

Multi-key homomorphic authenticators

 ISSN 1751-8709
 Received on 4th July 2018
 Revised 26th February 2019
 Accepted on 16th April 2019
 E-First on 26th June 2019
 doi: 10.1049/iet-ifs.2018.5341
 www.ietdl.org

 Dario Fiore¹, Aikaterini Mitrokotsa², Luca Nizzardo^{1,3} ✉, Elena Pagnin²
¹IMDEA Software Institute, Madrid, Spain

²Chalmers University of Technology, Gothenburg, Sweden

³Universidad Politécnica de Madrid, Spain

✉ E-mail: luca.nizzardo@imdea.org

Abstract: Homomorphic authenticators (HAs) enable a client to authenticate a large collection of data elements m_1, \dots, m_t and outsource them, along with the corresponding authenticators, to an untrusted server. At any later point, the server can generate a short authenticator $\sigma_{f,y}$ vouching for the correctness of the output y of a function f computed on the outsourced data, i.e. $y = f(m_1, \dots, m_t)$. The notion of HAs studied so far, however, only supports executions of computations over data authenticated by a single user. Motivated by realistic scenarios in which large datasets include data provided by multiple users, we study the concept of multi-key homomorphic authenticators. In a nutshell, multi-key HAs are like HAs with the extra feature of allowing the holder of public evaluation keys to compute on data authenticated under different secret keys. In this paper, we introduce and formally define multi-key HAs. Secondly, we propose a construction of a multi-key homomorphic signature based on standard lattices and supporting the evaluation of circuits of bounded polynomial depth. Thirdly, we provide a construction of multi-key homomorphic MACs based only on pseudorandom functions and supporting the evaluation of low-degree arithmetic circuits.

1 Introduction

The technological innovations offered by modern IT systems are changing the way digital data is collected, stored, processed, and consumed. As an example, think of an application where data is collected by some organisations (e.g. hospitals), stored and processed on remote servers (e.g. the Cloud) and finally consumed by other users (e.g. medical researchers) on other devices. On the one hand, this computing paradigm is very attractive, particularly as data can be shared and exchanged by multiple users. On the other hand, it is evident that in such scenarios, one may be concerned about security: while the users that collect and consume the data may trust each other (up to some extent), trusting the Cloud can be problematic for various reasons. More specifically, two main security concerns to be addressed are those about the *privacy* and *authenticity* of the data stored and processed in untrusted environments.

While it is widely known that privacy can be solved in such a setting using, e.g. homomorphic encryption [1], in this work, we focus on the orthogonal problem of providing authenticity of data during computation. Towards this goal, our contribution is on advancing the study of *homomorphic authenticators* (HAs), a cryptographic primitive that has been the subject of recent work [2–5].

HAs. Using an HA scheme, a user Alice can authenticate a collection of data items m_1, \dots, m_t using her secret key, and send the authenticated data to an untrusted server. The server can execute a program \mathcal{P} on the authenticated data and use a public evaluation key to generate a value $\sigma_{\mathcal{P},y}$ vouching for the correctness of $y = \mathcal{P}(m_1, \dots, m_t)$. Finally, a user Bob who is given the tuple $(\mathcal{P}, y, \sigma_{\mathcal{P},y})$ and Alice's verification key can use the authenticator to verify the authenticity of y as the output of the program \mathcal{P} executed on data authenticated by Alice. In other words, Bob can check that the server did not tamper with the computation's result. Alice's verification key can be either secret or public. In the former case, we refer to the primitive as *homomorphic message authentication codes (MACs)* [3, 6], while in the latter, we refer to it as *homomorphic signatures* [2]. One of the attractive features of HAs is that the authenticator $\sigma_{\mathcal{P},y}$ is *succinct*, i.e. much shorter than

\mathcal{P} 's input size. This means that the server can execute a program on a huge amount of data and convince Bob of its correctness by sending him only a short piece of information. As discussed in previous work (e.g. [3, 4, 7]), HAs provide a nice solution, with minimal communication and interaction, to the problem of delegating computations on outsourced data, and thus can be preferable to verifiable computation (VC) (more details on this comparison appear in Section 1.2).

Our contribution: multi-key HAs. Up to now, the notion of HAs has inherently been single-key, i.e. homomorphic computations are allowed only on data authenticated using the same secret key. This characteristic is obviously a limitation and prevents HA schemes from suiting scenarios where the data is provided (and authenticated) by multiple users. Consider the previously mentioned example of healthcare institutions which need to compute on data collected by several hospitals or even some remote-monitored patients. Similarly, it is often required to compute statistics for time-series data collected from multiple users, e.g. to monitor usage data in smart metering, clinical research or to monitor the safety of buildings. Another application scenario is in distributed networks of sensors. Imagine for instance a network of sensors where each sensor is in charge of collecting data about air pollution in a certain area of a city, it sends its data to a Cloud server, and then a central control unit asks the Cloud to compute on the data collected by the sensors (e.g. to obtain the average value of air pollution in a large area).

A trivial solution to address the problem of computing on data authenticated by multiple users is to use HAs in such a way that all data providers share the *same* secret authentication key. The desired functionality is obviously achieved since data would be authenticated using a single secret key. This approach, however, has several drawbacks. The first one is that users need to coordinate to agree on such a key. The second one is that in such a setting there would be no technical/legal way to differentiate between users (e.g. to make each user accountable for his/her duties) as any user can impersonate all the other ones. The third and more relevant reason is that sharing the same key exposes the overall system to way higher risks of attacks and makes disaster recovery more difficult: if a single user is compromised the whole

system is compromised too, and everything has to be reinitialised from scratch.

In contrast, this paper provides an innovative solution through the notion of *multi-key HAs*. This primitive guarantees that the corruption of one user affects the data of that user only, but does not endanger the authenticity of computations among the other (un-corrupted) users of the system. Moreover, the proposed system is dynamic, in the sense that compromised users can be assigned new keys and be easily reintegrated.

1.1 Overview of our results

Our contribution is mainly threefold. First, we elaborate a suitable definition of multi-key HAs. Second, we propose the first construction of a multi-key homomorphic signature (i.e. with public verifiability) which is based on standard lattices and supports the evaluation of circuits of bounded polynomial depth. Third, we present a multi-key homomorphic MAC that is based only on pseudorandom functions and supports the evaluation of low-degree arithmetic circuits. In spite of being less expressive and only secretly verifiable, this last construction is way more efficient than the signature scheme. In what follows, we elaborate more on our results.

Multi-key HAs: what are they? At a high level, multi-key HAs are like HAs with the additional property that one can execute a program \mathcal{P} on data authenticated using different secret keys. In multi-key HAs, Bob verifies using the verification keys of all users that provided inputs to \mathcal{P} . These features make multi-key HAs a perfect candidate for applications where multiple users gather and outsource data. Referring to our previous examples, using multi-key HAs each sensor can authenticate and outsource to the Cloud the data it collects; the Cloud can compute statistics on the authenticated data and provide the central control unit with the result along with a certificate vouching for its correctness.

An important aspect of our definition is a mechanism that allows the verifier to keep track of the users that authenticated the inputs of \mathcal{P} , i.e. to know which user contributed to each input wire of \mathcal{P} . To formalise this mechanism, we build on the model of *labelled data and programs* of Gennaro and Wichs [3] (see Section 3 for details). In terms of security, multi-key HAs allowed the adversary to corrupt users (i.e. to learn their secret keys); yet this knowledge should not help the adversary in tampering with the results of programs which involve inputs of honest (i.e. uncorrupted) users only. Our model allows us to handle compromised users in a similar way to what occurs with classical digital signatures: a compromised user could be banned by means of a certificate revocation, and could easily be re-integrated via a new key pair (Here we mean that this process does not add more complications than the ones already existing for classical digital signatures (e.g. relying on PKI mechanisms)). Thinking of the sensor network application, if a sensor in the field gets compromised, the data provided by other sensors remains secure, and a new sensor can be easily introduced in the system with new credentials.

Finally, we require multi-key HAs to be *succinct* in the sense that the size of authenticators is bounded by some fixed polynomial in $(\lambda, n, \log t)$, where λ is the security parameter, n is the number of users contributing to the computation and t is the total number of inputs of \mathcal{P} . Although such dependence on n may look undesirable, we stress that it is still meaningful in many application scenarios where n is much smaller than t . For instance, in the application scenario of healthcare institutions, a few hospitals can provide a large amount of data from patients.

A multi-key homomorphic signature for all circuits. After setting the definition of multi-key HAs, we proceed to construct multi-key HA schemes. Our first contribution is a multi-key homomorphic signature that supports the evaluation of Boolean circuits of depth bounded by a fixed polynomial in the security parameter. The scheme is proven secure based on the small integer solution (SIS) problem over standard lattices [8], and tolerates adversaries that corrupt users non-adaptively (Precisely, our ‘core’ scheme is secure against adversaries that make non-adaptive signing queries; this is upgraded to adaptive security via general

transformations). Our technique is inspired by the ones developed by Gorbunov *et al.* [4] to construct a (single-key) homomorphic signature. Our key contribution is on providing a new representation of the signatures that enables to homomorphically compute over them even if they were generated using different keys. Furthermore, our scheme enjoys an additional property, not fully satisfied by Gorbunov *et al.* [4]: every user can authenticate separately every data item m_i of a collection m_1, \dots, m_t , and the correctness of computations is guaranteed even when computing on not-yet-full datasets. Although it is possible to modify the scheme in [4] for signing data items separately, the security would only work against adversaries that query the whole dataset. In contrast, we prove our scheme to be secure under a stronger security definition where the adversary can *adaptively* query the various data items, and it can try to cheat by pretending to possess signatures on data items that it never queried (the so-called Type 3 forgeries). We highlight that the scheme in [4] is *not* secure under the stronger definition (with Type 3 forgeries) used in this paper, and we had to introduce new techniques to deal with this scenario. This new property is particularly interesting as it enables users to authenticate and outsource data items in a streaming fashion, without ever having to store the whole dataset. This is useful in applications where the dataset size can be very large or not fixed a priori.

A multi-key homomorphic MAC for low-degree circuits. Our second construction is a multi-key homomorphic MAC that supports the evaluation of arithmetic circuits whose degree d is at most polynomial in the security parameter, and whose inputs come from a small number n of users. For results of such computations, the corresponding authenticators have at the most size $s = \binom{n+d}{d}$

(Note that s can be bounded by $\text{poly}(n)$ for constant d , or by $\text{poly}(d)$ for constant n). Notably, the authenticator's size is completely independent of the total number of inputs of the arithmetic circuit. Compared to our multi-key homomorphic signature, this construction is only secretly verifiable (i.e. Bob has to know the secret verification keys of all users involved in the computation) and supports a class of computations that is less expressive; also its succinctness is asymptotically worse. In spite of these drawbacks, our multi-key homomorphic MAC achieves interesting features. From the theoretical point of view, it is based on very simple cryptographic tools: a family of pseudorandom functions. Thus, security relies only on one-way functions. On the practical side, it is particularly efficient: generating a MAC requires only one pseudo-random function evaluation and a couple of field operations; homomorphic computations boil down to additions and multiplications over a multi-variate polynomial ring $\mathbb{F}_p[X_1, \dots, X_n]$.

1.2 Related work

Homomorphic MACs and signatures. HAs have received a lot of attention in previous work focusing either on homomorphic signatures (publicly verifiable) or on homomorphic MACs (private verification with a secret key). The notion of homomorphic signatures was originally proposed by Johnson *et al.* [5]. The first schemes that appeared in the literature were homomorphic only for linear functions [9–13] and found important applications in network coding and proofs of retrievability. Boneh and Freeman [2] were the first to construct homomorphic signatures that can evaluate more than linear functions over signed data. Their scheme could evaluate bounded-degree polynomials, and its security was based on the hardness of the SIS problem in ideal lattices in the random oracle model. A few years later, Catalano *et al.* [14] proposed an alternative homomorphic signature scheme for bounded-degree polynomials. Their solution is based on multi-linear maps and bypasses the need for random oracles. More interestingly, the work by Catalano *et al.* [14] contains the first mechanism to verify signatures faster than the running time of the verified function. Recently, Gorbunov *et al.* [4] have proposed the first (levelled) fully homomorphic signature scheme that can evaluate arbitrary circuits of bounded polynomial depth over signed data. Some important advances have been also achieved in

the area of homomorphic MACs. Gennaro and Wichs [3] have proposed a fully homomorphic MAC based on fully homomorphic encryption. However, their scheme is not secure in the presence of verification queries. More efficient schemes have been proposed later [6, 7, 15] that are secure in the presence of verification queries and are more efficient at the price of supporting only restricted homomorphisms. Finally, we note that Agrawal *et al.* [16] considered a notion of *multi-key* signatures for network coding, and proposed a solution which works for linear functions only. Compared to this work, our contribution shows a full-fledged framework for multi-key HAs, and provides solutions that address a more expressive class of computations.

Verifiable computation (VC). Achieving correctness of outsourced computations is also the aim of the *verifiable delegation of computation* (VC) [17–22]. In this setting, a client wants to delegate the computation of a function f on an input x to an untrusted cloud-server. If the server replies with y , the client's goal is to verify the correctness of $y = f(x)$ spending less resources than those needed to execute f . As mentioned in previous work (e.g. [3, 4]) a crucial difference between VC and HAs is that in VC the verifier has to know the input of the computation – which can be huge – whereas in HAs one can verify by only knowing the function f and the result y . Moreover, although some results of VC could be re-interpreted to solve scenarios similar to the ones addressed by HAs, results based on VC would still present several limitations. For instance, using HAs the server can prove the correctness of $y = f(x)$ with a single message, without needing any special encoding of f from the delegator. Second, HAs come naturally with a composition property, which means that the outputs of some computations on authenticated data (which is already authenticated) can be fed as input for follow-up computations. This feature is of particular interest to parallelise and or distribute computations (e.g. MapReduce). Emulating this composition within VC systems is possible using certain non-interactive proof systems [23] but leads to complex statements and less natural realisations. The last advantage is that by using HAs, clients can authenticate various (small) pieces of data independently and without storing previously outsourced data. In contrast, most VC systems require clients to encode the whole input in ‘one-shot’, and often such encoding can be used in a single computation only.

Multi-client VC. Another line of work, closely related to ours is that on *multi-client VC* [24, 25]. This primitive, introduced by Choi *et al.* [24], aims to extend VC to the setting where inputs are provided by multiple users, and one of these users wants to verify the result's correctness. Choi *et al.* [24] proposed a definition and a multi-client VC scheme which generalises that of Gennaro *et al.* [20]. The solution in [24], however, does not consider malicious or colluding clients. This setting was addressed by Gordon *et al.* [25], where they provide a scheme with stronger security guarantees against a malicious server or an arbitrary set of malicious colluding clients.

It is interesting to notice that in the definition of multi-client VC, all the clients but the one who verifies can encode inputs independently of the function to be later executed on them. One may thus think that the special case in which the verifier provides no input yields a solution similar to the one achieved by multi-key HAs. However, a closer look at the definitions and the existing constructions of multi-client VC reveals three main differences. (i) In multi-client VC, to prove the correctness of an execution of a function f , the server has to wait for a message from the verifier which includes some encoding of f . This is not necessary for multi-key HAs where the server can directly prove the correctness of f on previously authenticated data with a single message and without any function's encoding. (ii) The communication between the server and the verifier is at least linear in the total number of inputs of f : this can be prohibitive in the case of computations on very large inputs (think of TBytes of data). In contrast, with multi-key HAs the communication between the server and the verifier is proportional only to the number of users and depends only logarithmically on the total number of inputs. (iii) In multi-client VC, an encoding of one input can be used in a single computation. Thus, if a user wants to first upload data on the server to later

execute many functions on it, then the user has to provide as many encodings as the number of functions to be executed. In contrast, multi-key HAs allowed one to encode (i.e. authenticate) every input only once and to use it for proving the correctness of computations an *unbounded* number of times.

1.3 Publication note

This paper includes all the results presented in [26] with additional details and complete proofs. In particular, compared to [26], in this version we additionally provide: the description of two generic transformations from weak to adaptive security (Section 4.4), a more detailed description of polynomial representations that could be employed in our MKHMac scheme (Section 5.1), and complete proofs for all the stated results (especially in Section 5).

2 Preliminaries

We collect here the notation and basic definitions used throughout the paper.

Notation. The Greek letter λ is reserved for the security parameter of the schemes. A function $\epsilon(\lambda)$ is said to be *negligible* in λ (denoted as $\epsilon(\lambda) = \text{negl}(\lambda)$) if $\epsilon(\lambda) = O(\lambda^{-c})$ for every constant $c > 0$. When a function can be expressed as a polynomial, we often write it as $\text{poly}(\cdot)$. For any $n \in \mathbb{N}$, we refer to $[n]$ as $[n] := \{1, \dots, n\}$. Moreover, given a set \mathcal{S} , the notation $s \xleftarrow{\$} \mathcal{S}$ stays for the process of sampling s uniformly at random from \mathcal{S} .

Definition 1 (Statistical distance): Let X, Y denote two random variables with support \mathcal{X}, \mathcal{Y} , respectively; the statistical distance between X and Y is defined as

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

If $\mathbf{SD}(X, Y) = \text{negl}(\lambda)$, we say that X and Y are statistically close and we write $X \stackrel{\text{stat}}{\approx} Y$. Below we recall the definition and basic facts about conditional min-entropy [27].

Definition 2 (Entropy): The min-entropy of a random variable X is defined as

$$\mathbf{H}_{\infty}(X) := -\log\left(\max_x \Pr[X = x]\right).$$

The (average-) conditional min-entropy of a random variable X conditioned on a correlated variable Y is defined as

$$\mathbf{H}_{\infty}(X|Y) := -\log\left(\mathbf{E}_{y \leftarrow Y} \left[\max_x \Pr[X = x | Y = y] \right]\right).$$

The optimal probability of an unbounded adversary guessing X when given a correlated value Y is $2^{-\mathbf{H}_{\infty}(X|Y)}$.

Lemma 1: Let X, Y be arbitrarily random variables where the support of Y lies in \mathcal{Y} . Then $\mathbf{H}_{\infty}(X|Y) \geq \mathbf{H}_{\infty}(X) - \log(|\mathcal{Y}|)$.

3 Multi-key HAs

In this section, we present our new notion of *Multi-Key HAs*. Intuitively, multi-key HAs extended the existing notions of homomorphic signatures [2] and homomorphic MACs [3] in such a way that one can homomorphically compute a program \mathcal{P} over data authenticated using different secret keys. For the sake of verification, in multi-key HAs the verifier needs to know the verification keys of all users that provided inputs to \mathcal{P} . Our definitions are meant to be general enough to be easily adapted to both the case in which verification keys are public and the one where verification keys are secret. In the former case, we call the primitive *multi-key homomorphic signatures*, whereas in the latter case we call it *multi-key homomorphic MACs*.

As already observed in previous work about HAs, it is important that an authenticator $\sigma_{\mathcal{P},y}$ does not authenticate a value y out of context, but only as of the output of a program \mathcal{P} executed on previously authenticated data. To formalise this notion, we build on the model of *labelled data and programs* of Gennaro and Wichs [3]. The idea of this model is that every data item is authenticated under a unique label ℓ . For example, in scenarios where the data is outsourced, such labels can be thought of as a way to index/identify the remotely stored data. A labelled program \mathcal{P} , on the other hand, consists of a circuit f where every input wire i has a label ℓ_i . Going back to the outsourcing example, a labelled program is a way to specify what portion of the outsourced data, one should execute a circuit f . More formally, the notion of labelled programs of [3] is recalled below.

Labelled programs [3]. A labelled program \mathcal{P} is a tuple $(f, \ell_1, \dots, \ell_n)$, such that $f: \mathcal{M}^n \rightarrow \mathcal{M}$ is a function of n variables (e.g. a circuit) and $\ell_i \in \{0, 1\}^*$ is a label for the i th input of f . Labelled programs can be composed as follows: given $\mathcal{P}_1, \dots, \mathcal{P}_t$ and a function $g: \mathcal{M}^t \rightarrow \mathcal{M}$, the composed program \mathcal{P}^* is the one obtained by evaluating g on the outputs of $\mathcal{P}_1, \dots, \mathcal{P}_t$, and it is denoted as $\mathcal{P}^* = g(\mathcal{P}_1, \dots, \mathcal{P}_t)$. The labelled inputs of \mathcal{P}^* are all the distinct labelled inputs of $\mathcal{P}_1, \dots, \mathcal{P}_t$ (all the inputs with the same label are grouped together and considered as a unique input of \mathcal{P}^*). Let $f_{id}: \mathcal{M} \rightarrow \mathcal{M}$ be the identity function and $\ell \in \{0, 1\}^*$ be any label. We refer to $\mathcal{I}_\ell = (f_{id}, \ell)$ as the identity program with the label ℓ . Note that a program $\mathcal{P} = (f, \ell_1, \dots, \ell_n)$ can be expressed as the composition of n identity programs $\mathcal{P} = f(\mathcal{I}_{\ell_1}, \dots, \mathcal{I}_{\ell_n})$.

We are using labelled programs to identify users. In our notion of multi-key HAs, one wishes to verify the outputs of computations executed on data authenticated under *different* keys. A meaningful definition of multi-key HAs thus, requires that the authenticators are not out of context also concerning the set of keys that contributed to the computation. To address this issue, we assume that every user has an identity id in some identity space ID , and that the user's keys are associated to id using any suitable mechanism (e.g. PKI). Next, to distinguish among data of different users and to identify to which input wires a certain user-contributed, we assume that the label space contains the set ID . Namely, in our model, a data item is assigned a label $\ell := (id, \tau)$, where id is a user's identity, and τ is a tag; this essentially identifies uniquely a data item of user id with index τ . For compatibility with previous notions of HAs, we assume that data items can be grouped in datasets, and one can compute only on data within the same dataset. In our definitions, a dataset is identified by an arbitrary string Δ (Although considering the dataset notion complicates the definition, it also provides some benefits, as we illustrate later in the constructions. For instance, when verifying for the same program \mathcal{P} over different datasets, one can perform some pre-computation that makes further verifications cheap).

Definition 3 (Multi-key HA): A multi-key HA scheme MKHAut consists of a tuple of PPT algorithms (Setup , KeyGen , Auth , Eval , Ver) satisfying the following properties: *authentication correctness*, *evaluation correctness*, *succinctness*, and *security*. The five algorithms work as follows:

$\text{Setup}(1^\lambda)$. The setup algorithm takes as input the security parameter λ and outputs some public parameters pp . These parameters include (at least) descriptions of a tag space \mathcal{T} , an identity space ID , the message space \mathcal{M} , and a set of admissible functions \mathcal{F} . Given \mathcal{T} and ID , the label space of the scheme is defined as their Cartesian product $\mathcal{L} := ID \times \mathcal{T}$. For a labelled program $\mathcal{P} = (f, \ell_1, \dots, \ell_t)$ with labels $\ell_i := (id_i, \tau_i) \in \mathcal{L}$, we use $id \in \mathcal{P}$ as a compact notation for $id \in \{id_1, \dots, id_t\}$. The pp are input to all the following algorithms, even when not specified.

$\text{KeyGen}(pp)$. The key generation algorithm takes as input the public parameters and outputs a triple of keys (sk, ek, vk) , where sk is a secret authentication key, ek is a public evaluation key and

vk is a verification key which could be either public or private (As mentioned earlier, the generated triple (sk, ek, vk) will be associated with an identity $id \in ID$. When this connection becomes explicit, we will refer to (sk, ek, vk) as $(sk_{id}, ek_{id}, vk_{id})$. $\text{Auth}(sk, \Delta, \ell, m)$. The authentication algorithm takes as input an authentication key sk , a dataset identifier Δ , a label $\ell = (id, \tau)$ for the message m , and it outputs an authenticator σ .

$\text{Eval}(f, \{(\sigma_i, \text{EKS}_i)\}_{i \in [t]})$. The evaluation algorithm takes as input a t -input function $f: \mathcal{M}^t \rightarrow \mathcal{M}$, and a set $\{(\sigma_i, \text{EKS}_i)\}_{i \in [t]}$ where each σ_i is an authenticator and each EKS_i is a set of evaluation keys (The motivation behind the evaluation-keys set EKS_i is that, if σ_i authenticates the output of a labelled program \mathcal{P}_i , then $\text{EKS}_i = \{ek_{id}\}_{id \in \mathcal{P}_i}$ should be the set of evaluation keys corresponding to identities involved in the computation of \mathcal{P}_i).

$\text{Ver}(\mathcal{P}, \Delta, \{vk_{id}\}_{id \in \mathcal{P}}, m, \sigma)$. The verification algorithm takes as input a labelled program $\mathcal{P} = (f, \ell_1, \dots, \ell_t)$, a dataset identifier Δ , the set of verification keys $\{vk_{id}\}_{id \in \mathcal{P}}$ corresponding to those identities id involved in the program \mathcal{P} , a message m , and an authenticator σ . It outputs 0 (reject) or 1 (accept).

Authentication correctness. Intuitively, a multi-key homomorphic authenticator has authentication correctness if the output of $\text{Auth}(sk, \Delta, \ell, m)$ verifies correctly for m as the output of the identity program \mathcal{I}_ℓ over the dataset Δ . More formally, a scheme MKHAut satisfies authentication correctness if, for all public parameters $pp \leftarrow \text{Setup}(1^\lambda)$, any key triple $(sk, ek, vk) \leftarrow \text{KeyGen}(pp)$, any label $\ell = (id, \tau) \in \mathcal{L}$, and any authenticator $\sigma \leftarrow \text{Auth}(sk, \Delta, \ell, m)$, we have $\text{Ver}(\mathcal{I}_\ell, \Delta, vk, m, \sigma)$ outputs 1 with all but a negligible probability.

