

Non-malleable Codes from Additive Combinatorics

Divesh Aggarwal*

Yevgeniy Dodis†

Shachar Lovett‡

May 13, 2013

Abstract

Non-malleable codes provide a useful and meaningful security guarantee in situations where traditional error-correction (and even error-detection) is impossible; for example, when the attacker can completely overwrite the encoded message. Informally, a code is non-malleable if the message contained in a modified codeword is either the original message, or a completely unrelated value. Although such codes do not exist if the family of “tampering functions” \mathcal{F} is completely unrestricted, they are known to exist for many broad tampering families \mathcal{F} . One such natural family is the family of tampering functions in the so called *split-state* model. Here the message m is encoded into two shares L and R , and the attacker is allowed to *arbitrarily* tamper with L and R *individually*. The split-state tampering arises in many realistic applications, such as the design of *non-malleable secret sharing schemes*, motivating the question of designing efficient non-malleable codes in this model.

Prior to this work, non-malleable codes in the split-state model received considerable attention in the literature, but were constructed either (1) in the random oracle model [14], or (2) relied on advanced cryptographic assumptions (such as non-interactive zero-knowledge proofs and leakage-resilient encryption) [22], or (3) could only encode 1-bit messages [12]. As our main result, we build the first efficient, multi-bit, information-theoretically-secure non-malleable code in the split-state model.

The heart of our construction uses the following new property of the inner-product function $\langle L, R \rangle$ over the vector space \mathbb{F}_p^n (for a prime p and large enough dimension n): if L and R are uniformly random over \mathbb{F}_p^n , and $f, g : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ are two arbitrary functions on L and R , then the joint distribution $(\langle L, R \rangle, \langle f(L), g(R) \rangle)$ is “close” to the convex combination of “affine distributions” $\{(U, aU + b) \mid a, b \in \mathbb{F}_p\}$, where U is uniformly random in \mathbb{F}_p . In turn, the proof of this surprising property of the inner product function critically relies on some results from additive combinatorics, including the so called *Quasi-polynomial Freiman-Ruzsa Theorem* (which was recently established by Sanders [25] as a step towards resolving the Polynomial Freiman-Ruzsa conjecture [17]).

*Department of Computer Science, New York University. Email: divesha@cs.nyu.edu.

†Department of Computer Science, New York University. Email: dodis@cs.nyu.edu.

‡Department of Computer Science, University of California at San Diego. Email: shachar.lovett@gmail.com.

1 Introduction

The problem of reliable storage/transmission of information is one of the oldest and fundamental problems of information theory. The basic problem can be abstracted as the question of designing an efficient way to encode/decode the message m , so that the resulted codeword $c = \text{Enc}(m)$ is “resilient” against some natural class of error or tampering functions \mathcal{F} . In more detail, one can imagine the attacker can choose an arbitrary (unknown) tampering function $f \in \mathcal{F}$ and modify the real codeword c into a corrupted codeword $c' = f(c)$, and the goal of a good coding scheme (Enc, Dec) is to protect against such tampering attacks. Depending on the richness of the tampering class \mathcal{F} , one can demand various security guarantees from such an encoding.

ERROR-CORRECTING CODES. The most desirable such guarantee would be *error-correction*, which demands that m can be correctly recovered (possibly, with high probability) from c' . This has led to the rich theory of *error-correcting codes*, which provide such error-correction for the natural family of functions \mathcal{F} which flip some (small) subset of the bits (or symbols) of the encoding. Still, as useful and natural error-correcting codes are, in some situations the tampering function $f \in \mathcal{F}$ might either exceed the maximum number of errors for reliable error-correction, or might even touch the entire codeword in some natural yet restricted way (see below). In such settings one must relax the notion of error-correction to some meaningful weaker notion.

ERROR-DETECTING CODES. One such notion is *error-detection*, which guarantees that the decoding of the corrupted codeword $c' = f(c)$ will almost never output some message $m' \neq m$, but is allowed to output a special symbol \perp when it detects some tampering which cannot be corrected reliably. For example, any (deterministic) code capable of correcting d Hamming errors much be able to reliably detect at least $2d$ errors. More interestingly, error-detecting codes allow one to possibly handle useful tampering classes \mathcal{F} where there is no hope for meaningful error-correction. One such class of tampering functions was considered by Cramer et al. [6] and consists of all functions $f_{\Delta}(c) = c + \Delta$ which add a fixed offset Δ to the codeword c in some appropriate group (e.g., such a function can flip every bit c when addition is \oplus). Notice, error-correction is indeed impossible here, since the attacker can simply choose a random offset Δ to completely erase any information about the original message m . More interestingly, although this class might seem somewhat artificial at the first glance, the authors showed that developing error-detecting codes — which they called *algebraic-manipulation detection (AMD) codes* — for this class has useful applications to the design of so called robust secret sharing schemes and robust fuzzy extractors [2, 9]. Finally, unlike error-correction codes, which can be deterministic, AMD codes must be probabilistic, since otherwise the attacker can set $\Delta = c_2 - c_1$ for two valid codewords c_1 and c_2 .

NON-MALLEABLE CODES. Unfortunately, even error-detecting codes are rather limited in some situations, since they cannot protect a natural tampering function $f(c)$ which simply overwrites the codeword c by another fixed (and valid) codeword c^* . This basic attack is quite natural both in the message transmission scenario (where the channel might simply block the original encoded message, and send a different message instead), and in the secure storage scenario (where the attacker might be able to format the hard-drive, for example). Until recently, it was believed that handling such “constant” tampering functions is impossible without having any secrets, and using tools from cryptography (such as signatures or message authentication codes) is essential for preventing more general tampering attacks. Fortunately, Dziembowski, Pietrzak and Wichs [14] recently showed that this belief is overly pessimistic, and introduced a natural and beautiful relaxation of error-detecting codes which they called *non-malleable codes* (with respect to a given family \mathcal{F}). Intuitively, such a non-malleable code ensures that the decoded message $m' = \text{Dec}(f(\text{Enc}(m)))$ is either (a) equal to m (tampering corrected); or (b) equal to \perp (tampering detected); or (c) completely “unrelated” to the

original message m .¹ Moreover, one can figure out which of the scenarios (a)-(c) happens by just looking at the function f (independent of the original message m , to ensure that the choice of the tampering (a)-(c) is not correlated with the message m). In other words, non-malleable codes aim to handle a much larger class of tampering functions \mathcal{F} at the expense of potentially allowing the attacker to replace a given message m by an unrelated message m' (and also necessarily allowing for a small “simulation error” ε).

The authors [14] also showed that non-malleable codes are still useful in many scenarios where the tampering capabilities of the attacker might be too strong for error-detection. For example, imagine a tamper-prone signature card storing a signing key sk and some “context information” α (e.g., the timestamp or some legal disclosure), which will return a signature σ of (α, β) when given an input message β . Imagine now the attacker would like to change α (that he knows) to some related value $\alpha' \neq \alpha$, in the hope of obtaining an “illegal” signature of (α', β) . If $m = (sk, \alpha)$ is encoded using a non-malleable code, then we are guaranteed that the signature σ' obtained by the attacker will either contain the correct value of α , or will not verify anyway, since changing α to α' will also force the attacker to change the signing key sk to a completely unrelated value sk' , making the resulting signature σ' (under sk') “useless”.

SPLIT-STATE MODEL. Given the elegance and utility of non-malleable codes, it is natural to understand the tampering families \mathcal{F} for which such codes exist. As the first observation, we cannot hope to include all possible tampering functions, since \mathcal{F} should not include “re-encoding functions” $f(c) = \text{Enc}(f'(\text{Dec}(c)))$ for any non-trivial function f' (as $m' = \text{Dec}(f(c)) = f'(m)$ is obviously related to m). On the other hand, [14] showed the following positive results. First, they showed a very promising (and surprising!) existence result for any family \mathcal{F} which is only slightly smaller than the family \mathcal{F}_{all} of all functions. Second, they showed an efficient non-malleable code for the family \mathcal{F}_{bit} of “individual” bit-tampering functions f . Although pretty restricted, \mathcal{F}_{bit} includes all constant functions $f(c) = c^*$ (something which cannot be error-detected), and all algebraic manipulation functions $f(c) = c + \Delta$ mentioned earlier.

This raises the question of finding a much larger family \mathcal{F} which is (1) general and realistic from the application point of view; but (2) naturally does not include the re-encoding function to avoid the impossibility. The authors [14] propose to solve this dilemma in the following very elegant way, by defining the so called *split-state model*. The model was originally proposed in the context of leakage-resilient cryptography [13, 7], but it also very natural from the perspective of tampering. Imagine that the encoded memory/state of the system is partitioned in several disjoint parts P_1, \dots, P_t , and the family \mathcal{F}_t of tampering functions consists of all functions $f = (f_1, \dots, f_t)$ where f_i is only applied to the data stored in the partition P_i . To put it differently, the message m is split into t shares s_1, \dots, s_t , and the attacker can arbitrarily tamper with each share independently² by changing it to $s'_i = f_i(s_i)$. Still, the decoded message $m' = \text{Dec}(s'_1, \dots, s'_t)$ is either equal to m , \perp or unrelated to m (as explained above).

As we can see, split-state tampering is very natural from the application point of view, especially when t is low and the shares s_1, \dots, s_t are stored in different parts of memory, or by different parties. Indeed, a non-malleable code w.r.t. \mathcal{F}_t can be viewed as a type of *non-malleable secret sharing* scheme. Recall, in traditional secret sharing schemes one primarily worries about the privacy of the secret m against a certain bounded coalition of shares s_i (which clearly cannot include *all* the t shares). *Robust* secret sharing schemes, considered by [6] (which used the AMD codes mentioned earlier), additionally ensure than a bounded coalition of players cannot maliciously modify their shares and

¹The formal definition (see Definition 2) is also quite clean and elegant, following the standard “simulation paradigm” for other such definitions.

²Of course, we allow $f_1 \dots f_t$ to be correlated, but each f_i can only look at s_i , and not at the other s_j 's.

cause the reconstruction of some secret $m' \neq m$. Once again, the coalition cannot include all t players. In contrast, a non-malleable secret sharing scheme, induced by a non-malleable code in the split-state model, provides the *non-malleability* of the secret m (as explained above) even if *all* t shares are individually modified, something which was never previously considered possible/meaningful in the secret sharing literature.

Coming back to the split-state model, it also overcomes the impossibility result mentioned earlier, since the decoding function will depend on *all* the shares s_1, \dots, s_t (something which is not allowed by the tampering function f). Moreover, since \mathcal{F}_t is indeed noticeably smaller than \mathcal{F}_{all} for $t > 1$, we know that non-malleable codes exist in the split-state model. In fact, the bit-wise tampering family \mathcal{F}_{bit} mentioned above can be viewed as an extreme setting of the split-state model, where each share s_i is only 1 bit (making it rather unrealistic for applications). In particular, it is clear that as t decreases, the tampering family \mathcal{F}_t becomes larger and larger (i.e., more realistic!), and the problem of building non-malleable codes w.r.t. \mathcal{F}_t correspondingly becomes harder and harder, becoming the hardest when $t = 2$. Hence, from now on we will concentrate on the most useful/ambitious case of only two partitions/shares (“left” and “right”), which we will denote by L and R in the sequel.

Summarizing the above discussion, this leads us to the main question of this work:

Main Question: *Build an efficient non-malleable code in the (two-partition) split-state model.*

KNOWN RESULTS. As we mentioned, this question is not new, and several partial results were known prior to our work. First, we already mentioned the existential result of [14] showing the existence of such non-malleable codes. Second, the work of [14] also gave an efficient construction in the random oracle model. Third, the work of Liu and Lysyanskaya [22] built an efficient *computationally-secure* non-malleable code in the split model (necessarily restricting the tampering functions f_1 and f_2 to be efficient as well). The construction assumes so called common reference string (CRS) which cannot be tampered, and also uses quite heavy tools from public-key cryptography, such as robust non-interactive zero-knowledge proofs [8] and leakage-resilient encryption [23]. Thus, given the clean information-theoretic definition of non-malleable codes, we believe it is important to construct such codes unconditionally.

