

Authentication in the Bounded Storage Model

Yevgeniy Dodis*
NYU

Willy Quach†
Northeastern

Daniel Wichs‡
Northeastern and NTT Research

June 1, 2022

Abstract

We consider the streaming variant of the Bounded Storage Model (BSM), where the honest parties can stream large amounts of data to each other, while only maintaining a small memory of size n . The adversary also operates as a streaming algorithm, but has a much larger memory size $m \gg n$. The goal is to construct unconditionally secure cryptographic schemes in the BSM, and prior works did so for symmetric-key encryption, key agreement, oblivious transfer and multiparty computation. In this work, we construct *message authentication and signatures* in the BSM.

First, we consider the symmetric-key setting, where Alice and Bob share a small secret key stored. Alice can authenticate arbitrarily many messages to Bob by streaming long authentication tags of size $k > m$, while ensuring that the tags can be generated and verified using only n bits of memory. We show a solution using local extractors (Vadhan; JoC '04), which allows for up to exponentially large adversarial memory $m = 2^{O(n)}$, and has tags of size $k = O(m)$.

Second, we consider the same setting as above, but now additionally require each individual tag to be small, of size $k \leq n$. We show a solution is still possible when the adversary's memory is $m = O(n^2)$, which is optimal. Our solution relies on a space lower bound for leaning parities (Raz; FOCS '16).

Third, we consider the public-key signature setting. A signer Alice initially streams a long verification key over an authentic channel, while only keeping a short signing key in her memory. A verifier Bob receives the streamed verification key and generates a short verification digest that he keeps in his memory. Later, Alice can sign arbitrarily many messages using her signing key by streaming large signatures to Bob, who can verify them using his verification digest. We show a solution for $m = O(n^2)$, which we show to be optimal. Our solution relies on a novel entropy lemma, of independent interest. We show that, if a sequence of blocks has sufficiently high min-entropy, then a large fraction of individual blocks must have high min-entropy. Naive versions of this lemma are false, but we show how to patch it to make it hold.

*E-mail: dodis@cs.nyu.edu. Partially supported by gifts from VMware Labs and Google, and NSF48 grants 1815546 and 2055578.

†E-mail: quach.w@northeastern.edu

‡E-mail: wichs@ccs.neu.edu. Daniel Wichs is partially supported by NSF grants CNS-1750795, CNS-2055510 and the Alfred P. Sloan Research Fellowship.

Contents

1	Introduction	1
1.1	Our Results	3
1.2	Our Techniques	4
2	Preliminaries	8
2.1	Information Theory	8
3	Definitions	10
3.1	Message Authentication Codes	10
3.2	Signatures	12
4	Block Entropy Lemma	13
5	MAC with Long Tags	16
5.1	Construction	16
6	MAC with Short Tags	20
6.1	Hardness of Learning Parities	20
6.2	Construction without Verification Queries	21
6.3	Domain Extension for MACs	23
6.4	Efficiently Recognizing Prior Tagging Queries	24
6.5	Upgrading to Security with Verification Queries	28
7	Public-Key Signatures	30
7.1	Set Key-Agreement Protocol	30
7.2	From Set Key Agreement to Signatures	34
7.3	Lower Bound for Signatures	36

1 Introduction

It is well known that there are strong restrictions on what cryptography can achieve without imposing any limits on the adversary’s resources. In particular, Shannon showed that Alice and Bob cannot communicate secretly over an insecure channel, unless they share a secret key that is at least as large as the total size of the messages exchanged. Similarly, Alice and Bob cannot ensure the authenticity of their communication over an insecure channel, unless they share a secret key that is at least as large as the total number of messages exchanged. In either case, no security is possible in the public-key setting, when Alice and Bob have no shared secrets to begin with. Traditionally, cryptography has focused on overcoming the above limitations by restricting the adversary Eve to run in polynomial time. Unfortunately, the security of such schemes relies on unproven computational hardness assumptions, and at the very least, requires that $P \neq NP$.

The *Bounded Storage Model (BSM)* [Mau92] offers an alternative by bounding the attacker’s space instead of her run time. The goal is to construct unconditionally secure cryptographic schemes in the BSM, without relying on any unproven assumptions. A long series of prior works [CM97, CCM98, Din01, HCR02, ADR02, DM02, Lu02, Vad04, DHRS07, Raz16, KRT17, Raz17, GRT18, GZ19, DQW21] showed how to achieve symmetric-key encryption, key agreement, oblivious transfer and multiparty computation in the BSM. In this work, we show how to achieve (symmetric-key) message authentication and (public-key) signatures in the BSM. We first begin by describing the BSM and survey what was done in prior work. We then turn to the problem of authentication in the BSM, and describe our results.

The Bounded Storage Model (BSM). In the bounded storage model (BSM), we restrict the storage capacity rather than the run-time of the adversary. Two distinct variants of the BSM exist in the literature, and we follow the terminology of [DQW21] to distinguish between them.

The default variant for this work will be the *streaming* BSM [Raz16, GZ19, DQW21]. In the streaming BSM, parties are modeled as streaming processes, that can send and receive large amounts of data between them by processing it one bit at a time, while only maintaining a small amount of memory. We denote the memory bound for the honest parties by n and the memory bound for the adversary by m . Typically, we assume the adversary has much higher resources than the honest parties, namely $m \gg n$. Although the honest parties only have a small amount of memory, they can stream large amounts of communication $k > m$, so as to overwhelm the memory capacity of the adversary and prevent it from storing the communication between them in its entirety. Still, the honest parties are restricted to remembering even less about the exchanged communication than the adversary can remember! Surprisingly, as we will see, this suffices to construct many powerful cryptographic primitives in the BSM.

We contrast the streaming BSM with an earlier variant, referred to as the *traditional BSM* [Mau92]. In the traditional BSM, a trusted third party (potentially, some natural process) broadcasts a long uniformly random string $X \in \{0, 1\}^k$ referred to as a *randomizer string*. The honest parties can store a subset of n physical locations of X , while the adversary can store any m bits of information about X . After that, the string X disappears, and the adversary gets unrestricted additional memory.

Most cryptographic schemes in the traditional BSM naturally translate to the streaming BSM by having one of the parties generate and stream the randomizer string X . However, looking forward, this will not be the case when it comes to authentication. The traditional BSM guarantees that the honest parties all get access to the same trusted randomizer string X , while in the streaming model, the adversary can tamper with any value X sent by one honest party to another. This makes the authentication problem in the streaming BSM more natural, but potentially harder than in the traditional BSM, since the latter already implicitly assumes a mechanism for authenticating X .

On the other hand, the traditional BSM has additional features that we don’t require in the streaming BSM: (a) in the traditional BSM, there is only one long randomizer string X used in each cryptographic operation, whereas in the streaming BSM, the parties can stream many long messages, (b) in the traditional BSM, the long string X is uniformly random, while in the streaming BSM, the streamed messages can have arbitrary structure, (c) in the traditional BSM, the honest parties can only access a small subset of locations in X , while in the streaming BSM, the parties read the entire communication while maintaining

small local memory, (d) in the traditional BSM, the adversary is only limited to remembering m bits of information about X and gets unlimited memory otherwise, while in the streaming BSM, the adversary observes communication in a streaming manner and is limited to m bits of memory overall. While some of our schemes do achieve some of these additional features, we view this as being of secondary importance and will only mention this in passing; our focus will be on the streaming BSM as our default model.

Prior Work in the BSM. The seminal work of Maurer [Mau92] introduced the traditional BSM. A series of papers [Mau92, DR02, ADR02, DM02, Lu02, Vad04] showed how to achieve symmetric-key encryption in the traditional BSM, where the honest parties share a short secret key that they can use to encrypt arbitrarily many messages (i.e., CPA security). Each encryption relies on a fresh randomizer string X . The main technical tool in all these works was nicely abstracted in the work of Vadhan [Vad04] as a locally computable (seeded) extractor, which can extract randomness from a long string X by only reading a few locations of X . Overall, these works show that symmetric-key encryption is possible even if the adversarial memory $m = 2^{O(n)}$ can be up to exponentially larger than honest user memory n , as long as the ciphertext size (in which we include the size of the randomizer string) can be sufficiently large $k = O(m)$.¹ In particular, if we also want the ciphertext size to be polynomial, then we have to restrict the adversarial memory $m = \text{poly}(n)$ to be some arbitrarily large polynomial.

The works of [CM97, CCM98, Din01, HCR02, DHRS07] showed that it is also possible to achieve public key agreement, oblivious transfer, and general multi-party computation with arbitrary corruption thresholds in the traditional BSM when $m = O(n^2)$, and this was shown to be optimal by [DM04].

All of the above results in the traditional BSM also carry over to the streaming BSM, by having one of the honest parties generate the randomizer string X . However, it turns out that one can often do better in the streaming BSM. The work of Raz [Raz16] proved memory lower bounds for learning parity, and used these to construct novel symmetric-key encryption in the streaming BSM with small ciphertexts of size $< n$, and with $m = O(n^2)$, which is tight. In particular, each individual ciphertext is small enough that it can even fully fit in honest user memory, but the adversary does not have enough memory to remember too many ciphertexts. Follow-up works [KRT17, Raz17, GRT18] generalized this result and showed that, for up to exponentially large $m = 2^{O(n)}$, one can achieve ciphertext size $O(m/n)$, improving by a factor of n over the prior results in the traditional BSM.² The work of [GZ19] studied key agreement and oblivious transfer in the streaming BSM, and showed how to use Raz’s lower-bound to construct new protocols that improve over the prior works in the traditional BSM by removing any correctness error and mildly reducing the number of rounds and the communication cost, while still requiring that the adversarial memory is bounded by $m = O(n^2)$. The recent work of [DQW21] noticed that the lower bound of $m = O(n^2)$ from [DM04] does not hold in the streaming BSM, and showed how to achieve key agreement, oblivious transfer, and general multiparty computation in the streaming BSM, even when the adversarial memory $m = 2^{O(n)}$ can be up to exponentially larger than honest user memory n , at the cost of growing the round and communication complexities polynomially with m .

The work of [DR02] also considered symmetric-key authentication in the traditional BSM, where parties have a short shared secret key and can use it to authenticate a large number of messages, by using a fresh (authentic) randomizer string X to authenticate each message. The construction is very simple – it uses a local extractor to extract a fresh secret key for a one-time message-authentication code (1-time MAC) from X , and then uses the 1-time MAC with this key to authenticate the message. Unfortunately, unlike all the previous results in the traditional BSM, this solution cannot immediately be ported to the streaming BSM: if one of the parties generates X and streams it, then the adversary can tamper with it. Indeed, this would result in a completely insecure scheme in the streaming BSM.

¹Throughout the introduction, we ignore polynomial factors in the security parameter.

²The above bound is for 1-bit messages and is optimal; if the adversary can store $> n$ ciphertexts under an n -bit key, then she can learn something about the encrypted messages via the Shannon lower bound.

1.1 Our Results

In this work, we study the problem of message authentication and signatures in the streaming BSM. We study three variants of the problem, as outlined below.

Symmetric-Key Authentication. We start with the symmetric-key setting, where Alice and Bob share a short secret key sk , that they store in their n -bit memory. To authenticate a message μ , Alice streams a long authentication tag σ to Bob who verifies the stream and either accepts or rejects.³ Even though the tag size $k = |\sigma|$ can be very large, namely $k > m$, we ensure that generating and verifying the tag can be done in a streaming manner using only n bits of memory. Our goal is to achieve unconditional security, where Alice and Bob can authenticate an arbitrarily large polynomial (or even sub-exponential) number of messages in this manner using their short secret key. We consider a streaming attacker Eve with m bits of memory. Eve sits on the channel between Alice and Bob, and can observe and modify all communication between them, subject only to keeping at most m bits of memory throughout this process. She performs a chosen-message attack, where she can adaptively choose many messages and observe the authentication tags honestly sent from Alice to Bob. Moreover, she can modify each of the messages and tags in transit and can observe whether Bob accepts or rejects. (We view the process of receiving and modifying the tags as a streaming algorithm, where Eve receives each tag as an input stream and produces a modified tag as an output stream. The streams need not be synchronized and Eve can wait to receive many/all of the bits of the input stream before outputting the first bit in the output stream.) Eve wins if she causes Bob to accept some arbitrary message that Alice did not authenticate.

We show how to solve this problem with up to exponentially large adversarial memory $m = 2^{O(n)}$ and tags of size $k = O(m)$. In particular, this means that if we also want the scheme to also be polynomially efficient (i.e., have polynomial-size tags), then we have to restrict the adversarial memory $m = \text{poly}(n)$ to be some arbitrarily large polynomial. We refer to Theorem 5.1 for a formal statement with parameters.

Symmetric-Key Authentication with Short Tags. Next, we consider the same problem as previously, but now also require that each individual tag σ is “short”, of size $k < n$. Moreover, each tag can be fully generated, stored and verified by the honest parties in n bits of memory, without needing to operate in the streaming model. We show how to solve this problem when the adversarial memory is $m = O(n^2)$, which is optimal.⁴ This result relies on Raz’s space lower bounds for learning parity [Raz16]. We refer to Theorem 6.1 for a formal statement with parameters.

Public-Key Signatures. Lastly, we consider the public-key signature setting, where the honest parties do not have any shared secrets to begin with. As in the standard signature setting, we have an initialization stage where Alice generates a public verification key vk and a secret signing key sk , and she broadcasts her verification key vk to other honest users over a *public but authentic* channel. In the BSM, we allow the verification key vk to be large, of size $> m$, and Alice generates and broadcasts vk in a streaming manner while maintaining n bits of memory, which also bounds the size of the secret signing key sk that remains in Alice’s memory at the end of the initialization stage. A verifier Bob receives and processes the streamed verification key vk using n bits of memory, and derives a short verification digest vd , which he keeps in his memory at the end of the initialization stage. Later, Alice can use her signing key sk to sign arbitrarily many messages by streaming large signatures over a *fully malicious* channel. Bob can verify the streamed signatures using his verification digest vd . The adversary Eve maintains m bits of memory throughout the entire process. During the initialization stage, Eve can observe (but not modify) the streamed verification key vk . Afterwards, Eve performs a chosen-message attack, where she can adaptively choose many messages and observe the streamed signatures sent from Alice to Bob. Moreover, she can modify each of the messages

³We typically think of the messages μ as small relative to n, m . However, all of our results also extend to support long messages that Alice and Bob receive in a streaming manner.

⁴Optimality follows from the standard authentication lower bound; if an adversary has enough memory to store n tags (generated using an n bit key) then it has enough information to forge tags of new messages.

and signatures in transit and can observe whether Bob accepts or rejects. Eve wins if she can cause Bob to accept some arbitrary message that Alice did not sign.

We give a solution to the above problem when $m = O(n^2)$, which we show to be optimal. The size of the verification key and each of the signatures is $O(m)$. Our proof of security relies on a new “block entropy lemma”, which we view as a result of independent interest, and discuss below.

Block Entropy Lemma. Consider a random variable $X = (X_1, \dots, X_k) \in (\{0, 1\}^b)^k$ consisting of k blocks of b bits each. Assume that X has a high min-entropy rate with $\mathbf{H}_\infty(X) \geq \alpha \cdot (kb)$ for some parameter α . We would intuitively like to say that this must imply that a large α_1 -fraction of the blocks X_i must individually have a high min-entropy rates $\mathbf{H}_\infty(X_i) \geq \alpha_2 \cdot b$. For example, if $\alpha = \Omega(1)$, then we may intuitively hope that the above should hold with some $\alpha_1 = \alpha_2 = \Omega(1)$.

Indeed, if we were to consider Shannon entropy instead of min-entropy, then $\mathbf{H}(X) \leq \sum_i \mathbf{H}(X_i)$, and therefore the above holds with (e.g.,) $\alpha_1 = \alpha_2 = \alpha/2$. Unfortunately, when it comes to min-entropy, such statements are false for any reasonable parameters. As an example, consider the distribution on X where we choose a random index $i^* \leftarrow [k]$, set $X_{i^*} = 0^b$, and for all other indices $i \neq i^*$ we choose $X_i \leftarrow \{0, 1\}^b$ uniformly at random. Then $\mathbf{H}_\infty(X) \geq (k-1) \cdot b$, but for each individual index $i \in [k]$, we have $\mathbf{H}_\infty(X_i) \leq \log(k)$ since $\Pr[X_i = 0^b] \geq 1/k$.

While the above example shows that the statement is false in its basic form, it also highlights a potential way to augment the statement to make it hold while preserving its intuitive meaning. The example shows that for every fixed index i chosen a-priori, X_i may have low min-entropy. However, we may be able to find a large set of good indices i after seeing X (e.g., to avoid $i = i^*$ in the previous example), such that each such X_i has high min-entropy even if we reveal that i is in this good set. We formalize this by showing that for any $\alpha_1 \approx \alpha_2 \approx \alpha/2$, with overwhelming probability over $x \leftarrow X$ there is a large set of “good” indices $\mathcal{I}(x)$ of size $|\mathcal{I}(x)| \geq \alpha_1 \cdot k$, such that the min-entropy of the i ’th block is high if we condition on $i \in \mathcal{I}(x)$ being a good index: $\mathbf{H}_\infty(X_i | i \in \mathcal{I}(x)) \geq \alpha_2 \cdot b$.

Our block-entropy lemma is somewhat related to previous lemmas of [NZ96, Vad04] showing that, if we (pseudo-)randomly sample a sufficiently large subset of locations $S \subseteq [k]$, then $X_S = \{X_i : i \in S\}$ has a high entropy rate, close to the entropy rate of X . However, these previous lemmas do not allow us to say anything about the entropy rates of individual blocks X_i . Our lemma is also somewhat related to a lemma of [DQW21], which shows that many individual *bits* of X have some non-trivial entropy. However, that lemma does not give good parameters when extended to blocks, and indeed, the natural attempt at extending it leads to a false statement as discussed above.

1.2 Our Techniques

We now discuss the high-level technical ideas behind each of our results. To avoid getting bogged down with detailed parameter choices, we will informally use the terms “short/small” to denote values of size $< n$ that can fit in honest user memory, and “long/big” to denote values of size $> m$ that cannot even fit in adversary memory.

Symmetric-Key Authentication. Our basic result for symmetric-key authentication is fairly simple. We rely on two building blocks:

- A (strong) local extractor Ext that takes as input a long source x and a short random seed , and produces a short output $\text{Ext}(x; \text{seed})$ by only reading a small subset of the locations in x , where the locations are determined by seed [Vad04]. In particular, it can be computed by reading x in a streaming manner using small memory. The extracted output is statistically close to uniform even given seed , as long as x has sufficiently high entropy.
- A streaming one-time message-authentication code $\sigma = \text{MAC}_{\text{sk}}(\mu)$ that allows us to authenticate a single long message μ using a small secret key sk , by reading μ in a streaming manner. The authentication tag is also small. This can be constructed easily using polynomial evaluation.

We use these building blocks to construct a symmetric-key authentication scheme in the BSM.

- The small shared secret key $\text{sk} = (\hat{\text{sk}}, \text{seed})$ consists of key $\hat{\text{sk}}$ for the streaming one-time MAC and seed for an extractor.
- To authenticate a message μ , Alice then streams a long random string x to Bob, and as she does so, also computes $r = \text{Ext}(x; \text{seed})$ and $\hat{\sigma} = \text{MAC}(\hat{\text{sk}}, (x||\mu))$ in a streaming manner using small memory. She then appends the short value $\psi = r \oplus \hat{\sigma}$ as the final component of the long tag $\sigma = (x, \psi)$
- To verify the tag $\sigma = (x, \psi)$ for message μ , Bob computes $r = \text{Ext}(x; \text{seed})$ and $\hat{\sigma} = \text{MAC}(\hat{\text{sk}}, (x, \mu))$ in a streaming manner using small memory. He then checks if $\psi = r \oplus \hat{\sigma}$ and, if so, accepts if $\hat{\sigma}$ is a valid tag for message $(x||\mu)$ under MAC (using secret $\hat{\text{sk}}$, and where the verification procedure is also performed in a streaming manner), and rejects otherwise.

In essence, Alice is “encrypting” the one-time MAC tags using a symmetric-key encryption in the BSM, and then uses a one-time MAC to also authenticate the BSM encryption randomness. Intuitively, the encryption ensures that even if the adversary sees many encrypted tags, she doesn’t learn much about the secret key of the one-time MAC. However, formalizing this is somewhat subtle.

Consider an adversary Eve to first passively observes q honestly generated tags $\sigma_1, \dots, \sigma_q$ from Alice in a streaming manner, while maintaining m bits of memory. Then Eve gets unlimited memory and observes an additional $(q+1)$ ’st tag σ and outputs σ' . She wins if Bob accepts σ' and $\sigma' \neq \sigma$. In our proof, we first switch the values ψ_i inside the q tags σ_i to uniformly random. We rely on the security of the strong local extractor to argue that the change is indistinguishable, *even* if we later reveal the entire secret key $\text{sk} = (\text{seed}, \hat{\text{sk}})$ and can therefore check if the adversary wins the game. Once we make this change, the one-time MAC key $\hat{\text{sk}}$ is only used to generate the $(q+1)$ ’st tag $\sigma = (x, \psi)$ for message μ and verify $\sigma' = (x', \psi')$ for message μ' . Therefore, we can rely on the security of the one-time MAC to argue that if $(\mu', \sigma') \neq (\mu, \sigma)$ then Bob rejects with overwhelming probability. The above argument easily generalizes to the case where Eve makes many verification queries, where she modifies tags from Alice and sees whether Bob accepts, as long as they are synchronous: each streamed verification query to Bob coincides with at most one streamed authentication query from Alice. Dealing with the general asynchronous case, where Eve’s verification queries can be arbitrarily overlapped with her authentication queries, requires more care. Nevertheless, we do manage to give a proof for the general asynchronous case, by relying on the basic argument outlined above and carefully partitioning the queries made by Eve.

We remark that, on a quick glance, it may seem that we could have applied the one-time MAC to just μ instead of the pair (x, μ) , and still argued that Eve cannot cause Bob to accept any $\sigma' = (\mu', x', \psi')$ where $\mu' \neq \mu$, which should suffice. However, this turns out to be false and the scheme would be completely insecure with this modification. In particular, Eve would be able to learn something about the secret key by modifying the tags from Alice to Bob while keeping the message intact and seeing whether or not he accepts. For example, each time Alice generates a tag $\sigma_i = (\mu_i, x_i, \psi_i)$, Eve could flip a bit of x_i and see if Bob accepts or rejects. This would eventually reveal the set of locations in x read by $\text{Ext}(x; \text{seed})$. Once Eve learns this set, extractor security is lost, and indeed Eve can fully recover sk and break the scheme. In our proof, by showing that Eve cannot cause Bob to accept any $\sigma' \neq \sigma$, even for the same message μ , we ensure that verification queries are useless. Our formal argument is more involved, and requires to carefully answer verification queries using low memory: our reduction cannot even store an entire tag.

