Identity-Based Format-Preserving Encryption

Mihir Bellare*
University of California San Diego
La Jolla, California
mihir@eng.ucsd.edu

Viet Tung Hoang†
Florida State University
Tallahassee, Florida
tvhoang@cs.fsu.edu

ABSTRACT

We introduce identity-based format-preserving encryption (IB-FPE) as a way to localize and limit the damage to format-preserving encryption (FPE) from key exposure. We give definitions, relations between them, generic attacks and two transforms of FPE schemes to IB-FPE schemes. As a special case, we introduce and cover identity-based tweakable blockciphers. We apply all this to analyze DFF, an FPE scheme proposed to NIST for standardization.

1 INTRODUCTION

Schemes for format-preserving encryption (FPE) have been standardized [19] and are in widespread use for the encryption of credit-card numbers. Towards limiting the damage from key exposure, we introduce identity-based FPE (IB-FPE). We provide a provable-security treatment involving definitions, attacks and two design paradigms. We apply this to analyze DFF [36], an FPE scheme proposed to NIST for standardization.

FPE. Format-preserving encryption (FPE) originates with [10, 13]. An FPE scheme F specifies a deterministic encryption function F.E : $\{0,1\}^{\mathsf{F.kl}} \times \mathsf{F.TS} \times \mathsf{F.Dom} \to \mathsf{F.Dom}$ that takes a F.kl-bit key J, a tweak T and a message X to return a ciphertext $Y = \mathsf{F.E}(J,T,X)$. There is a corresponding decryption function F.D : $\{0,1\}^{\mathsf{F.kl}} \times \mathsf{F.TS} \times \mathsf{F.Dom} \to \mathsf{F.Dom}$ such that the maps $\mathsf{F.E}(J,T,\cdot)$, $\mathsf{F.D}(J,T,\cdot)$ are permutations over F.Dom that are inverses of each other. What makes FPE special is that the domain F.Dom can be arbitrary and in particular very small. Some examples are $\mathsf{F.Dom} = \{0,1\}^8$ —encrypt a byte so that the ciphertext is also a byte— $\mathsf{F.Dom} = \mathbb{Z}_{10}^4$ —encrypt a 4 digit PIN so that the ciphertext is also four decimal digits— $\mathsf{F.Dom} = \mathbb{Z}_{10}^{16}$ —encrypt a 16-digit credit-card number so that the result is also a 16-digit credit-card number. FPE is motivated by legacy constraints which in many systems mandate that the ciphertext replace the plaintext, and must thus have the same "format" as the plaintext. Tweakable blockciphers [27] are the special case where $\mathsf{F.Dom} = \{0,1\}^{\mathsf{F.bl}}$ for some integer F.bl called the block length.

The canonical metric of security for an FPE scheme F is prp security [10, 26]. The game picks a challenge bit b and random key $J \in \{0,1\}^{\mathsf{F.kl}}$. For each tweak T it also lets $\Pi(T,\cdot)$ be a random permutation over F.Dom. The adversary $\mathcal{A}$ can ask for encryption under a tweak T and message X of its choice, being returned $\mathsf{F}(J,T,X)$ if $b=1$ or $\Pi(T,X)$ if $b=0$, and similarly for decryption.

FPE is not easy to build. Today the most practical approach is Feistel with strong —AES-based— round functions and a number of rounds $r \geq 8$. NIST SP 800-38G [19] standardizes two such schemes, FF1 ($r = 10$) and FF3 ($r = 8$). Recent attacks [7, 17] suggest that it would be good to increase the number of rounds when the inputs are very short, but this is largely orthogonal to our work.

Corporations offering FPE-based products include HPE Voltage, Verifone, Protegrity, Ingenico, Thales/Vormetric and Gemalto. Tens of millions of credit-cards have been encrypted with these products.

IB-FPE. We define an identity-based FPE (IB-FPE) scheme as a pair (F, KDF) consisting of a (base) FPE scheme F and an associated *key-derivation function* KDF. The latter takes a master key K and identity I to (deterministically) return a key $J = \mathsf{KDF}(K,I) \in \{0,1\}^{\mathsf{F.kl}}$ for I to use with F.

In the traditional usage of an FPE scheme F, an organization would have a single key K for F stored at many different devices (for example, point-of-sale terminals) that each encrypts directly under K. But each device is at some risk of compromise due to physical, insider or side-channel attacks. Compromise of even one device (which could be quite likely) then has the global consequence of exposure of K. IB-FPE allows us to localize, and thus limit, the damage from key exposure. With IB-FPE, we can associate an identity I to a device and delegate to it the derived key $J_I = \mathsf{KDF}(K,I)$, allowing the device to (effectively) encrypt under K without actually having K. (The master key K would be stored in a secure location, for example in secure hardware.) Compromise of device I would now have only local consequences, encryptions under J_I being compromised but (for an IB-FPE scheme meeting the definitions we will give) encryption under other identities remaining secure.

Another benefit of IB-FPE is to increase the lifetime of the key K. In practice it is recommended to limit the number of encryptions under a particular key, changing (rotating) the key periodically. With u identities each performing q encryptions, direct encryption with a traditional FPE scheme would result in uq encryptions under the base key. With IB-FPE, we have u key derivations under the master key and only q encryptions under each of u different

*Supported in part by NSF grants CNS-1526801 and CNS-1717640, ERC Project ERCC FP7/615074 and a gift from Microsoft corporation.
†Supported in part by NSF grant CICI-1738912 and First Year Assistant Professor Award of Florida State University.

CCS '17, October 30-November 3, 2017, Dallas, TX, USA

ACM ISBN 978-1-4503-4946-8/17/10…$15.00
https://doi.org/10.1145/3133956.3133995

derived keys. This structure can significantly increase the number uq of encryptions that can be safely performed [1, 34].

IB-FPE SECURITY. Let (F, KDF) be an IB-FPE scheme. We give a prp style definition of security called ib-prp. We also give two key-recovery security definitions called ib-kr-ai and ib-kr-ti. We show relations between them, summarized in Fig. 4.

While natural, these definitions are strong, in particular allowing selective opening attacks [5, 8, 18, 22] that make them hard to provably achieve. We also define non-adaptive versions, which continue to relate to each other as per Fig. 4, and which our schemes are shown to achieve.

THE DEFINITIONS. The ib-prp game picks a random challenge bit b and random master key K, and associates key $J_I = \mathsf{KDF}(K, I) \in \{0,1\}^{\mathsf{F.kl}}$ to identity I. The adversary gets oracle Enc taking identity I, tweak T and message X, and oracle Dec taking I, T and ciphertext Y. Initially, they respond with $\mathsf{F.E}(J_I, T, X)$ and $\mathsf{F.D}(J_I, T, Y)$, respectively. At any point, the adversary can either expose the key of I, querying $\text{Exp}(I)$ to get J_I, or switch I to challenge mode by querying $\text{Ch}(I)$, restricted, of course to not being able to do both for the same I. If I is switched to challenge mode, oracles Enc, Dec change in the $b = 0$ case, with $\text{Enc}(I, T, \cdot)$ and $\text{Dec}(I, T, \cdot)$ now becoming permutations that are random but consistent with prior replies.

Theoretical work has traditionally formalized only strong goals that represent the most desirable targets for security proofs, ib-prp in our case. But we also formalize weaker key-recovery security goals (ib-kr-ai and ib-kr-ti). Oracles in the games are like in the $b = 1$ case of ib-prp. The adversary returns a key J' and identity I'. In the ib-kr-ti (target identity) case, it wins if $J' = J_{I'}$ is the key for the identity it names, while in the ib-kr-ai (any identity) case, I' is ignored and it wins if $J' = J_I$ for any un-exposed challenge identity I. The motivation is that (1) We are interested not just in security proofs but in attacks, for which we want to make claims that are strong (violating ib-kr-ai or ib-kr-ti is much more damaging than violating ib-prp) as well as precise (which requires that key-recovery advantages be formalized), and (2) We might be able to prove better security (in terms of bounds on adversary advantage) for ib-kr-ai or ib-kr-ti than for ib-prp.

So far adversaries are adaptive in the sense that they can query Enc, Dec with I before deciding to expose I. We say that an adversary (whether ib-prp, ib-kr-ai or ib-kr-ti) is non-adaptive if its exposure decision for I does not depend on seeing encryptions or decryptions under I: if it queries $\text{Exp}(I)$, it has not previously queried $\text{Enc}(I, \cdot, \cdot)$ or $\text{Dec}(I, \cdot, \cdot)$.

Security in the face of exposure queries captures the above-mentioned application goal that the damage from compromise is local rather than global. (Encryption for an identity is secure even if the keys of other identities are known to the attacker.) Exposure is thus a central element of the framework, and is a powerful adversary capability even in the non-adaptive case. The definition adapts the classical one for IBE [14], differences being that our setting is symmetric (there is no public master key), encryption is deterministic, the goal is prp style security (rather than semantic security) and there are multiple challenge identities, not just one. In the adaptive case, the combination of these elements allows a selective opening attack [5, 8, 18, 22]. We stress that non-adaptive security, even if weaker than adaptive, is hardly a weak notion, and seems more than adequate for practice.

RELATIONS. It is clear that ib-kr-ai security (tightly) implies ib-kr-ti security. (If you can find the key for an identity you name, you can find a key for some identity.) Proposition 3.2 says that, conversely, ib-kr-ti *tightly* implies ib-kr-ai, because, given a candidate key, one can (under some conditions) test to see which identity it matches. We would expect that ib-prp implies ib-kr-ti (and thus, by the above, ib-kr-ai), and while Theorem 3.3, at the highest level, validates this, the truth it shows is more delicate. The difficulty is that in FPE the domain size can be small, and the reduction is parameterized to adjust. The relations, summarized in Fig. 4, hold in both the adaptive and non-adaptive cases.

ATTACKS. We give attacks on the security of *any* IB-FPE scheme (F, KDF), showing inherent limitations in achievable security. The attacks are strong (they violate non-adaptive ib-kr-ai, not just ib-prp) and rigorously analyzed (Theorems 4.1 and 4.2 provide and prove precise lower bounds on adversary advantage). Their implication is that for (F, KDF) to have k-bits of (even non-adaptive, ib-kr-ai) security, FPE scheme F must have $2k$-bit keys, regardless of the length of the master key and the choice of KDF. We call this the *double-key condition.*

The challenge with the attacks is to cover *all* IB-FPE schemes (F, KDF). We give two attacks, calling the first the *matching attack* and the second, which generalizes DP [20], the *exhaustive search attack*. Depending on the value of a quantity we define, called the *diversity* of the key-distribution function KDF, we are able to show that one or the other attack always has constant non-adaptive ib-kr-ai advantage with effort around $2^{\mathsf{F.kl}/2}$.

BUILDING IB-FPE SCHEMES. We now turn to constructing IB-FPE schemes that do as well as possible subject to the limitations uncovered by our attacks. Given that FPE schemes F (satisfying standard prp security) are hard to build, we want to leverage existing constructions of them. Accordingly, our approach is modular: taking as given a (base) FPE scheme F, we design key-derivation functions KDF for it and prove non-adaptive ib-prp security of (F, KDF) assuming the prp security of F and also possibly assuming something about KDF. We aim to make the master key of KDF as short as we can and to make KDF as efficient as we can. We also aim for instantiations of our key-derivation functions that use only a blockcipher, and moreover one that (like AES) has the same key and block length. (This is because practical FPE schemes already use such blockciphers, as Feistel round functions.) Below we first give a natural, standard-model key-derivation construction **PRF**. Then, to improve efficiency and get an analysis of DFF, we give and analyze an ideal-cipher model construction **Dbl**.

THE **PRF** CONSTRUCTION. We show in Section 5 that PRFs make good key-derivation functions: If $\mathsf{KDF} : \mathsf{KDF.MKS} \times \mathsf{KDF.IS} \to \{0,1\}^{\mathsf{F.kl}}$ is a PRF and base FPE scheme F is prp secure then IB-FPE scheme (F, KDF) is non-adaptive ib-prp secure. We call this the **PRF** construction of an IB-FPE scheme. Assuming $\mathsf{F.kl} = 2k$, the concrete reduction, as given by Theorem 5.1, implies that if KDF has k-bits of prf security and F has $2k$-bits of prp security then (F, KDF) has k bits of non-adaptive ib-prp security. Our attacks discussed above imply that the reduction is optimal.

For an instantiation we would like to base KDF solely on AES and achieve full 128-bit security with the master key being a (128-bit) AES key. Abstractly, assuming given a base FPE scheme F that has $2k$ bits of prp security with $\mathsf{F.kl} = 2k$, this means that we want to build $\mathsf{KDF} : \{0,1\}^k \times \mathsf{KDF.IS} \to \{0,1\}^{2k}$, with k bits of prf security, solely from a blockcipher $E : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}^k$ having k bits of prp-cpa security. This is a challenging goal, but we can reach it via DHT's new analysis [15] of the XOR prp-to-prf transform of BKR [9]. Our key-derivation function, shown in Fig. 8, has a computational cost of four invocations of the blockcipher E.

In summary, the **PRF** construction instantiated as above is an efficient way to generically turn an FPE scheme into an IB-FPE scheme with optimal security, a standard-model proof and a reasonable key-derivation cost of four blockcipher invocations. There are two motivations for the alternative key-derivation method that follows: (1) Our results about it will eventually yield an analysis of the DFF scheme proposed to NIST for standardization, and (2) It uses only two blockcipher invocations.

THE **Dbl** CONSTRUCTION. Letting F be the given prp-secure FPE scheme with $\mathsf{F.kl} = 2k$, our **Dbl** ("Double") construction of an IB-FPE scheme (F, KDF) lets $E : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}^k$ be a blockcipher and then defines key-derivation function $\mathsf{KDF} : \{0,1\}^k \times \mathsf{KDF.IS} \to \{0,1\}^{2k}$ by

$$\mathsf{KDF}(K, I) = E(K, \mathsf{M}_0(I)) \,\|\, E(K, \mathsf{M}_1(I)) , \qquad (1)$$

where $\mathsf{M}_0, \mathsf{M}_1 : \mathsf{KDF.IS} \to \{0,1\}^k$ are injective functions with disjoint ranges. We refer to M as an embedding scheme, and it parameterizes the construction. Theorem 6.1 implies that (F, KDF) has k-bits of ib-prp security assuming F has $2k$-bits of prp security and E is an ideal cipher. The double-key condition emanating from our attacks says that the analysis of Theorem 6.1 is optimal. Next we discuss some technical elements of the result.

One might have hoped to establish prf security of KDF in the ideal-cipher model and then apply our result about **PRF**, but, even in the ideal-cipher model, the key-derivation function KDF of Eq. (1) has only $k/2$ bits of prf security. Instead we give a direct analysis.

In practice we expect that $E = \mathsf{AES}$ will be used, not only by KDF, but also by F. To model this, we allow F to have oracle access to the *same* ideal cipher E that is used by KDF. This common use of the ideal primitive precludes a modular proof and makes the analysis more challenging. Given an ib-prp adversary $\mathcal{A}$ against F under KDF, the reduction aims to build a prp adversary $\overline{\mathcal{A}}$ and bound ϵ, the ib-prp advantage of $\mathcal{A}$ against F, KDF, as a function of $\overline{\epsilon}$, the prp advantage of $\overline{\mathcal{A}}$ against F. The natural approach is a hybrid argument. The difficulty is that, due to the structure of KDF, keys of different users are not statistically independent. If u is the number of users invoked by $\mathcal{A}$, the straightforward hybrid argument would incur a loss of $O(u/2^k)$ per hybrid step, resulting in a bound of the form $\epsilon \le u\overline{\epsilon} + \delta$ where $\delta = O(u^2/2^k)$. This would imply only $k/2$ bits of security for F under KDF, well short of what we want and believe to be true. Theorem 6.1 gives a different proof that includes a more sophisticated hybrid argument to obtain $\delta = O(u/2^k)$, which implies k-bit ib-prp security for (F, KDF), as desired.

IB-FPE FROM PRE-MASKING FPE. **Dbl** builds an IB-FPE scheme (F, KDF) assuming as given the base FPE scheme $\mathsf{F} : \{0,1\}^{2k} \times \mathsf{F.TS} \times \mathsf{F.Dom} \to \mathsf{F.Dom}$. We now ask if the assumption can be dropped. That is, we want to build a practical F from our blockcipher $E : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}^k$ so that, with KDF as in Eq. (1), IB-FPE scheme (F, KDF) can be shown to have k bits of ib-prp security assuming nothing more than ideality of E. The difficulty is that practical FPE schemes F are mostly Feistel-based, and Feistel (as we explain further in Section 7) notoriously lacks tight analyses showing prp security for small domains and number of rounds. However we show that the goal can be reached if we target key-recovery security rather than prp security.

Our results are quite general. We define a class of FPE schemes that we call pre-masking. This class includes Feistel-based schemes. The schemes use a blockcipher $E : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}^k$ but have $2k$-bit keys. Encryption and decryption do not have direct access to the key but can call an oracle that uses the key in conjunction with the blockcipher in a restricted way. (See Section 7 for the full definition.) Now, take any F in this class and adjoin the key-derivation function KDF of **Dbl** as per Eq. (1) to get IB-FPE scheme (F, KDF). Then Theorem 7.3 establishes that (F, KDF) has k bits of ib-kr-ti security.

SECURITY OF DFF. Two FPE schemes proposed to NIST for standardization, namely FF2 [35] —not standardized due to the attack of [20]— and DFF [36] —still under consideration— derive a subkey from the tweak and then encrypt under the subkey with an un-tweaked cipher. The authors highlight this method as providing a feature they call delegation, where knowledge of the subkey for one tweak would not impact security of encryption under another tweak. Our IB-FPE framework allows us to formalize this claim and evaluate the security, relative to it, of DFF.

We view $\mathsf{FF2} = (\mathsf{F}_{\mathsf{ff2}}, \mathsf{KDF}_{\mathsf{ff2}})$ and $\mathsf{DFF} = (\mathsf{FF}_{\mathsf{dff}}, \mathsf{KDF}_{\mathsf{dff}})$ as IB-FPE schemes with identity space the tweak space of the original scheme, and a tweak space that is trivial, consisting, say, of just the empty string. In FF2, the master key K is 128 bits and the key delegated to I is $J_I = \mathsf{KDF}_{\mathsf{ff2}}(K, I) = \mathsf{AES}(K, I)$. Since $\mathsf{F}_{\mathsf{ff2}}$ (accordingly) has 128-bit keys, our attacks from Section 4 say that FF2 has at most 64 bits of ib-kr-ti security, explaining the Dworkin and Perlner (DP) attack [20]. Arguing that a scheme with a 128-bit (master) key should provide 128-bits of security, NIST rejected FF2. Seeking 128 bits of security, DFF continues to have a 128-bit master key, but derived keys (and thus keys for $\mathsf{FF}_{\mathsf{dff}}$) are 256 bits long. Our attacks indicate that the security is at most 128 bits. The question relevant to standardization is whether it actually is 128, or significantly less.

