



HAL
open science

A Lattice-Based Enhanced Privacy ID

Nada El Kassem, Luís Fiolhais, Paulo Martins, Liqun Chen, Leonel Sousa

► **To cite this version:**

Nada El Kassem, Luís Fiolhais, Paulo Martins, Liqun Chen, Leonel Sousa. A Lattice-Based Enhanced Privacy ID. 13th IFIP International Conference on Information Security Theory and Practice (WISTP), Dec 2019, Paris, France. pp.15-31, <10.1007/978-3-030-41702-4_2>. <hal-03173905>

HAL Id: hal-03173905

<https://inria.hal.science/hal-03173905v1>

Submitted on 18 Mar 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

A Lattice-based Enhanced Privacy ID

Nada EL Kassem¹, Luís Fiolhais², Paulo Martins², Liqun Chen¹, and Leonel Sousa²

¹ University of Surrey, UK, {n.elkassem, liqun.chen@surrey.ac.uk}

² INESC-ID, Instituto Superior Técnico, Universidade de Lisboa, Portugal, {luis.azehas.fiolhais@tecnico.ulisboa.pt, paulo.sergio@ist.utl.pt, las@inesc-id.pt}

Abstract. The Enhanced Privacy ID (EPID) scheme is currently used for hardware enclave attestation by an increasingly large number of platforms that implement Intel Software Guard Extensions (SGX). However, the scheme currently deployed by Intel is supported on Elliptic Curve Cryptography (ECC), and will become insecure should a large quantum computer become available. As part of National Institute of Standards and Technology (NIST)’s effort for the standardisation of post-quantum cryptography, there has been a great boost in research on lattice-based cryptography. As this type of cryptography is more widely used, one expects that hardware platforms start integrating specific instructions that accelerate its execution. In this article, a new EPID scheme is proposed, supported on lattice primitives, that may benefit not only from future research developments in post-quantum cryptography, but also from instructions that may extend Intel’s Instruction Set Architecture (ISA) in the future. This paper presents a new security model for EPID in the Universal Composability (UC) framework. The proposed Lattice-based EPID (LEPID) scheme is proved secure under the new model. Experimentally compared with a closely related Lattice-based Direct Anonymous Attestation (DAA) (LDAA) scheme from related art, it is shown that the private-key size is reduced 1.5 times, and that signature and verification times are sped up up to 1.4 and 1.1 times, respectively, for the considered parameters, when LEPID is compared with LDAA. Moreover, the signature size compares favourably to LDAA for small and medium-sized communities.

1 Introduction

The Enhanced Privacy ID (EPID) scheme is a fundamental part of the security model underpinning Software Guard Extensions (SGX)’s functioning [9]. It gives the ability to attest that a hardware enclave was successfully established on an Intel platform.

EPID can be seen as Direct Anonymous Attestation (DAA) with different linkability requirements [5]. The DAA scheme was built having the Trusted Platform Module (TPM) standard in mind. In this context, the TPM holds a representation of the host machine state, and wishes to provide a verifier with a signature of the state representation, without revealing their identity. During

an offline phase, an issuer provisions the TPM and the host with membership credentials. Based on this cryptographic material, the TPM and the host jointly prove that they belong to the DAA community in zero-knowledge, while producing the above-mentioned signature. Unlike other privacy-preserving systems, like group signatures, DAA does not support the property of traceability, wherein a group manager can identify the signer from a given signature.

Alternatively, the DAA provides two approaches to prevent a malicious signer from abusing their anonymity. Firstly, when a private-key is leaked, anyone can check whether a specific DAA signature was created under this key or not. Secondly, two DAA signatures created by the same signer may or may not be linked from a verifier's point of view. The linkability is controlled by a parameter called *basename*. When the same *basename* is used by the same signer for two signatures, they are linked; otherwise they are not. However, there are situations where this model does not suffice to prevent malicious actions. For instance, should an attacker corrupt a TPM and obtain the private-key without ever publishing it, there is no way to revoke it. While this latter problem can be mitigated by having TPMs use the same *basename* whenever they access a certain service, this option removes the anonymity for all uses with the same *basename*.

EPID is a more general scheme than DAA and thus does not split signers into TPMs and hosts, but also targets the creation of anonymous signatures. An EPID scheme consists of an issuer, signers, verifiers and a revocation manager. Like with DAA, one can check whether a certain signature was generated by a leaked private-key. Nonetheless, the ability to link signatures with the same *basename* is removed. Instead, whenever a signer is corrupted, they may be revoked by including one of their signatures as part of a revocation list. As a result, EPID is capable of revoking corrupted signers from the system, even when their private-key is kept hidden, whilst providing maximum privacy for the platforms. Enhanced Privacy ID signatures can also be constructed on the top of group signatures that allow members of a group to anonymously sign messages on behalf of the group, with the added property that a group manager can revoke the credentials of a misbehaving or compromised group member.

A post-quantum EPID scheme has been proposed in [3] built on hash and pseudorandom functions. More concretely, the EPID credential corresponds to a hash-based signature generated by the issuer, and proofs-of-knowledge are constructed from the Multi-Party Computation (MPC) in the head technique from Ishai et al. [8]. While [3] achieves signature sizes in the order of MBs, execution times are not considered. The main goal of this article is not to outperform [8], but rather to ignite research on lattice-based EPID partially propelled by the National Institute of Standards and Technology (NIST)'s effort on post-quantum cryptography standardisation [16]. By basing our construction on lattices, future versions of EPID might leverage the research resulting from this standardisation process to improve their efficiency. Moreover, since post-quantum cryptography is still in its infancy, it might be useful for implementers to consider multiple security assumptions, to mitigate the effects of cryptanalysis against one of them.

Lattices have proven to be a flexible tool in constructing cryptographic schemes, with applications ranging from digital signatures to public-key encryption and zero-knowledge proofs, while offering post-quantum security [13,1,2]. One expects that as this type of cryptography matures, an increasing number of platforms exploiting EPID ship with accelerators for lattice-based constructs [15]. Herein, by building from a recently proposed DAA scheme [10], the range of cryptographic constructs supported by lattice-based cryptography is extended to EPID. The LEPID signature size compares favourably to Lattice-based DAA (LDAA) for small and medium-sized communities.

Organisation: The next section introduces the lattice-based hard problems and the two building blocks that support the proposed LEPID scheme, namely the LDAA scheme from [10] and the Zero Knowledge Proof of Knowledge (ZKPoK) of Ring-Learning With Errors (Ring-LWE) secrets from [2]. Section 3 presents a new security model for EPID in the UC framework. The novel LEPID scheme is proposed in Section 4 and proven secure in Section 5. The performance of the LEPID scheme is discussed in Section 6. Finally, Section 7 concludes the paper.

2 Preliminaries

Throughout this paper we will use the polynomial rings $\mathcal{R}_q = \mathbb{Z}_q[X]/\langle X^n + 1 \rangle$, where \mathbb{Z}_q is the quotient ring $\mathbb{Z}/q\mathbb{Z}$ and n a power of 2. We use names in bold, like \mathbf{a} , both to denote elements of \mathcal{R}_q and their coefficient embeddings in \mathbb{Z}_q^n . $\|\mathbf{a}\|_\infty$ represents the infinity norm of a polynomial \mathbf{a} , $\|\mathbf{a}\|_\infty = \max_i |a^i|$, and $\|\mathbf{a}\| = \sqrt{\sum_{i=1}^n (a^i)^2}$ where the a^i are the coefficients of \mathbf{a} . $\hat{A} = (\mathbf{a}_1, \dots, \mathbf{a}_m)$ denotes a vector where m is a positive integer and $\mathbf{a}_1, \dots, \mathbf{a}_m$ are polynomials. $\|\hat{A}\|_\infty$ denotes the infinity norm of \hat{A} , defined as $\|\hat{A}\|_\infty = \max_i \|\mathbf{a}_i\|_\infty$. B_{3d} represents the set of vectors $\mathbf{u} \in \{-1, 0, 1\}^{3d}$ having exactly d coordinates equal to -1, d coordinates equal to 0, and d coordinates equal to 1. We represent a challenge set by $\mathcal{C} = \{X^{c_v}, |c_v \in \{0, 1, \dots, 2n - 1\}\}$, where $\bar{\mathcal{C}}$ denotes the set of differences $\mathcal{C} - \mathcal{C}$ except 0. \mathcal{D}_s^h represents the discrete Gaussian distribution of standard deviation s , s.t. $\Pr_{\mathbf{x} \leftarrow \mathcal{D}_s^h} [\|\mathbf{x}\| > \sqrt{2}hs] \leq 2^{-h/4}$. We define the following rejection sampling algorithm from [14] to avoid the dependency of \mathbf{z} on the secret \mathbf{b} , $\text{rej}(\mathbf{z}, \mathbf{b}, \xi)$: Let $u \leftarrow [0, 1)$; if $u > 1/3 \exp\left(\frac{-2\langle \mathbf{z}, \mathbf{b} \rangle + \|\mathbf{b}\|^2}{2\xi^2}\right)$ return 0, else return 1, with ξ representing a standard deviation of some distribution.

