

Group Signatures with Message-Dependent Opening in the Standard Model

Benoît Libert and Marc Joye

Technicolor

975 Avenue des Champs Blancs
35576 Cesson-Sévigné Cedex, France
{benoit.libert,marc.joye}@technicolor.com

Abstract. Group signatures allow members of a group to anonymously sign messages in the name of this group. They typically involve an opening authority that can identify the origin of any signature if the need arises. In some applications, such a tracing capability can be excessively strong and it seems desirable to restrict the power of the authority. Sakai *et al.* recently suggested the notion of group signatures with message-dependent opening (GS-MDO), where the opening operation is made contingent on the knowledge of a trapdoor information – generated by a second authority – associated with the message. Sakai *et al.* showed that their primitive implies identity-based encryption (IBE). In the standard model, efficiently constructing such a system thus requires a structure-preserving IBE scheme, where the plaintext space is the source group \mathbb{G} (rather than the target group \mathbb{G}_T) of a bilinear map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. Sakai *et al.* used a structure-preserving IBE which only provides bounded collusion-resistance. As a result, their GS-MDO construction only provides a weak form of anonymity where the maximal number of trapdoor queries is determined by the length of the group public key. In this paper, we construct the first fully collusion-resistant IBE scheme that encrypts messages in \mathbb{G} . Using this construction, we obtain a GS-MDO system with logarithmic signature size (in the number N of group members) and prove its security in the standard model under simple assumptions.

Keywords. Group signatures, message-dependent opening, efficiency, collusion-resistance, structure-preserving cryptography.

1 Introduction

Group signatures are central anonymity-related primitives, suggested by Chaum and van Heyst [21], which allow users to sign messages while hiding their identity within a population they belong to. They notably find applications in trusted computing platforms, auction protocols, anonymous subscription systems or in mechanisms for protecting the privacy of commuters in public transportation. To prevent users from abusing the system, group signatures usually involve an opening authority (OA) which is capable of identifying the signer using some trapdoor information. Although the opening authority can remain most frequently offline, group members have no privacy at all against this all powerful entity that can spy on all signature generations and identify the signer every time. To address this problem, Sakai *et al.* [36] advocated the design of a special kind of group signatures, called *group signatures with message-dependent opening* (GS-MDO), where restrictions are placed on the power of the OA. In the GS-MDO primitive, opening authorities cannot open any signature on their own. In order to open a signature on a message M , they need both their private key *and* a message-specific trapdoor t_M generated by a separate authority called *admitter*.

While the notion of group signatures dates back to Chaum and van Heyst [21], truly scalable and secure solutions remained elusive until the construction put forth by Ateniese *et al.* [6]. For lack of well-understood definitions, the security of their scheme was analyzed w.r.t. a list of sometimes redundant properties. A suitable security model was studied later on by Bellare, Micciancio and Warinschi [7] in the setting of static groups, where previous properties were subsumed by two security notions named *full anonymity* and *full traceability*. The case of dynamically growing groups was independently considered by Bellare, Shi and Zhang [9] and Kiayias and Yung [30].

During the last decade, a number of practical schemes were analyzed (e.g., [6, 30, 13, 33, 22]) in

the random oracle model [8], which is known [19] to only provide heuristic arguments in terms of security. While theoretical standard model constructions were given under general assumptions [7, 9], they were “only” proofs of concept. Viable constructions were suggested for the first time by Boyen and Waters [15, 16] and Groth [24, 25] who took advantage of breakthrough results [23, 26] in the construction of non-interactive zero-knowledge (NIZK) and witness indistinguishable (NIWI) proofs. The most efficient standard model realizations to date rely on the Groth-Sahai methodology [26], which is tailored to specific languages involving elements in bilinear groups.

GROUP SIGNATURES WITH MESSAGE-DEPENDENT OPENING. Traditional group signature models allow opening authorities to identify the originator of every single signature. As discussed by Sakai *et al.* [36], it may be desirable to restrict this extremely high power in many real-life applications.

One way to address this problem is to use techniques from threshold cryptography and share the opening key among several distributed opening authorities (as considered in, e.g., [10]) in such a way that none of these can individually open signatures and hurt the privacy of group members. While this approach may be sufficient in some applications, it requires the distributed openers to run a joint opening protocol whenever they want to trace a signature back to its source. In applications where many signatures on the same message have to be opened, this may become impractical. For example, suppose that group signatures are used to verify anonymous access rights to a parking or to enhance the privacy of users in public transportation systems: by issuing a group signature on a message consisting of the current date and time, users can demonstrate that they hold a valid credential and paid the subscription without being linkable to their previous rides. If a crime is committed, the police may want to find out who used a given metro line during a specific time interval. This requires a mechanism allowing for the opening of all signatures generated for a given date-time message and only those. Running a distributed opening protocol for each individual signature may be a bottleneck in this scenario. The same is true when group signatures are used in auction protocols: if group members are bidders who anonymously sign their bids, the threshold opening approach entails a communication cost proportional to the number of winners who offered the highest amount.

The above use cases motivated Sakai *et al.* [36] to formalize the notion of *group signatures with message-dependent opening* (GS-MDO), which splits the role of the opening authority between two entities called *opener* and *admitter*. In order to identify the author of a signature on a message M , the opener needs both its opening key ok and a trapdoor t_M generated by the admitter for the message M : the opening operation must be approved by the admitter, depending on the content of the message. Importantly, neither entity is powerful enough to open a signature by itself. A crucial difference with the aforementioned threshold opening approach is that, once a trapdoor t_M has been released for a sensitive message M , the opener can trace all signatures on M without any further interaction with the admitter.

We believe this message-dependent opening property to be of interest even in the setting of a centralized opening authority. Indeed, it features a complementary property to that of traceable signatures [29]. These involve opening authorities which can release a user-specific trapdoor allowing anyone to trace all signatures issued by a misbehaving group member. The GS-MDO primitive is important when the tracing criterion is the signed message (which could contain keywords associated with an illegal transaction) instead of the group member’s identity. Both techniques could actually be used in conjunction: one could first use a message-specific trapdoor to identify all group members who signed a suspicious message before tracing all other signatures created by these members.

RELATED WORK. Sakai *et al.* [36] gave a general construction of GS-MDO and notably showed that it implies Identity-Based Encryption [37, 14] (IBE): in their specific construction, the trusted authority naturally serves as an admitter and message-specific trapdoors are nothing but IBE private keys associated with the message. They also pointed out that, in order to build an efficient GS-MDO system in the standard model with the current state of knowledge in the area, they need a form of structure-preserving IBE scheme. Recall that a cryptographic primitive is called *structure-preserving*

(see [24, 20, 3, 1, 2, 18, 4] for examples) if it handles objects – like ciphertexts or signatures – that only consist of elements from a group \mathbb{G} over which a bilinear map is efficiently computable and if the validity of these objects can be checked using pairing-product equations. The latter properties make the primitive compatible with the Groth-Sahai techniques [26], which is crucial when one seeks to prove security in the standard model.

The main difficulty is that no structure-preserving IBE scheme is available to date: all pairing-based schemes proceed either by XORing the message with a hashed Bilinear Diffie-Hellman key [14] or encrypting messages that live in the target group \mathbb{G}_T of the bilinear map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ (see, e.g., [11, 38]). In order to construct an efficient GS-MDO in the standard model, what we need is an IBE scheme that encrypts messages in the domain group \mathbb{G} . We call such a system *partially* structure-preserving since identities do not have to be group elements and private keys can be ordinary (non-structure-preserving) signatures. For lack of a fully collusion-resistant such IBE, Sakai *et al.* [36] used a variant of the k -resilient construction of Heng and Kurosawa [28]: in the latter, semantic security is only guaranteed against adversaries that obtain private keys for no more than an *a priori* bounded number of identities. Moreover, the master public key has linear size in the pre-determined upper bound k . As a consequence, the standard model GS-MDO realization of [36] only achieves a relaxed flavor of security: namely, anonymity against the opener is only guaranteed as long as the adversary obtains trapdoors for at most k distinct messages. Moreover, the group public key inherits the $\mathcal{O}(k)$ size of the underlying IBE system.

In the random oracle model, Ohara *et al.* [34] recently proposed a construction allowing for an unbounded number of trapdoor queries. However, for the time being, building a fully secure GS-MDO system in the standard model remains an open problem.

OUR CONTRIBUTION. In this paper, we describe a GS-MDO system with $\mathcal{O}(\log N)$ size signatures, where N is the number of group members, and prove its security in the standard model under simple, constant-size assumptions (*i.e.*, we do not use q -type assumptions where the number of input elements depends on the number of adversarial queries or other system-related parameters).

As a result of independent interest, we describe the first fully collusion-resistant pairing-based IBE scheme that allows encrypting messages in the source group \mathbb{G} . This property is useful when it comes to proving properties about IBE-encrypted data: for example, the techniques of Camenisch *et al.* [17] can be used in combination with Groth-Sahai proofs to provide evidence that an IBE-encrypted plaintext belongs to a public set. Our system proceeds by blinding the plaintext $M \in \mathbb{G}$ using a random mask obtained by multiplying a random subset $\prod_{i \in S} Z_i$ of public elements $(Z_1, \dots, Z_\ell) \in \mathbb{G}^\ell$, where ℓ is proportional to the security parameter. The ℓ -bit string K identifying the subset S (so that $K[i] = 1$ if and only if $Z_i \in S$) is in turn encoded in a bit-wise manner using a variant of the Waters IBE scheme, each bit $K[i]$ of K being encoded as an independent IBE ciphertext entirely comprised of elements in \mathbb{G} . A consequence of this bit-by-bit encoding is that we need $\mathcal{O}(\ell)$ group elements to encrypt one element $M \in \mathbb{G}$. Despite its relatively large ciphertext size, our construction suffices to provide $\mathcal{O}(\log N)$ size signatures.

If we naively plug our IBE scheme into the general GS-MDO construction of Sakai *et al.* [36], we obtain signatures consisting of $\mathcal{O}(\lambda)$ group elements (or $\mathcal{O}(\lambda^2)$ bits), where λ is the security parameter, as each signature includes an IBE ciphertext. Fortunately, we can obtain signatures of only $\mathcal{O}(\log N)$ group elements – which is substantially shorter since $\log N \ll \lambda$ for any group of polynomial cardinality N – by combining the bit-wise encoding of our IBE scheme with the technique used in the Boyen-Waters group signature [15]. In the latter, membership certificates consist of Waters signatures $(g^\omega \cdot (v_0 \cdot \prod_{j=1}^{\ell} v_j^{\text{id}[j]})^r, g^r)$ on the group members' identifiers $\text{id} \in \{0, 1\}^\ell$, where $\ell = \log N$, and each group signature contains commitments to the individual bits $\text{id}[j]$ of id as well as NIWI proofs showing that committed values are actually bits. Our idea is thus to encode each bit $\text{id}[j]$ of id using a structure-preserving identity-based bit encryption scheme where the receiver's identity is the message to be signed. In order to guarantee anonymity against the admitter, we follow [36] and

super-encrypt each IBE ciphertext under the opener’s public key using a CCA2-secure public-key cryptosystem. For groups of $N = 10^6$ users, we eventually obtain signatures of 68 kB at the 128-bit security level, which is approximately twice the signature length of the k -resilient scheme of [36].

ORGANIZATION. In the forthcoming sections, we first recall the syntax and the security definitions of group signatures with message-dependent opening in Section 2. Section 3 describes our structure-preserving IBE system and our GS-MDO scheme is detailed in Section 4.

2 Background

2.1 Bilinear Maps and Complexity Assumptions

We use bilinear maps $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ over groups of prime order p where $e(g, h) \neq 1_{\mathbb{G}_T}$ if and only if $g, h \neq 1_{\mathbb{G}}$. In these groups, we rely on two hardness assumptions that are both non-interactive and stated using a constant number of elements.