Let the domain be $\mathbb{Z}^n_{\mathsf{rdx}}$, the set of length n strings over alphabet $\mathbb{Z}_{\mathsf{rdx}}$ ($2 \le \mathsf{rdx}, n < 2^8$), and regard rdx, n as fixed. Let $k = 128$. Then the key-derivation function $\mathsf{KDF}_{\mathsf{dff}}$ of $\mathsf{DFF} = (\mathsf{FF}_{\mathsf{dff}}, \mathsf{KDF}_{\mathsf{dff}})$ can be viewed as obtained by applying our **Dbl** transform with $E = \mathsf{AES}$ and embedding scheme M defined by $\mathsf{M}_0(I) = [\mathsf{rdx}]^1 \| [|I|]^1 \| [n]^1 \| [I]^{13}$ and $\mathsf{M}_1(I) = [0]^3 \| [I]^{13}$, where $[x]^\ell$ denotes the encoding of x as an ℓ-byte string. Then (1) our results from Section 6 say that DFF has about 128 bits of ib-prp security if E is ideal and FPE scheme $\mathsf{FF}_{\mathsf{dff}}$ is assumed to have about 256 bits of prp security, and (2) Observing that $\mathsf{FF}_{\mathsf{dff}}$ is an E-based pre-masking FPE scheme, our results from Section 7 say that DFF has about 128 bits of ib-kr-ti security assuming only that E is ideal.

There is, however, a caveat that our analysis uncovers. For our results to apply, the functions $\mathsf{M}_0, \mathsf{M}_1$ defined above must be injective. This is not, strictly speaking, true for DFF, because the identity space is the set of all binary strings of length at most 13 bytes, and so, for example, $\mathsf{M}_1(001) = \mathsf{M}_1(01)$. It is true (our conditions on M are met) if we restrict identities further, for example to all have the same length, or so that no two represent, in binary, the same integer. For the general case we have neither a proof, nor an attack showing security to be significantly smaller than the desired 128 bits. In Section 8 we expand on this and also give the best attack we know for the general case. We would suggest that the embedding function used in DFF be changed to meet our conditions, so that our results would apply to validate security in the general case as well. For example, let identities be binary strings of at most 12 bytes, let $\mathsf{M}_0(I) = [0]^1 \| [\mathsf{rdx}]^1 \| [|I|]^1 \| [n]^1 \| [I]^{12}$ and $\mathsf{M}_1(I) = [1]^1 \| [\mathsf{rdx}]^1 \| [|I|]^1 \| [n]^1 \| [I]^{12}$.

We clarify that, as designed and presented in [35, 36], FF2 and DFF are FPE schemes targeting key delegation based on tweaks, not IB-FPE schemes. To translate findings above back to the original context, read "tweak" for "identity".

Discussion and related work. Identity-based cryptography was suggested by Shamir [33]. Identity-based encryption (IBE) was formalized and achieved by BF [14].

BHT [7] give message-recovery attacks on Feistel-based FPE schemes F, including the FF1 and FF3 standards [19] and $\mathsf{FF}_{\mathsf{dff}}$, in the case that the domain is tiny. DV [17] give small-domain attacks on FF3. FF1 and FF3 are not relevant for us. (Having 128 bit keys, they cannot, by our attacks, be base schemes for high-security IB-FPE.) For $\mathsf{F} = \mathsf{FF}_{\mathsf{dff}}$, the validity of Theorems 5.1 (for **PRF**) and 6.1 (for **Dbl**) is not affected, but to get the full possible k-bits of security for the IB-FPE scheme ($\mathsf{FF}_{\mathsf{dff}}$, KDF) from these results, one would have to increase the number of rounds in $\mathsf{FF}_{\mathsf{dff}}$ for tiny inputs. The BHT attacks do not contradict our proof of ib-kr-ti security of DFF because they are message-recovery attacks and do not succeed in key recovery.

Shuffle-based FPE schemes [21, 29, 31] are a possible choice in the role of F to obtain IB-FPE schemes via the **PRF** or **Dbl** constructions. For efficiency, however, schemes in practice, including $\mathsf{FF}_{\mathsf{dff}}$, have been Feistel based, so we have focused on the latter in considering instantiating F via pre-masking FPE schemes.

2 PRELIMINARIES

Notation and conventions. We let ε denote the empty string. If y is a string then $|y|$ denotes its length and $y[i]$ denotes its i-th bit for $1 \le i \le |y|$, and for $1 \le i \le j \le |y|$, let $y[i:j] = y[i] \cdots y[j]$. If X is a finite set, we let $x \leftarrow\$ X$ denote picking an element of X uniformly at random and assigning it to x. Algorithms may be randomized unless otherwise indicated. Running time is worst case. If A is an algorithm, we let $y \leftarrow A(x_1, \ldots; r)$ denote running A with random coins r on inputs $x_1, \ldots$ and assigning the output to y. We let $y \leftarrow\$ A(x_1, \ldots)$ be the result of picking r at random and letting $y \leftarrow A(x_1, \ldots; r)$. We let $[A(x_1, \ldots)]$ denote the set of all possible outputs of A when invoked with inputs $x_1, \ldots$.

We use the code based game playing framework of [11]. By $\Pr[\mathrm{G} \Rightarrow y]$ we denote the event that the execution of game G results in the game returning y. We write $\Pr[\mathrm{G}]$ as an abbreviation of $\Pr[\mathrm{G} \Rightarrow \mathsf{true}]$. In code of games, unless otherwise indicated, sets are assume initialized to empty, booleans to false, integers to 0 and anything else to $\perp$. We adopt the convention that the running time of an adversary refers to the worst-case execution time of the game with the adversary, so that the time for the execution of oracles to compute replies to oracle queries is included. This means that usually in reductions, adversary running time is roughly maintained.

If D, R are sets then Func(D, R) denotes the set of all functions from domain D to range R, and Perm(D) the set of all permutations on D.

PRFs and PRPs. Recall that the prf advantage of an adversary $\mathcal{A}$ against a family of functions $\mathsf{GG} : \mathsf{GG.Keys} \times \mathsf{GG.Dom} \to \mathsf{GG.Rng}$ is defined as $\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{GG}}(\mathcal{A}) = 2\Pr[\mathbf{G}^{\mathrm{prf}}_{\mathsf{GG}}(\mathcal{A})] - 1$, where game $\mathbf{G}^{\mathrm{prf}}_{\mathsf{GG}}(\mathcal{A})$ is shown in Fig. 1. Also the prp-cpa advantage of an adversary $\mathcal{A}$ against a family of permutations $\mathsf{GG} : \mathsf{GG.Keys} \times \mathsf{GG.Dom} \to \mathsf{GG.Dom}$ is defined as $\mathbf{Adv}^{\mathrm{prp\text{-}cpa}}_{\mathsf{GG}}(\mathcal{A}) = 2\Pr[\mathbf{G}^{\mathrm{prp\text{-}cpa}}_{\mathsf{GG}}(\mathcal{A})] - 1$, where game $\mathbf{G}^{\mathrm{prp\text{-}cpa}}_{\mathsf{GG}}(\mathcal{A})$ is shown in Fig. 1.

Ideal primitives. An *ideal primitive* is defined simply as a set of functions. An instance (meaning, a particular function) P will be picked at random in the games and provided as an oracle, to algorithms that need it and to the adversary. For example, the ideal primitive corresponding to a random oracle with domain D and range R is Func(D, R). Ideal ciphers are a bit more work since one must give access to both the map and its inverse. If K, D are sets then $\mathbf{IC}(\mathsf{K}, \mathsf{D})$ is the set of all maps $P : \mathsf{K} \times \mathsf{D} \times \{+, -\} \to \mathsf{D}$ with the property that $P(K, \cdot, +), P(K, \cdot, -) \in \mathrm{Perm}(\mathsf{D})$ are inverses of each other for every $K \in \mathsf{K}$. If $P \leftarrow\$ \mathbf{IC}(\mathsf{K}, \mathsf{D})$, and then P is provided as an oracle, we are in the ideal cipher model where one has oracle access to both the cipher $P(\cdot, \cdot, +)$ and its inverse $P(\cdot, \cdot, -)$. As an abbreviation, we let $\mathbf{IC}(k, n) = \mathbf{IC}(\{0,1\}^k, \{0,1\}^n)$, capturing ideal blockciphers with key length k and block length n.

3 FPE AND IB-FPE

We give definitions and basic results, including relations between notions, for FPE and IB-FPE.

FPE schemes. A *format-preserving encryption* (FPE) scheme F [10, 13] specifies a deterministic encryption algorithm $\mathsf{F.E} : \{0,1\}^{\mathsf{F.kl}} \times \mathsf{F.TS} \times \mathsf{F.Dom} \to \mathsf{F.Dom}$ together with a deterministic decryption algorithm $\mathsf{F.D} : \{0,1\}^{\mathsf{F.kl}} \times \mathsf{F.TS} \times \mathsf{F.Dom} \to \mathsf{F.Dom}$. Here $\{0,1\}^{\mathsf{F.kl}}$ is the keyspace, F.Dom is the domain and F.TS is the tweak space. For every key $J \in \{0,1\}^{\mathsf{F.kl}}$ and tweak $T \in \mathsf{T}$, the functions $\mathsf{F.E}(J, T, \cdot), \mathsf{F.D}(J, T, \cdot) \in \mathrm{Perm}(\mathsf{F.Dom})$ are permutations over F.Dom that are inverses of each other. We refer to F.kl as the key length. The scheme may have an associated ideal primitive F.IP, in which case F.E, F.D have oracle access to a function $P \in \mathsf{F.IP}$. Tweakable blockciphers [27] are a special case: FPE scheme F is a *tweakable blockcipher* if $\mathsf{F.Dom} = \{0,1\}^{\mathsf{F.bl}}$ for an integer F.bl called the blocklength.

FPE security. We recall the standard prp metric for an FPE scheme F [10, 13]. It coincides with the classic (strong) tweakable-prp metric of [26] in the case that F is a tweakable blockcipher. Let $\mathcal{A}$ be an adversary and define $\mathbf{Adv}^{\mathrm{prp}}_{\mathsf{F}}(\mathcal{A}) = 2\Pr[\mathbf{G}^{\mathrm{prp}}_{\mathsf{F}}(\mathcal{A})] - 1$, where game $\mathbf{G}^{\mathrm{prp}}_{\mathsf{F}}(\mathcal{A})$ is on the left in Fig. 2. The game picks a random challenge bit b and runs the adversary. The latter gets oracles Enc, Dec for

Game $\mathbf{G}^{\mathsf{prf}}_{\mathsf{GG}}(\mathcal{A})$	Game $\mathbf{G}^{\mathsf{prp\text{-}cpa}}_{\mathsf{GG}}(\mathcal{A})$
$b \leftarrow\!\$\, \{0,1\}; K \leftarrow\!\$\, \mathsf{GG.Keys}$	$b \leftarrow\!\$\, \{0,1\}; K \leftarrow\!\$\, \mathsf{GG.Keys}$
$b' \leftarrow\!\$\, A^{\textsc{Fn}}$; Return $(b' = b)$	$b' \leftarrow\!\$\, A^{\textsc{Fn}}$; Return $(b' = b)$
$\textsc{Fn}(X)$	$\textsc{Fn}(X)$
If $\mathrm{T}[X] \neq \bot$ then return $\mathrm{T}[X]$	If $\mathrm{ET}[X] \neq \bot$ then return $\mathrm{ET}[X]$
If $b = 0$ then $\mathrm{T}[X] \leftarrow\!\$\, \mathsf{GG.Rng}$	If $b = 0$ then $Y \leftarrow\!\$\, \{\, Y \in \mathsf{GG.Dom} : \mathrm{DT}[Y] = \bot \,\}$
Else $\mathrm{T}[X] \leftarrow \mathsf{GG}(K, X)$	Else $Y \leftarrow \mathsf{GG}(K, X)$
Return $\mathrm{T}[X]$	$\mathrm{ET}[X] \leftarrow Y$; $\mathrm{DT}[Y] \leftarrow X$; Return Y

Figure 1: Games defining PRF security (left) and PRP-CPA security (right) of GG.

Game $\mathbf{G}^{\mathsf{prp}}_{\mathsf{F}}(\mathcal{A})$	Game $\mathbf{G}^{\mathsf{prpa}}_{\mathsf{F}}(\mathcal{A})$
$b \leftarrow\!\$\, \{0,1\}$; $J \leftarrow\!\$\, \{0,1\}^{\mathsf{F.kl}}$; $P \leftarrow\!\$\, \mathsf{F.IP}$	$b \leftarrow\!\$\, \{0,1\}$; $J \leftarrow\!\$\, \{0,1\}^{\mathsf{F.kl}}$; $P \leftarrow\!\$\, \mathsf{F.IP}$; ch $\leftarrow$ false
$b' \leftarrow\!\$\, \mathcal{A}^{\textsc{Enc},\textsc{Dec},P}$; Return $(b = b')$	$b' \leftarrow\!\$\, \mathcal{A}^{\textsc{Enc},\textsc{Dec},\textsc{Ch},P}$; Return $(b = b')$
$\textsc{Enc}(T, X)$	$\textsc{Enc}(T, X)$
If $\mathrm{ET}[T, X] \neq \bot$ then return $\mathrm{ET}[T, X]$	If $\mathrm{ET}[T, X] \neq \bot$ then return $\mathrm{ET}[T, X]$
If $b = 0$ then $Y \leftarrow\!\$\, \{\, Y \in \mathsf{F.Dom} : \mathrm{DT}[T, Y] = \bot \,\}$	If (ch and $b = 0$) then $Y \leftarrow\!\$\, \{\, Y \in \mathsf{F.Dom} : \mathrm{DT}[T, Y] = \bot \,\}$
Else $Y \leftarrow \mathsf{F.E}^P(J, T, X)$	Else $Y \leftarrow \mathsf{F.E}^P(J, T, X)$
$\mathrm{ET}[T, X] \leftarrow Y$; $\mathrm{DT}[T, Y] \leftarrow X$; Return Y	$\mathrm{ET}[T, X] \leftarrow Y$; $\mathrm{DT}[T, Y] \leftarrow X$; Return Y
$\textsc{Dec}(T, Y)$	$\textsc{Dec}(T, Y)$
If $\mathrm{DT}[T, Y] \neq \bot$ then return $\mathrm{DT}[T, Y]$	If $\mathrm{DT}[T, Y] \neq \bot$ then return $\mathrm{DT}[T, Y]$
If $b = 0$ then $X \leftarrow\!\$\, \{\, X \in \mathsf{F.Dom} : \mathrm{ET}[T, X] = \bot \,\}$	If (ch and $b = 0$) then $X \leftarrow\!\$\, \{\, X \in \mathsf{F.Dom} : \mathrm{ET}[T, X] = \bot \,\}$
Else $X \leftarrow \mathsf{F.D}^P(J, T, Y)$	Else $X \leftarrow \mathsf{F.D}^P(J, T, Y)$
$\mathrm{ET}[T, X] \leftarrow Y$; $\mathrm{DT}[T, Y] \leftarrow X$; Return X	$\mathrm{ET}[T, X] \leftarrow Y$; $\mathrm{DT}[T, Y] \leftarrow X$; Return X
	$\textsc{Ch}()$
	ch $\leftarrow$ true

Figure 2: Games defining security of an FPE scheme F. Left: prp. Right: prpa.

encryption and decryption, and access to an instance **P** of the ideal primitive F.IP. It returns a bit b' and wins if $b' = b$. Enc takes a tweak T and message X and returns ciphertext Y, with Dec correspondingly taking tweak T and ciphertext Y to return message X. If $b = 1$, encryption and decryption are done using F with key J. If $b = 0$, each tweak is associated with a random permutation on F.Dom under which both encryption and decryption are done.

Letting $\mathcal{A}$ again be an adversary, we also define $\mathbf{Adv}^{\mathsf{prpa}}_{\mathsf{F}}(\mathcal{A}) = 2\Pr[\mathbf{G}^{\mathsf{prpa}}_{\mathsf{F}}(\mathcal{A})] - 1$, where game $\mathbf{G}^{\mathsf{prpa}}_{\mathsf{F}}(\mathcal{A})$ is on the right in Fig. 2. This captures what we call adaptive prp security, a notion we will find useful for proofs. Oracles Enc and Dec use F under key J until the adversary calls Ch to switch the game to challenge mode by setting flag ch to true. At that point, for each tweak, the associated permutation starts behaving randomly but consistent with the prior queries and (F-based) answers for that tweak. The prp notion corresponds to the special case where the first query is Ch(). We will exploit the following, which says that adaptivity can increase advantage by a factor of at most two in general.

Proposition 3.1. *Let F be an FPE scheme. Given a prpa adversary $\mathcal{A}_{\mathrm{prp}}$, we can build a prp adversary $\mathcal{A}_{\mathrm{prp}}$ of about the same running time, and making at most as many* Fn *queries, such that* $\mathbf{Adv}^{\mathsf{prpa}}_{\mathsf{F}}(\mathcal{A}_{\mathrm{prp}}) \leq 2 \cdot \mathbf{Adv}^{\mathsf{prp}}_{\mathsf{F}}(\mathcal{A}_{\mathrm{prp}})$.

Proof. The adversary $\mathcal{A}_{\mathrm{prp}}$ first picks a bit $a \leftarrow\!\$\, \{0,1\}$ and then runs $\mathcal{A}_{\mathrm{prp}}$. Before the latter calls Ch, the former always use its Enc/Dec oracles to reply to the Enc/Dec queries of $\mathcal{A}_{\mathrm{prp}}$. After $\mathcal{A}_{\mathrm{prp}}$ has called Ch to enter the challenge phase, if $a = 1$ then $\mathcal{A}_{\mathrm{prp}}$ continues to use its Enc/Dec oracles to reply to $\mathcal{A}_{\mathrm{prp}}$'s Enc/Dec queries. However, if $a = 0$ then $\mathcal{A}_{\mathrm{prp}}$ gives answers that are random but still consistent with prior queries and answers. When $\mathcal{A}_{\mathrm{prp}}$ outputs its guess b' then $\mathcal{A}_{\mathrm{prp}}$ outputs 1 if $b' = a$, and outputs 0 otherwise. Let c_{prp} be the challenge bit of game $\mathbf{G}^{\mathsf{prpa}}_{\mathsf{F}}(\mathcal{A}_{\mathrm{prp}})$. We claim that

$$\Pr[\mathbf{G}^{\mathsf{prp}}_{\mathsf{F}}(\mathcal{A}_{\mathrm{prp}}) \Rightarrow \mathsf{true} \mid c_{\mathrm{prp}} = 1] = \Pr[\mathbf{G}^{\mathsf{prpa}}_{\mathsf{F}}(\mathcal{A}_{\mathrm{prp}})] \quad (2)$$

$$\Pr[\mathbf{G}^{\mathsf{prp}}_{\mathsf{F}}(\mathcal{A}_{\mathrm{prp}}) \Rightarrow \mathsf{false} \mid c_{\mathrm{prp}} = 0] = \frac{1}{2} \ . \quad (3)$$

This is because (1) when $c_{\mathrm{prp}} = 0$, the answers for $\mathcal{A}_{\mathrm{prp}}$'s Enc and Dec queries are always simulated via an ideal family of permutations, meaning that whatever $\mathcal{A}_{\mathrm{prp}}$ receives is independent of a, but (2) when $c_{\mathrm{prp}} = 1$, the guess of $\mathcal{A}_{\mathrm{prp}}$ is incorrect if and only if $b' = a$. Subtracting Eq. (2) and Eq. (3) side by side we have

$$\mathbf{Adv}^{\mathsf{prp}}_{\mathsf{F}}(\mathcal{A}_{\mathrm{prp}}) = \frac{1}{2}\mathbf{Adv}^{\mathsf{prpa}}_{\mathsf{F}}(\mathcal{A}_{\mathrm{prp}})$$

as claimed. □

Proposition 3.1 says that prpa is an alternative, equivalent (up to a factor two in advantage) characterization of classic (strong) prp security for FPE schemes and tweakable blockciphers. For untweaked blockciphers, Desai and Miner [16] consider a notion of indistinguishable uniform permutation that is prpa with the adversary restricted to just one post-challenge encryption query and no decryption queries, showing it is equivalent to classic prp security up a factor two in advantage. Our proof extends theirs.

