

Non-interactive Anonymous Tokens with Private Metadata Bit

Foteini Baldimtsi^{*†}, Lucjan Hanzlik[†], Quan Nguyen^{*} and Aayush Yadav^{*}

^{*}George Mason University

Fairfax (USA)

Email: {foteini, qnguye31, ayadav5}@gmu.edu

[†]CISPA Helmholtz Center for Information Security

Saarbrücken (Germany)

Email: hanzlik@cispa.de

[‡]Mysten Labs

Palo Alto (USA)

Abstract—Anonymous tokens with private metadata bit (ATPM) have received increased interest as a method for anonymous user authentication while also allowing the issuer to embed trust signals inside the token that are only readable by the authority who holds the secret key. A drawback of all existing ATPM constructions is that they require interaction between the client and the issuer during the issuance process. In this work, we build the first *non-interactive* anonymous tokens (NIAT) with private metadata bit, inspired by the recent work of Hanzlik (EUROCRYPT ’23) on non-interactive blind signatures. We discuss how the property of non-interactivity during issuance allows for more efficient protocols that avoid the need for online signing. We construct an efficient NIAT scheme based on Structure-preserving Signatures on Equivalence Classes (SPS-EQ) and experimentally evaluate its performance. We also present an extension to our NIAT construction that allows the identification of clients who attempt to double-spend a token (i.e., present the same token twice) and argue that non-interactive schemes are uniquely positioned to offer this essential feature.

1. Introduction

Anonymous tokens are a powerful primitive that allow users to access services or resources securely while maintaining their privacy. Typically, an anonymous token system involves three types of parties: a *client*, an *issuer*, and a *redeemtee* (or verifier). The client engages in an issuance protocol with the issuer who first verifies the trustworthiness of the client and then issues them a token. Later, the client can present the token to a redeemtee, who verifies its authenticity to grant the client access to a service or resource, and additionally checks that the token was not redeemed before (double-spending detection). The basic properties satisfied by an anonymous token system include *unforgeability*, which ensures that a client cannot issue tokens on its own, and *anonymity* which prevents issuers and redeemtees to link issued tokens to any token later redeemed.

Anonymous tokens have numerous applications, including controlled access to content delivery networks (CDNs) without CAPTCHA solving [1], private web-browsing [2], private contact tracing [3], fraud detection [3] and private click measurement [4] just to name a few. Their importance is further highlighted by the fact that they have recently received significant attention from the industry with notable projects including Google’s Private State Tokens [5] and Cloudflare’s Privacy Pass [6]. Concurrently, there has also been a standardization effort through IETF [7] supported by large industry stakeholders such as Google, Apple, Cloudflare, and Fastly.

Given the increased interest in anonymous tokens, a number of protocols have been proposed in the literature [1], [8], [9], [3], [10], [11] satisfying a variety of properties and offering interesting tradeoffs. An important distinction between these is whether the scheme supports private or public token verification. In short, private verification simplifies double-spending detection but is limited to scenarios with a single issuer and redeemtee, whereas public verification allows broader application, such as in blockchain, but relies on centralized mechanisms for double-spending detection. A second distinction is whether the scheme uses public or private metadata. Public metadata enables applications like geographical tags, while private metadata (typically a single bit) allows issuers to secretly flag malicious users. The idea is that the issuer can encode a hidden bit within the token, in a way that the user cannot distinguish which of the two bits their token encodes, and only someone with a secret extraction key (typically the issuer) can extract the embedded bit on later receiving a token.

An anonymous token scheme with public verifiability and private metadata bit, is suitable for applications where there is a single, central issuing authority that has enough context to derive trust signals for users, but at the same time there exist many possible redeemtees¹. The redeemtees can run the public verification algorithm and then submit

1. The setting of private and public metadata, in addition to public verifiability is also considered in the IETF standardization process [7].

the already verified tokens to the central authority that will check for double-spending. Having a secret metadata bit embedded in the token, can enable more efficient double-spending detection by potentially only checking the flagged tokens. As another example, one can consider applications running on a blockchain where public verification ensures transparency and accessibility, while private metadata allows for the inclusion of sensitive information or trust signals that can be detected by designated parties (the issuer) without exposing them to the public.

The problem with existing solutions. A common drawback among all existing schemes for anonymous tokens is the requirement for an interactive issuance protocol. Under this paradigm, the client initiates the protocol by preparing the first message which they then submit to the issuer. When the client receives the issuer’s response, they use it to locally construct the final token. However, the interactive nature of this protocol can lead to various practical challenges including latency issues, especially in real-time applications, and scalability concerns, particularly when the issuer must handle a large volume of requests simultaneously where a relatively expensive computational task needs to be executed (i.e., a cryptographic signature operation). This leads us to the following natural question:

Can we design a non-interactive anonymous token scheme that avoids the need for performing the expensive signature operation in an online fashion?

Our approach. We affirmatively answer this question by proposing the first anonymous token protocol with *non-interactive* issuance. Our main observation is that at the core of many of the proposed protocols is a blind signature scheme on a randomly selected message [1], [8], [3], [10]. By definition, blind signatures require at least one round (i.e., two moves) of interaction between the signer/issuer and the client [12], [13], [14]. However, recent works [15], [16], [17] have demonstrated the feasibility of constructing non-interactive blind signatures (NIBS) as long as the message is random and does not come from a specific distribution or has a specific structure.

The core idea behind non-interactive blind signatures is that each client is associated with a public-secret key pair (pk_C, sk_C) such that the signer can asynchronously create partial signatures $psig$, called *presignatures*, given only the client’s public key pk_C . The client, using their secret key sk_C , can extract the final pair of a blinded signature σ and message m from the presignature. The resulting σ is a valid signature for message m , and can be verified given only the signer’s verification key vk . The message m will be an unpredictable message for both the receiver and the signer derived from $psig$ and sk_C .

In anonymous tokens, the signed message is effectively a random identifier, thus the NIBS constructions of [15], [16] could give rise to an anonymous token scheme with public verifiability. The constructions of [15], [16] could also be turned into partially blind signatures and thus also

support the embedding of public metadata. However, extending this framework to support private metadata such as the embedding of a hidden bit—a crucial property for multiple applications—presents additional challenges, as evidenced by prior work on anonymous tokens with private metadata bit from interactive blind signatures [18]. Additionally, adding protection against double-spending in the offline setting has been an open problem (both for NIBS and anonymous tokens).

Our Contributions. We summarize our contributions as follows:

1 Defining non-interactive anonymous tokens. We propose a formal framework for *Non-interactive Anonymous Tokens with Private Metadata Bit* (NIAT). Our definitions are inspired by existing ATPM definitions [8], [10] and NIBS definitions [15], [16]. In particular, we translate the interactive ATPM token generation algorithm into a issuance algorithm and a token generation algorithm. At a high level, the issuer now runs the issuance algorithm to create a *presignature* (with an embedded bit) on a random message that it sends over to the client who, in turn, is able to locally run its token generation algorithm with the presignature as input to obtain the final token. To facilitate this type of non-interactive issuance, we need to introduce a key generation algorithm for the client. We also extend the standard security ATPM properties of one-more unforgeability, unlinkability and metadata bit privacy to the non-interactive setting. Respectively, these require that a NIAT attacker (i) with oracle access to the issuer’s algorithm should not be able to create more valid tokens than the ones received through oracle queries; (ii) with oracle access to the client’s algorithm should not be able to link any specific token to the corresponding presignature; and (iii) with oracle access to the issuer should not be able to distinguish a presignature under the bit 0 from that under the bit 1. Finally, we also propose the notion of reusability of a NIAT scheme that characterizes the requirement that an issuer is able to issue multiple tokens under the same client public key.

2 NIAT from SPS-EQ. In Section 4, we give our main NIAT construction from Structure-preserving Signatures on Equivalence Classes (SPS-EQ) [19]. Our SPS-EQ construction builds upon the SPS-EQ NIBS construction of [15]² to also support metadata bit hiding and extraction, and is secure under the inverse and strong decisional Diffie-Hellman assumptions as well as the unforgeability of the underlying SPS-EQ scheme. In our evaluations, we leverage the fact that verification of the SPS-EQ signature requires computation of pairing operations, some of which can be aggregated.

2. At the moment, the NIBS of [15] seems to be the best starting point for the most efficient NIAT constructions. The lattice-based design of [16] provides post-quantum security but suffers from much larger parameters, while the concurrent RSA-based construction of [17] is less efficient computationally and size-wise and targets a specific setting where the client’s public key is taken from a standard PKI infrastructure. In our case, we allow clients to register custom keys and, therefore, can use the more efficient construction from [15] as a starting point.

③ **NIAT with double-spend identification.** Our Section 4 construction requires an “online” check during verification in order to avoid double-spending (that the same token was not previously presented). While this is in line with previous works on anonymous tokens with private metadata bit, the requirement for an “online” check during verification can turn out to be pretty cumbersome. For instance, in application scenarios where verification happens by multiple distinct verifiers, it is not easy to keep a synchronized record of all spent tokens across all verifiers.

We define and present the first anonymous token scheme with a private metadata bit that supports double-spend (DS) *identification*. This feature allows for “offline” verification—meaning tokens can be presented without immediate verification of their uniqueness. If a token is double-spent, it can be identified post-event, enabling detection and potential penalties. In Section 5 we provide an extension of our main NIAT scheme to also support public double-spending identification support.

④ **Evaluation.** In Section 6, we experimentally evaluate our main SPS-EQ based construction using a proof-of-concept implementation. In our implementation, token issuance takes about 0.9 ms, token generation takes 1.4 ms (amortized) and token redemption takes 1.2 ms (amortized). The amortized costs are due to the aforementioned ability to aggregate some pairing operations. Thus, the cost of verification can be amortized over the number of presignatures issued (for token generation) and the number of clients serviced (for token redemption). Lastly, for completeness we provide an asymptotic comparison of our scheme with existing ATPM schemes, although we note that a direct comparison is difficult given that most existing schemes either do not support public verification or private metadata bit hiding and all schemes are interactive.

1.1. Application: Improved CDN Access

A major advantage of our main construction is the fact that it does not require any client–issuer interaction during the issuance protocol. This can allow for **offline pre-computations** leading to *significant savings* in the issuance process, especially if an issuance server is responsible for a large number of requests.

Let us consider one of the most popular applications of anonymous tokens—that of privacy preserving access to CDNs, which were also the inspiration for the Privacy Pass protocol [1]. Currently, in a CDN system, there exists an attesting server which, on a client request, examines the request and presents some kind of challenge puzzle, such as a CAPTCHA, if the server believes the incoming request is from an untruthworthy origin. The client will solve the puzzle and respond with the solution to prove that they are a real human user. The attesting server will check for correctness of the solution and respond accordingly. It may forward the client’s request to the web server that the client is wishing to access, in order to fetch dynamic content, or it may serve the client with a cached version of the static web

page. This is a reasonable approach for CDNs to prevent bot requests and DDoS attacks, and has been deployed in practice for a long time. However, it does not offer privacy for users and their history of web browsing can be traced back to the clients. Privacy Pass anonymous tokens were deployed to address such privacy issues [1].

In the setting of Privacy Pass, when a client sends the CAPTCHA solution to the attesting server, they also include a number of blinded tokens for the issuer to sign. If the issuer determines that the client is an honest human user, by checking CAPTCHA solutions, it signs the received blinded tokens *on-the-fly* and sends them back to the client alongside the content in the response. Later, when the client wants to access some web servers that are protected by the same infrastructure, they can redeem the tokens to bypass CAPTCHAs while maintaining their anonymity due to the tokens’ unlinkability properties. To avoid token hoarding, where the clients may be able to request a large number of tokens and redistribute them for future attacks; [1] recommended issuing 30 tokens during an issuance session as the number is considered to offer a reasonable trade-off between usability, performance, and the token hoarding issue. However, it is well known that online signing is a costly operation which, if possible, should be avoided in network protocols where a server serves a large number of clients simultaneously. This concern was pivotal in the development of protocols like DNSSEC.

Leveraging our NIAT primitive, we propose an alternative system for token issuance that *eliminates the need for online interaction* between clients and issuers and circumvents the necessity for online signing. We describe the idea in Figure 1. This system includes two distinct phases: a one-time registration phase, and a request phase. During the registration phase, clients interact with the issuer to prove their legitimacy by solving some CAPTCHA puzzles. When responding to a challenge, a client will include their public key pk_C alongside the puzzle solution. After verifying a client’s solution, the issuer will register the public key pk_C , and is then able to proactively create a batch of presignatures for the newly registered client offline. We emphasize that this step is only required once per-user. Notably, an issuer only requires one short message from the client to process a large batch of presignatures. Once a client has registered with the issuer and has been deemed trustworthy, they may start requesting tokens. Since the issuer has the presignatures readily at-hand, the client does not need to wait for online signing processes to occur. This results in a significantly faster turnaround during the request phase. On the issuer’s side, they can create new presignatures offline and as long as they always have a new batch of presignatures prepared for all registered clients, the amortized cost for the issuance protocol no longer depends on the number of tokens requested. Rather, only client authentication and network delays incur the more significant costs during the request phase.

Public verifiability and embedding a private metadata bit are directly supported by instantiating the above described system with our NIAT protocol. This can allow clients to verify their tokens with multiple servers while also carrying

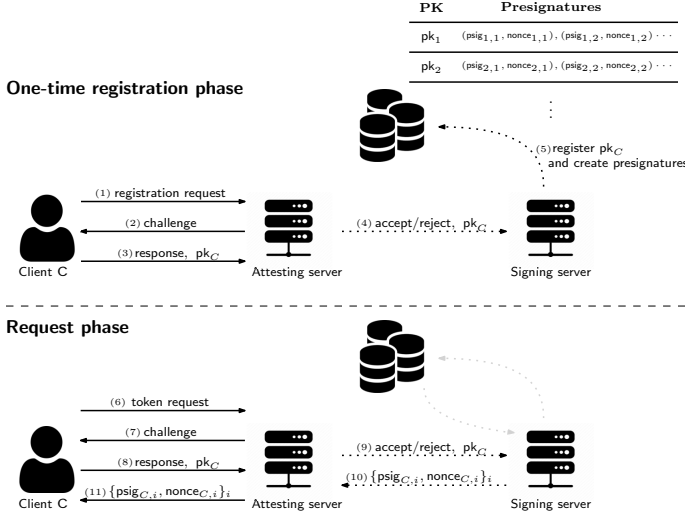


Figure 1. NIAT issuance for CDNs with offline signing.

trust signals that can be extracted by the attesting server.

1.2. Related Work

Privacy Pass [1] introduced the idea of one-time use anonymous tokens as a method to prevent DDoS attacks while enhancing privacy and ensuring a seamless user experience for those who access web servers protected by Cloudflare infrastructure. In the simplest setting, Privacy Pass is an interactive, two-move, client-initiated protocol which relies on Verifiable Oblivious Pseudorandom Functions (VOPRFs), and is secure in the random oracle model. The recent partially oblivious PRFs [20] may be used in place of VOPRFs for a more efficient and flexible instantiation of Privacy Pass. One disadvantage of the Privacy Pass protocol is that if a client is deemed malicious, they will not be issued tokens and the request would be dropped. This kind of feedback may inform malicious actors of their detection, which could be leveraged to refine their methods to circumvent bot-detection mechanisms.

To address this problem, Kreuter et al. proposed the idea of anonymous tokens with private metadata bit [8] in order to propagate trust signal from the issuers to the verifiers in a private way. They formalized the security properties of this primitive, and proposed a construction called PMBTokens which are essentially an extension of the Privacy Pass protocol to support the private metadata bit. Like Privacy Pass, PMBTokens are also based on VOPRFs and are secure in the random oracle model. Also like Privacy Pass, PMBTokens do not allow any parties other than the authoritative parties holding the secret issuing keys to verify the validity of tokens or to extract the private metadata bits in the case of PMBTokens. On the other hand, publicly verifiable anonymous tokens with private metadata bit [9] allow for public token verification under the issuer's public key, while the private metadata bit remains hidden from users and is only extractable with the issuer's secret

key. Publicly verifiable anonymous tokens have also been instantiated from RSA blind signatures [18], however they do not support private metadata.

Chase et al. [10], proposed ATHM which is an anonymous token protocol for the setting where the issuer and verifier are the same entity. ATHM is based on a symmetric key primitive, namely algebraic MACs, and is secure in the generic group model. In addition, they also revisited the definition of anonymous tokens and merged the two notions of token validity in [8]. Specifically, in the new definition, only tokens from which we can extract the embedded metadata bit successfully are considered valid tokens, and the private verification algorithm `AT.Verify` was removed for redundancy. ATHM can also be extended to support public metadata in tokens.

Anonymous Counting Tokens (ACTs) [21] are another variant of anonymous tokens where clients are not able to redeem more than one token per message with the same verifier. This security property can be achieved through two different approaches: either by allowing the issuer to detect repeated token issuance requests for the same message from the same client during issuance, or by allowing the verifier, during redemption, to detect that two different tokens for the same message were issued to the same client.

2. Preliminaries

Notation. We use λ to denote the security parameter. We use $\leftarrow \$$ to denote the output of a randomized algorithm, \leftarrow to denote output of a deterministic algorithm, and $:=$ for assignment. In all that follows, \mathbb{Z} is the ring of integers, and \mathbb{Z}_p denotes the set of integers modulo p . \mathbb{G} is a multiplicative group of prime order p . Lastly, for a vector \mathbf{x} , we write x_i as its i^{th} element.

