black-box reductions. In particular, we consider how to use our trapdoor one-way function to defend against key generation sabotage, and showcase a digital signature scheme that preserves existential unforgeability when *all* algorithms (including key generation, which was not considered to be under attack before) are subject to kleptographic attacks. Also we demonstrate that the classic Blum-Micali pseudorandom generator (PRG) using our “unforgeable” one-way function yields a backdoor-free PRG. Second, we generalize our immunizing technique for one way functions, and propose a new public immunization strategy to randomize the public parameters of a (backdoored) PRG. Since the previous result by Dodis, Ganesh, Golovnev, Juels, and Ristenpart [Eurocrypt '15] requires an honestly generated random key, construction of secure PRG in the complete subversion model was also open until our paper.

*University of Connecticut, acr@cse.uconn.edu

†University of Connecticut, qiang@cse.uconn.edu

‡Columbia University and Google, moti@cs.columbia.edu

§Virginia Commonwealth University, hszhou@vcu.edu

Contents

1	Introduction	3
1.1	Our contribution	3
1.2	Related work	4
2	Kleptographic Attacks and the Complete Subversion Model	5
3	Constructing One-Way Functions in the Complete Subversion Model	6
3.1	Formalizing kleptographic attacks on one way functions	6
3.1.1	Strong-forgability and unforgeability for OWF/TDOWF	6
3.1.2	Forgeability and strong-unforgeability for OWF	8
3.2	Eliminating backdoors	8
3.2.1	General feasibility results	9
3.2.2	Practical results in the split-program model	10
4	Constructing Signatures in the Complete Subversion Model	13
5	Constructing Pseudorandom Generator in the Complete Subversion Model	16
5.1	Preliminaries: backdoored and backdoor-free PRGs	16
5.2	Constructing backdoor-free PRG from strongly unforgeable OWP	17
5.3	General public immunization strategy for PRG	19
6	Conclusion and Open Problems	20
A	Omitted Proofs	22

1 Introduction

Consider the conventional use of a cryptographic primitive, such as an encryption scheme: To encrypt a plaintext, the encryptor simply runs an implementation of the encryption algorithm that obtained from a hardware or software provider with the plaintext as input to generate the corresponding ciphertext. Although the underlying algorithms are well-studied and proven secure, malicious implementations could still leak secret information exclusively to the provider/manufacturer without being noticed (through a covert channel that is built/constructed using an embedded backdoor). It is notable that *such leakage is possible even if the implementation produces “functionally and statistically clean” output that is indistinguishable from that of a faithful implementation*. While the underlying concept of kleptography was proposed by Young and Yung two decades ago [19, 20], the recent Snowden revelations [12, 16] provided striking real-world examples that awakened the security community to the seriousness of these issues. As a result, the topic has recently received renewed attention; see, e.g., [1, 2, 7, 14]. In particular, Bellare, Paterson, and Rogaway [2]¹ studied algorithm substitution attacks, with focus on symmetric key encryption. Soon, Dodis, Ganesh, Golovnev, Juels, and Ristenpart [7] studied the backdoored version of pseudorandom generator (PRG).

1.1 Our contribution

We continue this line of pursuit. Specifically, we are motivated to develop cryptographic schemes in *complete subversion model*, in which *all* algorithms of a scheme are potentially subverted by the adversary. We study the fundamental cryptographic primitives, (trapdoor) one-way functions (OWFs) in the complete subversion model, and apply these primitives to construct other cryptographic schemes such as digital signatures and PRG. Along the way, we identify novel generic defending strategies. We intend to stimulate a systematic study of **cliptography**, to provide a broader class of cryptographic building blocks and a larger set of defending strategies, eventually clipping out potential kleptographic attacks that arise from maliciously implemented components. We note that, prior to our work, only restricted type of kleptographic attacks can be addressed; see *Related work* in Section 1.2. In detail, we show the following:

- We study (trapdoor) one-way functions in the presence of kleptographic attacks. We first introduce *strong forgeability* that captures natural kleptographic attacks, namely: there is a specification that is proven secure for a (trapdoor) one-way function; the sabotaged function generation algorithm delivers a similar output distribution as the specification distribution; however with a pre-chosen backdoor, the adversary can invert the entire family of functions that are generated by this subverted algorithm. We then show that such (adversarial) objects can indeed be constructed from (trapdoor) one-way functions by showing that random padding, in particular, renders the cryptosystem vulnerable to those kind of attacks. We also provide a weaker notion, *forgeability* for one-way functions, that captures the case where the adversary only sets up the public parameters for which she keeps a backdoor.
- Our main goal is to provide defending mechanisms against kleptographic attacks. We construct *unforgeable* (trapdoor) one-way functions via a general transformation that “sanitizes” arbitrary OWFs by *randomizing the function index*. This transformation clips out all potential correlation between the function and the possible backdoor that the adversary may possess. Additionally, we introduce a *split-program* strategy to make the general method above applicable using only standard hash functions. In the split-program model, the function generation algorithm is composed of two parts: a (randomized) randomness generation algorithm RG that outputs a uniform bit string,

¹This paper won the 2015 PET award.

and a (deterministic) function generation algorithm dKG that converts such random string into the function index. Remark that our results even allow the sanitizing algorithm to be implemented by the adversary.

- Having unforgeable (trapdoor) one-way functions constructed, we next show the power of the primitives. In Section 4, we observe that unforgeable trapdoor one-way functions immediately give us a way to construct key generation algorithms (for digital signature schemes) against kelpographic attacks. We then showcase a concrete example of digital signature scheme in the complete subversion model. More concretely, we achieve this result by (1) using the unforgeable trapdoor one way permutation directly as a key generation algorithm, and then (2) instantiating the unique signature mechanism with the full domain hash. Previously, Bellare et al [2] demonstrated that unique signature scheme is secure against kelpographic attacks, *but assuming the key generation algorithm is honest*. Our result is the first digital signature scheme allowing the adversary to sabotage all algorithms.
- In Section 5, we investigate how to construct backdoor-free PRGs. Previously, Dodis et al. [7] investigated “backdoored PRG” in which the adversary sets up a PRG instance (i.e., the public parameter), and is able to distinguish the output from uniform with a backdoor. They then proposed immunizing strategies that one applies a keyed hash function to the output, *but assuming the key is uniformly generated, and is not known to the adversary* in the public parameter generation phase.

We construct backdoor-free PRGs in the complete subversion model. Our first construction is based on the classic Blum-Micali using our strongly unforgeable OWF and the Goldreich-Levin hardcore predicate [8]. In addition, in [7], Dodis et al show that it is impossible to have a public immunizing strategy for all PRGs by applying a public function to the PRG output. This is because there always exists a PRG (having the immunizing function built in) that reveals the secret to the adversary bit by bit in each iteration. We go beyond their impossibility via an alternative public immunizing strategy. Instead of randomizing the output of the PRG as considered in [7], we randomize the public parameter of PRG, which gives us another construction for PRG in the complete subversion model.

1.2 Related work

The concept of kleptography, subverting cryptographic algorithms by modifying its implementation to leak secrets covertly, was proposed by Young and Yung [19, 20]. They showed concrete examples that backdoors can be embed into the public keys of commonly used schemes; while the public keys look as normal to the users, the adversary is nevertheless capable of learning the secret keys. It may not be surprising that defending against such deliberate attacks is challenging and only limited feasibility results exist. We next briefly describe these existing results.

In [11], Juels and Guajardo suggested the following idea: the user and a trusted certificate authority (CA) jointly generate the public key, and the user has to prove to the CA that the public key is generated honestly. While in our paper, in contrast, the user does not have any secret, and every component is provided by the big brother.

Bellare et al. considered special cases of kelpographic attacks (without necessarily requiring protection from reverse engineering), that they call algorithm substitution attacks on symmetric key encryption [2] and public key encryption [1]. They first proposed a generic attack: a sabotaged randomized algorithm can leak the secret bit-by-bit by sampling biased output; the adversary who knows the backdoor can identify the leaked bits from the biased output. The attack is successful due to the effectiveness of covert subliminal channels [10, 17, 18] built using the backdoor. They then introduced a defending mechanism of using algorithms that has unique output for each input; Good examples of such algorithms include unique

ciphertext encryption algorithm (of encryption scheme), and unique signature generation algorithm of (digital signature scheme). Note that their defending mechanism implicitly assumes the key generation algorithm to be honest. We, on the other hand, cope with the issue left unfinished in their work, and deal with attacks under the generally accepted setting that a proper defense against kleptographic attacks should allow *all* the underlying algorithms to be subverted, and we give an affirmative answer under this *complete subversion model* (i.e., allow also the key generation algorithm to be backdoored).

Dodis et al [7] studied another type of special kleptographic attacks on pseudorandom generators, to formalize the study of the notorious Dual_EC PRG subversion [5, 15]. In their model, the adversary sets the public parameter and keeps some backdoor information (instead of providing an implementation), to break the PRG security. They first proved the equivalence of backdoored PRG and public key encryption with pseudorandom ciphertext. Then they proposed immunizing strategies that one applies a keyed hash function to the output. Note that the key is selected uniformly and not known to the adversary during the public parameter generation phase, thus in our terminology, they are not shown to be secure in the *complete subversion model*. Our results of strongly unforgeable OWF can be applied to construct a specific “backdoor-free” PRG following the classic Blum-Micali framework. Also our general immunizing strategy that randomizing the public parameter of backdoored PRG instead of randomizing PRG output is indeed effective. This indeed allows us to bypass an impossibility result that Dodis et al showed for general public immunization on the PRG output.

Other works suggest different angles of considerations about defending against mass surveillance. For example, in [14], the authors proposed a general framework of safeguarding the protocols by randomizing the incoming and outgoing messages via a trusted (reverse) firewall. Their results demonstrate that with a clean trusted random source, many tasks become reachable. Again, they rely on the firewall not to be subverted, thus not in the *complete subversion model*.

2 Kleptographic Attacks and the Complete Subversion Model

In this section, we explain the kleptographic attacks in more detail and introduce the complete-subversion model.

