Bounded KDM Security from iO and OWF

Antonio Marcedone[1], Rafael Pass*[1], and abhi shelat†[2]

[1]Cornell University, {marcedone,rafael}@cs.cornell.edu
[2]University of Virginia, abhi@virginia.edu

July 5, 2016

Abstract

To date, all constructions in the standard model (i.e., without random oracles) of Bounded Key-Dependent Message (KDM) secure (or even just circularly-secure) encryption schemes rely on specific assumptions (LWE, DDH, QR or DCR); all of these assumptions are known to imply the existence of collision-resistant hash functions. In this work, we demonstrate the existence of bounded KDM secure encryption assuming indistinguishability obfsucation for $P/poly$ and just one-way functions. Relying on the recent result of Asharov and Segev (STOC'15), this yields the first construction of a Bounded KDM secure (or even circularly secure) encryption scheme from an assumption that provably does not imply collision-resistant hash functions w.r.t. black-box constructions. Combining this with prior constructions, we show how to augment this Bounded KDM scheme into a Bounded CCA2-KDM scheme.

*Supported in part by NSF Award CNS-1217821, AFOSR Award FA9550-15-1-0262, a Microsoft Faculty Fellowship, and a Google Faculty Research Award.

†Supported in part by NSF grants CNS-0845811, TC-1111781, TC-0939718, a Microsoft Faculty Fellowship, an SAIC Faculty Award, and a Google Faculty Research Award.

1 Introduction

The notion of Key Dependent Message (KDM) security, introduced by Black, Rogaway and Shrimpton [BRS02], requires an encryption scheme to remain secure even if the attacker can request encryptions of functions of the the secret key, and more generally encryptions of different secret keys in use by different players. This notion generalizes *circular security* introduced by Camenish and Lysyanskaya [CL01] in which the adversary can request encryptions of the form $\mathsf{Enc}_{pk_i}(sk_{i+1 \mod N})$. Both circularly-secure and KDM-secure encryption schemes have various applications, such as anonymous credential schemes, the "bootstrapping" technique used to construct fully homomorphic encryption, and disk encryption in the cases where the key itself might be encrypted.

The original works of Black, Rogaway and Shrimpton [BRS02] and of Camenish and Lysyanskaya [CL01] provided construction of circularly-secure encryption and even "full" KDM security (where there is no bound on the class of functions) in the Random Oracle model. Subsequent results provided constructions in the standard model (i.e., without random oracles), which is the focus of this paper.

Circular security and KDM for simple functions In a breakthrough result, Boneh, Halevi, Hamburg and Ostrovsky [BHHO08], provided the first construction of circular-secure encryption in the standard model (without a random oracle); their construction is based on the DDH assumption. Subsequently, schemes which expanded the class of permissible KDM queries and which were based on different assumptions were presented: Applebaum, Cash, Peikert and Sahai [ACPS09] obtain KDM security for affine functions under the Learning With Error (LWE) assumption; Brakerski, Goldwasser and Kalai [BGK09] give a transformation to convert a KDM scheme (with some additional properties) into one that is secure w.r.t. a richer class of functions: applying such transformation to the [BHHO08] and [ACPS09] gives a scheme secure w.r.t. the class to functions that can be expressed as polynomials of bounded degree, and a second one where the class consists of functions expressed as Turing machines of logarithmic size description. Malkin, Teranishi, and Yung [MTY11] achieves KDM security w.r.t. modular arithmetic circuits of bounded degree but (unbounded) polynomial size, based on the Decisional Composite Residuosity assumption (DCR). Wee [Wee16] explains [BHHO08,BGK09,BG10] as instantiations of a common framework based on smooth projective hashing, but known constructions of such hashing are based on the DDH, QR and DCR assumptions.

Bounded KDM security Barak, Haitner, Hofheinz and Ishai [BHHI10] significantly expand the class of permissible functions by showing how to realize KDM secure encryption for any set of circuits of a-priori bounded size; this notion is referred to as Bounded KDM security. Roughly speaking, their construction shows how to utilize schemes that satisfy KDM-security w.r.t. affine functions (and additional properties, which are satisfied by the known constructions) to get KDM security w.r.t. *any* circuit of bounded size. Their constructions can be instantiated from schemes relying on either DDH or LWE. Applebaum [App14] also show how to use randomized encodings to amplify KDM security against a small class of functions to Bounded KDM security[1].

[1]Both [App14] and [BHHI10] discuss how to strengthen their schemes to achieve a notion called length-dependent KDM security, which is slightly stronger than Bounded KDM security in the sense that the functions queried by the adversary can have circuit size which grows polynomially in the length of their inputs and outputs. We choose to state our result using Bounded KDM security for simplicity of exposition, but our construction can be similarly adapted to achieve this stronger notion by padding the obfuscated circuits appropriately

Our results Summarizing, all known constructions in the standard model (i.e., without random oracles) of Bounded KDM secure, and even just circularly-secure encryption rely on specific assumptions (LWE, DDH, QR or DCR). This gives rise to the following natural question:

Can Bounded KDM encryption be based on general assumptions?

In fact, all assumptions under which Bounded KDM schemes can be constructed imply the existence of collision-resistant hash functions. An orthogonal, but related, question is thus:

Can Bounded KDM encryption be constructed from an assumption that does not imply collision-resistant hash functions?

In this work we address both of these questions *assuming the existence of indistinguishability obfuscation ($i\mathcal{O}$).* Roughly speaking, program obfuscation is a class of cryptographic primitives aimed at making programs "intelligible" while preserving their functionality: in particular, $i\mathcal{O}$ guarantees that the obfucations of two circuits of the size that compute the same function (although potentially very different) are computationally indistinguishable. Our key result shows:

Theorem 1 (Informally stated)**.** *Assume there exists an indistinguishability obfuscator for P/poly, and a family of one way functions, then there exists a Bounded KDM secure public key encryption scheme.*

Interpreting our results Although iO is seemingly stronger than all assumptions from which KDM security could previously be based, our construction relies on assumption of qualitatively different and more general nature (we make no number-theoretic or lattice-based assumptions).

By the recent beautiful result by Asharov and Segev [AS15], it is known that black-box construction of CRH from iO and OWF is not possible[2], and as such, the assumption we use are separated (at least w.r.t. oracle-aided black-box constructions) from the assumptions previously used. As such, our work also addresses the second italicized question[3]. Notably, by embedding (in the security reduction) the code of the functions that the adversary asks as queries inside obfuscated circuits, our construction circumvents the impossibility result of [HH09], which shows that KDM security is impossible to get from any hardness assumption, as long as the reduction's proof of security treats both the adversary and functions queried as black boxes.

CCA2-KDM Security Camenisch, Chandran and Shoup [CCS09] construct a CCA2-KDM secure encryption scheme by using a KDM-secure scheme for the function family, an NIZK proof system, a CCA2-secure encryption scheme, a strongly secure one time signature scheme, and applying the Naor-Yung construction [NY90]. By combining our Bounded-KDM construction with the known constructions of CCA2 encryption and NIZK from sub-exponentially secure iO, one-way functions and signatures, we construct bounded CCA2-KDM secure encryption.