2.1. Bilinear Pairings

Definition 2.1 (Bilinear pairings). *Given a security parameter λ , a bilinear group generator $\text{BG}(1^\lambda)$ returns a tuple $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, p, g_1, g_2)$, where $(\mathbb{G}_1, \mathbb{G}_2)$ are groups of the same prime order p with the generators g_1, g_2 respectively. Also, let \mathbb{Z}_p be the field of order p . A bilinear pairing is an efficiently computable map, $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$, satisfying the following properties:*

- **Bilinearity:** $\forall P \in \mathbb{G}_1, Q \in \mathbb{G}_2, a, b \in \mathbb{Z}_p,$

$$e(P^a, Q)^b = e(P, Q^b)^a = e(P, Q)^{ab}.$$

- **Non-degeneracy:** $e(g_1, g_2) \neq 1_{\mathbb{G}_T}$.

Bilinear pairings can be of a few types depending on whether there is an efficient isomorphism from \mathbb{G}_1 to \mathbb{G}_2 in both directions (Type 1), only one direction (Type 2), or in neither direction (Type 3) [22]. Type 3 pairings are the most efficient setting for a relevant security parameter and they are commonly deployed.

2.2. Hardness Assumptions

Definition 2.2 (Inverse decisional Diffie-Hellman). *For all PPT adversaries \mathcal{A} given tuple $(\mathbb{G}, g, g^\alpha, g^\beta)$, it is hard to decide whether $\beta = \alpha^{-1} \bmod p$ or $\beta \leftarrow \mathbb{Z}_p$.*

Definition 2.3 (k -decisional Diffie-Hellman). *For all PPT adversaries \mathcal{A} given tuple $(\mathbb{G}, g, g^\alpha, g^{\alpha^2}, \dots, g^{\alpha^k}, g^\beta)$, it is hard to decide whether $\beta = \alpha^{-1} \bmod p$ or $\beta \leftarrow \mathbb{Z}_p$.*

2.3. Structure-preserving Signatures on Equivalence Classes

Structure-preserving Signatures on Equivalence Classes (SPS-EQ) [19] are used to sign equivalence classes $[M]$ of message vectors $M \in (\mathbb{G}_i^*)^\ell$ for $\ell > 1$, having the equivalence relation $M, N \in \mathbb{G}_i^\ell : M \sim_{\mathcal{R}} N \Leftrightarrow \exists s \in \mathbb{Z}_p : M = N^s$. Let us now recall the definition for SPS-EQ [19], adapting the notation for multiplicative groups as in [15].

Definition 2.4 (SPS-EQ). *An SPS-EQ scheme consists of the following PPT algorithms:*

- $\text{Setup}(1^\lambda, \ell) \rightarrow \text{pp}_{\text{EQ}}$. On input the security parameter 1^λ , and the length of message vectors ℓ , it outputs the public parameters pp_{EQ} .
- $\text{KeyGen}(\text{pp}) \rightarrow (\text{pk}_{\text{EQ}}, \text{sk}_{\text{EQ}})$. On input the public parameters pp , it outputs a key pair $(\text{pk}_{\text{EQ}}, \text{sk}_{\text{EQ}})$.
- $\text{VerKey}(\text{sk}_{\text{EQ}}, \text{pk}_{\text{EQ}}) \rightarrow b \in \{0, 1\}$. On input a public-secret key pair $(\text{pk}_{\text{EQ}}, \text{sk}_{\text{EQ}})$, it deterministically returns a bit $b \in \{0, 1\}$.
- $\text{Sign}(\text{sk}_{\text{EQ}}, M) \rightarrow \sigma_{\text{EQ}}$. On input the secret key sk_{EQ} and a representative $M \in (\mathbb{G}_i^*)^\ell$, it outputs a signature σ_{EQ} for the equivalence class $[M]$.
- $\text{ChRep}(M, \sigma_{\text{EQ}}, \mu) \rightarrow \sigma'_{\text{EQ}}$. On input a representative M , a signature σ_{EQ} and a scalar μ , it returns an updated message-signature pair $(M', \sigma'_{\text{EQ}})$ where the new representative is $M' = M^\mu$ and σ'_{EQ} is the corresponding updated signature.
- $\text{Verify}(\text{pk}_{\text{EQ}}, M, \sigma_{\text{EQ}}) \rightarrow b \in \{0, 1\}$. On input a public key pk_{EQ} , a representative M , and a signature σ_{EQ} , it deterministically outputs a bit $b \in \{0, 1\}$.

An SPS-EQ scheme satisfies correctness, existentially unforgeable under adaptive chosen-message attacks (EUnf-CMA), and perfect signature adaptation under malicious keys. We recall the definitions from [19].

Definition 2.5 (Correctness). *An SPS-EQ scheme over \mathbb{G}_i^* is correct if for all security parameters $\lambda \in \mathbb{N}$, $\ell > 1$, $\text{pp}_{\text{EQ}} \leftarrow \text{EQ.Setup}(1^\lambda, \ell)$, $(\text{pk}_{\text{EQ}}, \text{sk}_{\text{EQ}}) \leftarrow \text{EQ.KeyGen}(\text{pp}_{\text{EQ}})$, $M \in (\mathbb{G}_i^*)^\ell$, such that $\sigma_{\text{EQ}} \leftarrow \text{EQ.Sign}(\text{sk}_{\text{EQ}}, M)$ and for all scalars $\mu \in \mathbb{Z}_p$ we have*

$$\text{EQ.VerKey}(\text{sk}_{\text{EQ}}, \text{pk}_{\text{EQ}}) = 1 \wedge$$

$$\Pr[\text{EQ.Verify}(\text{pk}_{\text{EQ}}, M, \sigma_{\text{EQ}}) = 1] = 1 \wedge$$

$$\Pr[\text{EQ.Verify}(\text{pk}_{\text{EQ}}, M^\mu, \text{EQ.ChRep}(\sigma_{\text{EQ}}, \mu)) = 1] = 1$$

Definition 2.6 (Existential unforgeability under chosen message attack). *Let $\text{Adv}_{\mathcal{A}, \ell}^{\text{EUnf-CMA}} =$*

Game $\text{EUnf-CMA}_{\mathcal{A}, \ell}(\lambda)$

```

1 :  $Q := \emptyset$ 
2 :  $\text{pp}_{\text{EQ}} \leftarrow \text{EQ.Setup}(1^\lambda, \ell)$ 
3 :  $(\text{pk}_{\text{EQ}}, \text{sk}_{\text{EQ}}) \leftarrow \text{EQ.KeyGen}(\text{pp}_{\text{EQ}})$ 
4 :  $(M^*, \sigma_{\text{EQ}}^*) \leftarrow \mathcal{A}^{\text{O}_{\text{Sign}}}(\text{pk}_{\text{EQ}})$ 
5 : return  $[M^*] \neq [M] \forall M \in Q \wedge$ 
       $\text{EQ.Verify}(\text{pk}_{\text{EQ}}, M^*, \sigma_{\text{EQ}}^*) = 1$ 

```

Oracle $\text{O}_{\text{Sign}}(M)$

```

1 :  $\sigma_{\text{EQ}} \leftarrow \text{EQ.Sign}(\text{sk}_{\text{EQ}}, M)$ 
2 :  $Q := Q \cup \{M\}$ 
3 : return  $\sigma_{\text{EQ}}$ 

```

Figure 2. The existential unforgeability (under chosen message attack) experiment for SPS-EQ.

$\Pr[\text{EUnf-CMA}_{\mathcal{A}, \ell}(\lambda) = \text{accept}]$ be the advantage of an PPT adversary \mathcal{A} in the EUnf-CMA experiment defined in Figure 2. An SPS-EQ scheme over $(\mathbb{G}_i^*)^\ell$ is existentially unforgeable under adaptive chosen-message attacks if for all $\ell > 1$, the advantage $\text{Adv}_{\mathcal{A}}^{\text{EUnf-CMA}} = \text{negl}(\lambda)$.

Definition 2.7 (Perfect signature adaptation under malicious keys). *For $\ell > 1$, an SPS-EQ scheme on $(\mathbb{G}_i^*)^\ell$ perfectly adapts signatures under malicious keys if for all tuples $(\text{pk}_{\text{EQ}}, M, \sigma_{\text{EQ}}, \mu)$ satisfying*

$$M \in \mathbb{G}_i^* \wedge \text{EQ.Verify}(\text{pk}_{\text{EQ}}, M, \sigma_{\text{EQ}}) = 1 \wedge \mu \in \mathbb{Z}_p$$

we have that the output of $\text{EQ.ChRep}(\sigma_{\text{EQ}}, \mu)$ is a uniformly random element in the space of signatures, conditioned on $\text{EQ.Verify}(\text{pk}_{\text{EQ}}, M^\mu, \sigma'_{\text{EQ}}) = 1$.

In this work, we will use the pairing-based construction of SPS-EQ due to [19]. For completeness, we recall their construction in Appendix D.

3. Non-interactive Anonymous Tokens

Formally, a non-interactive anonymous tokens (NIAT) scheme consists of the following polynomial time algorithms (values in light gray are required only for NIAT with double-spend protection):

- $\text{Setup}(1^\lambda) \rightarrow (\text{pp}, \text{aux})$. The setup algorithm takes a security parameter λ as input and outputs a set of common public parameters pp (and some auxiliary data aux). All of the remaining algorithms take pp as an input, but for notational clarity, we usually omit it as an explicit input.
- $\text{KeyGen}_I(\text{pp}) \rightarrow (\text{pk}_I, \text{sk}_I, \text{ek}_I)$. The issuer's key generation algorithm takes public parameters pp as input, and outputs a public key pk_I , a secret signing key sk_I , and an extraction key ek_I .
- $\text{KeyGen}_C(\text{pp}) \rightarrow (\text{pk}_C, \text{sk}_C)$. The client's key generation algorithm takes pp as input, and outputs a public-secret key pair $(\text{pk}_C, \text{sk}_C)$.
- $\text{Issue}(\text{sk}_I, \text{ek}_I, \text{pk}_C, b, \tau, \text{aux}) \rightarrow (\text{psig}, \text{nonce}, \text{aux})$. The issuer runs the probabilistic algorithm Issue

which takes as input the issuer's secret signing key sk_I , the extraction key ek_I , the client's public key pk_C , a private metadata bit $b \in \{0, 1\}$, and public metadata τ (and some auxiliary data aux). It then outputs a pre-signature $psig$, and a random nonce (and an updated aux).

- $\text{Obtain}(sk_C, pk_I, psig, nonce, \tau, id) \rightarrow (t, \sigma)$ or \perp . The probabilistic Obtain algorithm is run by the client to obtain the final token. It takes as input the client's secret key sk_C , the issuer's public key pk_I , a pre-signature $psig$, and nonce (and a verifier identifier id). It outputs a token (t, σ) if the algorithm runs successfully, or \perp otherwise.
- $\text{Verify}(pk_I, t, \sigma, \tau) \rightarrow b \in \{0, 1\}$. The deterministic public verification algorithm takes as input the issuer's public key pk_I , the token (t, σ) and the public metadata τ , and outputs a bit $b \in \{0, 1\}$.
- $\text{ReadBit}(ek_I, t) \rightarrow b \in \{0, 1\}$ or \perp . On input the issuer's extraction key and a tag t , the deterministic ReadBit algorithm outputs a bit $b \in \{0, 1\}$ or \perp ³.

Furthermore, for a NIAT with double-spend protection, we additionally define:

- $\text{DSIden}(pk_I, t, \sigma, \sigma', aux) \rightarrow (pk_C, \Pi)$ or (\perp, \perp) . The double-spend identification algorithm takes as input the issuer's public key pk_I , a tag t , two different presentations⁴ (σ, σ') with respect to t , and some auxiliary data aux , and outputs the embedded client public key and a proof of guilt (pk_C, Π) or (\perp, \perp) .
- $\text{DSVer}(pk_C, \Pi, aux) \rightarrow b \in \{0, 1\}$. The double-spend verification algorithm takes a client public key pk_C , a proof of guilt Π and the auxiliary data aux , and outputs a bit $b \in \{0, 1\}$.

3.1. Security Definitions

The reader will observe that our interface has two algorithms that could signal whether or not a token is valid. Namely, the Verify and ReadBit algorithms. However, as pointed out by [10], having two different notions of validity can potentially open up new attack vectors against an anonymous token protocol. For instance, a malicious client could forge tokens that yield a valid bit, but fail verification or vice-versa. This can be concerning in situations where distinct parties are performing ReadBit and Verify. To address this problem, we will insist that the correctness of the scheme hold only when the outcomes of ReadBit and Verify are consistent. In practice, this means that the ReadBit algorithm should only be run on tags whose corresponding tokens pass Verify. This unifies our two notions of validity and facilitates public verification, which is necessary for

3. Note that ReadBit is a deterministic function of the issuer's secret key and the tag.

4. We consider schemes where the tag t is unique during the execution of the Obtain algorithm (i.e. each tag t is uniquely determined by $(psig, nonce)$), while the σ part can be prepared during presentation and differ for every possible verifier.

applications where token verification is outsourced to a standalone server.

Definition 3.1 (Correctness). *A non-interactive anonymous tokens scheme is correct if for every security parameter $\lambda \in \mathbb{N}$ such that $pp \leftarrow \$ \text{Setup}(1^\lambda)$, $(pk_I, sk_I, ek_I) \leftarrow \$ \text{KeyGen}_I(pp)$, $(pk_C, sk_C) \leftarrow \$ \text{KeyGen}_C(pp)$, and for every $b \in \{0, 1\}$ and public metadata τ , such that $(psig, nonce) \leftarrow \$ \text{Issue}(sk_I, ek_I, pk_C, b, \tau)$ and $(t, \sigma) \leftarrow \$ \text{Obtain}(sk_C, pk_I, psig, nonce, \tau)$, we have*

$$\Pr[\text{Verify}(pk_I, t, \sigma, \tau) = 1 \wedge \text{ReadBit}(ek_I, t) = b] = 1$$

We also formalize the notion of *reusability* for NIAT based on a similar notion from the NIBS literature [16]. This property is not needed in the ATPM definition given by Chase et al. [10] because their scheme is interactive. More precisely, in an ATPM scheme, whenever a client initiates a token issuance session, the client sends to the issuer a query which suffices to uniquely identify that session. This is not possible in the non-interactive setting, as the only information that the issuer knows about a client is its public key pk_C . Reusability allows us to capture the meaningful property that an issuer should be able to issue multiple tokens to the same client's public key pk_C . In other words, an issuer can *reuse* a client's public key pk_C to issue tokens non-interactively.

Definition 3.2 (Reusability). *A non-interactive anonymous token scheme is reusable if for every security parameter $\lambda \in \mathbb{N}$, metadata bit $b \in \{0, 1\}$ and public metadata $\tau \in \mathcal{M}$ the following probability is negligible*

$$\Pr \left[\begin{array}{l} \text{nonce}_0 = \text{nonce}_1 \vee t_0 = t_1 : \\ \begin{array}{l} pp \leftarrow \$ \text{Setup}(1^\lambda) \\ (pk_I, sk_I, ek_I) \leftarrow \$ \text{KeyGen}_I(pp) \\ (pk_C, sk_C) \leftarrow \$ \text{KeyGen}_C(pp) \\ \forall i \in \{0, 1\} : (psig_i, nonce_i) \leftarrow \$ \text{Issue} \left(\begin{array}{l} sk_I, ek_I, \\ pk_C, b, \tau \end{array} \right) \\ \forall i \in \{0, 1\} : (t_i, \sigma_i) \leftarrow \$ \text{Obtain} \left(\begin{array}{l} sk_C, pk_I, \\ psig_i, nonce_i, \tau \end{array} \right) \end{array} \end{array} \right]$$

We now define one-more unforgeability for NIAT. A minor, but necessary, point of distinction from the one-more unforgeability definition of Chase et al. [10] is that the Issue oracle $\mathcal{O}_{\text{Issue}}$ (corresponding to their $\mathcal{O}_{\text{Sign}}$) accepts client public keys in place of the client's query message. More notably, in our setting, a forgery is considered valid if the public verification algorithm (and not the private bit extraction algorithm) accepts. This is to better align with applications of public verifiability, where it would be significantly more problematic for the Verify to pass and ReadBit to fail, than the other way around since, if Verify fails, ReadBit can trivially be forced to fail by requiring that ReadBit only be run if Verify passes.

As a result of this change in definition, we obtain a stronger security guarantee than the one considered in previous work [10] where the forgeries are fixed to the same bit that must, in turn, be extractable from said forgeries in order for the verifier to win. In particular, the previous definition

does not prevent against adversaries that are able to construct forgeries by manipulating the underlying bit. This is not a significant issue in schemes without public verifiability as the privacy of the bit ensures that (barring any side-channels) the adversary does not benefit from mauling the underlying bit. However, in the case of public verifiability, the ability to maul the underlying bit is still a risk as the adversary may still be able to pass verification.

