Journal of
CRYPTOLOGY

CrossMark

Garbling XOR Gates "For Free" in the Standard Model[*]

Benny Applebaum[†]

School of Electrical Engineering, Tel-Aviv University, Tel Aviv, Israel
benny.applebaum@gmail.com

Abstract. Yao's garbled circuit (GC) technique is a powerful cryptographic tool which allows to "encrypt" a circuit C by another circuit $\hat{C}$ in a way that hides all information except for the final output. Yao's original construction incurs a constant overhead in both computation and communication per gate of the circuit C (proportional to the complexity of symmetric encryption). Kolesnikov and Schneider (ICALP 2008) introduced an optimized variant that garbles XOR gates "for free" in a way that involves no cryptographic operations and no communication. This variant has become very popular and has lead to notable performance improvements. The security of the free-XOR optimization was originally proved in the random oracle model. Despite some partial progress (Choi et al., TCC 2012), the question of replacing the random oracle with a standard cryptographic assumption has remained open. We resolve this question by showing that the free-XOR approach can be realized in the standard model under the *learning parity with noise* (LPN) assumption. Our result is obtained in two steps:

1. We show that the random oracle can be replaced with a symmetric encryption, which remains secure under a combined form of related-key (RK) and key-dependent message (KDM) attacks.
2. We show that such a symmetric encryption can be constructed based on the LPN assumption.

As an additional contribution, we prove that the combination of RK and KDM security is nontrivial in the following sense: There exists an encryption scheme which achieves RK security and KDM security separately, but breaks completely at the presence of combined RK-KDM attacks.

1. Introduction

Yao's *garbled circuit* (GC) construction is an efficient transformation which maps any boolean circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$ together with secret randomness into a "garbled circuit" $\hat{C}$ along with n pairs of short k-bit keys (W_i^0, W_i^1) such that, for any (unknown)

---

input x, the garbled circuit $\hat{C}$ together with the n keys $W_x = (W_1^{x_1}, \ldots, W_n^{x_n})$ reveals $C(x)$ but gives no additional information about x. Yao's celebrated result shows that such a transformation can be based on the existence of any pseudorandom generator [13,44], or equivalently a one-way function [22].

Originally motivated by the problem of secure multiparty computation [21,44], the GC construction has found a diverse range of other applications to problems such as computing on encrypted data, parallel cryptography, verifiable computation, software protection, functional encryption, and key-dependent message security (see [5] for references). Despite its theoretical importance, GC was typically considered to be impractical due to a large computational and communication overhead, which is proportional to the circuit size. This belief was recently challenged by a fruitful line of works that optimizes the concrete efficiency of GC-based protocols up to a level that suits large-scale practical applications [23–25,30–32,35,38–40,42].

Among other improvements, most current implementations of GCs (e.g., [23,24,34, 40,42]) employ the so-called *free-XOR* optimization of Kolesnikov and Schneider [29]. While in Yao's original construction, every gate of the circuit C has a computational cost of few cryptographic operations (e.g., three or four applications of a symmetric primitive) and a communication cost of few ciphertexts, Kolesnikov and Schneider showed how to completely eliminate the communication and computational overhead of XOR gates. This optimization significantly improves the practical performance, especially for large- or medium-size circuits as demonstrated in [28,29,40].

As in many cases, this gain in efficiency requires stronger cryptographic assumptions. Unlike Yao's GC, which can be based on the existence of standard symmetric-key cryptography, the free-XOR optimization relies on a hash function H, which is modeled as a random oracle [9]. Due to the known limitations of the random oracle model [17], it is natural to ask:

> Is it possible to realize the free-XOR optimization in the standard model?

This question was raised in the original work of Kolesnikov and Schneider [29] and was further studied in [3,18]. In [29], it was conjectured that the full power of the random oracle is not really needed and that the function H can be instantiated with a *correlation-robust hash function* [26], a strong (yet seemingly realizable) version of a hash function which remains pseudorandom even when it is applied to linearly related inputs. Choi et al. [18] showed that the picture is actually more complex: Correlation robustness alone does not suffice for security (as demonstrated by an explicit counter-example in the random oracle model). Instead, one has to employ a stronger form of hash function, which, in addition to being correlation-robust, also satisfies some form of circular security [10,16]. While the existence of circular correlation-robust hash functions (a new primitive introduced by Choi et al. [18]) seems to be a reasonable assumption (significantly weaker than the existence of a random oracle), it is still unknown how to realize it based on a standard cryptographic assumption. This leaves open the problem of implementing the free-XOR optimization in the standard model.

1.1. *Our Contribution*

We resolve the above feasibility question by showing that the free-XOR optimization can be realized in the standard model under the learning parity with noise (LPN) as-

sumption [11,20]. This assumption, which can also be formulated as the intractability of decoding a random linear code, is widely studied by the coding and learning communities and was extensively employed in cryptographic constructions during the last two decades.

Specifically, we make the following contributions:

1. We introduce a new combined form of related-key (RK) and key-dependent message (KDM) attacks. Roughly speaking, in such an attack the adversary is allowed to see ciphertexts of the form $\mathsf{Enc}_{\phi(K)}(\psi(K))$ where K is the secret key and the functions ϕ and ψ are chosen by the adversary from some predefined function families. This notion of security, referred to as *RK-KDM security*, generalizes the previous definitions of semantic security under related-key attacks [3] and key-dependent message attacks [10,16]. In fact, as shown in Sect. 5, this is a *strict* generalization as there exists an encryption scheme which satisfies both RK security and KDM security separately, but fails to achieve the combined form of RK-KDM security.
2. We prove that the free-XOR construction is secure when instantiated with a semantically secure symmetric encryption scheme whose security is preserved under binary linear RK-KDM attacks. (Essentially, $\phi(K) = K \oplus \Delta_1$ and $\psi(K) = K \oplus \Delta_2$ for any fixed shift vectors Δ_1 and Δ_2.)
3. We show that the LPN-based symmetric encryption of [19] and its generalization [2] satisfies RK-KDM security with respect to binary linear functions. In fact, our proof provides a general template for proving RK-KDM security based on pseudorandomness and joint key/message homomorphism. This is similar to previous results along these lines [2,3,6,14].

Altogether, our proofs turn to be quite simple (which we consider as a virtue), short, and modular. This is due to the following choices:

Encryption Versus Hashing   The key point in which we deviate from [18,29] is the use of (randomized) *symmetric encryption*, as opposed to deterministic *hash function* (or some other pseudorandom primitive). Indeed, the GC construction essentially employs the hash function only as a "computational one-time pad", namely as a mean to achieve secrecy. Therefore, in terms of functionality, it seems best (i.e., more general) to abstract the underlying primitive as an encryption scheme. While this is true in general for the standard GC (cf. [4,32] and the recent discussion in [7]), this distinction becomes even more important in the context of the free-XOR variant. In this case, the underlying primitive should satisfy stronger notions of security (RKA and KDM), and this turns to be much easier for randomized encryption than for pseudorandom objects such as hash functions. (See also [3].) As a secondary gain, the new security definition that arises for symmetric encryption (RK-KDM semantic security) is natural and compatible with existing well-studied notions. In contrast, the analog definition of RK-KDM security for hash functions (*circular correlation robustness*) appears less natural as there is no obvious interpretation for the concepts of *message* and *key*.

GC as Randomized Encoding   It is important to distinguish between the garbled circuit transformation (i.e., the mapping from C to $\hat{C}$) and the secure function evaluation protocol, which is based on it. The distinction between the two, which is sometimes

blurred, can be formulated via the notion of *randomized encoding of functions* [27] as done in [4]. Our proofs follow this abstraction and show that the free-XOR technique yields computationally private randomized encoding. At this point, one can invoke, for example, the general theorem of [4] to derive a secure MPC protocol. Similarly, all other applications (cf. [1]) of randomized encoding can be obtained directly by invoking the reduction from RE to the desired task. This is the first modular treatment of the free-XOR variant.

1.2. *Discussion*

The main goal of this work is to provide a solid theoretical justification for the free-XOR heuristic. This is part of an ongoing effort of the theory community to explain the security of "real-world" protocols. Several such examples arise when trying to import random oracle-based protocols to the standard model. In this context, [17] suggested a two-step methodology: (1) "identify useful special-purpose properties of the random oracle" and (2) show that these properties "can be also provided by a fully specified function (or function ensemble)." In the context of the free-XOR technique, the first step was essentially taken by [18] who identified the extra need of "circular security," while the current paper completes the second step, which involves, in addition, some fine-tuning of step 1.

It should be emphasized that we do not suggest to replace the hash function with an LPN-based scheme in practical implementations (though we do not rule out such a possibility either). Still, we believe that the results of this work are useful even if one decides, due to efficiency considerations, to use a heuristic implementation. Specifically, viewing the primitive as an RK-KDM secure encryption scheme allows to rely on other heuristic solutions such as block ciphers, for which RKA and KDM security are well studied.

Other Related Works   The notions of key-dependent message security (aka circular security) and related-key attacks were introduced by [10,16] and [8]. Both notions were extensively studied (separately) during the last decade. Most relevant to this paper is our joint work with Harnik and Ishai [3]. This work introduces the notion of semantic security under related-key attacks, describes several constructions, and shows that protocols employing correlation-robust hash functions and their relatives (e.g., [26,37]) can be securely instantiated with RKA secure encryption schemes. In addition, [3] suggested to apply a similar modification to the free-XOR variant, which was believed to be secure when instantiated with correlation-robust hash functions [29]. As mentioned, the latter claim was found to be inaccurate, and therefore, the results of [3] cannot be used in the context of the free-XOR technique. (The other applications mentioned in [3] remain valid.)

Subsequent Work   Following our work, Böhl, Davies, and Hofheinz [15] constructed several RK-KDM public-key encryption schemes based on various intractability assumptions such as the decisional Diffie-Hellman (DDH) assumption, the learning with errors (LWE) assumption, quadratic residuosity and decisional Diffie-Hellman (QR+DDH) assumption, and the decisional composite residuosity (DCR) assumption. The proofs of

security follow the general template suggested here (as abstracted in Remark 3.7). Furthermore, some of the resulting schemes (the one based on DDH, LWE, and QR+DDH) support binary linear relations and can be therefore used for the free-XOR optimization. This further demonstrates the wide applicability of our approach.

Organization Following some preliminaries (Sect. 2), in Sect. 3 we define semantic security under RK-KDM attacks and describe an LPN-based implementation. Section 4 is devoted to the garbled circuit construction, including definitions (in terms of randomized encoding), a description of Yao's original construction and the free-XOR variant, and a proof of security that reduces the privacy of the free-XOR GC to the RK-KDM security of the underlying encryption. In Sect. 5, we describe an encryption scheme which is KDM secure and RKA secure but not RK-KDM secure, separating the latter notion from the formers. Finally, we end with a short conclusion in Sect. 6.

