

A preliminary version of this paper appears in the *ACM Computer and Communications Security (CCS) Conference, 2017*. This is the full version.

Defending Against Key Exfiltration: Efficiency Improvements for Big-Key Cryptography via Large-Alphabet Subkey Prediction

MIHIR BELLARE¹

WEI DAI²

September 2017

Abstract

Towards advancing the use of BIG keys as a practical defense against key exfiltration, this paper provides efficiency improvements for cryptographic schemes in the bounded retrieval model (BRM). We identify probe complexity (the number of scheme accesses to the slow storage medium storing the big key) as the dominant cost. Our main technical contribution is what we call the large-alphabet subkey prediction lemma. It gives good bounds on the predictability under leakage of a random sequence of blocks of the big key, as a function of the block size. We use it to significantly reduce the probe complexity required to attain a given level of security. Together with other techniques, this yields security-preserving performance improvements for BRM symmetric encryption schemes and BRM public-key identification schemes.

¹ Department of Computer Science & Engineering, University of California, San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. Email: mihir@eng.ucsd.edu. URL: <http://cseweb.ucsd.edu/~mihir/>. Supported in part by NSF grants CNS-1526801 and CNS-1717640, ERC Project ERCC FP7/615074 and a gift from Microsoft corporation.

² Department of Computer Science & Engineering, University of California, San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. Email: weidai@eng.ucsd.edu. URL: <http://cseweb.ucsd.edu/~weidai/>. Supported in part by a Powell Fellowship and grants of first author.

Contents

1	Introduction	3
2	Preliminaries	6
3	Large-Alphabet Subkey Prediction	8
3.1	The Problem	8
3.2	Subkey Prediction Theorem	9
3.3	Proof of Theorem 1	12
3.3.1	Proof of Lemma 7	16
3.3.2	Proof of Lemma 9 and Lemma 10	18
3.4	Multi-challenge Subkey Prediction	20
4	Big-Key Symmetric Encryption	21
4.1	Proof of Theorem 15	22
5	Big-Key Identification	25
5.1	Proof of Theorem 16	32
6	Acknowledgements	35
	References	35
A	Proofs of Lemmas 2, 3, and 4	37
B	A Rejection Sampling Lemma	39

1 Introduction

This paper is concerned with the threat of key exfiltration. This means attacker-planted malware on the key-storing system uses the system’s network connection to convey the key to a remote accomplice. A line of theoretical work has suggested a mitigation, called the Bounded Retrieval Model (BRM) [13, 12, 9, 2, 1], which involves using big keys. BKR [4] initiated an effort to take the BRM (they call it big-key cryptography) to practicality. We continue this effort. We identify probe complexity (the number of scheme accesses to the slow storage medium storing the big key) as the dominant cost. Our *large-alphabet subkey prediction lemma* allows us to minimize the probe complexity required to reach a given level of security, thereby optimizing storage usage. We use this to obtain efficiency improvements for big-key symmetric encryption [4]. We then provide an additional lemma on polynomial-evaluation entropy preservation, and use the two lemmas in conjunction to obtain efficiency improvements for the ADW big-key identification scheme [2].

LARGE-ALPHABET SUBKEY PREDICTION. Let $b \geq 2$ be an integer representing the *block size* in a storage system, for example $b = 32$ or $b = 64$ for words in memory, or $b = 8 \cdot 512$ (512 bytes) for a typical hard-disk drive. Let $q = 2^b$ be the *alphabet size*, and $[q] = \{0, 1, \dots, q-1\}$ the corresponding alphabet. Let $\mathbf{K} = (\mathbf{K}[0], \dots, \mathbf{K}[k-1]) \in [q]^k$ be a string over $[q]$ of length k , randomly chosen. It represents a (big) key stored in our storage system as a sequence of k blocks. We imagine that an adversary-chosen function $\text{Lk}: [q]^k \rightarrow [q]^\ell$ (representing adversary-implemented malware, and here called the leakage function) is applied to \mathbf{K} , and the result L (representing exfiltrated information, here called the leakage), is provided back to the adversary. Think of ℓ as somewhat smaller than k , for example $\ell \leq k/10$, so that the leakage, although not total, is certainly non-trivial. Despite this, we wish to make secure use of the big key, specifically to (repeatedly) extract “small” keys ($\tau \geq 1$ blocks, for a parameter τ) for use with conventional cryptography. In any such extraction, we make τ random but distinct probes $i_1, \dots, i_\tau \in [k] = \{0, 1, \dots, k-1\}$ into \mathbf{K} to determine $J = \mathbf{K}[i_1] \dots \mathbf{K}[i_\tau]$ as the τ -block short key. Given the leakage L and the probe positions i_1, \dots, i_τ , the adversary aims to predict (compute in its entirety) J . Two metrics (see Section 3 for precise definitions of what we discuss next) are of interest. First is the *subkey prediction advantage*

$$\mathbf{Adv}_{q,k,\tau}^{\text{skp}}(\ell), \tag{1}$$

defined as the maximum probability that an adversary can compute J , the maximum being over all leakage functions Lk returning ℓ blocks and over all adversary strategies. It is useful to let $k^* = kb$ denote the amount of storage occupied by the big key in bits, and, correspondingly, $\ell^* = \ell b$ the amount of allowed leakage in bits. (We will want to fix these and vary b , thereby defining k and ℓ .) Now, in usage, we would ask for a certain number s of bits of security (for example $s = 128$), meaning we want the subkey prediction advantage to be at most 2^{-s} , and then want to know the number τ of probes it takes to get there. This is the *probe complexity*,

$$\mathbf{Probes}_{k^*,\ell^*,s}(b) = \min \left\{ \tau : \mathbf{Adv}_{2^b,k^*/b,\tau}^{\text{skp}}(\ell^*/b) \leq 2^{-s} \right\}. \tag{2}$$

The probe complexity will be our cost in accesses to a potentially slow storage system, like a disk, and for efficiency of the overlying big-key scheme, we want to minimize it. To this end, Theorem 1 gives a good upper bound on the subkey prediction advantage, whence we obtain a good upper bound on the probe complexity. Next, we compare our bounds to prior ones, and discuss history and applications (to big-key cryptography and thus key exfiltration resistance).

PRIOR WORK AND COMPARISONS. ADW [3, Lemma A.3] is an elegant and general result that, as a special case, gives an upper bound on the subkey prediction advantage (and thus probe complexity) for all values of parameters we consider. The bounds, however are quite poor, so that, to get a

b	$s = 128$		b	$s = 512$	
	Us	ADW		Us	ADW
1	271	11532	1	971	46127
8	61	1584	8	219	6335
32	47	592	32	171	2366
64	45	434	64	165	1735
$8 \cdot 512$	43	287	$8 \cdot 512$	159	1146
$8 \cdot 4096$	43	285	$8 \cdot 4096$	158	1139

Figure 1: Fix the amount of storage we allocate to the big key at $k^* = 8 \cdot 10^{11}$ bits = 100 GBytes. Fix the amount of leakage at 10% of the length of the big key, $\ell^* = k^*/10 = 10$ GBytes. The first table considers security level $s = 128$, while the second considers $s = 512$. Each table then considers different block sizes b . (Once b is chosen, the length of the big key in blocks is $k = k^*/b$ and the length of the leakage in blocks is $\ell = \ell^*/b$.) The table entries show upper bounds on the probe complexity $\mathbf{Probes}_{k^*, \ell^*, s}(b)$. The “Us” column is our bound via Theorem 1, and “ADW” is what is obtained via [3, Lemma A.3]. The block sizes are chosen to represent common word or disk sector sizes in storage systems.

desired level of security, one needs a very large number of probes (we will see numbers in a bit), resulting in a significant loss of efficiency for the overlying big-key cryptography schemes. This lead BKR [4] (in their quest for practical big-key symmetric encryption) to formulate subkey prediction, and seek better bounds by direct analysis. They however *only considered the case $b = 1$* of a binary alphabet. They gave an example to show that there are non-obvious leakage functions that lead to better subkey prediction advantage than one might expect, making the problem of giving a (good) upper bound challenging. Via a combinatorial analysis, they showed that the worst case is when the pre-images of the outputs of the leakage function are approximate Hamming balls in the space of big keys, thereby deriving estimates (not quite upper bounds, something we rectify) on the subkey prediction advantage and probe complexity, for the case $b = 1$ ($q = 2$), that are much better than those obtained via ADW [2, Lemma A.3]. They posed the large alphabet ($b > 1$) case as an open question, asking, specifically, to give bounds on subkey prediction advantage and probe complexity, in the $b > 1$ case, that are better than the ones obtained via ADW [3, Lemma A.3]. (The motivation, as we will see later, was to improve efficiency of big-key symmetric encryption.) Our work answers this question, giving (good) upper bounds as a function of the block size b .

In usage, we would typically first decide on the amount of storage k^* (measured in bits) we allocate to the big key, for example $k^* = 8 \cdot 10^{11}$ bits = 100 GBytes. Next we would fix the amount of leakage ℓ^* (also measured in bits), for example $\ell^* = k^*/10 = 10$ GBytes, corresponding to 10% of the length of the big key. The block size b may be determined by the storage system (for example 512 bytes or 4096 bytes) or chosen to optimize security and efficiency as per our bounds. Once it is chosen, the length in blocks $k = k^*/b$ of the big key and $\ell = \ell^*/b$ of the leakage are determined. Now, for a given level s of security, we want to know the probe complexity $\mathbf{Probes}_{k^*, \ell^*, s}(b)$. Smaller (fewer probes into the likely slow storage system) is better. We tabulate results in Fig. 1. Our bounds emerge as substantially better than those obtained via ADW [3, Lemma A.3]. For example, for $s = 128$, the improvement ranges from a factor of 26 ($b = 8$) to a factor of 6.6 ($b = 8 \cdot 4096$). Below, we will see how this translates to efficiency improvements for big-key cryptography.

THE BRM. Assume (for concreteness of this discussion) that the primitive is symmetric encryp-

tion [4] (we will discuss other primitives later), and let \mathbf{K} denote the encryption key, k^* bits long. In the Bounded Retrieval Model (BRM) [13, 12, 9, 2, 1, 4], an adversary-chosen function Lk (representing adversary-implanted malware) is applied to \mathbf{K} , and the ℓ^* -bit result L (representing the exfiltrated information), is provided back to the adversary. Security would appear impossible, since Lk could be the identity function, so that $L = \mathbf{K}$, but the idea is that \mathbf{K} is big (for example $k^* = 100$ GBytes), while L is assumed to be somewhat smaller (like $\ell^* = k^*/10 = 10$ GBytes). In other words, the model assumes that the amount of data exfiltrated can be limited, say via network or system monitoring. Indeed, security product vendors such as McAfee [15] provide tools for this type of monitoring and detection.

If the scheme is poorly designed, the fact that the exfiltrated information is somewhat shorter than the key won't guarantee security. For example if the scheme applies SHA256 to \mathbf{K} to get a 256 bit key K and then uses AES256 to encrypt the data, then $\text{Lk}(\mathbf{K})$ can just return the 256 bit string $K = \text{SHA256}(\mathbf{K})$ and security is entirely compromised no matter how big is \mathbf{K} . The first requirement for a BRM (also called big-key) scheme is thus *leakage resilience*, meaning an adversary, given $L = \text{Lk}(\mathbf{K})$, still cannot violate security, and this must be true for *any* (adversary-chosen) function Lk that returns ℓ^* bits.

PROBE COMPLEXITY. Big keys may help for security, but it would be prohibitively costly to process a 100 GByte key for every encryption. The BRM addresses this via a condition that says that each encryption (or decryption) operation should only make a “small” number of probes into the big key \mathbf{K} , meaning have low probe complexity. Security in the presence of leakage is a difficult goal under any circumstances, but made even more so here by this requirement.

FROM BITS TO BLOCKS. Viewing the big key $\mathbf{K} = (\mathbf{K}[0], \dots, \mathbf{K}[k^*-1])$ as a sequence of bits, BKR encryption [4] begins by making some τ^* random probes $i_1, \dots, i_{\tau^*} \in [k^*]$ into \mathbf{K} to extract a τ^* -bit subkey $J = \mathbf{K}[i_1] \dots \mathbf{K}[i_{\tau^*}]$. It then applies a (randomized) hash function to J to get a key K for a conventional (AES-based) symmetric encryption scheme, and uses K to encrypt the data. Once J has been obtained, the computation, being symmetric cryptography operations, is quite efficient, but \mathbf{K} , being big, is likely stored on a slow medium like a hard drive, and so the encryption cost is dominated by the storage accesses needed to get J . For a subkey prediction advantage of $s = 128$ (based on which BKR show ind-cpa style security of their encryption scheme at the same security level), BKR will need $\tau^* = 271$ probes into the storage. (This is as per the $b = 1$ row of the first table in Fig. 1. BKR's subkey prediction lemma gives an estimate, not a bound, so we use our number, but numerically the two are almost the same.)

But (as BKR themselves point out), their scheme is making very poor use of storage by drawing only a bit of the big key per probe. Letting b be some appropriate block size determined by the storage system (for example $b = 8 \cdot 512$ bits = 512 Bytes), \mathbf{K} would actually be stored as a sequence of blocks, and a single probe into the storage can retrieve an entire block at about the same cost as retrieving a single bit. Indeed, a typical storage API does not even provide a way to directly access a bit, so an implementation of BKR would, for a probe for bit-position j , have to draw the block containing this bit position, take the corresponding bit, and throw the other bits away. A natural improvement (suggested by BKR) is to draw (and use) an entire block per probe. Thus, we now view the big key $\mathbf{K} = (\mathbf{K}[0], \dots, \mathbf{K}[k-1]) \in [2^b]^k$ as a sequence of blocks, corresponding to the way it is actually stored, where $k = k^*/b$ is the number of blocks. Now, making some τ probes $i_1, \dots, i_\tau \in [k]$ into \mathbf{K} , one obtains the subkey $J = \mathbf{K}[i_1] \dots \mathbf{K}[i_\tau]$. The rest of the encryption process is as before, and as we have already noted, is efficient, even though J could be a bit longer. Continuing to require a subkey prediction advantage of $s = 128$, the question is, what value of τ guarantees this? This is the question that BKR could not answer. It is answered by our large-alphabet subkey prediction lemma. Specifically, the first table of Fig. 1 gives values of τ for different choices of b .

For $b = 512$ Bytes, we see that $\tau = 43$. Recalling that BKR needed $\tau^* = 271$ probes, we have reduced the number of probes (storage accesses) by a factor of $271/43 \approx 6$, meaning offer a 6x speedup.

The price we pay (as alluded to above) is that J is longer, specifically, 271 bits for BKR and $43 \cdot 512 \approx 22$ KBytes for us. This means the hashing of J to obtain the AES key K takes longer, but modern hash functions are fast, and the time saved in storage accesses is more than the time lost in extra hashing [10, 11]. This is especially true since the hashing can be pipelined, taking advantage of the iterated structure of hash functions to process blocks incrementally as soon as they are retrieved rather than delaying hashing until after all blocks are retrieved.

BIG-KEY IDENTIFICATION. In a (public-key) identification scheme, a user (called the prover) has a secret key sk whose associated public key pk is held by the server (called the verifier). Access control is enforced by having the prover identify itself as the owner of sk via an interactive identification protocol. The Schnorr [18] and Okamoto [16] schemes are well-known examples, but they are of course conventional (small-key) schemes. Identification is an interesting target for a BRM scheme. Here it is the secret key sk that would be big (100 GBytes)— we want the public key pk to remain of conventional size. The usage we envision is hardware-assisted access control, where sk is on an auxiliary device like a USB stick that the user plugs into a possibly infected machine to identify herself (login) to the server across the network. The key being large, and reading from a USB being slow, the malware will have difficulty obtaining enough information about the key (10 GBytes) to violate BRM security, even after a significant number of usages (logins) by the user.

Identification in the BRM was first treated by ADW [2], who gave (asymptotic) security definitions and a clever scheme that involves combining multiple instances of the Okamoto scheme [16] in a compact way. We target making this scheme practical. The quest is meaningless in the absence of concrete security, for practicality is fundamentally about maximizing efficiency for a given level (eg. 128 bits) of security. A first and central step is thus a *concrete-security treatment*. We render the definitions of big-key identification (the goal is security against impersonation under active attack) concretely, then revisit the asymptotically-stated result of ADW [2] to render it, too, in concrete form. We note that for the ADW scheme, probe complexity dictates the computational cost of the two most costly phases of the protocol, the response phase and verification phase (as we will demonstrate in Fig. 14). Hence, improvements in probe complexity directly translate into improvements in efficiency. Towards lowering probe complexity for a given level of security, we first *improve the concrete security* of the reduction via a lemma on the entropy preservation of polynomial evaluation that improves bounds from ADW [2]. We then obtain further reductions in probe complexity, by *using our large-alphabet subkey prediction lemma* in place of ADW’s own [3, Lemma A.3]. The large-alphabet aspect here is crucial, for the scheme draws, from the big key, a value in \mathbb{Z}_p^m , where p is a prime of 512 bits long (for 128-bit security of the identification scheme), and $m \geq 2$ is an integer parameter, so probes need to return blocks of the (large) size $b = m \cdot \lceil \log_2(p) \rceil$. Putting it all together gives a reasonable-cost big-key identification scheme, and the first concrete rendition of the ADW big-key identification scheme. A preliminary implementation shows that with a pairing-friendly group of 512 bits, the execution of the protocol takes a few seconds.

2 Preliminaries

For n a positive integer, we let $[n] = \{0, 1, \dots, n - 1\}$, and $[1..n] = \{1, \dots, n\}$. We also use the notation \mathbb{Z}_n to denote the set $[n]$ in contexts where we use the underlying algebraic structure modulo n . If \mathbf{x} is a vector, then $|\mathbf{x}|$ denotes its length and $\mathbf{x}[i]$ denotes its i -th coordinate. We call \mathbf{x} an n -vector if $|\mathbf{x}| = n$. We number coordinates starting from 0. For example if $\mathbf{x} = (10, 0, 11)$

then $|\mathbf{x}| = 3$ and $\mathbf{x}[2] = 11$. We let ε denote the empty vector, which has length 0. If $0 \leq i \leq |\mathbf{x}| - 1$ then we let $\mathbf{x}[0..i] = (\mathbf{x}[0], \dots, \mathbf{x}[i])$, this being ε when $i = 0$. We say that \mathbf{x} is a vector over set S if all its coordinates belong to S . We let S^n denote the set of all n -vectors over S and we let S^* denote the set of all finite-length vectors over the set S . If S is a set then $|S|$ denotes its size. If $\tau \leq |S|$ is a positive integer, we let $S^{(\tau)}$ be the set of τ -vectors over S with distinct entries. Strings are treated as the special case of vectors over $\{0, 1\}$. Thus, if x is a string then $|x|$ is its length, $x[i]$ is its i -th bit, $x[0..i] = x[0] \dots x[i]$, ε is the empty string, $\{0, 1\}^n$ is the set of n -bit strings and $\{0, 1\}^*$ the set of all strings. For $\mathbf{K} \in [q]^k$ and $\mathbf{p} \in [k]^*$, we let $\mathbf{K}[\mathbf{p}] = (\mathbf{K}[\mathbf{p}[0]], \mathbf{K}[\mathbf{p}[1]], \dots, \mathbf{K}[\mathbf{p}[\mathbf{p} - 1]])$, this being ε when $\mathbf{p} = \varepsilon$.

If X is a finite set, we let $x \leftarrow_s X$ 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 $y \leftarrow A(x_1, \dots; r)$ denote running A with random coins r on inputs x_1, \dots and assigning the output to y . We let $y \leftarrow_s A(x_1, \dots)$ be the result of picking r at random and letting $y \leftarrow A(x_1, \dots; r)$. We let $[A(x_1, \dots)]$ denote the set of all possible outputs of A when invoked with inputs x_1, \dots .