Evaluation correctness. Intuitively, this property says that running the evaluation algorithm on signatures $(\sigma_1, \dots, \sigma_t)$ such that each σ_i verifies for m_i as the output of a labelled program \mathcal{P}_i over the dataset Δ , it produces a signature σ which verifies for $f(m_1, \dots, m_t)$ as the output of the composed program $f(\mathcal{P}_1, \dots, \mathcal{P}_t)$ over the dataset Δ . More formally, let us fix the public parameters $pp \leftarrow \text{Setup}(1^\lambda)$, a set of key triples $\{(sk_{id}, ek_{id}, vk_{id})\}_{id \in ID}$ for some $ID \subseteq ID$, a dataset Δ , a function $g: \mathcal{M}^t \rightarrow \mathcal{M}$, and any set of program/message/authenticator triples $\{(\mathcal{P}_i, m_i, \sigma_i)\}_{i \in [t]}$ such that $\text{Ver}(\mathcal{P}_i, \Delta, \{vk_{id}\}_{id \in \mathcal{P}_i}, m_i, \sigma_i) = 1$ for all $i \in [t]$. Let $m^* = g(m_1, \dots, m_t)$, $\mathcal{P}^* = g(\mathcal{P}_1, \dots, \mathcal{P}_t)$, and

$$\sigma^* = \text{Eval}(g, \{(\sigma_i, \text{EKS}_i)\}_{i \in [t]}), \quad (1)$$

where $\text{EKS}_i = \{ek_{id}\}_{id \in \mathcal{P}_i}$. Then, $\text{Ver}(\mathcal{P}^*, \Delta, \{vk_{id}\}_{id \in \mathcal{P}^*}, m^*, \sigma^*) = 1$ holds with all but a negligible probability.

Succinctness. A multi-key HA is said to be *succinct* if the size of every authenticator depends only logarithmically on the size of a dataset. However, we allow authenticators to depend on the number of keys involved in the computation. More formally, let $pp \leftarrow \text{Setup}(1^\lambda)$, $\mathcal{P} = (f, \ell_1, \dots, \ell_t)$ with $\ell_i = (id_i, \tau_i)$,

$$\{(sk_{id}, ek_{id}, vk_{id}) \leftarrow \text{KeyGen}(pp)\}_{id \in \mathcal{P}}, \quad (2)$$

And

$$\sigma_i \leftarrow \text{Auth}(sk_{id_i}, \Delta, \ell_i, m_i) \quad (3)$$

for all $i \in [t]$. A multi-key HA is said to be *succinct* if there exists a fixed polynomial p such that $|\sigma| = p(\lambda, n, \log t)$, where $\sigma = \text{Eval}(g, \{(\sigma_i, ek_{id_i})\}_{i \in [t]})$ and $n = |\{id \in \mathcal{P}\}|$.

Remark 1: Succinctness is one of the crucial properties that make multi-key HAs an interesting primitive. Without succinctness, a trivial multi-key HA construction is the one where Eval outputs the concatenation of all the signatures (and messages) given in input, and the verifier simply checks each

message–signature pair and recomputes the function by himself. Our definition of succinctness, where signatures can grow linearly with the number of keys but logarithmically in the total number of inputs, is also non-trivial, especially when considering settings where there are many more inputs than keys (in which case, the above trivial construction would not work). Another property that can make homomorphic signatures an interesting primitive is privacy – context-hiding – as considered in prior work. Intuitively, context-hiding guarantees that signatures do not reveal information on the original inputs. While we leave the study of context-hiding for multi-key HAs for future work, we note that a trivial construction that is context-hiding but not succinct can be easily obtained with the additional help of non-interactive zero-knowledge proofs of knowledge: the idea is to extend the trivial construction above by requiring the evaluator to generate a NIZK proof of knowledge of valid input messages and signatures that yield the public output. In this sense, we believe that succinctness is the most non-trivial property to achieve in homomorphic signatures, and this is what we focus on in this work.

Security. Intuitively, our security model for multi-key HAs guarantees that an adversary, without knowledge of the secret keys, can only produce authenticators that were either received from a legitimate user or verify correctly on the results of computations executed on the data authenticated by legitimate users. Moreover, we also give the adversary the possibility of corrupting users. In this case, it must not be able to cheat on the outputs of programs that get inputs from uncorrupted users only. In other words, our security definition guarantees that the corruption of one user affects the data of that user only, but does not endanger the integrity of computations among the other (un-corrupted) users of the system. We point out that preventing cheating on programs that involve inputs of corrupted users is inherently impossible in multi-key HAs, at least if general functions are considered. For instance, consider an adversary who picks the function $(x_1 + x_2 \bmod p)$, where x_1 is supposed to be provided by user Alice. If the adversary corrupts Alice, it can use her secret key to inject any input authenticated on her behalf and thus bias the output of the function at its will.

The formalisation of the intuitions illustrated above is more involved. For a scheme, MKHAut we define security via the following experiment between a challenger \mathcal{C} and an adversary \mathcal{A} ($\text{HomUF} - \text{CMA}_{\mathcal{A}, \text{MKHAut}}(\lambda)$):

Setup. \mathcal{C} runs $\text{Setup}(1^\lambda)$ to obtain the public parameters pp that are sent to \mathcal{A} .

Authentication queries. \mathcal{A} can adaptively submit queries of the form (Δ, ℓ, m) , where Δ is a dataset identifier, $\ell = (\text{id}, \tau)$ is a label in $\text{ID} \times \mathcal{T}$ and $m \in \mathcal{M}$ are messages of his choice. \mathcal{C} answers as follows:

- If (Δ, ℓ, m) is the first query for the dataset Δ , \mathcal{C} initialises an empty list $L_\Delta = \emptyset$ and proceeds as follows.
- If (Δ, ℓ, m) is the first query with identity id , \mathcal{C} generates keys $(\text{sk}_{\text{id}}, \text{ek}_{\text{id}}, \text{vk}_{\text{id}}) \xleftarrow{\$} \text{KeyGen}(\text{pp})$ (that are implicitly assigned to identity id), gives ek_{id} to \mathcal{A} and proceeds as follows.
- If (Δ, ℓ, m) is such that $(\ell, m) \notin L_\Delta$, \mathcal{C} computes $\sigma_\ell \xleftarrow{\$} \text{Auth}(\text{sk}_{\text{id}}, \Delta, \ell, m)$ (note that \mathcal{C} has already generated keys for the identity id), returns σ_ℓ to \mathcal{A} , and updates the list $L_\Delta \leftarrow L_\Delta \cup (\ell, m)$.
- If (Δ, ℓ, m) is such that $(\ell, \cdot) \in L_\Delta$ (which means that the adversary had already made a query (Δ, ℓ, m') for some message m'), then \mathcal{C} ignores the query.

Verification queries. \mathcal{A} is also given access to a verification oracle. Namely, the adversary can submit a query $(\mathcal{P}, \Delta, m, \sigma)$, and \mathcal{C} replies with the output of $\text{Ver}(\mathcal{P}, \Delta, \{\text{vk}_{\text{id}}\}_{\text{id} \in \mathcal{P}}, m, \sigma)$.

Corruption. The adversary has access to a corruption oracle. At the beginning of the game, the challenger initialises an empty list $L_{\text{corr}} = \emptyset$ of corrupted identities; during the game, \mathcal{A} can adaptively query identities $\text{id} \in \text{ID}$. If $\text{id} \notin L_{\text{corr}}$, then \mathcal{C} replies with the triple $(\text{sk}_{\text{id}}, \text{ek}_{\text{id}}, \text{vk}_{\text{id}})$ (i.e. generated using KeyGen if not

done before) and updates the list $L_{\text{corr}} \leftarrow L_{\text{corr}} \cup \text{id}$. If $\text{id} \in L_{\text{corr}}$, then \mathcal{C} replies with the triple $(\text{sk}_{\text{id}}, \text{ek}_{\text{id}}, \text{vk}_{\text{id}})$ assigned to id before.

Forgery. In the end, \mathcal{A} outputs a tuple $(\mathcal{P}^*, \Delta^*, m^*, \sigma^*)$. The experiment outputs 1 if the tuple returned by \mathcal{A} is a forgery (defined below), and 0 otherwise.

Definition 4 (Forgery): Consider an execution of the experiment $\text{HomUF} - \text{CMA}_{\mathcal{A}, \text{MKHAut}}(\lambda)$, where $(\mathcal{P}^*, \Delta^*, m^*, \sigma^*)$ is the tuple returned by the adversary at the end of the experiment, with $\mathcal{P}^* = (f^*, \ell_1^*, \dots, \ell_r^*)$. This is a forgery if

$$\text{Ver}(\mathcal{P}^*, \Delta^*, \{\text{vk}_{\text{id}}\}_{\text{id} \in \mathcal{P}^*}, m^*, \sigma^*) = 1, \quad (4)$$

for all $\text{id} \in \mathcal{P}^*$, we have that $\text{id} \notin L_{\text{corr}}$ (i.e. no identity involved in \mathcal{P}^* is corrupted), and either one of the following properties is satisfied:

Type 1: L_{Δ^*} has not been initialised during the game (i.e. the dataset Δ^* was never queried).

Type 2: For all $i \in [r]$, $\exists (\ell_i^*, m_i) \in L_{\Delta^*}$, but $m^* \neq f^*(m_1, \dots, m_r)$ (i.e. m^* is not the correct output of \mathcal{P}^* when executed over previously authenticated messages).

Type 3: There exists a label ℓ^* such that $(\ell^*, \cdot) \notin L_{\Delta^*}$ (i.e. \mathcal{A} never made a query with label ℓ^*).

We say that a multi-key HA scheme MKHAut is *secure* if for every PPT adversary \mathcal{A} , its advantage

$$\text{Adv}_{\text{MKHAut}, \mathcal{A}}^{\text{HomUF} - \text{CMA}}(\lambda) = \Pr [\text{HomUF} - \text{CMA}_{\mathcal{A}, \text{MKHAut}}(\lambda) = 1] \quad (5)$$

is negligible.

Remark 2 (Comparison with previous security definitions): Our security notion can be seen as the multi-key version of the one proposed by Gennaro and Wichs in [3] (in their model our Type 3 forgeries are called ‘Type 1’ as they do not consider multiple datasets). We point out that even in the special case of a single key, our security definition is stronger than the ones used in previous work [2, 4, 13, 14] (with the only exception of [3]). The main difference lies in our definition of Type 3 forgeries. The intuitive idea of this kind of forgeries is that an adversary who did not receive an authenticated input labelled by a certain ℓ^* cannot produce a valid authenticator for the output of a computation which has ℓ^* among its inputs. In [2, 4] these forgeries were not considered at all, as the adversary is assumed to query the dataset always in full. Other works [6, 13, 14] consider a weaker Type 3 notion, which deals with the concept of ‘well defined programs’, where the input wire labelled by the missing label ℓ^* is required to ‘contribute’ to the computation (i.e. it must change its outcome). The issue with such a Type 3 definition is that it may not be efficient to test if an input contributes to a function, especially if the admissible functions are general circuits. In contrast, our definition above is simpler and efficiently testable since it simply considers a Type 3 forgery one where the labelled program \mathcal{P}^* involves an un-queried input.

Multi-key homomorphic signatures. As previously mentioned, our definitions are general enough to be easily adapted to either case in which verification is secret or public. The only difference is whether the adversary is allowed to see the verification keys in the security experiment. When the verification is public, we call the primitive *multi-key homomorphic signatures*. More formally:

Definition 5 (Multi-key homomorphic signatures): A multi-key homomorphic signature is a multi-key HA in which verification keys are also given to the adversary in the security experiment.

Note that making verification keys public also allows to slightly simplify the security experiment by removing the verification oracle (the adversary can run the verification by itself). In the sequel, when referring to multi-key homomorphic signatures, we

adapt our notation to the typical one of digital signatures. Namely we denote $\text{Auth}(\text{sk}, \Delta, \ell, m)$ as $\text{Sign}(\text{sk}, \Delta, \ell, m)$, and call its outputs *signatures*.

3.1 Non-adaptive corruption queries

In our work, we consider a relaxation of the security definition in which the adversaries ask for corruptions in a non-adaptive way. More precisely, we say that an adversary \mathcal{A} makes *non-adaptive corruption* queries if for every identity id asked to the corruption oracle, id was not queried earlier in the game to the authentication oracle or the verification oracle. For this class of adversaries, it is easy to see that corruption queries are essentially of no help as the adversary can generate keys on its own. More precisely, we can prove the following proposition.

Proposition 1: MKHAut is secure against adversaries that do not make corruption queries if and only if MKHAut is secure against adversaries that make non-adaptive corruption queries.

Proof: The proof is rather simple, and we give here only a proof sketch. It is obvious to see that if a scheme is secure in the presence of non-adaptive corruptions, then the same scheme is secure when there is no corruption at all. For the opposite direction, we can show that for every adversary $\mathcal{A}_{\text{CORR}}$ that breaks the security of MKHAut by non-adaptive corruptions, there is an adversary \mathcal{B} that breaks the security of MKHAut in a game without corruptions. In particular, the reduction is tight, i.e. the success probability of \mathcal{B} is the same as that of $\mathcal{A}_{\text{CORR}}$.

\mathcal{B} simply runs the adversary $\mathcal{A}_{\text{CORR}}$ and forwards all the queries of $\mathcal{A}_{\text{CORR}}$ (except corruption queries) to its challenger. For every corruption query id , \mathcal{B} generates a tuple $(\text{sk}_{\text{id}}, \text{ek}_{\text{id}}, \text{vk}_{\text{id}}) \xleftarrow{\$} \text{KeyGen}(\text{pp})$ and returns the tuple to $\mathcal{A}_{\text{CORR}}$. If $\mathcal{A}_{\text{CORR}}$ comes up with a valid forgery, by definition this forgery does not involve any corrupted identity. Hence, the same forgery will be valid also in the game played by \mathcal{B} . It is not hard to see that the simulation is perfect since $\mathcal{A}_{\text{CORR}}$ only makes corruption queries for identities that were never asked to the authentication or verification oracles earlier in the game. Therefore, if $\mathcal{A}_{\text{CORR}}$ had an advantage ϵ in breaking the scheme in the game with non-adaptive corruptions, then \mathcal{B} has the same advantage to break the scheme without making corruptions at all. \square

3.2 Weakly-adaptive secure multi-key HAs

In our work, we also consider a weaker notion of security for multi-key HAs in which the adversary has to declare all the queried messages at the beginning of the experiment. More precisely, we consider a notion in which the adversary declares only the messages and the respective tags that will be queried, for every dataset and identity, without, however, need to specify the names of the datasets or of the identities. In a sense, the adversary \mathcal{A} is adaptive on identities and dataset names, but not on tags and messages. The definition is inspired by the one, for the single-key setting, of Catalano *et al.* [14].

To define the notion of weakly-adaptive security for multi-key HAs, we introduce here a new experiment $\text{Weak} - \text{HomUF} - \text{CMA}_{\mathcal{A}, \text{MKHAut}}$, which is a variation of the experiment $\text{HomUF} - \text{CMA}_{\mathcal{A}, \text{MKHAut}}$ (Definition 3) as described below.

Definition 6 (Weakly-secure multi-key HAs): In the security experiment $\text{Weak} - \text{HomUF} - \text{CMA}_{\mathcal{A}, \text{MKHAut}}$, before the setup phase, the adversary \mathcal{A} sends to the challenger \mathcal{C} a collection of sets of tags $\mathcal{T}_{i,k} \subseteq \mathcal{T}$ for $i \in [Q_{\text{id}}]$ and $k \in [Q_{\Delta}]$, where Q_{id} and Q_{Δ} are, respectively, the total number of distinct identities and datasets that will be queried during the game. Associated to every set $\mathcal{T}_{i,k}$, \mathcal{A} also sends a set of messages $\{\text{m}_\tau\}_{\tau \in \mathcal{T}_{i,k}}$. Basically the adversary declares, before key generation, all the messages and tags that it will query later on; however, \mathcal{A} is not required to specify identity and dataset names. Next, the adversary receives the public parameters from \mathcal{C} and can start the query-phase. Verification

queries are handled as in $\text{HomUF} - \text{CMA}_{\mathcal{A}, \text{MKHAut}}$. For authentication queries, \mathcal{A} can adaptively submit pairs (id, Δ) to \mathcal{C} . The challenger then replies with a set of authenticators $\{\sigma_\tau\}_{\tau \in \mathcal{T}_{i,k}}$, where indices i, k are such that id is the i th queried identity, and Δ is the k th queried a dataset.

An analogous security definition of weakly-secure multi-key homomorphic signatures is trivially obtained by removing a verification oracle.

In Section 4.4, full version of this paper, we present two generic transformations that turn weakly secure multi-key HA schemes into adaptive secure ones. Our first transformation holds in the standard model and works for schemes in which the tag space \mathcal{T} has polynomial size, while the second one avoids this limitation on the size of \mathcal{T} but holds in the random oracle model.

4 Our multi-key fully homomorphic signature

In this section, we present our construction of a multi-key homomorphic signature scheme that supports the evaluation of arbitrary circuits of bounded polynomial depth. The scheme is based on the SIS problem on standard lattices, a background of which is provided in the next section. Precisely, in Section 4.2, we present a scheme that is weakly-secure and supports a single dataset. Later, in Section 4.3, we discuss how to extend the scheme to handle multiple datasets, whereas the support of adaptive security can be obtained via the applications of our transformations in Section 4.4.

4.1 Lattices and SIS problem

We recall here notation and some basic results about lattices that are useful to describe our homomorphic signature construction.

For any positive integer q , we denote by \mathbb{Z}_q the ring of integers modulo q . Elements in \mathbb{Z}_q are represented as integers in the range $(-\frac{q}{2}, \frac{q}{2}]$. The absolute value of any $x \in \mathbb{Z}_q$ (denoted with $|x|$) is defined by taking the modulus q representative of x in $(-\frac{q}{2}, \frac{q}{2}]$,

i.e. take $y = x \bmod q$ and then set $|x| = |y| \in [0, \frac{q}{2}]$. Vectors and

matrices are denoted in bold. For any vector $\mathbf{u} := (u_1, \dots, u_n) \in \mathbb{Z}_q^n$, its infinity norm is $\|\mathbf{u}\|_\infty := \max_{i \in [n]} |u_i|$, and similarly, for a matrix $\mathbf{A} := [a_{i,j}] \in \mathbb{Z}_q^{n \times m}$ we write $\|\mathbf{A}\|_\infty := \max_{i \in [n], j \in [m]} |a_{i,j}|$.

SIS problem. For integer parameters n, m, q and β , the $\text{SIS}(n, m, q, \beta)$ problem [8] provides to an adversary \mathcal{P} a uniformly random $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, and requires \mathcal{A} to find a vector $\mathbf{u} \in \mathbb{Z}_q^{n \times m}$ such that $\mathbf{u} \neq \mathbf{0}$, $\|\mathbf{u}\|_\infty \leq \beta$, and $\mathbf{A} \cdot \mathbf{u} = \mathbf{0}$. More formally:

Definition 7 (SIS): Let $\lambda \in \mathbb{N}$ be the security parameter. For values $n = n(\lambda)$, $m = m(\lambda)$, $q = q(\lambda)$, $\beta = \beta(\lambda)$, defined as functions of λ , the $\text{SIS}(n, m, q, \beta)$ hardness assumption holds if for any PPT adversary \mathcal{A} we have

$$\Pr \left[\begin{array}{l} \mathbf{A} \cdot \mathbf{u} = \mathbf{0} \wedge \mathbf{u} \neq \mathbf{0} \wedge \|\mathbf{u}\|_\infty \leq \beta: \\ \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}, \mathbf{u} \leftarrow \mathcal{A}(1^\lambda, \mathbf{A}) \end{array} \right] \leq \text{negl}(\lambda).$$

For standard lattices, the SIS problem is known to be as hard as solving certain worst-case instances of lattice problems [8, 28–30], and is also implied by the hardness of learning with error (we refer any interested reader to the cited papers for the technical details about the parameters).

In our paper, we assume that for any $\beta = 2^{\text{poly}(\lambda)}$ there are some $n = \text{poly}(\lambda)$, $q = 2^{\text{poly}(\lambda)}$, with $q > \beta$, such that for all $m = \text{poly}(\lambda)$ the $\text{SIS}(n, m, q, \beta)$ hardness assumption holds. This parameter choice assures that hardness of worst-case lattice problems holds with sub-exponential approximation factors.

Trapdoors for lattices. The SIS problem is hard to solve for a random matrix \mathbf{A} . However, there is a way to sample a random \mathbf{A} together with a *trapdoor* such that SIS becomes easy to solve for that \mathbf{A} , given the trapdoor. Additionally, it has been shown that there exist ‘special’ (non-random) matrices \mathbf{G} for which SIS is easy

to solve as well. The following lemma from [31–34] summarises the above-known results (similar to a lemma in [35]):

Lemma 2: There exist efficient algorithms TrapGen , SamPre , Sam such that the following holds: given integers $n \geq 1, q \geq 2$, there exist some

$$m^* = m^*(n, q) = O(n \log q), \beta_{\text{sam}} = \beta_{\text{sam}}(n, q) = O(n\sqrt{\log q}) \quad (6)$$

such that for all $m \geq m^*$ and all k (polynomial in n) we have

1. $\text{Sam}(1^m, 1^k, q) \rightarrow U$ samples a matrix $U \in \mathbb{Z}_q^{m \times k}$ such that $\|U\|_\infty \leq \beta_{\text{sam}}$ (with probability 1).
2. For $(A, \text{td}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$, $A' \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$, $U \leftarrow \text{Sam}(1^m, 1^k, q)$, $V := AU$, $V' \xleftarrow{\$} \mathbb{Z}_q^{n \times k}$, $U' \leftarrow \text{SamPre}(A, V', \text{td})$, we have the following statistical indistinguishability (negligible in n):

$$A \stackrel{\text{stat}}{\approx} A' \quad \text{and} \quad (A, \text{td}, U, V) \stackrel{\text{stat}}{\approx} (A, \text{td}, U', V')$$

and $U' \leftarrow \text{SamPre}(A, V', \text{td})$ always satisfies $AU' = V'$ and $\|U'\|_\infty \leq \beta_{\text{sam}}$.

3. Given n, m, q as above, there is an efficiently and deterministically computable matrix $G \in \mathbb{Z}_q^{n \times m}$ and a deterministic polynomial-time algorithm G^{-1} that on input $V \in \mathbb{Z}_q^{n \times k}$ (for any integer k) outputs $R = G^{-1}(V)$ such that $R \in \{0, 1\}^{m \times k}$ and $GR = V$.

4.2 Our multi-key homomorphic signature for a single dataset

In this section, we present our multi-key homomorphic signature that supports the evaluation of Boolean circuits of bounded polynomial depth. Our construction is inspired by the (single-key) one of Gorbunov *et al.* [4], with the fundamental difference that in our case, we enable computations over data signed using different secret keys. Moreover, our scheme is secure against Type 3 forgeries. We achieve this via a new technique which consists of adding to every signature a component that specifically protects against this type of forgeries. We prove the scheme to be weakly-secure under the SIS hardness assumption.

Parameters. Before describing the scheme, we discuss how to set the various parameters involved. Let λ be the security parameter, and let $d = d(\lambda) = \text{poly}(\lambda)$ be the bound on the depth of the circuits supported by our scheme. We define the set of parameters used in our scheme $\text{Par} = \{n, m, q, \beta_{\text{SIS}}, \beta_{\text{max}}, \beta_{\text{init}}\}$ in terms of λ, d and of the parameters required by the trapdoor algorithm in Lemma 2: m^*, β_{sam} , where $m^* = m^*(n, q) := O(n \log q)$ and $\beta_{\text{sam}} := O(n\sqrt{\log q})$. More precisely, we set: $\beta_{\text{max}} := 2^{\omega(\log \lambda) d}$, $\beta_{\text{SIS}} := 2^{\omega(\log \lambda) \beta_{\text{max}}}$; $n = \text{poly}(\lambda)$; $q = \mathcal{O}(2^{\text{poly}(\lambda)}) > \beta_{\text{SIS}}$ is a prime (as small as possible) so that the $\text{SIS}(n, m', q, \beta_{\text{SIS}})$ assumption holds for all $m' = \text{poly}(\lambda)$;

$$m = \max \{m^*, n \log q + \omega(\log(\lambda))\} = \text{poly}(\lambda) \quad (7)$$

and, finally, $\beta_{\text{init}} := \beta_{\text{sam}} = \text{poly}(\lambda)$.

Construction. The PPT algorithms (Setup, KeyGen, Sign, Eval, Ver) which define our construction of multi-key homomorphic signatures work as follows:

Setup(1^λ). The setup algorithm takes as input the security parameter λ and generates the public parameters pp which include: the bound on the circuit depth d (which defines the class \mathcal{F} of functions supported by the scheme, i.e. Boolean circuits of depth d), the set $\text{Par} = \{n, m, q, \beta_{\text{SIS}}, \beta_{\text{max}}, \beta_{\text{init}}\}$, the set $\mathcal{U} = \{U \in \mathbb{Z}_q^{m \times m} : \|U\|_\infty \leq \beta_{\text{max}}\}$, the set $\mathcal{V} = \{V \in \mathbb{Z}_q^{n \times m}\}$, descriptions of the message space $\mathcal{M} = \{0, 1\}$, the tag space $\mathcal{T} = [T]$, and the identity space $\text{ID} = [C]$, for integers $T, C \in \mathbb{N}$. In this construction, the tag space is of polynomial size, i.e.