Definition 3.3 (One-more unforgeability). *Let $\text{Adv}_{\mathcal{A}}^{\text{OM-Unf}} = \Pr[\text{OM-Unf}_{\mathcal{A}}(\lambda) = \text{accept}]$ be the advantage of an adversary \mathcal{A} in the experiment defined in Figure 3. We say a NIAT scheme NIAT is one-more unforgeable if for any PPT adversary \mathcal{A} this advantage is negligible.*

Game OM-Unf $_{\mathcal{A}}(\lambda)$

```

1 : pp, aux  $\leftarrow$  Setup( $1^\lambda$ )
2 : (pkI, skI, ekI)  $\leftarrow$  KeyGenI(pp)
3 :  $\tau^*, \{(t_i, \sigma_i)\}_{i \in [N]} \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Issue}}, \mathcal{O}_{\text{Read}}}(\text{pp}, \text{pk}_I)$ 
4 : if  $\#\{t_i\}_{i \in [N]} \leq \ell$  then reject
5 : if  $\bigvee_{i \in [N]} \text{Verify}(\text{pk}_I, t_i, \sigma_i, \tau^*) \neq 1$  then reject
6 : accept

```

Oracle $\mathcal{O}_{\text{Issue}}(\text{pk}_C, b, \tau)$

```

1 : if  $\ell$  is uninitialized then  $\ell := 0$ 
2 : (psig, nonce, aux)  $\leftarrow$  Issue(skI, ekI, pkC, b,  $\tau$ , aux)
3 :  $\ell := \ell + 1$ 
4 : return (psig, nonce)

```

Oracle $\mathcal{O}_{\text{Read}}(t, \sigma, \tau)$

```

1 : return ReadBit(ekI, t)

```

Figure 3. The one-more unforgeability experiment for NIAT.

Next, we formalize NIAT unlinkability. The overall intuition here remains the same as in [10], except that the adversary is now equipped with a GenUser oracle $\mathcal{O}_{\text{GenUser}}$ that issues new public keys, and an Obtain oracle $\mathcal{O}_{\text{Obtain}}$ that returns the tokens on chosen (psig, nonce, τ) triples.

Definition 3.4 (κ -unlinkability). *Let $\text{Adv}_{\mathcal{A}, n}^{\text{UNLINK}} = \Pr[\text{UNLINK}_{\mathcal{A}, n}(\lambda) = \text{accept}]$ be the advantage of an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ for parameter $n \in \mathbb{N}$ in the experiment defined in Figure 4. We say a NIAT scheme NIAT is κ -unlinkable if for any PPT adversary \mathcal{A} and $n \in \mathbb{N}$, this advantage at most $\kappa/n + \text{negl}(\lambda)^5$.*

Concretely, the parameter n is the number of pairwise distinct $(\text{pk}_C, \text{psig}, \text{nonce})$ triples in the challenge set Q^* .

5. Intuitively, the probability κ/n quantifies the best strategy for an honest issuer that has κ choices for each private metadata b_i . In the case that the private metadata is a bit, ATPM (and thus also NIAT) schemes generally focus on the case where $\kappa = 2$.

Game UNLINK $_{\mathcal{A}, n}(\lambda)$

```

1 : pp, aux  $\leftarrow$  Setup( $1^\lambda$ )
2 : (pkI, state1)  $\leftarrow$   $\mathcal{A}_1(\text{pp}, \text{aux})$ 
3 : ( $Q^*, \tau, \text{id}, \text{state}_2$ )  $\leftarrow$   $\mathcal{A}_2^{\mathcal{O}_{\text{GenUser}}, \mathcal{O}_{\text{Obtain}}}(\text{state}_1)$ 
4 : foreach (pkC, psig, nonce)  $\in Q^*$ 
5 :   if (pkC,  $\cdot$ )  $\notin Q_{\text{usr}}$ 
6 :      $\vee (\text{pk}_C, \cdot, \text{nonce}) \in Q_{\text{obt}}$  then reject
7 : if  $|Q^*| \neq n$  then reject
8 : foreach  $i \in [n]$  do
9 :   outi  $\leftarrow$  Obtain(skCi, pkI, psigi, noncei,  $\tau$ , id)
10 :  if outi =  $\perp$  then reject
11 :  $\hat{i} \leftarrow [n]$ 
12 : Pick a random permutation  $\varphi : [n] \rightarrow [n]$ 
13 :  $i^* \leftarrow \mathcal{A}_3(\text{state}_2, \text{out}_{\hat{i}}, (\text{out}_{\varphi(i)})_{i \in [n] \setminus \{\varphi^{-1}(\hat{i})\}})$ 
14 : if  $i^* = \hat{i}$  then accept

```

Oracle $\mathcal{O}_{\text{GenUser}}(\text{pp})$

```

1 : if  $Q_{\text{usr}}$  is uninitialized then  $Q_{\text{usr}} := \emptyset$ 
2 : (pkC, skC)  $\leftarrow$  KeyGenC(pp)
3 :  $Q_{\text{usr}} := Q_{\text{usr}} \cup \{(\text{pk}_C, \text{sk}_C)\}$ 
4 : return pkC

```

Oracle $\mathcal{O}_{\text{Obtain}}(\text{pk}_C, \text{psig}, \text{nonce}, \tau, \text{id})$

```

1 : if  $Q_{\text{obt}}$  is uninitialized then  $Q_{\text{obt}} := \emptyset$ 
2 : if (pkC,  $\cdot$ )  $\notin Q_{\text{usr}}$  then return  $\perp$ 
3 : Retrieve skC from  $Q_{\text{usr}}$ 
4 : out  $\leftarrow$  Obtain(skC, pkI, psig, nonce,  $\tau$ , id)
5 : if out  $\neq \perp$  then
6 :    $Q_{\text{obt}} := Q_{\text{obt}} \cup \{(\text{pk}_C, \text{psig}, \text{nonce})\}$ 
7 : return out

```

Figure 4. The unlinkability experiment for NIAT.

Notice that we do not enforce any other constraints on the values inside Q^* . This allows for cases where some public keys can repeat. In particular, when $n = 2$, $\kappa = 1$, and $\text{pk}_{C_1} \neq \text{pk}_{C_2}$ (resp. $\text{pk}_{C_1} = \text{pk}_{C_2}$), we can view this experiment as being analogous to the stronger versions of (T)NIBS receiver (resp. nonce) blindness [16]. Thus we view our unlinkability definition as a generalization of the (T)NIBS blindness definitions.

Remark 3.5. An important requirement for unlinkability is that all (psig, nonce) pairs in Q^* must be under the same public metadata, otherwise the property is trivially broken. This also holds for the definition given in [10].

We then define the equivalent notion of privacy for the metadata bit in a NIAT.

Definition 3.6 (Privacy of metadata bit). *Let $\text{Adv}_{\mathcal{A}}^{\text{PMB}} = |\Pr[\text{PMB}_{\mathcal{A}, 1}(\lambda) = \text{accept}] - \Pr[\text{PMB}_{\mathcal{A}, 0}(\lambda) = \text{accept}]|$ be the advantage of an adversary \mathcal{A} in the experiment defined in Figure 5. We say a NIAT scheme NIAT preserves*

Game $\text{PMB}_{\mathcal{A}, \hat{b}}(\lambda)$
<hr/> 1 : $\text{pp}, \text{aux} \leftarrow \text{Setup}(1^\lambda)$ 2 : $(\text{pk}_I, \text{sk}_I, \text{ek}_I) \leftarrow \text{KeyGen}_I(\text{pp})$ 3 : $\text{flag} := \text{False}$ 4 : $b^* \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Issue}}, \mathcal{O}_{\text{Read}}, \mathcal{O}_{\text{Challenge}}}(\text{pp}, \text{pk}_I)$ 5 : if $b^* = \hat{b}$ then accept
Oracle $\mathcal{O}_{\text{Issue}}(\text{pk}_C, b, \tau)$
<hr/> 1 : $(\text{psig}, \text{nonce}, \text{aux}) \leftarrow \text{Issue}(\text{sk}_I, \text{ek}_I, \text{pk}_C, b, \tau, \text{aux})$ 2 : return (s)
Oracle $\mathcal{O}_{\text{Read}}(t, \sigma, \tau)$
<hr/> 1 : if $\text{flag} \wedge (\tau = \hat{\tau})$ then return \perp 2 : if $\text{Verify}(\text{pk}_I, t, \sigma, \tau) = 1$ then 3 : return $\text{ReadBit}(\text{ek}_I, t)$ 4 : return \perp
Oracle $\mathcal{O}_{\text{Valid}}(t, \sigma, \tau)$
<hr/> 1 : $b \leftarrow \text{ReadBit}(\text{ek}_I, t)$ 2 : return $(b \neq \perp \wedge \text{Verify}(\text{pk}_I, t, \sigma, \tau) = 1)$
Oracle $\mathcal{O}_{\text{Challenge}}(\text{pk}_C, \tau)$
<hr/> 1 : if flag then return \perp 2 : $\text{flag} := \text{True}, \hat{\tau} := \tau$ 3 : $(\text{psig}, \text{nonce}, \text{aux}) \leftarrow \text{Issue}(\text{sk}_I, \text{ek}_I, \text{pk}_C, \hat{b}, \hat{\tau}, \text{aux})$ 4 : return $(\text{psig}, \text{nonce})$

Figure 5. The privacy of metadata bit experiment for NIAT.

privacy of the metadata bit if for any PPT adversary \mathcal{A} , this advantage is negligible.

NIAT with double-spending identification security. For a NIAT scheme with double-spending identification security we define two additional properties: double-spending identification and exculpability. Double-spending identification guarantees that no adversary is able to present the same token twice without having their public key revealed. Formally, for oracles defined as above,

Definition 3.7 (Double-spending identification). *Let $\text{Adv}_{\mathcal{A}}^{\text{Iden}} = \Pr[\text{Iden}_{\mathcal{A}}(\lambda) = \text{accept}]$ be the advantage of an adversary \mathcal{A} in the experiment defined in Figure 6. We say a NIAT scheme NIAT achieves double-spending identification if for any PPT adversary \mathcal{A} , this advantage is negligible.*

We finally define exculpability which guarantees that an issuer cannot wrongly accuse an honest client of double-spending. Formally,

Definition 3.8 (Double-spending exculpability). *Let $\text{Adv}_{\mathcal{A}}^{\text{Excul}} = \Pr[\text{Excul}_{\mathcal{A}}(\lambda) = \text{accept}]$ be the advantage of an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in the experiment defined in*

Game $\text{Iden}_{\mathcal{A}}(\lambda)$
<hr/> 1 : $(\text{pp}, \text{aux}) \leftarrow \text{Setup}(1^\lambda)$ 2 : $(\text{pk}_I, \text{sk}_I, \text{ek}_I) \leftarrow \text{KeyGen}_I(\text{pp})$ 3 : $(t, \sigma, \sigma') \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Issue}}, \mathcal{O}_{\text{Read}}}(\text{pp}, \text{pk}_I)$ 4 : if $\sigma = \sigma' \vee \text{Verify}(\text{pk}_I, t, \sigma, \tau) \neq 1$ $\vee \text{Verify}(\text{pk}_I, t, \sigma', \tau) \neq 1$ then reject 5 : $(\text{pk}_C, \cdot) \leftarrow \text{DSIden}(\text{pk}_I, t, \sigma, \sigma', \text{aux})$ 6 : if $\text{pk}_C \in Q \setminus \{\perp\}$ then reject 7 : accept

Figure 6. The double-spending identification experiment for NIAT. Oracles $\mathcal{O}_{\text{Issue}}$ and $\mathcal{O}_{\text{Read}}$ are as defined in Figure 3, except $\mathcal{O}_{\text{Issue}}$ stores all pk_C queried by the adversary in a list Q .

Figure 7. We say a NIAT scheme NIAT achieves double-spending exculpability if for any PPT adversary \mathcal{A} , this advantage is negligible.

Game $\text{Excul}_{\mathcal{A}}(\lambda)$
<hr/> 1 : $(\text{pp}, \text{aux}) \leftarrow \text{Setup}(1^\lambda)$ 2 : $(\text{pk}_I, \text{state}_1) \leftarrow \mathcal{A}_1(\text{pp})$ 3 : $(\text{pk}_C, \Pi) \leftarrow \mathcal{A}_2^{\mathcal{O}_{\text{GenUser}}, \mathcal{O}_{\text{Obtain}}}(\text{state}_1)$ 4 : if $(\text{pk}_C, \cdot) \in Q_{\text{usr}} \wedge \text{DSVer}(\text{pk}_C, \Pi, \text{aux}) = 1$ then accept

Figure 7. The double-spend exculpability experiment for NIAT. Oracles $\mathcal{O}_{\text{GenUser}}$ and $\mathcal{O}_{\text{Obtain}}$ are as defined in Figure 4, except $\mathcal{O}_{\text{Obtain}}$ behaves as an honest user that accepts a given $(\text{psig}, \text{nonce})$ pair exactly once.

4. NIAT from Equivalence Class Signatures

We now present our NIAT construction from structure-preserving signatures on equivalence classes. At a high level, we extend the NIBS construction of [15] to additionally embed the private metadata bit into the issuer's presignature. This is a non-trivial exercise as the bit must also be extractable from the final token, under the issuer's secret key. We resolve this by hiding the bit inside the exponent of a random group element and sending the result as part of the presignature. Thus, when a client sends the re-randomized signature, the part corresponding to the hidden bit can be checked in the exponent using the issuer's secret key.

In all that follows, we will ignore the public metadata τ for simplicity, instead noting that (and as we show in Appendix F) all our constructions are extendable to include τ in a manner similar to [15].

4.1. Construction

Tools required. Our construction requires a hash function $H : \{0, 1\}^\lambda \rightarrow \mathbb{G}_1$ modeled as a random oracle, an SPS-EQ scheme $\text{EQ} = (\text{EQ.Setup}, \text{EQ.KeyGen}, \text{EQ.Sign},$

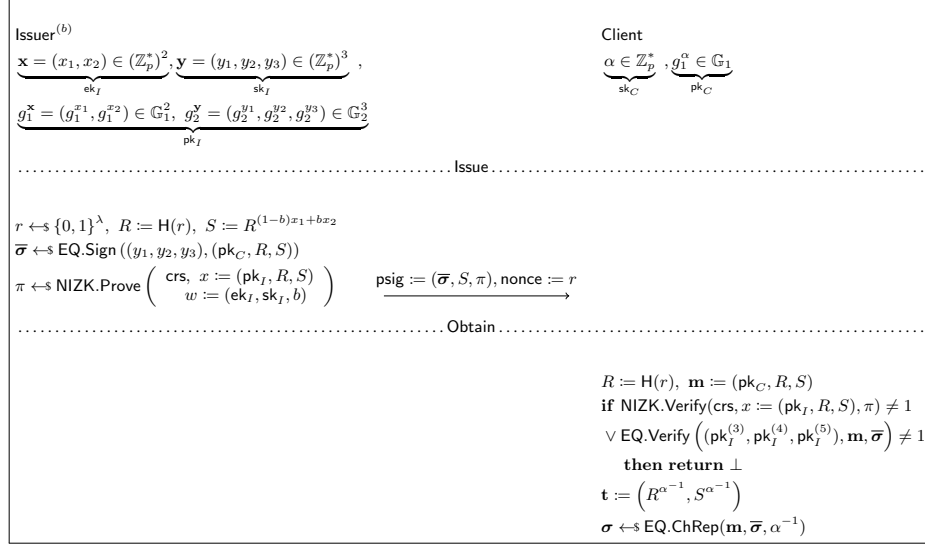


Figure 8. SPS-EQ construction for NIAT presignature issuance and token generation.

EQ.ChRep, EQ.Verify, EQ.VerKey), and a NIZK NIZK = (NIZK.Setup, NIZK.Prove, NIZK.Verify) for the following language:

Language \mathcal{L}_{iss}

Instance: Each instance x is interpreted as a collection of group elements pk_I and elements R, S .

Witness: Witness w consists of an integer vectors $\text{ek}_I := (x_1, x_2)$, $\text{sk}_I := (y_1, y_2, y_3)$ and a bit b .

Membership: w is a valid witness for x if the following are satisfied:

$$b \in \{0, 1\} \wedge S = R^{(1-b)x_1 + bx_2} \wedge$$

$$\text{pk}_I = (g_1^{x_1}, g_1^{x_2}, g_2^{y_1}, g_2^{y_2}, g_2^{y_3})$$

System setup. This algorithm sets up the generators for the bilinear groups \mathbb{G}_1 and \mathbb{G}_2 of prime order p . It also runs the NIZK setup to generate the crs for the language \mathcal{L}_{iss} .

Setup(1^λ)

- 1: $\text{pp}_{\text{EQ}} \leftarrow \text{EQ.Setup}(1^\lambda, \lambda)$
- 2: $\text{crs} \leftarrow \text{NIZK.Setup}(1^\lambda)$
- 3: **return** $\text{pp} := (\text{pp}_{\text{EQ}}, \text{crs})$

Key generation. The issuer's key generation algorithm samples a random integer vector \mathbf{x} and runs the SPS-EQ setup algorithm. The client's key generation algorithm samples a random value as the secret signing and extraction keys, and sets the public key with respect to this secret.

KeyGen_I(pp)

- 1: $\mathbf{x} \leftarrow (\mathbb{Z}_p^*)^2$
- 2: $\mathbf{g}_2^{\mathbf{y}}, \mathbf{y} \leftarrow \text{EQ.KeyGen}(\text{pp}_{\text{EQ}})$
- 3: **return** $\text{pk}_I := (g_1^{\mathbf{x}}, g_2^{\mathbf{y}}),$
 $\text{sk}_I := \mathbf{y}, \text{ek}_I := \mathbf{x}$

KeyGen_C(pp)

- 1: $\alpha \leftarrow \mathbb{Z}_p^*$
- 2: **return** $\text{pk}_C := g_1^\alpha,$
 $\text{sk}_C := \alpha$

ReadBit(ek_I, \mathbf{t})

- 1: **foreach** $b \in \{0, 1\}$ **do**
- 2: **if** $t_1^{\text{ek}_I^{(1+b)}} = t_2$ **then**
- 3: **return** b
- 4: **return** \perp

Presignature issuance and token generation. We show the presignature issuance and token generation protocols in detail in Figure 8. To issue a presignature, the issuer creates an equivalence class signature on $(\text{pk}_C, H(r), S)$, where group element $S \in \mathbb{G}_1$ embeds the metadata bit b . The presignature psig includes the signature $\bar{\sigma}$ and the bit embedding S , and the nonce is the uniformly chosen hash input r . We must additionally include a NIZK proof π in psig, proving $b \in \{0, 1\}$, the the knowledge of secret keys corresponding to the public keys and the correct computation of the bit embedding. In order to obtain a redeemable token, the client first verifies the NIZK proof and the equivalence class signature. It can then transform the representation of $\bar{\sigma}$ to the equivalence class signature on $(g_1, H(r)^{\alpha^{-1}}, S^{\alpha^{-1}})$. The final token is the tag $\mathbf{t} = (H(r)^{\alpha^{-1}}, S^{\alpha^{-1}})$ along with the transformed signature σ . We discuss the protocol for validation of this token next.