Classical cryptography assumes that all the algorithms are faithfully implemented. In reality, these algorithms could be implemented by an adversary (e.g., the Big brother); this potentially allows the adversary to gain advantage to influence the input/output behaviors of certain algorithms maliciously, and/or learn secret information that is not supposed to be leaked thru the algorithms. Kleptography studies how to apply these attacks on real-world cryptosystems but under the condition that the attacks will not be detected. Apparently, even the Big brother wants to “monitor” the system, he does not expect to get himself exposed (to get pressure from the whole society). Technically, for every cryptographic scheme, there exists a specification that is rigorously analyzed (proven secure) and designed by the experts; At the same time, the experts actively check whether implementations are honestly generated. Therefore, the adversary expects the output distributions of the subverted implementations should be computationally indistinguishable according to any publicly known input distributions.

Efforts have been made to defend against kleptographic attacks. However, the state-of-the-art of defending mechanisms are for limited scopes. For example, the key generation is assumed to be honest in [2]. Indeed, key generation, as shown in the original papers of Young and Yung, [19, 20], can be subjected to kleptographic attacks. This bring our attention to the feature (which makes the attack more devastating) of kleptography that the subversion attacks can be launched to *any* cryptographic algorithms. They include, e.g., the key generation algorithm, and even the defending algorithm (the immunizing strategy requires a honestly generated uniform seed and unknown to the adversary when she implements the algorithms [7].). This naturally motivates us to study cryptography in a model that

all algorithms are provided by the adversary (thus can be sabotaged); we name the model, **complete subversion model**. We will define security for each primitive we consider in this complete subversion model. More importantly, we make our first attempt to provide general defending mechanisms, i.e., preserve security of the underlying cryptographic primitive, in the setting that all algorithms can be subverted. We call this general paradigm **cliptography**.

Detection principle. As we pointed out above, the adversary in order to avoid being detected, provides an implementation that is consistent with the specification. In particular, for deterministic algorithms, and for each *public* input distribution, one can simply compute the function on given inputs, and compare with the outputs generated by the specification.

In more detail, for a deterministic algorithm F with its specification F_{SPEC} , a public input distribution \mathcal{D} , the testing simply generates samples x_1, \dots, x_q from \mathcal{D} independently for arbitrary polynomially large q , and test for each x_i , if it holds that $F(x_i) = F_{\text{SPEC}}(x_i)$. From the adversary's point of view, if she implements F in a way with a probability $1/p(\lambda)$ (for some polynomial p , and the security parameter λ), \mathcal{D} samples an input x such that $F(x) \neq F_{\text{SPEC}}(x)$, then there exists a PPT tester can notice this inconsistency using about $O(p^2(\lambda))$ samples. Thus, we can safely claim that for any deterministic algorithms with a public input distribution \mathcal{D} , the adversary has to faithfully implement the specification with an overwhelming probability.

Therefore, in our complete subversion model, we need to consider all *randomized* algorithms and deterministic algorithms with *hidden* input distributions chosen by the adversary to be subverted.

In this paper, when the deterministic algorithms are used, they indeed have a public input distribution. For the ease of presentation, we will use the specification directly (or omit those parts) for such deterministic algorithms in the definitions. We remark that our results are still in the complete subversion model even with those simplifications in the definition.

3 Constructing One-Way Functions in the Complete Subversion Model

3.1 Formalizing kleptographic attacks on one way functions

Before studying the security for one-way functions (OWF) and trapdoor one-way functions (TDOWF) in the presence of kleptographic attacks, we first recall the conventional definitions of OWF and TDOWF.

One-way function (OWF). A function family $\mathcal{F} = \{f_i : X_i \rightarrow Y_i\}_{i \in I}$ is *one-way* if there are PPT algorithms (KG, Eval) so that (i) KG, given a security parameter λ , outputs a function index i from $I_\lambda = I \cap \{0, 1\}^\lambda$; (ii) for $x \in X_i$, $\text{Eval}(i, x) = f_i(x)$; (iii) \mathcal{F} is one-way; that is, for any PPT algorithm \mathcal{A} , it holds that $\Pr[\mathcal{A}(i, y) \in f_i^{-1}(y) \mid i \leftarrow \text{KG}(\lambda); x \leftarrow X_i; y := f_i(x)] \leq \text{negl}(\lambda)$.

Trapdoor one way function (TDOWF). A function family $\mathcal{F} = \{f_i : X_i \rightarrow Y_i\}_{i \in I}$ is *trapdoor one-way* if there are PPT algorithms (KG, Eval, Inv) such that (i) KG, given a security parameter λ , outputs a function index and the corresponding trapdoor pair (i, t_i) from $I_\lambda \times T$, where $I_\lambda = I \cap \{0, 1\}^\lambda$, and T is the domain of t_i ; (ii) $\text{Eval}(i, x) = f_i(x)$ for $x \in X_i$; (iii) \mathcal{F} is one-way; and (iv) it holds that $\Pr[\text{Inv}(t_i, i, y) = x \mid i \leftarrow \text{KG}(\lambda); x \leftarrow X_i; y := f_i(x)] \geq 1 - \text{negl}(\lambda)$.

We note that, (trapdoor) one-way permutations can be defined similarly by setting $Y_i := X_i$. For simplicity, we often ignore Eval; for evaluating function with index i on input x , i.e., $\text{Eval}(i, x)$, we write it as $f_i(x)$.

3.1.1 Strong-forgeability and unforgeability for OWF/TDOWF

In this subsection, we define *strong-forgeability* to capture kleptographic attacks on one-way functions. Note that the complemented security notion, *unforgeability*, provides the guarantee that OWF is immune

to kleptographic attacks. Similarly, we can define strong-forgeability and the complemented notion, unforgeability, for trapdoor OWF. Next, let’s start with expressing the intuition of capturing kleptographic attacks on one-way functions.

We begin with a “laboratory specification” version of the OWF, $(\text{KG}_{\text{SPEC}}, \text{Eval}_{\text{SPEC}})$, which has been rigorously analyzed and certified (e.g., by the experts in the cryptography community).

The adversary then provides an alternate implementation. Note that Eval is deterministic on publicly known input distribution, and it is easy to test the correctness, thus we can consider it as honestly implemented (see detection principle in section 2). We therefore focus on KG (the algorithm which generates a function name). The goal of the adversary is to privately maintain some “backdoor information” z so that the subverted implementation of KG will output functions that can be inverted using z . In addition, the adversary must be sure that the output distributions of $\text{KG}(z)$ and that of specification function generation algorithm KG_{SPEC} are computationally indistinguishable, to avoid detection. Formally, we define *strongly-forgeable OWFs*; we note that this immediately allows us to define the complemented notion, *unforgeable OWFs*.

Definition 3.1. A one-way function family $\mathcal{F} = \{f_i : X_i \rightarrow Y_i\}_{i \in I}$ (with the specification function generation algorithm KG_{SPEC}) is δ -**strongly forgeable** if there exist PPT algorithms $(\text{BG}, \text{KG}, \text{Inv})$ so that given a backdoor z produced by backdoor generation algorithm BG, i.e., $z \leftarrow \text{BG}(\lambda)$, the function generation algorithm KG generates a function index i that is (1) invertible given z , and (2) computationally indistinguishable from that generated by the specification function generation algorithm KG_{SPEC} . That is, for every $z \leftarrow \text{BG}(\lambda)$, it holds that:

$$(1) \Pr[x' = x \mid i \leftarrow \text{KG}(\lambda, z); x \leftarrow X_i; y := f_i(x); x' \leftarrow \text{Inv}(z, i, y)] \geq \delta$$

$$(2) \{i \mid i \leftarrow \text{KG}(\lambda, z)\} \stackrel{c}{\approx} \{i \mid i \leftarrow \text{KG}_{\text{SPEC}}(\lambda)\}$$

Correspondingly, we say a one-way function family is **unforgeable** if it is not δ -strongly forgeable for any non-negligible function δ .

The notion of strongly-forgeable OWF is closely related to the conventional notion, TDOWF. See the following lemmas. These results state that if we want to use public key cryptography, we have to accept the possibility of kleptographic attacks on OWFs. For detailed proofs of the lemmas, we defer to Appendix A.

Lemma 3.1. A $(1 - \epsilon)$ -strongly forgeable OWF family is also a TDOWF family, where ϵ is a negligible function of the security parameter.

Next, we show how to construct a strongly-forgeable OWF from any TDOWF. This substantiates the folklore knowledge that sufficient random padding can render cryptosystems vulnerable to backdoor attacks, e.g., [19, 20]. Specifically, the random padding in the malicious implementation can be generated so that it encrypts the corresponding trapdoor using the backdoor as a key.

Lemma 3.2. One can construct a $(1 - \epsilon)$ -strongly forgeable OWF from a TDOWF, where ϵ is a negligible function on the security parameter.

The notions of strongly-forgeable OWF and TDOWF are similar in the sense that they both posit a secret that enables inversion of the OWF. Observe, however, that a strongly-forgeable OWF has a further (critical) property: the distribution of function names, is indistinguishable from a particular specification distribution. The same principle yields a notion of strongly-forgeable TDOWF. Moreover, the adversary can invert using the backdoor, without referring to the regular trapdoor. The formal definition is presented below.

Definition 3.2. A trapdoor one-way function family $\mathcal{F} = \{f_i : X_i \rightarrow Y_i\}_{i \in I}$ (with the specification algorithms $\text{KG}_{\text{SPEC}}, \text{Inv}_{\text{SPEC}}$ ²) is **δ -strongly forgeable** if there exist PPT adversarial algorithms $\mathcal{A} = (\text{BG}, \text{KG}, \text{Inv})$ so that given a backdoor z produced by backdoor generation algorithm BG , the function generation algorithm KG generates a function index i and the corresponding trapdoor t_i , with the following properties: (1) f_i is invertible given z without providing t_i , and (2) the output of KG is computationally indistinguishable from that generated by the specification function generation algorithm KG_{SPEC} . That is, for every $z \leftarrow \text{BG}(\lambda)$, it holds that

$$(1) \Pr[x' = x \mid (i, t_i) \leftarrow \text{KG}(\lambda, z); x \leftarrow X_i; y := f_i(x); x' \leftarrow \text{Inv}_{\text{SPEC}}(z, i, y)] \geq \delta$$

$$(2) \{(i, t_i) \mid (i, t_i) \leftarrow \text{KG}(\lambda, z)\} \stackrel{c}{\approx} \{(i, t_i) \mid (i, t_i) \leftarrow \text{KG}_{\text{SPEC}}(\lambda)\}$$

Correspondingly, we say a trapdoor one-way function family is **unforgeable** if it is not δ -strongly forgeable for any non-negligible function δ .

