Salvaging Weak Security Bounds for Blockcipher-Based Constructions

Thomas Shrimpton[1,2(✉)] and R. Seth Terashima[1,2]

[1] Department of Computer and Information Science and Engineering, University of Florida, Gainesville, USA
teshrim@ufl.edu, setht@qti.qualcomm.com
[2] Qualcomm Technologies, Inc., San Diego, USA

Abstract. The concrete security bounds for some blockcipher-based constructions sometimes become worrisome or even vacuous; for example, when a light-weight blockcipher is used, when large amounts of data are processed, or when a large number of connections need to be kept secure. Rotating keys helps, but introduces a "hybrid factor" m equal to the number of keys used. In such instances, analysis in the ideal-cipher model (ICM) can give a sharper picture of security, but this heuristic is called into question when cryptanalysis of the real-world blockcipher reveals weak keys, related-key attacks, etc.

To address both concerns, we introduce a new analysis model, the ideal-cipher model under key-oblivious access (ICM-KOA). Like the ICM, the ICM-KOA can give sharp security bounds when standard-model bounds do not. Unlike the ICM, results in the ICM-KOA are less brittle to current and future cryptanalytic results on the blockcipher used to instantiate the ideal cipher. Also, results in the ICM-KOA immediately imply results in the ICM *and* the standard model, giving multiple viewpoints on a construction with a single effort. The ICM-KOA provides a conceptual bridge between ideal ciphers and tweakable blockciphers (TBC): blockcipher-based constructions secure in the ICM-KOA have TBC-based analogs that are secure under standard-model TBC security assumptions. Finally, the ICM-KOA provides a natural framework for analyzing blockcipher key-update strategies that use the blockcipher to derive the new key. This is done, for example, in the NIST CTR-DRBG and in the hardware RNG that ships on Intel chips.

1 Introduction

When a secret-key cryptographic primitive $\mathcal{E}$ is based upon a blockcipher E, a security proof for $\mathcal{E}$ will typically appeal to the pseudorandom-permutation (PRP) assumption—namely, that no efficient adversary can distinguish between the input-output behavior of the secretly (and randomly) keyed blockcipher E_K, and that of a truly random permutation π with the same domain. When the proof states that the PRP-security of E is a tight upperbound for the security of $\mathcal{E}$, one can derive from it useful messages for practice; e.g., how many calls to the blockcipher should be allowed before changing its key. When the upperbound

© International Association for Cryptologic Research 2016
J.H. Cheon and T. Takagi (Eds.): ASIACRYPT 2016, Part I, LNCS 10031, pp. 429–454, 2016.
DOI: 10.1007/978-3-662-53887-6_16

is not tight, the usefulness of any such messages can be unclear. In particular, when there is no known attack on the security of $\mathcal{E}$ whose success probability approaches the upperbound evidenced in the security proof. Such gaps are common when the security proof uses a "hybrid argument".

As an example, consider the following self-rekeying version of counter-mode encryption. (This is similar to the NIST CTR-DRBG [9] that underlies Intel's hardware RNG [11,19].) Let $\mathsf{CTR}[E]_K^N(\cdot)$ denote counter-mode encryption (over n-bit blockcipher E) under key K and IV N. The scheme is initialized with a key K_1 that is random. To encrypt the i-th plaintext X_i, the scheme computes ciphertext $C_i \leftarrow \mathsf{CTR}[E]_{K_i}^0(X_i)$ using key K_i, and then computes a key K_{i+1} for the next encryption call via $K_{i+1} \leftarrow \mathsf{CTR}[E]_{K_i}^{\lceil |X_i|/n \rceil + 1}(0^k)$. The standard proof would show that the security of this construction is (roughly) upperbounded by m times the probability violating the PRP-security of E, where m is the number of strings X_i that are encrypted before the key is reinitialized to a fresh random, secret value. Such a bound can quickly become vacuous when the underlying blockcipher is lightweight and cannot be assumed to provide PRP-security comparable to blockciphers like AES, or in settings where frequent re-initialization (i.e., resetting to a fresh, random K_1) is difficult.

If this construction is analyzed instead in the ideal cipher model (ICM), the upperbound is considerably tighter, and nearly matched by an attack. This suggests that the multiplicative factor of m in the standard-model result isn't "real", but rather an artifact of the proof technique. On the other hand ICM analysis provides only a security heuristic, and seems particularly inappropriate when the underlying blockcipher is known to have obvious non-ideal behavior for certain "weak" keys, or to suffer from related-key attacks.

Yet for constructions like this one, the presence of weak blockcipher keys is unlikely to be a real issue for the security of the construction: intuitively, if the initial key K is *random*, then so should be the derived keys that follow it. Analysis in the ICM naturally captures this intuition, as the key K_i is (essentially) independent of keys $K_1, K_2, \ldots, K_{i-1}$, and of the ciphertexts $C_1, C_2, \ldots, C_i$ that the construction outputs.

Moreover, observe that the construction doesn't actually need to know the *value* of any of the keys. It could carry out its duties if its access to E was via an API that restricted it to refer to keys by handles, e.g., ask $(i, x,$ "return") and receive $E_{K_i}(x)$ in return, or $(i, x,$ "key") and cause the value $K_{i+1} = E_{K_i}(x)$ to be stored, receiving nothing in return. We refer to such an API as enforcing *key-oblivious access* (KOA) to E, and under this access model it is clear that the construction leaks nothing about the keys beyond what the blockcipher does. Said another way, the access model supports the intuition that if the initial key K_1 is secret, it and its successors remain so.

The ICM under key-oblivious access. We formalize all of this in a new model, the ICM under key-oblivious access (ICM-KOA). The construction has black-box access to the blockcipher via, roughly, the API just described. On the other hand, the adversary may query the ideal cipher freely, as in the traditional

ICM, capturing a real-world attacker's ability to compute (offline) blockcipher input-output pairs under any key it likes. Before we give more details about our formalism, let us explain what benefits it provides.

First, the ICM-KOA retains the power of ICM to give sharper bounds than those found under the standard-model PRP assumption. It can also expose important quantitative security distinctions among variants of a given blockcipher-based construction, where these would be hidden by a standard-model analysis. This may help to guide implementation decisions in practice. We also surface in our model the distinction between precomputation queries to the blockcipher, offline queries made to the blockcipher while attacking the construction, and online queries made to the construction under its secret keys.

Second, security results in the ICM-KOA imply comparable security results in the traditional ICM *and* results in the standard-model. The latter is possible precisely because the model guarantees that the blockcipher is called on random and secret keys. Thus a single effort yields multiple viewpoints on a given construction.

Third, while security proofs in this model are still heuristics, their value is more resilient to the discovery of weak keys and related-key attacks on the real blockcipher that is idealized. In fact, the formalism provides a clear path to analyzing the security of constructions when the blockcipher is modeled with explicit non-ideal behaviors. We leave this as interesting future work.

Finally, the ICM-KOA provides a conceptual bridge between ideal ciphers and tweakable blockciphers (TBC). This is pleasing because, intuitively, the strong-tweakable-PRP assumption suggests that a secure, secretly keyed TBC is computationally indistinguishable from an ideal cipher—both provide a *set* of random permutations (one permutation for each tweak or key, respectively). We show that blockcipher-based constructions that are secure in the ICM-KOA have TBC-based analogs that are secure in the standard model.

Decomposing constructions into modes and schedulers. We want our model to facilitate results for blockcipher-based constructions that may use many keys. So the ICM-KOA requires that constructions can be decomposed into two primitives, a *mode* $\mathcal{M}$ and a potentially stateful key-*scheduler* $\mathcal{S}$. Intuitively, the role of the mode is to affect the transformation of construction-inputs (e.g., plaintexts) into construction-outputs (e.g., ciphertexts), and the role of the scheduler is to determine what keys the mode must use during its execution. Many symmetric-key cryptographic primitives can be decomposed in this way, including encryption schemes and blockcipher-based PRFs, PRNGs, KDFs and MACs, whether or not rekeying strategies are applied to them.

Returning to our self-rekeying version of counter-mode encryption, we might decompose this into a mode $\mathcal{M}$ that, on input a key K_i and a string X, computes $C \leftarrow \mathsf{CTR}[E]^0_{K_i}(X)$; and a scheduler $\mathcal{S}$ that (effectively) computes $K_{i+1} \leftarrow \mathsf{CTR}[E]^{\lceil |X|/n \rceil + 1}_{K_i}(0^k)$. Each will be forced to be oblivious of the actual key values by our model.

Applying the ICM-KOA to constructions. Given a blockcipher-based construction that admits decomposition, we define what it means for the construction to produce outputs that are indistinguishable from some reference-behavior-oracle in the ICM-KOA. To be clear, we do not claim that this is, on its own, an intuitive security goal. It is a new tool that provides a means to obtain strong bounds in the ICM that are backed by a guarantee that keys are kept random and secret. And because of this guarantee, we gain simultaneous results in the standard model. We illuminate the usefulness of the ICM-KOA via two case studies.

First we consider the NIST-CTR-DRBG. As the name suggests, it is a deterministic random-bit generator based on running a blockcipher in CTR mode. A result by Shrimpton and Terashima [19] shows that the standard-model security is around $q^2/2^k$, where q is the number of calls the construction. For $k = 128$, this bound exceeds 2^{-40} when $q = 2^{44}$. This may seem safe; after all, this amounts to many terabytes of random bits. But the RNG has extremely high throughput—Intel reports 800 MB/s, which equates to 50 million queries per second—meaning the $q = 2^{44}$ limit in a little more than four days.

We analyze this in the ICM-KOA. For very little work, we recover the security bound from [19], and also get a much stronger bound in the ICM. The latter reveals the lack of a matching attack and shows that, barring cryptanalysis of AES *under random and secret keys*, we can permit on the order of 2^{70} queries before surpassing our 2^{-40} limit (assuming the adversary has resources for 2^{80} precomputation and 2^{80} offline queries). This translates to 750,000 years of runtime, and so is unlikely to be the limiting factor.

Next we consider three rekeying variants of CTR-mode, distinguished by how they choose IVs following a key change: (1) The IV is set to 0^n; (2) the upper bits of the IV are unique for each key; (3) The IV is chosen randomly. In each case, we use the same key scheduler that sets $K_i \leftarrow E_{K_1}(i)$ (for $i > 1$). In the standard model, these three schemes all have the same security bound. Our analysis in the ICM-KOA uncovers significant quantitative differences their security bounds; in particular, we show how (1) succumbs to precomputation for shorter key lengths while (2) and (3) resist such attacks.

Addressing hybrid-loss directly in the standard model. Another, arguably more natural approach to avoiding a factor of m hybrid-loss when analyzing a blockcipher-based construction that uses m keys is to generalize the PRP notion to an m-PRP notion [18]. Here the adversary must distinguish between the collection of oracles $E_{K_1}(\cdot), E_{K_2}(\cdot), \ldots, E_{K_m}(\cdot)$ for random keys $K_1, \ldots, K_m$, and the collection $\pi_1(\cdot), \pi_2(\cdot), \ldots, \pi_m(\cdot)$ of random permutations. If a construction uses no more than m blockcipher keys during the time that it is being attacked, reducing the construction's security to the blockcipher's m-PRP security can be done without a hybrid proof, and therefore does not incur a factor of m loss.

