

Efficient Non-Malleable Codes and Key-Derivation for Poly-Size Tampering Circuits

Sebastian Faust¹, Pratyay Mukherjee^{2*}, Daniele Venturi³, and Daniel Wichs^{4**}

¹ EPFL Switzerland

² Aarhus University

³ Sapienza University of Rome

⁴ Northeastern University

Abstract. Non-malleable codes, defined by Dziembowski, Pietrzak and Wichs (ICS '10), provide roughly the following guarantee: if a codeword c encoding some message x is tampered to $c' = f(c)$ such that $c' \neq c$, then the tampered message x' contained in c' reveals no information about x . Non-malleable codes have applications to immunizing cryptosystems against tampering attacks and related-key attacks.

One *cannot* have an *efficient* non-malleable code that protects against *all efficient* tampering functions f . However, in this work we show “the next best thing”: for any polynomial bound s given a-priori, there is an efficient non-malleable code that protects against all tampering functions f computable by a circuit of size s . More generally, for any family of tampering functions \mathcal{F} of size $|\mathcal{F}| \leq 2^s$, there is an efficient non-malleable code that protects against all $f \in \mathcal{F}$. The *rate* of our codes, defined as the ratio of message to codeword size, approaches 1. Our results are information-theoretic and our main proof technique relies on a careful probabilistic method argument using limited independence. As a result, we get an efficiently samplable family of efficient codes, such that a random member of the family is non-malleable with overwhelming probability. Alternatively, we can view the result as providing an efficient non-malleable code in the “common reference string” (CRS) model.

We also introduce a new notion of non-malleable key derivation, which uses randomness x to derive a secret key $y = h(x)$ in such a way that, even if x is tampered to a different value $x' = f(x)$, the derived key $y' = h(x')$ does not reveal any information about y . Our results for non-malleable key derivation are analogous to those for non-malleable codes. As a useful tool in our analysis, we rely on the notion of “leakage-resilient storage” of Davì, Dziembowski and Venturi (SCN '10) and, as a result of independent interest, we also significantly improve on the parameters of such schemes.

1 Introduction

Non-malleable codes were introduced by Dziembowski, Pietrzak and Wichs [18]. They provide meaningful guarantees on the integrity of an encoded message in

* Research supported by a European Research Commission Starting Grant (no. 279447), the CTIC and CFEM research center.

** Research supported by NSF grant 1314722.

the presence of tampering, even in settings where error-correction and error-detection may not be possible. Intuitively, a code (Enc, Dec) is non-malleable w.r.t. a family of tampering functions \mathcal{F} if the message contained in a codeword modified via a function $f \in \mathcal{F}$ is either the original message, or a completely unrelated value. For example, it should not be possible to just flip 1 bit of the message by tampering the codeword via a function $f \in \mathcal{F}$. More formally, we consider an experiment Tamper_x^f in which a message x is (probabilistically) encoded to $c \leftarrow \text{Enc}(x)$, the codeword is tampered to $c' = f(c)$ and, if $c' \neq c$, the experiment outputs the tampered message $x' = \text{Dec}(c')$, else it outputs a special value same^* . We say that the code is non-malleable w.r.t. some family of tampering functions \mathcal{F} if, for every function $f \in \mathcal{F}$ and every messages x , the experiment Tamper_x^f reveals almost no information about x . More precisely, we say that the code is ε -non-malleable if for every pair of messages x, x' and every $f \in \mathcal{F}$, the distributions Tamper_x^f and $\text{Tamper}_{x'}^f$ are statistically ε -close. The encoding/decoding functions are public and do not contain any secret keys. This makes the notion of non-malleable codes different from (but conceptually related to) the well-studied notions of non-malleability in cryptography, introduced by the seminal work of Dolev, Dwork and Naor [16].

Relation to Error Correction/Detection. Notice that non-malleability is a weaker guarantee than error correction/detection; the latter ensure that any change in the codeword can be corrected or at least detected by the decoding procedure, whereas the former does allow the message to be modified, but only to an unrelated value. However, when studying error correction/detection we usually restrict ourselves to limited forms of tampering which preserve some notion of distance (e.g., usually hamming distance) between the original and tampered codeword. (One exception is [12], which studies error-detection for more complex tampering.) For example, it is already impossible to achieve error correction/detection for the simple family of functions $\mathcal{F}_{\text{const}}$ which, for every constant c^* , includes a “constant” function f_{c^*} that maps all inputs to c^* . There is always some function in $\mathcal{F}_{\text{const}}$ that maps everything to a *valid* codeword c^* . In contrast, it is trivial to construct codes that are non-malleable w.r.t $\mathcal{F}_{\text{const}}$, as the output of a constant function is clearly independent of its input. The prior works on non-malleable codes, together with the results from this work, show that one can construct non-malleable codes for highly complex tampering-function families \mathcal{F} for which error correction/detection are unachievable.

Applications to Tamper-Resilience. The fact that non-malleable codes can be built for large and complex families of functions makes them particularly attractive as a mechanism for protecting memory against tampering attacks, known to be a serious threat for the security of cryptographic schemes [7,2,29,11]. As shown in [18], to protect a scheme with some secret state against memory-tampering, we simply encode the state via a non-malleable code and store the encoding in the memory instead of the original secret. One can show that if the code is non-malleable with respect to function family \mathcal{F} , the transformed system

is secure against tampering attacks carried out by any function in \mathcal{F} . See [18] for a discussion of the application of non-malleable codes to tamper resilience.

Limitations & Possibility. It is *impossible* to have codes that are non-malleable for *all* possible tampering functions. For any coding scheme (Enc, Dec) , there exists a tampering function $f_{\text{bad}}(c)$ that recovers $x = \text{Dec}(c)$, creates x' by (e.g.,) flipping the first bit of x , and outputs a valid encoding c' of x' . Notice that if Enc, Dec are *efficient*, then the function f_{bad} is efficient as well. Thus, it is also impossible to have an *efficient* code which is non-malleable w.r.t all *efficient* functions. Prior works [26,10,17,1,9,19] (discussed shortly) constructed non-malleable codes for several rich and interesting function families. In all cases, the families are restricted through their *granularity* rather than their computational *complexity*. In particular, these works envision that the codeword is split into several (possibly just 2) components, each of which can only be tampered independently of the others. The tampering function therefore only operates on a “granular” rather than “global” view of the codeword.

1.1 Our Contribution

In this work, we are interested in designing non-malleable codes for large families of functions which are only restricted by their “computational complexity” rather than “granularity”. As we saw, we cannot have a single efficient code that is non-malleable for all efficient tampering functions. However, we show the following positive result, which we view as the “next best thing”:

Main Result: For any polynomial bound $s = s(n)$ in the codeword size n , and any tampering family \mathcal{F} of size $|\mathcal{F}| \leq 2^s$, there is an efficient code of complexity $\text{poly}(s, \log(1/\varepsilon))$ which is ε -non-malleable w.r.t. \mathcal{F} . In particular, \mathcal{F} can be the family of all circuits of size at most s .

The code is secure in the information theoretic setting, and achieves *optimal rate* (message/codeword size) arbitrarily close to 1. It has a *simple* construction relying only on t -wise independent hashing.

The CRS Model. In more detail, if we fix some family \mathcal{F} of tampering functions (e.g., circuits of bounded size), our result gives us a *family* of efficient codes, such that, with overwhelming probability, a random member of the family is non-malleable w.r.t \mathcal{F} . Each code in the family is indexed by some hash function h from a t -wise independent family of hash functions \mathcal{H} . This result already shows the *existence* of efficient non-malleable codes with some small *non-uniform advice* to indicate a “good” hash function h .