Recently, an important step in this direction was taken by Dziembowski, Kazana and Obremski [12], who constructed a very elegant non-malleable code for 1-bit messages in the split-state model. Their construction is very simple. Both shares L and R lie in an n -dimensional vector space \mathbb{F}^n (for a large enough finite field \mathbb{F} , and of exponential-size). To encode 0, one chooses a random pair of orthogonal vectors L and R ($\langle L, R \rangle = 0$), and to encode 1 one chooses a random pair of non-orthogonal vectors L and R ($\langle L, R \rangle \neq 0$). Despite the simplicity of this construction, the security proof given by [12] was quite involved, and introduced several novel techniques, such as characterizing a given tampering function f_1 or f_2 as being “close” or “far” from a constant. Unfortunately, given the asymmetric nature of their construction (i.e., encodings of 0 and 1 are very different) and several other “bit-specific” proof techniques they use,³ it is unclear how to extend the proof (or even construction!) to the much more useful case of encoding longer than 1 bit messages.

To summarize, despite lots of partial progress, the question of constructing efficient, information-theoretically secure non-malleable codes for long messages was still open prior to our work.

OUR RESULT. As our main result, we resolve this open problem:

Theorem 1 *For any k and ε , there exists an efficient, information-theoretically secure ε -non-malleable code for encoding k -bit messages in the (two-partition) split-state model.*

³I.e., a special characterization of non-malleable codes for 1-bit messages.

As we discuss below, our code is very simple and efficient relative to the length N of the shares L and R (i.e., given N , our encoding and decoding are both very simple). On the other hand, the minimal length $N = \text{poly}(k, \log(1/\varepsilon))$ which is sufficient for our security proof is governed by the current state-of-the-art in additive combinatorics. We discuss this in more detail below and in Section 7, here only mentioning that the current provable bound is $N = \tilde{O}((k + \log(1/\varepsilon))^7)$ (which is very likely sub-optimal).

Our code is constructed in two steps. The first (and much simpler) step constructs a non-malleable code $(\text{Enc}', \text{Dec}')$ for an intermediate tampering family \mathcal{F}_{aff} consisting of all affine functions $f(y) = ay + b$ over some (sufficiently large) finite field \mathbb{F}_p of prime order, where $a, b \in \mathbb{F}_p$ are arbitrary constants. Notice, such \mathbb{F}_p -affine family is rather natural and again includes all constant functions (corresponding to $a = 0$), as well as all algebraic manipulation functions (corresponding to $a = 1$), potentially making our intermediate non-malleable code interesting in its own right. The actual code over the message space \mathcal{M} is constructed by building what we call an *affine-evasive* function $h : \mathbb{F}_p \rightarrow \mathcal{M} \cup \{\perp\}$. Informally, such functions not only send most field elements u to \perp , but also guarantee that $h(au + b) = \perp$ with high probability even conditioned on $h(u) = m$, for any message m and a, b where $(a, b) \neq (1, 0)$ and $a \neq 0$ (i.e., excluding the trivial identity and constant functions, respectively). As a result, the non-malleable code for \mathcal{F}_{aff} easily follows by setting $\text{Dec} = h$. Moreover, we give a construction of such affine-evasive functions h .

The second (and more involved) step can be seen as reducing the task of building a non-malleable code for the split-state model to the non-malleable code for the \mathbb{F}_p -affine function. In particular, we simply use the *inner product function over the n -dimensional vector space \mathbb{F}_p^n* (for a large enough n , discussed below) as our reduction. A bit more formally, $\text{Enc}(m)$ first computes the intermediate encoding $y \leftarrow \text{Enc}'(m)$ for the affine family above, and then picks random shares L and R whose inner product is y : $\langle L, R \rangle = y$. Thus, our construction is similar in spirit to the 1-bit construction of [12], except we treat all messages in a symmetric manner, and ensure that a random pair (L, R) decodes to \perp with high probability. We then show the soundness of our reduction from the split-state model to the \mathbb{F}_p -affine model, by showing the following key theorem about the “non-malleability” of the inner product function:

Theorem 2 (Informal) *Assume \mathbb{F}_p is a finite field of prime order, $n \geq \text{poly}(\log p)$, L and R are uniformly random over \mathbb{F}_p^n , and $f, g : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ are two arbitrary functions on L and R . Then, the joint distribution $(\langle L, R \rangle, \langle f(L), g(R) \rangle)$ is “close” to the convex combination of affine distributions $\{(U, aU + b) \mid a, b \in \mathbb{F}_p\}$, where U is uniformly random over \mathbb{F}_p .*

The formal statement appears in Theorem 3. Intuitively, though, the above result shows that the inner product function effectively maps the (seemingly) very powerful split-state tampering (given by arbitrary functions f and g) to a convex combination of much more basic affine functions $ay + b$ (which, in turn, are protected by our “inner” non-malleable code). Not surprisingly, the proof of Theorem 2 (or, more accurately, Theorem 3) forms the main technical contribution of our work, and may be of independent interest. It is detailed in Section 5, but crucially relies on Theorem 6, which in turn relies on several results from additive combinatorics. Theorem 6 can be seen as an improvement of the linearity test of [24] for functions $f : \mathbb{F}_p^n \mapsto \mathbb{F}_p^n$. The key ingredient resulting in this improvement is the so called the *Quasi-polynomial Freiman-Ruzsa Theorem*, which was recently established by Sanders [25] as a step towards resolving the Polynomial Freiman-Ruzsa conjecture [17]. We refer to Section 5.3 and Section 6 for more details on specific parameters and how they are used to establish Theorem 3, but mention that the (likely) sub-optimality of Sander’s result is the main reason for a relatively large dimension $n \approx \log^6 p = \tilde{O}((k + \log(1/\varepsilon))^6)$ of the vector space \mathbb{F}_p^n for our non-malleable encoding of k -bit messages, which leads to an even larger encoding length $N = n \log p = \tilde{O}((k + \log(1/\varepsilon))^7)$ (see also the “Conclusions” Section 7).

OTHER RELATED WORK. In addition to the already-mentioned results of [14, 22, 12], several recent works [3, 4, 5] either used or built various non-malleable codes, but none concentrated on the split-state model considered here.

The notion of non-malleability was introduced by the seminal paper of Dolev, Dwork and Naor [11], and has found many applications in cryptography. Traditionally, non-malleability is defined in the computational setting, but recently non-malleability has been successfully defined and applied in the information-theoretic setting (generally resulting in somewhat simpler and cleaner definitions than their computational counter-parts). For example, in addition to non-malleable codes studied in this work, the work of Dodis and Wichs [10] defined the notion of non-malleable extractors as a tool for building round-efficient privacy amplification protocols.

Finally, the study of non-malleable codes falls into a much larger cryptographic framework of providing counter-measures against various classes of tampering attacks. This work was pioneered by the early works of [19, 15, 18], and has since led to many subsequent models. Listing all such tampering models (which are not directly related to the study of non-malleable codes) is beyond the scope of this work, but we refer to [20, 22] for an excellent discussion of various such models.

2 Preliminaries

All logarithms are in base two. Unless stated otherwise, \mathbb{F}_p is a finite field of prime order p .

Distributions. Let D be a discrete distribution. We denote by $D[x]$ the probability it assigns to x , and by $X \sim D$ a random variable distributed according to D . For two distributions D, D' their statistical distance is $\Delta(D; D') = \frac{1}{2} \sum_x |D[x] - D'[x]|$. Let \mathcal{D} be a family of distributions. We denote by $\Delta(D; \mathcal{D})$ the infimum of $\Delta(D; D')$ over all $D' \in \mathcal{D}$.

A convex combination of distributions D_1, \dots, D_k is any distribution D for which

$$D[x] = \sum \alpha_i D_i[x],$$

for all x , where $\alpha_i \geq 0$ and $\sum \alpha_i = 1$.

The min-entropy of a distribution is $\mathbf{H}_\infty(D) = \min_x \log(D[x]^{-1})$. For a finite set S we denote by U_S the uniform distribution over S . Note that $\mathbf{H}_\infty(U_S) = \log |S|$. Moreover, if D is a distribution with min-entropy k then D is a convex combination of distributions uniform over sets of size 2^k .

We denote random variables by X, L, R . Let E be an event. We denote by $X|E$ the conditional random variable, conditioned on E holding. For a set S we shorthand $X|_S = X|[X \in S]$. When there is no chance of confusion, we use interchangeably a random variable to denote also its underlying distribution.

Inequalities on distributions far from uniform. We will need the following claims. Their proofs can be found in the appendix.

Claim 1 *Let $X = (X_1, X_2) \in \mathbb{F}_p \times \mathbb{F}_p$ be a random variable. Assume that for all $a, b \in \mathbb{F}_p$ not both zero, $\Delta(aX_1 + bX_2; U_{\mathbb{F}_p}) \leq \varepsilon$. Then $\Delta((X_1, X_2); U_{\mathbb{F}_p^2}) \leq \varepsilon p^2$.*

Claim 2 *Let $X \in \mathbb{F}_p$ be a random variable. Assume that $\Delta(X; U_{\mathbb{F}_p}) \geq \varepsilon$. Then if X' is an independent and i.i.d copy of X then*

$$\Pr[X = X'] \geq \frac{1 + \varepsilon^2}{p}.$$

Claim 3 Let $Z = (X, Y) \in \mathbb{F}_p^n \times \mathbb{F}_p^n$ be a random variable, and let $Z' = (X', Y')$ be an i.i.d copy of Z . Then

$$\Pr[\langle X, Y \rangle = \langle X', Y' \rangle] \leq \Pr[\langle X, Y \rangle = \langle X', Y \rangle].$$

Claim 4 Let $X_1, X_2, Y_1, Y_2 \in \mathcal{A}$ be random variables such that $\Delta((X_1, X_2); (Y_1, Y_2)) \leq \varepsilon$. Then, for any non-empty set $\mathcal{A}_1 \subseteq \mathcal{A}$, we have

$$\Delta(X_2 | X_1 \in \mathcal{A}_1; Y_2 | Y_1 \in \mathcal{A}_1) \leq \frac{2\varepsilon}{\Pr(X_1 \in \mathcal{A}_1)}.$$

The Hadamard extractor. The Hadamard extractor is one of the most basic two-source extractors, based on inner product. We would need the following folklore result. A proof can, for example, be found in [21].

Lemma 1 Let L and R be independent random variables over \mathbb{F}_p^n . If

$$\mathbf{H}_\infty(L) + \mathbf{H}_\infty(R) \geq (n+1) \log p + 2 \log \left(\frac{1}{\varepsilon} \right),$$

then

$$\Delta((L, \langle L, R \rangle); (L, U_{\mathbb{F}_p})) \leq \varepsilon \text{ and } \Delta((R, \langle L, R \rangle); (R, U_{\mathbb{F}_p})) \leq \varepsilon.$$

3 The joint probability distribution of $(\langle L, R \rangle, \langle f(L), g(R) \rangle)$

Let \mathbb{F}_p be a finite field of prime order. Let $L, R \in \mathbb{F}_p^n$ be uniform and independent. Let $f, g : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ be a pair of functions. We consider the following family of distributions

$$\phi_{f,g}(L, R) := (\langle L, R \rangle, \langle f(L), g(R) \rangle) \in \mathbb{F}_p^2$$

We characterize in this section the possible joint distributions of $\phi_{f,g}(L, R)$ over \mathbb{F}_p^2 for arbitrary functions f, g . In order to build intuition, let us first consider a few of possible distributions achievable this way.

- $f(L) = (a, 0, \dots, 0), g(R) = (1, 0, \dots, 0)$ for $a \in \mathbb{F}_p$. Then $\phi_{f,g}(L, R)$ has a distribution that is statistically very close to (U, a) where $U \in \mathbb{F}_p$ is uniform.
- $f(L) = aAL, g(R) = (A^T)^{-1}R$ for some $a \in \mathbb{F}_p$, and invertible matrix $A \in \mathbb{F}_p^{n \times n}$. Then $\phi_{f,g}(L, R)$ has a distribution that is statistically very close to (U, aU) where $X \in \mathbb{F}_p$ is uniform.