[2]They show that a CRHF cannot be constructed in a blackbox-manner from a one-way permutation and an indistinguishability obfuscator for all polynomial-sized *oracle-aided* circuits without exponential-loss in security. Such oracle-aided circuits can model most common uses of iO in cryptographic constructions such as puncturing in which the circuits that are obfuscated make oracle calls to the one-way permutation.

[3]In fact, combining our result with [AS15] directly rules black-box constructions of CRH from single-key BKDM security. On the other hand, it is not directly clear whether our final construction of multi-key BKDM falls into the class of oracle-aided circuits.

IND-CPA Security and Circular Security The dual problem of separating IND-CPA security from n-circular (and therefore KDM) security for $n > 1$ has also been open for a long time, and was solved assuming indistinguishability obfuscation and one way functions in [MO14, KRW15], and more recently relying on LWE in [KW16, AP16].

1.1 Proof Overview

Informally, the (N, L)-Bounded KDM security definition[4] states that no efficient adversary has non-negligible advantage in the following game:

1. The challenger generates a random bit b and N key pairs $(sk_1, PK_1), \ldots, (sk_N, PK_N)$ (where the secret keys have length k) and runs the adversary $\mathcal{A}$ on input the public keys.
2. $\mathcal{A}$ can adaptively make queries of the form (h, i), where $i \in \{1, \ldots, N\}$ and h is a circuit of size at most L, input size kN and output size k (representing a function from N secret keys to a k bit message). If $b = 1$, $\mathcal{A}$ receives an encryption $\mathsf{Enc}_{PK_i}(h(sk_1, \ldots, sk_n))$, and otherwise receives $\mathsf{Enc}_{PK_i}(0^k)$.
3. $\mathcal{A}$ halts and outputs a bit b'. $\mathcal{A}$ wins if $b = b'$.

The single-key case We start by giving an high level overview of our Bounded KDM secure scheme in the simpler case where $N = 1$. The secret key of our construction is just a string $s \in \{0,1\}^k$, while the public key consists of (p, K), where K is the key for an injective[5] one way function and $p = \mathrm{OWF}_K(s)$. To encrypt a message m, the ciphertext consists of the obfuscation of a program that on input x returns m if $\mathrm{OWF}_K(x) = p$ and $\perp$ otherwise. Decryption consists of running the obfuscated ciphertext program on input the secret key.

Informally, such a scheme should be IND-CPA secure because, if we treat the obfuscation as a black box, the only way to extract the message from a ciphertext (i.e. an obfuscated circuit) is to run the circuit on input the secret key, which is a sufficiently long uniformly random string. To argue the IND-CPA security of the scheme relying on an indistinguishability obfuscator, one can instead leverage a theorem from [BCP14]: informally, any adversary that distinguishes obfuscations of two circuits that differ on polynomially many inputs can be turned into an adversary that computes one input on which the two circuits differ. Therefore, an adversary distinguishing between encryptions of two different messages, i.e. only having different output on input an x such that $\mathrm{OWF}_K(x) = p$, can be turned into an adversary that computes such an x, effectively inverting the one way function.

To prove that the scheme is also KDM secure, the simulator needs to answer queries about a function h of a secret key s it does not know: this can be achieved by obfuscating a program that on input x, first checks whether $\mathrm{OWF}_K(x) = p$, and then returns either $h(x)$ if the check passes or $\perp$ otherwise. Since this new program is functionally equivalent to an honest encryption of $h(s)$ (as the one way function is injective[6] and therefore there is only one input s that passes the

[4] For simplicity, in this paper we assume that the message and key space of the encryption scheme are both $\{0,1\}^k$, where k is the security parameter.

[5] [BPW16] shows how to construct a family of one way functions where randomly sampled functions are injective with overwhelming probability. Their construction requires $i\mathcal{O}$, one way functions and q-wise independent hashing, as detailed in Section 2.

[6] To be more precise, the function is only injective with overwhelming probability. We will deal with this and other subtleties in the formal proof.

equality test, namely the secret key), indistinguishability obfuscation guarantees that no adversary will notice the difference. Moreover, since such a simulation does not require the secret key, we can later switch (in a standard hybrid argument) to a game where the public key is a pair (p, K) on which the simulator wants to invert the one way function, and prove security as in the IND-CPA case. This proof outline omits several subtle corner cases which complicate the formal proof.

The multi-key case We can extend the idea of computing $h(s)$ "on the fly" inside the ciphertext program when the correct secret key is given as input to the case where multiple secret keys are involved. The challenge is that the new ciphertext program is given as input only one of the secret keys and has to compute a function possibly depending on other independently generated keys[7].

We circumvent the problem by embedding in the simulated ciphertexts the relationship between secret keys $s_1, \ldots, s_n$ in the form a vector $\vec{r} = (s_1 \oplus s_1, s_2 \oplus s_1, \ldots, s_n \oplus s_1)$. Note that the vector itself is uniformly distributed (since the secret keys are, except for the first component $0^k = s_1 \oplus s_1$ which is left there for convenience of notation) and it allows the (simulated) ciphertext program, given one of the secret keys, to compute on the fly all the other secret keys (and therefore functions of them).

To reduce the security of the encryption scheme to the hardness of inverting the one way function, we have one last problem: the simulator has to compute the vector $\vec{r}$ without knowing the secret keys. Equivalently, we use the vector $\vec{r}$ to define the secret keys: the simulator will get a tuple (p_1, K_1) for which it has to find a preimage and sample a random $\vec{r}$, thus implicitly defining each secret key s_i as the string that satisfies $\mathrm{OWF}_{K_1}(s_i \oplus r_i) = p_1$. Note that this change does not modify the distribution of the secret keys and that the ciphertext programs will still be functionally equivalent to the ones in the real experiment. However, the simulator now cannot compute as public keys values (p_i, K_i) consistent with $p_i = \mathrm{OWF}_{K_i}(s_i)$. We therefore modify the original encryption scheme so that the public keys are also released in obfuscated form: the modified encryption scheme will have as public keys obfuscations of programs that have (p_i, K_i) embedded and on input x output 1 if $p_i \stackrel{?}{=} \mathrm{OWF}_{K_i}(x)$ and $\perp$ otherwise. The ciphertexts will be modified accordingly as obfuscations of programs that have the obfuscated public key PK embedded and on input x return the message if $PK(x) \stackrel{?}{=} 1$ and $\perp$ otherwise. In the simulation, these public key programs will be substituted with (functionally equivalent) obfuscated programs that output 1 iff $p_1 = \mathrm{OWF}_{K_1}(x \oplus r_i)$. This last modification allows the simulation to be completed without knowledge of any of the secret keys.

Lastly, as before, the same lemma from [BCP14] allows us to switch to a hybrid in which all the ciphertexts returned to the adversary are encryptions of 0; this implies the KDM security of the scheme.