Definition 1 ([13]). *The Decision Linear (DLIN) Problem in \mathbb{G} , is to distinguish the distributions $(g^a, g^b, g^{ac}, g^{bd}, g^{c+d})$ and $(g^a, g^b, g^{ac}, g^{bd}, g^z)$, with $a, b, c, d \xleftarrow{R} \mathbb{Z}_p$, $z \xleftarrow{R} \mathbb{Z}_p$. The Decision Linear assumption is the intractability of DLIN for any PPT distinguisher.*

Definition 2 ([14]). *The Decision 3-party Diffie-Hellman (D3DH) Problem in \mathbb{G} , is to distinguish the distributions $(g, g^a, g^b, g^c, g^{abc})$ and (g, g^a, g^b, g^c, g^z) , where $a, b, c, z \xleftarrow{R} \mathbb{Z}_p$.*

2.2 Groth-Sahai Proof Systems

Groth-Sahai (GS) proofs [26] can be based on the DLIN assumption, where they use prime order groups and a common reference string containing three vectors $\vec{f}_1, \vec{f}_2, \vec{f}_3 \in \mathbb{G}^3$, where $\vec{f}_1 = (f_1, 1, g)$, $\vec{f}_2 = (1, f_2, g)$ for some $f_1, f_2 \in \mathbb{G}$. To commit to $X \in \mathbb{G}$, one chooses $r, s, t \xleftarrow{R} \mathbb{Z}_p$ and computes $\vec{C} = (1, 1, X) \cdot \vec{f}_1^r \cdot \vec{f}_2^s \cdot \vec{f}_3^t$. In the soundness setting, we have $\vec{f}_3 = \vec{f}_1^{\xi_1} \cdot \vec{f}_2^{\xi_2}$ where $\xi_1, \xi_2 \in \mathbb{Z}_p$. Commitments $\vec{C} = (f_1^{r+\xi_1 t}, f_2^{s+\xi_2 t}, X \cdot g^{r+s+t(\xi_1+\xi_2)})$ are then extractable using $\beta_1 = \log_g(f_1)$, $\beta_2 = \log_g(f_2)$. In the witness indistinguishability (WI) setting, $\vec{f}_1, \vec{f}_2, \vec{f}_3$ are linearly independent and \vec{C} is a perfectly hiding commitment. Under the DLIN assumption, the two kinds of CRS are indistinguishable.

To commit to an exponent $x \in \mathbb{Z}_p$, the prover computes $\vec{C} = \vec{\varphi}^x \cdot \vec{f}_1^r \cdot \vec{f}_2^s$, where $r, s \xleftarrow{R} \mathbb{Z}_p$, using a CRS consisting of vectors $\vec{\varphi}, \vec{f}_1, \vec{f}_2$. In the perfect soundness setting, $\vec{\varphi}, \vec{f}_1, \vec{f}_2$ are linearly independent while, in the perfect WI setting, choosing $\vec{\varphi} = \vec{f}_1^{\xi_1} \cdot \vec{f}_2^{\xi_2}$ gives a perfectly hiding commitment.

To prove that committed variables satisfy a set of relations, the prover computes one commitment per variable and one proof element per relation. Such non-interactive witness indistinguishable (NIWI) proofs are available for pairing-product equations, which are equations of the form

$$\prod_{i=1}^n e(\mathcal{A}_i, \mathcal{X}_i) \cdot \prod_{i=1}^n \cdot \prod_{j=1}^n e(\mathcal{X}_i, \mathcal{X}_j)^{a_{ij}} = t_T, \quad (1)$$

for variables $\mathcal{X}_1, \dots, \mathcal{X}_n \in \mathbb{G}$ and constants $t_T \in \mathbb{G}_T$, $\mathcal{A}_1, \dots, \mathcal{A}_n \in \mathbb{G}$, $a_{ij} \in \mathbb{Z}_p$, for $i, j \in \{1, \dots, n\}$. Efficient NIWI proofs are also available for multi-exponentiation equations, which are of the form $\prod_{i=1}^m \mathcal{A}_i^{y_i} \cdot \prod_{j=1}^n \mathcal{X}_j^{b_j} \cdot \prod_{i=1}^m \cdot \prod_{j=1}^n \mathcal{X}_j^{y_i \gamma_{ij}} = T$, for variables $\mathcal{X}_1, \dots, \mathcal{X}_n \in \mathbb{G}$, $y_1, \dots, y_m \in \mathbb{Z}_p$ and constants $T, \mathcal{A}_1, \dots, \mathcal{A}_m \in \mathbb{G}$, $b_1, \dots, b_n \in \mathbb{Z}_p$ and $\gamma_{ij} \in \mathbb{G}$, for $i \in \{1, \dots, m\}, j \in \{1, \dots, n\}$.

In pairing-product equations, proofs for quadratic equations require 9 group elements whereas linear equations (*i.e.*, where $a_{ij} = 0$ for all i, j in equation (1)) only cost 3 group elements each. Linear multi-exponentiation equations of the type $\prod_{i=1}^m \mathcal{A}_i^{y_i} = T$ require 2 group elements.

2.3 Group Signatures with Message-Dependent Opening

We use the syntax of [36], which extends the static model of Bellare, Micciancio and Warinschi [7].

Keygen(λ, N): given a security parameter $\lambda \in \mathbb{N}$ and a maximal number of group members $N \in \mathbb{N}$, this algorithm outputs a group public key \mathbf{gpk} , a vector $\mathbf{gsk} = (\mathbf{gsk}[0], \dots, \mathbf{gsk}[N-1])$ of group members' private keys as well as private keys $\mathbf{msk}_{\text{ADM}}$ and ok for the admitter and the opener.

Sign: takes as input a message M , a private key $\mathbf{gsk}[i]$ and \mathbf{gpk} , it outputs a signature σ .

Verify: is a deterministic algorithm taking as input a signature σ , a message M and a group public key \mathbf{gpk} . It returns either 0 or 1.

TrapGen: is a possibly randomized algorithm that takes as input the admitter's private key $\mathbf{msk}_{\text{ADM}}$ and a message M . It outputs a trapdoor t_M allowing the OA to open all signatures on M .

Open: takes as input a message M , a valid signature σ w.r.t. \mathbf{gpk} , the opening authority's private key ok and a trapdoor t_M for the message M . It outputs $i \in \{0, \dots, N-1\} \cup \{\perp\}$, which is either the index of a group member or a symbol indicating an opening failure.

Definition 3. A GS-MDO scheme provides full traceability if, for any $\lambda \in \mathbb{N}$, any $N \in \text{poly}(\lambda)$ and any PPT adversary \mathcal{A} involved in the experiment hereafter, it holds that

$$\text{Adv}_{\mathcal{A}}^{\text{trace}}(\lambda) = \Pr[\text{Exp}_{\mathcal{A}}^{\text{trace}}(\lambda, N) = 1] \in \text{negl}(\lambda).$$

$\text{Exp}_{\mathcal{A}}^{\text{trace}}(n, N)$
 $(\mathbf{gpk}, \text{ok}, \mathbf{msk}_{\text{ADM}}, \mathbf{gsk}) \leftarrow \text{Keygen}(\lambda, N)$
 $\text{st} \leftarrow (\text{ok}, \mathbf{msk}_{\text{ADM}}, \mathbf{gpk})$; $\mathcal{C} \leftarrow \emptyset$; $K \leftarrow \varepsilon$; $\text{Cont} \leftarrow \text{true}$
while ($\text{Cont} = \text{true}$) **do**
 $(\text{Cont}, \text{st}, j) \leftarrow \mathcal{A}^{\text{Sign}(\mathbf{gsk}[\cdot], \cdot)}(\text{choose}, \text{st}, K)$
 if $\text{Cont} = \text{true}$ **then** $\mathcal{C} \leftarrow \mathcal{C} \cup \{j\}$; $K \leftarrow K \cup \{\mathbf{gsk}[j]\}$ **end if**
end while
 $(M^*, \sigma^*) \leftarrow \mathcal{A}^{\text{Sign}(\mathbf{gsk}[\cdot], \cdot)}(\text{guess}, \text{st})$
if $\text{Verify}(\mathbf{gpk}, M^*, \sigma^*) = 0$ **then** *Return* 0
if $\text{Open}(\mathbf{gpk}, \text{ok}, \text{TrapGen}(\mathbf{gpk}, \mathbf{msk}_{\text{ADM}}, M^*), M^*, \sigma^*) = \perp$ **then** *Return* 1
if $\exists j^* \in \{0, \dots, N-1\}$ *such that*
 $(\text{Open}(\mathbf{gpk}, \text{ok}, t_{M^*}, M^*, \sigma^*) = j^*) \wedge (j^* \notin \mathcal{C}) \wedge ((j^*, M^*) \text{ not queried by } \mathcal{A})$
 with $t_{M^*} \leftarrow \text{TrapGen}(\mathbf{gpk}, \mathbf{msk}_{\text{ADM}}, M^*)$
then *Return* 1
else *Return* 0

Definition 4. A GS-MDO scheme provides full anonymity against the admitter if, for any $\lambda \in \mathbb{N}$, any $N \in \text{poly}(\lambda)$ and any PPT adversary \mathcal{A} , the function

$$\text{Adv}_{\mathcal{A}}^{\text{anon-adm}}(\lambda) = |\Pr[\text{Exp}_{\mathcal{A}}^{\text{anon-adm}}(\lambda, N) = 1] - 1/2| \in \text{negl}(\lambda)$$

is a negligible function in the security parameter if the experiment proceeds as follows

$\text{Exp}_{\mathcal{A}}^{\text{anon-adm}}(\lambda, N)$
 $(\mathbf{gpk}, \text{ok}, \mathbf{msk}_{\text{ADM}}, \mathbf{gsk}) \leftarrow \text{Keygen}(\lambda, N)$
 $(\text{st}, j_0, j_1, M^*) \leftarrow \mathcal{A}^{\text{Ok}}(\text{choose}, \mathbf{gpk}, \mathbf{gsk}, \mathbf{msk}_{\text{ADM}})$
 $b \xleftarrow{R} \{0, 1\}$; $\sigma^* \leftarrow \text{Sign}(\mathbf{gpk}, \mathbf{gsk}[j_b], M^*)$
 $b' \leftarrow \mathcal{A}^{\text{Ok}}(\text{guess}, \text{st}, \sigma^*)$
Return 1 if $b' = b$ and 0 otherwise

In the above notation, \mathcal{O}_{ok} denotes an oracle that takes as input any adversarially chosen signature $\sigma \neq \sigma^*$ and uses ok and $\mathbf{msk}_{\text{ADM}}$ to determine and return the identity of the signer.

Definition 5. A GS-MDO scheme provides full anonymity against the opener if, for any $\lambda \in \mathbb{N}$, any $N \in \text{poly}(\lambda)$ and any PPT adversary \mathcal{A} , the function

$$\text{Adv}_{\mathcal{A}}^{\text{anon-oa}}(\lambda) = |\Pr[\text{Exp}_{\mathcal{A}}^{\text{anon-oa}}(\lambda, N) = 1] - 1/2| \in \text{negl}(\lambda)$$

is a negligible function in the security parameter if the experiment goes as follows

$\text{Exp}_{\mathcal{A}}^{\text{anon-oa}}(\lambda, N)$
 $(\text{gpk}, \text{ok}, \text{msk}_{\text{ADM}}, \text{gsk}) \leftarrow \text{Keygen}(\lambda, N)$
 $(\text{st}, j_0, j_1, M^*) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{msk}_{\text{ADM}}}}(\text{choose}, \text{gpk}, \text{gsk}, \text{ok})$
 $b \xleftarrow{R} \{0, 1\}; \quad \sigma^* \leftarrow \text{Sign}(\text{gpk}, \text{gsk}[j_b], M^*)$
 $b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{msk}_{\text{ADM}}}}(\text{guess}, \text{st}, \sigma^*)$
 Return 1 if $b' = b$ and 0 otherwise

In the above notation, $\mathcal{O}_{\text{msk}_{\text{ADM}}}(\cdot)$ is an oracle that returns trapdoors for arbitrary messages $M \neq M^*$ chosen by the adversary.

3 A Fully Collusion-Resistant Partially Structure-Preserving IBE

3.1 Intuition

The scheme is only partially structure-preserving in that identities are still encoded as binary strings and private keys are ordinary signatures (recall that, in any IBE, private keys are signatures on the corresponding identity, as mentioned in [14]) instead of structure-preserving ones. It can be seen as a variant of Waters' IBE [38] (see Appendix A for syntactic definitions) and builds on a consequence of the Leftover Hash Lemma [27]: namely, if $\ell > 2 \log_2(p)$ and $a_1, \dots, a_\ell \in_R \mathbb{Z}_p$ are uniformly distributed in \mathbb{Z}_p , then random subset sums $\sum_{i=1}^{\ell} \beta_i a_i$ with $(\beta_1, \dots, \beta_\ell) \in_R \{0, 1\}^\ell$ are statistically indistinguishable from uniformly random values in \mathbb{Z}_p .