We use the code-based game-playing framework [7] (see Fig. 2 for an example). By $\Pr[\mathbf{G}]$ we denote the probability that game \mathbf{G} returns true. Uninitialized boolean variables, sets and integers are assume initialized to false, the empty set and 0, respectively.

Following [6], our random oracle \mathbf{H} is variable range. This means it takes two inputs, x and \mathbf{Rng} , where the latter is a (description of) an efficiently sampleable set, and returns as output a random point in \mathbf{Rng} . In schemes, we (implicitly or explicitly) fix a unique description for each range set that is used. For example, $x \leftarrow_s \mathbf{H}(x, [k])$ will return a random element in $[k]$, and $\mathbf{p} \leftarrow_s \mathbf{H}(x, [k]^{(\tau)})$ will return a random τ -dimensional vector over $[k]$ with distinct entries.

HAMMING BALLS. Let $q \geq 2$ and $n \geq 1$ be integers. We define the weight of a n -vector v over $[q]$ to be

$$w(v) = \left| \{i \in [n] \mid v[i] \neq 0\} \right|,$$

the number of coordinates of v that are non-zero. Let $\mathcal{K} \subseteq [q]^n$ for some integer n , we define the weight of \mathcal{K} to be

$$w(\mathcal{K}) = \sum_{x \in \mathcal{K}} w(x),$$

the sum of weights of vectors in \mathcal{K} . For $0 \leq r \leq k$, the q -ary hamming ball of radius r over $[q]^k$ is the set

$$\mathbf{B}_{q,k}(r) = \left\{ v \in [q]^k : w(v) \leq r \right\}$$

of k -vectors over $[q]$ that have more at most r non-zero coordinates. We let $B_{q,k}(r)$ denote the size of the set $\mathbf{B}_{q,k}(r)$ and note that

$$B_{q,k}(r) = \sum_{i=0}^r (q-1)^i \binom{k}{i}.$$

For convenience of stating our results, we establish the following conventions: if $r > k$ then we let $B_{q,k}(r) = B_{q,k}(k) = q^k$, and if $k = 0$ then for all $r \geq 0$ we let $B_{q,k}(r) = 1$. We also define the function

$$\text{rd}_{q,k}(N) = \max \{ r \in [k+1] : B_{q,k}(r) \leq N \}$$

to return the largest radius r in the range $0 \leq r \leq k$ such that the ball $\mathbf{B}_{q,k}(r)$ has size at most N .

<p>Game $\mathbf{G}_{q,k,\tau}^{\text{skp}}(\mathcal{A}, \text{Lk})$</p> <p>$\mathbf{K} \leftarrow_{\\$} [q]^k; L \leftarrow \text{Lk}(\mathbf{K})$</p> <p>$\mathbf{p} \leftarrow_{\\$} [k]^{(\tau)}$</p> <p>$J \leftarrow_{\\$} \mathcal{A}(L, \mathbf{p})$</p> <p>Return $(J = \mathbf{K}[\mathbf{p}])$</p> <p>Game $\mathbf{G}_{k,\tau}^{\text{rskp}}(\mathcal{A}, \mathcal{K})$</p> <p>$\mathbf{K} \leftarrow_{\\$} \mathcal{K}$</p> <p>$\mathbf{p} \leftarrow_{\\$} [k]^{(\tau)}$</p> <p>$J \leftarrow_{\\$} \mathcal{A}(\mathbf{p})$</p> <p>Return $(J = \mathbf{K}[\mathbf{p}])$</p>	<p>$q \geq 2$: the alphabet size. A <i>block</i> is an element of $[q]$.</p> <p>k : the length in blocks of the big key</p> <p>$\tau \leq k$: the number probes into the big key \mathbf{K}</p> <p>\mathcal{A} : the adversary</p> <p>Lk : the leakage function, $\text{Lk}: [q]^k \rightarrow [q]^\ell$</p> <p>$\ell$: the length of the output of the leakage function, called the leakage length, in blocks</p> <p>$b \geq 1$: the block length, meaning $q = 2^b$. Theorem 1 does not assume q is a power of two, but it is in some applications.</p> <p>L : the leakage, an ℓ-vector over $[q]$ returned by Lk</p> <p>\mathbf{K} : the big key, a vector of length k over $[q]$</p> <p>\mathbf{p} : the probe vector, a τ-vector over $[k]$ all of whose coordinates are distinct</p> <p>k^* : the length of the big key in bits, $k^* = kb$</p> <p>ℓ^* : the length of the leakage in bits, $\ell^* = \ell b$</p> <p>ρ : the leakage rate, $\rho = \ell^*/k^* = \ell/k$</p>
--	---

Figure 2: **Top Left:** Subkey prediction game $\mathbf{G}_{q,k,\tau}^{\text{skp}}$. **Bottom Left:** Restricted subkey prediction game $\mathbf{G}_{k,\tau}^{\text{rskp}}$ used in Section 3.3. **Right:** Summary of quantities involved.

3 Large-Alphabet Subkey Prediction

Here we define the subkey prediction problem parameterized by alphabet size and give our results about it.

3.1 The Problem

We consider the subkey-prediction game, $\mathbf{G}_{q,k,\tau}^{\text{skp}}(\mathcal{A}, \text{Lk})$, shown on the top left of Fig. 2. (Ignore the game below it for now.) The quantities involved in the game, as well as associated ones, are summarized on the right of the same Figure. In the game, a k -vector \mathbf{K} over $[q]$, called the big key, is randomly chosen from $[q]^k$. Then, a random τ -vector \mathbf{p} is chosen from $[k]^{(\tau)}$, so that its coordinates are all distinct. (Recall that $[k]^{(\tau)}$, the set from which \mathbf{p} is selected in the game in Fig. 2, denotes the set of all τ -vectors over $[k]$ all of whose coordinates are distinct.) We refer to \mathbf{p} as the probe vector. Each of its coordinates is a probe, pointing to a location in the big key. Adversary \mathcal{A} is given the leakage $L = \text{Lk}(\mathbf{K})$ and the probe vector \mathbf{p} . Its goal is to predict (compute, output) $\mathbf{K}[\mathbf{p}] = (\mathbf{K}[\mathbf{p}[1]], \dots, \mathbf{K}[\mathbf{p}[\tau]])$, the τ -vector consisting of the coordinates of \mathbf{K} selected by the coordinates of the probe vector. The adversary returns J as its guess, and the game returns true if \mathcal{A} succeeds, meaning $J = \mathbf{K}[\mathbf{p}]$. We define the following advantage metrics:

$$\begin{aligned} \text{Adv}_{q,k,\tau}^{\text{skp}}(\mathcal{A}, \text{Lk}) &= \Pr \left[\mathbf{G}_{q,k,\tau}^{\text{skp}}(\mathcal{A}, \text{Lk}) \right], \\ \text{Adv}_{q,k,\tau}^{\text{skp}}(\text{Lk}) &= \max_{\mathcal{A}} \text{Adv}_{q,k,\tau}^{\text{skp}}(\mathcal{A}, \text{Lk}), \\ \text{Adv}_{q,k,\tau}^{\text{skp}}(\ell) &= \max_{\text{Lk}: [q]^k \rightarrow [q]^\ell} \text{Adv}_{q,k,\tau}^{\text{skp}}(\text{Lk}). \end{aligned}$$

The first advantage is the probability that the game outputs **true**, meaning the probability that the adversary successfully returns $\mathbf{K}[\mathbf{p}]$. The second advantage is obtained by maximizing the first one over all adversaries \mathcal{A} . Note that this is well-defined since here we consider all computationally unbounded adversaries. The third advantage is obtained by maximizing the second advantage over all leakage functions Lk that output ℓ blocks.

Now fix some big-key length k^* (in bits) and leakage length ℓ^* (in bits). Also fix an integer s representing the desired security. For any block length $b \geq 1$ such that b divides k^* and ℓ^* , we let

$$\mathbf{Probes}_{k^*, \ell^*, s}(b) = \min \left\{ \tau : \mathbf{Adv}_{2^b, k^*/b, \tau}^{\text{skp}}(\ell^*/b) \leq 2^{-s} \right\}. \quad (3)$$

Here, we have set the alphabet size to $q = 2^b$. The length k of the big key and ℓ of the leakage in blocks are determined, respectively, by $k = k^*/b$ and $\ell = \ell^*/b$. Then, $\mathbf{Probes}_{k^*, \ell^*, s}(b)$ is the smallest number of probes τ that will guarantee that $\mathbf{Adv}_{q, k, \tau}^{\text{skp}}(\ell)$ is at most 2^{-s} .

The subkey prediction game and problem formulated by BKR [4] differs in two ways. First, they had only considered the $q = 2$ case (that is, $b = 1$) of a binary alphabet. The large alphabet aspect of our treatment refers to the fact that our alphabet size is a parameter q that we view as quite large. In some applications, $q = 2^b$ where b is the block size of our storage medium, but Theorem 1 does not assume q is a power of two. The second difference with BKR [4] is that their probes $\mathbf{p}[1], \dots, \mathbf{p}[\tau]$ were random and independent, so in particular two of them might be the same, but ours are random subject to being distinct. This is important towards our being able to get a provable upper bound on the subkey prediction advantage, whereas BKR were only able to get (for their setting) an estimate or approximate upper bound.

Now our goal is to upper bound, as well as possible, the subkey prediction advantage $\mathbf{Adv}_{q, k, \tau}^{\text{skp}}(\ell)$ as a function of q, k, τ, ℓ . Thence we will obtain upper bounds on $\mathbf{Probes}_{k^*, \ell^*, s}(b)$.

3.2 Subkey Prediction Theorem

The bound in our subkey prediction theorem is the ratio of the sizes of two q -ary hamming balls. We refer to Section 2 for definitions.

Theorem 1 (Subkey-prediction bound). *Let q, k, ℓ, τ be integers with $q \geq 2$ and $\ell, \tau \leq k$. Let r be any integer in the range $0 \leq r \leq \text{rd}_{q, k}(q^{k-\ell})$. Then*

$$\mathbf{Adv}_{q, k, \tau}^{\text{skp}}(\ell) \leq \frac{B_{q, k-\tau}(r)}{B_{q, k}(r)}. \quad (4)$$

The theorem allows us to pick the parameter r arbitrarily in the given range, so for the best estimates we would pick a r that minimizes the ratio. We postpone the proof to first discuss how this compares to prior work and how to use it to get numerical bounds.

COMPARISON. BKR [4] give an upper bound we denote $G_{k, \tau}^{\text{bkr}}(2^{k-\ell})$ on the subkey prediction advantage in their setting. Recall that their setting differs from ours in two ways. First, $q = 2$ in their case. Second, in their game, the τ probes are random and independent, while in our game they are random but distinct. Their function $G_{k, \tau}^{\text{bkr}}(N)$ is a sum of $\text{rd}_{2, k}(N)$ terms. It is quite complex and it is hard to estimate numerically. BKR gave a simpler expression, that approximates $G_{k, \tau}^{\text{bkr}}(N)$, and that they use for numerical estimates, but this expression is not an upper bound, and thus it is not clear their numerical estimates are upper bounds either. Our bound, the ratio of the sizes of two q -ary Hamming balls, is simpler than the bound of BKR (this makes crucial use of the probes being distinct), and, we will see, more analytically tractable, even when $q = 2$. In

particular, we are able to upper bound $\min_r B_{q,k-\tau}(r)/B_{q,k}(r)$, subjected to $0 \leq r \leq \text{rd}_{q,k}(q^{k-\ell})$, quite nicely for numerical estimates, as discussed next.

TOOLS FOR DERIVING NUMERICAL BOUNDS. Theorem 1 upper bounds the subkey prediction advantage as the ratio of the sizes of two hamming balls. Below, we present tools to bound this ratio. First, we need some definitions. Let H_2 be the binary entropy function, defined for $x \in [0, 1]$ by $H_2(x) = -x \log_2(x) - (1-x) \log_2(1-x)$. We note that the value of $x \log_q(x)$ is taken to be 0 when $x = 0$. This ensures that H_2 is continuous over $[0, 1]$. More generally, for an integer $q \geq 2$ the q -ary entropy function is defined for $x \in [0, 1]$ by

$$\begin{aligned} H_q(x) &= x \log_q(q-1) - x \log_q(x) - (1-x) \log_q(1-x) \\ &= \frac{H_2(x)}{\log_2(q)} + x \log_q(q-1). \end{aligned}$$

We note that H_q attains its maximum at $x = 1 - 1/q$. We define its inverse function, $H_q^{-1} : [0, 1] \rightarrow [0, 1 - 1/q]$ to be such that $H_q^{-1}(H_q(x)) = x$ for any $x \in [0, 1 - 1/q]$. We define the following error function for $q \geq 2$ and $0 \leq r \leq k$,

$$\epsilon(q, k, r) = \log_q(e) \left(\frac{1}{12r} + \frac{1}{12(k-r)} - \frac{1}{12k+1} \right) + \frac{1}{2} \log_q \left(\frac{2\pi r(k-r)}{k} \right). \quad (5)$$

The following lemmas, the proofs of which are given in Appendix A, are key to deriving numerical bounds. The first gives both upper and lower bounds on the size of a Hamming ball.

Lemma 2. *Let k, q, r be integers with $q \geq 2$ and $0 \leq r \leq k$. Then,*

$$q^{kH_q(r/k) - \epsilon(q,k,r)} \leq B_{q,k}(r). \quad (6)$$

Additionally, if $0 \leq r \leq k(1 - 1/q)$,

$$B_{q,k}(r) \leq q^{kH_q(r/k)}. \quad (7)$$

The second lemma lower bounds the value of $\text{rd}_{q,k}(N)$.

Lemma 3. *Let N, q, k be positive integers such that $q \geq 2$ and $N \leq q^k$. Then,*

$$\left\lfloor H_q^{-1} \left(\frac{\log_q(N)}{k} \right) \cdot k \right\rfloor \leq \text{rd}_{q,k}(N).$$

The following provides a two-sided bound on H_q^{-1} :

Lemma 4. *Let $q \geq 2$ be an integer, and $x \in [0, 1]$ a real number. Then,*

$$\min(x, 1 - \frac{1}{q}) - \frac{1}{\log_2(q)} \leq H_q^{-1}(x) \leq x(1 - \frac{1}{q}).$$

These bounds are good when q is large.

DERIVING NUMERICAL BOUNDS. We now use the above to derive upper bounds for example parameter values. Let $b \geq 1$ be a block size, so that the alphabet has size $q = 2^b$. Fix some big-key length k^* (in bits) and leakage length ℓ^* (in bits) that are multiples of b , and let $k = k^*/b$ and $\ell = \ell^*/b$ be the big-key and leakage lengths, respectively, in blocks. We assume that τ and ℓ satisfy that $\ell \geq \tau$, as the below method only apply when this condition is met. We note that this is a reasonable

assumption for practical applications, as leakage length ℓ is usually large, and we are attempting to keep the probe complexity, τ , small. Now, suppose we have obtained some integer value r such that: (1) $r \leq \text{rd}_{q,k}(q^{k-\ell})$ and (2) $0 \leq r \leq (k - \tau)(1 - 1/q)$. Then, we use Equation (6) to lower bound $B_{q,k}(r)$. Given condition (2), we can use Equation (7) to upper bound the quantity $B_{q,k-\tau}(r)$. This results in an upper bound, denoted $\text{RatioBound}_{q,k,\ell,\tau}(r)$, for the ratio $B_{q,k-\tau}(r)/B_{q,k}(r)$:

$$\text{RatioBound}_{q,k,\tau}(r) = \frac{q^{(k-\tau)H_q(r/(k-\tau))}}{q^{kH_q(r/k) - \epsilon(q,k,r)}} .$$

Note that in the above expression, the terms $H_q(r/(k-\tau))$, $H_q(r/k)$ and $\epsilon(q,k,r)$ can be computed numerically for any given value of q, k, τ and r . Hence, deriving numerical upper bound for the ratio $B_{q,k-\tau}(r)/B_{q,k}(r)$ amounts to obtaining a value r satisfying the two conditions given above. We take r to be $r_{q,k,\ell}$, defined as

$$r_{q,k,\ell} = \left\lfloor H_q^{-1}\left(\frac{k-\ell}{k}\right) \cdot k \right\rfloor .$$

Here, we assume that a method of obtaining numerical lower bounds for $H_q^{-1}(x)$ is available¹. We now check the two conditions required. For condition (1), we know that $r_{q,k,\ell} \leq \text{rd}_{q,k}(q^{k-\ell})$ by Lemma 3 (taking $N = q^{k-\ell}$). For condition (2), note that by Lemma 4 and the assumption that $\ell \geq \tau$,

$$H_q^{-1}\left(\frac{k-\ell}{k}\right) \leq \frac{k-\ell}{k}(1 - 1/q) \leq \frac{k-\tau}{k}(1 - 1/q) .$$

Hence,

$$r_{q,k,\ell} = \left\lfloor H_q^{-1}\left(\frac{k-\ell}{k}\right) \cdot k \right\rfloor \leq (k - \tau)(1 - 1/q) .$$

We consider the quantity

$$\overline{\mathbf{Adv}}_{q,k,\tau}^{\text{skp}}(\ell) = \text{RatioBound}_{q,k,\ell,\tau}(r_{q,k,\ell}) . \quad (8)$$

We note that since $r = r_{q,k,\ell}$ satisfies condition (1) and (2), by Theorem 1 and above analysis,

$$\mathbf{Adv}_{q,k,\tau}^{\text{skp}}(\ell) \leq \frac{B_{q,k-\tau}(r_{q,k,\ell})}{B_{q,k}(r_{q,k,\ell})} \leq \overline{\mathbf{Adv}}_{q,k,\tau}^{\text{skp}}(\ell) .$$

Hence, $\overline{\mathbf{Adv}}_{q,k,\tau}^{\text{skp}}(\ell)$ is an upper bound for $\mathbf{Adv}_{q,k,\tau}^{\text{skp}}(\ell)$. Now, given a particular desired security level, s , we want to find the smallest τ such that $\overline{\mathbf{Adv}}_{q,k,\tau}^{\text{skp}}(\ell) \leq 2^{-s}$. We let

$$\overline{\mathbf{Probes}}_{k^*,\ell^*,b}(s) = \min \left\{ \tau \in [k + 1] : \overline{\mathbf{Adv}}_{q,k,\tau}^{\text{skp}}(\ell) \leq 2^{-s} \right\} .$$

Note that this is similar to the definition of $\mathbf{Probes}_{k^*,\ell^*,b}(s)$ (Equation (3)), only that $\mathbf{Adv}_{q,k,\tau}^{\text{skp}}(\ell)$ is replaced with $\overline{\mathbf{Adv}}_{q,k,\tau}^{\text{skp}}(\ell)$. Thence,

$$\mathbf{Probes}_{k^*,\ell^*,b}(s) \leq \overline{\mathbf{Probes}}_{k^*,\ell^*,b}(s) . \quad (9)$$

We note that $\overline{\mathbf{Probes}}_{k^*,\ell^*,b}(s)$ can be computed numerically by iteratively incrementing τ and computing $\overline{\mathbf{Adv}}_{q,k,\tau}^{\text{skp}}(\ell)$. Fig. 1 shows values of $\overline{\mathbf{Probes}}_{k^*,\ell^*,b}(s)$ for various practical values of k^* , ℓ^* , b and s .

PLOTS. For the left plot, we fix the following:

¹For example, this is available in mathematical software **Sage**. Also, when q is large, Lemma 4 provides a good lower bound for H_q^{-1} that is easily computed numerically.




Figure 3: Fix the big key length k^* to be 100 GBytes. The left graph plots (an upper bound on) $\mathbf{Probes}_{k^*, \rho k^*, 128}(32)$ as a function of the leakage rate ρ . The right graph plots (a lower bound on) $-\log_2(\mathbf{Adv}_{2^{32}, k, 47}^{\text{skp}}(\rho k))$ as a function of ρ , where $k = k^*/32$.