2. Preliminaries

We let $\circ$ denote string concatenation. Strings are often treated as vectors or matrices over the binary field $\mathbb{F}_2$, accordingly string *addition* is interpreted simply as bit-wise exclusive-or. When adding together two matrices $A_{n \times k}$ and $B_{N \times k}$ where $n < N$, we assume that the last $N - n$ missing rows of A are padded with zeroes. The same convention holds with respect to vectors (i.e., when $k = 1$).

2.1. *Randomized Functions*

We extensively use the abstraction of randomized functions, which can be seen as a special case of Maurer's Random Systems [36]. A *randomized function* is a two argument function $f : X \times R \to Y$ whose first input x is referred to as the *deterministic input* and the second input is referred to as the *random input*. For every deterministic input x, we think of $f(x)$ as the random variable induced by sampling $r \xleftarrow{R} R$ and computing $f(x; r) \in Y$. When a (randomized) algorithm A gets an oracle access to a randomized function f, we assume that A has control only on the deterministic input; namely, if A queries f with x, it gets as a result a fresh sample from $f(x)$. Note that A^f itself defines a randomized function. We say that $\{f_s\}_{s \in \{0,1\}^*}$ is a *collection of randomized functions* if f_s is a randomized function for every key s. By default, all the collections are efficiently computable in the sense that $f_s(x)$ can be sampled in time $\text{poly}(|s| + |x|)$. We note that a sequence of randomized functions $\{f_n\}_{n \in \mathbb{N}}$ can be viewed as a (degenerate) collection of randomized functions $\{f'_s\}_{s \in \{0,1\}^*}$ where $f'_s = f_{|s|}$. Under this convention, efficiency means that $f_n(x)$ should be computable in time $\text{poly}(n, |x|)$. Since the input length of f_n will always be polynomial in n, this boils down to standard $\text{poly}(n)$-time efficiency.

Indistinguishability A pair of randomized functions f, g is *equivalent* $f \equiv g$ if for every input x the random variables $f(x)$ and $g(x)$ are identically distributed. A pair $f = \{f_s\}$ and $g = \{g_s\}$ of collections of randomized functions is *computationally indistinguishable*, denoted by $f \stackrel{c}{\equiv} g$, if for every efficient adversary $\mathcal{A}$ it holds that

$$\left| \Pr_{s \xleftarrow{R} \{0,1\}^k} [\mathcal{A}^{f_s}(1^k) = 1] - \Pr_{s \xleftarrow{R} \{0,1\}^k} [\mathcal{A}^{g_s}(1^k) = 1] \right| < \varepsilon(k),$$

for some negligible function ε. We note that the key of the function s is chosen at random and then fixed across invocations, while the internal randomness of the function is refreshed in each oracle call.

Let $\{f_s\}$, $\{g_s\}$ and $\{h_s\}$ be collections of randomized functions. We will need the following standard facts (cf. [36]).

Fact 2.1. *If* $\{f_s\} \overset{c}{\equiv} \{g_s\}$ *and A is an efficient function, then the collections of randomized functions* $\{A^{f_s}\}_s$ *and* $\{A^{g_s}\}_s$, *which are indexed by s, are computationally indistinguishable.*

Fact 2.2. *If* $\{f_s\} \overset{c}{\equiv} \{g_s\}$ *and* $\{g_s\} \overset{c}{\equiv} \{h_s\}$, *then* $\{f_s\} \overset{c}{\equiv} \{h_s\}$.

3. RK-KDM Security

A pair of efficient probabilistic algorithms $(\mathsf{Enc}, \mathsf{Dec})$ is a *symmetric encryption scheme* over the message-space $\{0,1\}^*$ and key-space $\{0,1\}^k$ (where k serves as the security parameter) if for every message $M \in \{0,1\}^*$

$$\Pr_{s \xleftarrow{R} \{0,1\}^k} [\mathsf{Dec}_s(\mathsf{Enc}_s(M)) = M] = 1.$$

We also assume (WLOG) length regularity that messages of equal length M, M' are always encrypted by ciphertexts of equal length $|\mathsf{Enc}_s(M)| = |\mathsf{Enc}_s(M')|$.

Our security definitions are parameterized by a family of key-derivation and key-dependent-message functions (which are also indexed by the security parameter k)

$$\Phi_{\mathsf{RKA}} = \left\{ \phi : \{0,1\}^k \to \{0,1\}^k \right\}, \qquad \Psi_{\mathsf{KDM}} = \left\{ \psi : \{0,1\}^k \to \{0,1\}^* \right\}.$$

By default, we assume that Φ_{RKA} contains (at least) the identity function and that Ψ_{KDM} contains (at least) all constant functions $\psi_M : \{0,1\}^k \to M$ for every $M \in \{0,1\}^k$. The families Φ_{RKA} and Ψ_{KDM} determine the legal relations between the related-keys and the key-related messages. RK-KDM security is defined via the following pair of real/fake oracles Real_s and Fake_s, which are indexed by a key $s \in \{0,1\}^k$. For a query $(\phi \in \Phi_{\mathsf{RKA}}, \psi \in \Psi_{\mathsf{KDM}})$, the oracle Real_s returns a sample from the distribution $\mathsf{Enc}_{\phi(s)}(\psi(s))$, whereas the oracle Fake_s returns a sample from the distribution $\mathsf{Enc}_{\phi(s)}(0^{|\psi(s)|})$.

Definition 3.1. (*RK-KDM secure encryption*) A symmetric encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ is *semantically secure under related-key and key-dependent message attacks* (in short, *RK-KDM-secure*) with respect to $\Phi_{\mathsf{RKA}}, \Psi_{\mathsf{KDM}}$ if $\mathsf{Real}_s \overset{c}{\equiv} \mathsf{Fake}_s$ where $s \xleftarrow{R} \{0,1\}^k$.

Remark.

- **Relation to Previous Definitions** We note that the above definition generalizes semantic security under related-key attacks [3] and semantic security under key-dependent message attacks [10]. Indeed, the former notion is obtained by restricting Ψ_{KDM} to contain only constant functions, and the latter is obtained by letting Φ_{RKA} contain only the identity function. If both restrictions are applied simultaneously, the definition becomes identical to standard semantic security under chosen-plaintext attacks. On the other hand, as we show in Sect. 5, a scheme may satisfy both RKA security and KDM security (separately) without achieving the combined form of RK-KDM security.

- **Non-adaptivity** Definition 3.1 allows the adversary to choose its queries in a fully adaptive way. One may define a seemingly weaker nonadaptive variant in which the adversary has to specify all its queries at the beginning of the game. We note that this weaker variant suffices for the free-XOR application.

- **LIN RK-KDM Security** We will be interested in linear functions over $\mathbb{F}_2$. Namely, both Φ_{RKA} and Ψ_{KDM} contain functions of the form $s \mapsto s + \Delta$ for every $\Delta \in \mathbb{F}_2^k$. To be compatible with standard semantic security, we require that Ψ_{KDM} also contains all fixed functions. Using a compact notation, we can describe each function in Ψ_{KDM} by a message M and a bit σ and let $g_{M,\sigma} : s \mapsto (M + (\sigma \cdot s))$. If the length of M is larger than k, we assume that $(\sigma \cdot s)$ is padded with zeroes at the end. Hence, the adversary may ask for an encryption of the shifted key concatenated with some fixed message. We refer to this notion as *LIN RK-KDM* security.[1]

3.1. *LPN-based Construction*

We recall the learning parity with noise (LPN) problem, due to [11,20]. For a noise parameter $\varepsilon \in (0, \frac{1}{2})$, a positive integer k and a vector $s \in \mathbb{F}_2^k$ define a randomized function $\mathsf{LPN}_{\varepsilon,s}$, which ignores its input and in each invocation outputs a pair $(a, y = as + e) \in \mathbb{F}_2^k \times \mathbb{F}_2$ where $a \xleftarrow{R} \mathbb{F}_2^k$ is a fresh random vector and $e \xleftarrow{R} \mathsf{Ber}_\varepsilon$ is a fresh "error" bit, which takes the value 1 with probability ε. We view $\mathsf{LPN}_{\varepsilon,s}$ as an oracle that provides noisy evaluations of the linear function $f_s : x \mapsto s \cdot x$ with respect to random inputs. The LPN_ε assumption asserts that it is hard to learn the function (i.e., recover s) given polynomially many samples.

Assumption 3.2. *(*LPN_ε*) For every efficient adversary $\mathcal{A}$, the winning probability*

$$\Pr_{s \xleftarrow{R} \mathbb{F}_2^k} [\mathcal{A}^{\mathsf{LPN}_{\varepsilon,s}}(1^k) = s] \quad \text{is negligible in } k.$$

---

[1] A seemingly weaker definition of LIN RK-KDM security restricts the KDM family to functions $g_{M,\sigma} : s \mapsto (M + (\sigma \cdot s))$ where M and s are of the same length k. We note that a scheme that satisfies this notion can be trivially converted into a scheme that satisfies our definition (which supports M longer than s). This can be done by partitioning the long message M into t blocks $M_1, \ldots, M_t$ of length k each and concatenating the encryptions of these blocks. A query of the form $(f \in \Phi_{\mathsf{RKA}}, g_{M,\sigma})$ can then be emulated by a linear query $(f \in \Phi_{\mathsf{RKA}}, g_{M_1,1})$ and $t - 1$ fixed-message query $(f \in \Phi_{\mathsf{RKA}}, g_{M_i,0})$.

It is widely believed that LPN_ε is hard for any constant $\varepsilon \in (0, \frac{1}{2})$, and the best known algorithm runs in time $2^{\Theta(n/\log n)}$ [12]. In the following, we describe the LPN-based symmetric encryption scheme of [2], which is a variant of the scheme of [19]. We begin with few definitions.

Error Correcting Codes   A pair of efficient algorithms (Code, Cor) is a linear δ-error correcting codes with an expansion $L : \mathbb{N} \to \mathbb{N}$ if for every message length $\ell \in \mathbb{N}$ and codeword length $L = L(\ell) \in \mathbb{N}$ the followings hold:

- (Linearity) For every pair of messages $x, x' \in \mathbb{F}_2^\ell$, $\mathsf{Code}(x) + \mathsf{Code}(x') = \mathsf{Code}(x + x') \in \mathbb{F}_2^L$. Note that this means that $\mathsf{Code}(x) = Gx$ for some generating matrix $G \in \mathbb{F}_2^{L \times \ell}$. Furthermore, since Code is efficient, one can efficiently find such a generating matrix.
- (δ-error correction) For every message $x \in \mathbb{F}_2^\ell$ and every error vector $e \in \mathbb{F}_2^L$ of Hamming weight at most δL, we have that $\mathsf{Cor}(\mathsf{Code}(x) + e) = x$.

We note that the efficiency requirement implies that the expansion of the code $L(\ell)$ is polynomially bounded.