Overall, we can get security against adversaries with memory of size m where honest users only need memory of size $n = O(\log m)$ and where tags are of size $O(m)$. We refer to Section 5 for more details.

Symmetric-Key Authentication with Short Tags. Next, we turn to the same problem as in the previous paragraph, but additionally require that the authentication tags are small, of size $k < n$. Furthermore, they can be fully generated, stored, and verified inside honest user memory, without relying on the streaming model. Since the secret key is also small, of size $< n$, if an adversary can simultaneously remember n tags, then it can break security via the classical authentication lower bound. Therefore, the above setting necessitates bounding the adversary’s memory to $m = O(n^2)$.

Constructing this type of authentication scheme implies some sort of space lower bound on learning. In particular, the adversary observes many outputs of the (potentially randomized) authentication function with a secret key sk , but if its memory is bounded, it does not learn the function sufficiently well to authenticate new messages. Unlike the previous setting where each function output individually was long and could not be stored in memory, in this setting each output is short, but the adversary cannot store too many outputs. Lower bounds in this setting are highly non-trivial, and the first lower bound of this form was shown only recently in the celebrated work of Raz [Raz16], and subsequently extended by [KRT17, Raz17, GRT18]. In particular, Raz proved a space lower bound for learning random parities. Here, we choose a uniformly random $\mathbf{x} \leftarrow \mathbb{F}_2^n$. The adversary can get many samples (\mathbf{a}_i, b_i) where $\mathbf{a}_i \leftarrow \mathbb{F}_2^n$ and $b_i = \langle \mathbf{x}, \mathbf{a}_i \rangle$. If the adversary has $> n^2$ bits of memory and can remember $> n$ samples in full, then it can perform Gaussian elimination and recover \mathbf{x} . The work of Raz shows that, if the adversary’s memory is sufficiently smaller than n^2 (and the number of samples is smaller than exponential), then the adversary cannot recover \mathbf{x} or even distinguish the samples from uniformly random values.

Abstracting the above, we can think of Raz’s result as showing that the inner-product function $f_{\mathbf{x}}(\mathbf{a}) = \langle \mathbf{x}, \mathbf{a} \rangle$ is a “weak pseudorandom function” (weak PRF) in the BSM, in the sense that an adversary with sufficiently small memory cannot distinguish the outputs of the function on random inputs \mathbf{a}_i from uniformly random values. Unfortunately, it is not a-priori clear how to use a weak PRF in the BSM to solve the message authentication problem in the BSM. Indeed, even in the computational setting, we do not know of any “nice” construction of computational message authentication codes from a generic weak PRF (beyond just using the weak PRF to construct a PRG and applying [GGM84] to go from a PRG to a PRF, but this approach does not translate to the BSM). Instead, we rely on the specifics of the inner product function, rather than just generically relying on it being a weak PRF in the BSM. Notice that the inner-product function is essentially the same as Learning Parity with Noise (LPN), just without noise! Moreover, there is a long body of work trying to construct simple/efficient message authentication schemes from LPN [HB01, JW05, KPC⁺11, DKPW12]. In particular, the work of [DKPW12] abstracts out many of these ideas and shows how to construct message-authentication directly from any *key-homomorphic weak PRF*.

We rely on the fact that the inner-product function is essentially a key-homomorphic weak PRF in the BSM via Raz’s lower bound, and show how to adapt the ideas of [DKPW12] to construct a message-authentication scheme in the BSM with the desired parameters. Interestingly, we crucially rely on the linearity of the inner-product function in this construction, and do not know how to adapt subsequent space lower bounds for other functions (e.g., low-degree polynomials [GRT18]) to get analogous results using them. Still, adapting the work of [DKPW12] to the BSM requires a great deal of care to ensure that all the reductions use small memory. It is notably even non-trivial for the (low-memory) reduction to remember which tags have been previously issued by the tagging oracle during the unforgeability experiment so as to ensure that verification queries with these tags are correctly accepted. We develop new tricks to deal with this and other issues. We refer to Section 6 for more details.

Public-Key Signatures. In the public-key setting, we start with an initialization stage during which Alice streams a long verification key vk over an authentic channel. At the end of the initialization stage, Alice keeps a short signing key sk , while a verifier Bob keeps a short verification digest vd .

We first observe that if Alice and Bob could interact back-and-forth with each other during the initialization stage, we could solve the problem trivially by having them run a key-agreement protocol in the BSM [CM97, GZ19, DQW21], at the end of which they would have a shared secret key. Later, they could use the secret key to authenticate messages via a symmetric-key authentication schemes in the BSM that we constructed earlier. However, all such key-agreement protocols require at least two back-and-forth rounds of interaction, while in our case we only allow one-way communication from Alice to Bob.⁵ Unfortunately, it is easy to see that key agreement in one round is impossible, since nothing differentiates an honest Bob from the attacker Eve – if Bob knows Alice’s key, so can Eve.

⁵We view this as a crucial component of our model. Alice can broadcast her verification key to the world, obviously of who is listening and who will want to verify her signed messages in the future. There may potentially even be a large number of different verifiers and Alice should not need to know about them or interact with them.

Nevertheless, we show that a variant of key agreement that we call *set key agreement* is possible in one round! In a set key agreement scheme, Alice streams a long message \mathbf{vk} and outputs a small set of keys $\mathbf{sk} = (\mathbf{sk}_1, \dots, \mathbf{sk}_q)$, while Bob observes the stream and outputs some subset of Alice’s keys $\mathbf{vd} = (T, (\mathbf{sk}_i)_{i \in T})$ for $T \subseteq [q]$. Security requires that there is some shared key \mathbf{sk}_t for $t \in T$, meaning that the key is known to both Alice and Bob, such that \mathbf{sk}_t looks uniformly random even given Eve’s view of the protocol and all the other keys \mathbf{sk}_j for $j \in [q], j \neq t$. The index t of this “good” shared key is a random variable that depends on the entire protocol execution and on Eve’s view, and it is therefore not known to Alice or Bob.

We construct such set key agreement when the adversarial memory is sufficiently small $m = O(n^2)$, as follows. Alice streams a random string $x = (x_1, \dots, x_k)$ consisting of k blocks $x_i \in \{0, 1\}^b$, where b is very small (depending only on the security parameter). Alice chooses a small random subset $S_A \subseteq [k]$ of size $|S_A| = q$ and remember the blocks $x_i : i \in S_A$. Similarly, Bob chooses a small random subset $S_B \subseteq [k]$ of size $|S_B| = q$ and remember the blocks $x_i : i \in S_B$. We choose $k = O(m)$ and $q = O(\sqrt{m})$ carefully to ensure that Alice/Bob only use $n = \tilde{O}(\sqrt{m})$ bits of memory and that with high probability their sets have non-trivial intersection $|S_A \cap S_B|$ of roughly security parameter size. After Alice finishes streaming x , she then also streams her set $S_A = \{i_1, \dots, i_q\}$ along with q extractor seeds seed_j , and sets her keys to be the extracted values $\mathbf{sk}_j = \text{Ext}(x_{i_j}, \text{seed}_j)$ for $j \in [q]$. Bob computes the set $T = \{j : i_j \in S_B\}$ and sets $\mathbf{sk}_j = \text{Ext}(x_{i_j}; \text{seed}_j)$ for $j \in T$. We argue that security holds as long as Eve’s memory is $\leq k/2$. Eve needs to decide what to remember about x before she learns anything about S_A, S_B . We will use our new block entropy lemma to argue that there is a large fraction of locations i , such that the individual blocks x_i have high min-entropy conditioned on what Eve remembered about x , and therefore it is likely that some such location $i^* = i_t$ appears in $S_A \cap S_B$. The corresponding key $\mathbf{sk}_t = \text{Ext}(x_{i_t}; \text{seed}_t)$ will then be the “good key” which looks uniform given Eve’s view. As discussed in our description of the block entropy lemma, the set of locations i such that x_i has high min-entropy is not fixed, but rather a random variable that depends on x and on Eve’s view. Therefore, also the index t of the “good key” is such a random variable.

Once we have a set key agreement protocol, it is easy to build a signature scheme on top of it. To sign each message μ , Alice uses each of the secret keys \mathbf{sk}_i for $i \in [q]$ as a key for a symmetric-key authentication scheme in the BSM (from our first result), and sequentially authenticate the message μ using each of the q keys individually. The verifier Bob verifies the authentication tags corresponding to the indices $j \in T$. To argue security, we rely on the security of the symmetric-key authentication scheme in the location t corresponding to the “good” secret key. Note that, in the set key agreement protocol, we needed to choose $q = O(\sqrt{m})$ large to ensure that that S_A and S_B have a large overlap. This means that each of the keys \mathbf{sk}_i needs to be very small, since Alice’s storage is $n = q|\mathbf{sk}_i|$. However, this is not a problem since the symmetric-key authentication scheme in the BSM allowed us to make the keys as small as $O(\log m)$.⁶

We also give a lower bound to show that this type of public-key signatures in the BSM are impossible when $m > n^2$. We rely on a technical lemma due to Dziembowski and Maurer [DM04]. Translated to our context, the lemma addresses the scenario where Alice streams a long verification key \mathbf{vk} to Bob, after which she stores some n bit state \mathbf{sk} and Bob stores some n bit state \mathbf{vd} . The lemma says that there is some small value $E(\mathbf{vk})$ of size $m < n^2$ that Eve can store about \mathbf{vk} , such that, conditioned on $E(\mathbf{vk})$, the values \mathbf{sk} and \mathbf{vd} are statistically close to being independent. In particular, this means that if Eve samples a fresh \mathbf{sk}' conditioned on $E(\mathbf{vk})$, then she can sign messages using \mathbf{sk}' and Bob will accept the signatures with high probability. We refer to Section 7 for more details.

Block Entropy Lemma. Lastly, we give some intuition behind our block entropy lemma. We show that, for every fixed $x \in (\{0, 1\}^b)^k$ there is some set $\mathcal{I}(x)$ such that, if X has sufficiently high entropy $\mathbf{H}_\infty(X) \geq \alpha \cdot (bk)$, then:

- With overwhelming probability over $x \leftarrow X$, we have $|\mathcal{I}(x)| \geq \alpha_1 \cdot k$
- For all x and $i \in \mathcal{I}(x)$: $\mathbf{H}_\infty(X_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1}, i \in \mathcal{I}(X)) \geq \alpha_2 \cdot b$.

⁶This is also why we need to rely on the symmetric-key authentication scheme with long tags from our first result, rather than the one with short tags from our second result – in the latter, the secret keys would be of size $O(\sqrt{m})$ and hence Alice/Bob would need memory of size $O(m)$ exceeding that of Eve.

where the above holds for any α_1 and $\alpha_2 = \alpha - \alpha_1 - \lambda/b$, where λ is the security parameter (see Lemma 4.1 for a formal statement). We start by observing that the definition of min-entropy guarantees that for any x :

$$2^{-\alpha kb} \geq \Pr[X = x] = \prod_{i=1}^k \underbrace{\Pr[X_i = x_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1}]}_{2^{-e_i}} = 2^{-\sum_{i=1}^k e_i}.$$

By a simple averaging argument, there is an α_1 fraction of “good” indices $i \in \mathcal{I}(x) \subseteq [k]$ for which $e_i \geq (\alpha - \alpha_1) \cdot b$. But the fact that $\Pr[X_i = x_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1}]$ is small does not guarantee that $\mathbf{H}_\infty(X_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1})$ is large, since the concrete outcome x_i that we got may not have been the maximally likely one. However, let us additionally condition on the event that $i \in \mathcal{I}(X)$ is a good index. We show that, for each i , either (I) the conditional min-entropy is high $\mathbf{H}_\infty(X_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1}, i \in \mathcal{I}(X)) \geq \alpha - \alpha_1 - \lambda/b$ or (II) the probability $\Pr[i \in \mathcal{I}(X) | X_1 = x_1, \dots, X_{i-1} = x_{i-1}] \leq 2^{-\lambda}$ is low. By removing indices i for which (II) holds from $\mathcal{I}(x)$, we get our lemma. For a random $x \leftarrow X$, the probability that any index was removed is negligible.

2 Preliminaries

Notation. When X is a distribution, or a random variable following this distribution, we let $x \leftarrow X$ denote the process of sampling x according to the distribution X . If X is a set, we let $x \leftarrow X$ denote sampling x uniformly at random from X . We use the notation $[k] = \{1, \dots, k\}$. If $x \in \{0, 1\}^k$ and $i \in [k]$ then we let $x[i]$ denote the i 'th bit of x . If $s \subseteq [k]$, we let $x[s]$ denote the list of values $x[i]$ for $i \in s$.

2.1 Information Theory

Statistical Distance. Let X, Y be random variables with supports S_X, S_Y , respectively. We define their *statistical difference* as

$$\mathbf{SD}(X, Y) = \frac{1}{2} \sum_{u \in S_X \cup S_Y} |\Pr[X = u] - \Pr[Y = u]|.$$

We write $X \approx_\varepsilon Y$ to denote $\mathbf{SD}(X, Y) \leq \varepsilon$.

Predictability and Entropy. The *predictability* of a random variable X is $\mathbf{Pred}(X) \stackrel{\text{def}}{=} \max_x \Pr[X = x]$. The *min-entropy* of a random variable X is $\mathbf{H}_\infty(X) = -\log(\mathbf{Pred}(X))$. Following Dodis et al. [DORS08], we define the conditional predictability of X given Y as $\mathbf{Pred}(X|Y) \stackrel{\text{def}}{=} \mathbb{E}_{y \leftarrow Y} [\mathbf{Pred}(X|Y = y)]$ and the (average) conditional min-entropy of X given Y as: $\mathbf{H}_\infty(X|Y) = -\log(\mathbf{Pred}(X|Y))$. Note that $\mathbf{Pred}(X|Y)$ is the success probability of the optimal strategy for guessing X given Y .

Lemma 2.1 ([DORS08]). *For any random variables X, Y, Z where Y is supported over a set of size T we have $\mathbf{H}_\infty(X|Y, Z) \leq \mathbf{H}_\infty(X|Z) - \log T$.*

Lemma 2.2 ([DORS08]). *For any random variables X, Y , for every $\varepsilon > 0$ we have*

$$\Pr_{y \leftarrow Y} [\mathbf{H}_\infty(X|Y = y) \geq \mathbf{H}_\infty(X|Y) - \log(1/\varepsilon)] \geq 1 - \varepsilon.$$

Lemma 2.3. *If X and Y are independent conditioned on Z then $\mathbf{H}_\infty(X|Y) \geq \mathbf{H}_\infty(X|Y, Z) \geq \mathbf{H}_\infty(X|Z)$.*

Lemma 2.4. *If X and Y are independent conditioned on Z then $\mathbf{H}_\infty(X, Y|Z) \geq \mathbf{H}_\infty(X|Z) + \mathbf{H}_\infty(Y|Z)$.*

Extractors. We review the notion of randomness extractors and known parameters.

Definition 2.5 ((Strong, Average-Case) Seeded Extractor [NZ96]). *We say that an efficient function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^\ell$ is an (α, ε) -extractor if for all random variables (X, Z) such that X is supported over $\{0, 1\}^n$ and $\mathbf{H}_\infty(X|Z) \geq \alpha$ we have $\mathbf{SD}((Z, S, \text{Ext}(X; S)), (Z, S, U_\ell)) \leq \varepsilon$ where S, U_ℓ are uniformly random and independent bit-strings of length d, ℓ respectively.*

Theorem 2.6 ([ILL89]). *There exist (α, ε) -extractors with input length n and output length ℓ as long as $\alpha \geq \ell + 2 \log(1/\varepsilon)$. Furthermore, such an extractor can be computed in $O(n)$ time and space.*

Definition 2.7 (BSM Extractor [Vad04]). *We say that an efficient function $\text{BSMExt} : \{0, 1\}^k \times \{0, 1\}^d \rightarrow \{0, 1\}^\ell$ is an (n, m, ε) -BSM extractor if:*

- *Given seed $\in \{0, 1\}^d$ initially stored in memory, it is possible to compute $\text{BSMExt}(x; \text{seed})$ given streaming access to $x \in \{0, 1\}^k$ using at most n bits of total memory. Moreover, it can be done while only accessing at most n locations (chosen non-adaptively) in the string x .*
- *BSMExt is an (α, ε) -extractor (Definition 2.5) for $\alpha = k - m$.*

Note that a BSM Extractor gives a simple one-round protocol (n, m) -BSM protocol where Alice and Bob start with a uniformly random shared key key_0 of some small size d and derive a new shared key $\text{key}_1 \in \{0, 1\}^\ell$ of a larger size $\ell > d$. Alice just streams a random $x \in \{0, 1\}^k$ to Bob and both parties compute $\text{key}_1 = \text{BSMExt}(x; \text{key}_0)$. Security holds since the adversary can only store m -bits of information about x so it has $\alpha \geq k - m$ bits of entropy conditioned on the adversary's view, and key_0 acts as a random seed which is a-priori unknown to the adversary. Therefore $\text{key}_1 = \text{BSMExt}(x; \text{key}_0)$ is ε -close to uniform given the adversary's view of the protocol.

Theorem 2.8 ([Vad04]). *For any $m \geq \ell, \lambda$, there is a (n, m, ε) -BSM extractor $\text{BSMExt} : \{0, 1\}^k \times \{0, 1\}^d \rightarrow \{0, 1\}^\ell$ with $n = O(\ell + \lambda + \log m)$, $\varepsilon = 2^{-\Omega(\lambda)}$, $k = O(m + \lambda \log(\lambda))$, $d = O(\log m + \lambda)$.*

KL Divergence and Mutual Information. We recall the notions of KL divergence and (conditional) mutual information.

Definition 2.9 (KL Divergence). *For two discrete probability distributions P, Q with the same support $\text{Supp}(P)$, we define the KL divergence as*

$$D_{KL}(P||Q) = \sum_{x \in \text{Supp}(P)} P(x) \log \left(\frac{P(x)}{Q(x)} \right).$$

Definition 2.10 ((Conditional) Mutual Information). *If (X, Y, Z) denotes a triple of (potentially correlated) random variables with joint distribution $P_{(X,Y,Z)}$, we define the mutual information of X and Y conditioned on Z as:*

$$\mathbf{I}(X; Y|Z) = \mathbf{E}_Z D_{KL}(P_{(X,Y)|Z} || P_{X|Z} \times P_{Y|Z}),$$

where $P_{(X,Y)|Z}$ denotes the conditional distribution of (X, Y) given Z , and $P_{X|Z}$ and $P_{Y|Z}$ respectively denote the marginal distributions of X and Y given Z .

Lemma 2.11 (Pinsker's Inequality). *Let X, Y be random variables with distributions P_X, P_Y with the same support $\text{Supp}(X)$. We have:*

$$\mathbf{SD}(X, Y) \leq \sqrt{\frac{1}{2} D_{KL}(P_X || P_Y)}.$$

Lemma 2.12 (Generalized Chernoff Bound). *Let Y_1, \dots, Y_k be random variables such that $Y_i \in \{0, 1\}$ and for all y_1, \dots, y_{i-1} it holds that $\Pr[Y_i = 1 | Y_1 = y_1, \dots, Y_{i-1} = y_{i-1}] \leq \gamma$. Then, for any $\delta \in [0, 1]$, we have $\Pr[\sum_i Y_i \geq (1 + \delta)\gamma \cdot k] \leq \exp(-\frac{\delta^2}{3} \gamma k)$.*

We will use the following lemma (see e.g. [Hoe63] or [Din01, Corollary 3]).

Lemma 2.13. *Let k be an integer, and $S_A, S_B \subset [k]$ be two uniformly and independently sampled subset of $[k]$ of size q where $q \geq 2\sqrt{\lambda \cdot k}$. Then*

$$\Pr[|S_A \cap S_B| < \lambda] < e^{-\lambda/4}.$$

3 Definitions

In this section, we define message authentication codes (MACs) in Section 3.1 and signatures in Section 3.2 in the BSM. We discuss how we model honest parties and adversaries in these settings.

3.1 Message Authentication Codes

A message-authentication code (MAC) over a message space \mathcal{M} is a tuple of algorithms (KeyGen, MAC, Verify) with the following syntax:

- KeyGen(1^λ) \rightarrow sk: on input a security parameter λ , output a key sk.
- MAC(sk, μ): on input a key sk and a message $\mu \in \mathcal{M}$, output a tag σ .
- Verify(sk, μ, σ): on input a key sk, a message $\mu \in \mathcal{M}$ and a tag σ , output a bit b .

We define the following properties.

Definition 3.1 (Correctness). *We say that a MAC is correct if for all message $\mu \in \mathcal{M}$:*

$$\Pr[\text{Verify}(\text{sk}, \mu, \sigma) = 1 : \text{sk} \leftarrow \text{KeyGen}(1^\lambda), \sigma \leftarrow \text{MAC}(\text{sk}, \mu)] \geq 1 - \text{negl}(\lambda).$$

Definition 3.2 (Streaming MAC). *For an integer n , we say that a MAC is a streaming MAC with memory n if MAC and Verify can respectively be computed by streaming algorithms with memory n , given streaming access to μ and (μ, σ) respectively, and if KeyGen can be computed (in place) with memory n . In particular, MAC keys sk have size at most n .*

Definition 3.3 ((Selective) unforgeability under chosen message (and verification) attacks). *For an algorithm Adv, consider the following experiment:*

Experiment $\text{Exp}^{\text{uf-cmva}}(1^\lambda, \text{Adv})$:

1. Sample $\text{sk} \leftarrow \text{KeyGen}(1^\lambda)$.
2. Compute $(\mu^*, \sigma^*) \leftarrow \text{Adv}^{\text{MAC}(\text{sk}, \cdot), \text{Verify}(\text{sk}, \cdot, \cdot)}$.
3. Output 1 if $\text{Verify}(\text{sk}, \mu^*, \sigma^*) = 1$ output 0 otherwise.

We say that an adversary Adv is admissible if Adv did not query $\text{MAC}(\text{sk}, \cdot)$ on input μ^* . In the following, all the adversaries are assumed admissible.

We say that a MAC is $(\varepsilon, \mathcal{A})$ -unforgeable under chosen message and chosen verification queries attacks (uf-cmva) against a class \mathcal{A} of adversaries, if for all adversary $\text{Adv} \in \mathcal{A}$:

$$\Pr[\text{Exp}^{\text{uf-cmva}}(1^\lambda, \text{Adv}) = 1] \leq \varepsilon.$$

We simply say that a MAC is \mathcal{A} -uf-cmva-secure if it is $(\varepsilon, \mathcal{A})$ -uf-cmva-secure for some $\varepsilon = \text{negl}(\lambda)$.

