Hierarchical Attribute-based Signatures: Short Keys and Optimal Signature Length

Daniel Gardham and Mark Manulis

Surrey Centre for Cyber Security, University of Surrey, UK

`d.gardham@surrey.ac.uk, mark@manulis.eu`

Abstract. With Attribute-based Signatures (ABS) users can simultaneously sign messages and prove compliance of their attributes, issued by designated attribute authorities, with some verification policy. Neither signer's identity nor possessed attributes are leaked during the verification process, making ABS schemes a handy tool for applications requiring privacy-preserving authentication. Earlier ABS schemes lacked support for hierarchical delegation of attributes (across tiers of attribute authorities down to the signers), a distinct property that has made traditional PKIs more scalable and widely adoptable.

This changed recently with the introduction of Hierarchical ABS (HABS) schemes, where support for attribute delegation was proposed in combination with stronger privacy guarantees for the delegation paths (path anonymity) and new accountability mechanisms allowing a dedicated tracing authority to identify these paths (path traceability) and the signer, along with delegated attributes, if needed. Yet, current HABS construction is generic with inefficient delegation process resulting in sub-optimal signature lengths of order $O(k^2|\Psi|)$ where Ψ is the policy size and k the height of the hierarchy.

This paper proposes a direct HABS construction in bilinear groups that significantly improves on these bounds and satisfies the original security and privacy requirements. At the core of our HABS scheme is a new delegation process based on the length-reducing homomorphic trapdoor commitments to group elements for which we introduce a new delegation technique allowing step-wise commitments to additional elements without changing the length of the original commitment and its opening. While also being of independent interest, this technique results in shorter HABS keys and achieves the signature-length growth of $O(k|\Psi|)$ which is optimal due to the path-traceability requirement.

1 Introduction

Attribute-based Signatures, first introduced in [29] and [30], provide privacy-preserving mechanisms for authenticating messages. An ABS signature assures the verifier that the signer owns a set of attributes that satisfy the signing policy

without leaking their identity, nor the set of attributes used. Traditional ABS schemes considered two security properties, user privacy and unforgeability. Informally, a user is anonymous if an ABS signature does not leak their identity, nor the set of attributes used to satisfy the signing policy, while unforgeability requires that a signer cannot produce a signature conforming to a policy for which he does not own a set of suitable attributes. Later constructions [16, 14, 19] offered more advanced functionality with an additional property of traceability which holds signers accountable by allowing a dedicated tracing authority to identity them if required.

The vast majority of existing ABS schemes [16, 32, 31, 5, 14, 19, 13, 34] are non-interactive, in the standard model and is based on bilinear maps and Groth-Sahai proofs [20], with the exception of [22], which uses RSA setting and the random oracle model, and the recent schemes in [15, 37] which rely on lattices. Interactive ABS schemes, e.g. [26], where policies must be chosen by verifiers ahead of the signing phase have also been proposed. In general, signing polices can have varying levels of flexibility and range from threshold policies [29], to monotone boolean predicates [14, 19], and generalised circuits [34]. Typically, more restrictive policies allow for more efficient constructions. Policy-based Signatures (PBS) [5] can be viewed as a generalisation of ABS schemes, albeit their security is currently proven in a single-user setting without addressing stronger non-frameability requirement of more recent ABS schemes [14, 19, 12].

Hierarchical Attribute-based Signature and their Limitations. Hierarchical Attribute-based Signatures (HABS), recently introduced in [12], extend traditional ABS schemes by permitting controlled delegation of attributes from a root authority (RA) over possibly multiple intermediate authorities (IAs) down to the users. In this way HABS aims to close the gap between ABS and traditional PKIs where hierarchical delegation can be achieved at low cost. In HABS, IAs can delegate attributes to any authority in the scheme and users can acquire attributes from any authority in the hierarchy that is authorised to issue them. In addition to strong non-frameability property in a multi-user setting, the authors extend traditional ABS privacy guarantees to protect not only the identity of the signer but also the identities of all intermediate authorities in the delegation path, as part of their new path-anonymity property. Traditional traceability property of ABS schemes has also been extended to hold not only signers but also intermediate authorities accountable for their actions, through the new notion of path-traceability where a dedicated tracing authority can reveal the entire delegation path, along with delegated attributes.

We observe that the HABS scheme in [12] is generic, based on standard cryptographic primitives, i.e., public key encryption, one-time signature, tag-based signature, and non-interactive zero-knowledge proofs. Its delegation process is handled using a tag-based signature (TBS) where an authority at level i produces a TBS signature, using attribute as a tag, on the public key of authority j together with all public keys appearing previously in the delegation path. As part of its HABS signature the signer proves knowledge of each TBS at every delegation of each attribute that is required to satisfy the policy. Clearly this del-

egation process is highly inefficient. Not only does an additional signature need to be verified per delegation (and per attribute), the size of the signature grows linearly in the distance from the root authority. Thus, per attribute, verification of the delegation path is of order $O(k^2)$.

Other related work. Attribute-based signatures can be seen as a generalisation of group [11] and ring [33] signatures, in which case identities are viewed as attributes and policies can only contain disjunction over them. The notion of hierarchical delegation in these, more restricted, primitives have been explored in [36] and [28] respectively. Attribute delegation has been widely investigated in anonymous credentials [9, 10]. Maji et al. [30] give discussion that ACs are a more powerful primitive than ABS but with effiency drawbacks, as attribute acquisition typically requires expensive zero-knowledge proofs. Note that in HABS intermediate authorities may know each other, and so as discussed in [12], there is no need to hide their identities from each other during the delegation phase, which in turn helps to omit costly proofs and make this phase more efficient than in the case of ACs. Regardless, we note that ACs with hierarchical delegation have been proposed [4]. Further, a homomorphic ABS scheme [24] has been used to construct non-delegatable anonymous credentials. In this setting, a signer obtains attributes directly from the (multiple) root authorities where combining attributes from different issuers requires an online collaboration. Anonymous Proxy Signatures [17] also allow for verification of anonymous delegation paths back to a root authority. However, tasks that are delegated, when viewed as attributes, remain in the clear and are required for verification of the proxy signature. Homomorphic Signatures [37] have been claimed to be equivalent to Attribute-based Signatures, however this equality has been shown to hold in the weaker security setting that only considers a solitary user. In which, it is impossible to capture the notion of collusion and non-frameability. Finally, we note functional signatures [3, 8] also allow for delegation of signing rights. Here, however, keys are dependent on the function f and can only sign on messages that fall within the range. For an attribute-based scheme, this would require keys for each possible combination of attributes a user obtains.

Contribution. We address the suboptimal efficiency of the so-far only (generic) HABS construction [12] and propose a scheme with a completely new delegation mechanism which no longer relies on the consecutive issue of tag-based signatures from higher-level to lower-level authorities on the delegation path. The main novelty in our approach is a smart use of the length-reducing homomorphic trapdoor commitment scheme to multiple group elements from [20] which we extend with delegation capabilities. At a high level, at each delegation the issuing intermediate authority amends the current trapdoor opening such that the existing commitment incorporates the public key of the next-level authority or user to whom the attribute is delegated. With this new delegation mechanism we are able to significantly reduce the lengths of HABS keys and achieve the optimal growth of $O(k|\Psi|)$ for the length of HABS signatures, depending on the length k of the delegation path and size $|\Psi|$ of the signing policy. In particular, verifying delegation of an attribute along the path takes $O(k)$ steps (as opposed

to $O(k^2)$ in [12]). We use the original security model from [12] to show that our construction satisfies the required properties of path anonymity, path traceability, and non-frameability, in the standard model under standard assumptions in bilinear groups and an additional assumption which we justify using the generic group model [35]. Our efficiency improvement claims over [12] are reinforced in a detailed comparison between the two schemes.

2 HABS Model: Entities and Definitions

We start with the description of entities within the HABS ecosystem.

Attribute Authorities. The set of Attribute Authorities (AA) comprises the Root Authority (RA) and Intermediate Authorities (IAs). All AAs can delegate attributes to lower-level IAs and users. The RA is at the top of the hierarchy and upon setup, defines the universe of attributes $\mathbb{A}$. With its key pair (ask_0, apk_0), the RA can delegate a subset of attributes to IAs which hold their own key pairs (ask_i, apk_i), $i > 0$. IAs can further delegate/issue attributes to any end user (aka. signer). In this way a dynamically expandable HABS hierarchy can be established.

Users. Users join the scheme by creating their own key pair (usk, upk), and are issued attributes by possibly multiple AAs.

By Ψ we denote a predicate for some signing policy. A policy-conforming user can use usk to create a HABS signature, provided their issued set of attributes A satisfies the policy, i.e. $\Psi(A') = 1$ for some $A' \subseteq A$. Users are unable to delegate attributes further and thus can be viewed as the lowest tier of the hierarchy. To account for this, when an attribute is delegated to a user a dedicated symbol $\star$ will be used in addition to upk to mark the end of the delegation path.

Warrants. An IA or user, upon joining the HABS scheme, receives a *warrant* **warr** that consists of all their delegated attributes $a \in \mathbb{A}$ and a list of all AAs in each of the delegation paths. Warrants can be updated at any time, i.e. if the owner is issued a new attribute, by appending a new entry with the list of authorities on the delegation path. We use the notation $|\textbf{warr}|$ to denote the size of the warrant, i.e. the number of attributes stored in the warrant **warr**, and we use $|\textbf{warr}[\text{a}]|$ to denote the length of the delegation path of the attribute $a \in \mathbb{A}$. Upon signing, the user submits a reduced warrant for an attribute set $A' \subseteq A$ that satisfies $\Psi(A') = 1$.

Tracing Authority. The tracing authority (TA) is independent of the hierarchy. Upon receiving a valid HABS signature, it can identify the signer and all authorities on the delegation paths for attributes that the signer used to satisfy the signing policy. The tracing authority can output a publicly verifiable proof $\hat{\pi}$ that the path was identified correctly. The existence of such tracing authority improves the accountability of signers and IAs from possible misbehaviour.