But this may simply sweep problems under the rug: (1) it begs the question of how the m-PRP security of a given blockcipher relates to its PRP security (although we note that Hoang and Tessaro [12], building on the work of [18], have largely answered this question for key-alternating ciphers with independent

round keys) (2) it doesn't directly model interesting scenarios where the keys are themselves derived from the E using prior keys, particularly when, as with the NIST-RNG, the mode of operation is intertwined with key generation.

We explore this further in the full version of the paper. As one expects, the simplest result states that the m-PRP security of E falls somewhere between its PRP-security and m times that value. We go on to show that, under the assumption that a PRP-secure blockcipher E exists: (1) there is a related blockcipher for which these upper- and lowerbounds on its m-PRP security are tight; and (2) there is a related blockcipher that is PRP-secure but not m-PRP-secure, for sufficiently large values of m. (Of course, these distinctions are not binary, but the quantitative results are reasonable for modest m). These results are mainly of theoretical importance, as no real blockcipher will resemble the ones used to prove them.

But we also give a result that sheds some light on how much of a gap exists between any particular blockcipher's PRP security and m-PRP security. Given a PRP-adversary A for blockcipher E, the *best* m-PRP adversary $B[A]$ (that makes use of A in a black-box fashion) will have an advantage between $\mathbf{Adv}^{\mathrm{prp}}_{E}(A)$ and $m\mathbf{Adv}^{\mathrm{prp}}_{E}(A)$; moreover, its location on this continuum can be computed from $\mathbf{Adv}^{\mathrm{prp}}_{E}(A)$ and, interestingly, A's false-positive rate when distinguishing a keyed instance of E from a random permutation. When A's false-positive and false-negative rates are similar, then $B[A]$'s advantage scales with $\sqrt{m}$, rather than m. Again, see the full version of this paper for details.

Related Work. Abdalla and Bellare [1] were the first to rigorously study the security of rekeyed symmetric-encryption schemes, under various rekeying strategies. Concretely, they show that CBC-mode over an n-bit blockcipher, consistently rekeyed after $2^{n/3}$ blocks, can have meaningful security bounds up to about $2^{2n/3}$ total message blocks. (Specifically, they show that $2^{2n/3}$ one-block messages can be encrypted.) Our KOA modeling captures their rekeyed encryption schemes. As one example, they consider a rekeying strategy that computes $(K_{i+1}, L_{i+1}) = (E(L_i, 0), E(L_i, 1))$; we would say the scheduler $\mathcal{S}$ computes this (K_{i+1}, L_{i+1}), where L_i (resp. L_{i+1}) is the current (resp. next) scheduler state.

There are a number of works that analyaze secretly keyed constructions in the ICM. Kilian and Rogaway [14] proved that the DESX construction is a secure SPRP in the ICM. Dai et al. [10] leverage the ICM to prove the security of multiple encryption. Lee [17] uses the ICM to consider key-length extension offered by cascade encryption (aka multiple encryption) and xor-cascade encryption (of which DESX is a simple example). Recently there have been a line of nice papers on the security of key-alternating ciphers (aka xor-cascade encryption), including [2,7,8,15,16], that perform their analysis in the public-random-permutation model, which is derivative of the ICM. The randomized message-authentication code RMAC was analyzed in the ICM [13].

The classic "Luby-Rackoff Backwards" paper by Bellare, Krovetz and Rogaway [4] addresses the construction of beyond birthday-bound secure PRFs from PRPs, but they are unable to do so in the standard model because of hybrid

terms. Thus, their positive security results, which do show beyond-birthday-bound security of their constructions, are developed in the ICM, despite the presence of secret keys. It would be interesting to revisit their construction using the ICM-KOA.

Bellare, Boldyreva and Micali [3] consider multi-key security notions for public-key encryption, and show that, for left-or-right IND-CPA, the hybrid loss incurred by reducing from a multi-key instance to a single-key instance is inherent. Our discussion of the relationship between the PRP and m-PRP notions takes inspiration from that work, especially the construction of a cipher for which the bound is tight.

Bellare, Ristenpart and Tessaro [5] consider multi-instance (or multi-key) security notions, in which the attacker wins only if it breaks all of the instances. Their notions differ from ours, as it would suffice to break a single instance in our m-PRP notion.

Recent papers by Mouha and Luykx [18] and Hoang and Tessaro [12] consider the mutli-key security of key-alternating ciphers, demonstrating (in the random permutation model) that they do not suffer hybrid-like security losses. This work complements are own, which provides bounds for modes of operation that employ blockciphers with idealized behavior under random, secret keys.

Roadmap. Section 2 introduces the ICM with key-oblivious access. The central theorems are summarized up-front —that constructions (with certain properties) that are secure in the ICM-KOA are secure in both the ICM and standard models— and the bulk of the section is concerned with technical matters that support the formal theorem statements. The section ends by using the ICM-KOA framework to relate ideal ciphers and tweakable ciphers. Section 3 applies the results of Sect. 2 to various blockcipher-based constructions, including the NIST CTR-DRBG. Full proofs of all results are provided. Results on the relationship between the PRP and m-PRP standard-model notions will appear in the full version.

2 The ICM with Key-Oblivious Access

In this section, we formalize the notion of decomposing a construction into a mode (which carries out the cryptographic functionality) and a scheduler (which creates keys for the mode, as needed). We then define properties of modes and schedulers sufficient to imply results in both the standard model and the ICM. Roughly speaking:

- A mode and a scheduler constitute a *decomposition* of a construction if they preserve its black-box behavior.
- A mode is *compatible* with a scheduler if they query the underlying blockcipher on different points (and thus maintain an independence between keys and, e.g., ciphertexts).

– A decomposition has *dispersed inputs* if there are limits to how many blockcipher inputs an adversary can predict in advance.
– We quantify the computational resources consumed by the mode and scheduler using *mode efficiency*.

The first item and last items are straightforward, and the need for the second (in proofs) is intuitive after a moment's thought. Having dispersed inputs will help to make clear the impact of precomputation on security bounds. The coarser granularity of the standard model prevents it from benefiting from dispersed inputs, and we will demonstrate how this obscures the impact of precomputation.

The central theorems of this section, Theorems 1 and 2, have somewhat complicated statements. But, informally, they say the following:

Theorems 1 and 2, informally. If a decomposition (1) has these properties and (2) is difficult to distinguish from an appropriate reference oracle (e.g., an encryption oracle that returns random bits) when the underlying blockcipher is replaced by a random function that is inaccessible to the adversary, then the original construction is likewise hard to distinguish from the reference oracle *in both the standard model and in the ICM*.

We note that the "if" portion specifies indistinguishability when the blockcipher is treated as a random function that is inaccessible to the adversary. This isn't sweeping things under the rug: ICM-based proofs typically have to "decouple" the *actual* blockcipher used by the construction from the blockcipher available to the adversary using ad-hoc methods. Our informal theorem statement is merely surfacing this proof trick, and our model will allow us to enforce it cleanly.

The final significant contribution of this section is a result that uses the ICM-KOA framework to formalize a relationship between the ICM and TBCs.

2.1 Preliminaries

When X, Y are strings, $X \,\|\, Y$ is the concatenation of those strings, and $X \oplus Y$ is their bitwise exclusive-or. When $\mathcal{X}$ is a set, $X \xleftarrow{\$} \mathcal{X}$ means to sample uniformly from $\mathcal{X}$ and assign the result to X. When A is a randomized algorithm, then $X \xleftarrow{\$} A^{\mathcal{O}_1,\mathcal{O}_2,\cdots}(\sigma)$ means to provide A with oracle (black-box) access to $\mathcal{O}_1, \mathcal{O}_2, \ldots$ and input σ, and to assign the result of its execution to X. An *adversary* is a randomized algorithm. The notation $A^{\mathcal{O}_1,\mathcal{O}_2,\cdots} \Rightarrow b$ refers to the event that an algorithm A, when provided the indicated oracles (if any), ends its execution with output b.

Fix integers $k, n > 0$. A function family $E\colon \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$ is a blockcipher if, for all $K \in \{0,1\}^k$, the mapping $E_K(\cdot) = E(K,\cdot)$ is a permutation over $\{0,1\}^n$. We write $E_K^{-1}(\cdot)$ for the inverse of $E_K(\cdot)$. The set $\mathrm{Perm}(n)$ is the set of all permutations $\pi\colon \{0,1\}^n \to \{0,1\}^n$, and the set $\mathrm{BC}(k,n)$ is the set of all blockciphers $E\colon \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$.

If G is some game (in the sense of the game-playing framework of Bellare and Rogaway [6], where an adversary interacts with oracles) and $\mathcal{E}$ is some event, the notation $\Pr[\, G; \mathcal{C} \,]$ denotes the probability that the condition $\mathcal{C}$ will hold after G terminates.

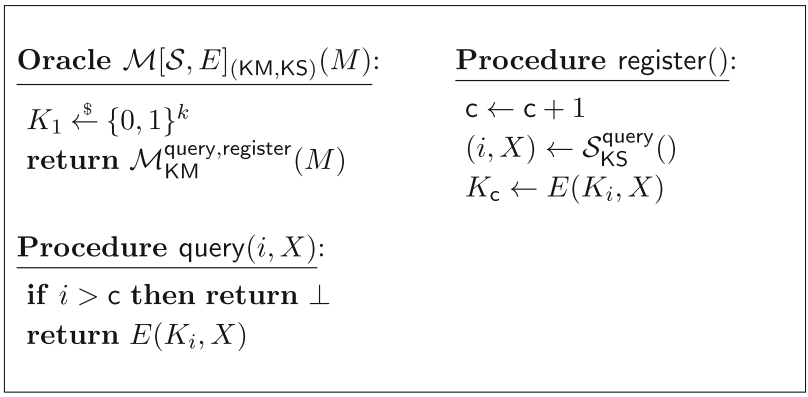

Fig. 1. A key-access manager exposes the query and register interfaces shown here. The oracle $\mathcal{M}[\mathcal{S},E]_{(\mathsf{KM},\mathsf{KS})}$, to which attackers will have oracle access in security experiments, uses these interfaces and a to implement the mode $\mathcal{M}$ of a given decomposition $\hat{\mathcal{E}} = (\mathcal{M},\mathcal{S},\mathcal{K})$. Here, c is initially 1.

2.2 Decompositions and Their Associated Notions

Let $\mathcal{E} : \mathcal{K}_{\mathcal{E}} \times \mathcal{D} \to \mathcal{R}$ be some scheme (e.g., CTR mode) that makes black-box use of a blockcipher $E : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$. We write $\mathcal{E}^E_K$ for the construction being keyed by $K \in \mathcal{K}_{\mathcal{E}}$, with E as a superscript to emphasize black-box access.

Our goal is to break $\mathcal{E}$ into a mode of operation and a key scheduler. A *decomposition* is a tuple $\hat{\mathcal{E}} = (\mathcal{M},\mathcal{S},\mathcal{K})$ of algorithms: a *mode* $\mathcal{M} : \mathcal{K}_{\mathcal{M}} \times \mathcal{D} \to \mathcal{R}$, a stateful but deterministic *scheduler* $\mathcal{S} : \mathcal{K}_{\mathcal{S}} \to \mathbb{N} \times \{0,1\}^n$, and a key-generation algorithm $\mathcal{K}$ that outputs values in $\mathcal{K}_{\mathcal{M}} \times \mathcal{K}_{\mathcal{S}}$. The mode $\mathcal{M}$ expects two oracles having the signatures of query and register, which are exposed as part of a *key-access manager* in Fig. 1. (Look ahead to World 1 of Fig. 3 for an illustration). The scheduler $\mathcal{S}$ expects oracle access to query, and is invoked by register.