The idea is to include a vector $(Z_1, \dots, Z_\ell) \in \mathbb{G}^\ell$ in the master public key. The message $M \in \mathbb{G}$ will be encrypted by choosing a random ℓ -bit string $K \in \{0, 1\}^\ell$ and multiplying M with a product of elements in the set $S = \{Z_i \mid K[i] = 1\}$. Then, each bit $K[i]$ of K will be individually encrypted using a variant of the Waters IBE. In the latter variant, an encryption of 1 will consist of a tuple $(C_{i,1}, C_{i,2}, C_{i,3}, C_{i,4}) = (g^{s_i}, H_{\mathbb{G}}(\text{ID})^{s_i}, g_1^{s_i/\omega_i}, g_2^{\omega_i})$, where $s_i, \omega_i \in_R \mathbb{Z}_p$. In an encryption of 0, the pair $(C_{i,3}, C_{i,4})$ is chosen uniformly in \mathbb{G}^2 . Upon decryption, the receiver can use his private key (d_1, d_2) to test whether the equality $e(C_{i,3}, C_{i,4}) = e(C_{i,1}, d_1)/e(C_{i,2}, d_2)$ holds. If it does, the receiver decodes the i -th bit of K as $K[i] = 1$. Otherwise, it sets $K[i] = 0$. The security of the resulting scheme can be proved under the D3DH assumption (instead of the DBDH assumption).

Although the latter scheme allows encrypting messages in the group \mathbb{G} , it still does not provide all the properties we need for the problem at hand. When it comes to proving that a ciphertext encrypts a message that coincides with the content of Groth-Sahai commitment, the difficulty is to prove that the equality $e(C_{i,3}, C_{i,4}) = e(C_{i,1}, d_1)/e(C_{i,2}, d_2)$ is *not* satisfied when $K[i] = 0$. For this reason, we need to modify the scheme as suggested in Section 3.2.

3.2 Construction

In order to be able to efficiently prove that a ciphertext and a Groth-Sahai commitment hide the same group element, we modify the scheme of Section 3.1 as follows. In the master public key, the element g_1 is replaced by a pair $(g_0, g_1) = (g^{\alpha_0}, g^{\alpha_1})$. The master secret key is twinned in the same way and now consists of $(g_2^{\alpha_0}, g_2^{\alpha_1})$. Likewise, each identity is assigned a private key of the form $(d_{0,1}, d_{0,2}, d_{1,1}, d_{1,2}) = (g_2^{\alpha_0} \cdot H_{\mathbb{G}}(\text{ID})^{r_0}, g^{r_0}, g_2^{\alpha_1} \cdot H_{\mathbb{G}}(\text{ID})^{r_1}, g^{r_1})$.

In the encryption algorithm, when the sender wants to “encrypt” a bit $K[i]$ of $K \in \{0, 1\}^\ell$, it generates $(C_{i,3}, C_{i,4})$ as $(C_{i,3}, C_{i,4}) = (g_{K[i]}^{s_i/\omega_i}, g_2^{\omega_i})$, so that the receiver can easily determine the value of $K[i]$ using his private key.

The modification will make it easier to prove equalities between the plaintext and a committed value. The reason is that the prover does not have to prove an inequality when $K[i] = 0$: he essentially has to prove statements of the form “ $(C_{i,3}, C_{i,4}) = (g_0^{s_i/\omega_i}, g_2^{\omega_i})$ OR $(C_{i,3}, C_{i,4}) = (g_1^{s_i/\omega_i}, g_2^{\omega_i})$ ”. Our construction of Groth-Sahai-compatible IBE thus goes follows.

Setup(λ) : Choose bilinear groups $(\mathbb{G}, \mathbb{G}_T)$ of prime order $p > 2^\lambda$. Then, do the following.

1. Choose $\alpha_0, \alpha_1 \xleftarrow{R} \mathbb{Z}_p$, $g \xleftarrow{R} \mathbb{G}$, $g_2 \xleftarrow{R} \mathbb{G}$ and set $g_0 = g^{\alpha_0}$, $g_1 = g^{\alpha_1}$.
2. Choose $u_0, u_1, \dots, u_L \xleftarrow{R} \mathbb{G}$, for a suitably large $L \in \text{poly}(\lambda)$. These will be used to implement a number-theoretic hash function $H_{\mathbb{G}} : \{0, 1\}^L \rightarrow \mathbb{G}$ such that any L -bit string $\tau = \tau[1] \dots \tau[L] \in \{0, 1\}^L$ is mapped to the value $H_{\mathbb{G}}(\tau) = u_0 \cdot \prod_{i=1}^L u_i^{\tau[i]}$.
3. Choose group elements $(Z_1, \dots, Z_\ell) \xleftarrow{R} \mathbb{G}^\ell$, where $\ell = 2\lceil \log_2(p) \rceil > 2\lambda$.

The master secret key is $\text{msk} := (g_2^{\alpha_0}, g_2^{\alpha_1})$ and the master public key is defined as

$$\text{mpk} = \left((\mathbb{G}, \mathbb{G}_T), p, g, g_0 = g^{\alpha_0}, g_1 = g^{\alpha_1}, g_2, \{u_i\}_{i=0}^L, \{Z_i\}_{i=1}^\ell \right)$$

Keygen(msk, ID) : given the master secret key $\text{msk} = (g_2^{\alpha_0}, g_2^{\alpha_1})$ and an identity $\text{ID} \in \{0, 1\}^L$, choose $r_0, r_1 \xleftarrow{R} \mathbb{Z}_p$ to compute and return

$$d_{\text{ID}} = (d_{0,1}, d_{0,2}, d_{1,1}, d_{1,2}) = (g_2^{\alpha_0} \cdot H_{\mathbb{G}}(\text{ID})^{r_0}, g^{r_0}, g_2^{\alpha_1} \cdot H_{\mathbb{G}}(\text{ID})^{r_1}, g^{r_1}).$$

Encrypt(mpk, ID, M) : to encrypt a message $M \in \mathbb{G}$, conduct the following steps.

1. Choose a random ℓ -bit string $K \xleftarrow{R} \{0, 1\}^\ell$, where $\ell = 2\log_2(p)$.
2. Choose $s_1, \dots, s_\ell \xleftarrow{R} \mathbb{Z}_p$ and $\omega_1, \dots, \omega_\ell \xleftarrow{R} \mathbb{Z}_p$.
3. Parse K as $K[1] \dots K[\ell] \in \{0, 1\}^\ell$. For $i = 1$ to ℓ , compute

$$C_{i,1} = g^{s_i} \quad C_{i,2} = H_{\mathbb{G}}(\text{ID})^{s_i} \quad C_{i,3} = g_{K[i]}^{s_i/\omega_i} \quad C_{i,4} = g_2^{\omega_i} \quad (2)$$

4. Then, compute $C_0 = M \cdot \prod_{i=1}^\ell Z_i^{K[i]}$.

Return the ciphertext $C = (C_0, \{(C_{i,1}, C_{i,2}, C_{i,3}, C_{i,4})\}_{i=1}^\ell) \in \mathbb{G}^{4\ell+1}$.

Decrypt($\text{mpk}, d_{\text{ID}}, C$) : parse C as $C = (C_0, \{(C_{i,1}, C_{i,2}, C_{i,3}, C_{i,4})\}_{i=1}^\ell)$.

1. For $i = 1$ to ℓ compute $\mu_b = e(C_{i,1}, d_{b,1})/e(C_{i,2}, d_{b,2})$ for each $b \in \{0, 1\}$. If there exists $b \in \{0, 1\}$ such that $\mu_b = e(C_{i,3}, C_{i,4})$, set $K[i] = b$. Otherwise, return \perp .
2. Compute and return $M = C_0 / (\prod_{i=1}^\ell Z_i^{K[i]})$.

Unlike the IBE system of Sakai *et al.* [36], the above scheme provides full collusion-resistance and the size of the master public key only depends on the security parameter and not on a pre-determined bound on the number of corrupted users.

Theorem 1. *The above IBE scheme provides IND-ID-CPA security under the D3DH assumption.*

Proof. To prove the result, we consider a sequence of games which begins with the real game and ends with a game where the adversary’s view is independent of the challenger’s bit $\beta \in \{0, 1\}$. For each i , we denote by S_i the event that the adversary wins in Game i and we define the adversary’s advantage as $\text{Adv}_i := |\Pr[S_i] - 1/2|$.

Game 0: This is the real attack game where the challenger generates a proper encryption of M_β , with $\beta \xleftarrow{R} \{0, 1\}$, in the challenge phase. The game ends with the adversary \mathcal{A} outputting $\beta' \in \{0, 1\}$ and we denote by S_0 the event that $\beta' = \beta$.

Game i ($1 \leq i \leq \ell$): In this game, the challenger generates the challenge ciphertext in a hybrid manner. Namely, for each $j \in \{1, \dots, \ell\}$, the challenger generates the ciphertext components $\{(C_{j,1}, C_{j,2}, C_{j,3}, C_{j,4})\}$ as follows.

- If $j \leq i$, it picks $s_j \xleftarrow{R} \mathbb{Z}_p$, computes $(C_{j,1}, C_{j,2}) = (g^{s_j}, H_{\mathbb{G}}(\text{ID})^{s_j})$ but chooses $(C_{j,3}, C_{j,4}) \xleftarrow{R} \mathbb{G}^2$ at random.
- If $j > i$, it runs the normal encryption algorithm and sets

$$(C_{j,1}, C_{j,2}, C_{j,3}, C_{j,4}) = (g^{s_j}, H_{\mathbb{G}}(\text{ID})^{s_j}, g_{K[j]}^{s_j/\omega_j}, g_2^{\omega_j})$$

for randomly chosen $s_j, \omega_j \xleftarrow{R} \mathbb{Z}_p$.

Game $\ell + 1$: This game is identical to Game ℓ with the difference that, in the challenge ciphertext, C_0 is chosen as a uniformly random $C_0 \xleftarrow{R} \mathbb{G}$ instead of being computed as $C_0 = M_{\beta} \cdot \prod_{j=1}^{\ell} Z_j^{K[j]}$.

For each $j \in \{1, \dots, \ell\}$, Lemma 1 shows that Game j is computationally indistinguishable from Game $j - 1$ if the D3DH assumption holds.

In Game ℓ , the ciphertext components $\{(C_{j,1}, C_{j,2}, C_{j,3}, C_{j,4})\}_{j=1}^{\ell}$ are completely uncorrelated to the string $K = K[1] \dots K[\ell] \in \{0, 1\}^{\ell}$ that is used to compute $C_0 = M_{\beta} \cdot \prod_{j=1}^{\ell} Z_j^{K[j]}$. For this reason, we argue that the adversary's view is statistically independent of M_{β} . This is easily seen by observing that the Leftover Hash Lemma implies that the two distributions

$$D_0 = \{(a, \langle a, z \rangle) \mid a \xleftarrow{R} \mathbb{Z}_p^{\ell}, z \xleftarrow{R} \{0, 1\}^{\ell}\} \quad D_1 = \{(a, w) \mid a \xleftarrow{R} \mathbb{Z}_p^{\ell}, w \xleftarrow{R} \mathbb{Z}_p\},$$

are statistically close when $\ell > 2 \log_2(p)$. Consequently, Game ℓ is statistically close to Game $\ell + 1$, where C_0 is replaced by a uniformly random group element in the challenge ciphertext. In the latter game, we clearly have $\Pr[S_{\ell+1}] = 1/2$ (and thus $\text{Adv}_{\ell+1} = 0$) since the challenge ciphertext is completely independent of M_{β} . \square

Lemma 1. *If the D3DH assumption holds, Game i is computationally indistinguishable from Game $i - 1$ for each $i \in \{1, \dots, \ell\}$. More precisely, if \mathcal{A} runs in time t and has significantly different advantages in Game i and Game $i - 1$, then there exists a PPT algorithm \mathcal{B} with running time $t + O(\varepsilon^{-2} \ln(\varepsilon^{-1}) \eta^{-1} \ln(\eta^{-1}))$ such that*

$$|\text{Adv}_i(\mathcal{A}) - \text{Adv}_{i-1}(\mathcal{A})| \leq 16 \cdot (L + 1) \cdot q \cdot \mathbf{Adv}^{\text{D3DH}}(\mathcal{B}),$$

where $\eta = 1/(4(L + 1)q)$ and q is the maximal number of private key queries. (The proof is given in Appendix B.)