Definition 1 (Hierarchical ABS Scheme [12]) *A scheme* HABS := (Setup, KGen, AttIssue, Sign, Verify, Trace, Judge) *consists of the following seven processes:*

- Setup(1^λ) is the initialisation process where based on some security parameter $\lambda \in \mathbb{N}$, the public parameters pp of the scheme are defined, and the root and tracing authority independently generate their own key pair, i.e. RA's (ask_0, apk_0) and TA's (tsk, tpk). In addition, RA defines the universe $\mathbb{A}$ of attributes, and a label $\star$ for users. We stress that due to dynamic hierarchy, the system can be initialised by publishing (pp, apk_0, tpk) with $\mathbb{A}$ and $\star$ contained in pp.
- KGen(pp) is a key generation algorithm executed independently by intermediate authorities and users. Each entity generates its own key pair, i.e., (ask_i, apk_i) for $i > 0$ or (usk, upk).
- AttIssue $(ask_i, \boldsymbol{warr}_i, A, \{apk_j | upk_j\})$ is an algorithm that is used to delegate attributes to an authority with apk_j or issue them to the user with upk. On input of an authority's secret key ask_i, $i \in \mathbb{N}_0$, its warrant $\boldsymbol{warr}_i$, a subset of attributes A from $\boldsymbol{warr}_i$, and the public key of the entity to which attributes are delegated or issued, it outputs a new warrant $\boldsymbol{warr}$ for that entity.
- Sign $((usk, \boldsymbol{warr}), \mathrm{m}, \Psi)$ is the signing algorithm. On input of the signer's usk and (possibly reduced) $\boldsymbol{warr}$, a message m and a predicate Ψ it outputs a signature σ.
- Verify $(apk_0, (\mathrm{m}, \Psi, \sigma))$ is a deterministic algorithm that outputs 1 if a candidate signature σ on a message m is valid with respect to the predicate Ψ and 0 otherwise.
- Trace $(tsk, apk_0, (\mathrm{m}, \Psi, \sigma))$ is an algorithm executed by the TA on input of its private key tsk and outputs either a triple $(upk, \boldsymbol{warr}, \hat{\pi})$ if the tracing is successful or $\perp$ to indicate its failure. Note that $\boldsymbol{warr}$ contains attributes and delegation paths that were used by the signer.
- Judge $(tpk, apk_0, (\mathrm{m}, \Psi, \sigma), (upk, \boldsymbol{warr}, \hat{\pi}))$ is a deterministic algorithm that checks a candidate triple $(upk, \boldsymbol{warr}, \hat{\pi})$ from the tracing algorithm and outputs 1 if the triple is valid and 0 otherwise.

A HABS scheme must have the *correctness* property ensuring that any signature σ generated based on an honestly issued warrant will verify and trace correctly. The output $(upk, \boldsymbol{warr}, \hat{\pi})$ of the tracing algorithm on such signatures will be accepted by the public judging algorithm with overwhelming probability.

2.1 Security Properties

Our security definitions resemble the requirements of *path anonymity*, *path traceability*, and *non-frameability* from [12]. We recall the associated game-based definitions assuming probabilistic polynomial time (PPT) adversaries interacting with HABS entities through the following set of oracles:

- O_{Reg} : $\mathcal{A}$ registers new IAs and users through this registration oracle, for which a key pair will be generated and added to *List*. The public key is given to the adversary. Initially, the entity is considered honest, and so the public key is also added to the list HU.

$$\begin{array}{ll}
\underline{O_{\mathrm{Reg}}(\,\cdot\,) \text{ with } (\,\cdot\,) = (\,i\,) \text{ and } i \notin HU} & \underline{O_{\mathrm{Att}}(\,\cdot\,)} \\
\end{array}$$

$O_{\mathrm{Reg}}(\,\cdot\,)$ with $(\,\cdot\,) = (\,i\,)$ and $i \notin HU$	$O_{\mathrm{Att}}(\,\cdot\,)$
1: $(sk_i, pk_i) \leftarrow \texttt{KGen}(\mathrm{pp})$	$(\,\cdot\,) = (i, \mathbf{warr}_i, a, \{apk_j|upk_j\})$
2: $List \leftarrow List \cup \{(i, pk_i, sk_i)\}$	1: $L := \{a|(a, pk_a, sk_a) \in List\}$
3: $HU \leftarrow HU \cup \{i\}$	2: **if** $i \in L \ \wedge j \in L$ **then**
4: **return** pk_i	3: $\mathbf{warr} \leftarrow \texttt{AttIssue}(ask_i,$
	$\mathbf{warr}_i, a, \{apk_j|upk_j\})$
$O_{\mathrm{Corr}}(\,\cdot\,)$ with $(\,\cdot\,) = (\,i\,)$	4: **return warr**
1: **if** $i \in HU$ **then**	5: **return** $\perp$
2: $HU \leftarrow HU - \{i\}$	
3: **return** sk_i from $List$	$O_{\mathrm{Sig}}(\,\cdot\,)$ with $(\,\cdot\,) = (i, \mathbf{warr}, \mathrm{m}, \Psi)$
	1: $A \leftarrow \{a| \ a \in \mathbf{warr}\}$
$O_{\mathrm{Tr}}(\,\cdot\,)$ with $(\,\cdot\,) = (\mathrm{m}, \Psi, \sigma)$	2: **if** $i \in HU \ \wedge \ \Psi(A)$ **then**
1: **return** $\texttt{Trace}(tsk, apk_0, (\mathrm{m}, \Psi, \sigma))$	3: $\sigma \leftarrow \texttt{Sign}((usk_i, \mathbf{warr}), \mathrm{m}, \Psi)$
	4: **return** σ
	5: **return** $\perp$

Fig. 1. Oracles used in the HABS security experiments.

- O_{Corr} : This oracle allows $\mathcal{A}$ to corrupt registered users or IAs. Upon input of a public key, the corresponding private key is given as output if it exists in $List$. The public key is removed from HU so the oracle keeps track of corrupt entities.
- O_{Att} : $\mathcal{A}$ uses this oracle to ask an attribute authority to delegate attributes to either an IA or to the user. In particular, the adversary has control over which attributes are issued and the oracle outputs a warrant **warr** if both parties are registered, otherwise it outputs $\perp$.
- O_{Sig} : $\mathcal{A}$ can ask for a HABS signature from a registered user. The adversary provides the warrant (and implicitly the attributes used), signing policy, message and the public key of the signer. If the attribute set satisfies the policy, and the public key is contained in HU then the signature will be given to $\mathcal{A}$, otherwise $\perp$ is returned.
- O_{Tr} : s $\mathcal{A}$ can ask the TA trace a HABS signature (provided by the adversary) to the output is returned. The TA does verification checks on the signature and upon failure, will return $\perp$.

Path Anonymity. This property extends the anonymity guarantees of traditional ABS schemes and hides not only the identity of the signer but also the identities of all intermediate authorities on delegation paths for attributes included into the signer's warrant. Path anonymity, as defined in Fig. 2, also ensures that signatures produced by the same signer remain unlinkable. The corresponding experiment requires the adversary to distinguish which warrant and private key were used in the generation of the challenge HABS signature σ_b. In the first phase, $\mathcal{A}_1$ generates a hierarchy of authorities and users, utilising

$$\mathbf{Exp}^{\text{pa-}b}_{HABS,\mathcal{A}}(\lambda)$$

1 : $(\text{pp}, ask_0, tsk) \leftarrow \mathtt{Setup}(1^\lambda)$

2 : $(st, (usk_0, \mathbf{warr}_0), (usk_1, \mathbf{warr}_1), m, \Psi) \leftarrow \mathcal{A}_1(\text{pp}, ask_0 : O_{\text{Reg}}, O_{\text{Corr}}, O_{\text{Tr}})$

3 : **if** $|\mathbf{warr}_0| = |\mathbf{warr}_1|$ **then**

4 : $\quad \sigma_0 \leftarrow \mathtt{Sign}((usk_0, \mathbf{warr}_0), m, \Psi), \ \sigma_1 \leftarrow \mathtt{Sign}((usk_1, \mathbf{warr}_1), m, \Psi)$

5 : $\quad$ **if** $\mathtt{Verify}(apk_0, (m, \Psi, \sigma_0)) = 1$ and $\mathtt{Verify}(apk_0, (m, \Psi, \sigma_1)) = 1$ **then**

6 : $\qquad b' \leftarrow \mathcal{A}_2(st, \sigma_b : O_{\text{Tr}})$

7 : $\qquad$ **return** $b' \ \wedge \ \mathcal{A}2$ did not query $O_{\text{Tr}}(tsk, (m, \Psi, \sigma_b))$

8 : **return** 0

Fig. 2. Path-Anonymity Experiment

the RA's secret key ask_0. If the warrants created by $\mathcal{A}_1$ are of the same size, a challenge HABS signature σ_b is produced on the randomly chosen user-warrant pair. In the second phase, with access to the tracing oracle, the adversary $\mathcal{A}_2$ must be able to guess the challenge bit b.

Definition 2 (Path Anonymity [12]) *A HABS scheme offers path anonymity if no PPT adversary adv can distinguish between* $\mathbf{Exp}^{\text{pa-0}}_{HABS,\mathcal{A}}$ *and* $\mathbf{Exp}^{\text{pa-1}}_{HABS,\mathcal{A}}$ *defined in Fig. 2, i.e., the following advantage is negligible in* λ:

$$\mathbf{Adv}^{\text{pa}}_{HABS,\mathcal{A}}(\lambda) = |\Pr[\mathbf{Exp}^{\text{pa-0}}_{HABS,\mathcal{A}}(\lambda) = 1] - |\Pr[\mathbf{Exp}^{\text{pa-1}}_{HABS,\mathcal{A}}(\lambda) = 1]|.$$

Non-frameability. This property, formalised in Fig. 3, captures the notion of unforgeability, i.e., that no PPT adversary can create a HABS signature without having an honestly issued warrant that satisfies the policy, and in particular, they cannot create one on behalf of an user for which the secret key is not known. The adversary wins if either he produces a valid HABS signature, or is able to perform delegation for at least one attribute on behalf of any honest authority anywhere in the delegation path.