- Blocksize $b = 32$ bits, so that $q = 2^{32}$.
- Leakage length $\ell^* = 8 \cdot 10^{10}$ bits = 10 GBytes, so that $\ell = \ell^*/32$.
- Desired security level $s = 128$ bits.

The left graph in Fig. 3 plots $\overline{\mathbf{Probes}}_{\ell^*/\rho, \ell^*, b}(s)$, upper bound for $\mathbf{Probes}_{\ell^*/\rho, \ell^*, s}(b)$, as a function of the leakage rate ρ . The left plot shows that the number of probes needed to maintain s bits of security increases faster once the leakage rate goes over 50%. Hence, for applications, it may be beneficial to use big keys that are big enough so that the leakage rate can be assumed to be less than 50%. For example, if 10 GBytes is the leakage bound, one might, for efficiency, target big key of size at least 20 GBytes.

For the right plot, we fix the following

- Blocksize $b = 32$ bits, so that $q = 2^{32}$.
- Big key length $k^* = 8 \cdot 10^{11}$ bits = 100 GBytes, so that $k = k^*/32$.
- Number of probes $\tau = 47$.

The number 47 has been chosen because, as per Fig. 1, it ensures $\mathbf{Adv}_{q, k, \tau}^{\text{skp}}(k/10) \leq 2^{-128}$. Now with b, k^*, τ (and thus also q, k) fixed, the right graph plots $-\log_2(\overline{\mathbf{Adv}}_{q, k, \tau}^{\text{skp}}(\rho \cdot k))$, lower bound for $-\log_2(\mathbf{Adv}_{q, k, \tau}^{\text{skp}}(\rho \cdot k))$, as a function of leakage rate ρ . The right plot in Fig. 3 demonstrates that, even though a scheme is designed for 10% leakage, security degrades gradually as the leakage rate goes over 10%.

3.3 Proof of Theorem 1

We follow the framework of the proof of BKR [4].

RESTRICTED SUBKEY PREDICTION. The proof involves consideration of a simpler game, called the restricted subkey prediction game, denoted \mathbf{G}^{rskp} and shown on the right in Fig. 2. Game \mathbf{G}^{rskp} is similar to game \mathbf{G}^{skp} , except that there is no leakage function Lk and leakage L . Instead, the big key \mathbf{K} is drawn from a restricted subset $\mathcal{K} \subseteq [q]^k$ of big keys. We define the following advantage

metrics:

$$\begin{aligned}\mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{A}, \mathcal{K}) &= \Pr \left[\mathbf{G}_{k,\tau}^{\text{rskp}}(\mathcal{A}, \mathcal{K}) \right], \\ \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}) &= \max_{\mathcal{A}} \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{A}, \mathcal{K}), \\ \mathbf{Adv}_{q,k,\tau}^{\text{rskp}}(N) &= \max_{\mathcal{K} \subseteq [q]^k, |\mathcal{K}|=N} \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}).\end{aligned}$$

The first advantage is the probability that the game outputs true, meaning the probability that the adversary successfully returns $\mathbf{K}[\mathbf{p}]$. The second advantage is obtained by maximizing the first one over all adversaries \mathcal{A} . The third advantage is obtained by maximizing the second advantage over all sets $\mathcal{K} \subseteq [q]^k$ that have size N . We note that the first two advantages do not have q in the subscript, which is due to the fact that \mathcal{K} encodes the value of q .

MONOTONE SETS. Let x, x' be vectors in $[q]^k$. We say that x dominates x' , or x' is dominated by x , written $x' \leq x$, if x' can be obtained by changing non-zero coordinates of x to 0. We let

$$\text{DS}_{q,k}(x) = \{x' \in [q]^k : x' \leq x\}$$

be the set of all x' dominated by x . A set $\mathcal{K} \subseteq [q]^k$ is *monotone* if

$$\bigcup_{x \in \mathcal{K}} \text{DS}_{q,k}(x) \subseteq \mathcal{K}.$$

That is, if $x \in \mathcal{K}$, and x' is dominated by x , then $x' \in \mathcal{K}$. For example, a Hamming ball in $[q]^k$, of any radius, is a monotone set.

SOME NOTATION. For integers $x, \tau \geq 0$, we let

$$x_{(\tau)} = \prod_{i=0}^{\tau-1} (x - i) = \prod_{j=x-\tau+1}^x j. \quad (10)$$

Notice that $x_{(\tau)} = 0$ if $\tau > x$. This can be seen because, if $\tau > x$, then, in the second product above, the starting value for j is ≤ 0 , and since $x \geq 0$, this means the term $j = 0$ is included in the product. Also when $\tau = 0$, the product has zero terms, and hence by convention takes value 1, meaning $x_{(0)} = 1$ for all $x \geq 0$. We use below the notation from Equation (10).

For a nonempty $\mathcal{K} \subseteq [q]^k$, we define the function

$$g_{k,\tau}(\mathcal{K}) = \frac{1}{|\mathcal{K}|} \sum_{x \in \mathcal{K}} \frac{(k - w(x))_{(\tau)}}{k_{(\tau)}}. \quad (11)$$

The following lemma says that if \mathcal{K} is monotone, then the restricted subkey prediction advantage for big keys drawn from \mathcal{K} can be expressed *exactly*, and in particular by the function of Equation (11).

Lemma 5. *Let q, τ, k be positive integers such that $\tau \leq k$ and $q \geq 2$. Let $\mathcal{K} \subseteq [q]^k$ be a non-empty monotone set. Then,*

$$\mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}) = g_{k,\tau}(\mathcal{K}).$$

Proof (of Lemma 5). Let \mathcal{A}_0 be the adversary that, on input \mathbf{p} , always returns the all-0 τ -vector. We claim that this adversary maximizes the advantage, meaning

$$\mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}) = \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}, \mathcal{A}_0).$$

This follows from the assumption that \mathcal{K} is monotone. Now, we compute the advantage of \mathcal{A}_0 . For $\mathbf{K} \in [q]^k$, let $Z(\mathbf{K})$ denote the set of all $\mathbf{p} \in [k]^{(\tau)}$ such that $\mathbf{K}[\mathbf{p}] = (0, \dots, 0)$. We have

$$\mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}, \mathcal{A}_0) = \frac{1}{|\mathcal{K}|} \sum_{\mathbf{K} \in \mathcal{K}} \frac{|Z(\mathbf{K})|}{|[k]^{(\tau)}|} = \frac{1}{|\mathcal{K}|} \sum_{\mathbf{K} \in \mathcal{K}} \frac{(k - w(\mathbf{K}))_{(\tau)}}{k_{(\tau)}} = g_{k,\tau}(\mathcal{K}).$$

□

We say that a set $\mathcal{K} \subseteq [q]^k$ is *sandwiched between hamming balls* if

$$B_{q,k}(r) \subseteq \mathcal{K} \subseteq B_{q,k}(r+1)$$

for $r = \text{rd}_{q,k}(|\mathcal{K}|)$. For N an integer such that $1 \leq N \leq q^k$, we define

$$G_{q,k,\tau}(N) = \frac{1}{N} \sum_{i=0}^{\text{rd}_{q,k}(N)} (q-1)^i \binom{k}{i} \frac{(k-i)_{(\tau)}}{k_{(\tau)}} + \left(1 - \frac{B_{q,k}(\text{rd}_{q,k}(N))}{N}\right) \frac{(k - (\text{rd}_{q,k}(N) + 1))_{(\tau)}}{k_{(\tau)}}. \quad (12)$$

The following says that if \mathcal{K} is monotone and sandwiched between Hamming balls, then the restricted subkey prediction advantage for big keys drawn from \mathcal{K} can be expressed *exactly*, and in particular by the function of Equation (12).

Lemma 6. *Let q, τ, k be positive integers such that $\tau \leq k$ and $q \geq 2$. Let $\mathcal{K} \subseteq [q]^k$ be a non-empty monotone set that is also sandwiched between hamming balls, i.e. $\mathbf{B}_{q,k}(r) \subseteq \mathcal{K} \subseteq \mathbf{B}_{q,k}(r+1)$ for $r = \text{rd}_{q,k}(|\mathcal{K}|)$. Then*

$$\mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}) = G_{q,k,\tau}(|\mathcal{K}|).$$

Proof (of Lemma 6). Let $N = |\mathcal{K}|$. By Lemma 5, we have

$$\mathbf{Adv}_{k,\tau}^{\text{skp}}(\mathcal{K}) = \frac{1}{N} \sum_{x \in \mathcal{K}} \frac{(k - w(x))_{(\tau)}}{k_{(\tau)}}.$$

Since $\mathbf{B}_{q,k}(r) \subseteq \mathcal{K} \subseteq \mathbf{B}_{q,k}(r+1)$. This means $\mathbf{B}_{q,k}(i) \subseteq \mathcal{K}$ for $i = 0, \dots, r$, and \mathcal{K} contains $N - B_{q,k}(r)$ vectors of weight $r+1$. Thus, the above equals

$$\frac{N - B_{q,k}(r)}{N} \frac{(k - r - 1)_{(\tau)}}{k_{(\tau)}} + \frac{1}{N} \sum_{i=0}^r (q-1)^i \binom{k}{i} \frac{(k-i)_{(\tau)}}{k_{(\tau)}} = G_{q,k,\tau}(N)$$

as claimed. □

Next, we show that monotone sets sandwiched between Hamming balls are the extremal cases for the restricted subkey prediction game, meaning that they maximize the restricted subkey prediction advantage. The following is analogous to [4, Lemmas 6,8]. We streamline their analysis and extend it to large alphabets.

Lemma 7. *Let q, k, N be positive integers. Suppose $q \geq 2$, $N \leq q^k$ and $\tau \leq k$. Then, there is a non-empty monotone set $\mathcal{K} \subseteq [q]^k$ of size N such that*

$$\mathbf{Adv}_{q,k,\tau}^{\text{rskp}}(N) = \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}).$$

Additionally, \mathcal{K} is also sandwiched between hamming balls, i.e. for $r = \text{rd}_{q,k}(N)$,

$$\mathbf{B}_{q,k}(r) \subseteq \mathcal{K} \subseteq \mathbf{B}_{q,k}(r+1).$$

The proof of Lemma 7 is deferred to Section 3.3.1. As a direct corollary of Lemma 6 and Lemma 7, we get the following result.

Corollary 8. Let q, τ, k be positive integers such that $\tau \leq k$ and $q \geq 2$. Then,

$$\mathbf{Adv}_{q,k,\tau}^{\text{rskp}}(N) = G_{q,k,\tau}(N). \quad (13)$$

Hence, from this point on, we identify the two functions $\mathbf{Adv}_{q,k,\tau}^{\text{rskp}}(\cdot)$ and $G_{q,k,\tau}(\cdot)$. Next, we observe a useful property of $G_{q,k,\tau}(N)$. In particular, it is decreasing in the domain $[1..q^k]$.

Lemma 9. Let q, τ, k be positive integers such that $\tau \leq k$ and $q \geq 2$. Let i, j be integers such that $1 \leq i \leq j \leq q^k$. Then,

$$G_{q,k,\tau}(i) \geq G_{q,k,\tau}(j).$$

We proceed to relate the restricted subkey-prediction game to the subkey-prediction game via the lemma below.

Lemma 10. Let ℓ, q, k, τ be integers such that $0 \leq \ell \leq k$, $q \geq 2$, and $1 \leq \tau \leq k$. Then,

$$\mathbf{Adv}_{q,k,\tau}^{\text{skp}}(\ell) \leq \mathbf{Adv}_{q,k,\tau}^{\text{rskp}}(q^{k-\ell}).$$

The proofs of Lemma 9 and Lemma 10 are deferred to Section 3.3.2. Finally, we give a way to bound the expression $G_{q,k,\tau}(N)$. In particular, we show that it is at most the ratio of two hamming balls of the same radius $\text{rd}_{q,k}(N)$; one with dimension $k - \tau$ and one with dimension k . Recall that BKR did not give concrete numerical upper bounds for their subkey-prediction advantage, only *estimates*. Due to assuming the uniqueness of probes, we are able to simplify our expression $G_{q,k,\tau}(N)$. In particular, we note that for non-negative integers k, i, τ such that $i, \tau \leq k$,

$$\binom{k}{i} \frac{(k-i)_{(\tau)}}{k_{(\tau)}} = \frac{k_{(i)}}{i_{(i)}} \cdot \frac{(k-i)_{(\tau)}}{k_{(\tau)}} = \frac{(k-\tau)_{(i)}}{i_{(i)}} = \binom{k-\tau}{i}. \quad (14)$$

This property allows us to prove the following lemma.

Lemma 11. Let N, q, k, τ, r be positive integers such that $q \geq 2$, $N \leq q^k$, $\tau \leq k$ and $r \leq \text{rd}_{q,k}(N)$. Then

$$G_{q,k,\tau}(N) \leq \frac{B_{q,k-\tau}(r)}{B_{q,k}(r)}.$$

Proof (of Lemma 11). By Lemma 9,

$$G_{q,k,\tau}(N) \leq G_{q,k,\tau}(B_{q,k}(r)).$$

By Equation (12) and Equation (14),

$$\begin{aligned} G_{q,k,\tau}(B_{q,k}(r)) &\leq \frac{1}{B_{q,k}(r)} \sum_{i=0}^{\text{rd}_{q,k}(B_{q,k}(r))} (q-1)^i \binom{k}{i} \frac{(k-i)_{(\tau)}}{k_{(\tau)}} \\ &= \frac{1}{B_{q,k}(r)} \sum_{i=0}^r (q-1)^i \binom{k-\tau}{i} \\ &= \frac{B_{q,k-\tau}(r)}{B_{q,k}(r)}. \end{aligned}$$

□

The proof of Theorem 1 follows directly.

Proof. (of Theorem 1). Note that when $r = 0$, Equation (4) is trivially true. Hence, we let $r \leq \text{rd}_{q,k}(N)$ be a positive integer. Then,

$$\mathbf{Adv}_{q,k,\tau}^{\text{skp}}(l) \leq \mathbf{Adv}_{q,k,\tau}^{\text{rskp}}(q^{k-l}) \quad (\text{Lemma 10})$$

$$= G_{q,k,\tau}(q^{k-l}) \quad (\text{Corollary 8})$$

$$\leq \frac{B_{q,k-\tau}(r)}{B_{q,k}(r)}. \quad (\text{Lemma 11})$$

□

3.3.1 Proof of Lemma 7

Proof (of Lemma 7). Let

$$\mathcal{T} = \left\{ \mathcal{K} \subseteq [q]^k : |\mathcal{K}| = N \text{ and } \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}) = \mathbf{Adv}_{q,k,\tau}^{\text{rskp}}(N) \right\}.$$

Let $\mathcal{K} \in \mathcal{T}$ be the minimal weight element, i.e. the element $\mathcal{K} \in \mathcal{T}$ that minimizes the value $w(\mathcal{K}) = \sum_{x \in \mathcal{K}} w(x)$. We will show that \mathcal{K} is a set satisfying the properties claimed in the lemma. We will prove the two properties separately, namely that \mathcal{K} is monotone and $B_{q,k}(r) \subseteq \mathcal{K} \subseteq B_{q,k}(r+1)$. We first claim that \mathcal{K} is monotone. The idea is to define a “shifting” operation for any set $\mathcal{K}' \subseteq [q]^k$ at a coordinate to increase $\mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}')$ while decreasing $w(\mathcal{K}')$. Seeking a contradiction, suppose \mathcal{K} is not monotone. Without loss of generality, suppose that for all pairs of $x \in \mathcal{K}$ and $y \notin \mathcal{K}$ such that $y \leq x$, we have that x and y differ only in the first component. We build another set \mathcal{K}' with the following properties.

1. $|\mathcal{K}'| = |\mathcal{K}|$
2. $w(\mathcal{K}') \leq w(\mathcal{K})$
3. $\mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}') \geq \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K})$

We first explain briefly explain the construction of \mathcal{K}' on the high level before giving the formal construction. Let $z \in [q]^{k-1}$. We will attempt to “swap” vectors of the form $\alpha||z$, for $\alpha \in [q]$, in and out of \mathcal{K} . The swapping is done in two cases. We define D_z to contain the α 's such that $\alpha||z \in \mathcal{K}$. First, if $0 \in D_z$ or $D_z = \emptyset$, no swapping will be done. Second, if $0 \notin D_z$ and $D_z \neq \emptyset$, then we will do the following. Let $\beta = \max D_z$. We will remove the element $\beta||z$ from \mathcal{K} and add the element $0||z$ to \mathcal{K} . After such operations are done for all $z \in [q]^{k-1}$, the resulting set will be \mathcal{K}' . Formally, the construction of \mathcal{K}' is given below. \mathcal{K}' is constructed from \mathcal{K} via the function $\phi : [q]^k \rightarrow [q]^k$, which is defined relative to the set B (set A is used in the later analysis). Sets A and B partition the set of strings of length $k-1$. Set A consists of z 's such that no swapping will be done. Set B consists of z 's such that swapping will be done. The formal definition for A, B, ϕ , and \mathcal{K}' is as

follows:

$$\begin{aligned}
A &= \left\{ z \in [q]^{k-1} : 0 \in D_z \text{ or } D_z = \emptyset \right\}, \\
B &= \left\{ z \in [q]^{k-1} : 0 \notin D_z \text{ and } D_z \neq \emptyset \right\}, \\
\phi(\alpha \| z) &= \begin{cases} 0 \| z & \text{if } z \in B \text{ and } \alpha = \max D_z \\ (\max D_z) \| z & \text{if } z \in B \text{ and } \alpha = 0 \\ \alpha \| z & \text{otherwise} \end{cases}, \\
\mathcal{K}' &= \left\{ \phi(x) : x \in \mathcal{K} \right\}.
\end{aligned}$$

By construction, we note that the swapping operation preserves the size of the set and only decreases its overall weight. Hence, $|\mathcal{K}'| = |\mathcal{K}|$ and $w(\mathcal{K}') \leq w(\mathcal{K})$. It remains to show property (3). Let \mathcal{A} be an adversary such that $\mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{A}, \mathcal{K}) = \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K})$. Consider the adversary \mathcal{A}' that behaves exactly as \mathcal{A} with the exception that it always guess 0 for the first position. More precisely, \mathcal{A}' does the following.

Adversary $\mathcal{A}'((s_1, \dots, s_\tau))$
 $J' \leftarrow \mathcal{A}((s_1, \dots, s_\tau))$
For $i \leftarrow 1, \dots, \tau$ do
 If $s_i = 1$ then $J'[i] \leftarrow 0$
Return J'

Let $P(\cdot)$ denote the probability function in game $\mathbf{G}_{k,\tau}^{\text{rskp}}(\mathcal{A}, \mathcal{K})$ and $P'(\cdot)$ the probability function in game $\mathbf{G}_{k,\tau}^{\text{rskp}}(\mathcal{A}', \mathcal{K}')$. We now define three events for both games $\mathbf{G}_{k,\tau}^{\text{rskp}}(\mathcal{A}, \mathcal{K})$, $\mathbf{G}_{k,\tau}^{\text{rskp}}(\mathcal{A}', \mathcal{K}')$, where $z \in [q]^{k-1}$.