Public verification (token redemption). The public verification algorithm simply verifies the equivalence class signature σ .

Verify($\text{pk}_I, \mathbf{t}, \sigma$)

- 1: **return** $\text{EQ.Verify}((\text{pk}_I^{(3)}, \text{pk}_I^{(4)}, \text{pk}_I^{(5)}), (g_1, t_1, t_2), \sigma)$

Bit extraction. The bit extraction proceeds by first calling the public verification algorithm. Then, the issuer simply checks which $b \in \{0, 1\}$ satisfies $t_1^{x_1+b} = t_2$ and returns that bit if one is found, otherwise it returns an error.

Correctness. Correctness of Issue follows from correctness of the SPS-EQ scheme EQ, and that of the NIZK scheme NIZK. Thus, $\bar{\sigma}$ is a valid signature on $(\text{pk}_C, H(r), S)$, and π is a valid NIZK for $b \in \{0, 1\}$ and the well-formedness of psig . Furthermore, correctness of EQ also guarantees the correctness of Obtain. Consequently, σ is a valid signature on $(g_1, H(r)^{\alpha^{-1}}, S^{\alpha^{-1}})$. Lastly, we have that

$$t_1^{x_1+b} = H(r)^{\frac{x_1+b}{\alpha}} = S^{\frac{1}{\alpha}} = t_2$$

where $\alpha =: \text{sk}_C$. Therefore, NIAT Construction 4 satisfies correctness.

Reusability. Let t_i be the tag created by the client in session i for $i \in \{0, 1\}$, and similarly let $(\bar{\sigma}_i, S_i, \pi_i) =: \text{psig}_i$ and $r_i =: \text{nonce}_i$ be the corresponding presignature and nonce. Then,

$$\begin{aligned} & \Pr[r_0 = r_1 \vee t_0 = t_1] \\ &= \Pr[r_0 = r_1] + \Pr[t_0 = t_1] \\ & \quad - \Pr[t_0 = t_1 \mid r_0 = r_1] \cdot \Pr[r_0 = r_1] \end{aligned}$$

Since, according to the definition of reusability, the metadata b is the same for both issuances, this probability is

$$\begin{aligned} &= \Pr[r_0 = r_1] + \Pr[t_0 = t_1] - 1 \cdot \Pr[r_0 = r_1] \\ &= \Pr[H(r_0)^{\alpha^{-1}} = H(r_1)^{\alpha^{-1}} \wedge S_0^{\alpha^{-1}} = S_1^{\alpha^{-1}}] \\ &= \Pr[H(r_0) = H(r_1)] \quad . \end{aligned}$$

Since each r_i is chosen uniformly from $\{0, 1\}^\lambda$, we have by collision resistance of H that the above probability is negligible. So, Construction 4 satisfies reusability.

Security. Below, we present our main theorems for NIAT security of Construction 4. Due to space constraints, we provide the proofs in Appendix C.

Theorem 4.1 (One-more unforgeability). *Construction 4 is one-more unforgeable (in the random oracle model) assuming NIZK satisfies zero knowledge and EQ is existentially unforgeable under adaptively chosen-message attacks.*

Remark 4.2 (Knowledge of secret keys). For the proof of κ -unlinkability of our construction, we will leverage the knowledge of secret keys (KOSK) model [23], [24] (also called the honest key model [15]). At a high level, this model requires the attacker to also specify the secret key sk_I corresponding to the public verification key pk_I in the κ -unlinkability experiment (Figure 4). A NIAT scheme that is secure in the KOSK can be transformed into a fully maliciously secure scheme simply by having the issuer provide a NIZK argument of knowledge (NIZKAoK) of the secret key sk_I as part of pk_I . Intuitively, in the security proof, the reduction can extract the secret key from the adversary's argument of knowledge, and then run the reduction for κ -unlinkability in the KOSK model.

Theorem 4.3 (κ -unlinkability). *Construction 4 is κ -unlinkable (in the random oracle and KOSK models) for $\kappa = 2$, assuming NIZK is an argument of knowledge, that inverse DDH assumption holds in \mathbb{G}_1 and EQ perfectly adapts signatures under a malicious signer.*

Theorem 4.4 (Privacy of metadata bit). *Construction 4 has private metadata bit (in the random oracle model) assuming NIZK satisfies zero-knowledge, that DDH assumption holds and EQ is existentially unforgeable under adaptively chosen-message attacks.*

4.2. Instantiation

We use the SPS-EQ scheme given by [19] to instantiate our protocol from Figure 8. We provide the expanded token issuance and generation protocol in Figure 14 of the Appendix.

Issuer's ZK proof. An essential part of the issuing protocol is the disjunctive OR-proof for consistency of the secret values x_1 and x_2 used in $S = R^{(1-b)x_1+bx_2}$ against the public key pk_I . Specifically, the OR-proof proves that $S = R^{x_1}$ (when $b = 0$), or R^{x_2} (when $b = 1$). Disjunctive proofs of two statements can be constructed by proving knowledge for one statement and simulating knowledge for the other, without revealing which of the two is actually proved and which is simulated [25]. We detail the full ZK protocol for the language \mathcal{L}_{iss} in Appendix D.

A note on batch verification. The main computational bottleneck in our protocol is due to the pairing operations required for signature verification during, both, token generation and redemption (verification). However, we observe that this cost can actually be amortized over the number of presignatures issued (resp., tokens redeemed) as some of the pairing operations are jointly verifiable. We will take advantage of this fact when we discuss our concrete implementation in Section 6. The precise verification check is explained with the concrete instantiation in Appendix D.

5. Double-spend Protection

While one-more unforgeability prevents a client from creating more than one tokens from a single presignature issuance, a NIAT system by itself does not preclude a client from attempting to redeem the same token twice. More importantly, even if an issuer were able to detect such a double-spending attempt by keeping track of all redeemed tokens, unlinkability seems to suggest that it may be hard to identify the offending client. However, we note that this is not the case—unlinkability with respect to a *fixed* client is defined for distinct nonces, and it turns out that a NIAT scheme that allows the public identification of a client attempting to redeem the same token more than once is perfectly within the bounds of Definition 3.4.




Figure 9. Modified SPS-EQ construction for NIAT presignature issuance and token generation.

5.1. Construction

We give a modification of Construction 4 that facilitates protection against double-spending by allowing a verifier to identify the offending client's public key. We utilize the fact that the presignature already embeds the client's public key pk_C (a unique property of our NIAT construction compared to interactive constructions). Due to space constraints, we omit algorithms which are identical to (or trivially extendable from) their counterparts in Construction 4.

Tools required. Our modified construction requires hash functions $H_{\mathbb{G}} : \{0, 1\}^\lambda \rightarrow \mathbb{G}_1$, $H_{\mathbb{Z}} : \{0, 1\}^\lambda \rightarrow \mathbb{Z}_p$, $H_{\mathbb{G} \rightarrow \mathbb{G}} : \mathbb{G}_1 \rightarrow \mathbb{G}_1$, an SPS-EQ scheme $\text{EQ} = (\text{EQ.Setup}, \text{EQ.Sign}, \text{EQ.ChRep}, \text{EQ.Verify}, \text{EQ.VerKey})$, $\text{NIZK}_{\text{iss}} = (\text{NIZK}_{\text{iss}}.\text{Setup}, \text{NIZK}_{\text{iss}}.\text{Prove}, \text{NIZK}_{\text{iss}}.\text{Verify})$ and $\text{NIZK}_{\text{obt}} = (\text{NIZK}_{\text{obt}}.\text{Setup}, \text{NIZK}_{\text{obt}}.\text{Prove}, \text{NIZK}_{\text{obt}}.\text{Verify})$ where the former is a NIZK for \mathcal{L}_{iss} defined previously, and the latter is a NIZK for the following language:

Language \mathcal{L}_{obt}
Instance: Each instance x is interpreted as group elements t_1, ϕ_1 and ϕ_2 and an integer id . Witness: Witness w consists of integers sk_C and \hat{r} . Membership: w is a valid witness for x if the following are satisfied:
$\phi_2 = g_1^{\text{id} \cdot (\text{sk}_C + \hat{r})} \cdot \phi_1^{\text{sk}_C + \hat{r}} \wedge t_1 = g_1^{(\text{sk}_C + \hat{r})^{-1}}$

Presignature issuance and token generation. In Figure 9, we show the presignature issuance and token generation protocols in detail, and highlight the differences from the

previous construction. The main distinction in issuance is the inclusion of the integer \hat{r} computed as $H_{\mathbb{Z}}(r)$ and the signature, which is now computed over $\mathbf{m} := (g_1, \text{pk}_C, g_1^{\hat{r}}, R, S)$. As we will see, this is the key to our double-spend identification mechanism.

To obtain a token, a client performs the same steps as before, except that instead of $\mathbf{m}^{\alpha^{-1}}$, it now computes the transformed signature over $\mathbf{m}^{(\alpha + \hat{r})^{-1}}$ with the corresponding tag $\mathbf{t} := (g_1^{(\alpha + \hat{r})^{-1}}, R^{(\alpha + \hat{r})^{-1}}, S^{(\alpha + \hat{r})^{-1}})$.

Now, if a client intends to redeem its token with some verifier with public identifier id , as we show in Figure 9, the client's Obtain algorithm must additionally compute an ElGamal ciphertext ϕ over id along with a proof π_o of the well-formedness of the ciphertext (see language \mathcal{L}_{obt}). Therefore, the final signature is the equivalence class signature s , the ciphertext ϕ and the proof π_o .

Public verification. The public verification algorithm verifies the equivalence class signature s and the ZK proof π_o .

$\text{DSIden}(\text{pk}_I, \mathbf{t}, \sigma, \sigma', \text{aux})$

- 1 : if $\text{Verify}(\text{pk}_I, \mathbf{t}, \sigma) \neq 1 \vee \text{Verify}(\text{pk}_I, \mathbf{t}, \sigma') \neq 1$
 $\vee \sigma = \sigma'$ then return (\perp, \perp)
- 2 : $\sigma := (s, \phi, \pi_o, \text{id}), \sigma' := (s', \phi', \pi_o', \text{id}')$
- 3 : $h := (\phi_2^{-1} \cdot \phi_2')^{(\text{id}' - \text{id})^{-1}}$
- 4 : **foreach** $(\text{pk}_C, \text{nonce}) \in \text{aux}$ **do**
- 5 : if $h = \text{pk}_C \cdot g_1^{H_{\mathbb{Z}}(\text{nonce})}$ **then**
- 6 : **return** $\text{pk}_C, \Pi := (\mathbf{t}, \sigma, \sigma', \text{nonce})$
- 7 : **return** (\perp, \perp)

Double-spend identification. Suppose that a client attempts to double-spend a token (i.e., one created under the same presignature and nonce pair) with another verifier with public identifier id' . In order to do so, it must create a fresh encryption ϕ' over id' and the corresponding proof π'_o with the same s and t . When a system detects that a double-spend of (t, s) has occurred, it can compute $(\phi_2^{-1} \cdot \phi'_2)^{(\text{id}' - \text{id})^{-1}} = g_1^{\alpha + \hat{r}}$ and then obtain the offending client's public key by checking against the available nonces. Note that we have assumed distinct id 's for ease of exposition, however, this is just as easily accomplished for the same verifier by hashing a timestamp into the identifier.

Notably, this step can be performed publicly by anyone and leads to a strong incentive against double-spending⁶. A simple optimization that offers significant asymptotic advantage over the linear search on aux is to instead store the value $\text{pk}_C \cdot g_1^{H_Z(\text{nonce})}$ value itself, which allows the verifier running DSIden to perform efficient lookup over aux .

Double-spend verification. The double-spend verification algorithm must essentially check the identification algorithm's work. It ensures that the accused pk_C and the corresponding nonce are a valid pair in the auxiliary data, and confirms that the DS identification equation holds.

$\text{DSVer}(\text{pk}_C, \Pi, \text{aux})$

```

1: if Verify( $\text{pk}_I, t, \sigma$ )  $\wedge$  Verify( $\text{pk}_I, t, \sigma'$ )  $\neq 1$  then
2:   return 0
3:  $\sigma := (s, \phi, \pi_o, \text{id}), \sigma' := (s', \phi', \pi'_o, \text{id}')$ 
4: return  $(\phi_2^{-1} \cdot \phi'_2)^{(\text{id} - \text{id}')^{-1}} = \text{pk}_C \cdot g_1^{H_Z(\text{nonce})}$ 
    $\wedge (\text{pk}_C, \text{nonce}) \in \text{aux}$ 

```

Security. Below, we present our main theorems for NIAT security of Construction 5. Due to space constraints, we provide the proofs in Appendix E.

Theorem 5.1 (One-more unforgeability). *Construction 5 is one-more unforgeable (in the random oracle model) assuming NIZK_{iss} satisfies zero knowledge, EQ is existentially unforgeable under adaptively chosen-message attacks and NIZK_{obt} is sound.*

Theorem 5.2 (κ -unlinkability). *Construction 5 is κ -unlinkable (in the random oracle and KOSK models) for $\kappa = 2$, assuming NIZK_{iss} is an argument of knowledge, that DDH, k -DDH assumption holds in \mathbb{G}_1 and EQ perfectly adapts signatures under a malicious signer NIZK_{obt} satisfies zero knowledge.*

Theorem 5.3 (Privacy of metadata bit). *Construction 5 has private metadata bit (in the random oracle model) assuming NIZK satisfies zero-knowledge, that DDH assumption holds and EQ is existentially unforgeable under adaptively chosen-message attacks.*

6. One could potentially alter the protocol to instead reveal the sk_C , but we consider this to be an extremely punitive measure given that DSIden is a public algorithm.

Theorem 5.4 (Double-spend identification). *Construction 5 satisfies double-spend identification (in the random oracle model) assuming NIZK_{obt} is sound.*

Theorem 5.5 (Double-spend exculpability). *Construction 5 satisfies double-spend exculpability assuming NIZK_{obt} is sound and the discrete log assumption holds in \mathbb{G}_1 .*

6. Implementation and Evaluation

We implemented our main NIAT Construction 4 in C++ using the `mcl` library [26]⁷. For our implementation, we chose the pairing-friendly BLS12-381 curve [27], [28]. All our experiments were done on a laptop equipped with an Apple M3 Pro chip and the reported numbers are the average of 1000 executions.

Storage and communication costs. An element in \mathbb{G}_1 occupies 48 bytes in its compressed representation, and similarly, an element in \mathbb{G}_2 occupies 96 bytes. The message from the issuer to the client consists of 3 elements in \mathbb{G}_1 , 1 element in \mathbb{G}_2 , 6 integers for the proof, and a nonce $r \in \{0, 1\}^\lambda$; and a token consists of 4 elements in \mathbb{G}_1 and 1 element in \mathbb{G}_2 , we followed the `minSig` approach (signatures in \mathbb{G}_1 and signing verification keys in \mathbb{G}_2) to minimize the token size. With this, we estimate that in terms of communication costs, a presignature issuance (to the client) needs around 464 bytes to be transferred over the wire, and the final token size is around 288 bytes. Therefore, in order to design a system that incorporates the offline signing protocol shown in Figure 1 where the server supports m clients, and each client will be issued n presignatures in a batch, the required total server storage would be calculated as $m \cdot (48 + 464n)$ bytes. We provide an estimation of storage needs for up to 10^8 clients in Figure 10 (left) for batch size $n = 30, 60$ and 100 . We further note that in order to prevent double spending, the system must anyway keep track of all redeemed tokens, which will require an additional storage capacity comparable to our initial storage estimation.




Figure 10. Storage estimates for an offline signing server storing batches of 30, 60 and 100 presignatures per client (left); and Issuer's amortized redemption cost per token (right). Roughly, our protocol requires 13.6 KB to store 30 tokens per client.

