Non-interactive Anonymous Tokens with Private Metadata Bit

Foteini Baldimtsi
foteini@gmu.edu
George Mason University and Mysten Labs
Fairfax, USA

Lucjan Hanzlik
hanzlik@cispa.de
CISPA Helmholtz Center for Information Security
Saarbrücken, Germany

Quan Nguyen
qnguye31@gmu.edu
George Mason University
Fairfax, USA

Aayush Yadav
ayadav5@gmu.edu
George Mason University
Fairfax, 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. However, all existing ATPM protocols require interaction during issuance as the client must send a blinded request and wait for the issuer to sign, introducing latency and scalability bottlenecks.

In this work, we present the first *Non-interactive Anonymous Token (NIAT) scheme with a private metadata bit*. Our design builds on structure-preserving signatures on equivalence classes (SPS-EQ) and recent advances in non-interactive blind signatures. With NIAT, tokens can be issued without any online interaction, enabling asynchronous pre-computation and drastically reducing issuer workload. We formalize NIAT security definitions, propose an efficient construction under standard assumptions 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.

CCS Concepts

• **Security and privacy → Privacy-preserving protocols**; *Public key (asymmetric) techniques.*

Keywords

Anonymous tokens; double-spending; non-interactive schemes; signatures on equivalence classes

ACM Reference Format:
Foteini Baldimtsi, Lucjan Hanzlik, Quan Nguyen, and Aayush Yadav. 2026. Non-interactive Anonymous Tokens with Private Metadata Bit. In *Proceedings of the 2026 ACM SIGSAC Conference on Computer and Communications Security (CCS '26), November 15–19, 2026, The Hague, Netherlands.* ACM, New York, NY, USA, 23 pages. https://doi.org/10.1145/3830454.3832714

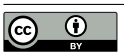

This work is licensed under a Creative Commons Attribution 4.0 International License.
CCS '26, The Hague, Netherlands
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2871-6/2026/11
https://doi.org/10.1145/3830454.3832714

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 *redemptee* (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 redemptee, who verifies its authenticity to grant the client access to a service or resource, and additionally checks that the token has not already been redeemed (double-spending detection). The basic properties satisfied by an anonymous token system are *unforgeability*, which ensures that a client cannot issue tokens on its own, and *anonymity* which prevents issuers and redemptees to link issued tokens to any token later redeemed.

Anonymous tokens have numerous applications, including access to content delivery networks (CDNs) without the need for CAPTCHA's [14], private web-browsing [2], private contact tracing [33], fraud detection [33] and private click measurement [35] just to name a few. Their significance is evidenced by the growing industry interest, with major initiatives such as Google's Private State Tokens [20] and Cloudflare's Privacy Pass [27]. At the same time, standardization efforts are underway within the IETF [25], backed by key stakeholders including Google, Apple, Cloudflare, and Fastly.

Given the increased interest in anonymous tokens, a number of protocols have been proposed in the literature [6, 11, 14, 26, 30, 33] satisfying a variety of properties and offering interesting tradeoffs. A first distinction between existing works is in regards to whether they support private or public token verification, i.e., whether the issuer and the redemptee are the same party (private setting) [11, 14, 26] or any third party can verify the issued tokens (public setting) [6, 33]. Having a single party for both issuance and redemption has the benefit of making double-spending detection much easier but it is restrictive for applications with multiple verifiers or verification of tokens in a public setting (such as a blockchain smart contract). On the other hand, public verification can enhance the application of anonymous tokens in broader systems by outsourcing signing or verification of tokens, but in order to detect double-spending, you still need to rely on centralized mechanisms.

A second distinction is whether the scheme allows for *public or private metadata*. Public metadata enables applications like geographical tags that are inherently visible to all verifiers and users. In contrast, anonymous tokens with private metadata bit (ATPM),

which is typically a single bit, enable asymmetric trust signaling. The issuer embeds a bit into the token so that the user (and ordinary verifiers) cannot distinguish which bit was set, while only entities holding a secret extraction key—possibly distinct from the secret key used for issuance—can later recover it. This asymmetry is essential in adversarial settings, since public trust signals would allow malicious users to immediately adapt their behavior.