In general, by choosing f, g as an arbitrary mix of the above, we can achieve nearly any convex combination of $\{(U, a) : a \in \mathbb{F}_p\}$ and $\{(U, aU) : a \in \mathbb{F}_p\}$, where U is uniform in \mathbb{F}_p . For a large number of choices of f, g , these are the only possible distributions of $\phi_{f,g}(L, R)$. The following, however, shows an example of f, g for which $\phi_{f,g}(L, R)$ has statistical distance about $1/p$ from any of these distributions.

- Fix $v \in \mathbb{F}_p^n$ with $\langle v, v \rangle = 1$. Let $f(L) = L + \langle L, v \rangle v, g(R) = R - \langle R, v \rangle v$. Then $\phi_{f,g}(L, R)$ is very close to being distributed as $(U, U + XY)$ where $U, X, Y \in \mathbb{F}_p$ are uniform and independent. Note that the distribution of XY is not uniform, as it is equal to zero with probability $2/p - 1/p^2$ instead of $1/p$.

We do not have a complete characterization of all possible distributions $\phi_{f,g}(L, R)$. However, our main technical result is that any such distribution is arbitrarily close to a convex combination of $(U, aU + b)$ where $a, b \in \mathbb{F}_p$ if n is large enough. Define \mathcal{D} to be the family of convex combinations of $\{(U, aU + b) : a, b \in \mathbb{F}_p\}$ where $U \in \mathbb{F}_p^n$ is uniform. This will be sufficient to analyze our construction of non-malleable codes.

Theorem 3 *There exists universal constants $c, c' > 0$ such that the following holds. For any finite field \mathbb{F}_p of prime order, and any $n > c' \log^6 p$, let $L, R \in \mathbb{F}_p^n$ be uniform, and fix $f, g : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$. Then*

$$\Delta(\phi_{f,g}(L, R) ; \mathcal{D}) \leq 2^{-cn^{1/6}}.$$

We give a proof of this theorem in Section 5.

4 Non-malleable Codes

Definitions. We first recall the definition of non-malleable codes from [14].

Definition 1 A coding scheme consists of two functions: a randomized encoding function $\text{Enc} : \mathcal{M} \mapsto \mathcal{C}$, and a deterministic decoding function $\text{Dec} : \mathcal{C} \mapsto \mathcal{M} \cup \{\perp\}$ such that, for each $m \in \mathcal{M}$, $\Pr(\text{Dec}(\text{Enc}(m)) = m) = 1$ (over the randomness of the encoding algorithm).

Definition 2 Let \mathcal{F} be some family of tampering functions. For each $f \in \mathcal{F}$, and $m \in \mathcal{M}$, define the tampering-experiment

$$\text{Tamper}_m^f := \left\{ \begin{array}{l} c \leftarrow \text{Enc}(m), \tilde{c} \leftarrow f(c), \tilde{m} = \text{Dec}(\tilde{c}) \\ \text{Output: } \tilde{m}. \end{array} \right\}$$

which is a random variable over the randomness of the encoding function Enc . We say that a coding scheme (Enc, Dec) is ε -non-malleable w.r.t. \mathcal{F} if for each $f \in \mathcal{F}$, there exists a distribution (corresponding to the simulator) D_f over $\mathcal{M} \cup \{\perp, \text{same}^*\}$, such that, for all $m \in \mathcal{M}$, we have that the statistical distance between Tamper_m^f and

$$\text{Sim}_m^f := \left\{ \begin{array}{l} \tilde{m} \leftarrow D_f \\ \text{Output: } m \text{ if } \tilde{m} = \text{same}^*, \text{ and } \tilde{m}, \text{ otherwise.} \end{array} \right\}$$

is at most ε . Additionally, D_f should be efficiently samplable given oracle access to $f(\cdot)$.

Our result. For any $\varepsilon > 0$, we give an encoding scheme from $\mathcal{M} = \{1, \dots, K\}$ to $\mathbb{F}_p^n \times \mathbb{F}_p^n$ (where $p = (\frac{K}{\varepsilon})^{\Theta(\log \log(K/\varepsilon))}$, and $n = \Theta(\log^6 p)$) that is ε -non-malleable with respect to the family of all functions in the split state model, i.e., all functions $(f, g) : \mathbb{F}_p^n \times \mathbb{F}_p^n \mapsto \mathbb{F}_p^n \times \mathbb{F}_p^n$, where f and g are functions from $\mathbb{F}_p^n \mapsto \mathbb{F}_p^n$, and $(f, g)(x, y) = (f(x), g(y))$, for all $x, y \in \mathbb{F}_p^n$. Our construction proceeds as follows.

- In Section 4.1, we construct an encoding scheme from \mathcal{M} to \mathbb{F}_p that is non-malleable with respect to the class of all affine functions over \mathbb{F}_p .
- In Section 4.2, we use Theorem 3 to argue that we can reduce the problem of constructing an encoding scheme from \mathcal{M} to $\mathbb{F}_p^n \times \mathbb{F}_p^n$ that is non-malleable in the split state model to the problem of constructing an encoding scheme from \mathcal{M} to \mathbb{F}_p that is non-malleable with respect to the class of all affine functions over \mathbb{F}_p . We then use the result of Section 4.1 to conclude the result.

For the subsequent sections, we denote by U a random variable distributed uniformly over \mathbb{F}_p .

4.1 A non-malleable encoding scheme with respect to affine functions

In this section, we will construct an encoding scheme from $\mathcal{M} = \{1, \dots, K\}$ to a finite field \mathbb{F}_p of prime order p , where $p = (\frac{K}{\varepsilon})^{\Theta(\log \log(K/\varepsilon))}$ that is ε -non-malleable with respect to the family of affine functions \mathcal{F}_{aff} over \mathbb{F}_p , i.e.,

$$\mathcal{F}_{\text{aff}} := \{f(y) = ay + b : a, b \in \mathbb{F}_p\}.$$

Construction. For our construction, we use affine-evasive functions, defined as follows: A surjective function $h : \mathbb{F}_p \mapsto \mathcal{M} \cup \{\perp\}$ is called (γ, δ) -*affine-evasive* if or any $a, b \in \mathbb{F}_p$ such that $a \neq 0$, and $(a, b) \neq (1, 0)$, and for any $m \in \mathcal{M}$,

- $\Pr(h(aU + b) \neq \perp) \leq \gamma$
- $\Pr(h(aU + b) \neq \perp \mid h(U) = m) \leq \delta$
- A uniformly random X such that $h(X) = m$ is efficiently samplable

Let $h : \mathbb{F}_p \mapsto \mathcal{M} \cup \{\perp\}$ be a (γ, δ) -affine-evasive function. The scheme is defined using h as follows: The decoding function $\text{Dec} : \mathbb{F}_p \mapsto \mathcal{M} \cup \{\perp\}$ is defined as $\text{Dec}(x) := h(x)$. The encoding function is defined as $\text{Enc}(m) = X$ where X is chosen at random from \mathbb{F}_p conditioned on the fact that $\text{Dec}(X) = m$.

Theorem 4 *Let $\mathcal{M} = \{1, \dots, K\}$ and let \mathbb{F}_p be a finite field. Let $\mathcal{F}_{\text{aff}}, \text{Enc} : \mathcal{M} \mapsto \mathbb{F}_p, \text{Dec} : \mathbb{F}_p \mapsto \mathcal{M} \cup \{\perp\}$ be as defined above. The scheme (Enc, Dec) is $(\gamma + \delta + \frac{1}{p})$ -non malleable w.r.t. \mathcal{F}_{aff} .*

We now give a proof of Theorem 4.

Simulator. For any function $f \in \mathcal{F}_{\text{aff}}$, we define the distribution D_f over $\mathcal{M} \cup \{\perp, \text{same}^*\}$ as the output of the following (efficient) sampling procedure:

1. Choose $x \leftarrow \mathbb{F}_p$.
2. If $f(x) = x$, then output same^* , else output $h(f(x))$.

The distribution D_f can thus be expressed as:

$$D_f = \begin{cases} \text{same}^* & \text{with prob. } \Pr_{x \leftarrow \mathbb{F}_p}(f(x) = x) \\ m' & \text{with prob. } \Pr_{x \leftarrow \mathbb{F}_p}(h(f(x)) = m', \text{ and } f(x) \neq x), \end{cases}$$

where $m' \in \mathcal{M} \cup \{\perp\}$.

Security Proof. Consider some $m \in \mathcal{M}$, and some $f \in \mathcal{F}_{\text{aff}}$ given by $f(y) = ay + b$ for some $a, b \in \mathbb{F}_p$. The random variable Tamper_m^f (abbreviated as $\text{Tamper}_m^{(a,b)}$) has the following distribution for all $m' \in \mathcal{M} \cup \{\perp\}$.

$$\Pr(\text{Tamper}_m^{(a,b)} = m') = \Pr(h(aU + b) = m' \mid h(U) = m) \quad (1)$$

The random variable corresponding to the simulator Sim_m^f (denoted as $\text{Sim}_m^{(a,b)}$) has the following distribution for all $m' \in \mathcal{M} \cup \{\perp\}$.

$$\Pr(\text{Sim}_m^{(a,b)} = m') = \begin{cases} \Pr(h(aU + b) = m' \wedge U \neq aU + b) & \text{if } m' \neq m \\ \Pr(U = aU + b \vee (h(aU + b) = m \wedge U \neq aU + b)) & \text{if } m' = m \end{cases} \quad (2)$$

Lemma 2 For any $m \in \mathcal{M}$, any $a, b \in \mathbb{F}_p$, and any (γ, δ) -affine evasive function h ,

$$\Delta \left(\text{Sim}_m^{(a,b)} ; \text{Tamper}_m^{(a,b)} \right) \leq \gamma + \delta + \frac{1}{p}.$$

Proof. If $(a, b) = (1, 0)$, then $\Pr(\text{Sim}_m^{(a,b)} = m) = \Pr(\text{Tamper}_m^{(a,b)} = m) = 1$, and so

$$\Delta \left(\text{Sim}_m^{(a,b)} ; \text{Tamper}_m^{(a,b)} \right) = 0.$$

Thus, we may assume $(a, b) \neq (1, 0)$. This implies that $\Pr(U = aU + b) \leq \frac{1}{p}$. Therefore,

$$\Delta \left(h(aU + b) ; \text{Sim}_m^{(a,b)} \right) \leq \frac{1}{p}.$$

If $a = 0$, then we have $\Delta \left(h(aU + b) ; \text{Tamper}_m^{(a,b)} \right) = 0$. So, we may also assume $a \neq 0$. We have by the definition of statistical distance that

$$\begin{aligned} \Delta \left(\text{Tamper}_m^{(a,b)} ; h(aU + b) \right) &= \frac{1}{2} \cdot \sum_{m' \in \mathcal{M}} \left| \Pr(\text{Tamper}_m^{(a,b)} = m') - \Pr(h(aU + b) = m') \right| \\ &\quad + \frac{1}{2} \cdot \left| \Pr(\text{Tamper}_m^{(a,b)} = \perp) - \Pr(h(aU + b) = \perp) \right|. \end{aligned}$$

Using the fact that

$$\Delta \left(\text{Tamper}_m^{(a,b)} ; h(aU + b) \right) \geq \left| \Pr(\text{Tamper}_m^{(a,b)} = \perp) - \Pr(h(aU + b) = \perp) \right|,$$

we get

$$\begin{aligned} \Delta \left(\text{Tamper}_m^{(a,b)} ; h(aU + b) \right) &\leq \sum_{m' \in \mathcal{M}} \left| \Pr(\text{Tamper}_m^{(a,b)} = m') - \Pr(h(aU + b) = m') \right| \\ &\leq \Pr(h(aU + b) \neq \perp \mid h(U) = m) + \Pr(h(aU + b) \neq \perp) \leq \gamma + \delta, \end{aligned}$$

where the last inequality makes use of the fact that h is (γ, δ) -affine evasive. Therefore, using the triangle inequality,

$$\Delta \left(\text{Sim}_m^{(a,b)} ; \text{Tamper}_m^{(a,b)} \right) \leq \gamma + \delta + \frac{1}{p}.$$

□

Remark: Note that the scheme (Enc, Dec) also achieves error-detection with respect to non-constant affine functions $\{f(y) = ay + b : a, b \in \mathbb{F}_p, a \neq 0\}$.