A natural first attempt at defining key-oblivious access to an ideal cipher E would be to choose set of keys $K_1, K_2, \ldots, K_m$ up front, and then give the mode $\mathcal{M}$ (e.g., CTR mode) being analyzed black-box access to some oracle $\mathcal{O}(i,X) := E(K_i,X)$ for $i \in [1..m]$. There would be no explicit scheduler, and the keys themselves would be independent of the blockcipher E. But we want to capture schemes that do use E to derive the keys. For example, the Intel RNG [11] and the Abdalla and Bellare [1] constructions mentioned in the introduction. Hence we surface a key scheduler $\mathcal{S}$ as an explicit component of the decomposition, and must provide it with some kind of access to E. We cannot provide $\mathcal{S}$ *unfettered* access to E, however. If we did, then we would not be able to argue that E is queried only under random (and secret) keys. Concretely, suppose $\mathcal{S}$ sets $K_i = E(C, E(C, K \oplus i))$, where C is some constant and K is some "master key"; this may be secure in the ICM, but if we instantiate E with DES and C is a one of the weak keys for DES, then we would have $K_i = K \oplus i$. The keys used by the mode of operation would be closely related, a scenario we wish to

Table 1. Symbols used in ICM-KOA security definitions.

Symbol	Upperbound for number of...
q	Adversary queries
m	Blockcipher keys used
σ	n-bit blocks per adversary query
μ	Key aliases used to encipher any given block
ν	Blocks enciphered using any given key alias

preclude. Thus we restrict the scheduler's access to E. Similar abuse from $\mathcal{M}$ must also be prevented.

The oracles in our key-access manager force both $\mathcal{S}$ and $\mathcal{M}$ to query the blockcipher via handles, values that are independent of the particular values of the keys. Moreover, when preparing to have a value assigned to the $m\mathsf{th}$ key K_m, the scheduler $\mathcal{S}$ can only request outputs of E under keys K_1 through $K_{(m-1)}$. Note that $\mathcal{S}$ is not allowed to "know" the resulting value of K_m: instead, $\mathcal{S}$ outputs a pair (i, X) and K_m is assigned $E(K_i, X)$. We also force $\mathcal{M}$ to query E using handles for keys.

We note that the syntax for both the mode $\mathcal{M}$ and the scheduler $\mathcal{S}$ provides them with what appear to be "master" keys KM and KS. This is to capture initial values (keys, IVs, etc.) provided to the blockcipher-based construction. We will not assume or demand that KM and KS are independent of each other, but allowing them to be distinct permits us to capture more general constructions.

Definition 1 (Decompositions of schemes). *Let $\mathcal{E} : \mathcal{K}_\mathcal{E} \times \mathcal{D} \to \mathcal{R}$ and $\hat{\mathcal{E}} = (\mathcal{M}, \mathcal{S}, \mathcal{K})$ be defined as above. For $K \in \mathcal{K}_\mathcal{M} \times \mathcal{K}_\mathcal{S}$, let $\mathcal{M}[\mathcal{S}, E]_K : \mathcal{D} \to \mathcal{R}$ be the procedure defined in Fig. 1; this procedure combines the mode of operation $\mathcal{M}$ with the key scheduler $\mathcal{S}$ and blockcipher E in the natural way. We say $\hat{\mathcal{E}}$ is a* faithful decomposition *of $\mathcal{E}$ if, for any adversary A and any $E \in \mathrm{BC}(k, n)$, $k = n$, $\Pr\left[A^{\mathcal{E}^E_{K'}, E, E^{-1}} \Rightarrow 1\right] = \Pr\left[A^{\mathcal{M}[\mathcal{S},E]_K, E, E^{-1}} \Rightarrow 1\right]$. The probabilities are over the choice of $K' \xleftarrow{\$} \mathcal{K}_\mathcal{E}$, $K \xleftarrow{\$} \mathcal{K}$ and the coins of A, $\mathcal{M}$, and $\mathcal{E}$.*

That is, the black-box behavior of $\mathcal{E}^E_{K'}$ must be identical to the black-box behavior of $\mathcal{M}[\mathcal{S}, E]_K$ (given the above distribution of keys) *for any blockcipher E* and *computationally unbounded* adversaries.

Note that by using blockcipher outputs as keys, this definition assumes for the sake of simplicity that the key size k is equal to the blocksize n (each key is the output of the blockcipher at some point). We note that our model could easily be extended to the case where $k \neq n$ by truncating or concatenating the keys produced, as required, at the expense of complicating notation. However, we will use both k and n in our definitions and security bounds in order to suggest how taking $k \neq n$ would impact our model and results.

Compatible modes. Our key-access manager formalism does not itself prevent a scheduler $\mathcal{S}$ from "cheating" by choosing non-random keys. For example, $\mathcal{S}$

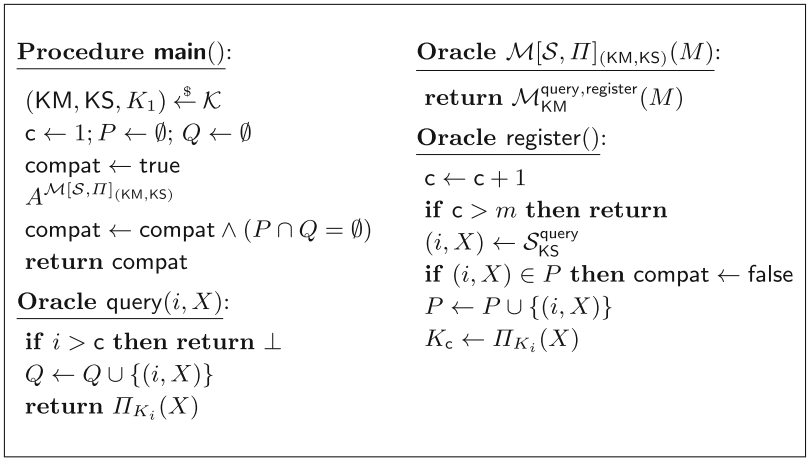

Fig. 2. Procedures and oracles for **Experiment** $\mathsf{COMPAT}^{\Pi}_{\hat{\mathcal{E}}}(A)$, where $\hat{\mathcal{E}} = (\mathsf{M}, \mathsf{S}, \mathcal{K})$. A mode $\mathcal{M}$ is m-compatible with a scheduler $\mathcal{S}$ if neither one queries the blockcipher on a point used to generate one of the first m keys.

could use its query oracle to search for a point (i, X) such that $E(K_i, X)$ ends in a zero, then output that point.

Informally, a scheduler $\mathcal{S}$ is *compatible* with a mode $\mathcal{M}$ if no adversary can cause either $\mathcal{S}$ or $\mathcal{M}$ to invoke query at a point (i, X) used to generate a key $K_j = E(K_i, X)$. This ensures that both the $\mathcal{S}$ and $\mathcal{M}$ are oblivious to the actual values of each key.

We'll show that as long as each key alias i is used significantly fewer than $2^{n/2}$ times, it follows that in both the ICM and the standard model there will be enough (computational) randomness in $E(K_i, X)$ for use as a cryptographic key. (This restriction results from the birthday paradox: since E is being used to generate keys, we need it to behave like a random function, rather than random permutation.)

Definition 2 (Compatible modes). *Let $\hat{\mathcal{E}} = (\mathcal{M}, \mathcal{S}, \mathcal{K})$ be a decomposition over an (k, n)-bit blockcipher, $k = n$, and set $K \xleftarrow{\$} \mathcal{K}$. Let m be a positive integer. Then $\mathcal{S}$ is* m-compatible *with $\mathcal{M}$ (with respect to $\mathcal{K}$) if for any keyed function $\Pi : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$, and any adversary A,* $\Pr\left[\mathsf{COMPAT}^{\Pi}_{\hat{\mathcal{E}}}(A) \Rightarrow \mathsf{true}\right] = 1$, *where Experiment* COMPAT *is defined in Fig. 2.*

Note that Π need not be a blockcipher. This generality is required to make some of our later reductions work, and does not appear to exclude interesting modes.

Some other, arguably more natural definitions fail to capture our goal of preventing cheating schedulers. For example, suppose we instead query $\mathcal{S}_{\mathsf{KS}}$ to obtain keys $(K_1, K_2, \ldots, K_m)$ and require that no adversary with access to E and E^{-1} be able to distinguish these keys from truly random values. This definition proves *too* strict, as it excludes schedulers that deterministically derive K_{i+1} from K_i.

It may then be tempting to instead allow schedulers to output keys directly (rather than (i, X) pairs), and task an adversary A to distinguish $\mathcal{M}[\mathcal{S}, E]_{(\mathsf{KM},\mathsf{KS})}$ from $\mathcal{M}[\$, E]_{\mathsf{KM},\mathsf{KS}}$, where $\$$ is a special oracle that samples and returns fresh random strings from $\{0,1\}^k$ on each invocation. This hides the keys from being directly observed by A, allowing K_{i+1} to depend on K_i deterministically. Such a definition, however, is too weak—it doesn't really depart from the familiar ICM. For example, if $\mathcal{S}_{\mathsf{KS}}$ sets $K_i = \mathsf{KS} \oplus i$ then the keys are not independent, yet A is unlikely to be able to exploit this (in the ICM). One of our goals is that our security definition should imply security in the standard model, so this candidate also isn't acceptable.

Dispersed inputs. The next two definitions are used to measure some important combinatorial properties of decompositions. We will require several symbols to define the relevant parameters, and so provide Table 1 for reference.

Definition 3 (Dispersed inputs). *Let k, n, μ and σ be non-negative integers, and let ϵ be positive. Let F be a uniformly random function mapping $\{0,1\}^k \times \{0,1\}^n$ to $\{0,1\}^n$. A decomposition $\hat{\mathcal{E}}$ over an (n,n)-bit blockcipher has $(q, \sigma, \mu, \epsilon)$-dispersed inputs if for any adversary A making q queries, each no longer than σn bits,*

$$\Pr\left[\, \mathsf{COMPAT}^{F}_{\hat{\mathcal{E}}}(A)\,;\, \max_{X} |\{i \mid (i,X) \in Q\}| > \mu \,\right] < \epsilon,$$

where Experiment COMPAT *is defined in Fig. 2, and Q refers to the final value of the set so named constructed during this experiment (i.e., the set of points submitted to the* query *oracle).*

The condition states that no single input is evaluated under more than μ key aliases except with probability ϵ. Small values of μ and ϵ limit the effectiveness of brute-force attacks by putting a cap on how many of the m keys can be attacked in parallel with a single blockcipher invocation.

Mode efficiency. A final definition is used to bound the computational work done by $\mathcal{M}$ and $\mathcal{S}$ given restrictions on an adversary.

Definition 4 (Mode efficiency). *Let $\hat{\mathcal{E}}$ be a decomposition over an (k,n)-bit blockcipher E, with $k = n$. Let* COMPAT *be the experiment defined in Fig. 2, and let A be any adversary making q queries, each of length at most σn bits. We say $\hat{\mathcal{E}}$ is (q, σ, m, ν)-efficient if after an execution of $\mathsf{COMPAT}^{E}_{\hat{\mathcal{E}}}(A)$, $\mathsf{c} < m$ and for each i, $|\{X \mid (i,X) \in P \cup Q\}| \leq \nu$. Here, c, P, and Q refer to the final values of the random variables constructed in the experiment's definition.*

That is, given such an adversary, the mode and scheduler will query the key manager using at most m key aliases, and will use each alias to encipher at most ν blocks.