Organization of the paper In the next section we go over the definition of KDM security and review some of the results required for our construction. Section 3 shows a constructions in the simpler case where only one secret key is used. Section 4 describes the more general construction where functions of more than one secret key might be used and considers some extensions and generalizations.

[7] Note that [BHHI10] solves the problem by embedding in their ciphertexts an encryption of the other secret keys under the appropriate public key, which is why circular security is required as an additional assumption for their underlying encryption scheme.

2 Preliminaries

Notation and Conventions If S is a finite set $s \leftarrow S$ is a uniformly random sample from S. If A is a randomized algorithm, $x \leftarrow A$ is the output of A on a uniformly random input tape.

Definition 2 (Injective OWF family (as stated in [BPW16]))**.** *Let l be a polynomially-bounded length function. An efficiently computable family of functions*

$$\mathcal{OWF} = \{\mathrm{OWF}_K : \{0,1\}^k \to \{0,1\}^* : K \in \{0,1\}^{l(k)}, k \in \mathbb{N}\}$$

associated with an efficient (probabilistic) key sampler $\mathcal{K}_{\mathcal{OWF}}$ is said to be an injective OWF family if it satisfies:

1. ***Injectiveness:*** *With overwhelming probability over the choice of $K \leftarrow \mathcal{K}_{\mathcal{OWF}}(1^k)$, the function OWF_K is injective*
2. ***One-wayness:*** *For any polysize inverter Adv there exists a negligible function $\mathrm{negl}(\cdot)$, such that for all $k \in \mathbb{N}$,*

$$\Pr\left[x \leftarrow \{0,1\}^k, K \leftarrow \mathcal{K}_{\mathcal{OWF}}(1^k) : \mathsf{Adv}(K, \mathrm{OWF}_K(x)) \stackrel{?}{=} x\right] \leq \mathrm{negl}(k)$$

[BPW16] shows how to construct injective one way functions assuming one way functions and indistinguishability obfuscation.

2.1 Bounded Key Dependent Message Security

Definition 3 (Public Key Encryption)**.** *A public key encryption scheme $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is a triple of PPT algorithms such that:*

- $\mathsf{Gen}(1^k)$ *is a randomized algorithm that takes as input a security parameter k and outputs a couple of strings $sk \in \mathcal{K}, pk$. For simplicity, we assume in the rest of the paper that the secret keys are exactly k bits long.*
- $\mathsf{Enc}(pk, m)$ *is a randomized algorithm that on input a public key and a message $m \in \mathcal{M}$ outputs a ciphertext c. Again, for simplicity we assume that $\mathcal{M} = \{0,1\}^k$.*
- $\mathsf{Dec}(sk, c)$ *is an algorithm that on input a secret key and a ciphertext c outputs a message m*

We require that $\forall k, \forall m \in \mathcal{M}, \forall (pk, sk) \leftarrow \mathsf{Gen}(1^k),\ \mathsf{Dec}(sk, \mathsf{Enc}(pk, m)) = m$.

Definition 4 (KDM Security w.r.t. $\mathcal{H}$)**.** *Let $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a public key encryption scheme with message space $\mathcal{M}$ and secret key space $\mathcal{K}$, where for simplicity we assume $\mathcal{M} = \mathcal{K} = \{0,1\}^k$. Fix a positive integer valued function $N = N(k) > 0$. Consider the following probabilistic experiment (i.e. a random variable) between a challenger and an adversary $\mathcal{A}$, parametrized by a bit b:*

$\mathbf{KDM}^b_{N,\mathcal{A}}(k)$*:*

- *The challenger runs $N = N(k)$ times $Gen(1^k)$ to get $(pk_1, sk_1), \ldots, (pk_N, sk_N)$ and runs the adversary $\mathcal{A}$ on input $\mathbf{pk} = (pk_1, \ldots, pk_N)$.*

- *The adversary can adaptively submit queries of the form (h, i), where $h : \mathcal{K}^N \to \mathcal{M}$ is a function (encoded as a circuit) and $i \in 1, \ldots, N$. If $b = 1$, the challenger answers these queries with $Enc(pk_i, h(\mathbf{sk}))$, otherwise with $Enc(pk_i, 0^k)$, where $\mathbf{sk} = (sk_1, \ldots, sk_N)$.*
- *The adversary stops and outputs a bit b', which is defined as the output of the game (i.e. the value of the random variable).*

The KDM advantage of $\mathcal{A}$ is defined as

$$\mathsf{Adv}^{KDM}_{PKE,N,\mathcal{A}}(k) \stackrel{\mathsf{def}}{=} |\Pr[\mathbf{KDM}^1_{N,\mathcal{A}}(k) = 1] - \Pr[\mathbf{KDM}^0_{N,\mathcal{A}}(k) = 1]|$$

We say that PKE is KDM secure with respect to a function class $\mathcal{H} = \{\mathcal{H}_k\}$ iff for every polynomial N and every PPT A that in the above game only queries the challenger with functions $h \in \mathcal{H}_k$, the advantage function $Adv^{KDM}_{\mathsf{PKE},N,\mathcal{A}}(k)$ is negligible in k.

Definition 5 (Bounded KDM Security)**.** *A public key encryption scheme PKE is said to be (N, L)-Bounded KDM secure if it is KDM secure with respect to the class $\mathcal{H} = \{\mathcal{H}_k\}$, where $\mathcal{H}_k$ consists of all functions $h : \mathcal{K}^{N(k)} \to \mathcal{M}$ that can be encoded as circuits[8] of size bounded by the polynomial function $L(k)$.*

Note that, for simplicity, we have denoted with $N(k)$ both the arity of the functions in $\mathcal{H}$ and the number of key pairs generated in the security experiment above. In general, the number of keys in the experiment might be higher than the arity of the functions in $\mathcal{H}$, and it is easy to extend our proofs to hold even in this case.

2.2 Indistinguishability Obfuscation

Definition 6 (Indistinguishability Obfuscation [GGH+13])**.** *Given a circuit class $\{\mathcal{C}_k\}$, a (uniform) PPT machine $i\mathcal{O}$ is called an indistinguishability obfuscator (iO) for $\{\mathcal{C}_k\}$ if it satisfies:*

Preserving Functionality: *For every $k \in \mathbb{N}$ and $C \in \mathcal{C}_k$,*

$$\Pr[C'(x) = C(x) | C' \leftarrow i\mathcal{O}(k, C)] = 1 \quad \forall x$$

Indistinguishability: *For any (not necessarily uniform) polynomial-size distinguisher $\mathcal{D}$, all security parameters k and all couples $C_0, C_1 \in \mathcal{C}_k$ such that $C_0(x) = C_1(x)$ for all inputs x, we have that*

$$\Big|\Pr[D(i\mathcal{O}(k, C_0)) = 1] - \Pr[D(i\mathcal{O}(k, C_1)) = 1]\Big| \leq \mathrm{negl}(k)$$