An affine-evasive function. For any set $S \subset \mathbb{Z}$, let $aS + b = \{as + b \mid s \in S\}$. By $S \bmod p \subseteq \mathbb{F}_p$, we denote the set of values of S modulo p .

We first define an affine-evasive set $S \subseteq \mathbb{F}_p$.

Definition 3 A set $S \subseteq \mathbb{F}_p$ is said to be (γ, ν) -affine-evasive if $|S| \leq \gamma p$, and for any $(a, b) \in \mathbb{F}_p^2 \setminus \{(1, 0)\}$, we have

$$|S \cap (aS + b \pmod{p})| \leq \nu |S|.$$

We claim that an affine-evasive function can be constructed from an affine-evasive set.

Claim 5 Let $S \subseteq \mathbb{F}_p$ be a (γ, ν) -affine-evasive set with $\nu \cdot K \leq 1$, and K divides $|S|$.⁴ Furthermore, let S be ordered such that for any i , the i -th element is efficiently computable. Then there exists a $(\gamma, \nu \cdot K)$ -affine-evasive function $h : \mathbb{F}_p \mapsto \mathcal{M} \cup \{\perp\}$.

Proof. Consider any fixed partition of S into K subsets S_1, \dots, S_K each of cardinality $|S|/K$. Let $h : \mathbb{F}_p \mapsto \mathcal{M} \cup \{\perp\}$ be defined as follows:

$$h(x) = \begin{cases} i & \text{if } x \in S_i \\ \perp & \text{otherwise.} \end{cases}$$

It is straightforward to see that h is a $(\gamma, \nu \cdot K)$ -affine-evasive function. The statement $\Pr(h(aU + b) \neq \perp) \leq \gamma$ is obvious by the definition of S , and the observation that $aU + b$ is uniform in \mathbb{F}_p .

Also, for any $m \in \mathcal{M}$, and for any $(a, b) \neq (1, 0)$, and $a \neq 0$,

$$\begin{aligned} \Pr(h(aU + b) \neq \perp | h(U) = m) &= \frac{\Pr(aU + b \in S \wedge U \in S_m)}{\Pr(U \in S_m)} \\ &\leq \frac{\Pr(aU + b \in S \wedge U \in S)}{|S|/K} \\ &= \frac{K}{|S|} \Pr(U \in S \cap (a^{-1}S - ba^{-1}) \pmod{p}) \\ &\leq \nu \cdot K. \end{aligned}$$

□

Next, we give a construction of an affine-evasive set.

An affine-evasive set We want to construct a large set $S \subset \mathbb{Z}_p$ so that for any $a, b \in \mathbb{Z}_p$, $|S \cap (aS + b)|$ is much smaller than $|S|$. We first argue that by choosing all elements of S small (as integers) it suffices to consider non-modular operations. In the following let $[p] = \{0, 1, \dots, p-1\}$ and we will consider $S \subset [p]$.

Lemma 3 Let $S \subset [[p^{1/2}/4]]$. Assume that for any $a, b, c \in \mathbb{Z}$ (except for $a = c, b = 0$) we have $|cS \cap (aS + b)| \leq \ell$ for some $\ell \geq 1$. Then also $|S \cap (a'S + b' \pmod{p})| \leq \ell$ for all $a', b' \in \mathbb{Z}_p$ (except for $(a', b') = (1, 0)$).

Proof. Assume there are $a', b' \in \mathbb{Z}_p$, $(a', b') \neq (1, 0)$, such that $|S \cap (a'S + b' \pmod{p})| > \ell$. We can find $x_1, x_2, y_1, y_2 \in S$ such that $x_1 = a'y_1 + b' \pmod{p}$, $x_2 = a'y_2 + b' \pmod{p}$, and $y_1 \neq y_2$. Subtracting these give $x_1 - x_2 = a'(y_1 - y_2) \pmod{p}$. This means we can write $a' = a/c \pmod{p}$ where $|a|, |c| \leq p^{1/2}/2$. That is, we have (at least) $\ell + 1$ solutions (x_i, y_i) with $i = 1, \dots, \ell + 1$ to

$$cx_i = ay_i + b \pmod{p}$$

where $b = b'c \pmod{p}$. Note that $|cx_i| \leq p/8$ and $|ay_i| \leq p/8$. Hence also $|b| \leq p/4$. So for $i = 1, \dots, \ell + 1$, $|cx_i - ay_i - b| \leq p/2$. But since $cx_i - ay_i - b = 0 \pmod{p}$ we must have that in fact $cx_i - ay_i - b = 0$ over the integers. Hence

$$cx_i = ay_i - b.$$

⁴The assumption K divides $|S|$ is just for simplicity.

That is, $|cS \cap (aS + b)| > \ell$. Contradiction to our assumption. \square

We now show how to construct an $S \subseteq [[p^{1/2}/4]]$ with such that $|cS \cap (aS + b)|$ is much smaller than $|S|$.

Let p_1, \dots, p_t be the first t primes. Let $M = \prod_{i=1}^t p_i$. By the chinese remainder theorem, for all $i \in [t]$, there exists a unique $M_i \in [M]$ be such that $M_i = 1 \pmod{p_i}$ and $M_i = 0 \pmod{p_j}$ for $j \neq i$. Define a set $S \subset \mathbb{Z}$ as

$$S = \left\{ \sum a_i M_i : a_i \in \{0, 1\} \right\}.$$

Note that we can ensure that $S \subset [[p^{1/2}/4]]$ by choosing $p > 4(Mt)^2$. This is satisfied for $p \geq t^{O(t)}$.

Lemma 4 *Let $a, b, c \in \mathbb{Z}$ be any values except for $a = c, b = 0$. Then $|cS \cap (aS + b)| \leq |S|^{0.92}$.*

Proof. Let $x, y \in S$ be an arbitrary solution of the equation $ax + b = cy$. Note that we can assume w.l.o.g that a, c are co-prime. Let $x = \sum x_i M_i$ and $y = \sum y_i M_i$ where $x_i, y_i \in \{0, 1\}$. Let p_i be a prime, $i \leq t$. We have that $ax_i + b = cy_i \pmod{p_i}$. We consider several cases.

1. If p_i divides a then $y_i = (b/c) \pmod{p_i}$ is fixed in all solutions.
2. If p_i divides c then $x_i = (-b/a) \pmod{p_i}$ is fixed in all solutions.
3. If p_i divides b but not a, c then $x_i a = y_i c \pmod{p_i}$. Any possible solution (with $x_i, y_i \in \{0, 1\}$) must satisfy $x_i = y_i$.
4. If p_i divides $b + a - c$ but not a, c , then $a(x_i - 1) = c(y_i - 1) \pmod{p_i}$. Any possible solution (with $x_i, y_i \in \{0, 1\}$) must satisfy $x_i = y_i$.
5. If p_i does not divide $a, b, c, b + a - c$ then $y_i = x_i(a/c) + (b/c) \pmod{p_i}$. In any possible solution, not both x_i, y_i can be $0 \pmod{p_i}$ because this implies that $b = 0 \pmod{p_i}$. Also, not both can be 1 since that implies $b + a - c = 0 \pmod{p_i}$. Thus, $y_i = 1 - x_i$.

We next use these observations to bound the number of solutions to $ax + b = cy$ with $x, y \in S$. Let $\varepsilon > 0$ be a parameter to be determined later.

1. If a is divisible by $\{p_i : i \in I\}$ with $|I| \geq \varepsilon t$. Then the value of $y_i, i \in I$ is fixed in all solutions. Hence there are at most $2^{(1-\varepsilon)t}$ solutions.
2. If c is divisible by $\{p_i : i \in I\}$ with $|I| \geq \varepsilon t$. Then the value of $x_i, i \in I$ is fixed in all solutions. Hence there are at most $2^{(1-\varepsilon)t}$ solutions.
3. If one or both of b and $b + a - c$ are divisible by $\{p_i : i \in I\}$ but a, c are not, with $|I| \geq (1/2 - \varepsilon)t$. Then $x_i = y_i, i \in I$ in all solutions (x, y) . Hence we can write $y = x + z$ with $z = \sum_{i \in [t] \setminus I} z_i M_i$ where $z_i \in \{-1, 0, 1\}$. Note that given the value of z , there is at most one rational solution for $ax + b = c(x + z)$. The number of possible values for z is $3^{(1/2 + \varepsilon)t}$, which bounds the number of solutions (x, y) .
4. Otherwise, a, c are divisible by at most εt primes each from p_1, \dots, p_t , and b or $b + a - c$ is divisible by at most $(1/2 - \varepsilon)t$ of the remaining set of primes. Let $\{p_i : i \in I\}$ be the set of primes which do not divide any of $a, b, c, a + b - c$, with $|I| \geq (1/2 - \varepsilon)t$. For any $i \in I$ we have $y_i = 1 - x_i$. Let $J = \sum M_i$. Then we can write $y = J - x + z$ with $z = \sum_{i \in [t] \setminus I} z_i M_i$ where $z_i \in \{-1, 0, 1\}$. As before, given the value of z there are at most one solution, hence the total number of solutions is bounded by $3^{(1/2 + \varepsilon)t}$.

To optimize we need to choose ε so that $2^{(1-\varepsilon)} = 3^{(1/2+\varepsilon)}$. Setting $\varepsilon = (\ln 2 - (1/2) \ln 3) / (\ln(3) + \ln(2)) \approx 0.08027$ gives a tight bound. Hence the number of solutions is bounded by

$$|cS \cap (aS + b)| \leq |S|^{1-\varepsilon} \leq |S|^{0.92}.$$

□

Using the above construction of an affine-evasive set, and Claim 5, we get a (γ, δ) -affine evasive function by choosing S such that $|S| = (\frac{K}{\delta})^{1/0.08}$, i.e., such that $\delta \leq \frac{K \cdot |S|^{0.92}}{|S|}$, and consequently $p \geq |S|^{O(\log |S|)}$. Also note that $\gamma \leq |S|/p \ll \delta$.

Corollary 1 *Let $\mathcal{M} = \{1, \dots, K\}$. There exists a universal constant ρ such that for any prime $p \geq (\frac{K}{\delta})^{\rho \log \log(K/\delta)}$, there exists a (δ, δ) -affine-evasive function $h : \mathbb{F}_p \mapsto \mathcal{M} \cup \{\perp\}$.*

Using this affine-evasive function in the decoding scheme, we obtain the following corollary using Lemma 2.

Corollary 2 *For any $\varepsilon > 0$, $\mathcal{M} = \{1, \dots, K\}$ and let $p \geq (\frac{4K}{\varepsilon})^{\rho \log \log(4K/\varepsilon)}$ be a prime. Then the scheme (Enc, Dec) is ε -non malleable w.r.t. \mathcal{F}_{aff} . In particular, for any $m \in \mathcal{M}$, any $a, b \in \mathbb{F}_p$,*

$$\Delta \left(\text{Sim}_m^{(a,b)} ; \text{Tamper}_m^{(a,b)} \right) \leq \varepsilon.$$

4.2 Non-malleable codes in the split-state model

Now we are in place to give an information-theoretically secure construction of non-malleable codes in the split-state model.

Construction. We construct an ε -non-malleable encoding scheme from $\mathcal{M} = \{1, \dots, K\}$ to $\mathbb{F}_p^n \times \mathbb{F}_p^n$, where \mathbb{F}_p is a finite field of prime order p such that $p \geq (\frac{4K}{\varepsilon})^{\rho \log \log(2K/\varepsilon)}$, and n chosen as $\left(\left\lceil \frac{2 \log p}{c} \right\rceil \right)^6$ (i.e., such that $2^{cn^{1/6}} \geq p^2$), where c is the constant from Theorem 3.