Definition 3 (Non-Frameability [12]) *A HABS scheme is non-frameable, if no PPT adversary* $\mathcal{A}$ *can win the experiment* $\mathbf{Exp}^{\text{nf}}_{\mathsf{HABS},\mathcal{A}}$ *defined in Fig. 3, i.e., the following advantage is negligible in* λ:

$$\mathbf{Adv}^{\text{nf}}_{HABS,\mathcal{A}}(\lambda) = |\Pr[\mathbf{Exp}^{\text{nf}}_{HABS,\mathcal{A}}(\lambda) = 1]|.$$

Path Traceability. This property, formalised in Fig. 4, ensures that any valid HABS signature can be traced (by the tracing authority) to the signer and the set of authorities that were involved in the issue of attributes used to produce the signature. The adversary is required to output either a HABS signature that verifies but cannot be traced, or one in which the tracing algorithm outputs a

$\boxed{\begin{array}{l}
\mathbf{Exp}^{\mathrm{nf}}_{HABS,\mathcal{A}}(\lambda) \\
\hline
1: \quad \mathrm{pp} \leftarrow \mathtt{Setup}(1^\lambda), ask_0 \leftarrow \mathtt{KGen}(1^\lambda), tsk \leftarrow \mathtt{TKGen}(1^\lambda) \\
2: \quad ((\sigma, m, \Psi),(upk_j, \mathbf{warr}, \hat{\pi})) \leftarrow \mathcal{A}(\mathrm{pp}, ask_0, tsk : O_{\mathrm{Att}}, O_{\mathrm{Sig}}, O_{\mathrm{Corr}}, O_{\mathrm{Reg}}) \\
3: \quad \mathbf{if}\ \mathtt{Verify}(apk_0,(\mathrm{m}, \Psi, \sigma))\ \wedge\ \mathtt{Judge}(tpk, apk_0, (\mathrm{m}, \Psi, \sigma), (upk_j, \mathbf{warr}, \hat{\pi}))\ \mathbf{then} \\
4: \qquad \mathbf{if}\ j \in HU \wedge \mathcal{A}\ \text{did not query } O_{\mathrm{Sig}}((usk_j, \mathbf{warr}), \mathrm{m}, \Psi)\ \mathbf{then},\quad \mathbf{return}\ 1 \\
5: \qquad \mathbf{if}\ \exists a.\ a \in \mathbf{warr} \implies (apk_0, apk_1, \ldots, apk_n, upk_j, \star) = \mathbf{warr}[a]\ \wedge \\
6: \qquad\qquad ((\exists i \in [0, n-1].\ \mathcal{A}\ \text{did not query } O_{\mathrm{Att}}(i, \cdot, a, apk_{i+1}) \wedge i \in HU)\ \vee \\
7: \qquad\qquad (\mathcal{A}\ \text{did not query } O_{\mathrm{Att}}(n, \cdot, a, upk_j)\ \wedge n \in HU)\ )\ \mathbf{then},\ \mathbf{return}\ 1 \\
8: \quad \mathbf{return}\ 0
\end{array}}$

Fig. 3. Non-Frameability Experiment

$\boxed{\begin{array}{l}
\mathbf{Exp}^{\mathrm{tr}}_{\mathsf{HABS},\mathcal{A}}(\lambda) \\
\hline
1: \quad \mathrm{pp} \leftarrow \mathtt{Setup}(1^\lambda), ask_0 \leftarrow \mathtt{KGen}(1^\lambda), tsk \leftarrow \mathtt{TKGen}(1^\lambda) \\
2: \quad ((\sigma, m, \Psi),(upk, \mathbf{warr}, \hat{\pi})) \leftarrow \mathcal{A}(\mathrm{pp}, tsk : O_{\mathrm{Att}}, O_{\mathrm{Corr}}, O_{\mathrm{Reg}}) \\
3: \quad \mathbf{if}\ \mathtt{Verify}(apk_0,(m, \Psi, \sigma))\ \mathbf{then} \\
4: \qquad \mathbf{if}\ \mathtt{Trace}(tsk,(m, \Psi, \sigma)) = \bot\ \mathbf{then},\quad \mathbf{return}\ 1 \\
5: \qquad \mathbf{if}\ \mathtt{Judge}(tpk, apk_0, (m, \Psi, \sigma), (upk, \mathbf{warr}, \hat{\pi}))\ \wedge \\
6: \qquad\qquad (\exists a.\ a \in \mathbf{warr} \implies (apk_0, apk_1, \ldots, apk_n, upk, \star) = \mathbf{warr}[a]\ \wedge \\
7: \qquad\qquad (\ (\exists i \in [0, n-1].\ i \in HU \wedge (i+1, apk_{i+1}, ask_{i+1}) \notin List) \vee \\
8: \qquad\qquad (n \in HU \wedge (\cdot, upk, usk) \notin List)\ )\ )\ \mathbf{then},\quad \mathbf{return}\ 1 \\
9: \quad \mathbf{return}\ 0
\end{array}}$

Fig. 4. Path-Traceability Experiment

warrant for which at least one IA or the user is unknown to the experiment, i.e., were not previously registered in *List*. The attribute-issuing oracle checks that both entities are registered to prevent the trivial attack where adversary asks the oracle to delegate to an unregistered entity.

Definition 4 (Path Traceability [12]) *A HABS scheme offers path traceability if no PPT adversary $\mathcal{A}$ can win the experiment $\mathbf{Exp}^{\mathrm{tr}}_{\mathsf{HABS},\mathcal{A}}$ defined in Fig. 4, i.e., the following advantage is negligible in λ:*

$$\mathbf{Adv}^{\mathrm{tr}}_{HABS,\mathcal{A}}(\lambda) = |\mathrm{Pr}[\mathbf{Exp}^{\mathrm{tr}}_{HABS,\mathcal{A}}(\lambda) = 1]|.$$

3 Our Short HABS Construction

We start with the description of the underlying hardness assumptions and building blocks.

3.1 Underlying Hardness Assumptions

In addition to widely used hardness assumptions in the asymmetric bilinear group setting generated by $\mathcal{BG}(1^\lambda)$, namely q-Srong Diffie-Hellman (q-SDH) [7], Symmetric eXternal Diffie-Hellman (SXDH) [1] and Simultaneous Decision LINear (SDLIN) [25] assumptions, which we do not recall here, our scheme requires the following interactive assumption, which we prove to hold in the generic group model [35, 23]. We equip the adversary with an oracle $\mathcal{O}_{X_0,Y_0}$ and assume that it is hard to produce group elements that satisfy the verification equations for an input that has not been queried to the oracle. We note that on input of (X,Y), the oracle can compute each component without knowledge the discrete logarithm of X or Y as it has access to r, s, t_1, t_2.

Assumption 1 *Let $pp := (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2) \leftarrow \mathcal{BG}(1^\lambda)$. The adversary $\mathcal{A}$ has access to the oracle $\mathcal{O}_{X_0,Y_0}(\cdot, \cdot)$ which on input (X, Y) returns $(A, B, T_1, T_2, T_3, T_4) := (h_2^{x_0 t_1} h_5^{x_0 t_2} h_2^{x t_1} h_5^{y t_1}, h_5^{y_0 t_1} h_4^{y_0 t_2} h_5^{x t_2} h_4^{y t_2}, f_1^{t_1}, f_2^{t_2}, g_2^{t_1}, g_2^{t_2})$ for $t_1, t_2 \leftarrow \mathbb{Z}_p^*$. We assume that for all p.p.t. adversaries $\mathcal{A}$, the following probability is negligible in λ.*

$$Adv_\mathcal{A}^{\texttt{Assump1}}(\lambda) :=$$
$$Pr\left[\begin{array}{l} x_0, y_0, r, s \leftarrow \mathbb{Z}_p^*; h_1 := g_1^s, h_2 := g_1^{s^2}, h_3 := g_1^r, h_4 := g_1^{r^2}, h_5 := g_1^{rs} \\ f_1 := g_2^s, f_2 := g_2^r, X_0 := h_1^{x_0}; Y_0 := h_3^{y_0}; \\ (\hat{A}, \hat{B}, \hat{T}_1, \hat{T}_2, \hat{T}_3, \hat{T}_4, \hat{X}, \hat{Y}) \leftarrow \mathcal{A}^{\mathcal{O}_{X_0,Y_0}}(pp, X_0, Y_0, h_1, h_2, h_3, h_4, h_5, f_1, f_2) : \\ e(\hat{A}, g_2) = e(X_0, \hat{T}_1 \hat{T}_2) e(\hat{X}\hat{Y}, \hat{T}_1) \ \wedge \ e(\hat{B}, g_2) = e(Y_0, \hat{T}_1 \hat{T}_2) e(\hat{X}\hat{Y}, \hat{T}_2) \\ \wedge \ e(g_1, \hat{T}_1 \hat{T}_2) = e(h_1, \hat{T}_3) e(h_3, \hat{T}_4) \quad \text{where } (\hat{X}, \hat{Y}) \text{ was not queried.} \end{array}\right]$$

Theorem 1. *Let $\mathcal{A}$ denote an adversary in the generic group model against Assumption 1. $\mathcal{A}$ has access to oracles for which he makes q_G group queries, q_P pairing queries, and q_O oracle queries. The probability ϵ of $\mathcal{A}$ winning the game for Assumption 1 is bounded by $\epsilon \leq 5(q_G + q_P + 6q_O + 11)^2/p$, where p is the prime order of the generic groups.*

Proof. See Appendix C.

3.2 Cryptographic Building Blocks

The following building blocks will be used in our HABS construction.

Tag-based Encryption. A TBE scheme has the same algorithms as a traditional public key encryption scheme, except that its encryption and decryption procedures take an extra tag t as input. A correctly formed TBE ciphertext C will fail to decrypt if the tag used as input to the decryption algorithm is different from the one used upon encryption. We adopt the TBE scheme from [25] which relies on the SDLIN assumption. It offers selective-tag witness-indistinguishable chosen-ciphertext (st-IND-CCA) security, where an adversary is unable to distinguish between two ciphertexts under the same tag t of their choosing. Kiltz

[27] showed that a st-IND-CCA `TBE` scheme combined with a strongly unforgeable one-time signature, where the one-time verification key is used as the tag, gives rise to an IND-CCA2 secure PKE. Our scheme uses this approach.