However, we can also efficiently sample a random member of the code family by sampling a random hash function h . Therefore, we find it most appealing to think of this result as providing a uniformly efficient *construction* of non-malleable code in the “common reference string (CRS)” model, where a random public string consisting of the hash function h is selected once and fixes the

non-malleable code. We emphasize that, although the family \mathcal{F} (e.g., circuits of bounded size) is fixed prior to the choice of the CRS, the attacker can choose the tampering function $f \in \mathcal{F}$ (e.g., a particular small circuit) adaptively depending on the choice of h .

We argue that it is unlikely that we can completely de-randomize our construction and come up with a fixed uniformly-efficient code which is non-malleable for all circuits of size (say) $s = O(n^2)$. In particular, this would require a circuit lower bound, showing that the function f_{bad} (described above) *cannot* be computed by a circuit of size $O(n^2)$.

Non-Malleable Key-Derivation. As an additional contribution, we introduce a new primitive called *non-malleable key derivation*. Intuitively, a function $h : \{0, 1\}^n \rightarrow \{0, 1\}^k$ is a non-malleable key derivation for tampering-family \mathcal{F} if it guarantees that for any tampering function $f \in \mathcal{F}$, if we sample uniform randomness $x \leftarrow \{0, 1\}^n$, the “derived key” $y = h(x)$ is statistically close to uniform even given $y' = h(f(x))$ derived from “tampered” randomness $f(x) \neq x$. Our positive results for non-malleable key derivation are analogous to those for non-malleable codes. One difference is that the rate k/n is now at most $1/2$ rather than 1, and we show that this is optimal.

While we believe that non-malleable key derivation is an interesting notion on its own (e.g., it can be viewed as a dual version of non-malleable extractors [15]), we also show it has useful applications for tamper resilience. For instance, consider some cryptographic scheme G using a uniform key in $y \leftarrow \{0, 1\}^k$. To protect G against tampering attacks, we can store a bigger key $x \leftarrow \{0, 1\}^n$ on the device and temporarily derive $y = h(x)$ each time we want to execute G . In the full version of this paper [20], we show that this approach protects any cryptographic scheme with a uniform key against one-time tampering attacks. The main advantage of using a non-malleable key-derivation rather than non-malleable codes is that the key x stored in memory is simply a uniformly random string with no particular structure (in contrast, the codeword in a non-malleable code requires structure).

In the full version, we also show how to use non-malleable key derivation to build a tamper-resilient stream cipher. Our construction is based on a PRG $\text{prg} : \{0, 1\}^k \rightarrow \{0, 1\}^{n+v}$ and a non-malleable key derivation function $h : \{0, 1\}^n \rightarrow \{0, 1\}^k$. For an initial key $s_0 \leftarrow \{0, 1\}^n$, sampled uniformly at random, the output of the stream cipher at each round $i \in [q]$ is $(s_i, x_i) := \text{prg}(h(s_{i-1}))$.

1.2 Our Techniques

Non-Malleable Codes. Our construction of non-malleable codes is incredibly simple and relies on t -wise independent hashing, where t is proportional to $s = \log |\mathcal{F}|$. In particular, if h_1, h_2 are two such hash functions, we encode a message x into a codeword $c = (r, z, \sigma)$ where r is randomness, $z = x \oplus h_1(r)$ and $\sigma = h_2(r, z)$. The security analysis, on the other hand, requires two independently interesting components. Firstly, we rely on the notion of *leakage-resilient*

encodings, proposed by Davì, Dziembowski and Venturi [14]. These provide a method to encode a secret in such a way that a limited form of leakage on the encoding does not reveal anything about the secret. One of our contributions is to significantly improve the parameters of the construction from [14] by using a fresh and more careful analysis, which gives us such schemes with an essentially optimal rate. Secondly, we analyze a simpler/weaker notion of *bounded* non-malleability, which intuitively guarantees that an adversary seeing the decoding of a tampered codeword can learn only a bounded amount of information on the encoded value. This notion of bounded non-malleability is significantly simpler to analyze than full non-malleability. Finally, we show how to carefully combine leakage-resilient encodings with bounded non-malleability to get our full construction of non-malleable codes. On a very high (and not entirely precise) level, we can think of h_1 above as providing “leakage resilience” and h_2 as providing “bounded non-malleability”.

We stress that the fact that t has to be proportional to s is not an artefact of our proof. In fact, one can see that whenever the hash function has seed size s , there is a family of 2^s functions that breaks the construction with probability 1: For each seed, just have a new function that decodes with that seed and encodes a related value. This shows that the t has to be proportional to $\log |\mathcal{F}|$.

Non-Malleable Key-Derivation. Our construction of non-malleable key-derivation functions is even simpler: a random t -wise independent hash function h already satisfies the definition with overwhelming probability, where t is proportional to $s = \log |\mathcal{F}|$. The analysis is again subtle and relies on a careful probabilistic method argument.

Similar to the case of non-malleable codes, the fact that t has to be proportional to s is necessary.

1.3 Related Works

Granular Tampering. Most of the earlier works on non-malleable codes focus on granular tampering models, where the tampering functions are restricted to act on individual components of the codeword independently. The original work of [18] gives an efficient construction for bit-tampering (i.e., the adversary can tamper with each bit of the codeword independently of every other bit). Very recently, Cheraghchi and Guruswami [9] gave a construction with improved rate and better efficiency for the same family. Choi *et al.* [10] considered an extended tampering family, where the tampering function can be applied to a small (logarithmic in the security parameter) number of blocks independently.

Perhaps the least granular and most general such model is the so-called *split-state* model, where the encoding consists of two parts L (left) and R (right), and the adversary can tamper L and R *arbitrarily but independently*. Starting with the random oracle construction of [18], a few other constructions of non-malleable split-state codes have been proposed, both in the computational setting [26,19] and in the information theoretic setting [17,1,9]. Notice that the family \mathcal{F}_{split}

of all split-state tampering functions (without restricting efficiency), has doubly exponential size $2^{O(2^{n/2})}$ in the codeword size n , and therefore it is not covered by our results, which can efficiently handle at most singly-exponential-size families $2^{\text{poly}(n)}$. On the other hand, the split-state model doesn't cover "computationally simple" functions, such as the function computing the XOR or the bit-wise inner-product of L, R . Therefore, although the works are technically orthogonal, we believe that looking at computational complexity may be more natural.

Global Tampering. The work of [18] gives an *existential* (inefficient) construction of non-malleable codes for doubly-exponential sized function families. More precisely, for any constant $0 < \alpha < 1$ and any family \mathcal{F} of functions of size $|\mathcal{F}| \leq 2^{2^{\alpha n}}$ in the codeword size n , there exists an inefficient non-malleable code w.r.t. \mathcal{F} ; indeed a completely random function gives such a code with high probability. The code is clearly not efficient, and this should be expected for such a broad result: the families \mathcal{F} can include all circuits of size (e.g.,) $s(n) = 2^{n/2}$, which means that the efficiency of the code must exceed $O(2^{n/2})$. Unfortunately, there is no direct way to "scale down" the result in [18] so as to get an efficient construction for singly-exponential-size families. (One can view our work as providing such "scaled down" result.) Moreover, the analysis only yielded a rate of at most $(1 - \alpha)/3 < 1/3$, and it was previously not known if such codes can achieve a rate close to 1, even for "small" function families. We note that [18] also showed that the probabilistic method construction can yield efficient non-malleable codes for large function families in the *random-oracle model*. However, this only considers function families that don't have access to the random-oracle. For example, one cannot interpret this as giving any meaningful result for tampering-functions with bounded complexity.