$T = \text{poly}(\lambda)$ while the identity space is essentially unbounded, i.e. we set $C = 2^\lambda$. Also, recall that \mathcal{F} and ID immediately define the label space $\mathcal{L} = \text{ID} \times \mathcal{F}$. The final output is $\text{pp} = \{d, \text{Par}, \mathcal{U}, \mathcal{V}, \mathcal{M}, \mathcal{T}, \text{ID}\}$. We assume that these public parameters pp are an input of all subsequent algorithms, and often omit them from the input explicitly.

KeyGen(pp). The key generation algorithm takes as input the public parameters pp and generates a key-triple $(\text{sk}, \text{ek}, \text{vk})$ defined as follows. First, it samples T random matrices $V_1, \dots, V_T \xleftarrow{\$} \mathcal{V}$. Second, it runs $(A, \text{td}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ to generate a matrix $A \in \mathbb{Z}_q^{n \times m}$ along with its trapdoor td . Then, it outputs $\text{sk} = (\text{td}, A, V_1, \dots, V_T)$, $\text{ek} = A$, $\text{vk} = (A, V_1, \dots, V_T)$.

Note that it is possible to associate the key-triple to identity $\text{id} \in \text{ID}$, when we need to stress this explicitly we write $(\text{sk}_{\text{id}}, \text{ek}_{\text{id}}, \text{vk}_{\text{id}})$. We also observe that the key generation process can be seen as the combination of two independent sub-algorithms (This splitting will be used to extend our multi-key homomorphic signature scheme from supporting a single dataset to support multiple datasets. This extension holds in the standard model and is described in Section 4.3.) KeyGen_1 and KeyGen_2 , where $\{V_1, \dots, V_T\} \leftarrow \text{KeyGen}_1(\text{pp})$ and $(A, \text{td}) \leftarrow \text{KeyGen}_2(\text{pp})$.

Sign(sk, ℓ, m). The signing algorithm takes as input a secret key sk , a label $\ell = (\text{id}, \tau)$ for the message m and it outputs a signature $\sigma := (m, z, I, U_{\text{id}}, Z_{\text{id}})$ where $I = \{\text{id}\}$, U_{id} is generated as $U_{\text{id}} \leftarrow \text{SamPre}(A, V_\ell - mG, \text{td})$ (using the algorithm SamPre from Lemma 2), $z = m$ and $Z_{\text{id}} = U_{\text{id}}$. The two latter terms are responsible for protection against Type 3 forgeries. Although they are redundant for fresh signatures, their value will become different from (m, U_{id}) during homomorphic operations, as we clarify later on. More generally, in our construction signatures are of the form $\sigma := (m, z, I, \{U_{\text{id}}\}_{\text{id} \in I}, \{Z_{\text{id}}\}_{\text{id} \in I})$ with $I \subseteq \text{ID}$ and $U_{\text{id}}, Z_{\text{id}} \in \mathcal{U}, \forall \text{id} \in I$.

Eval($f, \{(\sigma_i, \text{EKSi})\}_{i \in [t]}$). The evaluation algorithm takes as input a t -input function $f: \mathcal{M}^t \rightarrow \mathcal{M}$, and a set of pairs $\{(\sigma_i, \text{EKSi})\}_{i \in [t]}$, where each σ_i is a signature and each EKSi is a set of evaluation keys. In our description below we treat f as an arithmetic circuit over \mathbb{Z}_q consisting of addition and multiplication gates (We point out that considering f as an arithmetic circuit over \mathbb{Z}_q is enough to describe any Boolean circuits consisting of NAND gates as $\text{NAND}(m_1, m_2) = 1 - m_1 \cdot m_2$ holds for $m_1, m_2 \in \{0, 1\}$). Therefore, we only describe how to evaluate homomorphically a fan-in-2 addition (resp. multiplication) gate as well as a unary multiplication-by-constant gate.

Let g be a fan-in-2 gate with left input $\sigma_L := (m_L, z_L, I_L, U_L, Z_L)$ and right input $\sigma_R := (m_R, z_R, I_R, U_R, Z_R)$. To generate the signature $\sigma := (m, z, I, U, Z)$ on the gate's output one proceeds as follows. First set $I = I_L \cup I_R$. Second, 'expand' $U_L := \{U_L^{\text{id}}\}_{\text{id} \in I_L}$ as

$$U_L^{\text{id}} = \begin{cases} \mathbf{0} & \text{if } \text{id} \notin I_L \\ U_L^{\text{id}} & \text{if } \text{id} \in I_L \end{cases}, \quad \forall \text{id} \in I$$

where $\mathbf{0}$ denotes an $(m \times m)$ -matrix with all zero entries. Basically, we extend the set to be indexed overall identities in $I = I_L \cup I_R$ by inserting zero matrices for identities in $I \setminus I_L$. The analogous expansion process is applied to $U_R := \{U_R^{\text{id}}\}_{\text{id} \in I_R}$, $Z_L := \{Z_L^{\text{id}}\}_{\text{id} \in I_R}$ and $Z_R := \{Z_R^{\text{id}}\}_{\text{id} \in I_R}$, denoting the expanded sets $\{U_R^{\text{id}}\}_{\text{id} \in I}$, $\{Z_L^{\text{id}}\}_{\text{id} \in I}$ and $\{Z_R^{\text{id}}\}_{\text{id} \in I}$, respectively.

Next, depending on whether g is an addition or multiplication gate, one proceeds as follows:

Addition gate. If g is additive, compute $m = m_L + m_R$, $z = z_L + z_R$,

$$U = \{U_{\text{id}}\}_{\text{id} \in I} := \{U_L^{\text{id}} + U_R^{\text{id}}\}_{\text{id} \in I} \quad (8)$$

And

$$\mathbf{Z} = \{\mathbf{Z}_{id}\}_{id \in I} := \{\hat{\mathbf{Z}}_L^{id} + \hat{\mathbf{Z}}_R^{id}\}_{id \in I}. \quad (9)$$

If we refer to β_L and β_R as

$$\|\mathbf{U}_L\|_\infty := \max\{\|\mathbf{U}_L^{id}\|_\infty : id \in I_L\} \quad (10)$$

And

$$\|\mathbf{U}_R\|_\infty := \max\{\|\mathbf{U}_R^{id}\|_\infty : id \in I_R\}, \quad (11)$$

respectively, then for any fan-in-2 addition gate, it holds $\beta := \|\mathbf{U}\|_\infty = \beta_L + \beta_R$. The same noise growth applies to \mathbf{Z} .

Multiplication gate. If g is multiplicative, compute $m = m_L \cdot m_R$, $z = z_L + z_R$, define $\mathbf{V}_L = \sum_{id \in I_L} \mathbf{A}_{id} \mathbf{U}_{id} + m_L \mathbf{G}$, set

$$\mathbf{U} = \{\mathbf{U}_{id}\}_{id \in I} := \{m_R \hat{\mathbf{U}}_L^{id} + \hat{\mathbf{U}}_R^{id} \cdot \mathbf{G}^{-1}(\mathbf{V}_L)\}_{id \in I}$$

and $\mathbf{Z} = \{\mathbf{Z}_{id}\}_{id \in I} := \{\hat{\mathbf{Z}}_L^{id} + \hat{\mathbf{Z}}_R^{id}\}_{id \in I}$.

Letting β_L and β_R as defined before, then for any fan-in-2 multiplication gate it holds $\beta := \|\mathbf{U}\|_\infty = |m_R| \beta_L + m \beta_R$, while the noise growth of \mathbf{Z} is the same as in the addition gate.

Multiplication by the constant gate. Let g be a unary gate representing a multiplication by a constant $a \in \mathbb{Z}_q$, and let its single input signature be $\sigma_R := (m_R, z_R, I_R, \mathbf{U}_R, \mathbf{Z}_R)$. The output $\sigma := (m, z, I, \mathbf{U}, \mathbf{Z})$ is obtained by setting $m = a \cdot m_R \in \mathbb{Z}_q$, $z = z_R$, $I = I_R$, $\mathbf{Z} = \mathbf{Z}_R$, and $\mathbf{U}\{\mathbf{U}^{id}\}_{id \in I}$, where for all $id \in I$, $\mathbf{U}^{id} = a \cdot \mathbf{U}_R^{id}$ or, alternatively, $\mathbf{U}^{id} = \mathbf{U}_R^{id} \mathbf{G}^{-1}(a \cdot \mathbf{G})$. In the first case, the noise parameter becomes $\beta := \|\mathbf{U}\|_\infty = |a| \beta_L$ (thus a needs to be *small*), whereas in the second case it holds $\beta := \|\mathbf{U}\|_\infty = |m| \beta_L$, which is independent of a 's size.

$\text{Ver}(\mathcal{P}, \{\text{vk}_{id}\}_{id \in \mathcal{P}}, m, \sigma)$. The verification algorithm takes as input a labelled program $\mathcal{P} = (f, \ell_1, \dots, \ell_n)$, the set of the verification keys $\{\text{vk}_{id}\}_{id \in \mathcal{P}}$ of users involved in the program \mathcal{P} , a message m and a signature $\sigma = (m, z, I, \mathbf{U}, \mathbf{Z})$. It then performs three main checks and outputs 0 if at least one check fails. Otherwise it returns 1.

Firstly, it checks if the list of identities declared in σ corresponds to the ones in the labels of \mathcal{P} :

$$I = \{id : id \in \mathcal{P}\} \quad (12)$$

Secondly, from the circuit f (again seen as an arithmetic circuit) and the values $\{\mathbf{V}_{\ell_1}, \dots, \mathbf{V}_{\ell_n}\}$ contained in the verification keys, it computes two values \mathbf{V}^* and \mathbf{V}^+ proceeding gate by gate as follows. Given as left and right input matrices $\mathbf{V}_L^*, \mathbf{V}_R^*$ (resp. $\mathbf{V}_L^+, \mathbf{V}_R^+$), at every *addition gate* one computes $\mathbf{V}^* = \mathbf{V}_L^* + \mathbf{V}_R^*$ (resp. $\mathbf{V}^+ = \mathbf{V}_L^+ + \mathbf{V}_R^+$); at every *multiplication gate* one computes $\mathbf{V}^* = \mathbf{V}_R^* \mathbf{G}^{-1} \mathbf{V}_L^*$ (resp. $\mathbf{V}^+ = \mathbf{V}_L^+ + \mathbf{V}_R^+$). Every gate representing a multiplication by a constant $a \in \mathbb{Z}_q$, on input \mathbf{V}_R^* (resp. \mathbf{V}_R^+) outputs $\mathbf{V}^* = a \cdot \mathbf{V}_R^*$ (resp. $\mathbf{V}^+ = \mathbf{V}_R^+$). Note that the computation of \mathbf{V}^+ is essentially the computation of a linear function $\mathbf{V}^+ = \sum_{i=1}^t \gamma_i \cdot \mathbf{V}_{\ell_i}$, for some coefficients γ_i that depend on the structure of the circuit f .

Thirdly, the verification algorithm parses $\mathbf{U} = \{\mathbf{U}_{id}\}_{id \in I}$ and $\mathbf{Z} = \{\mathbf{Z}_{id}\}_{id \in I}$ and checks

$$\|\mathbf{U}\|_\infty \leq \beta_{\max} \quad \text{and} \quad \|\mathbf{Z}\|_\infty \leq \beta_{\max} \quad (13)$$

$$\sum_{id \in I} \mathbf{A}_{id} \mathbf{U}_{id} + m \cdot \mathbf{G} = \mathbf{V}^* \quad (14)$$

$$\sum_{id \in I} \mathbf{A}_{id} \mathbf{Z}_{id} + z \cdot \mathbf{G} = \mathbf{V}^+ \quad (15)$$

Finally, it is worth noting that the computation of the matrices \mathbf{V}^* and \mathbf{V}^+ can be precomputed (or performed offline), before seeing

the actual signature σ . In the multiple dataset extension of Section 4.3, this precomputation becomes particularly beneficial as the same $\mathbf{V}^*, \mathbf{V}^+$ can be re-used every time one wants to verify for the same labelled program \mathcal{P} (i.e. one can verify faster, in an amortised sense, than that of running f).

In the following paragraphs, we analyse the correctness, succinctness and the security of the proposed construction.

Noise growth and succinctness. First, we analyse the noise growth of the components \mathbf{U}, \mathbf{Z} in the signatures of our MKHSig construction. In particular, we need to show that when starting from ‘fresh’ signatures, in which the noise is bounded by β_{init} , and we apply an admissible circuit, then one ends up with signatures in which the noise is within the allowable amount β_{\max} .

An analysis similar to the one of Gorbunov *et al.* [4] applies also to our construction whenever the admissible functions are Boolean circuits of depth d composed only of NAND gates.

Let us first consider the case of the \mathbf{U} component of the signatures. At every NAND gate, if $\|\mathbf{U}_L\|_\infty, \|\mathbf{U}_R\|_\infty \leq \beta$, the noise of the resulting \mathbf{U} is at most $(m+1)\beta$. Therefore, if the circuit has depth d , the noise of the matrix \mathbf{U} at the end of the evaluation is bounded by

$$\|\mathbf{U}_R\|_\infty \leq \beta_{\text{init}} \cdot (m+1)^d \leq 2^{o(\log \lambda)^d} \leq \beta_{\max}. \quad (16)$$

For what regards the computation performed over the matrix \mathbf{Z} , we observe that we perform only additions (or identity functions) over them. This means that at every gate of any f , the noise in the \mathbf{Z} component at most doubles. Given that we consider depth- d circuits we have that $\|\mathbf{Z}\|_\infty \leq \beta_{\text{init}} \cdot 2^d \leq 2^{o(\log \lambda)^d} \leq \beta_{\max}$. Finally, by inspection, one can see that the size of every signature σ on a computation's output involving n users is at most $(1 + 2^d + n\lambda + 2n\beta_{\max})$, i.e. $O(n \cdot p(\lambda))$ for some fixed polynomial $p(\cdot)$.

Authentication correctness. Each fresh signature of a message m labelled by $\ell := (id, \tau)$ is of the form $\sigma := (m, z, m, I = \{id\}, \mathbf{U}, \mathbf{Z} = \mathbf{U})$. In this case, the labelled program used by Ver to verify σ is the identity program on m with a label $\ell = (id, \tau)$. We need to show that for $\mathcal{P} = \mathcal{F}_\ell$, and m the signature σ passes the four verification checks (12)–(15). It is trivial to check that (12) holds, since $I = \{id\} = \{id \in \mathcal{F}_\ell\}$. Moreover, (13) trivially follows from the fact that $\beta_{\text{init}} \leq \beta_{\max}$. For what concerns (14) we have

$$\sum_{id \in S} \mathbf{A}_{id} \mathbf{U}_{id} + m \cdot \mathbf{G} = \mathbf{A}_{id} \mathbf{U}_{id} + m \cdot \mathbf{G} = \mathbf{V}_\ell = \mathbf{V}^*$$

where the first equality holds because $I = \{id\}$, the second holds by the construction of \mathbf{U}_{id} in the Sign algorithm and the last one holds by the construction of the Ver algorithm. Note that, since here $z = m$ and $\mathbf{U}_{id} = \mathbf{Z}$, then (15) is trivially satisfied.

Evaluation correctness. We prove the correctness of the Eval algorithm for a generic fan-in-2 addition and multiplication gate g : by associativity, this proves the correctness of the algorithm on any circuit. Let m_L (resp. m_R) denote the left (resp. right) input wire of the considered gate g , let $\sigma_L := (m_L, z_L, I_L, \mathbf{U}_L, \mathbf{Z}_L)$ (resp. σ_R) be a signature that verifies correctly for m_L (resp. m_R) and program \mathcal{P}_L (resp. \mathcal{P}_R). Let $\sigma = (m, z, I = I_L \cup I_R, \mathbf{U}, \mathbf{Z})$ be the signature obtained after a gate evaluation.

It is immediate to see that the set I satisfies check (12). Moreover, from the noise growth analysis given above, it results that (13) is satisfied by both \mathbf{U} and \mathbf{Z} .

g additive gate: According to the definition of the algorithm Eval , we have (see equation below).

Below we focus on the proof for (14), and notice that the case of (15) is analogous

$$\begin{aligned}
\sum_{\text{id} \in I} A_{\text{id}} U_{\text{id}} + m \cdot G &= \sum_{\text{id} \in I} A_{\text{id}} U_L^{\text{id}} + \hat{U}_R^{\text{id}} + (m_L + m_R) \cdot G \\
&= \left(\sum_{\text{id} \in I_L} A_{\text{id}} U_L^{\text{id}} + m_L \cdot G \right) \\
&\quad + \left(\sum_{\text{id} \in I_R} A_{\text{id}} U_R^{\text{id}} + m_R \cdot G \right) \\
&= V_L + V_R = V^*
\end{aligned}$$

where the second equation holds because $\hat{U}_L^{\text{id}} = \mathbf{0}^{m \times m}$ for any $\text{id} \in I \setminus I_L$ (resp. $\hat{U}_L^{\text{id}} = \mathbf{0}^{m \times m}$ for any $\text{id} \in I \setminus I_L$), while the last equation follows from the assumption on the valid verification of σ_L (resp. σ_R) and from the definition of the Ver algorithm (w.r.t. the computation over the V_ℓ).

\mathcal{G} *multiplicative gate*: According to the definition of the algorithm Eval, we have (see (17)), where $V_L = \sum_{\text{id} \in I_L} A_{\text{id}} U_{\text{id}} + m_L G$. We focus on proving the case of (14); as the case of (15) follows from the same analysis of the additive gate shown above:

$$\begin{aligned}
\sum_{\text{id} \in I} A_{\text{id}} U_{\text{id}} + m \cdot G &= \sum_{\text{id} \in I} A_{\text{id}} (m_R U_L^{\text{id}} + \hat{U}_R^{\text{id}}) \cdot G^{-1}(V_L) \\
&= (m_L + m_R) \cdot G \\
&= m_R \cdot \left(\sum_{\text{id} \in I} A_{\text{id}} U_L^{\text{id}} + m_L G \right) \\
&\quad + \left(\sum_{\text{id} \in I} A_{\text{id}} \hat{U}_R^{\text{id}} \right) G^{-1}(V_L) \\
&= m_R V_L + \left(\sum_{\text{id} \in I} A_{\text{id}} \hat{U}_R^{\text{id}} \right) G^{-1}(V_L) \\
&= m_R G G^{-1}(V_L) \\
&\quad + \left(\sum_{\text{id} \in I} A_{\text{id}} \hat{U}_R^{\text{id}} \right) G^{-1}(V_L) \\
&= \left[\left(\sum_{\text{id} \in I_R} A_{\text{id}} U_R^{\text{id}} \right) + m_R \right] G^{-1}(V_L) \\
&= V_L + V_R = V^*
\end{aligned}$$

In the first two equations, we expand U and rearrange the terms grouping by $G^{-1}(V_L)$ and m_R , respectively. The third equation follows from the definition of V_L , while in the fourth one, we just rewrite V_L as $GG^{-1}(V_L)$: this allows to group by $G^{-1}(V_L)$ in the fifth equation. The last two equations, respectively, follow from the definition of V_L and V^* .

Security. The following theorem states the security of the scheme MKHSig.

Theorem 1: If the $\text{SIS}(n, m \cdot Q_{\text{id}}, q, \beta_{\text{SIS}})$ hardness assumption holds, $\text{MKHSig} = (\text{Setup}, \text{KeyGen}, \text{Sign}, \text{Eval}, \text{Ver})$ is a multi-key homomorphic signature weakly-adaptive secure against adversaries that make signing queries involving at most Q_{id} different identities and that make non-adaptive corruption queries.

$$\begin{aligned}
\sigma &= (m := m_L + m_R, I = I_L \cup I_R, U = \{U_{\text{id}}\}_{\text{id} \in I} := \{U_L^{\text{id}} + U_R^{\text{id}}\}_{\text{id} \in I}, \\
z &:= z_L + z_R, Z = \{Z_{\text{id}}\}_{\text{id} \in I} := \{Z_L^{\text{id}} + Z_R^{\text{id}}\}_{\text{id} \in I}).
\end{aligned}$$

$$\sigma = (m := m_L + m_R, I = I_L \cup I_R,$$

$$U = \{U_{\text{id}}\}_{\text{id} \in I} := \{m_R U_L^{\text{id}} + U_R^{\text{id}} \cdot G^{-1}(V_L)\}_{\text{id} \in I}, z := z_L + z_R, \quad (17)$$

$$Z = \{Z_{\text{id}}\}_{\text{id} \in I} := \{Z_L^{\text{id}} + Z_R^{\text{id}}\}_{\text{id} \in I},$$

Proof: Note that we can deal with the corruptions via our generic result of Proposition 1. Therefore, it is sufficient to prove the security against adversaries that make no corruptions. Moreover, since this scheme works for a single dataset note that Type 1 forgeries cannot occur.

For the proof, let us recall how the weakly-adaptive security experiment (Definition 6) works for our multi-key homomorphic signature scheme MKHSig. This is a game between an adversary \mathcal{A} and a challenger \mathcal{C} that has four main phases:

- i. \mathcal{A} declares an integer Q representing the number of different identities that it will ask in the signing queries. Moreover, for every $i \in [Q]$ \mathcal{A} sends to \mathcal{C} a set $\mathcal{T}_i \subseteq \mathcal{T} := \{\tau_1, \dots, \tau_T\}$ and a set of pairs $\{(m_\tau, \tau)\}_{\tau \in \mathcal{T}_i}$.
- ii. \mathcal{C} runs $\text{Setup}(1^\lambda)$ to obtain the public parameters and sends them to \mathcal{A} .
- iii. \mathcal{A} adaptively queries identities $\text{id}_1, \dots, \text{id}_Q$. When \mathcal{C} receives the query id_i it generates a key-triple $(\text{sk}_{\text{id}_i}, \text{ek}_{\text{id}_i}, \text{vk}_{\text{id}_i})$ by running $\text{KeyGen}(\text{pp})$, and for all labels $\ell = (\text{id}_i, \tau)$ such that $\tau \in \mathcal{T}_i$ it runs $\sigma_\tau^i \leftarrow \text{Sign}(\text{sk}_{\text{id}_i}, \ell, m_\tau)$. Then \mathcal{C} sends to \mathcal{A} : the public keys $\text{vk}_{\text{id}_i} := (A_{\text{id}_i}, \{V_\ell\}_{\tau \in \mathcal{T}_i})$ and $\text{ek}_{\text{id}_i} := (A_{\text{id}_i})$, and the signatures $\{\sigma_\tau^i\}_{\tau \in \mathcal{T}_i}$.
- iv. The adversary produces a forgery consisting of a labelled program $\mathcal{P}^* = (f^*, \ell_1^*, \dots, \ell_i^*)$ where $f^* \in \mathcal{F}, f^*: \mathcal{M}^i \rightarrow \mathcal{M}$, a message m^* and a signature σ^* .

\mathcal{A} wins the non-adaptive security game if $\text{Ver}(\mathcal{P}^*, \{\text{vk}_{\text{id}_i}\}_{\text{id}_i \in \mathcal{P}^*}, m^*, \sigma^*) = 1$ and one of the following conditions holds:

Type 2 forgery: There exist messages $m_{\ell_1^*}, \dots, m_{\ell_i^*}$ s.t. $m^* \neq f^*(m_{\ell_1^*}, \dots, m_{\ell_i^*})$ (i.e. m^* is not the correct output of \mathcal{P}^* when executed over previously signed messages).

Type 3 forgery: There exists at least one label $\ell^* = (\text{id}^*, \tau^*)$ that was not queried by \mathcal{A} .

Consider a variation of the above game obtained modifying phase (4.2) as follows:

(4.2') \mathcal{C} picks an instance $A \in \mathbb{Z}_q^{n \times m'}$ of the $\text{SIS}(n, m', q, \beta_{\text{SIS}})$ problem for $m' = m \cdot Q = \text{poly}(\lambda)$, and parse $A := A_{\text{id}_1} \parallel \dots \parallel A_{\text{id}_Q} \in \mathbb{Z}_q^{n \times m'}$ as the concatenation of Q different blocks of $n \times m$ matrices.

Next, when \mathcal{C} receives the i th query id_i from \mathcal{A} , it does the following:

- It samples a matrix $U_{\text{id}_i, \tau} \xleftarrow{\$} \mathcal{U}$ such that $\|U_{\text{id}_i, \tau}\| \beta_{\text{init}}$.
- For all $\ell := (\text{id}_i, \tau)$ with $\tau \in \mathcal{T}_i$, \mathcal{C} computes $V_\ell = A_{\text{id}_i} U_{\text{id}_i, \tau} + m_\tau \cdot G$.
- For all $\ell := (\text{id}_i, \tau)$ with $\tau \notin \mathcal{T}_i$, \mathcal{C} computes $V_\ell = A_{\text{id}_i} U_{\text{id}_i, \tau} + b_{i, \tau} \cdot G$, where $b_{i, \tau} \xleftarrow{\$} \{0, 1\}$.
- \mathcal{C} sends to \mathcal{A} the public keys $\text{vk}_{\text{id}_i} := (A_{\text{id}_i}, \{V_\ell\}_{\tau \in \mathcal{T}_i})$ and $\text{ek}_{\text{id}_i} := (A_{\text{id}_i})$, along with signatures $\{\sigma_\tau^i\}_{\tau \in \mathcal{T}_i}$ where $\sigma_\tau^i := (m_\tau, m_\tau, I: \{\text{id}_i\}, U_{\text{id}_i, \tau}, U_{\text{id}_i, \tau})$.