Definition 1 (The Ring Short Integer Solution Problem (Ring-SIS _{n,m,q,β})) [17]). Given m uniformly random elements $\mathbf{a}_i \in \mathcal{R}_q$ defining a vector $\hat{A} = (\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m)$, find a nonzero vector of polynomials $\hat{Z} \in \mathcal{R}_q^m$ of norm $\|\hat{Z}\|_\infty \leq \beta$ such that: $f_{\hat{A}}(\hat{Z}) = \sum_{i \in [m]} \mathbf{a}_i z_i = \mathbf{0} \in \mathcal{R}_q$. The Ring Inhomogeneous Short Integer Solution (Ring-ISIS _{n,m,q,β}) problem asks to find \hat{Z} of norm $\|\hat{Z}\|_\infty \leq \beta$, and such that: $f_{\hat{A}}(\hat{Z}) = \mathbf{y} \in \mathcal{R}_q$ for some uniform random polynomial \mathbf{y} .

Definition 2 (The Ring Learning With Error Problem (Ring-LWE)) [18]). Let χ be an error distribution defined over \mathcal{R} and $\mathbf{s} \leftarrow \mathcal{R}_q$ a uniformly

random ring element, the Ring-LWE distribution $A_{s,\chi}$ over $\mathcal{R}_q \times \mathcal{R}_q$ is sampled by choosing $\mathbf{a} \in \mathcal{R}_q$ uniformly at random, randomly choosing the noise $\mathbf{e} \leftarrow \chi$ and outputting $(\mathbf{a}, \mathbf{b}) = (\mathbf{a}, \mathbf{sa} + \mathbf{e} \bmod q) \in \mathcal{R}_q \times \mathcal{R}_q$. Let \mathbf{u} be uniformly sampled from \mathcal{R}_q . The decision problem of Ring-LWE asks to distinguish between $(\mathbf{a}, \mathbf{b}) \leftarrow A_{s,\chi}$ and (\mathbf{a}, \mathbf{u}) for a uniformly sampled secret $\mathbf{s} \leftarrow \mathcal{R}_q$. The search Ring-LWE problem asks to return the secret vector $\mathbf{s} \in \mathcal{R}_q$ given a Ring-LWE sample $(\mathbf{a}, \mathbf{b}) \leftarrow A_{s,\chi}$.

2.1 Lattice-based Direct Anonymous Attestation

The DAA scheme proposed in [10] can be split at a high level into three parts. In a first part, a TPM-host pair with identifier $\text{id} = (\text{id}_1, \dots, \text{id}_\ell) \in \{0, 1\}^\ell$ joins a DAA community. This consists of the TPM sampling small $\hat{X}_t = (\mathbf{x}_1, \dots, \mathbf{x}_m) \in \mathcal{R}_q^m$, and sending $\mathbf{u}_t = \hat{A}_t \hat{X}_t$ to the issuer, where $\hat{A}_t \in \mathcal{R}_q^m$ is part of the issuer's public-key. A signature proof of knowledge based on [12], showing that \mathbf{u} is well formed is also sent, along with a link token that prevents two TPMs from having the same secret-key. Using its private-key, the issuer then samples small $\hat{X}_h = (x_{m+1}, \dots, x_{3m}) \in \mathcal{R}_q^{2m}$ such that $\hat{A}_h \hat{X}_h = \mathbf{u} - \mathbf{u}_t$, where $\hat{A}_h = [\hat{A}_T | \hat{A}_0 + \sum_{i=1}^l \text{id}_i \hat{A}_i] \in \mathcal{R}_q^{2m}$, and $\mathbf{u} \in \mathcal{R}_q$, $\hat{A}_T \in \mathcal{R}_q^m$ and $\hat{A}_i \in \mathcal{R}_q^m \forall i \in \{0, \dots, l\}$ are part of the issuer's public-key. The vector \hat{X}_h is sent back to the host. After this process, the TPM and host own small key-shares satisfying

$$[X_t | X_h][\hat{A}_t | \hat{A}_h] = \mathbf{u}. \quad (1)$$

In a second part, the TPM and the host jointly generate a signature with respect to a message μ . The signature corresponds to a tuple $(\text{nym}, \text{bsn}, \pi)$, where nym is a link token, bsn is the basenome, and π is a signature-based proof:

$$\pi = \text{SPK} \left\{ \text{public} := \{\text{pp}, \text{nym}, \text{bsn}\}, \text{witness} := \{\hat{X} = (\mathbf{x}_1, \dots, \mathbf{x}_{3m}), \text{id}, \mathbf{e}\} : \right. \\ \left. \mathbf{u} = \hat{X}[\hat{A}_t | \hat{A}_h] \bmod q \wedge \|\hat{X}\|_\infty \leq \beta \wedge \text{nym} = \mathcal{H}(\text{bsn})\mathbf{x}_1 + \mathbf{e} \bmod q \wedge \|\mathbf{e}\|_\infty \leq \beta \right\}(\mu)$$

demonstrating not only (1) but also that $\text{nym} = \mathcal{H}(\text{bsn})\mathbf{x}_1 + \mathbf{e} \bmod q$, where \mathcal{H} is a random oracle mapping bsn to a polynomial and \mathbf{e} is small.

A final part deals with signature verification. First, π is verified. Then, the verifier iterates over the list of revoked private-keys, consisting of the elements $\mathbf{x}_1^{(i)}$ of the $\hat{X}_t^{(i)}$ in (1) of the corrupt signers. In the case that $\|\text{nym} - \mathcal{H}(\text{bsn})\mathbf{x}_1^{(i)}\|_\infty$ is small, the signature has been generated by the i -th revoked user and is rejected. Similarly, two signatures $(\text{nym}, \text{bsn}, \pi)$ and $(\text{nym}', \text{bsn}, \pi')$ having the same basenome are linked when $\|\text{nym} - \text{nym}'\|_\infty$ is small.

2.2 Zero Knowledge Proof of the Ring-LWE Secrets

The technique presented in [2] will herein be used to modify the LDAA and support the more effective revocation method of EPID. This techniques allows

one to efficiently prove in zero-knowledge possession of \mathbf{s} and \mathbf{e} , with $2\mathbf{s}$ and $2\mathbf{e}$ being short, such that $2\mathbf{y} = 2\mathbf{a}\mathbf{s} + 2\mathbf{e}$, for public \mathbf{a} and \mathbf{y} . Random $\mathbf{r}_s, \mathbf{r}_e \leftarrow \mathcal{D}_s$ are initially produced, and $\mathbf{t} = \mathbf{a}\mathbf{r}_s + \mathbf{r}_e$ is computed. A challenge $c = H(\mathbf{t}) \in \{0, 1, \dots, 2n - 1\}$ is generated and $\mathbf{s}_s = \mathbf{r}_s + X^c\mathbf{s}$, $\mathbf{s}_e = \mathbf{r}_e + X^c\mathbf{e}$ are outputted in response with probability $P(\mathbf{s}_s, \mathbf{s}_e)$, where P is chosen in a way that prevents \mathbf{s}_s and \mathbf{s}_e from depending on the prover's secret inputs.

3 UC based Security Model for EPID

The security model for the DAA [7] has been modified by replacing linkability with a revocation interface, adding the signature revocation check from [6], and introducing other modifications that results in a new EPID security model in the Universal Composability (UC) framework. Our new security definition is given in the UC model with respect to an ideal functionality $\mathcal{F}_{\text{EPID}}^l$. In UC, an environment \mathcal{E} should not be able to distinguish with a non-negligible probability between two worlds: the real world, where each party in the EPID protocol Π executes its assigned part of the protocol and the network is controlled by an adversary \mathcal{A} that communicates with \mathcal{E} ; and the ideal world, in which all parties forward their inputs to $\mathcal{F}_{\text{EPID}}^l$, which internally performs all the required tasks and creates the party's outputs. A protocol Π is said to securely realise $\mathcal{F}_{\text{EPID}}^l$ if, for every adversary \mathcal{A} performing an attack in the real world, there is an ideal world adversary \mathcal{S} that performs the same attack in the ideal world.

An EPID scheme should satisfy: *i*) unforgeability, i.e. no adversary can output a valid signature on a message μ without knowing the signer's secret key; *ii*) correctness, i.e. honestly generated signatures are always valid; and *iii*) anonymity, i.e. even for a corrupt issuer, no adversary can tell whether two honestly generated signatures were produced by the same signer. The UC framework allows us to focus on the analysis of a single protocol instance with a globally unique session identifier sid . $\mathcal{F}_{\text{EPID}}^l$ uses session identifiers of the form $sid = (\mathcal{I}, sid')$ for some issuer \mathcal{I} and a unique string sid' . In the procedures, functions `CheckTtdHonest` and `CheckTtdCorrupt` are used that return '1' when a key belongs to a honest signer that has produced no signature, and when a key belongs to a corrupt user such that there is no signature simultaneously linking back to the inputted key and another one, respectively; and return '0' otherwise. We label the checks that are done by the ideal functionality in roman numerals.