One-Time Signature. For the `OTS`, we use the strongly unforgeable BBS one-time signature scheme from [7]. It consist of three algorithms (`KeyGen`, `Sig`, `Ver`) with a verification key in $\mathbb{G}^2 \times \mathbb{Z}_p^*$ and the corresponding signing key in $\mathbb{Z}_p^{*2}$.

Groth-Sahai Proofs. We use the Groth-Sahai proof system [21] to construct the required non-interactive zero-knowledge proofs `NIZK`. A GS proof, which consists of five algorithms (`Setup`, `Prove`, `Verify`, `SimSetup`, `SimProve`), allows proving relations involving multi-linear, quadratic, and pairing-based equations. We use GS proofs in the asymmetric bilinear group setting with Type-3 curves [18], i.e., where there is no computable isomorphism between $\mathbb{G}_1$ and $\mathbb{G}_2$, in which case their security is based on the SXDH assumption [7].

Homomorphic Trapdoor Commitments to Group Elements. The key to our short HABS scheme is the length-reducing homomorphic trapdoor commitment scheme by Groth [20], adopted in our new delegation mechanism. With the HTC scheme, defined by four algorithms (`Setup`, `KeyGen`, `Commit`, `Trapdoor`), one can use the trapdoor key tk to open a constant-length commitment (c, d) to arbitrary group elements with respect to a commitment key. We observe that due to its construction this `HTC` scheme has an interesting property that allows a commitment to group elements step-wise, i.e. an opening (a_i, b_i) to elements $g_1, \ldots, g_i$ can be transformed into an opening (a_{i+1}, b_{i+1}) for an extended set of elements $g_1, \ldots, g_{i+1}$ without knowledge of the secret commitment key, i.e., without jeopardising the binding property for the already committed group elements. In our HABS scheme such step-wise extension of an initial commitment (c, d) produced by the root authority allows intermediate authorities, upon delegation, to embed public keys of next-level authorities or users, along with the delegated attribute, by providing appropriate modification to the opening of (c, d), that is without changing its value nor increasing its length. Proving ownership of a delegated attribute amounts to presenting an opening (a, b) for the commitment (c, d) to the attribute and the list of public keys of authorities on the delegation path, i.e., $apk_0, \ldots,$ upk, $\star$. In our scheme we use the asymmetric variant of Groth's `HTC` scheme with security based on the XDLIN assumption [1].

3.3 Specification of our HABS Scheme

We start with a high-level intuition behind our HABS construction and provide detailed specification in Figures 5, 6, and 7.

High-Level Overview. As part of the setup process public parameters pp of the scheme are generated. They include the security parameter λ, the description of bilinear groups $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$, the trapdoor key tk for the `HTC` scheme, the

```
Setup(λ)

 0 :   (𝔾₁, 𝔾₂, 𝔾_T, g₁, g₂, e, p) ← 𝓑𝓖(1^λ)      11 :   Sample m_r, m_s, n_r, n_s ← ℤ_p^*

 1 :   Sample r, s ← ℤ_p^*                          12 :   Sample a₀, b₀ ← 𝔾₁

 2 :   h₁ := g₁^s, h₂ := g₁^{s²}, h₃ := g₁^r,       13 :   Compute g_r ← g₂^{m_r}, g_s ← g₂^{m_s},
                                                              h_r ← g₂^{n_r}, h_s ← g₂^{n_s}
        h₄ := g₁^{r²}, h₅ := g₁^{rs}
                                                    14 :   Compute c := e(a₀, g_r)e(b₀, g_s)
 3 :   f₁ := g₂^s, f₂ := g₂^r
                                                              d := e(a₀, h_r)e(b₀, h_s)
 4 :   Define 𝓗₁ : 𝔸 → ℤ_p^*,
                                                    15 :   Compute Δ := m_r n_s − n_r m_s
        𝓗₂ : {0,1}^* → ℤ_p^*, 𝓗₃ : {0,1}^* → ℤ_p^*
                                                    16 :       α := n_s/Δ, β := −m_s/Δ,
 5 :   Sample g̃₁ ← 𝔾₁, g̃₂ ← 𝔾₂
                                                              γ := −n_r/Δ, δ := m_r/Δ
 6 :   Compute ζ ← e(g̃₁, g̃₂)
                                                    17 :   tk := (m_r, m_s, n_r, n_s, α, β, γ, δ)
 7 :   (tsk, tpk) ← TKGen
                                                    18 :   pp := (𝓖, c, d, a₀, b₀, tk, 𝓗₁, 𝓗₂,
 8 :   w₁ ← NIZK₁.Setup
                                                              𝓗₃, ζ, g̃₁, g̃₂, tpk, 𝔸, w₁, w₂)
 9 :   w₂ ← NIZK₂.Setup
                                                        for 𝓖 := (𝔾₁, 𝔾₂, 𝔾_T, g₁, g₂, e, p
10 :   Define attribute universe 𝔸
                                                              h₁, h₂, h₃, h₄, h₅, f₁, f₂)

                                                    19 :   return pp
```

Fig. 5. Setup algorithm of our HABS construction.

initial 'dummy' HTC commitment (c, d) with an opening (a_0, b_0), and the description of the attribute universe $\mathbb{A}$. The independent tracing authority TA generates the TBE key-pair $(tsk, tpk) := ((\eta_1, \eta_2), (V_1, V_2, V_3, V_4))$ with tpk included into pp. For simplicity we describe the setup phase as a single process involving computations performed by the RA and TA. We stress, however, that generation of $h_1, ..., h_5, f_1, f_2$ must be trusted in that no entity knows the corresponding exponents.

In our scheme, all attribute authorities and users generate their own private/public key pairs (ask_i, apk_i) and (usk, upk) respectively, of the form $\{(x, y), (X, Y, Z, \hat{Z})\}$. While only X and Y are used in the verification of attribute delegation which we prove in the signature, the components $(Z, \hat{Z})$ are used in the issuing phase. To ensure an authority creates a delegation that opens to (X, Y), we insist that the validity of a public key is checked prior to delegation. This is done by evaluating $e(XY, h_1) = e(Z, g_2)$ and $e(XY, h_3) = e(\hat{Z}, g_2)$. IAs and users obtain attributes from existing authorities at a higher level in the hierarchy. Ownership of a valid key-pair (ask_i, apk_i) allows authorities to delegate attributes further down the hierarchy and to the users, by manipulating the opening of the initial commitment (c, d).

With the trapdoor key tk an authority can create an opening (a_i, b_i) for (c, d) to the path that includes delegate's public key e.g. apk_j. Rather than opening directly, the issuer first creates randomisation tokens $T_1, T_2, T_3, T_4 \in \mathbb{G}_2$ and opens to these instead. It then uses T_1 and T_2 as *one-time* commitment keys to

KGen(pp)	AttIssue$(ask_i, \{apk_j | upk\}, att, a_i, b_i, \mathbf{warr})$
0 : Sample $x, y \leftarrow \mathbb{Z}_p^*$	0 : Parse $\{apk_j | upk\}$ as $(X_j, Y_j, Z_j, \hat{Z}_j)$
1 : $X := h_1^x, Y := h_3^y,$	1 : Verify $e(XY, h_1) = e(Z, g_2)$
2 : $Z := h_2^x h_5^y, \hat{Z} := h_5^x h_4^y$	$\quad$ and $e(XY, h_3) = e(\hat{Z}, g_2)$
3 : $pk := (X, Y, Z, \hat{Z}),$	2 : Sample $t_1, t_2 \leftarrow \mathbb{Z}_p^*$
4 : $sk := (pk, x, y)$	3 : $T_1 := f_1^{t_1}, T_2 := f_2^{t_2}, T_3 := g_2^{t_1}, T_4 := g_2^{t_2}$
5 : **return** (pk, sk)	4 : $\tilde{a} \leftarrow a_i^{m_r} b_i^{m_s} (h_2^{t_1} h_5^{t_2})^{-x_i} (Z_j h_1^{\mathcal{H}_1(att)})^{-t_1}$
	$\quad\;\, \tilde{b} \leftarrow a_i^{n_r} b_i^{n_s} (h_5^{t_1} h_4^{t_2})^{-y_i} (\hat{Z}_j h_3^{\mathcal{H}_1(att)})^{-t_2}$
TKGen(pp)	5 : $(a_{i+1}, b_{i+1}) := (\tilde{a}^\alpha \tilde{b}^\beta, \tilde{a}^\gamma \tilde{b}^\delta)$
	6 : $\mathbf{warr} = \mathbf{warr} \cup \{apk_i, T_1, T_2\}$
0 : Sample $\eta_1, \eta_2 \leftarrow \mathbb{Z}_p^*$	7 : **return** $(a_{i+1}, b_{i+1}, \mathbf{warr})$
1 : Compute $V_1 := g_1^{\eta_1}, V_2 := g_1^{\eta_2}$	
2 : Sample $V_3, V_4 \leftarrow \mathbb{G}_2$	
3 : $tpk := (V_1, V_2, V_3, V_4)$	
4 : $tsk := (tpk, \eta_1, \eta_2)$	
5 : **return** (tsk, tpk)	

Fig. 6. Key generation and issue of attributes in our HABS construction.

open to apk_j and the delegated attribute att, that is hashed into the message space using $g^{\mathcal{H}_1(att)}$. The randomisation tokens T_1 and T_2 are used to prevent forgeries (where the adversary combines multiple openings and in doing so, forges an opening to a new public key) whereas T_3 and T_4 are used to verify the well-formedness of T_1 and T_2, by evaluating $e(g_1, T_1 T_2) = e(h_1, T_4) e(h_3, T_3)$. The issuing authority updates the (possibly empty) warrant with opening (a_i, b_i), his public key apk_i and the randomisation tokens (T_1, T_2, T_3, T_4). As tk is part of public parameters, any IA in the hierarchy is able to perform the delegation procedure, where it receives (a_{i-1}, b_{i-1}) from its issue and generates (a_i, b_i) for the next delegation. When delegating to users, an issuing IA will open (c, d) to a designated element $\star \in \mathbb{G}$, in addition to the user's public key upk and the attribute att. A warrant contains the trapdoor opening and a list of all public keys of AAs that appear in the delegation path for any issued attribute.