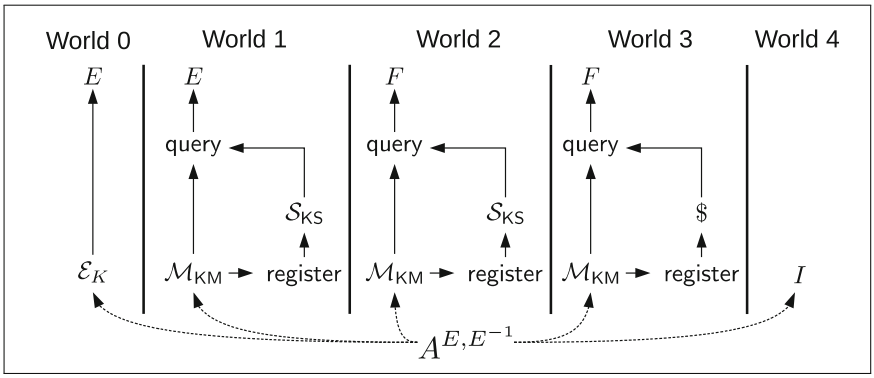

Fig. 3. Here, F is an ideal cipher and $\mathcal{E}$ is some cryptographic scheme based on a (concrete) blockcipher E that should be indistinguishable from some reference oracle $\mathcal{I}$. For example, $\mathcal{E}$ maybe an encryption scheme and $\mathcal{I}$ an oracle that returns a random string. From A's perspective, World 0 = World 1 if $\hat{\mathcal{E}} = (\mathcal{M}, \mathcal{S}, \mathcal{K})$ is a *decomposition* of $\mathcal{E}$; World 1 $\approx$ World 2 if $\hat{\mathcal{E}}$ has *dispersed inputs* and E is a PRP; World 2 $\approx$ World 3 if the scheduler $\mathcal{S}$ is *compatible* with the mode $\mathcal{M}$; World 3 $\approx$ World 4 if $\hat{\mathcal{E}}$ is indistinguishable from $\mathcal{I}$ in the ICM-KOA.

2.3 Generic Results About IND-KOA-ICM

We can now define what it means for a construction $\mathcal{E}$ to be indistinguishable from a reference oracle $\mathcal{I}$ in the ICM-KOA, the ICM, and the standard model. In general, we're interested in $\mathcal{I}$ that provide the desired idealized behavior of $\mathcal{E}$. For example, if $\mathcal{E}$ is an encryption algorithm, then we may want $\mathcal{I}$ to be the oracle that accepts a plaintext and outputs random bits.

We then show that ICM-KOA indistinguishability implies insecurity in both the ICM and the standard model, with a loss that is determined by the parameters of $\mathcal{E}$'s decomposition as surfaced by the efficiency and input-dispersion definitions. Figure 3 provides a graphical overview of how our key-access manager formalism will be used to argue indistinguishability of $\mathcal{E}$ and $\mathcal{I}$.

We emphasize that unlike most security definitions of this form, we do not claim that ICM-KOA indistinguishability offers an intuitive, compelling security goal on its own. Instead, it is a means to obtaining strong bounds in the ICM that are backed by a guarantee that keys are kept random and secret. And because of this guarantee, we gain simultaneous results in the standard model.

Definition 5 (ICM-KOA indistinguishability). *Let $\hat{\mathcal{E}} = (\mathcal{M}, \mathcal{S}, \mathcal{K})$ be a decomposition over an (k, n)-bit blockcipher, $k = n$, with $\mathcal{M}[\mathcal{S}, E]_K : \mathcal{D} \to \mathcal{R}$. Let $\mathcal{I} : \mathcal{D} \to \mathcal{R}$ be some reference scheme. Then the ICM-KOA-$\mathcal{I}$ advantage of an adversary A is*

$$\mathbf{Adv}^{koa\text{-}ind\text{-}\mathcal{I}}_{\hat{\mathcal{E}}}(A) = \Pr\left[A^{\mathcal{M}[F]_K, E, E^{-1}} \Rightarrow 1 \right] - \Pr\left[A^{\mathcal{I}, E, E^{-1}} \Rightarrow 1 \right].$$

Here, $F \xleftarrow{\$} Func(k+n, n)$ *and* $\mathcal{M}[F]_K$ *behaves identically to* $\mathcal{M}[\mathcal{S}, F]_K$ *(as defined in Fig. 1), except* register *assigns* $K_{\mathsf{c}} \xleftarrow{\$} \{0,1\}^k$ *instead of* $K_{\mathsf{c}} \leftarrow E_{K_i}(X)$.

Note that in this definition, the mode $\mathcal{M}$ does not interact with E, and so, without loss of generality, neither does A. ICM-KOA indistinguishability is only a useful notion for compatible decompositions with dispersed inputs, as these properties will allow us to "decouple" the ideal cipher used by the mode from the ideal cipher directly accessible by an adversary when proving results in the ICM.

Definition 6 (ICM indistinguishability). *Let* $\hat{\mathcal{E}} = (\mathcal{M}, \mathcal{S}, \mathcal{K})$ *be a decomposition over an* (k,n)*-bit blockcipher,* $k = n$*, where* $\mathcal{M}[\mathcal{S}, E]_K : \mathcal{D} \to \mathcal{R}$*. Let* $\mathcal{I} : \mathcal{D} \to \mathcal{R}$ *be some reference scheme (for example, an encryption algorithm with* $\mathcal{D} = \mathcal{R} = \{0,1\}^*$*). Then the ICM-IND-*$\mathcal{I}$ *advantage of an adversary* A *is*

$$\mathbf{Adv}^{icm\text{-}ind\text{-}\mathcal{I}}_{\hat{\mathcal{E}}}(A) = \Pr\left[A^{\mathcal{M}[\mathcal{S},E]_K, E, E^{-1}} \Rightarrow 1 \right] - \Pr\left[A^{\mathcal{I}, E, E^{-1}} \Rightarrow 1 \right],$$

where $K \xleftarrow{\$} \mathcal{K}$*, and* $E \xleftarrow{\$} \mathrm{BC}(k,n)$ *is an ideal cipher.*

Precomputation, offline and online queries. One benefit of the ICM-KOA model is that it can quantify the effectiveness of precomputation against specific modes. The following definition is general, but in it we have in mind $f_2 = E$, $f_3 = E^{-1}$ for some blockcipher E, while f_1 is an oracle for some blockcipher-based construction.

Definition 7 (Precomputation, offline, and online queries). *Let* A^{f_1, f_2, f_3} *be an adversary. We say* A *makes* q_P precomputation *queries,* q_E offline *queries, and* q online *queries if*

- A *makes* q_P *combined queries to* f_2 *and* f_3 *before making its first query to* f_1*,*
- *and afterwards makes a combined* q_E *queries to* f_2 *and* f_3*,*
- *while interleaving* q *queries to* f_1*.*

Relating the ICM-KOA and the ICM. We now give the first of our two main model-implication results. Namely, that security in the ICM-KOA implies security in the ICM.

Theorem 1 (ICM-KOA indistinguishability implies ICM indistinguishability). *Let* $\hat{\mathcal{E}} = (\mathcal{M}, \mathcal{S}, \mathcal{K})$ *be a decomposition over an* (k,n)*-bit blockcipher with* $k = n$*, and let* $\mathcal{I}$ *be some reference scheme. Fix a positive integer* c*. Let* A *be an adversary making* q_P *precomputation queries,* q_E *offline queries, and* q *online queries, the latter of at most* σn *bits each. Suppose*

1. $\mathcal{M}$ *is compatible with* $\mathcal{S}$*,*
2. $\hat{\mathcal{E}}$ *is* (q, σ, m, ν)*-efficient,*
3. $\hat{\mathcal{E}}$ *has* $(q, \sigma, \mu, \epsilon)$*-dispersed inputs, and*
4. *For any adversary* B *making* q *queries,* $\mathbf{Adv}^{koa\text{-}ind\text{-}\mathcal{I}}_{\hat{\mathcal{E}}}(B) \leq \delta$*.*

Further suppose[1] *that* $q_E + q_P < 2^n$. *Then*

$$\mathbf{Adv}^{icm\text{-}ind\text{-}\mathcal{I}}_{\hat{\mathcal{E}}}(A) \leq \delta + \frac{2q_E c\nu}{2^k(2^n - q_E - q_P)} + \frac{(q_E + q_P)m\nu}{2^{k+n}} + \frac{cm\nu^2}{2^n} + \frac{q_E(2\mu + c) + (q_P + m)\mu}{2^k} + \frac{m^{c+1}(1 + \nu^{c+1})}{2^{nc}(c+1)!} + 3\epsilon.$$

Although this general bound is complex, it simplifies substantially for various modes of operation. We will see this when we apply the general result to real constructions in Sect. 3. We note that the constant c can be chosen more-or-less arbitrarily to minimize the bound. This permits the possibility of "beyond birthday-bound security" when $c > 1$. (The $cm\nu^2/2^n$ term gives a birthday bound with respect to the amount of data ν processed with a single key, but $m\nu$ blocks are enciphered in total.) Before proving this theorem, we give the following useful lemma.

Lemma 1. (c-wise birthday bound). *Let c, q, and n be positive integers, with $c \leq q$. Let $X_1, \ldots, X_q$ be iid uniformly random n-bit strings. Then* $\Pr[\, \exists S \subseteq \{1, \ldots, q\}$ *s.t.* $|S| = c, X_j = X_i$ *for all* $i, j \in S \,] \leq \frac{q^c}{2^{n(c-1)}c!}$.

Proof. Fix some $x \in \{0,1\}^n$ and some c-sized index set $S \subseteq \{1, 2, \ldots, q\}$. Then $\Pr[\, \forall i \in S : x = X_i \,] = 2^{-cn}$. Since there are 2^n choices for x and $\binom{q}{c} < q^c/c!$ choices for S, a union bound provides us with the desired upper bound. □

Proof (Theorem 1). Let $F \xleftarrow{\$} \mathrm{Func}(k + n, n)$. Then $\Pr\left[\, A^{\mathcal{M}[F]_K} \Rightarrow 1 \,\right] - \Pr\left[\, A^{\mathcal{I}} \Rightarrow 1 \,\right] \leq \delta$, where $K \xleftarrow{\$} \mathcal{K}$ and $\mathcal{M}[F]_K$ is defined as Definition 5.

Game G1(A) (Fig. 4), which excludes the boxed statements, faithfully simulates $A^{\mathcal{M}[\mathcal{S},F]_K}$. In this figure, and for the remainder of the proof, F, E, and E^{-1} (without subscripts) refer to oracles, while F_K and E_K (with subscripts) refer to the lazily-defined functions the game builds to help implement these oracles. We've moved the calls to register to the start of the game, without loss of generality.

In G1(A), the behavior of F is independent of the behavior of E and E^{-1}. Consequently, the value of each key K_i is information theoretically hidden from the adversary; the adversary can at best learn information about whether two key aliases correspond to the same key.