$\mathcal{F}_{\text{EPID}}^l$ **Setup:** On input (SETUP, sid) from the issuer \mathcal{I} , $\mathcal{F}_{\text{EPID}}^l$ verifies that $(\mathcal{I}, sid') = sid$ and outputs (SETUP, sid) to \mathcal{S} . $\mathcal{F}_{\text{EPID}}^l$ receives from the simulator \mathcal{S} the algorithms `Kgen`, `sig`, `ver`, `identify` and `revoke`. These algorithms are responsible for generating keys for honest signers, creating signatures for honest signers, verifying the validity of signatures, checking whether a signature was generated by a given key, and updating the revocation lists respectively. $\mathcal{F}_{\text{EPID}}^l$ stores the algorithms, checks that the algorithms `ver`, `identify` and `revoke` are deterministic [`Check-I`], and outputs $(\text{SETUPDONE}, sid)$ to \mathcal{I} .

$\mathcal{F}_{\text{EPID}}^l$ **Join:**

1. JOIN REQUEST: On input (JOIN, $sid, jsid$) from a signer \mathcal{M}_i , create a join session $\langle jsid, \mathcal{M}_i, \text{request} \rangle$. Output (JOINSTART, $sid, jsid, \mathcal{M}_i$) to \mathcal{S} .
2. JOIN REQUEST DELIVERY: Proceed upon receiving delivery notification from \mathcal{S} by updating the session record to $\langle jsid, \mathcal{M}_i, \text{delivery} \rangle$. If \mathcal{I} or \mathcal{M}_i is honest and $\langle \mathcal{M}_i, \star, \star \rangle$ is already in Member List ML, output \perp [Check II]. Otherwise, output (JOINPROCEED, $sid, jsid, \mathcal{M}_i$) to \mathcal{I} .
3. JOIN PROCEED: Upon receiving an approval from \mathcal{I} , $\mathcal{F}_{\text{EPID}}^l$ updates the session record to $\langle jsid, sid, \mathcal{M}_i, \text{complete} \rangle$. Then it outputs (JOINCOMPLETE, $sid, jsid$) to \mathcal{S} .
4. KEY GENERATION: On input (JOINCOMPLETE, $sid, jsid, tsk$) from \mathcal{S} .
 - If the signer is honest, set $tsk = \perp$, else verify that the provided tsk is eligible by performing the following two checks that are described above: CheckTtdHonest(tsk)=1 [Check III]; CheckTtdCorrupt(tsk)=1 [Check IV].
 - Insert $\langle \mathcal{M}_i, tsk \rangle$ into Member List ML, and output JOINED.

$\mathcal{F}_{\text{EPID}}^l$ Sign:

1. SIGN REQUEST: On input (SIGN, $sid, ssid, \mathcal{M}_i, \mu, \mathbf{p}$) from the signer on a message μ with respect to \mathbf{p} , the ideal functionality aborts if \mathcal{I} is honest and no entry $\langle \mathcal{M}_i, \star \rangle$ exists in ML, else creates a sign session $\langle ssid, \mathcal{M}_i, \mu, \mathbf{p}, \text{request} \rangle$ and outputs (SIGNSTART, $sid, ssid, \mathcal{M}_i, l(\mu, \mathbf{p})$) to \mathcal{S} .
2. SIGN REQUEST DELIVERY: On input (SIGNSTART, $sid, ssid$) from \mathcal{S} , update the session to $\langle ssid, \mathcal{M}_i, \mu, \mathbf{p}, \text{delivered} \rangle$, and output (SIGNPROCEED, $sid, ssid, \mu, \mathbf{p}$) to \mathcal{M}_i .
3. SIGN PROCEED: On input (SIGNPROCEED, $sid, ssid$) from \mathcal{M}_i , $\mathcal{F}_{\text{EPID}}^l$ updates the records $\langle ssid, \mathcal{M}_i, \mu, \mathbf{p}, \text{delivered} \rangle$, and outputs (SIGNCOMPLETE, $sid, ssid, \text{KRL}, \text{SRL}$) to \mathcal{S} , where KRL and SRL represent the key and the signature revocation lists respectively.
4. SIGNATURE GENERATION: On input (SIGNCOMPLETE, $sid, ssid, \sigma, \text{KRL}, \text{SRL}$) from \mathcal{S} , if \mathcal{M}_i is honest then $\mathcal{F}_{\text{EPID}}^l$ will:
 - Ignore an adversary's signature σ , and generate the signature for a fresh or established tsk .
 - Check CheckTtdHonest(tsk)=1 [Check V], and store $\langle \mathcal{M}_i, tsk \rangle$ in DomainKeys.
 - Generate the signature $\sigma \leftarrow \text{sig}(tsk, \mu, \mathbf{p})$.
 - Check $\text{ver}(\sigma, \mu, \mathbf{p}, \text{KRL}, \text{SRL})=1$ [Check VI], and check $\text{identify}(\sigma, \mu, \mathbf{p}, tsk) = 1$ [Check VII].
 - Check that there is no signer other than \mathcal{M}_i with key tsk' registered in Members or DomainKeys such that $\text{identify}(\sigma, \mu, \mathbf{p}, tsk')=1$ [Check VIII].
 - For all $(\sigma^*, \mu^*, \mathbf{p}^*) \in \text{SRL}$, find all (tsk^*, \mathcal{M}^*) from Members and DomainKeys such that $\text{identify}(\sigma^*, \mu^*, \mathbf{p}^*, \star, tsk^*) = 1$
 - Check that no two distinct keys tsk^* trace back to σ^* .
 - Check that no pair (tsk^*, \mathcal{M}_i) was found.
 - If \mathcal{M}_i is honest, then store $\langle \sigma, \mu, \mathcal{M}_i, \mathbf{p} \rangle$ in Signed and output (SIGNATURE, $sid, ssid, \sigma, \text{KRL}, \text{SRL}$).

$\mathcal{F}_{\text{EPID}}^l$ **Verify**: On input (VERIFY, $sid, \mu, \mathbf{p}, \sigma, \text{KRL}, \text{SRL}$), from a party \mathcal{V} to check whether σ is a valid signature on a message μ with respect to \mathbf{p} , KRL and SRL, the ideal functionality does the following:

- Extract all pairs (tsk_i, \mathcal{M}_i) from the DomainKeys and ML, for which $\text{identify}(\sigma, \mu, \mathbf{p}, tsk_i) = 1$. Set $b = 0$ if any of the following holds:
 - More than one key tsk_i was found [Check IX].
 - \mathcal{I} is honest and no pair (tsk_i, \mathcal{M}_i) was found [Check X].
 - An honest \mathcal{M}_i was found, but no entry $\langle \star, \mu, \mathcal{M}_i, \mathbf{p} \rangle$ was found in Signed [Check XI].
 - There is a key $tsk^* \in \text{KRL}$, such that $\text{identify}(\sigma, \mu, \mathbf{p}, tsk^*) = 1$ and no pair (tsk, \mathcal{M}_i) for an honest \mathcal{M}_i was found [Check XII].
 - For matching tsk_i and $(\sigma^*, \mu^*, \mathbf{p}^*) \in \text{SRL}$, $\text{identify}(\sigma^*, \mu^*, \mathbf{p}^*, tsk_i) = 1$.
- If $b \neq 0$, set $b \leftarrow \text{ver}(\sigma, \mu, \mathbf{p}, \text{SRL}, \text{KRL})$. [Check XIII]
- Add $\langle \sigma, \mu, \mathbf{p}, \text{KRL}, \text{SRL}, b \rangle$ to VerResults, and output (VERIFIED, sid, b) to \mathcal{V} .

$\mathcal{F}_{\text{EPID}}^l$ **Revoke**: On input (tsk^*, KRL) , the ideal functionality replaces KRL with $\text{KRL} \cup tsk^*$. On input $(\sigma^*, \mu^*, \text{SRL})$, the ideal functionality replaces SRL with $\text{SRL} \cup \sigma^*$ after verifying σ^* .

4 The Proposed LEPID Scheme

The DAA scheme proposed in [10] is herein modified so as to support the security model described in Section 3. We give a general overview of the proposed Lattice-based EPID (LEPID) scheme in Subsection 4.1 before proceeding with the details in Subsection 4.2.