For FPE, we do not need to consider key-recovery security. We will for IB-FPE.

IB-FPE. A *key-derivation function* for FPE scheme F is a function $\mathsf{KDF} : \mathsf{KDF.MKS} \times \mathsf{KDF.IS} \to \{0,1\}^{\mathsf{F.kl}}$ that takes a *master key* K in the master-key space KDF.MKS and a *user identity* I in the identity-space KDF.IS to return a key $\mathsf{KDF}(K, I) \in \{0,1\}^{\mathsf{F.kl}}$ for F. An *identity-based FPE* (IB-FPE) scheme is a pair (F, KDF) consisting of a (base) FPE scheme F and a key-derivation function KDF for F. An IB-FPE scheme (F, KDF) is an *identity-based tweakable blockcipher* if F is a tweakable blockcipher.

The key-derivation function KDF may have an associated ideal primitive, denoted KDF.IP, in which case KDF has oracle access to a function $P \in \mathsf{KDF.IP}$. We require that F.IP = KDF.IP, meaning the ideal primitive of the key-derivation function is the same as that of the FPE scheme, and in games a single instance P of the ideal primitive will be used as the oracle for F.E, F.D and KDF. This is not only for simplicity but, more importantly, because the primitive in practice is often instantiated via the same cryptographic function, for example via AES.

IB-FPE SECURITY. Security requires that encryption under the key of some identity remains secure even if the adversary can obtain the keys of other identities. In terms of application and motivation, an identity might represent a point-of-sale terminal as discussed in Section 1, and thus our security requirement ensures that the damage from compromise of a terminal remains local, not affecting the security of encryption performed by other terminals. We give a prp style notion, ib-prp. We also give two variants of key-recovery security, ib-kr-ai and ib-kr-ti. The core notions are adaptive, but each has a corresponding non-adaptive version, obtained by restricting attention to non-adaptive adversaries as defined below. We establish relations between the notions as summarized in Fig. 4. The shown relations hold in both the adaptive and non-adaptive cases.

IB-PRP SECURITY. Let (F, KDF) be an IB-FPE scheme and $\mathcal{A}$ an adversary we call an ib-prp adversary. Define

$$\mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) = 2\Pr[\mathbf{G}^{\text{ib-prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A})] - 1,$$

where game $\mathbf{G}^{\text{ib-prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A})$ is on the left in Fig. 3. The game picks a random challenge bit b and runs the adversary. The latter gets oracles ENC, DEC for encryption and decryption, an expose oracle EXP, a challenge oracle CH and access to an instance P of the ideal primitive F.IP = KDF.IP. It returns a bit b' and wins if $b' = b$. XI is the set of exposed identities and ChI is the set of challenge identities. These sets stay disjoint throughout the game. Let us refer to an identity as neutral if it is in neither of these sets. All identities start neutral, since the sets XI, ChI are initialized to empty. Encryption oracle ENC takes an identity I, tweak T and message X and returns ciphertext Y, while decryption oracle DEC correspondingly taking identity I, tweak T and ciphertext Y to return message X. For neutral identities (and thus at the start of the game), these oracles behave honestly, meaning use F under keys derived via KDF under master key K, regardless of the value of the challenge bit b. Imagine the adversary querying these for a while. Adaptively, at any point in this process, it can either expose the key of an identity I via a EXP(I) query (this captures real-world compromise of the key of this identity), or switch I to challenge mode via a CH(I) query. If I is exposed, the encryption and decryption oracles for it continue to behave honestly. If I is switched to a challenge identity, then encryption and decryption continue to behave honestly if $b = 1$, but, if $b = 0$, they use, for any given tweak, a permutation that is random subject to being consistent with prior queries and replies for that identity and tweak.

IB-KR SECURITY. Let (F, KDF) be an IB-FPE scheme and $\mathcal{A}$ an adversary we call a ib-kr adversary. Define

$$\mathbf{Adv}^{\text{ib-kr-ti}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) = \Pr[\mathbf{G}^{\text{ib-kr-ti}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A})]$$
$$\mathbf{Adv}^{\text{ib-kr-ai}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) = \Pr[\mathbf{G}^{\text{ib-kr-ai}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A})] ,$$

where the games are defined (together, they differ on just one indicated line) on the right in Fig. 3. There is no challenge bit, and the encryption and decryption oracles are always honest, using F. Oracle EXP again allows key exposure. Choice of challenge identities is again adaptive, meaning an identity can be named as a challenge one after encryption and decryption queries, either to it or to other identities. The adversary returns a key and an identity. In the target-identity case (ib-kr-ti), it wins if the key it provides is the correct one for the identity it provides. In the all-identity (ib-kr-ai) case, the identity it provides is ignored, and the adversary wins if the key it provides is correct for some (any) identity. In both cases, of course, the adversary can only win if the identity for which it finds the key is not exposed.

KEY-DERIVATION FUNCTIONS. In designs of IB-FPE schemes we will of course want efficient key-derivation functions. But in analyses and for other conceptual purposes, it will be useful to consider key-derivation functions that are not efficient. In particular we define the *uniform key-derivation function* $\mathsf{U} = \mathbf{U}[\mathsf{F}, \mathsf{ID}]$, associated to F and a set ID of identities, to capture users having random, independent keys. Formally, let the master-key space $\mathsf{U.MKS} = \mathsf{Func}(\mathsf{ID}, \{0,1\}^{\mathsf{F.kl}})$ be the set of all functions from ID to $\{0,1\}^{\mathsf{F.kl}}$, so that a master key $K : \mathsf{ID} \to \{0,1\}^{\mathsf{F.kl}}$ is a function taking an identity and returning the key $K(I)$. Then the function $\mathsf{U} : \mathsf{U.MKS} \times \mathsf{ID} \to \{0,1\}^{\mathsf{F.kl}}$ is defined by $\mathsf{U}(K, I) = K(I)$. Picking K at random means the keys of different identities are random and independent.

NON-ADAPTIVE SECURITY. Let $\mathcal{A}$ be either a ib-prp or an ib-kr adversary. We say that it is *non-adaptive* if there is no identity I for which $\mathcal{A}$ makes both a EXP(I) query and a non-EXP(I) —that is, CH(I), ENC($I, \cdot, \cdot$) or DEC($I, \cdot, \cdot$)— query. Thus, the adversary must make its decision to expose the key of an identity I up front, without prior queries to ENC($I, \cdot, \cdot$) or DEC($I, \cdot, \cdot$). (The definition also excludes post EXP(I) queries ENC($I, \cdot, \cdot$), DEC($I, \cdot, \cdot$) and CH(I), but these are redundant anyway.) The security we prove for our constructions of IB-FPE schemes is restricted to non-adaptive adversaries as adaptivity allows selective-opening attacks (SOAs) [8, 18]. We elaborate on this below.

DISCUSSION. In the definition of security for identity-based encryption (IBE) [14], the adversary can pick its (single) challenge not

Game $\mathbf{G}^{\text{ib-prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A})$

$b \leftarrow_{\$} \{0,1\}$; $K \leftarrow_{\$} \mathsf{KDF.MKS}$
$\mathrm{XI} \leftarrow \emptyset$; $\mathrm{ChI} \leftarrow \emptyset$; $P \leftarrow_{\$} \mathsf{F.IP}$
For every $I \in \mathsf{KDF.IS}$ do $J_I \leftarrow \mathsf{KDF}^P(K, I)$
$b' \leftarrow_{\$} \mathcal{A}^{\textsc{Enc},\textsc{Dec},\textsc{Exp},\textsc{Ch},P}$; Return $(b = b')$

$\textsc{Enc}(I, T, X)$

If $\mathrm{ET}[I, T, X] \neq \bot$ then return $\mathrm{ET}[I, T, X]$
If $(I \in \mathrm{ChI}$ and $b = 0)$ then $Y \leftarrow_{\$} \{\, Y \in \mathsf{F.Dom} : \mathrm{DT}[I, T, Y] = \bot \,\}$
Else $Y \leftarrow \mathsf{F.E}^P(J_I, T, X)$
$\mathrm{ET}[I, T, X] \leftarrow Y$; $\mathrm{DT}[I, T, Y] \leftarrow X$; Return Y

$\textsc{Dec}(I, T, Y)$

If $\mathrm{DT}[I, T, Y] \neq \bot$ then return $\mathrm{DT}[I, T, Y]$
If $(I \in \mathrm{ChI}$ and $b = 0)$ then $X \leftarrow_{\$} \{\, X \in \mathsf{F.Dom} : \mathrm{ET}[I, T, X] = \bot \,\}$
Else $X \leftarrow \mathsf{F.D}^P(J_I, T, Y)$
$\mathrm{ET}[I, T, X] \leftarrow Y$; $\mathrm{DT}[I, T, Y] \leftarrow X$; Return X

$\textsc{Exp}(I)$

If $I \in \mathrm{ChI}$ then return $\bot$
$\mathrm{XI} \leftarrow \mathrm{XI} \cup \{I\}$; Return J_I

$\textsc{Ch}(I)$

If $I \in \mathrm{XI}$ then return $\bot$
$\mathrm{ChI} \leftarrow \mathrm{ChI} \cup \{I\}$

Game $\mathbf{G}^{\text{ib-kr-ti}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A})$ / $\mathbf{G}^{\text{ib-kr-ai}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A})$

$K \leftarrow_{\$} \mathsf{KDF.MKS}$
$\mathrm{XI} \leftarrow \emptyset$; $\mathrm{ChI} \leftarrow \emptyset$; $\mathrm{ChK} \leftarrow \emptyset$; $P \leftarrow_{\$} \mathsf{F.IP}$
For every $I \in \mathsf{KDF.IS}$ do $J_I \leftarrow \mathsf{KDF}^P(K, I)$
$(J', I') \leftarrow_{\$} \mathcal{A}^{\textsc{Enc},\textsc{Dec},\textsc{Exp},\textsc{Ch},P}$
Return $((J', I') \in \mathrm{ChK})$ // ib-kr-ti
Return $(\exists I : (J', I) \in \mathrm{ChK})$ // ib-kr-ai

$\textsc{Enc}(I, T, X)$

Return $\mathsf{F.E}^P(J_I, T, X)$

$\textsc{Dec}(I, T, Y)$

Return $\mathsf{F.D}^P(J_I, T, Y)$

$\textsc{Exp}(I)$

If $I \in \mathrm{ChI}$ then return $\bot$
$\mathrm{XI} \leftarrow \mathrm{XI} \cup \{I\}$; Return J_I

$\textsc{Ch}(I)$

If $I \in \mathrm{XI}$ then return $\bot$
$\mathrm{ChI} \leftarrow \mathrm{ChI} \cup \{I\}$
$\mathrm{ChK} \leftarrow \mathrm{ChK} \cup \{(J_I, I)\}$

Figure 3: Games defining security of an IB-FPE scheme (F, KDF)**. Left: ib-prp. Right: ib-kr-ti and ib-kr-ai.**

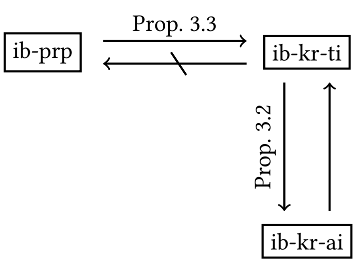

Figure 4: Relations between notions of IB-PRP security. An arrow A $\rightarrow$ B **is an implication: if an IB-PRP scheme meets** A **then it also meets** B**. A barred arrow** A $\nrightarrow$ B **is a separation: there exists an IB-PRP scheme meeting** A **but not** B**. Unannotated lines represent trivial relations. The relations hold in both the adaptive and non-adaptive settings.**

just while querying an exposure oracle (and, in the CCA case, a decryption oracle), but as a function of encryptions under identities of its choice. The latter is captured, trivially, by giving the adversary the master public key up front. Our setting is symmetric, so there is no master public key. As per the paradigm of [4], we accordingly give the adversary an encryption oracle. (To capture CCA, we also give it a decryption oracle. We continue of course to give the exposure oracle.) We allow multiple challenge identities, not just one. Starting encryption (and decryption) for an identity as honest and switching to challenge mode via $\textsc{Ch}$ captures an adaptive (encryption-dependent) choice of challenge identities to mirror IBE security. However, the presence of multiple challenges means that this effectively allows a SOA. SOAs are notoriously subtle, and security against them is known (at least for other primitives) to be hard to achieve [5, 8, 18, 22]. Correspondingly (and unsurprisingly) we find that we are unable to show our schemes meet our ib-prp definition for adaptive adversaries. We prove it, instead, for non-adaptive adversaries. These adversaries are still very powerful. (It is unclear that adaptivity is realistic or possible in practice.) We leave adaptive security as an open question.

The multi-user (mu) setting [2, 3] considers many users, having keys that are uniformly and independently distributed. Mu security of an FPE scheme F can be viewed as a special case of our setting, as follows. Let n be the number of users, and let $\mathsf{ID} = \{1, \ldots, n\}$. Let $\mathsf{U} = \mathbf{U}[\mathsf{F}, \mathsf{ID}]$ be the uniform key-derivation function for F over this set of identities, and consider the IB-FPE scheme (F, U). Let us call a ib-prp adversary $\mathcal{A}$ a mu adversary if it begins by querying all n identities to its $\textsc{Ch}$ oracle, and makes no $\textsc{Exp}$ queries. Then mu security of F is exactly ib-prp security of (F, U) relative to mu adversaries. In this way, certain results about IB-FPE will automatically imply results on the mu security of the base FPE scheme. Also, this lends a different perspective on IB-FPE, viewing it as a generalization of mu security in which keys of different users are not necessarily random and independent, key exposures are permitted and identities can be adaptively and optionally made challenge ones. We thank Stefano Tessaro for pointing out this connection and viewpoint to us.

A tweakable blockcipher [27] is the special case of an FPE scheme F in which $\mathsf{F.Dom} = \{0,1\}^{\mathsf{F.bl}}$ for some $\mathsf{F.bl}$. Mu security for tweakable blockciphers was considered in [25, 37]. The work of LLMM [25], Game $\mathbf{G}^{\text{fp}}_{\mathsf{F},d}(J, J')$

$P \leftarrow\$ \mathsf{F.IP}$; $T \leftarrow\$ \mathsf{F.TS}$
For $i = 1, \ldots, d$ do $X_i \leftarrow\$ \mathsf{F.Dom} \setminus \{X_1, \ldots, X_{i-1}\}$
$V \leftarrow (\mathsf{F.E}^P(J, T, X_1), \ldots, \mathsf{F.E}^P(J, T, X_d))$
$V' \leftarrow (\mathsf{F.E}^P(J', T, X_1), \ldots, \mathsf{F.E}^P(J', T, X_d))$
Return $(V = V')$

Figure 5: Game to define the false positive advantage of F on d random messages, for subkeys J and J'.

which is concurrent to, and independent of, ours, goes further to allow a key-derivation function so that they consider what in our language is effectively an identity-based tweakable blockcipher. Their definition of security, however, does not allow exposures and does not allow challenge identities to be adaptively determined. It is the special case of our ib-prp in which we restrict attention to what, above, we called mu adversaries.

False positive rate. Fix an IB-FPE scheme (F, KDF). In some settings, we have a F-key J and an identity I and want to test whether $J = \mathsf{KDF}(K, I)$. We don't have K, or the task is of course easy, but we do have access to an oracle $\mathsf{F.E}^P(\mathsf{KDF}(K, I), \cdot, \cdot)$. The strategy is pick some tweak T and inputs $X_1, \ldots, X_d$, and declare J correct if $\mathsf{F.E}^P(\mathsf{KDF}(K, I), T, X_i) = \mathsf{F.E}^P(J, T, X_i)$ for all $i \in \{1, \ldots, d\}$. This test is not always correct. There may be false positives, meaning it might accept even if $J \neq \mathsf{KDF}(K, I)$. Here we give definitions to quantify this. Consider game $\mathbf{G}^{\text{fp}}_{\mathsf{F},d}(J, J')$ defined in Fig. 5 associated to F, keys $J, J' \in \{0,1\}^{\mathsf{F.kl}}$ and integer d. Then define the *false positive advantage*

$$\mathbf{Adv}^{\text{fp}}_{\mathsf{F},d} = \max_{J \neq J'} \Pr[\mathbf{G}^{\text{fp}}_{\mathsf{F},d}(J, J')]$$

as the maximum, over all distinct keys $J, J' \in \{0,1\}^{\{0,1\}^{\mathsf{F.kl}}}$, of the probability that the game returns true.

We now compute this advantage for the case that F is ideal. Let $N = |\mathsf{F.Dom}|$ be the size of the domain. If $J \neq J'$ then $\mathsf{F.E}^P(J, T, \cdot)$ and $\mathsf{F.E}^P(J', T, \cdot)$ are independent random permutations, and hence

$$\mathbf{Adv}^{\text{fp}}_{\mathsf{F},d} = \frac{1}{N(N-1)\cdots(N-d+1)} \ . \tag{4}$$

The choice of d required to make the bound of Eq. (4) negligible is usually quite small. For example if $N = 2^{32}$ then setting $d = 9$ will be enough, by Eq. (4), to ensure a false positive advantage of only $\mathbf{Adv}^{\text{fp}}_{\mathsf{F},d} \leq 2^{-256}$.

When F is not ideal, the false positive advantage depends on the structure of F. It is easy to give artificial examples of F for which $\mathbf{Adv}^{\text{fp}}_{\mathsf{F},d}$ remains high even for large d, for example by having two distinct keys J, J' that induce the same encryption function on all tweaks, meaning $\mathsf{F.E}(J, T, X) = \mathsf{F.E}(J', T, X)$ for all T, X, in which case $\mathbf{Adv}^{\text{fp}}_{\mathsf{F},d} = 1$ for all d. Real and natural designs of FPE schemes, however, are not expected to have such anomalies, and so it is customary to assume that the false positive advantage is about the same as that of an ideal FPE with the same domain, meaning approximated by Eq. (4). We will do this in our estimates.