Clearly, if \mathbf{A} is a uniformly random matrix so is each block $\{\mathbf{A}_{i d_i}\}_{i \in [Q]}$.

Due to point (2) of Lemma 2, since $(\mathbf{A}_{i d_i} \mathbf{U}_{i d_i, \tau})$ is statistically indistinguishable from a random matrix, all the matrices \mathbf{V}_ℓ generated in (4.2') are statistically close to the ones generated in (4.2). Thus, the two games are statistically indistinguishable. At this point, we show that for every PPT adversary \mathcal{A} which produces a forgery in the modified game we can construct a PPT algorithm \mathcal{B} that solves the SIS($n, m \cdot Q, q, \beta_{\text{SIS}}$) problem. \mathcal{B} receives an SIS instance $\mathbf{A}_{i d_i} := (\mathbf{A}_{i d_1} | \dots | \mathbf{A}_{i d_Q}) \in \mathbb{Z}_q^{n \times mQ}$ and simulates the modified game to \mathcal{A} by acting exactly as the challenger \mathcal{C} described above. Then, once \mathcal{A} outputs its forgery, according to the forgery's type, \mathcal{B} proceeds as described below.

Type 2 forgeries. Let

$$(\mathcal{P}^* := (f^*, \ell_1^*, \dots, \ell_t^*), \mathbf{m}^*, \sigma^* := (\mathbf{m}^*, \mathbf{z}^*, \mathbf{I}^*, \mathbf{U}^*, \mathbf{Z}^*)) \quad (18)$$

be a Type 2 forgery produced by \mathcal{A} in the modified game. Moreover, let $\sigma = (\mathbf{m}, \mathbf{z}, \mathbf{I}, \mathbf{U}, \mathbf{Z})$ be the signature obtained by honestly applying Eval to the signatures corresponding to labels $\ell_1^*, \dots, \ell_t^*$ that were given to \mathcal{A} . Parse $\mathbf{U} := \{\mathbf{U}_{i d_i}\}_{i d_i \in \mathbf{I}}$ and notice that by the correctness of the scheme we have that

$$\mathbf{m} = f^*(\mathbf{m}_{\ell_1^*}, \dots, \mathbf{m}_{\ell_t^*}), \mathbf{I} = \{\text{id} : \text{id} \in \mathcal{P}^*\}, \quad (19)$$

and

$$\sum_{i d_i \in \mathbf{I}} \mathbf{A}_{i d_i} \mathbf{U}_{i d_i} + \mathbf{m} \cdot \mathbf{G} = \mathbf{V}^*.$$

Moreover, by definition of Type 2 forgery recall that $\mathbf{m}^* \neq f^*(\mathbf{m}_{\ell_1^*}, \dots, \mathbf{m}_{\ell_t^*})$ and that the tuple satisfies verification. In particular, the satisfaction of check (12) implies that $\mathbf{I} = \mathbf{I}^*$, while check (14) means

$$\sum_{i d_i \in \mathbf{I}^*} \mathbf{A}_{i d_i} \mathbf{U}_{i d_i}^* + \mathbf{m}^* \cdot \mathbf{G} = \mathbf{V}^*.$$

Combining the two equations above, we obtain $\sum_{i d_i \in \mathbf{I}} \mathbf{A}_{i d_i} \tilde{\mathbf{U}}_{i d_i} + \tilde{\mathbf{m}} \cdot \mathbf{G}$, where $\tilde{\mathbf{m}} = \mathbf{m} - \mathbf{m}^* \neq \mathbf{0}$ and, for all $i d_i \in \mathbf{I}$, $\tilde{\mathbf{U}}_{i d_i} = \mathbf{U}_{i d_i}^* - \mathbf{U}_{i d_i} \in \mathcal{U}$ such that $\|\tilde{\mathbf{U}}_{i d_i}\|_\infty \leq \beta_{\text{max}}$. Notice that there must exist at least one $i d_i \in \mathbf{I}$ for which $\mathbf{U}_{i d_i} \neq \mathbf{0}$.

Moreover, for all $i d_i \in \{\text{id}_1, \dots, \text{id}_Q\} \setminus \mathbf{I}$, define $\tilde{\mathbf{U}}_{i d_i} = \mathbf{0}$ and set

$$\tilde{\mathbf{U}} = \begin{pmatrix} \tilde{\mathbf{U}}_{i d_1} \\ \vdots \\ \tilde{\mathbf{U}}_{i d_Q} \end{pmatrix} \in \mathbb{Z}_q^{mQ \times m}. \text{ Then, we have } \mathbf{A} \tilde{\mathbf{U}} = \tilde{\mathbf{m}} \cdot \mathbf{G}.$$

Next \mathcal{B} samples $\mathbf{r} \xleftarrow{\$} \{0, 1\}^{mQ}$, set $\mathbf{s} = \mathbf{A} \mathbf{r} \in \mathbb{Z}_q^n$, and computes $\mathbf{r}' = \mathbf{G}^{-1}(\tilde{\mathbf{m}}^{-1} \cdot \mathbf{s})$, so that $\mathbf{r}' \in \{0, 1\}^m$ and $\tilde{\mathbf{m}} \cdot \mathbf{G} \mathbf{r}' = \mathbf{s}$. Finally, \mathcal{B} outputs $\mathbf{u} = \tilde{\mathbf{U}} \mathbf{r}' - \mathbf{r} \in \mathbb{Z}_q^{mQ}$. We conclude the proof by claiming that the vector \mathbf{u} returned by \mathcal{B} is a solution to the SIS problem for the matrix \mathbf{A} . To see this observe that

$$\mathbf{A}(\tilde{\mathbf{U}} \mathbf{r}' - \mathbf{r}) = (\mathbf{A} \tilde{\mathbf{U}}) \mathbf{r}' - \mathbf{A} \mathbf{r} = \tilde{\mathbf{m}} \cdot \mathbf{G} \cdot \mathbf{G}^{-1}(\tilde{\mathbf{m}}^{-1} \cdot \mathbf{s}) - \mathbf{s} = \mathbf{0}$$

and $\|\mathbf{u}\|_\infty \leq (2m + 1)\beta_{\text{max}} \leq \beta_{\text{SIS}}$.

It remains to show that $\mathbf{u} \neq \mathbf{0}$. We show that this is the case (i.e. $\tilde{\mathbf{U}} \mathbf{r}' \neq \mathbf{r}$) with overwhelming probability by using an entropy argument (the same argument used in [4]). In particular, this holds for any (worst case) choice of $\mathbf{A}, \tilde{\mathbf{U}}, \tilde{\mathbf{m}}$, and only based on the random choice of $\mathbf{r} \xleftarrow{\$} \{0, 1\}^{mQ}$. The intuition is that, even if $\mathbf{r}' = \mathbf{G}^{-1}(\tilde{\mathbf{m}}^{-1} \cdot \mathbf{s})$ depends on $\mathbf{s} = \mathbf{A} \mathbf{r}$, \mathbf{s} is too small to reveal much information about the random \mathbf{r} . More precisely, we have that $\mathbf{H}_\infty(\mathbf{r} | \mathbf{r}') \geq \mathbf{H}_\infty(\mathbf{r} | \mathbf{A} \mathbf{r})$ because \mathbf{r}' is chosen deterministically based on $\mathbf{s} = \mathbf{A} \mathbf{r}$. Due to Lemma 1, we have that

$\mathbf{H}_\infty(\mathbf{r} | \mathbf{A} \mathbf{r}) \geq \mathbf{H}_\infty(\mathbf{r}) - \log(|\mathcal{S}|)$, where \mathcal{S} is the space of all possible \mathbf{s} . Since $\mathbf{s} \in \mathbb{Z}_q^n$, $|\mathcal{S}| = q^n$, and then

$$\log(|\mathcal{S}|) = \log(q^n) = \log((2^{\log q})^n) = n \log(2^{\log q}) = n \log q. \quad (20)$$

Regarding $\mathbf{H}_\infty(\mathbf{r})$, since $\mathbf{H}_\infty(X) := -\log(\max_x \Pr[X = x])$, we have $\mathbf{H}_\infty(\mathbf{r}) = -\log(2^{-mQ}) = mQ \geq m$. Then

$$\mathbf{H}_\infty(\mathbf{r} | \mathbf{r}') \geq \mathbf{H}_\infty(\mathbf{r}) - \log(|\mathcal{S}|) \geq m - n \log q = \omega(\log \lambda)$$

Since we know that for random variables X, Y the optimal probability of an unbounded adversary guessing X given the correlated value Y is $2^{-\mathbf{H}_\infty(X|Y)}$, then

$$\Pr[\mathbf{r} = \tilde{\mathbf{U}} \mathbf{r}'] \leq 2^{-\mathbf{H}_\infty(\mathbf{r} | \mathbf{r}')} \leq 2^{-\omega(\log \lambda)} = \text{negl}(\lambda). \quad (21)$$

Type 3 forgery. Let

$$(\mathcal{P}^* := (f^*, \ell_1^*, \dots, \ell_t^*), \mathbf{m}^*, \sigma^* := (\mathbf{m}^*, \mathbf{z}^*, \mathbf{I}^*, \mathbf{U}^*, \mathbf{Z}^*)) \quad (22)$$

be a Type 3 forgery produced by \mathcal{A} in the modified game such that there exists (at least) one label $\ell_j^* = (\text{id}^*, \tau^*)$ such that $\text{id}^* = \text{id}_i$ but $\tau^* \notin \mathcal{T}_i$ (It is easy to see that the case in which id^* is new would imply the generation of a new $\mathbf{A}_{i d^*}$, which would make the verification equations hold with negligible probability (over the random choice of $\mathbf{A}_{i d^*}$). Actually, without loss of generality, we can assume that there is exactly one of such labels; if this is not the case, one could indeed redefine another adversary that makes more queries until it misses only this one. Note that for such a tag $\tau^* \notin \mathcal{T}_i$, \mathcal{B} simulated $\mathbf{V}_{i d_i, \tau^*} = \mathbf{A} \mathbf{U}_{i d_i, \tau^*} + b_{i, \tau^*} \mathbf{G}$ for a randomly chosen bit $b_{i, \tau^*} \xleftarrow{\$} \{0, 1\}$, that is perfectly hidden from \mathcal{A} .

By definition of Type 3 forgery, the tuple passes verification, and in particular check (15)

$$\sum_{i d_i \in \mathbf{I}^*} \mathbf{A}_{i d_i} \mathbf{Z}_{i d_i}^* + \mathbf{z}^* \cdot \mathbf{G} = \mathbf{V}^* = \sum_{i=1}^t \gamma_i \cdot \mathbf{V}_{\ell_i^*}$$

where the right-hand side of the equation holds by the construction of the verification algorithm.

Moreover, let $\sigma = (\mathbf{m}, \mathbf{z}, \mathbf{I}, \mathbf{U}, \mathbf{Z})$ be the signature obtained by honestly applying Eval to the signatures corresponding to labels $\ell_1^*, \dots, \ell_t^*$; in particular, for the specific, missing, label ℓ_j^* \mathcal{B} uses the values $\mathbf{U}_{i d_i, \tau^*}, b_{i, \tau^*}$ used to simulate $\mathbf{V}_{i d_i, \tau^*}$. Parsing $\mathbf{Z} := \{\mathbf{Z}_{i d_i}\}_{i d_i \in \mathbf{I}}$, notice that by correctness it holds $\mathbf{I} = \{\text{id} : \text{id} \in \mathcal{P}^*\}$ and

$$\sum_{i d_i \in \mathbf{I}} \mathbf{A}_{i d_i} \mathbf{Z}_{i d_i} + \mathbf{z} \cdot \mathbf{G} = \mathbf{V}^*$$

where $\mathbf{z} = \sum_{i=1, i \neq j}^t \gamma_i \mathbf{m}_i + \gamma_j b_{i, \tau^*}$. Now, the observation is that every $\gamma_i \leq 2^d < q$, i.e. $\gamma_i \neq 0 \pmod q$. Since b_{i, τ^*} is random and perfectly hidden to \mathcal{A} we have that with probability 1/2 it holds $\mathbf{z} \neq \mathbf{z}^*$.

Thus, if $\mathbf{z} \neq \mathbf{z}^*$, \mathcal{B} combines the equalities on \mathbf{V}^* to come up with an equation

$$\sum_{i d_i \in \mathbf{I}} \mathbf{A}_{i d_i} \tilde{\mathbf{Z}}_{i d_i} = \tilde{\mathbf{z}} \cdot \mathbf{G}$$

where $\tilde{\mathbf{z}} = \mathbf{z} - \mathbf{z}^* \neq \mathbf{0} \pmod q$ and, for all $i d_i \in \mathbf{I}$, $\tilde{\mathbf{Z}}_{i d_i} = \mathbf{Z}_{i d_i}^* - \mathbf{Z}_{i d_i} \in \mathcal{U}$ such that $\|\tilde{\mathbf{Z}}_{i d_i}\|_\infty \leq \beta_{\text{max}}$.

Finally, using the same technique as in the case of Type 2 forgeries, \mathcal{B} can compute a vector \mathbf{u} that is a solution of SIS with overwhelming probability, i.e. $\mathbf{A} \mathbf{u} = \mathbf{0}$.

Therefore, we have proven that if an adversary \mathcal{A} can break the MKHSig scheme with non-negligible probability, then \mathcal{C} can use

such an \mathcal{A} to break the SIS assumption for A with non-negligible probability as well. \square

Variant with unbounded tag space in the random oracle model. In this section, we show that the construction of multi-key homomorphic signatures of Section 4.2 can be easily modified to have short public keys and to support an unbounded tag space $\mathcal{T} = \{0, 1\}^*$. Note that once arbitrary tags are allowed, the scheme also allows handling *multiple datasets* for free. In fact, one can always extend tags to include the dataset name, i.e. simply redefine each tag τ as consisting of two substrings $\tau = (\Delta, \tau')$, where Δ is the dataset name and τ' the actual tag.

The idea of modifying the scheme to support an unbounded tag space is simple and was also suggested in [4] for their construction. Instead of sampling matrices $\{V_{id,1}, \dots, V_{id,T}\}$ in KeyGen , one can just choose a random string $r_{id} \xleftarrow{\$} \{0, 1\}^\lambda$ and define every $V_{id,\tau} := \hat{H}(r_{id}, \tau)$, where $\hat{H}: \{0, 1\}^* \rightarrow \mathcal{V}$ is a hash function chosen in Setup (modelled as a random oracle in the proof). In all the remaining algorithms, every time one needs $V_{id,\tau}$, this is obtained using H .

For this modified scheme, we also provide an idea of how the security proof of Theorem 1 has to be modified to account for these changes. The main change is the simulation of hash queries, which is done as follows:

Before phase (4.2), where \mathcal{A} declares its queries, \mathcal{B} simply answers every query $\hat{H}(r, \tau)$ with a randomly chosen $V \xleftarrow{\$} \mathcal{V}$. Afterwards, once \mathcal{A} has declared all its queries, \mathcal{B} chooses $r_{id_1}, \dots, r_{id_Q} \xleftarrow{\$} \{0, 1\}^\lambda$ and programs the random oracle so that, for all $\tau \in \mathcal{T}_i$, $\hat{H}(r_{id_i}, \tau) = V_{id_i,\tau}$ where $V_{id_i,\tau}$ is the same matrix generated in the phase (4.2) of the modified game. On the other hand, for all $\tau \notin \mathcal{T}_i$, $\hat{H}(r_{id_i}, \tau) = V_{id_i,\tau}$ where $V_{id_i,\tau} = A_{id}U_{id_i,\tau} + b_{i,\tau}G$ for a randomly chosen $U_{id_i,\tau} \xleftarrow{\$} \mathcal{U}$. All other queries $\hat{H}(r, \tau)$ where $r \neq r_{id_i}, \forall i \in [Q]$ are answered with random $V \xleftarrow{\$} \mathcal{V}$. With this simulation, it is not hard to see that, from \mathcal{A} 's forgery \mathcal{B} can extract a solution for SIS (except for some negligible probability that \mathcal{A} guesses one of r_{id_i} before seeing it).

4.3 From a single dataset to multiple datasets

In this section, we present a generic transformation to convert a single-dataset MKHSig scheme into a scheme that supports multiple datasets. The intuition behind this transformation is similar to the one employed in [4] and implicitly used in [12, 14], except that here we have to use additional techniques to deal with the multi-key setting. We combine a standard signature scheme NH.Sig (non-homomorphic) with a single dataset multi-key homomorphic signature scheme MKHSig' . The idea is that for every new dataset Δ , every user generates fresh keys of the multi-key homomorphic scheme MKHSig' and then uses the standard signature scheme NH.Sig to sign the dataset identifier Δ together with the generated public key. More precisely, in our transformation, we assume to start with (single-dataset) multi-key homomorphic signature schemes in which the key generation algorithm can be split into two independent algorithms: KeyGen_1 that outputs some public parameters related to the identity id , and KeyGen_2 which outputs the actual keys. Differently than [4], in our scheme, the signer does not need to sign the whole dataset at once, nor has to fix a bound N on the dataset size (unless such a bound is already contained in MKHSig').

In more details, let $\text{NH.Sig} = (\text{NH.KeyGen}, \text{NH.Sign}, \text{NH.Ver})$ be a standard (non-homomorphic) signature scheme, and let $\text{MKHSig}' = (\text{Setup}', \text{KeyGen}', \text{Sign}', \text{Eval}', \text{Ver}')$ be a single-dataset multi-key homomorphic signature scheme. We construct a multi-dataset multi-key homomorphic signature scheme $\text{MKHSig} = (\text{Setup}, \text{KeyGen}, \text{Sign}, \text{Eval}, \text{Ver})$ as follows:

$\text{Setup}(1^\lambda)$. The setup algorithm samples parameters of the single-dataset multi-key homomorphic signature scheme,

$\text{pp}' \leftarrow \text{Setup}'(1^\lambda)$, together with a description of a PRF $F: K \times \{0, 1\}^* \rightarrow \{0, 1\}^\rho$, and outputs $\text{pp} = (\text{pp}', F)$.

$\text{KeyGen}(\text{pp})$. The key generation algorithm runs NH.KeyGen to get $(\text{pk}_{id}^{\text{NH}}, \text{sk}_{id}^{\text{NH}})$, a pair of keys for the standard signature scheme. Also, it runs KeyGen_1 to generate user-specific public parameters pp_{id} , and chooses a seed K_{id} for the PRF F . The final output is the vector $(\text{sk}_{id}, \text{ek}_{id}, \text{vk}_{id})$, where $\text{sk}_{id} = (\text{sk}_{id}^{\text{NH}}, K_{id})$, $\text{ek}_{id} = (\text{pp}_{id})$ and $\text{vk}_{id} = (\text{pk}_{id}^{\text{NH}}, \text{pp}_{id})$.

$\text{Sign}(\text{sk}_{id}, \Delta, \ell, m)$. The signing algorithm proceeds as follows. First, it samples the keys of the single-dataset multi-key homomorphic signature scheme by feeding randomness $F_{K_{id}}(\Delta)$ to KeyGen_2 , i.e. it runs $\text{KeyGen}_2(\text{pp}; F_{K_{id}}(\Delta))$ to obtain the keys $(\text{sk}_{id}^\Delta, \text{ek}_{id}^\Delta, \text{vk}_{id}^\Delta)$ (Here we assume that a ρ -bits string is sufficient, otherwise, it can always be stretched using a PRG.). The algorithm then runs $\sigma' \leftarrow \text{Sign}'(\text{sk}_{id}^\Delta, \ell, m)$, and uses the non-homomorphic scheme to sign the concatenation of the public key vk_{id}^Δ and the dataset identifier Δ , i.e. $\sigma_{id}^\Delta \leftarrow \text{NH.Sign}(\text{sk}_{id}^{\text{NH}}, \text{vk}_{id}^\Delta \parallel \Delta)$. The output is the tuple $\sigma := (I = \{id\}, \sigma', \text{par}_\Delta)$, where $\text{par}_\Delta = \{(\text{ek}_{id}^\Delta, \text{vk}_{id}^\Delta, \sigma_{id}^\Delta)\}$. Note that the use of the PRF allows every signer (having the same K_{id}) to generate the same keys of the scheme MKHSig' on the same dataset Δ .

$\text{Eval}(f, \{(\sigma_i, \text{EKS}_i)\}_{i \in [t]})$. For each $i \in [t]$, the algorithm parses every signature as $\sigma_i := (I_i, \sigma'_i, \text{par}_{\Delta_i})$ with $\text{par}_{\Delta_i} = \{(\text{ek}_{id}^\Delta, \text{vk}_{id}^\Delta, \sigma_{id}^\Delta)\}_{id \in I_i}$ and sets $\text{EKS}'_i = \{(\text{ek}_{id}^\Delta)\}_{id \in I_i}$. It computes the signatures $\sigma' \leftarrow \text{Eval}'(f, \{(\sigma'_i, \text{EKS}'_i)\}_{i \in [t]})$, defines $I = \bigcup_{i=1}^t I_i$ and $\text{par}_\Delta = \bigcup_{i=1}^t \text{par}_{\Delta_i}$. The final output is $\sigma = (I, \sigma', \text{par}_\Delta)$.

$\text{Ver}(\mathcal{P}, \Delta, \{\text{vk}_{id}\}_{id \in \mathcal{P}}, m, \sigma)$. The verification algorithm begins by parsing the verification keys as $\text{vk}_{id} := (\text{pk}_{id}^{\text{NH}}, \text{pp}_{id})$ for each $id \in I$, and also the signature as $\sigma = (I, \sigma', \text{par}_\Delta)$ with $\text{par}_\Delta = \{(\text{ek}_{id}^\Delta, \text{vk}_{id}^\Delta, \sigma_{id}^\Delta)\}_{id \in I}$. Then, it proceeds with two main steps. First, for each $id \in I$, it verifies the standard signature σ_{id}^Δ on the public key of the single-dataset multi-key homomorphic scheme and the given dataset, i.e. it checks whether $\text{NH.Ver}(\text{pk}_{id}^{\text{NH}}, \text{vk}_{id}^\Delta \parallel \Delta, \sigma_{id}^\Delta) = 1, \forall id \in I$. If at least one of the previous equations is not satisfied, the algorithm returns 0. Otherwise it proceeds to the second check and returns the output of $\text{Ver}'(\mathcal{P}, \{\text{pp}_{id}, \text{vk}_{id}\}_{id \in \mathcal{P}}, m, \sigma)$.

Authentication correctness. The correctness of the scheme substantially follows from the correctness of the regular signature scheme NH.Sig , the single-dataset multi-key homomorphic scheme MKHSig' and the PRF F . More formally, a fresh signature σ on $(\text{sk}_{id}^\Delta, \Delta, \ell, m)$ is of the form $(I, \sigma', \text{par}_\Delta)$ where $\text{par}_\Delta = \{(\text{ek}_{id}^\Delta, \text{vk}_{id}^\Delta, \sigma_{id}^\Delta)\}$. By construction $\sigma_{id}^\Delta \leftarrow \text{NH.Sign}(\text{sk}_{id}^{\text{NH}}, \text{vk}_{id}^\Delta \parallel \Delta)$, and thus $\text{NH.Ver}(\text{pk}_{id}^{\text{NH}}, \text{vk}_{id}^\Delta \parallel \Delta, \sigma_{id}^\Delta)$ outputs 1 for the correctness of the NH.Sig . Similarly, being \mathcal{P} the identity program \mathcal{S}_ℓ , it holds that $\text{Ver}'(\mathcal{P}, \{\text{pp}_{id}, \text{vk}_{id}^\Delta\}, m, \sigma') = 1$ by the correctness of the MKHSig scheme.

Evaluation correctness. Evaluation correctness follows directly from the correctness of the evaluation algorithm Eval' of the single-dataset MKHSig scheme, the correctness of NH.Sig and of the PRF.

Security. Intuitively, the security of the scheme follows from two main observations. Firstly, no adversary can fake the keys of the single-dataset multi-key homomorphic signature scheme, due to the security of the standard signature scheme and the property of pseudo-random functions. Secondly, no adversary can tamper with the results of Eval for a specific dataset as this would correspond to breaking the security of the single-dataset multi-key homomorphic signature scheme.

Theorem 2: If F is a secure pseudo-random function, NH.Sig is an unforgeable signature scheme and MKHSig' is a secure single-

dataset multi-key homomorphic signature scheme, then the MKHSig scheme for multiple datasets described in Section 4.3 is secure against adversaries that make static corruptions of keys and produce forgeries as in Definition 4.

To prove Theorem 2, we define a sequence of hybrid games that we will use to bound the probability of winning in the original security game described in Section 3. Let Game 0 be the security game between the adversary \mathcal{A} and the challenger \mathcal{C} described in Section 3. Let us define two more hybrid games:

Game 1. This is the same as Game 0 except that every PRF instance $F_{k_{id}}(\cdot)$ is replaced with a truly random function $\mathcal{R}_{id}: \{0, 1\}^* \rightarrow \{0, 1\}^\rho$. Basically, for every dataset identifier Δ used in the game, Game 1 generates keys $(sk_{id}, ek_{id}, vk_{id})$ using KeyGen_2 using truly random coins as output by $\mathcal{R}_{id}(\Delta)$.