Computational cost. Due to the non-interactive nature of our protocol, we are able to cut down the first message from a client to the issuer. Consequently, the issuer does

7. Anonymized source code available at: <https://anonymous.4open.science/r/NIAT-94D7>.

	# Moves	Public Verification	Private Metadata Bit	Issue	Obtain	Readbit	Token size $ t + \sigma $
Const. 4*	1	✓	✓	$16\times$	$1e + 19\times$	$3e + 1\times$	$5G$
[9]	3	✓	✓	$7\times$	$11\times$	$6\times$	$3G + 4Z$
[10]	2	✗	✓	$11\times$	$17\times$	$1\times$	$2G + 1Z$
[8]	2	✗	✓	$12\times$	$15\times$	$4\times$	$2G + 1Z$
[1]	2	✗	✗	$7\times$	$2\times$	$1\times$	$1G$
[3]	2	✓	✗	$3\times$	$2e + 1\times$	$2e$	$1G + 1Z$

TABLE 1. COMPARISON OF ANONYMOUS TOKEN SCHEMES. \times REPRESENTS GROUP EXPONENTIATION AND e REPRESENTS PAIRING OPERATION. G AND Z STAND FOR GROUP ELEMENTS AND INTEGERS RESPECTIVELY. * INDICATES THE VERIFICATION COST IS AMORTIZED.

not need to maintain an open connection with a client for token issuance. To issue a token, the issuer performs 16 exponentiations (1 for bit embedding, 6 for signature generation and 9 for the proof computation) which can be done in an *offline* manner as part of a precomputation. After receiving the issuer’s message, the client at the other end computes 5 pairing operations (plus one precomputed pairing of $e(pk_C, pk_I^{(3)})$) to verify the SPS-EQ signature, performs 14 exponentiations to verify the proof, and 5 more to obtain the final token. Finally, at redemption, the verifier performs 5 pairing operations (plus one precomputed pairing of $e(g_1, pk_I^{(3)})$) to verify the SPS-EQ signature, and then extracts the embedded bit with 1 exponentiation. We summarize our results (after amortization) in Table 1 and present them as part of a comparison with related works.

Experimental benchmarks. From our experiments we found that an exponentiation in G_1 took approximately 0.047 ms, an exponentiation in G_2 took 0.088 ms, and a pairing operation took 0.445 ms on average. We further found that (i) the issuer took 0.84 ms to issue a token, (ii) the client verified the psig in 1.73 ms and the disjunctive NIZK proof in 0.69 ms, and finalized the token in another 0.34 ms, and (iii) during bit extraction, the issuer took 1.64 ms to verify a token and additional 0.06 ms to extract the bit. It is worth noting that since our scheme is non-interactive, there is, as such, no issuance “session” and these costs may incur asynchronously. Our results are summarized in Table 2. We remark that although the σ verification step is a required check before bit extraction during token redemption, these costs are reported separately because our verification algorithm can be done publicly and independently of the bit extraction computation.

	Operation	Time (in ms)
Issuer	Issue	0.94
	Read bit (σ verification)	1.98
	Read bit (extraction)	0.06
	Obtain ($\bar{\sigma}$ verification)	2.07
Client	Obtain (proof verification)	0.83
	Obtain (finalization)	0.34

TABLE 2. BENCHMARKS FOR OUR IMPLEMENTATION OF CONSTRUCTION 4.

Batched verification. We also ran benchmark tests for a batched version of signature verification as explained in Appendix D. In this batched version the client (resp. the issuer) is able to perform $n + 4$ (resp. $2n + 2$) pairings instead of $4n$, for a batch of n presignatures (resp. tokens).

As shown in Table 1, we calculated the amortized cost of our token generation to be approximately equivalent to 1 pairing (and 17 exponentiations).

Batch size	1	30	60	100
Time (in ms)	2.07	0.28	0.25	0.24

TABLE 3. CLIENT’S AMORTIZED RUNTIMES FOR BATCH VERIFICATION.

In Table 3 we show that actual costs amortized over 30, 60 and 100 presignatures⁸ based on experimental evaluations averaged over 1000 runs. A similar calculation gives the amortized cost of our token redemption (verification) to be approximately equivalent to 2 pairings, (plus 1 exponentiation for bit extraction). In Figure 10 (right) we show the actual amortized costs over up to 2^{12} token redemptions.

Estimates for double-spending protection. Although we did not implement the NIAT with double-spend protection, we do provide the estimates of the various *additional* costs incurred by the resulting extension over the base scheme. In particular, we estimate the cost of token issuance to be about 0.15 ms over the base scheme as it requires 2 extra exponentiations (1 more for the equivalence class signature and 1 more for the proof), and the overall presignature size should require about 2 additional integers in the proof. For amortized token generation, we estimate the additional cost to be about 0.6 ms over the base scheme as it requires 8 extra exponentiations (2 more for proof verification, 1 more for the tag, 2 for the proof and 3 for the encryption), and the overall token size to be about 3 additional G_1 (1 more in the tag and 2 for the ciphertext) elements and 4 additional integers (3 for the client’s proof and 1 for the verifier’s identifier). For amortized redemption, we estimate the additional cost to be about 0.4 ms over the base scheme as it requires about 1 additional pairing and 4 exponentiations for the proof verification. Aside from the additional cost of verification, the cost of bit extraction stays the same. Finally, double-spend identification requires 5 pairings and the cost of a lookup over aux.

References

- [1] A. Davidson, I. Goldberg, N. Sullivan, G. Tankersley, and F. Val-sorda, “Privacy pass: Bypassing internet challenges anonymously,” *Proceedings on Privacy Enhancing Technologies*, vol. 2018, no. 3, pp. 164–180, Jul. 2018.

8. The number of presignatures issued to a client in Privacy Pass is between 30 and 100.

- [2] Apple, “icloud private relay overview,” https://www.apple.com/icloud/docs/iCloud_Private_Relay_Overview_Dec2021.pdf, 2021, accessed: 2024-04-29.
- [3] T. Silde and M. Strand, “Anonymous tokens with public metadata and applications to private contact tracing,” in *FC 2022: 26th International Conference on Financial Cryptography and Data Security*, ser. Lecture Notes in Computer Science, I. Eyal and J. A. Garay, Eds., vol. 13411. Grenada: Springer, Cham, Switzerland, May 2–6, 2022, pp. 179–199.
- [4] J. Wilander, “Introducing private click measurement, pcm,” <https://webkit.org/blog/11529/introducing-private-click-measurement-pcm/>, 2021, accessed: 2024-04-29.
- [5] Google, “Private state tokens,” <https://developers.google.com/privacy-sandbox/protectations/private-state-tokens>, 2024, accessed: 2024-04-29.
- [6] T. Meunier, C. D. Rubin, and A. Faz-Hernández, “Privacy pass: upgrading to the latest protocol version,” <https://blog.cloudflare.com/privacy-pass-standard>, 2024, accessed: 2024-04-29.
- [7] IETF, “Privacy pass (privacypass),” <https://datatracker.ietf.org/wg/privacypass/about/>, 2024, accessed: 2024-04-29.
- [8] B. Kreuter, T. Lepoint, M. Orrù, and M. Raykova, “Anonymous tokens with private metadata bit,” in *Advances in Cryptology – CRYPTO 2020, Part I*, ser. Lecture Notes in Computer Science, D. Micciancio and T. Ristenpart, Eds., vol. 12170. Santa Barbara, CA, USA: Springer, Cham, Switzerland, Aug. 17–21, 2020, pp. 308–336.
- [9] F. Benhamouda, T. Lepoint, M. Orrù, and M. Raykova, “Publicly verifiable anonymous tokens with private metadata bit,” Cryptology ePrint Archive, Report 2022/004, 2022. [Online]. Available: <https://eprint.iacr.org/2022/004>
- [10] M. Chase, F. B. Durak, and S. Vaudenay, “Anonymous tokens with stronger metadata bit hiding from algebraic MACs,” in *Advances in Cryptology – CRYPTO 2023, Part II*, ser. Lecture Notes in Computer Science, H. Handschuh and A. Lysyanskaya, Eds., vol. 14082. Santa Barbara, CA, USA: Springer, Cham, Switzerland, Aug. 20–24, 2023, pp. 418–449.
- [11] M. Orrù, S. Tessaro, G. Zaverucha, and C. Zhu, “Oblivious issuance of proofs,” in *Advances in Cryptology – CRYPTO 2024, Part IX*, ser. Lecture Notes in Computer Science, L. Reyzin and D. Stebila, Eds., vol. 14928. Santa Barbara, CA, USA: Springer, Cham, Switzerland, Aug. 18–22, 2024, pp. 254–287.
- [12] D. Chaum, “Blind signature system,” in *Advances in Cryptology – CRYPTO’83*, D. Chaum, Ed. Santa Barbara, CA, USA: Plenum Press, New York, USA, 1983, p. 153.
- [13] D. Pointcheval and J. Stern, “Provably secure blind signature schemes,” in *Advances in Cryptology – ASIACRYPT’96*, ser. Lecture Notes in Computer Science, K. Kim and T. Matsumoto, Eds., vol. 1163. Kyongju, Korea: Springer Berlin Heidelberg, Germany, Nov. 3–7, 1996, pp. 252–265.
- [14] —, “Security arguments for digital signatures and blind signatures,” *Journal of Cryptology*, vol. 13, no. 3, pp. 361–396, Jun. 2000.
- [15] L. Hanzlik, “Non-interactive blind signatures for random messages,” in *Advances in Cryptology – EUROCRYPT 2023, Part V*, ser. Lecture Notes in Computer Science, C. Hazay and M. Stam, Eds., vol. 14008. Lyon, France: Springer, Cham, Switzerland, Apr. 23–27, 2023, pp. 722–752.
- [16] F. Baldimtsi, J. Cheng, R. Goyal, and A. Yadav, “Non-interactive blind signatures: Post-quantum and stronger security,” in *Advances in Cryptology – ASIACRYPT 2024, Part II*, ser. Lecture Notes in Computer Science, K.-M. Chung and Y. Sasaki, Eds., vol. 15485. Kolkata, India: Springer, Singapore, Singapore, Dec. 9–13, 2024, pp. 70–104.
- [17] L. Hanzlik, E. Paracucchi, and R. Zanotto, “Non-interactive blind signatures from RSA assumption and more,” in *Advances in Cryptology – EUROCRYPT 2025 – 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4–8, 2025, Proceedings, Part II*, ser. Lecture Notes in Computer Science, S. Fehr and P. Fouque, Eds., vol. 15602. Springer, 2025, pp. 365–394. [Online]. Available: https://doi.org/10.1007/978-3-031-91124-8_13
- [18] G. Amjad, K. Yeo, and M. Yung, “RSA blind signatures with public metadata,” Cryptology ePrint Archive, Report 2023/1199, 2023. [Online]. Available: <https://eprint.iacr.org/2023/1199>
- [19] G. Fuchsbaauer, C. Hanser, and D. Slamanig, “Structure-preserving signatures on equivalence classes and constant-size anonymous credentials,” *Journal of Cryptology*, vol. 32, no. 2, pp. 498–546, Apr. 2019.
- [20] N. Tyagi, S. Celi, T. Ristenpart, N. Sullivan, S. Tessaro, and C. A. Wood, “A fast and simple partially oblivious PRF, with applications,” in *Advances in Cryptology – EUROCRYPT 2022, Part II*, ser. Lecture Notes in Computer Science, O. Dunkelman and S. Dziembowski, Eds., vol. 13276. Trondheim, Norway: Springer, Cham, Switzerland, May 30 – Jun. 3, 2022, pp. 674–705.
- [21] F. Benhamouda, M. Raykova, and K. Seth, “Anonymous counting tokens,” in *Advances in Cryptology – ASIACRYPT 2023, Part II*, ser. Lecture Notes in Computer Science, J. Guo and R. Steinfeld, Eds., vol. 14439. Guangzhou, China: Springer, Singapore, Singapore, Dec. 4–8, 2023, pp. 245–278.
- [22] S. D. Galbraith, K. G. Paterson, and N. P. Smart, “Pairings for cryptographers,” Cryptology ePrint Archive, Report 2006/165, 2006. [Online]. Available: <https://eprint.iacr.org/2006/165>
- [23] S. Micali, K. Ohta, and L. Reyzin, “Accountable-subgroup multisignatures: Extended abstract,” in *ACM CCS 2001: 8th Conference on Computer and Communications Security*, M. K. Reiter and P. Samarati, Eds. Philadelphia, PA, USA: ACM Press, Nov. 5–8, 2001, pp. 245–254.
- [24] A. Boldyreva, “Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme,” in *PKC 2003: 6th International Workshop on Theory and Practice in Public Key Cryptography*, ser. Lecture Notes in Computer Science, Y. Desmedt, Ed., vol. 2567. Miami, FL, USA: Springer Berlin Heidelberg, Germany, Jan. 6–8, 2003, pp. 31–46.
- [25] I. Damgård, “On Σ -protocols,” <https://www.cs.au.dk/~ivan/Sigma.pdf>, 2010, accessed: 2025-03-20.
- [26] S. Mitsunari, “MCL,” <https://github.com/herumi/mcl>, 2024.
- [27] P. S. L. M. Barreto, B. Lynn, and M. Scott, “Constructing elliptic curves with prescribed embedding degrees,” in *SCN 02: 3rd International Conference on Security in Communication Networks*, ser. Lecture Notes in Computer Science, S. Cimato, C. Galdi, and G. Persiano, Eds., vol. 2576. Amalfi, Italy: Springer Berlin Heidelberg, Germany, Sep. 12–13, 2003, pp. 257–267.
- [28] S. Bowe, “Bls12-381: New zk-snark elliptic curve construction,” <https://electriccoin.co/blog/new-snark-curve/>, 2017, accessed: 2024-04-17.
- [29] A. Fiat and A. Shamir, “How to prove yourself: Practical solutions to identification and signature problems,” in *Advances in Cryptology – CRYPTO’86*, ser. Lecture Notes in Computer Science, A. M. Odlyzko, Ed., vol. 263. Santa Barbara, CA, USA: Springer Berlin Heidelberg, Germany, Aug. 1987, pp. 186–194.
- [30] L. Hanzlik and D. Slamanig, “With a little help from my friends: Constructing practical anonymous credentials,” in *ACM CCS 2021: 28th Conference on Computer and Communications Security*, G. Vigna and E. Shi, Eds. Virtual Event, Republic of Korea: ACM Press, Nov. 15–19, 2021, pp. 2004–2023.
- [31] Y. Dodis and A. Yampolskiy, “A verifiable random function with short proofs and keys,” in *PKC 2005: 8th International Workshop on Theory and Practice in Public Key Cryptography*, ser. Lecture Notes in Computer Science, S. Vaudenay, Ed., vol. 3386. Les Diablerets, Switzerland: Springer Berlin Heidelberg, Germany, Jan. 23–26, 2005, pp. 416–431.

- [32] D. Boneh, H. W. Montgomery, and A. Raghunathan, “Algebraic pseudorandom functions with improved efficiency from the augmented cascade,” in *ACM CCS 2010: 17th Conference on Computer and Communications Security*, E. Al-Shaer, A. D. Keromytis, and V. Shmatikov, Eds. Chicago, Illinois, USA: ACM Press, Oct. 4–8, 2010, pp. 131–140.

Appendix

1. More Preliminaries

In this section we provide some additional preliminaries essential to this work.

1.1. Zero-knowledge Proofs. We recall the standard definition of a Zero-knowledge (ZK) proof as an interactive protocol between a Prover \mathcal{P} and a Verifier \mathcal{V} . The Prover \mathcal{P} must convince the verifier \mathcal{V} that she knows a private witness w for a public instance $x \in \mathcal{L}$ such that $\mathcal{R}(x, w) = 1$, and \mathcal{V} gains no additional information. A ZK proof system consists of the following polynomial time algorithms:

- $\text{Setup}(1^\lambda) \rightarrow \text{crs}$. The setup algorithm takes as input the security parameter λ , and outputs a crs.
- $\text{Prove}(\text{crs}, x, w) \rightarrow \pi$. The prover algorithm takes as input a crs, an instance $x \in \mathcal{L}$, and a witness w . It outputs a proof π .
- $\text{Verify}(\text{crs}, x, \pi) \rightarrow \{0, 1\}$. The verification algorithm takes as input a crs, an instance x , and a proof π . It outputs 0 or 1.

Definition A.1 (Zero-knowledge proof). A *zero-knowledge proof* between \mathcal{P} and \mathcal{V} for an NP relation \mathcal{R} must satisfy the following properties:

- **Completeness:** If $\mathcal{R}(x, w) = 1$ and both players are honest, \mathcal{V} always accepts.
- **Soundness:** For every malicious and computationally unbounded \mathcal{P}^* , there is a negligible function $\epsilon(\cdot)$ such that if x is a false statement (i.e., $\forall w : \mathcal{R}(x, w) = 0$), after \mathcal{P}^* interacts with \mathcal{V} , $\Pr[\mathcal{V} \text{ accepts}] \leq \epsilon(|x|)$.

Moreover, a proof system is an *argument of knowledge* (AoK) if there exists a PPT extractor \mathcal{E} such that for every stateful PPT attacker \mathcal{A} , the following probability is negligible

$$\Pr \left[\begin{array}{l} \text{Verify}(\text{crs}, x, \pi) = 1 \\ \wedge \mathcal{R}(x, \mathcal{E}(\text{crs}, x, \pi)) = 0 \end{array} : (x, \pi) \leftarrow \mathcal{A}(1^\lambda, \text{crs}) \right]$$

- **Zero-knowledge:** For every malicious PPT adversary \mathcal{V}^* , there exists a PPT simulator \mathcal{S} and a negligible function $\epsilon(\cdot)$ such that for every distinguisher \mathcal{D} and $(x, w) \in R$, the following is negligible in $|x|$

$$|\Pr[\mathcal{D}(\text{View}_{\mathcal{V}^*}(x, w)) = 1] - \Pr[\mathcal{D}(\mathcal{S}) = 1]|$$

Composed statements. ZK proofs can be composed as follows: (1) AND composition $\pi_1 \wedge \pi_2$ can be constructed by sequential or parallel proving of underlying assertions, and (2) OR composition $\pi_1 \vee \pi_2$ can be constructed by proving knowledge for one statement and simulating

knowledge for the other, without revealing which of the two is actually proved and which is simulated [25].

Mixed statements. Let f and g be non-algebraic and algebraic relations with public instances y and z respectively. A ZK proof on a mixed statement has the generic form $\{(w) : f(y, w) = 1 \wedge g(z, w) = 1\}(y, z)$.