Chopped Noise Distribution   For constant $\varepsilon \in (0, 1)$, let $\mathsf{Ber}_\varepsilon^{t \times N}$ be the distribution over $t \times N$ binary matrices obtained by setting each entry to 1 independently with probability ε. For a constant $\varepsilon < \delta < 1$, we define the δ-"chopped" version of $\mathsf{Ber}_\varepsilon^{t \times N}$, denoted by $\mathsf{Ber}_{\varepsilon,\delta}^{t \times N}$, to be the distribution obtained by choosing $E \xleftarrow{R} \mathsf{Ber}_\varepsilon^{t \times N}$ and swapping each column of E whose hamming weight exceeds δt with the all zero column. By Chernoff bound, when N and t are polynomial in k and ε and δ are constants, the statistical distance between $\mathsf{Ber}_\varepsilon^{t \times N}$ and $\mathsf{Ber}_{\varepsilon,\delta}^{t \times N}$ is negligible in k.

Construction 3.3.   (LPN construction) *The scheme is parameterized with constants $0 < \varepsilon < \delta < \frac{1}{2}$, polynomially bounded functions $N = N(k)$, $\ell = \ell(k)$ and with an efficient linear δ-error correcting code (Code, Cor). We let $t = t(k)$ denote the length of a codeword, which corresponds to a message of length $\ell(k)$.*

- Key generation*: The private key of the scheme is a matrix S which is chosen uniformly at random from $\mathbb{F}_2^{k \times N}$.*
- Encryption*: To encrypt a message $M \in \mathbb{F}_2^{\ell \times N}$, choose a random $A \xleftarrow{R} \mathbb{F}_2^{t \times k}$ and a random noise matrix $E \xleftarrow{R} \mathsf{Ber}_{\varepsilon,\delta}^{t \times N}$. Output the ciphertext*

$$(A, A \cdot S + E + GM),$$

  where $G \in \mathbb{F}_2^{t \times \ell}$ is the generating matrix of the code.
- Decryption*: Given a ciphertext (A, Z) apply the correction algorithm Cor to each of the columns of the matrix $Z - AS$ and output the result.*

Observe that the correction algorithm never errs as E never contains a column whose Hamming weight is larger than δt. The scheme is also highly efficient. Encryption requires only cheap matrix operations, and decryption requires in addition to decode the code. It is shown in [2] that for proper choice of parameters, both encryption and

decryption can be done in quasilinear time in the message length (for sufficiently long message).[2] See [19] for a practical evaluation of similar LPN-based encryption schemes.

Construction 3.3 was proven to be semantically secure based on the intractability of the LPN_ε problem [2]. Security against KDM and RKA attacks with respect to linear functions was further proven in [2] and [3]. We now generalize these results and show that the scheme is LIN RK-KDM secure.

Theorem 3.4. *Assuming that* LPN_ε *is hard, the above construction is LIN RK-KDM secure.*

3.2. *Proof of Theorem 3.4*

Through this section, we keep the convention that $S \in \mathbb{F}_2^{k \times N}$ is a key, $\Delta \in \mathbb{F}_2^{k \times N}$ is a key-shift vector, $M \in \mathbb{F}_2^{\ell \times N}$ is a message, $b \in \{0, 1\}$ is a bit, and the pair $(A, Z) \in \mathbb{F}_2^{t \times k} \times \mathbb{F}_2^{t \times N}$ is a potential ciphertext. In addition, we let Enc denote the LPN encryption defined in Construction 3.3.

Recall that our goal is to prove that for a random key $S \xleftarrow{R} \mathbb{F}_2^{k \times N}$, the randomized functions

$$\mathsf{Real}_S : (\Delta, M, b) \mapsto \mathsf{Enc}_{S+\Delta}(M + bS)$$
$$\mathsf{Fake}_S : (\Delta, M, b) \mapsto \mathsf{Enc}_{S+\Delta}(0^{\ell \times N}),$$

are indistinguishable. This will be proven via a sequence of hybrids.

Let $\mathcal{R}_S$ be a randomized function, which ignores the key S and the given input and outputs a fresh uniformly chosen matrices $A \xleftarrow{R} \mathbb{F}_2^{t \times k}$ and $Z \xleftarrow{R} \mathbb{F}_2^{t \times N}$. (If $\mathcal{R}_S$ is applied to the same input more than once, it responds with independent answers.)

The following lemma shows that the LPN encryption scheme is not only semantically secure but also pseudorandom in the following sense:

Lemma 3.5. *Assuming that* LPN_ε *is hard,* $\{\mathsf{Enc}_S\} \stackrel{c}{\equiv} \{\mathcal{R}_S\}$, *where* $S \xleftarrow{R} \mathbb{F}_2^{k \times N}$.

The proof is implicit in [2], and we include it here for completeness.

Proof. Fix some $\varepsilon \in (0, \frac{1}{2})$. For polynomials $N, t = \mathrm{poly}(k)$ and $S \xleftarrow{R} \mathbb{F}_2^{k \times N}$, we define the randomized functions $\mathsf{LPN}_S^{t \times N}$ and $\mathcal{R}_S^{t \times N}$ which have no input (or equivalently ignore their input) as follows. In each call, $\mathsf{LPN}_S^{t \times N}$ samples a random matrix $A \xleftarrow{R} \mathbb{F}_2^{t \times k}$, a random noise matrix $E \xleftarrow{R} \mathsf{Ber}_\varepsilon^{t \times N}$, and outputs the pair $(A, A \cdot S + E)$. The function $\mathcal{R}_S^{t \times N}$ is defined similarly to $\mathcal{R}_S$, namely in each call it simply outputs a fresh random pair $A \xleftarrow{R} \mathbb{F}_2^{t \times k}$ and $Z \xleftarrow{R} \mathbb{F}_2^{t \times N}$. The well-known search-to-decision reduction of [11]

---

[2] This asymptotically fast implementation is based on efficient noise sampling algorithm, fast error correcting code (e.g., Spielman's codes [43]), and fast rectangular matrix multiplication. The latter requires $N, \ell > k^6$.

shows that, under the LPN_ε assumption,

$$\left\{ \mathsf{LPN}_S^{t \times N} \right\} \stackrel{\mathrm{c}}{\equiv} \left\{ \mathcal{R}_S^{t \times N} \right\}, \tag{1}$$

for $N = 1$ and any polynomial t. A standard hybrid argument allows to extend Eq. 1 to the case of an arbitrary polynomial N (and arbitrary polynomial t), as done in [2]. It remains to show that Eq. 1 implies the lemma.

Fix t, N to be the parameters from Construction 3.3, and let $G \in \mathbb{F}_2^{t \times \ell}$ be the generator matrix in use. Define an oracle-aided function $\mathcal{A}^{(\cdot)}$, which given $M \in \mathbb{F}_2^{\ell \times N}$ calls its oracle $\mathcal{O}$ to obtain a pair (A, R) and outputs $(A, R + GM)$.

For every S, we have that

$$\mathcal{R}_S \equiv \mathcal{A}^{\mathcal{R}_S^{t \times N}} \qquad \text{and} \qquad \mathcal{A}^{\mathsf{LPN}_S^{t \times N}} \stackrel{\mathrm{c}}{\equiv} \mathsf{Enc}_S.$$

The first part follows immediately from the definition of $\mathcal{A}$. To see the second part, note that the only difference between the two distributions is due to the fact that Enc_S uses the chopped noise distribution $\mathsf{Ber}_{\varepsilon,\delta}^{t \times N}$, whereas $\mathcal{A}^{\mathsf{LPN}_S^{t \times N}}$ uses the non-chopped distribution $\mathsf{Ber}_\varepsilon^{t \times N}$. The statistical distance between the two distributions is negligible in k, and therefore, a computationally bounded adversary (which makes only a polynomial number of calls to these distributions) cannot distinguish between Enc_S and $\mathcal{A}^{\mathsf{LPN}_S^{t \times N}}$ with more than negligible advantage.

By combining this with Eq. 1 and Fact 2.1, we have that for a random S

$$\mathcal{R}_S \equiv \mathcal{A}^{\mathcal{R}_S^{t \times N}} \stackrel{\mathrm{c}}{\equiv} \mathcal{A}^{\mathsf{LPN}_S^{t \times N}} \stackrel{\mathrm{c}}{\equiv} \mathsf{Enc}_S,$$

and the lemma follows by transitivity (Fact 2.2). $\qquad\qquad\qquad\qquad\qquad\square$

We will need the following key observation:

Lemma 3.6. *There exists an efficient oracle machine* $F^{(\cdot)} : (\Delta, M, b) \mapsto (A, Z)$ *such that*

$$\mathsf{Real}_S \equiv F^{\mathsf{Enc}_S} \quad and \quad F^{\mathcal{R}_S} \equiv \mathcal{R}_S,$$

for every $S \in \mathbb{F}_2^{k \times N}$.

Proof. We define F as follows: Given a query (Δ, M, b), the machine F calls the oracle with input M, gets back the answer (A', Z'), and outputs the pair $A = A' + GH$ and $Z = Z' + A\Delta$ where G is the generating matrix used in construction 3.3 and $H \in \mathbb{F}_2^{\ell \times k}$ is the matrix $\begin{pmatrix} b \cdot I_{k \times k} \\ 0^{\ell - k \times k} \end{pmatrix}$.

Fix a key S and a query (Δ, M, b), we will show that $F^{\mathsf{Enc}_S}(\Delta, M, b)$ is distributed identically to $\mathsf{Real}_S(\Delta, M, b)$. Let (A', Z') be a fresh sample from $\mathsf{Enc}_S(M)$. Clearly, $A = A' + GH$ is uniform in $\mathbb{F}_2^{t \times k}$ since A' is uniform. In addition, since $Z' = A' \cdot S +$

$E + G \cdot M$ where $E \xleftarrow{R} \mathsf{Ber}_{\varepsilon,\delta}^{t \times N}$, and since $A' = A + GH$, we can write Z as

$$(A + GH) \cdot S + E + G \cdot M + A\Delta = A \cdot (S + \Delta) + E + G \cdot (M + HS)$$
$$= A \cdot (S + \Delta) + E + G \cdot (M + bS),$$

where the first equality is due to linearity, and the second equality follows from the definition of H. It follows that (A, Z) is a fresh sample from $\mathsf{Enc}_{S+\Delta}(M + bS)$.

To prove that $F^{\mathcal{R}_S} \equiv \mathcal{R}_S$, it suffices to show that for any fixed query (Δ, M, b), the transformation from (A', Z') to (A, Z) is an affine invertible mapping. This follows immediately from the definition of F. □

We conclude that for $S \xleftarrow{R} \mathbb{F}_2^{k \times N}$,

$$\mathsf{Real}_S \equiv F^{\mathsf{Enc}_S} \stackrel{c}{\equiv} F^{\mathcal{R}_S} \equiv \mathcal{R}_S. \tag{2}$$

Indeed, the first and third transitions are due to Lemma 3.6, and the second transition is due to Lemma 3.5 and Fact 2.1.