We note that the same idea can be applied to construct other partially structure-preserving primitives. For example, it can be applied to selectively-secure attribute-based encryption schemes based on the Decision Bilinear Diffie-Hellman assumption [35]. We can also construct a structure-preserving variant of the Boneh-Boyen-Goh HIBE [12]: for a hierarchy of depth d , we have to rely on a slightly stronger assumption than [12], which posits the hardness of distinguishing $h^{(\alpha^{d+1})}$ from random given $(g, g^{\alpha}, \dots, g^{(\alpha^d)}, h) \in \mathbb{G}^{d+2}$.

3.3 Proving Properties about Encrypted Messages

Our solution retains the useful property of the scheme in [36] as it allows efficiently proving relations about the plaintext using the Groth-Sahai techniques. If $\vec{C}_M = (1, 1, M) \cdot \vec{f}_1^{r_M} \cdot \vec{f}_2^{s_M} \cdot \vec{f}_3^{t_M}$ denotes a Groth-Sahai commitment to a message $M \in \mathbb{G}$ which is also encrypted with the above IBE, the sender can proceed as follows to prove the equality between the committed message and the plaintext.

For each $i \in \{1, \dots, \ell\}$, the sender computes a commitment $\vec{C}_{K_i} = (1, 1, g^{K[i]}) \cdot \vec{f}_1^{r_{K[i]}} \cdot \vec{f}_2^{s_{K[i]}} \cdot \vec{f}_3^{t_{K[i]}}$

to the group element $K_i = g^{K[i]}$ and generates a non-interactive proof $\vec{\pi}_{K[i]}$ that $K[i] \in \{0, 1\}$. This is typically achieved by proving the equality $K[i]^2 = K[i] \pmod p$ with a proof $\vec{\pi}_{K[i]}$ consisting of 9 group elements. Next, the sender generates a commitment \vec{C}_{G_i} to the group element $G_i = g_{K[i]}$ and generates a non-interactive proof $\vec{\pi}_{G_i}$ that committed elements G_i and $K[i]$ satisfy the equality

$$G_i = g_1^{K[i]} \cdot g_0^{1-K[i]}$$

or, equivalently, $e(G_i, g) = e(g_1, K_i) \cdot e(g_0, K_i^{-1} \cdot g)$. The latter is a linear equation for which the proof $\vec{\pi}_{G_i}$ requires three group elements. Then, the sender generates a commitment \vec{C}_{Θ_i} to the auxiliary variable $\Theta_i = g^{s_i/\omega_i}$ and generate non-interactive proofs $\vec{\pi}_{\Theta_i,1}, \vec{\pi}_{\Theta_i,2}$ for the relations

$$e(\Theta_i, C_{i,4}) = e(C_{i,1}, g_2) \tag{3}$$

$$e(\Theta_i, G_i) = e(g, C_{i,3}). \tag{4}$$

Since (3) is a linear equation, $\vec{\pi}_{\Theta_i,1}$ only requires 3 group elements. On the other hand, (4) is quadratic, so that $\vec{\pi}_{\Theta_i,2}$ costs 9 group elements to prove.

Finally, the sender is left with proving the equality $e(C_0/M, g) = \prod_{i=1}^{\ell} e(Z_i, K_i)$, which is a linear equation whose proof $\vec{\pi}_{C_0}$ requires 3 group elements.

The whole NIWI proof $(\{\vec{C}_{K_i}, \vec{C}_{G_i}, \vec{C}_{\Theta_i}, \vec{\pi}_{K[i]}, \vec{\pi}_{G_i}, \vec{\pi}_{\Theta_i,1}, \vec{\pi}_{\Theta_i,2}\}_{i=1}^{\ell}, \vec{\pi}_{C_0})$ thus takes $35\ell + 3$ group elements overall.

In some cases, the above proof might have to be a NIZK (and not just NIWI) proof. In pairing-product equations, NIZK proofs are not known to always exist. Fortunately, we can solve this issue by introducing a constant number of extra variables, as we will see in Section 4.

4 A Fully Anonymous GS-MDO Scheme with Logarithmic-Size Signatures

Our construction departs from the general approach suggested in [36] in order to obtain shorter signatures. The signing algorithm of [36] proceeds by choosing two random session keys K^{PKE} and K^{IBE} which are separately encrypted using a CCA2-secure public-key encryption scheme and an IBE scheme, respectively. These two keys K^{PKE} and K^{IBE} are then used to hide the group member's credential in the fashion of nested multiple encryptions while adding a proof that the hidden value is a valid and properly encrypted credential. If we naively apply this approach using our IBE scheme, we will eventually obtain signatures consisting of $O(\lambda^2)$ bits, where λ is the security parameter.

To reduce the signature size to $O(\lambda \log N)$ bits (recall that $\log N \ll \lambda$ since the cardinality N of the group is assumed to be polynomial), we use a different approach. Instead of encrypting random session keys which conceal the group member's credential under two randomly generated session keys, we directly encrypt the bits of the group member's identity as if it were the session key K in the IBE scheme of Section 3.2. This allows reducing the number of bit-carrying IBE ciphertext components from $O(\lambda)$ to $O(\log N)$. In order to make sure that neither the admitter or the opening authority will be able to individually open any signature, we add a second encryption layer and additionally encrypt – under the admitter's public key using Kiltz's DLIN-based CCA2-secure encryption scheme [32] – the IBE ciphertext components which depend on the bits of the group member's identity.

The rest of the signing algorithm proceeds as in the Boyen-Waters group signature [15], by having the signer verifiably encrypt a two-level hierarchical signature [31], where the first-level (resp. second-level) message is the signer's identity (resp. the actual message). Like [15], we use a two-level hierarchical extension of Waters' signature [38].

4.1 Construction

Keygen(λ, N): given a security parameter $\lambda \in \mathbb{N}$ and $N = 2^\ell$,

1. Choose bilinear groups $(\mathbb{G}, \mathbb{G}_T)$ of prime order $p > 2^\lambda$, with $g \stackrel{R}{\leftarrow} \mathbb{G}$.
2. As a CRS for the Groth-Sahai NIWI proof system, select vectors $\mathbf{f} = (\vec{f}_1, \vec{f}_2, \vec{f}_3)$ such that $\vec{f}_1 = (f_1, 1, g) \in \mathbb{G}^3$, $\vec{f}_2 = (1, f_2, g) \in \mathbb{G}^3$, and $\vec{f}_3 = \vec{f}_1^{\xi_1} \cdot \vec{f}_2^{\xi_2}$, where $f_1 = g^{\beta_1}$, $f_2 = g^{\beta_2} \stackrel{R}{\leftarrow} \mathbb{G}$ and $\beta_1, \beta_2, \xi_1, \xi_2 \stackrel{R}{\leftarrow} \mathbb{Z}_p$. We also define the vector $\vec{\varphi} = \vec{f}_3 \cdot (1, 1, g)$.
3. Generate a master key pair $(\text{msk}_{\text{IBE}}, \text{mpk}_{\text{IBE}})$ for the identity-based key encapsulation scheme of Section 3.2¹. These consist of $\text{msk}_{\text{IBE}} = (g_2^{\alpha_0}, g_2^{\alpha_1})$ and

$$\text{mpk}_{\text{IBE}} = \left(g_0 = g^{\alpha_0}, g_1 = g^{\alpha_1}, g_2, \{u_i\}_{i=0}^L \right),$$

where $L \in \text{poly}(\lambda)$ denotes the length of (hashed) messages to be signed. For a message $M \in \{0, 1\}^L$, we define the function $H_U(M) \in \mathbb{G}$ as $H_U(M) = u_0 \cdot \prod_{i=1}^L u_i^{M[i]}$, where $M[i] \in \{0, 1\}$ denotes the i -th bit of M .

4. Generate a key pair (sk_W, pk_W) for a two-level hierarchical Waters signature. At level 1 (resp. level 2), messages will be of length ℓ (resp. L). This key pair consists of $sk_W = g^\omega$ and

$$pk_W = \left(e(g, g)^\omega, \{v_i\}_{i=0}^\ell, \{w_i\}_{i=0}^L \right),$$

where $\omega \in_R \mathbb{Z}_p$. Analogously to step 3, we denote by $H_W(M)$ the function that maps the message $M \in \{0, 1\}^L$ to $H_W(M) = w_0 \cdot \prod_{i=1}^L w_i^{M[i]}$, where $M[i] \in \{0, 1\}$ is the i -th bit of M .

5. For each $i \in \{0, \dots, N-1\}$ generate the private key $\text{gsk}[i]$ of member i as a Waters signature $\text{gsk}[i] = \left(g^\omega \cdot (v_0 \cdot \prod_{j=1}^\ell v_j^{\text{id}_i[j]})^r, g^r \right)$, with $r \stackrel{R}{\leftarrow} \mathbb{Z}_p$, on the message $\text{id}_i = \text{id}_i[1] \dots \text{id}_i[\ell] \in \{0, 1\}^\ell$ which is obtained as the binary expansion of $i \in \{0, \dots, N-1\}$. The private key $sk_W = g^\omega$ is not needed beyond this point and can be erased after the generation of the vector of private keys $\mathbf{gsk} = (\text{gsk}[0], \dots, \text{gsk}[N-1])$.
6. Generate a public key $(X, Y, U, V) = (g^{\beta_x}, g^{\beta_y}, g^{\beta_u}, g^{\beta_v})$, with random $\beta_x, \beta_y, \beta_u, \beta_v \stackrel{R}{\leftarrow} \mathbb{Z}_p$, for Kiltz's CCA2-secure encryption scheme.
7. Select a strongly unforgeable one-time signature scheme $\Sigma = (\mathcal{G}, \mathcal{S}, \mathcal{V})$.

The admitter's message specification key consists of $\text{msk}_{\text{ADM}} := \text{msk}_{\text{IBE}}$. The private key ok of the opening authority is defined as $\text{ok} := (\beta_x, \beta_y, \beta_u, \beta_v)$. The private key of member i is $\text{gsk}[i]$ while the group public key is defined to be

$$\text{gpk} := \left((\mathbb{G}, \mathbb{G}_T), p, g, \mathbf{f} = (\vec{f}_1, \vec{f}_2, \vec{f}_3), \text{mpk}_{\text{IBE}}, pk_W, (X, Y, U, V), \Sigma \right)$$

Sign($\text{gpk}, \text{gsk}[i], M$): to sign a message $M \in \{0, 1\}^L$ using the i -th group member's private key $\text{gsk}[i] = (S_{i,1}, S_{i,2}) = (g^\omega \cdot (v_0 \cdot \prod_{j=1}^\ell v_j^{\text{id}_i[j]})^r, g^r)$, generate a one-time signature key pair $(\text{SK}, \text{VK}) \leftarrow \Sigma.\mathcal{G}(\lambda)$ and do the following.

1. Generate a two-level Waters signature where the message is $\text{id}_i \in \{0, 1\}^\ell$ at the first level and $M \in \{0, 1\}^L$ at level 2. The signature consists of

$$\begin{aligned} (\Omega_1, \Omega_2, \Omega_3) &= \left(S_{i,1} \cdot (v_0 \cdot \prod_{j=1}^\ell v_j^{\text{id}_i[j]})^{r'} \cdot H_W(M)^s, S_{i,2} \cdot g^{r'}, g^s \right) \\ &= \left(g^\omega \cdot (v_0 \cdot \prod_{j=1}^\ell v_j^{\text{id}_i[j]})^{r''} \cdot H_W(M)^s, g^{r''}, g^s \right), \end{aligned}$$

where $r'' = r + r'$.

¹ Note that the $\{Z_i\}_{i=1}^\ell$ components are not needed here and can be discarded.

2. Generate a commitment \vec{C}_{H_V} to $H_V = v_0 \cdot \prod_{j=1}^{\ell} v_j^{\text{id}_i[j]}$. Then, for each $j \in \{1, \dots, \ell\}$, generate a commitment \vec{C}_{F_j} to $F_j = g^{\text{id}_i[j]}$ and generate a NIWI proof $\vec{\pi}_{H_V} \in \mathbb{G}^3$ that

$$e(H_V, g) \cdot \prod_{j=1}^{\ell} e(v_j, F_j)^{-1} = e(v_0, g) \quad (5)$$

Since (5) is a linear equation, $\vec{\pi}_{H_V}$ only requires 3 group elements.