The decoding function $\text{Dec}^* : \mathbb{F}_p^n \times \mathbb{F}_p^n \mapsto \mathcal{M} \cup \{\perp\}$ is defined using the Dec function (which was chosen to be an affine-evasive function h) from Section 4.1 as:

$$\text{Dec}^*(L, R) := \text{Dec}(\langle L, R \rangle) = h(\langle L, R \rangle).$$

The encoding function is defined as $\text{Enc}^*(m) := (L, R)$ where L, R are chosen uniformly at random from $\mathbb{F}_p^n \times \mathbb{F}_p^n$ conditioned on the fact that $h(L, R) = m$.

We will show that our scheme is ε -non-malleable with respect to the family of all functions $(f, g) : \mathbb{F}_p^n \times \mathbb{F}_p^n \mapsto \mathbb{F}_p^n \times \mathbb{F}_p^n$, where f and g are functions from $\mathbb{F}_p^n \mapsto \mathbb{F}_p^n$, and $(f, g)(x, y) = (f(x), g(y))$, for all $x, y \in \mathbb{F}_p^n$. Let us call this family of functions \mathcal{G} .

Theorem 5 *Let $\mathcal{M} = \{1, \dots, K\}$ and let $p \geq (\frac{4K}{\varepsilon})^{\rho \log \log(4K/\varepsilon)}$ be a prime. Let n be $\left(\left\lceil \frac{2 \log p}{c} \right\rceil \right)^6$. Let $\mathcal{G}, \text{Enc}^* : \mathcal{M} \mapsto \mathbb{F}_p^n \times \mathbb{F}_p^n, \text{Dec}^* : \mathbb{F}_p^n \times \mathbb{F}_p^n \mapsto \mathcal{M} \cup \{\perp\}$ be as defined above. Then the scheme $(\text{Enc}^*, \text{Dec}^*)$ is ε -non malleable w.r.t. \mathcal{G} .*

We now give a proof of Theorem 5.

Simulator. For any functions $f, g : \mathbb{F}_p^n \mapsto \mathbb{F}_p^n$, we define the distribution $D_{f,g}$ over $\mathcal{M} \cup \{\perp, \text{same}^*\}$ as the output of the following sampling procedure:

1. Choose $L, R \leftarrow \mathbb{F}_p^n$.
2. If $\langle f(L), g(R) \rangle = \langle L, R \rangle$, then output same^* , else output $h(\langle f(L), g(R) \rangle)$.

Note that this distribution is efficiently samplable given oracle access to f and g . The distribution $D_{f,g}$ can also be expressed as:

$$D_{f,g} = \begin{cases} \text{same}^* & \text{with prob. } \Pr_{L,R \leftarrow \mathbb{F}_p^n}(\langle f(L), g(R) \rangle = \langle L, R \rangle) \\ m' & \text{with prob. } \Pr_{L,R \leftarrow \mathbb{F}_p^n}(h(\langle f(L), g(R) \rangle) = m', \text{ and } \langle f(L), g(R) \rangle \neq \langle L, R \rangle), \end{cases}$$

where $m' \in \mathcal{M} \cup \{\perp\}$.

Security Proof. The random variable corresponding to the tampering experiment $\text{Tamper}_m^{(f,g)}$ has the following distribution for all $m' \in \mathcal{M} \cup \{\perp\}$.

$$\Pr(\text{Tamper}_m^{(f,g)} = m') = \Pr(h(\langle f(L), g(R) \rangle) = m' \mid h(\langle L, R \rangle) = m). \quad (3)$$

The random variable corresponding to the simulator $\text{Sim}_m^{(f,g)}$ has the following distribution for all $m' \in \mathcal{M} \cup \{\perp\}$.

$$\Pr(\text{Sim}_m^{(f,g)} = m') = \begin{cases} \Pr(h(\langle f(L), g(R) \rangle) = m' \wedge \overline{E}) & \text{if } m' \neq m \\ \Pr(E \vee (h(\langle f(L), g(R) \rangle) = m \wedge \overline{E})) & \text{if } m' = m \end{cases}, \quad (4)$$

where E is the event $\langle f(L), g(R) \rangle = \langle L, R \rangle$

From Theorem 3, we get that there exists a random variable (X, Y) taking values in $\mathbb{F}_p \times \mathbb{F}_p$ such that

$$\Delta(\langle L, R \rangle, \langle f(L), g(R) \rangle; X, Y) \leq \frac{1}{p^2}$$

and X, Y is a convex combination of $\{(U, aU + b) : a, b \in \mathbb{F}_p\}$, where U is uniformly distributed in \mathbb{F}_p . This implies that there exist $\{p_{a,b} : a, b \in \mathbb{F}_p\}$ such that $\sum_{a,b \in \mathbb{F}_p} p_{a,b} = 1$ and

$$\Pr(X = x, Y = y) = \sum_{a,b \in \mathbb{F}_p} p_{a,b} \Pr(U = x, aU + b = y),$$

for all $x, y \in \mathbb{F}_p$.

Using Claim 4 and that $\Delta(\langle L, R \rangle, \langle f(L), g(R) \rangle; X, Y) \leq \frac{1}{p^2}$, we get that

$$\Delta(\text{Tamper}_m^{(f,g)}; T) \leq \frac{2}{p} \quad \text{and} \quad \Delta(\text{Sim}_m^{(f,g)}; S) \leq \frac{1}{p^2},$$

where S and T are defined as follows for all $m' \in \mathcal{M} \cup \{\perp\}$:

$$\Pr(T = m') = \Pr(h(Y) = m' \mid h(X) = m)$$

$$\Pr(S = m') = \begin{cases} \Pr(h(Y) = m' \wedge Y \neq X) & \text{if } m' \neq m \\ \Pr(Y = X \vee (h(Y) = m \wedge Y \neq X)) & \text{if } m' = m \end{cases}.$$

The statistical distance between S and T is

$$\begin{aligned}
\Delta(S ; T) &= \frac{1}{2} \sum_{m' \in \mathcal{M} \cup \{\perp\}} \left| \Pr(S = m') - \Pr(T = m') \right| \\
&= \frac{1}{2} \sum_{m' \in \mathcal{M} \cup \{\perp\}} \left| \sum_{a,b \in \mathbb{F}_p} p_{a,b} \Pr(\text{Sim}_m^{(a,b)} = m') - \sum_{a,b \in \mathbb{F}_p} p_{a,b} \Pr(\text{Tamper}_m^{(a,b)} = m') \right| \\
&\leq \frac{1}{2} \sum_{m' \in \mathcal{M} \cup \{\perp\}} \sum_{a,b \in \mathbb{F}_p} p_{a,b} \left| \Pr(\text{Sim}_m^{(a,b)} = m') - \Pr(\text{Tamper}_m^{(a,b)} = m') \right| \\
&= \frac{1}{2} \sum_{a,b \in \mathbb{F}_p} p_{a,b} \sum_{m' \in \mathcal{M} \cup \{\perp\}} \left| \Pr(\text{Sim}_m^{(a,b)} = m') - \Pr(\text{Tamper}_m^{(a,b)} = m') \right| \\
&\leq \sum_{a,b \in \mathbb{F}_p} p_{a,b} \varepsilon / 2 = \varepsilon / 2,
\end{aligned}$$

where the last inequality follows from Corollary 2. Therefore, using triangle inequality,

$$\begin{aligned}
\Delta\left(\text{Tamper}_m^{(f,g)} ; \text{Sim}_m^{(f,g)}\right) &\leq \Delta\left(\text{Tamper}_m^{(f,g)} ; T\right) + \Delta(T ; S) + \Delta\left(S ; \text{Sim}_m^{(f,g)}\right) \\
&\leq \frac{\varepsilon}{2} + \frac{1}{p^2} + \frac{2}{p} \leq \varepsilon,
\end{aligned}$$

thus completing the proof of Theorem 5.

5 Proof of Theorem 3

We recall Theorem 3 for the convenience of the reader.

Theorem 3 *There exists universal constants $c, c' > 0$ such that the following holds. For any finite field \mathbb{F}_p of prime order, and any $n > c' \log^6 p$, let $L, R \in \mathbb{F}_p^n$ be uniform, and fix $f, g : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$. Then*

$$\Delta(\phi_{f,g}(L, R) ; \mathcal{D}) \leq 2^{-cn^{1/6}}.$$

We prove Theorem 3 in this section. Let us fix functions $f, g : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ and shorthand $\phi(L, R) = \phi_{f,g}(L, R)$. An important ingredient in the proof will be conditioning ϕ on various subsets of $\mathbb{F}_p^n \times \mathbb{F}_p^n$. We will use the following notation: for set $\mathcal{P} \subset \mathbb{F}_p^n \times \mathbb{F}_p^n$ let $\phi(L, R)|_{\mathcal{P}}$ denote the conditional distribution of $\phi(L, R)$ conditioned on $(L, R) \in \mathcal{P}$. Equivalently, it is the distribution of $\phi(L, R)$ for uniformly chosen $(L, R) \in \mathcal{P}$. We will typically be using this applied to product sets $\mathcal{P} = \mathcal{L} \times \mathcal{R}$ for $\mathcal{L}, \mathcal{R} \subseteq \mathbb{F}_p^n$.

We start with the following simple lemma, showing that it suffices to prove Theorem 3 for partitions of the ambient space.

Lemma 5 *Let $\mathcal{P} \subseteq \mathbb{F}_p^n \times \mathbb{F}_p^n$. Let $\mathcal{P}_1, \dots, \mathcal{P}_k$ be a partition of \mathcal{P} . Assume that for all $1 \leq i \leq k$,*

$$\Delta\left(\phi(L, R)|_{(L,R) \in \mathcal{P}_i} ; \mathcal{D}\right) \leq \varepsilon_i.$$

Then

$$\Delta\left(\phi(L, R)|_{(L,R) \in \mathcal{P}} ; \mathcal{D}\right) \leq \sum \varepsilon_i \frac{|\mathcal{P}_i|}{|\mathcal{P}|}.$$

Proof. The lemma follows immediately from the definitions. For all i let $D_i \in \mathcal{D}$ be such that $\Delta(\phi(L, R)|_{(L, R) \in \mathcal{P}_i}; D_i) \leq \varepsilon_i$. Let $p_i = |\mathcal{P}_i|/|\mathcal{P}|$ denote the probability that $(L, R) \in \mathcal{P}_i$ conditioned on $(L, R) \in \mathcal{P}$. Then $\phi(L, R)$ is $(\sum p_i \varepsilon_i)$ -close in statistical distance to $D \in \mathcal{D}$ given by $D[(a, b)] = \sum p_i D_i[(a, b)]$. \square

We next define a partition of $\mathbb{F}_p^n \times \mathbb{F}_p^n$ to which we will apply Lemma 5. Let $s = \lfloor \frac{n}{10} \rfloor$, and $t = \lfloor \frac{s^{1/6}}{c_1 \log p} \rfloor$, where c_1 is some constant that will be chosen later. Note that $s \gg t$. We choose the constant c' in the statement of Theorem 3 such that $t \geq 3$.

We first define a partition $\mathcal{L}_1, \dots, \mathcal{L}_a$ of \mathbb{F}_p^n based on f . Intuitively, \mathcal{L}_i for $1 \leq i < a$ will correspond to inputs on which f agrees with a popular linear function; and \mathcal{L}_a will be the remaining elements.

We define $\mathcal{L}_1, \dots, \mathcal{L}_a$ iteratively. For $i \geq 1$, given $\mathcal{L}_1, \dots, \mathcal{L}_{i-1}$, if there exists a linear map $A_i : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ for which

$$|\{x \in \mathbb{F}_p^n : f(x) = A_i x\} \setminus (\mathcal{L}_1 \cup \dots \cup \mathcal{L}_{i-1})| \geq p^{n-s},$$

then set \mathcal{L}_i to be $\{x \in \mathbb{F}_p^n : f(x) = A_i x\} \setminus (\mathcal{L}_1 \cup \dots \cup \mathcal{L}_{i-1})$. If no such linear map exists, set $a := i$, $\mathcal{L}_a := \mathbb{F}_p^n \setminus (\mathcal{L}_0 \cup \dots \cup \mathcal{L}_{a-1})$ and complete the process. Note we obtained a partition $\mathcal{L}_0, \dots, \mathcal{L}_a$ of \mathbb{F}_p^n with $a \leq p^s + 1$.