To complete the argument, we need two additional definitions. First, we define an oracle machine, which given an oracle $\mathcal{O}$ and an input (Δ, M, b) outputs a sample from $F^{\mathcal{O}}(\Delta, 0^{\ell \times N}, 0)$; namely, it replaces M, b with zeroes and proceeds as $F^{\mathcal{O}}$. By abuse of notation, we refer to this oracle as $F(\cdot, 0^{\ell \times N}, 0)$. Similarly, we let $\mathsf{Real}_S(\cdot, 0^{\ell \times N}, 0)$ denote the randomized function, which maps (Δ, M, b) to $\mathsf{Real}_S(\Delta, 0^{\ell \times N}, 0)$. Note that the latter is just an equivalent formulation of Fake_S. Moreover, we can write:

$$\mathcal{R}_S \equiv F(\cdot, 0^{\ell \times N}, 0)^{\mathcal{R}_S} \stackrel{c}{\equiv} F(\cdot, 0^{\ell \times N}, 0)^{\mathsf{Enc}_S(0^{\ell \times N})}$$
$$\equiv \mathsf{Real}_S(\cdot, 0^{\ell \times N}, 0) \equiv \mathsf{Fake}_S, \tag{3}$$

where the first and third transitions are due to Lemma 3.6, and the second transition is due to Lemma 3.5 and Fact 2.1. By combining Eq. 2 and Eq. 3 with Fact 2.2, we get that $\mathsf{Real}_S \stackrel{c}{\equiv} \mathsf{Fake}_S$, and Theorem 3.4 follows. □

Remark 3.7. (Abstraction) The proof of Theorem 3.4 provides a general template for proving RK-KDM security. Specifically, the properties needed are pseudorandomness (in the sense of Lemma 3.5) and key/message homomorphism (in the sense of Lemma 3.6). Indeed, observe that, apart from the proofs of Lemmas 3.5 and 3.6, the overall proof can be written in a fully generic form with no specific references to the LPN construction.

4. Yao's Garbled Circuit

4.1. *Definition*

Let $f = \{f_n\}_{n \in \mathbb{N}}$ be a polynomial-time computable function. In an abstract level, Yao's garbled circuit technique [45] constructs a randomized function $\hat{f} = \{\hat{f}_n\}_{n \in \mathbb{N}}$, which "encodes" f in the sense that for every x, the distribution $\hat{f}(x)$ reveals the value of $f(x)$

but no other additional information. We formalize this via the notion of *computationally private randomized encoding* from [4], while adopting the original definition from a nonuniform adversarial setting to the uniform setting (i.e., adversaries are modeled by probabilistic polynomial-time Turing machines).

Definition 4.1. (*Computational randomized encoding*) Let $f = \{f_n : \{0, 1\}^n \to \{0, 1\}^{\ell(n)}\}_{n \in \mathbb{N}}$ be an efficiently computable function and let $\hat{f} = \{\hat{f}_n : \{0, 1\}^n \times \{0, 1\}^{m(n)} \to \{0, 1\}^{s(n)}\}_{n \in \mathbb{N}}$ be an efficiently computable randomized function. We say that $\hat{f}$ is a *computational randomized encoding* of f (or encoding for short), if there exist an efficient recovery algorithm Rec and an efficient probabilistic simulator algorithm Sim that satisfy the following:

- **Perfect correctness.** For any n and any input $x \in \{0, 1\}^n$,

$$\Pr[\mathsf{Rec}(1^n, \hat{f}_n(x)) \neq f_n(x)] = 0,$$

  where the probability is taken over the internal randomness of $\hat{f}_n$.
- **Computational privacy.** The randomized function $\hat{f}_n(\cdot)$ is computationally indistinguishable from the randomized function $\mathsf{Sim}(1^n, f_n(\cdot))$.

Remark 4.2. The above definition uses n both as an input length parameter and as a cryptographic "security parameter" quantifying computational privacy. When describing the construction, it will be convenient to use a separate parameter k for the latter, where computational privacy will be guaranteed as long as $k = n^\epsilon$ for some constant $\epsilon > 0$. (An alternative definition which is parameterized by both the input length and the security parameter is discussed in "Appendix".) Furthermore, while it is convenient to define randomized encoding for a single function f, Yao's construction (as well as the free-XOR variant) actually provides an efficient *compiler* that maps the function f (represented as a Boolean circuit) into (circuit representations of) the encoding $\hat{f}$, the recovery algorithm Rec, and the simulator Sim. (See [5] for formal definition.) In this sense, the encoding is fully constructive.

4.2. *Yao's Construction and the Free-XOR Variant*

Let $f = \{f_n : \{0, 1\}^n \to \{0, 1\}^{\ell(n)}\}_{n \in \mathbb{N}}$ be a polynomial-time computable function computed by the uniform circuit family $\{C_n\}_{n \in \mathbb{N}}$. In the following, we describe Yao's construction and its free-XOR variant. Our notation and terminology borrow from previous presentations of Yao's construction in [4,33,38,41].

Double-Keyed Encryption Let $k = k(n)$ be a security parameter (by default, $k = n^\varepsilon$ for some constant $\varepsilon > 0$). We will employ a symmetric encryption scheme (E^2, D^2), which is keyed by a *pair* of k-bit keys K_1, K_2. Intuitively, this corresponds to a double-locked chest in the sense that decryption is possible only if one knows both keys. There are several ways to implement such an encryption scheme based on standard single-key symmetric encryption (Enc, Dec), and for simplicity, we choose to use

$$E^2_{K_1, K_2}(M) := (\mathsf{Enc}_{K_1}(R), \mathsf{Enc}_{K_2}(R + M)),$$
$$D^2_{K_1, K_2}(C_1, C_2) := \mathsf{Dec}_{K_1}(C_1) + \mathsf{Dec}_{K_2}(C_2) \qquad\qquad (4)$$

where R is a random string of length $|M|$. Other choices are also applicable under the LPN assumption.

The Original Construction   For each wire i of the circuit C_n, we assign a pair of keys: a 0-key $W_i^0 \in \{0, 1\}^k$ that represents the value 0, and a 1-key $W_i^1 \in \{0, 1\}^k$ that represents the value 1. For each of these pairs, we randomly "color" one key black and the other key white. This is done by choosing $r_i \overset{R}{\leftarrow} \{0, 1\}$ and by letting $c_i = r_i + b$ be the color of W_i^b. Fix some input x for f_n, and let $b_i = b_i(x)$ be the value of the ith wire induced by x. We refer to the key $W_i^{b_i}$ as the *active* key of the ith wire.

The encoding $\hat{f}_n(x)$ consists of three parts: (1) The active keys $W_i^{b_i}$ of the input wires together with their colors c_i; (2) For each gate, a propagation mechanism allows to translate the colored active keys of the incoming wires into the colored active keys of the outgoing wires. This mechanism is implemented via an encryption table (or "gate label") in which the keys of the outgoing wire are encrypted under the keys of the incoming wires. (3) For each output wire i, we also append the semantics of the coloring, i.e., the bit r_i. Altogether, one can propagate the values of the colored active keys $(W_i^{b_i}, c_i)$ from the inputs to the outputs, and at the end reveal, the values of the output wires by unmasking the colors c_i with r_i. Intuitively, privacy holds as for non-output wires the values of the colored active keys reveal nothing on their semantics b_i.

Free-XOR Gates   Consider a XOR gate with incoming wires i and j and outgoing wire ℓ. The "free-XOR" optimization modifies the above construction by making sure that the colored active key of the outgoing wire is simply the sum of the colored active keys of the incoming wires; namely,

$$\left(W_\ell^{b_\ell(x)}, c_\ell(x)\right) = \left(W_i^{b_i(x)}, c_i(x)\right) + \left(W_j^{b_j(x)}, c_j(x)\right), \qquad \text{for every input } x. \quad (5)$$

As a result, gate labels are not needed and XOR gates have no effect on the communication complexity of the encoding, and only a minor effect on the computational complexity.

To satisfy Eq. 5, we apply the following modifications. First, we set the zero-key W_ℓ^0 and coloring r_ℓ of a wire which outgoes a XOR gate to be the sum of the zero-keys and coloring of the incoming wires i and j, namely

$$W_\ell^0 = W_i^0 + W_j^0, \qquad\qquad r_\ell = r_i + r_j.$$

Second, instead of choosing the one-key at random, we will choose them based on the zero-key. That is, for every wire t, we let $W_t^1 = W_t^0 + s$ where s is a global (secret) shift vector. As a result, for every pair of values $(\alpha, \beta) \in \{0, 1\}^2$ for the input wires of a XOR gate, we have that
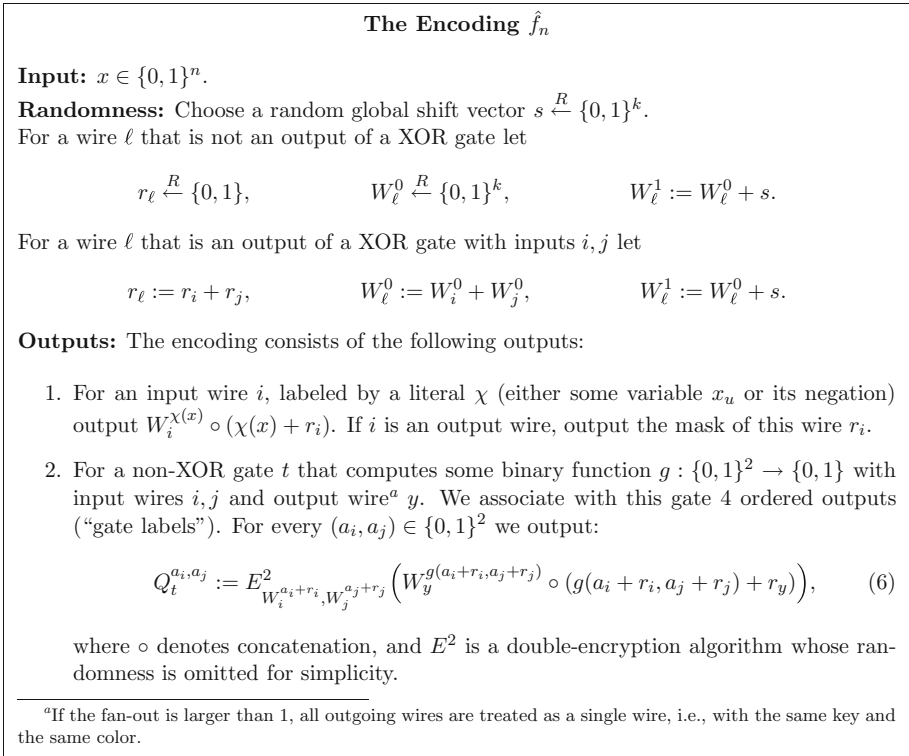

The Encoding $\hat{f}_n$

Input: $x \in \{0,1\}^n$.