Concurrent and Independent Work. In a concurrent and independent work, Cheraghchi and Guruswami [8] give two related results. Firstly, they improve the probabilistic method construction of [18] and show that, for families \mathcal{F} of size $|\mathcal{F}| \leq 2^{2^{\alpha n}}$, there exist (inherently inefficient) non-malleable codes with rate $1 - \alpha$, which they also show to be optimal. This gives the first characterization of the rate of non-malleable codes. Secondly, similar to our results, they use limited independence to construct efficient non-malleable codes when restricted to tampering families \mathcal{F} of size $|\mathcal{F}| \leq 2^{s(n)}$ for a polynomial $s(n)$. However, the construction of [8] is not "efficient" in the usual cryptographic sense: to get error-probability ε , the encoding and decoding procedures require complexity $\text{poly}(1/\varepsilon)$. If we set ε to be negligible, as usually desired in cryptography, then the encoding/decoding procedures would require super-polynomial time. In contrast, the encoding/decoding procedures in our construction have efficiency $\text{poly}(\log(1/\varepsilon))$, and therefore we can set ε to be negligible while maintaining polynomial-time encoding/decoding.

Other Approaches to Achieve Tamper Resilience. There is a vast body of literature that considers tampering attacks using other approaches besides

non-malleable codes. See, e.g., [5,22,24,4,21,25,3,23,27,30,6,13]. The reader is referred to (e.g.,) [18] for a more detailed comparison between these approaches and non-malleable codes.

2 Preliminaries

Notation. We denote the set of first n natural numbers, i.e. $\{1, \dots, n\}$, by $[n]$. Let X, Y be random variables with supports $S(X), S(Y)$, respectively. We define

$$\mathbf{SD}(X, Y) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{s \in S(X) \cup S(Y)} |\Pr[X = s] - \Pr[Y = s]|$$

to be their *statistical distance*. We write $X \approx_\varepsilon Y$ and say that X and Y are ε -statistically close to denote that $\mathbf{SD}(X, Y) \leq \varepsilon$. We let U_n denote the uniform distribution over $\{0, 1\}^n$. We use the notation $x \leftarrow X$ to denote the process of sampling a value x according to the distribution X . If f is a randomized algorithm, we write $f(x; r)$ to denote the execution of f on input x with random coins r . We let $f(x)$ denote a random variable over the random coins.

2.1 Definitions of Non-Malleable Codes

Definition 1 (Coding Scheme). A (k, n) -coding scheme consists of two functions: a randomized encoding function $\text{Enc} : \{0, 1\}^k \rightarrow \{0, 1\}^n$, and deterministic decoding function $\text{Dec} : \{0, 1\}^n \rightarrow \{0, 1\}^k \cup \{\perp\}$ such that, for each $x \in \{0, 1\}^k$, $\Pr[\text{Dec}(\text{Enc}(x)) = x] = 1$.

We now define non-malleability w.r.t. some family \mathcal{F} of tampering functions. The work of [18] defines a default and a strong version of non-malleability. The main difference is that, in the default version, the tampered codeword $c' \neq c$ may still encode the original message x whereas the strong version ensures that any change to the codeword completely destroys the original message. We only define the strong version below. We then add an additional strengthening which we call *super* non-malleability.

Definition 2 (Strong Non-Malleability [18]). Let (Enc, Dec) be a (k, n) -coding scheme and \mathcal{F} be a family of functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$. We say that the scheme is $(\mathcal{F}, \varepsilon)$ -non-malleable if for any $x_0, x_1 \in \{0, 1\}^k$ and any $f \in \mathcal{F}$, we have $\text{Tamper}_{x_0}^f \approx_\varepsilon \text{Tamper}_{x_1}^f$ where

$$\text{Tamper}_x^f \stackrel{\text{def}}{=} \left\{ \begin{array}{l} c \leftarrow \text{Enc}(x), c' := f(c), x' = \text{Dec}(c') \\ \text{Output same}^* \text{ if } c' = c, \text{ and } x' \text{ otherwise.} \end{array} \right\}. \quad (1)$$

For *super* non-malleable security (defined below), if the tampering manages to modify c to c' such that $c' \neq c$ and $\text{Dec}(c') \neq \perp$, then we will even give the attacker the tampered codeword c' in *full* rather than just giving $x' = \text{Dec}(c')$. We do not immediately see a concrete application of this strengthening, but it seems sufficiently interesting to define explicitly.

Definition 3 (Super Non-Malleability). Let (Enc, Dec) be a (k, n) -coding scheme and \mathcal{F} be a family of functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$. We say that the scheme is $(\mathcal{F}, \varepsilon)$ -super non-malleable if for any $x_0, x_1 \in \{0, 1\}^k$ and any $f \in \mathcal{F}$, we have $\text{Tamper}_{x_0}^f \approx_\varepsilon \text{Tamper}_{x_1}^f$ where:

$$\text{Tamper}_x^f \stackrel{\text{def}}{=} \left\{ \begin{array}{l} c \leftarrow \text{Enc}(x), c' := f(c) \\ \text{Output same}^* \text{ if } c' = c, \text{ output } \perp \text{ if } \text{Dec}(c') = \perp, \\ \text{and else output } c'. \end{array} \right\}. \quad (2)$$

3 Improved Leakage-Resilient Codes

We will rely on leakage-resilience as an important tool in our analysis. The following notion of leakage-resilient codes was defined by [14]. Informally, a code is leakage resilience w.r.t some leakage family \mathcal{F} if, for any $f \in \mathcal{F}$, “leaking” $f(c)$ for a codeword c does not reveal anything about the encoded value.

Definition 4 (Leakage-Resilient Codes [14]). Let $(\text{LREnc}, \text{LRDec})$ be a (k, n) -coding scheme. For a function family \mathcal{F} , we say that $(\text{LREnc}, \text{LRDec})$ is $(\mathcal{F}, \varepsilon)$ -leakage-resilient, if for any $f \in \mathcal{F}$ and any $x \in \{0, 1\}^k$ we have $\text{SD}(f(\text{LREnc}(x)), f(U_n)) \leq \varepsilon$.

The work of [14] gave a probabilistic method construction showing that such codes exist and can be efficient when the size of the leakage family $|\mathcal{F}|$ is singly-exponential. However, the rate k/n was at most some small constant ($< \frac{1}{4}$), even when the family size $|\mathcal{F}|$ and the leakage size ℓ are small. Here, we take the construction of [14] and give an improved analysis with improved parameters, showing that the rate can approach 1. In particular, the additive overhead of the code is very close to the leakage-amount ℓ , which is optimal. Our result and analysis are also related to the “high-moment crooked leftover hash lemma” of [28], although our construction is somewhat different, relying only on high-independence hash-functions rather than permutations.

Construction. Let \mathcal{H} be a t -wise independent function family consisting of functions $h : \{0, 1\}^v \rightarrow \{0, 1\}^k$. For any $h \in \mathcal{H}$ we define the $(k, n = k + v)$ -coding scheme $(\text{LREnc}_h, \text{LRDec}_h)$ where: (1) $\text{LREnc}_h(x) := (r, h(r) \oplus x)$ for $r \leftarrow \{0, 1\}^v$; (2) $\text{LRDec}_h((r, z)) := z \oplus h(r)$.

Theorem 1. Fix any function family \mathcal{F} consisting of functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$. With probability $1 - \rho$ over the choice of a random $h \leftarrow \mathcal{H}$, the coding scheme $(\text{LREnc}_h, \text{LRDec}_h)$ is $(\mathcal{F}, \varepsilon)$ -leakage-resilient as long as:

$$t \geq \log |\mathcal{F}| + \ell + k + \log(1/\rho) + 3 \quad \text{and} \quad v \geq \ell + 2 \log(1/\varepsilon) + \log(t) + 3.$$

For space reasons, the proof of Theorem 1 is deferred to the full version [20].

4 Non-Malleable Codes

We now construct a non-malleable code for any family \mathcal{F} of sufficiently small size. We will rely on leakage-resilience as an integral part of the analysis.