Adversary $\mathcal{A}^{\text{Enc},\text{Dec},\text{Exp},\text{Ch},P}_{\text{prp}}$

$(J, I) \leftarrow\$ \mathcal{A}^{\text{EncSim},\text{DecSim},\text{Exp},\text{Ch},P}_{\text{kr}}$; $T_I \leftarrow \arg\min_T |Q(I, T)|$
For $i = 1, \ldots, n$ do
  $X \leftarrow\$ \mathsf{F.Dom} \setminus Q(I, T_I)$; $Q(I, T_I) \leftarrow Q(I, T_I) \cup \{X\}$
  If $\text{Enc}(I, T_I, X) \neq \mathsf{F.E}(J, T_I, X)$ then return 0
Return 1

Subroutine $\text{EncSim}(I, T, X)$

$Y \leftarrow \text{Enc}(I, T, Y); Q(I, T) \leftarrow Q(I, T) \cup \{X\}$
Return Y

Subroutine $\text{DecSim}(I, T, Y)$

$X \leftarrow \text{Dec}(I, T, Y); Q(I, T) \leftarrow Q(I, T) \cup \{X\}$
Return X

Figure 6: Adversary $\mathcal{A}_{\text{prp}}$ for Proposition 3.3.

Equivalence of ib-kr notions. It is clear that ib-kr-ai tightly implies ib-kr-ti; we now prove the converse. Given an ib-kr-ai adversary $\mathcal{A}_{\text{ai}}$, one can construct an ib-kr-ai adversary $\mathcal{A}_{\text{ti}}$ by running the former to get a candidate (I, J), and then testing J for all identities in the challenge set ChI to find a matching identity. We will use $\mathbf{Adv}^{\text{fp}}_{\mathsf{F},d}$ defined above to account for the probability of false positive. The following shows that ib-kr-ti indeed tightly implies ib-kr-ai; the proof is in the full version [6].

Proposition 3.2. *Let* (F, KDF) *be an IB-FPE scheme. Suppose that we are given an ib-kr-ai adversary $\mathcal{A}_{\text{ai}}$ of q Ch queries. For a parameter $d \in \mathbb{N}$, we can construct an ib-kr-ti adversary $\mathcal{A}_{\text{ti}}$ of about the same running time plus qd calls to* F.E *such that*

$$\mathbf{Adv}^{\text{ib-kr-ai}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}_{\text{ai}}) \leq \mathbf{Adv}^{\text{ib-kr-ti}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}_{\text{ti}}) + q \cdot \mathbf{Adv}^{\text{fp}}_{\mathsf{F},d} \ .$$

Adversary $\mathcal{A}_{\text{ti}}$ uses the same number of queries as $\mathcal{A}_{\text{ai}}$, plus qd additional Enc *queries. Finally, if $\mathcal{A}_{\text{ai}}$ is non-adaptive, so is $\mathcal{A}_{\text{ti}}$.*

ib-prp implies ib-kr. One would expect that prpa security of an IB-FPE scheme implies its ib-kr-ti (and thus ib-kr-ai as per Proposition 3.2) security. A basic template for showing that indistinguishability style security implies key-recovery security is given in [32]. The kr adversary is executed to obtain a candidate key J'. To determine its challenge bit, the executing adversary now tests J' by seeing if encryption under it, on some "un-used" input, equals the output of the encryption oracle on the same input, where "un-used" means not already queried to the encryption oracle in the simulation of the kr-adversary. In the $b = 1$ case, the test will succeed. In the $b = 0$ case, the output of the encryption oracle is random and will equal the encryption under J' with a probability inverse in the size of the domain. Since the latter is usually large, the probability p of success in this case will be small.

Here we follow the basic template above. Given an ib-kr-ti adversary $\mathcal{A}_{\text{kr}}$, without loss of generality, we can assume that the adversary makes only a single Ch query, and it is the very last query before the adversary outputs its guess (I, J). We say that $\mathcal{A}_{\text{kr}}$ leaves at least v unused points if there is a tweak T such that $\mathcal{A}_{\text{kr}}$ makes at most $(|\mathsf{F.Dom}| - v)$ $\text{Enc}(I, T, \cdot)$ or $\text{Dec}(I, T, \cdot)$ queries.

Proposition 3.3 below shows that ib-prp implies ib-kr; the proof is in the full version [6].

Proposition 3.3. *Let* $(\mathsf{F}, \mathsf{KDF})$ *be an IB-FPE scheme and let* $d = |\mathsf{F.Dom}|$ *be the size of the domain of* F. *Suppose we are given an ib-kr adversary* $\mathcal{A}_{\rm kr}$ *leaving at least* v *unused points. Let* n *be an integer parameter satisfying* $1 \le n \le v$. *Then we build an adversary* $\mathcal{A}_{\rm prp}$ *(shown in Fig. 6) such that*

$$\mathbf{Adv}^{\text{ib-kr-ti}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}_{\rm kr}) \le \mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}_{\rm prp}) + p\ ,$$

where

$$p = 2^{\mathsf{F.kl}} \cdot \frac{(v-n)!}{v!}\ .$$

Adversary $\mathcal{A}_{\rm prp}$ *makes the same number of queries as* $\mathcal{A}_{\rm kr}$, *plus* n *additional* Enc *queries. The running time of* $\mathcal{A}_{\rm prp}$ *is that of* $\mathcal{A}_{\rm kr}$ *plus the time for* n *executions of* $\mathsf{F.E}$. *Finally, if* $\mathcal{A}_{\rm kr}$ *is non-adaptive, so is* $\mathcal{A}_{\rm prp}$.

If d is large then p can be easily made small. The difficult case is when d is small. To illustrate the quality of our bounds in this case, let us consider an example, namely $d = 10^4$, corresponding to the encryption of 4 decimal digits of a credit-card number. With DFF, the key length will be $\mathsf{F.kl} = 256$. Setting $v = d/2 = 5000$ and $n = 30$ we have

$$p = 2^{256} \cdot \frac{(5000-n)!}{5000!} \le 2^{256} \cdot 2^{-368}\ ,$$

which is tiny.

Ib-kr doesn't imply ib-prp. Conversely, we claim that ib-kr-ti (and thus ib-kr-ai due to Proposition 3.2) does not imply ib-prp. (This is the hatched arrow in Fig. 4.) This can be shown by counter-example. Thus, consider the FPE scheme F defined by $\mathsf{F.E}(J, T, X) = X$ for all J, T, X. Let ID be some non-empty set of identities and let $\mathsf{KDF} = \mathbf{U}[\mathsf{F}, \mathsf{ID}]$ be the associated uniform key-derivation function as defined above. IB-FPE scheme $(\mathsf{F}, \mathsf{KDF})$ is certainly not ib-prp secure. However an adversary $\mathcal{A}$ has $\mathbf{G}^{\text{ib-kr-ti}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) \le 2^{-\mathsf{F.kl}}$, making it ib-kr-ti secure if the key length of F is large.

4 ATTACKS

In this section we give generic non-adaptive attacks on *any* IB-FPE $(\mathsf{F}, \mathsf{KDF})$ that show inherent quantitative limits to the security that is achievable. Our attacks recover derived keys (meaning, have good advantage under our ib-kr-ai metric), not just violate ib-prp security. An important implication of these attacks is that for k bits of ib-kr-ai security, the key-length of F (which is the length of derived keys) must be at least $2k$-bits regardless of the length of the master key. The reason, roughly, is that collisions between derived keys can be exploited to violate security. These attacks are important to show that our constructions of IB-FPE schemes in later sections are optimal in security given the key lengths.

Overview and diversity. Let $(\mathsf{F}, \mathsf{KDF})$ be an IB-FPE scheme. For integer parameters $q, p \ge 1$, we will show that there is an attack (adversary) $\mathcal{A}$ that succeeds in key recovery with advantage $\mathbf{Adv}^{\text{ib-kr-ai}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) \ge \Omega(pq) \cdot 2^{-\mathsf{F.kl}}$. The adversary makes $O(q)$ Enc and Ch queries and has offline computation effort about the cost of $O(p)$ encryptions under $\mathsf{F.E}$. In particular, to allow p, q to reach $O(2^k)$, one must have $\mathsf{F.kl} \ge 2k$.

We define the *diversity* $\mathrm{KDiv}_{\mathsf{KDF}}(q)$ of KDF relative to q as the expected size of the set $\{\mathsf{KDF}^P(K, I_1), \dots, \mathsf{KDF}^P(K, I_q)\}$, where the expectation is over $P \leftarrow\$\ \mathsf{KDF.IP}$, $K \leftarrow\$\ \mathsf{KDF.MKS}$, and $I_1, \dots, I_q$ sampled uniformly without replacement from $\mathsf{KDF.IS}$ (that is, sampled uniformly and independently subject to being distinct). High diversity means that keys of different identities are largely distinct, while low diversity means keys of distinct identities frequently collide. We will give two, separate attacks. The first, called the matching attack (MA) works when the diversity is low. Specifically, it has a high (constant) ib-kr-ai advantage when $\mathrm{KDiv}_{\mathsf{KDF}}(q) \le q/4$. The second, called the exhaustive-search attack (ESA) works when the diversity is high. Specifically, it has ib-kr-ai advantage around $\Omega(pq) \cdot 2^{-\mathsf{F.kl}}$ when $\mathrm{KDiv}_{\mathsf{KDF}}(q) > q/4$. All cases for the diversity being covered, one or the other attack always applies to get ib-kr-ai advantage of the claimed form. The analyses of the attacks are made more difficult by the fact that F and KDF share the same instance of the ideal primitive.

The matching attack. Let $(\mathsf{F}, \mathsf{KDF})$ and q be given. We associate to them the *matching adversary* $\mathbf{MA}_q$ described in Fig. 7. In this attack, the adversary first samples without replacement q identities $I_1, \dots, I_q$, and picks a random $\ell \leftarrow\$\ \{1, \dots, q\}$. The goal of the adversary is to recover the key of some identity I_i, for $i \in \{1, \dots, q\} \backslash \{\ell\}$. To achieve this, it queries $\mathrm{Exp}(I_\ell)$ to get the key J_ℓ corresponding to I_ℓ, and outputs J_ℓ as its guess. The intuition is that if the set $\{J_1, \dots, J_q\}$ of keys for all identities involved is small (which happens on the average if the diversity is low) then J_ℓ is likely to equal J_i for some $i \ne \ell$, and the adversary wins. The cost of the attack is q queries to Enc and one query to Exp. The following theorem gives a precise lower bound on adversary advantage.

Theorem 4.1. *Let* $(\mathsf{F}, \mathsf{KDF})$ *be an IB-FPE scheme. Then for any* $q \in \mathbb{N}$ *we have*

$$\mathbf{Adv}^{\text{ib-kr-ai}}_{\mathsf{F},\mathsf{KDF}}(\mathbf{MA}_q) \ge \frac{1}{2} - \frac{\mathrm{KDiv}_{\mathsf{KDF}}(q)}{q}\ .$$

In particular if $\mathrm{KDiv}_{\mathsf{KDF}}(q) \le q/4$ then $\mathbf{Adv}^{\text{ib-kr-ai}}_{\mathsf{F},\mathsf{KDF}}(\mathbf{MA}_q) \ge 1/2 - 1/4 = 1/4$, meaning the ib-kr-ai advantage is very high.

Proof. Let K be the master key of KDF and let $P \leftarrow\$\ \mathsf{KDF.IP}$. From Markov's inequality,

$$\Pr\left[\left|\{\mathsf{KDF}^P(K, I_1), \dots, \mathsf{KDF}^P(K, I_q)\}\right| \ge q/2\right] \le \frac{\mathrm{KDiv}_{\mathsf{KDF}}(q)}{q/2}\ .$$

Let $S = \{\mathsf{KDF}^P(K, I_1), \dots, \mathsf{KDF}^P(K, I_q)\}$ and suppose $|S| \le q/2$ which occurs with probability at least $1 - 2 \cdot \mathrm{KDiv}_{\mathsf{KDF}}(q)/q$. We say that identity I_i is *bad* if there is some $j \in \{1, \dots, q\} \setminus \{i\}$ such that I_i and I_j have the same derived key, meaning $\mathsf{KDF}^P(K, I_i) = \mathsf{KDF}^P(K, I_j)$. Note that if there are at most r bad identities then the set S must have size at least $q - r$. Since we assumed $|S| \le q/2$, there are at least $q/2$ bad identities. Since we pick ℓ at random, the chance that I_ℓ is bad is at least $1/2$. Hence $\mathbf{Adv}^{\text{ib-kr-ai}}_{\mathsf{F},\mathsf{KDF}}(\mathbf{MA}_q) \ge 1/2 - \mathrm{KDiv}_{\mathsf{KDF}}(q)/q$ as claimed. □

The exhaustive search attack. Let $(\mathsf{F}, \mathsf{KDF})$ be given, as well as integer parameters p, q, d. We associate to them adversary $\mathbf{ESA}_{q,p,d}$

Adversary $\mathbf{MA}_q^{\text{Enc},\text{Dec},\text{Exp},\text{Ch},P}$

$S \leftarrow \emptyset$; $T \leftarrow\$ \mathsf{F.TS}$
$X \leftarrow\$ \mathsf{F.Dom}$; $\ell \leftarrow\$ \{1, \ldots, q\}$
For $i \leftarrow 1$ to q do
  $I_i \leftarrow\$ \mathsf{KDF.IS}\backslash S$; $S \leftarrow S \cup \{I_i\}$
For $i \in \{1, \ldots, q\}\backslash\{\ell\}$ do
  $C \leftarrow \text{Ch}(I_i, T, X)$
Pick arbitrary $I \in \{I_1, \ldots, I_q\}\backslash\{I_\ell\}$
Return $(I, \text{Exp}(I_\ell))$

Adversary $\mathbf{ESA}_{q,p,d}^{\text{Enc},\text{Dec},\text{Exp},\text{Ch},P}$

$S_1, S_2 \leftarrow \emptyset$; $T \leftarrow\$ \mathsf{F.TS}$
For $\ell \leftarrow 1$ to d do
  $X_\ell \leftarrow\$ \mathsf{F.Dom}\backslash S_1$; $S_1 \leftarrow S_1 \cup \{X_\ell\}$
For $i \leftarrow 1$ to q do
  $I_i \leftarrow\$ \mathsf{KDF.IS}\backslash S_2$; $S_2 \leftarrow S_2 \cup \{I_i\}$
  For $\ell \leftarrow 1$ to d do $V_\ell \leftarrow \text{Enc}(I_i, T, X_\ell)$
  $Z_i \leftarrow (V_1, \ldots, V_d)$
For $j \leftarrow 1$ to p do
  $J_j \leftarrow\$ \{0, 1\}^{\mathsf{F.kl}}$
  For $\ell \leftarrow 1$ do d do $U_\ell \leftarrow \mathsf{F.E}^P(J_j, T, X_\ell)$
  $Z \leftarrow (U_1, \ldots, U_\ell)$; $i \leftarrow \mathsf{Find}(Z, Z_1, \ldots, Z_q)$
  If $i > 0$ then $(\text{Ch}(I_i)$; Return $(I_i, J_j))$

Figure 7: Top: The matching attack. Bottom: The exhaustive search attack.

described in Fig. 7. Algorithm $\mathsf{Find}(Z, Z_1, \ldots, Z_q)$, used in the attack as a subroutine, returns an index i such that $Z = Z_i$ if $Z \in \{Z_1, \ldots, Z_q\}$, and 0 otherwise. The attack somewhat generalizes and extends the NIST/NSA attack on FF2 [20], and also resembles Biham's key-collision attack on DES [12]. Biham's attack can be viewed as a special case of ours, where the key-derivation function is the uniform one, the domain is large, and the parameters p and q are close to $2^{\mathsf{F.kl}/2}$. The main novelty is a rigorous analysis lower-bounding the ib-kr-ai advantage. The attack uses dq queries to Enc, q queries to Ch, and no Exp queries. The running time is that of dp executions of $\mathsf{F.E}$ plus p executions of Find. With appropriate data structures, the latter should cost about $O(p \log q)$. The value of d will be a small constant that, in estimates above, we absorbed into the big-oh.

The idea is as follows. The adversary picks distinct identities $I_1, \ldots, I_q$. Let $J'_i = \mathsf{KDF}^P(K, I_i)$, where K is the master key chosen in the overlying key-recovery game $\mathbf{G}_{\mathsf{F,KDF}}^{\text{ib-kr-ai}}(\mathbf{ESA}_{q,p,d})$. The adversary aims to find one of the target keys $J'_1, \ldots, J'_q$ via exhaustive search over the space of FPE keys. It picks at random p keys $J_1, \ldots, J_p$ from the key space $\{0,1\}^{\mathsf{F.kl}}$ of F. Now, for each i, j, it aims to test whether $J'_i = J_j$. If any such test succeeds, it can call $\text{Ch}(I_i)$, return (I_i, J_j) and win. If the tests are perfectly correct, then it wins with probability about $pm \cdot 2^{-\mathsf{F.kl}}$ where $m = |\{J'_1, \ldots, J'_q\}|$, and if the diversity is high, like $\geq q/4$, then this looks like the ib-kr-ai advantage we want. There are however several difficulties. One is that there is no reasonable way to do perfectly correct testing. We will handle this by using the false positive advantage $\mathbf{Adv}_{\mathsf{F},d}^{\mathrm{fp}}$ defined in Section 3. Another difficulty is the analysis. In particular, $\mathrm{KDiv}_{\mathsf{KDF}}(q)$ is an expectation taken over the choice of P, but the same P is used by the encryption algorithm in the tests, so we cannot use independence of the success and false-positive probabilities.

The following gives a lower bound on the ib-kr-ai advantage of the exhaustive search attack. The proof is in the full version [6] of this paper.

Theorem 4.2. *Let* $(\mathsf{F}, \mathsf{KDF})$ *be an IB-FPE scheme. Then for any* $p, q, d \in \mathbb{N}$ *such that* $pq \leq 2^{\mathsf{F.kl}}$ *we have*

$$\mathbf{Adv}_{\mathsf{F,KDF}}^{\text{ib-kr-ai}}(\mathbf{ESA}_{q,p,d}) \geq \frac{p \cdot \mathrm{KDiv}_{\mathsf{KDF}}(q)}{2^{\mathsf{F.kl}+1}} - pq \cdot \mathbf{Adv}_{\mathsf{F},d}^{\mathrm{fp}} \ . \qquad \square$$

We saw above that if $\mathrm{KDiv}_{\mathsf{KDF}}(q) \leq q/4$ then the matching attack already gives an attack with high (constant) ib-kr-ai advantage. The exhaustive search attack is effective in the complementary case where $\mathrm{KDiv}_{\mathsf{KDF}}(q) > q/4$. In this case, assuming $\mathbf{Adv}_{\mathsf{F},d}^{\mathrm{fp}}$ is negligible, Theorem 4.2 says the attack has ib-kr-ai advantage about $pq/2^{\mathsf{F.kl}+3}$. In particular $p = q \approx 2^{\mathsf{F.kl}/2}$ yields constant advantage, showing that k bits of security requires $\mathsf{F.kl} \geq 2k$.