Recall that the difference between $\mathcal{M}[F]_K$ and $\mathcal{M}[\mathcal{S}, F]_K$ is that the former's register procedure always assigns keys a uniformly random value that is independent of the other coins in the experiment. Hence, the oracle $\mathcal{M}[F]_K$ behaves identically to $\mathcal{M}[\mathcal{S}, F]_K$ until there is some query input (i, X) and some $\mathcal{S}$ output (j, X) with $K_i = K_j$.

Let us bound the probability of this happening during an execution of $A^{\mathcal{M}[F]_K}$. (The Fundamental Lemma of Game Playing implies that this probability is equal in both games; we are free to choose whichever best expedites

[1] The proof permits us to omit this final restriction by changing the first term in the bound to $2/2^k$.

the proof.) Fix one of the $m-1$ pairs (j, X) output by $\mathcal{S}$. As $\mathcal{M}$ and $\mathcal{S}$ are compatible, query never receives an input (j, X). Except with probability ϵ, there are at most μ aliases i such that query receives an input (i, X). For each such alias i, $\Pr[K_i = K_j] = 1/2^k$; hence, some such alias exists with probability at most $\mu/2^k$. Taking a union bound over the $m-1$ pairs (j, X) gives us $\Pr\left[A^{\mathcal{M}[F]^K} \Rightarrow 1\right] - \Pr\left[A^{\mathcal{M}[\mathcal{S},F]^K} \Rightarrow 1\right] \leq \frac{m\mu}{2^k} + \epsilon$.

In Game G1, the E and E^{-1} oracles behave independently of the others. However, in Game G2, which includes the boxed statements, the F and E oracles have been coupled together (turning F into a blockcipher). So $\Pr[\mathrm{G2}(A) \Rightarrow 1] = \Pr\left[A^{\mathcal{M}[\mathcal{S},E]^K, E, E^{-1}} \Rightarrow 1\right]$.

We therefore wish to bound $\Pr[\mathrm{G1}(A) \Rightarrow 1] - \Pr[\mathrm{G2}(A) \Rightarrow 1]$. The Fundamental Lemma of Game Playing allows us to do so by bounding the probability that one of the boolean "bad flags" of Fig. 4 is set during an execution of $\mathrm{G1}(A)$.

Let $\mathcal{C}_c$ be the event that for some key K, $|\{i : K_i = K\}| > c$. By Lemma 1, $\Pr[\mathrm{G1}(A) ; \mathcal{C}_c] \leq \frac{m^{c+1}}{2^{nc}(c+1)!}$.

Now, in Game $\mathrm{G1}(A)$, bad_1 is set on a particular query (K, X) to E only if the initial value for Y is in $\mathrm{Rng}(F_K)$:

$$\begin{aligned}\Pr[Y \in F_K \mid \neg\mathcal{C}_c] &= \sum_{K_i} \Pr[K = K_i \mid \neg\mathcal{C}_c] \Pr[Y \in F_K \mid K = K', \neg\mathcal{C}_c] \\ &\leq \sum_{K'} \frac{1}{2^k} \frac{|\mathrm{Dom}(F_{K'})|}{2^n - q_E - q_P} \leq \frac{c\nu}{2^k(2^n - q_E - q_P)}.\end{aligned}$$

Hence $\Pr[\mathrm{G1}(A) ; \mathsf{bad}_1 \mid \neg\mathcal{C}_c] \leq \frac{q_E c\nu}{2^k(2^n - q_E - q_P)}$. A symmetric argument shows the same bound applies to $\Pr[\mathrm{G1}(A) ; \mathsf{bad}_3 \mid \neg\mathcal{C}_c]$.

Similarly, bad_2 is set on a particular query (K, X) to E only if $X \in \mathrm{Dom}(F_K)$. Except with probability ϵ, There are at most μ key aliases i such that $X \in \mathrm{Dom}(F_{K_i})$. Hence, $\Pr[\mathrm{G1}(A) ; \mathsf{bad}_2] \leq \frac{q_E\mu}{2^k} + \epsilon$.

Note that bad_4 is only set if the adversary makes a query (K, Y) to E^{-1} for some $Y \in \mathrm{Rng}(F_K)$. Over the course of the game, the probability that there will exist some $Y' \in \{0,1\}^n$ with $|\{(K, X) : F_K(X) = Y'\}| > c$ is at most $\frac{(m\nu)^c}{2^{n(c-1)}}$; i.e., except with this probability, $|\{K' : Y \in \mathrm{Rng}(F_{K'})\}| \leq c$. (This follows from the fact that points in the range of each F_K are uniform and mutually independent; see Lemma 1). Thus $\Pr[\mathrm{G1}(A) ; \mathsf{bad}_4] \leq \frac{q_E c}{2^k} + \frac{(m\nu)^c}{2^{n(c-1)}}$.

To bound $\Pr[\mathrm{G1}(A) ; \mathsf{bad}_5]$, consider a query (i, X) to F. We sample a uniformly random $Y \xleftarrow{\$} \{0,1\}^n$ and set bad_5 if $Y \in \mathrm{Rng}(E_{K_i})$ or $Y \in \mathrm{Rng}(F_{K_i})$. Using an argument similar to that for our bound for bad_1, $\Pr[Y \in \mathrm{Rng}(E_{K_i})] \leq \frac{q_E + q_P}{2^{k+n}}$. Again fix a positive integer c. So as long as no key corresponds to more than c aliases, $Y \in \mathrm{Rng}(F_{K_i})$ with probability at most $c\nu/2^n$. Taking a union bound over each of $m\nu$ queries gives $\Pr[\mathrm{G1}(A) ; \mathsf{bad}_5 \mid \neg\mathcal{C}_c] \leq \frac{(q_E + q_P)m\nu}{2^{k+n}} + \frac{cm\nu^2}{2^n}$.

Finally, we need to bound $\Pr[\mathrm{G1}(A) ; \mathsf{bad}_6]$. This flag is set only if some E or E^{-1} query defines the point $E_K(X) = Y$ such that $K = K_i$ and $X = X'$, where

GAMES $G1$, $\boxed{G2}$

Procedure main(A):
for $j = 1$ to $m - 1$ **do**
  register()
$b \leftarrow A^{f,E,E^{-1}}$
return b

Oracle $F(i, X)$:
if $X \in \mathrm{Dom}\left(F_{K_i}\right)$ **then**
  return $F_{K_i}(X)$
$Y \xleftarrow{\$} \{0,1\}^n$
if $Y \in \mathrm{Rng}\left(E_{K_i}\right) \cup \mathrm{Rng}\left(F_{K_i}\right)$ **then**
  $\mathsf{bad}_5 \leftarrow \mathsf{true}$
  $\boxed{Y \xleftarrow{\$} \overline{\mathrm{Rng}\left(E_{K_i}\right) \cup \mathrm{Rng}\left(F_{K_i}\right)}}$
if $X \in \mathrm{Dom}\left(E_{K_i}\right)$ **then**
  $\mathsf{bad}_6 \leftarrow \mathsf{true}$
  $\boxed{Y \leftarrow E_{K_i}(X)}$
$F_{K_i}(X) \leftarrow Y$
return Y

Oracle query(i, X):
if $i >$ c **then return** $\perp$
return $F(i, X)$

Oracle register():
$(i, X) \leftarrow \mathcal{S}_{\mathsf{SK}}^{\mathsf{query}}$
$K_{\mathsf{c}+1} \leftarrow F(i, X)$
$\mathsf{c} \leftarrow \mathsf{c} + 1$

Oracle $E(K, X)$:
$Y \xleftarrow{\$} \overline{\mathrm{Rng}\left(E_K\right)}$
if $Y \in \mathrm{Rng}\left(F_K\right)$ **then**
  $\mathsf{bad}_1 \leftarrow \mathsf{true}$
  $\boxed{Y \xleftarrow{\$} \overline{\mathrm{Rng}\left(E_K\right) \cup \mathrm{Rng}\left(F_K\right)}}$
if $X \in \mathrm{Dom}\left(F_K\right)$ **then**
  $\mathsf{bad}_2 \leftarrow \mathsf{true}$
  $\boxed{Y \leftarrow F_K(X)}$
$E_K(X) \leftarrow Y$
return $E_K(X)$

Oracle $E^{-1}(K, Y)$:
$X \xleftarrow{\$} \overline{\mathrm{Dom}\left(E_K\right)}$
if $X \in \mathrm{Dom}\left(F_K\right)$ **then**
  $\mathsf{bad}_3 \leftarrow \mathsf{true}$
  $\boxed{X \xleftarrow{\$} \overline{\mathrm{Dom}\left(E_K\right) \cup \mathrm{Dom}\left(F_K\right)}}$
if $Y \in \mathrm{Rng}\left(F_{K_i}\right)$ **then**
  $\mathsf{bad}_4 \leftarrow \mathsf{true}$
  $\boxed{X \leftarrow F_K^{-1}(Y)}$
$E_K^{-1}(Y) \leftarrow X$
return $E_K^{-1}(Y)$

Oracle $f(M)$:
return $\mathcal{M}_{\mathsf{KM}}^{\mathsf{query},\mathsf{register}}(M)$

Fig. 4. In Game G2, A, $\mathcal{M}$, and $\mathcal{S}$ access the same blockcipher (directly, through query_E, and through query_F, respectively). In Game G1, the behavior of query_F is decoupled from E and query_E, in effect giving the scheduler $\mathcal{S}$ it's own blockcipher.

(i, X') is some (future) F-query. Let us first consider a precomputation query that defines $E_K(X) = Y$. Then bad_6 will be triggered by this precomputation query only if K is one of the at most μ keys under which X is queried. Hence, the probability that some precomputation query will define a point on E that triggers bad_6 is at most $q_P\mu/2^k$.