Construction. Let \mathcal{H}_1 be a family of hash functions $h_1 : \{0, 1\}^{v_1} \rightarrow \{0, 1\}^k$, and \mathcal{H}_2 be a family of hash functions $h_2 : \{0, 1\}^{k+v_1} \rightarrow \{0, 1\}^{v_2}$ such that \mathcal{H}_1 and \mathcal{H}_2 are both t -wise independent. For any $(h_1, h_2) \in \mathcal{H}_1 \times \mathcal{H}_2$, define $\text{Enc}_{h_1, h_2}(x) = (r, z, \sigma)$ where $r \leftarrow \{0, 1\}^{v_1}$ is random, $z := x \oplus h_1(r)$ and $\sigma := h_2(r, z)$. The codewords are of size $n := |(r, z, \sigma)| = k + v_1 + v_2$. Correspondingly define $\text{Dec}((r, z, \sigma))$ which first checks $\sigma \stackrel{?}{=} h_2(r, z)$ and if this fails, outputs \perp , else outputs $z \oplus h_1(r)$. Notice that, we can think of (r, z) as being a leakage-resilient encoding of x ; i.e., $(r, z) = \text{LREnc}_{h_1}(x; r)$.

Theorem 2. *For any function family \mathcal{F} , the above construction $(\text{Enc}_{h_1, h_2}, \text{Dec}_{h_1, h_2})$ is an $(\mathcal{F}, \varepsilon)$ -super non-malleable code with probability $1 - \rho$ over the choice of h_1, h_2 as long as:*

$$\begin{aligned} t &\geq t^* \quad \text{for some } t^* = O(\log |\mathcal{F}| + n + \log(1/\rho)) \\ v_1 &> v_1^* \quad \text{for some } v_1^* = 3 \log(1/\varepsilon) + 3 \log(t^*) + O(1) \\ v_2 &> v_1 + 3. \end{aligned}$$

For example, in the above theorem, if we set $\rho = \varepsilon = 2^{-\lambda}$ for “security parameter” λ , and $|\mathcal{F}| = 2^{s(n)}$ for some polynomial $s(n) = n^{O(1)} \geq n \geq \lambda$, then we can set $t = O(s(n))$ and the message length $k := n - (v_1 + v_2) = n - O(\lambda + \log n)$. Therefore the rate of the code k/n is $1 - O(\lambda + \log n)/n$ which approaches 1 as n grows relative to λ .

4.1 Proof of Theorem 2

Useful Notions. For a coding scheme (Enc, Dec) , we say that $c \in \{0, 1\}^n$ is *valid* if $\text{Dec}(c) \neq \perp$. For any function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$, we say that $c' \in \{0, 1\}^n$ is δ -heavy for f if $\Pr[f(\text{Enc}(U_k)) = c'] \geq \delta$. Define

$$H_f(\delta) = \{c' \in \{0, 1\}^n : c' \text{ is } \delta\text{-heavy for } f\}.$$

Notice that $|H_f(\delta)| \leq 1/\delta$.

Definition 5 (Bounded-malleable). *We say that a coding scheme (Enc, Dec) is $(\mathcal{F}, \delta, \tau)$ -bounded-malleable if for all $f \in \mathcal{F}, x \in \{0, 1\}^k$ we have*

$$\Pr[c' \neq c \wedge c' \text{ is valid} \wedge c' \notin H_f(\delta) \mid c \leftarrow \text{Enc}(x), c' = f(c)] \leq \tau,$$

where the probability is over the randomness of the encoding.

Intuition. The above definition says the following. Take any message $x \in \{0, 1\}^k$, tampering function $f \in \mathcal{F}$ and do the following: choose $c \leftarrow \text{Enc}(x)$, set $c' = f(c)$, and output: (1) **same*** if $c' = c$, (2) \perp if c' is not valid, (3) c' otherwise. Then, with probability $1 - \tau$ the output of the above experiment takes on one of the values: $\{\text{same}^*, \perp\} \cup H_f(\delta)$. Therefore, the output of the above tampering experiment only leaks a bounded amount of information about c ; in particular it leaks at most $\ell = \lceil \log(1/\delta + 2) \rceil$ bits. Furthermore the “leakage” on c is independent of the choice of the code, up to knowing which codewords are valid and which are δ -heavy. In particular, in our construction, the “leakage” only depends on the choice of h_2 but *not* on the choice of h_1 . This will allow us to then rely on the fact that $\text{LREnc}_{h_1}(x; r) = (r, h_1(r) \oplus x)$ is a leakage-resilient encoding of x to argue that the output of the above experiment is the same for x as for a uniformly random value. We formalize this intuition below.

From Bounded-Malleable to Non-Malleable. For any “tampering function” family \mathcal{F} consisting of functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$, any $\delta > 0$, and any $h_2 \in \mathcal{H}_2$ we define the “leakage function” family $\mathcal{G} = \mathcal{G}(\mathcal{F}, h_2, \delta)$ which consists of the functions $g_f : \{0, 1\}^{k+v_1} \rightarrow H_f(\delta) \cup \{\text{same}^*, \perp\}$ for each $f \in \mathcal{F}$. The functions are defined as follows:

- $g_f(c_1)$: Compute $\sigma = h_2(c_1)$. Let $c := (c_1, \sigma), c' = f(c)$. If c' is not valid output \perp . Else if $c' = c$ output **same***. Else if $c' \in H_f(\delta)$ output c' . Lastly, if none of the above cases holds, output \perp .

Notice that the notion of “ δ -heavy” and the set $H_f(\delta)$ are completely specified by h_2 and do not depend on h_1 . This is because the distribution $\text{Enc}_{h_1, h_2}(U_k)$ is equivalent to $(U_{k+v_1}, h_2(U_{k+v_1}))$ and therefore c' is δ -heavy if and only if $\Pr[f(U_{k+v_1}, h_2(U_{k+v_1})) = c'] \geq \delta$. Therefore the family $\mathcal{G} = \mathcal{G}(\mathcal{F}, h_2, \delta)$ is fully specified by \mathcal{F}, h_2, δ . Also notice that $|\mathcal{G}| = |\mathcal{F}|$ and that the output length of the functions g_f is given by $\ell = \lceil \log(|H_f(\delta)| + 2) \rceil \leq \lceil \log(1/\delta + 2) \rceil$.

Lemma 1. *Let \mathcal{F} be any function family and let $\delta > 0$. Fix any h_1, h_2 such that $(\text{Enc}_{h_1, h_2}, \text{Dec}_{h_1, h_2})$ is $(\mathcal{F}, \delta, \varepsilon/4)$ -bounded-malleable and $(\text{LREnc}_{h_1}, \text{LRDec}_{h_1})$ is $(\mathcal{G}(\mathcal{F}, h_2, \delta), \varepsilon/4)$ -leakage-resilient, where the family $\mathcal{G} = \mathcal{G}(\mathcal{F}, h_2, \delta)$, with size $|\mathcal{G}| = |\mathcal{F}|$, is defined above, and the leakage amount is $\ell = \lceil \log(1/\delta + 2) \rceil$. Then $(\text{Enc}_{h_1, h_2}, \text{Dec}_{h_1, h_2})$ is $(\mathcal{F}, \varepsilon)$ -non-malleable.*