Game 2. This is the same as Game 1 except that the challenger rejects any adversary's forgery $(\mathcal{P}^*, \Delta^*, m^*, \sigma^*)$, where Δ^* had not been queried during the authentication query phase (i.e. L_{Δ^*} was not initialised). Note that in Game 2 the adversary can no longer make Type 1 forgeries.

Let $G_i(\mathcal{A})$ denote the event that Game i , executed with an adversary \mathcal{A} , outputs 1. The proof of Theorem 2 is obtained by showing that any PPT adversary has a negligible probability of distinguishing between Game 0 and Game 1 (Lemma 3), and between Game 1 and Game 2 (Lemma 4). Finally, we show that \mathcal{A} has a negligible probability of winning in the final game (Lemma 5).

Lemma 3: For every PPT \mathcal{A} there is a PPT distinguisher \mathcal{D} such that $|\Pr[G_0(\mathcal{A})] - \Pr[G_1(\mathcal{A})]| \leq Q_{id} \cdot \text{Adv}_{F, \mathcal{D}}^{\text{PRF}}(\lambda)$, where n is the number of distinct identities queried by \mathcal{A} during authentication queries.

Proof: The proof of Lemma 3 is rather simple and proceeds as follows. Given that the only difference between Game 0 and Game 1 is in the random coins used in the algorithm KeyGen_2 , the probability of \mathcal{A} winning in Game 1 is the same as in Game 0, a part from an additive factor that comes from distinguishing the PRF instance $F_{k_{id}}(\cdot)$ from a truly random function $\mathcal{R}_{id}(\cdot)$. Since we assume \mathcal{A} performs at most Q_{id} distinct PRF instances during the game, a fairly standard hybrid argument assures that the difference between the two games is bound by a $Q_{id} \cdot \text{Adv}_{F, \mathcal{D}}^{\text{PRF}}(\lambda)$ factor. \square

Lemma 4: For every PPT \mathcal{A} there is a PPT forger \mathcal{F} such that $|\Pr[G_1(\mathcal{A})] - \Pr[G_2(\mathcal{A})]| \leq Q_{id} \cdot \text{Adv}_{\text{NH.Sig}, \mathcal{F}}^{\text{UF-CMA}}(\lambda)$, where $\text{Adv}_{\text{NH.Sig}, \mathcal{F}}^{\text{UF-CMA}}(\lambda)$ is the advantage of \mathcal{F} breaking the existential unforgeability of the standard (non-homomorphic) signature scheme and Q_{id} is the number of distinct identities queried by \mathcal{A} .

Proof: For the proof of Lemma 4, we notice that the outputs of Game 1 and Game 2 coincide on every forgery $(\mathcal{P}^*, \Delta^*, m^*, \sigma^*)$ except when \mathcal{A} produces a candidate Type 1 forgery, i.e. when Δ^* has never been queried during the game. Let Bad denote the event in which the adversary outputs a valid Type 1 forgery, then $|\Pr[G_1(\mathcal{A})] - \Pr[G_2(\mathcal{A})]| = \Pr[G_1(\mathcal{A}) \wedge \text{Bad}]$. In what follows we bound this probability by showing that for every adversary \mathcal{A} that wins in Game 1 by returning a Type 1 forgery, there is a forger \mathcal{F} that produces an existential forgery against the scheme NH.Sig .

Observe that the forgery $(\mathcal{P}^*, \Delta^*, m^*, \sigma^*)$ passes verification Ver of the multi-dataset MKHSig and that \mathcal{P}^* contains at least one identity. Fix any identity $\bar{id} \in \mathcal{P}^*$ and let it be the k th identity that \mathcal{A} queried in the game. The crucial observation is that since Δ^* has never been queried, the forgery contains a signature $\sigma_{\bar{id}}^*$, which was not generated by the challenger and yet verifies correctly under the public key $pk_{\bar{id}}^{\text{NH}}$. Therefore, for any such adversary \mathcal{A} , we can build a PPT \mathcal{F} which wins in the unforgeability game of NH.Sig with non-negligible probability. The forger \mathcal{F} works as follows. First, \mathcal{F}

samples a random index $k^* \xleftarrow{\$} [Q_{id}]$, which represents its guess for the index of the identity id on which \mathcal{A} will make the forgery (i.e. k^* is a guess for the above k). Second, \mathcal{F} simulates Game 1 to \mathcal{A} by generating the keys of the NH.Sig scheme for all identities except the k^* th one, for which \mathcal{F} invokes its own oracles for the simulation. The forger \mathcal{F} also generates all the keys of the multi-key homomorphic scheme necessary in the simulation.

Then, let $(\mathcal{P}^*, \Delta^*, m^*, \sigma^*)$ be the Type 1 forgery returned by \mathcal{A} , where $\sigma^* = (I^*, \sigma', \text{par}_{\Delta^*})$ and $\text{par}_{\Delta^*} = \{(ek_{id}^{\Delta^*}, vk_{id}^{\Delta^*}, \sigma_{id}^{\Delta^*})\}_{id \in I^*}$. If the k^* th identity id is in \mathcal{P}^* then $\sigma_{id}^{\Delta^*}$ is such that $\text{NH.Ver}(pk_{id}^{\text{NH}}, (vk_{id}^{\Delta^*} | \Delta), \sigma_{id}^{\Delta^*}) = 1$, i.e. the signature $\sigma_{id}^{\Delta^*}$ verifies correctly on the message $(vk_{id}^{\Delta^*} | \Delta^*)$. Since no message $(\cdot | \Delta^*)$ was queried by \mathcal{F} to its oracle, the pair $((vk_{id}^{\Delta^*} | \Delta), \sigma_{id}^{\Delta^*})$ is an existential forgery for NH.Sig . In case the k^* th identity id is not in \mathcal{P}^* , the forger \mathcal{F} aborts the simulation.

In conclusion, if $\Pr[G_1(\mathcal{A}) \wedge \text{Bad}] > \epsilon$, then \mathcal{F} has an advantage $> \epsilon/Q_{id}$ in breaking the security of NH.Sig , where the factor $1/Q_{id}$ is the probability that \mathcal{F} guesses correctly the index k^* , i.e. $\Pr[k = k^*]$. Since by assumption NH.Sig is an unforgeable signature scheme, we derive that

$$\left| \Pr[G_1(\mathcal{A})] - \Pr[G_2(\mathcal{A})] \right| = \Pr[G_1(\mathcal{A}) \wedge \text{Bad}] \leq Q_{id} \cdot \text{Adv}_{\text{NH.Sig}, \mathcal{F}}^{\text{UF-CMA}}(\lambda), \quad (23)$$

which concludes the proof. \square

Lemma 5: For every PPT adversary \mathcal{A} there is a PPT algorithm \mathcal{B} such that $\Pr[G_2(\mathcal{A})] Q_{\Delta} \cdot \text{Adv}_{\text{MKHSig}, \mathcal{B}}^{\text{HomUF-CMA}}(\lambda)$, where Q_{Δ} is the number of distinct datasets queried by \mathcal{A} .

Proof: The intuition of the proof of Lemma 5 is the following: since Game 2 rules out all Type 1 forgeries, the only way in which \mathcal{A} can win Game 2 is to produce a forgery on a single, previously queried dataset. In this case, we will reduce the security of Game 2 to the standard security game of the single-dataset multi-key homomorphic signature.

We have already noticed that $G_2(\mathcal{A})$ happens if and only if \mathcal{A} produces a Type 2 or Type 3 forgery against the multi-dataset multi-client homomorphic signature scheme, where the pseudo-random function F is replaced by a random function \mathcal{R} . The latter change assures that for every new queried dataset the fresh generated keys for the MKHSig' scheme (obtained from KeyGen_2) have the same distribution as the keys generated in the standard security game of MKHSig'.

In what follows we show that every Type 2 or Type 3 forgery produced by \mathcal{A} in Game 2 can be used by an adversary \mathcal{B} to win the security game for the single dataset scheme MKHSig' with a forgery of the same type for one of the Q_{Δ} instances of the single-dataset MKHSig' scheme that were generated during Game 2. Fix a dataset $\bar{\Delta}^*$ and let it be the k th dataset queried by \mathcal{A} in Game 2. The adversary \mathcal{B} proceeds as follows. It samples a random index $k^* \xleftarrow{\$} [Q_{\Delta}]$ (which represents its guess for k), and then simulates Game 2 to \mathcal{A} by choosing the keys of the MKHSig' on its own for all datasets $k \neq k^*$, whereas for the k^* th dataset \mathcal{B} does the simulation forwarding the queries to its challenger.

Let $(\mathcal{P}^*, \Delta^*, m^*, \sigma^*)$ be a valid Type 2 (resp. Type 3) forgery returned by \mathcal{A} , where $\sigma^* = (I^*, \sigma', \text{par}_{\Delta^*})$. If Δ^* is not the k^* th queried dataset, \mathcal{B} aborts. Otherwise, Δ^* is the k^* th queried dataset, i.e. $\Delta^* = \bar{\Delta}^*$, and the second entry of $\sigma^* = (I^*, \sigma', \text{par}_{\Delta^*})$ satisfies $\text{Ver}'(\mathcal{P}^*, \bar{\Delta}^*, m^*, \sigma^*) = 1$. The last equations imply that $(\mathcal{P}^*, \bar{\Delta}^*, m^*, \sigma^*)$ is a valid Type 2 (resp. Type 3) forgery against the MKHSig' scheme.

To conclude, we notice that if $\Pr[G_2(\mathcal{A})] > \epsilon$, then \mathcal{B} has at least ϵ/Q_{Δ} advantage in breaking the security of MKHSig', where the factor $1/Q_{\Delta} = \Pr[k = k^*]$ is the probability of \mathcal{B} guessing

correctly the index k^* of the database Δ^* used in by \mathcal{A} in the forgery.

Combining the results of the three lemmas above, we obtain the following bound:

$$\Pr [G(\mathcal{A})] \leq Q_{\text{id}} \cdot \text{Adv}_{F, \mathcal{D}}^{\text{PRF}}(\lambda) + Q_{\text{id}} \cdot \text{Adv}_{\text{NH.Sig}, \mathcal{F}}^{\text{UF-CMA}}(\lambda) + Q_{\Delta} \cdot \text{Adv}_{\text{MKHSig}, \mathcal{D}}^{\text{HomUF-CMA}}(\lambda)$$

which proves the statement of Theorem 2, since all the addends are negligible by assumption. \square

4.4 Generic transformations from weak to adaptive security

In this section, we present two generic transformations that turn weakly secure multi-key HA schemes into adaptively secure ones.

Our first transformation holds in the standard model and works for schemes in which the size of the tag space \mathcal{T} is at most polynomial in the security parameter. The second transformation avoids the limitation on the size of the tag space (i.e. it works for schemes that can handle tags that are binary strings of virtually arbitrary length), but holds in the random oracle model. Both transformations rely on the notion of homomorphic trapdoor functions (HTDFs) introduced in [4].

In what follows, we first recall the notion of HTDFs and subsequently describe our transformations. Moreover, it is straightforward to see that both transformations extend to the case of multi-key homomorphic signatures.

Homomorphic trapdoor functions [4]. The concept of homomorphic trapdoor functions was recently introduced by Gorbunov *et al.* in [4]. Intuitively, a homomorphic trapdoor function f is defined by a public key pk which allows computing $v = f_{\text{pk},x}(u)$ on input u and index x . The homomorphic property enables anyone with knowledge of triples $\{u_i, x_i, v_i = f_{\text{pk},x_i}(u_i)\}_{i \in [t]}$ and a t -input function g to generate values u^*, v^* such that $v^* = f_{\text{pk},g(x_1, \dots, x_t)}(u^*)$. More formally:

Definition 8 (Homomorphic trapdoor functions): A homomorphic trapdoor function (HTDF) is a tuple of five PPT algorithms

$\text{HTDF} = (\text{HTDF}.\text{KeyGen}, f, \text{Inv}, \text{HTDF}.\text{Eval}^{\text{in}}, \text{HTDF}.\text{Eval}^{\text{out}})$ which are defined as follows:

$\text{HTDF}.\text{KeyGen}(1^\lambda)$. The key generation algorithm takes as input a security parameter λ and outputs a key-pair (pk, sk) . The public key implicitly defines the index space \mathcal{X} , the input space \mathcal{U} , the output space \mathcal{V} and some efficiently samplable input distribution $\mathcal{D}_{\mathcal{U}}$ over \mathcal{U} . Testing membership in the sets $\mathcal{U}, \mathcal{V}, \mathcal{X}$ and sampling uniformly at random in \mathcal{V} are required to be efficiently doable.

$f_{\text{pk},x}(u)$. This is a deterministic function, indexed by pk and $x \in \mathcal{X}$. It takes as input an element $u \in \mathcal{U}$ and outputs an element $v \in \mathcal{V}$. Note that the function f is not required to be injective.

$\text{Inv}_{\text{sk},x}(v)$. This is a probabilistic algorithm which is indexed by sk and $x \in \mathcal{X}$. It takes as input an element $v \in \mathcal{V}$ and outputs an element $u \in \mathcal{U}$.

$\text{HTDF}.\text{Eval}_{\text{pk}}^{\text{in}}(g, (x_1, u_1), \dots, (x_t, u_t))$. The input evaluation algorithm takes as input a function $g: \mathcal{X}^t \rightarrow \mathcal{X}$ along with a collection of pairs $\{(x_i, u_i)\}_{i=1}^t \in \mathcal{X} \times \mathcal{U}$ and outputs a value $u^* \in \mathcal{U}$.

$\text{HTDF}.\text{Eval}_{\text{pk}}^{\text{out}}(g, v_1, \dots, v_t)$. The output evaluation algorithm takes as input a function $g: \mathcal{X}^t \rightarrow \mathcal{X}$ and elements $\{v_i\}_{i=1}^t \in \mathcal{V}$, and outputs a value $v^* \in \mathcal{V}$.

Additionally, an HTDF is required to satisfy the properties of correctness, the distributional equivalence of inversion and security, as described below.

The correctness of evaluation procedure: Let $(\text{pk}, \text{sk}) \leftarrow \text{HTDF}.\text{KeyGen}(1^\lambda)$, $x_1, \dots, x_t \in \mathcal{X}$, $g: \mathcal{X}^t \rightarrow \mathcal{X}$ and $y = g(x_1, \dots, x_t)$. For all $i \in [t]$, let $u_i \in \mathcal{U}$ and $v_i \in \mathcal{V}$ such that $v_i = f_{\text{pk},x_i}(u_i)$. If $u^* := \text{HTDF}.\text{Eval}_{\text{pk}}^{\text{in}}(g, (x_1, u_1), \dots, (x_t, u_t))$ and

$v^* := \text{HTDF}.\text{Eval}_{\text{pk}}^{\text{out}}(g, v_1, \dots, v_t)$, then it must be $u^* \in \mathcal{U}$ and $f_{\text{pk},y}(u^*) = v^*$.

Distributional equivalence of inversion: Let $(\text{pk}, \text{sk}) \leftarrow \text{HTDF}.\text{KeyGen}(1^\lambda)$, $x \in \mathcal{X}$, $u \xleftarrow{\$} \mathcal{D}_{\mathcal{U}}$, $v = f_{\text{pk},x}(u)$, $v' \xleftarrow{\$} \mathcal{V}$ and $u' \xleftarrow{\$} \text{Inv}_{\text{sk},x}(v')$. Then, we require that

$$(\text{pk}, \text{sk}, x, u, v) \stackrel{\text{stat}}{\approx} (\text{pk}, \text{sk}, x, u', v')$$

HTDF security. For security, HTDFs are required to satisfy a property similar to claw-freeness. Intuitively, it should be difficult for any adversary to come up with $u, u' \in \mathcal{U}$ and $x, x' \in \mathcal{X}$ such that $x \neq x'$ and $f_{\text{pk},x}(u) = f'_{\text{pk},x'}(u')$. More formally, for any PPT adversary \mathcal{A} , it holds

$$\Pr \left[\begin{array}{l} f_{\text{pk},x}(u) = f'_{\text{pk},x'}(u') \\ u, u' \in \mathcal{U}, \quad x, x' \in \mathcal{X}, \quad x \neq x' \end{array} \middle| (\text{pk}, \text{sk}) \leftarrow \text{HTDF}.\text{KeyGen}(1^\lambda) \right. \\ \left. (u, u', x, x') \leftarrow \mathcal{A}(1^\lambda, \text{pk}) \right] \leq \text{negl}(\lambda).$$

Our transformation for polynomially-large tag spaces. Before showing the transformation in details, we provide a general intuition of the approach. First, we assume to start with a weakly-secure multi-key HA scheme MKHAut' and a homomorphic trapdoor function HTDF . The idea of the transformation is that HTDFs are essentially *homomorphic chameleon hash functions*. Hence, one can follow the standard hash-and-sign approach based on chameleon hashes, which is known to provide a weak-to-adaptive security transformation for digital signatures. To authenticate a message m with the new adaptive secure scheme MKHAut one chooses $u \xleftarrow{\$} \mathcal{D}_{\mathcal{U}}$ from the input distribution $\mathcal{D}_{\mathcal{U}}$ of F , computes $v \leftarrow f_{\text{pk},m}(u)$ and finally authenticates v using MKHAut' thus obtaining an authenticator σ . The final authenticator is $\sigma = (v, u, \bar{\sigma})$. To homomorphically evaluate a function $g: \mathcal{M}^t \rightarrow \mathcal{M}$ over such authenticators one first computes

$$u^* \leftarrow \text{HTDF}.\text{Eval}_{\text{pk}}^{\text{in}}(g, (m_1, u_1), \dots, (m_t, u_t)) \quad (24)$$

and

$$v^* \leftarrow \text{HTDF}.\text{Eval}_{\text{pk}}^{\text{out}}(g, v_1, \dots, v_t). \quad (25)$$

Next, the idea is to use MKHAut' to vouch for the correctness of v^* as the output of a function $g'(\dots) = \text{HTDF}.\text{Eval}_{\text{pk}}^{\text{out}}(g, \dots)$. Namely, one executes the homomorphic evaluation algorithm of MKHAut' over $\{\bar{\sigma}_i\}_{i \in [t]}$ with function g' . In terms of security, the intuition is that to break the security of the new scheme one has to either find a claw in the HTDF, or break the security of MKHAut' to cheat on a false $v' \neq v^*$. It is also worth observing that, due to the property of HTDFs, MKHAut' is essentially used to authenticate data – v – that is random and independent of the real message m .

In terms of efficiency, we note that instantiating our transformations with our scheme and the HTDF of [4] come with an overhead. Slightly more in detail, in the HTDF of [4], an output $v = f_{\text{pk},x}(u)$ for a bit x has size $\text{poly}(\lambda)$, and then authenticating v with our scheme means authenticating every bit of v which results in $\text{poly}(\lambda)$ authenticators, one for each bit. Although concretely such overhead may be significant, we stress that asymptotically succinctness is still preserved as our final authenticators are still of size $\text{poly}(\lambda)$.

We describe our transformation more formally below.

Our first transformation. Let

$$F = (\text{HTDF}.\text{KeyGen}, f, \text{Inv}, \text{HTDF}.\text{Eval}^{\text{in}}, \text{HTDF}.\text{Eval}^{\text{out}}) \quad (26)$$

be an HTDF with index space \mathcal{X} , input space \mathcal{U} (with distribution $\mathcal{D}_{\mathcal{U}}$), and output space \mathcal{V} . Let

$$\text{MKHAut}' = (\text{Setup}', \text{KeyGen}', \text{Auth}', \text{Eval}', \text{Ver}') \quad (27)$$

be a (weakly-secure) multi-key homomorphic authentication scheme. Without loss of generality, we will assume the message space of MKHAut' to be \mathcal{V} . We also assume that the tag space \mathcal{T} of MKHAut' is of polynomial size, i.e. $T := |\mathcal{T}| = \text{poly}(\lambda)$. Using these tools, we construct an adaptive secure HA

$$\text{MKHAut} = (\text{Setup}, \text{KeyGen}, \text{Auth}, \text{Eval}, \text{Ver}) \quad (28)$$

with message space $\mathcal{M} = \mathcal{X}$. The scheme works as follows:

Setup(1^λ). Generate $\text{pp}' \leftarrow \text{Setup}'(1^\lambda)$ and the keys $(\text{pk}_h, \text{sk}_h) \leftarrow \text{HTDF}.\text{KeyGen}(1^\lambda)$, and output $\text{pp} = (\text{pp}', \text{pk}_h)$.

KeyGen(pp). Generate $(\text{sk}', \text{ek}', \text{vk}') \leftarrow \text{KeyGen}'(\text{pp}')$, and output $\text{sk} = \text{sk}'$, $\text{ek} = \text{ek}'$, and $\text{vk} = \text{vk}'$.

Auth($\text{sk}, \Delta, \ell, \text{m}$). Sample $u \xleftarrow{\$} \mathcal{D}_{\mathcal{M}}$, compute $v = f_{\text{pk}_h, \text{m}}(u)$, and finally compute $\bar{\sigma} \leftarrow \text{Auth}'(\text{sk}', \Delta, \ell, v)$ and output $\sigma := (v, u, \text{m}, \bar{\sigma})$.

Eval($g, \{\sigma_i, \text{EKS}_i\}_{i \in [t]}$). Given the function $g: \mathcal{M}^t \rightarrow \mathcal{M}$, we define a corresponding function $g': \mathcal{V}^t \rightarrow \mathcal{V}$ as $g'(v_1, \dots, v_t) = \text{HTDF}.\text{Eval}_{\text{pk}_h}^{\text{out}}(g, v_1, \dots, v_t)$. For all $i \in [t]$, parse every authenticator $\sigma_i = (v_i, u_i, m_i, \bar{\sigma}_i)$. Then, the evaluation algorithm computes $\text{m} = g(m_1, \dots, m_t)$,

$v \leftarrow \text{HTDF}.\text{Eval}_{\text{pk}_h}^{\text{out}}(g, v_1, \dots, v_t)$,

$u \leftarrow \text{HTDF}.\text{Eval}_{\text{pk}_h}^{\text{in}}(g, (m_1, u_1), \dots, (m_t, u_t))$ and

$\bar{\sigma} \leftarrow \text{Eval}'(g', \{\bar{\sigma}_i, \text{EKS}_i\}_{i=1}^t)$. Finally, it outputs $\sigma = (v, u, \text{m}, \bar{\sigma})$.

Ver($\mathcal{P} := (g, \ell_1, \dots, \ell_t), \Delta, \{\text{vk}_{\text{id}}\}_{\text{id} \in \mathcal{P}}, \text{m}, \sigma$). Parse

$\sigma := (v, u, \text{m}, \bar{\sigma})$, and output 1 if both $f_{\text{pk}_h, \text{m}}(u) = v$ and

$\text{Ver}'(\mathcal{P}' := (g', \ell_1, \dots, \ell_t), \Delta, \{\text{vk}_{\text{id}}\}_{\text{id} \in \mathcal{P}}, v, \bar{\sigma}) = 1$ hold.

Otherwise, output 0.

It is easy to see that the correctness of MKHAut follows from the correctness of both the MKHAut' and the HTDF .

Theorem 3: If \mathbb{F} is a secure HTDF , MKHAut' is a weakly secure multi-key HA whose tag space \mathcal{T} is such that $|\mathcal{T}| = \text{poly}(\lambda)$, then the scheme MKHAut defined above is an adaptively secure multi-key HA.

The proof proceeds by contradiction. Namely, assuming the existence of an adversary \mathcal{A} which breaks the adaptive security of MKHAut we show that one can break either the security of the HTDF or the weak security of MKHAut' . To see this, let \mathcal{A} be an adversary that wins in the adaptive security game described in Definition 3, and let $(\mathcal{P}^* := (g^*, \ell_1^*, \dots, \ell_t^*), \Delta^*, \text{m}^*, \sigma^* := (v^*, u^*, \text{m}^*, \bar{\sigma}^*))$ be its forgery. We observe that σ^* can be of either one of the following types:

Type 1: As per Definition 3, the list L_{Δ^*} has never been initialised during the game.

Type 2: $\forall i \in [t]$ we have $(\ell_i^*, m_i) \in L_{\Delta^*}$, but $\text{m}^* \neq g^*(m_1, \dots, m_t)$. If we let $\{\sigma_i \leftarrow \text{Auth}(\text{sk}, \Delta^*, \ell_i^*, m_i)\}_{i \in [t]}$ be the authenticators given to \mathcal{A} during the game, and let

$$\begin{aligned} \hat{\sigma} &:= (\hat{v}, \hat{u}, \hat{\text{m}}, \hat{\bar{\sigma}}) \leftarrow \text{Eval}(g^*, \{\sigma_i, \text{EKS}_i\}_{i \in [t]}), \\ \text{m} &:= g^*(m_1, \dots, m_t) \end{aligned} \quad (29)$$

There are two sub-cases which can occur: either (a) $v^* \neq \hat{v}$, or (b) $v^* = \hat{v}$.

Type 3: As per Definition 3, there exists a label ℓ^* such that $(\ell^*, \cdot) \notin L_{\Delta^*}$.

We show that for every \mathcal{A} producing forgeries of Types 1, 2(a) or 3 there exists an adversary that breaks the weak security of MKHAut' , whereas for every \mathcal{A} producing a forgery of Type 2(b) it is possible to break the security of \mathbb{F} . More precisely, we provide a detailed reduction \mathcal{B} , which exploits any Type 1 or Type 2(a) forgery made by \mathcal{A} to break the security of the MKHAut' scheme. We then notice that a similar reduction \mathcal{B}' holds for Type 3

forgeries. Theorem 3 is proven by combining the results of Lemmas 6–8 given below.

