

Bounded Key-Dependent Message Security

Boaz Barak^{1,*}, Iftach Haitner², Dennis Hofheinz^{3,**}, and Yuval Ishai^{4,***}

¹ Princeton University
boaz@cs.princeton.edu

² Microsoft Research
iftach@microsoft.com

³ Karlsruhe Institute of Technology
Dennis.Hofheinz@kit.edu

⁴ Technion and UCLA
yuvali@cs.technion.ac.il

Abstract. We construct the first public-key encryption scheme that is proven secure (in the standard model, under standard assumptions) even when the attacker gets access to encryptions of arbitrary efficient functions of the secret key. Specifically, under either the DDH or LWE assumption, and for arbitrary but fixed polynomials L and N , we obtain a public-key encryption scheme that resists key-dependent message (KDM) attacks for up to $N(k)$ public keys and functions of *circuit size* up to $L(k)$, where k denotes the size of the secret key. We call such a scheme *bounded KDM secure*. Moreover, we show that our scheme suffices for one of the important applications of KDM security: ability to securely instantiate symbolic protocols with axiomatic proofs of security.

We also observe that any fully homomorphic encryption scheme that additionally enjoys circular security and circuit privacy is *fully KDM secure* in the sense that its algorithms can be independent of the polynomials L and N as above. Thus, the recent fully homomorphic encryption scheme of Gentry (STOC 2009) is fully KDM secure under certain non-standard hardness assumptions.

Finally, we extend an impossibility result of Haitner and Holenstein (TCC 2009), showing that it is impossible to prove KDM security against a family of query functions that contains exponentially hard pseudo-random functions if the proof makes only a *black-box* use of the query function and the adversary attacking the scheme. This shows that the non-black-box use of the query function in our proof of security is inherent.

Keywords: KDM/cliQUE/circular security; fully homomorphic encryption; formal security.

* Supported by NSF grants CNS-0627526, CCF-0426582 and CCF-0832797, US-Israel BSF grant 2004288 and Packard and Sloan fellowships.

** Part of work performed while at CWI and supported by an NWO Veni grant.

*** Supported by BSF grant 2008411, ISF grant 1310/06, and NSF grants 0830803, 0716835, 0627781.

1 Introduction

An encryption scheme is *key-dependent message (KDM) secure* if it is secure even against an attacker who has access to encryptions of messages that depend on the secret key. This strong notion of security, introduced by Black et al. [6], tries to capture scenarios where there could be correlations between the secret key and the encrypted messages. At a first glance, it may seem that such correlations only arise from bugs or errors on part of the protocol designer, and hence achieving such a strong security is not of much importance. It turns out, however, that such attacks naturally occur when considering complex systems. For example, in some popular disk encryption utilities, the disk encryption key can end up being stored in the page file, and thus is encrypted along with the disk content [7]. In addition, Camenisch and Lysyanskaya [9] showed that schemes with a certain restricted form of KDM security known as “circular security” are useful for constructing *Anonymous Credential Systems*. Finally, and perhaps most importantly, KDM security naturally arises as the right notion when one wishes to securely instantiate symbolic protocols with an axiomatic proof of formal security (see Section 6).

For a while, building a KDM-secure encryption scheme in the standard model, under any well studied hardness assumption, seemed too hard a nut to crack. The only scheme that was shown to resist any kind of KDM attacks was given by Black et al. [6] in the random-oracle model. Yet, in recent years KDM-secure encryption schemes were given for some non-trivial families of functions. This line of work started with the works of Halevi and Krawczyk [18] and Hofheinz and Unruh [19], who gave private-key encryption schemes secure against significantly restricted classes of KDM queries. Concretely, [18] prove security against arbitrary but fixed KDM queries that are known in advance, and against KDM queries that do not depend on certain “protected” parts of the key. The constructions from [19] obtain statistical KDM security in the presence of sufficiently few (arbitrary) KDM queries, as well as a stateful KDM-secure scheme in which KDM queries may only depend on the current state (but not on previous states).

A major step was taken by Boneh, Halevi, Hamburg, and Ostrovsky [7] who presented, under the *decisional Diffie-Hellman (DDH)* assumption, a public-key encryption scheme that is $N(k)$ -circular secure for every polynomial N , and in fact is secure against the more general family of attacks allowing the adversary access to encryptions of arbitrary affine functions of the vector of $N(k)$ secret keys. Applebaum, Cash, Peikert, and Sahai [3] presented more efficient schemes that are secure against a similar family of key-dependent attacks, whose security is based on different assumptions: the *learning parity with noise (LPN)* assumption in the secret-key case and the *learning with errors (LWE)* assumption in the public-key case. In a recent independent work, Brakerski, Goldwasser, and Kalai [8] presented a transformation from a KDM secure scheme satisfying a certain property (in particular satisfied by the DDH and LWE based schemes of [7, 3]) into a scheme that is KDM secure with respect to a larger class of functions. While their transformation cannot be used to achieve security against

all circuits of size $p(n)$, it has the benefit of depending only on the number of functions in the class, and being independent of their circuit size or number of keys. In particular they achieve KDM security with respect to the class of constant degree polynomials and any polynomial number of keys.

Despite the above progress, the families of functions for which KDM security was achieved prior to our work (in the standard model, under standard assumptions) was still quite restricted. In particular, these families were not sufficiently rich for several of the applications of KDM security in the context of complex systems and formal protocols. A partial explanation for this rather limited success was recently given by Haitner and Holenstein [17], who showed the impossibility of obtaining KDM security based on standard assumptions and using standard techniques. (In Section 1.2, we will describe their results in more detail, since we will later extend them to our case of bounded KDM security.)

1.1 Our Results

Our main result is the following:

Theorem 1 (Informal). *Under the DDH or LWE assumption, for any given polynomials $L = L(k)$ and $N = N(k)$, there exists a public-key encryption scheme that is KDM-secure with respect to the class of circuits of size $L(k)$, and for $N(k)$ independent keys, where k denotes the size of the keys.*

We call such a scheme a *bounded KDM-secure* encryption scheme. (This is in contrast with a *fully-* or *unbounded-KDM* scheme, where the circuit size and the number of keys can be an arbitrarily large polynomial in the security parameter, independent of the scheme’s complexity.) We argue that this is the first encryption scheme (under standard cryptographic assumptions) that handles a rich enough function class to capture most “real life” KDM attacks.

The original motivation for KDM security was to securely instantiate symbolic cryptographic protocols that have a formal proof of security in some axiomatic system. As further evidence for the usefulness of bounded KDM security, we show that our notion is strong enough for this application:

Theorem 2 (Informal). *Let P be a symbolic protocol with operations such as public-key encryption and digital signatures. Then, instantiating P with a bounded KDM-secure¹ encryption scheme provides a computationally sound implementation.*

This yields the first soundness result without restrictions (such as assuming protocols without key-cyclic expressions) in the standard model.

Finally, we show that the above positive results are tight, by extending an impossibility result of Haitner and Holenstein [17] in the following sense:

¹ Actually, the precise notion we use is *length-dependent* KDM security (see Definition 5). This is a slight strengthening of bounded KDM security, and our scheme satisfies this stronger notion as well.

Theorem 3 (Informal). *An encryption scheme cannot be proven to be KDM-secure against a family of functions that contains exponentially hard pseudorandom functions, if the proof of security only accesses the query function and the adversary attacking the scheme in a black-box manner (i.e., as oracles).*

Remarks. We note the following points about our result:

1. *Efficiency.* Our scheme, although polynomial time, is not practically efficient as it uses the garbled circuit construction and its ciphertext length is at least L , where L is a bound on the circuit size of the KDM function. There are more efficient candidate KDM-secure cryptosystems if one is willing to settle for non-standard assumptions or the random oracle model.
2. *Full KDM security.* Although we only prove our scheme to be bounded KDM secure, it is of course possible that it is KDM secure with respect to *any* efficient KDM function. In fact, there seems to be an interesting obstacle to any KDM attack on our scheme. Suppose that we instantiate the scheme to be secure with respect to KDM functions of size k^3 . Now suppose that there is a successful KDM attack against it, and for simplicity assume the attack consists of getting one encryption of $h(sk)$ where h is some efficiently computable function. Then the success of this attack implies that either DDH is false (assuming we instantiate our scheme from the DDH assumption), or that h has no circuit of size k^3 . Hence, a proof that this construction is insecure against a polynomial-time KDM attacker will provably demonstrate that either DDH is false, or that $\mathbf{P} \not\subseteq \mathbf{Size}(k^2)$ (we lose a factor of k because h has a k -bit output). The latter is a widely believed fact, but its proof would be considered a major breakthrough in complexity theory. (Also, it is not at all clear how to derive such a conclusion directly from the DDH assumption—typically in cryptography we need to use *subexponential* hardness assumptions to get such a condition.) More generally, a successful attack is some way to *certify* that h is hard—even though it is easy in time k^3 to generate a random function outside of $\mathbf{Size}(k^2)$, it is not at all clear how to generate such a function along with a publicly verifiable certificate of hardness.
3. *Black-box-ness.* Our scheme makes a non-black-box use of the KDM function h , where Theorem 3 shows that this use is *inherent*.