Proof. For any $x_0, x_1 \in \{0, 1\}^k$ and any $f \in \mathcal{F}$:

$$\begin{aligned} \text{Tamper}_{x_0}^f &= \left\{ \begin{array}{l} c \leftarrow \text{Enc}_{h_1, h_2}(x_0), c' := f(c) \\ \text{Output} : \text{same}^* \text{ if } c' = c, \perp \text{ if } \text{Dec}_{h_1, h_2}(c') = \perp, \\ c' \text{ otherwise.} \end{array} \right\} \\ &\stackrel{\text{stat}}{\approx}_{\varepsilon/4} \left\{ \begin{array}{l} c_1 \leftarrow \text{LREnc}_{h_1}(x_0) \\ \text{Output} : g_f(c_1) \end{array} \right\} \quad (3) \\ &\stackrel{\text{stat}}{\approx}_{\varepsilon/4} \left\{ \begin{array}{l} c_1 \leftarrow \text{LREnc}_{h_1}(U_k) \\ \text{Output} : g_f(c_1) \end{array} \right\} \quad (4) \end{aligned}$$

$$\stackrel{\text{stat}}{\approx}_{\varepsilon/4} \left\{ \begin{array}{l} c_1 \leftarrow \text{LREnc}_{h_1}(x_1) \\ \text{Output} : g_f(c_1) \end{array} \right\} \quad (5)$$

$$\stackrel{\text{stat}}{\approx}_{\varepsilon/4} \left\{ \begin{array}{l} c \leftarrow \text{Enc}_{h_1, h_2}(x_1), c' := f(c) \\ \text{Output} : \text{same}^* \text{ if } c' = c, \perp \text{ if } \text{Dec}_{h_1, h_2}(c') = \perp, \\ c' \text{ otherwise.} \end{array} \right\} \quad (6)$$

$$= \text{Tamper}_{x_1}^f$$

Eq. (3) and Eq. (6) follows as $(\text{Enc}_{h_1, h_2}, \text{Dec}_{h_1, h_2})$ is an $(\mathcal{F}, \delta, \varepsilon/4)$ -bounded-malleable code, and Eq. (4) and Eq. (5) follow as the code $(\text{LREnc}_{h_1}, \text{LRDec}_{h_1})$ is $(\mathcal{G}(\mathcal{F}, \delta), \varepsilon/4)$ -leakage-resilient.

We can use Theorem 1 to show that $(\text{LREnc}_{h_1}, \text{LRDec}_{h_1})$ is $(\mathcal{G}(\mathcal{F}, h_2, \delta), \varepsilon/4)$ -leakage-resilient with overwhelming probability. Therefore, it remains to show that our construction is $(\mathcal{F}, \delta, \tau)$ -bounded-malleable, which we do below.

Analysis of Bounded-Malleable Codes. We now show that the code $(\text{Enc}_{h_1, h_2}, \text{Dec}_{h_1, h_2})$ is bounded-malleable with overwhelming probability. As a very high-level intuition, if a tampering function f can often map valid codewords to other valid codewords (and many different ones), then it must guess the output of h_2 on many different inputs. If the family \mathcal{F} is small enough, it is highly improbable that it would contain some such f . For more detailed intuition, we show that the following two properties hold for any message x and any function f with overwhelming probability: (1) there is at most some “small” set of q valid codewords c' that we can hit by tampering some encoding of x via f (2) for each such codeword c' which is not in δ -heavy, the probability of landing in c' after tampering an encoding of x cannot be higher than 2δ . This shows that the total probability of tampering an encoding of x and landing in a valid codeword which not δ -heavy is at most $2q\delta$, which is small. Property (1) roughly follows by showing that f would need to “predict” the output of h_2 on q different inputs, and property (2) follows by using “leakage-resilience” of h_1 to argue that we cannot distinguish an encoding of x from an encoding of a random message, for which the probability of landing in c' is at most δ .

Lemma 2. *For any function family \mathcal{F} , any $\delta > 0$, the code $(\text{Enc}_{h_1, h_2}, \text{Dec}_{h_1, h_2})$ is $(\mathcal{F}, \delta, \tau)$ -bounded-malleable with probability $1 - \psi$ over the choice of h_1, h_2 as long as:*

$$\begin{aligned} \tau &\geq 2(\log |\mathcal{F}| + k + \log(1/\psi) + 2)\delta \\ t &\geq \log |\mathcal{F}| + n + k + \log(1/\psi) + 5 \\ v_1 &\geq 2 \log(1/\delta) + \log(t) + 4 \quad \text{and} \quad v_2 \geq v_1 + 3. \end{aligned}$$

Proof. Set $q := \lceil \log |\mathcal{F}| + k + \log(1/\psi) + 1 \rceil$. For any $f \in \mathcal{F}, x \in \{0, 1\}^k$ define the events $E_1^{f, x}$ and $E_2^{f, x}$ over the random choice of h_1, h_2 as follows:

1. $E_1^{f, x}$ occurs if there exist at least q distinct values $c'_1, \dots, c'_q \in \{0, 1\}^n$ such that each c'_i is valid and $c'_i = f(c_i)$ for some $c_i \neq c'_i$ which encodes the message x (i.e., $c_i = \text{Enc}_{h_1, h_2}(x; r_i)$ for some r_i).

2. $E_2^{f,x}$ occurs if there exists some $c' \in \{0, 1\}^n \setminus H_f(\delta)$ such that

$$\Pr_{r \leftarrow \{0,1\}^{v_1}} [f(\text{Enc}_{h_1, h_2}(x; r)) = c'] \geq 2\delta.$$

Let $E_1 = \bigvee_{f,x} E_1^{f,x}$, $E_2 = \bigvee_{f,x} E_2^{f,x}$ and $\text{BAD} = E_1 \vee E_2$. Assume (h_1, h_2) are any hash functions for which the event BAD does *not* occur. Then, for every $f \in \mathcal{F}, x \in \{0, 1\}^k$:

$$\begin{aligned} & \Pr[f(C) \neq C \wedge f(C) \text{ is valid} \wedge f(C) \notin H_f(\delta)] \\ &= \sum_{c': c' \text{ valid and } c' \notin H_f(\delta)} \Pr[f(C) = c' \wedge C \neq c'] < 2q\delta \leq \tau, \end{aligned} \quad (7)$$

where $C = \text{Enc}_{h_1, h_2}(x; U_{v_1})$ is a random variable. Eq. (7) holds since (1) given that E_1 does not occur, there are fewer than q values c' that are valid and for which $\Pr[f(C) = c' \wedge C \neq c'] > 0$, and (2) given that E_2 does not occur, for any $c' \notin H_f(\delta)$, we also have $\Pr[f(C) = c' \wedge C \neq c'] \leq \Pr[f(C) = c'] < 2\delta$.

Therefore, if the event BAD does not occur, then the code is $(\mathcal{F}, \delta, \tau)$ -bounded-malleable. This means:

$$\begin{aligned} & \Pr_{h_1, h_2} [(\text{Enc}_{h_1, h_2}, \text{Dec}_{h_1, h_2}) \text{ is not } (\mathcal{F}, \delta, \tau)\text{-bounded-malleable}] \\ & \leq \Pr[\text{BAD}] \leq \Pr[E_1] + \Pr[E_2] \end{aligned}$$

So it suffices to show that $\Pr[E_1]$ and $\Pr[E_2]$ are both bounded by $\psi/2$, which we do next.

Claim. $\Pr[E_1] \leq \psi/2$.

Proof. Fix some message $x \in \{0, 1\}^k$ and some function $f \in \mathcal{F}$. Assume that the event $E_1^{f,x}$ occurs for some choice of hash functions (h_1, h_2) . Then there must exist some values $\{r_1, \dots, r_q\}$ such that: if we define $c_i := \text{Enc}(x; r_i), c'_i := f(c_i)$ then $c'_i \neq c_i, c'_i$ is valid, and $|\{c'_1, \dots, c'_q\}| = q$. The last condition also implies $|\{c_1, \dots, c_q\}| = q$. However, it is possible that $c_i = c'_j$ for some $i \neq j$. We claim that we can find a subset of at least $s := \lceil q/3 \rceil$ of the indices such that the $2s$ values $\{c_{a_1}, \dots, c_{a_s}, c'_{a_1}, \dots, c'_{a_s}\}$ are all distinct. To do so, notice that if we want to keep some index i corresponding to values c_i, c'_i , we need to take out at most two indices j, k if $c'_j = c_i$ or $c_k = c'_i$.⁵ To summarize, if $E_1^{f,x}$ occurs, then (by re-indexing) there is some set $R = \{r_1, \dots, r_s\} \subseteq \{0, 1\}^{v_1}$ of size $|R| = s$ satisfying the following two conditions:

- (1) If we define $c_i := \text{Enc}(x; r_i), c'_i \neq c_i$ and c'_i is valid meaning that $c'_i = (r'_i, z'_i, \sigma'_i)$ where $\sigma' = h_2(r', z')$.