WIN : The game returns true
ONE : $1 \in \{s_1, \dots, s_\tau\}$
 S_z : $(\mathbf{K}[1..k] = z)$, ONE and $(\forall i, s_i \neq 1 : J'[i] = \mathbf{K}[s_i])$

Note that $P(\text{ONE}) = P'(\text{ONE})$, and $P(\text{WIN} | \neg \text{ONE}) = P'(\text{WIN} | \neg \text{ONE})$. We claim that $P(\text{WIN} | \text{ONE}) \leq P'(\text{WIN} | \text{ONE})$. If so we have

$$\begin{aligned}
\mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{A}, \mathcal{K}) &= P(\text{WIN}) \\
&= P(\text{WIN} | \text{ONE}) \cdot P(\text{ONE}) + P(\text{WIN} | \neg \text{ONE}) \cdot P(\neg \text{ONE}) \\
&= P(\text{WIN} | \text{ONE}) \cdot P'(\text{ONE}) + P'(\text{WIN} | \neg \text{ONE}) \cdot P'(\neg \text{ONE}) \\
&\leq P'(\text{WIN} | \text{ONE}) \cdot P'(\text{ONE}) + P'(\text{WIN} | \neg \text{ONE}) \cdot P'(\neg \text{ONE}) \\
&= P'(\text{WIN}) = \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{A}', \mathcal{K}').
\end{aligned}$$

So now we need to show that $P(\text{WIN} \mid \text{ONE}) \leq P'(\text{WIN} \mid \text{ONE})$. We have

$$\begin{aligned} P(\text{WIN} \mid \text{ONE}) &= \sum_{z \in [q]^{k-1}} P(\text{WIN} \mid S_z) \cdot P(S_z) \\ &= \sum_{z \in [q]^{k-1}} P(\text{WIN} \mid S_z) \cdot P'(S_z) \end{aligned} \quad (15)$$

$$\begin{aligned} &\leq \sum_{z \in [q]^{k-1}} P'(\text{WIN} \mid S_z) \cdot P'(S_z) \\ &= P'(\text{WIN} \mid \text{ONE}) . \end{aligned} \quad (16)$$

Equation (15) is true because $P(S_z) = P'(S_z)$ for all $z \in [q]^{k-1}$, since the swapping operation do not change the last $k-1$ component of any vector. Next, we argue the validity of Equation (16). Let $z \in [q]^{k-1}$ such that $P(S_z) \neq 0$ (and hence $P'(S_z) \neq 0$), which means that there is some $\alpha \in [q]$ such that $\alpha \parallel z \in \mathcal{K}$. For any $z \in [q]^{k-1}$, consider the sets $U_z = \{\alpha \in [q] : \alpha \parallel z \in \mathcal{K}\}$ and $V_z = \{\alpha \in [q] : \alpha \parallel z \in \mathcal{K}'\}$. Note that $P(\text{WIN} \mid S_z) \leq 1/|U_z|$ and $P'(\text{WIN} \mid S_z) \leq 1/|V_z|$. Additionally, we note that $|U_z| = |V_z|$, and V_z always contains 0. Since \mathcal{A}' always guess 0 for the first component, we have $P'(\text{WIN} \mid S_z) = 1/|V_z|$. Therefore,

$$P(\text{WIN} \mid S_z) \leq \frac{1}{|U_z|} = \frac{1}{|V_z|} = P'(\text{WIN} \mid S_z) .$$

Next, we show that \mathcal{K} must be sandwiched between two hamming balls. We first claim that $\mathbf{B}_{q,k}(r) \subseteq \mathcal{K}$. Seeking a contradiction, suppose that $\mathbf{B}_{q,k}(r) \not\subseteq \mathcal{K}$. Let x' be a point in $\mathbf{B}_{q,k}(r) \setminus \mathcal{K}$ of minimal Hamming weight. Let x be a point in $\mathcal{K} \setminus \mathbf{B}_{q,k}(r)$ of maximal Hamming weight. We claim that $w(x) > w(x')$, otherwise $\mathbf{B}_{q,k}(r) \subseteq \mathcal{K}$. Let \mathcal{K}' be obtained by removing x from \mathcal{K} and then adding x' , i.e. $\mathcal{K}' = (\mathcal{K} \setminus \{x\}) \cup \{x'\}$. Because x' was minimal in Hamming weight and x was maximal in Hamming weight, the set \mathcal{K}' continues to be monotone, and it has size N . Also $g_{k,\tau}(\mathcal{K}) < g_{k,\tau}(\mathcal{K}')$ because $w(x) > w(x')$. Hence, by Lemma 5

$$\mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}) = g_{k,\tau}(\mathcal{K}) < g_{k,\tau}(\mathcal{K}') = \mathbf{Adv}_{q,k,\tau}^{\text{rskp}}(\mathcal{K}') .$$

This contradicts the assumption that $\mathbf{Adv}_{q,k,\tau}^{\text{rskp}}(N) = \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K})$. Hence, it must be that $\mathbf{B}_{q,k}(r) \subseteq \mathcal{K}$. Now suppose $\mathcal{K} \not\subseteq \mathbf{B}_{q,k}(r+1)$. Let x' be a point in $\mathbf{B}_{q,k}(r+1) \setminus \mathcal{K}$. Such a point exists because we know that $N < B_{q,k}(r+1)$. It must be that $w(x') = r+1$ since $\mathbf{B}_{q,k}(r) \subseteq \mathcal{K}$. Let x be a point in $\mathcal{K} \setminus \mathbf{B}_{q,k}(r+1)$ of maximal Hamming weight. Note that $w(x) > r+1 = w(x')$. Let \mathcal{K}' be obtained by removing x from \mathcal{K} and then adding x' , meaning $\mathcal{K}' = (\mathcal{K} \setminus \{x\}) \cup \{x'\}$. The set \mathcal{K}' continues to be monotone, and it has size N . Also $g_{k,\tau}(\mathcal{K}) < g_{k,\tau}(\mathcal{K}')$ because $w(x) > w(x')$. Hence, by Lemma 5,

$$\mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}) = g_{k,\tau}(\mathcal{K}) < g_{k,\tau}(\mathcal{K}') = \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K}') .$$

This contradicts the assumption that $\mathbf{Adv}_{q,k,\tau}^{\text{rskp}}(N) = \mathbf{Adv}_{k,\tau}^{\text{rskp}}(\mathcal{K})$. Hence, it must be that $\mathcal{K} \subseteq \mathbf{B}_{q,k}(r+1)$. \square

3.3.2 Proof of Lemma 9 and Lemma 10

To prove Lemma 9 and Lemma 10, we recall the notion of *discrete concavity*. Suppose $F: [1..M] \rightarrow \mathbb{R}$. We say that F is concave if $F(a+1) - F(a) \leq F(b+1) - F(b)$ for all $a, b \in [1..M]$ satisfying

$a \geq b$. Now suppose t, m are integers with $1 \leq m \leq t$. Then we let

$$S(M, m, t) = \left\{ (x_1, \dots, x_m) \in [1..M]^m : x_1 + \dots + x_m = t \right\}.$$

Define $F^m: [1..M]^m \rightarrow \mathbb{R}$ by $F^m(x_1, \dots, x_m) = F(x_1) + \dots + F(x_m)$. We use the following lemma proved by [4].

Lemma 12 ([4]). *Suppose $F: [1..M] \rightarrow \mathbb{R}$ is concave. Suppose $1 \leq m \leq t$ are integers such that m divides t and $t/m \in [1..M]$. Then*

$$\max_{(x_1, \dots, x_m) \in S(M, m, t)} F^m(x_1, \dots, x_m) = m \cdot F(t/m) \quad \square$$

Lemma 13. *The function, $F_{q,k,\tau}: [1..q^k] \rightarrow \mathbb{R}$, defined below, is concave.*

$$F_{q,k,\tau}(N) = \frac{N}{q^k} \cdot \mathbf{Adv}_{q,k,\tau}^{\text{rskp}}(N).$$

Proof. Let N_0, N_1 be two integers such that $q^k \geq N_0 \geq N_1 \geq 1$. Consider, for $i = 0, 1$,

$$\Delta_i = F_{q,k,\tau}(N_i + 1) - F_{q,k,\tau}(N_i),$$

For $i = 0, 1$, we let r_i be defined as follows. If $N_i = B_{q,k}(r)$ for some r , then we take r_i to be the value such that $B_{q,k}(r_i) = N_i$. Otherwise, we let $r_i = \text{rd}_{q,k}(N_i) + 1$. Note that we can now express Δ_i in terms of r_i as follow (via Equation (12)),

$$\Delta_i = q^k \cdot \frac{(k - r_i)_{(\tau)}}{k_{(\tau)}}.$$

Since $N_0 \geq N_1$, we note that $r_0 \geq r_1$. Therefore, we have $\Delta_0 \leq \Delta_1$ and that $F_{q,k,\tau}$ is concave. \square

We first prove Lemma 9 using Lemma 13.

Proof (of Lemma 9). Note that,

$$G_{q,k,\tau}(N) = \frac{q^k \cdot F_{q,k,\tau}(N)}{N}.$$

We let $\Delta_i = q^k \cdot F_{q,k,\tau}(i + 1) - q^k \cdot F_{q,k,\tau}(i)$ for all $i = 0, \dots, q^k - 1$. We define $\Delta_0 = q^k \cdot F_{q,k,\tau}(1) = q^k \cdot G_{q,k,\tau}(1)$. Hence, by construction $G_{q,k,\tau}(i) = (\sum_{j=0}^{i-1} \Delta_j) / i$. Note that since $F_{q,k,\tau}(\cdot)$ is concave in the domain $[1..q^k]$, the sequence $\Delta_1, \dots, \Delta_{q^k-1}$ is non-increasing, meaning that $\Delta_i \geq \Delta_j$ whenever $1 \leq i \leq j \leq q^k - 1$. Additionally, we check that $\Delta_0 = q^k$ and $\Delta_1 \leq q^k$, hence $\Delta_0 \geq \Delta_1$. Therefore, the partial averages of the sequence $\Delta_0, \dots, \Delta_{q^k-1}$,

$$\left(\sum_{j=0}^{i-1} \Delta_j \right) / i = G_{q,k,\tau}(i),$$

is non-increasing as claimed. \square

Lastly, we prove Lemma 10 using Lemma 12 and 13.

<p style="margin: 0;">Game $\mathbf{G}_{q,k,\tau,t}^{\text{mcskp}}(\mathcal{A}, \text{Lk})$</p> <p style="margin: 0;">$\mathbf{K} \leftarrow_s [q]^k; L \leftarrow \text{Lk}(\mathbf{K})$</p> <p style="margin: 0;">For $i \in [t]$ do $\mathbf{p}_i \leftarrow_s [k]^{(\tau)}$</p> <p style="margin: 0;">$J \leftarrow_s \mathcal{A}(L, \mathbf{p}_0, \dots, \mathbf{p}_{t-1})$</p> <p style="margin: 0;">Return $(\exists i \in [t] : J = \mathbf{K}[\mathbf{p}_i])$</p>

Figure 4: Multi-challenge subkey prediction game $\mathbf{G}_{q,k,\tau,t}^{\text{mcskp}}$.

Proof (of Lemma 10). Let $M = q^k$, $m = q^\ell$ and $t = q^k$. We note that the leakage function $\text{Lk} : [q]^k \rightarrow [q]^\ell$ defines a partition of $[q]^k$ into q^ℓ sets, with each set being $\text{Lk}^{-1}(L)$ for some $L \in [q]^\ell$. Hence, we can expand $\Pr[\mathbf{G}_{q,k,\tau}^{\text{skp}}(\mathcal{A}, \text{Lk})]$ by conditioning on the value of L . Suppose $[q]^\ell = \{L_1, \dots, L_m\}$. We let $N_i = |\text{Lk}^{-1}(L_i)|$. We derive

$$\begin{aligned}
& \text{Adv}_{q,k,\tau}^{\text{skp}}(\ell) \\
&= \max_{\text{Lk}} \left(\sum_L \frac{|\text{Lk}^{-1}(L)|}{q^k} \cdot \max_{\mathcal{A}} \Pr[\mathbf{G}_{q,k,\tau}^{\text{skp}}(\mathcal{A}, \text{Lk}) \mid \text{Lk}(\mathbf{K}) = L] \right) \\
&= \max_{\text{Lk}} \left(\sum_L \frac{|\text{Lk}^{-1}(L)|}{q^k} \cdot \text{Adv}_{k,\tau}^{\text{rskp}}(\text{Lk}^{-1}(L)) \right) \\
&\leq \max_{(N_1, \dots, N_m) \in S(M, m, t)} \sum_{i=1}^m F_{q,k,\tau}(N_i) \\
&= \max_{(N_1, \dots, N_m) \in S(M, m, t)} F_{q,k,\tau}^m(N_1, \dots, N_m) \\
&= m \cdot F_{q,k,\tau}(2^{k-\ell}) \tag{17}
\end{aligned}$$

$$= m \cdot \frac{q^{k-\ell}}{q^k} \cdot \text{Adv}_{q,k,\tau}^{\text{rskp}}(q^{k-\ell}) = \text{Adv}_{q,k,\tau}^{\text{rskp}}(q^{k-\ell}). \tag{18}$$

Equation (17) is justified since $F_{q,k,\tau}$ is concave and $t/m = 2^{k-\ell}$. Equation (18) is by definition of F and because $m = q^\ell$. \square

3.4 Multi-challenge Subkey Prediction

Here, we present an extension of $\mathbf{G}_{q,k,\tau}^{\text{skp}}$ with multiple challenges, $\mathbf{G}_{q,k,\tau,t}^{\text{mcskp}}$ (Fig. 4), the multi-challenge subkey prediction game. Note that [4] considers the multi-challenge version directly. However, we only need this extension in the proof of Theorem 15.

Let q, k, τ, t, ℓ be positive integers such that $q \geq 2$, $k \geq \tau$, $k \geq \ell$, $t \geq 1$. We define the following advantages associated with the game $\mathbf{G}_{q,k,\tau,t}^{\text{mcskp}}$, analogously to the advantages associated with $\mathbf{G}_{q,k,\tau}^{\text{skp}}$.

$$\begin{aligned}
\text{Adv}_{q,k,\tau,t}^{\text{mcskp}}(\mathcal{A}, \text{Lk}) &= \Pr \left[\mathbf{G}_{q,k,\tau,t}^{\text{mcskp}}(\mathcal{A}, \text{Lk}) \right], \\
\text{Adv}_{q,k,\tau,t}^{\text{mcskp}}(\text{Lk}) &= \max_{\mathcal{A}} \text{Adv}_{q,k,\tau,t}^{\text{skp}}(\mathcal{A}, \text{Lk}), \\
\text{Adv}_{q,k,\tau,t}^{\text{mcskp}}(\ell) &= \max_{\text{Lk}: [q]^k \rightarrow [q]^\ell} \text{Adv}_{q,k,\tau,t}^{\text{mcskp}}(\text{Lk}).
\end{aligned}$$

Game $\mathbf{G}_{\text{KEY}}^{\text{key}}(\mathcal{A})$	ROR(ℓ)
$b \leftarrow_{\$} \{0, 1\}; \mathbf{K} \leftarrow_{\$} [q]^k$	$R \leftarrow_{\$} \{0, 1\}^r$
$(\text{Lk}, \sigma) \leftarrow_{\$} \mathcal{A}^{\text{H}}()$	If $(b = 0)$ then $K \leftarrow_{\$} \{0, 1\}^{\kappa}$
$L \leftarrow \text{Lk}^{\text{H}}(\mathbf{K})$	Else $K \leftarrow \text{KEY}^{\text{H}}(\mathbf{K}, R)$
$b' \leftarrow_{\$} \mathcal{A}^{\text{ROR,H}}(L, \sigma)$	Return (R, K)
Return $(b' = b)$	$\text{H}(x, \text{Rng})$
	If not $T[x, \text{Rng}]$ then $T[x, \text{Rng}] \leftarrow_{\$} \text{Rng}$
	Return $T[x, n]$

Figure 5: Game for defining the security of a big-key key encapsulation algorithm KEY: $\{0, 1\}^k \times \{0, 1\}^r \rightarrow \{0, 1\}^{\kappa}$.

Lemma 14. *Let q, k, τ, t be positive integers such that $q \geq 2, k \geq \tau, k \geq \ell, t \geq 1$. Then,*

$$\mathbf{Adv}_{q,k,\tau,t}^{\text{mcskp}}(\ell) \leq t \cdot \mathbf{Adv}_{q,k,\tau}^{\text{skp}}(\ell) .$$

Proof. Let $\text{Lk} : [q]^k \rightarrow [q]^\ell$ be any leakage function. Let \mathcal{A} be a $\mathbf{G}_{q,k,\tau,t}^{\text{mcskp}}$ adversary. We construct $\mathbf{G}_{q,k,\tau}^{\text{skp}}$ adversary \mathcal{A}' such that

$$\mathbf{Adv}_{q,k,\tau}^{\text{skp}}(\mathcal{A}', \text{Lk}) \geq \frac{1}{t} \cdot \mathbf{Adv}_{q,k,\tau,t}^{\text{mcskp}}(\mathcal{A}, \text{Lk}) , \quad (19)$$

which implies the lemma by maximizing over all \mathcal{A} and Lk . \mathcal{A}' is defined as follows.

Adversary $\mathcal{A}'(L, \mathbf{p}')$
 $j \leftarrow_{\$} [t]; \mathbf{p}_j \leftarrow \mathbf{p}'$
For $i \in [t] - \{j\}$ do $\mathbf{p}_i \leftarrow [k]^{(\tau)}$
 $J' \leftarrow \mathcal{A}(L, \mathbf{p}_0, \dots, \mathbf{p}_{t-1})$
Return J'

Let E_1 be the event that \mathcal{A} succeed in the game $\mathbf{G}_{q,k,\tau,t}^{\text{mcskp}}(\mathcal{A}', \text{Lk})$, i.e. $J' = \mathbf{K}[\mathbf{p}_\alpha]$ for some $\alpha \in [t]$. Note that α is random variable that is well-defined given E_1 (in case $J' = \mathbf{K}[\mathbf{p}_\alpha]$ for multiple $\alpha \in [t]$, we can take the smallest one). We note that since \mathcal{A}' simulates the multi-challenge game for \mathcal{A} perfectly, $\Pr[E_1] = \mathbf{Adv}_{q,k,\tau,t}^{\text{mcskp}}(\mathcal{A}, \text{Lk})$. Let $E_2 \subseteq E_1$ be the event that \mathcal{A}' also guesses the correct α , i.e. $j = \alpha$ in the game $\mathbf{G}_{q,k,\tau,t}^{\text{mcskp}}(\mathcal{A}', \text{Lk})$. We note that $\Pr[E_2] = \frac{1}{t} \cdot \Pr[E_1]$, since j is independently uniform in $[t]$ and the distribution of $(\mathbf{p}_0, \dots, \mathbf{p}_{t-1})$ does not depend on the value of j . Notice that $\mathbf{Adv}_{q,k,\tau}^{\text{skp}}(\mathcal{A}', \text{Lk}) \geq \Pr[E_2] = \frac{1}{t} \cdot \Pr[E_1] = \frac{1}{t} \cdot \mathbf{Adv}_{q,k,\tau,t}^{\text{mcskp}}(\mathcal{A}, \text{Lk})$. \square

4 Big-Key Symmetric Encryption

In [4], Big-Key symmetric encryption schemes are constructed modularly from Big-Key encapsulation schemes. In this section, we present a block-based big key encapsulation scheme that is more efficient than achieved previously.