4.1 High Level Description of the LEPID Scheme

The first part of the DAA protocol described in Subsection 2.1 is herein mirrored, with the exception that the TPM and the host are fused into a single signer. In particular, the issuer makes one further polynomial \mathbf{b} available in Procedure 1. When requesting to join a DAA community in Procedure 2, the signer with identifier $\text{id} = (\text{id}_1, \dots, \text{id}_\ell) \in \{0, 1\}^\ell$ samples a small $\hat{X}_t = (\mathbf{x}_1, \dots, \mathbf{x}_{m+1}) \in \mathcal{R}_q^{m+1}$ and sends $\mathbf{u}_t = [b|\hat{A}_\mathcal{I}]\hat{X}_t \bmod q$ to the issuer, along with a link token $\text{nym}_\mathcal{I} = \mathcal{H}(\text{bsn}_\mathcal{I})\mathbf{x}_1 + \mathbf{e}_\mathcal{I}$ and a zero-knowledge proof $\pi_{\mathbf{u}_t}$ from [10] showing that \mathbf{u}_t is well formed. Upon receiving this message, the issuer uses $\text{nym}_\mathcal{I}$ to check that no other signer has the same \mathbf{x}_1 , verifies $\pi_{\mathbf{u}_t}$ and samples small $\hat{X}_h = [\hat{X}_{h_1}|\hat{X}_{h_2}] = (\mathbf{y}_2, \dots, \mathbf{y}_{2m+1}) \in \mathcal{R}_q^m \times \mathcal{R}_q^m$ such that $\hat{A}_h\hat{X}_h = \mathbf{u} - \mathbf{u}_t \bmod q$, with $\hat{A}_h = [\hat{A}_\mathcal{I}|\hat{A}_0 + \sum_{i=1}^l \text{id}_i\hat{A}_i] \in \mathcal{R}_q^{2m}$. \hat{X}_h is sent back to the signer, that updates their key as $\hat{X} = (\mathbf{x}_1, \forall_{i=(2, \dots, m+1)} \mathbf{x}_i := \mathbf{x}_i + \mathbf{y}_i, \forall_{i=(m+2, \dots, 2m+1)} \mathbf{x}_i := \mathbf{y}_i)$ in Procedure 3.

Signatures are generated in Procedures 4 and 5 as in Subsection 2.1 for the DAA, but the basename is always chosen at random, generating link tokens $\text{nym} = \mathbf{p}\mathbf{x}_1 + \mathbf{e} \bmod q$ for a uniformly random \mathbf{p} , and the proof-of-knowledge π

is as described in the Appendix of the full version of this paper [11]. In particular, this allows one to maintain linkability in the case of leaked private-keys, whilst maintaining full anonymity. In addition, when signing a message, the signer is presented with a list of signatures from revoked users and proves in zero-knowledge that their underlying \mathbf{x}_1 was not used to produce any of those signatures. We achieve this by firstly randomising the $(\text{nym}_i^* = \mathbf{p}_i^* \mathbf{f}_i + \mathbf{l}_i, \mathbf{p}_i^*)$ pairs from the list of revoked signatures, where \mathbf{f}_i corresponds to the \mathbf{x}_1 polynomial of the i -th revoked user and \mathbf{l}_i has small norm, as

$$\mathbf{d}_i = \text{nym}_i^* \mathbf{q}_i + \mathbf{l}_i''' \quad (2)$$

$$\mathbf{o}_i = \mathbf{p}_i^* \mathbf{q}_i + \mathbf{l}_i' \quad (3)$$

for small \mathbf{q}_i , \mathbf{l}_i''' and \mathbf{l}_i' sampled from a Gaussian distribution. Note that $\mathbf{d}_i = \mathbf{o}_i \cdot \mathbf{f}_i + \mathbf{e}_i$ for a small \mathbf{e}_i . The signature includes not only \mathbf{d}_i and \mathbf{o}_i , but also $\mathbf{k}_i = \mathbf{o}_i \mathbf{x}_1 + \mathbf{l}_i''$ along with a zero-knowledge proof of the construction of \mathbf{d}_i , \mathbf{o}_i and \mathbf{k}_i . This zero-knowledge proof is an adaptation of the one described in Subsection 2.2, the details of which can be found in the Appendix of the full version of this paper [11].

Signature verification in Procedure 6 is similar to that of the DAA, with the difference that now the proof of the shape of \mathbf{d}_i , \mathbf{o}_i and \mathbf{k}_i is verified, and the norm of $\mathbf{d}_i - \mathbf{k}_i$ is assessed to ascertain whether the \mathbf{x}_1 used to produce the signature under verification is the same as the one used to produce the i -th revoked signature. Finally, the community revocation manager may revoke users by updating the list of revoked private-keys (KRL) or the list of signatures of revoked users (SRL) using Procedure 7.

4.2 Detailed Description of the LEPID Scheme

We now present our LEPID scheme in detail. We start by recalling some standard functionalities that are used in the UC model of the DAA [7]:

- \mathcal{F}_{CA} is a common certificate authority functionality that is available to all parties.
- \mathcal{F}_{CRS} is a common reference string functionality that provides participants with all system parameters.
- $\mathcal{F}_{\text{auth}}^*$ is a special authenticated communication functionality that provides an authenticated channel between the issuer and the signer.

The LEPID scheme includes the Setup, Join, Sign, Verify and Revoke procedures that are as follows.

Procedure 1 (Setup). \mathcal{F}_{CRS} creates the system parameters: $\text{sp} = (\lambda, t, q, n, m, \mathcal{R}_q, \beta, \ell, r, s, \xi)$, where λ, t are positive integer security parameters, β is a positive real number such that $\beta < q$, ℓ is the length of the users' identifiers, and r, s and ξ represent standard deviations of Gaussian distributions.

Upon input $(\text{SETUP}, \text{sid})$, where sid is a unique session identifier, the issuer first checks that $\text{sid} = (\mathcal{I}, \text{sid}')$ for some sid' , then creates its key pair.

The Issuer's public key is $\text{pp} = (\text{sp}, \mathbf{b}, \hat{A}_{\mathcal{I}}, \hat{A}_0, \hat{A}_1, \dots, \hat{A}_\ell, \mathbf{u}, \mathcal{H}_0, \mathcal{H}, H)$, where $\hat{A}_{\mathcal{I}}, \hat{A}_i (i = 0, 1, \dots, \ell) \in \mathcal{R}_q^m$, $\mathbf{b}, \mathbf{u} \in \mathcal{R}_q$, $\mathcal{H}_0 : \{0, 1\}^* \rightarrow \{1, 2, 3\}^t$, $\mathcal{H} : \{0, 1\}^* \rightarrow \mathcal{R}_q$, and $H : \{0, 1\}^* \rightarrow \{0, 1, 2, \dots, 2n - 1\}$. The Issuer's private key is $\hat{T}_{\mathcal{I}}$, which is the trapdoor of $\hat{A}_{\mathcal{I}}$ with $\|\hat{T}_{\mathcal{I}}\|_\infty \leq \beta$.

The issuer initialises the Member List $\text{ML} \leftarrow \emptyset$. The issuer proves that his secret key is well formed in $\pi_{\mathcal{I}}$, and registers the key $(\hat{T}_{\mathcal{I}}, \pi_{\mathcal{I}})$ with \mathcal{F}_{CA} and outputs $(\text{SETUPDONE}, \text{sid})$.

Procedure 2 (Join Request). On input query $(\text{JOIN}, \text{sid}, \text{jsid}, \mathcal{M})$, the signer \mathcal{M} forwards $(\text{JOIN}, \text{sid}, \text{jsid})$ to \mathcal{I} , who replies by sending $(\text{sid}, \text{jsid}, \rho, \text{bsn}_{\mathcal{I}})$ back to the signer, where ρ is a uniform random nonce $\rho \leftarrow \{0, 1\}^\lambda$, and $\text{bsn}_{\mathcal{I}}$ is the issuer's base name. The signer \mathcal{M} proceeds as follows:

1. It checks that no such entry exists in its storage.
2. It samples a private key: $\mathbf{x}_1 \leftarrow \mathcal{D}_s$ and $(\mathbf{x}_2, \dots, \mathbf{x}_{m+1}) \leftarrow \mathcal{D}_r^m$. Let $\hat{X}_t = (\mathbf{x}_1, \dots, \mathbf{x}_{m+1})$ correspond to \mathcal{M} 's secret key with the condition $\|(\mathbf{x}_2, \dots, \mathbf{x}_{m+1})\|_\infty \leq \beta/2$ and $\|\mathbf{x}_1\|_\infty \leq \beta$. \mathcal{M} stores its key as (sid, \hat{X}_t) , and computes the corresponding public key $\mathbf{u}_t = [\mathbf{b}|\hat{A}_{\mathcal{I}}]\hat{X}_t \pmod q$, a link token $\text{nym}_{\mathcal{I}} = \mathcal{H}(\text{bsn}_{\mathcal{I}})\mathbf{x}_1 + \mathbf{e}_{\mathcal{I}} \pmod q$ for some error $\mathbf{e}_{\mathcal{I}} \leftarrow \mathcal{D}_s$ such that $\|\mathbf{e}_{\mathcal{I}}\|_\infty \leq \beta$, and generates a signature based proof:

$$\begin{aligned} \pi_{\mathbf{u}_t} = \text{SPK} \left\{ \text{public} := \{\text{pp}, \mathbf{u}_t, \text{bsn}_{\mathcal{I}}, \text{nym}_{\mathcal{I}}\}, \text{witness} := \{\hat{X}_t = (\mathbf{x}_1, \dots, \mathbf{x}_{m+1}), \right. \\ \left. \mathbf{e}_{\mathcal{I}}\}, \mathbf{u}_t = [\mathbf{b}|\hat{A}_{\mathcal{I}}]\hat{X}_t \pmod q \wedge \|\hat{X}_t/\mathbf{x}_1\|_\infty \leq \beta/2 \wedge \|\mathbf{x}_1\|_\infty \leq \beta \right. \\ \left. \wedge \text{nym}_{\mathcal{I}} = \mathcal{H}(\text{bsn}_{\mathcal{I}})\mathbf{x}_1 + \mathbf{e}_{\mathcal{I}} \pmod q \wedge \|\mathbf{e}_{\mathcal{I}}\|_\infty \leq \beta \right\}(\rho). \end{aligned}$$

3. It sends $(\text{nym}_{\mathcal{I}}, \mathbf{u}_t, \pi_{\mathbf{u}_t})$ to the issuer by giving $\mathcal{F}_{\text{auth}}^*$ an input $(\text{SEND}, \text{nym}_{\mathcal{I}}, \pi_{\mathbf{u}_t}, \text{sid}, \text{jsid})$.