Upon signing, the user first generates an OTS key-pair $(otssk, otsvk) := \{(k_1, k_2), (K_1, K_2, \kappa)\}$ and an opening to $\mathcal{H}_3(otsvk)$ by modifying the opening (a_0, b_0) using his public key upk and the trapdoor key tk. The reduced **warr** along with upk are encrypted in a TBE ciphertext under the TA's public key tpk and tag $\mathcal{H}_3(otsvk)$. The signing policy Ψ is modelled as a monotone span program, with labelling function ρ that maps rows from $\mathbf{S}$ to the attribute set $\mathbb{A}$. The signer proves that this set satisfies Ψ by computing a vector $\mathbf{z}$ such that $\mathbf{zS} = [1, 0, ..., 0]$, where any non-zero entry $\mathbf{z}_i$ implies $\rho(i) \in \mathbf{warr}$. A NIZK proof π is then computed using Groth-Sahai framework with witness $(upk, \mathbf{warr}, \mathbf{z}, \tilde{r}, \tilde{s})$ for the following relation:

$$\big((a,b),\mathbf{warr},upk,\mathbf{z}\big), \big(\Psi,otsvk,apk_0,C,tpk\big): \mathbf{z}S = [1,0,...,0]$$
$$\wedge \ (\forall i. \ z_i \neq 0 \implies att_i = \rho(i) \ \wedge \ (apk_{i_1},...,apk_{i_n},apk_{i_{n+1}} := upk) \in C$$
$$\wedge \ c^{|\mathbf{warr}|+1} = e(a,g_r)e(b,g_s)e(X,g_2^{\mathcal{H}_3(otsvk)})$$
$$\Pi_i \Pi_{n=0}^{k} e(X_{i_n},T_{1,i_n}T_{2,i_n})e(X_{i_{n+1}}Y_{i_{n+1}}g_1^{\mathcal{H}_1(att)},T_{1,i_n})$$
$$\wedge \ d^{|\mathbf{warr}|+1} = e(a,h_r)e(b,h_s)e(Y,g_2^{\mathcal{H}_3(otsvk)})$$
$$\Pi_i \Pi_{n=0}^{k} e(Y_{i_n},T_{1,i_n}T_{2,i_n})e(X_{i_{n+1}}Y_{i_{n+1}}g_1^{\mathcal{H}_1(att)},T_{2,i_n})$$
$$\wedge \ e(g_1,\Pi_i \Pi_{n=0}^{k}T_{1,i_n}T_{2,i_n}) = e(h_1,\Pi_i \Pi_{n=0}^{n}T_{4,i_n})e(h_3,\Pi_i \Pi_{n=0}^{k}T_{3,i_n}).$$

The message m and policy Ψ are then bound to this proof and ciphertext by hashing $\mathcal{H}_2(\pi,\mathrm{C},\Psi,m)$, before an OTS signature σ_o is produced with $otssk$. The resulting signature is verified with respect to the public parameters of the scheme, and the RA's public key apk_0 by verifying the OTS signature and the NIZK proof.

As part of the tracing procedure, executed by TA with knowledge of tsk, the ciphertext C is decrypted to obtain the warrant $\mathbf{warr}$, signer's public key upk, and the opening (a,b). A publicly verifiable NIZK proof $\hat{\pi}$ is created with witness tsk for the statement $(otsvk,\mathrm{C},tpk,(apk_0,\mathbf{warr}))$ and relation:

$$\mathtt{TBE.Dec}(tsk,C,\mathcal{H}_3(otsvk)) = (upk,\mathbf{warr},a,b).$$

We give a detailed construction for the Groth-Sahai proofs NIZK$_1$ and NIZK$_2$ in Appendix D.

3.4 Security Analysis

In this section we prove that our construction meets HABS security properties of path anonymity, non-frameability and path traceability.

Lemma 1 *The* HABS *construction from Fig. 5, 6 and 7 offers path anonymity, if SXDH, SDLIN and q-SDH hold in $\mathcal{G}$.*

Proof. We follow a game-based approach and show that the advantage of the PPT adversary $\mathcal{A}$ in the path-anonymity experiment for the HABS construction from Fig.5, 6 and 7, is bounded by the advantages of the constructed adversaries for the underlying primitives. We assume that adversary $\mathcal{A}$ asks n user registration queries and the probability for sampling one of these users is $1/n$.

Game G_0: This game is defined to the be the experiment $\mathbf{Exp}^{\mathrm{pa}\text{-}b}_{\mathtt{HABS},\mathcal{A}}(\lambda)$ in Fig. 2, where the 2-stage adversary $\mathcal{A} = (\mathcal{A}_1,\mathcal{A}_2)$ is required to distinguish between the HABS signatures $\sigma_0 = (\sigma_o^0,\mathrm{C}_0,\pi_0,otsvk_0)$ and $\sigma_1 = (\sigma_o^1,\mathrm{C}_1,\pi_1,otsvk_1)$.
Game G_1: We define the game G_1 as G_0 where the check "$\mathcal{A}_2$ did not query $O_{\mathrm{Tr}}(m,\Psi,\sigma_b)$" is enforced by the O_{Tr} oracle available to $\mathcal{A}_2$, which aborts the game if this is the case. The probability from G_0 to G_1 is preserved.
Game G_2: The game G_2 is obtained from G_1 where, on the output of O_{Tr}, we replace the NIZK$_2$ proof $\hat{\pi}$ with $\hat{\pi}'$ from the simulator NIZK$_2$.SimProve. We also replace Setup by SimSetup for NIZK$_2$. This prevents the case where $\mathcal{A}$ may

$$
\begin{array}{ll}
\underline{\texttt{Sign}(usk, m, \Psi, \{att_j, a_j, b_j, \mathbf{warr}_j\}_{j \in J})} & \underline{\texttt{Verify}(pk, \sigma, m, \Psi)} \\[4pt]
0: \quad (k_1, k_2, k_3) \leftarrow \mathbb{Z}_p^* & 0: \quad \text{Parse } \sigma \text{ as } (ots, \pi, otsvk) \\[4pt]
1: \quad otsvk := (\tilde{g}_2^{k_1}, \tilde{g}_2^{k_2}, k_3) & 1: \quad H \leftarrow \mathcal{H}_2(\pi||C||\Psi||m||otsvk) \\[4pt]
2: \quad \text{Compute } \mathbf{z} \text{ s.t. } \mathbf{z}S = [1, 0, ..., 0] & 2: \quad \textbf{return } \texttt{NIZK.Verify}(\pi) \\[4pt]
3: \quad T_1 := f_1^{t_1}, T_2 := f_2^{t_2}, & \qquad \land \; e(\sigma_o, \tilde{g}_2^{k_1} \cdot \tilde{g}_2^{H} \cdot \tilde{g}_2^{k_2 k_3}) = \zeta \\[4pt]
\qquad\quad T_3 := g_2^{t_1}, T_4 := g_2^{t_2} & \\[4pt]
4: \quad a' \leftarrow a_i^{m_r} b_i^{m_s} (h_2^{t_1} h_5^{t_2})^{-x} h_1^{-t_1 \mathcal{H}_3(otsvk)} & \underline{\texttt{Trace}(tsk, \sigma, m, \Psi)} \\[4pt]
\qquad\; b' \leftarrow a_i^{n_r} b_i^{n_s} (h_5^{t_1} h_4^{t_2})^{-y} h_3^{-t_2 \mathcal{H}_3(otsvk)} & 0: \quad \textbf{if } \texttt{Verify}(\sigma, m, \Psi) = 1 \textbf{ then} \\[4pt]
5: \quad (a', b') := (\tilde{a}^{\alpha} \tilde{b}^{\beta}, \tilde{a}^{\gamma} \tilde{b}^{\delta}) & 1: \quad \mathbf{warr} \leftarrow \texttt{TBE.Dec}(tsk, C, t) \\[4pt]
6: \quad (a, b) = (a' \cdot \Pi a_j, b' \cdot \Pi b_j) & \qquad\;\; \text{for } t = \mathcal{H}_1(otsvk) \\[4pt]
7: \quad C \leftarrow \texttt{TBE.Enc}(tpk, \mathbf{warr}, upk, & 2: \quad \hat{\pi} \leftarrow \texttt{NIZK}_2.\texttt{Prove}(tsk : \\[4pt]
\qquad\qquad\qquad a, b, \{\mathcal{H}_1(otsvk)\}) & \qquad\quad (otsvk, C, tpk, (apk_0, \mathbf{warr}))) \\[4pt]
8: \quad \pi \leftarrow \texttt{NIZK}_1.\texttt{Prove}((upk, \mathbf{z}, \mathbf{warr}, a, b) : & 3: \quad \textbf{return } (\mathbf{warr}, \hat{\pi}) \\[4pt]
\qquad\quad (C, otsvk, tpk, apk_0, \Psi) \in \mathcal{R}) & \\[4pt]
9: \quad H \leftarrow \mathcal{H}_2(\pi||C||\Psi||m||otsvk) & \underline{\texttt{Judge}(tpk, \mathbf{warr}, \sigma)} \\[4pt]
10: \quad \sigma_o \leftarrow \tilde{g}_1^{1/(k_1 + H + k_2 k_3)} & 0: \quad \texttt{Verify}(ask_0, (\sigma, m, \Psi)) \\[4pt]
11: \quad \textbf{return } (\sigma_o, C, \pi, otsvk) & \qquad\quad \land \; \texttt{NIZK}_2.\texttt{Verify}(\pi_2)
\end{array}
$$

Fig. 7. Sign, Verify, Trace and Judge algorithms of our HABS construction.