Non-interactive zero-knowledge. Public-coin interactive ZK proofs can be made non-interactive in the random oracle model using the Fiat-Shamir heuristic [29].

2. Anonymous Tokens with Private Metadata Bit

We recall the anonymous tokens (AT) interface defined by Chase et al. [10]. Formally, it is comprised of three polynomial time algorithms Setup , KeyGen_I and ReadBit , and a probabilistic polynomial time (PPT) interactive token issuance protocol between Client and Issuer.

- $\text{Setup}(1^\lambda) \rightarrow \text{pp}$. The setup algorithm takes as input the security parameter λ and outputs the protocol’s public parameters pp .
- $\text{KeyGen}_I(\text{pp}) \rightarrow (\text{pk}_I, \text{sk}_I)$. The issuer’s key generation algorithm takes the public parameters pp as input and generates the public key pk_I and secret key sk_I of the issuer.
- $\langle \text{Client}(\text{pk}_I, \tau), \text{Issuer}(\text{sk}_I, b, \tau) \rangle \rightarrow (t, \sigma) \text{ or } \perp$. This is the interactive token issuance protocol between Client and Issuer where the client’s input is the issuer’s public key pk_I and some public metadata $\tau \in \mathcal{M}$, and the issuer’s input is its secret key sk_I , the hidden metadata bit $b \in \{0, 1\}$ and the public metadata τ . It can be broken down into three algorithms.
 - $\text{ClientQuery}(\text{pk}_I, \tau) \rightarrow (\text{query}, \text{state})$. Client initiates the protocol by sending query to the Issuer.
 - $\text{IssueToken}(\text{sk}_I, b, \tau, \text{query}) \rightarrow \text{resp}$. Issuer replies with resp.
 - $\text{ClientObtain}(\text{state}, \text{resp}) \rightarrow (t, \sigma)$. Client locally computes the token from the Issuer’s response.
- $\text{ReadBit}(\text{sk}_I, t, \sigma, \tau) \rightarrow \{0, 1\} \text{ or } \perp$. The Issuer’s ReadBit algorithm takes as input the issuer’s secret key sk_I , a token (t, σ) , the public metadata τ and outputs $b \in \{0, 1\}$ or \perp .

Correctness of an AT scheme holds if for any given execution of the interactive protocol, the output of the Issuer’s ReadBit algorithm is the same as the embedded metadata bit b . In terms of security, an AT scheme must satisfy *one-more unforgeability*, *unlinkability*, and *privacy of the metadata bit*. One-more unforgeability ensures that for an issuer running ℓ times, on some fixed (b, m) pair, an adversary should not be able to produce more than ℓ tokens with pairwise distinct tags. Unlinkability ensures that a malicious issuer can not link tokens to any given client,

and more generally to the corresponding issuing sessions. Finally, privacy of the metadata bit ensures that an adversary can do no better than guess the hidden metadata bit with trivial probability. We refer the reader to [10] for the formal security definitions.

2.1. Tag-based Equivalence Class Signatures. We now recall the definition of Tag-based Equivalence Class Signatures (TB-EQS) [30] that modifies the SPS-EQ definition additionally support the inclusion of an auxiliary tag τ such that the tag for a signature on a given message does not change under changing representations.

Definition A.2 (TB-EQS). A TB-EQS scheme consists of the following PPT algorithms:

- $\text{Setup}(1^\lambda, \ell) \rightarrow \text{pp}_{\text{TEQ}}$. On input the security parameter 1^λ , and the length of message vectors ℓ , it outputs the public parameters pp_{TEQ} .
- $\text{KeyGen}(\text{pp}_{\text{TEQ}}) \rightarrow (\text{pk}_{\text{TEQ}}, \text{sk}_{\text{TEQ}})$. On input the public parameters pp_{TEQ} , it outputs a key pair $(\text{pk}_{\text{TEQ}}, \text{sk}_{\text{TEQ}})$.
- $\text{VerKey}(\text{sk}_{\text{TEQ}}, \text{pk}_{\text{TEQ}}) \rightarrow b \in \{0, 1\}$. On input a public-secret key pair $(\text{pk}_{\text{TEQ}}, \text{sk}_{\text{TEQ}})$, it deterministically returns a bit $b \in \{0, 1\}$.
- $\text{Sign}(\text{sk}_{\text{TEQ}}, M, \tau) \rightarrow \sigma_{\text{TEQ}}$. On input the secret key sk_{TEQ} , a representative $M \in (\mathbb{G}_i^*)^\ell$ and a tag $\tau \in \{0, 1\}^*$, it outputs a signature σ_{TEQ} for the equivalence class $[M]$.
- $\text{ChRep}(M, \sigma_{\text{TEQ}}, \mu) \rightarrow \sigma'_{\text{TEQ}}$. On input a message M , signature σ_{TEQ} and a scalar μ , it returns an updated message-signature pair $(M', \sigma'_{\text{TEQ}})$ where the new representative is $M' = M^\mu$ and σ'_{TEQ} is the corresponding updated signature.
- $\text{Verify}(\text{pk}_{\text{TEQ}}, M, \tau, \sigma_{\text{TEQ}}) \rightarrow b \in \{0, 1\}$. On input a public key pk_{TEQ} , a representative $M \in (\mathbb{G}_i^*)^\ell$, a tag $\tau \in \{0, 1\}^*$, and a signature σ_{TEQ} , it deterministically outputs a bit $b \in \{0, 1\}$.

Definition A.3 (Existential unforgeability under chosen message attack). Let $\text{Adv}_{\mathcal{A}, \ell}^{\text{EUnf-CMA}} = \Pr[\text{EUnf-CMA}_{\mathcal{A}, \ell}(\lambda) = \text{accept}]$ be the advantage of an PPT adversary \mathcal{A} in the EUnf-CMA experiment defined in Figure 11. A TB-EQS scheme over $(\mathbb{G}_i^*)^\ell$ is existentially unforgeable under adaptive chosen-message attacks if for all $\ell > 1$, the advantage $\text{Adv}_{\mathcal{A}}^{\text{EUnf-CMA}} = \text{negl}(\lambda)$.

Definition A.4 (Perfect signature adaptation under malicious keys). For $\ell > 1$, an TB-EQS scheme on $(\mathbb{G}_i^*)^\ell$ perfectly adapts signatures under malicious keys if for all tuples $(\text{pk}_{\text{TEQ}}, M, \sigma_{\text{TEQ}}, \mu, \tau)$ satisfying

$$M \in \mathbb{G}_i^* \wedge \text{TEQ.Verify}(\text{pk}_{\text{TEQ}}, M, \tau, \sigma_{\text{TEQ}}) = 1 \wedge \mu \in \mathbb{Z}_p^*$$

we have that the output of $\text{TEQ.ChRep}(\sigma_{\text{TEQ}}, \mu)$ is a uniformly random element in the space of signatures, conditioned on $\text{TEQ.Verify}(\text{pk}_{\text{TEQ}}, M^\mu, \tau, \sigma'_{\text{TEQ}}) = 1$.

Game $\text{EUnf-CMA}_{\mathcal{A}, \ell}(\lambda)$

```

1 :  $Q := \emptyset$ 
2 :  $(\text{pp}_{\text{TEQ}}) \leftarrow \text{TEQ.Setup}(1^\lambda, \ell)$ 
3 :  $(\text{pk}_{\text{TEQ}}, \text{sk}_{\text{TEQ}}) \leftarrow \text{TEQ.KeyGen}(\text{pp}_{\text{TEQ}})$ 
4 :  $(M^*, \sigma_{\text{TEQ}}^*, \tau^*) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Sign}}}(\text{pk}_{\text{TEQ}})$ 
5 : return  $([M^*], \tau^*) \neq ([M], \tau) \forall (M, \tau) \in Q \wedge$ 
       $\text{TEQ.Verify}(\text{pk}_{\text{TEQ}}, M^*, \tau^*, \sigma_{\text{TEQ}}^*) = 1$ 

```

Oracle $\mathcal{O}_{\text{Sign}}(M, \tau)$

```

1 :  $\sigma_{\text{TEQ}} \leftarrow \text{TEQ.Sign}(\text{sk}_{\text{TEQ}}, M, \tau)$ 
2 :  $Q := Q \cup \{(M, \tau)\}$ 
3 : return  $\sigma_{\text{TEQ}}$ 

```

Figure 11. The existential unforgeability (under chosen message attack) experiment for TB-EQS.

3. Security Proofs for Section 4

THEOREM 4.1 (ONE-MORE UNFORGEABILITY). Construction 4 is one-more unforgeable (in the random oracle model) assuming NIZK satisfies zero knowledge and EQ is existentially unforgeable under adaptively chosen-message attacks.

Pf. We proceed through a series of hybrids.

- Hybrid \mathbf{H}_0 : This is the original NIAT one-more unforgeability experiment $\text{OM-Unf}_{\mathcal{A}}$ (see Definition 3.3) defined with respect to Construction 4.
- Hybrid \mathbf{H}_1 : This is the same as \mathbf{H}_0 except, when there is a collision on an adversaries random oracle query to H , the challenger aborts. For a PPT adversary making $q_H = \text{poly}(\lambda)$ random oracle queries, it can be shown by a simple application of the union bound that the adversary's advantage in \mathbf{H}_0 and that in \mathbf{H}_1 differs by at most $q_H/2^\lambda$.
- Hybrid \mathbf{H}_2 : This is the same as \mathbf{H}_1 except, the challenger simulates the proof π without a witness. If NIZK satisfies zero knowledge, then \mathbf{H}_1 and \mathbf{H}_2 are indistinguishable to the adversary.

Let $\text{Adv}_{\mathcal{A}}^j$ be the advantage of a PPT adversary \mathcal{A} in hybrid \mathbf{H}_j . Then, we claim that $\text{Adv}_{\mathcal{A}}^2$ must be negligible. In order to do so, we give a reduction \mathcal{B} such that if \mathcal{A} wins the game in \mathbf{H}_2 , \mathcal{B} can break the existential unforgeability (under adaptively chosen-message attack) of the underlying SPS-EQ scheme EQ. In particular, the reduction does the following:

- It sets $(\text{pk}_I^{(3)}, \text{pk}_I^{(4)}, \text{pk}_I^{(4)}) := \text{pk}_{\text{EQ}}$ given by the EUnf-CMA challenger \mathcal{C} . It additionally samples x_1 and x_2 from \mathbb{Z}_p and sets the rest of pk_I in the usual way.
- On receiving a Issue query (pk_C, b) , it samples r , and computes S in the usual way. It then queries \mathcal{C} for the signature $\bar{\sigma}$ on $(\text{pk}_C, H(r), S)$ and returns

$\text{psig} := (\bar{\sigma}, S, \pi)$ and $\text{nonce} := r$, where π is the simulated NIZK proof.

- On receiving a Read query (t, σ) it first checks if $\text{Verify}(\text{pk}_I, t, \sigma) = 1$ and returns \perp if it is not. Otherwise for each $b \in \{0, 1\}$ it checks whether $t_1^{x_{1+b}} = t_2$. If such a b is found, it outputs it. Otherwise it outputs \perp .
- Finally when \mathcal{A} outputs its forgery $\{(t^{(i)}, \sigma^{(i)})\}_{i \in [N]}$, \mathcal{B} uniformly chooses a pair (t^*, σ^*) and outputs it as its forgery.

For a successful adversary \mathcal{A} , each $\sigma^{(i)}$ must be a signature on a unique class represented by $(g_1, t_1^{(i)}, t_2^{(i)})$. However, since \mathcal{B} can not determine which of the N forgeries is with respect to a fresh equivalence class, it uniformly chooses one in the last step. Therefore, $\text{Adv}_{\mathcal{A}}^2 = \frac{\ell}{N-\ell} \cdot \text{Adv}_{\mathcal{B}}^{\text{EUnf-CMA}}$, which is negligible for any PPT adversary \mathcal{A} . \square

THEOREM 4.3 (κ -UNLINKABILITY). Construction 4 is κ -unlinkable (in the random oracle and KOSK models) for $\kappa = 2$, assuming NIZK is sound, that inverse DDH assumption holds in \mathbb{G}_1 and EQ perfectly adapts signatures under a malicious signer.

Pf. We proceed through a series of hybrids.

- Hybrid \mathbf{H}_0 : This is the original NIAT unlinkability experiment $\text{UNLINK}_{\mathcal{A}, n}$ (see Definition 3.4) defined with respect to Construction 4.
- Hybrid \mathbf{H}_1 : This is the same as \mathbf{H}_0 except, we program the random oracle such that for any *fresh* query $r \in \{0, 1\}^\lambda$, we set $H(r) = g_1^{\nu_r}$ for $\nu_r \leftarrow \mathbb{Z}_p$. We must also keep track of each (r, ν_r) pair and answer any repeated random oracle queries accordingly. This clearly affects no change from the adversary's point of view, so \mathbf{H}_0 and \mathbf{H}_1 are indistinguishable.
- Hybrid \mathbf{H}_2 : This is the same as \mathbf{H}_1 except, on every Obtain query $(\text{pk}_C, \text{psig} := (\bar{\sigma}, S, \pi), \text{nonce} := r)$, if $\text{EQ.Verify}((\text{pk}_I^{(3)}, \text{pk}_I^{(4)}, \text{pk}_I^{(5)}), (\text{pk}_C, H(r), S)) \neq 1$ it sets $\text{out} := (\perp, \perp)$. Otherwise, it sets t appropriately, and uses sk_I to compute $\sigma \leftarrow \text{EQ.Sign}((\text{sk}_I^{(1)}, \text{sk}_I^{(2)}, \text{sk}_I^{(3)}, (g_1, t_1, t_2)))$. Since EQ perfectly adapts signatures under malicious keys, hybrids \mathbf{H}_1 and \mathbf{H}_2 are indistinguishable to the adversary.
- Hybrid \mathbf{H}_3 : This is the same as \mathbf{H}_2 except, the challenger parses each $(\text{pk}_C^{(i)}, \text{psig}^{(i)}, \text{nonce}^{(i)})$ triple in Q^* such that for each $i \in [n]$, $\text{psig}^{(i)} = (\bar{\sigma}^{(i)}, S^{(i)}, \pi^{(i)})$. Then for any i , if $\text{EQ.Verify}((\text{pk}_I^{(3)}, \text{pk}_I^{(4)}, \text{pk}_I^{(5)}), (\text{pk}_C^{(i)}, H(r^{(i)}), S^{(i)})) \neq 1$ it sets $\text{out}_i := (\perp, \perp)$ for all i ; otherwise, it sets each $t^{(i)}$ appropriately, and uses sk_I to compute $\sigma^{(i)} \leftarrow \text{EQ.Sign}((\text{sk}_I^{(1)}, \text{sk}_I^{(2)}, \text{sk}_I^{(3)}, (g_1, t_1^{(i)}, t_2^{(i)})))$. Since EQ perfectly adapts signatures under malicious keys, hybrids \mathbf{H}_2 and \mathbf{H}_3 are indistinguishable to the adversary.

- Hybrid \mathbf{H}_4 : This is the same as \mathbf{H}_3 except, on every Obtain query $(\text{pk}_C, \text{psig} := (\bar{\sigma}, S, \pi), \text{nonce} := r)$, if $\text{EQ.Verify}((\text{pk}_I^{(3)}, \text{pk}_I^{(4)}, \text{pk}_I^{(5)}), (\text{pk}_C, H(r), S)) \neq 1$ the challenger sets $\text{out} := (\perp, \perp)$. Otherwise, it uses its knowledge of sk_I to extract b from S , and then sets $t_1 := H(r)^{\text{sk}_C^{-1}}$ as usual, and $t_2 := t_1^{(1-b)\text{ek}_I^{(1)} + b\text{ek}_I^{(2)}}$. By soundness of NIZK, hybrids \mathbf{H}_3 and \mathbf{H}_4 are indistinguishable to the adversary with all but negligible probability.
- Hybrid \mathbf{H}_5 : This is the same as \mathbf{H}_4 except, the challenger parses each $(\text{pk}_C^{(i)}, \text{psig}^{(i)}, \text{nonce}^{(i)})$ triple in Q^* such that for each $i \in [n]$, $\text{psig}^{(i)} = (\bar{\sigma}^{(i)}, S^{(i)}, \pi^{(i)})$. Then, if $\text{EQ.Verify}((\text{pk}_I^{(3)}, \text{pk}_I^{(4)}, \text{pk}_I^{(5)}), (\text{pk}_C^{(i)}, H(r^{(i)}), S^{(i)})) \neq 1$ it sets $\text{out}^{(i)} := (\perp, \perp)$ for every i ; otherwise, it uses its knowledge of sk_I to extract $b^{(i)}$ from $S^{(i)}$, and then sets $t_1^{(i)} := H(r^{(i)})^{1/\text{sk}_C^{(i)}}$ as usual, and $t_2^{(i)} := (t_1^{(i)})^{(1-b^{(i)})\text{ek}_I^{(1)} + b^{(i)}\text{ek}_I^{(2)}}$. By soundness of NIZK, hybrids \mathbf{H}_4 and \mathbf{H}_5 are indistinguishable to the adversary with all but negligible probability.
- Hybrid \mathbf{H}_6 : This is the same as \mathbf{H}_5 except, on every Obtain query $(\text{pk}_C, \text{psig} := (\bar{\sigma}, S, \pi), \text{nonce} := r)$, instead of setting t_1 as $H(r)^{\text{sk}_C^{-1}}$, it sets it to g_1^ρ for some $\rho \leftarrow \mathbb{Z}_p$ of its choice. The signature is now computed with respect to this new t_1 . We argue that \mathbf{H}_6 is indistinguishable from \mathbf{H}_5 if the inverse DDH assumption holds in \mathbb{G}_1 . In particular, for each user in Q_{usr} , the reduction algorithm \mathcal{B} instantiates a new inverse DDH challenger over \mathbb{G}_1 and receives the corresponding challenge (g_1^α, g_1^β) . It sets $\text{pk}_C := g_1^\alpha$ and $\text{sk}_C := \perp$. On any valid Obtain query, it looks up g_1^β value corresponding to the pk_C , and ν_r corresponding the nonce r . It then sets $t_1 := (g_1^\beta)^{\nu_r}$. Now observe that if $\beta = \alpha^{-1}$, the reduction simulates \mathbf{H}_5 to \mathcal{A} and otherwise simulates \mathbf{H}_6 . Thus the adversary's advantage in \mathbf{H}_5 and that in \mathbf{H}_6 differs by at most $\text{Adv}_{\mathcal{B}}^{\text{invDDH}}$.
- Hybrid \mathbf{H}_7 : This is the same as \mathbf{H}_6 except, after parsing each $(\text{pk}_C^{(i)}, \text{psig}^{(i)}, \text{nonce}^{(i)})$ triple in Q^* such that for each $i \in [n]$, $\text{psig}^{(i)} = (\bar{\sigma}^{(i)}, S^{(i)}, \pi^{(i)})$, instead of setting $t_1^{(i)}$ as $H(r^{(i)})^{\text{sk}_C^{(i)-1}}$ for each $i \in [n]$, the challenger sets it to $g_1^{\rho^{(i)}}$ for $\rho^{(i)} \leftarrow \mathbb{Z}_p$ of its choice. Each signature is now computed with respect to this new $t_1^{(i)}$. Indistinguishability between hybrids \mathbf{H}_6 and \mathbf{H}_7 follow similarly to the previous argument.