KEY ENCAPSULATION SCHEMES. A (symmetric, Big-Key) encapsulation schemes, on input a big key \mathbf{K} and a random string R , returns a (short) key K . The string R encapsulates the short key K in the sense that any party holding the big key \mathbf{K} can derive K from R . The security of a key

Algorithm $\mathsf{XKEY}_{q,k,\kappa,\tau,r}^{\mathsf{H}}(\mathbf{K}, R) // \mathbf{K} \in [q]^k, |R| = r$
 $\mathbf{p} \leftarrow \mathsf{H}(R, [k]^{(\tau)}); J \leftarrow \mathbf{K}[\mathbf{p}]; K \leftarrow \mathsf{H}(R||J, \kappa); \text{Return } K$

Figure 6: Encapsulation algorithm XKEY . Given a length- k big-key \mathbf{K} and a length- r selector R , the algorithm returns a length- κ subkey K . The value τ specifies the number of unique probes used.

encapsulation schemes is captured by $\mathbf{G}_{\mathsf{KEY}}^{\text{key}}(\mathcal{A})$ (Fig. 5). In this game, a big key \mathbf{K} is randomly sampled. The goal of the two-stage adversary \mathcal{A} is to guess whether the real-or-random oracle, ROR, is returning real keys, derived using key encapsulation scheme KEY from randomly sampled R , or randomly sampled keys that is independent of R . In its first stage, \mathcal{A} gets access to H and chooses a leakage function Lk and state σ . Next, the game computes $L \leftarrow \mathsf{Lk}^{\mathsf{H}}(\mathbf{K})$ and run the second stage of \mathcal{A} with inputs L, σ and oracles ROR and H . \mathcal{A} wins the game if it successfully guesses the bit b . We define the following advantage of \mathcal{A} against key encapsulation scheme KEY

$$\mathbf{Adv}_{\mathsf{KEY}}^{\text{key}}(\mathcal{A}) = 2 \cdot \Pr \left[\mathbf{G}_{\mathsf{KEY}}^{\text{key}}(\mathcal{A}) \right] - 1.$$

OUR CONSTRUCTION. Our random oracle model construction is given in Fig. 6.

Theorem 15. *Let $k, b, \kappa, \tau, r \geq 1$ be integers. Let $q = 2^b$. Let $\mathsf{KEY} = \mathsf{XKEY}_{q,k,\kappa,\tau,r}$ be the big-key encapsulation scheme associated to them as per Fig. 6. Let \mathcal{A} be an adversary making at most t queries to its ROR oracle and leaking $\ell \cdot b$ bits. Assume the number of H queries made by \mathcal{A} in its first stage, plus the number made by the oracle leakage function Lk that it outputs in this stage, is at most q_1 , and the number of H queries made by \mathcal{A} in its second stage is at most q_2 . Then*

$$\mathbf{Adv}_{\mathsf{KEY}}^{\text{key}}(\mathcal{A}) \leq q_2 \cdot t \cdot \mathbf{Adv}_{q,k,\tau}^{\text{skp}}(\ell) + \frac{t \cdot (2q_1 + t - 1)}{2^{r+1}}. \quad (20)$$

The proof of Theorem 15 is deferred to Section 4.1.

SAMPLING UNIQUE PROBES. In XKEY , we have outsourced the sampling of the unique probes to the variable-range random oracle. We note that sampling from $[k]^{(\tau)}$ can be done via rejection sampling efficiently. For example, per Lemma 20 in Appendix B, it holds with all but $2^{-3\tau}$ probability that 4τ samples from $[k]$ contains τ unique probes (hence for parameters involved in Fig. 1, the failure probability is less than 2^{-129} since $\tau \geq 43$).

SYMMETRIC ENCRYPTION SCHEMES. To obtain a (big-key) symmetric encryption scheme, one can plug our XKEY construction directly into the (big-key) symmetric encryption scheme (in Fig. 7) by BKR. For security, we omit the details here and appeal to [4, Theorem 13].

EFFICIENCY. Let $k^* = 8 \cdot 10^{11} = 100$ GBytes, and $\ell^* = 10$ GBytes. Using $b = 8 \cdot 512 = 512$ Bytes, our XKEY makes roughly the same number of H queries compared to [4] but makes significantly less access into the big key \mathbf{K} (43 vs. 271, Fig. 1). In practical instantiations where \mathbf{K} is stored on slow storage medium (e.g. hard disk), this translate to 6x improvement in efficiency.

4.1 Proof of Theorem 15

Proof (of Theorem 15). Consider the games, $\mathbf{G}_0, \dots, \mathbf{G}_3$ defined in Fig. 8. Let $\mathsf{KEY} = \mathsf{XKEY}_{q,k,\kappa,\tau,r}$. We note that game \mathbf{G}_0 , with the boxed code included, simulates the game $\mathbf{G}_{\mathsf{KEY}}^{\text{key}}(\mathcal{A})$ exactly for

Algorithm SE.Enc^H(K, <i>M</i>) $R \leftarrow_{\$} \{0, 1\}^r$; $K \leftarrow \text{KEY}^H(\mathbf{K}, R)$ $C \leftarrow \text{SE.Enc}(K, M)$; $\overline{C} \leftarrow (R, C)$ Return \overline{C}	Algorithm SE.Dec^H(K, <i>M</i>) $(R, C) \leftarrow \overline{C}$ $K \leftarrow \text{KEY}^H(\mathbf{K}, R)$ $M \leftarrow \text{SE.Dec}(K, C)$ Return M
--	---

Figure 7: Big-Key Symmetric Encryption Scheme [4, Section 5], **SE**, using a standard symmetric key encryption scheme **SE** and a key encapsulation mechanism **KEY**.

Game G₀ G₁ $\mathbf{K} \leftarrow_{\$} [q]^k$ For $j \leftarrow 1, \dots, t$ do $R[j] \leftarrow_{\$} \{0, 1\}^r$; $K[j] \leftarrow_{\$} \{0, 1\}^\kappa$ $P[j] \leftarrow_{\$} [k]^{(\tau)}$ For $i \leftarrow 1, \dots, t$ do If $R[i] = R[j]$ then $\text{bad} \leftarrow \text{true}$; $\boxed{K[j] \leftarrow K[i]}$ stage $\leftarrow 1$; $(\text{Lk}, \sigma) \leftarrow_{\$} \mathcal{A}^{\text{H}_0}()$; $L \leftarrow_{\$} \text{Lk}^{\text{H}_0}(\mathbf{K})$ stage $\leftarrow 2$; $b' \leftarrow_{\$} \mathcal{A}^{\text{ROR}, \text{H}_0}(\sigma, L)$ Return $(b' = 1)$ <u>H₀(<i>x</i>, Rng)</u> If not $T[x, \text{Rng}]$ then $T[x, \text{Rng}] \leftarrow_{\$} \text{Rng}$ If stage = 1 then For $j \in [t]$ do If $x = R[j]$ and $\text{Rng} = [k]^{(\tau)}$ then $\text{bad} \leftarrow \text{true}$; $\boxed{T[x, \text{Rng}] \leftarrow P[j]}$ If $x = R[j] \parallel J[j]$ and $\text{Rng} = \{0, 1\}^\kappa$ then $\text{bad} \leftarrow \text{true}$; $\boxed{T[x, \text{Rng}] \leftarrow K[j]}$ If stage = 2 then For $j \in [t]$ do If $x = R[j]$ and $\text{Rng} = [k]^{(\tau)}$ then $T[x, \text{Rng}] \leftarrow P[j]$ If $x = R[j] \parallel J[j]$ and $\text{Rng} = \{0, 1\}^\kappa$ then $T[x, \text{Rng}] \leftarrow K[j]$ Return $T[x, \text{Rng}]$	Game G₂ G₃ $\mathbf{K} \leftarrow_{\$} [q]^k$ For $j \leftarrow 1, \dots, t$ do $R[j] \leftarrow_{\$} \{0, 1\}^r$; $K[j] \leftarrow_{\$} \{0, 1\}^\kappa$ $P[j] \leftarrow_{\$} [k]^{(\tau)}$ stage $\leftarrow 1$; $(\text{Lk}, \sigma) \leftarrow_{\$} \mathcal{A}^{\text{H}_1}()$; $L \leftarrow_{\$} \text{Lk}^{\text{H}_1}(\mathbf{K})$ stage $\leftarrow 2$; $b' \leftarrow_{\$} \mathcal{A}^{\text{ROR}, \text{H}_1}(\sigma, L)$ Return $(b' = 1)$ <u>H₁(<i>x</i>, Rng)</u> If not $T[x, \text{Rng}]$ then $T[x, \text{Rng}] \leftarrow_{\$} \text{Rng}$ If stage = 2 then For $j \in [t]$ do If $x = R[j]$ and $\text{Rng} = [k]^{(\tau)}$ then $T[x, \text{Rng}] \leftarrow P[j]$ If $x = R[j] \parallel J[j]$ and $\text{Rng} = \{0, 1\}^\kappa$ then $\text{bad} \leftarrow \text{true}$; $\boxed{T[x, \text{Rng}] \leftarrow K[j]}$ Return $T[x, \text{Rng}]$ <hr/> <u>ROR()</u> $j \leftarrow j + 1$; Return $(R[j], K[j])$
---	---

Figure 8: Games **G₀**, ..., **G₃**. All games share the same procedure ROR shown on the bottom of the middle column.

$b = 1$ case and outputs true when \mathcal{A} outputs 1. Similarly, we note that **G₃**, without the boxed code, simulates the game $\mathbf{G}_{\text{KEY}}^{\text{key}}(\mathcal{A})$ exactly for $b = 0$ case and outputs true when \mathcal{A} outputs 1. Hence,

$$\text{Adv}_{\text{KEY}}^{\text{key}}(\mathcal{A}) = \Pr[\mathbf{G}_0] - \Pr[\mathbf{G}_3]. \quad (21)$$

<p><u>Game \mathbf{G}_4</u> $\mathbf{K} \leftarrow_{\\$} [q]^k$ For $j \leftarrow 1, \dots, t$ do $R[j] \leftarrow_{\\$} \{0, 1\}^r$; $K[j] \leftarrow_{\\$} \{0, 1\}^\kappa$ $P[j] \leftarrow_{\\$} [k]^{(\tau)}$ stage $\leftarrow 1$; $(\text{Lk}, \sigma) \leftarrow_{\\$} \mathcal{A}^{\text{H}_2}()$; $L \leftarrow_{\\$} \text{Lk}^{\text{H}_2}(\mathbf{K})$ stage $\leftarrow 2$; $b' \leftarrow_{\\$} \mathcal{A}^{\text{ROR}, \text{H}_2}(\sigma, L)$ Return ($b' = 1$)</p> <hr/> <p><u>ROR()</u> $j \leftarrow j + 1$; Return ($R[j], K[j]$)</p>	<p><u>$\text{H}_2(x, \text{Rng})$</u> If not $T_0[x, \text{Rng}]$ then $T_0[x, \text{Rng}] \leftarrow_{\\$} \text{Rng}$ If stage = 2 then For $j \in [t]$ do If $x = R[j]$ and $\text{Rng} = [k]^{(\tau)}$ then $T_0[x, \text{Rng}] \leftarrow P[j]$ If $x = R[j] \parallel J[j]$ and $\text{Rng} = \{0, 1\}^\kappa$ then bad \leftarrow true Return $T_0[x, \text{Rng}]$</p>
---	--

Figure 9: Game \mathbf{G}_4 . Note that T_0 is a table obtained via coin-fixing.

<p><u>Adversary $\mathcal{B}(L, \mathbf{p}_0, \dots, \mathbf{p}_{t-1})$</u> $i \leftarrow 0$; For $j \in [t]$ do $R[j] \leftarrow_{\\$} \{0, 1\}^r$ $P[j] \leftarrow_{\\$} \mathbf{p}_j$ $b' \leftarrow_{\\$} \mathcal{A}^{\text{ROR}, \text{H}_3}(L)$ $\alpha \leftarrow_{\\$} [i]$ Return J_α</p> <hr/> <p><u>ROR()</u> $j \leftarrow j + 1$; Return ($R[j], K[j]$)</p>	<p><u>$\text{H}_3(x, \text{Rng})$</u> If not $T_0[x, \text{Rng}]$ then $T_0[x, \text{Rng}] \leftarrow_{\\$} \text{Rng}$ If stage = 2 then For $j \in [t]$ do If $x = R[j]$ and $\text{Rng} = [k]^{(\tau)}$ then $T_0[x, \text{Rng}] \leftarrow P[j]$ If $\text{Rng} = \{0, 1\}^\kappa$ then $R_i \parallel J_i \leftarrow x$; $i \leftarrow i + 1$ Return $T_0[x, \text{Rng}]$</p>
---	---

Figure 10: Subkey prediction adversary \mathcal{B} .

We will proceed to bound $\Pr[\mathbf{G}_0]$. Note that \mathbf{G}_0 and \mathbf{G}_1 are identical-until-bad. Hence, via the Fundamental Lemma of Game Playing [7]

$$\begin{aligned} \Pr[\mathbf{G}_0] &= \Pr[\mathbf{G}_1] + (\Pr[\mathbf{G}_0] - \Pr[\mathbf{G}_1]) \\ &\leq \Pr[\mathbf{G}_1] + \Pr[\mathbf{G}_1 \text{ sets bad}] . \end{aligned} \tag{22}$$

Next, we claim that

$$\Pr[\mathbf{G}_1 \text{ sets bad}] \leq \frac{t(t-1)}{2^{r+1}} + \frac{t \cdot q_1}{2^r} . \tag{23}$$

First, there is at most $t(t-1)/2^{r+1}$ probability of collision in the r -bit values $R[1], \dots, R[t]$ by the birthday bound. Next, H_0 sets bad only when stage = 1, and there are exactly q_1 H_0 -queries when stage = 1. We note that each H_0 -query when stage = 1 has at most $t/2^r$ probability of setting bad since there are at most t distinct values for $R[1], \dots, R[t]$.

We proceed to bound $\Pr[\mathbf{G}_1]$. We note that \mathbf{G}_2 , with the boxed code included, is equivalent to \mathbf{G}_1 . Furthermore, \mathbf{G}_3 , without the boxed code, is identical to \mathbf{G}_2 until bad is set. Hence,

$$\begin{aligned} \Pr[\mathbf{G}_1] &= \Pr[\mathbf{G}_2] = \Pr[\mathbf{G}_3] - (\Pr[\mathbf{G}_2] - \Pr[\mathbf{G}_3]) \\ &\leq \Pr[\mathbf{G}_3] + \Pr[\mathbf{G}_3 \text{ sets bad}] . \end{aligned} \tag{24}$$

Lastly, we claim that

$$\Pr[\mathbf{G}_3 \text{ sets bad}] \leq q_2 \cdot t \cdot \mathbf{Adv}_{q,k,\tau}^{\text{skp}}(\ell) , \quad (25)$$

Notice that the theorem follows from Equations (21), (22), (23), (24), and (25). It remains to show Equation (25). The justification of Equation (25) involves two step. First, we argue that there is some fixing of the coins of $\mathcal{A}, \mathbf{H}_1, \text{Lk}$, which results in a deterministic leakage function Lk' , an adversary \mathcal{A}' , and partial \mathbf{H} table T_0 such that

$$\Pr[\mathbf{G}_3 \text{ sets bad}] \leq \Pr[\mathbf{G}_4 \text{ sets bad}] , \quad (26)$$

where \mathbf{G}_4 is given in Fig. 9. Next, we show that

$$\Pr[\mathbf{G}_4 \text{ sets bad}] \leq q_2 \cdot \mathbf{Adv}_{q,k,\tau,t}^{\text{mcskp}}(\mathcal{B}, \text{Lk}') \leq q_2 \cdot t \cdot \mathbf{Adv}_{q,k,\tau}^{\text{skp}}(\ell) , \quad (27)$$

by constructing a multi-challenge subkey prediction adversary \mathcal{B} , which is given in Fig. 10. \mathcal{B} will embed the probes given, $\mathbf{p}_0, \dots, \mathbf{p}_{t-1}$ into the \mathbf{H} response and run \mathcal{A}' . It will guess, at random, one of the \mathbf{H} queries of \mathcal{A}' of the form $(R||J, \{0, 1\}^\kappa)$. Hence, if \mathbf{G}_4 sets bad, then with at least $\frac{1}{q_2}$ probability, \mathcal{B} succeeds. The second part of Equation (27) follows from Lemma 14. This justifies Equation (25) and concludes the proof of the theorem. \square

5 Big-Key Identification

IDENTIFICATION SCHEMES. An identification scheme ID specifies the following:

- Via $\text{prm} \leftarrow_{\$} \text{ID.Pg}$, parameter generation algorithm ID.Pg generates parameter prm , which is a common input to all other algorithms.
- Via $(\text{sk}, \text{vk}, \text{hlp}) \leftarrow_{\$} \text{ID.Kg}(\text{prm})$, key generation algorithm ID.Kg is run by the prover to generate secret key sk , corresponding verification key vk and a string hlp called the help string. The last is information that, conceptually, can be viewed as part of the public verification key vk , meaning public and available to the adversary, but to keep the verification key small, hlp is stored by the prover along with sk .
- Via $(\text{com}, \text{st}) \leftarrow_{\$} \text{ID.Com}(\text{prm})$, commitment algorithm ID.Com is run by the prover to generate its first message com , called the commitment, along with state information st that it saves.
- Via $\text{chl} \leftarrow_{\$} \{0, 1\}^{\text{ID.Chl}}$, the verifier generates a random challenge chl to return to the prover.
- Via $\text{rsp} \leftarrow \text{ID.Rsp}(\text{prm}, \text{hlp}, \text{sk}, \text{st}, \text{chl})$, deterministic response algorithm ID.Rsp is run by the prover to generate its response rsp .
- Via $d \leftarrow \text{ID.Vrf}(\text{prm}, \text{vk}, \text{com}, \text{chl}, \text{rsp})$, deterministic decision algorithm ID.Vrf returns a boolean decision d for the verifier to accept or reject.

In the ROM, algorithms may have oracle access to the random oracle \mathbf{H} . This syntax is non-asymptotic, in that there is no explicit security parameter. Correctness requires that

$$\Pr[\text{Execute}_{\text{ID}}(\text{prm}, \text{vk}, \text{sk}, \text{hlp})] = 1$$

for all $\text{prm} \in [\text{ID.Pg}]$ and $(\text{sk}, \text{vk}, \text{hlp}) \in [\text{ID.Kg}(\text{prm})]$, where

$$\begin{array}{l} \text{Game } \text{Execute}_{\text{ID}}(\text{prm}, \text{vk}, \text{sk}, \text{hlp}) \\ (\text{com}, \text{st}) \leftarrow_{\$} \text{ID.Com}(\text{prm}) \\ \text{chl} \leftarrow_{\$} \{0, 1\}^{\text{ID.Chl}} \end{array}$$

<p><u>Game $\mathbf{G}_{\text{ID}, \ell}^{\text{imp}}(\mathcal{A})$</u></p> <p>$\text{prm} \leftarrow_{\\$} \text{ID.Pg}; s \leftarrow 0$ $(\text{sk}, \text{vk}, \text{hlp}) \leftarrow_{\\$} \text{ID.Kg}(\text{prm})$ $\text{st} \leftarrow_{\\$} \mathcal{A}.\text{Setup}^{\text{Leak}_{\ell}, \text{Prover.H}}(\text{prm}, \text{vk}, \text{hlp})$ $(\text{com}, \text{st}') \leftarrow_{\\$} \mathcal{A}.\text{Com}^{\text{H}}(\text{st});$ $\text{chl} \leftarrow_{\\$} \{0, 1\}^{\text{ID.Chl}}$ $\text{rsp} \leftarrow_{\\$} \mathcal{A}.\text{Rsp}^{\text{H}}(\text{prm}, \text{hlp}, \text{sk}, \text{st}', \text{chl})$ $d \leftarrow_{\\$} \text{ID.Vrf}^{\text{H}}(\text{prm}, \text{vk}, \text{com}, \text{chl}, \text{rsp})$ Return d</p> <hr/> <p><u>Leak$_{\ell}(f)$</u></p> <p>$L \leftarrow_{\\$} f(\text{sk}); s \leftarrow s + L$ If $s \leq \ell$ then return L else return \perp</p>	<p><u>Prover(i, args)</u></p> <p>If $\text{pst}[i] = \perp$ then // Commit $(\text{pcom}[i], \text{pst}[i]) \leftarrow_{\\$} \text{ID.Com}(\text{prm});$ return $\text{pcom}[i]$ If $\text{prsp}[i] = \perp$ then // Response $\text{prsp}[i] \leftarrow_{\\$} \text{ID.Rsp}^{\text{H}}(\text{prm}, \text{hlp}, \text{sk}, \text{pst}[i], \text{args})$ Return $\text{prsp}[i]$ Return \perp</p> <hr/> <p><u>H(x, Rng)</u></p> <p>If $T[x, \text{Rng}] = \perp$ then $T[x, \text{Rng}] \leftarrow_{\\$} \text{Rng}$ Return $T[x, \text{Rng}]$</p>
--	--

Figure 11: Game defining security of identification scheme ID under pre-impersonation leakage.

```

rsp ← ID.Rsp(prm, hlp, sk, st, chl)
d ← ID.Vrf(prm, vk, com, chl, rsp)
Return d

```

SECURITY OF IDENTIFICATION SCHEMES. We give definitions allowing concrete-security assessments. The core definition is that of adversary advantage. The notion captured is security against impersonation under active attack [14, 5] in the further presence of leakage on the secret key [2].