3. Choose $s_1, \dots, s_{\ell} \xleftarrow{R} \mathbb{Z}_p$ and $\omega_1, \dots, \omega_{\ell} \xleftarrow{R} \mathbb{Z}_p$. For $j = 1$ to ℓ , compute

$$\begin{aligned} C_{j,1} &= g^{s_j} & C_{j,2} &= H_U(M)^{s_j} \\ C_{j,3} &= g_{\text{id}_i[j]}^{s_j/\omega_j} & C_{j,4} &= g_2^{\omega_j} \end{aligned} \quad (6)$$

Then, encrypt $C_{j,3}$ using Kiltz's encryption scheme, by choosing $z_{j,1}, z_{j,2} \xleftarrow{R} \mathbb{Z}_p$ and computing

$$\Psi_j = (\Psi_{j,1}, \Psi_{j,2}, \Psi_{j,3}, \Psi_{j,4}, \Psi_{j,5}) = (X^{z_{j,1}}, Y^{z_{j,2}}, C_{j,3} \cdot g^{z_{j,1}+z_{j,2}}, (g^{\text{VK}} \cdot U)^{z_{j,1}}, (g^{\text{VK}} \cdot V)^{z_{j,2}})$$

The next step will be to prove that the ciphertexts $\{\Psi_j\}_{j=1}^{\ell}$ encrypt $\{C_{j,3}\}_{j=1}^{\ell}$ such that $\{(C_{j,1}, C_{j,2}, C_{j,3}, C_{j,4})\}_{j=1}^{\ell}$ are of the form (6) with $\text{id}_i[j] \in \{0, 1\}$.

4. To generate NIZK proofs for the next statements, generate commitments $\vec{C}_{\theta} = \vec{\varphi}^{\theta} \cdot \vec{f}_1^{r\theta} \cdot \vec{f}_2^{s\theta}$, as well as \vec{C}_{Γ_1} and \vec{C}_{Γ_2} to the variables

$$\theta = 1, \quad \Gamma_1 = g^{\theta}, \quad \Gamma_2 = g_2^{\theta} \quad (7)$$

and a non-interactive proof $\vec{\pi}_{\Gamma}$ for the three equalities (7), which requires 9 group elements (3 for each equation). Then, for each $j \in \{1, \dots, \ell\}$, generate Groth-Sahai commitments \vec{C}_{G_j} , \vec{C}_{Θ_j} , $\vec{C}_{z_{j,1}}$, $\vec{C}_{z_{j,2}}$ to the variables $G_j = g_1^{\text{id}_i[j]} \cdot g_0^{1-\text{id}_i[j]}$, $\Theta_j = g^{s_j/\omega_j}$, $Z_{j,1} = g^{z_{j,1}}$ and $Z_{j,2} = g^{z_{j,2}}$. Then, generate NIZK proofs $\vec{\pi}_j$, $\vec{\pi}_{G_j}$, $\vec{\pi}_{\Theta_j}$, $\{\vec{\pi}_{\Psi_{j,k}}\}_{k=1}^3$ for the relations

$$e(F_j, F_j) = e(g, F_j) \quad (8)$$

$$e(G_j, g) = e(g_1, F_j) \cdot e(g_0, F_j^{-1} \cdot g) \quad (9)$$

$$e(\Theta_j, C_{j,4}) = e(C_{j,1}, g_2) \quad (10)$$

$$e(\Psi_{j,1}, g) = e(X, Z_{j,1}) \quad (11)$$

$$e(\Psi_{j,2}, g) = e(Y, Z_{j,2}) \quad (12)$$

$$e(\Psi_{j,3}, g) = e(\Theta_j, G_j) \cdot e(g, Z_{j,1} \cdot Z_{j,2}) \quad (13)$$

This is done by proving that

$$e(F_j, F_j) = e(g, F_j) \quad (14)$$

$$e(G_j, g) = e(g_1, F_j) \cdot e(g_0, F_j^{-1} \cdot g) \quad (15)$$

$$e(\Theta_j, C_{j,4}) = e(C_{j,1}, \Gamma_2) \quad (16)$$

$$e(\Psi_{j,1}, \Gamma_1) = e(X, Z_{j,1}) \quad (17)$$

$$e(\Psi_{j,2}, \Gamma_1) = e(Y, Z_{j,2}) \quad (18)$$

$$e(\Psi_{j,3}, \Gamma_1) = e(\Theta_j, G_j) \cdot e(\Gamma_1, Z_{j,1} \cdot Z_{j,2}) \quad (19)$$

Note that relation (8) guarantees that each $\text{id}_i[j]$ is indeed a bit. Relations (14) and (19) are quadratic equation and thus require 9 elements each whereas 12 elements suffice for relations (15)-(18). Note that the same variable $\theta \in \mathbb{Z}_p$ can be re-used for each $j \in \{1, \dots, \ell\}$, so that (7) only needs to be proved once.

5. Generate a commitment \vec{C}_{Ω_1} to Ω_1 with a NIWI proof $\vec{\pi}_W \in \mathbb{G}^3$ that variables (Ω_1, H_V) satisfy the verification equation

$$e(g, g)^\omega \cdot e(H_W(M), \Omega_3) = e(\Omega_1, g) \cdot e(H_V, \Omega_2^{-1}) \quad (20)$$

of the two-level Waters signature.

6. Finally, use SK to generate a one-time signature σ_{ots} on the entire set of commitments and NIWI/NIZK proofs in order to achieve anonymity in the CCA2 sense.

The whole signature σ consists of

$$\sigma = \left(\text{VK}, \vec{C}_{H_V}, \vec{C}_\theta, \vec{C}_{\Gamma_1}, \vec{C}_{\Gamma_2}, \vec{\pi}_\Gamma, \vec{\pi}_{H_V}, \vec{\pi}_W, \{\vec{C}_{F_j}, (C_{j,1}, C_{j,2}, C_{j,4}, \Psi_j), \right. \\ \left. \vec{C}_{G_j}, \vec{\pi}_{G_j}, \vec{\pi}_{\Theta_j}, \vec{C}_{\Theta_j}, \vec{C}_{Z_{j,1}}, \vec{C}_{Z_{j,2}}, \vec{\pi}_j, \{\vec{\pi}_{\Psi_{j,k}}\}_{k=1}^3\}_{j=1}^\ell, \vec{C}_{\Omega_1}, \Omega_2, \Omega_3, \sigma_{ots} \right)$$

Verify(gpk, M , σ): parse σ as above. Return 1 if and only if: (i) σ_{ots} is a valid one-time signature on the entire bundle; (ii) $\{\Psi_j\}_{j=1}^\ell$ are all valid ciphertexts for Kiltz's cryptosystem (*i.e.*, by testing if $e(\Psi_{j,4}, X) = e(\Psi_{j,1}, g^{\text{VK}} \cdot U)$ and $e(\Psi_{j,5}, Y) = e(\Psi_{j,2}, g^{\text{VK}} \cdot V)$); (iii) $e(C_{j,1}, H_U(M)) = e(g, C_{j,2})$ for each $j \in \{1, \dots, \ell\}$; (iv) All proofs properly verify.

TrapGen(gpk, $\text{msk}_{\text{ADM}}, M$): given the admitter's key $\text{msk}_{\text{ADM}} = (g_2^{\alpha_0}, g_2^{\alpha_1})$ and a message $M \in \{0, 1\}^\mathcal{L}$, choose $r_0, r_1 \xleftarrow{R} \mathbb{Z}_p$ to compute and return

$$t_M = (t_{0,1}, t_{0,2}, t_{1,1}, t_{1,2}) = (g_2^{\alpha_0} \cdot H_U(M)^{r_0}, g^{r_0}, g_2^{\alpha_1} \cdot H_U(M)^{r_1}, g^{r_1}). \quad (21)$$

Open(gpk, ok, t_M, M, σ): return \perp if σ is not a valid group signature w.r.t. gpk and M . Otherwise, parse t_M as in (21). For $i = 1$ to ℓ , do the following.

1. Decrypt $\Psi_j = (\Psi_{j,1}, \Psi_{j,2}, \Psi_{j,3}, \Psi_{j,4}, \Psi_{j,5})$ using $\text{ok} = (\beta_x, \beta_y, \beta_u, \beta_v)$ to obtain $C_{j,3} \in \mathbb{G}$.
2. Use t_M to determine the bit $\text{id}[i] \in \{0, 1\}$ for which the equalities (6) are satisfied.

Return the identifier $\text{id} = \text{id}[1] \dots \text{id}[\ell] \in \{0, 1\}^\ell$.

Overall, each signature consists of $53\ell + 35$ group elements if the scheme is instantiated with Groth's discrete-logarithm-based one-time signature [24]. For groups of $N \approx 10^6$ members (which can accommodate the population of a city), we can set $\ell = 20$ and obtain signatures of 68 kB at the 128-bit security level assuming that each group element has a 512-bit representation. In comparison, the k -resilient system of Sakai *et al.* [36] already requires signatures of 32 kB for the same security level. While less efficient than the random-oracle-based realization of [34], our scheme is not unrealistically expensive for practical applications.

4.2 Security

The traceability of the scheme relies on the standard CDH assumption whereas the anonymity properties rely on the D3DH and DLIN assumptions. In the proof of anonymity against the admitter, we also need to assume that the one-time signature is strongly unforgeable [5], which is implied by the DLIN assumption in Groth's scheme [24]. Since the CDH assumption is implied by both D3DH and DLIN, we only need two assumptions to prove the following result (as detailed in Appendix C).

Theorem 2. *The scheme provides full traceability as well as full anonymity against the opener and the admitter assuming that: (i) Σ is a strongly unforgeable one-time signature; (ii) The DLIN and D3DH assumption both hold in \mathbb{G} .*