3.1.2 Forgeability and strong-unforgeability for OWF

The notion of strong forgeability models an attack where the adversary may provide a subverted implementation of the defining algorithms. In many cases, it may also be interesting to consider a weaker form of attack that the adversary simply provides the function index, this is similar to the notion of backdoored PRG that the adversary sets up the public parameters (as in the Dual_EC PRG example [5]). This is a weaker definition in that the adversary only have to generate a backdoor that can be helpful for inverting one single function instead of a family of functions (as in the case of strongly-forgeable OWF). It is easy to see that this notion is equivalent to the standard notion of TDOWF if the KG_{SPEC} outputs the same distribution of the function index as the regular generation algorithm of the TDOWF, thus we only consider this notion for OWF. However, putting into the context of kleptography, it is suggested that when the experts recommend standard for OWF, TDOWF should not be a good candidate.

More importantly, we will later consider unforgeability for OWF, thus the reverse of the weaker notion will yield a stronger definition for unforgeability, and we will give a generic construction that converts any OWF into a strongly unforgeable OWF, which means we can destroy the trapdoor structure generically.

Definition 3.3. A one-way function family $\mathcal{F} = \{f_i : X_i \rightarrow Y_i\}_{i \in I}$ (with the specification of function generation algorithm KG_{SPEC}) is **δ -forgeable** if there exist PPT adversarial algorithms $\mathcal{A} = (\text{KG}, \text{Inv})$ so that

$$1. \Pr[x' = x \mid (i, z) \leftarrow \text{KG}(\lambda); x \leftarrow X_i; y \leftarrow f_i(x); x' \leftarrow \text{Inv}(z, i, y)] \geq \delta$$

2. The distribution of i is indistinguishable from the output distribution of KG_{SPEC} .

Correspondingly, we say a one-way function family is **strongly-unforgeable** if it is not δ -forgeable for any non-negligible function δ .

3.2 Eliminating backdoors

In this section, we discuss methods for safeguarding OWF generation against kleptographic attacks. We first present a general approach that immunizes any OWF generation procedure. We prove that hashing the function index is sufficient to eliminate potential backdoor information. However, in many cases of interest, function indices have specific algebraic structure. It is not clear in general how one can guarantee

²Following the detection principle, the inversion algorithm of the TDOWF is deterministic and with a public input distribution, thus we omit the indistinguishability condition for this algorithm, considering it is honestly implemented as Inv_{SPEC} , and focus only on KG .

that a public hash function has such “structure preserving” properties. In order to apply our approach in more realistic settings, we propose a “split-program” model in which the function generation algorithm is necessarily composed of two parts: a random string generation algorithm RG that outputs random bits r , and a deterministic function index generation algorithm dKG which uses r to generate the index.

Before going to the details of our feasibility results, we remark that the definitions of (strongly) unforgeable OWFs can be found in Definitions 3.1 and 3.3. The definitions essentially insist the sabotaged KG looks the same as the specification KG_{SPEC} which is rigorously analyzed. To put it another way, invertibility in the backdoored mode will be restricted by the indistinguishability condition for the function index.

3.2.1 General feasibility results

We will show below that randomizing the function index (that is, the relationship between names and functions) can provide satisfactory immunization against possibly sabotaged function generation. The intuition behind this idea is that according to the definition of forgeable OWF, each backdoor that frequently appears can only be useful for inverting a sparse subset of one way functions (i.e., the range of $\text{KG}(z)$ is exponentially sparse for every z , otherwise, one can use such backdoor to break the one-wayness of the functions generated by KG_{SPEC}). Thus, randomizing the function index will map the function index to a “safe” domain, and destroys the possible correlation with any selected backdoor. That said, it is difficult for the adversary to arrange a backdoor that works for a disturbed subset of function index (after hashing), even if she knows the immunizing strategy.

Constructing strongly-unforgeable OWFs. Given any OWF family $\mathcal{F} := (\text{KG}_{\mathcal{F}}, \text{Eval}_{\mathcal{F}}) := \{f_i\}_{i \in I}$ (which might be forgeable) that is secure if $\text{KG}_{\text{SPEC}}(\lambda)$ outputs uniform i from I_λ , we will construct a strongly unforgeable OWF family $\mathcal{G} := (\text{KG}_{\mathcal{G}}, \text{Eval}_{\mathcal{G}}) = \{g_i\}$. We assume a (public) hash function family $H_\lambda = \{h_\lambda : I_\lambda \rightarrow I_\lambda\}$ modeled as a random oracle (for each λ), and that h is randomly sampled from H_λ . The key generation algorithm $\text{KG}_{\mathcal{G}}$ is given as $(h, \text{KG}_{\mathcal{F}})$ and is defined as $h \circ \text{KG}_{\mathcal{F}}$,³ and $\text{Eval}_{\mathcal{G}}$ is the same as $\text{Eval}_{\mathcal{F}}$. (thus for each i , $g_i(\cdot) = f_{h(i)}(\cdot)$). See also the pictorial illustration in Fig 1.




Figure 1: Immunization strategy for OWF.

Remark 3.4. *In the case that we can build hash functions that hash directly onto the index space I_λ , then we can use the above immunization strategy. However, in practice, it is not clear if we can always easily build hash functions on index space. To address this issue, in Section 3.2.2, we introduce a new framework called “split-program model”. There, we can use standard hash function, e.g., SHA-256 to work on the random bits directly.*

Theorem 3.5. *The OWF family \mathcal{G} defined above is strongly unforgeable in the random oracle model.*

Proof. Suppose that \mathcal{G} is δ -forgeable for a non-negligible function δ , and let $(\mathcal{A}_{\text{KG}}, \mathcal{A}_{\text{Inv}})$ be the adversarial algorithms of Definition 3.3. We will construct a simulator \mathcal{S} which will break the one-way security of \mathcal{F} .

Suppose $(f_{\tilde{i}}, y)$ are the challenges \mathcal{S} receives from the one way security challenger \mathcal{C} , where $y = f_{\tilde{i}}(x)$ for a randomly selected x . Then (i.) \mathcal{S} first randomly samples a bit b to decide whether to embed j into

³ Note that the hash function can be implemented by the adversary, since it is deterministic and with a public input distribution (the output distribution of $\text{KG}_{\mathcal{F}}$). We can also interpret the function generator of the strongly-unforgeable OWF \mathcal{G} as a single algorithm $h \circ \text{KG}_{\mathcal{F}}$, and then we insist it outputs both i, \tilde{i} , for checking using the specifications h_{SPEC} and KG_{SPEC} .

the answers to the random oracle queries from \mathcal{A}_{KG} or \mathcal{A}_{Inv} . (W.l.o.g., we assume all the random oracle queries are different.) (ii.) \mathcal{S} runs \mathcal{A}_{KG} . Suppose \mathcal{A}_{KG} makes q_1 random oracle queries $Q_1 = \{i_1, \dots, i_{q_1}\}$. If $b = 0$, \mathcal{S} randomly selects an index $t_1 \in \{1, \dots, q_1\}$. When answering the random oracle queries i_1, \dots, i_{q_1} from \mathcal{A}_{KG} , \mathcal{S} answers \tilde{i} for $h(i_{t_1})$; for all other queries, \mathcal{S} answers with random elements from the index set I_λ . If $b \neq 0$, \mathcal{S} answers all these queries using random elements from I_λ . \mathcal{S} maintains a list for the query-answer pairs. \mathcal{A}_{KG} outputs a pair (i, z) .

If $[b = 0 \wedge i \neq i_{t_1}]$, \mathcal{S} aborts; otherwise, \mathcal{S} runs \mathcal{A}_{Inv} with inputs (i, y, z) . Assuming \mathcal{A}_{Inv} asks q_2 random oracle queries, \mathcal{S} sets \tilde{i} as $h(i)$ (even if i is not asked) and for all others queries, \mathcal{S} answers with random elements from I_λ . \mathcal{A}_{Inv} outputs x' .

If $[b = 1 \wedge i \in Q_1]$, \mathcal{S} aborts; otherwise, it returns x' as his answer to \mathcal{C} .

Probabilistic analysis. Now we bound the success probability of \mathcal{S} . Let us use W to denote the event that \mathcal{S} aborts, \overline{W}_1 to denote the event that $b = 0 \wedge i \neq i_{t_1}$, and W_2 to denote the event that $b = 1 \wedge i \in Q_1$. We have

$$\Pr[x' = x] = \Pr[x' = x|W] \Pr[W] + \Pr[x' = x|\overline{W}] \Pr[\overline{W}] \geq \Pr[x' = x|\overline{W}] \Pr[\overline{W}]$$

We first bound $\Pr[\overline{W}]$ as $1 - \Pr[W]$, we have $\Pr[W] = \Pr[W_1 \vee W_2] \leq \Pr[W_1] + \Pr[W_2]$. Assuming $\Pr[i \in Q_1] = \eta$, we bound $\Pr[W_1]$ as follows:

$$\begin{aligned} \Pr[W_1] &= \Pr[b = 0 \wedge i \neq i_t] = \Pr[b = 0] \Pr[i \neq i_{t_1}] \\ &= \frac{1}{2} (\Pr[i \neq i_{t_1} | i \in Q_1] \Pr[i \in Q_1] + \Pr[i \neq i_{t_1} | i \notin Q_1] \Pr[i \notin Q_1]) \\ &= \frac{1}{2} \left[\left(1 - \frac{1}{q_1}\right) \eta + (1 - \eta) \right] \end{aligned}$$

While $\Pr[W_2] = \Pr[b = 1] \Pr[i \in Q_1] = \eta/2$, we have: $\Pr[W] \leq \frac{1}{2} \left(1 - \frac{1}{q_1}\right) \eta + \frac{1}{2} \leq 1 - \frac{1}{2q_1}$. Thus we can derive that $\Pr[\overline{W}] \geq 1/(2q_1)$.

Furthermore, conditioned on \mathcal{S} not aborting, the input distributions of the adversaries (\mathcal{A}_{KG} , \mathcal{A}_{Inv}) are identical to that of Definition 3.3. By definition of δ -forgeability, $\Pr[\mathcal{A}_{\text{Inv}}(\tilde{i}, f_{\tilde{i}}(x), z) = x] \geq \delta$, thus $\Pr[x' = x|\overline{W}] \geq \delta$.