5 THE PRF CONSTRUCTION

We give a modular approach to build IB-FPE schemes. Given a prpa-secure FPE scheme F we set KDF to a PRF to get an ib-prp-secure IB-FPE scheme $(\mathsf{F}, \mathsf{KDF})$. Then we instantiate KDF to get IB-FPE schemes with security matching our attacks.

The **PRF** construction. Theorem 5.1 below proves ib-prp security of $(\mathsf{F}, \mathsf{KDF})$ assuming prpa security of F and prf security of KDF. The different resource metrics for $\mathcal{A}$ referred to below were defined in Section 3. The proof, which is in the full version [6], is a standard hybrid argument.

Theorem 5.1. *Let* $(\mathsf{F}, \mathsf{KDF})$ *be an IB-FPE scheme. Suppose we are given a non-adaptive ib-prp adversary* $\mathcal{A}$ *whose* Enc, Dec, Ch *queries involve at most* u *different identities, with at most* q_1 *queries to* Enc, Dec *per identity. Assume* $\mathcal{A}$ *makes* q_e *queries to* Exp. *Then we can construct a prpa adversary* $\overline{\mathcal{A}}$ *of* q_1 Enc/Dec *queries, and a prf adversary* $\mathcal{B}$ *making* $u + q_e$ *queries to its* Fn *oracle, such that*

$$\mathbf{Adv}_{\mathsf{F,KDF}}^{\text{ib-prp}}(\mathcal{A}) \leq u \cdot \mathbf{Adv}_{\mathsf{F}}^{\text{prpa}}(\overline{\mathcal{A}}) + 2 \cdot \mathbf{Adv}_{\mathsf{KDF}}^{\text{prf}}(\mathcal{B}) \ . \qquad (5)$$

The running time of $\overline{\mathcal{A}}$ *and* $\mathcal{B}$ *is about the same as that of* $\mathcal{A}$ *plus the time for* q_1 *executions of* $\mathsf{F.E}$. $\square$

From Theorem 5.1, one can obtain an IB-FPE scheme in the standard model, by setting F to a standard-model FPE scheme such as the Sometimes-Recurse shuffle [29]. This answers in the affirmative the theoretical question of whether IB-FPE is achievable in the standard model.

Tightness of bound. Suppose $\mathsf{F.kl} = 2k$ and F has ideal behavior. Then we would expect $\mathbf{Adv}_{\mathsf{F}}^{\text{prp}}(\overline{\mathcal{A}}) \approx q_1/2^{2k}$, corresponding to exhaustive key search being the best attack on prp security, and consequently from Proposition 3.2, $\mathbf{Adv}_{\mathsf{F}}^{\text{prpa}}(\overline{\mathcal{A}}) \lesssim 2q_1/2^{2k}$. Similarly assuming KDF has optimal prf security and $|\mathsf{KDF.MKS}| \geq 2^k$, we would expect $\mathbf{Adv}_{\mathsf{KDF}}^{\text{prf}}(\mathcal{B}) \approx (q_e + u)/2^k$. Then the bound of Eq. (7) becomes

$$\mathbf{Adv}_{\mathsf{F,KDF}}^{\text{ib-prp}}(\mathcal{A}) \lesssim \frac{2uq_1}{2^{2k}} + \frac{2(q_e + u)}{2^k} \ . \qquad (6)$$

KDF[E](K, I)
$J_0 \leftarrow E_K(I \,\|\, 00) \oplus E_K(I \,\|\, 01)$; $J_1 \leftarrow E_K(I \,\|\, 10) \oplus E_K(I \,\|\, 11)$
Return $J_0 \,\|\, J_1$

Figure 8: Key-derivation function KDF[E]**, where** $E : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}^k$ **is a blockcipher.**

This allows u, q_1, q_e to reach $O(2^k)$, which as per our attacks means the bound from Theorem 6.1 is essentially tight.

INSTANTIATING KDF. Recall that we want to use *only* AES as our cryptographic primitive. Thus one needs to show how to instantiate KDF from a blockcipher $E : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}^k$ such that KDF achieves k-bits of prf security assuming that E has k bits of prp-cpa security. This is non-trivial, and as a stepping stone, we first aim to achieve a good PRF $F : \{0,1\}^k \times \{0,1\}^{k-1} \to \{0,1\}^k$. BKR [9] suggest that one can build F by way of

$$F_K(x) = E_K(x \,\|\, 0) \oplus E_K(x \,\|\, 1) \ .$$

The following result by DHT [15] confirms that F indeed has k-bit prf security.

LEMMA 5.2. *[15] Let $E : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}^k$ be a blockcipher. Let $F : \{0,1\}^k \times \{0,1\}^{k-1} \to \{0,1\}^k$ be constructed by $E_K(x) = F_K(x \,\|\, 0) \oplus F_K(x \,\|\, 1)$. Then for any prf adversary A making $q \le 2^{k-5}$ queries to* FN*, we can construct an adversary $\overline{\mathcal{A}}$ of about the same running time and $2q$ oracle queries such that*

$$\mathbf{Adv}^{\mathrm{prf}}_F(\mathcal{A}) \le \mathbf{Adv}^{\mathrm{prp\text{-}cpa}}_E(\overline{\mathcal{A}}) + \frac{1.5q + 3\sqrt{q}}{2^k} \ . \qquad \square$$

We then can construct a key-derivation function KDF[E] : $\{0,1\}^k \times \{0,1\}^{k-2} \to \{0,1\}^{2k}$ by

$$\mathsf{KDF}[E](K, I) = F_K(I \,\|\, 0) \,\|\, F_K(I \,\|\, 1)$$

for any k-bit master key K and $(k-2)$-bit identity I. The key-derivation function KDF[E] can be expressed in terms of E as in Fig. 8. Proposition 5.3 below shows that KDF[E] also has k-bit prf security.

PROPOSITION 5.3. *Let $E : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}^k$ be a blockcipher. Let* KDF[E] *: $\{0,1\}^k \times \{0,1\}^{k-2} \to \{0,1\}^{2k}$ be as in Fig. 8. Then for any adversary $\mathcal{A}$ that makes $q \le 2^{k-6}$ queries, we can construct another adversary $\overline{\mathcal{A}}$ of about the same running time and $4q$ oracle queries such that*

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{KDF}[E]}(\mathcal{A}) \le \mathbf{Adv}^{\mathrm{prp\text{-}cpa}}_E(\overline{\mathcal{A}}) + \frac{3q + 5\sqrt{q}}{2^k} \ .$$

PROOF. Without loss of generality, suppose that $\mathcal{A}$ does not repeat a prior query. We first reduce the prf security of KDF to the prf security of F, by constructing an adversary $\mathcal{B}$ attacking F. Adversary $\mathcal{B}$ runs $\mathcal{A}$. When the latter queries FN(I), the former queries $I \,\|\, 0$ and $I \,\|\, 1$ to its oracle to get answer Z_0 and Z_1, and returns $Z_0 \,\|\, Z_1$ to $\mathcal{A}$. When $\mathcal{A}$ finally outputs a bit b', $\mathcal{B}$ outputs the same bit. Let a and b be the challenge bit in game $\mathbf{G}^{\mathrm{prf}}_{\mathsf{KDF}}(\mathcal{A})$ and $\mathbf{G}^{\mathrm{prf}}_F(\mathcal{B})$ respectively. Then

$$\Pr[\mathbf{G}^{\mathrm{prf}}_{\mathsf{KDF}[E]}(\mathcal{A}) \mid a = 1] = \Pr[\mathbf{G}^{\mathrm{prf}}_F(\mathcal{B}) \mid b = 1], \text{ and}$$
$$\Pr[\mathbf{G}^{\mathrm{prf}}_{\mathsf{KDF}[E]}(\mathcal{A}) \mid a = 0] = \Pr[\mathbf{G}^{\mathrm{prf}}_F(\mathcal{B}) \mid b = 0] \ .$$

Adding the equations above side by side we have

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{KDF}[E]}(\mathcal{A}) \le \mathbf{Adv}^{\mathrm{prf}}_F(\mathcal{B}) \ .$$

Adversary $\mathcal{B}$ has about the same running time as $\mathcal{A}$, and makes $2q \le 2^{k-5}$ oracle queries. Using Lemma 5.2, one can construct another adversary $\overline{\mathcal{A}}$ of about the same running time as $\mathcal{B}$, and $4q$ oracle queries such that

$$\begin{aligned}\mathbf{Adv}^{\mathrm{prf}}_F(\mathcal{B}) &\le \mathbf{Adv}^{\mathrm{prp\text{-}cpa}}_E(\overline{\mathcal{A}}) + \frac{3q + 3\sqrt{2q}}{2^k} \\ &\le \mathbf{Adv}^{\mathrm{prp\text{-}cpa}}_E(\overline{\mathcal{A}}) + \frac{3q + 5\sqrt{q}}{2^k} \ .\end{aligned}$$

Putting all this together we get the claimed result. $\square$

DISCUSSION. The KDF construction above uses 4 blockcipher calls. Alternatively, one might consider using Iwata's CENC method [23] that makes only 3 blockcipher calls. Specifically, let $G : \{0,1\}^k \times \{0,1\}^{k-2} \to \{0,1\}^k$ be constructed via

$$G_K(I) = (Z \oplus E(K, I \,\|\, 01)) \,\|\, (Z \oplus E(K, I \,\|\, 10)) \ ,$$

for any k-bit master key K and $(k-2)$-bit identity I, where $Z = E(K, I \,\|\, 00)$. IMV [24] claim that G has k-bit prf security, but their analysis is based on a combinatorial result by Patarin [30] whose proof is very involved with some unproven claims [15]. We therefore use the KDF[E] construction above, as it has a rigorous proof.

6 THE DBL CONSTRUCTION

In Section 5, we followed the natural route to building IB-FPE in which the key-derivation function KDF is a PRF, and showed that one can instantiate KDF using four calls to an underlying blockcipher. In this section, we'll consider how to build a faster key-derivation function KDF for a class of FPE schemes F that we call *square*. It is in fact an abstraction of the key-derivation function of the proposed DFF standard [36]. The key-derivation function makes just two calls to the underlying blockcipher. Interestingly, it has poor (only birthday-bound) prf security, but we'll give a dedicated analysis to justify the (non-adaptive) ib-prp security of the IB-FPE scheme.

We use a *common ideal primitive* framework. All schemes use a common instance of a single ideal primitive—an ideal cipher $\mathbf{IC}(k, k)$ in which the key and block length are the same. In particular we allow the starting square FPE scheme F to use $P \in \mathbf{IC}(k, k)$, and then define KDF using the *same* P. This is because, for efficiency and implementation ease, we aim for all final constructions to be (only) AES-based.

The analysis is made challenging by two elements. First is to not only prove security, but with a good bound. Second is that the ideal primitive being in common means there can be queries made to it (directly, or indirectly via other oracles) by both the key-derivation function and the encryption and decryption functions, so we cannot use independence in a straightforward way.

$\mathsf{KDF}^P(K, I)$
$J_0 \leftarrow P(K, \mathsf{M}_0(I), +)$; $J_1 \leftarrow P(K, \mathsf{M}_1(I), +)$; $J \leftarrow J_0 \| J_1$
Return J

Figure 9: Key-derivation function $\mathsf{KDF} = \mathbf{Dbl}[k, \mathsf{M}]$ associated to embedding scheme M, where $P \in \mathbf{IC}(k, k)$.

SQUARE FPE SCHEMES. Let F be an FPE scheme. We say that it is *square* if there is an integer $k \geq 1$ such that $\mathsf{F.kl} = 2k$ and $\mathsf{F.IP} = \mathbf{IC}(k, k)$. That is, the ideal primitive associated to F is the ideal cipher with key and block length both the same value k, and moreover keys for the scheme are of length $2k$. DFF is underlain by such a square scheme [36]. (In contrast, FF2 was not.) The term "square" refers to the fact that the key space has size 2^{2k}, the square of the size 2^k of what, below, will be the master key space of the IB-FPE scheme, which is crucial for getting high security due to the attacks from Section 4.

THE **Dbl** CONSTRUCTION. Let F be a square FPE scheme with $\mathsf{F.kl} = 2k$. We first define embedding schemes, and then a key-derivation function KDF to turn F into an IB-FPE scheme $(\mathsf{F}, \mathsf{KDF})$.

An *embedding scheme* M specifies a pair of functions $\mathsf{M}_0, \mathsf{M}_1 : \mathsf{M.IS} \to \{0,1\}^k$ satisfying two conditions: (1) Both M_0 and M_1 are injective and (2) the two maps have disjoint images, meaning $\mathsf{M}_0(I_1) \neq \mathsf{M}_1(I_2)$ for all $I_1, I_2 \in \mathsf{M.IS}$. We refer to $\mathsf{M}_0, \mathsf{M}_1$ as the embedding functions, and $\mathsf{M.IS}$ as the identity space, of M.

Now we define the key-derivation function $\mathsf{KDF} = \mathbf{Dbl}[k, \mathsf{M}]$ to construct an IB-FPE scheme $(\mathsf{F}, \mathsf{KDF})$. We let $\mathsf{KDF.IS} = \mathsf{M.IS}$, meaning the identity space is that of M. We let $\mathsf{KDF.mkl} = k$, so that a master key is a k-bit string. Then the key-derivation function $\mathsf{KDF}^P : \{0,1\}^k \times \mathsf{M.IS} \to \{0,1\}^{2k}$ is as specified in Fig. 9. The key for identity I is the result of applying the ideal cipher, keyed with the master key K, to $\mathsf{M}_0(I)$ and $\mathsf{M}_1(I)$, and concatenating these k-bit strings to get a $2k$-bit key. The key-derivation function, the encryption algorithm and the decryption algorithm all have access to $P \in \mathbf{IC}(k, k)$. We stress that, as discussed above for this common ideal primitive framework, the key derivation uses the *same* instance of the ideal cipher as encryption and decryption.

RESISTANCE TO ATTACKS. Let F be a square FPE scheme with $\mathsf{F.kl} = 2k$. We consider how well the attacks of Section 4 do against F under $\mathsf{KDF} = \mathbf{Dbl}[k, \mathsf{M}]$. First, we claim that our choice of KDF renders the matching attack entirely ineffective. Indeed, since $P(K, \cdot, +)$ is a permutation, the keys for distinct identities will be distinct. Thus for any $K \in \{0,1\}^k$ and any identities $I_1, \ldots, I_q \in \mathsf{M.IS}$, the set $\{\mathsf{KDF}^P(K, I_1), \ldots, \mathsf{KDF}^P(K, I_q)\}$ will have size exactly q. So its expected size, which is our diversity metric $\mathrm{KDiv}_{\mathsf{KDF}}(q)$ from Section 4, will equal q. Not only does Theorem 4.1 become vacuous, but, looking at the attack in Fig. 7, we see that it will have ib-kr-ai advantage zero, because the key returned by $\mathrm{Exp}(I_\ell)$ will not equal any of the keys corresponding to the other identities. This shows a benefit of using a block cipher as the tool in key derivation for KDF. Had we used even a random oracle, the matching attack would have had at least some success.

The exhaustive search attack does have a non-trivial ib-kr-ai advantage. We noted above that $\mathrm{KDiv}_{\mathsf{KDF}}(q) = q$. Assuming the false positive advantage $\mathbf{Adv}^{\mathrm{fp}}_{\mathsf{F},d}$ is negligible, recall that Theorem 4.2 says that the ib-kr-ai advantage of the exhaustive search attack is about $p \cdot \mathrm{KDiv}_{\mathsf{KDF}}(q) \cdot 2^{-\mathsf{F.kl}-1} = pq \cdot 2^{-\mathsf{F.kl}-1} = pq \cdot 2^{-2k-1}$, where q is the number of adversary Enc queries and p is roughly its running time. So the ib-kr-ai advantage stays below 1 as long as p and q each stay below 2^k. This means we have k-bits of security against this attack, and explains the choice of square schemes.

In summary, $\mathsf{KDF} = \mathbf{Dbl}[k, \mathsf{M}]$ has been designed so that the attacks we gave in Section 4 are not threats to the security of F under KDF, in particular because $\mathsf{F.kl} = 2k$ while $\mathsf{KDF.mkl} = k$. However, this does *not* guarantee security, since there may well be other attacks. Moreover, KDF has only birthday-bound prf security, and thus using Theorem 5.1 gives us only $k/2$-bits of ib-prp security for $(\mathsf{F}, \mathsf{KDF})$. The main purpose of this section is to supply proof-based evidence of k-bit security.

GOALS AND NAIVE REDUCTION. The assumption we make is that the given square FPE scheme F satisfies prpa security. (This is equivalent to conventional prp security as per Proposition 3.1.) Our goal is thus to reduce the ib-prp security of $(\mathsf{F}, \mathbf{Dbl}[k, \mathsf{M}])$ to the prpa security of F. As F is defined in the ideal-cipher model, this involves something somewhat novel, namely a reduction in the ideal cipher model. (Usually, in idealized models, one directly proves bounds on adversary advantage rather than giving reductions.) Given a non-adaptive ib-prp adversary $\mathcal{A}$ we aim to build another adversary $\overline{\mathcal{A}}$ and bound $\mathbf{Adv}^{\mathrm{ib\text{-}prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A})$ as a function of $\mathbf{Adv}^{\mathrm{prpa}}_{\mathsf{F}}(\overline{\mathcal{A}})$ and the resources of $\mathcal{A}$, in particular the number u of users (identities) queried. $\overline{\mathcal{A}}$ will simulate $\mathcal{A}$'s P oracle.

The natural approach is a hybrid argument. The naive way of doing this, however, will incur a loss of $u^2/2^k$ in the advantage. This is undesirable since we want to show security up to $u \approx 2^k$, not $u \approx 2^{k/2}$. In more detail, the i-th hybrid game would let the keys of the first i identities be random, and the rest be specified via KDF as per Fig. 9 ($0 \leq i \leq u$). Adversary $\overline{\mathcal{A}}$ would pick i at random to play the role of its single user, aiming to simulate the other identities for $\mathcal{A}$. Let J_i denote the key (underlying the single identity queried) in $\overline{\mathcal{A}}$'s game. The difficulty is that, for the simulation to be correct in the case that $\overline{\mathcal{A}}$'s challenge bit is 1, the $j = u - i + 1$ keys $J_i, \ldots, J_u$ must be consistent with the structure imposed by KDF, meaning be formed by taking $2j$ distinct, random k bit strings and concatenating them in pairs. But J_i is random since $\overline{\mathcal{A}}$ is in the prpa game, and while $\overline{\mathcal{A}}$ can pick $J_{i+1}, \ldots, J_u$ to have the desired structure, this leaves a probability $\epsilon = O(u/2^k)$ that J_i will not have a consistent structure. Specifically, regardless of how $\overline{\mathcal{A}}$ picks distinct $J_{i+1}, \ldots, J_u$, the chance that one of those is J_i is $\epsilon = (u-i)/2^k = O(u/2^k)$. This means a loss of ϵ in each hybrid step, meaning, when $\overline{\mathcal{A}}$ picks i, its advantage is the difference in probabilities from the $(i+1)$-th and i-th hybrid games plus ϵ. When we sum over all hybrids (corresponding to the random choice of i), we get a $u\epsilon$ loss. What we want instead is a reduction with a loss that is $O(u/2^k)$ *globally*. This is what we will provide below, thereby showing security matching our attacks.