Applications to formal security. A central motivation for the study of KDM security lies in the connection between formal and computational cryptography. In formal cryptography (starting with [12, 13, 23]), cryptographic operations like encryption or digital signatures are abstracted as symbolic operators that (only) adhere to natural rules. Given such rules, a simple calculus enables machine-assisted security analysis.

It was proven by Adão et al. [2] that *fully* KDM-secure encryption schemes imply computational soundness for *arbitrary* symbolic protocols. We reconsider their proof and show (Theorem 9) that bounded KDM security of the type that we achieve suffices. Hence, our combined results give the first encryption scheme (under standard cryptographic assumptions) whose security implication can be verified using formal security methods.

We stress that the clique security achieved by [7, 3] only enables to apply these formal methods to a very limited class of applications. For more details see Section 6.

1.2 Our Techniques

We now give an informal overview of the proof of Theorem 1. The following exposition focuses on a scheme that is secure against a *single-key* KDM attack. That is, there is only one public/private key pair (pk, sk) of length k , and the attacker can obtain encryptions of messages of the form $h(sk)$ for an arbitrary function h of circuit complexity at most $L(k)$. (Here $L = L(k)$ is an arbitrary fixed polynomial which affects the complexity of the encryption and decryption, but not the complexity of key generation.) The multiple-key case raises some additional subtleties that we ignore for the moment.

Recall that a *homomorphic* public-key encryption scheme is a public-key encryption scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ that also has an additional algorithm Eval for evaluating functions on an encrypted message. Concretely, Eval takes the public key pk , an encryption of a message M , and a description of a function h from some family \mathcal{H} , and outputs a string from which $h(M)$ can be efficiently decrypted using the secret key sk . Our starting point is the following observation: a sufficiently strong homomorphic encryption is in fact also KDM-secure (with respect to the same class of functions \mathcal{H}), where “sufficiently strong” means that the scheme satisfies the following additional properties:

1. Self-referential (i.e., 1-circular) security: $\text{Enc}_{pk}(sk)$ is indistinguishable from $\text{Enc}_{pk}(0^k)$.
2. Strong function privacy: For every $h \in \mathcal{H}$ and plaintext M , $\text{Eval}_{pk}(h, \text{Enc}_{pk}(M))$ is indistinguishable from $\text{Enc}_{pk}(h(M))$, even against a distinguisher that knows the secret key.

The basic idea for proving the KDM-security of such a scheme is that a distinguisher between $\text{Enc}_{pk}(h(sk))$ and $\text{Enc}_{pk}(h(0^k))$ can be used to distinguish between $\text{Enc}_{pk}(sk)$ and $\text{Enc}_{pk}(0^k)$ by simply running Eval with the function h (and thus the “KDM queries” are useless). When turning this idea into a proof one sees that it is crucial that function privacy hold even with respect to a distinguisher that knows the secret key.

This observation already implies that Gentry’s recent breakthrough fully homomorphic encryption scheme [15] is fully KDM-secure, assuming that it is circular-secure (an assumption which is anyway necessary in Gentry’s case to get a truly fully homomorphic encryption, where the public key does not grow with the depth of the circuit).² Since all *natural* candidates for public-key encryption schemes are not known to be 1-circular insecure, we find this observation interesting, as the assumption of circular security seems cleaner and more conservative than assuming full KDM security. (In particular, it is more easily “falsifiable” in the sense of Naor [25].)

² As Gentry notes, if one assumes his scheme is circular-secure then it also enjoys strong statistical function privacy.

In fact, it turns out that it suffices to have only *weak* function privacy, requiring that $\text{Eval}_{pk}(h, \text{Enc}_{pk}(M))$ be indistinguishable from $\text{Eval}_{pk}(h', \text{Enc}_{pk}(M))$ for h, h', M such that $h(M) = h'(M)$ (again indistinguishability is with respect to attackers who know the secret key).³ See Theorem 5 for the details.

The latter observation suggests an approach to get KDM security for circuits of size L under standard assumptions. Consider any two-message protocol for evaluating a *universal function* with security against semi-honest parties. Such a protocol takes an input M from a receiver and a circuit h from a sender, and delivers the output $h(M)$ to the receiver. Given any such protocol and a standard public-key encryption (PKE), one can construct a homomorphic scheme with *weak* function privacy as follows. The public key is the public key pk of the PKE. The encryption of M under pk is a triple (C, pk', C') where C is an encryption of M under pk , pk' is the receiver's first message in the protocol on input M , and C' is an encryption under pk of the secret randomness sk' used to generate pk' (which is needed to recover the output of the protocol). The algorithm $\text{Eval}((C, pk', C'), h)$ returns the sender's response to pk' on input h along with C' . Given sk , the output of Eval can be used to decrypt $h(M)$ by first recovering sk' and then computing the receiver's output in the protocol.

The advantage of this approach is that it can be instantiated under standard assumptions by using Yao's protocol [28]. More concretely, a secure two-message protocol for the universal function can be obtained by combining Yao's garbled circuit construction and any *two-message oblivious transfer* (OT) [27, 14] protocol.⁴ Unlike the alternative of using *fully* homomorphic encryption, however, this protocol has the caveat that the communication must grow with the size of h , and hence (weak) function privacy can only hold with respect to the class \mathcal{H} of all circuits with some a-priori size bound L .

A more subtle problem is that of making the homomorphic scheme constructed in the above way circular-secure. Indeed, encrypting the secret key of the homomorphic scheme with its own public key results in a circular dependency between the underlying PKE and the two-message protocol: the secret key sk of the PKE is encrypted using the "public key" pk' of the protocol, whereas the "secret key" sk' of the protocol is encrypted using the public key pk of the PKE. Even if the PKE is circular-secure, it is not clear that this property will be respected by the above construction.

Our way to handle this difficulty is by introducing a new notion that we call "targeted encryption", which is aimed towards resolving the above dependency

³ This is indeed a weaker notion since if $y = h(M) = h'(M)$ then one can see that strong function privacy implies that $\text{Eval}(h, \text{Enc}(M)) \approx \text{Enc}(y) \approx \text{Eval}(h', \text{Enc}(M))$. Intuitively, weak function privacy allows Eval to map ciphertexts from one domain to a different domain, while this is ruled out by strong function privacy.

⁴ Two-message OT is a protocol comprised of one message from the receiver and one message from the sender, where the receiver has an input selection bit s and the sender has a pair of input strings X_0, X_1 . In the end of the protocol the receiver only learns X_s and the sender learns nothing about s . Here we need k parallel instances of OT, where the receiver has a k -bit input selection vector s and the sender has k pairs of strings.

when applied to a two-message protocol based on Yao’s technique. Targeted encryption can be viewed as a circular-secure extension of both public-key encryption and two-message OT. Loosely speaking, one can think of this as an OT protocol where the receiver has *no secret information* apart from the input selection vector s . This may look strange at first, and indeed it can be shown to be inherently at odds with the standard notion of OT, which requires that the sender learn nothing about s . But it turns out that one can obtain a meaningful relaxation of the above notion that is strong enough for our purposes. We then show that both the schemes of Boneh et al. [7] and Applebaum et al. [3] can be used to construct targeted encryption. The key property we use is that both schemes enjoy KDM security against *affine functions*, and in fact this is proven by giving a public algorithm to compute an encryption of any affine function of the secret key. We show that such an algorithm implies targeted encryption.

As mentioned above, multiple-key security adds some additional difficulties. In particular, targeted encryption on its own does not seem sufficient for *multiple key* security, and to handle this case we need to appeal to the multiple-key circular security of the underlying schemes.