We alternatively say that a MAC is $(\varepsilon, \mathcal{A})$ -selectively unforgeable under chosen message and chosen verification queries attack (suf-cmva) if any adversary $\text{Adv} \in \mathcal{A}$ that declare μ^* before the beginning of experiment makes the associated experiment output 1 with probability at most ε .

Next, we say that a MAC is $(\varepsilon, \mathcal{A})$ -(selectively) unforgeable under chosen message attacks ((s)uf-cma) if any adversary $\text{Adv} \in \mathcal{A}$ that only make tagging queries $\text{MAC}(\text{sk}, \cdot)$ makes the associated experiment output 1 with probability at most ε .

Last, we say that a MAC is \mathcal{A} -(s)uf-cm(v)a if it is $(\varepsilon, \mathcal{A})$ -(s)uf-cm(v)a for some $\varepsilon = \text{negl}(\lambda)$.

Definition 3.4 (Indistinguishability under chosen message attacks). We say that a MAC satisfies $(\varepsilon, \mathcal{A})$ -indistinguishability under chosen-message attacks (ind-cma) if for all $\text{Adv} \in \mathcal{A}$:

$$\left| \Pr[\text{Adv}^{\text{MAC}(\text{sk}, \cdot)} = 1] - \Pr[\text{Adv}^{\text{MAC}(\text{sk}, 0)} = 1] \right| \leq \varepsilon$$

where the probabilities are over the randomness of $\text{sk} \leftarrow \text{KeyGen}(1^\lambda)$.

We say that a MAC is \mathcal{A} -ind-cma if it is $(\varepsilon, \mathcal{A})$ -ind-cma for some $\varepsilon = \text{negl}(\lambda)$.

Next, we define the class of adversaries we consider in this work.

Definition 3.5 (Streaming Adversaries for MACs). We say that an adversary Adv in any of the security experiments above is streaming with memory m if it is a streaming algorithm that accesses oracles MAC and Verify (whenever applicable) in a streaming manner using total memory at most m . Namely, it provides the inputs to oracles MAC and Verify as streams of bits and receives the outputs of MAC as streams. The calls to MAC and Verify are not required to be synchronous: for instance Adv can start a oracle call to Verify while streaming the input to MAC or receiving a tag in a streaming manner.

We restrict adversaries to have at most one concurrent call to MAC at any time, and at most one concurrent call to Verify at any time during the experiment. In other words, the calls to MAC are done in a sequential manner, and similarly for Verify .

We will usually refer to Q (resp. Q_T, Q_V) as the number of total (resp. tagging, verification) queries made by Adv .

The class of adversaries above notably covers man-in-the-middle adversaries who have the ability alter a tag being currently streamed, and observe either the output of the verifier, or use it as a final forgery.

We argue that the restriction to the concurrency of oracle calls above is reasonable. Notably, in a setting where an authenticator streams tags to a single verifier and where a man-in-the-middle possibly tampers the communication between them, we believe it is reasonable to assume that the authenticator certifies messages one message at a time, and the verifier only verifies one streamed tag at a time. In such a setting, adversaries would satisfy the conditions of Definition 3.5.

Remark 3.6 (Adversaries with Unbounded Temporary Memory). In the traditional BSM, adversaries have access to unbounded temporary memory, but have to eventually compress that memory to a state of size at most m . We can similarly define a class of adversaries stronger than Definition 3.5, who, given a streamed tag output by $\text{MAC}(\text{sk}, \cdot)$, can process any single streamed bit using unbounded temporary memory, as long as they compress their state down to some m -bit state before the next bit is streamed. In particular, such adversaries can compute bounded-size non-uniform advice on their own.

We will use the following MAC secure against unbounded adversaries, as long as only one tag is given out.

Claim 3.7 (One-Time Information-Theoretic MACs). Let \mathcal{A} be the set of (potentially unbounded) algorithms that make at most 1 tagging query and at most Q_V verification queries. Then for all integer k and constant $c > 0$, there exists a streaming MAC with memory $n = O(\lambda + \log k)$ and message space $\{0, 1\}^k$ which is $(k \cdot Q_V / 2^{c\lambda}, \mathcal{A})$ -uf-cmva-secure. Furthermore, the MAC produces tags of size n .

Proof. This follows as noting that polynomial evaluation acts as an approximate universal hash function, and therefore induces a MAC. In more details:

$$\text{MAC}(a, \mu) := \sum_{i \in k} \mu_i a^i$$

where the sum is over a field \mathbb{F} of size at least $2^{c\lambda + \lceil \log k \rceil}$, and we are interpreting the bits of μ as field elements, and the secret key is $a \leftarrow \mathbb{F}$. Given μ as a streamed input, MAC can be computed in a streaming manner using memory $O(\lambda + \lceil \log k \rceil) = O(|\mathbb{F}|)$.

Security follows as for any $\mu, \mu^* \in \{0, 1\}^k, \mu \neq \mu^*$ and any $x \in \mathbb{F}$,

$$P_{\mu, \mu^*}(a) := \sum_{i \leq k} (\mu_i - \mu_i^*) a^i = x$$

with probability at most $k/|\mathbb{F}|$ over the randomness of $a \leftarrow \mathbb{F}$: this is because P is of degree k and can only take a given value at most d times. In particular, for all $\mu \neq \mu^* \in \{0, 1\}^k$, and $\sigma, \sigma' \in \mathbb{F}$:

$$\Pr_a [\text{MAC}(a, \mu) = \sigma \wedge \text{MAC}(a, \mu^*) = \sigma'] \leq \frac{k}{|\mathbb{F}|},$$

that is, $\{\text{MAC}(a, \cdot) : \{0, 1\}^k \rightarrow \mathbb{F}\}_{a \in \mathbb{F}}$ is a $k/|\mathbb{F}|$ -universal hash function. Therefore the probability that any adversary Adv which makes at most 1 tagging query and at most Q_V verification queries succeeds in the uf-cmva experiment is at most $\frac{kQ_V}{2^{c\lambda}}$. \square

3.2 Signatures

A *streaming* signature scheme over a message space \mathcal{M} is a tuple of algorithms ($\text{KeyGen}, \text{Sign}, \text{Verify}$) with the following syntax:

- $\text{KeyGen}(1^\lambda) \rightarrow (\text{vk}, \text{sk})$: on input a security parameter λ , stream a verification key vk and store signing key sk .
- $\text{KeyReceive}(1^\lambda, \text{vk})$: on input a security parameter λ and a streamed verification key vk , output a verification state vd .
- $\text{Sign}(\text{sk}, \mu)$: on input a signing key sk and a (potentially streamed) message $\mu \in \mathcal{M}$, output a (potentially streamed) signature σ .
- $\text{Verify}(\text{vd}, \mu, \sigma)$: on input a verification state vd , a (potentially streamed) message $\mu \in \mathcal{M}$ and a (potentially streamed) signature σ , output a bit $b \in \{0, 1\}$.

Definition 3.8 (Correctness). *We say that a signature scheme is correct if for all message $\mu \in \mathcal{M}$:*

$$\Pr[\text{Verify}(\text{vd}, \mu, \sigma) = 1 : \text{sk} \leftarrow \text{KeyGen}(1^\lambda), \text{vd} \leftarrow \text{KeyReceive}(1^\lambda, \text{vk}), \sigma \leftarrow \text{Sign}(\text{sk}, \mu)] \geq 1 - \text{negl}(\lambda).$$

Definition 3.9 (Streaming Signature). *For an integer n , we say that a streaming signature scheme has memory n if $\text{KeyGen}, \text{KeyReceive}, \text{Sign}$ and Verify can respectively be computed by streaming algorithms with memory n , given streaming access to μ and (μ, z) respectively. We additionally require that sk and vd have size at most n .*

Definition 3.10 (Unforgeability under chosen message attacks). *For an algorithm Adv , consider the following experiment:*

Experiment $\text{Exp}^{\text{uf-cmva}}(1^\lambda, \text{Adv})$:

1. Sample $(\text{vk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda)$, $\text{vd} \leftarrow \text{KeyReceive}(1^\lambda, \text{vk})$, and compute $\text{view}_{\text{Adv}} \leftarrow \text{Adv}(\text{vk})$.
2. Compute $(\mu^*, \sigma^*) \leftarrow \text{Adv}^{\text{Sign}(\text{sk}, \cdot), \text{Verify}(\text{vd}, \cdot, \cdot)}(\text{view}_{\text{Adv}})$.
3. Output 1 if $\text{Verify}(\text{vd}, \mu^*, \sigma^*) = 1$ output 0 otherwise.

We say that an adversary Adv is admissible if Adv did not query $\text{Sign}(\text{sk}, \cdot)$ on input μ^* . In the following, all the adversaries are assumed admissible.

We say that a signature scheme is $(\varepsilon, \mathcal{A})$ -unforgeable under chosen message attacks (uf-cma) against a class \mathcal{A} of adversaries, if for all adversary $\text{Adv} \in \mathcal{A}$:

$$\Pr[\text{Exp}^{\text{uf-cmva}}(1^\lambda, \text{Adv}) = 1] \leq \varepsilon.$$

We simply say that a signature scheme is \mathcal{A} -uf-cma-secure if it is $(\varepsilon, \mathcal{A})$ -uf-cmva-secure for some $\varepsilon = \text{negl}(\lambda)$.

We define variants of unforgeability, namely selective and/or without verification queries ((s)uf-cm(v)a)-security), in a similar way than in Section 3.1.

We consider a similar class of streaming adversaries for signatures, as in Definition 3.5.

Definition 3.11 (Streaming Adversaries for Signatures). *We say that an adversary Adv in any of the security experiments above is streaming with memory m if it is a streaming algorithm that accesses oracles KeyGen , MAC and Verify (whenever applicable) in a streaming manner using memory at most m . Namely, such adversaries first observe the stream vk produced by KeyGen (but not sk). Then, it provides the inputs to oracles Sign and Verify as streams of bits and receives the outputs of Sign as streams. The calls to Sign and Verify are not required to be synchronous: for instance Adv can start a oracle call to Verify while streaming an input to Sign or receiving a tag in a streaming manner. We add the restriction that there are at most one concurrent call to Sign at any time, and at most one concurrent call to Verify at any time during the experiment.*

We will refer to Q (resp. Q_S, Q_V) as the number of total (resp. signing, verification) queries made by Adv .

As in Remark 3.6, one can consider adversaries with temporary unbounded memory for signatures.

4 Block Entropy Lemma

In this section, we describe and prove our block entropy lemma (Lemma 4.1), which we will use to build signature schemes in the streaming BSM Section 7. We also present an alternative form of the lemma in Corollary 4.2.

Lemma 4.1. *Let $X = (X_1, \dots, X_k) \in \{0, 1\}^{kb}$ be a random variable with blocks $X_i \in \{0, 1\}^b$, such that $\mathbf{H}_\infty(X) \geq \alpha \cdot (kb)$ for some $\alpha > 0$. Then, there are some parameters $\alpha_1, \alpha_2, \varepsilon$ instantiated below and some set $\text{BAD} \subseteq \{0, 1\}^{kb}$ such that the following holds:*

1. $\Pr[X \in \text{BAD}] \leq \varepsilon$.
2. For all $x = (x_1, \dots, x_k)$ there is a set $\mathcal{I}(x) \subseteq [k]$ such that:
 - (a) if $x \notin \text{BAD}$ then $|\mathcal{I}(x)| \geq \alpha_1 \cdot k$,
 - (b) for all $i \in \mathcal{I}(x)$: $\mathbf{H}_\infty(X_i \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1}, i \in \mathcal{I}(X)) \geq \alpha_2 \cdot b$.

Furthermore, the values x_1, \dots, x_i fully determine whether or not $i \in \mathcal{I}(x)$.

The parameters $\alpha_1, \alpha_2, \varepsilon$ can be set according to either of the following two options:

- (i) For any $\alpha_1 > 0, \gamma > 0$, we can set $\alpha_2 = (\alpha - \alpha_1 - \log(1/\gamma)/b), \varepsilon = k \cdot \gamma$.
For example, for any $\rho > 0$, we can set $\alpha_1 = \alpha_2 = \alpha/2 - \rho$ and $\varepsilon = k \cdot 2^{-2\rho b}$.
- (ii) For any $\beta > 0, \gamma > 0, \delta \in [0, 1]$, we can set: $\alpha_1 = \beta - (1 + \delta)\gamma, \alpha_2 = (\alpha - \beta - \log(1/\gamma)/b), \varepsilon = \exp\left(-\frac{\delta^2}{3}\gamma k\right)$. For example, for any $\rho > 0$ such that $b \geq \log(1/(2\rho))/(2\rho)$, we can set:
 $\alpha_1 = \alpha_2 = \alpha/2 - \rho$ and $\varepsilon = e^{-\frac{2}{3}\rho k}$.

Proof. Fix any $x = (x_1, \dots, x_k)$. Then, by the definition of min-entropy, we have:

$$2^{-\alpha \cdot (kb)} \geq \Pr[X = x] = \prod_{i=1}^k \underbrace{\Pr[X_i = x_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1}]}_{\stackrel{\text{def}}{=} 2^{-e_i}} = 2^{-\sum_{i=1}^k e_i}.$$

Let $\beta > 0$ be some parameter and define the set $\mathcal{H}(x) \subseteq [k]$ via $\mathcal{H}(x) := \{i : e_i \geq (\alpha - \beta)b\}$. Intuitively, these are the indices i where the i 'th block takes on a low probability value x_i conditioned on previous blocks, and therefore the indices where we should expect to find entropy. Note that x_1, \dots, x_i fully determines whether $i \in \mathcal{H}(x)$. By an averaging argument, it must be the case that $|\mathcal{H}(x)| \geq \beta \cdot k$.

Now, let us consider the random variable $\mathcal{H}(X)$. We claim that for every $i \in [k]$ and every $\gamma > 0$, at least one of the following two conditions must hold:

- (I) Either $\Pr[i \in \mathcal{H}(X) | X_1 = x_1, \dots, X_{i-1} = x_{i-1}] \leq \gamma$,
- (II) or $\mathbf{H}_\infty(X_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1}, i \in \mathcal{H}(X)) \geq (\alpha - \beta) \cdot b - \log(\gamma)$.

Intuitively, the above says that if we condition on any choice of the first $i - 1$ blocks then either (I) it is unlikely for the i 'th block to take on any value whose a-priori probability is too small, or (II) the i 'th block has high entropy even if we condition on it taking some such value. In particular, if (I) does not hold, then

$$\begin{aligned} & \max_z \Pr[X_i = z | X_1 = x_1, \dots, X_{i-1} = x_{i-1}, i \in \mathcal{H}(X)] \\ &= \max_z \frac{\Pr[X_i = z \wedge i \in \mathcal{H}(X) | X_1 = x_1, \dots, X_{i-1} = x_{i-1}]}{\Pr[i \in \mathcal{H}(X) | X_1 = x_1, \dots, X_{i-1} = x_{i-1}]} \\ &\leq \max_{z \in \mathcal{Z}_i(x)} \frac{\Pr[X_i = z | X_1 = x_1, \dots, X_{i-1} = x_{i-1}]}{\gamma} \\ &\quad \text{where } \mathcal{Z}_i(x) = \{z : \Pr[X_i = z | X_1 = x_1, \dots, X_{i-1} = x_{i-1}] \leq 2^{-(\alpha - \beta) \cdot b}\} \\ &\leq \frac{2^{-(\alpha - \beta) \cdot b}}{\gamma} \end{aligned}$$

The second inequality follows by noting that $i \in \mathcal{H}(X) \Leftrightarrow X_i \in \mathcal{Z}_i(x)$ whenever $X_1 = x_1, \dots, X_{i-1} = x_{i-1}$.

Now define $\mathcal{A}(x) = \{i : \Pr[i \in \mathcal{H}(X) | X_1 = x_1, \dots, X_{i-1} = x_{i-1}] \leq \gamma\}$, which corresponds to the indices i for which case (I) holds. Let $\mathcal{I}(x) = \mathcal{H}(x) \setminus \mathcal{A}(x)$, which corresponds to the set of indices for which case (II) holds *and* they are in $\mathcal{H}(x)$. Then, for any $i \in \mathcal{I}(x)$:

$$\mathbf{H}_\infty(X_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1}, i \in \mathcal{I}(X)) = \mathbf{H}_\infty(X_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1}, i \in \mathcal{H}(X)) \geq (\alpha - \beta)b - \log(1/\gamma),$$

where the second first equality follows from the fact that x_1, \dots, x_{i-1} fully determine whether $i \in \mathcal{A}(x)$, and the second inequality follows from the definition of condition (II). This proves part 2.(b) of the lemma with $\alpha_2 = (\alpha - \beta - \log(1/\gamma))/b$.

We first prove the lemma for the first parameter choice (i). Define $\text{BAD} = \{x : |\mathcal{I}(x)| < \beta \cdot k\}$. Clearly this ensures that part 2.(a) of the lemma is satisfied with $\alpha_1 = \beta$. Therefore, we are left to prove part 1 of the lemma. Note that $\mathcal{I}(X) = \mathcal{H}(X) \setminus \mathcal{A}(X)$ and $|\mathcal{H}(x)| \geq \beta k$. Intuitively, the only way we can have $X \in \text{BAD}$ is if there is some index i where the i 'th block X_i was unlikely to take on any low-probability value conditioned on the prior blocks, but it did so anyway, and this is unlikely to occur. Formally

$$\begin{aligned} \Pr[X \in \text{BAD}] &= \Pr[|\mathcal{I}(X)| < \beta k] \\ &\leq \Pr[|\mathcal{H}(X) \cap \mathcal{A}(X)| \geq 1] \\ &\leq \sum_i \Pr[i \in \mathcal{H}(X) \cap \mathcal{A}(X)] \\ &\leq \sum_i \max_{\{(x_1, \dots, x_{i-1}) : i \in \mathcal{A}(x)\}} \Pr[i \in \mathcal{H}(X) | X_1 = x_1, \dots, X_{i-1} = x_{i-1}] \leq k \cdot \gamma. \end{aligned}$$

The second line follows from the fact that $\mathcal{I}(X) = \mathcal{H}(X) \setminus \mathcal{A}(X)$ and $|\mathcal{H}(x)| \geq \beta k$ for all x . The second to last inequality follows by noting that (x_1, \dots, x_{i-1}) fully determine whether $i \in \mathcal{A}(x)$, and the last inequality follows by noting that the definition of $\mathcal{A}(x)$ guarantees that for any x_1, \dots, x_{i-1} for which $i \in \mathcal{A}(x)$ we have $\Pr[i \in \mathcal{H}(X) : X_1 = x_1, \dots, X_{i-1} = x_{i-1}] \leq \gamma$. Therefore this proves part 1 of the lemma with $\varepsilon = k \cdot \gamma$.

We now modify the above to get parameter choice (ii). Pick any $\delta > 0$ and define $\text{BAD} = \{x : |\mathcal{I}(x)| \leq (\beta - (1 + \delta)\gamma)k\}$. Clearly this ensures that part 2.(a) of the lemma is satisfied with $\alpha_1 = \beta - (1 + \delta)\gamma$. Therefore, we are left to prove part 1. Note that $\mathcal{I}(X) = \mathcal{H}(X) \setminus \mathcal{A}(X)$ and $|\mathcal{H}(x)| \geq \beta k$. Intuitively, the only way we can have $X \in \text{BAD}$ is if there are many indices i where the i 'th block X_i was unlikely to take on any low-probability value conditioned on the prior blocks, but it did so anyway, and this is unlikely to occur too many times via a Chernoff bound. Formally:

$$\begin{aligned} \Pr[X \in \text{BAD}] &= \Pr[|\mathcal{I}(X)| \leq (\beta - (1 + \delta)\gamma)k] \\ &\leq \Pr[|\mathcal{H}(X) \cap \mathcal{A}(X)| \geq (1 + \delta)\gamma k] \\ &= \Pr\left[\sum_i Y_i \geq (1 + \delta)\gamma \cdot k\right] \\ &\leq \exp\left(-\frac{\delta^2}{3}\gamma k\right) \end{aligned}$$

where $Y_i = 1$ if $i \in \mathcal{H}(X) \cap \mathcal{A}(X)$. Note that, if we condition on any choice of $X_1 = x_1, \dots, X_{i-1} = x_{i-1}$, then the values of Y_1, \dots, Y_{i-1} are fully determined, and also whether or not $i \in \mathcal{A}(X)$ is fully determined. If $i \notin \mathcal{A}(X)$ then $\Pr[Y_i = 1 | X_1 = x_1, \dots, X_{i-1} = x_{i-1}] = 0$. If $i \in \mathcal{A}(X)$ then $\Pr[Y_i = 1 | X_1 = x_1, \dots, X_{i-1} = x_{i-1}] = \Pr[i \in \mathcal{H}(X) : X_1 = x_1, \dots, X_{i-1} = x_{i-1}] \leq \gamma$, by the definition of $\mathcal{A}(X)$. In particular, this shows that for any choice of y_1, \dots, y_{i-1} we have

$$\Pr[Y_i = 1 | Y_1 = y_1, \dots, Y_{i-1} = y_{i-1}] \leq \max_{(x_1, \dots, x_{i-1})} \Pr[Y_i = 1 | X_1 = x_1, \dots, X_{i-1} = x_{i-1}] \leq \gamma.$$

Therefore, the last inequality follows via the (generalized) Chernoff bound (Lemma 2.12). \square

We also show the following corollary, which may prove to be a useful form of the above lemma of independent interest. However, in our work, we end up only relying on the above lemma without needing this corollary. The corollary says that we can think of the set of good indices $I = \mathcal{I}(X)$ as a random variable correlated with X .

Corollary 4.2. *Let $X = (X_1, \dots, X_k)$ be a random variable with blocks $X_i \in \{0, 1\}^b$, such that $\mathbf{H}_\infty(X) \geq \alpha \cdot (kb)$ for some $\alpha > 0$. Then, for any $\alpha_1, \alpha_2, \varepsilon$ satisfying the conditions of Lemma 4.1, there are some random variables X' and I such that:*

- $\mathbf{SD}((X, I), (X', I)) \leq \varepsilon$.
- $I = \{I_1 < I_2 < \dots < I_\ell\} \subseteq [k]$ for $\ell = \lceil \alpha_1 \cdot k \rceil$.
- For all $j \in [\ell]$: $\mathbf{H}_\infty(X'_{I_j}) \geq \mathbf{H}_\infty(X'_{I_j} | I_j) \geq \mathbf{H}_\infty(X'_{I_j} | X'_1, \dots, X'_{I_j-1}, I_j) \geq \alpha_2 \cdot b - 1$.