In the following, we set $Q = \text{poly}(\lambda)$ to be the total number of queries made by \mathcal{A} , similarly $Q_{\text{id}}, Q_{\mathcal{T}}, Q_{\Delta}$ will be, respectively, the number of distinct identities, tags and datasets queried by \mathcal{A} during the authentication-query phase (note that $Q_{\text{id}}, Q_{\mathcal{T}}, Q_{\Delta} \leq Q$).

Lemma 6: If \mathcal{A} is an adversary that has a non-negligible advantage ϵ against the security of MKHAut by producing a Type 1 (respectively Type 2) forgery, then it is possible to build an adversary \mathcal{B} that has an advantage ϵ against the weak security of MKHAut' by producing a Type 1 (respectively Type 2) forgery.

Proof: We want here to build a reduction \mathcal{B} , which uses a Type 1 or 2(a) forgery by \mathcal{A} to break the security of MKHAut' . Let $\mathcal{T} := \{\tau_1, \dots, \tau_T\}$, where $T = \text{poly}(\lambda)$.

At the beginning \mathcal{B} samples $Q_{\text{id}} \cdot T \cdot Q_{\Delta}$ elements

$$\{v_{j,k}^i\}_{i \in [Q_{\text{id}}], j \in [T], k \in [Q_{\Delta}]} \xleftarrow{\$} \mathcal{V}$$

and chooses a pair of keys $(\text{pk}_h, \text{sk}_h)$ for \mathbb{F} . Following the weak-security definition, \mathcal{B} sends $\{(v_{j,k}^i, \tau_j)\}_{(i,j,k) \in [Q_{\text{id}}] \times [T] \times [Q_{\Delta}]}$ to its challenger \mathcal{C} and gets back the public parameters pp' of MKHAut' . Next, \mathcal{B} sets $\text{pp} = (\text{pp}', \text{pk}_h)$, and gives pp to \mathcal{A} . Afterwards, \mathcal{A} starts making authentication queries $(\Delta, \ell := (\text{id}, \tau), \text{m})$. For all such queries, let $\Delta_1, \dots, \Delta_{Q_{\Delta}}$ be all the distinct dataset names queried, in order, by \mathcal{A} . Similarly, let $\text{id}_1, \dots, \text{id}_{Q_{\text{id}}}$ be the distinct identities queried, in order, by \mathcal{A} in the whole game.

To every authentication query $(\Delta, \ell := (\text{id}, \tau), \text{m})$ such that $\Delta = \Delta_k$ (i.e. Δ is the k th queried dataset), $\text{id} = \text{id}_i$ (i.e. id is the i th queried identity) and $\tau = \tau_j \in \mathcal{T}$, \mathcal{B} answers as follows. First, it queries \mathcal{C} with the pair (id, Δ) and obtains the authenticators $\{\bar{\sigma}_{1,k}^i, \dots, \bar{\sigma}_{T,k}^i\}$ corresponding to the values $\{v_{1,k}^i, \dots, v_{T,k}^i\}$. If (id, Δ) were already queried to \mathcal{C} , then \mathcal{B} uses the values obtained earlier. Second, \mathcal{B} uses the inversion property of HTDF s to compute $u_{j,k}^i \leftarrow \text{Inv}_{\text{sk}_h, \text{m}}(v_{j,k}^i)$, and finally sends to \mathcal{A} the authenticator $\sigma := (v_{j,k}^i, u_{j,k}^i, \text{m}, \bar{\sigma}_{j,k}^i)$. By the distributional equivalence of inversion property of HTDF s this simulation is statistically close to the real game. For what regards verification queries, any of them from \mathcal{A} is sent to the challenger and the answer is then sent back to \mathcal{A} itself. When \mathcal{A} produces a forgery $(\mathcal{P}^* := (g^*, \ell_1^*, \dots, \ell_t^*), \Delta^*, \text{m}^*, \sigma^* := (v^*, u^*, \text{m}^*, \bar{\sigma}^*))$ \mathcal{B} proceeds as follows:

Type 1: If this is the case, it is easy to see that $(\mathcal{P}' := ((g^*)', \ell_1^*, \dots, \ell_t^*), \Delta^*, v^*, \sigma^*)$ is a Type 1 forgery for MKHAut' .

Type 2(a): Since in this case it holds $v^* \neq \hat{v}$, the tuple $(\mathcal{P}' := ((g^*)', \ell_1^*, \dots, \ell_t^*), \Delta^*, v^*, \sigma^*)$ is a Type 2 forgery for MKHAut' .

Therefore, if \mathcal{A} has a non-negligible probability ϵ of coming up with a Type 1 or Type 2(a) forgery, then \mathcal{B} has a probability ϵ of coming up with a Type 1 or Type 2(a) forgery against MKHAut' . \square

Lemma 7: If \mathcal{A} is an adversary that has a non-negligible advantage ϵ against the security of MKHAut by producing a Type 3 forgery, then it is possible to build an adversary \mathcal{B}' that has a non-negligible advantage $\epsilon / (Q_{\text{id}} \cdot T \cdot Q_{\Delta})$ against the weak security of MKHAut' by producing a Type 3 forgery.

Proof: The reduction \mathcal{B}' works similarly to the one in Lemma 6. The game starts with \mathcal{B}' sampling $Q_{\text{id}} \cdot T \cdot Q_{\Delta}$ elements $\{(v_{j,k}^i)\}_{i \in [Q_{\text{id}}], j \in [T], k \in [Q_{\Delta}]}$ $\xleftarrow{\$} \mathcal{V}$. Next \mathcal{B}' chooses a triple $(\bar{i}, \bar{j}, \bar{k}) \xleftarrow{\$} [Q_{\text{id}}] \times [T] \times [Q_{\Delta}]$, and generates keys $(\text{pk}_h, \text{sk}_h)$ for the HTDF function \mathbb{F} .

Next, \mathcal{B}' sends to its challenger \mathcal{C} the values

$$\{(v_{j,k}^i, \tau_j)\}_{(i,j,k) \in ([Q_{\text{id}}] \times [T] \times [Q_{\Delta}]) \setminus \{\bar{i}, \bar{j}, \bar{k}\}}$$

Basically, this is the same as in the proof of Lemma 6 except that \mathcal{B}' does not send any value related to the triple $(\bar{i}, \bar{j}, \bar{k})$. Intuitively, this can be interpreted as \mathcal{B}' betting on the dataset $\Delta_{\bar{k}}$ and the label $\ell = (\text{id}_{\bar{i}}, \tau_{\bar{j}})$ over which the adversary \mathcal{A} will make the Type 3 forgery.

\mathcal{B}' answers the queries of \mathcal{A} exactly as in Lemma 6 except if \mathcal{A} makes a query $(\Delta, \ell := (\text{id}, \tau), m)$ such that $\Delta = \Delta_{\bar{k}}$, $\text{id} = \text{id}_{\bar{i}}$ and $\tau = \tau_{\bar{j}}$. If this case occurs \mathcal{B}' aborts.

Let $(\mathcal{P}^* := (g^*, \ell_1^*, \dots, \ell_t^*), \Delta^*, m^*, \sigma^* := (v^*, u^*, m^*, \sigma^*))$ be \mathcal{A} 's forgery, and following our assumption, let it be of Type 3. Namely, there exists a label $\ell^* = (\text{id}^*, \tau^*) \in \{\ell_1^*, \dots, \ell_t^*\}$ such that $i \in [Q_{\text{id}}], j \in [Q_{\mathcal{T}}], k \in [Q_{\Delta}]$. If $\Delta^* \neq \Delta_{\bar{k}}$, or $\text{id}^* \neq \text{id}_{\bar{i}}$ or $\tau^* \neq \tau_{\bar{j}}$ then \mathcal{B}' aborts. Otherwise, it returns $(\mathcal{P}' := ((g^*)', \ell_1^*, \dots, \ell_t^*), \Delta^*, v^*, \sigma^*)$ as a forgery for MKHAut'.

Let bad be the event that \mathcal{B}' aborts in the simulation. If bad does not happen \mathcal{B}' has the same advantage as \mathcal{A} in producing a Type 3 forgery.

Therefore, what is left to show is an estimation of $\Pr[\neg \text{bad}]$. However, it is not hard to see that, before bad occurs, the random choice of $(\bar{i}, \bar{j}, \bar{k}) \xleftarrow{\$} [Q_{\text{id}}] \times [T] \times [Q_{\Delta}]$ is perfectly hidden to \mathcal{A} . Hence, we have that $\Pr[\neg \text{bad}] = (1/Q_{\text{id}} \cdot T \cdot Q_{\Delta})$ and thus \mathcal{B}' has an advantage $(\epsilon/Q_{\text{id}} \cdot T \cdot Q_{\Delta})$, i.e. non-negligible as $(Q_{\text{id}} \cdot T \cdot Q_{\Delta}) = \text{poly}(\lambda)$. \square

Lemma 8: If \mathcal{A} is an adversary that has a non-negligible advantage ϵ against the security of MKHAut by producing a Type 2(b), then it is possible to build an adversary \mathcal{B}'' which breaks the security of the HTDFF with non-negligible probability ϵ .

Proof: The algorithm \mathcal{B}'' takes as input a public key pk_h of the homomorphic trapdoor function F , and then starts simulating the adaptive security game to \mathcal{A} by sampling all the secrets on its own – i.e. \mathcal{B}'' runs Setup and KeyGen – and giving $\text{pp} = (\text{pp}', \text{pk}_h)$ to \mathcal{A} . It is easy to see that since \mathcal{B}'' has all the secret keys, it can perfectly simulate all queries to \mathcal{A} .

Let $(\mathcal{P}^* := (g^*, \ell_1^*, \dots, \ell_t^*), \Delta^*, m^*, \sigma^* := (v^*, u^*, m^*, \sigma^*))$ be \mathcal{A} 's forgery that, by the assumption in this lemma, is of Type 2(b). Namely, if we let $\{\sigma_i \leftarrow \text{Auth}(\text{sk}, \Delta^*, \ell_i^*, m_i)\}_{i \in [t]}$ be the authenticators given to \mathcal{A} in the experiment for dataset Δ^* and labels $\ell_1^*, \dots, \ell_t^*$, and parse each $\sigma_i := (v_i, u_i, m_i, \bar{\sigma}_i)$, it is the case that $m^* \neq \hat{m}$ and $v^* = \hat{v}$, where $\hat{v} \leftarrow \text{HTDFF.Eval}_{\text{pk}_h}^{\text{out}}(g, v_1, \dots, v_t)$ and $\hat{m} = g(m_1, \dots, m_t)$. Let $\hat{u} \leftarrow \text{HTDFF.Eval}_{\text{pk}_h}^{\text{in}}(g, (m_1, u_1), \dots, (m_t, u_t))$ and note that by the correctness of the HTDFF, we have that $f_{\text{pk}_h, \hat{m}}(\hat{u}) = \hat{v}$, where $\hat{m} = g(m_1, \dots, m_t)$. Moreover, since the forgery verifies correctly, we also have $f_{\text{pk}_h, m^*}(u^*) = v^*$.

One can see that \mathcal{B}'' has values $\hat{u}, u^* \in \mathcal{U}$, $m, \hat{m} \in \mathcal{X}$ such that $m \neq \hat{m}$ and $f_{\text{pk}_h, m^*}(u^*) = f_{\text{pk}_h, \hat{m}}(\hat{u})$, i.e. a tuple that breaks the security of the HTDFF.

Therefore, if \mathcal{A} has a non-negligible advantage ϵ in producing a forgery of Type 2(b), then \mathcal{B}'' has a non-negligible probability ϵ of breaking the security of the HTDFF. \square

Our transformation for arbitrary tag spaces. Here we present our second transformation for converting a multi-key homomorphic signature from weak to adaptive security, now without limitations on the size of the tag space. We present first the intuition behind this transformation; the formal detailed description will follow. The transformation is quite similar to the previous one; we assume again to start with a weakly-secure multi-key HA scheme MKHAut' and a homomorphic trapdoor function HTDFF. Also, however, we also use two hash functions $H: \{0, 1\}^* \rightarrow \{0, 1\}^n$ and $H': \{0, 1\}^* \rightarrow \{0, 1\}^{n'}$ that in the proof are modelled as random oracles. The first hash function is used to hash the dataset identifiers, i.e. to compute $H(\Delta) = \hat{\Delta}$, whereas the second one is used to obtain tags $\hat{\tau} = H'(\hat{\Delta}, \text{id}, \tau)$. This way the scheme MKHAut' is required to support dataset identifiers that are n -bits strings and tags that are n' -bits strings. Intuitively, modelling these hash

functions as random oracles 'breaks' the adaptivity of the adversary on the choice of these values.

Our transformation.

Let $F = (\text{HTDFF}, \text{KeyGen}, f, \text{Inv}, \text{HTDFF}, \text{Eval}^{\text{in}}, \text{HTDFF}, \text{Eval}^{\text{out}})$ be an HTDFF with index space \mathcal{X} , input space \mathcal{U} (with distribution $\mathcal{D}_{\mathcal{U}}$), and output space \mathcal{V} . Let $H: \{0, 1\}^* \rightarrow \{0, 1\}^n$, $H': \{0, 1\}^* \rightarrow \{0, 1\}^{n'}$ be two efficiently computable hash functions. Let $\text{MKHAut}' = (\text{Setup}', \text{KeyGen}', \text{Auth}', \text{Eval}', \text{Ver}')$ be a weakly-secure multi-key HA whose message space, without loss of generality, is \mathcal{V} . Moreover, MKHAut' has tag space $\mathcal{T}' = \{0, 1\}^{n'}$ and it supports dataset identifiers represented as binary strings of n bits. Using these tools we construct an adaptively secure HA scheme $\text{MKHAut} = (\text{Setup}, \text{KeyGen}, \text{Auth}, \text{Eval}, \text{Ver})$ with message space $\mathcal{M} = \mathcal{X}$, and tag space $\mathcal{T} = \{0, 1\}^*$. The scheme works as follows:

Setup(1^λ). Generate $\text{pp}' \leftarrow \text{Setup}'(1^\lambda)$ and the keys $(\text{pk}_h, \text{sk}_h) \leftarrow \text{HTDFF.KeyGen}(1^\lambda)$, choose two hash functions $H: \{0, 1\}^* \rightarrow \{0, 1\}^n$, $H': \{0, 1\}^* \rightarrow \{0, 1\}^{n'}$ and output $\text{pp} = (\text{pp}', \text{pk}_h, H, H')$.

KeyGen(pp). Generate $(\text{sk}', \text{ek}', \text{vk}') \leftarrow \text{KeyGen}'(\text{pp}')$, and output $\text{sk} = \text{sk}'$, $\text{ek} = \text{ek}'$ and $\text{vk} = \text{vk}'$.

Auth($\text{sk}, \Delta, \ell, m$). Evaluate $\hat{\Delta} \leftarrow H(\Delta)$ and $\hat{\tau} \leftarrow H'(\hat{\Delta}, (\text{id}, \tau))$ and set $\ell = (\text{id}, \hat{\tau})$. Sample $u \xleftarrow{\$} \mathcal{D}_{\mathcal{U}}$, compute $v \leftarrow f_{\text{pk}_h, m}(u)$ and finally, compute $\bar{\sigma} \leftarrow \text{Auth}'(\text{sk}', \hat{\Delta}, \hat{\tau}, v)$ and output $\sigma := (v, u, \bar{\sigma})$.

Eval($g, \{\sigma_i, \mathcal{H}_i\}_{i \in [t]}$). This algorithm is identical to one of the previous transformation.

Ver($\mathcal{P} := (g, \ell_1, \dots, \ell_t), \Delta, \{\text{vk}_{\text{id}}\}_{\text{id} \in \hat{\mathcal{P}}}, v, \sigma$). Evaluate $\hat{\Delta} = H(\Delta)$ and, for all $i \in [t]$ compute $\hat{\tau}_i = H'(\hat{\Delta}, \ell_i)$ and set $\ell_i = (\text{id}_i, \hat{\tau}_i)$. Parse $\sigma = (v, u, \bar{\sigma})$ and output 1 if both $f_{\text{pk}_h, m}(u) = v$ and $\text{Ver}'(\hat{\mathcal{P}} := (g', \hat{\ell}_1, \dots, \hat{\ell}_t), \hat{\Delta}, \{\text{vk}_{\text{id}}\}_{\text{id} \in \hat{\mathcal{P}}}, v, \sigma) = 1$. Else output 0.

As for the previous transformation, the correctness of the scheme immediately follows from the correctness of MKHAut', F and the hash functions.

Theorem 4: If F is a secure HTDFF, MKHAut' is a weakly secure multi-key HA and the hash functions H and H' are modelled as random oracles, then the HA MKHAut defined above is adaptively secure in the random oracle model.

Proof: The analysis of the different kinds of forgeries is the same as the one presented at the beginning of the proof of Theorem 3. However, in the case of an arbitrary tag space, the reduction needs to be adapted since it is not possible to make a query for every element of the tag space during the authentication-query phase (given that space now is exponentially large), contrarily to what we did in the proof of Theorem 3. We bypass this problem moving to the random oracle model.

As we did for the polynomially bounded case, we show here that for every \mathcal{A} producing forgeries of Types 1, 2(a) or 3 there exists an adversary that breaks the weak security of MKHAut', whereas for every \mathcal{A} producing a forgery of Type 2(b) it is possible to break the security of F . As before, for what regards MKHAut', we show how to build a reduction \mathcal{B} , which uses any Type 1 or Type 2(a) forgeries made by \mathcal{A} to break MKHAut' and that a similar reduction \mathcal{B}' can do the same with Type 3 forgeries.

In the following proofs, we again set $Q = \text{poly}(\lambda)$ to be the total number of queries made by \mathcal{A} and $Q_{\text{id}}, Q_{\mathcal{T}}, Q_{\Delta}$ to be, respectively, the number of distinct identities, tags and datasets which are queried by \mathcal{A} during the authentication-query phase. \square

Lemma 9: If \mathcal{A} is an adversary that has a non-negligible advantage ϵ against the security of MKHAut by producing a Type 1 (respectively Type 2(a)) forgery, then it is possible to build an adversary \mathcal{B} that has an advantage $\epsilon - Q/2^n$ against the weak security of $I = \{\text{id}_1, \dots, \text{id}_n\}$ by producing a Type 1 (respectively Type 2) forgery.

Proof: We want to build a reduction \mathcal{B} , which uses a Type 1 or Type 2(a) forgery by \mathcal{A} against MKHAut to break the security of MKHAut'. We recall that Q_{id} , Q_{Δ} , $Q_{\mathcal{T}}$, respectively, are the number of different identities, dataset, and labels which are queried by \mathcal{A} during the signing-query phase. For any $(i, k) \in [Q_{id}] \times [Q_{\Delta}]$, \mathcal{B} chooses random strings $\{\hat{\tau}_{j,k}^i\}_{j \in [Q_{\mathcal{T}}]} \xleftarrow{\$} \{0, 1\}^n$ and sets $T_{i,k} := \{\hat{\tau}_{i,k}^j : j \in [Q_{\mathcal{T}}]\}$; then, \mathcal{B} samples $Q_{id} \cdot Q_{\mathcal{T}} \cdot Q_{\Delta}$ elements $\{(v_{j,k}^i)_{(i,j,k) \in [Q_{id}] \times [Q_{\mathcal{T}]} \times [Q_{\Delta}]}\} \xleftarrow{\$} \mathcal{V}$. Moreover, the reduction chooses a pair of keys (pk_h, sk_h) for \mathbb{F} .

Following the definition of weak-security, the reduction \mathcal{B} sends $\{(v_{j,k}^i, \hat{\tau}_{j,k}^i)\}_{(i,j,k) \in [Q_{id}] \times [Q_{\mathcal{T}]} \times [Q_{\Delta}]}$ to the challenger \mathcal{C} and gets back the public parameters pp' of MKHAut. Next, \mathcal{B} sets $pp = (pp', pk_h, H, H')$, sends pp to \mathcal{A} and initialises a counter $cnt = 0$. When \mathcal{A} starts making its queries, \mathcal{B} answers as described below.

First of all, let $\Delta_1, \dots, \Delta_{Q_{\Delta}}$ be all the distinct dataset identifiers queried, in order, by \mathcal{A} in authentication queries; similarly, let $id_1, \dots, id_{Q_{id}}$ be the distinct identities queried, in order, by \mathcal{A} in authentication queries too.

For hash queries to H or H' , \mathcal{B} prepares random strings $\hat{\Delta}_1, \dots, \hat{\Delta}_{Q_{\Delta}} \xleftarrow{\$} \{0, 1\}^n$ and answer queries as follows:

$H(\Delta)$: If $\Delta = \Delta_k$, i.e. Δ is the k th distinct dataset identifier queried in authentication, then \mathcal{B} answers with $\hat{\Delta}_k$. Otherwise, \mathcal{B} samples a random $\hat{\Delta} \xleftarrow{\$} \{0, 1\}^n$ and replies with $\hat{\Delta}$.

$H'(\hat{\Delta}, id, \tau)$: If $\hat{\Delta} = \hat{\Delta}_k$ (i.e. $\hat{\Delta}$ is the k th one of the Q_{Δ} random strings initially chosen) and $id = id_i$ (i.e. id is the i th distinct identity queried in authentication), then \mathcal{B} increments $cnt \leftarrow cnt + 1$ and replies with $\hat{\tau}_{cnt,k}^i$. Otherwise, if $\hat{\Delta} \notin \{\hat{\Delta}_1, \dots, \hat{\Delta}_{Q_{\Delta}}\}$ or $id \notin \{id_1, \dots, id_{Q_{id}}\}$, then \mathcal{B} replies with a randomly chosen $\hat{\tau} \xleftarrow{\$} \{0, 1\}^n$.

We define bad' be the event for which \mathcal{A} makes a hash query $H'(\hat{\Delta}, id, \tau)$ such that $\hat{\Delta} = \hat{\Delta}_k$, for some k , but Δ_k was never returned earlier as the output of H . It is straightforward to see that $\Pr[bad'] \leq Q/2^n$ is negligible.

When \mathcal{A} makes an authentication query of the form $(\Delta, \ell := (id, \tau), m)$, \mathcal{B} answers as follows. First, makes the corresponding hash queries $\hat{\Delta} = H(\Delta)$ and $\hat{\tau} = H'(\hat{\Delta}, (id, \tau))$, if not already done. Then, notices that it must be $\hat{\Delta} = \hat{\Delta}_k$ for some $k \in [Q_{\Delta}]$ and $\hat{\tau} = \hat{\tau}_{j,k}^i$, for some proper i, j . Then \mathcal{B} can proceed as in the proof of Lemma 6 to compute $u_{j,k}^i$ and finally return $\sigma_{j,k}^i$. The remaining part of the simulation proceeds the same as in Lemma 6. \square

Lemma 10: If \mathcal{A} is an adversary that has a non-negligible advantage ϵ against the security of MKHAut by producing a Type 3 forgery, then it is possible to build an adversary \mathcal{B}' that has a non-negligible advantage $\epsilon/(Q_{id} \cdot Q_{\mathcal{T}} \cdot Q_{\Delta}) - Q/2^n$ against the weak security of MKHAut' by producing a Type 3 forgery.

Proof: The reduction \mathcal{B}' works similarly to the one in Lemma 7 except that \mathcal{B}' also simulates hash queries as in Lemma 9.

The game starts with \mathcal{B}' sampling $Q_{id} \cdot Q_{\mathcal{T}} \cdot Q_{\Delta}$ elements $v_{j,k}^i \xleftarrow{\$} \mathcal{V}$ and strings $\hat{\tau}_{j,k}^i \xleftarrow{\$} \{0, 1\}^n$ uniformly at random for all $i \in [Q_{id}], j \in [Q_{\mathcal{T}}], k \in [Q_{\Delta}]$. Next \mathcal{B}' chooses a triple $(\hat{i}, \hat{j}, \hat{k}) \xleftarrow{\$} [Q_{id}] \times [Q_{\mathcal{T}}] \times [Q_{\Delta}]$, generates keys (pk_h, sk_h) for the homomorphic trapdoor function \mathbb{F} , and sends to its challenger \mathcal{C} the values $\{(v_{j,k}^i, \hat{\tau}_{j,k}^i)\}_{(i,j,k) \in ([Q_{id}] \times [Q_{\mathcal{T}]} \times [Q_{\Delta}]) \setminus (\hat{i}, \hat{j}, \hat{k})}$. Basically, this is the same as in the proof of Lemma 9 except that \mathcal{B}' does not send any value related to the triple $(\hat{i}, \hat{j}, \hat{k})$. Intuitively, this can be interpreted as \mathcal{B}' betting on the dataset $\Delta_{\hat{k}}$ and the label $\ell = (id_{\hat{i}}, \tau_{\hat{j}})$ over which the adversary \mathcal{A} will make the Type 3 forgery.

\mathcal{B}' answers the queries of \mathcal{A} exactly as in Lemma 9 except if \mathcal{A} makes a query $(\Delta, \ell := (id, \tau), m)$ such that $H(\Delta) = \hat{\Delta}_{\hat{k}}$, $id = id_{\hat{i}}$ and $H'(\hat{\Delta}_{\hat{k}}, (id, \tau)) = \hat{\tau}_{\hat{j}, \hat{k}}^{\hat{i}}$. If this case occurs \mathcal{B}' aborts.