References

1. M. Abe, G. Fuchsbauer, J. Groth, K. Haralambiev, M. Ohkubo. Structure-Preserving Signatures and Commitments to Group Elements. In *Crypto'10*, LNCS 6223, pp. 209–236, 2010.
2. M. Abe, J. Groth, K. Haralambiev, M. Ohkubo. Optimal Structure-Preserving Signatures in Asymmetric Bilinear Groups. In *Crypto'11*, LNCS 6841, pp. 649–666, 2011.
3. M. Abe, K. Haralambiev, M. Ohkubo. Signing on Elements in Bilinear Groups for Modular Protocol Design. Cryptology ePrint Archive: Report 2010/133, 2010.
4. M. Abe, K. Haralambiev, M. Ohkubo. Group to Group Commitments Do Not Shrink. In *Eurocrypt'12*, LNCS 7237, pp. 301–317, 2012.
5. J.-H. An, Y. Dodis, T. Rabin. On the Security of Joint Signature and Encryption. In *Eurocrypt'02*, LNCS 2332, pp. 83–107, 2002.
6. G. Ateniese, J. Camenisch, M. Joye, G. Tsudik. A practical and provably secure coalition-resistant group signature scheme. In *Crypto'00*, LNCS 1880, pp. 255–270, 2000.
7. M. Bellare, D. Micciancio, B. Warinschi. Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In *Eurocrypt'03*, LNCS 2656, pp. 614–629, 2003.
8. M. Bellare, P. Rogaway. Random Oracles are Practical: A Paradigm for Designing Efficient Protocols. In *1st ACM Conference on Computer and Communications Security*, pp. 62–73, ACM Press, 1993.
9. M. Bellare, H. Shi, C. Zhang. Foundations of group signatures: The case of dynamic groups. In *CT-RSA'05*, LNCS 3376, pp. 136–153, 2005.
10. V. Benjumea, S. G. Choi, J. Lopez, M. Yung. Fair traceable multi-group signatures. In *Financial Cryptography 2008*, LNCS 5143, pp. 231–246. Springer, 2008.
11. D. Boneh, X. Boyen. Efficient Selective-ID Secure Identity-Based Encryption Without Random Oracles. In *Eurocrypt'04*, LNCS 3027, pp. 223–238, 2004.
12. D. Boneh, X. Boyen, E.-J. Goh. Hierarchical identity based encryption with constant size ciphertext. In *Eurocrypt'05*, LNCS 3494, pp. 440–456, 2005.
13. D. Boneh, X. Boyen, H. Shacham. Short Group Signatures. In *Crypto'04*, LNCS 3152, pp. 41–55. Springer, 2004.
14. D. Boneh, M. Franklin. Identity-Based Encryption from the Weil Pairing. In *SIAM Journal of Computing* 32(3), pp. 586–615, 2003, earlier version in *Crypto'01*, LNCS 2139, pp. 213–229, 2001.
15. X. Boyen, B. Waters. Compact Group Signatures Without Random Oracles. In *Eurocrypt'06*, LNCS 4004, pp. 427–444, Springer, 2006.
16. X. Boyen, B. Waters. Full-Domain Subgroup Hiding and Constant-Size Group Signatures. In *PKC'07*, LNCS 4450, pp. 1–15, 2007.
17. J. Camenisch, R. Chaabouni, a. shelat. Efficient Protocols for Set Membership and Range Proofs. In *Asiacrypt'08*, LNCS 5350, pp. 234–252, Springer, 2008.
18. J. Camenisch, K. Haralambiev, M. Kohlweiss, J. Lapon, V. Naessens. Structure Preserving CCA Secure Encryption and Applications. In *Asiacrypt'11*, LNCS 7073, pp. 89–106, 2011.
19. R. Canetti, O. Goldreich, and S. Halevi. The random oracle methodology, revisited. *Journal of the ACM*, 51(4):557–594, 2004.
20. J. Cathalo, B. Libert, M. Yung. Group Encryption: Non-Interactive Realization in the Standard Model. In *Asiacrypt'09*, LNCS 5912, pp. 179–196, 2009.
21. D. Chaum, E. van Heyst. Group Signatures. In *Eurocrypt'91*, LNCS 547, pp. 257–265, Springer, 1991.
22. C. Delerablée, D. Pointcheval. Dynamic Fully Anonymous Short Group Signatures. In *Vietcrypt'06*, LNCS 4341, pp. 193–210, Springer, 2006.
23. J. Groth, R. Ostrovsky, and A. Sahai. Perfect non-interactive zero knowledge for NP. In *Eurocrypt'06*, LNCS 4004, pp. 339–358. Springer, 2006.
24. J. Groth. Simulation-Sound NIZK Proofs for a Practical Language and Constant Size Group Signatures. In *Asiacrypt'06*, LNCS 4284, pp. 444–459, Springer, 2006.
25. J. Groth. Fully anonymous group signatures without random oracles. In *Asiacrypt'07*, LNCS 4833, pp. 164–180. Springer, 2007.
26. J. Groth, A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *Eurocrypt'08*, LNCS 4965, pp. 415–432, 2008.
27. J. Håstad, R. Impagliazzo, L. Levin, M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, vol. 28(4), pp. 1364–1396, 1999.
28. S.-H. Heng, K. Kurosawa. k-Resilient Identity-Based Encryption in the Standard Model. In *CT-RSA'04*, LNCS 2964, pp. 67–80, 2004.
29. A. Kiayias, Y. Tsiounis, and M. Yung. Traceable signatures. In *Eurocrypt'04*, LNCS 3027, pp. 571–589, 2004.
30. A. Kiayias, M. Yung. Secure scalable group signature with dynamic joins and separable authorities. *International Journal of Security and Networks (IJSN)* Vol. 1, No. 1/2, pp. 24–45, 2006.
31. E. Kiltz, A. Mityagin, S. Panjwani, B. Raghavan. Append-Only Signatures. In *ICALP'05*, LNCS 3580, pp. 434–445, 2005.
32. E. Kiltz. Chosen-ciphertext security from tag-based encryption. In *TCC'06*, LNCS 3876, pp. 581–600, 2006.

33. L. Nguyen, R. Safavi-Naini. Efficient and Provably Secure Trapdoor-Free Group Signature Schemes from Bilinear Pairings. In *Asiacrypt'04*, LNCS 3329, pp. 372–386. Springer-Verlag, 2004.
34. K. Ohara, Y. Sakai, K. Emura, G. Hanaoka. A Group Signature Scheme with Unbounded Message-Dependent Opening. In *AsiaCCS 2013*, ACM Press, 2013.
35. A. Sahai, B. Waters. Fuzzy Identity-Based Encryption. In *Eurocrypt'05*, LNCS 3494, pp. 457–473. Springer, 2005.
36. Y. Sakai, K. Emura, G. Hanaoka, Y. Kawai, T. Matsuda, and K. Omote. Group Signatures with Message-Dependent Opening. In *5th International Conference on Pairing-Based Cryptography (Pairing 2012)*, LNCS series, Springer, 2013.
37. A. Shamir. Identity-Based Cryptosystems and Signature Schemes. In *Crypto'84*, LNCS 196, pp. 47–53, 1984.
38. B. Waters. Efficient identity-based encryption without random oracles. In *Eurocrypt'05*, LNCS 3494, pp. 114–127. Springer, 2005.

A Definitions for Identity-Based Encryption

Definition 6 ([14]). *An IBE scheme is a tuple (**Setup**, **Keygen**, **Encrypt**, **Decrypt**) of efficient algorithms such that:*

- **Setup** takes as input a security parameter $\lambda \in \mathbb{N}$ and outputs a master public key mpk and a matching master secret key msk .
- **Keygen** takes as input an identity ID and a master secret key msk . It outputs a private key d_{ID} for the identity ID .
- **Encrypt** takes as input the master public key mpk , an identity ID and a message m and outputs a ciphertext C .
- **Decrypt** takes as input the master public key mpk , a decryption key d_{ID} and a ciphertext C and outputs a message M .

*Correctness requires that, for any $\lambda \in \mathbb{N}$, any outputs (mpk, msk) of **Setup**(λ), any plaintext M and any ID , if $d_{\text{ID}} \leftarrow \text{Keygen}(\text{msk}, \text{ID})$, we have $\text{Decrypt}(\text{mpk}, d_{\text{ID}}, \text{Encrypt}(\text{mpk}, \text{ID}, M)) = M$.*

The standard security notion captures the semantic security of messages encrypted under some identity, even when the adversary has corrupted polynomially-many other identities.

Definition 7. [14] *An IBE system is semantically secure (or IND-ID-CPA secure) if no PPT adversary \mathcal{A} has non-negligible advantage in this game:*

1. *The challenger generates a master key pair $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(\lambda)$ and gives mpk to \mathcal{A} .*
2. *\mathcal{A} issues a number of key extraction queries for identities ID of its choice. The challenger responds with $d_{\text{ID}} \leftarrow \text{Keygen}(\text{msk}, \text{ID})$.*
3. *When the adversary \mathcal{A} decides that phase 2 is over, it chooses distinct equal-length messages M_0, M_1 and an identity ID^* that has never been queried to the key extraction oracle at step 2. The challenger flips a coin $d \stackrel{R}{\leftarrow} \{0, 1\}$ and returns a challenge ciphertext $C^* = \text{Encrypt}(\text{mpk}, \text{ID}, M_d^*)$.*
4. *\mathcal{A} issues new queries but cannot ask for the private key of ID^* .*
5. *\mathcal{A} finally outputs a bit $d' \in \{0, 1\}$ and wins if $d' = d$. \mathcal{A} 's advantage is defined as the distance $\text{Adv}^{\text{ind-id-cpa}}(\mathcal{A}) := |\Pr[d' = d] - 1/2|$.*

In k -resilient IBE schemes [28], the adversary is restricted to make private key extraction queries on at most k distinct identities. In this paper, we consider the standard definition where the maximal number of private key queries is not fixed in advance.

B Proof of Lemma 1

Proof. Let us assume that there exists $i \in \{1, \dots, \ell\}$ for which a PPT adversary \mathcal{A} can tell Game i apart from Game $i-1$. We show how to build an algorithm \mathcal{B} that takes in an instance (g, g^a, g^b, g^c, T) of the D3DH problem and uses its interaction with \mathcal{A} to decide if $T = g^{abc}$ or $T \in_R \mathbb{G}$.

To this end, algorithm \mathcal{B} prepares the master public key mpk by randomly choosing $\gamma_0, \gamma_1 \xleftarrow{R} \mathbb{Z}_p$ and setting $g_0 = (g^a)^{\gamma_0}$, $g_1 = (g^a)^{\gamma_1}$ as well as $g_2 = g^b$. Note that this implicitly defines $\alpha_0 = a \cdot \gamma_0$ and $\alpha_1 = a \cdot \gamma_1$. Next, \mathcal{B} chooses random values $\nu \xleftarrow{R} \{0, \dots, L\}$, $\rho_0, \rho_1, \dots, \rho_L \xleftarrow{R} \{0, \dots, \zeta - 1\}$ and $\delta_0, \delta_1, \dots, \delta_L \xleftarrow{R} \mathbb{Z}_p$, with $\zeta = 2q$ and where q is the maximal number of private key queries throughout the game. These are used to define

$$\begin{aligned} u_0 &= g^{\delta_0} \cdot (g^b)^{\nu \cdot \zeta - \rho_0} \\ u_i &= g^{\delta_i} \cdot (g^b)^{-\rho_i}, \end{aligned} \quad i \in \{1, \dots, L\}, \quad (22)$$

so that any L -bit identity $\text{ID} = \text{ID}[1] \dots \text{ID}[L] \in \{0, 1\}^L$ has a hash value $H_{\mathbb{G}}(\text{ID}) = u_0 \cdot \prod_{i=1}^L u_i^{\text{ID}[i]}$ that can be written $H_{\mathbb{G}}(\text{ID}) = g^{J_2(\text{ID})} \cdot (g^b)^{J_1(\text{ID})}$ if we define the functions

$$J_1(\text{ID}) = \nu \cdot \zeta - \rho_0 - \sum_{i=1}^L \rho_i \cdot \text{ID}[i], \quad J_2(\text{ID}) = \delta_0 - \sum_{i=1}^L \delta_i \cdot \text{ID}[i].$$

The generation of mpk is completed by having \mathcal{B} choose $Z_1, \dots, Z_\ell \xleftarrow{R} \mathbb{G}$ at random.

Whenever \mathcal{A} queries an identity ID for private key extraction, \mathcal{B} uses the same strategy as in the security proofs of [11, 38]. Namely, it first evaluates the function $J_1(\text{ID})$. If $J_1(\text{ID}) = 0$, it aborts and outputs a random bit. Otherwise, it chooses $r_0, r_1 \xleftarrow{R} \mathbb{Z}_p$ and computes

$$\begin{aligned} &(d_{0,1}, d_{0,2}, d_{1,1}, d_{1,2}) \\ &= \left(H_{\mathbb{G}}(\text{ID})^{r_0} \cdot (g^a)^{-\gamma_0 \cdot J_2(\text{ID})}, g^{r_0} \cdot (g^a)^{-\gamma_0 / J_1(\text{ID})}, H_{\mathbb{G}}(\text{ID})^{r_1} \cdot (g^a)^{-\gamma_1 \cdot J_2(\text{ID})}, g^{r_1} \cdot (g^a)^{-\gamma_1 / J_1(\text{ID})} \right) \end{aligned}$$

which can be written $(g_2^{\gamma_0 \cdot a} \cdot H_{\mathbb{G}}(\text{ID})^{\tilde{r}_0}, g^{\tilde{r}_0}, g_2^{\gamma_1 \cdot a} \cdot H_{\mathbb{G}}(\text{ID})^{\tilde{r}_1}, g^{\tilde{r}_1})$ if we define $\tilde{r}_0 = r_0 - \gamma_0 \cdot a / J_1(\text{ID})$ and $\tilde{r}_1 = r_1 - \gamma_1 \cdot a / J_1(\text{ID})$. The 4-uple $d_{\text{ID}} = (d_{0,1}, d_{0,2}, d_{1,1}, d_{1,2})$ thus forms a valid private key and is returned to \mathcal{A} .

When \mathcal{A} decides to enter the challenge phase, it chooses messages $M_0, M_1 \in \mathbb{G}$ and a target identity ID^* . At this point, \mathcal{B} aborts and outputs a random bit in the event that $J_1(\text{ID}^*) \neq 0$. Otherwise (*i.e.*, if $J_1(\text{ID}^*) = 0$), \mathcal{B} chooses a bit $\beta \xleftarrow{R} \{0, 1\}$ as well as a random ℓ -bit string $K \xleftarrow{R} \{0, 1\}^\ell$ and generates the challenge ciphertext as follows.