\mathcal{I} , upon receiving $(\text{SENT}, \text{nym}_{\mathcal{I}}, \pi_{\mathbf{u}_t}, \text{sid}, \text{jsid}, \mathcal{M})$ from $\mathcal{F}_{\text{auth}}^*$, verifies the proof $\pi_{\mathbf{u}_t}$ and makes sure that the signer $\mathcal{M} \notin \text{ML}$. \mathcal{I} stores $(\text{jsid}, \text{nym}_{\mathcal{I}}, \pi_{\mathbf{u}_t}, \mathcal{M})$, and generates the message $(\text{JOINPROCEED}, \text{sid}, \text{jsid}, \text{id}, \pi_{\mathbf{u}_t})$, for some identity $\text{id} \in \{0, 1\}^\ell$ assigned to \mathcal{M} , and not used before by any joined member.

Procedure 3 (Join Proceed). If the signer chooses to proceed with the Join session, the message $(\text{JOINPROCEED}, \text{sid}, \text{jsid})$ is sent to the issuer, who performs as follows:

1. It checks the record $(\text{jsid}, \text{nym}_{\mathcal{I}}, \text{id}, \mathcal{M}, \pi_{\mathbf{u}_t})$. For all $\text{nym}'_{\mathcal{I}}$ from the previous Join records, the issuer checks whether $\|\text{nym}_{\mathcal{I}} - \text{nym}'_{\mathcal{I}}\|_\infty \leq 2\beta$ holds; if yes, the issuer further checks if $\mathbf{u}_t = \mathbf{u}'_t$. If the equality $\mathbf{u}_t = \mathbf{u}'_t$ holds, the issuer will jump to Step 4 returning $\hat{X}_h = \hat{X}'_h$, if not the issuer will abort. Note that this double check will make sure that no two EPID keys will include the same \mathbf{x}_1 value.
2. For all $\text{nym}^*_{\mathcal{I}}$ in the Issuer's Revocation record IR , the issuer checks whether the equation

$$\|\text{nym}_{\mathcal{I}} - \text{nym}^*_{\mathcal{I}}\|_\infty \leq 2\beta$$

holds, if yes the issuer aborts.

3. It calculates the vector of polynomials $\hat{A}_h = [\hat{A}_{\mathcal{I}} | \hat{A}_0 + \sum_{i=1}^{\ell} \text{id}_i \hat{A}_i] \in \mathcal{R}_q^{2m}$.
4. It samples, using the issuer's private key $\hat{T}_{\mathcal{I}}$, a preimage $\hat{X}_h = [\hat{X}_{h_1} | \hat{X}_{h_2}] = (\mathbf{y}_2, \dots, \mathbf{y}_{2m+1}) \in \mathcal{D}_r^m \times \mathcal{D}_s^m$ of $\mathbf{u} - \mathbf{u}_t$ such that $\hat{A}_h \hat{X}_h = \mathbf{u}_h = \mathbf{u} - \mathbf{u}_t \pmod{q}$ and $\|\hat{X}_{h_1}\|_{\infty} \leq \beta/2$ and $\|\hat{X}_{h_2}\|_{\infty} \leq \beta$.
5. The issuer adds $(\text{nym}_{\mathcal{I}}, \text{id}, \mathcal{M}, \pi_{\mathbf{u}_t})$ to his data base, and sends $(\text{sid}, \text{jsid}, \hat{X}_h)$ to \mathcal{M} via $\mathcal{F}_{\text{auth}}^*$.

When \mathcal{M} receives the message $(\text{sid}, \text{jsid}, \hat{X}_h)$, it checks that the equations $\hat{A}_h \hat{X}_h = \mathbf{u}_h \pmod{q}$ and $\mathbf{u} = \mathbf{u}_t + \mathbf{u}_h$ are satisfied with $\|\hat{X}_{h_1}\|_{\infty} \leq \beta/2$ and $\|\hat{X}_{h_2}\|_{\infty} \leq \beta$. It stores $(\text{sid}, \mathcal{M}, \text{id}, \hat{X}_h, \mathbf{u}_t)$ and outputs $(\text{JOINED}, \text{sid}, \text{jsid})$. \mathcal{M} then computes $\hat{X} = (\mathbf{x}_1, \forall_{i=(2, \dots, m+1)} \mathbf{x}_i := \mathbf{x}_i + \mathbf{y}_i, \forall_{i=(m+2, \dots, 2m+1)} \mathbf{x}_i := \mathbf{y}_i)$, where $\|\hat{X}\|_{\infty} \leq \beta$.

Procedure 4 (Sign Request). Upon input $(\text{SIGN}, \text{sid}, \text{ssid}, \mathcal{M}, \mu)$, the signer does the following:

1. It makes sure to have a Join record $(\text{sid}, \text{id}, \hat{X}, \mathcal{M})$.
2. It generates a sign entry $(\text{sid}, \text{ssid}, \mu)$ in its record.
3. Finally it outputs $(\text{SIGNPROCEED}, \text{sid}, \text{ssid}, \mu)$.

Procedure 5 (Sign Proceed). When \mathcal{M} gets permission to proceed for ssid , the signer proceeds as follows:

1. It retrieves the records $(\text{sid}, \text{id}, \pi_{\mathbf{u}_t})$ and $(\text{sid}, \text{ssid}, \mu)$.
2. \mathcal{M} samples a random polynomial \mathbf{p} and computes the polynomial $\text{nym} = \mathbf{p}\mathbf{x}_1 + \mathbf{e} \pmod{q}$, for an error term $\mathbf{e} \leftarrow \mathcal{D}_s$ such that $\|\mathbf{e}\|_{\infty} \leq \beta$. \mathcal{M} then generates a signature based knowledge proof π .

$$\begin{aligned} \pi &= \text{SPK} \left\{ \text{public} := \{\text{pp}, \text{nym}, \mathbf{p}\}, \right. \\ &\quad \text{witness} := \{\hat{X} = (\mathbf{x}_1, \dots, \mathbf{x}_{2m+1}), \text{id}, \mathbf{e}\} : \\ &\quad \left. [\hat{\mathbf{b}} | \hat{A}_h] \hat{X} = \mathbf{u} \wedge \|\hat{X}\|_{\infty} \leq \beta \wedge \text{nym} = \mathbf{p}\mathbf{x}_1 + \mathbf{e} \wedge \|\mathbf{e}\|_{\infty} \leq \beta \right\}(\mu). \end{aligned}$$

The details of the proof π are presented in the Appendix of the full version of this paper [11].

3. The signer proves that it is not using any of the keys that produced a revoked signature $(\sigma_i^*, \mathbf{p}_i^*, \text{nym}_i^*)$ in the signature revocation list (more details about the proof can be found in the Appendix of the full version of this paper [11]).
 - Let $\text{nym}_i^* = \mathbf{p}_i^* \mathbf{f}_i + \mathbf{l}_i$, where $(\mathbf{f}_i, \mathbf{l}_i)$ were used before to create nym_i^* by some \mathcal{M}_i^* that generated a revoked signature $\sigma_i^* \in \text{SRL}$. \mathcal{M} proceeds as follows:
 - $\mathbf{q}_i, \mathbf{l}'_i, \mathbf{l}''_i, \mathbf{l}'''_i \leftarrow \mathcal{D}_s$
 - $\mathbf{o}_i = \mathbf{p}_i^* \mathbf{q}_i + \mathbf{l}'_i, \mathbf{k}_i = \mathbf{o}_i \mathbf{x}_1 + \mathbf{l}''_i, \mathbf{d}_i = \text{nym}_i^* \mathbf{q}_i + \mathbf{l}'''_i$
 - $\mathbf{r}_{x_1}, \mathbf{r}_e, \mathbf{r}_{q_i}, \mathbf{r}_{l'_i}, \mathbf{r}_{l''_i}, \mathbf{r}_{l'''_i} \leftarrow \mathcal{D}_s$
 - $\mathbf{t}_{\text{nym}} = \mathbf{p}\mathbf{r}_{x_1} + \mathbf{r}_e, \mathbf{t}_{o_i} = \mathbf{p}_i^* \mathbf{r}_{q_i} + \mathbf{r}_{l'_i},$
 $\mathbf{t}_{k_i} = \mathbf{o}_i \mathbf{r}_{x_1} + \mathbf{r}_{l''_i}, \mathbf{t}_{d_i} = \text{nym}_i^* \mathbf{r}_{q_i} + \mathbf{r}_{l'''_i}.$
 - Calculates the challenge $c_v = H(\mathbf{t}_{\text{nym}} | \mathbf{t}_{o_i} | \mathbf{t}_{k_i} | \mathbf{t}_{d_i} | \mu) \in \{0, 1, 2, \dots, 2n-1\}$.
 - The following responses are computed:

- $\mathbf{s}_{x_1} = \mathbf{r}_{x_1} + X^{c_v} \mathbf{x}_1$, $\mathbf{s}_e = \mathbf{r}_e + X^{c_v} \mathbf{e}$, $\mathbf{s}_{q_i} = \mathbf{r}_{q_i} + X^{c_v} \mathbf{q}_i$,
 $\mathbf{s}_{l'_i} = \mathbf{r}_{l'_i} + X^{c_v} \mathbf{l}'_i$, $\mathbf{s}_{l''_i} = \mathbf{r}_{l''_i} + X^{c_v} \mathbf{l}''_i$, $\mathbf{s}_{l'''_i} = \mathbf{r}_{l'''_i} + X^{c_v} \mathbf{l}'''_i$.

Abort if any of these rejection samples outputs 1:

- $\text{rej}(\mathbf{s}_{x_1}, X^{c_v} \mathbf{x}_1, \xi)$, $\text{rej}(\mathbf{s}_e, X^{c_v} \mathbf{e}, \xi)$, $\text{rej}(\mathbf{s}_{q_i}, X^{c_v} \mathbf{q}_i, \xi)$,
 $\text{rej}(\mathbf{s}_{l'_i}, X^{c_v} \mathbf{l}'_i, \xi)$, $\text{rej}(\mathbf{s}_{l''_i}, X^{c_v} \mathbf{l}''_i, \xi)$ or $\text{rej}(\mathbf{s}_{l'''_i}, X^{c_v} \mathbf{l}'''_i, \xi)$.

4. Finally, \mathcal{M} outputs $\sigma = (\pi, \text{nym}, \mathbf{o}_i, \mathbf{k}_i, \mathbf{d}_i, \mathbf{s}_{x_1}, \mathbf{s}_e, \mathbf{s}_{q_i}, \mathbf{s}_{l'_i}, \mathbf{s}_{l''_i}, \mathbf{s}_{l'''_i}, c_v, \text{KRL}, \text{SRL})$.

Procedure 6 (Verify). Let KRL denotes the revocation list with all the rogue signer's secret keys \mathbf{x}_1^* . Upon input (VERIFY, $sid, \sigma, \mu, \text{KRL}, \text{SRL}$), the verifier proceeds as follows:

1. It checks the zero-knowledge proof regarding the statement: $\{[b] \hat{A}_h \hat{X} = \mathbf{u} \wedge \|\hat{X}\|_\infty \leq \beta \wedge \text{nym} = \mathbf{p} \mathbf{x}_1 + \mathbf{e} \pmod{q} \wedge \|\mathbf{e}\|_\infty \leq \beta\}$.
2. For all $\mathbf{x}_1^* \in \text{KRL}$, if $\|\mathbf{p} \mathbf{x}_1^* - \text{nym}\|_\infty \leq \beta$ the verifier outputs 0.
3. For all $\sigma_i^* = (\pi_{\text{nym}_i^*}, \text{nym}_i^*, \mathbf{p}_i^*) \in \text{SRL}$, the verifier
 - (a) computes:
$$\begin{aligned} - \mathbf{t}'_{k_i} &= \mathbf{o}_i \mathbf{s}_{x_1} + \mathbf{s}_{l''_i} - X^{c_v} \mathbf{k}_i, \mathbf{t}'_{d_i} = \text{nym}_i^* \mathbf{s}_{q_i} + \mathbf{s}_{l'''_i} - X^{c_v} \mathbf{d}_i, \\ \mathbf{t}'_{o_i} &= \mathbf{p}_i^* \mathbf{s}_{q_i} + \mathbf{s}_{l'_i} - X^{c_v} \mathbf{o}_i, \mathbf{t}'_{\text{nym}} = \mathbf{p} \mathbf{s}_{x_1} + \mathbf{s}_e - X^{c_v} \text{nym}. \end{aligned}$$
 - (b) checks $c_v \stackrel{?}{=} H(\mathbf{t}'_{\text{nym}} | \mathbf{t}'_{o_i} | \mathbf{t}'_{k_i} | \mathbf{t}'_{d_i} | \mu)$ and that all the following norms satisfy $\|\mathbf{s}_{x_1}\|_\infty, \|\mathbf{s}_e\|_\infty, \|\mathbf{s}_{q_i}\|_\infty, \|\mathbf{s}_{l'_i}\|_\infty, \|\mathbf{s}_{l''_i}\|_\infty, \|\mathbf{s}_{l'''_i}\|_\infty \leq \beta + \sqrt{n} \beta$.
4. For all $\sigma_i^* = (\pi_{\text{nym}_i^*}, \text{nym}_i^*, \mathbf{p}_i^*)$, the verifier checks $2\|\mathbf{d}_i - \mathbf{k}_i\| < \Gamma$, where Γ is a function of β . If $2\|\mathbf{d}_i - \mathbf{k}_i\| < \Gamma$ the verifier outputs 0, otherwise 1.

Procedure 7 (Revoke). On input (Revoke, $sid, \mathbf{x}_1^*, \text{KRL}$) or (Revoke, $sid, \sigma^*, \mu^*, \text{SRL}$), the revocation manager adds \mathbf{x}_1^* to KRL or σ^* to SRL after verifying σ^* .

5 A Sketched Security Proof for LEPID

In this section, we provide a sketch of the security proof of the LEPID scheme. A detailed security proof is presented in the Appendix of the full version of this paper [11]. A variant of the sequence of games of [7] is presented, showing that no environment \mathcal{E} can distinguish the real world protocol Π with an adversary \mathcal{A} , from the ideal world $\mathcal{F}_{\text{EPID}}^l$ with a simulator \mathcal{S} . Starting with the real world protocol game, we change the protocol game by game in a computationally indistinguishable way, finally ending with the ideal world protocol.

Game 1. This is the real world protocol.

Game 2. An entity C is introduced, that receives all inputs from the honest parties and simulates Π for them. This is equivalent to Game 1.

Game 3. C is split into \mathcal{F} and \mathcal{S} . \mathcal{F} behaves as an ideal functionality, receiving all inputs and forwarding them to \mathcal{S} , who simulates the real world protocol for honest parties. \mathcal{S} sends the outputs to \mathcal{F} , who forwards them to \mathcal{E} . This game is similar to Game 2, but with a different structure.

Game 4. \mathcal{F} now behaves differently in the setup interface. It stores the algorithms for the issuer \mathcal{I} , and checks that the structure of sid is correct for an honest \mathcal{I} , aborting if not. In case \mathcal{I} is corrupt, \mathcal{S} extracts the secret key for

\mathcal{I} and proceeds in the setup interface on behalf of \mathcal{I} . Clearly \mathcal{E} will notice no change.

Game 5. \mathcal{F} now performs the verification and key revocation checks instead of forwarding them to \mathcal{S} . There are no protocol messages and the outputs are exactly as the real world protocol. However, the verification algorithm that \mathcal{F} uses does not contain any key or signature revocation checks. \mathcal{F} can perform this check separately, so the outcomes are equal.

Game 6. \mathcal{F} stores in its records the members that have joined. If \mathcal{I} is honest, \mathcal{F} stores the secret key tsk , extracted from \mathcal{S} , for corrupt platforms. \mathcal{S} always has enough information to simulate the real world protocol except when the issuer is the only honest party. In this case, \mathcal{S} does not know who initiated the join, and so cannot make a join query with \mathcal{F} on the signer's behalf. Thus, to deal with this case, \mathcal{F} can safely choose any corrupt signer and put it into Members. The identities of signers are only used for creating signatures for honest signers, so corrupted signers do not matter. In the case that the signer is already registered in Members, \mathcal{F} would abort the protocol, but \mathcal{I} will have already tested this case before continuing with the query JOINPROCEED. Hence \mathcal{F} will not abort. Thus in all cases, \mathcal{F} and \mathcal{S} can interact to simulate the real world protocol.

Game 7. (Anonymity). In this game, \mathcal{F} creates anonymous signatures for honest platforms by running the algorithms defined in the setup interface. Let us start by defining Game 7.k.k'. In this game \mathcal{F} handles the first k' signing inputs of \mathcal{M}_i for $i < k$ using algorithms, and subsequent inputs are forwarded to \mathcal{S} who creates signatures as before. We note that Game 7.0.0=Game 6. For increasing k' , Game 7.k.k' will be at some stage equal to Game 7.k + 1.0, this is because there can only be a polynomial number of signing queries to be processed. Therefore, for large enough k and k' , \mathcal{F} handles all the signing queries of all signers, and Game 7 is indistinguishable from Game 7.k.k'. To prove that Game 7.k.k' + 1 is indistinguishable from Game 7.k.k', suppose that there exists an environment that can distinguish a signature of an honest party using $tsk = \mathbf{x}_1$ from a signature using a different $tsk^j = \mathbf{x}_1^j$, then the environment can solve the Decision Ring -LWE Problem.