We next define a partition based on g to elements whose output is too popular; and the rest. For $y \in \mathbb{F}_p^n$ let $g^{-1}(y) = \{x \in \mathbb{F}_p^n : g(x) = y\}$ be the set of pre-images of y . Define

$$\mathcal{R}_0 := \{x \in \mathbb{F}_p^n : |g^{-1}(g(x))| \geq p^t\}.$$

and set $\mathcal{R}_1 := \mathbb{F}_p^n \setminus \mathcal{R}_0$. We define the following partition of $\mathbb{F}_p^n \times \mathbb{F}_p^n$:

$$\{\mathcal{P}_0, \dots, \mathcal{P}_a\} = \{\mathbb{F}_p^n \times \mathcal{R}_0, \mathcal{L}_1 \times \mathcal{R}_1, \dots, \mathcal{L}_a \times \mathcal{R}_1\}.$$

We will argue that for any part, either its probability is small, or the joint distribution of $\phi(L, R)$ conditioned on (L, R) belonging to it, is close to \mathcal{D} . We then apply Lemma 5 to obtain a proof of Theorem 3.

5.1 g is close to constant

We first analyze the distribution conditioned of $(L, R) \in \mathbb{F}_p^n \times \mathcal{R}_0$, that is on inputs x for which $g(x)$ has many preimages.

Lemma 6 $\Delta(\phi(L, R)|_{\mathbb{F}_p^n \times \mathcal{R}_0}; \mathcal{D}) \leq p^{-(t-1)/2}$.

Proof. Let $Y = \{y \in \mathbb{F}_p^n : |g^{-1}(y)| \geq p^t\}$. We can decompose \mathcal{R}_0 as the disjoint union over $y \in Y$ of $g^{-1}(y)$. By Lemma 5 it suffices to prove the lemma conditioned on $R \in g^{-1}(y)$ for all $y \in Y$. Fix such a $y \in Y$ and let $R_y = R|_{g(R)=y}$ denote the conditional random variable. Since by assumption $|g^{-1}(y)| \geq p^t$ and $L \in \mathbb{F}_p^n$ is uniform, using Lemma 1

$$\Delta(\langle (L, R_y), L \rangle; (U, L)) \leq p^{-(t-1)/2},$$

where $U \in \mathbb{F}_p$ is uniform independent of L, R_y . In particular,

$$\Delta(\langle (L, R_y), \langle f(L), g(R_y) \rangle \rangle; (U, \langle f(L), y \rangle)) \leq p^{-(t-1)/2}.$$

This concludes the proof since $(U, \langle f(L), y \rangle)$ is in the convex combination of $\{(U, a) : a \in \mathbb{F}_p\}$ which is contained in \mathcal{D} . \square

5.2 f is close to linear

Fix $1 \leq i < a$. We analyze in this subsection the joint distribution for $(L, R)|_{\mathcal{L}_i \times \mathcal{R}_1}$. Let $A : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ be a linear map so that for all $x \in \mathcal{L}_i$, $f(x) = Ax$.

Lemma 7 *If $|\mathcal{L}_i \times \mathcal{R}_1| \geq p^{2n-2s}$ then*

$$\Delta(\phi(L, R)|_{\mathcal{L}_i \times \mathcal{R}_1} ; \mathcal{D}) \leq 2p^{-s}.$$

Proof. Let $L' \in \mathcal{L}_i, R' \in \mathcal{R}_1$ be uniform and independent. Note that

$$\langle f(L'), g(R') \rangle = \langle AL', g(R') \rangle = \langle L', A^T g(R') \rangle.$$

If $(\langle L', R' \rangle, \langle f(L'), g(R') \rangle)$ is p^{-s} -close to $U_{\mathbb{F}_p^2}$ we are done since the uniform distribution is in \mathcal{D} . If not, then by Claim 1 there exist $a, b \in \mathbb{F}_p$, not both zero, such that

$$\Delta(\langle L', aR' + bA^T g(R') \rangle ; U_{\mathbb{F}_p}) \geq p^{-2-s}.$$

Now, by assumption, L' is uniform over a set of size at least p^{n-s} . Assume that $\mathbf{H}_\infty(aR' + bA^T g(R')) = k \log p$. Then, using Lemma 1 gives

$$\Delta(\langle L', aR' + bA^T g(R') \rangle ; U_{\mathbb{F}_p}) \leq p^{-(k-s-1)/2}.$$

This means that $k \leq 3s+4 \leq 4s$. So, there exist $y \in \mathbb{F}_p^n$ and a subset $\mathcal{R}'_1 \subset \mathcal{R}_1$ of size $|\mathcal{R}'_1| \geq |\mathcal{R}_1| \cdot p^{-4s}$ such that

$$ax + bA^T g(x) = y \quad \forall x \in \mathcal{R}'_1.$$

We clearly cannot have $b = 0$ since $ax = y$ can hold only for one value of x . So, as $b \neq 0$ we can rewrite (and rename the constants for convenience) as

$$A^T g(x) = a_1 x + y_1 \quad \forall x \in \mathcal{R}'_1.$$

Let $\mathcal{R}_2 = \mathcal{R}_1 \setminus \mathcal{R}'_1$. We repeat this process with \mathcal{R}_1 replaced by \mathcal{R}_2 to get a set $\mathcal{R}'_2 \subset \mathcal{R}_2$ of size $|\mathcal{R}'_2| \geq |\mathcal{R}_2| \cdot p^{-4s}$ and $y_2 \in \mathbb{F}_p^n$ such that

$$A^T g(x) = a_2 x + y_2 \quad \forall x \in \mathcal{R}'_2.$$

We continue this process to get $\mathcal{R}_3, \dots, \mathcal{R}_b$ until $|\mathcal{R}_b| < p^{-s}|\mathcal{R}_1|$ or until $(L, R)|_{\mathcal{L}_i \times \mathcal{R}_b}$ is p^{-s} close to $U_{\mathbb{F}_p \times \mathbb{F}_p}$. Note that for $j < b$ we have $|\mathcal{R}'_j| \geq |\mathcal{R}_j| \cdot p^{-4s} \geq |\mathcal{R}_1| p^{-5s}$.

Consider the partition of $\mathcal{L}_i \times \mathcal{R}_1$ as $\{\mathcal{L}_i \times \mathcal{R}'_1, \dots, \mathcal{L}_i \times \mathcal{R}'_{b-1}, \mathcal{L}_i \times \mathcal{R}_b\}$. We argue next that all the partitions, except for perhaps the last one, induce distributions very close to \mathcal{D} .

Claim 6 *For $1 \leq j < b$,*

$$\Delta(\phi(L, R)|_{\mathcal{L}_i \times \mathcal{R}'_j} ; \mathcal{D}) \leq p^{-s}.$$

Proof. Let $L^* \in \mathcal{L}_i$ and $R^* \in \mathcal{R}'_j$ be independent and uniform. We know that $\langle f(L^*), g(R^*) \rangle = \langle L^*, A^T g(R^*) \rangle = a_j \langle L^*, R^* \rangle + \langle L^*, y_j \rangle$. Moreover, we know that $|\mathcal{L}_i \times \mathcal{R}'_j| \geq |\mathcal{L}_i \times \mathcal{R}_1| p^{-5s} \geq p^{2n-7s}$. So by Lemma 1 we have that

$$\Delta(\langle L^*, R^* \rangle, L^* ; U, L^*) \leq p^{(n-7s-1)/2} \leq p^{-s}$$

where the last inequality follows from our assumption that $n \geq 10s$. So

$$\Delta(\langle L^*, R^* \rangle, \langle f(L^*), g(R^*) \rangle ; U, a_j U + X) \leq p^{-s}$$

where $U \in \mathbb{F}_p$ is uniform and $X \in \mathbb{F}_p$ is independent from U and distributed like $\langle L^*, y_j \rangle$. As this distribution is in \mathcal{D} this conclude the proof. \square

For all $j < b$ we have that the joint distribution of $\phi(L, R)|_{\mathcal{L}_i \times \mathcal{R}'_j}$ is p^{-s} close to \mathcal{D} . Also, we know that either $\frac{|\mathcal{L}_i \times \mathcal{R}_b|}{|\mathcal{L}_i \times \mathcal{R}_1|} \leq p^{-s}$; or that $(L, R)|_{\mathcal{L}_i \times \mathcal{R}_b}$ is p^{-s} close to $U_{\mathbb{F}_p \times \mathbb{F}_p}$, which implies $\Delta(\phi(L, R)|_{\mathcal{L}_i \times \mathcal{R}_b} ; \mathcal{D}) \leq p^{-s}$. Hence, the lemma follows by Lemma 5. \square

5.3 f is far from linear and g is far from constant

The last partition we need to analyze is $\mathcal{L}_a \times \mathcal{R}_1$, corresponding to the case where f is far from linear and g is far from constant. For this, we need the following result that can be seen as a generalization of the linearity test from [24] and that is discussed and proved in Section 6.

Theorem 6 *Let p be a prime, and $n \in \mathbb{N}$. For any $\varepsilon = \varepsilon(n, p) > 0$, $\gamma_1 = \gamma_1(n, p) \leq 1$, $\gamma_2 = \gamma_2(n, p) \geq 1$, the following is true. For any function $f : \mathbb{F}_p^n \mapsto \mathbb{F}_p^n$, let $\mathcal{A} \subseteq \{(x, f(x)) : x \in \mathbb{F}_p^n\} \subseteq \mathbb{F}_p^{2n}$. If $|\mathcal{A}| \geq \gamma_1 \cdot |\mathbb{F}_p^n|$ and there exists some set \mathcal{B} such that $|\mathcal{B}| \leq \gamma_2 \cdot p^n$, and*

$$\Pr_{a, a' \in \mathcal{A}} [a - a' \in \mathcal{B}] \geq \varepsilon,$$

then there exists a linear map $M : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ such that

$$\Pr_{(x, f(x)) \in \mathcal{A}} [f(x) = Mx] \geq p^{-O(\log^6(\frac{\gamma_2}{\gamma_1 \varepsilon}))}.$$

We will now show that, $\phi(L, R)|_{\mathcal{L}_a \times \mathcal{R}_1}$ is close to uniform over $\mathbb{F}_p \times \mathbb{F}_p$.

Lemma 8 *If $|\mathcal{L}_a \times \mathcal{R}_1| \geq p^{2n-t}$, then*

$$\Delta(\phi(L, R)|_{\mathcal{L}_a \times \mathcal{R}_1} ; U_{\mathbb{F}_p \times \mathbb{F}_p}) \leq p^{-t}.$$

In particular,

$$\Delta(\phi(L, R)|_{\mathcal{L}_a \times \mathcal{R}_1} ; \mathcal{D}) \leq p^{-t}.$$

Proof. Let $L' \in \mathcal{L}_a, R' \in \mathcal{R}_1$ be uniform and independent. We assume that $\phi(L', R')$ is not p^{-t} -close to $U_{\mathbb{F}_p \times \mathbb{F}_p}$, as otherwise the result trivially holds. Then, by Claim 1 there exist $a, b \in \mathbb{F}_p$, not both zero, so that $\Delta(a\langle L', R' \rangle + b\langle f(L'), g(R') \rangle ; U_{\mathbb{F}_p}) \geq p^{-t-2}$. Define functions $F, G : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^{2n}$ as follows

$$F(x) = (x, f(x)), \quad G(y) = (ay, bg(y)).$$

We have that $\Delta(\langle F(L'), G(R') \rangle ; U_{\mathbb{F}_p}) \geq p^{-t-2}$. Applying Claim 2, we get that for (L'', R'') i.i.d to (L', R') we have

$$\Pr[\langle F(L'), G(R') \rangle = \langle F(L''), G(R'') \rangle] \geq \frac{1}{p} + \frac{1}{p^{2t+5}}.$$

Applying Claim 3 with $X = F(L'), Y = G(R'), X' = F(L''), Y' = G(R'')$ we get that

$$\Pr[\langle F(L') - F(L''), G(R') \rangle = 0] \geq \frac{1}{p} + \frac{1}{p^{2t+5}}.$$

Define

$$\mathcal{B} := \left\{ \alpha \in \mathbb{F}_p^{2n} : \Pr[\langle \alpha, G(R') \rangle = 0] \geq \frac{1}{p} + \frac{1}{p^{2t+6}} \right\}.$$

Let $B \in \mathcal{B}$ be uniform. Then $\Delta(\langle B, G(R') \rangle, U_{\mathbb{F}_p}) \geq \frac{1}{p^{2t+6}}$. Also, since $g(y)$ has at most p^t preimages for any $y \in \mathbb{F}_p^n$, $G(R')$ has min-entropy at least $\log(|\mathcal{R}_1|p^{-t}) \geq (n-2t)\log p$. Hence, by Lemma 1, we have $\mathbf{H}_\infty(B) \leq (n+6t+13) \cdot \log p$, which implies $|\mathcal{B}| \leq p^{n+6t+13}$. Furthermore, we have that

$$\Pr[\langle F(L') - F(L''), G(R') \rangle = 0] \leq \Pr[F(L') - F(L'') \in \mathcal{B}] + \frac{1}{p} + \frac{1}{p^{2t+6}}.$$

So we must have that

$$\Pr[F(L') - F(L'') \in \mathcal{B}] \geq \frac{1}{p^{2t+5}} - \frac{1}{p^{2t+6}} \geq \frac{1}{p^{2t+6}}.$$

Thus, using Theorem 6, we get that there exists a linear map $M : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ for which

$$\Pr_{x \in \mathbb{F}_p^n} [Mx = f(x)] \geq p^{-O(t^6 \log^6 p)}.$$

This violates the definition of \mathcal{L}_a whenever $s \geq C(t^6 \log^6 p)$ for a big enough constant C .⁵ \square

5.4 Putting things together

In this section, we combine the results of Lemmas 6, 7, and 8, and use Lemma 5 to conclude the proof of Theorem 3.