To show our negative result (Theorem 3), we employ the techniques of Haitner and Holenstein [17]. Concretely, they showed that an encryption scheme cannot be proved to be KDM-secure against the family of *all* functions, if the proof of security only accesses the query function and the adversary attacking the scheme in a black-box manner (i.e., as oracles).⁵ Here we extend this result to *every* family of functions that contains exponentially hard pseudorandom functions. There was no prior scheme that was shown (under a standard assumption) to be secure with respect to such a family, although many of the applications of KDM security require that the KDM function can be a cryptographic primitive such as a signature, a hash function, etc.

2 Preliminaries

Notation. For $n \in \mathbb{N}$, let $[n] := \{1, \dots, n\}$. Throughout the paper, $k \in \mathbb{N}$ denotes the security parameter. For a finite set X , we denote by $x \leftarrow X$ the process of sampling x uniformly from X . For a probabilistic algorithm A , we denote by $y \leftarrow A(x)$ the process of running A on input x and with uniform randomness, and assigning y the result. If A runs in time polynomial in the security parameter k , then A is a PPT machine. (We always assume that k can be efficiently computed from the input to the algorithm even if it not explicitly given.) A function $f : \mathbb{N} \rightarrow [0, 1]$ is *negligible* iff $\forall c \in \mathbb{N} \exists k_0 \in \mathbb{N} \forall k > k_0 : |f(k)| < k^{-c}$. We say f is *overwhelming* iff $1 - f$ is negligible. Two collections $X = (X_k)_{k \in \mathbb{N}}$ and $Y = (Y_k)_{k \in \mathbb{N}}$ of random variables are *computationally indistinguishable*, written

⁵ They also showed that it is impossible to prove (in a black-box way) that a trapdoor-permutation based scheme is KDM-secure against a family of t -wise independent hash functions, for t that is longer than the ciphertext size (here a non-black-box access to the query function is allowed).

$X \overset{c}{\approx} Y$, iff for every nonuniform polynomial-time distinguisher D , we have that $\Pr [D(1^k, X_k) = 1] - \Pr [D(1^k, Y_k) = 1]$ is negligible. We use \circ for concatenation.

Encryption schemes. A public-key encryption (PKE) scheme with message space $\mathcal{M} = \mathcal{M}_k$ and secret key space $\mathcal{K} = \mathcal{K}_k$, consists of three algorithms $(\text{Gen}, \text{Enc}, \text{Dec})$ — Key generation $\text{Gen}(1^k)$ outputs a public key pk and a secret key $sk \in \mathcal{K}_k$. Encryption $\text{Enc}_{pk}(M)$ takes a public key pk and a message $M \in \mathcal{M}_k$, and outputs a ciphertext C . Decryption $\text{Dec}_{sk}(C)$ takes a secret key sk and a ciphertext C , and outputs a message M . For correctness, we require $\text{Dec}_{sk}(C) = M$ for all $M \in \mathcal{M}_k$, all (pk, sk) in the range of $\text{Gen}(1^k)$, and all C in the range of $\text{Enc}_{pk}(M)$. For simplicity, we will assume from now on that both the key space and the message space are $\{0, 1\}^k$. Our definitions and results, however, can be easily adapted to the case of messages of arbitrary length.

2.1 Garbled Circuits

An essential building block of our KDM secure encryption scheme is Yao’s garbled circuit construction, attributed to [28]. Informally, the variant of this construction on which we rely transforms any circuit h with k input bits along with k pairs of random keys $(K_{1,0}, K_{1,1}), \dots, (K_{k,0}, K_{k,1})$ into a “garbled circuit” GC such that the following properties hold:

- For any input $x \in \{0, 1\}^k$ and any choice of $2k$ keys, the output $h(x)$ can be efficiently decoded (without knowing h) from GC and the k keys K_{i,x_i} corresponding to x .
- GC together with the k keys corresponding to x computationally hide all information about h other than the size of h and $h(x)$.
- GC alone computationally hides all information about h other than its size,

where the last two properties hold with respect to a random choice of the keys and a random execution of the transformation. The existence of a construction satisfying the above requirements is formally captured by the following theorem. See the full version of this paper [5] for a derivation of this theorem from the literature.

Theorem 4 (Garbled circuits). *Suppose that one-way functions exist. Then there is a pair of polynomial-time randomized algorithms $(\text{Garble}, \text{GCEval})$ that for security/input parameter k , output parameter m , and circuit size parameter s satisfy the following:*

Syntax. *Garble takes a $2k$ key tuple $\overline{K} = \{K_{i,b}\}_{i \in [k], b \in \{0,1\}}$, where $K_{i,b} \in \{0, 1\}^k$, and a size s circuit describing a function $h : \{0, 1\}^k \rightarrow \{0, 1\}^m$, and outputs a “garbled circuit” GC . GCEval takes an input $x \in \{0, 1\}^k$, a k key tuple, and a garbled circuit GC and outputs $y \in \{0, 1\}^m$.*

Correctness. *We require that if $GC = \text{Garble}(\overline{K}, h)$ then $\text{GCEval}(\overline{K}_x, GC) = h(x)$, where we define $\overline{K}_x = \{(x_i, K_{i,x_i})\}_{i \in [k]}$.*⁶

⁶ For ease of notation we assume that the input x is included in the description of \overline{K}_x . This is needed to guarantee correctness even when a pair of keys happen to be identical. Alternatively, we could avoid giving x as input to GCEval by either settling for statistical correctness or allowing the keys to be correlated.

Security against receiver. For every polynomials $s(k), m(k)$, every $x \in \{0, 1\}^k$ and every $h, h' : \{0, 1\}^k \rightarrow \{0, 1\}^{m(k)}$ of size $s(k)$ such that $h(x) = h'(x)$, if \overline{K} is chosen at random then

$$\overline{K}_x \circ \text{Garble}(\overline{K}, h) \stackrel{c}{\approx} \overline{K}_x \circ \text{Garble}(\overline{K}, h')$$

Security against outsiders. For every polynomials $s(k), m(k)$ and every $h, h' : \{0, 1\}^k \rightarrow \{0, 1\}^{m(k)}$ of size $s(k)$, if \overline{K} is chosen at random then

$$\text{Garble}(\overline{K}, h) \stackrel{c}{\approx} \text{Garble}(\overline{K}, h')$$

2.2 Key-Dependent Message Security

Loosely speaking, the notion of *key-dependent message (KDM) security* gives an adversary access to encryptions of messages of the form $h(sk)$, where $h : \mathcal{K} \rightarrow \mathcal{M}$ is a function that the adversary can choose from some family. The formal definition below is taken from Black et al. [6] and allows the function to depend on some $N = N(k)$ secret keys. While handling multiple keys is important for the application to formal cryptography (see Section 6), much of the technical challenge is already manifested in the case $N = 1$, and so the reader may wish to focus on this case initially.

Definition 1 (KDM security). Let $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ be a *public-key encryption scheme with message space \mathcal{M} and secret key space \mathcal{K}* . Let $\overline{pk} := (\overline{pk}_1, \dots, \overline{pk}_N)$ and $\overline{sk} := (\overline{sk}_1, \dots, \overline{sk}_N)$ be *public, resp., secret key vectors, where $N = N(k) > 0$ is a positive-valued function*. Let A be a PPT machine. Let

- $\text{Real}_{\overline{pk}, \overline{sk}}$ be the oracle that on input a function $h : \mathcal{K}^N \rightarrow \mathcal{M}$ (encoded as a circuit) and $\mu \in [N]$ returns $C \leftarrow \text{Enc}(\overline{pk}_\mu, h(\overline{sk}))$, and
- $\text{Fake}_{\overline{pk}}$ be the oracle that on input h, μ as above returns $C \leftarrow \text{Enc}(\overline{pk}_\mu, 0^k)$.

The KDM advantage of A is

$$\text{Adv}_{\text{PKE}, A}^{\text{KDM}}(k) := \Pr \left[A^{\text{Real}_{\overline{pk}, \overline{sk}}(\cdot, \cdot)}(\overline{pk}) = 1 \right] - \Pr \left[A^{\text{Fake}_{\overline{pk}}(\cdot, \cdot)}(\overline{pk}) = 1 \right]$$

where $(\overline{pk}_i, \overline{sk}_i) \leftarrow \text{Gen}(1^k)$ for $i \in [N]$ in both probabilities. We say that PKE is KDM secure with respect to a function class \mathcal{H} iff for every polynomial N and every PPT A that only queries its oracle with functions $h \in \mathcal{H}$, the advantage function $\text{Adv}_{\text{PKE}, A}^{\text{KDM}}$ is negligible in the security parameter. PKE is fully KDM secure iff PKE is KDM secure with respect to the class \mathcal{H} that consists of all functions.