Schemes supporting public verifiability with a private metadata bit are especially relevant because they allow one central authority to issue tokens with embedded trust signals while many verifiers can independently check token validity without contacting the issuer. Public verifiability ensures scalability and deployment in federated environments, while the private bit enables asymmetric trust signaling by allowing risk or eligibility information to be embedded in the token during issuance. The Privacy Pass RFC [15] explicitly mandates that publicly verifiable token schemes support both public and private metadata.

Importantly, we observe that the extraction capability for the private bit need not be restricted to the issuer alone. The issuer may delegate extraction rights to selected ("premium") verifiers that require stronger trust assessment, while ordinary verifiers only perform public validity checks. The combination of public verifiability and hidden trust signals is particularly useful in settings such as blockchain systems, where verification must remain open and decentralized, yet certain eligibility or compliance signals must be visible only to designated parties.

The problem with existing solutions. A common drawback among all existing schemes for anonymous tokens is the requirement for an interactive issuance protocol. In this model, the client submits a blinded request, waits for the issuer to sign online, and only then obtains usable tokens. The interactive nature of issuance leads to practical challenges. In particular, it introduces latency in the system as each issuance requires a full round trip before tokens can be used, which can be particularly problematic in real-time applications. It also raises scalability concerns since the issuer must execute a relatively expensive cryptographic signing operation at request time and potentially handle millions of such online signing queries simultaneously.

Beyond issuance, existing systems also typically require some form of online coordination during verification in order to prevent double-spending. In practice, this often means maintaining a synchronized record of previously redeemed tokens across verifiers. Such synchronization can become cumbersome in decentralized or multi-verifier settings, where different verifiers may not share a common online state or may experience delays in propagating spent-token information.

These challenges motivate the central question of this work:

> *Can we design an anonymous token scheme that eliminates interactive issuance while still supporting public verifiability, private metadata bits, and offers robust double-spend protection?*

Our approach. We affirmatively answer this question by proposing the first Non-interactive Anonymous Token (NIAT) scheme with a private metadata bit. Our main observation is that at the heart of many of the proposed protocols is a blind signature scheme on a *randomly selected* message [11, 14, 26, 33]. By definition, blind signatures require at least one round (i.e., two moves) of interaction between the signer/issuer and the client [12, 31, 32]. However, a series of recent works [3, 21, 22] have demonstrated the feasibility of constructing non-interactive blind signatures (NIBS) as long as the signed message is random and does not come from a specific distribution or has a specific structure.

The core idea behind NIBS is that each client is associated with a public-secret key pair $(\mathsf{pk}_C, \mathsf{sk}_C)$ 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 protocols, the signed message is effectively a random identifier, thus the NIBS constructions of [3, 21] could give rise to an anonymous token scheme with public verifiability. The constructions of [3, 21] 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 [1].

Interestingly, the NIBS-based structure also opens the door to practical offline double-spending identification. Since tokens are derived from presignatures tied to client public keys, the resulting construction can admit mechanisms for identifying double-spenders upon repeated token redemption, similarly to classical e-cash systems.

Our Contributions. We summarize our contributions as follows:

❶ **Defining non-interactive anonymous tokens.** We propose a formal framework for NIAT in Section 3. In particular, we define one-more unforgeability, public-validity consistency, unlinkability, metadata bit privacy, and reusability (namely the property that an issuer is able to issue multiple tokens for the same client public key).[1] Crafting the right definitions for NIAT requires extreme care, and we consider this definitional framework to be a central contribution of our work. In particular, the subtleties introduced by non-interactivity and public verifiability give rise to security and modeling challenges beyond prior work [11].

❷ **NIAT from SPS-EQ.** In Section 4, we give our main NIAT construction from structure-preserving signatures on equivalence classes (SPS-EQ) [18]. Our SPS-EQ construction builds upon the SPS-EQ NIBS construction of [21] 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]

[1] We note that this property is not needed for the interactive schemes of prior work.

[2] At the moment, the NIBS of [21] seems to be the best starting point for the most efficient NIAT constructions. The lattice-based design of [3] provides post-quantum security but suffers from much larger parameters, while the concurrent RSA-based

❸ **NIAT with double-spend identification.** For the first time, we define and present an 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 requiring an immediate online check for uniqueness at redemption time. If a token is later redeemed more than once, the double-spender can be identified post-event, enabling detection and potential penalties.