⁵ In other words, if we take any set of tuples $\{(c_i, c'_i)\}$ such that all the left components are distinct $c_i \neq c_j$ and all the right components are distinct $c'_i \neq c'_j$, but there may be common values $c_i = c'_j$, then there is a subset of at least 1/3 of the tuples such that all left and right components in this subset are mutually distinct.

$$(2) \quad |\{c_1, \dots, c_s, c'_1, \dots, c'_s\}| = 2s.$$

Therefore we have:

$$\begin{aligned} \Pr[E_1^{f,x}] &\leq \Pr_{h_1, h_2} [\exists R \subseteq \{0, 1\}^{v_1}, |R| = s, R \text{ satisfies (1) and (2)}] \\ &\leq \sum_R \Pr_{h_1, h_2} [R \text{ satisfies (1) and (2)}] \\ &\leq \sum_{R=\{r_1, \dots, r_s\}} \max_{h_1, \sigma_1, \dots, \sigma_s} \Pr_{h_2} \left[\forall i, c'_i \text{ valid} \left| \begin{array}{l} c_i := (r_i, z_i = h_1(r_i) \oplus x, \sigma_i), \\ c'_i := f(c_i), c'_i \neq c_i \\ |\{c_1, \dots, c_s, c'_1, \dots, c'_s\}| = 2s \end{array} \right. \right] \\ &\leq \binom{2^{v_1}}{s} 2^{-sv_2} \leq \left(\frac{e2^{v_1}}{s} \right)^s 2^{-sv_2} \leq 2^{s(v_1-v_2)} \leq 2^{q(v_1-v_2)/3} \leq 2^{-q}, \quad (8) \end{aligned}$$

where Eq. (8) follows from the fact that, even if we condition on any choice of the hash function h_1 which fixes $z_i = h_1(r_i) \oplus x$, and any choice of the s values $\sigma_i = h_2(r_i, z_i)$, which fixes $c_i := (r_i, z_i = h_1(r_i) \oplus x, \sigma_i)$, $c'_i := f(c_i)$ such that $c'_i \neq c_i$ and $|\{c_1, \dots, c_s, c'_1, \dots, c'_s\}| = 2s$, then the probability that $h_2(r'_i, z'_i) = \sigma'_i$ for all $i \in [s]$ is at most 2^{-sv_2} . Here we use the fact that \mathcal{H}_2 is t -wise independent where $t \geq q \geq 2s$. Now, we calculate

$$\Pr[E_1] \leq \sum_{f \in \mathcal{F}} \sum_{x \in \{0,1\}^k} \Pr[E_1^{f,x}] \leq |\mathcal{F}| 2^{k-q} \leq \psi/2,$$

where the last inequality follows from the assumption, $q = \lceil \log |\mathcal{F}| + k + \log(1/\psi) + 1 \rceil$.

Claim. $\Pr[E_2] \leq \psi/2$.

Proof. For this proof, we will rely on the leakage-resilience property of the code $(\text{LREnc}_{h_1}, \text{LRDec}_{h_1})$ as shown in Theorem 1. First, let us write:

$$\begin{aligned} \Pr[E_2] &= \Pr_{h_1, h_2} \left[\exists (f, x, c') \in \mathcal{F} \times \{0, 1\}^k \times \{0, 1\}^n \setminus H_f(\delta) : \right. \\ &\quad \left. \Pr[f(\text{Enc}_{h_1, h_2}(x; U_{v_1})) = c'] \geq 2\delta \right] \\ &\leq \Pr_{h_1, h_2} \left[\exists (f, x, c') \in \mathcal{F} \times \{0, 1\}^k \times \{0, 1\}^n \setminus H_f(\delta) : \right. \\ &\quad \left. \left| \begin{array}{l} \Pr[f(\text{Enc}_{h_1, h_2}(x; U_{v_1})) = c'] \\ - \Pr[f(\text{Enc}_{h_1, h_2}(U_k; U_{v_1})) = c'] \end{array} \right| \geq \delta \right] \end{aligned} \quad (9)$$

since, for any $c' \notin H_f(\delta)$, we have $\Pr[f(\text{Enc}_{h_1, h_2}(U_k; U_{v_1})) = c'] < \delta$ by definition. Notice that we can write $\text{Enc}_{h_1, h_2}(x; r) = (c_1, c_2)$ where $c_1 = \text{LREnc}_{h_1}(x; r)$, $c_2 = h_2(c_1)$. We will now rely on the leakage-resilience of the code $(\text{LREnc}_{h_1}, \text{LRDec}_{h_1})$ to bound the above probability by $\psi/2$. In fact, we show that the above holds even if we take the probability over h_1 only, for a worst-case choice of h_2 .

Let us fix some choice of h_2 and define the family $\mathcal{G} = \mathcal{G}(h_2)$ of leakage-functions $\mathcal{G} = \{g_{f,c'} : \{0,1\}^{k+v_1} \rightarrow \{0,1\} \mid f \in \mathcal{F}, c' \in \{0,1\}^n\}$ with output size $\ell = 1$ bits as follows:

- $g_{f,c'}(c_1)$: Set $c = (c_1, c_2 = h_2(c_1))$. If $f(c) = c'$ output 1, else output 0.

Notice that the size of the family \mathcal{G} is $2^n|\mathcal{F}|$ and the family does not depend on the choice of h_1 . Therefore, continuing from inequality (9), we get:

$$\begin{aligned}
\Pr[E_2] &\leq \Pr_{h_1, h_2} \left[\exists (f, x, c') \in \mathcal{F} \times \{0,1\}^k \times \{0,1\}^n \setminus H_f(\delta) : \right. \\
&\quad \left. \left| \begin{array}{l} \Pr[f(\text{Enc}_{h_1, h_2}(x; U_{v_1})) = c'] \\ - \Pr[f(\text{Enc}_{h_1, h_2}(U_k; U_{v_1})) = c'] \end{array} \right| \geq \delta \right] \\
&\leq \max_{h_2} \Pr_{h_1} \left[\exists (g_{f,c'}, x) \in \mathcal{G}(h_2) \times \{0,1\}^k : \right. \\
&\quad \left. \left| \begin{array}{l} \Pr[g_{f,c'}(\text{LEnc}_{h_1}(x; U_{v_1})) = 1] \\ - \Pr[g_{f,c'}(\text{LEnc}_{h_1}(U_k; U_{v_1})) = 1] \end{array} \right| \geq \delta \right] \\
&= \max_{h_2} \Pr_{h_1} \left[\exists (g_{f,c'}, x) \in \mathcal{G}(h_2) \times \{0,1\}^k : \right. \\
&\quad \left. \left| \begin{array}{l} \Pr[g_{f,c'}(\text{LEnc}_{h_1}(x; U_{v_1})) = 1] \\ - \Pr[g_{f,c'}(U_{k+v_1}) = 1] \end{array} \right| \geq \delta \right] \\
&\leq \max_{h_2} \Pr_{h_1} [(\text{LEnc}_{h_1}, \text{LRDec}_{h_1}) \text{ is not } (\mathcal{G}(h_2), \delta)\text{-Leakage-Resilient}] \\
&\leq \psi/2,
\end{aligned}$$

where the last inequality follows from Theorem 1 by the choice of parameters.