"extract" tsk from $\texttt{NIZK}_2$ proofs. Thus, for all future O_{Tr} oracle calls we use the simulated $\texttt{NIZK}_2$ proof. The probability that $\mathcal{A}$ can distinguish between these two games is bounded by the advantage of the zero-knowledge adversary $\mathcal{B}_{nizk2}$ for $\texttt{NIZK}_2$. For our instantiation of GS proofs, this is reduced to the SXDH assumption [21].

Game G_3: Let G_3 be the game obtained from G_2 where we replace the proof π_b from the challenge signature $\sigma_b = (\sigma_{o,b}, C_b, \pi_b, otsvk_b)$ with the simulated proof π_b' by calling $\texttt{NIZK}_1.\texttt{Sim}$ on $(C_b, otsvk_b, tpk, apk_0, \Psi)$. Additionally, we replace $\texttt{NIZK}_1.\texttt{Setup}$ by $\texttt{NIZK}_1.\texttt{SimSetup}$. The probability that $\mathcal{A}$ can distinguish between games G_2 and G_3 is bounded by the advantage of the zero-knowledge adversary $\mathcal{B}_{nizk1}$ for $\texttt{NIZK}_1$ proof. Similarly, this property is implied by SXDH.

Game G_4: Game G_4 only differs from game G_3 in that we abort if $\mathcal{A}_2$ queries $O_{\text{Tr}}(m, \Psi, (\sigma_o, C_b, \pi, otsvk_b))$. The adversary $\mathcal{A}$ is only able to distinguish between these games if it can produce a valid OTS signature σ_o for the message (C_b, π, m, Ψ) and public key $otsvk_b$, without knowledge of the secret key $otssk_b$. Thus, the capabilities of the adversary $\mathcal{A}$ to distinguish between these two games is bounded by the advantage of the adversary $\mathcal{B}_{ots}$ against the strong unforgeability of the OTS scheme, which is reduced to the q-SDH assumption [7].

Game G_5: Game G_5 is defined to be G_4, except we additionally do a check for any queries $\mathcal{A}_2$ makes do not contain the challenge ciphertext, that is $O_{\text{Tr}}(m, \Psi,$

$(\sigma_o, C_b, \pi, otsvk))$. If so the game is aborted. The output of O_{Tr} is for G_4 and G_5 is the same, as the oracle returns $\perp$ if the tag $otsvk_b$ for C is different from $otsvk$ received as input. Hence, the probability is preserved.

Game G_6: The game G_6 is the same as G_5, except that we move the OTS key generation from the signature generation phase into the setup of the experiment. Note that only one key pair needs to be created in this game since the adversary only sees the challenge signature. This step is necessary to utilise the st-IND-CCA property of the TBE scheme. The probability is unchanged from game G_5 to G_6.

Game G_7: Let G_7 be the game obtained from G_6 where the TBE ciphertext C_b from the challenge signature $\sigma_b = (\sigma_o^b, C_b, \pi'_b, otsvk_b)$ is replaced with the C_0. The adversary $\mathcal{A}$ is unable to query a ciphertext $C' \neq C_b$ for the same tag $\mathcal{H}_3(otsvk)$ as a result of game G_4. Further, any query to the oracle for a tag $t' \neq \mathcal{H}_3(otsvk)$ will also fail as decryption of C_b is dependent on the correct tag. Thus, the ability of the adversary $\mathcal{A}_2$ to distinguish between the ciphertexts C_0 and C_b is bounded by the advantage of the st-IND-CCA adversary $\mathcal{B}_{ind}$. For out instantiation, this property of TBE is implied by SDLIN [25].

The experiment G_7 provides $\mathcal{A}$ with the same challenge signature independent of b that $\mathcal{A}$ is asked to guess. Additionally, due to the zero-knowledge property of NIZK_2 used in G_1, $\mathcal{A}$ does not have access to tsk. Therefore, the probability that the adversary wins game G_7 is $1/2$ and hence the advantage of $\mathcal{A}$ to win this experiment is 0. □

Lemma 2 *The HABS construction from Fig. 5, 6 and 7 is non-frameable, if $\mathcal{H}_1, \mathcal{H}_2$ and $\mathcal{H}_3$ are second-preimage resistant hash functions, and q-SDH, SXDH and Assumption 1 hold in $\mathcal{G}$.*

Proof. See Appendix A.

Lemma 3 *The HABS construction from Fig. 5, 6 and 7 offers path traceability, if SXDH, SDLIN and Assumption 1 hold in $\mathcal{G}$.*

Proof. See Appendix B.

Theorem 2. *The HABS scheme in Figs. 5, 6 and 7 offers path-anonymity, non-frameability and path-traceability if $\mathcal{H}_1, \mathcal{H}_2$ and $\mathcal{H}_3$ are second-preimage resistant hash functions and SXDH, SDLIN, q-SDH and Assumption 1 hold in $\mathcal{G}$.*

Proof. The result follows from Lemmas 1, 2 and 3. □

4 Efficiency Comparison

We first compare the warrant sizes for our scheme and [12]. For a single attribute, an authority at level 1 (with respect to the root authority at level 0) has a warrant size of 6 group elements (of the form $\mathbb{G}_1^2 \times \mathbb{G}_2^4$). Further delegation to level 2 adds a further 6 elements in $\mathbb{G}_1^2 \times \mathbb{G}_2^4$. A delegation from the root authority contains the opening (a, b) (as part of the 6 elements) which is updated by

15

subsequent delegations, however the warrant must now also contain the issuers public key (X_i, Y_i) in $\mathbb{G}_1^2$. This generalises, and for a user at level k, the size of the warrant is $6k$ for a single attribute. Likewise, if we extend the number of attributes in the warrant to $|A|$, each of which has a delegation path of length k, then the warrant has $6k|A|$ group elements.

In contrast, for a single attribute issued to a level-1 entity, the warrant in [12] contains $7\lceil\frac{12+2m}{m-2}\rceil$ group elements, where m is the size of the message space used in the TBS instantiation. A level-2 delegation increases this to $7\lceil\frac{24+4m}{m-2}\rceil+2m+12$ elements, and this generalises for a single attribute that is issued to a level k entity to $7\lceil\frac{k(12+2m)}{m-2}\rceil+(k-1)(2m+12)$ elements. Similarly, a warrant that contains $|A|$ attributes adds a linear factor of $|A|$ to this term. To give a concrete comparison, a user with 3 attributes at level 4 of the hierarchy would have a warrant containing 72 group elements in our scheme, as opposed to 208 elements for an optimal choice of m (i.e., $m = 10$) in the scheme from [12]. Since m would be chosen in advance during the setup phase, the warrant would unlikely reach its optimal bound and for any suboptimal choice of m, the warrant grows linearly in this parameter.

Next, in Table 1 we compare the sizes of public keys (of users and authorities) and the lengths of signatures generated by our scheme and [12]. By β we denote the size of the span program representing the policy Ψ. As before k is the maximum length of a delegation path, $|\Psi|$ is the number of attributes needed to satisfy the signing policy, and m is the size of the message space for the TBS scheme used in [12].

		Dragan et al. [12]		This Work		
		$\mathbb{G}$	$\mathbb{Z}_p$	$\mathbb{G}_1$	$\mathbb{G}_2$	$\mathbb{Z}_p$
Public Keys	upk	14	-	4	-	-
	apk	$12+2m$	-	4	-	-
Sig.	ots	3	1	2	1	1
	C	$\frac{5(6+m)k(k-1)|\Psi|}{(m-2)} + 110$	-	$6(2k-1)|\Psi|+12$	$4(2k-1)|\Psi|+8$	-
	π	$\frac{28(6+m)k(k-1)|\Psi|}{(m-2)} + 18$	2β	8	$2k|\Psi|+8$	β

Table 1. Comparison of key and signature sizes.

In addition to being more efficient and shorter than [12], our scheme, in fact, produces HABS signatures of *optimal* length, from the asymptotic point of view. The need to provide path traceability, where the TA must be able to reveal the entire delegation path along with delegated attributes from a valid HABS signature implies the $O(k|\Psi|)$ growth of its length. This means that in order to reduce this bound path-tracability property would need to be relaxed.

Finally, our scheme brings a few other efficiency improvements. The use of Type-3 pairings results in fewer group elements and the possibility to achieve the same level security for a smaller choice of the prime p [18], which would give rise to smaller groups and faster operations than in the symmetric setting. In addition, we can adopt batch verification techniques available for Groth-Sahai proofs [6] to speed up the computations.

5 Conclusion

We proposed a direct construction of Hierarchical Attribute-based Signatures (HABS) with a new delegation process based on length-reducing homomorphic trapdoor commitments. Our HABS scheme significantly reduces the lengths of warrants, public keys and signatures in comparison to the so-far only known (generic) HABS construction. Moreover, due to the need to support the path-traceability requirement, our HABS scheme achieves optimal signature length growth of $O(k|\Psi|)$ for delegations paths of size k and signing policies of size $|\Psi|$. Our technique of step-wise embedding of new group elements into the homomorphic trapdoor commitment can be considered to be of independent interest, e.g., it could add support for delegation to other privacy-preserving signature schemes that rely on homomorphic trapdoor commitments, e.g. [2].