Our approach leverages the structure of NIAT tokens, which are derived from presignatures associated with client public keys. This naturally enables mechanisms for identifying misbehaving clients upon repeated token redemption, similarly to classical e-cash systems. To support this functionality, token presentations additionally incorporate verifier-specific identifiers, which allow repeated use of the same token across different verification contexts to be detected and linked to the offending client.

Our construction of NIAT with double-spend identification is presented in Section 5. We also note that supporting double-spend identification also introduces new definitional and modeling challenges, particularly regarding the interaction between unlinkability and verifier-specific information which are discussed in Section 3.

We note that the private metadata bit and the double-spending mechanism *reinforce each other in both directions*. Assume a central log or server that records all spent tokens, periodically resetting according to the token validity period. Information from detected double-spending can guide the issuer to set the private bit adaptively in future issuances, for example by flagging clients who double-spend frequently or in suspicious contexts. Conversely, the private bit can influence how double-spending checks are applied. "Premium" verifiers authorized to read the bit may perform stricter verification or consult the spent-token log before granting access when a token carries a low-trust flag.

❹ **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 previous 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 [14]. Currently, in a CDN system, there exists an attestation 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 unstruthworthy origin. The client will solve the puzzle and respond with the solution to prove that they are a real human user. The attestation 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 [14].

In the setting of Privacy Pass, when a client sends the CAPTCHA solution to the attestation 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; [14] 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 [24].

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

construction of [22] 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 [21] as a starting point.

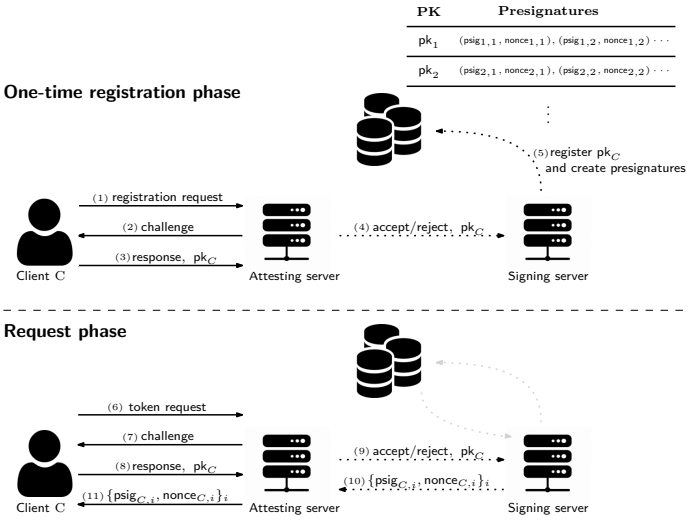

Figure 1: NIAT issuance for CDNs with offline signing.

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. In Section 6, paragraph "Offline pre-computation", we provide some relevant estimate on the advantage of pre-computation versus online token issuance.

Beyond non-interactivity, our design offers additional properties to CDNs. Existing deployments, including Privacy Pass, rely on binary and immediate issuance decisions: a client either receives tokens or is blocked. Such decisions provide direct feedback to adversaries, enabling probing attacks and preventing the accumulation of longer-term evidence about client behavior. This limits the ability of the CDN to apply nuanced trust assessment over time. Our support of private metadata on top of public verifiability, resolves this problem as clients can redeem tokens across multiple CDN servers without contacting the issuer, while tokens carry hidden trust signals extractable only by authorized ("premium") servers. Finally, our support for offline double-spending detection further strengthens the properties of anonymous CDN access by allowing a delayed check for double-spending and thus faster client access.

1.2 Related Work

Privacy Pass [14] 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 [34] may be used in place of VOPRFs for a more efficient and flexible instantiation of Privacy Pass. One disadvantage of the Privacy Pass protocol and other similar anonymous token protocols [33] 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 [26] 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 [6] 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 [1], however they do not support private metadata.

Chase et al. [11], 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 [26]. 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) [7] 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 $\coloneqq$ 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 ith element.