KEY USAGE METRIC. When invoked with a particular key J, the algorithms $\mathsf{F.E}$, $\mathsf{F.D}$ of the square FPE scheme will invoke their ideal cipher instance P with certain keys. Specifically there is a set $T(J) \subseteq$

$\{0,1\}^k$ such that all P-queries of F.E and F.D only use keys in $T(J)$, regardless of the inputs to F.E, F.D and responses to oracle queries. We let F.nk be the maximum, over all J, of the size of $T(J)$. This may sound complicated but it is really simple because typical constructions will evaluate the ideal cipher only on some fixed number of keys related to J. For example, for $\mathsf{F} = \mathsf{FF}_{\mathsf{dff}}$, we have F.nk = 1. That is, there is only one blockcipher key used in the construction. We define this because our bounds will depend on it.

Reduction theorem. We now reduce the non-adaptive ib-prp security of our constructed IB-FPE scheme to the prpa security of the underlying FPE scheme. (The latter can be further reduced to its conventional prp security via Proposition 3.1.) The following theorem gives a good bound, where the global loss (the second term in the bound) is only $O(q/2^k)$ over and above the inevitable linear loss from the hybrid argument, where q is linear (not quadratic) in the number of queries that $\mathcal{A}$ makes to its different oracles. The quality of the bound is the same as that of Theorem 5.1, despite the low prf security of $\mathbf{Dbl}[k, \mathsf{M}]$.

Theorem 6.1. *Let* F *be a square FPE scheme with* F.kl = $2k$. *Let* KDF = $\mathbf{Dbl}[k, \mathsf{M}]$ *be a key-derivation function as per Fig. 9. Suppose we are given a non-adaptive adversary* $\mathcal{A}$ *whose* Enc, Dec, Ch *queries involve at most* u *different identities, with at most* q_1 *queries to* Enc, Dec *per identity. Assume* $\mathcal{A}$ *makes* q_e *queries to* Exp *and* p *queries to* IP. *The proof constructs an adversary* $\overline{\mathcal{A}}$ *such that*

$$\begin{aligned} &\mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) \\ \le\quad & u \cdot \mathbf{Adv}^{\text{prpa}}_{\mathsf{F}}(\overline{\mathcal{A}}) + \frac{8u + 8q_e + 2p + 2u \cdot \mathsf{F.nk} - 6}{2^k} . \quad (7) \end{aligned}$$

Adversary $\overline{\mathcal{A}}$ *makes at most* q_1 *queries to* Enc, Dec *and* p *queries to* IP. *Its running time is about the same as that of* $\mathcal{A}$. □

Starting above, we may use IP as the name of the game procedure that implements the ideal primitive instance. Where Fig. 3 gives $\mathcal{A}$ oracles Enc, Dec, Exp, Ch, P, we would now give it oracles Enc, Dec, Exp, Ch, IP, with IP(x) defined to simply return $P(x)$ in the games of Fig. 3. The reason it helps to name the procedure is that in our proofs it will be programmed, and not always set to P. It will also be useful to define the key-derivation function $\overline{\mathsf{KDF}} : \mathrm{Perm}(\{0,1\}^k) \times \mathsf{M.IS} \to \{0,1\}^{2k}$ by

$$\overline{\mathsf{KDF}}(\overline{\pi}, I) = \overline{\pi}(\mathsf{M}_0(I)) \,\|\, \overline{\pi}(\mathsf{M}_1(I)) \quad (8)$$

for all $\overline{\pi} \in \mathrm{Perm}(\{0,1\}^k)$ and all $I \in \mathsf{M.IS}$. We prove Theorem 6.1 by invoking lemmas that will follow.

Proof of Theorem 6.1. Let $N = u + q_e$. Let $\overline{\mathsf{KDF}}$ be the key derivation function defined by Eq. (8). Using Lemma 6.2 and then Lemma 6.3 we have

$$\begin{aligned} \mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) &\le \mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\overline{\mathsf{KDF}}}(\mathcal{A}) + \frac{2p + 2u \cdot \mathsf{F.nk}}{2^k} \\ &\le u \cdot \mathbf{Adv}^{\text{prpa}}_{\mathsf{F}}(\overline{\mathcal{A}}) + \frac{8N - 6 + 2p + 2u \cdot \mathsf{F.nk}}{2^k} , \end{aligned}$$

where $\overline{\mathcal{A}}$ is the adversary given by Lemma 6.3. □

Lemmas. The first lemma allows a move to a setting where key derivation no longer uses the ideal primitive P that is used by F, but instead generates keys independently, although still with the same distribution as that of the prescribed key-derivation scheme. This lemma holds for both adaptive and non-adaptive adversaries $\mathcal{A}$.

Lemma 6.2. *Let* F *be a square FPE scheme with* F.kl = $2k$. *Let* KDF = $\mathbf{Dbl}[k, \mathsf{M}]$ *be the key-derivation function of Fig. 9. Let* $\overline{\mathsf{KDF}}$ *be the key derivation function defined by Eq. (8). Let* $\mathcal{A}$ *be an adversary. Then*

$$\mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) \le \mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\overline{\mathsf{KDF}}}(\mathcal{A}) + \frac{2p + 2u \cdot \mathsf{F.nk}}{2^k} \quad (9)$$

where p *is the number of* IP *queries of* $\mathcal{A}$ *and* u *is the number of different identities involved in the* Enc, Dec *queries of* $\mathcal{A}$. □

Note that the reduction does not change the adversary. Our claim is that the ib-prp advantage of $\mathcal{A}$ with respect to the original key-derivation scheme is bounded by its ib-prp advantage with respect to the newly-defined key-distribution scheme plus an error term that is linear in the resources.

Proof of Lemma 6.2. Consider games G_0 and G_1 of Fig. 10. They optimistically imagine that key generation works as per $\overline{\mathsf{KDF}}$, picking a random permutation $\overline{\pi}$ and using it to specify the user keys. The notation $(L, W, s) \leftarrow x$ in the code for IP means this oracle parses its query x as a triple consisting of a key $L \in \{0,1\}^k$, an input $W \in \{0,1\}^k$, and a sign $s \in \{+,-\}$. If L equals the master key K, the bad flag is set to true, and game G_0, which includes the boxed code, corrects by setting $P(K,\cdot,+)$ to $\overline{\pi}$ and its inverse $P(K,\cdot,-)$ to $\overline{\pi}^{-1}$. Game G_1, however, does not include the boxed code. The result is that game G_0 is using KDF for key generation while game G_1 is using $\overline{\mathsf{KDF}}$. Thus

$$\mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) = 2\Pr[\mathrm{G}_0] - 1 \quad (10)$$

$$\mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\overline{\mathsf{KDF}}}(\mathcal{A}) = 2\Pr[\mathrm{G}_1] - 1 . \quad (11)$$

Games $\mathrm{G}_0, \mathrm{G}_1$ are identical-until-bad, so by the Fundamental Lemma of Game Playing [11] we have

$$\begin{aligned} 2\Pr[\mathrm{G}_0] - 1 &= 2\Pr[\mathrm{G}_1] - 1 + 2 \cdot (\Pr[\mathrm{G}_0] - \Pr[\mathrm{G}_1]) \\ &\le 2\Pr[\mathrm{G}_1] - 1 + 2\Pr[\mathrm{G}_1 \text{ sets bad}] . \quad (12) \end{aligned}$$

Queries to IP may be made directly by the adversary, and there are p such. However, such queries may also be made by the F.E and F.D algorithms when invoked in Enc and Dec queries. But we know that for any J the total number of different keys that $\mathsf{F.E}^{\mathrm{IP}}(J,\cdot,\cdot)$ and $\mathsf{F.D}^{\mathrm{IP}}(J,\cdot,\cdot)$ ever use in their oracle queries is limited to F.nk. Since a total of u identities is involved across the Enc and Dec queries we have

$$\Pr[\mathrm{G}_1 \text{ sets bad}] \le \frac{p + u \cdot \mathsf{F.nk}}{2^k} . \quad (13)$$

Putting together Equations (10)—(13) yields Eq. (9). □

The next lemma bounds the ib-prp advantage of a non-adaptive adversary $\mathcal{A}$ relative to $\overline{\mathsf{KDF}}$ via the prpa advantage of a constructed adversary $\overline{\mathcal{A}}$ against the FPE scheme F. This uses a hybrid argument, but done in such a way that the global loss from the structure of the key-derivation scheme remains linear (not quadratic) in the resources.

Game $\boxed{\mathrm{G}_0}$, G_1

$b \leftarrow\$ \{0,1\}$; $K \leftarrow\$ \{0,1\}^k$; $\mathrm{XI} \leftarrow \emptyset$; $\mathrm{ChI} \leftarrow \emptyset$
$P \leftarrow\$ \mathbf{IC}(k,k)$; $\overline{\pi} \leftarrow\$ \mathrm{Perm}(\{0,1\}^k)$
For every $I \in \mathsf{M.IS}$ do
  $J_{I,0} \leftarrow \overline{\pi}(\mathsf{M}_0(I))$; $J_{I,1} \leftarrow \overline{\pi}(\mathsf{M}_1(I))$; $J_I \leftarrow J_{I,0} \| J_{I,1}$
$b' \leftarrow\$ \mathcal{A}^{\text{Enc},\text{Dec},\text{Exp},\text{Ch},\text{IP}}$
Return $(b = b')$

$\text{Enc}(I, T, X)$

If $\mathrm{ET}[I,T,X] \neq \bot$ then return $\mathrm{ET}[I,T,X]$
If ($I \in \mathrm{ChI}$ and $b = 0$) then $Y \leftarrow\$ \{\, Y \in \mathsf{F.Dom} : \mathrm{DT}[I,T,Y] = \bot \,\}$
Else $Y \leftarrow \mathsf{F.E}^{\mathrm{IP}}(J_I, T, X)$
$\mathrm{ET}[I,T,X] \leftarrow Y$; $\mathrm{DT}[I,T,Y] \leftarrow X$; Return Y

$\text{Dec}(I, T, Y)$

If $\mathrm{DT}[I,T,Y] \neq \bot$ then return $\mathrm{DT}[I,T,Y]$
If ($I \in \mathrm{ChI}$ and $b = 0$) then $X \leftarrow\$ \{\, X \in \mathsf{F.Dom} : \mathrm{ET}[I,T,X] = \bot \,\}$
Else $X \leftarrow \mathsf{F.D}^{\mathrm{IP}}(J_I, T, Y)$
$\mathrm{ET}[I,T,X] \leftarrow Y$; $\mathrm{DT}[I,T,Y] \leftarrow X$; Return X

$\text{Exp}(I)$

If $I \in \mathrm{ChI}$ then return $\bot$
$\mathrm{XI} \leftarrow \mathrm{XI} \cup \{I\}$; Return J_I

$\text{Ch}(I)$

If $I \in \mathrm{XI}$ then return $\bot$
$\mathrm{ChI} \leftarrow \mathrm{ChI} \cup \{I\}$

$\mathrm{IP}(x)$

$(L, W, s) \leftarrow x$
If $(L = K)$ then
  bad $\leftarrow$ true ; $\boxed{P(K,\cdot,+) \leftarrow \overline{\pi} \text{ ; } P(K,\cdot,-) \leftarrow \overline{\pi}^{-1}}$
$y \leftarrow P(x)$; Return y

Figure 10: Games for proof of Lemma 6.2.

Lemma 6.3. *Let* F *be a square FPE scheme with* $\mathsf{F.kl} = 2k$. *Let* $\overline{\mathsf{KDF}}$ *be the key derivation function defined by Eq. (8). Let* $\mathcal{A}$ *be a non-adaptive adversary whose* Enc, Dec, Ch *queries involve at most* u *different identities, with at most* q_1 *queries to* Enc, Dec *per identity. Assume* $\mathcal{A}$ *makes* q_e *queries to* Exp *and* p *queries to* IP. *The proof constructs an adversary* $\overline{\mathcal{A}}$ *such that*

$$\mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\overline{\mathsf{KDF}}}(\mathcal{A}) \le u \cdot \mathbf{Adv}^{\text{prpa}}_{\mathsf{F}}(\overline{\mathcal{A}}) + \frac{8u + 8q_e - 6}{2^k} \ . \tag{14}$$

Adversary $\overline{\mathcal{A}}$ *makes at most* q_1 *queries to* Enc, Dec *and* p *queries to* IP. *Its running time is about the same as that of* $\mathcal{A}$. □

Proof. Let $N = u + q_e$. Let $I_1, \ldots, I_u$ denote the identities involved in $\mathcal{A}$'s Enc, Dec, Ch queries. Since $\mathcal{A}$ is non-adaptive, these are distinct from the identities, denoted $I_{u+1}, \ldots, I_N$, in $\mathcal{A}$'s Exp queries. To be more precise, since this is how the proof makes crucial use of the non-adaptivity assumption on $\mathcal{A}$, the sets $\{I_1, \ldots, I_u\}$ and $\{I_{u+1}, \ldots, I_N\}$ are disjoint.

We would like to use a hybrid argument in which I_g is viewed as the target for $\overline{\mathcal{A}}$. The difficulty is that the keys of different identities are not independent so we cannot simulate the keys of non-target identities without knowing the target key, and the latter is of course denied to us in the reduction. We could move to a game with random, independent keys, but this would result in an additive security loss involving terms like $N^2/2^k$. The following argument keeps the loss to $N/2^k$.

Consider games $\mathrm{G}_{2,g}, \mathrm{G}_{3,g}$ of Fig. 11, where $g \in [0..u]$ is an associated parameter. Here $J_i = J_{i,0} \| J_{i,1}$ is the key associated to I_i for $i \in [1..N]$. Rather than specify the keys via a permutation $\overline{\pi}$ as prescribed by KDF, we consider sampling them directly, meaning the $2N$ k-bit strings $J_{i,j}$ for $i \in [1..N]$ and $j \in \{0,1\}$ are random subject to being distinct. The games do this, but not quite. The key $J_g = J_{g,0} \| J_{g,1}$ is treated differently, being sampled uniformly at random, independently of other keys. If $J_{g,0}, J_{g,1}$ coincide with some other $J_{i,j}$ or with each other, the distribution is incorrect. Game $\mathrm{G}_{2,g}$, which includes the boxed code, corrects, re-sampling this key to obey the distinctness rule, but game $\mathrm{G}_{3,g}$, which excludes the boxed code, does not correct. The former reflects the real game, the latter the one conducive to doing our hybrid because non-target keys can be sampled without knowing the target key. (It is important that we did not overkill by asking all keys to be independent of each other in the second game, for this would incur the quadratic security loss.) While we have discussed J_i as associated to I_i, the identities to be queried are not known upfront, and the allocation of an index $\mathbf{v}[I]$ to identity I is made dynamically at the time identity I is first queried to Enc or Dec. Queries to Exp are answered directly, simply revealing the created keys. The games do not pick a challenge bit, instead returning true when the output b' of $\mathcal{A}$ is 1, and false otherwise. When $g = u$, all Enc, Dec queries are answered via F, and when $g = 0$ they are answered randomly but consistently with prior replies, so that

$$\mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) = \Pr[\mathrm{G}_{2,u}] - \Pr[\mathrm{G}_{2,0}] \ . \tag{15}$$

For each g, the two games $\mathrm{G}_{2,g}, \mathrm{G}_{3,g}$ are identical-until-bad, so by the Fundamental Lemma of Game Playing [11] we have

$$\begin{aligned} &\Pr[\mathrm{G}_{2,u}] - \Pr[\mathrm{G}_{2,0}] \\ &= \Pr[\mathrm{G}_{3,u}] + (\Pr[\mathrm{G}_{2,u}] - \Pr[\mathrm{G}_{3,u}]) \\ &\qquad - \Pr[\mathrm{G}_{3,0}] + (\Pr[\mathrm{G}_{3,0}] - \Pr[\mathrm{G}_{2,0}]) \\ &\le \Pr[\mathrm{G}_{3,u}] - \Pr[\mathrm{G}_{3,0}] + \Pr[\mathrm{G}_{3,u} \text{ sets bad}] \\ &\qquad + \Pr[\mathrm{G}_{3,0} \text{ sets bad}] \ . \end{aligned} \tag{16}$$

In game $\mathrm{G}_{3,g}$, the set D has size $2N - 2$ at the time of the test "$J_{g,0} \in D$," so bad is set here with probability $(2N-2)/2^k$. Similarly the test involving $J_{g,1}$ sets bad with probability at most $(2N-1)/2^k$, so

$$\forall\, g \in [0..u] \ : \ \Pr[\mathrm{G}_{3,g} \text{ sets bad}] \le \frac{4N-3}{2^k} \ . \tag{17}$$

Game $\boxed{G_{2,g}}$, $G_{3,g}$ $(0 \le g \le u)$

$P \leftarrow\!\!\$\ \mathbf{IC}(k,k)$; $\mathrm{ChI} \leftarrow \emptyset$; $N \leftarrow u + q_e$; $c \leftarrow 0$; $e \leftarrow u$
$D \leftarrow \emptyset$; $R \leftarrow \{0,1\}^k$
For $i \in [1..N] \setminus \{g\}$ and $j \in \{0,1\}$ do
  $J_{i,j} \leftarrow\!\!\$\ R$; $R \leftarrow R \setminus \{J_{i,j}\}$; $D \leftarrow D \cup \{J_{i,j}\}$
$J_{g,0} \leftarrow\!\!\$\ \{0,1\}^k$; $J_{g,1} \leftarrow\!\!\$\ \{0,1\}^k$
If $(J_{g,0} \in D)$ then bad $\leftarrow$ true ; $\boxed{J_{g,0} \leftarrow\!\!\$\ R \text{ ; } R \leftarrow R \setminus \{J_{g,0}\}}$
If $(J_{g,1} \in D \cup \{J_{g,0}\})$ then bad $\leftarrow$ true ; $\boxed{J_{g,1} \leftarrow\!\!\$\ R}$
$b' \leftarrow\!\!\$\ \mathcal{A}^{\textsc{Enc},\textsc{Dec},\textsc{Exp},\textsc{Ch},P}$
Return $(b' = 1)$

$\textsc{Enc}(I, T, X)$
If $\mathrm{ET}[I,T,X] \ne \bot$ then return $\mathrm{ET}[I,T,X]$
If $(\mathbf{v}[I] = \bot)$ then $c \leftarrow c+1$; $\mathbf{v}[I] \leftarrow c$
$J \leftarrow J_{\mathbf{v}[I],0} \| J_{\mathbf{v}[I],1}$
If $(I \in \mathrm{ChI}$ and $\mathbf{v}[I] > g)$ then $Y \leftarrow\!\!\$\ \{ Y \in \mathsf{F.Dom} : \mathrm{DT}[I,T,Y] = \bot \}$
Else $Y \leftarrow \mathsf{F.E}^P(J, T, X)$
$\mathrm{ET}[I,T,X] \leftarrow Y$; $\mathrm{DT}[I,T,Y] \leftarrow X$; Return Y