Let ID be an identification scheme. Let ℓ be an integer representing a bound (in bits) on the leakage. Let \mathcal{A} be an *impersonation adversary*, made up of component algorithms $\mathcal{A}.\text{Setup}$, $\mathcal{A}.\text{Com}$, and $\mathcal{A}.\text{Rsp}$. We associate to these the game of Fig. 11. First, the parameters and keys are generated. Next, $\mathcal{A}.\text{Setup}$ is run with access to a leakage oracle Leak_{ℓ} a prover oracle Prover and the random oracle H . The leakage oracle takes input a function Lk from the adversary and returns leakage $L = \text{Lk}(\text{sk})$. This oracle can be called adaptively and any number of times, its code ensuring that the total number of bits returned to the adversary does not exceed ℓ . The prover oracle allows an active attack in which the adversary, playing the role of a dishonest verifier, can generate prover instances and interact with them. The commitment and state of instance i are produced by the game and stored as $\text{pcom}[i]$ and $\text{pst}[i]$, respectively. If instance i has been activated, meaning $\text{pst}[i] \neq \perp$, then the adversary can submit, via args , a challenge of its choice, and obtain response $\text{prsp}[i]$. After exiting this setup phase, the adversary turns into a dishonest prover, aiming to convince the honest verifier to accept. It produces its commitment via $\mathcal{A}.\text{Com}$, receives a random challenge chl , and produces its response via $\mathcal{A}.\text{Rsp}$. The game returns the boolean decision d of the verifier's decision function. We define the leakage impersonation advantage of \mathcal{A} against ID to be

$$\text{Adv}_{\text{ID}, \ell}^{\text{imp}}(\mathcal{A}) = \Pr \left[\mathbf{G}_{\text{ID}, \ell}^{\text{imp}}(\mathcal{A}) \right].$$

GROUPS. We fix a *bilinear group description* $\mathcal{G} = (G, G_T, g, \mathbf{e}, p)$, where

- $p \geq 3$ is a prime number that will be the order of the groups
- G, G_T are (cyclic) groups of order p

<p>Game $\mathbf{G}_{\mathcal{G}}^{\text{cdh}}(\mathcal{A})$</p> <p>$(G, G_T, g, \mathbf{e}, p) \leftarrow \mathcal{G}$</p> <p>$x, y \leftarrow_{\\$} [p]$; $h \leftarrow_{\\$} \mathcal{A}(G, g^x, g^y)$</p> <p>Return $(h = g^{xy})$</p> <hr/> <p>Game $\mathbf{G}_{\mathcal{G}}^{\text{dl}}(\mathcal{A})$</p> <p>$(G, G_T, g, \mathbf{e}, p) \leftarrow \mathcal{G}$</p> <p>$x \leftarrow_{\\$} [p]$; $x' \leftarrow_{\\$} \mathcal{A}(G, g^x)$</p> <p>Return $(x = x')$</p>	<p>Game $\mathbf{G}_{p,m,k,\tau}^{\text{pskp}}(\mathcal{A}, \text{Lk})$</p> <p>For $i \in [k]$ do $\text{sk}[i] \leftarrow_{\\$} \mathbb{Z}_p^m$</p> <p>$\mathbf{p} \leftarrow_{\\$} [k]^{(\tau)}$; $e \leftarrow_{\\$} \mathbb{Z}_p$</p> <p>For $j \in [m]$ do</p> <p style="padding-left: 20px;">$\text{sk}^*[j] = \sum_{i=0}^{\tau-1} (\text{sk}[\mathbf{p}[i]][j])e^i$</p> <p>$L \leftarrow \text{Lk}(\text{sk})$; $\overline{\text{sk}} \leftarrow_{\\$} \mathcal{A}(\mathbf{p}, e, L)$</p> <p>Return $(\text{sk}^* = \overline{\text{sk}})$.</p>
---	--

Figure 12: **Left:** Games $\mathbf{G}_{\mathcal{G}}^{\text{cdh}}$ and $\mathbf{G}_{\mathcal{G}}^{\text{dl}}$ defining the security of CDH and DL problems in \mathcal{G} . **Right:** Game $\mathbf{G}_{p,m,k,\tau}^{\text{pskp}}(\mathcal{A}, \text{Lk})$. Where $\text{Lk} : [q]^k \rightarrow [q]^\ell$ is a leakage function. $[k]^{(\tau)}$ contains the set of τ -dimensional vectors over $[k]$ with distinct entries.

- $g \in G$ is a generator of G
- $\mathbf{e} : G \times G \rightarrow G_T$ is an efficiently computable, non-degenerate bilinear map. This means that (1) $e(g^a, g^b) = \mathbf{e}(g, g)^{ab}$ for all $a, b \in [p]$, and (2) $\mathbf{e}(g, g)$ is not the identity element of G_T .

We will base security on the assumed hardness of the CDH (Computational Diffie-Hellman) and DL (Discrete Logarithm) problems in G . The definitions are based on games \mathbf{G}^{cdh} and \mathbf{G}^{dl} in Fig. 12, associated to \mathcal{G} and an adversary \mathcal{A} . We define the following CDH and DL advantages:

$$\begin{aligned} \text{Adv}_{\mathcal{G}}^{\text{cdh}}(\mathcal{A}) &= \Pr[\mathbf{G}_{\mathcal{G}}^{\text{cdh}}(\mathcal{A})] \\ \text{Adv}_{\mathcal{G}}^{\text{dl}}(\mathcal{A}) &= \Pr[\mathbf{G}_{\mathcal{G}}^{\text{dl}}(\mathcal{A})]. \end{aligned}$$

Hardness of CDH of course implies hardness of DL. Quantitatively, given \mathcal{A} , one can construct \mathcal{A}' with similar running time such that

$$\text{Adv}_{\mathcal{G}}^{\text{dl}}(\mathcal{A}) \leq \text{Adv}_{\mathcal{G}}^{\text{cdh}}(\mathcal{A}').$$

ADW IDENTIFICATION SCHEME. We present a variant of ADW’s identification scheme [2], which uses a random oracle to derive the challenges (as considered in [2] without analysis). The scheme $\text{ID} = \text{ADW}[\mathcal{G}, k, m, \tau, r]$ is parameterized by a bilinear group description \mathcal{G} and positive integers k, m, τ, r . We require that $m \geq 2$ and $k \geq \tau \geq 1$. Here k is the number of blocks of the secret key, where each block is an m -dimensional vector over \mathbb{Z}_p , and τ is the number of probes that algorithms make into the secret key. The parameter r determines the challenge length, meaning we set $\text{ID.chl} = r$. The algorithms ID.Pg , ID.Kg , ID.Com , ID.Rsp , ID.Vrf are given in Fig. 13.

Intuitively, the scheme consists of k generalized Okamoto identification scheme [16, 2], and one instance of BLS signature scheme [8]. Each block of the secret key (in \mathbb{Z}_p^m) is a secret key for a generalized Okamoto identification scheme of dimension m . The public keys, $\text{pk}[0], \dots, \text{pk}[k-1]$, of the k Okamoto’s identification schemes, are signed using the BLS signature scheme under signing key s , yielding signatures $\sigma[0], \dots, \sigma[k-1]$. The public verification key of the identification scheme, consists only of the verification key, vk , of the BLS signature scheme. During identification, a random τ instances out of k instances is chosen (via H by the verifier) and compressed via polynomial evaluation to sk^*, pk^* , and σ^* by the prover. During response phase, the prover, in addition to answering the challenge from the Okamoto identification scheme, needs to transmit pk^* and σ^* to the

<u>ID.Pg()</u> For $i \in [m]$ do $g_i \leftarrow_s G$ Return (g_0, \dots, g_{m-1})	<u>ID.Com(prm)</u> $y \leftarrow_s (\mathbb{Z}_p)^m$ $a \leftarrow_s \prod_{j=0}^{m-1} g_i^{y[j]}$ Return (a, y)	<u>Derive^H(R)</u> $\mathbf{p} \leftarrow H(R, [k]^{(\tau)})$ $e \leftarrow_s H(0 R, [p]); c^* \leftarrow_s H(1 R, [p])$ Return (\mathbf{p}, e, c^*)
<u>ID.Kg^H(prm)</u> $s \leftarrow_s \mathbb{Z}_p; \mathbf{vk} \leftarrow g^s$ For $i \in [k]$ do $\mathbf{sk}[i] \leftarrow_s (\mathbb{Z}_p)^m$ $\mathbf{pk}[i] \leftarrow \prod_{j=0}^{m-1} (g_i)^{\mathbf{sk}[i][j]}$ $\sigma[i] \leftarrow_s (H(i, G)\mathbf{pk}[i])^s$ $\mathbf{hlp} \leftarrow (\mathbf{pk}, \sigma)$ Return $(\mathbf{sk}, \mathbf{vk}, \mathbf{hlp})$	<u>ID.Rsp^H(prm, hlp, sk, st, chl)</u> $(\mathbf{p}, e, c^*) \leftarrow \text{Derive}^H(\text{chl})$ For $j \in [m]$ do $\mathbf{sk}^*[j] \leftarrow \sum_{i=0}^{\tau-1} (\mathbf{sk}[\mathbf{p}[i]][j])e^i$ $\mathbf{pk}^* \leftarrow \prod_{i=0}^{\tau-1} \mathbf{pk}[\mathbf{p}[i]]^{e^i}$ $\sigma^* \leftarrow \prod_{i=0}^{\tau-1} \sigma[\mathbf{p}[i]]^{e^i}$ For $j \in [m]$ do $z \leftarrow y[j] + c^* \cdot \mathbf{sk}^*[j]$ Return $(\mathbf{pk}^*, \sigma^*, z)$	<u>ID.Vrf^H(prm, vk, com, chl, rsp)</u> $a \leftarrow \text{com}$ $(\mathbf{p}, e, c^*) \leftarrow \text{Derive}^H(\text{chl})$ $(\mathbf{pk}^*, \sigma^*, z) \leftarrow \text{rsp}$ $A \leftarrow (\prod_{i=0}^{m-1} g_i^{z[i]} = a(\mathbf{pk}^*)^{c^*})$ $h_1 \leftarrow \mathbf{e}(\mathbf{pk}^* \prod_{i=0}^{\tau-1} H(\mathbf{p}[i], G)^{e^i}, \mathbf{vk})$ $h_2 \leftarrow \mathbf{e}(\sigma^*, g)$ $B \leftarrow (h_1 = h_2)$ Return $(A \wedge B)$

Figure 13: Algorithms of identification scheme $\text{ID} = \text{ADW}[\mathcal{G}, k, m, \tau, r]$ associated to bilinear group description $\mathcal{G} = (G, G_T, g, \mathbf{e}, p)$ and parameters k, m, τ, r satisfying $m \geq 2$ and $k \geq \tau \geq 1$. Here H is a variable range function, meaning $H(\cdot, \text{Rng})$ returns outputs in the set (described by) Rng . In addition, algorithms $\text{Kg}, \text{Com}, \text{Rsp}, \text{Vrf}$ also takes prm as argument.

verifier. We note that the signing key, s , of the underlying signature scheme *must not* be visible to the attacker. This signing key is simply be discarded after Kg . (However, we note that, as ADW has pointed out, there are advanced uses of this key such as updating the big secret key.) The correctness of $\text{ID} = \text{ADW}[\mathcal{G}, k, m, \tau, r]$ is checked as follows. Let $\text{prm} \in [\text{ID.Pg}]$ and $(\mathbf{sk}, \mathbf{vk}, \mathbf{hlp}) \in [\text{ID.Kg}(\text{prm})]$. We claim that, during a honest execution of the protocol ($\text{Execute}_{\text{ID}}(\text{prm}, \mathbf{sk}, \mathbf{vk}, \mathbf{hlp})$), the flags A, B in ID.Vrf will both be set to true. A is set to true because

$$\begin{aligned}
\prod_{i=0}^{m-1} g_i^{z[i]} &= \prod_{i=0}^{m-1} g^{y[i] + c^* \cdot \mathbf{sk}^*[i]} \\
&= \prod_{i=0}^{m-1} g^{y[i]} \cdot \left(\prod_{i=0}^{m-1} g^{\mathbf{sk}^*[i]} \right)^{c^*} = a \cdot \mathbf{pk}^{*c^*} .
\end{aligned}$$

B is set to true because

$$\begin{aligned}
\mathbf{e}(\mathbf{pk}^* \prod_{i=0}^{\tau-1} H(\mathbf{p}[i], G)^{e^i}, \mathbf{vk}) &= \mathbf{e}\left(\prod_{i=0}^{\tau-1} \mathbf{pk}[\mathbf{p}[i]]^{e^i} \prod_{i=0}^{\tau-1} H(\mathbf{p}[i], G)^{e^i}, g^s\right) \\
&= \mathbf{e}\left(\prod_{i=0}^{\tau-1} ((\mathbf{pk}[\mathbf{p}[i]]H(\mathbf{p}[i], G))^s)^{e^i}, g\right) \\
&= \mathbf{e}\left(\prod_{i=0}^{\tau-1} (\sigma[\mathbf{p}[i]])^{e^i}, g\right) = \mathbf{e}(\sigma^*, g) .
\end{aligned}$$

Hence, $\Pr[\text{Execute}_{\text{ID}}(\text{prm}, \mathbf{sk}, \mathbf{vk}, \mathbf{hlp})] = 1$, and ID satisfies correctness.

Computation cost					
	Kg	Com	Chl	Rsp	Vrf
Mult G	$k \cdot m$	$m - 1$	0	2τ	$m + \tau$
Exp G	$k(m + 1) + 1$	m	0	$2\tau - 2$	$\tau + m + 1$
Mult \mathbb{Z}_p	0	0	0	m	1
Exp \mathbb{Z}_p	m	0	0	2τ	τ
\mathbf{e} eval	0	0	0	0	1
Communication cost					
G	-	1	0	2	0
\mathbb{Z}_p	-	0	0	m	0
$\{0, 1\}^r$	-	0	1	0	0

Example parameters		
m	τ (Us)	τ (ADW)
2	718	3951
4	349	2397
8	245	1996
16	201	1840
32	180	1771
64	169	1739

Figure 14: **Left:** Table illustrating computation and communication cost of different operations of the identification scheme $\text{ADW}_{\mathcal{G},k,m,\tau,r}$. **Chl** here represents the challenge phase of the protocol. **Right:** Example parameters for ADW scheme to achieve 128-bit security. The schemes uses group of size p such that $2^{511} < p < 2^{512}$, and we impose a bound on the leakage of 10% on a big-key of size $100 \text{ GB} = 8 \times 10^{11}$ bits. For each value of m on the left column, we look the value of τ needed to achieve 128-bit security for the identification scheme, both using our bound and using ADW’s bound.

EFFICIENCY. As pointed out in [2], the identification scheme has nice efficiency properties. First, the public key (verification key) is very short (one group element). Second, the communication costs of all phases are very small. The bulk of communication happens in the response phase, which outputs 2 group elements and m elements from \mathbb{Z}_p . Third, the scheme has probe complexity depending τ , which can be made small while preserving security. In particular, during each run of the protocol, only τ locations of the secret-key will be accessed (each location consist of m elements of \mathbb{Z}_p). Fig. 14 demonstrates the computation and communication costs of different operations. Note that very small values of τ makes the scheme insecure. The crux of the security analysis amounts to giving a lower bound of τ for a desired security level. Here is where we make significant *concrete security* improvements over ADW.

CONCRETE-SECURITY ANALYSIS. Before we present the theorem stating the concrete security of the ADW identification scheme, we first need to define the following special subkey prediction game. The game $\mathbf{G}_{p,m,k,\tau}^{\text{pskp}}(\text{Lk}, \mathcal{A})$ (Fig. 12) captures a particular type of subkey prediction game in which the subkey is interpreted as a tuple of polynomials. In this game, the adversary \mathcal{A} needs to predict the value of these polynomials at a random point e , which is given to \mathcal{A} . We define the following prediction advantage

$$\mathbf{Adv}_{p,m,k,\tau}^{\text{pskp}}(\ell) = \max_{\mathcal{A}, \text{Lk}: (\mathbb{Z}_p^m)^k \rightarrow (\mathbb{Z}_p^m)^\ell} \Pr \left[\mathbf{G}_{p,m,k,\tau}^{\text{pskp}}(\mathcal{A}, \text{Lk}) \right].$$

We state a theorem which captures the *concrete security* of the ADW identification scheme. The theorem streamlines the original analysis of ADW to a precise relation of advantages, which allows us to instantiate parameters of practical sizes.

Theorem 16. *Let $\mathcal{G} = (G, G_T, g, \mathbf{e}, p)$ be a group with efficient pairngs. Let $\text{ID} = \text{ADW}_{\mathcal{G},k,m,\tau,r}$ be the ADW identification scheme shown in Fig. 13. Let $\mathcal{A} = (\mathcal{A}.\text{Setup}, \mathcal{A}.\text{Com}, \mathcal{A}.\text{Rsp})$ be a leakage*

impersonation adversary. Let q denote the number of \mathbf{H} queries plus the number of Prover queries that $\mathcal{A}.\text{Setup}$ and $\mathcal{A}.\text{Com}$ makes. Fig. 16 and Fig. 16 gives two adversaries \mathcal{A}_{cdh} and \mathcal{A}_{dl} such that

$$\mathbf{Adv}_{\text{ID},\ell}^{\text{imp}}(\mathcal{A})^2 \leq \mathbf{Adv}_{\mathcal{G}}^{\text{cdh}}(\mathcal{A}_{\text{cdh}}) + m \cdot \mathbf{Adv}_{\mathcal{G}}^{\text{dl}}(\mathcal{A}_{\text{dl}}) + \mathbf{Adv}_{p,m,k,\tau}^{\text{pskp}}(\ell + k/m) + \frac{q}{2^r} + \frac{1}{p}. \quad (28)$$

Additionally, let t_1 be the running time of $\mathcal{A}.\text{Setup}$, t_2 be the running time of $\mathcal{A}.\text{Com}$, t_3 be the running time of $\mathcal{A}.\text{Rsp}$, and let t_4 be the running time of ID.Kg . We have that the running time of \mathcal{A}_{cdh} and \mathcal{A}_{dl} is approximately $t_1 + t_2 + 2 \cdot t_3 + t_4$.

The proof of Theorem 16 is given in Section 5.1. The following lemma relates $\mathbf{Adv}_{p,m,k,\tau}^{\text{pskp}}(\ell + k/m)$ to the large-alphabet subkey prediction advantage (as bounded in Section 3.3).