Combing these facts, we conclude that if \mathcal{G} is δ -forgeable for some non-negligible δ , then there exists an algorithm \mathcal{S} that breaks the one-way security of \mathcal{F} with probability at least $\delta/(2q_1)$ (which is non-negligible). This completes the proof. \square

Remark 3.6. We may also prove a similar result in the standard model, that the random oracle is instantiated with a pseudorandom function PRF. However, in this case, the adversary KG can only have oracle access to the PRF, but the Inv can have full access to the PRF key. Still, this result will not be in the complete subversion model that the PRF requires a trusted key. Thus we leave as an open problem that how to establish a general immunizing result in the standard model. Furthermore, the above proof works even if the distribution of i is different with the output distribution of KG_{SPEC} , as long as it still belongs to the index set.

3.2.2 Practical results in the split-program model

Indices (names) of a one-way function family may have structure. For example, for OWF based on discrete logarithm, $f_{g,p}(x) = g^x \bmod p$, the function index consists of an algebraically meaningful pair (p, g) , where p is a prime and g a random generator. This would require that the hash function in the general immunization method above maps (g, p) to (g', p') ; note that (g', p') is another algebraically meaningful pair with the same structure. Furthermore, for a TDOWF, the hash function needs to map the

function/trapdoor pair to another function/trapdoor pair. It is not clear in general how one can guarantee that a public hash function has such “structure preserving” properties.

To address this problem, we propose a split-program model in which every function generation algorithm is composed of two algorithms, a random string generation algorithm RG that outputs a uniform ℓ -bit random string r , and a deterministic function index generation algorithm dKG that transforms the randomness r into a function index i . In this model, as dKG is deterministic with a public input distribution, we can consider it to be honestly implemented and we can focus on “cleaning up” the randomness generated by RG.⁴

Remark 3.7. *It is not hard to see, the split-program model is quite general and can be applied to most practical algorithms. To see this, the user gets the source code of the implementation, which makes calls to some API for generating randomness (e.g., `rand()`) whenever necessary. The user can hook up the interface with the calls to the API with the separate program RG provided by the big brother. In principle, one can always augment a randomized KG algorithm to output the function index i together with the randomness r used to generate i , thus the RG can be implemented as this augmented KG, and discards the function index i from the output.*

We first rephrase the standard OWF/TDOWF definitions in the split-program model.

Definition 3.8. *A function family \mathcal{F} is one way in the split-program model if there exist a pair of algorithms (RG, dKG) where (i.) RG, given a security parameter λ , outputs a uniform $\ell(\lambda)$ -bit string r ; (ii.) dKG is deterministic: given the randomness r it outputs a function index $i \in I_\lambda$; and (iii.) \mathcal{F} is one-way under this procedure for generating i .*

Similarly, we can define a TDOWF family in the split program model (In this case, dKG outputs a function index together with a trapdoor). For the ease of presentation, we often use the pair of algorithms (RG, dKG) to represent the OWF/TDOWF family \mathcal{F} in the split-program model. Next we define δ -forgeable OWF in the split-program model by modifying Definition 3.3; we immediately have the complemented security notion of strongly-unforgeable OWF in the split-program model.⁵

Definition 3.9. *A one-way function family $\mathcal{F} = \{f_i : X_i \rightarrow Y_i\}_{i \in I}$ (with the specification version of function generation algorithm $\text{RG}_{\text{SPEC}}, \text{dKG}_{\text{SPEC}}$) is **δ -forgeable** if there exist PPT algorithms (RG, dKG, Inv) such that:*

1. $\Pr[x' = x \mid (r, z) \leftarrow \text{RG}(\lambda); i \leftarrow \text{dKG}_{\text{SPEC}}(\lambda, r, z); x \leftarrow X_i; y := f_i(x); x' \leftarrow \text{Inv}(z, i, y)] \geq \delta$.
2. *The distribution of r is indistinguishable from the output distribution of RG_{SPEC} .*

*Correspondingly, we say a one-way function family is **strongly-unforgeable** in the split-program model if it is not δ -forgeable in the split-program model for any non-negligible function δ .*

⁴ The split-program model essentially forces the adversary to concentrate the parts which may potentially contain backdoors of a malicious implementation into RG. This gives us more flexibility to apply the sanitizing strategy. Furthermore, conceptually, the immunizing strategy itself (e.g., Fig 1) can be seen as a bigger piece of implementation in the split-program model, i.e., the hash function and the actual KG algorithm should be individually implemented and checked.

⁵Note that it is easy to provide a more general definition to capture the fact dKG is also implemented by the adversary, we can simply require the indistinguishability condition hold for the joint output distribution of RG, dKG. However as pointed out in section 2, it is fine for us to consider them as honestly implemented for simplicity, as it is deterministic and with a public input distribution.

Strongly-unforgeable OWF in the split-program model. Given a OWF family $\mathcal{F} := (\text{RG}_{\mathcal{F}}, \text{dKG}_{\mathcal{F}}, \text{Eval}_{\mathcal{F}})$ (in the split program model) whose RG_{SPEC} outputs uniform bits, and a public hash function $h(\cdot)$ randomly selected from a hash family $H : \{0, 1\}^{\ell(\lambda)} \rightarrow \{0, 1\}^{\ell(\lambda)}$ which is modeled as a random oracle, we construct an strongly-unforgeable OWF family \mathcal{G} . $\mathcal{G} := (\text{RG}_{\mathcal{G}}, \text{dKG}_{\mathcal{G}}, \text{Eval}_{\mathcal{G}})$ can be described as $(h \circ \text{RG}_{\mathcal{F}}, \text{dKG}_{\mathcal{F}}, \text{Eval}_{\mathcal{F}})$, i.e., the function index i is generated by $\text{dKG}_{\mathcal{F}}$ using $h(r)$ as randomness, where $r \leftarrow \text{RG}_{\mathcal{F}}$. See Figure 2 below:⁶

$$\boxed{\text{RG}_{\mathcal{F}}} \rightarrow r \quad r \rightarrow \boxed{h} \rightarrow \tilde{r} \quad \tilde{r} \rightarrow \boxed{\text{dKG}_{\mathcal{F}}} \rightarrow \tilde{i} \quad x \rightarrow \boxed{\text{Eval}_{\mathcal{F}}(\tilde{i}, \cdot)} \rightarrow y$$

Figure 2: Immunization strategy for OWF in the split-program model.

The intuition is analogous to that captured in the proof of Theorem 3.5: For a backdoor z that frequently appears in the output of \mathcal{A}_{RG} , the “bad” set of randomness (i.e., for which the dKG algorithm outputs a function that can be efficiently inverted using z) should be sparse in $\{0, 1\}^{\ell(\lambda)}$. (Otherwise, one can break the one way security of \mathcal{F} by simply running \mathcal{A}_{KG} to get z .) In this case, the probability that a uniform random string falls into the “bad” set that z is useful for inverting is negligible. Hashing the random bits will then break the delicate connection between the backdoor and the function index that will be generated using the cleaned randomness. Specifically, it will be challenging for the adversary to design an efficient connection between a backdoor z and the “scrambled” sets of functions backdoored by z .

Theorem 3.10. *The OWF family \mathcal{G} described above is strongly unforgeable in the split-program model if H is modeled as a random oracle.*

Proof. Assume \mathcal{G} is δ -forgeable in the split program model, i.e., there exist PPT adversaries $(\mathcal{A}_{\text{RG}}, \mathcal{A}_{\text{dKG}}, \mathcal{A}_{\text{Inv}})$ satisfying Definition 3.9. Then we can construct a simulator \mathcal{S} that breaks the one-way security of \mathcal{F} in the split-program model.

Suppose r^* is the randomness and $y = f_i(x)$ is the challenge value (for a randomly chosen x) received from the one-way security challenger \mathcal{C} , where $i = \text{dKG}(r^*)$.

\mathcal{S} first chooses a random bit b and runs \mathcal{A}_{RG} , which asks random oracle queries $\{r_1^0, \dots, r_{q_0}^0\} := Q_0$. If $b = 0$, \mathcal{S} randomly selects $t_0 \in \{1, \dots, q_0\}$, and answers $h(r_{t_0}^0) = r^*$; all other queries are answered with uniform ℓ -bit strings. If $b \neq 0$, \mathcal{S} answers all queries with random strings. \mathcal{S} maintains a list for the query-answer pairs.

If $[b = 0 \wedge r \neq r_{t_0}^0]$, \mathcal{S} aborts; otherwise, it sets $h(r) = r^*$ and it runs \mathcal{A}_{KG} with inputs r^* , and expects \mathcal{A}_{KG} to output i . \mathcal{S} then runs \mathcal{A}_{Inv} with inputs (i, y, z) and receives the response x' .

If $b = 1 \wedge r \in Q_0$, \mathcal{S} aborts; otherwise, it sends x' to the challenger \mathcal{C} as his answer.

Similar to the proof of Theorem 3.5, let \overline{W} denote the event that \mathcal{S} aborts: we can bound $\Pr[\overline{W}] \leq 1 - \frac{1}{q_0}$. Thus,

$$\Pr[x' = x] \geq \Pr[x' = x | \overline{W}] \Pr[\overline{W}] \geq \frac{\Pr[x' = x | \overline{W}]}{2q_0}.$$

While following Definition 3.9,

$$\Pr[x' = x | \overline{W}] = \Pr[\mathcal{A}_{\text{Inv}}(z, i, f_i(x)) = x] \geq \delta.$$

To summarize, in the split-program model, if \mathcal{G} is δ -forgeable, \mathcal{S} will break the one way security of \mathcal{F} with probability at least $\delta/(2q_0)$. \square

⁶We emphasize that in the split-program model, we can use any regular hash function such as SHA-256.