Finally, observe that the final $(t^{(i)}, \sigma^{(i)})$ pairs are all independent of their presignatures and nonces. So the best adversarial strategy is to create some n_b responses with bit b for each bit ($n_0 + n_1 = n$), read the bit b_i of out_i and output a random value i^* from the set of all n_{b_i} indices where the embedded bit was equal to b_i . The probability

that the adversary wins is

$$\sum_{b \in \{0,1\}} \Pr[b_{i^*} = b_i] \cdot \Pr[i^* = \hat{i}] = \sum_{b \in \{0,1\}} \frac{n_b}{n} \cdot \frac{1}{n_b} = \frac{2}{n}.$$

Let $\text{Adv}_{\mathcal{A}}^j$ be the advantage of a PPT adversary \mathcal{A} in hybrid \mathbf{H}_j . Then, $\text{Adv}_{\mathcal{A}}^7 \leq \frac{2}{n}$. \square

THEOREM 4.4 (PRIVACY OF METADATA BIT). Construction 4 has private metadata bit (in the random oracle model) assuming NIZK satisfies zero-knowledge, that DDH assumption holds and EQ is existentially unforgeable under adaptively chosen-message attacks.

Pf. We proceed through a series of hybrids.

- Hybrid \mathbf{H}_0 : This is the original NIAT metadata bit privacy experiment $\text{PMB}_{\mathcal{A}, \hat{b}}$ for $\hat{b} \in \{0, 1\}$ (see Definition 3.6) defined with respect to Construction 4.
- Hybrid \mathbf{H}_1 : This is the same as \mathbf{H}_0 except, we program the random oracle such that for any *fresh* query $r \in \{0, 1\}^\lambda$, we set $H(r) = g_1^{\nu_r}$ for $\nu_r \leftarrow \mathbb{Z}_p$. We must also keep track of each (r, ν_r) pair and answer any repeated random oracle queries accordingly. This clearly affects no change from the adversary's point of view, so \mathbf{H}_0 and \mathbf{H}_1 are indistinguishable.
- Hybrid \mathbf{H}_2 : This is the same as \mathbf{H}_1 except, the challenger simulates the proof π without a witness. If NIZK satisfies zero knowledge, then \mathbf{H}_1 and \mathbf{H}_2 are indistinguishable to the adversary.
- Hybrid \mathbf{H}_3 : This is the same as \mathbf{H}_2 except, on receiving the challenge query, the challenger runs the issue algorithm with $\rho \leftarrow \mathbb{Z}_p$ instead of \hat{b} . In particular, it computes $S := g_1^\rho$. We argue that \mathbf{H}_3 is indistinguishable from \mathbf{H}_2 if the DDH assumption holds in \mathbb{G}_1 . In particular, the reduction algorithm \mathcal{B} instantiates a DDH challenger over \mathbb{G}_1 and receives the corresponding challenge $(g_1^\alpha, g_1^\beta, g_1^\gamma)$. If $\hat{b} = 0$, it sets $\text{pk}_I^{(1)} := g_1^\alpha$ and $\text{ek}_I^{(1)} := \perp$ otherwise if $\hat{b} = 1$ it sets $\text{pk}_I^{(2)} := g_1^\alpha$ and $\text{ek}_I^{(2)} := \perp$. On any valid Issue query (pk_C, b) , it answers as before if $b \neq \hat{b}$. Otherwise, \mathcal{B} sets R as before, sets $S := (g_1^\alpha)^{\nu_r}$ and computes the rest of the presignature using this S . On receiving a valid Read query, the reduction first verifies the signature under pk_I , and outputs \perp if it fails. Otherwise, it checks if $t_1^{\text{ek}_I^{(2-\hat{b})}} = t_2$ (note that \mathcal{B} knows $\text{ek}_I^{(2-\hat{b})}$) and outputs $1 - \hat{b}$ if the check passes, and \hat{b} otherwise. Finally, on receiving the Challenge query for pk_C , \mathcal{B} samples the nonce r as usual, and then programs $H(r) := g_1^\beta$, sets $S := g_1^\gamma$ and performs rest of the computation as before. Now, notice that if an adversary is able to distinguish the two hybrids, then either \mathcal{B} breaks DDH or \mathcal{A} managed to forge a SPS-EQ signature σ on some tag t such that \mathcal{B} answers the Read query incorrectly on (t, σ) . This happens if the underlying “bit” is invalid, but \mathcal{B} answered with \hat{b} . It follows that the

adversary's advantage in distinguishing hybrids \mathbf{H}_2 and \mathbf{H}_3 is equal to $\text{Adv}_{\mathcal{B}}^{\text{DDH}} + \text{Adv}_{\mathcal{B}}^{\text{EUnf-CMA}}$.

Let $\text{Adv}_{\mathcal{A}, \hat{b}}^j$ be the advantage of a PPT adversary \mathcal{A} in hybrid \mathbf{H}_j with respect to \hat{b} . Then,

$$\begin{aligned} & \left| \Pr[\mathbf{H}_3^{\mathcal{A}, 1}(\lambda) = \text{accept}] - \Pr[\mathbf{H}_3^{\mathcal{A}, 0}(\lambda) = \text{accept}] \right| \\ &= |\text{Adv}_{\mathcal{A}, 1}^3 - \text{Adv}_{\mathcal{A}, 0}^3| \end{aligned}$$

which is zero. This proves the theorem. \square

4. Full Instantiation of Construction 4

We begin this section by recalling the SPS-EQ scheme from [19], which we use to instantiate our NIAT protocol. We will then give the expanded version, along with the full zero knowledge proof for the language \mathcal{L}_{iss} .

4.1. SPS-EQ Construction. A signature on a message $\mathbf{m} = (m_1, m_2, \dots, m_n) \in (\mathbb{G}_1^*)^n$ is a triple $(Z, Y_1, Y_2) \in \mathbb{G}_1 \times \mathbb{G}_1 \times \mathbb{G}_2$. The secret key of the signer is a tuple of n elements in \mathbb{Z}_p^* , and the public key is an n -tuple from \mathbb{G}_2^* . Importantly, the signature can be adapted to another signature on message \mathbf{m}^μ without the secret key of the signer. We explain the full construction in Figure 12.

4.2. The Concrete SPS-EQ Component. In Figure 14, we expand the token issuance and generation algorithms by instantiating with the aforementioned SPS-EQ of [19].

4.3. The ZK Proof. Recall that the issuer's ZK proof consists of a proof of knowledge of discrete log of the elements of pk_I corresponding to the secret key, along with the proof that $S = R^{(1-b)x_1 + bx_2}$ for x_1, x_2 in the issuer's secret key. We detail the full (interactive) ZK proof protocol for the language \mathcal{L}_{iss} in Figure 13. As previously mentioned, we can apply the Fiat-Shamir heuristic [29] to transform the proof into a NIZK for our construction.

- **Prove.** The NIZK prover algorithm takes as input a common reference string crs , the statement $x = (S, T_0, T_1, U, V, W)$, where $T_0 = \text{pk}_I^{(1)} = g_1^{x_1}$, $T_1 = \text{pk}_I^{(2)} = g_1^{x_2}$, $U = \text{pk}_I^{(3)} = g_2^{y_1}$, $V = \text{pk}_I^{(4)} = g_2^{y_2}$ and $W = \text{pk}_I^{(5)} = g_2^{y_3}$, and the witness $w = (\text{sk}_I := (x_1, x_2, y_1, y_2, y_3), b)$. It does the following depending on the bit b :

- 1) If $b = 0$, the issuer (prover) samples integers $z_0, z_u, z_v, z_w, c_1, a_1$, and computes commitments $\tilde{S}_0 := R^{z_0}$, $\tilde{S}_1 := R^{a_1} \cdot S^{-c_1}$, $\tilde{T}_0 := g_1^{z_0}$, $\tilde{T}_1 := g_1^{a_1} \cdot T_1^{-c_1}$, $\tilde{U} := g_2^{z_u}$, $\tilde{V} := g_2^{z_v}$ and $\tilde{W} := g_2^{z_w}$. It then computes the challenge $c := H(g_1, g_2, \text{pk}_C, S, T_0, T_1, U, V, W, \tilde{S}_0, \tilde{S}_1, \tilde{T}_0, \tilde{T}_1, \tilde{U}, \tilde{V}, \tilde{W})$, followed by $a_u := z_u + cy_1$, $a_v := z_v + cy_2$, $a_w := z_w + cy_3$, $c_0 := c - c_1$, $a_0 := z_0 + c_0 x_1$.

KeyGen ($1^\lambda, \ell$)	ChRep ($\text{pk}_{\text{EQ}}, \mathbf{m}, \sigma_{\text{EQ}}, \mu$)
1 : $\mathbf{x} \leftarrow \mathbb{Z}_p^*$ ^{ℓ} 2 : return $\text{pk}_{\text{EQ}} := (g_2^{\mathbf{x}} = (g_2^{x_1}, g_2^{x_2}, \dots, g_2^{x_\ell}),$ $\text{sk}_{\text{EQ}} := \mathbf{x} = (x_1, x_2, \dots, x_\ell)$	1 : if $\text{Verify}(\text{pk}_{\text{EQ}}, \mathbf{m}, \sigma_{\text{EQ}}) \neq 1$ then 2 : return \perp 3 : $\psi \leftarrow \mathbb{Z}_p^*$ 4 : return $\sigma' := (Z^{\psi \cdot \mu}, Y_1^{1/\psi}, Y_2^{1/\psi})$
Sign ($\text{sk}_{\text{EQ}}, \mathbf{m}$)	Verify ($\text{pk}_{\text{EQ}}, \mathbf{m}, \sigma_{\text{EQ}}$)
1 : Parse $\text{sk}_{\text{EQ}} := (x_1, x_2, \dots, x_\ell)$ 2 : $v \leftarrow \mathbb{Z}_p$ 3 : $Z := \left(\prod_{i=1}^{\ell} m_i^{x_i} \right)^v$ 4 : $Y_1 := g_1^{1/v}, Y_2 := g_2^{1/v}$ 5 : return $\sigma := (Z, Y_1, Y_2)$	1 : Parse $\text{pk}_{\text{EQ}} := (g_2^{x_1}, g_2^{x_2}, \dots, g_2^{x_\ell})$ 2 : return $\prod_{i=1}^{\ell} e(M_i, g_2^{x_i}) = e(Z, Y_2)$ $\wedge e(Y_1, g_2) = e(g_1, Y_2)$

Figure 12. SPS-EQ scheme from [19].

Statement: $S = R^{(1-b)x_1 + bx_2} \wedge T_0 = \text{pk}_I^{(1)} = g_1^{x_1} \wedge T_1 = \text{pk}_I^{(2)} = g_1^{x_2}$ $\wedge U = \text{pk}_I^{(3)} = g_2^{y_1} \wedge V = \text{pk}_I^{(4)} = g_2^{y_2} \wedge W = \text{pk}_I^{(5)} = g_2^{y_3}$	
$b=0$, and witness $w = (x_1, y_1, y_2, y_3)$	$b=1$, and witness $w = (x_2, y_1, y_2, y_3)$
\mathcal{P} samples $z_0, z_u, z_v, z_w, c_1, a_1 \leftarrow \mathbb{Z}_p^*$ \mathcal{P} computes $\tilde{S}_0 := R^{z_0}, \tilde{S}_1 := R^{a_1} \cdot S^{-c_1},$ $\tilde{T}_0 := g_1^{z_0}, \tilde{T}_1 := g_1^{a_1} \cdot T_1^{-c_1},$ $\tilde{U} := g_2^{z_u}, \tilde{V} := g_2^{z_v}, \tilde{W} := g_2^{z_w}$ and sends $(\tilde{S}_0, \tilde{S}_1, \tilde{T}_0, \tilde{T}_1, \tilde{U}, \tilde{V}, \tilde{W})$ to \mathcal{V} .	\mathcal{P} samples $z_1, z_u, z_v, z_w, c_0, a_0 \leftarrow \mathbb{Z}_p^*$ \mathcal{P} computes $\tilde{S}_0 := R^{a_0} \cdot S^{-c_0}, \tilde{S}_1 := R^{z_1},$ $\tilde{T}_0 := g_1^{a_0} \cdot T_0^{-c_0}, \tilde{T}_1 := g_1^{z_1},$ $\tilde{U} := g_2^{z_u}, \tilde{V} := g_2^{z_v}, \tilde{W} := g_2^{z_w}$ and sends $(\tilde{S}_0, \tilde{S}_1, \tilde{T}_0, \tilde{T}_1, \tilde{U}, \tilde{V}, \tilde{W})$ to \mathcal{V} .
\mathcal{V} samples $c \leftarrow \mathbb{Z}_p$ and sends c to \mathcal{P} .	
\mathcal{P} computes $a_u := z_u + cy_1, a_v := z_v + cy_2,$ $a_w := z_w + cy_3, c_0 := c - c_1,$ $a_0 := z_0 + c_0x_1$ and sends $(c_0, c_1, a_u, a_v, a_w, a_0, a_1)$ to \mathcal{V} .	\mathcal{P} computes $a_u := z_u + cy_1, a_v := z_v + cy_2,$ $a_w := z_w + cy_3, c_1 := c - c_0,$ $a_1 := z_1 + c_1x_2$ and sends $(c_0, c_1, a_u, a_v, a_w, a_0, a_1)$ to \mathcal{V} .
\mathcal{V} verifies $c = c_0 + c_1 \wedge \tilde{S}_0 \cdot S^{c_0} = R^{a_0} \wedge \tilde{S}_1 \cdot S^{c_1} = R^{a_1} \wedge \tilde{T}_0 \cdot T_0^{c_0} = g_1^{a_0}$ $\wedge \tilde{T}_1 \cdot T_1^{c_1} = g_1^{a_1} \wedge \tilde{U} \cdot U^c = g_2^{a_u} \wedge \tilde{V} \cdot V^c = g_2^{a_v} \wedge \tilde{W} \cdot W^c = g_2^{a_w}$	

Figure 13. The ZK Proof Protocol for the language \mathcal{L}_{iss} .

- 2) If $b = 1$, the issuer (prover) samples integers $z_1, z_u, z_v, z_w, c_0, a_0$, and computes commitments $\tilde{S}_0 := R^{a_0} \cdot S^{-c_0}, \tilde{S}_1 := R^{z_1}, \tilde{T}_0 := g_1^{a_0} \cdot T_0^{-c_0}, \tilde{T}_1 := g_1^{z_1}, \tilde{U} := g_2^{z_u}, \tilde{V} := g_2^{z_v}$ and $\tilde{W} := g_2^{z_w}$. It then computes the challenge $c := H(g_1, g_2, \text{pk}_C, S, T_0, T_1, U, V, W, \tilde{S}_0, \tilde{S}_1, \tilde{T}_0, \tilde{T}_1, \tilde{U}, \tilde{V}, \tilde{W})$, followed by $a_u := z_u + cy_1, a_v := z_v + cy_2, a_w := z_w + cy_3, c_1 := c - c_0, a_1 := z_1 + c_1x_2$. Finally, it outputs the proof $\pi = (c_0, c_1, a_u, a_v, a_w, a_0, a_1)$.
- **Verify.** The NIZK verification algorithm takes the crs, the statement $x = (S, T_0, T_1, U, V, W)$ and the proof π . We use a small modification to the Fiat-Shamir transformation similar to [10] for efficient verification. Specifically, the client (verifier) parses $\pi = (c_0, c_1, a_u, a_v, a_w, a_0, a_1)$ and computes $c = c_0 + c_1, \tilde{S}_0 := R^{a_0} \cdot S^{-c_0}, \tilde{S}_1 := R^{a_1} \cdot S^{-c_1}, \tilde{T}_0 := g_1^{a_0} \cdot T_0^{-c_0}, \tilde{T}_1 := g_1^{a_1} \cdot T_1^{-c_1}, \tilde{U} := g_2^{a_u} \cdot U^{-c}, \tilde{V} := g_2^{a_v} \cdot V^{-c}$ and $\tilde{W} := g_2^{a_w} \cdot W^{-c}$. Finally, the client checks that $c := H(g_1, g_2, \text{pk}_C, S, T_0, T_1, U, V, W, \tilde{S}_0, \tilde{S}_1, \tilde{T}_0, \tilde{T}_1, \tilde{U}, \tilde{V}, \tilde{W})$.