The remaining part of the simulation is basically the same as that of Lemma 7. Namely, as long as the guess of \mathcal{B}' is correct (which is the case with probability $1/(Q_{id} \cdot Q_{\mathcal{T}} \cdot Q_{\Delta})$) and there is no other Bad event in the simulation of hash queries (which is the case with non-negligible probability $Q/2^n$ as in Lemma 9), \mathcal{B}' can use \mathcal{A} 's Type 3 forgery to output a Type 3 forgery for the weakly-secure scheme. \square

Lemma 11: If \mathcal{A} is an adversary that has a non-negligible advantage ϵ against the security of MKHAut by producing a Type 2(b) forgery, then it is possible to build an adversary \mathcal{B}'' which breaks the security of the HTDF \mathbb{F} with non-negligible probability ϵ .

It is not hard to see that this case of the proof is basically identical to that in Lemma 8, and thus we omit the proof.

5 Our multi-key homomorphic MAC from OWFs

In this section, we describe our construction of a multi-key HA with private verification keys and supporting the evaluation of low-degree arithmetic circuits. More precisely, for a computation represented by an arithmetic circuit of degree d and involving inputs from n distinct identities, the final authenticator has size $\binom{n+d}{d}$, that is bounded by $\text{poly}(n)$ (for constant d) or by $\text{poly}(d)$ (for constant n). Essentially, the authenticators of our scheme grow with the degree of the circuit and the number of distinct users involved in the computation, whereas their size remains *independent* of the total number of inputs/users. This property is particularly desirable in contexts that involve a small set of users each of which contributes to several inputs.

Although our multi-key homomorphic MAC supports less expressive computations than our homomorphic signatures of Section 4, the scheme comes with two main benefits. First, it is based on a simple, general assumption: it relies on pseudo-random functions and thus is secure only assuming the existence of one-way functions (OWF). Second, the scheme is very intuitive and efficient: fresh MACs essentially consist only of two \mathbb{F}_p field elements (where p is a prime of λ bits) and an identity identifier; after evaluation, the authenticators consist of $\binom{n+d}{d}$ elements in \mathbb{F}_p , and homomorphic operations are simply additions and multiplications in the multi-variate polynomial ring $\mathbb{F}_p[X_1, \dots, X_n]$.

We describe the five algorithms of our scheme MKHMac below. We note that our solution is presented for a single data set only. However, since it admits labels that are arbitrarily long strings it is straight-forward to extend the scheme for handling multiple data sets: simply redefine each tag τ as consisting of two substrings $\tau = (\Delta, \tau')$, where Δ is the dataset name and τ' the actual tag.

Setup(1^λ). The setup algorithm generates a λ -bit prime p and let the message space be $\mathcal{M} := \mathbb{F}_p$. The set of identities is $ID = [C]$ for some integer bound $C \in \mathbb{N}$, while the tag space consists of arbitrary binary strings, i.e. $\mathcal{T} = \{0, 1\}^*$. The set \mathcal{F} of admissible functions is made up of all arithmetic circuits whose degree d is bounded by some polynomial in the security parameter. The setup algorithm outputs the public parameters pp which include the descriptions of $\mathcal{M}, ID, \mathcal{T}, \mathcal{F}$ as in Section 3, as well as the description of a PRF family $F: \mathcal{X} \times \{0, 1\}^* \rightarrow \mathbb{F}_p$ with seed space \mathcal{K} . The public parameters also define the authenticator space. Each authenticator σ consists of a pair (I, Y) where $I \subseteq ID$ and y is in the C -variate polynomial ring $\mathbb{F}_p[X_1, \dots, X_C]$. More precisely, if C is set up as a very large number (e.g. $C = 2^\lambda$) the polynomials y can still live in some smaller sub-rings of $\mathbb{F}_p[X_1, \dots, X_C]$.

KeyGen(pp). The key generation algorithm picks a random $x \xleftarrow{\$} \mathbb{F}_p^*$, a PRF seed $K \xleftarrow{\$} \mathcal{K}$, and outputs (sk, ek, vk) , where $sk = ek(x, K)$ and ek is void.

$\text{Auth}(\text{sk}_{\text{id}}, \ell, m)$. To authenticate the message m with the label $\ell = (\text{id}, \tau) \in \text{ID} \times \mathcal{T}$, the authentication algorithm parses $\text{sk}_{\text{id}} = (x_{\text{id}}, K_{\text{id}})$ and produces an authenticator $\sigma = (\text{I}, y)$, where $\text{I} \subseteq \text{ID}$ and $y \in \mathbb{F}_p[X_{\text{id}}] \subset \mathbb{F}_p[X_1, \dots, X_C]$. The set I is simply $\{\text{id}\}$. The polynomial y is a degree-1 polynomial in the variable X_{id} such that $y(0) = m$ and $y(x_{\text{id}}) = F(K_{\text{id}}, \ell)$. Note that the coefficients of $y(X_{\text{id}}) = y_0 + y_{\text{id}}X_{\text{id}} \in \mathbb{F}_p[X_{\text{id}}]$ can be efficiently computed with the knowledge of x_{id} by setting $y_0 = m$ and $y_{\text{id}} = (F(K_{\text{id}}, \ell) - m/x_{\text{id}})$. Moreover, Y can be compactly represented by only giving the coefficients $y_0, y_{\text{id}} \in \mathbb{F}_p$.

$\text{Eval}(f, \{\sigma_k\}_{k \in \text{ID}})$. Given a t -input arithmetic circuit $f: \mathbb{F}_p^t \rightarrow \mathbb{F}_p$, and the t authenticators $\{\sigma_k := (\text{I}_k, y_k)\}_k$, the evaluation algorithm outputs $\sigma = (\text{I}, y)$ obtained in the following way. First, it determines all the identities involved in the computation by setting $\text{I} = \bigcup_{k=1}^t \text{I}_k$. Then every polynomial y_k is ‘expanded’ into a polynomial \hat{y}_k , defined on the variables X_{id} corresponding to all the identities in I . This is done using the canonical embedding $\mathbb{F}_p[X_{\text{id}}: \text{id} \in \text{I}_k] \rightarrow \mathbb{F}_p[X_{\text{id}}: \text{id} \in \text{I}]$. It is worth noticing that the terms of y_k that depend on variables in $\text{I} \setminus \text{I}_k$ have coefficient 0. Next, let $\hat{f}: \mathbb{F}_p[X_{\text{id}}: \text{id} \in \text{I}]^t \rightarrow \mathbb{F}_p[X_{\text{id}}: \text{id} \in \text{I}]$ be the arithmetic circuit corresponding to the given f , i.e. \hat{f} is the same as f except that additions (resp. multiplications) in \mathbb{F}_p are replaced by additions (resp. multiplications) over the polynomial ring $\mathbb{F}_p[X_{\text{id}}: \text{id} \in \text{I}]$. Finally, y is obtained as $y = \hat{f}(\hat{y}_1, \dots, \hat{y}_t)$.

$\text{Ver}(\mathcal{P}, \{\text{vk}_{\text{id}}\}_{\text{id} \in \mathcal{P}}, m, \sigma)$. Let $\mathcal{P} = (f, \ell_1, \dots, \ell_t)$ be a labelled program, where f is a degree- d arithmetic circuit and every label is of the form $\ell_k = (\text{id}_k, \tau_k)$. Let $\sigma = (\text{I}, y)$ where $\text{I} = \{\text{id}_1, \dots, \text{id}_n\}$ with $\text{id}_i \neq \text{id}_j$ for $i \neq j$. The verification algorithm outputs 1 (accept) if and only if the authenticator satisfies the following three checks. Otherwise, it outputs 0 (reject)

$$\{\overline{\text{id}}_1, \dots, \overline{\text{id}}_n\} = \{\text{id}: \text{id} \in \mathcal{P}\}, \quad (30)$$

$$y(0, \dots, 0) = m, \quad (31)$$

$$y(x_{\overline{\text{id}}_1}, \dots, x_{\overline{\text{id}}_n}) = f(F(K_{\text{id}_1}, \ell_1), \dots, F(K_{\text{id}_t}, \ell_t)) \quad (32)$$

In the following sections remainder of the section, we discuss the efficiency and succinctness of our MKHMac (Section 5.2) and prove the correctness of our scheme (Section 5.3). Section 5.4 is devoted to concluding with the security analysis of the proposed MKHMac scheme.

5.1 Polynomials representations

In the sequel, we adopt the multi-index notation for multivariate polynomials. We consider an element $f \in \mathbb{F}_p[X_1, \dots, X_n]$ as a finite sum of terms of the form $c \cdot X^e$ where $c \in \mathbb{F}_p \setminus \{0\}$ is a non-zero coefficient, $e = (e_1, \dots, e_n) \in \mathbb{N}^n$ is an exponent vector and X^e is a shorthand for $X_1^{e_1}, \dots, X_n^{e_n}$.

Here we recall two main methods to represent multivariate polynomials: the *completely dense representation* and the *sparse representation*. Both representations require an ordering on the set of unknowns.

In the completely dense representation, a polynomial $f \in \mathbb{F}_p[X] = \mathbb{F}_p[X_1, \dots, X_n]$ is stored as a $(\prod_{i=1}^n d_i)$ -dimensional array with entries in \mathbb{F}_p , such that the entry at index $e = (e_1, \dots, e_n)$ is the coefficient of the term X^e in f . Setting $d > \max\deg(f)$, the size of the dense representation is $O(d^n \log p)$. The completely dense representation allows constant-time access to any coefficient of f , and using Kronecker substitution it is possible to easily apply all ‘fast’ methods known for dense univariate arithmetic.

The sparse representation (which corresponds to the default representation of polynomials in most computer algebra systems) of a multivariate polynomial $f \in \mathbb{F}_p[X_1, \dots, X_n]$ is a list of t coefficient-exponent couples $((c_1, e_1), \dots, (c_t, e_t)) \in (\mathbb{F}_p \times \mathbb{N}^n)^t$,

where t denotes the total number of distinct terms (monomials) in f . Setting $d > \max\deg(f)$, the size of the sparse representation is $O(n \log d)$. The sparse representation is the most compact polynomial representation, however, its small size entails higher operation costs. We recall the costs relevant to our scheme in Section 5.

5.2 Efficiency and succinctness

Below we discuss the efficiency and the succinctness of our MKHMac scheme. We focus first on succinctness and later we discuss the cost of running Eval on our authenticators (the cost of the other algorithms can be easily extracted from their description).

Succinctness. Let us consider the case of an authenticator σ which was obtained after running Eval on a circuit of degree d and taking inputs from n distinct identities. Note that every σ consists of two elements: a set $\text{I} \subseteq [\text{C}]$ and a polynomial $y \in \mathbb{F}_p[X_{\overline{\text{id}}}: \overline{\text{id}} \in \text{I}]$.

For the set I , it is easy to see that $|\text{I}| = n$ and I can be represented with $n \log C$ bits. The other part of the authenticator, y , is instead an n -variate polynomial in $\mathbb{F}_p[X_{\overline{\text{id}}_1}, \dots, X_{\overline{\text{id}}_n}]$ of degree d . Since the circuit degree is d , the maximum number of coefficients of y is $\binom{n+d}{d}$. More precisely, the total size of y depends on the

particular representation of the multi-variate polynomial y , which is chosen for implementation. In Section 5.1, we discuss some possible representations (further details can also be found in [36]). For example, when employing the *sparse representation* of polynomials, the size of y is bounded by $O(n \log d)$, where t is the number of non-zero coefficients in y (note that in the worst case, a polynomial $y \in \mathbb{F}_p[X_{\overline{\text{id}}_1}, \dots, X_{\overline{\text{id}}_n}]$ of degree d has at most

$t = \binom{n+d}{d}$ non-zero coefficients). Thus, setting $\log C \approx \log p \approx \lambda$,

we have that the size in bits of the authenticator σ is $|\sigma| \leq \lambda n + \lambda \binom{n+d}{d}$. Ignoring the security parameter, we have that $|\sigma| = \text{poly}(n)$ when d is constant, or $|\sigma| = \text{poly}(d)$ when n is constant.

Efficiency of Eval. In what follows, we discuss the cost of computing additions and multiplications over authenticators in our MKHMac scheme. Let $\sigma^{(i)} = (\text{I}^{(i)}, y^{(i)})$, to $i = 1, 2$ be two authenticators and consider the operation $\sigma = \text{Eval}(g, \sigma^{(1)}, \sigma^{(2)})$, where g is a fan-in-2 addition or multiplication gate. In both cases, the set I of identities of $\sigma = (\text{I}, y)$ is obtained as the union $\text{I} = \text{I}^{(1)} \cup \text{I}^{(2)}$ that can be computed in time $O(n)$, where $n = |\text{I}|$, assuming the sets $\text{I}^{(1)}, \text{I}^{(2)}$ are ordered. Regarding the computation of y from $y^{(1)}$ and $y^{(2)}$, one has to first embed each y_i into the ring $\mathbb{F}_p[X_{\overline{\text{id}}}: \overline{\text{id}} \in \text{I}]$, and then evaluate addition (resp. multiplication) over $\mathbb{F}_p[X_{\overline{\text{id}}}: \overline{\text{id}} \in \text{I}]$. Again, the costs of these operations depend on the adopted representation [36, 37].

Using the *sparse representation* of polynomials, expanding a y having t non-zero coefficients into an n -variate polynomial \hat{y} requires time at most $O(tn)$. To give an idea, such expansion indeed consists simply into inserting zeros in the correct positions of the exponent vectors of every non-zero monomial term of y . On the other hand, the complexity of operations (additions and multiplications) on polynomials using the *sparse representation* is usually estimated in terms of the number of monomial comparisons. The cost of such comparisons depends on the specific monomial ordering chosen, but is usually $O(n \log d)$, where n is the total number of variables and d is the maximum degree. Given two polynomials in sparse representation having t_1 and t_2 non-zero terms, respectively, addition costs about $O(t_1 t_2)$ monomial comparisons (if the monomial terms are stored in sorted order the cost of addition drops to $O(t_1 + t_2)$), while multiplication requires to add (merge) t_2 intermediate products of t_1 terms each, and can be performed with $O(t_1 t_2 \log t_2)$ monomial comparisons [37].

5.3 Correctness

Regarding the correctness of the proposed MKHMac, by construction, each fresh authenticator $\sigma = (I, y)$ of a message m labelled by $\ell := (id, \tau)$ is of the form $I = \{id\}$ and $y(X_{id}) := y_0 + y_{id}X_{id} = m + (F(K_{id}, \ell) - m/x_{id})X_{id}$. Thus the set I satisfies (16) since $\{id: id \in \mathcal{I}_\ell\} = \{id\}$. The two last verification checks (31) and (32) are automatically granted for the identity program \mathcal{F}_ℓ because $y(0) = y_0 = m$ and $y(x_{id}) = m + (F(K_{id}, \ell) - m/x_{id})x_{id} = F(K_{id}, \ell)$.

Evaluation correctness. The correctness of the Eval algorithm essentially comes from the structure of the multi-variate polynomial ring. Nevertheless, we will prove the evaluation correctness property for generic fan-in-2 addition and multiplication gates (by associativity, this is the inductive step to prove evaluation correctness for any circuit f).

For $i = 1, 2$, let $(\mathcal{P}_i, m_i, \sigma^{(i)})$ be a triple such that $\text{Ver}(\mathcal{P}_i, \{vk_{id}\}_{id \in \mathcal{P}_i}, m_i, \sigma^{(i)}) = 1$, where $\mathcal{P}_i = (f_i, \ell_1^{(i)}, \dots, \ell_{t_i}^{(i)})$. We need to show that for every fan-in-2 gate g we have $\text{Ver}(\mathcal{P}, \{vk_{id}\}_{id \in \mathcal{P}}, m, \sigma) = 1$, where $\mathcal{P} = g(\mathcal{P}_1, \mathcal{P}_2)$, $m = g(m_1, m_2)$ and $\sigma = \text{Eval}(g, \sigma^{(1)}, \sigma^{(2)})$.

To see this, let us further expand on the meaning of assuming that $\text{Ver}(\mathcal{P}_i, \{vk_{id}\}_{id \in \mathcal{P}_i}, m_i, \sigma^{(i)}) = 1$. For every $i = 1, 2$ let $I^{(i)} = \{\bar{id}_1^{(i)}, \dots, \bar{id}_{n_i}^{(i)}\}$, where all identities are distinct. Then verification means that the following equations hold:

$$\{\bar{id}_1^{(i)}, \dots, \bar{id}_{n_i}^{(i)}\} = \{id: id \in \mathcal{P}_i\} \quad (33)$$

$$y^{(i)}(0, \dots, 0) = m_i \quad (34)$$

$$y^{(i)}(\bar{id}_1^{(i)}, \dots, \bar{id}_{n_i}^{(i)}) = f_i(F(K_{id_1^{(i)}}, \ell_1^{(i)}), \dots, F(K_{id_{t_i}^{(i)}}, \ell_{t_i}^{(i)})) \quad (35)$$

$$\{\bar{id}_1^{(i)}, \dots, \bar{id}_{n_i}^{(i)}\} = \{id: id \in \mathcal{P}_i\} \quad (36)$$

$$y^{(i)}(0, \dots, 0) = m_i \quad (37)$$

$$y^{(i)}(\bar{id}_1^{(i)}, \dots, \bar{id}_{n_i}^{(i)}) = f_i(F(K_{id_1^{(i)}}, \ell_1^{(i)}), \dots, F(K_{id_{t_i}^{(i)}}, \ell_{t_i}^{(i)})) \quad (38)$$

Let $I = I^{(1)} \cup I^{(2)}$, it is trivial to check that (16) of the verification holds for I and $\mathcal{P} = g(\mathcal{P}_1, \mathcal{P}_2)$ since

$$I = I^{(1)} \cup I^{(2)} = \{id: id \in \mathcal{P}_1\} \cup \{id: id \in \mathcal{P}_2\} = \{id: id \in \mathcal{P}\},$$

where the second equality derives from (36).

Let $I = \{\bar{id}_1, \dots, \bar{id}_n\}$, denote by $\hat{y}^{(i)}$ the n -variate polynomial obtained from $y^{(i)}$ after the expansion process described in Eval. We need to prove that if (31) and (32) hold true for $y^{(i)}$ they also for its expansion $\hat{y}^{(i)}$ (concerning the same labelled program \mathcal{P}_i). Equation (37) and the fact that the expansion process does not affect the constant term of the polynomial trivially imply that $\hat{y}^{(i)}(0) = m_i$; whereas

$$\hat{y}^{(i)}(x_{\bar{id}_1}, \dots, x_{\bar{id}_n}) = y(x_{\bar{id}}, id \in I^{(i)}) \quad (39)$$

$$= f_i(F(K_{id_1^{(i)}}, \ell_1^{(i)}), \dots, F(K_{id_{t_i}^{(i)}}, \ell_{t_i}^{(i)})) \quad (40)$$

holds because of the polynomial expansion in Eval sets to 0 the coefficients of all terms that depend on the *new* variables.

Consider the n -variate polynomial $y = g(\hat{y}^{(1)}, \hat{y}^{(2)})$, where g is a generic gate. It is left to show that y verifies (31) and (32). In what follows, we prove correctness for a fan-in-2 *addition* gate; the

case of a fan-in-2 *multiplication* gate is analogous to this one. Check (31) is straightforward since

$$y(0, \dots, 0) = \hat{y}^{(1)}(0, \dots, 0) + \hat{y}^{(2)}(0, \dots, 0) = m_1 + m_2 = m \quad (41)$$

holds because of (37) and the property of addition over the polynomial ring $\mathbb{F}_p[X_{id} : id \in I]$. For check (32) we have

$$\begin{aligned} y(x_{\bar{id}_1}, \dots, x_{\bar{id}_n}) &= \hat{y}^{(1)}(x_{\bar{id}_1}, \dots, x_{\bar{id}_{n_1}}) \\ &\quad + \hat{y}^{(2)}(x_{\bar{id}_1}, \dots, x_{\bar{id}_{n_2}}) \\ &= y^{(1)}(x_{\bar{id}_1}, \dots, x_{\bar{id}_{n_1}}) \\ &\quad + y^{(2)}(x_{\bar{id}_1}, \dots, x_{\bar{id}_{n_2}}) \\ &= f_1(F(K_{id_1^{(1)}}, \ell_1^{(1)}), \dots, F(K_{id_{t_1}^{(1)}}, \ell_{t_1}^{(1)})) \\ &\quad + f_2(F(K_{id_1^{(2)}}, \ell_1^{(2)}), \dots, F(K_{id_{t_2}^{(2)}}, \ell_{t_2}^{(2)})) \\ &= f(F(K_{id_1}, \ell_1), \dots, F(K_{id_t}, \ell_t)) \end{aligned}$$

where the first equality follows by the construction of y and by the property of addition over $\mathbb{F}_p[X_{id} : id \in ID]$. The second equality is derived from the definition of polynomial expansion. The third equality follows immediately from (39), while the last one holds by the composition property of labelled programs, i.e. $\mathcal{P} = g(\mathcal{P}_1, \mathcal{P}_2)$.

5.4 Security

In what follows, we prove the security of our scheme against adversaries that make static corruptions, and produce forgeries according to the following restrictions.

Definition 9 (Weak forgery): Consider an execution of the experiment described in Section 3, $\text{HomUF} - \text{CMA}_{\mathcal{A}, \text{MKHMac}}(\lambda)$, where $(\mathcal{P}^*, \Delta^*, m^*, \sigma^*)$ is the tuple returned by the adversary at the end of the experiment, with $\mathcal{P}^* = (f^*, \ell_1^*, \dots, \ell_{t^*}^*)$, Δ^* a dataset identifier, $m^* \in \mathcal{M}$ and σ^* an authenticator. First, we say that the labelled program \mathcal{P}^* is well-defined on a list L if either one of the following two cases occurs:

1. There exists $i \in [t]$ such that $(\ell_i^*, \cdot) \notin L$ (i.e. \mathcal{A} never made a query with label ℓ_i^*), and $f^*(\{m_j\}_{(\ell_j, m_j) \in L \cup \{\bar{m}_j\}_{(\ell_j, \cdot) \notin L}})$ outputs the same value for all possible choices of $\bar{m}_j \in \mathcal{M}$.
2. L contains the tuples $(\ell_i^*, m_1), \dots, (\ell_i^*, m_t)$, for some messages m_1, \dots, m_t .

Then we say that $(\mathcal{P}^*, \Delta^*, m^*, \sigma^*)$ is a weak forgery if $\text{Ver}(\mathcal{P}^*, \Delta^*, \{vk_{id}\}_{id \in \mathcal{P}^*}, m^*, \sigma^*) = 1$ and either one of the following conditions is satisfied:

Type 1: L_{Δ^*} was not initialised during the game (i.e. Δ^* was never queried).

Type 2: \mathcal{P}^* is well-defined on L_{Δ^*} but $m^* \neq f^*(\{m_j\}_{(\ell_j, m_j) \in L_{\Delta^*}} \cup \{0\}_{\ell_j \notin L_{\Delta^*}})$ (i.e. m^* is not the correct output of \mathcal{P}^* when executed over previously authenticated messages).

Type 3: \mathcal{P}^* is not well-defined on L_{Δ^*} .

Although Definition 9 is weaker than our Definition 4, we stress that the above definition still protects the verifier from adversaries that try to cheat on the output of a computation. In more details, the difference between Definitions 9 and 4 is the following: if f^* has an input wire that has never been authenticated during the game (a Type 3 forgery in Definition 4), but f^* is *constant* concerning such input wire, then the above definition does not consider it a forgery. The intuitive reason why such a relaxed definition still makes sense is that ‘irrelevant’ inputs would not help in any case the adversary to cheat on the output of f^* . Definition 9 is essentially the multi-key version of the forgery definition used in previous (single-key) homomorphic MAC works, e.g. [6]. As discussed in [13] testing whether a program is well-defined may not be doable in

polynomial time in the most general case (i.e. every class of functions). However, in [15] it is shown how this can be done efficiently via a probabilistic test in the case of arithmetic circuits of a degree d over a finite field of the order p such that $d/p < 1/2$. Finally, we notice that for our MKHMac Type 1 forgery cannot occur as the scheme described here supports only one dataset (As noted at the beginning of the section the extension to multiple datasets is straightforward given that the tags are arbitrary strings. When such extension is applied, it is easy to see that Type 1 forgeries are Type 3 ones in the underlying scheme.)

Theorem 5: If F is a pseudo-random function, then the multi-key homomorphic MAC described in Section 5 is secure against adversaries that make static corruptions of keys and produce forgeries as in Definition 9.

Note that we can deal with the corruptions via our generic result of Proposition 1. Therefore, it is sufficient to prove the security against adversaries that make no corruptions. The proof is done via a chain of games following this (intuitive) path. First, we rule out adversaries that make Type 3 forgeries. Intuitively, this can be done as the adversary has never seen one of the inputs of the computation, and in particular, an input which can change the result. Second, we replace every PRF instance with a truly random function. Note that at this point, the security of the scheme is information theoretic. Third, we change the way to answer verification queries that are candidates to be Type 2 forgeries. Finally, we observe that in this last game, the adversary gains no information on the secret keys x_i and thus has a negligible probability of making a Type 2 forgery.

More formally, we prove Theorem 5 via a sequence of hybrid games that we will use to bound the probability of winning in the original security game described in Section 3. Let Game 0 be the security game between the adversary \mathcal{A} and the challenger \mathcal{C} described in Section 3, where forgeries follow Definition 9 and there are no corruptions. We define the following chain of hybrid games:

Game 1: This game proceeds as Game 0 except that the challenger answers with 0 (reject) all those verification queries (\mathcal{P}, m, y) that constitute a Type 3 forgery, i.e. such that they verify correctly and \mathcal{P} is not well-defined on the list L .