- For each $j \in \{1, \dots, i-1\}$, \mathcal{B} sets $(C_{j,1}, C_{j,2}, C_{j,3}, C_{j,4}) = (g^{s_j}, H_{\mathbb{G}}(\text{ID})^{s_j}, \tilde{C}_{j,3}, \tilde{C}_{j,4})$ for randomly chosen $s_j, \omega_j \xleftarrow{R} \mathbb{Z}_p$ and $\tilde{C}_{j,3}, \tilde{C}_{j,4} \xleftarrow{R} \mathbb{G}$.
- For each $j \in \{i+1, \dots, \ell\}$, \mathcal{B} sets $(C_{j,1}, C_{j,2}, C_{j,3}, C_{j,4}) = (g^{s_j}, H_{\mathbb{G}}(\text{ID})^{s_j}, g_{K[j]}^{s_j / \omega_j}, g_2^{\omega_j})$ for randomly chosen $s_j, \omega_j \xleftarrow{R} \mathbb{Z}_p$.
- For $j = i$, \mathcal{B} picks $\omega_i \xleftarrow{R} \mathbb{Z}_p$ and sets $(C_{i,1}, C_{i,2}, C_{i,3}, C_{i,4}) = (g^c, (g^c)^{J_2(\text{ID}^*)}, T^{\gamma_{K[i]} / \omega_i}, g^{\omega_i})$.

Finally, \mathcal{B} computes $C_0 = M_\beta \cdot \prod_{j=1}^\ell Z_j^{K[j]}$ and provides the adversary with the challenge ciphertext $C = (C_0, \{(C_{j,1}, C_{j,2}, C_{j,3}, C_{j,4})\}_{j=1}^\ell)$.

We remark that, if $T = g^{abc}$, the challenge ciphertext C is distributed as in Game $i-1$ as $(C_{i,1}, C_{i,2}, C_{i,3}, C_{i,4})$ can be written

$$(C_{i,1}, C_{i,2}, C_{i,3}, C_{i,4}) = \left(g^c, H_{\mathbb{G}}(\text{ID}^*)^c, g^{ac \cdot \gamma_{K[i]} / \tilde{\omega}_i}, (g^b)^{\tilde{\omega}_i} \right) = \left(g^c, H_{\mathbb{G}}(\text{ID}^*)^c, g_{K[i]}^{c / \tilde{\omega}_i}, g_2^{\tilde{\omega}_i} \right).$$

where $\tilde{\omega}_i = \omega_i / b$. In contrast, if $T \in_R \mathbb{G}$, then the pair $(C_{i,3}, C_{i,4})$ is uniformly distributed in \mathbb{G}^2 , which means that $(C_{i,1}, C_{i,2}, C_{i,3}, C_{i,4})$ has the same distribution as in Game i .

At this stage, the adversary's probability may be correlated with the probability that the simulator \mathcal{B} has to abort (*i.e.*, because \mathcal{A} queries the private key of an identity ID for which $J_1(\text{ID}) = 0$

or because $J_1(\text{ID}^*) \neq 0$ in the challenge phase). As in [38], one way to address this problem is to introduce an artificial abort step in order to guarantee that \mathcal{B} always aborts with the maximal probability, no matter which particular set of queries is made by \mathcal{A} .

Namely, with $\zeta = 2q$, the same analysis as [38] shows that \mathcal{B} 's probability not to abort for any set of queries is at least $\eta = 1/(4(L+1)q)$.

At the end of the game, \mathcal{B} considers the sequence of identities $(\text{ID}_1, \dots, \text{ID}_q, \text{ID}^*)$ chosen by \mathcal{A} during the game and estimates the probability that this choice causes the simulation to abort. This process does not require to run \mathcal{A} again but rather involves repeatedly sampling vectors $(\rho_0, \rho_1, \dots, \rho_L) \stackrel{R}{\leftarrow} \mathbb{Z}_\zeta^{L+1}$ and evaluate $J_1(\text{ID}_1), \dots, J_1(\text{ID}_q)$ and $J_1(\text{ID}^*)$ accordingly. When the estimated probability η' is obtained after $O(\varepsilon^{-2} \ln(\varepsilon^{-1}) \eta^{-1} \ln(\eta^{-1}))$ samples, if $\eta' > \eta$, \mathcal{B} artificially aborts and outputs a random bit with probability $1 - \eta/\eta'$. With probability η/η' , it continues.

After the artificial abort step, if the simulator \mathcal{B} did not naturally or artificially abort, it outputs 1 if \mathcal{A} successfully guesses $\beta' = \beta$ and 0 otherwise. We now argue that \mathcal{B} has non-negligible advantage as a D3DH distinguisher if \mathcal{A} can distinguish Game i from Game $i - 1$. Indeed, depending on the distribution of T , \mathcal{B} is playing either Game $i - 1$ or Game i with \mathcal{A} . Using the same analysis as in [38], we find that, if the difference $|\text{Adv}_{i-1} - \text{Adv}_i|$ between \mathcal{A} 's advantage functions in Game $i - 1$ and Game i is ε , then \mathcal{B} can break the D3DH assumption with probability $\varepsilon/(16(L+1)q)$. \square

C Security Proofs for the GS-MDO Scheme

C.1 Anonymity

Theorem 3. *The scheme provides anonymity against the opener if the DLIN and D3DH assumptions both hold in \mathbb{G} .*

Proof. The proof uses a sequence of games where the first game is the real game and the last one leaves no advantage to the adversary. In Game i , W_i denotes the event that the adversary \mathcal{A} wins.

Game 0: This game is the actual attack game of Definition 5. Namely, the challenger performs the setup of the system and generates the group public key \mathbf{gpk} , the vector \mathbf{gsk} of members' private keys as well as the opening key ok and the admitter's private key msk_{ADM} . The adversary \mathcal{A} is given \mathbf{gpk} , \mathbf{gsk} and ok and it is also allowed to invoke the oracle $\mathcal{O}_{\text{msk}_{\text{ADM}}}(\cdot)$ an arbitrary but polynomial number of times. In the challenge phase, \mathcal{A} chooses two indices $j_0, j_1 \in \{0, \dots, N-1\}$ and a message M^* that has never been submitted to $\mathcal{O}_{\text{msk}_{\text{ADM}}}(\cdot)$. The challenger replies by returning a challenge signature σ^* on the message M^* using the private key $\mathbf{gsk}[j_b]$, where $b \stackrel{R}{\leftarrow} \{0, 1\}$. In the second phase, \mathcal{A} is allowed further queries to $\mathcal{O}_{\text{msk}_{\text{ADM}}}(\cdot)$ (with the restriction of not querying M^*) and eventually outputs a bit $b' \in \{0, 1\}$. We denote by W_0 the event that $b' = b$.

Game 1: In this game, the perfectly sound Groth-Sahai CRS $\mathbf{f} = (\vec{f}_1, \vec{f}_2, \vec{f}_3)$ is replaced by a CRS for the NIWI setting, where $\vec{f}_3 = \vec{f}_1^{\xi_2} \cdot \vec{f}_2^{\xi_2} \cdot (1, 1, g)^{-1}$. Under the DLIN assumption, this change should not be noticeable to \mathcal{A} and we can write $|\Pr[W_1] - \Pr[W_0]| \leq 2 \cdot \text{Adv}^{\text{DLIN}}(\mathcal{B})$.

Game 2: We modify the generation of the challenge signature σ^* . In the latter, we replace the actual NIZK proofs for equations (10)-(13) by simulated proofs which are produced using the simulation trapdoor (ξ_1, ξ_2) of the CRS $\mathbf{f} = (\vec{f}_1, \vec{f}_2, \vec{f}_3)$. To this end, \mathcal{B} uses the variable assignment $\Gamma_1 = \Gamma_2 = 1_{\mathbb{G}}$ and

$$\Theta_j = Z_{j,1} = Z_{j,2} = 1_{\mathbb{G}} \quad j \in \{1, \dots, \ell\},$$

which amounts to set $\theta = 0$ and immediately allows generating proofs for (16)-(19). In order to generate a simulated proof $\vec{\pi}_\Gamma$ for the relations (7), \mathcal{B} uses the trapdoor (ξ_1, ξ_2) to open to 1 a commitment \vec{C}_θ which is initially generated as a commitment to 0. We remark that, in Game 2, the proofs for relations (16)-(19) are computed without using the IBE encryption exponents $\{(s_j, \omega_j)\}_{j=1}^\ell$. Moreover, since simulated NIZK proofs are perfectly indistinguishable from actual proof on a perfectly hiding CRS, \mathcal{A} 's view remains unchanged and we thus have $\Pr[W_2] = \Pr[W_1]$.

We now consider a subsequence of games where the IBE ciphertexts contained in the challenge signature σ^* have a hybrid distribution. We define Game 3.0 to be identical to Game 2 for simplicity.

Game 3.i ($0 \leq i \leq \ell$): This game is a hybrid game where the challenger \mathcal{B} generates the IBE ciphertext components $\{(C_{j,1}, C_{j,2}, C_{j,3}, C_{j,4})\}$ as follows at step 3 of the signature generation algorithm.

- If $j \leq i$, \mathcal{B} chooses $s_j \xleftarrow{R} \mathbb{Z}_p$ and computes $(C_{j,1}, C_{j,2}) = (g^{s_j}, H_{\mathbb{G}}(M^*)^{s_j})$ normally. However, it chooses $(C_{j,3}, C_{j,4}) \xleftarrow{R} \mathbb{G}^2$ at random.
- If $j > i$, \mathcal{B} sets $(C_{j,1}, C_{j,2}, C_{j,3}, C_{j,4}) = (g^{s_j}, H_{\mathbb{G}}(M^*)^{s_j}, g_{\text{id}_{j_b}^{s_j/\omega_j}}, g_2^{\omega_j})$ for random $s_j, \omega_j \xleftarrow{R} \mathbb{Z}_p$, where $\text{id}_{j_b} \in \{0, 1\}^\ell$ is the identity of one of the two candidates $\text{id}_{j_0}, \text{id}_{j_1}$ chosen by \mathcal{A} .

Using exactly the same strategy as in the proof of Lemma 1, we can easily prove that, for each $i \in \{0, \dots, \ell - 1\}$, Game 3.i is computationally indistinguishable from Game 3.(i+1) if the D3DH assumption holds in \mathbb{G} . We have $|\Pr[W_{3,i}] - \Pr[W_{3,(i+1)}]| \leq 16 \cdot q \cdot (n + 1) \cdot \mathbf{Adv}^{\text{D3DH}}(\mathcal{B})$. The proof is essentially identical to the proof of Lemma 1 and omitted.

In Game 3. ℓ , we observe that, in the challenge signature σ^* , all IBE ciphertexts have been tampered with and $\{(C_{j,3}, C_{j,4})\}_{j=1}^\ell$ are just a set of random group elements. This implies that σ^* is devoid of any information about the signer's identity id_{j_b} . Indeed, the bits of id_{j_b} are only used to compute H_V , F_j and G_j . However, these group elements are hidden within perfectly hiding Groth-Sahai commitments and the proof of relations (5), (8), (9) and (20) are perfectly NIWI proofs. It follows that $\mathbf{gsk}[j_0]$ and $\mathbf{gsk}[j_1]$ are equally likely candidates as witnesses for the generation of these proofs. We thus conclude that $\Pr[W_{3,\ell}] = 1/2$. \square

Theorem 4. *The scheme provides anonymity against the admitter assuming that the DLIN assumption holds in \mathbb{G} and that Σ is a strongly unforgeable one-time signature.*

Proof. We proceed again with a sequence of games that begins with the actual attack game and ends with a game where even an unbounded adversary has no advantage. In Game i , we denote by W_i the event that the adversary \mathcal{A} wins.

Game 0: This game is the real attack game described in Definition 4. In this game, the challenger \mathcal{B} generates the group public key \mathbf{gpk} , the vector \mathbf{gsk} of private keys as well as private keys \mathbf{ok} and $\mathbf{msk}_{\text{ADM}}$ for the opener and the admitter, respectively. The adversary \mathcal{A} receives \mathbf{gpk} , \mathbf{gsk} and $\mathbf{msk}_{\text{ADM}}$ at the very beginning of the game. It is further granted access to a signature opening oracle $\mathcal{O}_{\mathbf{ok}}(\cdot)$ which it can invoke on polynomially many occasions. In the challenge phase, \mathcal{A} chooses two indices $j_0, j_1 \in \{0, \dots, N - 1\}$ and a message M^* . At this point, \mathcal{B} generates a challenge signature σ^* on M^* using $\mathbf{gsk}[j_b]$, where $b \xleftarrow{R} \{0, 1\}$ is chosen at random. In the second phase, \mathcal{A} can make additional queries to $\mathcal{O}_{\mathbf{ok}}(\cdot)$ (with the restriction of not asking for the opening of σ^*) and eventually outputs a bit $b' \in \{0, 1\}$. We denote by W_0 the event that $b' = b$.