Randomness: Choose a random global shift vector $s \stackrel{R}{\leftarrow} \{0,1\}^k$.
For a wire ℓ that is not an output of a XOR gate let

$$r_\ell \stackrel{R}{\leftarrow} \{0,1\}, \qquad W_\ell^0 \stackrel{R}{\leftarrow} \{0,1\}^k, \qquad W_\ell^1 := W_\ell^0 + s.$$

For a wire ℓ that is an output of a XOR gate with inputs i, j let

$$r_\ell := r_i + r_j, \qquad W_\ell^0 := W_i^0 + W_j^0, \qquad W_\ell^1 := W_\ell^0 + s.$$

Outputs: The encoding consists of the following outputs:

1. For an input wire i, labeled by a literal χ (either some variable x_u or its negation) output $W_i^{\chi(x)} \circ (\chi(x) + r_i)$. If i is an output wire, output the mask of this wire r_i.

2. For a non-XOR gate t that computes some binary function $g : \{0,1\}^2 \to \{0,1\}$ with input wires i, j and output wire[a] y. We associate with this gate 4 ordered outputs ("gate labels"). For every $(a_i, a_j) \in \{0,1\}^2$ we output:

$$Q_t^{a_i, a_j} := E^2_{W_i^{a_i + r_i}, W_j^{a_j + r_j}} \left(W_y^{g(a_i + r_i, a_j + r_j)} \circ (g(a_i + r_i, a_j + r_j) + r_y) \right), \qquad (6)$$

where $\circ$ denotes concatenation, and E^2 is a double-encryption algorithm whose randomness is omitted for simplicity.

---
[a]If the fan-out is larger than 1, all outgoing wires are treated as a single wire, i.e., with the same key and the same color.

Fig. 1. The encoding $\hat{f}_n(x; (W, r, s))$ of the function $f_n(x)$. We assume that wires and gates of the circuit that computes f_n are numbered according to some topological order. The double-encryption algorithm $E^2_{K_1, K_2}(M)$ is defined based on a standard encryption (Enc, Dec) as in Eq. 4 .

$$W_\ell^{\alpha + \beta} = W_i^\alpha + W_j^\beta.$$

Hence, one can derive the colored active key $(W_\ell^{b_\ell(x)}, r_\ell + b_\ell(x))$ of the output wire by XOR-ing the colored active keys $(W_i^{b_i(x)}, r_i + b_i(x))$, $(W_j^{b_j(x)}, r_j + b_j(x))$ of the input wires, as required. A formal description of the encoding is given in Fig. 1.

Our main result shows that, assuming LIN RK-KDM security, the free-XOR variant gives rise to a valid computational encoding:

Theorem 4.3. (Main) *If the underlying symmetric encryption scheme* (Enc, Dec) *is LIN RK-KDM secure, then the randomized function $\hat{f}$, as defined in Fig. 1, is a randomized encoding of the function f.*

The proof of the theorem is deferred to Sect. 4.3 (correctness) and 4.4 (privacy).

4.3. *Correctness*

The following lemma shows that the encoding is correct.

Lemma 4.4. (Correctness) *There exists an efficient recovery algorithm* Rec *such that for every* $x \in \{0, 1\}^n$, *it holds that*

$$\Pr[\mathsf{Rec}(1^n, \hat{f}_n(x)) \neq f_n(x)] = 0,$$

where the probability is taken over the internal randomness of $\hat{f}_n$.

Proof. Let $\alpha = \hat{f}_n(x; (r, W, s))$ for some input $x \in \{0, 1\}^n$ and coins $(r, W, s) \in \{0, 1\}^{m(n)}$. The recovery algorithm traverses the circuit in topological order from inputs to outputs, and for each wire y, it recovers the active key $W_y^{b_y}$ together with its color $c_y = (b_y(x) + r_y)$ as follows.

If y is an input wire, then the value $W_y^{b_y} \circ c_y$ is given as part of α. Otherwise, assume that the wire y outgoes a gate t whose incoming wires are i and j (for which we already computed the desired values). If t is a XOR gate, then we let

$$W_y^{b_y} = W_y^{b_i+b_j} = W_i^{b_i} + W_i^{b_j}, \text{ and } c_y = (b_i + b_j) + r_y$$
$$= (b_i + b_j) + (r_i + r_j) = c_i + c_j.$$

If t is not a XOR gate, then we use the colors c_i, c_j of the active keys of the input wires to select the *active* label $Q_t^{c_i,c_j}$ of the gate t (and ignore the other 3 *inactive* labels of this gate). Consider this label as in Eq. (6); recall that this cipher was "double-encrypted" under the key $W_i^{c_i-r_i} = W_i^{b_i}$ and the key $W_j^{c_j-r_j} = W_j^{b_j}$. Since we have already computed the values $c_i, c_j, W_i^{b_i}$ and $W_j^{b_j}$, we can decrypt the label $Q_t^{c_i,c_j}$ (by applying the decryption algorithm D^2) and recover the value

$$W_y^{g(b_i,b_j)} \circ (g(b_i, b_j) + r_y) = W_y^{b_y} \circ (c_y),$$

where g is the function that the gate t computes, and therefore, $b_y = g(b_i, b_j)$.

Finally, once we have the colors of an output wire y, we can recover its value b_y by XOR-ing c_y with the mask r_y, which is given explicitly as part of α. □

4.4. *Privacy*

Computational privacy is slightly more subtle. The free-XOR optimization correlates the key pairs via the global shift s. This introduces two form of dependencies: (1) The four ciphertexts of every gate are encrypted under *related-keys*, and (2) the keys (of the incoming wires) which are used to encrypt the gate labels are correlated with the content of the labels (i.e., the keys of the outgoing wires). We show that if the underlying encryption $(\mathsf{Enc}, \mathsf{Dec})$ is RK-KDM secure with respect to linear functions, then the encoding is indeed private.

Lemma 4.5. (Privacy) *There exists an efficient simulator* Sim *such that*

$$\hat{f}_n(\cdot) \overset{c}{\equiv} \mathsf{Sim}(1^n, f_n(\cdot)).$$

To prove the lemma, we define an oracle-aided algorithm $H^{\mathcal{O}}(x)$ such that (1) when the oracle $\mathcal{O}$ is the real RK-KDM oracle (with respect to linear queries), the distribution of $H^{\mathcal{O}}(x)$ is identical to the distribution $\hat{f}_n(x)$, and (2) when the oracle $\mathcal{O}$ is the fake RK-KDM oracle, the distribution $H^{\mathcal{O}}(x)$ can be efficiently sampled based on the output $f_n(x)$ and therefore can be used as a simulator $\mathsf{Sim}(1^n, f_n(x))$. The indistinguishability of the two oracles implies that the simulator's output is computationally indistinguishable from the encoding's distribution $\hat{f}_n(x)$.

The Algorithm $H^{(\cdot)}(x)$. Let $k = k(n)$, $x \in \{0, 1\}^n$ be the input. We assume that H is given an oracle access to a randomized function $\mathcal{O}_s$ where $s \xleftarrow{R} \{0, 1\}^k$ will play the role of the secret global shifts. We will assume that $\mathcal{O}_s$ has the same interface as Real_s and Fake_s, namely given a pair of linear functions (ϕ, ψ), the oracle outputs a ciphertext of Enc. For every wire ℓ, we define the following values:

1. If ℓ is not an output of a XOR gate, choose a random active key $W_\ell^{b_\ell} \xleftarrow{R} \{0, 1\}^k$ and a random color bit $c_\ell \xleftarrow{R} \{0, 1\}$.
2. If the wire ℓ is an output of a XOR gate, set the active key to be $W_\ell^{b_\ell} := W_i^{b_i} + W_j^{b_j}$ and set its color to $c_\ell = c_i + c_j$ where i and j are the incoming wires.
3. If ℓ is an input wire, output the colored active key $W_\ell^{b_\ell} \circ c_\ell$; if it is an output wire, output $r_\ell = c_\ell - b_\ell(x)$.
4. The inactive key $W_\ell^{b_\ell+1}$ is unknown, but it can be written as a linear function of the master-key s, i.e., $\phi_\ell : s \mapsto s + W_\ell^{b_\ell}$.

For every (non-XOR) gate t with input wires i, j and output wire y, we do the following:

5. Output the active label

$$Q_t^{c_i,c_j} := E_{W_i^{b_i},W_j^{b_j}}^2 (W_y^{b_y} \circ c_y) \tag{7}$$

6. Compute the inactive labels as follows. For every $(\alpha, \beta) \neq (0, 0)$, choose $R_{\alpha,\beta} \xleftarrow{R} \{0, 1\}^{k+1}$ and define the linear function $\psi_{\alpha,\beta}$ which maps s to the value

$$\Big((W_y^{b_y} + s \cdot g(b_i + \alpha, b_j + \beta) + b_y) \circ (g(c_i + \alpha + r_i, c_j + \beta + r_j) + r_y)\Big) + R_{\alpha,\beta},$$

where g is the function that the gate computes, and $b_i = b_i(x)$, $r_i = b_i + c_i$, $b_j = b_j(x)$, $r_j = b_j + c_j$ and $b_y = b_y(x)$, $r_y = b_y + c_y$. Now, output

$$Q_t^{c_i+1,c_j} := \Big(\mathcal{O}(\phi_i, \psi_{1,0}), \mathsf{Enc}_{W_j^{b_j}}(R_{1,0})\Big)$$

$$Q_t^{c_i+1,c_j+1} := \Big(\mathcal{O}(\phi_i, \psi_{1,1}), \mathcal{O}(\phi_j, R_{1,1})\Big)$$

$$Q_t^{c_i,c_j+1} := \Big(\mathsf{Enc}_{W_i^{b_i}}(R_{0,1}), \mathcal{O}(\phi_j, \psi_{0,1})\Big), \tag{8}$$

where in the second equation, we let the string $R_{1,1}$ represent the constant function $s \mapsto R_{1,1}$.