2.3 Extractability Obfuscation

Definition 7 (Weak Extractability Obfuscation [BCP14])**.** *A uniform transformation $\mathcal{O}$ is a weak extractability obfuscator for a class of circuits $\mathcal{C} = \{\mathcal{C}_k\}$ if the following holds. For every PPT adversary $\mathcal{A}$ and polynomial $p(k)$, there exists a PPT algorithm E and polynomials $p_E(k), t_E(k)$*

[8] Recall that we assume for simplicity $\mathcal{M} = \mathcal{K} = \{0,1\}^k$.

for which the following holds. For every polynomial $d(k)$, for all sufficiently large k, and every pair of circuits $C_0, C_1 \in \mathcal{C}_k$ differing on at most $d(k)$ inputs, and every auxiliary input z,

$$\Pr[b \leftarrow \{0,1\}; \tilde{C} \leftarrow \mathcal{O}(1^k, C_b) : \mathcal{A}(1^k, \tilde{C}, C_0, C_1, z) = b] \geq \frac{1}{2} + \frac{1}{p(k)}$$

$$\Rightarrow \Pr[x \leftarrow E(1^k, C_0, C_1, z) : C_0(x) \neq C_1(x)] \geq \frac{1}{p_E(k)},$$

and the runtime of E is $t_E(k, d(k))$.

Lemma 8 ([BCP14])**.** *Let $i\mathcal{O}$ be an indistinguishability obfuscator for P/poly. Then $i\mathcal{O}$ is also a weak extractability obfuscator for P/poly.*

3 (1,L)-Bounded KDM Construction

In this section, we consider the simpler case where functions queried by the adversary can involve only one public key at a time.

3.1 Scheme description

The scheme is parametrized over a polynomial function $L(k)$ (which is a bound on the size of the circuits for which we can prove the Bounded KDM security of the scheme).

Π_L :

Key Generation: The algorithm $\mathsf{Gen}(1^k)$ generates a random secret key $s \leftarrow \{0,1\}^k$ and a key for an injective one way function $K \leftarrow \mathcal{K}_{\mathcal{OWF}}(1^k)$. It outputs s as the secret key and the couple (p, K) where $p \leftarrow \mathrm{OWF}_K(s)$ as the public key.

Encryption: The algorithm $\mathsf{Enc}((p, K), m)$ on input a public key (p, K) and a message $m \in \{0,1\}^k$ outputs an obfuscated circuit $C \leftarrow i\mathcal{O}(G_{p,K,m}(\cdot))$ (the circuit $G_{p,K,m}$ is described in figure 1).

Decryption: The algorithm $\mathsf{Dec}(s, \overline{C})$ on input a secret key $s \in \{0,1\}^k$ and a ciphertext $\overline{C} \in \mathcal{P}$ outputs $m' = \overline{C}(s)$.

It can be verified that correctness of the Obfuscator implies correctness of the encryption scheme. The following theorem argues that the scheme achieves Bounded KDM Security.

Theorem 9. *If $i\mathcal{O}$ is an indistinguishability obfuscator for P/poly and $\mathcal{OWF}$ is a family of injective one way functions, then for any polynomial function $L(\cdot)$ the encryption scheme $\Pi_L = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ described above is $(1, L)$-Bounded KDM secure.*

Proof. The proof proceeds by a hybrid argument. Assume by contradiction that there exists an adversary $\mathcal{A}$ such that $\mathsf{Adv}^{KDM}_{\Pi_L,1,\mathcal{A}}(k)$ is non negligible in k, i.e. there exists a polynomial p such that $\mathsf{Adv}^{KDM}_{\Pi_L,1,\mathcal{A}}(k) > \frac{1}{p(k)}$ for infinitely many k. We define the random variable $\mathbf{KDM}^{Hyb}_{1,\mathcal{A}}(k)$ exactly as $\mathbf{KDM}^{1}_{1,\mathcal{A}}(k)$, but where queries $(h, 1)$ by the adversary[9] are answered by returning as the ciphertext

[9]since there is only one public key, in the rest of the theorem we will just refer to the query for a function h and implicitly assume $i = 1$

$G_{p,K,m}$:
Constants: $p \in \{0,1\}^*, m \in \{0,1\}^k, K \in \{0,1\}^{l(k)}$
Inputs: $x \in \{0,1\}^k$

1. If $(\mathrm{OWF}_K(x) \stackrel{?}{=} p)$, then output m; else output $\perp$.

$G'_{p,K,h}$:
Constants: $p \in \{0,1\}^*, h \in \{0,1\}^L, K \in \{0,1\}^{l(k)}$
Inputs: $x \in \{0,1\}^k$

1. If $(\mathrm{OWF}_K(x) \stackrel{?}{=} p)$, then output $h(x)$; else output $\perp$.

All the circuits are padded to a specific length $\ell(k, L)$ which will be specified in the proof.

Figure 1: Circuits used in the encryption of the (1-L)-Bounded KDM scheme

an obfuscation $i\mathcal{O}(G'_{p,K,h})$, where (p, K) is the public key generated in the first step of the game and G' is described in figure 1.

By definition of advantage and the triangular inequality, it must be that at least one of the following two inequalities hold for infinitely many k:

$$|\Pr[\mathbf{KDM}^1_{1,\mathcal{A}}(k) = 1] - \Pr[\mathbf{KDM}^{Hyb}_{1,\mathcal{A}}(k) = 1]| > \frac{1}{2p(k)}$$

$$|\Pr[\mathbf{KDM}^{Hyb}_{1,\mathcal{A}}(k) = 1] - \Pr[\mathbf{KDM}^0_{1,\mathcal{A}}(k) = 1]| > \frac{1}{2p(k)}$$

However, the next two lemmas will prove that the quantities on the right hand side of the above inequalities are both negligible, which is a contradiction and therefore proves the claim. In the following, for brevity, we will denote $\mathbf{KDM}^b_{1,\mathcal{A}}(k)$ for $b = 0, 1, Hyb$ as Z_b.

To make sure we can rely on the security of the $i\mathcal{O}$ in the lemmas below, we set $\ell(k, L)$ to be an upper bound on the size of the circuits $G_{p,K,m}$ and $G'_{p,K,h}$ of figure 1. □

Lemma 10. $|\Pr[Z_1 = 1] - \Pr[Z_{Hyb} = 1]| < \mathrm{negl}(k)$

Proof. To prove this lemma, we rely on the security of the $i\mathcal{O}$ obfuscator using a hybrid argument. Let $q(k)$ be a (polynomial) upper bound on the number of queries that $\mathcal{A}$ makes. We consider a series of hybrid games: for $j = 0, \ldots, q(k)$ define the random variable H_j as an interactive experiment where the first and third step (i.e. the key generation phase and the output of the game) are defined as in $\mathbf{KDM}^1_{1,\mathcal{A}}(k)$, while the queries are handled as follows. The first $q(k) - j$ queries made by $\mathcal{A}$ are answered with $i\mathcal{O}(G_{p,K,h(s)})$ (where h is the function the adversary queried), i.e. according to what would happen in game Z_1; instead, the last j queries are answered with $i\mathcal{O}(G'_{p,K,h})$ (i.e. according to Z_{Hyb}). Since H_0 has the same distribution as Z_1 and $H_{q(k)}$ has the same distribution as Z_{Hyb}, to prove the claim it is enough to show that for all j, $|\Pr[H_j = 1] - \Pr[H_{j+1} = 1]| < \mathrm{negl}(k)$.