Proof. Consider the partition $\mathcal{P}_0, \dots, \mathcal{P}_a$ of $\mathbb{F}_p^n \times \mathbb{F}_p^n$ as defined earlier. In the following, let p_i denote $\frac{|\mathcal{P}_i|}{p^{2n}}$. Note that if for any α, β, i , we have a statement of the form: If $p_i \geq \alpha$, then $\Delta(\phi(L, R)|_{\mathcal{P}_i}; \mathcal{D}) \leq \beta$. Then this statement implies that

$$\Delta(\phi(L, R)|_{\mathcal{P}_i}; \mathcal{D}) \cdot p_i \leq \alpha + \beta \cdot p_i.$$

Thus, using Lemma 5, and the results of Lemmas 6, 7, and 8, we get that

$$\begin{aligned} \Delta(\phi_{f,g}(L, R); \mathcal{D}) &\leq \Delta(\phi(L, R)|_{\mathcal{P}_0}; \mathcal{D}) \cdot p_0 + \sum_{i=1}^{a-1} \Delta(\phi(L, R)|_{\mathcal{P}_i}; \mathcal{D}) \cdot p_i \\ &\quad + \Delta(\phi(L, R)|_{\mathcal{P}_a}; \mathcal{D}) \cdot p_a \\ &\leq \frac{1}{p^{(t-1)/2}} \cdot p_0 + \sum_{i=1}^{a-1} \left(\frac{1}{p^{2s}} + \frac{2}{p^s} \cdot p_i \right) \\ &\quad + \left(\frac{1}{p^t} + \frac{1}{p^t} \cdot p_a \right) \\ &\leq \frac{1}{p^{(t-1)/2}} \sum_{i=0}^a p_i + \frac{p^s}{p^{2s}} + \frac{1}{p^t} \\ &\leq \frac{2}{p^{(t-1)/2}} \leq 2^{-cn^{1/6}}, \end{aligned}$$

for some constant c . \square

⁵The constant C here determines the choice of the constant c_1 used while defining the parameter t .

6 Generalized linearity testing

We now take a detour and prove Theorem 6 that generalizes the linearity test from [24] for large fields of prime order. The linearity test in [24] for checking whether a function $f : \mathbb{F}_p^n \mapsto \mathbb{F}_p^n$ does the following: It picks $x, x' \in \mathbb{F}_p^n$ uniformly at random and accept if and only if $f(x - x') = f(x) - f(x')$. Clearly, this test always accepts if f is linear, and it was shown for $p = 2$ that the test rejects with high probability if f is sufficiently far from linear. More precisely, it was shown that for any ε , if $\Pr_{x, x' \in \mathbb{F}_p^n}(f(x) - f(x') = f(x - x')) \geq \varepsilon$, then there exists a matrix $M \in \mathbb{F}_p^{n \times n}$ such that $\Pr(f(x) = Mx) \geq \varepsilon'$. The dependence of ε' on ε in the proof of [24] was exponential.

We show here a more general and improved result that we stated in Section 5.3. The key difference between this proof and the proof of [24] is the use of a recent result by Sanders [25].

Theorem 6 *Let p be a prime, and $n \in \mathbb{N}$. For any $\varepsilon = \varepsilon(n, p) > 0$, $\gamma_1 = \gamma_1(n, p) \leq 1$, $\gamma_2 = \gamma_2(n, p) \geq 1$, the following is true. For any function $f : \mathbb{F}_p^n \mapsto \mathbb{F}_p^n$, let $\mathcal{A} \subseteq \{(x, f(x)) : x \in \mathbb{F}_p^n\} \subseteq \mathbb{F}_p^{2n}$. If $|\mathcal{A}| \geq \gamma_1 \cdot |\mathbb{F}_p^n|$ and there exists some set \mathcal{B} such that $|\mathcal{B}| \leq \gamma_2 \cdot p^n$, and*

$$\Pr_{a, a' \in \mathcal{A}}[a - a' \in \mathcal{B}] \geq \varepsilon$$

, then there exists a linear map $M : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ such that

$$\Pr_{(x, f(x)) \in \mathcal{A}}[f(x) = Mx] \geq p^{-O(\log^6(\frac{\gamma_2}{\gamma_1 \varepsilon}))}.$$

This result improves the linearity test from [24] in several ways (i) The dependence of ε' on ε is only quasi-polynomial instead of exponential. (ii) This result is proven for any finite field of prime order. While the ideas of [24] generalize for larger fields, it results in an exponential dependence of ε' on p in addition to that on ε (iii) The linearity test is a special case of our result since we can obtain it by setting $\mathcal{B} = \mathcal{A} = \{(x, f(x)) : x \in \mathbb{F}_p^n\}$ (and hence, $\gamma_1 = \gamma_2 = 1$).

For giving a proof of this theorem, we need the following results from additive combinatorics. First we introduce some notation. Let $\mathcal{A}' \subset \mathbb{F}_p^n$ be a set. We denote by $\mathcal{A}' - \mathcal{A}' = \{a - a' | a, a' \in \mathcal{A}'\}$ the difference set of \mathcal{A}' . We denote by $\text{span}(\mathcal{A}')$ the linear subspace over \mathbb{F}_p spanned by \mathcal{A}' .

The following result is due to Balog, Szemerédi, and Gowers [1, 16]. The current formulation is from a survey of Viola [26], Theorem 3.1. The statement given in [26] is for the case when the field is $p = 2$, and $\mathcal{A} = \mathcal{B}$, but the proof is essentially the same.

Lemma 9 *Let $\mathcal{A}, \mathcal{B} \subseteq \mathbb{F}_p^n$. If $\Pr_{a, a' \in \mathcal{A}}[a - a' \in \mathcal{B}] \geq \varepsilon$ then there exists $\mathcal{A}' \subseteq \mathcal{A}$ of size $|\mathcal{A}'| \geq (\varepsilon/3) \cdot |\mathcal{A}|$ such that $|\mathcal{A}' - \mathcal{A}'| \leq \frac{6^8 |\mathcal{B}|^4}{\varepsilon^8 |\mathcal{A}|^3}$.*

The following result is of Sanders [25].

Lemma 10 *Let $\mathcal{A}' \subset \mathbb{F}_p^n$ be a set such that $|\mathcal{A}' - \mathcal{A}'| \leq K|\mathcal{A}'|$. Then there exists $\mathcal{A}'' \subseteq \mathcal{A}'$ such that $|\mathcal{A}''| \geq p^{-O(\log^6 K)} |\mathcal{A}'|$ such that $|\text{span}(\mathcal{A}'')| \leq |\mathcal{A}'|$.*

Finally, we need the following fact in linear algebra. Its proof can be found e.g. in [26], Lemma 5.1.

Lemma 11 *Let $f : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ be a function. Let $\mathcal{A}'' \subset \mathbb{F}_p^n \times \mathbb{F}_p^n$ be a set such that*

$$\mathcal{A}'' \subseteq \{(x, f(x)) : x \in \mathbb{F}_p^n\}.$$

Assume furthermore that

$$\varepsilon p^n \leq |\mathcal{A}''| \leq |\text{span}(\mathcal{A}'')| \leq \frac{p^n}{\varepsilon}.$$

Then there exists a linear map $M : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ such that

$$\Pr_{(x, f(x)) \in \mathcal{A}''} [f(x) = Mx] \geq \frac{\varepsilon^3}{2p}.$$

Now we have the tools to complete the proof of Theorem 6.

Proof. First, we apply Lemma 9. We get that there exists a set $\mathcal{A}' \subset \mathcal{A}$ of size $|\mathcal{A}'| \geq \Omega(\varepsilon \gamma_1 p^n)$ for which $|\mathcal{A}' - \mathcal{A}'| = O((\gamma_2^4 / \gamma_1^3 \varepsilon^8) p^n)$. Applying Lemma 10 we get that there exists a subset $\mathcal{A}'' \subset \mathcal{A}'$ of size $|\mathcal{A}''| \leq p^{-O(\log^6(\frac{\gamma_2}{\gamma_1 \varepsilon}))} |\mathcal{A}'| = p^{-O(\log^6(\frac{\gamma_2}{\gamma_1 \varepsilon}))} \cdot p^n$ for which $|\text{span}(\mathcal{A}'')| \leq |\mathcal{A}'|$. Applying Lemma 11 we get that there exists a linear map $M : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ for which $\Pr_{(x, f(x)) \in \mathcal{A}''} [Mx = f(x)] \geq p^{-O(\log^6(\frac{\gamma_2}{\gamma_1 \varepsilon}))}$, which implies $\Pr_{(x, f(x)) \in \mathcal{A}} [Mx = f(x)] \geq p^{-O(\log^6(\frac{\gamma_2}{\gamma_1 \varepsilon}))}$. \square

7 Conclusions and Open Problems

We give an encoding scheme for k -bit messages to $\mathbb{F}_p^n \times \mathbb{F}_p^n$ that is ε -non-malleable in the split state model. The messages are encoded in a space of size $O(n \log p)$. For our security proof, which is based on Theorem 3, we need n to be $\Omega(\log^6 p)$, and p is $2^{\Omega((k + \log 1/\varepsilon) \log(k + \log 1/\varepsilon))}$, and thus the size of the space in which messages are encoded is $O((k + \log 1/\varepsilon)^7 \log^7(k + \log 1/\varepsilon))$. We believe that there is a possibility of reducing the size of both p and n .

The choice of p is governed by the construction of an affine-evasive set in Section 4.1. It might very well be possible that we can find an affine-evasive set $S \subset \mathbb{F}_p$ where $p, |S|$ are polynomially related. In particular, if we pick the set S as a random subset of size \sqrt{p} from \mathbb{F}_p , then it is not difficult to show that this set is affine-evasive with high probability for $p = 2^{2(k + \log 1/\varepsilon)}$. However, we cannot choose the set randomly for our encoding scheme since the encoding and decoding algorithm have to be efficient.

Also, although the proof does not show this, the encoding scheme might actually be secure for a much smaller value of n (conceivably, even for n being a big enough constant.) In particular, if the bound on Sanders' result (Lemma 10) were to be improved, it would immediately imply the security of our scheme for a smaller value of n . It should however be noted that the security does not follow immediately for the case when n is a constant even if we assume the well-known polynomial Freiman-Ruzsa conjecture [17], which is a stronger assumption than what was proven by Sanders. Assuming the PFR conjecture, we can show that the coding scheme is secure for n being $\Theta(\log p)$. This follows immediately by an alternative version of Lemma 10 that follows from the PFR conjecture.