References

1. M. Abe, M. Chase, B. David, M. Kohlweiss, R. Nishimaki, and M. Ohkubo. Constant-size structure-preserving signatures: Generic constructions and simple assumptions. *Journal of Cryptology*, pages 833–878, 2016.
2. M. Abe, K. Haralambiev, and M. Ohkubo. Signing on elements in bilinear groups for modular protocol design. *IACR Cryptology ePrint Archive*, page 133, 2010.
3. M. Backes, S. Meiser, and D. Schröder. Delegatable functional signatures. In *PKC (1)'16*, pages 357–386, 2016.
4. M. Belenkiy, J. Camenisch, M. Chase, M. Kohlweiss, A. Lysyanskaya, and H. Shacham. Randomizable Proofs and Delegatable Anonymous Credentials. In *CRYPTO 2009*, pages 108–125. LNCS 5677.
5. M. Bellare and G. Fuchsbauer. Policy-based signatures. In *PKC 2014*, pages 520–537, 2014. LNCS 8383.
6. O. Blazy, G. Fuchsbauer, M. Izabachène, A. Jambert, H. Sibert, and D. Vergnaud. Batch groth–sahai. In *Applied Cryptography and Network Security*, pages 218–235, 2010.
7. D. Boneh and X. Boyen. Short Signatures Without Random Oracles. In *EUROCRYPT 2004*, pages 56–73. LNCS 3027, 2004.
8. E. Boyle, S. Goldwasser, and I. Ivan. Functional signatures and pseudorandom functions. In *PKC 2014*, pages 501–519, 2014.
9. J. Camenisch and A. Lysyanskaya. A signature scheme with efficient protocols. In *SCN 2002*, SCN 2002, pages 268–289. LNCS 2576, 2003.
10. D. Chaum. Security without identification: Transaction systems to make big brother obsolete. *Commun. ACM*, 28(10):1030–1044, 1985.
11. D. Chaum and E. van Heyst. Group signatures. In *EUROCRYPT 1991*, LNCS, pages 257–265. Springer, 1991.
12. C.-C. Drăgan, D. Gardham, and M. Manulis. Hierarchical attribute-based signatures. In *Cryptology and Network Security*, pages 213–234, 2018.

13. A. El Kaafarani and E. Ghadafi. Attribute-based signatures with user-controlled linkability without random oracles. In *Cryptog. and Coding*, pages 161–184, 2017.
14. A. El Kaafarani, E. Ghadafi, and D. Khader. Decentralized traceable attribute-based signatures. In *CT-RSA 2014*, pages 327–348. LNCS 8366, 2014.
15. A. El Kaafarani and S. Katsumata. Attribute-based signatures forunbounded circuits in the rom and efficient instantiations from lattices. In M. Abdalla and R. Dahab, editors, *Public-Key Cryptography – PKC 2018*, pages 89–119, 2018.
16. A. Escala, J. Herranz, and P. Morillo. Revocable attribute-based signatures with adaptive security in the standard model. In *AFRICACRYPT 2011*, pages 224–241. LNCS 6737, 2011.
17. G. Fuchsbauer and D. Pointcheval. Anonymous proxy signatures. In *SCN*, volume 5229 of *LNCS*, pages 201–217. Springer, 2008.
18. S. D. Galbraith, K. G. Paterson, and N. P. Smart. Pairings for cryptographers. *Discrete Applied Mathematics*, 156(16):3113–3121, 2008.
19. E. Ghadafi. Stronger security notions for decentralized traceable attribute-based signatures and more efficient constructions. In *CT-RSA 2015*, pages 391–409. LNCS 9048, 2015.
20. J. Groth. Homomorphic Trapdoor Commitments to Group Elements. Cryptology ePrint Archive, Report 2009/007, 2009.
21. J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *EUROCRYPT 2008*, pages 415–432, 2008. LNCS 4965.
22. J. Herranz. Attribute-based Signatures from RSA. *TCS*, 527:73–82, 2014.
23. T. Jager and A. Rupp. The semi-generic group model and applications to pairing-based cryptography. In *ASIACRYPT 2010*, pages 539–556. LLNC 6477.
24. N. Kaaniche, M. Laurent, P.-O. Rocher, C. Kiennert, and J. Garcia-Alfaro. PCS, A Privacy-Preserving Certification Scheme. In *Data Privacy Management, Cryptocurrencies and Blockcain Technology*, pages 239–256, 2017.
25. S. A. Kakvi. Efficient fully anonymous group signatures based on the groth group signature scheme. *Master's Thesis, University College London*, 2010.
26. D. Khader, L. Chen, and J. H. Davenport. Certificate-free attribute authentication. In *Cryptography and Coding*, pages 301–325, 2009.
27. E. Kiltz. Chosen-ciphertext security from tag-based encryption. In *Theory of Cryptography*, pages 581–600, 2006.
28. Ł. Krzywiecki, M. Sulkowska, and F. Zagórski. Hierarchical ring signatures revisited – unconditionally and perfectly anonymous schnorr version. In *Security, Privacy, and Applied Cryptography Engineering*, pages 329–346, 2015.
29. J. Li, M. H. Au, W. Susilo, D. Xie, and K. Ren. Attribute-based Signature and Its Applications. In *ACM ASIACCS 2010*, pages 60–69. ACM, 2010.
30. H. K. Maji, M. Prabhakaran, and M. Rosulek. Attribute-based signatures. In *CT-RSA 2011*, pages 376–392, 2011.
31. T. Okamoto and K. Takashima. Decentralized attribute-based signatures. In *PKC 2013*, pages 125–142. LNCS 7778, 2013.
32. T. Okamoto and K. Takashima. Efficient attribute-based signatures for non-monotone predicates in the standard model. In *PKC 2011*, pages 35–52. LNCS 6571, 2011.
33. R. L. Rivest, A. Shamir, and Y. Tauman. How to leak a secret. In *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 552–565. Springer, 2001.
34. Y. Sakai. Practical attribute-based signature schemes for circuits from bilinear map. *IET Information Security*, pages 184–193(9), 2018.
35. V. Shoup. Lower bounds for discrete logarithms and related problems. In *Advances in Cryptology — EUROCRYPT '97*, pages 256–266, 1997.

36. M. Trolin and D. Wikström. Hierarchical Group Signatures. In *ICALP 2005*, volume 3580 of *LNCS*, pages 446–458. Springer, 2005.
37. R. Tsabary. An equivalence between attribute-based signatures and homomorphic signatures, and new constructions for both. In *TCC (2)'17*, pages 489–518, 2017.

A   Proof of Theorem 1

Proof. A challenger $\mathcal{B}$ interacts with an adversary $\mathcal{A}$ in the game for Assumption 1. In the generic group model, $\mathcal{A}$ is given random encodings of group elements by defining 3 functions $\xi_i : \mathbb{G}_i \to \{0,1\}^*$ for $i = 1, 2, T$. Since the only property the adversary can test for is equality, it uses oracles to compute all group operations. It interacts with oracles using ξ-encodings of the group elements only.

The challenger keeps 3 lists, L_1, L_2, L_T, which contain the pair $(K, \kappa) \in L_i$, where K is the output of the map ξ_i, and κ is a multivariate polynomial in $\mathbb{Z}_p[X_0, Y_0, X_i, Y_i, A_i, B_i, T_{1,i}, T_{2,i}]$. The lists are initiated with $\xi_1(g_1), \xi_1(h_1), \xi_1(h_2)$, $\xi_1(h_3), \xi_1(h_4), \xi_1(h_5), \xi_2(g_2), \xi_2(f_1), \xi_2(f_2), \xi_1(X_0), \xi_1(Y_0)$.

- **Group Operations.** Oracle $\mathcal{O}_i$ allows the adversary to compute addition and subtraction in group $\mathbb{G}_i$. Upon input $\mathcal{O}_i(M, N, \pm)$, the challenger, who controls the oracle, checks that $K_M := \xi_i(M), K_N := \xi_i(N) \in L_i$ and identifies the pairs (K_M, κ_M) and (K_N, κ_N). If this step fails the oracle returns $\bot$. It computes the polynomial $\kappa_{M \pm N} := \kappa_M \pm \kappa_N$ and checks if $\kappa_{M \pm N} \in L_i$. If it exists, $\mathcal{B}$ returns the corresponding encoding K from the pair $(K, \kappa_{M \pm N})$, otherwise $\mathcal{B}$ selects a new random encoding K distinct from all other encodings and updates L_i with the pair $(K, \kappa_{M \pm N})$.

- **Pairing Operations.** The oracle $\mathcal{O}_P$ takes as input random encodings $(K_1, K_2) \in \mathbb{G}_1 \times \mathbb{G}_2$. $\mathcal{B}$ searches the lists L_1 and L_2 for the corresponding polynomials κ_1 and κ_2 (and returns $\bot$ if this fails). It then searches list L_T for the term $\kappa_1 \cdot \kappa_2$ in the second component. If the search is successful, it returns the corresponding encoding K, otherwise it generates a new random encoding distinct from all others in L_T, returns this, and updates the list with $(K, \kappa_1 \cdot \kappa_2)$.

- **AttIssue Oracle.** Upon input of random encodings (X, Y), the challenger looks up the corresponding polynomial κ_X and κ_Y in L_1. If the search fails, return $\bot$. Otherwise, it selects $t_1, t_2 \leftarrow \mathbb{Z}_p^*$ and computes the polynomials for the tuple $(A, B, T_1, T_2, T_3, T_4) := (h_2^{x_0 t_1} h_5^{x_0 t_2} h_2^{x t_1} h_5^{y t_1}, h_5^{y_0 t_1} h_4^{y_0 t_2} h_5^{x t_2} h_4^{y t_2}, f_1^{t_1}, f_2^{t_2}, g_2^{t_1}, g_2^{t_2}) \in \mathbb{G}_1^2 \times \mathbb{G}_2^4$. If any of $\kappa_A, \kappa_B, \kappa_{T_1}, \kappa_{T_2}$ exist in L_1 and L_2 respectively, then return the paired encoding K for that polynomial. Otherwise, generate and return a new random encoding distinct from all elements in the list.

Eventually, $\mathcal{A}$ stops and outputs the tuple $(\hat{A}, \hat{B}, \hat{X}, \hat{Y}, \hat{T_1}, \hat{T_2}, \hat{T_3}, \hat{T_4}) \in \mathbb{G}_1^4 \times \mathbb{G}_2^4$, where $(\hat{X}, \hat{Y})$ was not queried. Let $(K_A, K_B, K_{T_1}, K_{T_2})$ be the polynomials that represent the elements in the tuple. We assume SXDH [1], is hard in $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$

---

[1] We note that we are not introducing additional complexity assumptions here as our construction for `HABS` already assumes SXDH for the instantiation of Groth-Sahai proofs.

and therefore no isomorphism exists and aim to find a contradiction.