Examples of KDM function classes. The following examples of KDM function classes will be important for us.

Clique/circular security. Let \mathcal{S}_N consist of all functions $h_i : (\{0, 1\}^k)^N \rightarrow \{0, 1\}^k$ for $i \in [N]$, where $h_i(\overline{sk}_1, \dots, \overline{sk}_N) = \overline{sk}_i$. Thus, KDM security with respect to \mathcal{S}_N allows the adversary to obtain encryptions $\text{Enc}_{\overline{pk}_i}(\overline{sk}_j)$ for every $i, j \in [N]$. This was called “clique security” by Boneh et al. [7] who gave a scheme that is KDM secure with respect to \mathcal{S}_N for every N that is polynomial in the security parameter. (See Applebaum

et al. [3] for another construction.) Security with respect to \mathcal{S}_N implies “ N -circular security”. This notion, defined by [9] states that for independently generated N key pairs $(pk_1, sk_1), \dots, (pk_N, sk_N)$, the vector of N encryptions $\text{Enc}_{pk_1}(sk_2), \text{Enc}_{pk_2}(sk_3), \dots, \text{Enc}_{pk_N}(sk_1)$ is indistinguishable from $\text{Enc}_{pk_1}(0^k), \dots, \text{Enc}_{pk_N}(0^k)$.

Bounded security. Let $\mathcal{C}_{N,L}$ consist of all functions $h : (\{0, 1\}^k)^N \rightarrow \{0, 1\}^k$ that can be described with circuits of size at most L . We say that a scheme is (N, L) bounded KDM secure, if it is KDM secure with respect to $\mathcal{C}_{N(k),L(k)}$, where k denotes both the security parameter and the secret key size.⁷

Full (unbounded) security. Full KDM security is equivalent to requiring that a scheme is KDM secure with respect to $\mathcal{C}_{N,L}$ for every polynomials, in the security parameter, N and L . Note that this definition seems like the best one should look for, since a PPT adversary cannot generate circuits (i.e., queries) of superpolynomial size.

Finally, we say that a scheme has *single-key KDM security*, if in the KDM attack above the number of keys N is restricted to being 1. This notion makes sense with respect to bounded/unbounded security, where in the case of or clique or circular security it is equivalent to “self reference security” — the adversary has access to $\text{Enc}_{pk}(sk)$.

2.3 KDM Security from Homomorphic Encryption

In this section we observe that one can get KDM security from a certain kind of homomorphic encryption schemes.

Definition 2 (Homomorphic encryption). Let $\mathcal{H} = \{\mathcal{H}_k\}$ be a sequence of sets of Boolean circuits. A tuple of algorithms $\xi = (\text{Gen}, \text{Enc}, \text{Dec}, \text{Eval})$ is a homomorphic encryption scheme with respect to \mathcal{H} , if $(\text{Gen}, \text{Enc}, \text{Dec})$ is a public key encryption scheme, and in addition for every $(pk, sk) \leftarrow \text{Gen}(1^k)$, $h \in \mathcal{H}_k$ and message $M \in \mathcal{M}$

$$\text{Dec}_{sk}(\text{Eval}_{pk}(h, \text{Enc}_{pk}(M))) = h(M)$$

We say that ξ satisfies strong (statistical) function-privacy if for every $h \in \mathcal{H}_k$, pk in the range of $\text{Gen}(1^k)$ and $M \in \mathcal{M}$, $\text{Eval}_{pk}(h, \text{Enc}_{pk}(M)) \stackrel{s}{\approx} \text{Enc}_{pk}(h(M))$.

We say that ξ satisfies weak (statistical) function-privacy if for every $h, h' \in \mathcal{H}_k$, pk in the range of $\text{Gen}(1^k)$ and $M \in \mathcal{M}$, if $h(M) = h'(M)$ then $\text{Eval}_{pk}(h, \text{Enc}_{pk}(M)) \stackrel{s}{\approx} \text{Eval}_{pk}(h', \text{Enc}_{pk}(M))$.⁸

⁷ Requiring the secret key to be at most k prevents trivialities such as making the key so big that $L(k)$ -sized circuits don't have time to read it. In fact, our scheme will satisfy a slightly stronger notion which is that the key generation algorithm will be completely independent of N and L , see Definition 5.

⁸ One can naturally define *computational* versions of weak and strong function privacy, in which case one needs to allow the distinguisher to get the secret key as part of the input, and in some applications also the randomness used to generate this secret key.

We say that a scheme is *fully homomorphic* if (1) for every polynomial $s = s(k)$ it is homomorphic with respect to the family $\mathcal{H} = \{\mathcal{H}_k\}$, where \mathcal{H}_k is the set of all Boolean circuits of size at most $s(k)$, and (2) the running time (and hence also output size) of both the encryption and decryption algorithm is a fixed polynomial in the security parameter k . It was a longstanding open problem to come up with even a plausible candidate for such a scheme, until this was achieved this year by Gentry [15], who gave such a candidate based on ideal lattices.⁹ If a scheme satisfies only (1) (but not necessarily (2)) then we say that it is *size-dependent* homomorphic encryption. There is a trivial construction of a size dependent homomorphic encryption: just have Eval concatenate the circuit to the ciphertext. Using Yao's garbled circuit construction and two-message OT one can get a size-dependent homomorphic encryption with weak function privacy. In contrast, *strong* function privacy for this class \mathcal{H} implies condition (2).

As mentioned in Section 1.2, we observe that a homomorphic encryption scheme with respect to a class \mathcal{H} that is strongly function-private and is circular secure, is also KDM secure with respect to the same class. This already implies that Gentry's scheme is fully KDM secure under certain assumptions that do not refer to *full* KDM security (i.e., hardness of a certain bounded-distance decoding problem on ideal lattices, a sparse subset sum problem, and assuming the scheme is *circular* secure). Moreover, for this application we can relax the condition to *weak* function-privacy:

Theorem 5. *Suppose that there is a homomorphic encryption scheme with respect to a class \mathcal{H} that is weakly function private and 1-circular secure. Then there is a single-key KDM-secure scheme with respect to the same class \mathcal{H} .*

Proof Sketch. Let $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Eval})$ be the homomorphic encryption scheme. Our encryption scheme $(\text{Gen}', \text{Enc}', \text{Dec}')$ will be as follows:

Key Generation. $\text{Gen}'(1^k)$ runs $(pk, sk) \leftarrow \text{Gen}(1^k)$ and outputs the same secret key sk and as public key the concatenation of pk and $C = \text{Enc}_{pk}(sk)$.

Encryption. $\text{Enc}'_{pk,C}(M)$ outputs $\text{Eval}_{pk}(\text{const}^M, C)$, where const^M is the constant function that always outputs M .

Decryption. We have $\text{Dec}' = \text{Dec}$.

Correctness follows easily from the homomorphic property. For security, consider a KDM attacker, that queries an oracle with h and gets back $\text{Enc}'_{pk,C}(h(sk)) = \text{Eval}_{pk}(\text{const}^{h(sk)}, C)$. We proceed by a hybrid argument. Suppose that the oracle was changed so that it returned $\text{Eval}_{pk}(h, C)$. Since C is an encryption of sk , and obviously $\text{const}^{h(sk)}(sk) = h(sk)$, and the scheme satisfies weak function-privacy this will not change the attacker's output distribution (Since we need the secret key to compute $\text{const}^{h(sk)}$, we will need here to use the fact that weak function-privacy holds even with respect to distinguishers that know the secret key).

⁹ The fully homomorphic version of Gentry's scheme requires three assumptions: hardness of a certain bounded-distance decoding problem on ideal lattices, hardness of a sparse version of subset sum, and circular security of his basic ideal-lattice based scheme.

The new oracle, however, can be simulated by the attacker on its own (since it does not use the secret key at all, but only h and C). Hence, we complete the proof by appealing to the circular security of the encryption, to argue that C might as well be an encryption of “junk”. ■

As a corollary, assuming the *circular* security of a version of Paillier’s cryptosystem [26, 11], the homomorphic PKE construction from [20] yields a KDM-secure encryption scheme with respect to the class of branching programs of a bounded (polynomial) *length*, but unbounded (polynomial) *size*. In other words, the length of the ciphertexts should only depend on the length of branching programs computing the KDM function but not on their size. Compared to the alternative based on the circular-secure version of Gentry’s scheme, the conclusion is much weaker but the assumption is different (and seemingly more conservative).