Proof. Given X , we define random variable X', I as follows:

- If $X \notin \text{BAD}$ then $X' = X$, $I = \hat{\mathcal{I}}(X)$, where $\hat{\mathcal{I}}(x)$ consists of the first ℓ elements of $\mathcal{I}(x)$.
- Else if $X \in \text{BAD}$ then $X' = U_{kb}$, $I = \{1, \dots, \ell\}$, where U_{kb} is uniformly random over $\{0, 1\}^{kb}$.

It is easy to see that $\mathbf{SD}((X, I), (X', I)) \leq \Pr[X \in \text{BAD}] \leq \varepsilon$ and that $|I| = \ell$.

To prove the last part of the lemma, define B to be the event that $X \in \text{BAD}$. Then:

$$\begin{aligned}
\mathbf{Pred}(X'_{I_j} \mid X'_1, \dots, X'_{I_j-1}, I_j) &= \max_{\mathcal{A}} \Pr[\mathcal{A}(X'_1, \dots, X'_{I_j-1}, I_j) = X'_{I_j}] \\
&\leq \max_{\mathcal{A}} \Pr[\mathcal{A}(X'_1, \dots, X'_{j-1}, j) = X'_j \wedge B] + \max_{\mathcal{A}} \Pr[\mathcal{A}(X'_1, \dots, X'_{I_j-1}, I_j) = X'_{I_j} \wedge \neg B] \\
&\leq 2^{-b} \cdot \varepsilon + \max_{\mathcal{A}} \Pr[\mathcal{A}(X_1, \dots, X_{I_j-1}, I_j) = X_{I_j} \wedge I_j \in \mathcal{I}(X)] \\
&\leq 2^{-b} \cdot \varepsilon + \max_{(i, x_1, \dots, x_i)} \Pr[X_i = x_i \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1}, i \in \mathcal{I}(X)] \\
&\leq 2^{-b} \cdot \varepsilon + 2^{-\alpha_2 \cdot b} \\
&\leq 2^{-\alpha_2 \cdot b + 1}.
\end{aligned}$$

Therefore $\mathbf{H}_\infty(X'_{I_j} \mid X'_1, \dots, X'_{I_j-1}, I_j) = -\log(\mathbf{Pred}(X'_{I_j} \mid X'_1, \dots, X'_{I_j-1}, I_j)) \geq \alpha_2 \cdot b - 1$. \square

5 MAC with Long Tags

In this section, we build a streaming MAC in the streaming BSM where the size of tags grow with the memory bound of the adversary. More precisely, we prove the following theorem:

Theorem 5.1. *For all integers m, λ , there exists a streaming MAC with memory $n = O(\lambda + \text{polylog}(m, \lambda))$ (Definition 3.2) which can authenticate messages of length up to 2^λ , and which is $(2^{-\lambda}, \mathcal{A})$ -uf-cmva-secure (Definition 3.3), where \mathcal{A} is the set of streaming adversaries with memory m that make a total number of at most $Q = 2^\lambda$ oracle queries in the unforgeability experiment (Definition 3.5). Furthermore the (streamed) tags are of size $O(m + \lambda \log \lambda)$.*

5.1 Construction

Construction. Let $|\mu|$ be a parameter. Let m be an integer.

Let $n = O(\lambda + \text{polylog}(m, \lambda, \mu))$, $k = \Theta(m + \lambda \log \lambda)$ and $\ell \leq \max(m, O(\lambda + \log(k + |\mu|)))$. Let $\text{Ext} : \{0, 1\}^k \times \{0, 1\}^d \rightarrow \{0, 1\}^\ell$ be a $(n, m, 2^{-\Omega(\lambda)})$ -BSM extractor (Definition 2.7 and Theorem 2.8),

Let $c \geq 3$ be a constant. Let $(\text{KeyGen}, \text{MAC}, \text{Verify})$ be a one-time, information-theoretic MAC with message space $\{0, 1\}^{k+|\mu|}$ that can be evaluated in a streaming manner using memory $O(\lambda + \log(k + |\mu|))$ and tag size at most ℓ , which is $\frac{(k+|\mu|)Q_V}{2^{c\lambda}}$ -secure against adversaries that make at most one tagging query and Q_V verification queries (which exists by Claim 3.7).

We define the following algorithms:

- $\overline{\text{KeyGen}}(1^\lambda)$: Sample a seed $\text{seed} \leftarrow \{0, 1\}^d$ for Ext and a key sk for the one-time MAC. Output:

$$\overline{\text{sk}} = (\text{sk}, \text{seed}).$$

- $\overline{\text{MAC}}(\text{sk}, \mu)$: On input $\mu \leftarrow \{0, 1\}^{|\mu|}$ and $\overline{\text{sk}} = (\text{sk}, \text{seed})$, sample $x \leftarrow \{0, 1\}^k$ and output:

$$\overline{\sigma} = (x, \text{Ext}(x, \text{seed}) \oplus \text{MAC}(\text{sk}, (x \parallel \mu))).$$

We consider here that x is generated and output in a streaming manner, while $\text{Ext}(x; \text{seed})$ and $\text{MAC}(\text{sk}, (x \parallel \mu))$ are computed in a streaming manner.

- $\overline{\text{Verify}}(\mu, \text{sk}, \sigma)$: on a streamed input $(\mu, \overline{\sigma} = (x, \psi))$, compute

$$\sigma := \psi \oplus \text{Ext}(x; \text{seed}),$$

and output $\text{Verify}(\text{sk}, \sigma)$.

Claim 5.2 (Correctness). *Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is correct. Then $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$ is correct.*

Claim 5.3 (Efficiency). $\overline{\text{KeyGen}}$ can be computed (in place) with memory $O(n)$, and MAC and Verify can be computed in a streaming manner using memory $O(n)$.

Proof. Computing KeyGen, as well as the size of $\overline{\text{sk}}$ follow directly by Theorem 2.8 and Claim 3.7, which gives a (one-time, information-theoretic) MAC with tag size $O(\lambda + \log(k + |\mu|))$ where keys can be sampled efficiently. Instantiating Ext with $k = O(m + \lambda \log \lambda)$ and $n \geq \Omega(\lambda + \log(k + |\mu|)) = \Omega(\lambda + \text{polylog}(m, \lambda, |\mu|))$ gives the claim.

As for computing MAC and Verify in a streaming manner, given seed and $\overline{\text{sk}}$, one can sample and output x in a streaming manner. Then one can compute both $\text{Ext}(x, \text{seed})$ and $\text{MAC}(\overline{\text{sk}}, (x, \mu))$ given x as a streaming input, using memory $O(n)$. \square

Theorem 5.4 (Security). Suppose Ext is an $(n, m, 2^{-\Omega(\lambda)})$ -BSM extractor (Definition 2.7 and Theorem 2.8) where $4m \geq 2\lambda$, and that $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is a $(k + |\mu|)/2^{c\lambda}$ one-time information-theoretic MAC (over messages of length $k + |\mu|$).

Let \mathcal{A} be the set of streaming adversaries running in space $m/4$ and that make at most Q_T tagging and Q_V verification queries in the uf-cmva experiment (Definition 3.5). Let $\varepsilon \leq \frac{1}{2^\lambda} + Q_T \cdot 2^{-\Omega(\lambda)} + Q_V \cdot (k + |\mu|)/2^{c\lambda}$. Then $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$ is $(\varepsilon, \mathcal{A})$ -uf-cmva-secure.

In particular, there exists a constant c' such that if $Q_T, Q_V \leq 2^{c'\lambda}$, then $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$ is $(2^{-\Omega(\lambda)}, \mathcal{A})$ -uf-cmva-secure.

Proof. Let Adv be a streaming adversary against the uf-cmva-experiment for $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$ with advantage ε . Without loss of generality, we assume that Adv queries its final forgery to the verification oracle during the experiment. Let Q_T (resp. Q_V) be the number of tagging queries (resp. verification queries) it makes.

For $j \in [Q_T]$, consider the j th tagging query, which on input message μ_j , outputs a stream $\overline{\sigma}_j = (x_j, \psi_j)$. Define the j th tagging epoch as the part of the uf-cmva experiment taking place while the part of response to the j th tagging query $x_j \in \{0, 1\}^k$ is being streamed to the adversary. Note that by assumption on Adv, who makes sequential calls to the tagging oracle, these epochs are disjoint (Definition 3.5).

For $i \in [Q_V]$ and $j \in [Q_T]$, we say that the i th verification query is *dominated* by the j th tagging query if the two following conditions hold:

- (1) Denoting the input to the i th verification query $(\mu'_i, \overline{\sigma}'_i)$ where $\overline{\sigma}'_i = (x'_i, \psi'_i)$, at least $k - m/2 + 1$ bits of $x'_i \in \{0, 1\}^k$ are streamed during the j th tagging epoch;
- (2) The output of the i th verification query is given *after* the output of the j th tagging query has been fully streamed to the adversary.

We will use the following properties:

- Because of condition (1) and $k \geq m$, any given verification query is dominated by at most one tagging query. Similarly, because of condition (2), a tagging query can only dominate the ongoing verification query at the time the vector x_j finishes being streamed to the adversary (if such a verification query exists).
- One can decide whether the i th verification query is dominated by the j th tagging query in a streaming manner during the experiment: this can be done by counting the number of input bits streamed by the adversary to the verification oracle while the output of the j th tagging is being streamed.

We consider the following hybrid experiments for $0 \leq i \leq Q_V$. Hybrid H_0 is defined as the uf-cmva experiment for $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$.

Let i where $1 \leq i \leq Q_V$. Let $j(i)$ be the last tagging epoch that ends before the i th verification query output (with the convention $j(i) = 0$ if no such tagging epoch exists, and $j(Q_V + 1) = Q_T + 1$).

Hybrid H_i . We modify how certain verification queries and tagging queries are handled.

We change how the i th verification query is handled.

- If the i th verification query is dominated by the j th tagging query for some $j \in [Q_T]$, let $\bar{\sigma}_j = (x_j, \psi_j)$ denote the response issued to the j th tagging query on input message μ_j . The verification query accepts if it is called on input $(\mu_j, \bar{\sigma}_j)$, and rejects otherwise. Note that this process is well-defined thanks to the properties of domination listed above.
- Otherwise (namely, if the i th verification query is not dominated by any tagging query), reject the query regardless of the input.

We additionally change how the tagging queries with indices $j \leq j(i)$ are handled. If the j th tagging query is not dominating the i th verification query, answer the query instead with $\bar{\sigma} = (x_j, \psi_j)$ where $x_j \leftarrow \{0, 1\}^k$ and $\psi_j \leftarrow \{0, 1\}^\ell$ are sampled uniformly at random.

Note that hybrid H_0 corresponds to the `uf-cmva` experiment, and hybrid H_{Q_V} corresponds to a hybrid where only trivial queries are accepted (namely, ones corresponding to the outputs of dominating tagging queries), so that `Adv` never succeeds in the `uf-cmva` experiment in that hybrid (recall that, without loss of generality, `Adv` calls the verification oracle on its final forgery).

To prove that, for $0 \leq i \leq Q_V - 1$, the output of `Adv` in hybrid H_i is indistinguishable from its output in H_{i+1} , we consider the following sub-hybrids indexed by j , where $j(i) + 1 \leq j \leq j(i+1)$, with the convention $H_i = H_i^{0,1}$ (we don't consider any sub-hybrids if $j(i) + 1 > j(i+1)$).

Hybrid $H_i^{j,0}$. We change how verification queries are handled, if the j th tagging query is not dominating any verification query. The experiment samples ahead of time $x_j^* \leftarrow \{0, 1\}^k$. It replies to the j th tagging query for message μ_j with $\bar{\sigma}_j = (x_j^*, \psi_j)$, where $\psi_j = \text{Ext}(x_j^*; \text{seed}) \oplus \text{MAC}(\text{sk}, (x_j^* \parallel \mu_j))$. On any verification query that is not dominated by the j th tagging query, with input $(\mu, (x_j^*, \psi))$ for any μ, ψ , the experiment aborts. Verification queries are otherwise handled as in hybrid H_i .

Hybrid $H_i^{j,1}$. We change how the j th tagging query is handled. Namely, if the j th tagging query is not dominating the $(i+1)$ st verification query, answer the j th tagging query with input message μ_j with (x_j, ψ_j) where $x_j \leftarrow \{0, 1\}^k$ and $\psi_j \leftarrow \{0, 1\}^\ell$.

Note that the answers to tagging queries only differ between hybrids $H_i^{j,0}$ and $H_i^{j,1}$ when either j is dominating the i th query, or when $j(i) + 1 \leq j \leq j(i+1)$. Verification queries are handled as in hybrid H_i .

Claim 5.5. *Suppose $Q_V Q_T \leq 2^{m/4-\lambda}$. Then for all i, j where $0 \leq i \leq Q_V - 1$ and $j(i) + 1 \leq j \leq j(i+1) - 1$, the outputs of hybrids $H_i^{j-1,1}$ and $H_i^{j,0}$ are within statistical distance $\frac{1}{Q_T \cdot 2^\lambda}$.*

Proof. We argue that the aborting condition only occurs with small probability. Let (x_j^*, ψ_j) be the output of the j th tagging query. Fix any verification query, with input $(\mu, \bar{\sigma} = (x, \psi))$, that is not dominated by the j th tagging query. Let x_{stream} denote the part of x being streamed as an input to the verification during the j th tagging epoch (which can potentially be an empty string). By definition of domination $|x_{\text{stream}}| \leq k - m/2 + 1$. Furthermore, denoting by $\text{state}_{\text{Adv},j}$ the internal state of the adversary at the end of the j th tagging epoch, we have $|\text{state}_{\text{Adv},j}| = m/4$. Thus, $\mathbf{H}_\infty(x_j^* | x) \geq \mathbf{H}_\infty(x_j^* | x_{\text{stream}}, \text{state}_{\text{Adv},j}) \geq m/4 - 1$, so that $x = x_j^*$ with probability at most $2^{-m/4+1}$ over the randomness of x_j^* alone.⁷ An union bound over the Q_V verification queries concludes. \square

Claim 5.6. *Suppose `Ext` is a (n, m, δ) -BSM extractor. Then, for all i, j where $0 \leq i \leq Q_V - 1$ and $j(i) + 1 \leq j \leq j(i+1)$, the outputs of the experiment with `Adv` in hybrids $H_i^{j,0}$ and $H_i^{j,1}$ are within statistical distance δ .*

⁷For verification queries such that x finishes being streamed before the end of the j th tagging epoch, we actually have $\mathbf{H}_\infty(x_j^* | x) \geq m/2 - 1$, as $\text{state}_{\text{Adv},j}$ doesn't influence the input x .

Proof. Fix j , where $j(i) + 1 \leq j \leq j(i + 1)$. Denote the output of the j th tagging query $\bar{\sigma}_j = (x_j, \psi_j)$, where $x_j \in \{0, 1\}^k$. We suppose that the j th tagging query is not dominating the $(i + 1)$ th verification query; otherwise hybrids H_i^j and H_i^{j+1} are identical. Observe that, in both hybrids, the outputs of verification queries of index $i' < i$ are computed as equality tests to its associated dominating query, if such a dominating query exists, and always rejects otherwise. In particular, the experiment up to x_j being fully streamed can be computed without any knowledge of sk nor seed .

Consider the state of the experiment $\text{state}_{\text{Exp}}$ at the moment x_j has finished being streamed to the adversary, but no bit of ψ_j has been streamed yet. $\text{state}_{\text{Exp}}$ can fully be described by (i) the internal state of the adversary, of size at most $m/4$; (ii) the state of the ongoing verification query (if it exists).

We distinguish two cases, according to whether j th tagging query dominates the i th verification query.

Suppose the j th tagging query does not dominate any query. By construction of the hybrids, the ongoing verification does not keep any information about x_j , as it just checks equality with its dominated tagging query. Therefore $\mathbf{H}_\infty(x_j | \text{state}_{\text{Exp}}) \geq k - m/4$.

Suppose the j th tagging query dominates the i th verification query. We claim that by definition of domination, the state of the ongoing verification query at that moment can be described using (at most) $m/2$ bits. This is because at least $k - m/2 + 1$ input bits of x have already been streamed to the dominated verification query during the j th epoch. Therefore, the state of the verification query can be described using (i) one bit, representing whether the first $k - m/2 + 1$ input bits to the verification query match x_j , and (ii) the remaining last bits of x_j that the further input bits to the verification query needs to match; there are at most $m/2 - 1$ of them. We obtain: $\mathbf{H}_\infty(x_j | \text{state}) \geq k - m/2 - m/4 \geq k - m$.

In any case, observing that no extractor seed is needed to execute the experiment up to that point, it can be sampled independently of x_j and state then, so that using BSM-extractor security (Definition 2.7), we obtain:

$$(\text{Ext}(x_j; \text{seed}), \text{state}, \text{seed}) \approx_\delta (u, \text{state}, \text{seed})$$

where $u \leftarrow \{0, 1\}^\ell$ is uniformly random, independent of state and seed. □

Claim 5.7. *Suppose (KeyGen, MAC, Verify) is a ε_{MAC} -secure one-time information-theoretic MAC. Then for all i where $0 \leq i \leq Q_V - 1$, the outputs of the experiments with Adv in hybrids $H_i^{j(i+1),1}$ and H_{i+1} are within statistical distance ε_{MAC} .*

Proof. The only difference between the two hybrids is how the $(i + 1)$ th verification query is answered. In both hybrids, the responses (x_j, ψ_j) to all tagging queries made before the response to the $(i + 1)$ th verification query is given are uniformly random, except for at most one tagging query, the one dominating the $(i + 1)$ th verification query (if it exists), which outputs tag $\bar{\sigma} = (x, \psi)$ where $x \leftarrow \{0, 1\}^k$ and $\psi = \text{Ext}(x; \text{seed}) \oplus \text{MAC}(\text{sk}, (x || \mu))$.

In particular, if

$$\left| \Pr \left[\text{Exp}_{H_i^{j(i+1),1}}^{\text{uf-cmva}}(1^\lambda, \text{Adv}) = 1 \right] - \Pr \left[\text{Exp}_{H_{i+1}}^{\text{uf-cmva}}(1^\lambda, \text{Adv}) = 1 \right] \right| = \varepsilon$$

then Adv makes the $(i + 1)$ th verification query accept $H_i^{j(i+1),1}$, but reject in H_{i+1} , with probability ε . This corresponds to an accepting input different from the dominating tagging query (if it exists). We show that this leads to an attack on the one-time security of (KeyGen, MAC, Verify) with success probability ε .

The reduction generates seed itself. On a tagging query with input μ , it samples $x \leftarrow \{0, 1\}^k$ and streams it to the adversary, and stores x in full. When x finishes being streamed, it determines whether the tagging query was dominating the $(i + 1)$ th verification query. If not, it samples and streams to Adv a uniform ψ . If the tagging query dominates the $(i + 1)$ th verification query, the reduction calls its MAC oracle with input $(x || \mu)$, receives $\text{MAC}(\text{sk}, (x || \mu))$, computes $\text{Ext}(x; \text{seed})$ itself, and streams to Adv the end of the tag $\psi = \text{Ext}(x; \text{seed}) \oplus \text{MAC}(\text{sk}, (x || \mu))$. Verification queries prior to the $(i + 1)$ th are handled as in hybrid $H_i^{(j+1),1}$.

The reduction outputs as a forgery the input to the $(i + 1)$ th verification query from Adv. This corresponds, with probability ε , to an accepting input in $H_i^{j(i+1),1}$ but a rejecting input in H_{i+1} , that is, an input different

from the output of the MAC oracle (if such a MAC oracle query is made). In other words, this constitutes a forgery on the MAC. Last, there exists at most one dominating tagging query for the $(i + 1)$ th verification query, and therefore the reduction makes at most one MAC call. \square

Wrapping up, in hybrid H_{Q_V} , only inputs corresponding to dominating tagging queries are accepted by the verification queries, and therefore no adversary produces forgeries in this hybrid. Combining Claim 5.5, Claim 5.6 and Claim 5.7, we obtain

$$\left| \Pr \left[\text{Exp}_{H_{Q_V}}^{\text{uf-cmva}}(1^\lambda, \text{Adv}) = 1 \right] - \Pr \left[\text{Exp}_{H_0}^{\text{uf-cmva}}(1^\lambda, \text{Adv}) = 1 \right] \right| \leq \frac{1}{2^\lambda} + Q_T \cdot \delta + Q_V \varepsilon_{\text{MAC}}.$$

\square

6 MAC with Short Tags

In this section, we build a streaming MAC in the streaming BSM where the size of tags does *not* grow with the memory bound of the adversary. In particular, honest parties can store entire tags. More precisely, we prove the following theorem:

Theorem 6.1 (MAC with Short Tags). *For all m, λ there exists a streaming MAC with memory $n = O(\lambda\sqrt{m} + \lambda^2)$ (Definition 3.2) which can authenticate messages of length up to $2^{O(n)}$, which is $(2^{-\lambda}, \mathcal{A})$ -uf-cmva-secure (Definition 3.3), where \mathcal{A} is the set of streaming adversaries with memory m that make at most $Q = 2^\lambda$ oracle queries in the unforgeability experiment (Definition 3.5). Furthermore the resulting MAC has tags of size $O(n)$.*

To do so, we first present the main hardness result from [Raz16] in Section 6.1. We present a construction with small message space and no verification queries in Section 6.2. We show how to modify the construction how to handle larger messages in Section 6.3, and how to efficiently detect whether a tag has been previously issued by the tagging oracle in Section 6.4. Then we show how to ensure security against adversaries that have access to verification queries in Section 6.5.

6.1 Hardness of Learning Parities

We recall here the main theorem from [Raz16], which will be the core component of our constructions in this section.

Theorem 6.2 ([Raz16]). *Let $k \geq 1$ be an integer. There exist constants $C, \alpha > 0$ such that the following holds. Let Adv be an algorithm which takes as input a stream $(\mathbf{a}_1, b_1), \dots, (\mathbf{a}_Q, b_Q)$, where $\mathbf{a}_i \in \{0, 1\}^k$, $\mathbf{s} \leftarrow \{0, 1\}^k$ and $b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle \in \mathbb{Z}_2$ and $Q \leq 2^{\alpha k}$, and outputs a vector $\mathbf{s}' \in \mathbb{Z}_2^k$. Suppose that Adv uses at most Ck^2 memory. Then $\Pr[\mathbf{s}' = \mathbf{s}] \leq O(2^{-\alpha k})$.*

Futhermore the same conclusion holds if Adv is given access to unbounded temporary memory between receiving elements (\mathbf{a}, b) , as long as its state is compressed down to Ck^2 bits whenever an element (\mathbf{a}, b) is received. We refer to such adversaries as adversaries with memory Ck^2 and unbounded temporary memory (Remark 3.6).