We start by noting that the forgery must be a solution to the verification equations under the restriction of the check on (T_1, T_2, T_3, T_4). Thus we have the following:

$$K_A - (SK_{T_1} + RK_{T_2})SX_0 - (S\hat{X} + R\hat{Y})SK_{T_1} = 0 \bmod p \quad (1_A)$$
$$K_B - (SK_{T_1} + RK_{T_2})RY_0 - (S\hat{X} + R\hat{Y})RK_{T_2} = 0 \bmod p \quad (1_B)$$

$$K_A = P_0^A + P_{1,i}^A X_i + P_{2,i}^A Y_i + P_{3,i}^A A_i + P_{4,i}^A B_i + P_{5,j}^A H_j + P_{6,l}^A F_k + P_7^A SX_0 + P_8^A RY_0$$
$$K_B = P_0^B + P_{1,i}^B X_i + P_{2,i}^B Y_i + P_{3,i}^B A_i + P_{4,i}^B B_i + P_{5,j}^B H + P_{6,l}^B F + P_7^B SX_0 + P_8^B RY_0$$
$$K_{T_j} = P_0^{T_j} + P_{1,i}^{T_j} ST_i^{(1)} + P_{2,i}^{T_j} RT_i^{(2)} + P_{3,i}^{T_j} T_i^{(1)} + P_{4,i}^{T_j} T_i^{(2)} + P_5^{T_j} F_1 + P_6^{T_j} F_2$$

Substituting these and equating coefficients immediately gives:

$$P_0^A = 0, \; P_{1,i}^A = 0, \; P_{2,i}^A = 0, \; P_{4,i}^A = 0, \; P_{5,j}^A = 0, \; P_{6,l}^A = 0, P_7^A = 0, P_8^A = 0,$$
$$P_0^B = 0, \; P_{1,i}^B = 0, \; P_{2,i}^B = 0, \; P_{3,i}^B = 0, \; P_5^B = 0, \; P_6^B = 0, P_7^B = 0, P_8^B = 0,$$
$$P_0^{T_j} = 0, \; P_{1,i}^{T_j} = 0, \; P_{2,i}^{T_j} = 0, \; P_{3,i}^{T_j} = 0, \; P_{4,i}^{T_3} = 0, P_5^{T_j} = 0, P_6^{T_j} = 0,$$
$$\forall i \in \{1, ..., q_O\} \text{ and } i \neq j.$$

Simplifying:

$$P_{3,i}^A A_i - (SP_{1,i}^{T_1} T_{1,i} + RP_{2,i}^{T_2} T_{2,i})SX_0 - (S\hat{X} + R\hat{Y})SP_{1,i}^{T_1} T_{1,i} = 0 \bmod p \quad (2_A)$$
$$P_{3,i}^B B_i - (SP_{1,i}^{T_1} T_{1,i} + RP_{2,i}^{T_2} T_{2,i})RY_0 - (S\hat{X} + R\hat{Y})RP_{1,i}^{T_1} T_{2,i} = 0 \bmod p \quad (2_B)$$

$$P_{3,i}^A \big((ST_{1,i} + RT_{2,i})SX_0 - (SX_i + RY_i)ST_{1,i}\big) -$$
$$(SP_{1,i}^{T_1} T_{1,i} + RP_{2,i}^{T_2} T_{2,i})SX_0 - (S\hat{X} + R\hat{Y})SP_{1,i}^{T_1} T_{1,i} = 0 \bmod p \quad (3_A)$$

$$P_{4,i}^B \big((ST_{1,i} + RT_{2,i})RY_0 - (SX_i + RY_i)ST_{2,i}\big) -$$
$$(SP_{1,i}^{T_1} T_{1,i} + RP_{2,i}^{T_2} T_{2,i})RY_0 - (S\hat{X} + R\hat{Y})RP_{1,i}^{T_1} T_{2,i} = 0 \bmod p \quad (3_B)$$

Therefore we can make the equality:

$$P_{3,i}^A = P_{1,i}^{T_1} = P_{2,i}^{T_2} = P_{4,i}^B \qquad \forall i \in \{1, ..., q_O\}$$

Which then simplifies $(3_A), (3_B)$ to:

$$P_{3,i}^A (SX_i + RY_i)ST_{1,i} - (S\hat{X} + R\hat{Y})SP_{3,i}^A T_{1,i} = 0 \bmod p \quad (4_A)$$
$$P_{4,i}^B (SX_i + RY_i)ST_{2,i} - (S\hat{X} + R\hat{Y})SP_{4,i}^B T_{2,i} = 0 \bmod p \quad (4_B)$$

These equations can only hold if $i \neq 0$ for exactly one i, thus $\hat{X} = X_i$ and $\hat{Y} = Y_i$ for some $i \in \{1, ..., q_O\}$, which is a contradiction to our assumption that $(\hat{X}, \hat{Y})$ had not been queried by the adversary. Hence, our assumption holds in the generic group model.

Analysis of $\mathcal{B}$'s Simulation. The game is initialised with $(x_0, y_0, r, s) \in \mathbb{Z}_p^*$. $\mathcal{A}$ can only tell it is in a simulation if for any two polynomials $\kappa_{i,j} \neq \kappa_{i,j'}$ in L_i, we have $\kappa_{i,j}(x_0, y_0, r, s, t_{1,i}, t_{2,i}) = \kappa_{i,j'}(x_0, y_0, r, s, t_{1,i}, t_{2,i})$. This can only occur

if the randomly sampled instantiation of elements coincide with the roots of any of the polynomials. We use Lemma 1 in [35] to bound this probability above by the maximum total degree of the polynomial divided by p.

$$\epsilon \leq \binom{q_G + q_P + 6q_O + 11}{2} \frac{5}{p} \leq \frac{5(q_G + q_P + 6q_O + 11)^2}{p}.$$

To conclude, the advantage of an adversary against Assumption 1 is bounded in the GGM by $\epsilon \leq 5(q_G + q_P + 4q_O + 11)^2/p$. $\qquad\square$

B Construction of NIZK Proofs

Following the structure given in [21], we first define:

Step 0. Sample $R_1 \leftarrow \mathrm{Mat}_{|\Psi| \times 1}(\mathbb{Z}_p)$ and $R_2 \leftarrow \mathrm{Mat}_{|\Psi| \times 3}(\mathbb{Z}_p)$.

Step 1. Commit to $\mathbf{z}$ with tpk by computing, for every $z_i \in \mathbf{z}$.

Step 2. Next construct proofs for $\sum_i z_i S_{i,1} = 1$ and $\sum_i z_i S_{i,j} = 0$ for $j \neq 1$. These are linear equations. For each equation, compute: $\phi_j^{(1)} := R_1^T S_j \in \mathbb{Z}_p$.

Step 3. We next link the commitments $(c(X), c(Y))$ to public keys (X, Y) to the attributes $a_i \in \Psi$.

Case 1: Attribute a_i is a delegated attribute. For each public key of each attribute a_i in policy, we need to prove the multi-scalar multiplication equations $\hat{X} = X^{s_i}$ and $\hat{Y} = Y^{s_i}$.

Case 2: Attribute a_i is not an issued attribute. Simulate commitments to (X, Y) by running $\mathtt{TBE.Enc}(1_{\mathbb{G}})$ and prove the multi-scalar multiplication equations $\hat{X} = X^{s_i}$ and $\hat{Y} = Y^{s_i}$ where s_i is committed to in Step 1. For each a_i not issued, we require $4k$ commitments to $1_{\mathbb{G}}$ where k is the depth of the delegation path.

Step 4. Use the commitment for s_i from step 1, the commitments to (X_i, Y_i) from step 3 or 4 and R_1' is defined to be R_1 with an appended 0-row. The proof ϕ is computed as:

$$\phi_i^{(2)} := R_i^{(1)} X_i$$
$$\phi_i^{(3)} := R_i^{(2)} Y_i$$

Step 6. Show the correct opening $(a, b) := (\tilde{a} \cdot a_0^{|\Psi|-|A|}, \tilde{b} \cdot b_0^{|\Psi|-|A|})$ to $(c^{|\Psi|+1}, d^{|\Psi|+1})$ by proving the following 2 product pairing equations $(\mathcal{A}_1 \cdot \Gamma_1 \mathcal{A}_2) = c^{|\Psi|+1}$ and $(\mathcal{A}_1 \cdot \Gamma_2 \mathcal{A}_2) = d^{|\Psi|+1}$. The vectors $\mathcal{A}_1$ contain $apk \in \mathbf{warr}$. and contains $\mathcal{A}_2$ randomisers $T \in \mathbf{warr}$. The sparse matrices Γ_1 and Γ_2 are constructed as follows, where $1 \leq i, j \leq 2n|\Psi|$.

$$\Gamma_{i,j}^{(1)} := \begin{cases} 1 & \text{if } i \in 2\mathbb{N}_0 + 1 \text{ and } j = i+2 \text{ or } j = i+3 \\ 0 & \text{otherwise} \end{cases}$$

$$\Gamma_{i,j}^{(2)} := \begin{cases} 1 & \text{if } i \in 2\mathbb{N}_0 \text{ and } j = i+2 \text{ or } j = i+3 \\ 0 & \text{otherwise} \end{cases}$$

The vector of group elements $\mathcal{A}_1, \mathcal{A}_2 \in \mathbb{G}_1^{2n|\Psi|+2} \times \mathbb{G}_2^{2n|\Psi|+2}$ is given by X_0, Y_0 and $T_{1,0}$, $T_{2,0}$ in the first position, for $\mathcal{A}_1$ and $\mathcal{A}_2$ respectively, and $g_2^{\mathcal{H}_3(otsvk)}$ in position $2n|\Psi|+2$ for both $\mathcal{A}_2$. The proofs can be computed as such:

$$\phi^{(4)} := R^T \Gamma \iota_2(\mathcal{A}_1) + (R^T \Gamma_1 S - T^T)\mathbf{v}$$
$$\phi^{(5)} := S^T \Gamma^T \iota_1(\mathcal{A}_2) + T\mathbf{u}$$

Output the proof as $\pi := (\phi^{(1)}, \phi^{(2)}, \phi^{(3)}, \phi^{(4)}, \phi^{(5)})$.