3 Targeted Encryption

The main tool we use to realize our KDM secure scheme is a new notion we call *targeted encryption*. This is a variant of a public key encryption scheme that has the following curious property: the encryption algorithm gets, apart from the message x , two additional inputs: an index $i \in [k]$ (where k is the bit length of the secret key), and a bit b . The decryption algorithm successfully retrieves x if the i^{th} bit of the secret key is b , but otherwise gets no information about x .¹⁰

Definition 3 (Targeted encryption). *An targeted encryption scheme TES consists of a tuple of algorithms (TGen, TEnc, TDec) such that on security parameter k , TGen outputs a pair (pk, sk) with $sk = (sk_1, \dots, sk_k) \in \{0, 1\}^k$ and:*

Targeted decryption. *For every message $x \in \{0, 1\}^n$ and index $i \in [k]$,*

$$\text{TDec}_{sk}(\text{TEnc}_{pk,i,sk_i}(x)) = x .$$

I.e., it is possible for a sender, given (i, b) , to encrypt a message x such that the following hold: if the i^{th} bit of the secret key is b , then the receiver decrypts this message successfully.

(Statistical) security against receiver. *For every $x, x' \in \{0, 1\}^n$ and index $i \in [k]$,*

$$\text{TEnc}_{pk,i,1-sk_i}(x) \stackrel{s}{\approx} \text{TEnc}_{pk,i,1-sk_i}(x') .$$

I.e., if the i^{th} bit of the secret key is not b , then the receiver gets no information about the message x .¹¹

¹⁰ We do not actually require a targeted encryption to also have a standard (“untargeted”) encryption algorithm, that always succeeds although this can easily be achieved by, say, concatenating two encryptions using parameters $i, 0$ and $i, 1$. Later, to achieve multiple-key security, we will need to assume such an algorithm with particular properties, see Section 5.

¹¹ For our purposes we can relax this notion to computational indistinguishability with respect to distinguishers that get the secret key as additional input.

Security against outsiders. For every $x, x' \in \{0, 1\}^n$, index $i \in [k]$, and $b \in \{0, 1\}$,

$$pk \circ \text{Enc}_{pk,i,b}(x) \stackrel{c}{\approx} pk \circ \text{Enc}_{pk,i,b}(x').$$

I.e., outsiders, who do not know the secret key, get no information about the encryption, even if the i^{th} bit of sk does equal b .

The next theorem states that targeted encryption scheme can be obtained from either the DDH or the LWE assumptions.

Theorem 6. Suppose that (1) the DDH Assumption holds, or (2) the LWE assumption holds (with certain parameters),¹² then there exists a targeted encryption scheme.

Theorem 6 is proven by showing that targeted encryption is implied by both the work of Boneh et al. [7] and the work of Applebaum et al. [3] (see the full version [5] of this paper for the formal proof). The idea of the proof is as follows. Both works give schemes that are KDM secure with respect to affine functions over \mathbb{Z}_q^k for some number q , where k being the secret key size. Their proofs,¹³ however, actually give the following stronger homomorphic property: there exists an algorithm Eval that gets the public key and an affine function $h : \mathbb{Z}_q^k \rightarrow \mathbb{Z}_q^k$, and outputs an encryption of $h(sk)$ that is statistically indistinguishable from $\text{Enc}_{pk}(h(sk))$. Note that this is a property that indeed immediately implies KDM security for affine functions. We will use this property to get the following targeted encryption scheme: to encrypt a message $x \in \{0, 1\}^n$ so that it can only be decrypted if the i^{th} bit of the key is b , we view x as an element inside \mathbb{Z}_q (using some natural embedding for large enough q , where if n is too large, we encrypt x in chunks) and choose a random $r \in \mathbb{Z}_q$ and use the encryption of scheme to encrypt $r \cdot (sk_i - b) + x$. Note that this is an affine function of sk ,¹⁴ and its value is independent of x if $sk_i \neq b$, but is equal to x otherwise. Some complications arise from the fact that in [7] the group is actually given “in the exponent”, where in [3] the key is not a bit string, but rather a vector in \mathbb{Z}_q^k . Nevertheless, these issues can be easily handled in both cases.

Discussion — Targeted encryption and oblivious transfer. Recall that in a (one out of two) *oblivious transfer* (OT) protocol, a sender holds a pair of values (x_0, x_1) , and a receiver has a bit b . At the end of the protocol, the receiver learns x_b and learns nothing about x_{1-b} , while the sender learns nothing about b . A *two-message* OT protocol is one that consists of only two messages — the first from the receiver and the second from the sender. It is easy to see that any two-message OT implies a public-key cryptosystem (with the first message

¹² The exact group for DDH and parameters for LWE are inherited from the assumptions [7, 3]; one important note is that we need to assume LWE for a prime modulus that is polynomial in the security parameter.

¹³ In [7]’s case, the above is true for what they call their “expanded” scheme.

¹⁴ Indeed this is the function $h(sk) = \langle \bar{r}, sk \rangle + x'$, where $\bar{r}_i = r$, $\bar{r}_j = 0$ for $j \neq i$, and $x' = x - b \cdot r$.

being the public key); in addition, almost all popular candidates for public-key cryptosystems imply two-message OT protocols.

A targeted encryption scheme can be thought of as a type of “self-referential” OT where the receiver’s input selection bits are equal to the secret information it keeps after its first message (i.e., the secret key). It does not satisfy, however, the standard notion of OT, since the sender is not guaranteed to learn *nothing* about this secret key (although the “security against outsiders” property does imply that the sender cannot recover it completely). We note that it is possible (though we do not need to use this fact in this paper) to transform an OT with such a guarantee into a full-fledged OT, using the techniques of [16, 10].

4 Our Bounded KDM Secure Construction

Let k be the security parameter. Let $\text{TES} = (\text{TGen}, \text{TEnc}, \text{TDec})$ be a targeted encryption scheme. We will construct the following PKE scheme $\text{bKDM} = (\text{Gen}, \text{Enc}, \text{Dec})$ that is parameterized over polynomials N and L .

Key generation. $\text{Gen}(1^k)$ samples and outputs $(pk, sk) \leftarrow \text{TGen}(1^k)$.

Encryption. $\text{Enc}_{pk}(M)$ chooses $2k$ random strings $\overline{K} = (K_{i,b})_{(i,b) \in [k] \times \{0,1\}}$ and computes the garbled circuit transformation on \overline{K} and the constant function const^M that outputs M on *every* input $x \in \{0,1\}^k$. We use S , which is some fixed polynomial $S(N, L)$ to be specified later, as the size parameter for the garbled circuit transformation. Let GC be the resulting output. Enc also computes for every $(i, b) \in [k] \times \{0,1\}$ the value $\tilde{K}_{i,b} = \text{TEnc}_{pk,i,b}(K_{i,b})$ and outputs $C := (GC, (\tilde{K}_{i,b})_{(i,b) \in [k] \times \{0,1\}})$ as the ciphertext.

Decryption. $\text{Dec}_{sk}(GC, (\tilde{K}_{i,b})_{(i,b) \in [k] \times \{0,1\}})$ parses $sk = (sk_1, \dots, sk_k) \in \{0,1\}^k$ and computes the value $K_i = \text{TDec}_{sk}(\tilde{K}_{i,sk_i})$ for every $i \in [k]$. Then, it outputs the result of evaluating the garbled circuit GC on K_1, \dots, K_k .

It is easily verified that the decryption will indeed output $\text{const}^M(sk) = M$. We would like to emphasize that key generation does *not* depend on L or N , only encryption does. Hence, we can generate and distribute keys even without knowing L and N in advance.

4.1 Single-Key Security of the Construction

We now show that bKDM is KDM secure for the case of a single key (i.e., $N = 1$). In Section 5, we show that if the underlying targeted encryption scheme is circular secure (when suitably interpreted as a PKE scheme), bKDM actually is secure for an arbitrary number of keys.