Constructing unforgeable TDOWFs in the split-program model. More interestingly, we can apply the same method to immunize a TDOWF in the split-program model, whose RG_{SPEC} outputs uniform bits r and whose (deterministic) dKG_{SPEC} , given r , outputs a function index and a trapdoor pair. Similar to the OWF case, hashing the randomness ensures that the resulting function index together with the corresponding trapdoor will be “safe.” However, as pointed out in Section 3.1, it is preferable to consider strong forgeability for TDOWF (since a forgeable TDOWF can be seen as itself), thus we adapt Definition 3.2 to the split-program model:⁷

Definition 3.11. A trapdoor one-way function family $\mathcal{F} = \{f_i : X_i \rightarrow Y_i\}_{i \in I}$ (with the specification algorithms $\text{RG}_{\text{SPEC}}, \text{dKG}_{\text{SPEC}}$) is δ -strongly forgeable if there exist PPT adversarial algorithms $\mathcal{A} = (\text{BG}, \text{RG}, \text{dKG}, \text{Inv})$ so that given a backdoor z produced by backdoor generation algorithm BG , the random string generation algorithm RG generates an $\ell(\lambda)$ -bit random string r , and based on such r the deterministic function generation algorithm dKG generates a function index i and the corresponding trapdoor t_i , with the following properties: (1) f_i is invertible given z without providing t_i , and (2) the output distribution of RG is computationally indistinguishable from that generated by the specification algorithms RG_{SPEC} . That is, for every $z \leftarrow \text{BG}(\lambda)$, it holds that:

- (1) $\Pr[x' = x \mid z \leftarrow \text{BG}; r \leftarrow \text{RG}(\lambda, z); (i, t_i) \leftarrow \text{dKG}_{\text{SPEC}}(\lambda, r); x \leftarrow X_i; y := f_i(x); x' \leftarrow \text{Inv}(i, y, z)] \geq \delta$
- (2) $\{r \mid z \leftarrow \text{BG}; r \leftarrow \text{RG}(\lambda, z)\} \stackrel{c}{\approx} \{r \mid r \leftarrow \text{RG}_{\text{SPEC}}(\lambda)\}$

Correspondingly, we say a TDOWF family is **unforgeable** if it is not δ -strongly forgeable for any non-negligible δ .

The immunization methodology for OWF discussed before can also be used for immunizing TDOWF in the split program model. That is, given any TDOWF family \mathcal{F} whose RG_{SPEC} outputs uniform bits r , a hash function h randomly selected from a hash family $H_\lambda : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ can be used to determine an unforgeable TDOWF family \mathcal{G} , see Figure 3 below:

$$\boxed{\text{RG}(z)} \rightarrow r \quad r \rightarrow \boxed{h} \rightarrow \tilde{r} \quad \tilde{r} \rightarrow \boxed{\text{dKG}} \rightarrow (i, t_i) \quad x \rightarrow \boxed{\text{Eval}(i, \cdot)} \rightarrow y$$

Figure 3: Immunization strategy for TDOWF in the split-program model. The modified key generation algorithm c

Theorem 3.12. The TDOWF family \mathcal{G} described above is unforgeable in the split-program model if H is modeled as a random oracle.

The proof of Theorem 3.12 is similar to that of Theorem 3.10; we can think (i, t_i) as an extended form of function index.

4 Constructing Signatures in the Complete Subversion Model

As an immediate application, in this section, we will demonstrate how to use unforgeable OWFs and unforgeable TDOWFs as fundamental building blocks to construct backdoor-free digital signature schemes in the complete subversion model.

We build on recent works of Bellare et al. [1, 2], who studied a family of *algorithm-substitution attacks* (ASAs) on cryptographic primitives. In their ASAs, an encryption algorithm is implemented by a

⁷Again, since dKG is deterministic with public input distributions, for simplicity we consider they are honestly implemented as dKG_{SPEC} , even though they are implemented by the adversary.

saboteur who may have backdoors embedded; then the sabotaged implementation judiciously samples randomness in order to generate the outputs (e.g., ciphertexts) in a way that leaks the secret bit-by-bit via a steganographic channel.

To defend against the ASAs, Bellare et al provided successful mechanisms of using deterministic algorithms (e.g., signing) with unique outputs. However, we note that the ASAs considered in [1, 2] are restricted in the sense that the adversary is not allowed to launch the ASAs on *all* algorithms; In particular, the key generation algorithm is assumed to be implemented honestly and the adversary is not allowed to attack on it.

We are interested in securing cryptographic schemes in the *complete subversion model* where the adversary is allowed to launch algorithm-substitution attacks on *all* components of the schemes. We demonstrate below how our unforgeable TDOWF can be applied to defend against the ASAs, even if the key generation algorithm has been sabotaged. Together with the results of [1, 2], we then show a concrete example of a digital signature scheme, achieving security in the case that *all* the algorithms are sabotaged.

Digital signature scheme in the complete subversion model. We can define the existential unforgeability in the setting of complete subversion⁸ in a way that the public key is generated using the adversarial implementation, and the signing/verification algorithm can also be implemented by the adversary, while all the other steps are the same as the standard definition.

Definition 4.1. We say a signature scheme $(\text{KG}, \text{Sign}, \text{Verify})$ (with the specification version of key generation algorithm KG_{SPEC} , signing algorithm $\text{Sign}_{\text{SPEC}}$ and verification algorithm $\text{Verify}_{\text{SPEC}}$) is **existentially unforgeable in the complete subversion setting**, if for all PPT adversary $\mathcal{A} = (\mathcal{A}_{\text{INITIAL}}, \mathcal{A}_{\text{FORGE}})$, the following properties hold:

- The adversary \mathcal{A} wins the unforgeability game defined below, with no more than a negligible probability.

UNFORGEABILITY GAME between PPT adversary $\mathcal{A} = (\mathcal{A}_{\text{INITIAL}}, \mathcal{A}_{\text{FORGE}})$ and a challenger \mathcal{C} :

1. $\mathcal{A}_{\text{INITIAL}}$ provides an implementation of $\text{KG}, \text{Sign}, \text{Verify}$ to the challenger (which may contain backdoor information z).
2. \mathcal{C} queries KG to learn the key pair (pk, sk) , and sends the public key pk to $\mathcal{A}_{\text{FORGE}}$.
3. $\mathcal{A}_{\text{FORGE}}$, having the backdoor z from $\mathcal{A}_{\text{INITIAL}}$, asks signing queries for arbitrarily chosen messages m_1, \dots, m_q ; for each m_i from $\mathcal{A}_{\text{FORGE}}$, the challenger \mathcal{C} queries Sign with input (sk, m_i) to learn the corresponding signature σ_i , and return σ_i to $\mathcal{A}_{\text{FORGE}}$.
4. $\mathcal{A}_{\text{FORGE}}$ returns a message-signature pair (m^*, σ^*) to the challenger \mathcal{C} ; now \mathcal{C} queries Verify with input (pk, m^*, σ^*) to learn output b .
5. $\mathcal{A} = (\mathcal{A}_{\text{INITIAL}}, \mathcal{A}_{\text{FORGE}})$ wins the game if $m^* \notin \{m_1, \dots, m_q\}$ and $b = 1$.

- In the above game the output distribution of KG, Sign and Verify are indistinguishable from the output distribution of $\text{KG}_{\text{SPEC}}, \text{Sign}_{\text{SPEC}}$ and $\text{Verify}_{\text{SPEC}}$ ⁹ respectively.

⁸This definition may be weaker than the general definition of surveillance security defined in [2] that the output distribution of subverted algorithms looks the same as that of the specification even to the adversary who has the backdoor. However, preserving its standard security in this complete subversion model (we think) is good enough for each primitive to be used as normal.

⁹The checking for the Verify algorithm is a bit subtle since the tester can not generate message-signature pairs himself for a given public key to to run the specification. However, in the security game in definition 4.1, if a deterministic Verify algorithm is consistent with the the specification on the “forgery” returned by the adversary, then we can safely assume it is honestly implemented; otherwise, the challenger can notice the difference.

Construction. Following the result of [2], if the key generation algorithm is honest, a unique signature scheme [9, 13] remains existentially unforgeable against the algorithm substitution attack. In this case, the signing and verification algorithms are both deterministic and the detection condition “forces” the adversary to honestly implement these two. To obtain a *complete* subversion-secure solution from a unique signature scheme, we still need to “upgrade” the key generation algorithm so that it is secure against subversion attack.

We may consider the following approach to obtain a subversion-secure key generation algorithm: initially generate an unforgeable OWF f , and then randomly select a secret key sk and compute the public key $pk = f(sk)$. Since the OWF is unforgeable, it seems that we may be able to force the adversarial key generation algorithm to output a “safe” OWF. However, it is not clear how to show that the key pairs are well distributed.

We here consider an alternative approach; we use an unforgeable TDOWF, $(KG_{\mathcal{F}}, \text{Inv}_{\mathcal{F}})$. More concretely, we use the key generation algorithm of unforgeable TDOWF to generate a function index i together with the corresponding trapdoor t_i , and set the index i as the public key, and the trapdoor t_i as the secret key. To be compatible with the unique signature scheme, we choose to instantiate the unique signature scheme from the full domain hash construction [3, 6]. Details of the construction can be found below:

- Key generation $(pk, sk) \leftarrow \text{KG}(\lambda)$:
compute $(i, t_i) \leftarrow \text{KG}_{\mathcal{F}}(\lambda)$, and set $pk := i$ and $sk := t_i$;
- Signature generation $\sigma \leftarrow \text{Sign}(sk, m)$:
upon receiving message m , compute $\sigma := \text{Inv}_{\mathcal{F}}(sk, i, h(m))$; here $sk = t_i$;
- Signature verification $b := \text{Verify}(pk, m, \sigma)$:
upon receiving message-signature pair (m, σ) , if $f_i(\sigma) = h(m)$ then set $b := 1$, otherwise set $b := 0$; here $pk = i$.

Theorem 4.2. *Given an unforgeable TDOWF \mathcal{F} , the full domain hash signature scheme is existentially unforgeable (under complete subversion) in the random oracle model.*

Proof. First, since the Sign, Verify and the hash algorithms we use are all deterministic, and in the unforgeability game, they are all with public input distributions, as explained in section 2, we can assume they are honestly implemented, i.e., consistent with the specifications.

We will show that if there is a subversion attack, then one can break the unforgeability of \mathcal{F} . Suppose $(\mathcal{A}_{\text{INITIAL}}, \mathcal{A}_{\text{FORGE}})$ are the adversaries who break the existential unforgeability of the full domain hash signature scheme in the complete subversion model with a non-negligible probability δ , we will construct \mathcal{S} to simulate the adversaries $(\text{BG}, \text{KG}, \text{Inv})$ as in Definition 3.2.