Notice that, until the $(q(k) - j)^{th}$ query is answered by the challenger, the "state" of the game has the same distribution both in H_j and H_{j+1}. Here the "state" of the game consists of the

keys sampled by the challenger in the first step of the game, the view of the adversary $\mathcal{A}$ and its internal state; it depends on the random choices made by both $\mathcal{A}$ and the challenger up to the point where the $(q(k)-j)^{th}$ query is made (but not answered). We can therefore consider the state $\bar{t}$ that maximizes (over the choice of t) the quantity $|\Pr[H_j = 1|t] - \Pr[H_{j+1} = 1|t]|$ subject to the constraint that the one way function key K chosen by the challenger defines an injective one way function. Let $\bar{v}(k)$ be this maximum value. It is not hard to see that, in order for the lemma to hold, it is enough to prove that this $v(\bar{k})$ is negligible in k. In fact, let T denote the set of all possible game states (as defined above), T_1 be the set of states such that the key K chosen by the challenger as part of the public key defines an injective one way function, and $T_2 = T \setminus T_1$. Since the one way function family is injective, $\Pr[t \in T_2]$ is negligible. Therefore we have that:

$$|\Pr[H_j = 1] - \Pr[H_{j+1} = 1]| \leq |\bar{v}(k)| \Pr[t \in T_1] + \Pr[t \in T_2] = |\bar{v}(k)| + \text{negl}(k)$$

Therefore, we are left with proving that $\bar{v}(k)$ is negligible. Assume by contradiction it is not, and let $s, (p, K)$ and h be the sampled keys and the $(q(k)-j)^{th}$ function queried by $\mathcal{A}$ according to state $\bar{t}$. We will build a (non-uniform) distinguisher $\mathcal{D}$ which will distinguish between obfuscations of $G_{p,K,h(s)}$ and $G'_{p,K,h}$, thus contradicting the security of the obfuscator. First of all notice that, since $\bar{t}$ was conditioned on K being injective, the two circuits are functionally equivalent. $\mathcal{D}$ works as follows: it receives as a non uniform advice the state $\bar{t}$ (which has size polynomial in k) and as an input an obfuscation O of one of the two circuits. $\mathcal{D}$ continues the simulation of the experiment from state $\bar{t}$ as the challenger, using O as an answer to the $(q(k)-j)^{th}$ query of the adversary and then answering all other queries according to Z_{Hyb}. When $\mathcal{A}$ halts and outputs a bit b, $\mathcal{D}$ does the same. It holds that, when $O \leftarrow i\mathcal{O}(G_{p,K,h(s)})$, the view of $\mathcal{A}$ is distributed as in H_j, while when $O \leftarrow i\mathcal{O}(G'_{p,K,h})$, it is distributed as in H_{j+1}. Therefore

$$\left|\Pr[\mathcal{D}^{\bar{t}}(i\mathcal{O}(k, G_{p,K,h(s)})) = 1] - \Pr[\mathcal{D}^{\bar{t}}(i\mathcal{O}(k, G'_{p,K,h})) = 1]\right| =$$
$$\left|\Pr[H_j = 1 \mid \bar{t}] - \Pr[H_{j+1} = 1 \mid \bar{t}]\right| = |\bar{v}(k)|$$

which contradicts the security of $i\mathcal{O}$ and concludes the proof. □

Lemma 11. $|\Pr[Z_{Hyb} = 1] - \Pr[Z_0 = 1]| < \text{negl}(k)$

Proof. The proof is by contradiction of the one-wayness property of the function family leveraging Lemma 8. Let $q(k)$ be a (polynomial) upper bound on the number of queries that $\mathcal{A}$ makes. We consider a series of hybrid games: for $j = 0, \ldots, q(k)$ define the random variable H_j as an interactive experiment where the first and third step (i.e. the key generation phase and the output of the game) are defined as in $\mathbf{KDM}^0_{1,\mathcal{A}}(k)$, while the queries are handled as follows. The first $q(k) - j$ queries made by $\mathcal{A}$ are answered with $i\mathcal{O}(G'_{p,K,h})$ (where h is the function the adversary queried), i.e. according to what would happen in game Z_{Hyb}; instead, the last j queries are answered with $i\mathcal{O}(G_{p,K,0^k})$ (i.e. according to Z_0). Since H_0 has the same distribution as Z_{Hyb} and $H_{q(k)}$ has the same distribution as Z_0, to prove the claim it is enough to show that for all j, $|\Pr[H_j = 1] - \Pr[H_{j+1} = 1]| < \text{negl}(k)$.

Assume by contradiction that there exist a specific index j and an adversary $\mathcal{A}$ that can distinguish between H_j and H_{j+1} with non negligible probability $a(k)$. Note that, as in the previous lemma, the view of the adversary in games H_j and H_{j+1} has the same distribution up to the point where $\mathcal{A}$ makes the $(q(k)-j)^{th}$ query. We will use such an adversary to build an adversary B that

breaks the one wayness of $\mathcal{OWF}$. $\mathcal{B}$ takes as input randomly chosen function key $K \leftarrow \mathcal{K}_{\mathcal{OWF}}$ and the image p of the function on a random input, and has to compute a preimage x such that $\mathrm{OWF}_K(x) = p$.

By Lemma 8, the existence of an adversary C that distinguishes (with non negligible probability) between obfuscations of two circuits that differ on only one input implies the existence of a polynomial time algorithm E that computes the input on which they are different with overwhelming probability. $\mathcal{B}$ proceeds in two stages: first, it simulates for $\mathcal{A}$ an experiment similar to H_{j+1}, using its own input (p, K) as the public key and up to the point where the $(q(k) - j)^{th}$ query for a function $\bar{h}$ is asked. Let t be the state of the adversary $\mathcal{A}$ (including its view) at this point in the simulation. Note that this simulation is possible because knowledge of the preimage x is not necessary to compute obfuscations of the programs $G'_{p,K,h}$ that are returned as answers to the queries. As a second step, $\mathcal{B}$ can run the algorithm E (given by Lemma 8) on input $(1^k, G'_{p,K,\bar{h}}, G_{p,K,0^k}, t)$, which runs in polynomial time, and return its output.

We now analyze the success of $\mathcal{B}$ in two steps: we define a property of the states t sampled by $\mathcal{B}$ (a state satisfying the property will be called a "good" state), and show that $\mathcal{B}$ samples a good state with non negligible probability. Second, we show there exists an algorithm E that succeeds with noticeable probability conditioned on the fact that t is good.