Claim 4.6. *The randomized functions $\hat{f}_n$ and H^{Real_s} for $s \overset{R}{\leftarrow} \{0,1\}^k$ are identically distributed.*

Proof. We prove a stronger claim: for every $x \in \{0,1\}^n$ even if the encoding and the hybrid $H^{\mathsf{Real}_s}(x)$ output their internal coins (including the ones used by the oracle Real_s), the two experiments are identically distributed. First, it is not hard to verify that the values s, W_ℓ^0, r_ℓ and $W_\ell^1 = W_\ell^0 + s$ are identically distributed in both experiments. When these values are fixed, the active labels are also identically distributed. Finally, by substituting ϕ_i, $\psi_{\alpha,\beta}$ in Eq. 8, it follows that the inactive labels are also distributed exactly as in $\hat{f}(x)$. □

Let us move to the case where the oracle $\mathcal{O}$ is instantiated with the oracle Fake_s for $s \overset{R}{\leftarrow} \{0,1\}^k$. By the RK-KDM security of the scheme $(\mathsf{Enc}, \mathsf{Dec})$ and Fact 2.1, we get that

Claim 4.7. *The randomized functions $\{H^{\mathsf{Real}_s}\}_s$ and $\{H^{\mathsf{Fake}_s}\}_s$ are computationally indistinguishable.*

Finally, we define the simulator, which is just an equivalent description of $H^{\mathsf{Fake}_s}(x)$:

The Simulator Sim. Given $z = f_n(x)$, for some $x \in \{0,1\}^n$, the simulator mimics the first three steps of H which can be computed based on the value of the output wires $f_n(x)$ (without knowing x itself). However, instead of virtually setting inactive keys in the forth step, the simulator chooses a random shift vector $s \overset{R}{\leftarrow} \{0,1\}^k$ and sets $W_\ell^{1+b_\ell} = W_\ell^{b_\ell} + s$ for every wire ℓ. Then, the simulator computes the active labels exactly as in Eq. 7. Note that all these computations can be done without knowing x (or $b_i(x)$). To compute the inactive labels, the simulator mimics the distribution of $H^{\mathsf{Fake}_s}(x)$: It chooses $R_{1,0}, R_{1,1}, R_{0,1} \overset{R}{\leftarrow} \{0,1\}^{k+1}$ and computes

$$Q_t^{c_i+1,c_j} := \left(\mathsf{Enc}_{W_i^{b_i+1}}(0^{k+1}), \mathsf{Enc}_{W_j^{b_j}}(R_{1,0}) \right)$$

$$Q_t^{c_i+1,c_j+1} := \left(\mathsf{Enc}_{W_i^{b_i+1}}(0^{k+1}), \mathsf{Enc}_{W_j^{b_j+1}}(0^{k+1}) \right)$$

$$Q_t^{c_i,c_j+1} := \left(\mathsf{Enc}_{W_i^{b_i}}(R_{0,1}), \mathsf{Enc}_{W_j^{b_j+1}}(0^{k+1}) \right). \tag{9}$$

Indeed, all these ciphertexts can be computed directly since the inactive keys (and the global shift s) are known.

Claim 4.8. *The randomized functions $\mathsf{Sim}(f_n(\cdot))$ and $H^{\mathsf{Fake}_s}(\cdot)$ for $s \overset{R}{\leftarrow} \{0,1\}^k$ are identically distributed.*

Proof. Again, a stronger claim holds: For every $x \in \{0,1\}^n$ even if the simulator and the algorithm $H^{\mathsf{Fake}_s(\cdot)}(x)$ output their internal coins, the two experiments are identically distributed. First, it is not hard to verify that the values s, W_ℓ^0, r_ℓ and $W_\ell^1 = W_\ell^0 + s$ are

identically distributed in both experiments. When these values are fixed, the active labels are also identically distributed. Finally, the inactive labels as defined by the simulator (Eq. 9) are computed exactly as they are computed by $H^{\mathsf{Fake}_s(\cdot)}(x)$ (i.e., as defined in Eq. 8 when the oracle $\mathsf{Fake}_s(\cdot)$ is being used).   □

The proof of Lemma 4.5 follows from Claims 4.6–4.8 and Fact 2.2.

5. Separating RK-KDM from RKA & KDM

Recall that LIN RKA security corresponds to $(\Phi_{\mathsf{RKA}}, \Psi_{\mathsf{KDM}})$ RK-KDM security where Φ_{RKA} contains all linear functions (over the binary field) and Ψ_{KDM} contains the identity function. Similarly, LIN KDM security corresponds to the complementary case where Ψ_{KDM} contains all linear (and fixed) functions, and Φ_{RKA} contains the identity function.

We describe a symmetric encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$, which is semantically secure under linear related-key attacks and semantically secure under linear key-dependent message attacks but does not achieve linear RK-KDM security. In fact, one can fully recover the secret key via a combined LIN RK-KDM attack. Our counter-example is based on a pair of symmetric encryption schemes. The first scheme $(\mathsf{RE}, \mathsf{RD})$ is LIN RKA secure but can be completely broken via LIN KDM attacks, and the second scheme $(\mathsf{KE}, \mathsf{KD})$ is LIN KDM secure but can be broken via LIN RK attacks. Both schemes are based on the LPN-based encryption of Construction 3.3 instantiated with $N = 1$. Through this section, we denote the LPN encryption scheme by $(\mathsf{PE}, \mathsf{PD})$ ("P" stands for *parity*).

5.1. *Achieving RKA Security & KDM Insecurity*

We define the scheme $(\mathsf{RE}, \mathsf{RD})$ identically to the LPN construction (Construction 3.3) except that if the prefix of a plaintext M is *equal* to the key S, then the corresponding ciphertext will be M itself (unencrypted). Formally[3],

$$
\mathsf{RE}_S(M) := \begin{cases} M & \text{if } M_{[1:k]} = S \\ \mathsf{PE}_S(M) & \text{otherwise.} \end{cases}, \quad
\mathsf{RD}_S(C) := \begin{cases} C & \text{if } C_{[1:k]} = S \\ \mathsf{PD}_S(M) & \text{otherwise.} \end{cases}
$$

It is not hard to prove that $(\mathsf{RE}, \mathsf{RD})$ is secure under linear related-key attacks, but is completely insecure at the presence of linear key-dependent message attacks.

Lemma 5.1. *Under the* LPN *assumption, the scheme* $(\mathsf{RE}, \mathsf{RD})$ *is secure against linear related-key attacks.*

---

[3] The decryption RD may err with negligible probability due to the possibility that some message M, whose prefix does not equal to the key S, will be mapped to a ciphertext $\mathsf{PE}_S(M)$ whose prefix equals to the key. This can be handled in several ways, e.g., by modifying the encryption algorithm so that such event never happens. We prefer the current version (with negligible error probability) for simplicity.

Proof. Recall that in a LIN RK attack on an encryption algorithm E, the adversary makes queries of the form (Δ, M) and attempts to distinguish between the real oracle $E\mathsf{Real}_S$ which returns $E_{S+\Delta}(M)$ and the fake oracle $E\mathsf{Fake}_S$ which returns $E_{S+\Delta}(0^{|M|})$. The view of an adversary $\mathcal{A}$ that breaks the LIN RKA security of $(\mathsf{RE}, \mathsf{RD})$ is identical to the view of an adversary who breaks the LIN RKA security of the LPN-based scheme $(\mathsf{PE}, \mathsf{PD})$, as long as the adversary does not make a *revealing query* of the form (Δ, M) where $S + \Delta$ equals to the k-bit prefix of M. Hence, it suffices to show that the probability of asking a revealing query is negligible. Indeed, this must be the case as a revealing query (Δ, M) can be used to recover the key by XOR-ing Δ with the k-bit prefix of the message $M_{[1:k]}$.

We proceed with a formal argument. Our goal is to prove that $\mathsf{RERel}_S \overset{c}{\equiv} \mathsf{REFake}_S$. First, we show that RERel_S and PERel_S are indistinguishable. Assume, toward a contradiction, that there exists some adversary $\mathcal{A}$, which distinguishes RERel_S from PERel_S with noticeable advantage ε. We construct an adversary $\mathcal{B}^{\mathsf{PERel}_S}$, which outputs S with noticeable probability ε/t where t is the number of queries that $\mathcal{A}$ makes. Clearly, such an adversary contradicts the LIN RKA security of the LPN scheme. The adversary $\mathcal{B}$ simply chooses a random $i \in [t]$ and halts before making the i-th query (Δ, M) with the output $\Delta + M_{[1:k]}$. To analyze the success probability of $\mathcal{B}$, we note that: (a) conditioned on not asking a revealing query, the oracles RERel_S and PERel_S are identically distributed; (b) hence, under our assumption, $\mathcal{A}$ makes a revealing query with probability at least ε; (c) therefore, with probability ε/t, the adversary $\mathcal{B}$ halts just before the first revealing query, and in this case, it outputs the key S.

A similar argument shows that REFake_S is indistinguishable from PEFake_S, and, since $\mathsf{PERel}_S \overset{c}{\equiv} \mathsf{PEFake}_S$, we conclude, by Fact 2.2, that $\mathsf{RERel}_S \overset{c}{\equiv} \mathsf{REFake}_S$ and the scheme is LIN RKA secure. $\qquad\square$

5.2. *Achieving KDM Security & RKA Insecurity*

The second scheme $(\mathsf{KE}, \mathsf{KD})$ is obtained by modifying the LPN construction $(\mathsf{PE}, \mathsf{PD})$ as follows. The key $S \in \{0, 1\}^k$ is augmented with an index $i \in \{1, \ldots, k\}$. A plaintext M will be encrypted by the triple $(\mathsf{PE}_S(M), i, S_i)$, i.e., in addition to the ciphertext $\mathsf{PE}_S(M)$, we leak a single bit of the key S_i whose location i is determined by another (public) part of the key. Formally,

$$\mathsf{KE}_{S,i}(M) := (\mathsf{PE}_S(M), i, S_i), \qquad \mathsf{KD}_S(C_1, C_2, C_3) := \mathsf{PD}_S(C_1)$$

Below we show that the scheme is LIN KDM secure. In fact, it will be useful to prove KDM security with respect to a slightly richer family of "extended linear functions" which contains functions of the form $\psi_{M,T} : S \to M + TS$ for every $M \in \mathbb{F}_2^\ell$ and matrix $T \in \mathbb{F}_2^{\ell \times k}$.

Lemma 5.2. *Under the* LPN *assumption, the scheme* $(\mathsf{KE}, \mathsf{KD})$ *is secure against extended linear key-dependent message attacks.*

Proof. Recall that in an extended LIN KDM attack on an encryption algorithm E, the adversary makes queries of the form (M, T) and attempts to distinguish between the real

oracle $E\mathsf{Real}_S$ which returns $E_S(M + TS)$ and the fake oracle $E\mathsf{Fake}_S$ which returns $E_S(0^{|M|})$. Our goal is to show that the scheme $\mathsf{KE}_{S,i}$ is LIN KDM secure. Formally, we should support functions which map the combined key $(S \circ i) \in \{0, 1\}^{k + \lceil \log(k) \rceil}$ (viewed as a single long vector) into messages of the form $M + T \cdot (S \circ i)$, where $M \in \mathbb{F}_2^{\ell}$ and $T \in \mathbb{F}_2^{\ell \times k + \lceil \log(k) \rceil}$, and (by abuse of notation) we identify the index $i \in [k]$ with its canonical representation as a string of length $\lceil \log(k) \rceil$. Observe that since i is public, any linear function in $(S \circ i)$ can be efficiently translated into a linear function in S of the form $M' + T'S$ where $M' \in \mathbb{F}_2^{\ell}$ and $T' \in \mathbb{F}_2^{\ell \times k}$, and so it suffices to focus on such functions.