Lemma 17. *Let p, m, k, τ, ℓ be positive integers, then*

$$\mathbf{Adv}_{p,m,k,\tau}^{\text{pskp}}(\ell) \leq \sqrt{\mathbf{Adv}_{pm,k,\tau}^{\text{skp}}(\ell) + \frac{\tau}{p}}.$$

We note that with Lemma 17, we can bound the term $\mathbf{Adv}_{p,m,k,\tau}^{\text{pskp}}(\ell)$ for any value p, m, k, τ, ℓ . Hence, the only term that is not explicitly bounded on the right-hand side of Equation (28) are $\mathbf{Adv}_{\mathcal{G}}^{\text{cdh}}(\mathcal{A})$ and $m \cdot \mathbf{Adv}_{\mathcal{G}}^{\text{dl}}(\mathcal{A})$, which can be assumed to be small when the CDH and DL problems are suspected to be hard in group G .

COMPARISON WITH ADW'S ANALYSIS. Our analysis of ADW's identification scheme improves upon the original analysis in the following ways. First, we analyze the scheme in which the challenge is generated using a random oracle directly. (The construction that uses a random oracle to derive the challenge is mentioned to be secure in [2] with no proof.) Second, while ADW's analysis is offered in the asymptotic case, we state and prove a reduction that gives *concrete security*, which lead to practical instantiation of parameters. The reduction gives a bound of the impersonation advantage in terms of three dominating quantities: CDH and DL advantages in \mathcal{G} , and a special form of subkey-prediction advantage under polynomial compression, $\mathbf{Adv}_{p,m,k,\tau}^{\text{pskp}}(\text{Lk}, \mathcal{A})$. Hence, giving a good numerical bound of the impersonation advantage amounts to bounding $\mathbf{Adv}_{p,m,k,\tau}^{\text{pskp}}(\text{Lk}, \mathcal{A})$. Here is where we make significant improvements: we use the large-alphabet subkey prediction lemma (Theorem 1) as well as a tighter polynomial-evaluation entropy preservation lemma (Lemma 17) to give significantly better concrete bounds. The comparison of parameters can be found in Fig. 14.

PARAMETER INSTANTIATION. We give an example instantiation of the ADW identification scheme with 128-bits security. First, we find a pairing friendly group G with symmetric pairing $\mathbf{e} : G \times G \rightarrow G_T$. Because of the square-root loss of security, we need 256-bit of security for CDH and DL in G . Hence, G needs to be of size roughly 512 bits. We consider $\mathcal{G} = (G, G_T, g, \mathbf{e}, p)$, where p is a prime of roughly 512 bits ($2^{511} < p < 2^{512}$). We represent elements in \mathbb{Z}_p using exactly 512 bits. We pick a big key size of 100 GB, i.e. $k^* = 8 \cdot 10^{11}$. For a choice of $m \geq 2$, we have that the block size in bits is $b = m \cdot 512$. We let $k = k^*/b$ be the size of the big key in blocks. We fix a leakage rate of 10%. By Theorem 16 and Lemma 17, to achieve 128-bit security for the identification scheme, we need 512 bits of security from $\mathbf{Adv}_{pm,k,\tau}^{\text{skp}}(\ell + \frac{k}{m})$. Hence, we need

$$\tau = \mathbf{Probes}_{k^*, \ell^* + k^*/m, s}(m \cdot 512)$$

probes. Values of $\mathbf{Probes}_{k^*, \ell^* + k^*/m, s}(m \cdot 512)$ versus various values of m is shown in Fig. 14 using both our bound and ADW's bound.

ENTROPY PRESERVATION UNDER POLYNOMIAL EVALUATION. Lemma 17 relates the prediction advantage to the large-alphabet subkey prediction advantage. Note that our bound is quantitatively better than [2, Corollary A.1]. In particular, we prove $\frac{1}{2}$ rate entropy preservation while ADW proves a rate of $\frac{1}{3}$. Before proving the lemma, we define the following quantities for jointly distributed random variables (X, Y) . Let X be a random variable, the *prediction* and *collision* probability of X is defined, respectively, to be

$$\text{Pred}(X) = \max_x \Pr[X = x], \quad \text{CP}(X) = \Pr[X = X'],$$

where X' is an independent random variable that is identically distributed to X . Additionally, suppose that (X, Y) are jointly distributed, we define the *conditional* prediction and collision probability of X given Y , respectively, to be

$$\widetilde{\text{Pred}}(X | Y) = \mathbf{E}_Y[\text{Pred}(X | Y)], \quad \widetilde{\text{CP}}(X | Y) = \mathbf{E}_Y[\text{CP}(X | Y)].$$

We note that $\text{Pred}(X | Y)$ and $\text{CP}(X | Y)$ are random variables in Y . We need the following well-known lemma,

Lemma 18. *Let (X, Y) be jointly distributed random variables, then*

$$\widetilde{\text{CP}}(X | Y) \leq \widetilde{\text{Pred}}(X | Y) \leq \sqrt{\widetilde{\text{CP}}(X | Y)}.$$

Proof. For each value y of the random variable Y , we consider the probability mass function of the random variable $X | Y = y$, $P_{X|Y=y}(\cdot)$. We note that

$$\text{Pred}(X | Y = y) = \max_x P_{X|Y=y}(x),$$

$$\text{CP}(X | Y = y) = \sum_x P_{X|Y=y}(x)^2.$$

First, we derive that

$$\text{CP}(X | Y = y) \geq \left(\max_x P_{X|Y=y}(x) \right) \cdot \sum_x P_{X|Y=y}(x) = \text{Pred}(X | Y = y).$$

Taking expectation over y sampled from Y on both sides of the above equation, we obtain that $\widetilde{\text{Pred}}(X | Y) \leq \widetilde{\text{CP}}(X | Y)$. Next, we note that viewing $\text{Pred}(X | Y = y)$, $\sqrt{\text{CP}(X | Y = y)}$ as 1- and 2-norms of $P_{X|Y=y}(\cdot)$ respectively, we have $\text{Pred}(X | Y = y) \leq \sqrt{\text{CP}(X | Y = y)}$. Hence, by the above property and Jensen's inequality

$$\begin{aligned} \widetilde{\text{Pred}}(X | Y) &= \mathbf{E}_Y[\text{Pred}(X | Y)] \\ &\leq \mathbf{E}_Y[\sqrt{\text{CP}(X | Y)}] \\ &\leq \sqrt{\mathbf{E}_Y[\text{CP}(X | Y)]} \\ &= \widetilde{\text{CP}}(X | Y). \end{aligned}$$

□

Proof (of Lemma 17). Let \mathcal{A} be any adversary and $\text{Lk} : [q]^k \rightarrow [q]^\ell$ be a leakage function. Consider the sample space defined by the experiment $\mathbf{G}_{p,m,k,\tau}^{\text{pskp}}(\mathcal{A}, \text{Lk})$ (all the coins used by the experiment and adversary \mathcal{A}). We consider all the variables used inside $\mathbf{G}_{p,m,k,\tau}^{\text{pskp}}(\mathcal{A}, \text{Lk})$ as random variables (e.g. sk^* and $L = \text{Lk}(\mathbf{K})$). We note that

$$\Pr \left[\mathbf{G}_{p,m,k,\tau}^{\text{pskp}}(\mathcal{A}, \text{Lk}) \right] \leq \widetilde{\text{Pred}}(\text{sk}^* \mid \mathbf{p}, L, e) .$$

Furthermore, by Lemma 18,

$$\widetilde{\text{Pred}}(\text{sk}^* \mid \mathbf{p}, L, e) \leq \sqrt{\widetilde{\text{CP}}(\text{sk}^* \mid \mathbf{p}, L, e)} .$$

We now need to bound $\widetilde{\text{CP}}(\text{sk}^* \mid \mathbf{p}, L, e)$. To compute this quantity. We consider another independent execution of $\mathbf{G}_{p,m,k,\tau}^{\text{pskp}}(\mathcal{A}, \text{Lk})$, where the variables in the second execution is denoted with $'$, e.g. sk' . We restrict to the event that $\text{Lk}(\text{sk}) = \text{Lk}(\text{sk}')$ and $\mathbf{p} = \mathbf{p}'$. We define polynomials p_1, \dots, p_j , which are functions of $\text{sk}, \text{sk}', \mathbf{p}$, $p_j(x) = \sum_{i=0}^{\tau-1} (\text{sk}[\mathbf{p}[i]][j] - \text{sk}'[\mathbf{p}[i]][j])x^i$. Notice that these polynomials are of degree at most τ . If $\text{sk} \neq \text{sk}'$, then at least one of p_j is a non-zero polynomial, and has at most τ roots. Hence, if $\text{sk} \neq \text{sk}'$, over a independently uniform e , the probability that $p_j(e) = 0$ is at most $\frac{\tau}{p}$ when p_j is not the zero polynomial. Finally, we derive that

$$\begin{aligned} \widetilde{\text{CP}}(\text{sk}^* \mid \mathbf{p}, L, e) &= \mathbf{E}_{\mathbf{p}, L, e} \left[\text{CP}(\text{sk}^* \mid \mathbf{p}, L, e) \right] \\ &\leq \mathbf{E}_{\mathbf{p}, L, e} \left[\Pr [\text{sk}[\mathbf{p}] = \text{sk}'[\mathbf{p}] \mid \mathbf{p}, L, e] \right. \\ &\quad \left. + \Pr [\forall j \in [m] : p_j(e) = 0 \mid \text{sk}[\mathbf{p}] \neq \text{sk}'[\mathbf{p}], \mathbf{p}, L, e] \right] \\ &= \widetilde{\text{CP}}(\text{sk}[\mathbf{p}] \mid \mathbf{p}, L) + \mathbf{E}_e \left[\Pr [\forall j \in [m] : p_j(e) = 0] \right] \\ &\leq \widetilde{\text{Pred}}(\text{sk}[\mathbf{p}] \mid \mathbf{p}, L) + \frac{\tau}{p} \\ &\leq \mathbf{Adv}_{pm,k,\tau}^{\text{pskp}}(\ell) + \frac{\tau}{p} . \end{aligned}$$

□

5.1 Proof of Theorem 16

We follow the proof technique used by [2].

Proof. (of Theorem 16). Let $\text{ID} = \text{ADW}_{\mathcal{G},k,m,\tau,r}$ be the ADW identification scheme. The reduction is very similar to the reduction from [2, Appendix B.5]. **Rewind** attempts to run a given leakage impersonation adversary \mathcal{A} twice with two different programmed challenges that only differ in the element c^* (R and e stay the same). **Rewind** takes an algorithm **Gen** that generates $(\text{prm}, \text{vk}, \text{sk}, \text{hlp}, T)$, where T is the table used by **H**. **Rewind** simulates **H** for \mathcal{A} using H_{Rewind} as described by the code. **Rewind** returns the success status of the rewinding process, along with the two responses of the two executions $(\text{rsp}_1, \text{rsp}_2)$, plus the honest response (rsp^*) and the honestly generated and compressed secrete key (sk^*) . Let $x \in \{1, 2\}$, we use $\Pr[\text{Rewind}^x(\text{Gen}, \mathcal{A})]$ to denote the probability that the first component of the output of Rewind^x is true. First, using the well-known rewind technique [5], we will argue that

$$\Pr[\text{Rewind}^1(\text{Gen}, \mathcal{A})] \geq \mathbf{Adv}_{\text{ID},\ell}^{\text{imp}}(\mathcal{A})^2 - \frac{1}{p} . \quad (29)$$

<p>Game $\text{Rewind}^1(\text{Gen}, \mathcal{A})$ $\text{Rewind}^2(\text{Gen}, \mathcal{A})$</p> <p>$(\text{prm}, \text{vk}, \text{sk}, \text{hlp}, T) \leftarrow \text{Gen}(); s \leftarrow 0$ $\text{st} \leftarrow_s \mathcal{A}.\text{Setup}^{\text{Leak}, \text{Prover}, \text{H}_{\text{Rewind}}}(\text{prm}, \text{vk}, \text{hlp})$ $(\text{com}, \text{st}') \leftarrow \mathcal{A}.\text{Com}^{\text{H}_{\text{Rewind}}}(\text{st}); \text{chl} \leftarrow \{0, 1\}^{\text{ID}.\text{Chl}}$ $T_1[0 \parallel \text{chl}, [p]] \leftarrow_s [p]; T_1[1 \parallel \text{chl}, [p]] \leftarrow \perp$ $T_1[\text{chl}, [k]^{(\tau)}] \leftarrow_s [k]^{(\tau)}; T_2 \leftarrow T_1$ If $C[\text{chl}]$ then $\text{bad} \leftarrow \text{true}$ $T[0 \parallel \text{chl}, [p]] \leftarrow \perp; T[1 \parallel \text{chl}, [p]] \leftarrow \perp; T[\text{chl}, [k]^{(\tau)}] \leftarrow \perp$ $\text{rsp}_1 \leftarrow \mathcal{A}.\text{Rsp}^{\text{H}_{\text{Rewind}}[T_1]}(\text{st}', \text{chl})$ $\text{rsp}_2 \leftarrow \mathcal{A}.\text{Rsp}^{\text{H}_{\text{Rewind}}[T_2]}(\text{st}', \text{chl})$ For $j \in [m]$ do $y[j] \leftarrow 0$ $(\text{pk}^*, \sigma^*, z^*) \leftarrow \text{ADW}.\text{Rsp}^{\text{H}_{\text{Rewind}}[T_2]}(\text{prm}, \text{hlp}, \text{vk}, \text{sk}, y, \text{chl})$ $(\mathbf{p}, c, e) \leftarrow \text{Derive}^{\text{H}}(\text{chl})$ For $j \in [m]$ do $\text{sk}^*[j] \leftarrow \sum_{i=0}^{\tau-1} (\text{sk}[\mathbf{p}[i]][j])e^i$ $A \leftarrow \text{Vrf}^{\text{H}_{\text{Rewind}}[T_1]}(\text{prm}, \text{vk}, \text{com}, \text{chl}, \text{rsp}_1)$ $B \leftarrow \text{Vrf}^{\text{H}_{\text{Rewind}}[T_2]}(\text{prm}, \text{vk}, \text{com}, \text{chl}, \text{rsp}_2)$ $C \leftarrow (T_1[1 \parallel \text{chl}, [p]] \neq T_2[1 \parallel \text{chl}, [p]])$ Return $(A \wedge B \wedge C, \text{rsp}_1, \text{rsp}_2, \text{rsp}^*, \text{sk}^*)$</p> <p><u>Gen()</u> $\text{prm} \leftarrow_s \text{Pg}(); (\text{vk}, \text{sk}, \text{hlp}) \leftarrow_s \text{Kg}(\text{prm})$ Return $(\text{prm}, \text{vk}, \text{sk}, \text{hlp}, \perp)$</p>	<p><u>$\text{H}_{\text{Rewind}}[T'](x, \text{Rng})$</u> If $\text{Rng} = [p]$ then $b \parallel x \leftarrow x$ $C[x] \leftarrow \text{true}$ If $T[x, \text{Rng}]$ then return $T[x, \text{Rng}]$ If T' then If not $T'[x, \text{Rng}]$ then $T'[x, \text{Rng}] \leftarrow_s \text{Rng}$ Return $T'[x, \text{Rng}]$ Else if not $T[x, \text{Rng}]$ then $T[x, \text{Rng}] \leftarrow_s \text{Rng}$ Return $T[x, \text{Rng}]$</p> <p><u>Prover(i, args)</u> If $\text{pst}[i] = \perp$ then // Commit $(\text{pcom}[i], \text{pst}[i]) \leftarrow_s \text{ID}.\text{Com}(\text{prm})$ Return $\text{pcom}[i]$ Else If $\text{prsp}[i] = \perp$ then // Response $\text{prsp}[i] \leftarrow_s \text{ID}.\text{Rsp}^{\text{H}_{\text{Rewind}}}(\text{prm}, \text{hlp}, \text{sk}, \text{pst}[i], \text{args})$ Return $\text{prsp}[i]$ Return \perp</p>
--	--

Figure 15: Game Rewind^1 and Rewind^2 (boxed). The oracle Leak is the same as the one given in Fig. 11.

We now justify Equation (29). We consider the event that the flags A, B, C are all set to true. Notice that the marginal probability that A is true and the marginal probability that B is true are both exactly $\Pr[\mathbf{G}_{\text{ID}, \ell}^{\text{imp}}(\mathcal{A})]$. We partition the random tape for $\mathbf{G}_{\text{ID}, \ell}^{\text{imp}}(\mathcal{A})$ into two parts: the random tape that is used upto right before $\mathcal{A}.\text{Rsp}$ is run, and the rest of the tape that is used after $\mathcal{A}.\text{Rsp}$ starts its execution. Let T be a random variable denoting the first part of the random tape. For any value of T , say t , we let $\mathbf{G}(t)$ be the game $\mathbf{G}_{\text{ID}, \ell}^{\text{imp}}(\mathcal{A})$ with the first part of random tape fixed to t . We have that

$$\begin{aligned}
\Pr[\text{Rewind}^1(\text{Gen}, \mathcal{A})] &= \Pr[A \wedge B \wedge C] \\
&= \mathbf{E}_T[\Pr[A \wedge B \wedge C \mid T]] \\
&\geq \mathbf{E}_T[\Pr[A \wedge B \mid T] - \Pr[\neg C \mid T]] \\
&= \mathbf{E}_T[\Pr[\mathbf{G}(T)]^2] - \frac{1}{p} \\
&\geq \Pr[\mathbf{G}_{\text{ID}, \ell}^{\text{imp}}(\mathcal{A})]^2 - \frac{1}{p},
\end{aligned}$$

where at the last step we used Jensen's inequality and the convexity of squaring. This justifies

<p><u>Adversary $\mathcal{A}_{\text{cdh}}(\mathcal{G}, v, h)$</u></p> <p>$(G, G_T, g, \mathbf{e}, p) \leftarrow \mathcal{G}$ $(t, \text{rsp}_1, \text{rsp}_2, \text{rsp}^*, \text{sk}^*) \leftarrow \text{Rewind}^2(\text{Gen}_{\text{cdh}}, \mathcal{A})$ $(\text{pk}_1^*, \sigma_1^*, z^{(1)}) \leftarrow \text{rsp}_1$ $(\text{pk}_2^*, \sigma_2^*, z^{(2)}) \leftarrow \text{rsp}_2$ $(\text{pk}^*, \sigma^*, z^*) \leftarrow \text{rsp}^*$ $\hat{\sigma} \leftarrow ((\sigma_1^*)^{c_1^*} / (\sigma_2^*)^{c_2^*}) \cdot \sigma^{*c_2^* - c_1^*}$ $\omega = \sum_{j=1}^m \gamma_j (z_j^{(1)} - z_j^{(2)} - x_j^*(c_1^* - c_2^*))$ $s' \leftarrow (\hat{\sigma})^{1/\omega}$ Return s'</p> <p><u>Gen_{cdh}()</u></p> <p>$(G, G_T, g, \mathbf{e}, p) \leftarrow \mathcal{G}$ For $j \in [m]$ do $\gamma_j \leftarrow \mathbb{Z}_p; g_j \leftarrow h^{\gamma_j}$ $\text{prm} = (g_0, \dots, g_{m-1}, g); \text{vk} \leftarrow v$ For $i \in [m]$ do $\text{sk}[i] \leftarrow [p]^m; \text{pk}[i] \leftarrow \prod_{j=0}^{m-1} g_j^{\text{sk}[i][j]}$ $\beta_i \leftarrow [p]; \sigma[i] \leftarrow \text{vk}^{\beta_i}$ $T[i, G] \leftarrow g^{\beta_i} / \text{pk}[i]$ $\text{hlp} \leftarrow (\text{pk}, \sigma)$ Return $(\text{prm}, \text{vk}, \text{sk}, \text{hlp}, T)$</p>	<p><u>Adversary $\mathcal{A}_{\text{dl}}(\mathcal{G}, X)$</u></p> <p>$(t, \text{rsp}_1, \text{rsp}_2, \text{rsp}^*, \text{sk}^*) \leftarrow \text{Rewind}^2(\text{Gen}_{\text{dl}}, \mathcal{A})$ $(\text{pk}_1^*, \sigma_1^*, z^{(1)}) \leftarrow \text{rsp}_1$ $(\text{pk}_2^*, \sigma_2^*, z^{(2)}) \leftarrow \text{rsp}_2$ $(\text{pk}^*, \sigma^*, z^*) \leftarrow \text{rsp}^*$ For $j = 1, \dots, m$ do $\hat{\text{sk}}^*[j] \leftarrow (z^{(1)}[j] - z^{(2)}[j]) / (c_1^* - c_2^*)$ $x \leftarrow (\sum_{i \in [m] - \{\rho\}} x_i (\hat{\text{sk}}^*[i] - \text{sk}^*[i])) / (\text{sk}^*[\rho] - \hat{\text{sk}}^*[\rho])$ Return x</p> <p><u>Gen_{dl}()</u></p> <p>$\rho \leftarrow \\$_{[m]}; g_\rho \leftarrow X$ For $j \in [m] - \{\rho\}$ do $x_j \leftarrow \\$_{\mathbb{Z}_p}; g_j \leftarrow g^{x_j}$ $\text{prm} = (g_0, \dots, g_{m-1}, g)$ $(\text{pk}, \text{sk}, \text{hlp}) \leftarrow \text{ADW.Kg}(\text{prm})$ Return $(\text{prm}, \text{pk}, \text{sk}, \text{hlp}, \perp)$</p>
--	--

Figure 16: **Left:** Adversary \mathcal{A}_{cdh} . **Right:** Adversary \mathcal{A}_{dl} .