Denote with T the distribution on the states of $\mathcal{A}$ obtained by running H_{j+1} up to the point where the $(q(k) - j)^{th}$ query for a function $\bar{h}$ is asked. A state $t \leftarrow T$ (containing the public key (p, K) and the $(q(k) - j)^{th}$ query $\bar{h}$) is said to be "good" if all the following holds:

1. K denotes an injective function
2. $\bar{h}$ is such that $\forall x, OWF_K(x) = p \Rightarrow \bar{h}(x) \neq 0^k$
3. $\Big|\Pr[H_j = 1|t] - \Pr[H_{j+1} = 1|t]\Big| > \frac{a(k)}{2}$. Here $\Pr[H_j = 1|t]$ denotes the probability that the value of the random variable H_j is 1 given that after the $(q(k) - j)^{th}$ query is asked $\mathcal{A}$ is in state t.

Denote with T_g the set of good states, with T_1 the set of states that do not satisfy condition 1, with T_2 the states that satisfy condition 1 but not condition 2, and with T_3 the states that satisfy conditions 1 and 2 but not 3. Note that T_g, T_1, T_2, T_3 are a partition of T. First, although $\mathcal{B}$ executes $\mathcal{A}$ using its own input (p, K) as the public key, when this input is randomly sampled $(x \leftarrow \{0,1\}^k; K \leftarrow \mathcal{K}_{\mathcal{OWF}}(1^k); p \leftarrow \mathrm{OWF}_K(x))$ the distribution of t obtained by $\mathcal{B}$ is exactly T. To argue that $t \leftarrow T$ is good with non negligible probability, assume by contradiction it was not. We have that, by a union bound

$$\begin{aligned} a(k) = &\Big|\Pr[H_j = 1 \mid t \in T_g] - \Pr[H_{j+1} = 1 \mid t \in T_g]\Big| \Pr[t \in T_g] + \\ &\sum_{i=1}^{3} \Big|\Pr[H_j = 1 \mid t \in T_i] - \Pr[H_{j+1} = 1 \mid t \in T_i]\Big| \Pr[t \in T_i] \leq (*) \end{aligned}$$

We note that $\Pr[t \in T_1]$ is negligible because the $\mathcal{OWF}$ is injective. Moreover, it is not hard to prove that $\Big|\Pr[H_j = 1 \mid t \in T_2] - \Pr[H_{j+1} = 1 \mid t \in T_2]\Big|$ is negligible as well: in fact, if $\bar{h}(x) = 0^k$ then $G'_{p,K,\bar{h}}$ and $G_{p,K,0^k}$ would be functionally equivalent, and therefore their obfuscations computationally indistinguishable (because of the security of $i\mathcal{O}$), so $\mathcal{A}$ in an execution from state

t would only be able to distinguish between them (and therefore between H_j and H_{j+1}) with negligible probability. Moreover, since for $t \in T_3$ condition 3 is not satisfied, $\Big|\Pr[H_j = 1 \mid t \in T_3] - \Pr[H_{j+1} = 1 \mid t \in T_3]\Big| \leq a(k)/2$ from which

$$(*) \leq \Big|\Pr[H_j = 1 \mid t \in T_g] - \Pr[H_{j+1} = 1 \mid t \in T_g]\Big| \Pr[t \in T_g] + \frac{a(k)}{2} + negl(k)$$

Therefore, if $\Pr[t \in T_g]$ was negligible, then $a(k)$ would be bounded by a negligible function , which is a contradiction.

For the second step we prove that, conditioned on $t \in T_g$, $\mathcal{B}$ inverts with non negligible probability. Note that, on good states, there is exactly one x such that $OWF_K(x) = p$, and moreover it holds that $\bar{h}(x) \neq 0^k$. Therefore $G'_{p,K,\bar{h}}$ and $G_{p,K,0^k}$ will differ on input x and have the same output on all others. Under this condition, $\mathcal{B}$ inverts the function iff algorithm E (given by Lemma 1) is successful, which happens with overwhelming probability as long as we can prove that as long as we can prove that there is an adversary C such that for each $t \in T_g$, C distinguishes obfuscations of $G'_{p,K,\bar{h}}$ and $G_{p,K,0^k}$ with non negligible probability. Consider the following adversary $C(O, G'_{p,K,\bar{h}}, G_{p,K,0^k}, t))$:

1. Resume running $\mathcal{A}$ from the saved state t, answering its first (i.e. $(q(k)-j)^{th}$) query with O.
2. Answer all subsequent queries with obfuscations of $G_{p,K,0^k}$.
3. When A halts outputs a bit b', halt and output the same bit.

Note that when $O \leftarrow i\mathcal{O}(G'_{p,K,h})$, the output of C has the same distribution as $H_j|t$, while if instead C is run on an obfuscation $O \leftarrow i\mathcal{O}(G'_{p,K,0^k})$, the view of $\mathcal{A}$ is consistent with $H_{j+1}|t$.

Therefore, if $t \in T_g$, the advantage of C in distinguishing obfuscations of the two circuits is equal to $\Big|\Pr[H_j = 1 \mid t] - \Pr[H_{j+1} = 1 \mid t]\Big| > \frac{a(k)}{2}$ which is non negligible. This proves that algorithm E exists and has overwhelming success probability on good states. Since we have also proven that $\mathcal{B}$ samples good states with non negligible probability, we can conclude that it has non negligible probability of inverting the $\mathcal{OWF}$, which contradicts its one-wayness. □

4 (N,L)-Bounded KDM Construction

The scheme is parametrized over polynomial functions $L(k), N(k)$ (which are bounds on the size of the circuits and number of keys for which we can prove the Bounded KDM security of the scheme).

$\Pi_{N,L}$**:**

Key Generation: The algorithm $\mathsf{Gen}(1^k)$ samples a random secret key $s \leftarrow \{0,1\}^k$ and a key for an injective one way function $K \leftarrow \mathcal{K}_{\mathcal{OWF}}(1^k)$. Then it computes $p \leftarrow \mathrm{OWF}_K(s)$. It outputs s as the secret key, and the program $PK(\cdot) \leftarrow i\mathcal{O}(F_{p,K}(\cdot))$ as the public key (the circuit $F_{p,K}$ is described in figure 2).

Encryption: The algorithm $\mathsf{Enc}(PK, m)$ on input a public key PK (which is interpreted as an obfuscated program) and a message $m \in \{0,1\}^k$ outputs an obfuscated circuit $C \leftarrow i\mathcal{O}(G_{PK,m}(\cdot))$ (the circuit $G_{PK,m}$ is described in figure 3).

$F_{p,K}$:
Constants: $p \in \{0,1\}^*, K \in \{0,1\}^{l(k)}$
Inputs: $x \in \{0,1\}^k$

1. if $(\mathrm{OWF}_K(x) \stackrel{?}{=} p)$, then output 1; else output $\perp$.