Putting it All Together. Lemma 1 tells us that for any $\delta > 0$ and any function family \mathcal{F} :

$$\begin{aligned}
&\Pr[(\text{Enc}_{h_1, h_2}, \text{Dec}_{h_1, h_2}) \text{ is not } (\mathcal{F}, \varepsilon)\text{-super-non-malleable}] \\
&\leq \Pr[(\text{Enc}_{h_1, h_2}, \text{Dec}_{h_1, h_2}) \text{ is not } (\mathcal{F}, \delta, \varepsilon/4)\text{-bounded-malleable}] \quad (10) \\
&\quad + \Pr[(\text{LEnc}_{h_1}, \text{LRDec}_{h_1}) \text{ is not } (\mathcal{G}(\mathcal{F}, h_2, \delta), \varepsilon/4)\text{-leakage-resilient}], \quad (11)
\end{aligned}$$

where $\mathcal{G} = \mathcal{G}(\mathcal{F}, h_2, \delta)$ is of size $|\mathcal{G}| = |\mathcal{F}|$ and consists of function with output size $\ell = \lceil \log(1/\delta + 2) \rceil$.

Let us set $\delta := (\varepsilon/8)(\log |\mathcal{F}| + k + \log(1/\rho) + 3)^{-1}$. This ensures that the first requirement of Lemma 2 is satisfied with $\tau = \varepsilon/4$. We choose $t^* = O(\log |\mathcal{F}| + n + \log(1/\rho))$ such that $\log(1/\delta) \leq \log(1/\varepsilon) + \log(t^*) + O(1)$. Notice that the leakage amount of \mathcal{G} is $\ell = \lceil \log(1/\delta + 2) \rceil \leq \log(1/\varepsilon) + \log(t^*) + O(1)$. With v_1, v_2 as in Theorem 2, we satisfy the remaining requirements of Lemma 2 (bounded-malleable codes) and Theorem 1 (leakage-resilient codes) to ensure that the probabilities (10), (11) are both bounded by $\rho/2$, which proves our theorem.

Experiment $\text{Real}_h(f)$ vs. $\text{Sim}_h(f)$	
Experiment $\text{Real}_h(f)$: Sample $x \leftarrow U_n$. If $f(x) = x$: Output $(h(x), \text{same}^*)$. Else Output $(h(x), h(f(x)))$.	Experiment $\text{Sim}_h(f)$: Sample $x \leftarrow U_n; y \leftarrow U_k$ If $f(x) = x$: Output (y, same^*) . Else Output $(y, h(f(x)))$.

Fig. 1. Experiments defining a non-malleable key derivation function h

5 Non-malleable Key-derivation

In this section we introduce a new primitive, which we name non-malleable key derivation. Intuitively a function h is a non-malleable key derivation function if $h(x)$ is close to uniform even given the output of h applied to a related input $f(x)$, as long as $f(x) \neq x$.

Definition 6 (Non-malleable Key-Derivation). *Let \mathcal{F} be any family of functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$. We say that a function $h : \{0, 1\}^n \rightarrow \{0, 1\}^k$ is an $(\mathcal{F}, \varepsilon)$ -non-malleable key derivation function if for every $f \in \mathcal{F}$ we have $\text{SD}(\text{Real}_h(f); \text{Sim}_h(f)) \leq \varepsilon$ where $\text{Real}_h(f)$ and $\text{Sim}_h(f)$ denote the output distributions of the corresponding experiments described in Fig. 1.*

Note that the above definition can be interpreted as a dual version of the definition of non-malleable extractors [15].⁶ The theorem below states that by sampling a function h from a set \mathcal{H} of t -wise independent hash functions, we obtain a non-malleable key derivation function with overwhelming probability.

Theorem 3. *Let \mathcal{H} be a $2t$ -wise independent function family consisting of functions $h : \{0, 1\}^n \rightarrow \{0, 1\}^k$ and let \mathcal{F} be some function family as above. Then with probability $1 - \rho$ over the choice of a random $h \leftarrow \mathcal{H}$, the function h is an $(\mathcal{F}, \varepsilon)$ -non-malleable key-derivation function as long as:*

$$n \geq 2k + \log(1/\varepsilon) + \log(t) + 3 \quad \text{and} \quad t > \log(|\mathcal{F}|) + 2k + \log(1/\rho) + 5.$$

Proof. For any $h \in \mathcal{H}$ and $f \in \mathcal{F}$, define a function $h_f : \{0, 1\}^n \rightarrow \{0, 1\}^k \cup \text{same}^*$ such that if $f(x) = x$ then $h_f(x) = \text{same}^*$ otherwise $h_f(x) = h(f(x))$. Fix a function family \mathcal{F} . Now, taking probabilities (only) over the choice of h , let BAD be the event that h is not an $(\mathcal{F}, \varepsilon)$ -non-malleable-key-derivation function. Then:

$$\Pr[\text{BAD}] = \Pr_{h \leftarrow \mathcal{H}} \left[\exists f \in \mathcal{F} : \text{SD}(\text{Real}_h(f), \text{Sim}_h(f)) > \varepsilon \right]$$

⁶ The duality comes from the fact that the output of a non-malleable extractor is close to uniform even given a certain number of outputs computed with related seeds (whereas for non-malleable key derivation the seed is unchanged but the input can be altered).

$$\begin{aligned}
&= \Pr_{h \leftarrow \mathcal{H}} \left[\exists f \in \mathcal{F} : \mathbf{SD}((h(X), h_f(X)), (U_k, h_f(X))) > \varepsilon \right] \\
&\leq \sum_{f \in \mathcal{F}} \Pr_{h \leftarrow \mathcal{H}} \left[\sum_{y \in \{0,1\}^k} \sum_{y' \in \{0,1\}^k \cup \text{same}^*} \left| \Pr[h(X) = y \wedge h_f(X) = y'] \right. \right. \\
&\quad \left. \left. - \Pr[U_k = y \wedge h_f(X) = y'] \right| > 2\varepsilon \right] \\
&\leq \sum_{f \in \mathcal{F}} \Pr_{h \leftarrow \mathcal{H}} \left[\exists y \in \{0,1\}^k, y' \in \{0,1\}^k \cup \text{same}^* : \right. \\
&\quad \left. \left| \Pr[h(X) = y \wedge h_f(X) = y'] \right. \right. \\
&\quad \left. \left. - \Pr[U_k = y \wedge h_f(X) = y'] \right| > 2^{-2k} \varepsilon \right] \\
&\leq \sum_{f \in \mathcal{F}} \sum_{y \in \{0,1\}^k} \sum_{y' \in \{0,1\}^k \cup \text{same}^*} \Pr_{h \leftarrow \mathcal{H}} \left[\left| \Pr[h(X) = y \wedge h_f(X) = y'] \right. \right. \\
&\quad \left. \left. - 2^{-k} \Pr[h_f(X) = y'] \right| > 2^{-2k} \varepsilon \right] \quad (12)
\end{aligned}$$

Fix f, y, y' . For every $x \in \{0,1\}^n$, define a random variable C_x over the choice of $h \leftarrow \mathcal{H}$, such that

$$C_x = \begin{cases} 1 - 2^{-k} & \text{if } h(x) = y \wedge h_f(x) = y' \\ -2^{-k} & \text{if } h(x) \neq y \wedge h_f(x) = y' \\ 0 & \text{otherwise.} \end{cases}$$

Notice that each C_x is 0 on expectation. However, the random variables C_x are not even pairwise independent.⁷ In the full version of this paper [20], we prove the following lemma about the variables C_x .

Lemma 3. *There exists a partitioning of $\{0,1\}^n$ into four disjoint subsets $\{A_j\}_{j=1}^4$, such that for any $A > 0$ and for all $j = 1, \dots, 4$:*

$$\Pr \left[\left| \sum_{x \in A_j} C_x \right| > A \right] < K_t \left(\frac{t}{A} \right)^t,$$

where $K_t \leq 8$.