The first $j \leq k'$ signing queries on behalf of \mathcal{M}_k are handled by \mathcal{F} using the algorithms, and subsequent inputs are then forwarded to \mathcal{S} as before. Now suppose that \mathcal{F} outputs the tuples $(nym^j, \mathbf{p}^j, \mathbf{o}_i^j, \mathbf{k}_i^j, \mathbf{d}_i^j, \mathbf{s}_{x_1}^j, \mathbf{s}_e^j, \mathbf{s}_{q_i}^j, \mathbf{s}_{l_i}^j, \mathbf{s}_{l_i''}^j, \mathbf{s}_{l_i'''}^j, c_v^j, \text{SRL})$ for $j \leq k'$, with $nym^j = \mathbf{p}^j \mathbf{x}_1 + e^j$, for an error term $e_j \leftarrow \mathcal{D}_s$, and the remaining proofs are honestly generated. The $j = k' + 1$ -th query for \mathcal{M}_k is as follows: $(nym^S, \mathbf{p}^S, \mathbf{o}_i^S, \mathbf{k}_i^S, \mathbf{d}_i^S, \mathbf{s}_{x_1}^S, \mathbf{s}_e^S, \mathbf{s}_{q_i}^S, \mathbf{s}_{l_i}^S, \mathbf{s}_{l_i''}^S, \mathbf{s}_{l_i'''}^S, c_v^S, \mu^S, \text{SRL})$. \mathcal{S} is challenged to decide if $(nym^S, \mathbf{p}^S, \mathbf{o}_i^S, \mathbf{k}_i^S, \mathbf{d}_i^S, \mathbf{s}_{x_1}^S, \mathbf{s}_e^S, \mathbf{s}_{q_i}^S, \mathbf{s}_{l_i}^S, \mathbf{s}_{l_i''}^S, \mathbf{s}_{l_i'''}^S, c_v^S, \mu^S, \text{SRL})$ is chosen from a Ring LWE distribution for some secret \mathbf{x}_1 or uniformly at random. \mathcal{S} proceeds in simulating the signer without knowing the secret \mathbf{x}_1 . \mathcal{S} can answer all the H queries, as \mathcal{S} is controlling \mathcal{F}_{CRS} . \mathcal{S} sets: $\mathbf{t}_{k_i}^S = \mathbf{o}_i^S \mathbf{s}_{x_1}^S + \mathbf{s}_{l_i''}^S - X^{c_v^S} \mathbf{k}_i^S$; $\mathbf{t}_{d_i}^S = nym_i^S \mathbf{s}_{q_i}^S + \mathbf{s}_{l_i'''}^S - X^{c_v^S} \mathbf{d}_i^S$; $\mathbf{t}_{o_i}^S = \mathbf{p}_i^S \mathbf{s}_{q_i}^S + \mathbf{s}_{l_i}^S - X^{c_v^S} \mathbf{o}_i^S$; $\mathbf{t}_{nym}^S = \mathbf{p}^S \mathbf{s}_{x_1}^S + \mathbf{s}_e^S - X^{c_v^S} nym^S$; and, finally, $c_v^S := H(\mathbf{t}_{nym}^S | \mathbf{t}_{o_i}^S | \mathbf{t}_{k_i}^S | \mathbf{t}_{d_i}^S | \mu^S)$. For $i > k' + 1$, \mathcal{S} outputs the tuples $(nym^j, \mathbf{p}^j, \mathbf{o}_i^j, \mathbf{k}_i^j, \mathbf{d}_i^j, \mathbf{s}_{x_1}^j, \mathbf{s}_e^j, \mathbf{s}_{q_i}^j, \mathbf{s}_{l_i}^j, \mathbf{s}_{l_i''}^j, \mathbf{s}_{l_i'''}^j, c_v^j, \mu^j, \text{SRL})$, with

$\text{nym}^j = \mathbf{p}^j \mathbf{x}_1^j + \mathbf{e}^j \pmod q$, for some freshly generated secret \mathbf{x}_1^j and error term $\mathbf{e}^j \leftarrow \mathcal{D}_s$. For each case, \mathcal{M}_k can provide a simulated proof as follows. \mathcal{S} sets $\mathbf{t}_{k_i}^j = \mathbf{o}_i^j \mathbf{s}_{x_1}^j + \mathbf{s}_{l_i'}^j - X^{c_v^j} \mathbf{k}_i^j$; $\mathbf{t}_{d_i}^j = \text{nym}_i^* \mathbf{s}_{q_i}^j + \mathbf{s}_{l_i''}^j - X^{c_v^j} \mathbf{d}_i^j$; $\mathbf{t}_{o_i}^j = \mathbf{p}_i^* \mathbf{s}_{q_i}^j + \mathbf{s}_{l_i'}^j - X^{c_v^j} \mathbf{o}_i^j$; $\mathbf{t}_{\text{nym}}^j = \mathbf{p}^j \mathbf{s}_{x_1}^j + \mathbf{s}_e^j - X^{c_v^j} \text{nym}^j$; and, finally, $c_v^j := H(\mathbf{t}_{\text{nym}}^j | \mathbf{t}_{o_i}^j | \mathbf{t}_{k_i}^j | \mathbf{t}_{d_i}^j | \mu^j)$.

Thus, any distinguisher between Game 7.k.k' and Game 7.k.k' + 1 can solve the Decision Ring LWE Problem.

Game 8. \mathcal{F} now no longer informs \mathcal{S} about the message and \mathbf{p} that are being signed. If the signer \mathcal{M} is honest, then \mathcal{S} can learn nothing about the message μ and \mathbf{p} . Instead, \mathcal{S} knows only the leakage $l(\mu, \mathbf{p})$. To simulate the real world, \mathcal{S} chooses a pair (μ', \mathbf{p}') such that $l(\mu', \mathbf{p}') = l(\mu, \mathbf{p})$. An environment \mathcal{E} observes no difference, and thus Game 8=Game 7.

Game 9. If \mathcal{I} is honest, then \mathcal{F} now only allows members that joined to sign. An honest signer will always check whether it has joined before signing in the real world protocol, so there is no difference for honest signers. Therefore Game 9=Game 8.

Game 10. When storing a new $tsk = \mathbf{x}_1$, \mathcal{F} checks $\text{CheckTskCorrupt}(tsk) = 1$ or $\text{CheckTskHonest}(tsk) = 1$. We want to show that these checks will always pass. In fact, valid signatures always satisfy $\text{nym} = \mathbf{p} \mathbf{x}_1 + \mathbf{e}$ where $\|\mathbf{x}_1\|_\infty \leq \beta$ and $\|\mathbf{e}\|_\infty \leq \beta$. By the unique Shortest Vector Problem, there exists only one tuple $(\mathbf{x}_1, \mathbf{e})$ such that $\|\mathbf{x}_1\|_\infty \leq \beta$ and $\|\mathbf{e}\|_\infty \leq \beta$ for small enough β . Thus, $\text{CheckTskCorrupt}(tsk)$ will always give the correct output. Also, due to the large min-entropy of discrete Gaussians the probability of sampling $\mathbf{x}_1' = \mathbf{x}_1$, and thus of having a signature already using the same $tsk = \mathbf{x}_1$, is negligible, which implies that $\text{CheckTskHonest}(tsk)$ will give the correct output with overwhelming probability. Hence Game 10=Game 9.

Game 11. (Completeness). In this game, \mathcal{F} checks that honestly generated signatures are always valid. This is true as sig algorithm always produces signatures passing through verification checks. Those signatures satisfy $\text{identify}(tsk, \sigma, \mu, \mathbf{p}) = 1$, which is checked via nym . \mathcal{F} also makes sure, using its internal records Members and DomainKeys that honest users are not sharing the same secret key tsk . If there exists a key $tsk' = \mathbf{x}_1'$ in Members and DomainKeys such that $\|\text{nym} - \mathbf{p} \mathbf{x}_1'\|_\infty \leq \beta$, then this breaks search Ring-LWE.

Game 12. Check-IX is added to ensure that there are no multiple tsk tracing back to the same signature. Since there exists only one pair $(\mathbf{x}_1, \mathbf{e}_{\mathcal{I}})$, $\|\mathbf{x}_1\|_\infty \leq \beta$, $\|\mathbf{e}_{\mathcal{I}}\|_\infty \leq \beta$, satisfying $\text{nym}_{\mathcal{I}} = \mathcal{H}(\text{bsn}_{\mathcal{I}}) \mathbf{x}_1 + \mathbf{e}_{\mathcal{I}}$, two different signers cannot share the same \mathbf{x}_1 , thus any valid signature traces back to a single tsk .