$F'_{p,r,K}$:
Constants: $p \in \{0,1\}^*, r \in \{0,1\}^k, K \in \{0,1\}^{l(k)}$
Inputs: $x \in \{0,1\}^k$

1. if $(\mathrm{OWF}_K(x \oplus r) \stackrel{?}{=} p)$, then output 1; else output $\perp$.

All the circuits are padded to a specific length $\ell_{PK}(k, N, L)$ which will be specified in the proof.

Figure 2: Circuits used in the key generation of the (N-L)-Bounded KDM scheme

$G_{PK,m}$:
Constants: $PK \in \{0,1\}^{\ell'_{PK}(k,N,L,i\mathcal{O})}, m \in \{0,1\}^k$
Inputs: $x \in \{0,1\}^k$

1. If $(PK(x) \stackrel{?}{=} 1)$, then output m; else output $\perp$.

$G'_{PK,\mathbf{r},h,i}$:
Constants: $PK \in \{0,1\}^{\ell'_{PK}(k,N,L,i\mathcal{O})}, \mathbf{r} = (r_1, \ldots, r_N) \in \{0,1\}^{k \cdot N}, h \in \mathcal{H}, i \in \{1, \ldots, N\}$
Inputs: $x \in \{0,1\}^k$

1. If $(PK(x \oplus r_i) \stackrel{?}{=} 1)$, then output $h(x \oplus r_i \oplus r_1,\ x \oplus r_i \oplus r_2,\ \ldots\ ,\ x \oplus r_i \oplus r_N)$; else output $\perp$.

All the circuits are padded to a specific length $\ell_{\mathsf{Enc}}(k, N, L, i\mathcal{O})$ which will be specified in the proof.

Figure 3: Circuits used in the encryption of the (N-L)-Bounded KDM scheme

Decryption: The algorithm $\mathsf{Dec}(s, \overline{C})$ on input a secret key $s \in \{0,1\}^k$ and a ciphertext $\overline{C} \in \mathcal{P}$ outputs $m' = \overline{C}(s)$.

It can be easily verified that correctness of the Obfuscator implies correctness of the encryption scheme. The following theorem argues that the scheme achieves Bounded KDM Security.

Theorem 12. *If $i\mathcal{O}$ is an indistinguishability obfuscator for P/poly and there exists a family of injective $\mathcal{OWF}$, then for any polynomial function L and any $N \in \mathbb{N}$ the encryption scheme $\Pi_{N,L} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ described above is (N, L)-Bounded KDM secure.*

Proof. The proofs proceeds by an hybrid argument, with a very similar structure to Theorem 9. Given any adversary $\mathcal{A}$ define the following random variables:

Z_1**:** this is the same as $\mathbf{KDM}^1_{\Pi_{N,L},N,\mathcal{A}}(k)$.

Z_2**:** this is the same as the previous one, but the public keys are generated as follows: the challenger samples $s_1, \ldots, s_N \leftarrow \{0,1\}^k$ and $K \leftarrow \mathcal{K}_{\mathcal{OWF}}$; then it computes $p \leftarrow \mathrm{OWF}_K(s_1)$, $\vec{r} \leftarrow (0^k, s_1 \oplus s_2, s_1 \oplus s_3, \ldots, s_1 \oplus s_n)$ and sets $PK_i \leftarrow i\mathcal{O}(F'_{p,r_i,K})$ for all i.

Z_{Hyb}**:** this is the same as the previous one, but queries (h, i) by the adversary are answered by returning as the ciphertext an obfuscation $i\mathcal{O}(G'_{PK_1,\vec{r},h,i})$. (Note that knowing the secret keys is not needed to simulate this step).

Z_3**:** this is the same as Z_0, but the public keys are generated (as in Z_2) as follows: the challenger samples $s_1, \ldots, s_N \leftarrow \{0,1\}^k$ and $K \leftarrow \mathcal{K}_{\mathcal{OWF}}$; then it computes $p \leftarrow \mathrm{OWF}_K(s_1)$, $\vec{r} \leftarrow (0^k, s_1 \oplus s_2, s_1 \oplus s_3, \ldots, s_1 \oplus s_n)$ and sets $PK_i \leftarrow i\mathcal{O}(F'_{p,r_i,K})$ for all i.

Z_0**:** this is the same as $\mathbf{KDM}^0_{\Pi_{N,L},N,\mathcal{A}}(k)$.

Again, to prove that Z_1 and Z_0 are indistinguishable, it is enough to show that any pair of adjacent games Z_i, Z_j are also indistinguishable, i.e. for all $\mathcal{A}$, $|\Pr[Z_i = 1] - \Pr[Z_j = 1]| < negl(k)$.

To prove that Z_1 is indistinguishable from Z_2, we rely again on a hybrid argument: for $j = 0, \ldots, N(k)$, we define $Z_{1,j}$ to be the same as Z_1, but where the public keys are sampled as follows. The challenger samples $s_1, \ldots, s_{N(k)} \leftarrow \{0,1\}^k$ and $K_1, \ldots, K_{N(k)} \leftarrow \mathcal{K}_{\mathcal{OWF}}$, then it computes $p_i \leftarrow \mathrm{OWF}_K(s_i)\ \forall i$ and $\vec{r} \leftarrow (0^k, s_1 \oplus s_2, s_1 \oplus s_3, \ldots, s_1 \oplus s_{N(k)})$; finally, it samples $PK_i \leftarrow i\mathcal{O}(F_{p_i,K_i})$ if $i \leq N(k) - j$ and $PK_i \leftarrow i\mathcal{O}(F'_{p_1,r_i,K_1})$ if $i > N(k) - j$. Note that $Z_{1,0}$ has the same distribution as Z_1, and $Z_{1,N(k)}$ has the same distribution as Z_2: an adversary distinguishing between the two with non negligible probability must also distinguish between $Z_{1,j}$ and $Z_{1,j+1}$ with non negligible probability for at least one value of j. But, similarly as in the proof of lemma 11 and because $\mathcal{OWF}$ is an injective family, such an adversary must distinguish between $Z_{1,j}$ and $Z_{1,j+1}$ even conditioned on the fact that K_1 and K_j both denote injective functions. But under this assumption, the two versions of PK_j given to the adversary in the two hybrids (i.e. $i\mathcal{O}(F_{p_j,K_j})$ and $i\mathcal{O}(F'_{p_1,r_j,K_1})$) are functionally equivalent, and therefore (by the security of the $i\mathcal{O}$) computationally indistinguishable, which implies that no such efficient adversary can exist.

The proof that Z_2 is indistinguishable from Z_{Hyb} and the one that Z_{Hyb} is indistinguishable from Z_3 are analogous to the ones of lemma 10 and 11. Note that in such proofs it is only necessary to condition on K_1 being injective (all other function keys do not affect the view of the adversary), and in the second proof $\mathcal{B}$ receives (p_1, K_1) as the challenge to invert and will sample $\vec{r}$ uniformly:

conditioned on K_1 being injective, any input x on which two ciphertexts encrypted under public key PK_i differ can be transformed into a preimage $x \oplus r_i$ for p_1, K_1.