BG first runs $\mathcal{A}_{\text{INITIAL}}$ to receive the backdoor z and the implementation KG_0 for the signature key generation. KG simply runs KG_0 and outputs a function index $i := pk$ and the corresponding trapdoor t_i . \mathcal{S} discards t_i and sends pk to $\mathcal{A}_{\text{FORGE}}$. Suppose $y = f_i(x)$ is the challenge that Inv receives from the unforgeable TDOWF challenger, Inv (constructed by \mathcal{S}) will feed $\mathcal{A}_{\text{FORGE}}$ with z and runs it as follows to try to invert.

W.l.o.g, we assume $\mathcal{A}_{\text{FORGE}}$ asks a random oracle query for a message before she asks for the signing query, (if not, the simulator can ask instead of her), and $\mathcal{A}_{\text{FORGE}}$ asks the random oracle query for her final forgery m^* .

Suppose $\mathcal{A}_{\text{FORGE}}$ asks q random oracle queries m_1, \dots, m_q . Inv randomly choose $j \in \{1, \dots, q\}$, and answers the query with $y = h(m_j)$; He then chooses $q - 1$ random elements $\sigma_1, \dots, \sigma_{j-1}, \sigma_{j+1}, \dots, \sigma_q$, and answers the random oracle queries as $h(m_k) = f_i(\sigma_k)$ for $k \neq j$. Inv maintains a list.

When $\mathcal{A}_{\text{FORGE}}$ asks a signing query m_k , if $m_k = m_j$, Inv aborts, otherwise, Inv checks the list and returns the corresponding σ_k . (Note that Sign algorithm can only have one output for each message, thus the above procedure perfectly simulates the signing oracle implemented by the subverted Sign algorithm). $\mathcal{A}_{\text{FORGE}}$ outputs m^*, σ^* . If $m^* \neq m_j$, Inv aborts, otherwise, Inv outputs σ^* .

Let us use W to denote the event that Inv aborts. Following the classic proof of security of the full domain hash signature scheme (e.g., [6]), $\Pr[W] \leq 1 - \text{poly}(\frac{1}{q})$, thus Inv successfully inverts with probability at least $\delta_0 = \text{poly}(\frac{1}{q})\delta$ (as we assume $\mathcal{A}_{\text{FORGE}}$ has δ advantage, and the Verify has to be honestly implemented to avoid detection). This means \mathcal{F} is δ_0 -forgeable, which contradicts our condition. \square

5 Constructing Pseudorandom Generator in the Complete Subversion Model

As mentioned in the Introduction, our goal is to stimulate a *systematic* study of Cliptography. Having studied the fundamental backdoor-free building blocks, unforgeable OWFs, and unforgeable TDOWFs, we intend to mimic the classic footprints of constructing cryptographic primitives from OWF/TDOWF, and provide solutions to other important backdoor-free building blocks. As our first example, we will show an interesting connection between our notion of unforgeable OWF and the notion of backdoor-free PRG, recently studied by Dodis et al. [7]. Next we first review the basic notions of PRG under subversion attacks. We then provide a specific solution based on the Blum-Micali PRG; this result can be viewed as a mimic of the classic result of Blum-Micali construction in cliptography. Furthermore, we examine how to extend the applicability of our general sanitizing strategy for OWF/TDOWF to more settings. We will demonstrate a general method of public immunizing strategy for PRG. We remark that, all algorithms in our backdoor-free PRG construction, including the sanitizing function (which can be part of the KG algorithm in the specification), can be subverted. Thus we provide the first PRG constructions secure in the complete subversion model.

5.1 Preliminaries: backdoored and backdoor-free PRGs

We adopt the definition from [7], that a pseudorandom generator consists of a pair of algorithms (KG, PRG), where KG outputs a public parameter pk and $\text{PRG} : \{0, 1\}^* \times \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell \times \{0, 1\}^\ell$ takes the public parameter pk and an ℓ -bit random seed s as input; it returns a state $s_1 \in \{0, 1\}^\ell$ and an output string $r_1 \in \{0, 1\}^\ell$. PRG may be iteratively executed; in the i -th iteration, it takes the state from the previous iteration s_{i-1} as the seed and generates the current state s_i and output r_i . We use PRG^q to denote the result of q iterations of PRG with outputs r_1, \dots, r_q (each $r_i \in \{0, 1\}^\ell$).

In a backdoored PRG, the algorithms (in particular KG) are implemented by the adversary (represented by $\mathcal{A}_{\text{INITIAL}}$), outputs a public parameter pk together with a backdoor sk . The output distribution $\text{PRG}(pk, \mathcal{U})$ is still pseudorandom, where \mathcal{U} is the uniform distribution; however, with the corresponding backdoor sk , the adversary (represented by $\mathcal{A}_{\text{DIST}}$) is able to break the PRG security (e.g., the adversary can distinguish the output from a uniform string).

We will rephrase the definition for backdoored PRG in our setting—as in the definition of a forgeable OWF—there exist “specification” versions of the algorithms. In particular, the parameter generation algorithm KG_{SPEC} is with the requirement that the distribution of the adversarially generated public parameter must be indistinguishable from the output distribution of KG_{SPEC} . It is easy to see that the output distribution of $\text{PRG}(pk, s)$ for a uniformly chosen s is pseudorandom even pk is generated by $\mathcal{A}_{\text{INITIAL}}$. (Otherwise, one can easily distinguish the output distribution of $\mathcal{A}_{\text{INITIAL}}$ from KG_{SPEC} .) Additionally, as the PRG algorithm is deterministic, and its input distribution is public, we may assume that the adversary implements it honestly as the PRG_{SPEC} to avoid easy detection so that we can focus on the KG algorithm.

The formal definitions are presented as follows:

Backdoored PRG. We re-phrase the definition of (q, δ) -**backdoored PRG**¹⁰ as follows: We define a backdoored PRG game (see Figure 4) with a PPT adversary $\mathcal{A} = (\mathcal{A}_{\text{INITIAL}}, \mathcal{A}_{\text{DIST}})$ such that (i) the pk distribution is indistinguishable from that generated by KG_{SPEC} ; and (ii) the adversary wins the backdoored PRG game with probability δ , i.e., $\Pr[b = b'] - \frac{1}{2} \geq \delta$.

$$\begin{array}{l}
 (pk, sk) \leftarrow \mathcal{A}_{\text{INITIAL}} \\
 s \leftarrow \{0, 1\}^\ell \\
 r_1^0, \dots, r_q^0 \leftarrow \text{PRG}^q(pk, s) \\
 r_1^1, \dots, r_q^1 \leftarrow \{0, 1\}^{\ell \cdot q} \\
 b \leftarrow \{0, 1\} \\
 b' \leftarrow \mathcal{A}_{\text{DIST}}(pk, sk, r_1^b, \dots, r_q^b)
 \end{array}$$

Figure 4: The backdoored PRG game

Backdoor-free PRG. Then we say that a PRG is q -**backdoor free** if, in the above backdoored PRG game, for all PPT adversaries $\mathcal{A} = (\mathcal{A}_{\text{INITIAL}}, \mathcal{A}_{\text{DIST}})$, whenever pk is indistinguishable from the specification distribution, the advantage is negligible, i.e., $|\Pr[b = b'] - \frac{1}{2}| \leq \text{negl}(\lambda)$.

Remark 5.1. *The generation of the seed s is out of the scope of this paper (same as [7]), since even the specification of pseudorandom generators do not cover this part. Our techniques can guarantee a received implementation is as good as the specification. Of course, it would be an important open question to consider the random seed generation for practice.*

5.2 Constructing backdoor-free PRG from strongly unforgeable OWP

In this subsection, we provide constructions for backdoor-free PRG based on strongly unforgeable one-way permutation. We start with a basic solution based on a (simplified) Blum-Micali PRG, and then extend it to a full-fledged solution. Before going to the details of our constructions, we recall the classic generic construction of Goldreich-Levin (GL), yielding a hardcore predicate [8] for any OWF f . We suppose the input x of f is divided into two halves $x = (x_1, x_2)$ and define the bit $B(x) = \langle x_1, x_2 \rangle$; $B(x)$ is hard to predict given $x_1, f(x_2)$, assuming that f is one-way. Moreover, if there is a PPT algorithm that predicts $B(x)$ with significant advantage δ given $x_1, f(x_2)$, then there is a PPT algorithm I that inverts f with probability $\text{poly}(\delta)$.

Basic construction. We will show that given a strongly unforgeable one-way permutation (OWP) family \mathcal{F} with algorithms $(\text{KG}_{\mathcal{F}}, \text{Eval}_{\mathcal{F}})$, the classic Blum-Micali PRG [4] (using the GL hardcore predicate) is 1-backdoor free. Our basic construction (KG, PRG) is as follows:

- Parameter generation algorithm $pk \leftarrow \text{KG}(\lambda)$:
compute $i \leftarrow \text{KG}_{\mathcal{F}}(\lambda)$ and set $pk := i$;
- Bit string generation algorithm $(s', b) \leftarrow \text{PRG}(pk, s)$:
upon receiving s and pk , where $pk = i$, $s = s_1 || s_2$ and $|s_1| = |s_2| = \ell$, compute the following:
 $s'_1 := s_1, s'_2 := f_i(s_2)$ (or $s'_2 := \text{Eval}_{\mathcal{F}}(i, s_2)$), and $s' = s'_1 || s'_2, b := \langle s_1, s_2 \rangle$.

¹⁰We ignore the running time of the adversary here for simplicity. Also, according to the detection principle in section 2, the PRG algorithm is deterministic and with a user selected input distribution, thus we can treat it as honest implemented.

We can show in the lemma below that the basic construction above is a 1-backdoor free PRG. The intuition is that in the (simplified) Blum-Micali PRG, a distinguisher can be transformed into an OWF inverter (following the GL proof), thus an adversary who can build a backdoor for this PRG implies that she has the ability to make \mathcal{F} (forgeable), which violates the strong unforgeability of \mathcal{F} .

Lemma 5.2. *Given a strongly unforgeable one way permutation family \mathcal{F} , the basic construction above is 1-backdoor free.*

Proof. The specification KG_{SPEC} of the simplified Blum-Micali PRG outputs a random function index from the corresponding index set (by simply running the key generation specification of \mathcal{F}).