Continuing from Eq. (12), we get:

$$\begin{aligned}
&\Pr_{h \leftarrow \mathcal{H}} \left[\left| \Pr[h(X) = y \wedge h_f(X) = y'] - 2^{-k} \Pr[h_f(X) = y'] \right| > 2^{-2k} \varepsilon \right] \\
&= \Pr_{h \leftarrow \mathcal{H}} \left[\left| \sum_{x \in \{0,1\}^n} C_x \right| > 2^{n-2k} \varepsilon \right] \quad (13)
\end{aligned}$$

$$\leq \sum_{j=1}^4 \Pr_{h \leftarrow \mathcal{H}} \left[\left| \sum_{x \in A_j} C_x \right| > 2^{n-2k-2} \varepsilon \right] < 4K_t \left(\frac{t}{2^{n-2k-2} \varepsilon} \right)^t. \quad (14)$$

⁷ For example if $f(x) = f(x')$ and $C_x = 0$ then $C_{x'} = 0$ as well.

Eq. (13) follows from the definitions of the variables C_x and Eq. (14) follows by applying Lemma 3 to the sum. Combining Eq. (12) and Eq. (14), we get $\Pr[\text{BAD}] < |\mathcal{F}|2^{2k} \left[4K_t \left(\frac{t}{2^{n-2k-2\varepsilon}} \right)^t \right]$. In particular, it holds that $\Pr[\text{BAD}] \leq \rho$ as long as:

$$n \geq 2k + \log(1/\varepsilon) + \log(t) + 3 \quad \text{and} \quad t > \log(|\mathcal{F}|) + 2k + \log(1/\rho) + 5.$$

Optimal Rate of Non-Malleable Key-Derivation. We can define the rate of a key derivation function $h : \{0, 1\}^n \rightarrow \{0, 1\}^k$ as the ratio k/n . Notice that our construction achieves rate arbitrary close to $1/2$. We claim that this is *optimal* for non-malleable key derivation. To see this, consider a tampering function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ which is a permutation and never identity: $f(x) \neq x$. In this case the joint distribution $(h(X), h(f(X)))$ is ε -close to $(U_k, h(f(X)))$ which is ε -close to the distribution (U_k, U'_k) consisting of $2k$ random bits. Since all of the randomness in $(h(X), h(f(X)))$ comes from X , this means that X must contain at least $2k$ bits of randomness, meaning that $n > 2k$.

Acknowledgements

We thank Ivan Damgård for useful discussions at the early stages of this research.

References

1. Divesh Aggarwal, Yevgeniy Dodis, and Shachar Lovett. Non-malleable codes from additive combinatorics. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:81, 2013.
2. Ross Anderson, Markus Kuhn, and England U. S. A. Tamper resistance — a cautionary note. In *In Proceedings of the Second Usenix Workshop on Electronic Commerce*, pages 1–11, 1996.
3. Benny Applebaum, Danny Harnik, and Yuval Ishai. Semantic security under related-key attacks and applications. In *ICS*, pages 45–60, 2011.
4. Mihir Bellare and David Cash. Pseudorandom functions and permutations provably secure against related-key attacks. In *CRYPTO*, pages 666–684, 2010.
5. Mihir Bellare and Tadayoshi Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In *EUROCRYPT*, pages 491–506, 2003.
6. Mihir Bellare, Kenneth G. Paterson, and Susan Thomson. RKA security beyond the linear barrier: IBE, encryption and signatures. In *ASIACRYPT*, pages 331–348, 2012.
7. Dan Boneh, Richard A. DeMillo, and Richard J. Lipton. On the importance of eliminating errors in cryptographic computations. *J. Cryptology*, 14(2):101–119, 2001.
8. Mahdi Cheraghchi and Venkatesan Guruswami. Capacity of non-malleable codes. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:118, 2013.
9. Mahdi Cheraghchi and Venkatesan Guruswami. Non-malleable coding against bit-wise and split-state tampering. *IACR Cryptology ePrint Archive*, 2013:565, 2013.

10. Seung Geol Choi, Aggelos Kiayias, and Tal Malkin. BiTR: Built-in tamper resilience. In *ASIACRYPT*, pages 740–758, 2011.
11. Jean-Sébastien Coron, Antoine Joux, Ilya Kizhvatov, David Naccache, and Pascal Paillier. Fault attacks on rsa signatures with partially unknown messages. In Christophe Clavier and Kris Gaj, editors, *CHES*, volume 5747 of *Lecture Notes in Computer Science*, pages 444–456. Springer, 2009.
12. Ronald Cramer, Yevgeniy Dodis, Serge Fehr, Carles Padró, and Daniel Wichs. Detection of algebraic manipulation with applications to robust secret sharing and fuzzy extractors. In Nigel P. Smart, editor, *EUROCRYPT*, volume 4965 of *Lecture Notes in Computer Science*, pages 471–488. Springer, 2008.
13. Ivan Damgård, Sebastian Faust, Pratyay Mukherjee, and Daniele Venturi. Bounded tamper resilience: How to go beyond the algebraic barrier. In *ASIACRYPT (2)*, pages 140–160, 2013.
14. Francesco Davì, Stefan Dziembowski, and Daniele Venturi. Leakage-resilient storage. In Juan A. Garay and Roberto De Prisco, editors, *SCN*, volume 6280 of *Lecture Notes in Computer Science*, pages 121–137. Springer, 2010.
15. Yevgeniy Dodis and Daniel Wichs. Non-malleable extractors and symmetric key cryptography from weak secrets. In *STOC*, pages 601–610, 2009.
16. Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM J. Comput.*, 30(2):391–437, 2000.
17. Stefan Dziembowski, Tomasz Kazana, and Maciej Obremski. Non-malleable codes from two-source extractors. In *CRYPTO (2)*, pages 239–257, 2013.
18. Stefan Dziembowski, Krzysztof Pietrzak, and Daniel Wichs. Non-malleable codes. In *ICS*, pages 434–452, 2010.
19. Sebastian Faust, Pratyay Mukherjee, Jesper Buus Nielsen, and Daniele Venturi. Continuous non-malleable codes. In *TCC*, 2014. To appear.
20. Sebastian Faust, Pratyay Mukherjee, Daniele Venturi, and Daniel Wichs. Efficient non-malleable codes and key-derivation for poly-size tampering circuits. *IACR Cryptology ePrint Archive*, 2013:702, 2013.
21. Sebastian Faust, Krzysztof Pietrzak, and Daniele Venturi. Tamper-proof circuits: How to trade leakage for tamper-resilience. In *ICALP (1)*, pages 391–402, 2011.
22. Rosario Gennaro, Anna Lysyanskaya, Tal Malkin, Silvio Micali, and Tal Rabin. Algorithmic tamper-proof (ATP) security: Theoretical foundations for security against hardware tampering. In *TCC*, pages 258–277, 2004.
23. Vipul Goyal, Adam O’Neill, and Vanishree Rao. Correlated-input secure hash functions. In *TCC*, pages 182–200, 2011.
24. Yuval Ishai, Manoj Prabhakaran, Amit Sahai, and David Wagner. Private circuits II: Keeping secrets in tamperable circuits. In *EUROCRYPT*, pages 308–327, 2006.
25. Yael Tauman Kalai, Bhavana Kanukurthi, and Amit Sahai. Cryptography with tamperable and leaky memory. In *CRYPTO*, pages 373–390, 2011.
26. Feng-Hao Liu and Anna Lysyanskaya. Tamper and leakage resilience in the split-state model. In *CRYPTO*, pages 517–532, 2012.
27. Krzysztof Pietrzak. Subspace LWE. In *TCC*, pages 548–563, 2012.
28. Ananth Raghunathan, Gil Segev, and Salil P. Vadhan. Deterministic public-key encryption for adaptively chosen plaintext distributions. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT*, volume 7881 of *Lecture Notes in Computer Science*, pages 93–110. Springer, 2013.
29. Sergei P. Skorobogatov and Ross J. Anderson. Optical fault induction attacks. pages 2–12. Springer-Verlag, 2002.
30. Hoeteck Wee. Public key encryption against related key attacks. In *Public Key Cryptography*, pages 262–279, 2012.