In this paper we only consider adversaries that have to compress their state down between every bit streamed, even though Theorem 6.2 more generally holds even against adversaries that can process whole samples $(\mathbf{a}_i, b_i) \in \mathbb{Z}_2^k \times \mathbb{Z}_2$ using unbounded memory. Looking ahead, Theorem 6.2 holding against adversaries with unbounded temporary memory will be crucial in the transformation of Section 6.4.

Using that the inner product is a strong extractor, we obtain the following.

Lemma 6.3 (Pseudorandomness of parities (1)). *Suppose $\alpha \cdot k \geq 3\lambda$. Let \mathcal{O}_s be an oracle which, given $\mathbf{s} \in \{0, 1\}^k$, samples $\mathbf{a} \in \{0, 1\}^k$ and outputs $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle)$. Then for any adversary Adv with memory at most Ck^2 , using unbounded temporary memory, and making at most $2^{\alpha k}$ oracle calls:*

$$\left| \Pr_{\mathbf{s} \leftarrow \{0, 1\}^k} \left[\text{Adv}^{\mathcal{O}_s(\cdot)}(\mathbf{A}, \mathbf{A}\mathbf{s}) = 1 \right] - \Pr_{\mathbf{s} \leftarrow \{0, 1\}^k} \left[\text{Adv}^{\mathcal{O}_s(\cdot)}(\mathbf{A}, \mathbf{b}) = 1 \right] \right| \leq 2^{-2\lambda},$$

where $\mathbf{A} \leftarrow \{0, 1\}^{\lambda \times k}$ and $\mathbf{b} \leftarrow \{0, 1\}^\lambda$.

Last, a standard hybrid argument proves the following.

Corollary 6.4 (Pseudorandomness of parities (2)). *Suppose $\alpha \cdot k \geq 3\lambda$. Let \mathcal{O}_S be an oracle which, given $\mathbf{S} \in \{0, 1\}^{k \times \ell}$, samples $\mathbf{a} \in \{0, 1\}^k$ and outputs $(\mathbf{a}, \mathbf{a}^T \mathbf{S}) \in \{0, 1\}^k \times \{0, 1\}^{1 \times \ell}$; and let \mathcal{O}_R be an oracle which samples $(\mathbf{a}, \mathbf{B}) \leftarrow \{0, 1\}^k \times \{0, 1\}^{1 \times \ell}$. Then for any adversary Adv with memory at most Ck^2 , using unbounded temporary memory, and making at most $Q \leq 2^{\alpha k}$ oracle calls:*

$$\left| \Pr_{\mathbf{s} \leftarrow \{0, 1\}^k} \left[\text{Adv}^{\mathcal{O}_S(\cdot)}(1^\lambda) = 1 \right] - \Pr_{\mathbf{s} \leftarrow \{0, 1\}^k} \left[\text{Adv}^{\mathcal{O}_R(\cdot)}(1^\lambda) = 1 \right] \right| \leq \ell Q / 2^{2\lambda}.$$

6.2 Construction without Verification Queries

We start with a construction of MAC with short tags that is unforgeable given access to tagging queries, but not to verification queries.

Construction. Let $k = \Theta(\lambda)$ be an integer.

- **KeyGen**(1^λ): Sample $\mathbf{K}_1 \leftarrow \{0, 1\}^{k \times \lambda}$, $\mathbf{K}_2 \leftarrow \{0, 1\}^k$. Output $\mathbf{K} = (\mathbf{K}_1, \mathbf{K}_2)$.
- **MAC**(\mathbf{K}, μ): on input $\mu \in \{0, 1\}^\lambda$, compute $\mathbf{K}_\mu = (\mathbf{K}_1 \cdot \mu) + \mathbf{K}_2 \in \{0, 1\}^k$. Sample $\mathbf{X} \leftarrow \{0, 1\}^{\lambda \times k}$. Output:

$$\sigma = (\mathbf{X}, \mathbf{X} \cdot \mathbf{K}_\mu) \in \{0, 1\}^{\lambda \times k} \times \{0, 1\}^\lambda.$$
- **Verify**($\mathbf{K}, \mu, \sigma = (\mathbf{X}, \mathbf{Y})$): Check that $\text{rank}(\mathbf{X}) = \lambda$, otherwise reject. Compute $\mathbf{K}_\mu = (\mathbf{K}_1 \cdot \mu) + \mathbf{K}_2$, and accept if and only if $\mathbf{Y} = \mathbf{X} \cdot \mathbf{K}_\mu$.

Claim 6.5 (Correctness). *Suppose $k \geq C\lambda$ for a large enough constant C . Then (KeyGen, MAC, Verify) is correct.*

Proof. This follows as $k \geq \lambda + O(\log \lambda)$ so that $\Pr_{\mathbf{X}}[\text{rank}(\mathbf{X}) \neq \lambda] \leq \text{negl}(\lambda)$. \square

Claim 6.6 (Efficiency). *MAC and Verify can be implemented by streaming algorithms with memory $O(\lambda k)$. Furthermore KeyGen can be computed (in place) with memory $O(\lambda k)$, and the size of the keys and tags are also $O(\lambda k)$.*

Theorem 6.7 (Security). *Let $\mathcal{A} = \mathcal{A}(O(k^2), Q)$ be any class of streaming adversaries such that there exists a constant C' such that algorithms in \mathcal{A} run with memory at most $C'k^2$, using unbounded temporary memory, and that make at most $Q \leq 2^{\alpha k} / \lambda$ oracle calls (Definition 3.5, and where α is defined in Theorem 6.2).*

(KeyGen, MAC, Verify) is $(\frac{\lambda Q + 1}{2^\lambda}, \mathcal{A})$ -unforgeable against tagging queries (uf-cma, Definition 3.3) and $(\frac{\lambda Q}{2^{2\lambda}}, \mathcal{A})$ -message-hiding (ind-cma, Definition 3.4).

Proof. We first prove that the construction satisfies selective unforgeability (suf-cma). Let $\mu^* \in \{0, 1\}^\lambda$ be the message the adversary is committing to forge a MAC on.

We consider the following hybrid experiments:

Hybrid H_0 . This is the suf-cma experiment (Definition 3.3).

Hybrid H_1 . The tagging queries in the suf-cma experiment are instead answered by sampling $(\mathbf{X}, \mathbf{Y}) \leftarrow \{0, 1\}^{\lambda \times k} \times \{0, 1\}^\lambda$.

Lemma 6.8. *for any adversary Adv with memory at most Ck^2 and unbounded temporary memory which makes at most $Q \leq 2^{\alpha k}/\lambda$ tagging queries, we have*

$$\left| \Pr \left[\text{Exp}_{H_0}^{\text{suf-cma}}(1^\lambda, \text{Adv}) = 1 \right] - \Pr \left[\text{Exp}_{H_1}^{\text{suf-cma}}(1^\lambda, \text{Adv}) = 1 \right] \right| \leq \lambda Q / 2^{2\lambda},$$

where $\text{Exp}_{H_b}^{\text{suf-cma}}$ refers to the suf-cma experiment in Hybrid b , and C denotes the constant in Corollary 6.4.

Proof. For any Adv for $\text{Exp}^{\text{suf-cma}}$ with memory Ck^2 , unbounded temporary memory, and at most $Q \leq 2^{\alpha k}$ tagging queries such that

$$\left| \Pr \left[\text{Exp}_{H_0}^{\text{suf-cma}}(1^\lambda, \text{Adv}) = 1 \right] - \Pr \left[\text{Exp}_{H_1}^{\text{suf-cma}}(1^\lambda, \text{Adv}) = 1 \right] \right| = \varepsilon$$

we build a distinguisher between \mathcal{O}_S and \mathcal{O}_R (defined in Corollary 6.4 with $\ell = \lambda$), that runs with memory $O(Ck^2)$, uses unbounded temporary memory, calls the oracle \mathcal{O} λQ times and has advantage ε .

Our reduction samples $\mathbf{K}_{\mu^*} \leftarrow \{0, 1\}^k$. For every tagging query from Adv with message $\mu \neq \mu^*$ (where μ^* is the message for the forgery declared in advance by Adv in the suf-cma experiment), the reduction queries its oracle λ times and obtains $(\mathbf{X}, \mathbf{Y}) \in \{0, 1\}^{\lambda \times k} \times \{0, 1\}^\lambda$, where either $\mathbf{Y} = \mathbf{X}\mathbf{K}_1 \in \{0, 1\}^{\lambda \times \lambda}$ and $\mathbf{K}_1 \leftarrow \{0, 1\}^{k \times \lambda}$, or $\mathbf{Y} \leftarrow \{0, 1\}^{\lambda \times \lambda}$. It replies to the tagging query with $(\mathbf{X}, \mathbf{X}\mathbf{K}_{\mu^*} + \mathbf{Y} \cdot (\mu - \mu^*)) \in \{0, 1\}^{\lambda \times k} \times \{0, 1\}^\lambda$. Upon receiving a forgery $(\mu^*, (\mathbf{X}^*, \mathbf{Y}^*))$ from the adversary, the reduction checks whether the forgery is valid by testing $\mathbf{Y}^* \stackrel{?}{=} \mathbf{X}^* \cdot \mathbf{K}_{\mu^*}$. It outputs 1 if the forgery is valid, and 0 otherwise.

First, note that the distinguisher makes at most λQ oracle queries to \mathcal{O} , and runs in memory $Ck^2 + O(\lambda k + \lambda^2) = C'k^2$ for some constant C' (independent of Adv, and not accounting for unbounded temporary memory), as $\alpha k \geq 2\lambda$.

If the oracle outputs are such that $\mathbf{Y} = \mathbf{X}\mathbf{K}_1$, then the view of Adv corresponds to the original suf-cma experiment from H_0 , noting that even given \mathbf{K}_1 , \mathbf{K}_{μ^*} is uniformly random over the random choice of \mathbf{K}_2 alone.

If $\mathbf{Y} \leftarrow \{0, 1\}^{\lambda \times \lambda}$, then for all $\mu \neq \mu^*$, $\mathbf{Y} \cdot (\mu - \mu^*)$ is uniformly random in over the choice of \mathbf{Y} alone, so that $\mathbf{X}\mathbf{K}_{\mu^*} + \mathbf{Y} \cdot (\mu - \mu^*)$ is uniform. Therefore the view of Adv corresponds to hybrid H_1 .

Corollary 6.4 implies that the advantage of the reduction is at most $\lambda Q / 2^{2\lambda}$ and therefore so is the distinguishing advantage ε of Adv. \square

We now claim that the probability of outputting a valid forgery in H_1 is negligible, even for an unbounded adversary. This is because the view of the adversary is now completely independent of \mathbf{K}_{μ^*} . Now, given any full-rank \mathbf{X} , $\mathbf{X}\mathbf{K}_{\mu^*}$ is uniformly random over the choice of \mathbf{K}_{μ^*} alone (where we use $\lambda \leq k$), so that the probability of outputting $(\mathbf{X}, \mathbf{X} \cdot \mathbf{K}_{\mu^*})$ is at most $2^{-\lambda}$ over the randomness of \mathbf{K}_{μ^*} .

Overall, this shows that for any adversary with memory at most $O(k^2)$ (and unbounded temporary memory), which makes at most $Q \leq 2^{\alpha k}/\lambda$ tagging queries, its advantage is at most

$$\Pr \left[\text{Exp}^{\text{suf-cma}}(1^\lambda, \text{Adv}) \right] \leq \frac{\lambda Q + 1}{2^{2\lambda}}.$$

Next, we claim that (KeyGen, MAC, Verify) is in fact uf-cma-secure, namely adaptively secure::

$$\Pr \left[\text{Exp}^{\text{uf-cma}}(1^\lambda, \text{Adv}) \right] \leq \frac{\lambda Q + 1}{2^\lambda}.$$

This follows from a simple guessing argument: a reduction to the suf-cma security of the construction guesses μ^* in advance, and succeeds only if the guess is correct, which incurs an advantage loss of 2^λ .

Moving on to message-hiding, we argue that the MAC is ind-cma against any adversary with memory $O(k^2)$, unbounded temporary memory, that makes at most $Q \leq 2^{\alpha k}/\lambda$ tagging queries. This follows from a similar argument as suf-cma security, using Corollary 6.4 with secret $\mathbf{S} = \mathbf{K}_2$, which gives that the view of any such adversary is $\lambda Q / 2^{2\lambda}$ -close to one where all tagging queries are answered with uniformly random tags. \square

6.3 Domain Extension for MACs

Next, we show how to extend the message space of any MAC, which preserves uf-cma and ind-cma security, as well as efficiency. The transformation is very similar to [DKPW12, Section 3.2], with the main difference being that the reductions are adapted to the low memory regime. In particular, our reductions, even given unbounded temporary memory, cannot afford to record all the tagging queries made by the adversary.

Construction. Let n, m, ℓ be integers such that $\lambda + \log \ell = O(m)$. Let $\mathcal{A} = \mathcal{A}(O(m), Q)$ be any class of streaming adversaries such that there exists a constant C' such that algorithms in \mathcal{A} run with memory at most $C'm$ making at most Q oracle calls (Definition 3.5). Let $(\text{KeyGen}, \text{MAC}, \text{Verify})$ be a MAC with message space $\{0, 1\}^\lambda$ satisfying:

- $(\varepsilon, \mathcal{A})$ -uf-cma-security (Definition 3.3);
- $(\varepsilon, \mathcal{A})$ -ind-cma-security (Definition 3.4);
- $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is a streaming MAC with memory $O(n)$ (Definition 3.2), with tags of size n .

Let $\mathcal{H} = \{h : \{0, 1\}^\ell \rightarrow \{0, 1\}^\lambda\}$ be a family $2^{-\lambda}$ -universal hash function for some integer ℓ . Let m be an integer such that $O(\lambda + \log \ell) = O(m)$. We define the following MAC $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$ with message space $\{0, 1\}^\ell$.

- $\widetilde{\text{KeyGen}}(1^\lambda)$: Sample $\text{sk} \leftarrow \text{KeyGen}$ and h from the family of universal hash functions, and output

$$\widetilde{\text{sk}} = (\text{sk}, h).$$

- $\widetilde{\text{MAC}}(\widetilde{\text{sk}}, \mu)$: On input $\mu \in \{0, 1\}^\ell$ and $\widetilde{\text{sk}} = (\text{sk}, h)$, output

$$\widetilde{\sigma} = \text{MAC}(\text{sk}, h(\mu)).$$

- $\widetilde{\text{Verify}}(\widetilde{\text{sk}}, \mu, \widetilde{\sigma})$: Parsing $\widetilde{\text{sk}} = (\text{sk}, h)$, output:

$$\text{Verify}(\text{sk}, h(\mu), \widetilde{\sigma}).$$

Claim 6.9 (Correctness). *Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is correct. Then $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$ is correct.*

Claim 6.10 (Efficiency). *Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is a streaming MAC with memory $O(n)$ (Definition 3.2), with tags of size n . Then $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$ is a streaming MAC with memory $O(n + \lambda + \log \ell)$. Furthermore $\widetilde{\text{MAC}}$ produces tags of size n .*

Proof. Taking as the universal hash function the function $h_a(\cdot) = \text{MAC}(a, \cdot)$ from Claim 3.7 with $\mathbb{F} = \{0, 1\}^{\lambda + \lceil \log \ell \rceil}$, we have that the description size of h is $O(\lambda + \log \ell)$ and h can be computed in a streaming manner with memory $O(\lambda + \log \ell)$, and that the resulting universal hash function is $2^{-\lambda}$ -universal. \square

Theorem 6.11 (Security). *Let $\mathcal{H} = \{h : \{0, 1\}^\ell \rightarrow \{0, 1\}^\lambda\}$ be a $2^{-\lambda}$ -universal hash function, and suppose $\lambda + \log \ell = O(m)$. Let $\mathcal{A} = \mathcal{A}(O(m), Q)$ be any class of streaming adversaries such that there exists a constant C' such that algorithms in \mathcal{A} run with memory at most $C'm$, using unbounded temporary memory, and make at most Q oracle calls (Definition 3.5). Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is $(\varepsilon, \mathcal{A})$ -uf-cma-secure (Definition 3.3) and $(\varepsilon, \mathcal{A})$ -ind-cma-secure (Definition 3.4).*

Then $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$ is $(O(Q(\varepsilon + 2^{-\lambda})), \mathcal{A})$ -uf-cma-secure and $(2\varepsilon, \mathcal{A})$ -ind-cma-secure.

Proof. We first prove that $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$ is $(2\varepsilon + Q/2^\lambda, \mathcal{A})$ -uf-cma-secure. Let Adv be a streaming adversary with memory m and unbounded temporary memory against the uf-cma security of $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$, which makes at most Q queries, and with advantage ε' . Let $\mu_1, \dots, \mu_Q \in \{0, 1\}^\ell$ be the messages queried by Adv to the tagging oracle, and let $\mu^* \notin \{\mu_1, \dots, \mu_Q\}$ be the message associated to its forgery.

Let us denote by BAD the event that there exists some $i \in [Q]$ such that $h(\mu^*) = h(\mu_i)$. Let us prove that $\Pr[\text{BAD}] \leq Q(\varepsilon + 2^{-\lambda})$. To do so, we build an adversary against the ind-cma-security of $(\text{KeyGen}, \text{MAC}, \text{Verify})$ with success probability $\Pr[\text{BAD}]/Q - 2^{-\lambda}$. The reduction samples $h \leftarrow \mathcal{H}$, and answers tagging queries for message μ_i by relaying the output of the ind-cma tagging oracle on input $h(\mu_i)$. It additionally chooses a random query $i^* \in [Q]$, and outputs 1 if $h(\mu^*) = h(\mu_{i^*})$, and 0 otherwise. Note that the reduction runs in memory $(m + O(\lambda + \log \ell))$ (and unbounded temporary memory). If the reduction is given a real tagging oracle $\text{MAC}(\text{sk}, \cdot)$, then the reduction outputs 1 with probability at least $\Pr[\text{BAD}]/Q$, over the random choice of his guess i^* . If the reduction is given oracle access to $\text{MAC}(\text{sk}, 0)$, then the reduction succeeds with probability $2^{-\lambda}$ by universality of h . Overall, this gives

$$\Pr[\text{BAD}] \leq Q(\varepsilon + 2^{-\lambda}).$$

Conditioned on $\neg\text{BAD}$, we claim that there exists a memory $m + O(\lambda + \log \ell)$ attack on the uf-cma-security of $(\text{KeyGen}, \text{MAC}, \text{Verify})$ (that uses unbounded memory if Adv does). A reduction samples $h \leftarrow \mathcal{H}$, answers the tagging queries from Adv for message μ_i by first computing $h(\mu_i)$ and querying $\text{MAC}(\text{sk}, \cdot)$ on message $h(\mu_i)$. Given the final forgery (μ^*, σ) of Adv , the reduction outputs $(h(\mu^*), \sigma)$ as a forgery for $(\text{KeyGen}, \text{MAC}, \text{Verify})$. Conditioned on $\neg\text{BAD}$, the reduction is an admissible adversary with memory $m + O(\lambda + \log \ell)$ for the uf-cma experiment, which succeeds whenever Adv succeeds. By $(\varepsilon, \mathcal{A})$ -uf-cma-security of $(\text{KeyGen}, \text{MAC}, \text{Verify})$, the probability that the reduction succeeds is at most ε , so that

$$\varepsilon' \leq \varepsilon + Q(\varepsilon + 2^\lambda).$$

Next, we prove that $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$ is $(2\varepsilon, \mathcal{A})$ -ind-cma-secure. Let Adv be a streaming adversary with memory m against the ind-cma-security of $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$, which uses unbounded temporary memory, makes at most Q queries, and which has advantage ε' . We claim that:

$$\left| \Pr[\text{Adv}^{\widetilde{\text{MAC}}(\text{sk}, \cdot)} = 1] - \Pr[\text{Adv}^{\text{MAC}(\text{sk}, 0)} = 1] \right| \leq \varepsilon.$$

This follows by $(\varepsilon, \mathcal{A})$ -ind-cma security of $(\text{KeyGen}, \text{MAC}, \text{Verify})$, where a reduction with memory $m + O(\lambda + \log \ell)$ (and uses unbounded temporary memory if Adv does) samples $h \leftarrow \mathcal{H}$ and computes on every tagging query $h(\mu)$ in a streaming manner using $O(\lambda + \log \ell)$ memory. Similarly, we have:

$$\left| \Pr[\text{Adv}^{\widetilde{\text{MAC}}(\text{sk}, g(0))} = 1] - \Pr[\text{Adv}^{\text{MAC}(\text{sk}, 0)} = 1] \right| \leq \varepsilon,$$

and therefore $\varepsilon \leq 2\varepsilon'$, where we use that $O(\lambda + \log \ell) = O(m)$. \square

6.4 Efficiently Recognizing Prior Tagging Queries

In this section, we show how to efficiently recognize tags previously issued by tagging queries. In particular, we wish to do so without having to store *all* the previously issued tags, as our reductions have limited memory. Looking ahead, this will be crucial in Section 6.5 in order to answer trivially accepting verification queries only using small memory. In terms of techniques, this is the only step where we require that our reductions can run using unbounded temporary memory.

We first define variants of both unforgeability under chosen message attacks (uf-cma) and message-hiding (ind-cma), where the adversary is given the additional ability to check whether a message-tag pair has previously been issued by the tagging oracle during the respective experiments. Note that these are stronger security notions than standard uf-cma and ind-cma in our setting, as our adversaries don't have enough memory to store all previously issued message-tag pairs in general.

Definition 6.12 (Special Unforgeability under Chosen Message Attacks). *For an algorithm Adv , consider the following experiment:*

Experiment $\text{Exp}^{\text{uf-cmva-special}}(1^\lambda, \text{Adv})$:

1. Sample $\text{sk} \leftarrow \text{KeyGen}(1^\lambda)$.
2. Compute $(\mu^*, \sigma^*) \leftarrow \text{Adv}^{\text{MAC}(\text{sk}, \cdot), \text{isSame}(\cdot, \cdot)}$.
3. Output 1 if $\text{Verify}(\text{sk}, \mu^*, \sigma^*) = 1$ output 0 otherwise,

where the oracle isSame keeps track of all prior oracle queries made by Adv to the tagging oracle $\text{MAC}(\text{sk}, \mu) \rightarrow \sigma$, and $\text{isSame}(\mu, \sigma)$ outputs 1 if (μ, σ) corresponds to such a prior query, and 0 otherwise.