First, it is easy to see that the OWF function family $\mathcal{G} = \{g_i\}$, where $g_i(x_1||x_2) := x_1||f_i(x_2)$, is strongly unforgeable if \mathcal{F} is strongly unforgeable.

If the above basic construction is a $(1, \delta)$ -backdoored PRG for some non-negligible δ , there exist PPT adversaries $(\mathcal{A}_{\text{INITIAL}}, \mathcal{A}_{\text{DIST}})$ such that (i.) $\mathcal{A}_{\text{INITIAL}}$ can output a pair (pk, sk) , where pk is indistinguishable from a randomly sampled public parameter; and (ii.) $\mathcal{A}_{\text{DIST}}$, with the backdoor sk , can distinguish $s_1||f_i(s_2)||B(s)$ (where $B(s) = \langle s_1, s_2 \rangle$) from a uniform $(2\ell + 1)$ -bit string. It is not hard to see that we can then construct algorithms (KG, Inv) that break the strong unforgeability of \mathcal{G} .

In particular, KG runs $\mathcal{A}_{\text{INITIAL}}$ and outputs the received key pair (pk, sk) , where pk here corresponds to a function index i that is indistinguishable from a random index, and sk corresponds to the backdoor z . Inv receives a challenge $y = g_i(x)$ and the backdoor z ; it first constructs an algorithm \mathcal{A}_P . \mathcal{A}_P selects a random bit b , and runs $\mathcal{A}_{\text{DIST}}(pk, sk, y||b)$. By the definition of $\mathcal{A}_{\text{DIST}}$, if $b = B(s)$, $\mathcal{A}_{\text{DIST}}$ (with sk) will output 0 with probability $1/2 + \delta$. It is easy to see that \mathcal{A}_P can predict the GL hardcore predicate B with advantage $\delta/2$, following the GL proof [8], there exists another algorithm $I^{\mathcal{A}_P}(pk, sk, \cdot)$ that can invert y with probability $\delta' = \text{poly}(\delta/2)$. Inv runs $I^{\mathcal{A}_P}(pk, sk, i, y)$ and recovers x' ; Inv then outputs x' if it is a valid pre-image of y . It follows that $\Pr[x' = x] \geq \text{poly}(\delta/2) = \delta'$ and \mathcal{G} will be δ' -forgeable for a non-negligible δ' , thus contradicts our assumption. \square

Full-fledged construction. We now extend our basic construction via iterations to show that the full fledged Blum-Micali PRG construction, using our strongly unforgeable OWF, achieves a q -backdoor free PRG for any $q = \text{poly}(\lambda)$. Our full-fledged construction (KG, PRG) ¹¹ is as follows:

- Parameter generation algorithm $pk \leftarrow \text{KG}(\lambda)$:
compute $i \leftarrow \text{KG}_{\mathcal{F}}(\lambda)$ and set $pk := i$;
- Bit string generation algorithm $(s', r) \leftarrow \text{PRG}(pk, s)$:
upon receiving s and pk where $pk = i$, $s = s_1||s_2$, and $|s_1| = |s_2| = \ell$, compute the following:
 - let $s_1^0 := s_1$ and $s_2^0 := s_2$;
 - for $j = 1, \dots, \ell'$,
 $b_j := \langle s_1^{j-1}, s_2^{j-1} \rangle$;
 $s_1^j := s_1^{j-1}$; $s_2^j := f_i(s_2^{j-1})$; $s^j := s_1^j||s_2^j$;
 - $s' = s^{\ell'} = s_1||f_i^{\ell'}(s_2)$; and $r = b_1 \dots b_{\ell'}$.

Please see Figure 5 for pictorial illustration.

¹¹PRG^q can be defined in a straightforward manner that runs PRG for q iterations, each iteration outputs ℓ' bits and updates the state for next iteration.




Figure 5: One iteration of BM-PRG

Theorem 5.3. *The full fledged construction above is q -backdoor free (for any polynomially large q), if the underlying OWP family \mathcal{F} is strongly unforgeable.*

Proof sketch. Following Lemma 5.2, $s_1 || f_i^j(s_2) || b_j$ is pseudorandom, i.e., it is indistinguishable from “ u_1, \dots, u_ℓ, v_1 ”, even to the adversaries $(\mathcal{A}_{\text{INITIAL}}, \mathcal{A}_{\text{DIST}})$ who set i , where $\{u_i\}, v_1$ are all random bits. While b_{j+1} is only related to $s_1, f_i^j(s_2)$, it follows that to the adversary $\mathcal{A}_{\text{DIST}}$ (who has the backdoor), $b_1 \dots b_{\ell'}$ satisfies the next-bit unpredictability. Following the classic reduction from pseudorandomness to next-bit unpredictability, we can conclude that $b_1 \dots b_{\ell'}$ is indistinguishable from uniform bits from $\{0, 1\}^{\ell'}$, even to $\mathcal{A}_{\text{DIST}}$. (This can be shown via the hybrid argument.) Then, inductively, we can conclude that r_1, \dots, r_q are indistinguishable from $\ell' \cdot q$ uniform bits. \square

Remark 5.4. *It is easy to see that if starting from an OWF(not necessarily unforgeables), the full fledged construction can be easily modified by replacing $f_i(\cdot)$ with $f_{h(i)}(\cdot)$.*

5.3 General public immunization strategy for PRG

An impossibility result about public immunization of a backdoored PRG was presented in [7]. However, we observe that this impossibility result only applies to an immunization procedure that operates on the *output* of the backdoored PRG . The application of strongly unforgeable OWF to backdoor-free PRG shown above inspires us to consider a new, general immunizing strategy for backdoored PRG . We suggest that—similar to the procedure above for eliminating backdoors in OWFs—one can randomize the public parameters to sanitize the PRG.¹² The intuition for this strategy to be effective in the setting of PRG is similar: if a specification KG_{SPEC} that outputs a uniform pk from its domain, no single backdoor can be used to break the security for large amount of public parameters; otherwise, one can use this trapdoor to break the PRG security of the specification.

Consider a PRG implementation (KG, PRG) (in which the KG algorithm might be backdoored), which is proven secure if the KG_{SPEC} outputs uniformly from its range PP . Let h be randomly selected from a hash family $H : PP \rightarrow PP$ which is modeled as a random oracle. Then we can construct a backdoor-free PRG, $(h \circ \text{KG}, \text{PRG})$, i.e., applying the hash to the given public parameter to derive the actual pk that will be fed into the PRG algorithm (alternatively, we can think as the new deterministic pseudo randomness generation algorithm to be defined as $\text{PRG}(h(pk), \cdot)$). Note that in order for this method to work, we must insist that pk can not be null, and is indeed used by the PRG algorithm, as in the case of, e.g., Dual_EC PRG. Also see the pictorial illustration in Fig 6.




Figure 6: Public immunization strategy for PRG.

¹²To interpret this results, since the solution of [7] requires a trusted seed/key generation and apply the function to the PRG output, thus part of the PRG algorithm can not be subverted. It follows that the construction of PRG in the complete subversion model was still open until our solution. In contrast, our sanitizing strategy does not require any secret, and even the deterministic hash function can be implemented by the adversary as part of the KG algorithm.

Theorem 5.5. Assume $(\text{KG}_{\text{SPEC}}, \text{PRG}_{\text{SPEC}})$ is a pseudorandom generator if KG_{SPEC} outputs pk randomly from its domain. Given any implementation (KG, PRG) , hashing the public parameters as described above, i.e., $(h \circ \text{KG}, \text{PRG}_{\text{SPEC}})$ yields a q -backdoor-free pseudorandom generator in the random oracle model for any polynomially large q .

Proof sketch. Suppose $(h \circ \text{KG}, \text{PRG}_{\text{SPEC}})$ is a (q, δ) -backdoored PRG. Then there exist a pair of adversaries $(\mathcal{A}_{\text{INITIAL}}, \mathcal{A}_{\text{DIST}})$ that can win the backdoored PRG game defined in Figure 4 with advantage δ . We will transform these adversaries into an adversary \mathcal{S} that breaks the security of $(\text{KG}_{\text{SPEC}}, \text{PRG}_{\text{SPEC}})$.

Suppose the challenger of the specification version of PRG selects the parameter pk^* and the challenge string r^* is either $\text{PRG}(pk^*, s)$ (for a uniform $s \in \{0, 1\}^\ell$) or a uniform string from $\{0, 1\}^{\ell' \cdot q}$ for some ℓ' .

The reduction follows the proof of Theorem 3.5: \mathcal{S} attempts to embed pk^* into the answers to the random oracle queries. In particular, if $\mathcal{A}_{\text{INITIAL}}$ outputs pk , \mathcal{S} wishes to answer the random oracle query about pk using pk^* , i.e., $h(pk) = pk^*$. We then proceed with a similar probabilistic analysis.

It is easy to see that if $\mathcal{A}_{\text{DIST}}$ can distinguish r^* from a uniform string, then \mathcal{S} will be able to distinguish the output of the specification version from random which violates the PRG security of the specification. \square

Remark 5.6. There are several points we would like to stress:

- If the public parameter contains only random elements from a group, e.g., the *Dual_EC PRG*, we may simply encode them into bits and use the regular hash functions like *SHA-256*, directly and convert the resulting bits back to a group element;
- If the public parameters are structured elements, or the KG_{SPEC} does not output a uniform distribution, we can work on the split-program model that forces the adversarial implementation to explicitly isolate its generation of randomness, and make the randomness public as part of the public parameter.
- If we treat the immunizing method as part of the KG algorithms, i.e., $h \circ \text{KG}$ is a single algorithm, we can let the adversary sets both (pk, \tilde{pk}) as the public parameter, regular user simply uses pk as the actual parameter and the crypto experts can check the validity of it. Similarly, in the split-program model, we can let the big brother set the original randomness r and sanitized randomness \tilde{r} , together with \tilde{pk} as the public parameters.

6 Conclusion and Open Problems

We initiate the systematic study of defending mechanism against kleptographic attacks of cryptographic primitives when *all* algorithms are subject to the subversion, we call **cliptography**.

We start from the fundamental primitives of (trapdoor) one way functions. We formalize the notions of forgeable (trapdoor) OWF to capture the possibility of embedding backdoors, in particular, into the function generation algorithms, and show how to launch such attacks. More interestingly, we suggest a general sanitization method employing randomization of the function index to destroy the possibility of embedding backdoor information. To instantiate our method in practice, we propose a split-program model in which the function generation algorithm consists of two components, a randomized component RG and a deterministic component dKG; here, the first component RG generates a uniform random string for the second component dKG, and then the second component generates an index based on such random string. In such split-program model, we can directly apply our general method of immunizing one way function generation to the randomness generated by RG.