Acknowledgments: We thank Oded Regev, Tom Sanders, and Terence Tao for useful discussions, especially related to Section 5.3 of the paper. We would also like to thank Stefan Dziembowski, Tomasz Kazana, and Maciej Obremski for sharing their recent work on non-malleable codes for 1-bit messages [12].

References

- [1] A. Balog and E. Szemerédi. A statistical theorem for set addition. *Combinatorica*, 14(3):263–268, 1994.
- [2] X. Boyen, Y. Dodis, J. Katz, R. Ostrovsky, and A. Smith. Secure remote authentication using biometric data. In R. Cramer, editor, *Advances in Cryptology—EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 147–163. Springer-Verlag, 2005.
- [3] H. Chabanne, G. Cohen, J. Flori, and A. Patey. Non-malleable codes from the wire-tap channel. In *Information Theory Workshop (ITW), 2011 IEEE*, pages 55–59. IEEE, 2011.
- [4] H. Chabanne, G. Cohen, and A. Patey. Secure network coding and non-malleable codes: Protection against linear tampering. In *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, pages 2546–2550. IEEE, 2012.
- [5] S. G. Choi, A. Kiayias, and T. Malkin. Bitr: built-in tamper resilience. In *Advances in Cryptology—ASIACRYPT 2011*, pages 740–758. Springer, 2011.
- [6] R. Cramer, Y. Dodis, S. Fehr, C. Padro, and D. Wichs. Detection of algebraic manipulation with applications to robust secret sharing and fuzzy extractors. In *EUROCRYPT 2008*, April 2008. To Appear.
- [7] F. Davì, S. Dziembowski, and D. Venturi. Leakage-resilient storage. In J. A. Garay and R. D. Prisco, editors, *SCN*, volume 6280 of *Lecture Notes in Computer Science*, pages 121–137. Springer, 2010.
- [8] A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero knowledge. In *Advances in Cryptology—CRYPTO 2001*, pages 566–598. Springer, 2001.
- [9] Y. Dodis, J. Katz, L. Reyzin, and A. Smith. Robust fuzzy extractors and authenticated key agreement from close secrets. In C. Dwork, editor, *Advances in Cryptology—CRYPTO 2006*, volume 4117 of *LNCS*, pages 232–250. Springer-Verlag, 20–24 Aug. 2006.
- [10] Y. Dodis and D. Wichs. Non-malleable extractors and symmetric key cryptography from weak secrets. In M. Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, pages 601–610, Bethesda, MD, USA, 2009. ACM.
- [11] D. Dolev, C. Dwork, and M. Naor. Nonmalleable cryptography. *SIAM*, 30:391–437, 2000.
- [12] S. Dziembowski, T. Kazana, and M. Obremski. Non-malleable codes from two-source extractors. In *Advances in Cryptology—CRYPTO 2013*. Springer, 2013.
- [13] S. Dziembowski and K. Pietrzak. Leakage-resilient cryptography. In *49th Symposium on Foundations of Computer Science*, pages 293–302, Philadelphia, PA, USA, Oct. 25–28 2008. IEEE Computer Society.
- [14] S. Dziembowski, K. Pietrzak, and D. Wichs. Non-malleable codes. In A. C.-C. Yao, editor, *ICS*, pages 434–452. Tsinghua University Press, 2010.
- [15] R. Gennaro, A. Lysyanskaya, T. Malkin, S. Micali, and T. Rabin. Algorithmic Tamper-Proof (ATP) security: Theoretical foundations for security against hardware tampering. In M. Naor, editor, *First Theory of Cryptography Conference — TCC 2004*, volume 2951 of *LNCS*, pages 258–277. Springer-Verlag, Feb. 19–21 2003.
- [16] T. Gowers. A new proof of szemerédi’s theorem for arithmetic progression of length four. *Geom. Func. Anal.*, 8(3):529–551, 1998.
- [17] B. Green. Finite field models in additive number theory. *Surveys in Combinatorics*, pages 1–29, 2005.
- [18] Y. Ishai, M. Prabhakaran, A. Sahai, and D. Wagner. Private circuits II: Keeping secrets in tamperable circuits. In S. Vaudenay, editor, *Advances in Cryptology—EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 308–327. Springer-Verlag, 2006.
- [19] Y. Ishai, A. Sahai, and D. Wagner. Private circuits: Securing hardware against probing attacks. In D. Boneh, editor, *Advances in Cryptology—CRYPTO 2003*, volume 2729 of *LNCS*. Springer-Verlag, 2003.
- [20] Y. T. Kalai, B. Kanukurthi, and A. Sahai. Cryptography with tamperable and leaky memory. In *Advances in Cryptology—CRYPTO 2011*, pages 373–390. Springer, 2011.

- [21] C.-J. Lee, C.-J. Lu, S.-C. Tsai, and W.-G. Tzeng. Extracting randomness from multiple independent sources. *Information Theory, IEEE Transactions on*, 51(6):2224–2227, 2005.
- [22] F.-H. Liu and A. Lysyanskaya. Tamper and leakage resilience in the split-state model. In *Advances in Cryptology–CRYPTO 2012*, pages 517–532. Springer, 2012.
- [23] M. Naor and G. Segev. Public-key cryptosystems resilient to key leakage. In S. Halevi, editor, *Advances in Cryptology - CRYPTO 2009*, volume 5677 of *LNCS*, pages 18–35. Springer-Verlag, 2009.
- [24] A. Samorodnitsky. Low-degree tests at large distances. In *ACM symposium on Theory of computing*, pages 506–515. ACM, 2007.
- [25] T. Sanders. On the bogolyubov-ruzsa lemma, anal. *PDE*, 5:627–655, 2012.
- [26] E. Viola. Selected results in additive combinatorics: An exposition. *Theory of Computing Library, Graduate Surveys series*, 3:1–15, 2011.

8 Appendix

In this section, we prove the claims we stated in Section 2. The proofs use the following simple inequality: $(\sum_{i=1}^n |x_i|)^2 \leq n \sum_{i=1}^n |x_i|^2$.

Claim 1 *Let $X = (X_1, X_2) \in \mathbb{F}_p \times \mathbb{F}_p$ be a random variable. Assume that for all $a, b \in \mathbb{F}_p$ not both zero, $\Delta(aX_1 + bX_2 ; U_{\mathbb{F}_p}) \leq \varepsilon$. Then $\Delta((X_1, X_2) ; U_{\mathbb{F}_p^2}) \leq \varepsilon p^2$.*

Proof. Let $\omega = e^{2\pi i/p}$. Then for any $a, b \in \mathbb{F}_p$ not both zero,

$$\mathbf{E}[\omega^{aX_1 + bX_2}] = \sum_{c \in \mathbb{F}_p} \omega^c \Pr[aX_1 + bX_2 = c] = \sum_{c \in \mathbb{F}_p} \omega^c \left(\Pr[aX_1 + bX_2 = c] - \frac{1}{p} \right).$$

Hence $|\mathbf{E}[\omega^{aX_1 + bX_2}]| \leq \varepsilon$. Let $p_{c,d} = \Pr[X = (c, d)]$. Then by Parseval's identity

$$\sum_{c,d \in \mathbb{F}_p} \left(p_{c,d} - \frac{1}{p} \right)^2 = \sum_{(a,b) \neq (0,0)} \mathbf{E}_X[\omega^{aX_1 + bX_2}]^2 \leq \varepsilon^2 p^2,$$

and

$$\left(\sum_{c,d \in \mathbb{F}_p} \left| p_{c,d} - \frac{1}{p} \right| \right)^2 \leq p^2 \sum_{c,d \in \mathbb{F}_p} \left(p_{c,d} - \frac{1}{p} \right)^2 \leq \varepsilon^2 p^4.$$

□

Claim 2 *Let $X \in \mathbb{F}_p$ be a random variable. Assume that $\Delta(X ; U_{\mathbb{F}_p}) \geq \varepsilon$. Then if X' is an independent and i.i.d copy of X then*

$$\Pr[X = X'] \geq \frac{1 + \varepsilon^2}{p}.$$

Proof. Let $p_x = \Pr[X = x]$ for $x \in \mathbb{F}_p$. Then

$$\Pr[X = X'] - \frac{1}{p} = \sum_{x \in \mathbb{F}_p} \left(p_x - \frac{1}{p} \right)^2 \geq \frac{1}{p} \left(\sum_{x \in \mathbb{F}_p} \left| p_x - \frac{1}{p} \right| \right)^2 \geq \frac{\varepsilon^2}{p}.$$

□

Claim 3 Let $Z = (X, Y) \in \mathbb{F}_p^n \times \mathbb{F}_p^n$ be a random variable, and let $Z' = (X', Y')$ be an i.i.d copy of Z . Then

$$\Pr[\langle X, Y \rangle = \langle X', Y' \rangle] \leq \Pr[\langle X, Y \rangle = \langle X', Y \rangle].$$

Proof. We would use the following identity: for a random variable $R \geq 0$ we have $\mathbf{E}[R]^2 \leq \mathbf{E}[R^2]$. We would actually prove a stronger inequality. For any $a \in \mathbb{F}_p$,

$$\Pr[\langle X, Y \rangle = \langle X', Y' \rangle = a] \leq \Pr[\langle X, Y \rangle = \langle X', Y \rangle = a].$$

Fix $a \in \mathbb{F}_p$ and define $f(x, y) = 1_{\langle x, y \rangle = a}$. Then

$$\begin{aligned} \Pr[\langle X, Y \rangle = \langle X', Y' \rangle = a] &= \Pr[\langle X, Y \rangle = a]^2 = (\mathbf{E}_{X, Y} f(X, Y))^2 \\ &\leq \mathbf{E}_Y (\mathbf{E}_X f(X, Y))^2 = \Pr[\langle X, Y \rangle = \langle X', Y \rangle = a]. \end{aligned}$$

□

Claim 4 Let $X_1, X_2, Y_1, Y_2 \in \mathcal{A}$ be random variables such that $\Delta((X_1, X_2); (Y_1, Y_2)) \leq \varepsilon$. Then, for any non-empty set $\mathcal{A}_1 \subseteq \mathcal{A}$, we have

$$\Delta(X_2 \mid X_1 \in \mathcal{A}_1; Y_2 \mid Y_1 \in \mathcal{A}_1) \leq \frac{2\varepsilon}{\Pr(X_1 \in \mathcal{A}_1)}.$$

Proof.

$$\begin{aligned} \Delta(X_2 \mid X_1 \in \mathcal{A}_1; Y_2 \mid Y_1 \in \mathcal{A}_1) &= \frac{1}{2} \sum_{x \in \mathcal{A}} \left| \Pr(X_2 = x \mid X_1 \in \mathcal{A}_1) - \Pr(Y_2 = x \mid Y_1 \in \mathcal{A}_1) \right| \\ &\leq \frac{1}{2} \sum_{x \in \mathcal{A}} \left(\left| \frac{\Pr(X_2 = x \wedge X_1 \in \mathcal{A}_1)}{\Pr(X_1 \in \mathcal{A}_1)} - \frac{\Pr(Y_2 = x \wedge Y_1 \in \mathcal{A}_1)}{\Pr(Y_1 \in \mathcal{A}_1)} \right| \right. \\ &\quad \left. + \left| \frac{\Pr(Y_2 = x \wedge Y_1 \in \mathcal{A}_1)}{\Pr(Y_1 \in \mathcal{A}_1)} - \frac{\Pr(X_2 = x \wedge X_1 \in \mathcal{A}_1)}{\Pr(X_1 \in \mathcal{A}_1)} \right| \right) \\ &\leq \frac{\varepsilon}{\Pr(X_1 \in \mathcal{A}_1)} + \frac{\varepsilon \cdot \sum_{x \in \mathcal{A}} \Pr(Y_1 \in \mathcal{A}_1 \wedge Y_2 = x)}{\Pr(Y_1 \in \mathcal{A}_1) \cdot \Pr(X_1 \in \mathcal{A}_1)} \\ &= \frac{2\varepsilon}{\Pr(X_1 \in \mathcal{A}_1)}. \end{aligned}$$

□