We say that a MAC satisfies $(\varepsilon, \mathcal{A})$ -special unforgeability under chosen message attacks against a class \mathcal{A} of adversaries, if for all adversary $\text{Adv} \in \mathcal{A}$:

$$\Pr[\text{Exp}^{\text{uf-cmva-special}}(1^\lambda, \text{Adv}) = 1] \leq \varepsilon.$$

Definition 6.13 (Special Indistinguishability under Chosen Message Attacks). We say that a MAC satisfies $(\varepsilon, \mathcal{A})$ special indistinguishability under chosen-message attacks if for all $\text{Adv} \in \mathcal{A}$:

$$\left| \Pr[\text{Adv}^{\text{MAC}(\text{sk}, \cdot), \text{isSame}(\cdot, \cdot)} = 1] - \Pr[\text{Adv}^{\text{MAC}(\text{sk}, 0), \text{isSame}(\cdot, \cdot)} = 1] \right| \leq \varepsilon,$$

where the oracle isSame keeps track of all prior oracle queries made by Adv to the tagging oracle $\text{MAC}(\text{sk}, \mu) \rightarrow \sigma$, and $\text{isSame}(\mu, \sigma)$ outputs 1 if (μ, σ) corresponds to such a prior query, and 0 otherwise, and where the probabilities are over the randomness of $\text{sk} \leftarrow \text{KeyGen}(1^\lambda)$.

Next, we show that any MAC secure against adversaries with unbounded temporary memory⁸ remains secure given access an isSame oracle, up to some minor change in the construction.

Construction. Let n, m be integers, $C > 6$ be a constant, and $(\text{KeyGen}, \text{MAC}, \text{Verify})$ be a MAC with message space $\{0, 1\}^{\ell+Cn}$ satisfying:

- $(\varepsilon, \mathcal{A})$ -uf-cma-security (Definition 3.3);
- $(\varepsilon, \mathcal{A})$ -ind-cma-security (Definition 3.4);
- $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is a streaming MAC with memory $O(n)$ (Definition 3.2), with tags of size n .

Consider the following MAC $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$:

- $\widetilde{\text{KeyGen}}(1^\lambda)$: Sample $\text{sk} \leftarrow \text{KeyGen}$ and output $\widetilde{\text{sk}} = \text{sk}$.
- $\widetilde{\text{MAC}}(\widetilde{\text{sk}}, \mu)$: On input $\mu \in \{0, 1\}^\ell$ and $\widetilde{\text{sk}} = \text{sk}$, sample $r \leftarrow \{0, 1\}^{Cn}$, and output

$$\widetilde{\sigma} = (r, \text{MAC}(\text{sk}, (\mu||r))).$$

- $\widetilde{\text{Verify}}(\widetilde{\text{sk}}, \mu, \widetilde{\sigma})$: Parsing $\widetilde{\sigma} = (r, \sigma)$, output:

$$\text{Verify}(\text{sk}, (\mu||r), \sigma).$$

One can check that $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$ directly inherits correctness and security of $(\text{KeyGen}, \text{MAC}, \text{Verify})$. Furthermore, if $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is a streaming MAC with memory n , then $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$ is a streaming MAC with memory $(C + 1)n$.

Next, we show that the construction above further achieves the special versions of security (Definitions 6.12 and 6.13).

⁸Recall these correspond to adversaries that have access to unbounded temporary memory to treat the stream it receives, as long as their internal state is compressed to some bounded-size state between every bit sent.

Claim 6.14. *Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is a streaming MAC satisfying uf-cma-security against adversaries with memory m' , using unbounded temporary memory, making at most $Q_T \leq 2^{n/3}$ tagging queries, and such that $m' \geq n$. Then $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$ satisfies special unforgeability under chosen message attacks (Definition 6.12) against adversaries with memory $m = m'/3$, using unbounded temporary memory, and making at most Q_T tagging queries and $Q_O \leq 2^n$ queries to the isSame oracle.*

Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is ind-cma-secure against adversaries with memory m which use unbounded temporary memory. Then $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$ satisfies special indistinguishability under chosen message attacks (Definition 6.13) against adversaries with memory $m' = m/3$, using unbounded temporary memory, and making at most $Q_O \leq 2^n$ queries to the isSame oracle.

Proof. Let Adv be an adversary with memory m (using unbounded temporary memory) against the special unforgeability of $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$. We build an adversary Adv', which uses unbounded temporary memory, for the plain uf-cma unforgeability experiment of $(\widetilde{\text{KeyGen}}, \widetilde{\text{MAC}}, \widetilde{\text{Verify}})$ as follows.

Adv' keeps track of a set S of message-tag pairs, which is initialized to an empty set at the beginning of the experiment. On a tagging oracle query $\text{MAC}(\text{sk}, \mu)$ from Adv, Adv' does the following:

1. It queries its own tagging oracle $\text{MAC}(\text{sk}, \cdot)$ (from the uf-cma experiment) on input μ , receives a tag σ , and relays σ to Adv. Let state be the internal state of Adv after that query: note that $\text{bitlength}(\text{state}) \leq m$ by assumption on Adv.
2. It adds (μ, σ) to the set S .
3. It updates the set S as follows, using unbounded temporary memory. For all secret keys $\text{sk}' \in \{0, 1\}^n$, Adv' emulates 2^n independent continuations of the special uf-cma experiment (from Definition 6.12), with adversary Adv which is initialized with state and using signing key sk' (using some independent randomness coins to run these continuations). In particular Adv' answers tagging queries from Adv using its knowledge of sk' . If any element (μ, σ) in S has *not* been queried as input to the isSame oracle in any of these $2^n \cdot 2^n = 2^{2n}$ continuations of the experiment, then Adv' removes (μ, σ) from S .

The remaining elements of S are then compressed in the following manner: Adv' stores the internal memory state of Adv following the (μ, σ) query, and, for every remaining element in S , stores the execution and the query where Adv queried the element to isSame. This consists of the corresponding secret $\text{sk}' \in \{0, 1\}^n$, the index of the continuation $\iota \in [2^n]$, along with the index of the isSame query $q \in [Q_O]$ from Adv (but not counting the randomness used coins).

On a special oracle query $\text{isSame}(\mu, \sigma)$ from Adv, Adv' answers to Adv 1 if $(\mu, \sigma) \in S$, and 0 otherwise.

Lemma 6.15. *Suppose Adv wins the special uf-cma experiment with probability ε , and $Q_T \leq 2^{n/3}$. Then Adv' wins the uf-cma experiment with probability $\varepsilon - 2^{-\Omega(n)}$.*

Proof. We consider hybrids experiments, indexed by $i \in \{0, \dots, Q_T\}$ where Q_T denotes the number of tagging queries made by Adv.

Hybrid 0 corresponds to the special unforgeability experiment (Definition 6.12), where the oracle isSame keeps track of all the tagging queries made by Adv so far.

In Hybrid $i \in [Q_T]$, we change how the queries with index $j \leq i$ are handled. Namely, instead of keeping track of all tagging queries, the hybrid adversary now sets up a set S and updates S for all tagging queries of index $j \leq i$ as specified in Steps 1-3 from the description of Adv' above. The reduction fully stores the tagging queries (μ, σ) of index $j > i$. An oracle call to $\text{isSame}(\mu, \sigma)$ is answered with 1 if (μ, σ) belongs either to S or to the calls stored by the hybrid isSame oracle (which keeps track of tagging queries with indices $j > i$).

The outputs of Hybrids i and $i + 1$ differ only if there exists some isSame query on input (μ, σ) corresponding to tagging query $j \leq i + 1$, but (μ, σ) has been removed from the set S maintained by Adv' after the $(i + 1)$ th tagging query.

Given a state of Adv state , denote by $\rho = \rho(\text{sk}, \text{state}, \mu, \sigma)$ the probability that Adv makes some isSame query on input (μ, σ) in a fresh continuation of the game while initialized with state and using signing key sk .

The probability that Adv' removes (μ, σ) from S while Adv makes some isSame query on (μ, σ) is $(1 - \rho)^{2^n} \cdot \rho$. An union bound over the at most $i \leq Q_T$ prior tagging queries therefore shows that the probabilities of winning in Hybrids i and $i + 1$ differ by at most $i(1 - \rho)^{2^n} \cdot \rho \leq Q_T(1 - \rho)^{2^n} \cdot \rho$.

Overall, the advantage of Adv' is at least $\varepsilon - Q_T^2(1 - \rho)^{2^n} \cdot \rho \geq \varepsilon - 2^{\Omega(n)}$. □

Next, we prove that Adv' runs in low memory, and more specifically, that S remains relatively small throughout its execution. We will use the following lemma:

Lemma 6.16. *Let X be a random variable with min-entropy at least k . Let $A = (A_1, A_2)$ be a potentially unbounded and randomized stateful adversary. Consider the following experiment.*

1. *Sample independently $x_1, \dots, x_Q \leftarrow X$.*
2. *Compute $\text{state}' = A_1(x_1, \dots, x_Q; \text{coins})$, where coins denotes the internal randomness of A . Let $m' := \text{bitlength}(\text{state}')$, that is, $\text{state}' \in \{0, 1\}^{m'}$.*
3. *$A_2(\text{state}; r)$ outputs (x'_1, \dots, x'_ℓ) .*

We say that A wins if there exists an injection $f : [\ell] \rightarrow [Q]$ such that $x'_i = x_{f(i)}$ for all $i \in [\ell]$.

Then the probability that A wins is at most $2^{m'} \cdot (Q \cdot 2^{-k})^\ell$.

Proof. Fix any value $\text{state} \in \{0, 1\}^m$. Then for any internal randomness coins and index $i \in [\ell]$, the probability that the value x'_i output by $A_2(\text{state}; \text{coins})$ matches any of the x_j sampled in Step 1 is at most $Q \cdot 2^{-k}$ over the randomness of $x_1, \dots, x_Q \leftarrow X$ alone. Therefore the probability that $A_2(\text{state}; \text{coins})$ wins is at most $(Q \cdot 2^{-k})^\ell$, again over the randomness of $x_1, \dots, x_Q \leftarrow X$ alone. By union bound over the 2^m possible values of state , the probability that A wins the experiment is at most $2^m \cdot (Q \cdot 2^{-k})^\ell$. □

Lemma 6.17. *The algorithm Adv' uses memory at most $m' = m + (m/n) \cdot (2n + \log Q_O)$ except with negligible probability, where Q_O denotes the number of isSame oracle calls from Adv .*

Proof. Fix any tagging query $i \in [Q_T]$. Let ℓ denote the number of elements in S after resolution of the i th tagging query. Given all the previous tagging queries $x_j = \widetilde{\text{MAC}}(\mu_j)$, $j \leq i$, that have min-entropy at least Cn (by construction, over the randomness of $r \leftarrow \{0, 1\}^{Cn}$), and the internal state of Adv' (along with some independent randomness coins used to run the continuations of the experiment) represents the ℓ elements of S using a state of size $m' = m + \ell(2n + \log Q_O)$ using the representation from Step 3, which by Lemma 6.16 happens with probability at most

$$2^{m'} \cdot (i \cdot 2^{-k})^\ell \leq 2^{m'} \cdot (Q_T \cdot 2^{-k})^\ell = 2^{m - \ell \cdot (C-4)n},$$

where we use $Q_O, Q_T \leq 2^n$. Therefore, the probability that Adv' stores in S more than $\ell \geq m/n$ elements is at most $2^{-(C-5)m}$ which is negligible. Using that $m \geq n$ and an union bound over the number of tagging queries $Q_T \leq 2^n$, the probability that Adv' stores at any point more than $\ell = m/n$ elements in S is at most $2^{-(C-6)m}$, which concludes. □

The analysis for the ind-cma setting is identical. □

6.5 Upgrading to Security with Verification Queries

We now show how to build, starting with a MAC satisfying both special uf-cma -security (Definition 6.12, namely security without verification queries but with isSame queries) and special ind-cma -security (i.e. the MAC is message-hiding with isSame queries), a MAC satisfying uf-cmva -security (that is, security with both tagging and verification queries). The transformation is very similar to [DKPW12, Section 3.1], with the difference that our reductions are adapted to the low-memory setting.

We will use a family of pairwise independent hash functions which can be efficiently evaluated in a streaming manner.

Lemma 6.18 (Efficient Pairwise Independent Hash Functions). *Let n, λ be integers. There exists a family of pairwise independent hash functions $\mathcal{H} = \{h : \{0, 1\}^n \rightarrow \{0, 1\}^\lambda\}$ where functions have description size $O(n + \lambda)$, and can be evaluated in a streaming manner using $O(n + \lambda)$ bits of memory.*

Proof. We take as the family of pairwise independent hash functions the family of affine function over $\mathbb{F} = \mathbb{F}_{\max(2^n, 2^\lambda)}$, that is $\{h_{a,b} : x \rightarrow ax + b\}_{a,b \in \mathbb{F}}$ (and truncating the output if $n > \lambda$, and padding x if $\lambda > n$), its keys have size $\max(n, \lambda) + \lambda = O(n + \lambda)$ bits. Furthermore, $h_{a,b}$ can be computed with memory $O(n + \lambda)$, as being a multiplication and an addition over $\mathbb{F}_{\max(2^n, 2^\lambda)}$. \square

Construction. Let n, m be integers such that $\lambda \leq O(m)$, and n an integer such that $n \leq O(k^2)$. Let $\mathcal{A} = \mathcal{A}(O(m), Q)$ be any class of streaming adversaries such that there exists a constant C' such that algorithms in \mathcal{A} run with memory at most $C'm$ making at most Q oracle calls (Definition 3.5). Let $(\text{KeyGen}, \text{MAC}, \text{Verify})$ be a MAC with message space $\mathcal{M} = \{0, 1\}^{|\mu|+\lambda}$ that satisfies:

- $(\varepsilon, \mathcal{A})$ - special uf-cma -security (Definition 6.12);
- $(\varepsilon, \mathcal{A})$ special ind-cma -security (Definition 6.13);
- $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is a streaming MAC with memory $O(n)$ (Definition 3.2), with tags of size n .

Let $\mathcal{H} = \{h : \{0, 1\}^n \rightarrow \{0, 1\}^\lambda\}$ be a family of pairwise independent hash functions from Lemma 6.18. Consider the following MAC $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$.

- $\overline{\text{KeyGen}}(1^\lambda)$: Sample $\text{sk} \leftarrow \text{KeyGen}$, and sample $h : \{0, 1\}^n \rightarrow \{0, 1\}^\lambda$ from a family of pairwise independent hash functions. Output

$$\overline{\text{sk}} = (\text{sk}, h)$$

- $\overline{\text{MAC}}(\overline{\text{sk}}, \mu)$: Parse $\overline{\text{sk}} = (\text{sk}, h)$. Sample $r \leftarrow \{0, 1\}^\lambda$ and compute $\sigma = \text{MAC}(\text{sk}, (\mu||r))$. Output

$$\overline{\sigma} = (\sigma, h(\sigma) \oplus r)$$

- $\overline{\text{Verify}}(\overline{\text{sk}}, m, \overline{\sigma} = (\sigma, t))$: Parse $\overline{\text{sk}} = (\text{sk}, h)$. Compute $r = t \oplus h(\sigma)$. Output $\text{Verify}(\text{sk}, (\sigma||r), z)$.

Claim 6.19 (Correctness). *Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is correct. Then $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$ is correct.*

Claim 6.20 (Efficiency). *Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is a streaming MAC with memory $O(n)$ (Definition 3.2), with tags of size n . Then $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$ is a streaming MAC with memory $O(n + \lambda)$ Furthermore $\overline{\text{MAC}}$ produces tags of size $n + \lambda$.*

Proof. Taking the pairwise independent hash function as a random affine function over $\mathbb{F}_{\max(2^n, 2^\lambda)}$ (and truncating the output if $n \geq \lambda$), its keys have size $\max(n, \lambda) + \lambda = O(n + \lambda)$ bits, and can be computed with memory $O(n + \lambda)$, as being a multiplication and an addition over $\mathbb{F}_{\max(2^n, 2^\lambda)}$. This makes the key and tag size, as well as the memory cost of $\overline{\text{MAC}}$ and $\overline{\text{Verify}}$ $O(n + \lambda)$. \square

Theorem 6.21 (Security). *Suppose that $\lambda, n \leq O(m)$. Let $\mathcal{A} = \mathcal{A}(O(m), Q)$ be any class of streaming adversaries such that there exists a constant C' such that algorithms in \mathcal{A} run with memory at most $C'm$ making at most Q oracle calls (Definition 3.5), and that do not query the verification oracle on an input provided by the tagging oracle. Suppose that $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is $(\varepsilon, \mathcal{A})$ -special-uf-cma-secure (Definition 6.12) and $(\varepsilon, \mathcal{A})$ -special-ind-cma-secure (Definition 6.13). Then $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$ is $(2Q_T Q_V (\varepsilon + \frac{1}{2\lambda}), \mathcal{A})$ -uf-cmva-secure (Definition 3.3) where Q_T and Q_V respectively denote the number of tagging and verification queries made by adversaries.*

Proof. Let Adv be an adversary in the uf-cmva experiment for $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$, running in memory $O(m)$ and having success probability ε' . Respectively denoting by Q_T and Q_V the number of tagging and verification queries made by Adv , we either build:

- an adversary with memory $O(m)$ for the special-uf-cma experiment, with success probability $\frac{\varepsilon'}{2Q_T Q_V}$; or
- an adversary with memory $O(m)$ for the special-ind-cma experiment, with success prob $\frac{\varepsilon'}{2Q_T Q_V} - \frac{1}{2\lambda}$.

Without loss of generality, we assume Adv makes a verification query corresponding to its forgery.

Our reduction guesses the first verification query $i^* \in [Q_V]$ submitted by Adv that verifies but was not produced by a tagging query. Note that conditioned on this guess i^* being correct, all previous verification queries $i < i^*$ can be answered as follows, using the `isSame` oracle: accept the verification query on input (μ, σ) if and only if `isSame` (μ, σ) outputs 1 (which indicates that (μ, σ) has been previously queried during the experiment). If Adv makes Q_V verification queries, then with probability ε'/Q_V such a verification query is a valid tag that has not previously been queried to the tagging oracle.

Let such a verification query be (μ^*, σ^*, t^*) . Let $r^* = t^* \oplus h(\sigma^*)$. We consider two cases.

- Case 1: r^* has not been used for any previous tagging queries (where, for all tagging queries, r is defined as $r := t \oplus h(\sigma)$). Let ε_{FF} denote the probability that Case 1 occurs conditioned on the guess i^* being correct. Then any such query provides a forgery for the underlying MAC in the uf-cma experiment, with message $(\mu^* || r^*)$, with a straightforward reduction that outputs the guessed verification query as its forgery. This gives an attack on the special-uf-cma experiment with success probability ε_{FF}/Q_V (where we use the `isSame` oracle to answer prior verification queries). Furthermore, if Adv runs with memory $O(m)$, then so does our reduction by assumption on n, λ and k , and uses the same number of tagging queries as Adv .
- Case 2: r^* has been used for a previous tagging query. Let ε_C denote the probability that Case 2 occurs conditioned on the guess i^* being correct. We show that this induces an attack on the ind-cma security of the original MAC.

Our reduction to the special ind-cma security of $(\text{KeyGen}, \text{MAC}, \text{Verify})$ samples a pairwise independent hash function h . By Lemma 6.18 this can be done in memory $O(n + \lambda)$. It answers tagging queries μ from Adv by sampling $r \leftarrow \{0, 1\}^\lambda$, querying the tagging oracle provided by the ind-cma oracle on $(\mu || r)$, obtaining a tag σ and outputting $(\sigma, h(\sigma) \oplus r)$, and answers the verification queries before i^* using the `isSame` oracle. It picks a random tagging query $j^* \in [Q_T]$; denote it by $(\mu_{j^*}, \sigma_{j^*}, t_{j^*})$. It checks that (1) the tagging query j^* is made before the verification query i^* , and (2) we have $r^* = t^* \oplus h(\sigma^*) = r_{j^*} = t_{j^*} \oplus h(\sigma_{j^*})$. It outputs 1 if these two conditions hold, and 0 otherwise.

Note that if Adv runs with memory $O(m)$, our reduction runs in memory $O(m)$ by assumption on n, λ and k and uses the same number of tagging queries as Adv .

If our reduction to the special ind-cma experiment is given oracle access to $\text{MAC}(\text{sk}, \cdot)$, the i^* th query of Adv is a collision on r^* with some tagging query with probability at least ε_C/Q_V , and therefore our reduction outputs 1 with probability at least $\varepsilon_C/(Q_T Q_V)$.

If our reduction to the special ind-cma experiment is given oracle access to $\text{MAC}(\text{sk}, 0)$, the view of Adv is identical to receiving as answers to tagging queries for μ :

$$(\sigma \leftarrow \text{MAC}(\text{sk}, 0), t \leftarrow \{0, 1\}^\lambda),$$

notably without sampling r nor explicitly sampling the pairwise independent hash function h .

Instead, the choice of the pairwise independent hash function implicitly defines $r = t \oplus h(\sigma)$, and the adversary manages to produce a valid verification query only if $t^* \oplus h(\sigma^*) = t \oplus h(\sigma)$, which, over the random choice of $h \leftarrow \mathcal{H}$, happens with probability at most $1/2^\lambda$. In particular, the reduction outputs 1 with probability at most $1/2^\lambda$.

Overall, the advantage of our reduction is at least $\frac{\varepsilon_C}{Q_T Q_V} - \frac{1}{2^\lambda}$.

Last, observe that $\varepsilon_{FF} + \varepsilon_C = \varepsilon'$, so that either $\varepsilon_{FF} \geq \varepsilon'/2$ or $\varepsilon_C \geq \varepsilon'/2$. In particular, if $\varepsilon_{FF} \geq \varepsilon'/2$ then $\varepsilon_{FF}/Q_V \geq 2Q_T (\varepsilon + 1/2^\lambda) \geq \varepsilon$, and if $\varepsilon_C \geq \varepsilon'/2$ then $\frac{\varepsilon_C}{Q_T Q_V} - \frac{1}{2^\lambda} \geq \varepsilon$, which concludes the proof. \square

Combining the constructions and transformations from Sections 6.2 and 6.3 with $k = O(\sqrt{(m)})$, $n = O(\lambda\sqrt{m} + m)$, $\ell = |\mu| = O(k\lambda)$, we obtain Theorem 6.1.

Remark 6.22 (Security against Adversaries with Temporary Unbounded Memory). Our resulting MAC of Theorem 6.1 also inherits security against adversaries with unbounded temporary memory. In particular security holds even given non-uniform advice of size m .

7 Public-Key Signatures

In this section, we show how to build signature schemes in the streaming BSM. More precisely, we prove the following:

Theorem 7.1. *For all m, λ , there exists a streaming signature scheme with memory $n = \tilde{O}(\lambda^3 + \sqrt{m\lambda})$ (Definition 3.9) which can authenticate messages of length up to 2^λ , and which is $(2^{-\Omega(\lambda)}, \mathcal{A})$ -uf-cmva-secure (Definition 3.10) where \mathcal{A} is the set of streaming adversaries with memory m that make a total number of at most $Q = 2^\lambda$ oracle queries in the unforgeability experiment (Definition 3.11). Furthermore, (streamed) signatures have size $\tilde{O}(m + \lambda^{3/2})$.*

To do so, we first present a *set key-agreement protocol* in Section 7.1, using our block entropy lemma Lemma 4.1 as our main technical tool. Then we show how to upgrade such a protocol to a signature scheme in Section 7.2. Last, we show that the quadratic gap between the adversary's memory bound and the honest users is optimal (up to poly(λ) factors), in Section 7.3 (Theorem 7.11).