We will essentially reduce the extended LIN KDM security of $\mathsf{KE}_{S,i}$ with $S \xleftarrow{R} \{0, 1\}^k$, $i \xleftarrow{R} [k]$ to the security of $\mathsf{PE}_{S'}$ with 1-bit shorter key $S' \xleftarrow{R} \{0, 1\}^{k-1}$. The extended LIN KDM security of the latter is proven in [2, Thm. 8]. The reduction uses a sequence of hybrids.

For an index $i \in [k]$ and a bit $\sigma \in \{0, 1\}$, we define an oracle-aided randomized function $\mathcal{A}_{i,\sigma}^{(\cdot)}$ as follows. Given a KDM query $(M, T) \in \mathbb{F}_2^{\ell} \times \mathbb{F}_2^{\ell \times k}$, the algorithm $\mathcal{A}_{i,\sigma}$ does the following: (1) defines the matrix $T_{-i} \in \mathbb{F}_2^{\ell \times k-1}$ by removing the i-th column T_i of T; (2) queries its oracle with (M, T_{-i}) and obtains a ciphertext $(A' \in \mathbb{F}_2^{t \times k-1}, Z' \in \mathbb{F}_2^t)$; (3) samples a random column $a_i \in \mathbb{F}_2^t$ and outputs the matrix $A = (A'_{[1:i-1]}|a_i|A'_{[i:k-1]})$, the vector $Z = Z' + a_i \cdot \sigma + G \cdot T_i \cdot \sigma$ and the pair (i, σ). (Recall that G is the generating matrix of the error correcting code used in the LPN construction.) We claim that

$$\mathsf{KEReal}_{S,i} \equiv \mathcal{A}_{i,\sigma}^{\mathsf{PEReal}_{S'}} \qquad \text{whenever } S = (S'_{[1:i-1]}, \sigma, S'_{[i:k-1]}). \tag{10}$$

Indeed, assume that $\mathcal{A}_{i,\sigma}$ has an oracle access to $\mathsf{PEReal}_{S'}$. Then, on a query (M, T_{-i}), the oracle responds with a fresh ciphertext

$$(A' \xleftarrow{R} \mathbb{F}_2^{t \times k-1}, Z' = AS' + E + G(M + T_{i-1}S')),$$

where $E \xleftarrow{R} \mathsf{Ber}_{\varepsilon,\delta}^t$ is a fresh noise vector. By linearity, it follows that the modified ciphertext $(A, Z, (i, \sigma))$ computed by $\mathcal{A}_{i,\sigma}$ satisfies

$$Z = A'S' + E + G(M + T_{-i}S') + a_i \cdot \sigma + G \cdot T_i \cdot \sigma = AS + E + G(M + TS).$$

Since a_i is chosen at random, $(A, Z, (i, \sigma))$ is a fresh sample from $\mathsf{PE}_S(M + TS)$ as required.

We now claim that, for randomly chosen $S' \xleftarrow{R} \{0, 1\}^{k-1}$ and every $i \in [k]$, $\sigma \in \{0, 1\}$,

$$\mathcal{A}_{i,\sigma}^{\mathsf{PEReal}_{S'}} \stackrel{c}{\equiv} \mathcal{A}_{i,\sigma}^{\mathsf{PEFake}_{S'}} \stackrel{c}{\equiv} \mathcal{A}_{i,\sigma}^{\mathcal{R}_{S'}} \equiv (\mathcal{R}_S, i, \sigma), \tag{11}$$

where $\mathcal{R}_S$ is a randomized function which ignores the key S and the given input, and outputs a fresh uniformly chosen matrices $A \xleftarrow{R} \mathbb{F}_2^{t \times |S|}$ and $Z \xleftarrow{R} \mathbb{F}_2^t$ and the notation $(\mathcal{R}_S, i, \sigma)$ refers to the oracle which ignores its query and returns (A, Z, i, σ) where

$(A, Z) \overset{R}{\leftarrow} \mathcal{R}_S$. The first transition of Eq. 11 follows from the security of the parity-based encryption PE against (extended) LIN KDM attacks ([2, Thm. 8]), the second transition follows from the pseudorandomness of PE (Lemma 3.5), and the last transition follows by noting that if the oracle answers (A', Z') are uniform, then so are the converter outputs (A, Z).

Next, we define another converter $\mathcal{B}_{i,\sigma}$ which acts similarly to $\mathcal{A}$, except that it computes the vector Z by $Z' + a_i \cdot \sigma$. We claim that, for randomly chosen $S' \overset{R}{\leftarrow} \{0, 1\}^{k-1}$ and every $i \in [k], \sigma \in \{0, 1\}$,

$$(\mathcal{R}_S, i, \sigma) \equiv \mathcal{B}_{i,\sigma}^{\mathcal{R}_{S'}} \overset{c}{\equiv} \mathcal{B}_{i,\sigma}^{\mathsf{PEFake}_{S'}} \equiv \mathsf{KEFake}_{S,i}, \tag{12}$$

where $S = (S'_{[1:i-1]}, \sigma, S'_{[i:k-1]})$. Indeed, the first transition follows by noting that $\mathcal{B}$ maps the uniform samples to uniform samples, the second transition is due to the pseudorandomness of PE, and the last transition follows by noting that $\mathcal{B}$ maps an encryption of zero $(A' \overset{R}{\leftarrow} \mathbb{F}_2^{t \times k-1}, Z' = AS' + E)$ under PE'_S into a fresh encryption of zero $(A, Z = AS + E)$ under KE_S. The lemma now follows by combining Eq. 10, 11, and 12 with Fact 2.2. $\qquad\square$

On the other hand, one can fully recover the key S via an RKA by shifting the index i through all possible indices in $\{1, \ldots, k\}$. Note that this attack is oblivious to the messages encrypted; in particular, all the attacker needs is the ability to obtain, for any choice of Δ, a ciphertext $\mathsf{KE}_{(S,i)+\Delta}(M)$ where the message M may be arbitrary and possibly unknown (e.g., chosen by the oracle).

5.3. Counter-Example: RKA+KDM $\not\Rightarrow$ RK-KDM

Our counter-example is defined via the following double-encryption:

$$\mathsf{Enc}_{S_1, S_2}(M) := \mathsf{KE}_{S_2}(\mathsf{RE}_{S_1}(M)), \qquad \mathsf{Dec}_{S_1, S_2}(C) := \mathsf{RD}_{S_1}(\mathsf{KD}_{S_2}(C)),$$

where $S_1 \in \{0, 1\}^k$ and S_2 is the concatenation of a vector $S'_2 \in \{0, 1\}^k$ and an index $i \in \{1, \ldots, k\}$.

Lemma 5.3. *Under the LPN assumption, the scheme (Enc, Dec) satisfies the followings:*

1. *Security under linear related-key attacks.*
2. *Security under linear key-dependent message attacks.*
3. *The secret key can be fully recovered via a LIN RK-KDM attack.*

Proof. (1) We show that any double-encryption Enc, whose inner encryption RE is LIN RKA secure, is also LIN RKA secure. For an encryption E, let $E\mathsf{Real}_S$ and $E\mathsf{Fake}_S$ be the real/fake RKA oracles as defined in Lemma 5.1. We define an oracle-aided randomized function $\mathcal{A}_{S_2}^{\mathcal{O}}$ as follows: Given a LIN RKA query with shift vector $\Delta = (\Delta_1, \Delta_2)$ and message M, the function $\mathcal{A}_{S_2}^{\mathcal{O}}$ outputs a sample from $\mathsf{KE}_{S_2+\Delta_2}(\mathcal{O}(\Delta_1, M))$. It fol-

lows that, for random S_1 and every S_2,

$$\mathsf{EncReal}_{S_1,S_2} \equiv \mathcal{A}_{S_2}^{\mathsf{REReal}_{S_1}} \stackrel{\mathrm{c}}{\equiv} \mathcal{A}_{S_2}^{\mathsf{REFake}_{S_1}} \equiv \mathsf{EncFake}_{S_1,S_2},$$

where the first and third transitions follow from the definition of $\mathcal{A}$ and the second transition is due to the LIN RKA security of RE.

(2) We will need the following observation, which follows from the linear structure of the LPN-based encryption PE. For every key S_1 and internal randomness r, the inner encryption $\mathsf{RE}_{S_1}(X; r)$ can be written as an (extended) linear mapping $\psi_{M,T} : X \to M + TX$ where M and T can be computed based on S_1 and r via some efficiently computable mapping ρ. Using this observation, we show that the double-encryption Enc inherits (extended) LIN KDM security from the outer encryption KE.

Formally, let $E\mathsf{Real}_S$ and $E\mathsf{Fake}_S$ be the real/fake KDM oracles for an encryption E defined in Lemma 5.2. Let $\mathcal{A}_{S_1}^{\mathcal{O}}$ be an oracle-aided randomized function, which, given an extended LIN KDM query $\psi_{M,T}$, samples randomness r for the inner encryption RE, computes $(M', T') = \rho(S_1, r)$, and queries the oracle $\mathcal{O}$ with the composed linear function $\psi : S_2 \to \psi_{M',T'}(\psi_{M,T}(S_1, S_2))$. It is not hard to see that ψ is indeed an extended linear function, and for random (S_1, S_2)

$$\mathsf{EncReal}_{S_1,S_2} \equiv \mathcal{A}_{S_1}^{\mathsf{KEReal}_{S_2}} \stackrel{\mathrm{c}}{\equiv} \mathcal{A}_{S_1}^{\mathsf{KEFake}_{S_2}},$$

where the first transition is due to the definition of $\mathcal{A}$ (and holds for every (S_1, S_2)) and the second transition follows from the security of KE.

To complete the proof, define an oracle-aided randomized function $\mathcal{B}_{S_2}^{\mathcal{O}}$, which given a LIN KDM query $\psi_{M,T}$ outputs $\mathcal{O}(\mathsf{Enc}_{S_2}(0^{|M|}))$. For random (S_1, S_2), we have that

$$\mathcal{A}_{S_1}^{\mathsf{KEFake}_{S_2}} \equiv \mathcal{B}_{S_1}^{\mathsf{KEFake}_{S_2}} \stackrel{\mathrm{c}}{\equiv} \mathcal{B}_{S_1}^{\mathsf{KEReal}_{S_2}} \equiv \mathsf{EncFake}_{S_1,S_2},$$

and item (2) follows.

(3) We show that, given an access to the real LIN RK-KDM oracle $\mathsf{EncReal}_{S_1,S_2}$, it is possible to fully recover the key (S_1, S_2). First, use RKA queries to fully recover the key S_2 via the attack described in Sect. 5.2. Second, in order to recover S_1, apply a KDM query to obtain an encryption C of (S_1, S_2) and use the decryption algorithm KD_{S_2} to decrypt the ciphertext C. We claim that the resulting value is simply (S_1, S_2). Indeed, by the definition of RE, we have that

$$C = \mathsf{Enc}_{S_1,S_2}(S_1, S_2) = \mathsf{KE}_{S_2}(\mathsf{RE}_{S_1}(S_1, S_2)) = \mathsf{KE}_{S_2}(S_1, S_2)$$

and therefore $\mathsf{KD}_{S_2}(C) = (S_1, S_2)$ and the lemma follows. $\square$

6. Conclusion

We defined a new combined form of RK-KDM security, proved that such an encryption scheme can be realized based on the LPN assumption and showed that the free-XOR

technique can be securely instantiated with it. Altogether, our results enable a realization of the free-XOR optimization in the standard model under a well-studied cryptographic assumption.