We then pursue the possibility of building cliptography from unforgeable (trapdoor) OWFs. In particular, we show how to construct a signature scheme and a pseudorandom generator that preserves its

security in the complete subversion model. These are done by using our unforgeable trapdoor OWF and strongly unforgeable OWF as the key/parameter generation algorithm. Finally, we show how to apply our immunizing technique directly to the setting of PRG and present a general public immunizing strategy for PRG.

Many important problems remain open about defending against kleptographic attacks, and in general against mass surveillance. The immediate open questions left by our paper would be: to what extent the results and techniques developed in this paper can be used to build cliptography. In particular, can we construct other cryptographic primitives that preserves their security in the setting that all algorithms are subject to kleptographic attacks? There are also many other related open questions, to name a few: (1) how to sanitize the OWF generation in the standard model? (2) how to eliminate the subliminal channel in general? (3). Can we show some results in the complete subversion model so that we can preserve the security for multiparty protocols? and finally, (4) how to design an accountable mechanism for surveillance, such that key escrow is provided but not abused?

References

- [1] M. Bellare and V. T. Hoang. Resisting randomness subversion: Fast deterministic and hedged public-key encryption in the standard model. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 627–656. Springer, Apr. 2015. [3](#), [4](#), [13](#), [14](#)
- [2] M. Bellare, K. G. Paterson, and P. Rogaway. Security of symmetric encryption against mass surveillance. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 1–19. Springer, Aug. 2014. [3](#), [4](#), [5](#), [13](#), [14](#), [15](#)
- [3] M. Bellare and P. Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In U. M. Maurer, editor, *EUROCRYPT’96*, volume 1070 of *LNCS*, pages 399–416. Springer, May 1996. [15](#)
- [4] M. Blum and S. Micali. How to generate cryptographically strong sequences of pseudo random bits. In *23rd FOCS*, pages 112–117. IEEE Computer Society Press, Nov. 1982. [17](#)
- [5] S. Checkoway, R. Niederhagen, A. Everspaugh, M. Green, T. Lange, T. Ristenpart, D. J. Bernstein, J. Maskiewicz, H. Shacham, and M. Fredrikson. On the practical exploitability of dual EC in TLS implementations. In *Proceedings of the 23rd USENIX Security Symposium, San Diego, CA, USA, August 20-22, 2014.*, pages 319–335, 2014. [5](#), [8](#)
- [6] J.-S. Coron. On the exact security of full domain hash. In M. Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 229–235. Springer, Aug. 2000. [15](#), [16](#)
- [7] Y. Dodis, C. Ganesh, A. Golovnev, A. Juels, and T. Ristenpart. A formal treatment of backdoored pseudorandom generators. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 101–126. Springer, Apr. 2015. [3](#), [4](#), [5](#), [16](#), [17](#), [19](#)
- [8] O. Goldreich and L. A. Levin. A hard-core predicate for all one-way functions. In *21st ACM STOC*, pages 25–32. ACM Press, May 1989. [4](#), [17](#), [18](#)
- [9] S. Goldwasser and R. Ostrovsky. Invariant signatures and non-interactive zero-knowledge proofs are equivalent (extended abstract). In E. F. Brickell, editor, *CRYPTO’92*, volume 740 of *LNCS*, pages 228–245. Springer, Aug. 1993. [15](#)

- [10] N. J. Hopper, J. Langford, and L. von Ahn. Provably secure steganography. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 77–92. Springer, Aug. 2002. 4
- [11] A. Juels and J. Guajardo. RSA key generation with verifiable randomness. In D. Naccache and P. Paillier, editors, *PKC 2002*, volume 2274 of *LNCS*, pages 357–374. Springer, Feb. 2002. 4
- [12] J. Larson, N. Perlroth, and S. Shane. Revealed: The NSA’s secret campaign to crack, undermine internet security. Pro-Publica, 2013. <http://www.propublica.org/article/the-nsas-secret-campaign-to-crack-undermine-internet-encryption>. 3
- [13] A. Lysyanskaya. Unique signatures and verifiable random functions from the DH-DDH separation. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 597–612. Springer, Aug. 2002. 15
- [14] I. Mironov and N. Stephens-Davidowitz. Cryptographic reverse firewalls. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 657–686. Springer, Apr. 2015. 3, 5
- [15] NIST. Special publication 800-90: Recommendation for random number generation using deterministic random bit generators. National Institute of Standards and Technology, 2012. <http://csrc.nist.gov/publications/PubsSPs.html>. 5
- [16] N. Perlroth, J. Larson, and S. Shane. N.S.A. able to foil basic safeguards of privacy on web. The New York Times, 2013. <http://www.nytimes.com/2013/09/06/us/nsa-foils-much-internet-encryption.html>. 3
- [17] G. J. Simmons. The prisoners’ problem and the subliminal channel. In D. Chaum, editor, *CRYPTO’83*, pages 51–67. Plenum Press, New York, USA, 1983. 4
- [18] G. J. Simmons. A secure subliminal channel (?). In H. C. Williams, editor, *CRYPTO’85*, volume 218 of *LNCS*, pages 33–41. Springer, Aug. 1986. 4
- [19] A. Young and M. Yung. The dark side of “black-box” cryptography, or: Should we trust capstone? In N. Kobitz, editor, *CRYPTO’96*, volume 1109 of *LNCS*, pages 89–103. Springer, Aug. 1996. 3, 4, 5, 7
- [20] A. Young and M. Yung. Kleptography: Using cryptography against cryptography. In W. Fumy, editor, *EUROCRYPT’97*, volume 1233 of *LNCS*, pages 62–74. Springer, May 1997. 3, 4, 5, 7

A Omitted Proofs

Lemma 3.1. *A $(1 - \epsilon)$ -strongly forgeable OWF family is also a TDOWF family, where ϵ is a negligible function of the security parameter.*

Proof. Suppose $\mathcal{F} := (\text{BG}_{\mathcal{F}}, \text{KG}_{\mathcal{F}}, \text{Inv}_{\mathcal{F}})$ is a $(1 - \epsilon)$ strongly-forgeable OWF family. We can define $(\text{KG}, \text{Eval}, \text{Inv})$ for the TDOWF family as follows. The generation algorithm KG is as follows: first run $\text{BG}_{\mathcal{F}}(\lambda)$ and receive a string z ; then run $\text{KG}_{\mathcal{F}}$ with input (λ, z) and receive a function index i ; finally output (i, z) . The inversion algorithm Inv for the TDOWF is simply $\text{Inv}_{\mathcal{F}}$, and the evaluation algorithm is defined as $\text{Eval}(i, x) = f_i(x)$.

We can easily see that $(\text{KG}, \text{Eval}, \text{Inv})$ is invertible once trapdoor is given: since \mathcal{F} is a $(1 - \epsilon)$ strongly-forgeable OWF family, it holds that by definition, $\Pr[\text{Inv}_{\mathcal{F}}(z, i, y) = x \mid x \leftarrow X_i; y := f_i(x)] \geq 1 - \epsilon$; therefore, $\Pr[\text{Inv}(i, y, z) = x \mid x \leftarrow X_i; y := f_i(x)] \geq 1 - \epsilon$.

We can also show the one-wayness: without z , no PPT algorithm can invert y for a random x ; otherwise assume $(\text{KG}, \text{Eval}, \text{Inv})$ is not one-way, then there exists an adversary \mathcal{A} who for $i \leftarrow \text{KG}_{\mathcal{F}}$, can invert $y := f_i(x)$ with non-negligible probability. We note that the specification function generation algorithm is one-way, i.e., for $i \leftarrow \text{KG}_{\text{SPEC}}$, no one can invert $y := f_i(x)$ except negligible probability. Now, one can distinguish the output distribution of $\text{KG}_{\mathcal{F}}(z)$ from the output distribution of KG_{SPEC} , simply by trying inversion using \mathcal{A} . \square

Lemma 3.2. *One can construct a $(1 - \epsilon)$ -strongly forgeable OWF from a TDOWF, where ϵ is a negligible function on the security parameter.*

Proof. Consider a TDOWF $\mathcal{F} = \{f_i\}$ with the associated algorithms $(\text{KG}_{\mathcal{F}}, \text{Inv}_{\mathcal{F}})$. Assuming the trapdoors can be represented using $\ell(\lambda)$ bits, we construct a strongly forgeable OWF family $\mathcal{G} = \{g_{i,r}\}$, where $g_{i,r}(x) = f_i(x) || r$ and $r \in \{0, 1\}^{\ell(\lambda)}$.

The specification version of the sfunction generation algorithm KG_{SPEC} is defined as follows: run the $\text{KG}_{\mathcal{F}}$ algorithm and receive a function index/trapdoor pair (i, t_i) ; then discard t_i and sample randomly $r \leftarrow \{0, 1\}^{\ell(\lambda)}$; finally output (i, r) . It is easy to see that $g_{i,r}$ is one way because f_i is one way (without t_i).

While for the backdoored implementations, BG first outputs a random key k for a symmetric key encryption scheme $\text{SE} = (\text{SE.Enc}, \text{SE.Dec})$ which is assumed to be a pseudorandom permutation (PRP). $\text{KG}(k)$ is defined as follows: it first runs $\text{KG}_{\mathcal{F}}$, and receives an index i together with the corresponding trapdoor t_i ; the second part \tilde{r} is generated by encrypting t_i using k , i.e., $\tilde{r} = \text{SE.Enc}(k, t_i)$. $\text{KG}(k)$ outputs (i, \tilde{r}) . $(g_{i,\tilde{r}}(x) = f_i(x) || \text{SE.Enc}(k, t_i))$. It is easy to see that with the backdoor k , one can define the Inv as follows: it first decrypts \tilde{r} using k to retrieve t_i , and then inverts $f_i(x)$ by running $\text{Inv}_{\mathcal{F}}$ with t_i as an input.

Furthermore, since SE.Enc is modeled as a PRP, the distributions of (i, r) returned by KG_{SPEC} and (i, \tilde{r}) returned by $\text{KG}(k)$ for any k are computationally indistinguishable. \square