Now let us consider an offline query that defines $E_K(X) = Y$. Except with probability ϵ, there are at most μ key aliases i that will be used to encipher X; the probability that one of these μ keys will be K is at most $\frac{\mu}{2^k}$. Hence, the probability that some offline query will define a point on E that triggers

bad_6 is at most $q_E\mu/2^k$. Therefore $\Pr[\,\mathrm{G1}(A)\,;\,\mathsf{bad}_6\,] \leq \mu(q_E+q_P)/2^k+\epsilon$. The Fundamental Lemma of Game-Playing gives us:

$$\begin{aligned}
&\Pr[\,\mathrm{G1}(A)\Rightarrow 1\,]-\Pr[\,\mathrm{G2}(A)\Rightarrow 1\,]\\
&\qquad\leq \Pr[\,\mathsf{bad}_1\vee\mathsf{bad}_3\vee\mathsf{bad}_5\mid\neg\mathcal{C}_c\,]+\Pr[\,\mathcal{C}_c\,]\\
&\qquad\quad+\Pr[\,\mathsf{bad}_2\vee\mathsf{bad}_4\vee\mathsf{bad}_6\,]\\
&\qquad\leq \frac{2q_Ec\nu}{2^k(2^n-q_E-q_P)}+\frac{(q_E+q_P)m\nu}{2^{k+n}}+\frac{cm\nu^2}{2^n}+\frac{m^{c+1}}{2^{nc}c+1!}\\
&\qquad\quad+\frac{2q_E\mu}{2^k}+\frac{q_Ec}{2^k}+\frac{(m\nu)^c}{2^{n(c-1)}c!}+\frac{q_P\mu\epsilon}{2^k}+3\epsilon\\
&\qquad= \frac{2q_Ec\nu}{2^k(2^n-q_E-q_P)}+\frac{(q_E+q_P)m\nu}{2^{k+n}}+\frac{cm\nu^2}{2^n}\\
&\qquad\quad+\frac{q_E(2\mu+c)+q_P\mu}{2^k}+\frac{m^{c+1}(1+\nu^{c+1})}{2^{nc}(c+1)!}+3\epsilon.
\end{aligned}$$

Collecting our results completes the proof. □

Relating the ICM-KOA to the standard model. We now move on to a standard-model analogue. The indistinguishability advantage definition is the same, except now A has an implicit description of E rather than oracle access:

Definition 8 (Standard model indistinguishability). *Let $\mathcal{E}:\mathcal{K}\times\mathcal{D}\to\mathcal{R}$ be a scheme over an (n,n)-bit blockcipher and let $I:\mathcal{D}\to\mathcal{R}$ be some oracle. Let E be an (n,n)-bit blockcipher. We define standard model indistinguishability advantage of an adversary A (with respect to $\mathcal{E}$ and $\mathcal{I}$) as:* $\mathbf{Adv}^{ind\text{-}\mathcal{I}}_{\mathcal{E};E}(A) = \Pr\left[\,A^{\mathcal{M}[\mathcal{S},E]_K}\Rightarrow 1\,\right]-\Pr\left[\,A^{\mathcal{I}}\Rightarrow 1\,\right]$, *where $K\xleftarrow{\$}\mathcal{K}$ is a random key and E is an (n,n)-bit blockcipher.*

We now give the second of our two main model-implication results. Namely, that security in the ICM-KOA implies security in the standard model.

Theorem 2 (ICM-KOA indistinguishability implies standard model indistinguishability). *Let $\mathcal{E}$ be an (k,n)-bit blockcipher-based scheme, and let $\hat{\mathcal{E}}=(\mathcal{M},\mathcal{S},\mathcal{K})$ be a decomposition of $\mathcal{E}$. Suppose*

1. *$\mathcal{M}$ is compatible with $\mathcal{S}$,*
2. *$\hat{\mathcal{E}}$ is (q,σ,m,ν)-efficient,*
3. *For any adversary B' making q queries, $\mathbf{Adv}^{koa\text{-}ind\text{-}\mathcal{I}}_{\hat{\mathcal{E}}}(B')\leq\delta$.*

Then for any adversary A running in time t and making q queries, each at most σn bits in length, there exists some adversary B running in time $t'\approx t$ and making ν queries such that $\mathbf{Adv}^{ind\text{-}\mathcal{I}}_{\mathcal{E};E}(A)\leq m\mathbf{Adv}^{\mathrm{prf}}_{E}(B)+\frac{m^2}{2^k}+\delta$.

This theorem relates ICM-KOA security to the PRF security of the underlying blockcipher. This implies a relationship between ICM-KOA security and PRP security via the PRP-PRF switching lemma, at the expense of an additional $m\sigma^2/2^{n+1}$ term. This term beats the birthday bound by a factor of m.

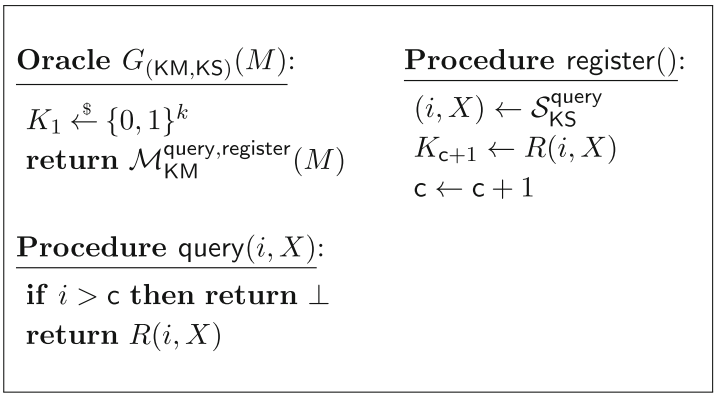

Fig. 5. Replacing E with a random function R

Proof (Theorem 2). We will use a game-playing proof. First A's oracle will transition from $\mathcal{M}[\mathcal{S}, E]_K$ into G, where references to $E_{K_i}(X)$ are replaced with $R(i, X)$ for some random function R (see Fig. 5).

This transition will itself involve a sequence of games. Define the oracle G_ℓ to be identical $\mathcal{M}[\mathcal{S}; E]_K$ for $K \xleftarrow{\$} \mathcal{K}$, except that query and register compute $R(i, X)$ in place of $E(K_i, X)$ when $i < \ell$. This gives us

$$\Pr\left[A^{\mathcal{M}[\mathcal{S},E]_K} \Rightarrow 1 \right] - \Pr\left[A^{G} \Rightarrow 1 \right] \leq \sum_{j=0}^{m-1} \left(\Pr\left[A^{G_{j+1}} \Rightarrow 1 \right] - \Pr\left[A^{G_j} \Rightarrow 1 \right] \right).$$

Now in G_{j+1}, we have $K_{j+1} = R(i, X)$ for some $i \leq j$, where the compatibility condition ensures that this is the only time R is evaluated at the point (i, X). Consequently, K_{j+1} is uniformly distributed and independent of the other coins of the experiment. It can therefore be freely discarded and replaced with some other value draw from this distribution without affecting the black-box behavior of G_{j+1}. Therefore from A we can construct a PRF adversary B_j with the property $\mathbf{Adv}_E^{\mathrm{prf}}(B_j) = \Pr\left[A^{G_{j+1}} \Rightarrow 1 \right] - \Pr\left[A^{G_j} \Rightarrow 1 \right]$. This is accomplished by having B_j^f simulate G_j for A, but using its own oracle to set $\mathsf{query}(j+1, \cdot) = f(\cdot)$. So B_j^f in behaves identically to either G_j (when f is E_K) or G_{j+1} (when f is a random function). We note that B_j makes at most ν queries and has roughly the same running time as A.

Setting B to be the B_j with maximal advantage ($1 \leq j \leq m$) gives us $\Pr\left[A^{\mathcal{M}[\mathcal{S},E]_K} \Rightarrow 1 \right] - \Pr\left[A^{G} \Rightarrow 1 \right] \leq m\mathbf{Adv}_E^{\mathrm{prf}}(B)$.

We observe that the G and $\mathcal{M}[F]$ differ in behavior only when $K_i = K_j$ for some $i \neq j$, which happens with probability at most $m^2/2^k$. Hence, $\Pr\left[A^{G} \Rightarrow 1 \right] - \Pr\left[A^{\mathcal{M}[F]} \Rightarrow 1 \right] < m^2/2^k$.

Finally, by hypothesis $\Pr\left[\, A^{\mathcal{M}[F]} \Rightarrow 1 \,\right] - \Pr\left[\, A^{\mathcal{I}} \Rightarrow 1 \,\right] \leq \delta$. Combining these results provides the desired bound. □

2.4 Connection to TBC-based Constructions

A tweakable blockcipher $\widetilde{E}$ is a (strong) TPRP if a keyed instance of $\widetilde{E}$ is computationally indistinguishable from an ideal cipher. This suggests that there ought to be some formal relationship between TBCs and the ideal cipher model, but the fact that TBCs are a keyed construction means the two objects cannot be directly compared. However, the key managers we have introduced *are* keyed constructions that mediate access between modes of operation and an underlying cipher. They thus offer a means of bridging the conceptual gap between TBCs and ideal ciphers: specifically, the following theorem states that any mode of operation secure in the ICM-KOA can be transformed into a TBC-based construction secure in the standard model. In the following theorem statement, ε denotes the empty string.

Theorem 3 (Decompositions imply TBC-based constructions). *Let $\mathcal{E}$ be a scheme over a (k, n)-bit blockcipher, and fix a decomposition $\hat{\mathcal{E}} = (\mathcal{M}, \mathcal{S}, \mathcal{K})$. Let be $\widetilde{E} : \{0,1\}^k \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$ be an n-bit TBC. Sample $K \xleftarrow{\$} \{0,1\}^k$ and $(\mathsf{KM}, \mathsf{KS}) \xleftarrow{\$} \mathcal{K}$.*

Define an oracle $\mathcal{F}\langle\widetilde{E}_K\rangle_{\mathsf{KM}}$ as follows: On input M, the output of $\mathcal{F}\langle\widetilde{E}_K\rangle_{\mathsf{KM}}$ is the value returned by the oracle $\mathcal{M}[S,E]_{(\mathsf{KM},\varepsilon)}(M)$ in Fig. 1 when (1) the register *procedure is replaced by a procedure* register-nop *that does nothing, and (2) the* query *procedure is modified so that, on input (i, X), it returns $\widetilde{E}_K(i, X)$.*[2] *(This assumes that the maximum number of key aliases permitted by the mode is at most $|\mathcal{T}|$.) For any adversary A running in time t and making q queries, each of length at most σn bits, there exists some adversary B making $m\nu$ queries and running in time $t' \approx t$ such that*

$$\Pr\left[\, A^{\mathcal{F}\langle\widetilde{E}_K\rangle_{\mathsf{KM}}} \Rightarrow 1 \,\right] - \Pr\left[\, A^{\mathcal{I}} \Rightarrow 1 \,\right] \leq \mathbf{Adv}^{\widetilde{\mathrm{prp}}}_{\widetilde{E}}(B) + \frac{m\nu^2}{2^n} + \frac{m^2}{2^k} + \delta$$

where $K \xleftarrow{\$} \mathcal{K}$.

Proof. Let $\Pi \xleftarrow{\$} \mathrm{BC}(k, n)$ be an ideal cipher and $F \xleftarrow{\$} \mathrm{Func}(k + n, n)$ be a random function. By a standard reduction argument, there exists some adversary B with the stated resources such that $\Pr\left[\, A^{\mathcal{F}\langle\widetilde{E}_K\rangle_{\mathsf{KM}}} \Rightarrow 1 \,\right] - \Pr\left[\, A^{\mathcal{F}\langle\Pi\rangle_{\mathsf{KM}}} \Rightarrow 1 \,\right] \leq \mathbf{Adv}^{\widetilde{\mathrm{prp}}}_{\widetilde{E}}(B)$. By the m applications of the Switching Lemma, $\Pr\left[\, A^{\mathcal{F}\langle\Pi\rangle_{\mathsf{KM}}} \Rightarrow 1 \,\right] - \Pr\left[\, A^{\mathcal{F}\langle F_K\rangle_{\mathsf{KM}}} \Rightarrow 1 \,\right] \leq m\nu^2/2^n$. Finally, note that $\mathcal{F}\langle F_K\rangle_{\mathsf{KM}}$ and $\mathcal{F}[F]_{(\mathsf{KM},\varepsilon)}$ behave identically unless the m random keys generated by the latter oracle's register procedure are not pairwise distinct, an event that happens with probability $m^2/2^k$. Collecting results completes the proof. □

[2] With these changes, the parameter E is unused.

3 ICM-KOA Analysis of Constructions

We now put the ICM-KOA to work, using it to analyze example blockcipher-based constructions. We begin with the NIST-CTR-DRBG, as used in Intel's recent hardware random-number generator [11], whose standard-model security bounds [19] can become quite weak when an adversary is co-located on the same physical machine, due to the rate at which such an adversary can make queries. The weakness of these bounds is do to a hybrid-factor loss. Our ICM-KOA analysis yields considerably better bounds, and suggests that the multiplicative loss in the standard-model isn't "real".