$\textsc{Dec}(I, T, Y)$
If $\mathrm{DT}[I,T,Y] \ne \bot$ then return $\mathrm{DT}[I,T,Y]$
If $(\mathbf{v}[I] = \bot)$ then $c \leftarrow c+1$; $\mathbf{v}[I] \leftarrow c$
$J \leftarrow J_{\mathbf{v}[I],0} \| J_{\mathbf{v}[I],1}$
If $(I \in \mathrm{ChI}$ and $\mathbf{v}[I] > g)$ then $X \leftarrow\!\!\$\ \{ X \in \mathsf{F.Dom} : \mathrm{ET}[I,T,X] = \bot \}$
Else $X \leftarrow \mathsf{F.D}^P(J, T, Y)$
$\mathrm{ET}[I,T,X] \leftarrow Y$; $\mathrm{DT}[I,T,Y] \leftarrow X$; Return X

$\textsc{Exp}(I)$
If $(\mathbf{v}[I] = \bot)$ then $e \leftarrow e+1$; $\mathbf{v}[I] \leftarrow e$
$J \leftarrow J_{\mathbf{v}[I],0} \| J_{\mathbf{v}[I],1}$; Return J

$\textsc{Ch}(I)$
If $(\mathbf{v}[I] = \bot)$ then $c \leftarrow c+1$; $\mathbf{v}[I] \leftarrow c$
$\mathrm{ChI} \leftarrow \mathrm{ChI} \cup \{I\}$

Adversary $\overline{\mathcal{A}}^{\textsc{Enc},\textsc{Dec},\textsc{Ch},\mathrm{IP}}$

$\mathrm{ChI} \leftarrow \emptyset$; $N \leftarrow u + q_e$; $c \leftarrow 0$; $e \leftarrow u$
$g \leftarrow\!\!\$\ [1..u]$; $R \leftarrow \{0,1\}^k$
For $i \in [1..N] \setminus \{g\}$ and $j \in \{0,1\}$ do
  $J_{i,j} \leftarrow\!\!\$\ R$; $R \leftarrow R \setminus \{J_{i,j}\}$
$J_{g,0} \leftarrow \bot$; $J_{g,1} \leftarrow \bot$
$b' \leftarrow\!\!\$\ \mathcal{A}^{\textsc{EncSim},\textsc{DecSim},\textsc{ExpSim},\textsc{ChSim},\mathrm{IP}}$
Return b'

$\textsc{EncSim}(I, T, X)$
If $\mathrm{ET}[I,T,X] \ne \bot$ then return $\mathrm{ET}[I,T,X]$
If $(\mathbf{v}[I] = \bot)$ then $c \leftarrow c+1$; $\mathbf{v}[I] \leftarrow c$
$J \leftarrow J_{\mathbf{v}[I],0} \| J_{\mathbf{v}[I],1}$
If $(I \in \mathrm{ChI}$ and $\mathbf{v}[I] > g)$ then $Y \leftarrow\!\!\$\ \{ Y \in \mathsf{F.Dom} : \mathrm{DT}[I,T,Y] = \bot \}$
If $(\mathbf{v}[I] = g)$ then $Y \leftarrow \textsc{Enc}(T, X)$
If $(I \notin \mathrm{ChI}$ or $\mathbf{v}[I] < g)$ then $Y \leftarrow \mathsf{F.E}^{\mathrm{IP}}(J, T, X)$
$\mathrm{ET}[I,T,X] \leftarrow Y$; $\mathrm{DT}[I,T,Y] \leftarrow X$; Return Y

$\textsc{DecSim}(I, T, Y)$
If $\mathrm{DT}[I,T,Y] \ne \bot$ then return $\mathrm{DT}[I,T,Y]$
If $(\mathbf{v}[I] = \bot)$ then $c \leftarrow c+1$; $\mathbf{v}[I] \leftarrow c$
$J \leftarrow J_{\mathbf{v}[I],0} \| J_{\mathbf{v}[I],1}$
If $(I \in \mathrm{ChI}$ and $\mathbf{v}[I] > g)$ then $X \leftarrow\!\!\$\ \{ X \in \mathsf{F.Dom} : \mathrm{ET}[I,T,X] = \bot \}$
If $(\mathbf{v}[I] = g)$ then $X \leftarrow \textsc{Dec}(T, Y)$
If $(I \notin \mathrm{ChI}$ or $\mathbf{v}[I] < g)$ then $X \leftarrow \mathsf{F.D}^P(J, T, Y)$
$\mathrm{ET}[I,T,X] \leftarrow Y$; $\mathrm{DT}[I,T,Y] \leftarrow X$; Return X

$\textsc{ExpSim}(I)$
If $(\mathbf{v}[I] = \bot)$ then $e \leftarrow e+1$; $\mathbf{v}[I] \leftarrow e$
$J \leftarrow J_{\mathbf{v}[I],0} \| J_{\mathbf{v}[I],1}$; Return J

$\textsc{ChSim}(I)$
If $(\mathbf{v}[I] = \bot)$ then $c \leftarrow c+1$; $\mathbf{v}[I] \leftarrow c$
If $(\mathbf{v}[I] = g)$ then $\textsc{Ch}(I)$
$\mathrm{ChI} \leftarrow \mathrm{ChI} \cup \{I\}$

Figure 11: Games and adversary for proof of Lemma 6.3.

Using Equations (15), (16) and (17), we have

$$\mathbf{Adv}^{\text{ib-prp}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) \le \Pr[G_{3,u}] - \Pr[G_{3,0}] + \frac{8N-6}{2^k} . \quad (18)$$

We use a hybrid argument to bound $\Pr[G_{3,u}] - \Pr[G_{3,0}]$. Consider adversary $\overline{\mathcal{A}}$ of Fig. 11. It picks g at random from $[0..u]$, and then picks keys $J_{i,j}$ for $i \ne g$ to be random but distinct. It then runs $\mathcal{A}$. It simulates the latter's $\textsc{Enc}, \textsc{Dec}, \textsc{Exp}, \textsc{Ch}$ oracles with the shown sub-routines $\textsc{EncSim}, \textsc{DecSim}, \textsc{ExpSim}, \textsc{ChSim}$, respectively. For IP, it directly uses its own IP oracle. The ability to do the latter is important and is why we needed Lemma 6.2 to remove all uses of the ideal primitive other than those made by F. In answering $\textsc{Enc}, \textsc{Dec}$ queries of $\mathcal{A}$ for an identity I_i, it uses F under the keys it has created if $i < g$, forwards the queries to its own $\textsc{Enc}, \textsc{Dec}$ oracles if $i = g$ —so that J_g is identified with the hidden key in game $\mathbf{G}^{\text{prp}}_{\mathsf{F}}(\overline{\mathcal{A}})$— and answers randomly if $i > g$, all this adjusted to take into account whether or not the identity is in ChI. A $\textsc{Exp}(I)$ query of $\mathcal{A}$ can be answered because $\mathbf{v}[I] \ne g$ so $\overline{\mathcal{A}}$ created, and has, the relevant key, and can return it. In answering a $\textsc{Ch}(I)$ query, $\overline{\mathcal{A}}$ calls its own $\textsc{Ch}$ oracle with I if $\mathbf{v}[I] = g$. We have

$$\mathbf{Adv}^{\text{prp}}_{\mathsf{F}}(\overline{\mathcal{A}}) = \frac{1}{u} \cdot \sum_{i=1}^{u} \Pr[G_{3,i}] - \Pr[G_{3,i-1}]$$
$$= \frac{1}{u} \left(\Pr[G_{3,u}] - \Pr[G_{3,0}] \right) . \quad (19)$$

Equations (18) and (19) imply Eq. (14). □

7 PRE-MASKING-BASED IB-FPE

While Theorem 6.1 shows that if we adjoin $\mathbf{Dbl}[k, \mathsf{M}]$ to an ideal square FPE F, the resulting IB-FPE scheme $(\mathsf{F}, \mathbf{Dbl}[k, \mathsf{M}])$ has k-bit ib-prp security, we'd like to have some provable guarantees if F is concretely instantiated from the base FPE scheme of DFF. However, while the Feistel structure of DFF seems to have very strong empirical security, it's notoriously hard to give even a satisfactory prp bound on Feistel networks on small domains. Let us now elaborate on the reason of this difficulty. Recall that in an information-theoretic proof for prp security of Feistel (such as the classic Luby-Rackoff result [28]), all current techniques can only give a bound based on the number of queries of the adversary, but not its running time. However, for a r-round balanced Feistel network on domain $\{0,1\}^{2n}$, by a simple counting argument, if $r < 2^n$, there *is* an adversary (of astronomical running time) that makes only 2^n ENC queries and wins with advantage very close to 1. But in our setting, n can be any number greater than 3, whereas in practice, r is often at most 36.

Given the huge obstacle in proving ib-prp security as described above, we turn into ib-kr-ti security. We now give a class of square FPE constructions that we call *pre-masking* FPE, such that for *any* F in this class, $(\mathsf{F}, \mathbf{Dbl}[k, \mathsf{M}])$ has nearly k-bit ib-kr-ti security. Members of this class use an ideal cipher $P \in \mathbf{IC}(k,k)$ (which will be instantiated via AES), but we make no other hardness assumption. This class includes the FPE scheme of DFF, and thus justifies the design choice of DFF. We warn that we only claim ib-kr security, and a pre-masking FPE therefore might be subject to different attacks. Thus our security guarantee here doesn't contradict the message-recovery attacks of BHT [7] on Feistel-based FPE schemes, including DFF. (These attacks, however, are easily put out of reach by increasing the number of rounds on small inputs.) Likewise, our security claim for pre-masking FPEs does not contract the recent message-recovery attack of Durak and Vaudenay [17] that exploits a bug in the design of round functions of FF3.

PRE-MASKING FPE. Let F be a square FPE scheme, meaning F has key-length $\mathsf{F.kl} = 2k$ and its ideal primitive is $\mathsf{F.IP} = \mathbf{IC}(k,k)$. We say that F is *pre-masking* if it additionally specifies algorithms F.EC, F.DC (we call them encode and decode) such that

$$\mathsf{F.E}^P(J,T,X) = \mathsf{F.EC}^{\mathsf{Round}^P(J,\cdot)}(T,X)$$
$$\mathsf{F.D}^P(J,T,Y) = \mathsf{F.DC}^{\mathsf{Round}^P(J,\cdot)}(T,Y)\ ,$$

where we have defined

$\mathsf{Round}^P(J,U)$
$J_1 \leftarrow J[1:k]\ ;\ J_2 \leftarrow J[k+1:2k]$
Return $P(J_1, U \oplus J_2, +)$.

That is, F.E and F.D use the $2k$-bit key J in a limited way, through Round. The latter takes a k-bit input and implements Rivest's classical DESX construction on top of the ideal-cipher instance P, but omits the post-whitening (meaning that the output is not XOR'ed with J_2). Note the encoding and decoding functions do not have direct access to the key J; they can only access it indirectly through queries to $\mathsf{Round}^P(J,\cdot)$. As an example, if $\mathsf{F.Dom} = \{0,1\}^{2n}$ and $\mathsf{F.TS} = \{0,1\}^t$ for $n + t \le k - 4$, a 10-round Feistel-based pre-masking FPE scheme can be built as in Fig. 12.

The efficiency improvement we obtain (due to dropping the post-whitening in DESX) is based on the fact that Round only calls the forward direction of the ideal cipher.

SECURITY ANALYSIS. As a stepping stone in obtaining the ib-kr-ti security of a pre-masking FPE scheme F, we consider security of the following FPE scheme $\overline{\mathsf{F}}$. Informally, the scheme $\overline{\mathsf{F}}$ is a blockcipher, implementing the DESX variant on top of AES. That is, FPE scheme $\overline{\mathsf{F}}$ has $\overline{\mathsf{F}}.\mathsf{Dom} = \{0,1\}^k$ and $\overline{\mathsf{F}}.\mathsf{TS} = \{\varepsilon\}$. Its encryption scheme $\overline{\mathsf{F}}.\mathsf{E}^P(J,T,X)$ returns $\mathsf{Round}(J,X)$, and the decryption scheme is defined accordingly.

$\mathsf{F.EC}^f(T,X)$
$L \leftarrow X[1:n]; R \leftarrow X[n+1:2n]$
$\ell \leftarrow k - n - t$
For $i = 1$ to 10 do
  $V \leftarrow f([i]_\ell \,\|\, T \,\|\, R)$
  $L' \leftarrow R; R \leftarrow L \oplus V[1:n]; L \leftarrow L'$
Return $L \,\|\, R$

$\mathsf{F.DC}^f(T,Y)$
$L \leftarrow Y[1:n]; R \leftarrow Y[n+1:2n]$
$\ell \leftarrow k - n - t$
For $i = 10$ downto 1 do
  $V \leftarrow f([i]_\ell \,\|\, T \,\|\, R)$
  $L' \leftarrow R; R \leftarrow L \oplus V[1:n]; L \leftarrow L'$
Return $L \,\|\, R$

Figure 12: A 10-round Feistel-based pre-masking FPE scheme F. Here $[i]_\ell$ denotes the ℓ-bit encoding of a number $i \in \{1,\ldots,10\}$. The oracle $f : \{0,1\}^k \to \{0,1\}^k$ is implemented as $\mathsf{Round}^P(J,\cdot)$.

In Lemma 7.1 below, we'll reduce the ib-kr-ti security of F to the ib-kr-ti security of $\overline{\mathsf{F}}$, both relative to $\mathbf{Dbl}[k,\mathsf{M}]$. The constructed adversary however makes no DEC query in attacking $\overline{\mathsf{F}}$. This restriction is crucial, because in $\overline{\mathsf{F}}$, there's pre-whitening but no post-whitening of the output of $P(J_1,\cdot,+)$.

LEMMA 7.1. *Let* F *be a pre-masking FPE scheme of* $\mathsf{F.kl} = 2k$ *and* $\overline{\mathsf{F}}$ *be as described above. Let* KDF *be the key-derivation function* $\mathbf{Dbl}[k,\mathsf{M}]$. *Suppose that we are given an adversary* $\mathcal{A}$ *whose* ENC/DEC *queries involve at most q calls to* Round. *Assume* $\mathcal{A}$ *makes* q_e *queries to* EXP *and p queries to* IP. *Then we can construct an adversary* $\overline{\mathcal{A}}$ *of the same number of* IP *and* EXP *queries such that*

$$\mathbf{Adv}^{\text{ib-kr-ti}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) \le \mathbf{Adv}^{\text{ib-kr-ti}}_{\overline{\mathsf{F}},\mathsf{KDF}}(\overline{\mathcal{A}})\ .$$

Adversary $\overline{\mathcal{A}}$ *makes at most q* ENC *queries and no* DEC *query.*

PROOF. Adversary $\overline{\mathcal{A}}$ runs $\mathcal{A}$ and shares the EXP and P oracles with it. When $\mathcal{A}$ wants to encrypt (I,T,X), adversary $\overline{\mathcal{A}}$ runs $\mathsf{F.EC}^{\text{ENC}(I,\varepsilon,\cdot)}(T,X)$, where ENC is $\overline{\mathcal{A}}$'s own encryption oracle. Likewise, when $\mathcal{A}$ wants to decrypt (I,T,Y), adversary $\overline{\mathcal{A}}$ runs $\mathsf{F.DC}^{\text{ENC}(I,\varepsilon,\cdot)}(T,Y)$. Finally, when $\mathcal{A}$ outputs its guessed key, adversary $\overline{\mathcal{A}}$ returns the same output. Hence game $\mathbf{G}^{\text{ib-kr-ti}}_{\overline{\mathsf{F}},\mathsf{KDF}}(\overline{\mathcal{A}})$ coincides with game $\mathbf{G}^{\text{ib-kr-ti}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A})$, and thus

$$\mathbf{Adv}^{\text{ib-kr-ti}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) \le \mathbf{Adv}^{\text{ib-kr-ti}}_{\overline{\mathsf{F}},\mathsf{KDF}}(\overline{\mathcal{A}})\ .$$

□

Next, we bound the ib-kr-ti security of $\overline{\mathsf{F}}$ relative to $\mathbf{Dbl}[k,\mathsf{M}]$, but the adversary is forbidden from calling DEC. The analysis is challenging, because there's no post-whitening of the output of $P(J_1,\cdot,+)$ in the encryption scheme of $\overline{\mathsf{F}}$, yet the adversary can still query $P(\cdot,\cdot,-)$. The proof is in the full version [6]. We note that if q is small, say $q \le 2^k/k^3$, then in Lemma 7.2 the blowup $k/\lg(k)$ can be reduced to $\frac{3k}{k-\lg(q)}$. However, for the practical choice $k = 128$, the blowup $k/\lg(k)$ is smaller than 19 and the bound in Lemma 7.2 is already satisfactory.

LEMMA 7.2. *Let* $\overline{\mathsf{F}}$ *be as described above and let* KDF *be the key-derivation function* $\mathbf{Dbl}[k, \mathsf{M}]$. *Assume that* $k \geq 16$. *Then for any adversary* $\mathcal{A}$ *that makes at most* $q \leq 2^{k-2}$ *queries to* ENC, *no query to* DEC, $q_e \leq 2^{k-3}$ *queries to* EXP, *and* p *queries to* IP,

$$\mathbf{Adv}^{\text{ib-kr-ti}}_{\overline{\mathsf{F}},\mathsf{KDF}}(\mathcal{A}) \leq \frac{2q(p+1)}{2^{2k}} + \frac{4(p+1)k}{2^k \cdot \lg(k)} + \frac{q+q_e+p+5}{2^k} \quad . \qquad \square$$

Combining Lemma 7.1 and Lemma 7.2, we immediately obtain the following result.

THEOREM 7.3. *Let* F *be a pre-masking FPE scheme of* $\mathsf{F.kl} = 2k$ *and let* KDF *be the key-derivation function* $\mathbf{Dbl}[k, \mathsf{M}]$. *Assume that* $k \geq 16$. *Suppose that we are given an adversary* $\mathcal{A}$ *whose* ENC/DEC *queries involve at most* q *calls to* Round. *Assume* $\mathcal{A}$ *makes* q_e *queries to* EXP *and* p *queries to* IP. *Then*

$$\mathbf{Adv}^{\text{ib-kr-ti}}_{\mathsf{F},\mathsf{KDF}}(\mathcal{A}) \leq \frac{2q(p+1)}{2^{2k}} + \frac{4(p+1)k}{2^k \cdot \lg(k)} + \frac{q+q_e+p+5}{2^k} \quad . \qquad \square$$

We note that the results of Lemma 7.2 and Theorem 7.3 hold even for adaptive adversaries if the ideal cipher is programmable. If the ideal cipher is non-programmable then these results only work for non-adaptive adversaries.