7.1 Set Key-Agreement Protocol

Given any parameters m, λ and ℓ , we define additional parameters: $b = 8\lambda(\ell + 2)$, $k = \max(\lceil 4m/b \rceil + 4, 64\lambda)$, $q = \lceil 2\sqrt{k\lambda} \rceil$, which guarantees $k \geq 2m/b + 2q$. Let $\text{Ext} : \{0, 1\}^b \times \{0, 1\}^d \rightarrow \{0, 1\}^\ell$ be a $(\ell + 2\lambda, 2^{-\lambda})$ -seeded extractor with some seed-length d , that can be computed using $O(b)$ space, as guaranteed by Theorem 2.6. Consider the following “set key agreement” protocol between Alice and Bob, with one round of communication from Alice to Bob.

- $\overline{\text{KeyGen}}$: Alice stream a value $\text{vk} = (x = (x_1, \dots, x_k), \mathcal{S}_A, (\text{seed}_1, \dots, \text{seed}_q))$ generated as follows.
 - She chooses a uniformly random subset $\mathcal{S}_A \subseteq [k]$ of size $|\mathcal{S}_A| = q$ and stores it in memory as an ordered tuple $\mathcal{S}_A = (i_1^A, \dots, i_q^A)$.
 - She streams a uniformly random value $x = (x_1, \dots, x_k) \leftarrow (\{0, 1\}^b)^k$ and additionally stores the values $x_{\mathcal{S}_A} = (x_{i_1^A}, \dots, x_{i_q^A})$ in memory.
 - She streams her set \mathcal{S}_A .
 - She chooses q random extractor seeds $\text{seed}_j \leftarrow \{0, 1\}^d$ and computes $\text{sk}_j = \text{Ext}(x_{i_j^A}; \text{seed}_j)$ for $j \in [q]$. She sends $(\text{seed}_1, \dots, \text{seed}_q)$ and stores $\text{sk} = (\text{sk}_1, \dots, \text{sk}_q)$ in memory.

Her final outputs consists of the secret key $\text{sk} = (\text{sk}_1, \dots, \text{sk}_q) \in (\{0, 1\}^\ell)^q$ stored in memory.

- $\overline{\text{KeyReceive}}$: Bob processes the stream $\mathbf{vk} = (x = (x_1, \dots, x_k), \mathcal{S}_A, (\text{seed}_1, \dots, \text{seed}_q))$ as follows:
 - He chooses a uniformly random subset $\mathcal{S}_B \subseteq [k]$ of size $|\mathcal{S}_B| = q$ and stores it in memory.
 - As he receives the stream x , he additionally stores $x_{\mathcal{S}_B}$ in memory.
 - When he receives $\mathcal{S}_A = (i_1^A, \dots, i_q^A)$ he stores it in memory.
 - When he receives $(\text{seed}_1, \dots, \text{seed}_q)$ he computes $\mathcal{T} = \{j \in [q] : i_j^A \in \mathcal{S}_B\}$. For each $j \in \mathcal{T}$ he computes $\text{sk}_j = \text{Ext}(x_{i_j^A}; \text{seed}_j)$. He stores the value $(\mathcal{T}, \text{sk}_{\mathcal{T}} = (\text{sk}_j)_{j \in \mathcal{T}})$ in memory.

Bob's final outputs consists of $\mathbf{vd} = (\mathcal{T}, \text{sk}_{\mathcal{T}} = (\text{sk}_j)_{j \in \mathcal{T}})$ stored in memory.

Lemma 7.2. *The procedures $\overline{\text{KeyGen}}$ and $\overline{\text{KeyReceive}}$ can be computed by streaming algorithms with memory $\tilde{O}(bq) = \tilde{O}(\lambda^2 \ell + \lambda \sqrt{\ell \cdot \lambda})$. Furthermore sk and \mathbf{vd} have size $\tilde{O}(q \cdot \ell) = \tilde{O}(\sqrt{m\ell} + \lambda \ell)$.*

Lemma 7.3 (Set Key-Agreement Lemma). *Let Eve be a streaming attacker with m bits of memory, who observes the stream \mathbf{vk} and outputs view_{Eve} in the above protocol. Let $\text{sk} = (\text{sk}_1, \dots, \text{sk}_q)$ be Alice's output and let $\mathbf{vd} = (\mathcal{T}, \text{sk}_{\mathcal{T}})$ be Bob's output in the protocol. Then there is some index $t \in \mathcal{T} \cup \{\perp\}$ defined as a random variable depending on the entire protocol execution and the view of Eve, and some $\varepsilon = 2^{-\Omega(\lambda)}$ such that:*

$$(\text{view}_{Eve}, \mathcal{T}, t, \text{sk}_{-t} = (\text{sk}_j : j \in [q] \setminus \{t\}), \text{sk}_t) \approx_{\varepsilon} (\text{view}_{Eve}, \mathcal{T}, t, \text{sk}_{-t} = (\text{sk}_j : j \in [q] \setminus \{t\}), u)$$

where $u \leftarrow \{0, 1\}^{\ell}$ is random and independent of all other values, and we define $\text{sk}_t := \perp$ if $t = \perp$.

Proof. At the end of the protocol execution, we select the index t as follows:

1. If $|\mathcal{S}_A \cap \mathcal{S}_B| < \lambda$ then set $t = \perp$. (We refer to this event as \perp_0 .)
2. Choose a uniformly random $\mathcal{V} \subseteq \mathcal{S}_A \cap \mathcal{S}_B$ such that $|\mathcal{V}| = \lambda$. Let $\mathcal{W} = \mathcal{S}_A \setminus \mathcal{V}$.
3. Let view_{Eve}^0 denote the view of Eve immediately after processing the x component of the stream but before receiving \mathcal{S}_A .
4. Let $\text{aux} = (\text{view}_{Eve}^0, \mathcal{W}, x_{\mathcal{W}})$, and let AUX be a random variable corresponding to aux . Let $X = (U_{kb} | \text{AUX} = \text{aux})$ be the random variable corresponding to choosing $x \leftarrow (\{0, 1\}^b)^k$ from the uniform distribution, conditioned on $\text{AUX} = \text{aux}$. If, for the given value of aux , we have $\mathbf{H}_{\infty}(X) < kb/2$ then set $t = \perp$. (We refer to this event as \perp_1 .)
5. Let the sets $\text{BAD} \subseteq (\{0, 1\}^b)^k$ and $\mathcal{I}(x) \subseteq [k]$ be the sets from Lemma 4.1 defined with respect to the distribution X with $\alpha = 1/2$, $\alpha_1 = \alpha_2 = 1/8$. Let x be the value sent during the protocol execution. If $\mathcal{I}(x) \cap \mathcal{V} = \emptyset$ then set $t = \perp$. (We refer to this event as \perp_2 .)
6. Let i^* be the smallest value in $\mathcal{I}(x) \cap \mathcal{V} \subseteq \mathcal{S}_A = \{i_1^A, \dots, i_q^A\}$ and let t^* be the value such $i_{t^*}^A = i^*$. Set $t = t^*$.

We show a sequence of hybrid distributions that are statistically close.

Hybrid 0. This is the distribution on the left-hand side of the Lemma

$$(\text{view}_{Eve}, \mathcal{T}, t, \text{sk}_{-t} = (\text{sk}_j : j \in [q] \setminus \{t\}), \text{sk}_t)$$

Hybrid 1. We now change how we select the sets $\mathcal{S}_A, \mathcal{S}_B, \mathcal{V}, \mathcal{W}$.

In hybrid 0, $\mathcal{S}_A, \mathcal{S}_B \subseteq [k]$ are chosen uniformly at random with $|\mathcal{S}_A| = |\mathcal{S}_B| = q$. Then, if $|\mathcal{S}_A \cap \mathcal{S}_B| > \lambda$, we choose $\mathcal{V} \subseteq \mathcal{S}_A \cap \mathcal{S}_B$ of size $|\mathcal{V}| = \lambda$ uniformly at random and define $\mathcal{W} = \mathcal{S}_A \setminus \mathcal{V}$.

In hybrid 1, we instead choose $\mathcal{W} \subseteq [k]$ of size $|\mathcal{W}| = q - \lambda$ uniformly at random. Then we select $\mathcal{V} \subseteq [k] \setminus \mathcal{W}$ of size $|\mathcal{V}| = \lambda$ uniformly at random. We define $\mathcal{S}_A = \mathcal{V} \cup \mathcal{W}$. Then we choose $\mathcal{S}_B \subseteq [k]$ of size $|\mathcal{S}_B| = q$ uniformly at randomly subject to $\mathcal{V} \subseteq [k]$.

Note that the above change ensures that, in hybrid 1, $\mathcal{V} \subseteq \mathcal{S}_A \cap \mathcal{S}_B$ with $|\mathcal{V}| = \lambda$, and therefore the event \perp_0 never occurs.

Hybrids 1 is distributed identically to hybrid 0 if we condition on \perp_0 not occurring in hybrid 0. Therefore the statistical distance between the hybrids is bounded by the probability that \perp_0 occurs, meaning that $|\mathcal{S}_A \cap \mathcal{S}_B| < \lambda$, in hybrid 0. This is bounded by $\varepsilon_1 \leq 2^{-\lambda/4} = 2^{-\Omega(\lambda)}$ by Lemma 2.13 as $q \geq 2\sqrt{k\lambda}$.

Hybrid 2. In hybrid 1, if the event \perp_1 occurs (i.e., \mathbf{aux} is such hat $\mathbf{H}_\infty(X) < kb/2$) then we set $t = \perp$ and $\mathbf{sk}_t = \perp$. In hybrid 2, if \perp_1 occurs then we set $t = \perp$ and choose $\mathbf{sk}_t \leftarrow \{0, 1\}^\ell$ uniformly at random.

Hybrids 1 and 2 are distributed identically as long as \perp_1 does not occur, and so the statistically distance between them is bounded by the probability that $\mathbf{H}_\infty(X) < kb/2$. Note that $X = (U_{kb} | \mathbf{AUX} = \mathbf{aux})$. Furthermore, we have:

$$\mathbf{H}_\infty(U_{kb} | \mathbf{AUX}) \geq \mathbf{H}_\infty(U_{kb} | \mathcal{W}) - m - (q - \lambda) \cdot b \geq kb - m - (q - \lambda) \cdot b \geq kb/2 + \lambda,$$

where the first inequality follows from Lemma 2.1 since $\mathbf{aux} = (\mathbf{view}_{Eve}^0, \mathcal{W}, x_{\mathcal{W}})$ where $|\mathbf{view}_{Eve}^0| = m$ and $|x_{\mathcal{W}}| = (q - \lambda) \cdot b$, and the second inequality follows since \mathcal{W} is chosen independently of $x \leftarrow U_{kb}$ and hence does not reduce entropy. Therefore, by Lemma 2.2:

$$\Pr_{\mathbf{aux} \leftarrow \mathbf{AUX}} \left[\underbrace{\mathbf{H}_\infty(U_{kb} | \mathbf{AUX} = \mathbf{aux})}_{\mathbf{H}_\infty(X)} < kb/2 \right] \leq 2^{-\lambda},$$

and so the statistical distance between the hybrids is $\varepsilon_2 \leq 2^{-\lambda}$.

Hybrid 3. In hybrid 3, if \perp_1 does not occur and $x \in \text{BAD}$, then we define $t = \perp$ and choose $\mathbf{sk}_t \leftarrow \{0, 1\}^\ell$ uniformly at random. We refer to this event as $\perp_{1.5}$.

Hybrids 1 and 2 are distributed identically unless: \perp_1 does not occur and $x \in \text{BAD}$. But, if we fix any \mathbf{aux} for which \perp_1 does not occur, then by Lemma 4.1 we can bound the probability that $x \in \text{BAD}$ by $k \cdot 2^{-b/4} = 2^{-\Omega(\lambda)}$. Therefore, the statistical distance between the hybrids is $\varepsilon_3 \leq 2^{-\Omega(\lambda)}$.

Hybrid 4. In hybrid 3, if \perp_1 and $\perp_{1.5}$ don't occur but $|\mathcal{I}(x) \cap \mathcal{V}| = \emptyset$ then we set $t = \perp$ and $\mathbf{sk}_t = \perp$. In hybrid 4, if \perp_1 and $\perp_{1.5}$ don't occur but $|\mathcal{I}(x) \cap \mathcal{V}| = \emptyset$, then we set $t = \perp$ and choose $\mathbf{sk}_t \leftarrow \{0, 1\}^\ell$ uniformly at random.

Hybrids 3 and 4 are distributed identically unless: \perp_1 and $\perp_{1.5}$ don't occur but $|\mathcal{I}(x) \cap \mathcal{V}| = \emptyset$. Let us fix any \mathbf{aux}, x such that $\perp_1, \perp_{1.5}$ do not occur. This also fixes $\mathcal{I}(x)$. Furthermore, by Lemma 4.1, we have $|\mathcal{I}(x)| \geq k/8$. Moreover, $\mathcal{I}(x) \subseteq [k] \setminus \mathcal{W}$ since, for $i \in \mathcal{W}$, the values X_i are completely fixed by \mathbf{aux} and hence have entropy 0. On the other hand \mathcal{V} is uniformly random over $[k] \setminus \mathcal{W}$ with $|\mathcal{V}| = \lambda$, and independent of $\mathcal{I}(x)$. Therefore the probability that $\mathcal{V} \cap \mathcal{I}(x) = \emptyset$ is bounded by $\varepsilon_4 \leq (7/8)^\lambda = 2^{-\Omega(\lambda)}$, which also bounds the statistical distance between these hybrids.

Hybrid 5. We undo the change introduced in hybrid 3. That is, we no longer check if $x \in \text{BAD}$ and take any special action if it is.

Hybrids 4 and 5 are statistically close for the same reason hybrids 2 and 3 are statistically close, with distance $\varepsilon_5 \leq 2^{-\Omega(\lambda)}$.

Hybrid 6. In hybrid 6, we now always choose $\text{sk}_t \leftarrow \{0, 1\}^\ell$ uniformly random.

We argue that hybrids 5 and 6 are statistically close by the security of the extractor. Let us fix and condition on any choice of values

$$(\text{aux} = (\text{view}_{Eve}^0, \mathcal{W}, x_{\mathcal{W}}), \mathcal{V}, \mathcal{S}_B, i^*, x_1, \dots, x_{i^*-1}, (\text{seed}_j)_{j \in [q] \setminus t^*})$$

chosen during the experiment, subject to \perp_1, \perp_2 not occurring. We define i^* as being the smallest value in $\mathcal{V} \cap \mathcal{I}(x)$ and t^* is defined as the value such $i_{t^*}^A = i^*$, which is fixed once i^* and $\mathcal{S}_A = \mathcal{V} \cup \mathcal{W} = \{i_1^A, \dots, i_q^A\}$ are fixed. The above values also fix $\text{sk}_{\mathcal{W}} := (\text{sk}_j = \text{Ext}(x_{i_j^A}; \text{seed}_j)_{i_j^A \in \mathcal{W}})$.

If either \perp_1 occurs or \perp_2 occurs then the hybrids are identical, and therefore it suffices to only show that they are statistically close for any fixed choice of the values as above for which \perp_1, \perp_2 do not occur.

Note that, by Lemma 4.1, fixing x_1, \dots, x_{i^*-1} also fixes whether $i \in \mathcal{I}(x)$ for all $i \leq i^*$ and therefore, once we fix all of the above, conditioning on i^* being the smallest value in $\mathcal{V} \cap \mathcal{I}(x)$ is equivalent to just conditioning on $i^* \in \mathcal{I}(x)$. Therefore, the distribution of X_{i^*} (i.e., the i^* block of x) conditioned on the above is equivalent to $(X_{i^*} | X_1 = x_1, \dots, X_{i^*-1} = x_{i^*-1}, i^* \in \mathcal{I}(x))$, and By Lemma 4.1, it has min-entropy $\mathbf{H}_\infty(X_{i^*}) \geq b/8$.⁹

Let $\text{SK}_{\mathcal{V}^-} := (\text{SK}_j = \text{Ext}(X_{i_j^A}; \text{seed}_j)_{i_j^A \in \mathcal{V} \setminus \{i^*\}})$. Then

$$\mathbf{H}_\infty(X_{i^*} | \text{SK}_{\mathcal{V}^-}) \geq \mathbf{H}_\infty(X_{i^*}) - (\lambda - 1) \cdot \ell \geq b/8 - (\lambda - 1) \cdot \ell \geq \ell + 2\lambda.$$

Let SEED_{t^*} to be a random variable for seed_{t^*} . Then

$$(\text{SK}_{t^*} = \text{Ext}(X_{i^*}; \text{SEED}_{t^*}), \text{SEED}_{t^*}, \text{SK}_{\mathcal{V}^-}) \approx_{\varepsilon_6} (U_\ell, \text{SEED}_{t^*}, \text{SK}_{\mathcal{V}^-})$$

where $\varepsilon_6 \leq 2^{-\lambda}$, by the security of the extractor.

Now observe that, conditioned on the fixed values, the outputs of hybrid 5 and 6 are completely defined given the additional values $\text{SEED}_{t^*}, \text{SK}_{\mathcal{V}^-}$ and either $\text{SK}_{t^*} = \text{Ext}(X_{i^*}; \text{SEED}_{t^*})$ in hybrid 5 or U_ℓ in hybrid 6. In particular, everything else in the hybrid is defines as follows:

- \mathcal{T}, t are completely determined by the fixed values $\mathcal{S}_A = \mathcal{V} \cup \mathcal{W} = \{i_1^A, \dots, i_q^A\}, \mathcal{S}_B, i^*$.
- view_{Eve} is defined in terms of $\text{view}_{Eve}^0, \mathcal{S}_A, \mathcal{S}_B, (\text{seed}_j)_{j \in [q] \setminus t^*}, \text{SEED}_{t^*}$.
- sk_{-t} consists of $\text{SK}_{\mathcal{V}^-}$ and $\text{sk}_{\mathcal{W}} := (\text{sk}_j = \text{Ext}(x_{i_j^A}, \text{seed}_j)_{i_j^A \in \mathcal{W}})$, where the latter only depends on the fixed values $\mathcal{W}, x_{\mathcal{W}}, (\text{seed}_j)_{j \in [q] \setminus t^*}$.
- And the last component in the hybrid is SK_{t^*} in hybrid 5 and U_ℓ in hybrid 6.

Therefore, the statistical distance between the hybrids is bounded by $\varepsilon_6 \leq 2^{-\lambda}$.

Hybrid 7. Undo the change from Hybrid 1 and select $\mathcal{S}_A, \mathcal{S}_B, \mathcal{V}, \mathcal{W}$ as in Hybrid 0.

Hybrids 6 and 7 are statistically close for the same reason hybrids 0 and 1 are statistically close, with distance $\varepsilon_7 \leq 2^{-\Omega(\lambda)}$.

This hybrid is equivalent to the right-hand side distribution of the Lemma.

$$(\text{view}_{Eve}, \mathcal{T}, t, \text{sk}_{-t} = (\text{sk}_j : j \in [q] \setminus \{t\}), u)$$

where $u \leftarrow \{0, 1\}^\ell$ is random and independent of all other values.

Combining the above hybrids, the Lemma holds with $\varepsilon \leq \sum_{i=1}^7 \varepsilon_i = 2^{-\Omega(\lambda)}$. □

⁹For the rest, of the argument we will condition on all the fixed values implicitly and will not write this conditioning explicitly.

7.2 From Set Key Agreement to Signatures

Construction. Let $\ell, m, |\mu|$ be parameters. Let $\mathcal{A} = \mathcal{A}(O(m), Q)$ be any class of *non-uniform* streaming algorithms such that there exists a constant C' such that algorithms in \mathcal{A} run with memory at most $C'm$ making at most Q oracle calls (Definition 3.5). Let $(\text{KeyGen}, \text{MAC}, \text{Verify})$ be a streaming MAC with memory ℓ and message space $\{0, 1\}^{|\mu|}$ satisfying $(\varepsilon, \mathcal{A})$ -uf-cmva-security (Definition 3.3) for some $\varepsilon > 0$, such that signing keys (of size at most ℓ) are uniformly random.

Let b, k, q , be the parameters instantiated in Section 7.1, and let $\text{Ext} : \{0, 1\}^b \times \{0, 1\}^d \rightarrow \{0, 1\}^\ell$ be a $(\ell + 2\lambda, 2^{-\lambda})$ -seeded extractor.

We define a streaming signature scheme $(\overline{\text{KeyGen}}, \overline{\text{KeyReceive}}, \overline{\text{Sign}}, \overline{\text{Verify}})$ as follows.

- $\overline{\text{KeyGen}}(1^\lambda)$: stream vk and store $\text{sk} = (\text{sk}_1, \dots, \text{sk}_q)$, as defined in Section 7.1.
- $\overline{\text{KeyReceive}}(1^\lambda, \text{vk})$: on input a streamed verification key vk , store $\text{vd} = (\mathcal{T}, \text{sk}_\mathcal{T} = (\text{sk}_j)_{j \in \mathcal{T}})$ as specified in Section 7.1.
- $\overline{\text{Sign}}(\text{sk}, \mu)$: on input a (potentially streamed) message μ , compute and output (potentially in a streaming manner):

$$\sigma = \{\sigma_j = \text{MAC}(\text{sk}_j, \mu)\}_{j \in [q]}.$$

- $\overline{\text{Verify}}(\text{vd}, \mu, \sigma)$: on input $\text{vd} = (\mathcal{T}, \{\text{sk}_j\}_{j \in \mathcal{T}})$, a (potentially streamed) message μ , and a (potentially streamed) signature $\sigma = \{\sigma_j\}_{j \in \mathcal{T}}$, output

$$\bigwedge_{i \in \mathcal{T}} \text{Verify}(\text{sk}_j, \mu, \sigma_j),$$

with the convention that it outputs 1 if $\mathcal{T} = \emptyset$.