The proof that Z_3 is indistinguishable from Z_0 is analogous to the one that Z_1 is indistinguishable from Z_2. To make sure we can rely on the security of the $i\mathcal{O}$ in all the above hybrids, we first set $\ell_{PK}(k, N, L)$ to be an upper bound on the maximum size of the circuits $F_{p,K}$ and $F'_{p,r,K}$ of figure 2 (which depends on k, N, L). Then we can define $\ell'_{PK}(k, N, L, i\mathcal{O})$ (which in turns determines the size of $G_{PK,m}$ and $G'_{PK,\mathbf{r},h,i}$) to be a bound on the size of the output of $i\mathcal{O}$ on input a circuit of length $\ell_{PK}(k, N, L)$. At last, we set $\ell_{\mathsf{Enc}}(k, N, L, i\mathcal{O})$ to be an upper bound on the size of circuits $G_{PK,m}$ and $G'_{PK,\mathbf{r},h,i}$. Notice that all these functions are polynomially bounded. □

Combining the above theorem with the construction of a family of one way permutations by Bitansky, Paneth and Wichs [BPW16] gives the following corollary:

Corollary 13. *If there exists an indistinguishability obfuscator for P/poly, and a family of one way functions, then for any polynomial function L and any $N \in \mathbb{N}$ there exists a (N, L)-Bounded KDM secure public key encryption scheme.*

Avoiding reliance on a specific OWF As an additional interesting result, we note that the scheme can be further simplified so that its security does not depend on a specific (injective) one way function, but rather on the existence of an (injective) one way functions which can be computed by a circuit whose size is below an explicitly specified bound. The idea behind this construction is the fact that in our encryption scheme the public key is just an obfuscation of a point function, and therefore an obfuscated public key is indistinguishable from the obfuscation of a program (padded to an appropriate size) which directly checks if its input x is equal to the secret s (as opposed to checking whether $\mathrm{OWF}_K(x) = s$) and therefore does not have to internally compute the one way function.

We can leverage this fact and design a new encryption scheme where the key generation algorithm chooses a secret key s uniformly at random and outputs as the public key the obfuscation of a program that returns 1 iff its input $x \stackrel{?}{=} s$ (encryption and decryption algorithms are unchanged). To prove security, we argue that the public keys of this new scheme are indistinguishable from the ones of $\Pi_{N,L,i\mathcal{O}}$ (when instantiated with a secure OWF) and therefore reduce to the latter's Bounded KDM security.

Bounded KDM-CCA2 security Our construction can also be used to construct Bounded KDM-CCA2 security in which, informally, the KDM adversary may also ask CCA2 queries on any ciphertexts except for the ones received as answers to KDM queries (see [CCS09] for a formal definition). Camenish, Chandran and Shoup [CCS09] show a generic transformation for any KDM secure encryption scheme into a KDM-CCA2 secure one (w.r.t. the same family of functions) by applying the Naor-Yung paradigm [NY90]. Their transformation thus requires an NIZK proof system and a CCA2-secure (normal) encryption scheme and a strongly secure one time signature. Combining our construction with their Theorem 1, and the construction of NIZK, CCA2 and signatures from sub-exponentially secure iO and one-way functions from [SW14], we get the following corollary:

Corollary 14. *If there exists a sub-exponentially secure indistinguishability obfuscator for P/poly, and a family of one way functions, then for any polynomial function L and any $N \in \mathbb{N}$ there exists a (N, L)-Bounded KDM-CCA2 secure public key encryption scheme.*

References

[ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO'09*. Springer, 2009.

[AP16] Navid Alamati and Chris Peikert. Three's compromised too: Circular insecurity for any cycle length from (ring-) lwe. Technical report, Cryptology ePrint Archive, Report 2016/110, 2016.

[App14] Benny Applebaum. Key-dependent message security: Generic amplification and completeness. *Journal of Cryptology*, 27(3), 2014.

[AS15] Gilad Asharov and Gil Segev. Limits on the power of indistinguishability obfuscation and functional encryption. In *FOCS'15*. IEEE, 2015.

[BCP14] Elette Boyle, Kai-Min Chung, and Rafael Pass. On extractability obfuscation. In *TCC'14*. Springer, 2014.

[BG10] Zvika Brakerski and Shafi Goldwasser. Circular and leakage resilient public-key encryption under subgroup indistinguishability. In *CRYPTO'10*. Springer, 2010.

[BGK09] Zvika Brakerski, Shafi Goldwasser, and Yael Kalai. Circular-secure encryption beyond affine functions. Technical report, Citeseer, 2009.

[BHHI10] Boaz Barak, Iftach Haitner, Dennis Hofheinz, and Yuval Ishai. Bounded key-dependent message security. In *EUROCRYPT'10*. Springer, 2010.

[BHHO08] Dan Boneh, Shai Halevi, Mike Hamburg, and Rafail Ostrovsky. Circular-secure encryption from decision diffie-hellman. In *CRYPTO'08*. Springer, 2008.

[BPW16] Nir Bitansky, Omer Paneth, and Daniel Wichs. Perfect structure on the edge of chaos. In *TCC'16*. Springer, 2016.

[BRS02] John Black, Phillip Rogaway, and Thomas Shrimpton. Encryption-scheme security in the presence of key-dependent messages. In *Selected Areas in Cryptography*, volume 2595. Springer, 2002.

[CCS09] Jan Camenisch, Nishanth Chandran, and Victor Shoup. A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In *EUROCRYPT'09*, 2009.

[CL01] Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *EUROCRYPT'01*. Springer, 2001.

[GGH+13] Shelly Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Anant Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS'13*. IEEE, 2013.

[HH09] Iftach Haitner and Thomas Holenstein. On the (im) possibility of key dependent encryption. In *TCC'09*. Springer, 2009.

[KRW15] Venkata Koppula, Kim Ramchen, and Brent Waters. Separations in circular security for arbitrary length key cycles. In *TCC'15*. Springer, 2015.

[KW16] Venkata Koppula and Brent Waters. Circular security counterexamples for arbitrary length cycles from lwe. Technical report, Cryptology ePrint Archive, Report 2016/117, 2016.

[MO14] Antonio Marcedone and Claudio Orlandi. Obfuscation$\rightarrow$(ind-cpa security$\not\rightarrow$ circular security). In *SCN'14*. Springer, 2014.

[MTY11] Tal Malkin, Isamu Teranishi, and Moti Yung. Efficient circuit-size independent public key encryption with kdm security. In *EUROCRYPT'11*. Springer, 2011.

[NY90] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC'90*. ACM, 1990.

[SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: Deniable encryption, and more. In *STOC '14*. ACM, 2014.

[Wee16] Hoeteck Wee. Kdm-security via homomorphic smooth projective hashing. In *PKC'16*, 2016.