Equation (29). Second, we argue that

$$\Pr[\text{Rewind}^1(\text{Gen}, \mathcal{A})] - \Pr[\text{Rewind}^2(\text{Gen}, \mathcal{A})] \leq \Pr[\text{Rewind}^1(\text{Gen}, \mathcal{A}) \text{ sets bad}] = \frac{q}{2^r}. \quad (30)$$

This is because the size of the table C is upper-bounded by the number of queries that $\mathcal{A}.\text{Setup}$ and $\mathcal{A}.\text{Com}$ makes to H and Derive , which is q . Next, we attempt to bound $\Pr[\text{Rewind}^2(\text{Gen}, \mathcal{A})]$. We define the following events in the game $\text{Rewind}^2(\text{Gen}, \mathcal{A})$.

$$E : A \wedge B \wedge C \wedge \left(\frac{(\text{pk}_1^*)^{(c_1^*)}}{(\text{pk}_2^*)^{(c_2^*)}} = (\text{pk}^*)^{c_1^* - c_2^*} \right),$$

$$\bar{E} : A \wedge B \wedge C \wedge \left(\frac{(\text{pk}_1^*)^{(c_1^*)}}{(\text{pk}_2^*)^{(c_2^*)}} \neq (\text{pk}^*)^{c_1^* - c_2^*} \right).$$

Notice that per construction of the events,

$$\Pr[\text{Rewind}^2(\text{Gen}, \mathcal{A})] = \Pr[E] + \Pr[\bar{E}]. \quad (31)$$

Consider \mathcal{A}_{cdh} (Fig. 16) and \mathcal{A}_{dl} (Fig. 16), which attempts to break CDH and DL problems, respectively, using Rewind^2 . We will show the following (in)equalities

$$\Pr[\bar{E}] = \mathbf{Adv}_{\mathcal{G}}^{\text{cdh}}(\mathcal{A}_{\text{cdh}}), \quad (32)$$

and

$$\Pr[E] \leq m \cdot \mathbf{Adv}_{\mathcal{G}}^{\text{dl}}(\mathcal{A}_{\text{dl}}) + \mathbf{Adv}_{p,m,k,\tau}^{\text{pskp}}\left(\ell + \frac{k}{m}\right). \quad (33)$$

This part of the analysis follows from [2, Appendix B.5] and we restate their derivation here. Assume E or \bar{E} , since the signatures verifies, for $w = \prod_{i=0}^{\tau} \mathbf{H}(\mathbf{p}[i], G)^{e^i}$, we have

$$\sigma^* = (\mathbf{pk}^* w)^s, \quad \sigma_1^* = (\mathbf{pk}_1^* w)^s, \quad \sigma_2^* = (\mathbf{pk}_2^* w)^s.$$

If \bar{E} happens, the following two values are distinct

$$\frac{(\sigma_1)^{c_1^*}}{(\sigma_2)^{c_2^*}} = \left(w^{(c_1^* - c_2^*)} \frac{(\mathbf{pk}_1^*)^{c_1^*}}{(\mathbf{pk}_2^*)^{c_2^*}} \right)^s, \quad (\sigma^*)^{c_1^* - c_2^*} = \left(w^{c_1^* - c_2^*} (\mathbf{pk})^{c_1^* - c_2^*} \right)^s.$$

Hence, the value $\hat{\sigma}$ computed by \mathcal{A}_{cdh} is

$$\hat{\sigma} = \left(\frac{(\mathbf{pk}_1^*)^{c_1^*}}{(\mathbf{pk}_2^*)^{c_2^*}} \cdot \frac{1}{(\mathbf{pk}^*)^{c_1^* - c_2^*}} \right)^s = (g^w)^s.$$

Therefore, \mathcal{A}_{cdh} can compute g^s and solve the CDH problem that it was given. This concludes the proof for Equation (32). If E happens, then we claim that $\mathbf{sk}^* = \hat{\mathbf{sk}}$ with probability at most $\mathbf{Adv}_{p,k,k,\tau}^{\text{pskp}}(\ell + k/m)$. This is true per definition of $\mathbf{Adv}_{p,k,k,\tau}^{\text{pskp}}(\ell + k/m)$. Notice that if $\mathbf{sk}^* \neq \hat{\mathbf{sk}}$, then with probability $1/m$, \mathcal{A}_{dl} can solve the DL problem that it embedded into the parameters. This is because \mathcal{A}_{dl} has two representation of \mathbf{pk}^* in the basis g_0, \dots, g_{m-1} , namely \mathbf{sk}^* and $\hat{\mathbf{sk}}$. This concludes the proof of Equation (33). Notice that Equations (29), (31), (30), (32), and (33) together implies the theorem. Finally, notice that \mathcal{A}_{cdh} and \mathcal{A}_{dl} has roughly the running time of Rewind^2 and ADW.Kg , which is about $t_1 + t_2 + 2t_3 + t_4$. \square

6 Acknowledgements

We thank the anonymous CCS 2018 reviewers for their helpful comments. We thank Yevgeniy Dodis and Daniel Wichs for helpful discussions.

References

- [1] J. Alwen, Y. Dodis, M. Naor, G. Segev, S. Walfish, and D. Wichs. Public-key encryption in the bounded-retrieval model. In H. Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 113–134. Springer, Heidelberg, May 2010. 3, 5
- [2] J. Alwen, Y. Dodis, and D. Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In S. Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 36–54. Springer, Heidelberg, Aug. 2009. 3, 4, 5, 6, 26, 27, 29, 30, 31, 32, 35
- [3] J. Alwen, Y. Dodis, and D. Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. Cryptology ePrint Archive, Report 2009/160, 2009. <http://eprint.iacr.org/2009/160>. 3, 4, 6
- [4] M. Bellare, D. Kane, and P. Rogaway. Big-key symmetric encryption: Resisting key exfiltration. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 373–402. Springer, Heidelberg, Aug. 2016. 3, 4, 5, 9, 12, 14, 19, 20, 21, 22, 23

- [5] M. Bellare and A. Palacio. GQ and Schnorr identification schemes: Proofs of security against impersonation under active and concurrent attacks. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 162–177. Springer, Heidelberg, Aug. 2002. 26, 32
- [6] M. Bellare, B. Poettering, and D. Stebila. From identification to signatures, tightly: A framework and generic transforms. In J. H. Cheon and T. Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 435–464. Springer, Heidelberg, Dec. 2016. 7
- [7] M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006. 7, 24
- [8] D. Boneh, B. Lynn, and H. Shacham. Short signatures from the Weil pairing. In C. Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 514–532. Springer, Heidelberg, Dec. 2001. 27
- [9] D. Cash, Y. Z. Ding, Y. Dodis, W. Lee, R. J. Lipton, and S. Walfish. Intrusion-resilient key exchange in the bounded retrieval model. In S. P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 479–498. Springer, Heidelberg, Feb. 2007. 3, 5
- [10] F. Chen, D. A. Koufaty, and X. Zhang. Understanding intrinsic characteristics and system implications of flash memory based solid state drives. In *ACM SIGMETRICS Performance Evaluation Review*, volume 37, pages 181–192. ACM, 2009. 6
- [11] <https://www.cryptopp.com/benchmarks.html>, 2015. Accessed: 2017-05-18. 6
- [12] G. Di Crescenzo, R. J. Lipton, and S. Walfish. Perfectly secure password protocols in the bounded retrieval model. In S. Halevi and T. Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 225–244. Springer, Heidelberg, Mar. 2006. 3, 5
- [13] S. Dziembowski. Intrusion-resilience via the bounded-storage model. In S. Halevi and T. Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 207–224. Springer, Heidelberg, Mar. 2006. 3, 5
- [14] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In A. M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, Aug. 1987. 26
- [15] McAfee. Stop data exfiltration, 2015. August 2015. <https://www.mcafee.com/us/resources/solution-briefs/sb-quarterly-threats-aug-2015-1.pdf>. 5
- [16] T. Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In E. F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 31–53. Springer, Heidelberg, Aug. 1993. 6, 27
- [17] H. Robbins. A remark on stirling's formula. *The American Mathematical Monthly*, 62(1):26–29, 1955. 37
- [18] C.-P. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991. 6

A Proofs of Lemmas 2, 3, and 4

We need the following version of Stirling's approximation of $n!$.

Lemma 19. [17] For any $n \in \mathbb{Z}^+$,

$$\sqrt{2\pi n} \left(\frac{n}{e}\right)^n e^{\frac{1}{12n+1}} \leq n! \leq \sqrt{2\pi n} \left(\frac{n}{e}\right)^n e^{\frac{1}{12n}}.$$

We first prove Lemma 2.

Proof (of Lemma 2). We first show the lower bound Equation (6). Notice that by definition of $H_q(r/k)$,

$$q^{kH_q(r/k)} = (q-1)^r (r/k)^{-r} (1-r/k)^{r-k}.$$

Hence, by Lemma 19,

$$\begin{aligned} B_{q,k}(r) &= \sum_{i=0}^r (q-1)^i \binom{k}{i} \\ &\geq (q-1)^r \frac{k!}{r!(k-r)!} \\ &\geq (q-1)^r \frac{\sqrt{2\pi k} \left(\frac{k}{e}\right)^k e^{\frac{1}{12k+1}}}{\sqrt{2\pi r} \left(\frac{r}{e}\right)^r e^{\frac{1}{12r}} \sqrt{2\pi(k-r)} \left(\frac{k-r}{e}\right)^{k-r} e^{\frac{1}{12(k-r)}}} \\ &= q^{kH_q(r/k)} \frac{\sqrt{k} e^{\frac{1}{12k+1}}}{\sqrt{2\pi r(k-r)} e^{\frac{1}{12r}} e^{\frac{1}{12(k-r)}}} \\ &= q^{kH_q(r/k) - \epsilon(k,r)}. \end{aligned}$$

Now, we assume that $r \leq k - k/q$ and derive the upper bound, Equation (7).

$$\begin{aligned} \frac{B_{q,k}(r)}{q^{kH_q(r/k)}} &= \frac{\sum_{i=0}^r (q-1)^i \binom{k}{i}}{(q-1)^r (r/k)^{-r} (1-r/k)^{r-k}} \\ &= \sum_{i=0}^r \binom{k}{i} (q-1)^{i-r} (r/k)^r (1-r/k)^{k-r} \\ &= \sum_{i=0}^r \binom{k}{i} (q-1)^i (1-r/k)^k \left(\frac{r/k}{(q-1)(1-r/k)}\right)^r \\ &\leq \sum_{i=0}^r \binom{k}{i} (q-1)^i (1-r/k)^k \left(\frac{r/k}{(q-1)(1-r/k)}\right)^i \\ &= \sum_{i=0}^r \binom{k}{i} (r/k)^i (1-r/k)^{k-i} \\ &\leq \sum_{i=0}^k \binom{k}{i} (r/k)^i (1-r/k)^{k-i} \\ &= 1, \end{aligned}$$

where the first inequality is by the fact that $r/k \leq (q-1)(1-r/k)$ if $r \leq k - k/q$. □

Lemma 3 follows from Lemma 2.

Proof (of Lemma 3). Per definition of $\text{rd}_{q,k}(N)$, it suffices to show that

$$B_{q,k}(r) \leq N ,$$

for $r = \lfloor H_q^{-1}(\log_q(N)/k) \cdot k \rfloor$. Per definition of H_q^{-1} , $r \leq (1 - 1/q) \cdot k$. Hence, we can apply Equation (7) and obtain

$$B_{q,k}(r) \leq q^{kH_q(r/k)} \leq q^{kH_q(H_q^{-1}(\log_q(N)/k))} = N .$$

□

Lastly, we prove Lemma 4.

Proof (of Lemma 4). We first show the lower bound that

$$\min(x, 1 - \frac{1}{q}) - \frac{1}{\log_2(q)} \leq H_q^{-1}(x) . \quad (34)$$

Note that this is trivially true if the left-hand side of Equation (34) is negative. Hence, we suppose that the left-hand side of Equation (34) is non-negative. As noted before, H_q is increasing in the domain $[0, 1 - 1/q]$. Additionally, note that $\min(x, 1 - 1/q) - 1/\log_2(q) \leq 1 - 1/q$. Hence, it suffices to show

$$H_q\left(\min(x, 1 - \frac{1}{q}) - \frac{1}{\log_2(q)}\right) \leq x . \quad (35)$$

We consider two cases. Case 1, $x \leq (1 - 1/q)$. Case 2, $(1 - 1/q) \leq x \leq 1$. We claim that both cases follow from the equation below, which holds for $x \in [\log_2(q), 1]$.

$$H_q\left(x - \frac{1}{\log_2(q)}\right) \leq x . \quad (36)$$

Case 1 is directly implied by Equation (36). For case 2, note that the left-hand side of Equation (34) always evaluate to $1 - 1/q - 1/\log_2(q)$. Hence, by Equation (36), $H_q(1 - 1/q - 1/\log_2(q)) \leq 1 - 1/q \leq x$. Finally, we justify Equation (36). Recall that $H_q(x) = H_2(x)/\log_2(q) + x \log_q(q-1)$. We compute

$$\begin{aligned} H_q\left(x - \frac{1}{\log_2(q)}\right) &= \frac{H_2\left(x - \frac{1}{\log_2(q)}\right)}{\log_2(q)} + \left(x - \frac{1}{\log_2(q)}\right) \log_q(q-1) \\ &\leq \frac{1}{\log_2(q)} + x \log_q(q) \cdot \frac{q-1}{q} - \frac{\log_q(q-1)}{\log_2(q)} \\ &= \frac{1}{\log_2(q)} + x - x \log_q\left(\frac{q}{q-1}\right) - \frac{\log_q(q-1)}{\log_2(q)} \\ &= x + \frac{1}{\log_2(q)} \left(1 - \log_q\left(\frac{q^{(x \log_2(q))}}{(q-1)^{(x \log_2(q)-1)}}\right)\right) \\ &\leq x + \frac{1}{\log_2(q)} (1 - \log_q(q)) \\ &= x . \end{aligned}$$

Next, we show the upper bound that

$$H_q^{-1}(x) \leq x(1 - \frac{1}{q}). \quad (37)$$

Similar to the lower bound we just obtained, we note that it suffices to show $H_q(x(1 - \frac{1}{q})) \geq x$. Let us define, for $x \in [0, 1]$:

$$f(x) = \frac{x}{q} \log_q\left(\frac{x}{q}\right) - x \log_q(x) - (1 - x + \frac{x}{q}) \log_q(1 - x + \frac{x}{q}).$$

We will show that $H_q(1(1 - 1/q)) = x + f(x)$. The derivation is as follows.

$$\begin{aligned} & H_q(x(1 - 1/q)) \\ &= x(1 - 1/q) \log_q(q - 1) - x(1 - 1/q) \log_q(x(1 - 1/q)) \\ &\quad - (1 - x + x/q) \log_q(1 - x + x/q) \\ &\geq x \log_q(q - 1) - x/q \log_q(q - 1) \\ &\quad - x(1 - 1/q)(\log_q(x) + \log_q(1 - 1/q)) \\ &\quad - (1 - x + x/q) \log_q(1 - x + x/q) \\ &= x \log_q(q - 1) - x/q \log_q(q - 1) \\ &\quad - x \log_q(x) - x \log_q(1 - 1/q) + x/q \log_q(x) + x/q \log_q(1 - 1/q) \\ &\quad - (1 - x + x/q) \log_q(1 - x + x/q) \\ &= x \left(\log_q(q - 1) + \log_q(q/(q - 1)) \right) - x \log_q(x) \\ &\quad - x/q \left(\log_q(q - 1) + \log_q(1/x) + \log_q(q/(q - 1)) \right) \\ &\quad - (1 - x + x/q) \log_q(1 - x + x/q) \\ &= x + x/q \log_q(x/q) - x \log_q(x) - (1 - x + x/q) \log_q(1 - x + x/q) \\ &= x + f(x). \end{aligned}$$

Lastly, we show that $f(x) \geq 0$ for any $x \in [0, 1]$. First, we check that $f(0) = f(1) = 0$. Next, check that the second derivative of f ,

$$f''(x) = \frac{q - 1}{x(qx - q - x)} \leq 0,$$

is at most 0 for any $x \in [0, 1]$. We omit the details of the derivative computation here. Hence, f is concave over the domain $[0, 1]$, with $f(0) = f(1) = 0$. Thence, $f(x) \geq 0$ for all $x \in [0, 1]$. \square

B A Rejection Sampling Lemma

We prove the following lemma, which allows one to efficiently sample from $[k]^{(\tau)}$, for appropriately constrained integers k, τ , via rejection sampling.

Lemma 20. Let τ, k, c be positive integers. Suppose $k \geq 2(c+1) \cdot \tau^2$. Let $x_1, \dots, x_t \leftarrow^s [k]$ be i.i.d samples, where

$$t = \tau + \left\lceil \frac{c\tau}{\log_2(k) - \log_2((c+1)\tau^2)} \right\rceil \leq \tau + c\tau.$$

Let $S = \{x_1, \dots, x_t\}$. Then

$$\Pr[|S| < \tau] \leq 2^{-c\tau}.$$

Proof. Let $\delta = t - \tau$. Since $k \geq 2(c+1) \cdot \tau^2$, we have $\log_2(k) - \log_2((c+1)\tau^2) \geq 1$. Hence, $\delta \leq c \cdot \tau$. Define $S_i = \{x_1, \dots, x_i\}$. Hence, $S_0 = \emptyset$ and $S_t = S$. Suppose that $|S| < \tau$, then there exists at least δ positions i such that $x_i \in S_{i-1}$. Since x_i is a independent uniform sample from $[k]$, the probability that $x_i \in S_{i-1}$ is $|S_{i-1}|/k$, which is at most τ/k . Hence,

$$\begin{aligned} \Pr[|S| < \tau] &\leq \binom{\tau + \delta}{\delta} \left(\frac{\tau}{k}\right)^\delta \\ &\leq \binom{(c+1)\tau}{\delta} \left(\frac{\tau}{k}\right)^\delta \\ &\leq \left(\frac{(c+1)\tau^2}{k}\right)^\delta. \end{aligned}$$

Hence,

$$\log_2(\Pr[|S| < \tau]) \leq \delta \log_2\left(\frac{(c+1)\tau^2}{k}\right) \leq -c \cdot \tau.$$

□