Claim 7.4 (Correctness). *Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is correct. Then $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$ is correct.*

Claim 7.5 (Efficiency). *Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ be a streaming MAC with memory ℓ (Definition 3.2). Then $\overline{\text{KeyGen}}$ and $\overline{\text{KeyReceive}}$ can be computed by streaming algorithms with memory $\tilde{O}(\lambda^2 \ell + \lambda \sqrt{\ell \cdot \lambda})$, and sk and vd have size $\tilde{O}(\sqrt{m\ell} + \lambda\ell)$. $\overline{\text{Sign}}$ and $\overline{\text{Verify}}$ can be computed by streaming algorithms with memory $O(\sqrt{m\ell} + \lambda\ell)$.*

Theorem 7.6 (Security). *Let $\mathcal{A}_{\text{MAC}} = \mathcal{A}(O(m), Q)$ be any class of non-uniform¹⁰ streaming algorithms such that there exists a constant C' such that algorithms in \mathcal{A}_{MAC} run with memory at most $C'm$ making at most Q oracle calls (Definition 3.5), and \mathcal{A}_{Sig} be its signature counterpart (Definition 3.11). Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is $(\varepsilon, \mathcal{A}_{\text{MAC}})$ -uf-cmva-secure (Definition 3.3). Then $(\overline{\text{KeyGen}}, \overline{\text{KeyReceive}}, \overline{\text{Sign}}, \overline{\text{Verify}})$ is $(2\varepsilon, \mathcal{A}_{\text{Sig}})$ -uf-cmva-secure (Definition 3.10).*

Proof. Let $\text{Adv} \in \mathcal{A}_{\text{MAC}}$ be a streaming adversary with memory m for the uf-cmva experiment for $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$. Consider the following hybrid experiments.

Hybrid H_0 . This corresponds to the uf-cmva experiment $\text{Exp}^{\text{uf-cmva}}(1^\lambda, \text{Adv})$.

Hybrid H_1 . We change how $\overline{\text{KeyReceive}}$ in Step 1 of $\text{Exp}^{\text{uf-cmva}}(1^\lambda, \text{Adv})$ is computed. vd is now computed as $(\mathcal{T}, \text{sk}_{-t}, u)$ where $u \leftarrow \{0, 1\}^\ell$. In other words, the secret key sk_t is replaced by uniformly random, where t is given by Lemma 7.3.

Lemma 7.7. *Suppose b, k, q are instantiated as in Section 7.1. Then the advantage of Adv decreases by at most ε between H_0 and H_1 :*

$$\left| \Pr \left[\text{Exp}_{H_1}^{\text{uf-cmva}}(1^\lambda, \text{Adv}) = 1 \right] - \Pr \left[\text{Exp}_{H_0}^{\text{uf-cmva}}(1^\lambda, \text{Adv}) = 1 \right] \right| \leq \varepsilon,$$

where, for $b \in \{0, 1\}$, $\text{Exp}_{H_b}^{\text{uf-cmva}}$ denotes the uf-cmva experiment in Hybrid H_b .

¹⁰Looking ahead, this will not affect our final result, because our base MAC is secure against non-uniform adversaries.

Proof. Denote by view_{Eve} the state of Adv after Step 1 of $\text{Exp}^{\text{uf-cmva}}(1^\lambda, \text{Adv})$. Let $\text{sk} = (\text{sk}_j)_{j \in [q]}$ and let $\text{vd} = (\mathcal{T}, (\text{sk}_j)_{\mathcal{T}})$ be the respective outputs of $\overline{\text{KeyGen}}$ and $\overline{\text{KeyReceive}}$ in Step 1.

The lemma follows as the output of $\text{Exp}^{\text{uf-cmva}}(1^\lambda, \text{Adv})$ can be computed using $(\text{view}_{Eve}, \mathcal{T}, t, \text{sk}_{-t} = (\text{sk}_j)_{j \neq t}, \text{sk}_t)$ alone (with the convention that $\text{Verify}(\perp, \mu, \sigma) = 1$ for all μ, σ). Namely, run Adv starting from Step 2 with state view_{Eve} . Signing queries are answered using the knowledge of $\text{sk}_{-t}, \text{sk}_t$ and t . Verification queries as well as the final forgery are answered using the knowledge of $\text{sk}_{\mathcal{T} \setminus \{t\}}, \mathcal{T}, \text{sk}_t$ and t .

By Lemma 7.3:

$$(\text{view}_{Eve}, \mathcal{T}, t, \text{sk}_{-t} = (\text{sk}_j)_{j \neq t}, \text{sk}_t) \approx_\varepsilon (\text{view}_{Eve}, \mathcal{T}, t, \text{sk}_{-t} = (\text{sk}_j)_{j \neq t}, u).$$

Now, computing the output of $\text{Exp}^{\text{uf-cmva}}(1^\lambda, \text{Adv})$ using $(\text{view}_{Eve}, \mathcal{T}, t, \text{sk}_{-t} = (\text{sk}_j)_{j \neq t}, \text{sk}_t)$ corresponds to $\text{Exp}_{H_0}^{\text{uf-cmva}}(1^\lambda, \text{Adv})$, and using $(\text{view}_{Eve}, \mathcal{T}, t, \text{sk}_{-t} = (\text{sk}_j)_{j \neq t}, u)$ corresponds to $\text{Exp}_{H_1}^{\text{uf-cmva}}(1^\lambda, \text{Adv})$, and the lemma follows. \square

Lemma 7.8. *Suppose $(\text{KeyGen}, \text{MAC}, \text{Verify})$ is $(\varepsilon, \mathcal{A})$ -uf-cmva-secure. Then*

$$\Pr[\text{Exp}_{H_1}^{\text{uf-cmva}}(1^\lambda, \text{Adv}) = 1] \leq \varepsilon.$$

Proof. Let Adv be an adversary for $(\overline{\text{KeyGen}}, \overline{\text{KeyReceive}}, \overline{\text{Sign}}, \overline{\text{Verify}})$ with memory m having advantage ε' in $\text{Exp}_{H_1}^{\text{uf-cmva}}$. We build a (non-uniform) adversary against the uf-cmva-security of $(\text{KeyGen}, \text{MAC}, \text{Verify})$ with memory $m + O(q\ell)$ and advantage ε' as follows.

Our reduction computes $\overline{\text{KeyGen}}$, streams vk to Adv and stores a secret key $(\text{sk}_i)_{i \in [q]}$. It computes $\text{vd} = (\mathcal{T}, \text{sk}_{\mathcal{T}})$. It receives the (inefficiently computable) index t defined in Lemma 7.3 as non-uniform advice: note that t only depends on the execution of the set key agreement and the resulting view of the adversary, and is therefore independent of the unforgeability experiment for $(\text{KeyGen}, \text{MAC}, \text{Verify})$.

To answer signing queries on message μ , the reduction computes for all $j \in [q], j \neq t$: $\sigma_j = \text{Sign}(\text{sk}_j, \mu)$, and makes a signing query μ to the MAC oracle (which implicitly uses $\text{MAC.sk} = \text{sk}_t$), thus obtaining σ_t . It forwards $(\sigma_j)_{j \in [q]}$ to Adv (using its knowledge of t).

To answer verification queries with message μ and signature $\sigma = (\sigma_j)_{j \in [q]}$, the reduction computes for all $j \in \mathcal{T}, j \neq t$: $b_j \leftarrow \text{Verify}(\text{sk}_j, \mu, \sigma_j)$ and makes a verification query with input (μ, σ_j) to Verify oracle, obtaining a bit b_t . It outputs $\bigwedge_{j \in \mathcal{T}} b_j$. The final output of the experiment is computed identically.

Our reduction runs in memory $m + q\ell$, makes the same number of queries Q as Adv , and succeeds if Adv successfully produces a forgery. Therefore its advantage is at least ε' which is at most ε by uf-cmva-security of $(\text{KeyGen}, \text{KeyReceive}, \text{Sign}, \text{Verify})$. \square

Overall, the advantage of Adv in the uf-cmva experiment for $(\overline{\text{KeyGen}}, \overline{\text{MAC}}, \overline{\text{Verify}})$ is at most

$$\Pr[\text{Exp}_{H_0}^{\text{uf-cmva}}(1^\lambda, \text{Adv}) = 1] \leq 2\varepsilon.$$

\square

We instantiate $(\text{KeyGen}, \text{MAC}, \text{Verify})$ with our construction in Section 5, which has uniformly random MAC keys, using $|\mu| = 2^\lambda$. Recall that this construction is secure against non-uniform adversaries (and in fact, adversaries with unbounded temporary memory, see ??). Setting $\ell = O(\lambda + \text{polylog}(m, \lambda))$, gives $n = \tilde{O}(\lambda^3 + \sqrt{m\lambda})$. Combined with the following observation with $\tau = m$, we obtain Theorem 7.1.

Remark 7.9 (Optimizing the Communication Cost). The signatures for our scheme consist of $[q]$ independent copies of $\text{MAC}(\text{sk}_j, \mu)$, $j \in [q]$, where MAC is our base MAC. Using our scheme from Section 5, and noting that (1) security only relies on security of a single of the copies, and (2) tags of the form (x, ψ) can be computed given x and sk , we can instead compute our signatures as $(x, (\psi_j)_{j \in [q]})$, namely reusing the x part across the different copies of the base MAC, while preserving correctness and security. This makes our signatures of size $\tilde{O}(\tau + q\lambda)$, where τ is the size of the tags from MAC .

Remark 7.10 (Weaker Notions of Security). Our construction constructs a uf-cmva -secure signature scheme starting from any uf-cmva -secure MAC. We note that this extends to weaker notions of security, namely starting from a (selectively)-unforgeable MAC with signing (and verification) queries, one obtains a signature scheme satisfying the same notion of security.

7.3 Lower Bound for Signatures

We show here that any streaming signature with memory n , namely, where all the procedures can be run in a streaming manner with memory n , can only be secure against streaming adversaries with memory $m = O(n^2)$.

Theorem 7.11. *Suppose $(\text{KeyGen}, \text{KeyReceive}, \text{Sign}, \text{Verify})$ is a streaming signature with memory n (Definition 3.9). Let $\mathcal{A} = \mathcal{A}(m)$ be the set of streaming adversaries running with memory m (Definition 3.11). Suppose that $(\text{KeyGen}, \text{KeyReceive}, \text{Sign}, \text{Verify})$ is $(\varepsilon, \mathcal{A})$ -suf-cma-secure (Definition 3.10). Then $m = O(n^2)$.*

Proof. We consider the following adversary Adv with memory m . Let $q = m/2n$, and $i \in [q]$ be an index to be determined later.

- 0 It declares an arbitrary message $\mu \in \mathcal{M}$ for its forgery.
- 1. On input vk , it executes i copies of KeyReceive , obtaining $\text{vd}_1, \dots, \text{vd}_i$, which it stores.
- 1' It samples a secret key $\tilde{\text{sk}}$ as follows. It runs $\text{KeyGen}(1^\lambda)$, and processes the stream vk with q copies of KeyReceive . If the i outputs of KeyReceive are equal to $\text{vd}_1, \dots, \text{vd}_i$, it stores the key $\tilde{\text{sk}}$ output by KeyGen , and repeats the whole step otherwise.
- 2. It outputs as a forgery $(\mu, \text{Sign}(\tilde{\text{sk}}, \mu))$.

First, note that such an adversary is streaming, and uses memory m . This is because Step 1 can be computed in a streaming manner using memory $in \leq qn = m/2$, and Step 2 uses an additional $qn = m/2$ memory. Step 3 uses memory n , but Adv does not need to store $(\text{vd}_1, \dots, \text{vd}_i)$ at this point. So the total memory cost of Adv is $\max(m, n) = m$. Note also that Adv does not query any of the oracles in the suf-cma game.

We now argue that such an adversary has high success probability.

Claim 7.12. *We have:*

$$\Pr \left[\text{Exp}^{\text{suf-cma}}(1^\lambda, \text{Adv}) = 1 \right] \geq 1 - \sqrt{\frac{n}{2(q+1)}} - \text{negl}(\lambda).$$

In particular if $q = \Omega(n)$, then Adv has a constant probability of outputting a forgery, so that Claim 7.12 implies Theorem 7.11.

To prove the claim, our central tool is the following lemma, previously used in [DQW21].

Lemma 7.13. *Let SK and VD denote random variables corresponding to the secret key and states respectively output by KeyGen and KeyReceive stored by the challenger in Step 1 of the suf-cma experiment, and VD_j , $j \leq i$ and $\tilde{\text{SK}}$ denote the quantities vd_j and $\tilde{\text{sk}}$ computed by Adv in Step 1 and 1', respectively. Then there exists some $i \in [q]$ such that*

$$\text{SD} \left((\text{SK}, \text{VD}, \text{VD}_1, \dots, \text{VD}_i), (\tilde{\text{SK}}, \text{VD}, \text{VD}_1, \dots, \text{VD}_i) \right) \leq \sqrt{\frac{n}{q+1}}.$$

Proof. This follows from the proofs from [DM04, DQW21]. We recall it for completeness.

Let $\text{VD}, \text{VD}_1, \dots, \text{VD}_q$ denote independent random variables denoting the outputs of $(q+1)$ independent executions of KeyReceive on input $\text{vk} \leftarrow \text{KeyGen}$, which outputs SK . The random variables $\text{VD}, \text{VD}_1, \dots, \text{VD}_q$ are symmetric ([DM04]) with respect to the random variable SK , namely

$$\Pr_{\text{SK}, \text{VD}_{i_1}, \dots, \text{VD}_{i_w}} (\text{SK} = x, \text{VD}_{i_1} = z_1, \dots, \text{VD}_{i_w} = z_w) = \Pr_{\text{SK}, \text{VD}_{i'_1}, \dots, \text{VD}_{i'_w}} (\text{SK} = x, \text{VD}_{i'_1} = z_1, \dots, \text{VD}_{i'_w} = z_w)$$

for all $w \leq q$, and all sets of distinct indices (i_1, \dots, i_w) and (i'_1, \dots, i'_w) . Combined with the chain rule, there exists $i \in [q]$ such that:

$$\mathbf{I}(\text{SK}; \text{VD} | \text{VD}_1, \dots, \text{VD}_i) \leq \frac{\mathbf{I}(\text{SK}; \text{VD}_0, \dots, \text{VD}_q)}{q+1} \leq \frac{\mathbf{H}(\text{SK})}{q+1} \leq \frac{n}{q+1}, \quad (1)$$

as SK has size at most n . This gives:

$$\begin{aligned} & \mathbf{SD} \left((\text{SK}, \text{VD}, \{\text{VD}_j\}_{j \leq i}), (\widetilde{\text{SK}}, \text{VD}, \{\text{VD}_j\}_{j \leq i}) \right) \\ & \leq \mathbf{E}_{\{\text{VD}_j\}} \mathbf{SD} \left(((\text{SK}, \text{VD}) | \{\text{VD}_j\}_{j \leq i}), (\widetilde{\text{SK}}, \text{VD}) | \{\text{VD}_j\}_{j \leq i} \right) \\ & \leq \mathbf{E}_{\{\text{VD}_j\}} \sqrt{\frac{1}{2} \cdot D_{KL} \left(P_{(\text{SK}, \text{VD}_0) | \{\text{VD}_j\}_{j \leq i}} \| P_{\widetilde{\text{SK}} | \{\text{VD}_j\}_{j \leq i}} \times P_{\text{VD}_0 | \{\text{VD}_j\}_{j \leq i}} \right)} \\ & \leq \sqrt{\frac{1}{2} \cdot \mathbf{E}_{\{\text{VD}_j\}} D_{KL} \left(P_{(\text{SK}, \text{VD}_0) | \{\text{VD}_j\}_{j \leq i}} \| P_{\widetilde{\text{SK}} | \{\text{VD}_j\}_{j \leq i}} \times P_{\text{VD}_0 | \{\text{VD}_j\}_{j \leq i}} \right)} \\ & = \sqrt{\frac{1}{2} \cdot \mathbf{I}(\text{SK}; \text{VD}_0 | \text{VD}_1, \dots, \text{VD}_i)} \\ & \leq \sqrt{\frac{n}{2(q+1)}}, \end{aligned}$$

where P_X denotes the distribution associated to X , and where the second inequality follows by Pinsker's inequality (Lemma 2.11), the third by concavity of the square root function, the equality by Definition 2.10, and the last inequality by Eq. (1). \square

We set $i \leq q$ to be the index from Lemma 7.13. Note that it is not clear whether the index $i \in [q]$ can be computed with memory m ; instead, we provide it to Adv as non-uniform advice.

Last, to prove Claim 7.12, we note that the output of $\text{Exp}^{\text{suf-cma}}(1^\lambda, \text{Adv})$ can be computed as a function of $(\text{sk}, \text{vd}, \text{vd}_1, \dots, \text{vd}_i)$. But on input $(\text{sk}, \text{vd}, \text{vd}_1, \dots, \text{vd}_i)$, where $\text{sk} \leftarrow \text{KeyGen}$ is the key stored by the challenger in Step 1, the probability that the resulting experiment outputs 1 is $1 - \text{negl}(\lambda)$ by correctness of $(\text{KeyGen}, \text{KeyReceive}, \text{Sign}, \text{Verify})$. Now these two differ by at most $\sqrt{\frac{n}{2(q+1)}}$ by Lemma 7.13, which concludes the proof. \square

Remark 7.14 (Non-uniformity of the Lower Bound). Theorem 7.11 rules out signatures that are secure against *non-uniform* adversaries with bounded memory. We observe here that the result extends to rule out uniform adversaries as follows: we define an adversary Adv' identical to Adv except that it guesses the index $i \in [q]$ from Lemma 7.13 instead of receiving it as non-uniform advice. The experiment $\text{Exp}^{\text{suf-cma}}(1^\lambda, \text{Adv})$ outputs 1 at least when the guess for i is correct, and therefore Adv' has success probability at least

$$\Pr \left[\text{Exp}^{\text{suf-cma}}(1^\lambda, \text{Adv}') = 1 \right] \geq \frac{1}{q} \cdot \left(1 - \sqrt{\frac{n}{2(q+1)}} - \text{negl}(\lambda) \right)$$

which is non-negligible as long as $q = \Omega(n)$ and $1/q$ is non-negligible. This for instance applies to $n = \text{poly}(\lambda)$ and $q = \Theta(n)$.

References

- [ADR02] Y. Aumann, Yan Zong Ding, and M.O. Rabin. Everlasting security in the bounded storage model. *IEEE Transactions on Information Theory*, 48(6):1668–1680, 2002.
- [CCM98] Christian Cachin, Claude Crépeau, and Julien Marcil. Oblivious transfer with a memory-bounded receiver. In *39th FOCS*, pages 493–502. IEEE Computer Society Press, November 1998.
- [CM97] Christian Cachin and Ueli M. Maurer. Unconditional security against memory-bounded adversaries. In Burton S. Kaliski Jr., editor, *CRYPTO’97*, volume 1294 of *LNCS*, pages 292–306. Springer, Heidelberg, August 1997.
- [DHRS07] Yan Zong Ding, Danny Harnik, Alon Rosen, and Ronen Shaltiel. Constant-round oblivious transfer in the bounded storage model. *Journal of Cryptology*, 20(2):165–202, April 2007.
- [Din01] Yan Zong Ding. Oblivious transfer in the bounded storage model. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 155–170. Springer, Heidelberg, August 2001.
- [DKPW12] Yevgeniy Dodis, Eike Kiltz, Krzysztof Pietrzak, and Daniel Wichs. Message authentication, revisited. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 355–374. Springer, Heidelberg, April 2012.
- [DM02] Stefan Dziembowski and Ueli M. Maurer. Tight security proofs for the bounded-storage model. In *34th ACM STOC*, pages 341–350. ACM Press, May 2002.
- [DM04] Stefan Dziembowski and Ueli M. Maurer. On generating the initial key in the bounded-storage model. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 126–137. Springer, Heidelberg, May 2004.
- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008.
- [DQW21] Yevgeniy Dodis, Willy Quach, and Daniel Wichs. Speak much, remember little: Cryptography in the bounded storage model, revisited. Cryptology ePrint Archive, Report 2021/1270, 2021. <https://ia.cr/2021/1270>.
- [DR02] Yan Zong Ding and Michael O. Rabin. Hyper-encryption and everlasting security. In *Proceedings of the 19th Annual Symposium on Theoretical Aspects of Computer Science*, STACS ’02, page 1–26, Berlin, Heidelberg, 2002. Springer-Verlag.
- [GGM84] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. On the cryptographic applications of random functions. In G. R. Blakley and David Chaum, editors, *CRYPTO’84*, volume 196 of *LNCS*, pages 276–288. Springer, Heidelberg, August 1984.
- [GRT18] Sumegha Garg, Ran Raz, and Avishay Tal. Extractor-based time-space lower bounds for learning. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *50th ACM STOC*, pages 990–1002. ACM Press, June 2018.
- [GZ19] Jiaxin Guan and Mark Zhandary. Simple schemes in the bounded storage model. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 500–524. Springer, Heidelberg, May 2019.
- [HB01] Nicholas J. Hopper and Manuel Blum. Secure human identification protocols. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 52–66. Springer, Heidelberg, December 2001.

- [HCR02] Dowon Hong, Ku-Young Chang, and Heuisu Ryu. Efficient oblivious transfer in the bounded-storage model. In Yuliang Zheng, editor, *ASIACRYPT 2002*, volume 2501 of *LNCS*, pages 143–159. Springer, Heidelberg, December 2002.
- [Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- [ILL89] R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, STOC '89, page 12–24, New York, NY, USA, 1989. Association for Computing Machinery.
- [JW05] Ari Juels and Stephen A. Weis. Authenticating pervasive devices with human protocols. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 293–308. Springer, Heidelberg, August 2005.
- [KPC⁺11] Eike Kiltz, Krzysztof Pietrzak, David Cash, Abhishek Jain, and Daniele Venturi. Efficient authentication from hard learning problems. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 7–26. Springer, Heidelberg, May 2011.
- [KRT17] Gillat Kol, Ran Raz, and Avishay Tal. Time-space hardness of learning sparse parities. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *49th ACM STOC*, pages 1067–1080. ACM Press, June 2017.
- [Lu02] Chi-Jen Lu. Hyper-encryption against space-bounded adversaries from on-line strong extractors. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 257–271. Springer, Heidelberg, August 2002.
- [Mau92] Ueli M. Maurer. Conditionally-perfect secrecy and a provably-secure randomized cipher. *Journal of Cryptology*, 5(1):53–66, January 1992.
- [NZ96] Noam Nisan and David Zuckerman. Randomness is linear in space. *J. Comput. Syst. Sci.*, 52(1):43–52, 1996.
- [Raz16] Ran Raz. Fast learning requires good memory: A time-space lower bound for parity learning. In Irit Dinur, editor, *57th FOCS*, pages 266–275. IEEE Computer Society Press, October 2016.
- [Raz17] Ran Raz. A time-space lower bound for a large class of learning problems. In Chris Umans, editor, *58th FOCS*, pages 732–742. IEEE Computer Society Press, October 2017.
- [Vad04] Salil P. Vadhan. Constructing locally computable extractors and cryptosystems in the bounded-storage model. *Journal of Cryptology*, 17(1):43–77, January 2004.