Game 1: This game is like Game 0 with the difference that \mathcal{B} aborts if \mathcal{A} queries the opening oracle $\mathcal{O}_{\mathbf{ok}}(\cdot)$ on input of a valid signature for which the built-in one-time verification key \mathbf{VK} coincides with the one-time verification key \mathbf{VK}^* used in the challenge phase (we can assume w.l.o.g. that the latter is chosen at the outset of the game). It is easy to see that, if the latter event occurs, \mathcal{A} can be turned into a forger that breaks the strong unforgeability of the one-time signature Σ . Indeed, even if this event occurs before the challenge phase (which is unlikely since \mathbf{VK}^* is independent of \mathcal{A} 's view until that point), it implies that \mathcal{A} is able to forge a signature without having seen a single one-time signature. We can thus write $|\Pr[W_1] - \Pr[W_0]| \leq \mathbf{Adv}^{\text{ots-suf}}(\mathcal{B})$.

Game 2: In this game, the perfectly sound Groth-Sahai CRS $\mathbf{f} = (\vec{f}_1, \vec{f}_2, \vec{f}_3)$ is traded for a perfectly WI CRS, where $\vec{f}_3 = \vec{f}_1^{\kappa_2} \cdot \vec{f}_2^{\kappa_2} \cdot (1, 1, g)^{-1}$. Under the DLIN assumption, this change should not noticeably affect \mathcal{A} 's behavior and we can write $|\Pr[W_2] - \Pr[W_1]| \leq 2 \cdot \mathbf{Adv}^{\text{DLIN}}(\mathcal{B})$.

Game 3: We now change the generation of the challenge signature σ^* . Namely, we replace its underlying NIZK proofs for equations (10)-(13) by simulated proofs which are produced using the trapdoor (ξ_1, ξ_2) as in the proof of Theorem 3. To this end, \mathcal{B} defines $\Gamma_1 = \Gamma_2 = 1_{\mathbb{G}}$ and

$$\Theta_j = Z_{j,1} = Z_{j,2} = 1_{\mathbb{G}} \quad j \in \{1, \dots, \ell\},$$

which allows generating proofs for the relations (16)-(19) since they form witnesses satisfying these equalities. When it comes to simulating a fake proof $\vec{\pi}_r$ for the relations (7), \mathcal{B} uses the trapdoor (ξ_1, ξ_2) which allows trapdoor opening to 1 a commitment \vec{C}_θ to 0 (recall that \vec{C}_θ is a perfectly hiding commitment from Game 2 onwards). In this game, the proofs for (17)-(19) are computed without using the encryption exponents $\{(z_{j,1}, z_{j,2})\}_{j=1}^\ell$ at step 2 of the signing algorithm. Also, since simulated NIZK proofs have the same distribution as genuine proofs on a perfectly hiding CRS \mathbf{f} , \mathcal{A} 's has not changed since Game 2 and we thus have $\Pr[W_3] = \Pr[W_2]$.

We now consider a subsequence of games where, in the challenge signature σ^* , the ciphertexts of Kiltz's tag-based encryption scheme are gradually replaced by encryptions of random group elements. For convenience, we define Game 4.0 to be identical to Game 3.

Game 4.i ($0 \leq i \leq \ell$): This game is a hybrid game where the challenger \mathcal{B} generates the ciphertexts $\{\Psi_j = (\Psi_{j,1}, \Psi_{j,2}, \Psi_{j,3}, \Psi_{j,4}, \Psi_{j,5})\}_{j=1}^\ell$ as follows at step 3 of the signing algorithm.

- If $j \leq i$, \mathcal{B} computes Ψ_j as an encryption of a random group element.
- If $j > i$, \mathcal{B} follows the usual signing algorithm and computes Ψ_j as a tag-based encryption of $C_{j,3}$ (i.e., the third component of the IBE ciphertext) with respect to the tag VK^* .

Since we do not use the encryption exponents $\{(z_{j,1}, z_{j,2})\}_{j=1}^\ell$ as witnesses to generate NIZK proofs and we ruled out the case of \mathcal{A} recycling VK^* in a query to $\mathcal{O}_{\text{ok}}(\cdot)$, we can meaningfully appeal to the selective-tag security of Kiltz's tag-based cryptosystem when it comes to arguing that Game 4.i is computationally indistinguishable from Game 4.(i+1). By repeating the proof of [32] (we omit the details here), it is straightforward to show that $|\Pr[W_{4.i}] - \Pr[W_{4.(i+1)}]| \leq \text{Adv}^{\text{DLIN}}(\mathcal{B})$.

In Game 4. ℓ , we observe that the challenge signature σ^* carries no information about the signer's identity id_{j_b} . Indeed, all tag-based encryption ciphertexts $\{\Psi_j\}_{j=1}^\ell$ have been replaced by encryptions of random group elements. This implies that the bits of id_{j_b} are only involved in the computation H_V , F_j and G_j as well as in proofs for the relations (5), (8), (9) and (20). However, these are perfectly NIWI proofs and only perfectly hiding commitments to H_V and $\{(F_j, G_j)\}_{j=1}^\ell$ are given out. We thus find that $\Pr[W_{4.\ell}] = 1/2$. \square

C.2 Traceability

Theorem 5. *The scheme provides full traceability under the CDH assumption in \mathbb{G} .*

Proof. We prove that, if an adversary \mathcal{A} can defeat the traceability property of the scheme with non-negligible advantage, we can build a chosen-message adversary \mathcal{B} against the Waters signature, which contradicts the CDH assumption (as proved in [38]).

Our forger \mathcal{B} inputs a public key $PK_W = (g, g^a, g^b, \{w_i\}_{i=0}^L)$ for the Waters signature scheme, for which the underlying secret key is $SK_W = g^{ab}$, and interacts with a chosen-message challenger.

Setup: To prepare the group public key \mathbf{gpk} , our forger \mathcal{B} sets $e(g, g)^\omega = e(g^a, g^b)$ (which implicitly defines $\omega = ab$) at step 4 of the group key generation algorithm. It also picks a random ℓ -bit string $\text{id}^\dagger = \text{id}^\dagger[1] \dots \text{id}^\dagger[\ell] \stackrel{R}{\leftarrow} \{0, 1\}^\ell$, hoping that \mathcal{A} 's fake group signature will open to id^\dagger . It completes the generation of the pk_W part of \mathbf{gpk} by choosing $\delta_0, \dots, \delta_L \stackrel{R}{\leftarrow} \mathbb{Z}_p$ and $\rho_0, \dots, \rho_L \stackrel{R}{\leftarrow} \mathbb{Z}_p$ subject to the constraint $\rho_0 + \sum_{j=1}^\ell \rho_j \cdot \text{id}^\dagger[j] = 0 \pmod p$ in order to define $v_j = g^{\delta_j} \cdot (g^b)^{\rho_j}$ for $j = 1$ to ℓ . By doing so, \mathcal{B} is able to generate valid private keys $\mathbf{gsk}[\text{id}]$ for all users id such that $\text{id} \neq \text{id}^\dagger$ (here, user

indexes are identified with their binary representation, which are their ℓ -bit identifier) as follows: a private key for $\text{id} = \text{id}[1] \dots \text{id}[\ell] \in \{0, 1\}^\ell$ is obtained by choosing $r \xleftarrow{R} \mathbb{Z}_p$ and computing

$$\text{gsk}[\text{id}] = (H_V(\text{id})^r \cdot (g^a)^{-J_2(\text{id})/J_1(\text{id})}, g^r \cdot (g^a)^{-1/J_1(\text{id})})$$

where $J_1(\text{id}) = \rho_0 + \sum_{j=1}^{\ell} \rho_j \cdot \text{id}[j]$, $J_2(\text{id}) = \delta_0 + \sum_{j=1}^{\ell} \delta_j \cdot \text{id}[j]$. In the unlikely event that $J_1(\text{id}) = 0$ for some $\text{id} \neq \text{id}^\dagger$, \mathcal{B} aborts. Note that $\text{gsk}[\text{id}]$ is a valid private key for id as it can be written $\text{gsk}[\text{id}] = (g^{ab} \cdot H_V(\text{id})^{\tilde{r}}, g^{\tilde{r}})$ if we define $\tilde{r} = r - a/J_1(\text{id})$. On the other hand, \mathcal{B} lacks a private key $\text{gsk}[\text{id}^\dagger]$ for the expected target member id^\dagger . Still, with probability $1/N$, \mathcal{A} will create a forgery that opens to id^\dagger , in which case it will refrain from requesting $\text{gsk}[\text{id}^\dagger]$.

Finally, \mathcal{B} generates a perfectly sound Groth-Sahai CRS $\mathbf{f} = (\vec{f}_1, \vec{f}_2, \vec{f}_3)$ for which it retains the extraction trapdoor $(\beta_1, \beta_2) = (\log_g(f_1), \log_g(f_2))$.

Queries: If the adversary chooses to corrupt the private key $\text{gsk}[\text{id}^\dagger]$, \mathcal{B} halts and declares failure since its choice of the expected target user id^\dagger was necessarily incorrect. For any other private key query, \mathcal{B} has the private key $\text{gsk}[\text{id}]$ at disposal and can consistently answer the query.

For all signing queries (id, M) involving users $\text{id} \neq \text{id}^\dagger$, \mathcal{B} can generate a valid signature by faithfully running the signing algorithm since it knows $\text{gsk}[\text{id}]$. We thus focus on signing queries involving id^\dagger . For each of these, our forger \mathcal{B} invokes its own chosen-message challenge which replies by returning a Waters signature $(\Omega_1, \Omega_3) = (g^{ab} \cdot (w_0 \cdot \prod_{j=1}^L w_j^{M[j]})^s, g^s)$ on the message $M = M[1] \dots M[L]$. From (Ω_1, Ω_3) , \mathcal{B} can form a valid 2-level hierarchical signature $(\tilde{\Omega}_1, \tilde{\Omega}_2, \tilde{\Omega}_3)$ by choosing $r \xleftarrow{R} \mathbb{Z}_p$ and setting $(\tilde{\Omega}_1, \tilde{\Omega}_2, \tilde{\Omega}_3) = (\Omega_1 \cdot g^{r \cdot (\delta_0 + \sum_{j=1}^{\ell} \delta_j \cdot \text{id}^\dagger[j])}, g^r, \Omega_3)$. Having constructed $(\tilde{\Omega}_1, \tilde{\Omega}_2, \tilde{\Omega}_3)$, \mathcal{B} can faithfully run steps 2-6 of the real signing algorithm and obtain a well-formed group signature.

Output: When \mathcal{A} halts, it outputs a message M^* and a valid group signature σ^* which presumably opens to an uncorrupted user that did not sign the message M^* . At this point, \mathcal{B} uses the extraction trapdoor (β_1, β_2) to open all the commitments contained in σ^* . Having extracted $\{F_j\}_{j=1}^{\ell}$ from their respective commitments, it can decode the identity id^* of the involved group member and aborts in the event that $\text{id}^* \neq \text{id}^\dagger$. With probability $1/N = 1/2^\ell$, we have $\text{id}^* = \text{id}^\dagger$ (which also implies that \mathcal{B} did not have to abort during the querying phase) and \mathcal{B} can forge a Waters signature as follows. Having extracted a valid hierarchical signature $(\Omega_1^*, \Omega_2^*, \Omega_3^*)$ from σ^* , \mathcal{B} obtains a Waters signature by computing $(\Omega_1', \Omega_3') = (\Omega_1^*/\Omega_2^{*\delta_0 + \sum_{j=1}^{\ell} \delta_j \cdot \text{id}^*[j]}, \Omega_3^*)$. The latter is easily seen to be a valid forgery since, as long as $\text{id}^* = \text{id}^\dagger$, \mathcal{B} never had to ask for a Waters signature on M^* to its own challenger.

Since \mathcal{B} 's probability to correctly predict id^* is at least $1/N$, if \mathcal{A} has advantage ε , the result of [38] implies that \mathcal{B} can be used to break the CDH assumption with probability $\varepsilon/(4 \cdot N \cdot q \cdot (L+1))$. \square