4.4. Batch verification. Let us now explain the joint verification check that reduces the overall number of pairing computations required for batch verifications. Firstly notice that for both, the client and the verifier, the pairing computation for $e(\text{pk}_C, \text{pk}_I^{(3)})$ and $e(g_1, \text{pk}_I^{(3)})$ respectively, can be precomputed. This already reduces the number of pairing computations per presignature/token by one. For an issuing authority with key pair $(\text{sk}_I, \text{pk}_I)$, let n be the number of presignatures issued to some client with public key pk_C , and n' be the number of tokens redeemed. Then the batch verification check for the client is given by

$$\prod_{i=1}^n e(Z_i, Y_{2,i}) \stackrel{?}{=} e(\text{pk}_C, \text{pk}_I^{(3)})^n \cdot e\left(\prod_{i=1}^n R_i, \text{pk}_I^{(4)}\right) \cdot e\left(\prod_{i=1}^n S_i, \text{pk}_I^{(5)}\right) \quad (1)$$

$$\wedge e\left(\prod_{i=1}^n Y_{1,i}, g_2\right) \stackrel{?}{=} e\left(g_1, \prod_{i=1}^n Y_{2,i}\right) \quad (2)$$

Similarly, the batch verification check for the verifier is given by replacing n by n' in (2)⁹ and linearly computing the other pairing check. So that instead of $5n$ (resp. $5n'$) pairings, the client (resp. the verifier) performs $n + 4$ (resp. $3n' + 2$) pairings.

5. Security Proofs for Section 5

Correctness, reusability of Construction 5 are easily extendable from that of the previous protocol (see Appendix C); and privacy of metadata bit is also proven identically. One-more unforgeability of Construction 5 can similarly be extended from that of the previous protocol with the additional requirement of soundness of NIZK_{obt} . Due to space constraints, these proofs are thus omitted. We now provide sketches for the unlinkability, double-spend identification and exculpability of our construction, and defer the proofs to the full version of this article.

THEOREM 5.2 (κ -UNLINKABILITY). Construction 5 is κ -unlinkable (in the random oracle and KOSK models) for $\kappa = 2$, assuming NIZK_{iss} is sound, that DDH, k -DDH assumption holds in \mathbb{G}_1 and EQ perfectly adapts signatures under a malicious signer NIZK_{obt} satisfies zero knowledge.

Pf sketch. Recall the proof of Theorem 4.3 (Appendix C). The idea there was to make the final token effectively independent of $\text{sk}_C =: \alpha$, $\text{psig} =: (\bar{\sigma}, R, S, \pi)$ and $\text{nonce} =: r$, and then use the perfect adaptation of the SPS-EQ signature scheme in order to simulate a valid token. This was facilitated by our use of the random oracle model along with assumptions from the DDH family that allowed us to replace $H(r)^{\alpha^{-1}}$ in the tag with a random group element in \mathbb{G}_1 . The approach here remains the same, however we now need to deal with several additional elements in the final token.

Notice that due to zero-knowledge of NIZK_{obt} , the proof π_o is simulatable. Next, we can

9. At verification, we are cautious to not batch first part of the pairing check similar to (1) as the client does not produce any ZK proof for its computations.

replace $\mathbf{t} = (g_1^{(\alpha+\hat{r})^{-1}}, R^{(\alpha+\hat{r})^{-1}}, S^{(\alpha+\hat{r})^{-1}})$ by $(g_1^\rho, (g_1^\varepsilon)^\rho, (g_1^{\varepsilon \cdot ((1-b)x_1 + bx_2)})^\rho)$ where we can might hope to again use inverse DDH argue indistinguishability, but this does not fully work here. Instead, we make the observation that each $g_1^{(\alpha+\hat{r})^{-1}}$ is the output of the Dodis-Yampolskiy PRF [31] which is adaptively secure under the k -DDH assumption [32] for $\text{poly}(\lambda)$ sized domains. This, in particular means that we must ensure that the range of $H_{\mathbb{Z}}(r)$ is of size $\text{poly}(\lambda)$. However, we remark that this does not influence security of our scheme in any way, although resuability will now hold only with high (and not all but negligible) probability.

Lastly, viewing ϕ as ElGamal encryption of $g_1^{\text{id} \cdot (\alpha+\hat{r})}$ (times some extra randomness), we replace it with the encryption of a random value and argue indistinguishability by IND-CPA security of the encryption scheme (which, in turn, follows from DDH). At a high level, we first set each $H_{\mathbb{G} \rightarrow \mathbb{G}}(r) := g_1^{r_r}$. Then, for a query of the form (\cdot, r, id) , we create the encryption ϕ as $\phi_1 := g_1^{r_r}$, and $\phi_2 := \mu \cdot \phi_1^\alpha$ for message $\mu := \text{pk}_C^{\text{id}} \cdot g_1^{(\text{id} + \nu_r) \cdot \hat{r}}$. Next, we can replace $\phi_1^\alpha = g_1^{\alpha \cdot \nu_r}$ with g_1^γ for $\gamma \leftarrow \mathbb{Z}_p$, so that $\phi_2 = \mu \cdot g_1^\gamma$ is indistinguishable from uniform. Consequently, the final token is now independent of pk_C , psig and nonce as desired. \square

THEOREM 5.4 (DOUBLE-SPEND IDENTIFICATION). Construction 5 satisfies double-spend identification (in the random oracle model) assuming NIZK_{obt} is sound.

Pf sketch. If NIZK_{obt} is sound, then it follows that the adversary knows (sk_C, \hat{r}) (resp. (sk'_C, \hat{r}')) such that $t_1 = g_1^{(\text{sk}_C + \hat{r})^{-1}}$ (resp. $t_1 = g_1^{(\text{sk}'_C + \hat{r}')^{-1}}$) and $\phi_2 = g_1^{\text{id} \cdot (\text{sk}_C + \hat{r})}$ (resp. $\phi_2 = g_1^{\text{id}' \cdot (\text{sk}'_C + \hat{r}')}$). $H_{\mathbb{G} \rightarrow \mathbb{G}}(t_1)^{\text{sk}_C + \hat{r}}$ (resp. $\phi_2' = g_1^{\text{id}' \cdot (\text{sk}'_C + \hat{r}')} \cdot H_{\mathbb{G} \rightarrow \mathbb{G}}(t_1)^{\text{sk}'_C + \hat{r}'}$) with respect to (\mathbf{t}, σ) (resp. (\mathbf{t}, σ')). Now since \mathbf{t} is common between both tokens, we have $(\text{sk}_C + \hat{r}) = (\text{sk}'_C + \hat{r}')$ (all computations are mod p) so that $(\phi_2^{-1} \cdot \phi_2')^{(\text{id} - \text{id}')^{-1}}$ simplifies to $g_1^{\text{sk}_C + \hat{r}} = g_1^{\text{sk}'_C + \hat{r}'}$, such that at least one of $(g_1^{\text{sk}_C}, \text{nonce})$ and $(g_1^{\text{sk}'_C}, \text{nonce}')$ is in aux with $\hat{r} =: H_{\mathbb{Z}}(\text{nonce})$ (similarly, \hat{r}'). Thus, with all but negligible probability, $\text{DSIden}(\text{pk}_I, \mathbf{t}, \sigma, \sigma') \neq \perp$. \square

THEOREM 5.5 (DOUBLE-SPEND EXCULPABILITY). Construction 5 satisfies double-spend exculpability assuming NIZK_{obt} is sound and the discrete log assumption holds in \mathbb{G}_1 .

Pf sketch. Suppose an adversary accuses some (honest) client pk_C and provides a proof of guilt $\Pi := (\mathbf{t}, \sigma, \sigma', \text{nonce})$ where $\sigma := (s, \phi, \pi_o, \text{id})$ (similarly σ') such that $(\phi_2^{-1} \cdot \phi_2')^{(\text{id} - \text{id}')^{-1}} = \text{pk}_C \cdot g_1^{H_{\mathbb{Z}}(\text{nonce})}$ and, both, the proof and signature verify. Having queried $\mathcal{O}_{\text{Issue}}$ at most once per $(\text{psig}, \text{nonce})$ pair (w.l.o.g. suppose the query corresponds to the token (\mathbf{t}, σ)) then, such an adversary must have to create a satisfying proof π'_o . However, in order to do so, it either learns sk_C —in which case, we can reduce to the hardness of discrete log—or it is able to create a

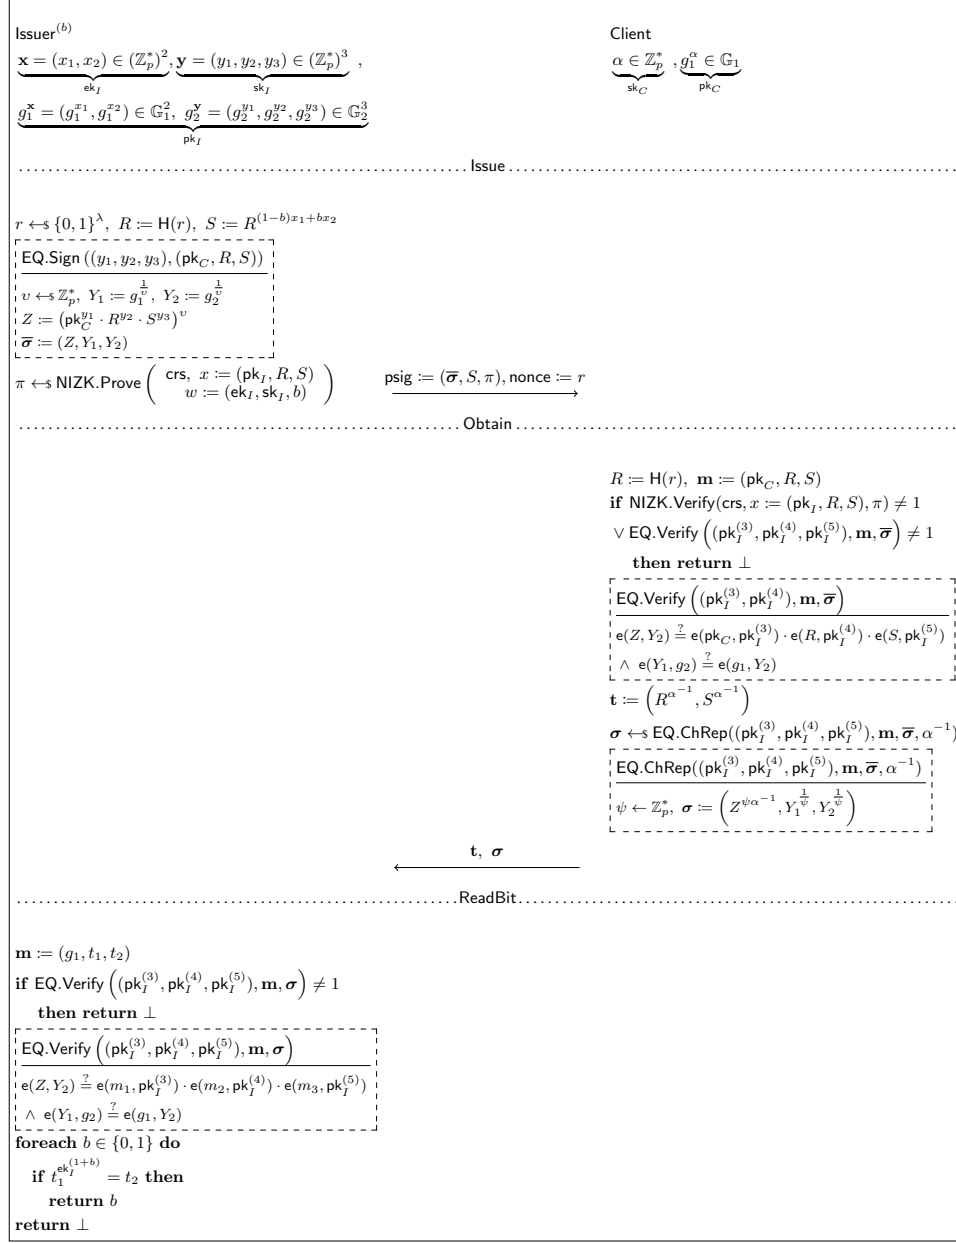


Figure 14. Expanded SPS-EQ construction for NIAT presignature issuance and token generation using [19].

proof without sk_C , and we can reduce to the soundness of NIZK_{obt} . \square

6. NIAT with Public Attributes

We now show that our NIAT constructions can be modified to also include public attributes. The construction is almost identical to our Construction 4. The only notable change is that the issuer now uses a tag-based equivalence class signature scheme instead of a structure-preserving signature.

6.1. Construction. Tools required. Our construction requires a hash function $H : \{0, 1\}^\lambda \rightarrow \mathbb{G}_1$ modeled as a random oracle, a TB-EQS scheme $\text{TEQ} = (\text{TEQ.Setup}, \text{TEQ.KeyGen}, \text{TEQ.Sign}, \text{TEQ.ChRep}, \text{TEQ.Verify}, \text{TEQ.VerKey})$, and a NIZK $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ for \mathcal{L}_{iss} .

Presignature issuance and token generation. The presignature issuance and token generation protocols are presented in detail in Figure 15. The difference here from construction 4 is the use of the public attribute τ as input to the TB-EQS scheme.

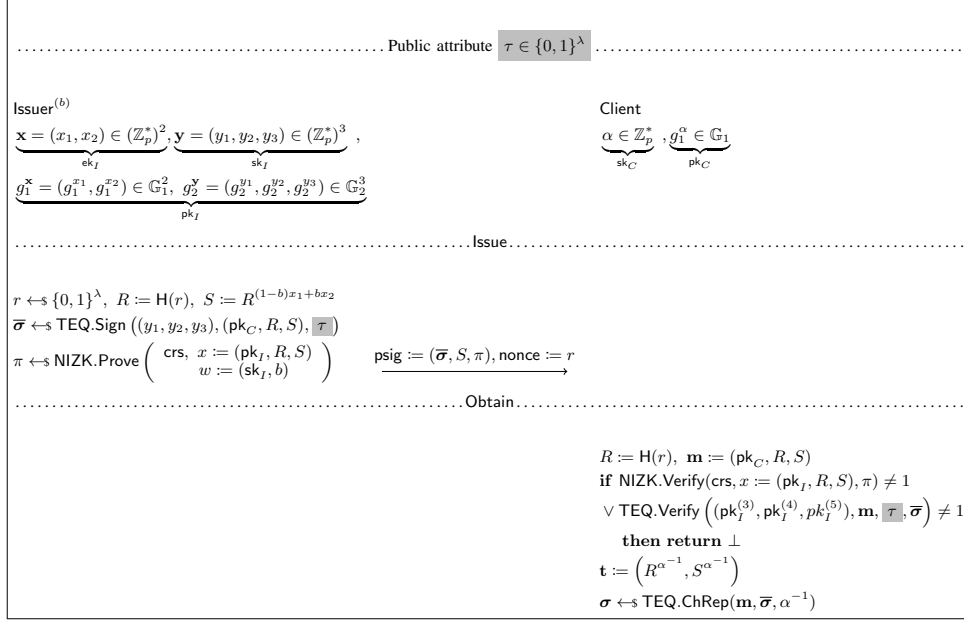


Figure 15. TB-EQS construction for NIAT (with public attributes) presignature issuance and token generation.

Public verification (token redemption). The public verification algorithm simply verifies the tag-based equivalence class signature σ for the public attribute τ .

Proof. Proof is omitted as it is nearly identical to that of Theorem 4.4. \square

$\text{Verify}(\text{pk}_I, \mathbf{t}, \sigma, \tau)$

1 : return $\text{TEQ.Verify}((\text{pk}_I^{(3)}, \text{pk}_I^{(4)}, \text{pk}_I^{(5)}), (g_1, t_1, t_2), \tau, \sigma)$

Correctness, reusability and all security properties follow similarly from that of construction 4.

Security.

Theorem A.5 (One-more unforgeability). *Construction F is one-more unforgeable (in the random oracle model) assuming NIZK satisfies zero knowledge and TEQ is existentially unforgeable under adaptively chosen-message attacks.*

Proof. Proof is omitted as it is nearly identical to that of Theorem 4.1 but with respect to EUnf-CMA security of the TB-EQS scheme. \square

Theorem A.6 (κ -unlinkability). *Construction F is κ -unlinkable (in the random oracle and KOSK models) for $\kappa = 2$, assuming NIZK is an argument of knowledge, that inverse DDH assumption holds in \mathbb{G}_1 and TEQ perfectly adapts signatures under a malicious signer.*

Proof. Proof is omitted as it is nearly identical to that of Theorem 4.3 but with respect to perfect signature adaptation of the TB-EQS scheme. \square

Theorem A.7 (Privacy of metadata bit). *Construction F has private metadata bit (in the random oracle model) assuming NIZK satisfies zero-knowledge, that DDH assumption holds and EQ is existentially unforgeable under adaptively chosen-message attacks.*