The new definition of RK-KDM security further motivates the study of security under related-key and key-dependent attacks. Specifically, in light of our counter-example, it is natural to ask whether LIN RK-KDM security can be constructed based on some combination of an RKA secure scheme and a KDM secure scheme, or better yet, based on more general assumptions (e.g., CPA-secure encryption scheme). It will also be interesting to find additional applications of RKA/KDM secure primitives.

Acknowledgements

Appendix: Alternative Definition to Computationally Private Randomized Encodings

It might be natural to define the computational encoding of a function family $f = \{f_n : \{0, 1\}^n \to \{0, 1\}^{\ell(n)}\}_{n \in \mathbb{N}}$ to be a function family indexed with two parameters n and k where k is a security parameter. In this definition, every function f_n is represented by a function family $\hat{f}_{n,k}$. Computational privacy requires that for every n, the randomized functions $\mathsf{Sim}(1^n, 1^k, f_n(\cdot))$ and $\hat{f}_{n,k}(\cdot)$ (indexed by k) are computationally indistinguishable. In this case, the parameters m, s, as well as the running time of the simulator, the recovery algorithm, and the time it takes to compute $\hat{f}_{n,k}$ are all functions of the input length n and the security parameter k. This definition looks more appealing than Definition 4.1, since it allows to tune the privacy of a specific function with fixed input length and enables a closer security analysis. However, this definition is meaningless, since we can set a null encoding for small k's and use a perfectly private, perfectly correct encoding for $k = \exp(n)$. Such an encoding satisfies computational privacy (the distribution ensembles are computationally indistinguishable for large enough k's) and is also uniform since it can be computed in time $\mathrm{poly}(n, k)$. This paradox can be prevented by demanding that n, k are polynomially related. Since in this paper we focus on high-level security and do not care about security in concrete terms (i.e., the specific running time vs. success probability rate), we avoid such complications and use the simple single-parameter definition given above. However, Definition 4.1 also enables to tune the privacy of a specific function f_n by arbitrarily augmenting f_n into an infinite family of functions (this can be done by simply padding the input). We also stress that our construction uses a security parameter (which we set to $n^{O(1)}$ to satisfy the single-parameter definition) and thus allows a closer analysis and a more natural way to tune the privacy of a specific function (see Sect. 4).

Bibliographic Note The above remark was written originally for [4] and was eventually omitted. We thank Yuval Ishai and Eyal Kushilevitz for allowing us to include it here.

References

[1] B. Applebaum, Randomly encoding functions: A new cryptographic paradigm (invited talk), in *ICITS*, pp. 25–31 (2011)

[2] B. Applebaum, D. Cash, C. Peikert, A. Sahai, Fast cryptographic primitives and circular-secure encryption based on hard learning problems, in *Advances in Cryptology—CRYPTO 2009*, pp. 595–618 (2009)

[3] B. Applebaum, D. Harnik, Y. Ishai, Semantic security under related-key attacks and applications, in *ICS*, pp. 45–60 (2011)

[4] B. Applebaum, Y. Ishai, E. Kushilevitz, Computationally private randomizing polynomials and their applications. *Comput. Complex.* **15**(2), 115–162 (2006)

[5] B. Applebaum, Y. Ishai, E. Kushilevitz, How to garble arithmetic circuits, in *Proc. 52nd FOCS*, pp. 120–129 (2011)

[6] M. Bellare, D. Cash, Pseudorandom functions and permutations provably secure against related-key attacks, in *Advances in Cryptology—CRYPTO 2010*, pp. 666–684 (2010)

[7] M. Bellare, V.T. Hoang, P. Rogaway, Foundations of garbled circuits, in *CCS '12*, pp. 784–796 (2012)

[8] M. Bellare, T. Kohno, A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications, in *Advances in Cryptology—EUROCRYPT 2003*, pp. 491–506 (2003)

[9] M. Bellare, P. Rogaway, Random oracles are practical: A paradigm for designing efficient protocols, in *First ACM Conference on Computer and Communications Security*, pp. 62–73 (1993)

[10] J. Black, P. Rogaway, T. Shrimpton, Encryption-scheme security in the presence of key-dependent messages, in *SAC '02*, pp. 62–75 (2002)

[11] A. Blum, M. Furst, M. Kearns, R.J. Lipton, Cryptographic primitives based on hard learning problems, in *Advances in Cryptology—CRYPTO 1993*, pp. 278–291 (1993)

[12] A. Blum, A. Kalai, H. Wasserman, Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM* **50**(4), 506–519 (2003)

[13] M. Blum, S. Micali, How to generate cryptographically strong sequences of pseudo-random bits. *SIAM J. Comput.* **13**, 850–864 (1984)

[14] D. Boneh, S. Halevi, M. Hamburg, R. Ostrovsky, Circular-secure encryption from decision diffie-hellman, in *Advances in Cryptology—CRYPTO 2008*, pp. 108–125 (2008)

[15] F. Böhl, G.T. Davies, D. Hofheinz, Encryption schemes secure under related-key and key-dependent message attacks, in *Public Key Cryptography*, pp. 483–500 (2014)

[16] J. Camenisch, A. Lysyanskaya, An efficient system for non-transferable anonymous credentials with optional anonymity revocation, in *Advances in Cryptology—EUROCRYPT 2001*, pp. 93–118 (2001)

[17] R. Canetti, O. Goldreich, S. Halevi, The random oracle methodology, revisited. *J. ACM* **51**(4), 557–594 (2004)

[18] S.G. Choi, J. Katz, R. Kumaresan, H.S. Zhou, On the security of the "free-XOR" technique, in *TCC '12*, pp. 39–53 (2012)

[19] H. Gilbert, M.J.B. Robshaw, Y. Seurin, How to encrypt with the LPN problem. in *Automata, Languages and Programming, 35th International Colloquium, ICALP '08*, pp. 679–690 (2008)

[20] O. Goldreich, H. Krawczyk, M. Luby, On the existence of pseudorandom generators. *SIAM J. Comput.* **22**(6), 1163–1175 (1993)

[21] O. Goldreich, S. Micali, A. Wigderson, How to play ANY mental game, in *Proc. 19th STOC*, pp. 218–229 (1987)

[22] J. Håstad, R. Impagliazzo, L.A. Levin, M. Luby, A pseudorandom generator from any one-way function. *SIAM J. Comput.* **28**(4), 1364–1396 (1999)

[23] W. Henecka, S. Kögl, A.R. Sadeghi, T. Schneider, I. Wehrenberg, TASTY: Tool for automating secure two-party computations, in *CCS 10'*, pp. 451–462 (2010)

[24] Y. Huang, D. Evans, J. Katz, L. Malka, Faster secure two-party computation using garbled circuits, in *USENIX Security Symposium*, pp. 539–554 (2011).

[25] Y. Huang, C.H Shen, D. Evans, J. Katz, A. Shelat, Efficient secure computation with garbled circuits, in *ICISS '11*, pp. 28–48 (2011)

[26] Y. Ishai, J. Kilian, K. Nissim, E. Petrank, Extending oblivious transfers efficiently, in *Advances in Cryptology—CRYPTO 2003*, pp. 145–161 (2003)

[27] Y. Ishai, E. Kushilevitz, Randomizing polynomials: A new representation with applications to round-efficient secure computation, in *Proc. 41st FOCS*, pp. 294–304 (2000)

[28] V. Kolesnikov, A.R. Sadeghi, T. Schneider, Improved garbled circuit building blocks and applications to auctions and computing minima, in *CANS*, pp. 1–20 (2009)

[29] V. Kolesnikov, T. Schneider, Improved garbled circuit: Free XOR gates and applications, in *Automata, Languages and Programming, 35th International Colloquium, ICALP '08*, pp. 486–498 (2008)

[30] B. Kreuter, A. Shelat, C.H. Shen, Billion-gate secure computation with malicious adversaries, in *Security '12: Proceedings of the 21st USENIX conference on Security Symposium*, pp. 14–14 (2012)

[31] Y. Lindell, B. Pinkas, N. Smart, Implementing two-party computation efficiently with security against malicious adversaries, in *SCN '08*, pp. 2–20 (Sep 2008)

[32] Y. Lindell, B. Pinkas, An efficient protocol for secure two-party computation in the presence of malicious adversaries, in *Advances in Cryptology—EUROCRYPT 2007*, pp. 52–78 (2007)

[33] Y. Lindell, B. Pinkas, A proof of security of yao's protocol for two-party computation. *J. Cryptol.* **22**(2), 161–188 (2009)

[34] L. Malka, J. Katz, Vmcrypt—Modular software architecture for scalable secure computation, in *CCS '11*, pp. 715–724 (2011)

[35] D. Malkhi, N. Nisan, B. Pinkas, Y. Sella, Fairplay—A secure two-party computation system, in *Proc. of 13th USENIX Security Symposium*, pp. 287–302 (2004)

[36] U.M. Maurer, Indistinguishability of random systems, in *Advances in Cryptology—EUROCRYPT 2002*, pp. 110–132 (2002)

[37] M. Naor, B. Pinkas, Oblivious transfer with adaptive queries, in *Advances in Cryptology—CRYPTO 1999*, pp. 573 –590 (1999)

[38] M. Naor, B. Pinkas, R. Sumner, Privacy preserving auctions and mechanism design, in *Proc. 1st ACM Conference on Electronic Commerce*, pp. 129–139 (1999)

[39] J.B. Nielsen, C. Orlandi, LEGO for two-party secure computation, in *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009*, pp. 368–386 (2009)

[40] B. Pinkas, T. Schneider, N. Smart, S. Williams, Secure two-party computation is practical, in *Advances in Cryptology—ASIACRYPT 2009*, pp. 250–267 (2009)

[41] P. Rogaway, The Round Complexity of Secure Protocols. Ph.D. thesis, MIT (June 1991)

[42] A. Shelat, C.H. Shen, Two-output secure computation with malicious adversaries, in *Advances in Cryptology—EUROCRYPT 2011*, pp. 386–405 (2011)

[43] D.A. Spielman, Linear-time encodable and decodable error-correcting codes, in *Proc. 27th STOC*, pp. 388–397 (1995)

[44] A.C. Yao, Theory and application of trapdoor functions, in *Proc. 23rd FOCS*, pp. 80–91 (1982)

[45] A.C. Yao, How to generate and exchange secrets, in *Proc. 27th FOCS*, pp. 162–167 (1986)