A MATCHING ATTACK. In Lemma 7.2, at the first glance, the blowup $k/\lg(k)$ looks like an artifact of the analysis. However, Proposition 7.4 shows that it's inherent by demonstrating a matching key-recovery attack. The proof, which is in the full version [6] is non-trivial. In both Lemma 7.2 and Proposition 7.4, the term $k/\lg(k)$ comes from some balls-into-bins phenomena.

PROPOSITION 7.4. *Let* $\overline{\mathsf{F}}$ *be as described above. Let* KDF *be the key-derivation function* $\mathbf{Dbl}[k, \mathsf{M}]$. *Assume that* $k \geq 128$. *Let* $r = \lfloor k/9\lg(k) \rfloor$ *and let* $q = r\lfloor 2^k/9r^2 \rfloor \approx 2^k \lg(k)/k$. *Then we can construct a non-adaptive adversary* $\mathcal{A}$ *making at most* $q+r$ *queries* ENC *queries and* $q+r$ *queries to* IP, *a single query to* CH, *and no query to* EXP *or* DEC *query, yet*

$$\mathbf{Adv}^{\text{ib-kr-ti}}_{\overline{\mathsf{F}},\mathsf{KDF}}(\mathcal{A}) \geq \frac{(1 - 5 \cdot 2^{-k/9})qr}{2^{k+1}} \geq \frac{1}{19} \quad . \qquad \square$$

8 SECURITY ANALYSIS FOR DFF

Here we discuss how to cast DFF as an IB-FPE scheme obtained via the **Dbl** transform and apply the results of Sections 6 and 7 to validate its security, as long as (1) the tweak (identity) space is appropriately restricted and (2) the radix and input length are fixed. Limitation (1) arises because, over the full tweak (identity) space, the M_1 embedding function is not injective: even for a fixed radix and input length, two tweaks may have derived keys with the same second halves. This does not, as far as we know, give rise to a damaging attack (we give below the best attack we could find) but it can be viewed as a design weakness. We suggest modifications to the embedding that restore injectivity and allow our results to apply. Limitation (2) means that a (master) key is used for just one choice of radix and tweak. To prove security for varying radix and input lengths would require that we use the broader definition of FPE from [10] in which the domain is a union of slices, in our case a slice being associated to a choice of radix and tweak.

DFF AS IB-FPE. We first briefly explain how to view DFF [36] as an IB-FPE scheme $(\mathsf{F}, \mathsf{KDF}_{\mathsf{dff}})$. (See the full version [6] for the complete specification.) The DFF specification allows different choices of radix rdx and input length n, but here we fix both, so that $\mathsf{F.Dom} = \mathbb{Z}^n_{\mathsf{rdx}}$. F has 256-bit keys and tweak space the singleton set $\{\varepsilon\}$. The algorithm itself is a 10-round Feistel network. The identity space $\mathsf{KDF}_{\mathsf{dff}}.\mathsf{IS}$ is the set of all $I \in \{0,1\}^*$ such that $|I|$ is at most 13 bytes. The underlying blockcipher $E : \{0,1\}^{128} \times \{0,1\}^{128} \to \{0,1\}^{128}$ is AES. Let $[x]^b$ denote the representation of x as a b-byte string. The embedding scheme $\mathsf{M} = (\mathsf{M}_0, \mathsf{M}_1)$ is specified via $\mathsf{M}_0(I) = [\mathsf{rdx}]^1 \parallel [|I|]^1 \parallel [n]^1 \parallel [I]^{13}$ and $\mathsf{M}_1(I) = [0]^3 \parallel [I]^{13}$. Note that M_1 is *not* injective: for example, $\mathsf{M}_1(00) = \mathsf{M}_1(000)$.

SECURITY OVER RESTRICTED IDENTITY SPACES. If the radix and input length are fixed, and one restricts the identities to a subset $S \subset \mathsf{KDF}_{\mathsf{dff}}.\mathsf{IS}$ such that no two strings in S correspond to the same integer in binary, then the embedding functions M_0 and M_1 above are injective and have disjoint images. Under these restrictions, our results in Sections 6 and 7 apply, and DFF has k-bit non-adaptive ib-prp security, and k-bit adaptive ib-kr-ti security.

SECURITY OVER THE FULL IDENTITY SPACE. The non-injectivity of M_1 allows an adversary to get the second half of the subkey of an identity I without querying $\mathrm{EXP}(I)$, by picking another identity I' such that $[I']^{13} = [I]^{13}$, and calling $\mathrm{EXP}(I')$. Note that for any $I' \in \mathsf{KDF}_{\mathsf{dff}}.\mathsf{IS}$, there are up to 104 other identities $I \in \mathsf{KDF}_{\mathsf{dff}}.\mathsf{IS}$ such that $[I']^{13} = [I]^{13}$. This leads to the following non-adaptive ib-kr-ti attack $\mathsf{KR}_{p,d}$. The adversary $\mathsf{KR}_{p,d}$ picks identities $I_0 = \varepsilon, I_1 = 0, I_2 = 0^2, \ldots, I_{104} = 0^{104}$. Note that $[I_0]^{13} = \cdots = [I_{104}]^{13}$. It first queries $J \leftarrow \mathrm{EXP}(I_0)$, and let $R \leftarrow J[k+1:2k]$. Note that for any $i \leq 104$, R is also the right half of the subkey of identity I_i. The adversary now picks p candidates subkeys $J_1, \ldots, J_p$ such that $J_j[k+1:2k] = L$. Now, for every $i \in \{1, \ldots, 104\}$ and $j \in \{1, \ldots, p\}$, it aims to test whether J_j is the subkey of I_i by comparing $\mathrm{ENC}(I_i, \varepsilon, \cdot)$ and $\mathsf{F.E}(J_j, \varepsilon, \cdot)$ on d messages. The code of the adversary is given in Fig. 13. Proposition 8.1 below shows that this attack achieves ib-kr-ti advantage about $104p/2^{129} - 104p \cdot \mathbf{Adv}^{\mathsf{fp}}_{\mathsf{F},d}$, where the false positive advantage $\mathbf{Adv}^{\mathsf{fp}}_{\mathsf{F},d}$ was defined in Section 3. The proof of Proposition 8.1 is in the full version [6].

PROPOSITION 8.1. *Let* $(\mathsf{F}, \mathsf{KDF}_{\mathsf{dff}})$ *be as above. Then for any* $p, d \in \mathbb{N}$ *such that* $p \leq 2^{128}/104$ *we have*

$$\mathbf{Adv}^{\text{ib-kr-ti}}_{\mathsf{F},\mathsf{KDF}_{\mathsf{dff}}}(\mathsf{KR}_{p,d}) \geq \frac{104p}{2^{129}} - 104p \cdot \mathbf{Adv}^{\mathsf{fp}}_{\mathsf{F},d} \quad . \qquad \square$$

DISCUSSION. While the attack $\mathsf{KR}_{p,d}$ above is impractical and does not significantly affect the 128-bit security claim of DFF, our results, at least, offer no proof that a better attack is not possible. Furthermore, that the right halves of the keys of two different identities can coincide does not feel right. Accordingly, we recommend fixing this. If rdx is fixed, this could be done by setting $\mathsf{M}_1(I) = [0]^1 \parallel [|I|]^1 \parallel [n]^1 \parallel [I]^{13}$. Alternatively one could restrict the identities as mentioned above. If rdx cannot be viewed as fixed and we

Adversary $\mathsf{KR}_{p,d}^{\textsc{Enc},\textsc{Dec},\textsc{Exp},\textsc{Ch},P}$

For $i \leftarrow 0$ to 104 do $I_i \leftarrow 0^i$
$J \leftarrow \textsc{Exp}(I_0)$; $R \leftarrow J[k+1 : : 2k]$; $S \leftarrow \emptyset$
For $\ell \leftarrow 1$ to d do $X_\ell \leftarrow\!\!\$\; \mathsf{F.Dom} \backslash S$; $S \leftarrow S \cup \{X_\ell\}$
For $i \leftarrow 1$ to 104 do
  For $\ell \leftarrow 1$ to d do $V_\ell \leftarrow \textsc{Enc}(I_i, \varepsilon, X_\ell)$
  $Z_i \leftarrow (V_1, \ldots, V_d)$
For $j \leftarrow 1$ to p do
  $L_j \leftarrow\!\!\$\; \{0,1\}^{128}$; $J_j \leftarrow L_j \,\|\, R$
  For $\ell \leftarrow 1$ do d do $U_\ell \leftarrow \mathsf{F.E}^P(J_j, T, X_\ell)$
  $Z \leftarrow (U_1, \ldots, U_\ell)$; $i \leftarrow \mathsf{Find}(Z, Z_1, \ldots, Z_{104})$
  If $i > 0$ then ($\textsc{Ch}(I_i)$; Return (I_i, J_j))

Figure 13: The attack $\mathsf{KR}_{p,d}$ on the IB-FPE scheme $(\mathsf{F}, \mathsf{KDF}_{\mathsf{dff}})$.

want a more natural space of identities, we would suggest to let identities be binary strings of at most 12 bytes, let $\mathsf{M}_0(I) = [0]^1 \| [\mathsf{rdx}]^1 \| [|I|]^1 \| [n]^1 \| [I]^{12}$ and $\mathsf{M}_1(I) = [1]^1 \| [\mathsf{rdx}]^1 \| [|I|]^1 \| [n]^1 \| [I]^{12}$. All these choices ensure the embedding functions satisfy our conditions so that our results in Sections 6 and 7 apply.

REFERENCES

[1] Michel Abdalla and Mihir Bellare. 2000. Increasing the Lifetime of a Key: a Comparative Analysis of the Security of Re-keying Techniques. In *ASIACRYPT 2000 (LNCS)*, Tatsuaki Okamoto (Ed.), Vol. 1976. Springer, Heidelberg, 546–559.

[2] Mihir Bellare, Alexandra Boldyreva, and Silvio Micali. 2000. Public-Key Encryption in a Multi-user Setting: Security Proofs and Improvements. In *EUROCRYPT 2000 (LNCS)*, Bart Preneel (Ed.), Vol. 1807. Springer, Heidelberg, 259–274.

[3] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. 1996. Pseudorandom functions revisited: The cascade construction and its concrete security. In *37th FOCS*. IEEE Computer Society Press, 514–523.

[4] Mihir Bellare, Anand Desai, Eric Jokipii, and Phillip Rogaway. 1997. A Concrete Security Treatment of Symmetric Encryption. In *38th FOCS*. IEEE Computer Society Press, 394–403.

[5] Mihir Bellare, Rafael Dowsley, Brent Waters, and Scott Yilek. 2012. Standard Security Does Not Imply Security against Selective-Opening. In *EUROCRYPT 2012 (LNCS)*, David Pointcheval and Thomas Johansson (Eds.), Vol. 7237. Springer, Heidelberg, 645–662.

[6] Mihir Bellare and Viet Tung Hoang. 2017. Identity-Based Format-Preserving Encryption. Cryptology ePrint Archive. (2017). Full version of this paper.

[7] Mihir Bellare, Viet Tung Hoang, and Stefano Tessaro. 2016. Message-Recovery Attacks on Feistel-Based Format Preserving Encryption. In *ACM CCS 16*, Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi (Eds.). ACM Press, 444–455.

[8] Mihir Bellare, Dennis Hofheinz, and Scott Yilek. 2009. Possibility and Impossibility Results for Encryption and Commitment Secure under Selective Opening. In *EUROCRYPT 2009 (LNCS)*, Antoine Joux (Ed.), Vol. 5479. Springer, Heidelberg, 1–35.

[9] Mihir Bellare, Ted Krovetz, and Phillip Rogaway. 1998. Luby-Rackoff Backwards: Increasing Security by Making Block Ciphers Non-invertible. In *EUROCRYPT'98 (LNCS)*, Kaisa Nyberg (Ed.), Vol. 1403. Springer, Heidelberg, 266–280.

[10] Mihir Bellare, Thomas Ristenpart, Phillip Rogaway, and Till Stegers. 2009. Format-Preserving Encryption. In *SAC 2009 (LNCS)*, Michael J. Jacobson Jr., Vincent Rijmen, and Reihaneh Safavi-Naini (Eds.), Vol. 5867. Springer, Heidelberg, 295–312.

[11] Mihir Bellare and Phillip Rogaway. 2006. The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs. In *EUROCRYPT 2006 (LNCS)*, Serge Vaudenay (Ed.), Vol. 4004. Springer, Heidelberg, 409–426.

[12] Eli Biham. 2002. How to decrypt or even substitute DES-encrypted messages in 2^{28} steps. *Inform. Process. Lett.* 84, 3 (2002), 117–124.

[13] John Black and Phillip Rogaway. 2002. Ciphers with Arbitrary Finite Domains. In *CT-RSA 2002 (LNCS)*, Bart Preneel (Ed.), Vol. 2271. Springer, Heidelberg, 114–130.

[14] Dan Boneh and Matthew K. Franklin. 2001. Identity-Based Encryption from the Weil Pairing. In *CRYPTO 2001 (LNCS)*, Joe Kilian (Ed.), Vol. 2139. Springer, Heidelberg, 213–229.

[15] Wei Dai, Viet Tung Hoang, and Stefano Tessaro. 2017. Information-theoretic Indistinguishability via the Chi-Squared Method. In *CRYPTO 2017*. Springer, 497–523.

[16] Anand Desai and Sara Miner. 2000. Concrete security characterizations of PRFs and PRPs: Reductions and applications. In *ASIACRYPT 2000 (LNCS)*, Tatsuaki Okamoto (Ed.), Vol. 1976. Springer, Heidelberg, 503–516.

[17] F. Betül Durak and Serge Vaudenay. 2017. Breaking and Repairing the FF3 Format Preserving Encryption over Small Domain. In *CRYPTO 2017*. Springer, 679–707.

[18] Cynthia Dwork, Moni Naor, Omer Reingold, and Larry J. Stockmeyer. 1999. Magic Functions. In *40th FOCS*. IEEE Computer Society Press, 523–534.

[19] Morris Dworkin. 2016. Recommendation for Block Cipher Modes of Operation: Methods for Format-Preserving Encryption. *NIST Special Publication 800-38G* (Mar. 2016). http://dx.doi.org/10.6028/NIST.SP.800-38G.

[20] Morris Dworkin and Ray Perlner. 2015. Analysis of VAES3 (FF2). Cryptology ePrint Archive, Report 2015/306. (2015). http://eprint.iacr.org/2015/306.

[21] Viet Tung Hoang, Ben Morris, and Phillip Rogaway. 2012. An Enciphering Scheme Based on a Card Shuffle. In *CRYPTO 2012 (LNCS)*, Reihaneh Safavi-Naini and Ran Canetti (Eds.), Vol. 7417. Springer, Heidelberg, 1–13.

[22] Dennis Hofheinz, Vanishree Rao, and Daniel Wichs. 2016. Standard Security Does Not Imply Indistinguishability Under Selective Opening. In *TCC 2016-B, Part II (LNCS)*, Martin Hirt and Adam D. Smith (Eds.), Vol. 9986. Springer, Heidelberg, 121–145. https://doi.org/10.1007/978-3-662-53644-5_5

[23] Tetsu Iwata. 2006. New Blockcipher Modes of Operation with Beyond the Birthday Bound Security. In *FSE 2006 (LNCS)*, Matthew J. B. Robshaw (Ed.), Vol. 4047. Springer, Heidelberg, 310–327.

[24] Tetsu Iwata, Bart Mennink, and Damian Vizár. 2016. CENC is Optimally Secure. Cryptology ePrint Archive, Report 2016/1087. (2016). http://eprint.iacr.org/2016/1087.

[25] Jooyoung Lee, Atul Luykx, Bart Mennink, and Kazuhiko Minematsu. 2017. Connecting tweakable and multi-key blockcipher security. *Designs, Codes and Cryptography* (Mar 2017). https://doi.org/10.1007/s10623-017-0347-3

[26] Moses Liskov, Ronald L. Rivest, and David Wagner. 2002. Tweakable Block Ciphers. In *CRYPTO 2002 (LNCS)*, Moti Yung (Ed.), Vol. 2442. Springer, Heidelberg, 31–46.

[27] Moses Liskov, Ronald L. Rivest, and David Wagner. 2011. Tweakable Block Ciphers. *Journal of Cryptology* 24, 3 (July 2011), 588–613.

[28] Michael Luby and Charles Rackoff. 1988. How to construct pseudorandom permutations from pseudorandom functions. *SIAM J. Comput.* 17, 2 (1988).

[29] Ben Morris and Phillip Rogaway. 2014. Sometimes-Recurse Shuffle - Almost-Random Permutations in Logarithmic Expected Time. In *EUROCRYPT 2014 (LNCS)*, Phong Q. Nguyen and Elisabeth Oswald (Eds.), Vol. 8441. Springer, Heidelberg, 311–326. https://doi.org/10.1007/978-3-642-55220-5_18

[30] Jacques Patarin. 2010. Introduction to Mirror Theory: Analysis of Systems of Linear Equalities and Linear Non Equalities for Cryptography. Cryptology ePrint Archive, Report 2010/287. (2010). http://eprint.iacr.org/2010/287.

[31] Thomas Ristenpart and Scott Yilek. 2013. The Mix-and-Cut Shuffle: Small-Domain Encryption Secure against N Queries. In *CRYPTO 2013, Part I (LNCS)*, Ran Canetti and Juan A. Garay (Eds.), Vol. 8042. Springer, Heidelberg, 392–409. https://doi.org/10.1007/978-3-642-40041-4_22

[32] Phillip Rogaway and Mihir Bellare. 2007. Robust computational secret sharing and a unified account of classical secret-sharing goals. In *ACM CCS 07*, Peng Ning, Sabrina De Capitani di Vimercati, and Paul F. Syverson (Eds.). ACM Press, 172–184.

[33] Adi Shamir. 1984. Identity-Based Cryptosystems and Signature Schemes. In *CRYPTO'84 (LNCS)*, G. R. Blakley and David Chaum (Eds.), Vol. 196. Springer, Heidelberg, 47–53.

[34] Thomas Shrimpton and R. Seth Terashima. 2016. Salvaging Weak Security Bounds for Blockcipher-Based Constructions. In *ASIACRYPT 2016, Part I (LNCS)*, Jung Hee Cheon and Tsuyoshi Takagi (Eds.), Vol. 10031. Springer, Heidelberg, 429–454. https://doi.org/10.1007/978-3-662-53887-6_16

[35] Joachim Vance. 2011. VAES3 scheme for FFX: An addendum to âĂIJThe FFX Mode of Operation for Format Preserving Encryption. Submission to NIST. (May 2011).

[36] Joachim Vance and Mihir Bellare. 2015. Delegatable Feistel-based Format Preserving Encryption Mode. Submission to NIST. (Nov 2015).

[37] Ping Zhang and Honggang Hu. 2016. On the Provable Security of the Tweakable Even-Mansour Cipher Against Multi-Key and Related-Key Attacks. Cryptology ePrint Archive, Report 2016/1172. (2016). http://eprint.iacr.org/2016/1172.