Theorem 7. *If $\text{TES} = (\text{TGen}, \text{TEnc}, \text{TDec})$ is a targeted encryption scheme, then for every polynomial L , bKDM instantiated with $S(N, L) = L$ is $(1, L)$ -bounded KDM secure.*

Proof. Fix $N = 1$, an arbitrary L and a PPT adversary A on bKDM ’s bounded KDM security. In order to keep the notations simple, we concentrate on the

single query case (i.e., the attacker only asks a single key related query). The multi query case, however, easily follows from the same lines. We proceed in games. Let X_i be A 's output in Game i , and write $X_i \approx X_j$ as a shorthand for $\Pr[X_i = 1] - \Pr[X_j = 1] \in \text{negl}$. See Table 1 for an overview of all games used in the proof. In all the following games, (sk, pk) are chosen at random using TGen and the oracles get $h : \{0, 1\}^k \rightarrow \{0, 1\}^k$ as input.

Game 0 is the real KDM game. Namely, the oracle $\text{Real}_{pk, sk}$ returns the ciphertext $\text{Enc}_{pk}(h(sk))$. Recall that this is computed by (1) choosing a random $2k$ key tuple \overline{K} , (2) encrypting the keys using TEnc to obtain a tuple of ciphertexts \tilde{K} where $\tilde{K}_{i,b} = \text{TEnc}_{pk}(K_{i,b})$ for every $(i, b) \in [k] \times \{0, 1\}$, and (3) computing $GC = \text{Garble}(\overline{K}, \text{const}^{h(sk)})$. Ciphertext is $C := (GC, \tilde{K})$.

In **Game 1** the oracle sets $\tilde{K}_{i,b}$ to $\text{TEnc}_{pk}(0^k)$, instead of $\text{TEnc}_{pk}(K_{i,b})$, for every (i, b) with $sk_i \neq b$. (Note that we still use the original \overline{K} in the garbled circuit construction.) Since GC is independent from the random coins used to encrypt \tilde{K} , the “security against receiver” property of TES yields that $X_0 \approx X_1$.

In **Game 2** the oracle uses h instead of $\text{const}^{h(sk)}$ in the garbled circuit construction (i.e., it computes $GC = \text{Garble}(\overline{K}, h)$). Since $h(sk) = \text{const}^{h(sk)}(sk)$ and only the keys $\overline{K}_{sk} = (\overline{K}_{i, sk_i})_i$ are used outside the garbled circuit construction, the security against receiver of the garbled circuit construction yields that $X_1 \approx X_2$. We note that the only role of the secret key in this game, is for deciding which elements of \tilde{K} are replaced by encryptions of 0^k .

In **Game 3** we go back to using the original \tilde{K} (also for the (i, b) with $b \neq sk_i$). Again, the “security against receiver” property of TES implies that $X_2 \approx X_3$. Note that in this game the encryption oracle does not use the secret key *at all*.

We define **Game 4** to be the variant of **Game 3** in which we set $\tilde{K}_{i,b} = \text{TEnc}_{pk}(0^k)$ for every i, b (i.e., we ignore the value of \overline{K} for this part). Since the secret key is never used in either Game 3 or Game 4, the “security against

Table 1. Overview of the games used in the proof of Theorem 7. We use boxes to highlight the component that changed from the previous hybrid, and note in the remark the justification for the fact that the hybrid is indistinguishable from the previous one.

Game	Oracle needs	\tilde{K}_{i, sk_i}	$\tilde{K}_{i, 1-sk_i}$	Function in GC	Remark
0	pk, sk	$\text{TEnc}(K_{i, sk_i})$	$\text{TEnc}(K_{i, 1-sk_i})$	$\text{const}^{h(sk)}$	Real KDM game
1	pk, sk	$\text{TEnc}(K_{i, sk_i})$	$\text{TEnc}(0^k)$	$\text{const}^{h(sk)}$	TES's sec. ag. recv.
2	pk, sk	$\text{TEnc}(K_{i, sk_i})$	$\text{TEnc}(0^k)$	h	GC 's sec. ag. recv.
3	pk	$\text{TEnc}(K_{i, sk_i})$	$\text{TEnc}(K_{i, 1-sk_i})$	h	TES's sec. ag. recv.
4	pk	$\text{TEnc}(0^k)$	$\text{TEnc}(0^k)$	h	TES's sec. ag. outs.
5	pk	$\text{TEnc}(0^k)$	$\text{TEnc}(0^k)$	const^{0^k}	GC 's sec. ag. outs.
6	pk	$\text{TEnc}(K_{i, sk_i})$	$\text{TEnc}(K_{i, 1-sk_i})$	const^{0^k}	TES's sec. ag. outs. fake KDM game

outsiders” property of the TES implies that $X_3 \approx X_4$. Note that in this game, the vector \overline{K} is independent of \tilde{K} .

In **Game 5** we change h to $const^{0^k}$ in the garbled circuit construction. Since no information on the key vector \overline{K} , except for the garbled circuit itself, is given in both oracles, the “security against outsiders” property of the garbled circuit construction implies that $X_4 \approx X_5$. (Note that we need to use the “security against outsiders” and not the “security against receiver” property, since obviously we have no guarantee that $h(sk) = const^{0^k}(sk)$.)

We define **Game 6** to be the game in which we go back to using the real \tilde{K} . Since the oracles do not use the secret key, we get that $X_5 \approx X_6$. Observe that the encryption oracle is exactly the Fake oracle as per Definition 1, and hence we have completed the proof. ■

5 Multiple Key Security

While the notion of KDM security for a single key is challenging and elegant, many of the applications actually require KDM security in the presence of arbitrarily (polynomially) many keys. Hence, let now the number of keys N be an arbitrary polynomial in the security parameter. We will prove that our scheme bKDM from Section 4 is (N, L) -bounded KDM secure, but now under an additional assumption, and with different parameters.

Complication and central idea. Recall the proof of Theorem 7. There, we have first substituted the function $const^{h(sk)}$ that is evaluated by GC by the KDM query function h itself. By the secrecy against receiver property of the garbled circuit, we could argue that this change goes unnoticed by the receiver. This modification was a crucial step in our proof, since it allowed to construct the garbled circuit without knowing sk . Recall that in multiple secret keys case, the security is defined with respect to N public and secret keys pairs $((\overline{pk}_1, \overline{sk}_1), \dots, (\overline{pk}_N, \overline{sk}_N))$, and the adversary gets encryptions of a query function $h = h(\overline{sk})$ under arbitrary \overline{pk}_μ for $\mu \in [N]$. Hence, we cannot simply substitute $const^{h(\overline{sk})}$ with h directly (the secrecy against receiver property of the garbled circuit would not help in this case, since we cannot claim that $h(\overline{sk}_\mu) = const^{h(\overline{sk})}$).¹⁵ Instead, we will substitute $const^{h(\overline{sk})}$ with a function h' for which $h'(\overline{sk}_\mu) = h(\overline{sk})$. This function h' contains an encryption of \overline{sk} under the receiver’s public key \overline{pk}_μ . In this, we will have to interpret bKDM’s underlying targeted encryption scheme TES as a *circular-secure* encryption scheme. (Circular security is required to guarantee that we can later replace these encryptions of secret keys with 0^k -encryptions.) Since our targeted encryption scheme instance is based on the clique-secure encryption schemes of [7, 3], it already has this property. The remaining part of the proof follows the proof of Theorem 7.

Definition 4 (Augmented targeted encryption). An augmented targeted encryption scheme $ATES = (\text{TGen}, \text{TEnc}, \text{TDec}, \text{Enc}, \text{Dec})$ is a targeted

¹⁵ $h(\overline{sk}_\mu)$ is not even well defined; h is expecting a vector of (secret) keys as input, and not a single key.

encryption scheme $(\text{TGen}, \text{TEnc}, \text{TDec})$, complemented by PPT algorithms Enc, Dec for (un-targeted) encryption and decryption. We require that $(\text{TGen}, \text{Enc}, \text{Dec})$ is a public-key encryption scheme with message space $\mathcal{M} \subseteq \{0, 1\}^k$. In particular, $\text{Dec}_{sk}(\text{Enc}_{pk}(M)) = M$ for all $(pk, sk) \leftarrow \text{TGen}(1^k)$ and $M \in \{0, 1\}^k$.

We say that ATES is circular secure if $(\text{TGen}, \text{Enc}, \text{Dec})$ is. We stress that our both TES instances from Theorem 6 are circular secure augmented targeted encryption schemes with the natural encryption and decryption algorithms from Boneh et al. [7] and Applebaum et al. [3] respectively. The following theorem implies our main result (i.e., Theorem 1). We provide a proof in the full version [5] of this paper.