Next, we give an example of when the standard-model fails to surface quantitative differences between the security of closely related schemes. In particular, we consider various rekeying and nonce-choice strategies for CTR mode. Although these schemes yield similar bounds in the standard model, we show that the best-possible black-box attacks tell quite a different story. These results are of particular importance when CTR is built over a lightweight blockciphers, where the standard-model security bounds for all of the strategies suggest that problems may arise quickly. Our ICM-KOA analysis (and the implied ICM results) offers a different viewpoint on these concerns, and identifies the best strategies from among the choices.

3.1 Analysis of NIST CTR-DRBG Generation Algorithm

As the name suggests, CTR-DRBG is a deterministic random-bit generator based on running a blockcipher in CTR mode. Here, we analyze its generation algorithm[3], specializing for the sake of simplicity to the case where AES-128 is used (so $n = k = 128$), and where 128 bits are requested on each invocation. This case is of special interest because these parameters are used inside of Intel's hardware random number generator.

Concretely, we consider the scheme $\mathsf{ISK\text{-}RNG} : \{0,1\}^{2n} \times \{0,1\}^0 \to \{0,1\}^n$ over an (n,n)-bit blockcipher defined in Fig. 6. The system maintains an initially random internal state (K, IV), and on each query computes $(R, K, \mathsf{IV}) \leftarrow (E_K(\mathsf{IV}), E_K(\mathsf{IV}+1), E_K(\mathsf{IV}+2))$, updating the state, and returns R. In order to decompose this into a model, we need the mode and scheduler to share the IV portion of the state. This is accomplished by using the initial IV as part of both the mode and scheduler key (these keys are not required to be independent).

We define $\mathsf{Rand} : \{0,1\}^0 \to \{0,1\}^n$ to be the oracle that on each query samples $R \xleftarrow{\$} \{0,1\}^n$ and then returns R.

Stronger than standard-model results desirable. A result by Shrimpton and Terashima [19] shows, as one might expect, that the standard-model security bound for q queries includes an $\mathcal{O}(q\mathbf{Adv}^{\mathrm{prp}}_E(B))$ term, where B is an adversary making three queries. However, B also has time t to run, where t is sufficient time to evaluate E on $3q$ inputs. Hence even if B conducts a naïve brute-force

[3] The specification also includes algorithms for, e.g., reseeding.

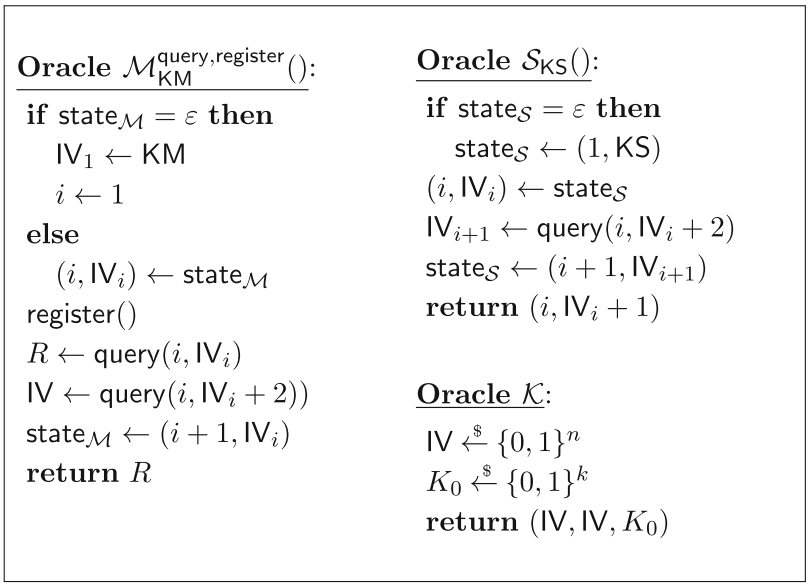

Fig. 6. The NIST CTR-DRBG decomposes into the mode and scheduler described above. The key-generation algorithm $\mathcal{K}$ ensures $\mathsf{KM} = \mathsf{KS}$.

attack, $\mathbf{Adv}_E^{\mathrm{prp}}(B) \approx 3q/2^k$. So the security bound becomes roughly $q^2/2^k$. For $k = 128$, this bound exceeds 2^{-40} when $q = 2^{44}$.

This may seem safe; after all, this amounts to many terabytes of random bits. But the RNG has extremely high throughput—Intel reports 800 MB/s, which equates to 50 million queries per second. This means an attacker who shares a physical machine with his target can reach the $q = 2^{44}$ limit in a little more than four days.

The following lemma provides a security bound for the ISK-RNG in the ICM-KOA. For very little work, we recover the security bound of Shrimpton and Terashima [19], and immediately also get a much stronger bound in the ICM. The ICM bound reveals the lack of a matching attack, and shows that barring cryptanalysis of AES *under random and secret keys*, we can permit on the order of 2^{70} queries before surpassing our 2^{-40} limit (assuming the adversary has resources for 2^{80} precomputation and 2^{80} offline queries). This translates to 750,000 years of ISK-RNG runtime, and so is unlikely to be the limiting factor.

Lemma 2. *For any positive integers μ and any adversary A making at most q online queries, ISK-RNG is $(q, 0, q, 3)$-efficient, has $(q, 0, c, \epsilon)$-dispersed inputs, and* $\mathbf{Adv}^{\textit{koa-ind}\text{-}\mathsf{Rand}}_{\textit{ISK-RNG}}(A) \leq \delta$, *where* $\delta = \frac{5q^2}{2^{2n}}$ *and* $\epsilon = \delta + \frac{(3q)^3}{2^{3n}3!}$.

Proof. If A makes q queries (0 bits each), the RNG will make three queries using each of q distinct key aliases. Hence $\hat{\mathcal{E}}$ is $(q, 0, q, 3)$-efficient.

Let $R : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$ be an oracle that samples and returns a fresh random string on each query (so R may return different outputs on

the same input). Consider Experiment $\mathsf{COMPAT}^R_{\hat{\mathcal{E}}}(A)$. Let $(K_i, \mathsf{IV}_i)_{i=1}^q$ be the sequence of keys and IVs generated during this experiment. Then the probability that there exists some string $x \in \{0,1\}^n$ that is enciphered under more than c key aliases is less than $\frac{(3q)^{c+1}}{2^{nc}(c+1)!}$.

Let $F \xleftarrow{\$} \mathrm{Func}(k+n, n)$. Then Experiment $\mathsf{COMPAT}^F_{\hat{\mathcal{E}}}(A)$ proceeds identically to $\mathsf{COMPAT}^R_{\hat{\mathcal{E}}}(A)$ unless an F-query is repeated; i.e., unless there exists $i < j$ such that $K_j = K_i$ and $\mathsf{IV}_j \in \{\mathsf{IV}_i + \ell \,:\, -2 \le \ell \le 2\}$. The probability that this happens (which is identical in both games, but easier to compute with respect to the R oracle), is less than $\frac{q^2}{2^k}\left(\frac{5}{2^n}\right)$. Therefore $\hat{\mathcal{E}}$ has $(q, 0, c, \epsilon)$-dispersed inputs for $\epsilon = \frac{5q^2}{2^{k+n}} + \frac{(3q)^{c+1}}{2^{nc}(c+1)!}$.

Finally, we need to bound $\Pr\left[A^{\mathcal{M}[F]_K} \Rightarrow 1\right] - \Pr\left[A^{\mathsf{Rand}} \Rightarrow 1\right]$. As before $\Pr\left[A^{\mathcal{M}[F]_K} \Rightarrow 1\right] - \Pr\left[A^{\mathcal{M}[R]_K} \Rightarrow 1\right] \le \frac{5q^2}{2^{k+n}}$, and $\Pr\left[A^{\mathcal{M}[R]_K} \Rightarrow 1\right] - \Pr\left[A^{\mathsf{Rand}} \Rightarrow 1\right] = 0$. □

Combining this result with Theorem 2 and immediately gives the following results:

Corollary 1. *Let A be an adversary making q queries and running in time t. Then there exists an adversary B making 3 queries and running in time $t' \approx t$ such that* $\mathbf{Adv}^{\textit{ind}\text{-}\mathsf{Rand}}_{\textit{ISK-RNG}[E]}(A) \le q\mathbf{Adv}^{\mathrm{prf}}_E(B) + \frac{q^2}{2^n} + \frac{5q^2}{2^{2n}}$.

Note that up to a small constant factor, we've recovered, essentially the security bound from [19]. But we can do better:

Corollary 2. *Let A be an adversary making q_P precomputation queries, q_E offline queries, and q online queries, where $q_E + q_P < 2^{n-1}$. Then*

$$\mathbf{Adv}^{\textit{icm-ind}\text{-}\mathsf{Rand}}_{\textit{ISK-RNG}}(A) \le \frac{20q^2 + 24q_E + 3q(q_E + q_P) + 19q^3}{2^{2n}} + \frac{20q + 6q_E + 2q_P}{2^n}$$

Here we have set $c = 2$ for the sake of notational cleanliness.

Taking $q_E = q_P = 2^{80}$ allows the upper bound to stay below 2^{-40} even when $q = 2^{70}$, a substantial improvement over the previous $q = 2^{44}$ (which only applied to attackers with $q_P = 2^{44}$). This is a significantly stronger result than we could obtain in the standard model, and it retains the standard model's strength of only relying on random, secret keys. A brute-force attack on the key would obtain about the same success rate.

3.2 Analysis of CTR-mode Variants

We consider three variants on CTR mode, distinguished by how they choose IVs following a key change: (1) The IV is set to 0^n; (2) the upper bits of the IV are unique for each key; (3) The IV is chosen randomly. In each case, we use the same key scheduler that sets $K_i \leftarrow E_{K_1}(i)$ (for $i > 1$). See Fig. 7. For simplicity, we consider the case where the key changes with each message. This models a situation where the counter state is retained between messages with the same

key. The loss of adaptivity within the lifetime of a given key does not hamper a chosen-plaintext adversary in this context because the nature of CTR mode permits him to compute what a ciphertext would have been with a different plaintext. The variants are distinguished by the choice of $\mathsf{iv\text{-}gen} : \mathbb{N} \to \{0,1\}^n$, which on input i outputs some IV_i. Define the reference scheme $\mathcal{R}[\mathsf{iv\text{-}gen}]$ to be the stateful function that on its $i\mathsf{th}$ query M, computes $\mathsf{IV} \leftarrow \mathsf{iv\text{-}gen}(i)$, samples $C \xleftarrow{\$} \{0,1\}^{|M|}$, and returns (IV, C).