Game 2: This is the same as Game 1 except that every PRF instance $F(K_{id}, \cdot)$ is replaced with a truly random function $\mathcal{R}_{id}: \{0, 1\}^* \rightarrow \mathbb{F}_p$. Basically, for every label $\ell = (id, \tau)$ used in the game, while in Game 0 and Game 1 one generates the value r_ℓ as $r_\ell = F(K_{id}, \ell)$, in Game 2 the value r_ℓ is taken uniformly at random in \mathbb{F}_p .

Game 3: This is the same as Game 2 except that the challenger proceeds differently when answering those verification queries (\mathcal{P}, m, σ) such that $\mathcal{P} = (f, \ell_1, \dots, \ell_t)$ is well-defined on L . If one equation between (16) and (31) of the verification algorithm is not satisfied, the challenger immediately answers with 0 (reject), without executing the check of (32). Otherwise, if both (16) and (31) are satisfied, \mathcal{C} proceeds as follows. For every $j \in [t]$ such that $(\ell_j, \cdot) \notin L$, \mathcal{C} generates a ‘dummy’ polynomial $y_j \in \mathbb{F}_p[X_{id}: id \in \ell_j]$ (e.g. taking its two coefficients randomly), whereas for all $i \in [t]$ such that $(\ell_i, m_i) \in L$, \mathcal{C} fetches the authenticators y_i generated previously during the authentication queries phase of the game. The challenger then computes $\text{Eval}(f, \sigma_1, \dots, \sigma_t)$ to obtain the output $\hat{\sigma} = (\hat{I}, \hat{y})$, and answers the verification query as follows:

- If $y = \hat{y}$ (over $\mathbb{F}_p[X_{id}: id \in I = \hat{I}]$), output 1 (accept).
- If $y \neq \hat{y}$, compute the polynomial $z = y - \hat{y} \in \mathbb{F}_p[X_{id}: id \in I]$ and output 1 (accept) if $z(x_{id_1}, \dots, x_{id_n}) = 0 \pmod p$, and 0 (reject) otherwise.

Note that the immediate rejection of queries that do not satisfy (16) or (31) does not change the outcome of the verification queries, and thus does not affect the distribution of the game. This change is introduced only to ease the security analysis of the game.

Intuitively, this kind of queries can be immediately rejected without using any secret; in other words, the outcome of these queries reveals zero information about the secret keys. Also, the second change does not affect the distribution of Game 3 compared to Game 2. This is slightly trickier to observe and is addressed in Lemma 14.

Let $G_i(\mathcal{A})$ denote the event that Game i , executed with an adversary \mathcal{A} , outputs 1. The proof of Theorem 5 is obtained by showing that any PPT adversary has a negligible probability of distinguishing between every consecutive pair of games, as well as it has a negligible probability of winning in Game 3. In what follows, Q_{ver} (resp. Q_{id}) represents the number of verification (resp. new identity) queries to \mathcal{C} and is polynomial in the security parameter, i.e. $Q_{\text{ver}}, Q_{id} \sim \text{poly}(\lambda)$. Theorem 5 is proven in the following lemmas:

Lemma 12: For every PPT \mathcal{A} there is a PPT \mathcal{B} such that $|\Pr[G_0(\mathcal{A})] - \Pr[G_1(\mathcal{A})]| \leq Q_{\text{ver}}(\Pr[G_1(\mathcal{B})] + d/p)$, where Q_{ver} is the number of verification queries made by \mathcal{A} .

Proof: The only case in which Game 0 and Game 1 differ is when the adversary \mathcal{A} makes a verification query that is a valid Type 3 forgery. Let us call this event Bad , and note that $|\Pr[G_0(\mathcal{A})] - \Pr[G_1(\mathcal{A})]| = \Pr[G_0(\mathcal{A}) \wedge \text{Bad}]$. In what follows, we bound this probability by showing that for every adversary \mathcal{A} that wins in Game 0 by returning a Type 3 forgery there is an adversary \mathcal{B} that, in the same game, produces a Type 2 forgery with all but a negligible probability. The proof is an adaptation to the multi-key setting of Proposition 2 in [15].

For every \mathcal{A} , we construct \mathcal{B} as follows. First of all, \mathcal{B} chooses randomly an index $j^* \xleftarrow{\$} [Q_{\text{ver}}]$ as a guess that \mathcal{A} will make a Type 3 forgery in the j^* th verification query. Next, \mathcal{B} runs \mathcal{A} and answers its queries as follows: \mathcal{B} forwards to its challenger all the authentication queries of \mathcal{A} , and all verification queries (\mathcal{P}, m, σ) such that \mathcal{P} is well-defined on the list L . The j th verification query, where \mathcal{P} is not well-defined, is answered with 0 (reject) if $j \neq j^*$. Otherwise, if $j = j^*$ \mathcal{B} does the following. Parse $\mathcal{P} = (f, \ell_1, \dots, \ell_t)$; for all $i \in [t]$ such that $(\ell_i, m_i) \in L$, \mathcal{B} fetches the message m_i from the list and set $\tilde{m}_i = m_i$, and for all $i \in [t]$ such that $(\ell_i, \cdot) \notin L$, \mathcal{B} chooses random messages $\tilde{m}_i \xleftarrow{\$} \mathbb{F}_p$. Then \mathcal{B} computes $\tilde{m} = f(\tilde{m}_1, \dots, \tilde{m}_t)$ and tests if $\tilde{m} = m$. If this is the case, then \mathcal{B} aborts. Otherwise, for all $i \in [t]$, such that $(\ell_i, \cdot) \notin L$, \mathcal{B} queries (ℓ_i, \tilde{m}_i) to its challenger, and finally returns (\mathcal{P}, m, σ) as its forgery.

To complete the proof, we first observe that $\Pr[\tilde{m} = m] = d/p$ over the random choices of $\tilde{m}_i \xleftarrow{\$} \mathbb{F}_p$ by the Schwartz-Zippel lemma [38, 39], since f is not well defined and assuming that d is a bound on f 's degree. Then since p is taken $> 2^2$ and $d = \text{poly}(\lambda)$ it holds that d/p is negligible. Finally, if \mathcal{B} does not abort, \mathcal{A} returns a Type 3 forgery and the guess on j^* is correct, it is easy to see that the forgery returned by \mathcal{B} is of Type 2, i.e. \mathcal{B} would win in Game 1. Therefore, $\Pr[G_1(\mathcal{B})] > \Pr[G_0(\mathcal{A}) \wedge \text{Bad}] / Q_{\text{ver}} - d/p$, which concludes the proof. \square

Lemma 13: For every PPT \mathcal{A} there is a PPT \mathcal{D} such that $|\Pr[G_1(\mathcal{A})] - \Pr[G_2(\mathcal{A})]| \leq Q_{id} \cdot \text{Adv}_{F, \mathcal{D}}^{\text{PRF}}(\lambda)$, where Q_{id} is the number of distinct identities queried by \mathcal{A} through both authentication and verification queries.

The proof is rather simple and we only give an intuition on how it works. Given that the difference between Game 1 and Game 2 is solely in the generation of the values r_ℓ , the probability of \mathcal{A} winning is the same in the two games, a part from an additive factor that comes from distinguishing the PRF instance $F(K_{id}, \cdot)$ from a truly random function. Since there are at most Q_{id} distinct PRF instances that are used in the game, by a fairly standard hybrid argument, we have that the difference between the two games is bound by the $Q_{id} \cdot \text{Adv}_{F, \mathcal{D}}^{\text{PRF}}(\lambda)$ factor.

Lemma 14: For every PPT \mathcal{A} we have $\Pr [G_2(\mathcal{A})] = \Pr [G_3(\mathcal{A})]$.

Proof: To prove the lemma, we show that the outcome of every verification query made by \mathcal{A} is the same in both Game 2 and Game 3. Let us consider only those queries for which the change mentioned in Game 3 applies, i.e. those in which \mathcal{P} is well-defined on L .

As already pointed out in the description of Game 3, rejecting immediately queries that do not satisfy (16) or (31) does not introduce any change to the outcome of the game.

As before, let $I = \{i\hat{d}_1, \dots, i\hat{d}_n\}$ and $\hat{I} = \{i\hat{d}_1, \dots, i\hat{d}_n\}$. Consider queries that satisfy (16) and (31). By correctness of the scheme, the authenticating polynomial y computed by \mathcal{C} verifies correctly for \mathcal{P} . This in particular means that the following equation holds $\hat{y}(x_{i\hat{d}_1}, \dots, x_{i\hat{d}_n}) = f(F(K_{i\hat{d}_1}, \ell_1), \dots, F(K_{i\hat{d}_n}, \ell_n))$. Since checking (32) for y means checking $y(x_{i\hat{d}_1}, \dots, x_{i\hat{d}_n}) = f(F(K_{i\hat{d}_1}, \ell_1), \dots, F(K_{i\hat{d}_n}, \ell_n))$ and by hypothesis y satisfies (16), i.e. $\hat{I} = I$, it is easy to see that $z(x_{i\hat{d}_1}, \dots, x_{i\hat{d}_n}) = 0$ is an equivalent check as $z = y - \hat{y}$. \square

Lemma 15: For every PPT \mathcal{A} we have $\Pr [G_3(\mathcal{A})] \leq Q_{\text{Ver}} \cdot \frac{d}{p}$,

where Q_{Ver} is the number of verification queries made by \mathcal{A} and d the (maximal) degree of the arithmetic circuits queried by \mathcal{A} in verification.

Proof: To prove Lemma 15, we define a sequence of hybrid games departing from Game 3, that we call Game (3, 0), ..., Game (3, Q_{Ver}), where Q_{Ver} denotes the total number of verification queries performed by the adversary. For $i = 0$ to Q_{Ver} , define Game (3, i) as Game 3 except that the first i verification queries that are a candidate to be Type 2 forgeries (i.e. they satisfy all Type 2 requirements except that verification accepts) are answered with 0, whereas the remaining ($Q_{\text{Ver}} - i$) queries are answered as stated in Game 3. Clearly, Game (3, 0) is identical to Game 3, i.e. $\Pr [G_3(\mathcal{A})] = \Pr [G_{(3,0)}(\mathcal{A})]$, while

$$\Pr [G_{(3, Q_{\text{Ver}})}(\mathcal{A})] = 0 \quad (42)$$

since in Game (3, Q_{Ver}), the challenger replies 0 to any verification query.

To proceed with the proof, we analyse the difference between each consecutive pair of games. For every $i \in [Q_{\text{Ver}}]$, the only case in which Game (3, i) may differ from Game (3, $i - 1$) is in the i th verification query. Let Bad be the event in which the i th verification query submitted by \mathcal{A} is a *valid* Type 2 forgery; then $|\Pr [G_{(3, i-1)}(\mathcal{A})] - \Pr [G_{(3, i)}(\mathcal{A})]| \leq \Pr [\text{Bad}]$. In what follows, we bound the probability that Bad happens.

Notice that Bad occurs in those verification queries (\mathcal{P}, m, σ) with $\sigma = (I, y)$ that satisfy the Type 2 requirements, pass the verification (16) and (31), and for which (a) $y \neq \hat{y}$ and (b) $z(x_{i\hat{d}_1}, \dots, x_{i\hat{d}_n}) = y(x_{i\hat{d}_1}, \dots, x_{i\hat{d}_n}) - \hat{y}(x_{i\hat{d}_1}, \dots, x_{i\hat{d}_n}) = 0 \pmod{p}$, where $I = \{i\hat{d}_1, \dots, i\hat{d}_n\}$. It is easy to see that condition (a) is immediately implied by the fact that the query is a Type 2 candidate. Indeed, if $y = \hat{y}$ then the query would be accepted and could *not* be a Type 2 forgery. To see this, let $m = f(\{m_j\}_{(\ell_j, m_j) \in L_\Delta} \cup \{0\}_{\ell_j \notin L_\Delta})$. By definition of Type 2 forgery $m \neq \hat{m}$ and by the correctness of the scheme $\hat{y}_0 = \hat{m}$. Hence, combining the previous two equations with (31) we get $y_0 = m \neq \hat{m} = \hat{y}_0$ which implies $y \neq \hat{y}$.

Now, we claim that in the execution of Game (3, i), $\Pr [\text{Bad}] = \Pr [z(x_{i\hat{d}_1}, \dots, x_{i\hat{d}_n}) = 0] \leq d/p$.

If the values $x_{i\hat{d}}$ were randomly selected in \mathbb{F}_p at the moment of Bad, the bound follows immediately from the Schwartz–Zippel lemma [38, 39], using that z is a non-zero polynomial of degree at most d . It remains to show that the $x_{i\hat{d}}$'s can indeed be selected at the moment of Bad, and to do this we show that before the happening of an event Bad, the adversary has no information about

$x_{i\hat{d}}$, $i\hat{d} \in I$. First, we observe that every authentication query $(\ell = (i\hat{d}, \tau), m)$ does not reveal information about $x_{i\hat{d}}$: the only value of the authenticator σ which depends on $x_{i\hat{d}}$ is the degree-1 coefficient of the authenticator y given in reply to the query; however this coefficient is uniformly distributed in \mathbb{F}_p due to the randomness of r_ℓ . Second, all the previous $i - 1$ verification queries in Game (3, i) are answered *without* the use of any secret $x_{i\hat{d}}$, $i\hat{d} \in [C]$: they are either answered without checking condition (b), or are directly rejected (without evaluating (b)) as per definition of Game (3, i). This means that before submitting the i th Type 2 verification query, the adversary has gained *no* information about the $x_{i\hat{d}}$, $i\hat{d} \in [C]$, and thus about the secret $x_{i\hat{d}}$, $i\hat{d} \in I$. Hence, $\Pr [\text{Bad}] = d/p$, which implies that for every $i \in [Q_{\text{Ver}}]$ it holds that $|\Pr [G_{(3, i-1)}(\mathcal{A})] - \Pr [G_{(3, i)}(\mathcal{A})]| \leq d/p$. By summing up for all the Q_{Ver} hybrids, we get

$$\left| \Pr [G_3(\mathcal{A})] - \Pr [G_{(3, Q_{\text{Ver}})}(\mathcal{A})] \right| \leq Q_{\text{Ver}} \cdot \frac{d}{p}.$$

By also considering (42), the above formula gives the bound stated in Lemma 15.

Combining the results of the lemmas we obtain that for every PPT \mathcal{A} , it holds that

$$\text{Adv}_{\text{MKHAut}, \mathcal{A}}^{\text{HomUF-CMA}}(\lambda) \leq (Q_{\text{Ver}} + 1) \cdot Q_{\text{id}} \cdot \text{Adv}_{F, \mathcal{D}}^{\text{PRF}}(\lambda) + \frac{(Q_{\text{Ver}}^2 + 2 \cdot Q_{\text{Ver}}) \cdot d}{p}$$

which proves Theorem 5. \square

6 Conclusions

In this paper, we introduced the concept of multi-key HAs, a cryptographic primitive that enables an untrusted third party to execute a function f on data authenticated using different secret keys to obtain a value certifying the correctness of f 's result, which can be checked with knowledge of corresponding verification keys. In addition to providing suitable definitions, we also propose two constructions: one which is publicly verifiable and supports general Boolean circuits, and a second one that is secretly verifiable and supports low-degree arithmetic circuits.

Although our work does not address directly the problem of privacy, extensions of our results along this direction are possible, and we leave the details to future investigation. A first extension is defining a notion of context-hiding for multi-key HAs. Similarly to the single key setting, this property should guarantee that authenticators do not reveal non-trivial information about the computation's inputs. The second extension has to do with preventing the Cloud from learning the data over which it computes. In this case, we note that multi-key HAs can be executed on top of homomorphic encryption following an approach similar to that suggested in [40]. Finally, an interesting problem left open by our work is to find multi-key HA schemes where authenticators have size independent of the number of users involved in the computation.

7 Acknowledgments

This work was partially supported by the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement no 608743, the VR grant PRECIS no 621-2014-4845. This work was also partially supported by the COST Action IC1306 through an STSM grant to Elena Pagnin. The research of Dario Fiore and Luca Nizzardo is partially supported by the Spanish Ministry of Economy under project references TIN2015-70713-R (DEDETIS), RTC-2016-4930-7 (DataMantium), and by the Madrid Regional Government under project N-Greens (ref. S2013/ICE-2731). Dario Fiore was also supported by a Juan de la Cierva fellowship from the Spanish Ministry of Economy.

8 References

- [1] Gentry, C.: 'Fully homomorphic encryption using ideal lattices'. 41st ACM STOC, Bethesda, MD, USA, May/June 2009, pp. 169–178
- [2] Boneh, D., Freeman, D.M.: 'Homomorphic signatures for polynomial functions'. EUROCRYPT 2011, Tallinn, Estonia, May 2011 (LNCS, **6632**), pp. 149–168
- [3] Gennaro, R., Wichs, D.: 'Fully homomorphic message authenticators'. ASIACRYPT 2013, Part II, Bangalore, India, December 2013 (LNCS, **8270**), pp. 301–320. doi: 10.1007/978-3-642-42045-0_16
- [4] Gorbunov, S., Vaikuntanathan, V., Wichs, D.: 'Leveled fully homomorphic signatures from standard lattices'. 47th ACM STOC, Portland, OR, USA, June 2015, pp. 469–477
- [5] Johnson, R., Molnar, D., Song, D.X., et al.: 'Homomorphic signature schemes'. CT-RSA 2002, San Jose, CA, USA, February 2002 (LNCS, **2271**), pp. 244–262
- [6] Catalano, D., Fiore, D.: 'Practical homomorphic MACs for arithmetic circuits'. EUROCRYPT 2013, Athens, Greece, May 2013 (LNCS, **7881**), pp. 336–352. doi: 10.1007/978-3-642-38348-9_21
- [7] Backes, M., Fiore, D., Reischuk, R.M.: 'Verifiable delegation of computation on outsourced data'. ACM CCS 13, Berlin, Germany, November 2013, pp. 863–874
- [8] Micciancio, D., Regev, O.: 'Worst-case to average-case reductions based on Gaussian measures'. 45th FOCS, Rome, Italy, October 2004, pp. 372–381
- [9] Boneh, D., Freeman, D., Katz, J., et al.: 'Signing a linear subspace: signature schemes for network coding'. PKC 2009, Irvine, CA, USA, March 2009 (LNCS, **5443**), pp. 68–87
- [10] Catalano, D., Fiore, D., Warinschi, B.: 'Adaptive pseudo-free groups and applications'. EUROCRYPT 2011, Tallinn, Estonia, May 2011 (LNCS, **6632**), pp. 207–223
- [11] Catalano, D., Fiore, D., Warinschi, B.: 'Efficient network coding signatures in the standard model'. PKC 2012, Darmstadt, Germany, May 2012 (LNCS, **7293**), pp. 680–696
- [12] Catalano, D., Fiore, D., Nizzardo, L.: 'Programmable hash functions go private: constructions and applications to (homomorphic) signatures with shorter public keys'. CRYPTO 2015, Part II, Santa Barbara, CA, USA, August 2015 (LNCS, **9216**), pp. 254–274. doi: 10.1007/978-3-662-48000-7_13
- [13] Freeman, D.M.: 'Improved security for linearly homomorphic signatures: A generic framework'. PKC 2012, Darmstadt, Germany, May 2012 (LNCS, **7293**), pp. 697–714
- [14] Catalano, D., Fiore, D., Warinschi, B.: 'Homomorphic signatures with efficient verification for polynomial functions'. CRYPTO 2014, Part I, Santa Barbara, CA, USA, August 2014b (LNCS, **8616**), pp. 371–389. doi: 10.1007/978-3-662-44371-2_21
- [15] Catalano, D., Fiore, D., Gennaro, R., et al.: 'Generalizing homomorphic MACs for arithmetic circuits'. PKC 2014, Buenos Aires, Argentina, March 2014a, (LNCS, **8383**), pp. 538–555. doi: 10.1007/978-3-642-54631-0_31
- [16] Agrawal, S., Boneh, D., Boyen, X., et al.: 'Preventing pollution attacks in multi-source network coding'. PKC 2010, Paris, France, May 2010 (LNCS, **6056**), pp. 161–176
- [17] Benabbas, S., Gennaro, R., Vahlis, Y.: 'Verifiable delegation of computation over large datasets'. CRYPTO 2011, Santa Barbara, CA, USA, August 2011 (LNCS, **6841**), pp. 111–131
- [18] Chung, K.-M., Kalai, Y., Vadhan, S.P.: 'Improved delegation of computation using fully homomorphic encryption'. CRYPTO 2010, Santa Barbara, CA, USA, August 2010 (LNCS, **6223**), pp. 483–501
- [19] Fiore, D., Gennaro, R.: 'Publicly verifiable delegation of large polynomials and matrix computations, with applications'. ACM CCS 12, Raleigh, NC, USA, October 2012, pp. 501–512
- [20] Gennaro, R., Gentry, C., Parno, B.: 'Non-interactive verifiable computing: outsourcing computation to untrusted workers'. CRYPTO 2010, Santa Barbara, CA, USA, August 2010 (LNCS, **6223**), pp. 465–482
- [21] Goldwasser, S., Kalai, Y.T., Rothblum, G.N.: 'Delegating computation: interactive proofs for muggles'. 40th ACM STOC, Victoria, British Columbia, Canada, May 2008, pp. 113–122
- [22] Parno, B., Raykova, M., Vaikuntanathan, V.: 'How to delegate and verify in public: verifiable computation from attribute-based encryption'. TCC 2012, Taormina, Sicily, Italy, March 2012 (LNCS, **7194**), pp. 422–439
- [23] Bitansky, N., Canetti, R., Chiesa, A., et al.: 'Recursive composition and bootstrapping for SNARKS and proof-carrying data'. 45th ACM STOC, Palo Alto, CA, USA, June 2013, pp. 111–120
- [24] Choi, S.G., Katz, J., Kumaresan, R., et al.: 'Multi-client non-interactive verifiable computation'. TCC 2013, Tokyo, Japan, March 2013 (LNCS, **7785**), pp. 499–518. doi: 10.1007/978-3-642-36594-2_28
- [25] Gordon, S.D., Katz, J., Liu, F.-H., et al.: 'Multi-client verifiable computation with stronger security guarantees'. TCC 2015, Part II, Warsaw, Poland, March 2015 (LNCS, **9015**), pp. 144–168. doi: 10.1007/978-3-662-46497-7_6
- [26] Fiore, D., Mitrokotsa, A., Nizzardo, L., et al.: 'Multi-key homomorphic authenticators'. Int. Conf. on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, 2016, pp. 499–530
- [27] Dodis, Y., Ostrovsky, R., Reyzin, L., et al.: 'Fuzzy extractors: How to generate strong keys from biometrics and other noisy data', *SIAM J. Comput.*, 2008, **38**, (1), pp. 97–139
- [28] Ajtai, M.: 'Generating hard instances of lattice problems (extended abstract)'. 28th ACM STOC, Philadelphia, PA, USA, May 1996, pp. 99–108
- [29] Micciancio, D.: 'Almost perfect lattices, the covering radius problem, and applications to Ajtai's connection factor', *SIAM J. Comput.*, 2004, **34**, (1), pp. 118–169. doi: 10.1137/S0097539703433511. Preliminary version in STOC 2002
- [30] Micciancio, D., Peikert, C.: 'Hardness of SIS and LWE with small parameters'. CRYPTO 2013, Part I, Santa Barbara, CA, USA, August 2013 (LNCS, **8042**), pp. 21–39. doi: 10.1007/978-3-642-40041-4_2
- [31] Ajtai, M.: 'Generating hard instances of the short basis problem'. Automata, Languages and Programming, 1999, pp. 1–9
- [32] Alwen, J., Peikert, C.: 'Generating shorter bases for hard random lattices', *Theory Comput. Syst.*, 2011, **48**, (3), pp. 535–553
- [33] Gentry, C., Peikert, C., Vaikuntanathan, V.: 'Trapdoors for hard lattices and new cryptographic constructions'. 40th ACM STOC, Victoria, British Columbia, Canada, May 2008, pp. 197–206
- [34] Micciancio, D., Peikert, C.: 'Trapdoors for lattices: simpler, tighter, faster, smaller'. EUROCRYPT 2012, Cambridge, UK, April 2012 (LNCS, **7237**), pp. 700–718
- [35] Boneh, D., Gentry, C., Gorbunov, S., et al.: 'Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits'. EUROCRYPT 2014, Copenhagen, Denmark, May 2014 (LNCS, **8441**), pp. 533–556. doi:10.1007/978-3-642-55220-5_30
- [36] van der Hoeven, J., Lecerf, G.: 'On the bit-complexity of sparse polynomial and series multiplication', *J. Symb. Comput.*, 2013, **50**, pp. 227–254
- [37] Geddes, K.O., Czapor, S.R., Labahn, G.: '*Algorithms for computer Algebra*' (Kluwer Academic Publishers, Norwell, MA, USA, 1992). ISBN 0-7923-9259-0
- [38] Schwartz, J.T.: 'Fast probabilistic algorithms for verification of polynomial identities', *J. ACM*, 1980, **27**, pp. 701–717
- [39] Zippel, R.: 'Probabilistic algorithms for sparse polynomials', *E. W. Ng, EUROSM '79*, 1979, **72**, pp. 216–226
- [40] Fiore, D., Gennaro, R., Pastro, V.: 'Efficiently verifiable computation on encrypted data'. ACM CCS 14, Scottsdale, AZ, USA, November 2014, pp. 844–855