Theorem 8. *If $\text{ATES} = (\text{TGen}, \text{TEnc}, \text{TDec}, \text{Enc}, \text{Dec})$ is a circular secure augmented targeted encryption scheme, then for every polynomials L and N , bKDM instantiated with a suitable polynomial $S(N, L)$ is bounded KDM secure.*

6 Application to Formal Cryptography

One of the main motivations to study KDM security lies in the connection between formal and computational cryptography. In formal cryptography (starting with [12, 13, 23]), cryptographic operations like encryption or digital signatures are abstracted as symbolic operators that (only) adhere to natural rules like $D_K(E_K(M)) = M$ for symmetric encryption and decryption operators E and D . A simple calculus like this enables machine-assisted security analysis (e.g., [21, 22]). It is not a priori clear, however, that security properties proved in the symbolic calculus also hold for the computational implementation of the protocol.

Computational soundness. Abadi and Rogaway [1] were the first to relate the formal and computational views on cryptography. Specifically, they showed that every symbolically proven property also holds in the computational world, assuming a suitable computational implementation. This is usually referred to as a soundness result, and suitable computational implementations are dubbed sound. In order to provide computational soundness in this sense in face of a passive adversary, an encryption scheme essentially needs to be IND-CPA secure.

Key-cyclic expressions. There is a technical nuisance, however, that limits the generality and expressivity of [1]’s approach. Namely, the soundness result only holds for protocols that do *not* contain key-cyclic expressions. That is, only protocols in which no expressions with cyclic dependencies of encryption keys (such as $E_{K_1}(E_{K_2}(K_1))$) appear are considered. This is for the following reason: in the symbolic setting, the natural deduction rules explicitly require secret keys for decryption. Hence, the encrypted plaintexts in such expressions are secret by definition in the symbolic world (i.e., there is no formal way to apply, say, D_{K_2} on the ciphertext $E_{K_1}(E_{K_2}(K_1))$). On the other hand, key-dependent messages like the one above, are not modeled in standard (computational) security

experiments for encryption schemes.¹⁶ Hence, there is an asymmetry between symbolic and computational setting, and any soundness result that connects symbolic encryption and *standard* computational encryption notions has to exclude key-cyclic expressions.

Soundness from Bounded KDM Security. It was informally claimed by Black et al. [6], and formally proven by Adão et al. [2], that *fully* KDM-secure encryption schemes imply computational soundness for *arbitrary* symbolic protocols. Since we can only achieve bounded KDM security against arbitrary circuits up to a certain size, we ask whether bounded KDM security suffices for computational soundness of arbitrary symbolic protocols. The answer we give is essentially affirmative.

To do so, we introduce the following slight strengthening of bounded KDM security:

Definition 5 (Length-dependent bounded KDM security). *A PKE scheme with message space $\mathcal{M} \subseteq \{0, 1\}^*$ is N -key length-dependent bounded KDM secure, if it is KDM secure with respect to the circuit class of all $h : (\{0, 1\}^k)^N \rightarrow \{0, 1\}^{\lfloor \sqrt{|h|} \rfloor}$, where $|h|$ is the circuit size of h .*

That is, length-dependent KDM secure schemes are secure against larger KDM queries if longer messages are encrypted. We stress that our scheme **bKDM** from Section 4 is N -key length-dependent bounded KDM secure, if we choose L suitably (e.g., $L = |M|^{2.1}$) during encryption. Namely, **bKDM**'s key generation algorithm does not depend on N or L , and the proofs of Theorems 7 and 8 do not use that L is fixed.

Theorem 9 (Following [1, 2]: Bounded KDM security implies soundness). *Let **bKDM** be an N -key length-dependent KDM secure PKE scheme, and let P be a symbolic protocol with N parties in the setting of Adão et al. [2]. Then **bKDM** provides a computationally sound implementation of P .*

We stress that the choice of symbolic setting [2] was made only for simplicity. We provide a proof outline in the full version [5] of this paper.

Application of our results. Theorem 9 can be instantiated with our scheme **bKDM** from Section 4. (As argued above, **bKDM** actually is N -key length-dependent bounded KDM secure.) This yields the first encryption scheme that provides soundness under a standard computational assumption.

Relation to circular security and extensibility. For extremely simple calculi that only feature public-key encryption, along with a few syntactic operations like pairings of terms, already clique security may enable soundness. (This is so since all key-dependent encryptions that can possibly occur in a symbolic

¹⁶ Some subsequent soundness results (e.g., [4, 24]) consider an active adversary and require IND-CCA security. We stress that does not change the technical complications regarding key-cyclic expressions.

protocol can be traced back to simple encryptions of the form $E_{K_i}(K_j)$, assuming a suitable way to encrypt longer terms in chunks.) Nevertheless, we stress that our scheme bKDM allows much richer classes of calculi. For instance, the above soundness proof also works in the presence of signatures, so that terms of the form $E_{K_i}(\text{sig}(K_j), K_\ell)$ may occur. (The crucial observation is that we can suitably pad, e.g., signatures such that the signing algorithm can be expressed as a length-dependent KDM circuit.) On the other hand, clique security, or even security against a polynomial number of arbitrary but predetermined KDM functions is *not* sufficient to treat such richer classes of calculi.

7 Extending Haitner and Holenstein's Impossibility Result

In this section we observe that a result of Haitner and Holenstein [17], showing that there is no KDM-secure scheme with a proof of security which makes a black-box use of both the adversary and the KDM function, can be extended to rule out not just *full* KDM security but also *bounded* KDM security. The idea is simple: while this result used a *random* function h for the KDM function, a *pseudorandom* function could work just as well.

The following definition is adopted from [17].

Definition 6 (Cryptographic games). *A cryptographic game is a (possibly inefficient) random system Γ , where on security parameter k , Γ interacts with an attacker A and may output 1. We define the game value of such an interaction, denoted $\Gamma_A(1^k)$, as the probability that Γ outputs 1 in the end of the interaction with A , where the probability is taken over the random coins of Γ and A . A cryptographic game is non-interactive if it consists of two messages, from Γ to A and back.*

Examples:

OWF. The security of a one-way function f is equivalent to requiring that the value of the following game is negligible for any efficient A . On security parameter k , the system Γ selects a random $x \in \{0, 1\}^k$ and sends $y = f(x)$ to the adversary. Γ outputs 1 if A outputs $x' \in f^{-1}(y)$.

DDH. The security of the DDH hardness assumption is equivalent to requiring that the value of the following game is at most negligibly bounded from $\frac{1}{2}$ for any efficient adversary. Let G be an appropriate DDH group (e.g., Z_p^* for some prime p) and let g be a generator in the group. The system Γ chooses a random bit b , sends the tuple (g^x, g^y, g^z) to the challenger A , where x and y are random exponents, and $z = x \cdot y$ if $b = 0$ and a random value otherwise. Γ outputs 1 iff A has guessed b correctly.

Definition 7 (Strongly-black-box reductions). *An encryption scheme (Enc, Dec) has a δ -strongly-black-box reduction from its KDM security to a cryptographic game Γ with respect to a family of query function \mathcal{Q} , if there exists an*

oracle-aided algorithm R with the following guarantee: Let A be an efficient adversary that breaks the KDM security of the scheme using query functions from \mathcal{Q} with advantage $\epsilon_A = \epsilon_A(k)$ (i.e., On security parameter k , A distinguishes between encryptions of functions of the secret key and encryptions of garbage with advantage $\epsilon_A(k)$). Then the value of $\Gamma_{RA}(1^k) \geq \delta(k, \epsilon_A(k))$. Here, we require that R treat the query functions it gets from A as black boxes — all it can do is to query them on arbitrary chosen inputs.

Informally, we say that a proof for the KDM security of a scheme is strongly-black-box with respect to a game Γ and a family of query function \mathcal{Q} , if the value of $\delta(k, \epsilon_A(k))$ for every non-negligible $\epsilon_A(k)$ is considered a “break” of Γ (i.e., $\delta(k, \epsilon_A(k)) > \frac{1}{2} + \text{negl}$ for the DDH game). We remark that *all* known KDM constructions in the literature have strongly-black-box reductions with respect to the relevant hardness assumption (e.g., DDH) and the class of query functions they are secure against.