Game 13. (Unforgeability). To prevent accepting signatures that were issued by the use of join credentials not issued by an honest issuer, \mathcal{F} further adds Check-X. This is due to the unforgeability of Boyen signatures [4].

Game 14. (Unforgeability). Check-XI is added to \mathcal{F} , preventing the forging of signatures with honest tsk and credentials. If a valid signature is given on a message that the signer has never signed, the proof could not have been

simulated. \mathbf{x}_1 would be extracted and Ring-LWE would be broken. So Game 14=Game 13.

Game 15. Check-XII is added to \mathcal{F} , ensuring that honest signers keys are not being revoked. If an honest signer is simulated by means of the Ring-LWE problem instance and a proper key KRL is found, it must be the secret key of the target instance. This is equivalent to solving the search Ring-LWE problem.

Game 16. \mathcal{F} now performs signature based revocation when verifying signatures. \mathcal{F} checks that there is no $(\sigma^*, \text{nym}^*, \mathbf{p}^*) \in \text{SRL}$ such that for some matching tsk_i and $(\sigma^*, \mu^*, \mathbf{p}^*) \in \text{SRL}$, we have $\text{identify}(\sigma^*, \mu^*, \mathbf{p}^*, tsk_i) = 1$. By the soundness of the proof presented in the Appendix of the full version of this paper [11], this check will always pass with overwhelming probability. \square

6 Experimental Results

Let $q \geq 2$ represents an integer modulus such that $q = \text{poly}(n)$. For correctness, we require the main hardness parameter n , to be large enough (e.g., $n \geq 100$) and $q > \beta$ as both being at least a small polynomial in n . We also let $m = O(\log q)$ as in [17]. A concrete choice of parameters can be as follows: $n = 512$, $l = 32$, $q = 8380417$, $m = 24$, and $\beta = 275$.

Both LDAA and LEPID were implemented in C, emulating all entities in a single machine. The code was compiled with gcc 4.8.5 with the `-O3` and `-march=native` flags and executed on an Intel i9 7900X CPU with 64GB running at 3.3 GHz operated by CentOS 7.5. The obtained experimental results can be found in Table 1. Note that the measured times for signing and verification do not take into account transfer times between the entities or object creation and destruction.

By construing the signer as a single entity instead of two as in the LDAA, the proposed LEPID scheme achieves a reduction of the private-key size of 1.5 times. While the comparison in signatures sizes between both schemes yields favourably for the proposed LEPID scheme with a small amount of rejected users, as the number of users in the SRL increases, its signature size increases linearly at a rate of 18kB per rejected user (9 polynomials and an integer). When the SRL contains 500 users, the LEPID signature size closely matches that of the LDAA scheme. Should LEPID signing be implemented on a device with limited

Scheme	Private-key (kB)	Signature (MB)	Signing Time (s)	Verification Time (s)
LDAA	147	847	541	129
LEPID (no revoked users)	100	836	361	114
LEPID (100 users in SRL)	100	838	371	117
LEPID (500 users in SRL)	100	845	372	119
LEPID (1000 users in SRL)	100	854	374	121

Table 1: Experimental results for the proposed LEPID and LDAA [10] for $n = 512$, $q = 8380417$, $l = 32$, $m = 24$ and $\beta = 256$ obtained on an Intel i9 7900X

computational resources like the TPM, its constrained memory resources and the cost of data transfer might limit its application to small and medium-sized communities. In particular, if one considers a revocation rate of 0.1%, LEPID signatures will compare favourably in size to LDAA signatures for communities with fewer than 500,000 users.

The signing time in the LEPID scheme is dominated by the signature based knowledge proof π . The addition of the SRL, and consequently of 13 polynomial multiplications per rejected user, shows no meaningful impact in the final signing time, where LEPID maintains a speedup of 1.4 over the LDAA scheme. Likewise, in the verification time, the additional 2 polynomial multiplications per rejected user incurred by the SRL are negligible compared to the verification of π . Hence, the proposed LEPID scheme achieves a speedup of 1.1 when compared with the LDAA scheme across both small and medium rejection lists. For the computational complexity introduced by the SRL to be meaningful, the number of rejected users must be in the order of millions. Once more, the proposed LEPID scheme shows improved signature and verification times for small and medium communities when compared with the LDAA.

7 Conclusion

While EPID plays a determinant role in the security of SGX, the scheme currently deployed by Intel will become insecure in the event that a large-scale quantum-computer is produced. Herein, a novel EPID scheme is proposed, supported on lattice-based security assumptions, and achieving presumed quantum resistance. A security model for EPID is presented for the first time in the UC framework, and the proposed scheme is proven secure under this model. When compared with a closely related LDAA scheme from related art, the proposed LEPID achieves a reduction in the private-key size of 1.5 times, and of the signature and verification times of 1.4 and 1.1 times, respectively, when no users have been revoked. It is furthermore shown, experimentally, that the overhead introduced by the more effective revocation method of LEPID is minimal for small to medium-sized communities. Finally, it is expected that the proposed LEPID may benefit from theoretical developments and hardware accelerators that result from the increased interest that lattice-based cryptography has gathered in the last few years.

Acknowledgements. This research was supported by European Unions Horizon 2020 research and innovation programme under grant agreement No. 779391 (FutureTPM), and by national funds through Fundação para a Ciência e a Tecnologia (FCT) with references UID/CEC/50021/2019 and FCT Grant No. SFRH/BD/145477/2019.

References

1. Carsten Baum, Ivan Damgård, Sabine Oechsner, and Chris Peikert. Efficient commitments and zero-knowledge protocols from ring-sis with applications to lattice-based threshold cryptosystems. *IACR Cryptology ePrint Archive*, 2016:997, 2016.

2. Fabrice Benhamouda, Jan Camenisch, Stephan Krenn, Vadim Lyubashevsky, and Gregory Neven. Better zero-knowledge proofs for lattice encryption and their application to group signatures. In *ASIACRYPT (1)*, pages 551–572. Springer, 2014.
3. Dan Boneh, Saba Eskandarian, and Ben Fisch. Post-quantum epid signatures from symmetric primitives. In *Cryptographers' Track at the RSA Conference*, pages 251–271. Springer, 2019.
4. Xavier Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In *International Workshop on Public Key Cryptography*, pages 499–517. Springer, 2010.
5. Ernie Brickell, Jan Camenisch, and Liqun Chen. Direct anonymous attestation. In *Proceedings of the 11th ACM Conference on Computer and Communications Security, CCS '04*, pages 132–145, New York, NY, USA, 2004. ACM.
6. Jan Camenisch, Liqun Chen, Manu Drijvers, Anja Lehmann, David Novick, and Rainer Urian. One tpm to bind them all: fixing tpm2.0 for provably secure anonymous attestation. *Proceedings of IEEE S&P 2017*, 2017.
7. Jan Camenisch, Manu Drijvers, and Anja Lehmann. Universally composable direct anonymous attestation. In *Public-Key Cryptography – PKC 2016*, volume 9615 of *LNCS*, pages 234–264. Springer, 2016.
8. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 21–30. ACM, 2007.
9. Simon Johnson, Vinnie Scarlata, Carlos Rozas, Ernie Brickell, and Frank Mckeen. Intel® software guard extensions: Epid provisioning and attestation services. *White Paper*, 1:1–10, 2016.
10. Nada Kassem, Liqun Chen, Rachid El Bansarkhani, Jan Camenisch Ali El Kaafarani, Patrick Hough, Paulo Sérgio Alves Martins, , and Leonel Sous. More efficient, provably-secure direct anonymous attestation from lattices. In *Future Generation Computer Systems*, 2019.
11. Nada EL Kassem, Luis Fiolhais, Paulo Martins, Liqun Chen, and Leonel Sousa. A lattice-based enhanced privacy id. *Cryptology ePrint Archive*, Report 2019/1366, 2019. <https://eprint.iacr.org/2019/1366>.
12. San Ling, Khoa Nguyen, Damien Stehlé, and Huaxiong Wang. Improved zero-knowledge proofs of knowledge for the isis problem, and applications. In *Public-Key Cryptography–PKC 2013*, pages 107–124. Springer, 2013.
13. Vadim Lyubashevsky. *Towards practical lattice-based cryptography*. University of California, San Diego, 2008.
14. Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, volume 7237 of *LNCS*, 2012.
15. Hamid Nejatollahi, Nikil D Dutt, Indranil Banerjee, and Rosario Cammarota. Domain-specific accelerators for ideal lattice-based public key protocols. *IACR Cryptology ePrint Archive*, 2018:608, 2018.
16. National Institute of Standards and Technology. Post-quantum cryptography standardization, 1 2017. <https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Post-Quantum-Cryptography-Standardization>.
17. Chris Peikert et al. A decade of lattice cryptography. *Foundations and Trends® in Theoretical Computer Science*, 10(4):283–424, 2016.
18. Oded Regev. The learning with errors problem (invited survey). In *2010 IEEE 25th Annual Conference on Computational Complexity*, pages 191–204. IEEE, 2010.