Theorem 4. *Fix positive integers σ, q, and b with $q < \sigma < 2^b$ and $b < n$. Let $\mathsf{const}(i) = 0^n$, let $\mathsf{unique}(i) = \langle i \rangle_b 0^{n-b}$ (where $\langle i \rangle_b$ is a b-bit encoding of i), and let $\mathsf{rand}(i)$ sample and return $R \xleftarrow{\$} \{0,1\}^n$ on each invocation. Let A be an adversary making q online queries, each at most σn bits long, q_P precomputation queries, and q_E offline queries. Then:*

$$(1)\quad \mathbf{Adv}^{ind\text{-}\mathcal{R}[\mathsf{const}]}_{\mathsf{CTR}[\mathsf{const}]}(A) \le \frac{4q_E\sigma}{2^k(2^n - q_E - q_P)} + \frac{(q_E+q_P)q\sigma}{2^{k+n}} + \frac{2q\sigma^2}{2^n} + \frac{2q_E(q+1) + q_P q + 2q^2}{2^k} + \frac{q^3(1+\sigma^3)/6}{2^{2n}}$$

$$(2)\quad \mathbf{Adv}^{ind\text{-}\mathcal{R}[\mathsf{unique}]}_{\mathsf{CTR}[\mathsf{unique}]}(A) \le \frac{4q_E\sigma}{2^k(2^n - q_E - q_P)} + \frac{(q_E+q_P)q\sigma}{2^{k+n}} + \frac{2q\sigma^2}{2^n} + \frac{6q_E + 2q_P + 2q}{2^k} + \frac{q^3(1+\sigma^3)/6}{2^{2n}}$$

$$(3)\quad \mathbf{Adv}^{ind\text{-}\mathcal{R}[\mathsf{rand}]}_{\mathsf{CTR}[\mathsf{rand}]}(A) \le \frac{4q_E\sigma}{2^k(2^n - q_E - q_P)} + \frac{(q_E+q_P)q\sigma + (q\sigma)^2}{2^{k+n}} + \frac{2q\sigma^2}{2^n} + \frac{6q_E + 2q_P}{2^k} + \frac{q^3(1+4\sigma^3)/6}{2^{2n}}$$

Proof. Each decomposition is $(q, \sigma, q+1, \sigma)$-efficient. Sample $F \xleftarrow{\$} \mathrm{Func}(k+n, n)$. Let $\mathsf{iv\text{-}gen} \in \{\mathsf{const}, \mathsf{unique}, \mathsf{rand}\}$. Let bad be the event that during an execution $A^{\mathsf{CTR}[\mathsf{iv\text{-}gen}][F]}$, $\mathsf{CTR}[\mathsf{iv\text{-}gen}][F]$ repeats a query to F. Barring this event, the outputs of $\mathsf{CTR}[\mathsf{iv\text{-}gen}][F]$ are independent and uniformly random (with the possible exception of the IV component). Therefore $\Pr\left[\, A^{\mathsf{CTR}[\mathsf{iv\text{-}gen}][F]} \Rightarrow 1 \,\right] - \Pr\left[\, A^{\mathcal{R}[\mathsf{iv\text{-}gen}]} \Rightarrow 1 \,\right] \le \Pr[\,\mathsf{bad}\,]$. We want to find an upper bound δ for $\Pr[\,\mathsf{bad}\,]$, and do so for each method of generating the IV. Specifically,

- When $\mathsf{iv\text{-}gen} = \mathsf{const}$, $\Pr[\,\mathsf{bad}\,] \le \Pr[\,\exists i \ne j \,:\, K_i = K_j\,] \le q^2/2^k$
- When $\mathsf{iv\text{-}gen} = \mathsf{unique}$, $\Pr[\,\mathsf{bad}\,] = 0$ because regardless of what value the keys have, the inputs never repeat.
- When $\mathsf{iv\text{-}gen} = \mathsf{rand}$, any two queries to F collide with probability $1/2^{k+n}$ because both keys and IVs are uniform and independent. There are fewer than $(q\sigma)^2$ pairs of queries, so $\Pr[\,\mathsf{bad}\,] < (q\sigma)^2/2^{k+n}$.

To apply Theorem 1 (with $c = 2$), we need to measure how much each variant disperses its inputs.

- CTR[const] has $(q, \sigma, q+1, 0)$-dispersed inputs because 0^n is evaluated under each of the $q+1$ keys.
- CTR[unique] has $(q, \sigma, 2, 0)$-dispersed inputs because each input is guaranteed to be used at most twice (including once by the scheduler).
- CTR[rand] has $(q, \sigma, c, (q\sigma)^{c+1}/2^{nc}(c+1)!)$. The argument here follows that of Lemma 1, except each that we are interested in the probability that $x \in \{X_i, X_i + 1, \ldots, X_i + (\sigma - 1)\}$, instead of $x = X_i$, where X_i plays the role of IV_i.

Plugging these values into Theorem 1 gives us the previously stated bounds. □

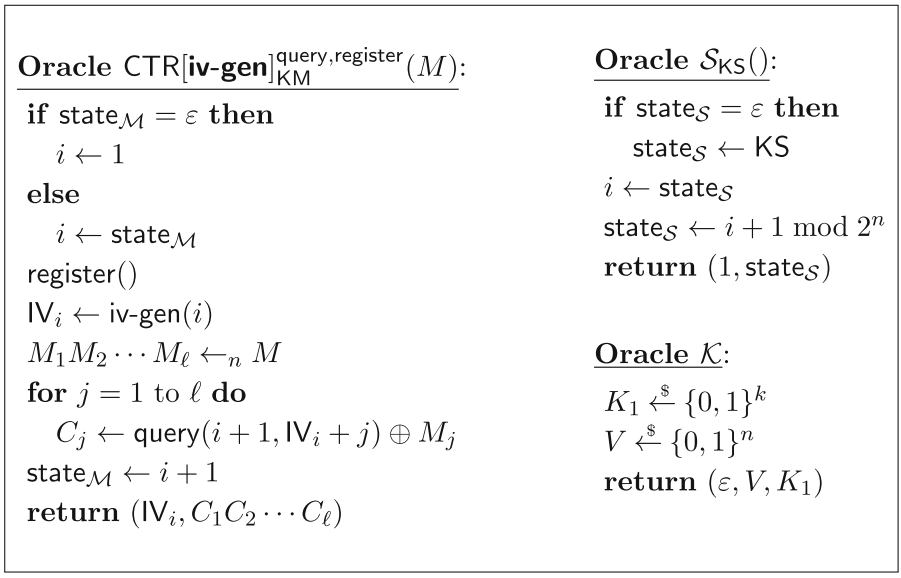

Fig. 7. A general decomposition of CTR parameterized by the IV selection function, iv-gen.

Interpretation. Assume $q_P \gg q_E, q$. Using the const IV generation function permits $\sigma = 2^{n/3}, q = 2^{n/3}$ (up to constants) as long as $2^{k-n/3} \gg q_P$. This allows on the order of $2^{2n/3}$ n-bit blocks of data to be securely encrypted, beating the birthday bound. However, the constraint on q_P may be worrisome for, e.g., $n = 64$, $k = 80$, which is only secure against adversaries for which $q_P \ll 2^{59}$. Using a predictable IV amplifies the effectiveness of precomputation because the adversary knows what precomputations will likely be helpful (in this case, finding preimages of $E_K(0^n)$). On the other hand, unique and rand also permit $\sigma = q = 2^{n/3}$, but the $\mathcal{O}(q_P q/2^k)$ term is now $\mathcal{O}(q_P/2^k)$. Precomputation is no longer nearly as much of a threat.

This $\mathcal{O}(q_P q/2^k)$ term for const corresponds to the following attack: Precompute $Y = E_K(0^n)$ for q_P arbitrary keys K, and store each K in a hash table using Y as the hash table key. Encrypt the string 0^{2n} q times, and perform a hash table lookup of the first n bits of the ciphertext. This recovers the key if it happened to be one of the q_P values used during precomputation. False positives can be all but eliminated by verifying the second n bits of the ciphertext.

References

1. Abdalla, M., Bellare, M.: Increasing the lifetime of a key: a comparative analysis of the security of re-keying techniques. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 546–559. Springer, Heidelberg (2000). doi:10.1007/3-540-44448-3_42
2. Andreeva, E., Bogdanov, A., Dodis, Y., Mennink, B., Steinberger, J.P.: On the indifferentiability of key-alternating ciphers. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 531–550. Springer, Heidelberg (2013). doi:10.1007/978-3-642-40041-4_29
3. Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: security proofs and improvements. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 259–274. Springer, Heidelberg (2000). doi:10.1007/3-540-45539-6_18
4. Bellare, M., Krovetz, T., Rogaway, P.: Luby-Rackoff backwards: increasing security by making block ciphers non-invertible. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 266–280. Springer, Heidelberg (1998). doi:10.1007/BFb0054132
5. Bellare, M., Ristenpart, T., Tessaro, S.: Multi-instance security and its application to password-based cryptography. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 312–329. Springer, Heidelberg (2012). doi:10.1007/978-3-642-32009-5_19
6. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006). doi:10.1007/11761679_25
7. Bogdanov, A., Knudsen, L.R., Leander, G., Standaert, F.-X., Steinberger, J.P., Tischhauser, E.: Key-alternating ciphers in a provable setting: encryption using a small number of public permutations (extended abstract). In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 45–62. Springer, Heidelberg (2012). doi:10.1007/978-3-642-29011-4_5
8. Chen, S., Steinberger, J.: Tight security bounds for key-alternating ciphers. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 327–350. Springer, Heidelberg (2014). doi:10.1007/978-3-642-55220-5_19
9. Recommendation for random number generation using deterministic random bit generators. National Institute of Standards and Technology, NIST Special Publication 800–90A, U.S. Department of Commerce, January 2012
10. Dai, Y., Lee, J., Mennink, B., Steinberger, J.P.: The security of multiple encryption in the ideal cipher model. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8616, pp. 20–38. Springer, Heidelberg (2014). doi:10.1007/978-3-662-44371-2_2
11. Hamburg, M., Kocher, P., Marson, M.E.: Analysis of Intel's Ivy Bridge digital random number generator (2012). http://www.cryptography.com/public/pdf/Intel_TRNG_Report_20120312.pdf

12. Hoang, V.T., Tessaro, S.: Key-alternating ciphers and key-length extension: exact bounds and multi-user security. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9814, pp. 3–32. Springer, Heidelberg (2016). doi:10.1007/978-3-662-53018-4_1
13. Jaulmes, É., Joux, A., Valette, F.: On the security of randomized CBC-MAC beyond the birthday paradox limit a new construction. In: Daemen, J., Rijmen, V. (eds.) FSE 2002. LNCS, vol. 2365, pp. 237–251. Springer, Heidelberg (2002). doi:10.1007/3-540-45661-9_19
14. Kilian, J., Rogaway, P.: How to protect DES against exhaustive key search (an analysis of DESX). J. Cryptology **14**(1), 17–35 (2001)
15. Lampe, R., Patarin, J., Seurin, Y.: An asymptotically tight security analysis of the iterated even-mansour cipher. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 278–295. Springer, Heidelberg (2012). doi:10.1007/978-3-642-34961-4_18
16. Lampe, R., Seurin, Y.: Security analysis of key-alternating feistel ciphers. Cryptology ePrint Archive, Report 2014/151 (2014). http://eprint.iacr.org/2014/151
17. Lee, J.: Towards key-length extension with optimal security: cascade encryption and xor-cascade encryption. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 405–425. Springer, Heidelberg (2013). doi:10.1007/978-3-642-38348-9_25
18. Mouha, N., Luykx, A.: Multi-key security: the even-mansour construction revisited. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9215, pp. 209–223. Springer, Heidelberg (2015). doi:10.1007/978-3-662-47989-6_10
19. Shrimpton, T., Terashima, R.S.: A provable-security analysis of Intel's secure key RNG. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 77–100. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46800-5_4