Definition 8 (Pseudorandom functions (PRF)). An ensemble of functions $\mathcal{F} = \{\mathcal{F}_k = \{f: \{0, 1\}^{m(k)} \mapsto \{0, 1\}^{\ell(k)}\}\}$ is pseudorandom, if, on security parameter k , an efficient adversary cannot distinguish with more than negligible advantage between a random $f \in \mathcal{F}_k$, and a truly random function defined on the same input and output domains. Here, the adversary may only access the function as a black box. The ensemble is α -exponential hard for a constant $\alpha > 0$, if no adversary that runs in time 2^{n^α} wins in the above game with advantage greater than 2^{-k^α} .

Theorem 10. Let (Enc, Dec) be a δ -strongly-black-box reduction from its KDM security to a non-interactive cryptographic game Γ with respect to a family of query functions $\mathcal{Q} = \{\mathcal{Q}_k\}$.¹⁷ Assume that \mathcal{Q}_k contains the family of functions $\mathcal{G}_k = \{g_k(x) = f(x, 0^{t(k)-k}) \oplus r: f \in \mathcal{F}_{t(k)}, r \in \{0, 1\}^{\ell(k)}\}$, where $t(k) \geq 2k$ and $\mathcal{F}_{t(k)} = \{f: \{0, 1\}^{t(k)} \mapsto \{0, 1\}^{\ell(k)}\}$ is an α -exponential hard PRF with $t(k)^\alpha \geq 2k$. Then, there exists an efficient algorithm A with $\Gamma_A \geq \delta(k, 1 - 2^{-k}) - 2^{-k}$.

In particular, giving such a strongly-black-box reduction implies that either the class of query function considered is weak (does not contain exponentially hard PRF), or the game Γ can be efficiently broken with probability $\delta(k, 1 - 2^{-k}) - 2^{-k}$.

Proof (sketch). The proof is similar to the proof of [17, Theorem 5]. Consider the following (inefficient) adversary A for breaking the KDM security of (Enc, Dec) with respect to \mathcal{Q} . On security parameter k , choose a random $g \in \mathcal{G}_k$ and make a KDM query to obtain a ciphertext C . Then check (via exhaustive search) if there exists $sk \in \{0, 1\}^k$ such that $\text{Dec}_{sk}(C) = g(sk)$. If positive return 1, otherwise return 0. It is easy to verify that A breaks the KDM security with advantage $1 - 2^{-k}$ (the probability that a decryption of random ciphertext C equals to $g(sk)$)

¹⁷ Theorem 10 can be shown to hold also against all natural interactive games (see [17] for details). For the sake of simplicity, however, we choose to focus here on the non-interactive case.

for some sk , is bounded by $\sum_{sk} \Pr[\text{Dec}_{sk}(C) = g(sk)] = \sum_{sk} \Pr_r[\text{Dec}_{sk}(C) = f(sk) \oplus r] \leq 2^k \cdot 2^{-2k} = 2^{-k}$. More interestingly, we notice that the probability that R^A sends a ciphertext $C = \text{Enc}_{sk}(g(sk))$ to A without previously making the query $g(sk)$ is bounded by 2^{-2k} . Assume otherwise, then R^A is an algorithm that runs in time $\text{poly} \cdot 2^k$ and breaks the security of \mathcal{F} . It follows that we can emulate the execution of R^A : throughout the execution, keep track of all queries that R makes to g , and let T denote the list of queries. When R queries A on a ciphertext C , act as the inefficient A above, but *only* with respect to the secret keys in T . The above observation yields that we emulate R^A with error bounded by 2^{-2k} . ■

References

- [1] Abadi, M., Rogaway, P.: Reconciling two views of cryptography (the computational soundness of formal encryption). *Journal of Cryptology* 15(2), 103–127 (2002)
- [2] Adão, P., Bana, G., Herzog, J., Scedrov, A.: Soundness of formal encryption in the presence of key-cycles. In: di Vimercati, S.d.C., Syverson, P.F., Gollmann, D. (eds.) *ESORICS 2005*. LNCS, vol. 3679, pp. 374–396. Springer, Heidelberg (2005)
- [3] Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In: Halevi, S. (ed.) *CRYPTO 2009*. LNCS, vol. 5677, pp. 595–618. Springer, Heidelberg (2009)
- [4] Backes, M., Pfitzmann, B., Waidner, M.: A composable cryptographic library with nested operations (extended abstract). In: *ACM CCS 2003*, pp. 220–230 (2003); Full version in *IACR Cryptology ePrint Archive 2003/015* (January 2003), <http://eprint.iacr.org/>
- [5] Barak, B., Haitner, I., Hofheinz, D., Ishai, Y.: Bounded key-dependent message security. *Cryptology ePrint Archive, Report 2009/511* (2009), <http://eprint.iacr.org/>
- [6] Black, J., Rogaway, P., Shrimpton, T.: Encryption-scheme security in the presence of key-dependent messages. In: Nyberg, K., Heys, H.M. (eds.) *SAC 2002*. LNCS, vol. 2595, pp. 62–75. Springer, Heidelberg (2003)
- [7] Boneh, D., Halevi, S., Hamburg, M., Ostrovsky, R.: Circular-secure encryption from decision diffie-hellman. In: Wagner, D. (ed.) *CRYPTO 2008*. LNCS, vol. 5157, pp. 108–125. Springer, Heidelberg (2008)
- [8] Brakerski, Z., Goldwasser, S., Kalai, Y.: Circular-secure encryption beyond affine functions. *Cryptology ePrint Archive, Report 2009/485* (2009), <http://eprint.iacr.org/>
- [9] Camenisch, J., Lysyanskaya, A.: An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In: Pfitzmann, B. (ed.) *EUROCRYPT 2001*. LNCS, vol. 2045, p. 93. Springer, Heidelberg (2001)
- [10] Crépeau, C., Kilian, J.: Weakening security assumptions and oblivious transfer. In: Goldwasser, S. (ed.) *CRYPTO 1988*. LNCS, vol. 403, pp. 2–7. Springer, Heidelberg (1990)
- [11] Damgård, I., Jurik, M.J.: A generalisation, a simplification and some applications of paillier’s probabilistic public-key system. In: Kim, K.-c. (ed.) *PKC 2001*. LNCS, vol. 1992. Springer, Heidelberg (2001)
- [12] Dolev, D., Yao, A.C.: On the security of public key protocols. *IEEE Transactions on Information Theory* 29(2), 198–208 (1983)

- [13] Even, S., Goldreich, O.: On the security of multi-party ping-pong protocols. In: FOCS 1983 (1983)
- [14] Even, S., Goldreich, O., Lempel, A.: A randomized protocol for signing contracts. *Commun. ACM* (1985)
- [15] Gentry, C.: Fully homomorphic encryption using ideal lattices. In: STOC (2009)
- [16] Goldreich, O., Levin, L.A.: A hard-core predicate for all one-way functions. In: STOC (1989)
- [17] Haitner, I., Holenstein, T.: On the (im)possibility of key dependent encryption. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 202–219. Springer, Heidelberg (2009)
- [18] Halevi, S., Krawczyk, H.: Security under key-dependent inputs. In: ACM CCS 2007 (2007)
- [19] Hofheinz, D., Unruh, D.: Towards key-dependent message security in the standard model. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 108–126. Springer, Heidelberg (2008)
- [20] Ishai, Y., Paskin, A.: Evaluating branching programs on encrypted data. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 575–594. Springer, Heidelberg (2007)
- [21] Kemmerer, R., Meadows, C., Millen, J.: Three systems for cryptographic protocol analysis. *Journal of Cryptology* 7(2), 79–130 (1994)
- [22] Lowe, G.: Breaking and fixing the Needham-Schroeder public-key protocol using FDR. In: Margaria, T., Steffen, B. (eds.) TACAS 1996. LNCS, vol. 1055. Springer, Heidelberg (1996)
- [23] Merritt, M.: Cryptographic Protocols. PhD thesis, Georgia Institute of Technology (1983)
- [24] Micciancio, D., Warinschi, B.: Soundness of formal encryption in the presence of active adversaries. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 133–151. Springer, Heidelberg (2004)
- [25] Naor, M.: On cryptographic assumptions and challenges. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 96–109. Springer, Heidelberg (2003)
- [26] Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, p. 223. Springer, Heidelberg (1999)
- [27] Rabin, M.: How to exchange secrets by oblivious transfer. Technical Report TR-81, Harvard Aiken Computation Laboratory (1981)
- [28] Yao, A.C.: How to generate and exchange secrets. In: FOCS 1986 (1986)