2.1 Bilinear Pairings

Definition 2.1 (Bilinear pairings)**.** Given a security parameter λ, a bilinear group generator $\mathsf{BG}(1^\lambda)$ returns a tuple $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \mathsf{e}, p, g_1, g_2)$, where $\mathbb{G}_1$ and $\mathbb{G}_2$ 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, $\mathsf{e} : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$, satisfying the following properties:

- **Bilinearity:** $\forall h_1 \in \mathbb{G}_1, h_2 \in \mathbb{G}_2,\ a, b \in \mathbb{Z}_p$,
$$e(h_1^a, h_2)^b = \mathsf{e}(h_1, h_2^b)^a = \mathsf{e}(h_1, h_2)^{ab}\ .$$
- **Non-degeneracy:** $\mathsf{e}(g_1, g_2) \neq 1_{\mathbb{G}_T}$.

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) [19]. 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 $\beta = \alpha^{-1} \mod p$ or $\beta \leftarrow\!\$\ \mathbb{Z}_p$.

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

2.3 Structure-preserving Signatures on Equivalence Classes

Structure-preserving Signatures on Equivalence Classes (SPS-EQ) [18] are used to sign equivalence classes $[\mathbf{m}]$ of message vectors $\mathbf{m} \in (\mathbb{G}_i^*)^\ell$ for $\ell > 1$, having the equivalence relation $\mathbf{m}, \mathbf{n} \in \mathbb{G}_i^\ell$: $\mathbf{m} \sim_{\mathscr{R}} \mathbf{n} \Leftrightarrow \exists s \in \mathbb{Z}_p : \mathbf{m} = \mathbf{n}^s$. We recall the definition for SPS-EQ [18], adapting the notation for multiplicative groups as in [21].

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

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

An SPS-EQ scheme satisfies correctness, existentially unforgeable under adaptive chosen-message attacks (EUnf-CMA), and perfect signature adaptation under malicious keys defined as follows:

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

$$\mathsf{EQ.VerKey}(\mathsf{sk}_{\mathsf{EQ}}, \mathsf{pk}_{\mathsf{EQ}}) = 1 \wedge$$
$$\Pr[\mathsf{EQ.Verify}(\mathsf{pk}_{\mathsf{EQ}}, \mathbf{m}, \sigma_{\mathsf{EQ}}) = 1] = 1 \wedge$$
$$\Pr\left[\mathsf{EQ.Verify}\left(\mathsf{pk}_{\mathsf{EQ}}, \mathbf{m}^\mu, \mathsf{EQ.ChRep}(\sigma_{\mathsf{EQ}}, \mu)\right) = 1\right] = 1$$

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

$Q \coloneqq \varnothing$
$\mathsf{pp}_{\mathsf{EQ}} \leftarrow \mathsf{EQ.Setup}(1^\lambda, \ell)$
$(\mathsf{pk}_{\mathsf{EQ}}, \mathsf{sk}_{\mathsf{EQ}}) \leftarrow\!\$\ \mathsf{EQ.KeyGen}(pp_{\mathsf{EQ}})$
$(\mathbf{m}^*, \sigma^*_{\mathsf{EQ}}) \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{Sign}}}(\mathsf{pk}_{\mathsf{EQ}})$
return $[\mathbf{m}^*] \neq [\mathbf{m}]\ \forall \mathbf{m} \in Q \wedge$ $\mathsf{EQ.Verify}(\mathsf{pk}_{\mathsf{EQ}}, \mathbf{m}^*, \sigma^*_{\mathsf{EQ}}) = 1$

Oracle $\mathcal{O}_{\mathsf{Sign}}(\mathbf{m})$

$\sigma_{\mathsf{EQ}} \leftarrow\!\$\ \mathsf{EQ.Sign}(\mathsf{sk}_{\mathsf{EQ}}, \mathbf{m})$
$Q \coloneqq Q \cup \{\mathbf{m}\}$
return σ_{EQ}

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

Definition 2.6 (Existential unforgeability under chosen message attack)**.** Let $\mathsf{Adv}^{\mathsf{EUnf\text{-}CMA}}_{\mathcal{A},\ell} = \Pr\left[\mathsf{EUnf\text{-}CMA}_{\mathcal{A},\ell}(\lambda) = \mathbf{accept}\right]$ 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 $\mathsf{Adv}^{\mathsf{EUnf\text{-}CMA}}_{\mathcal{A}} = \mathsf{negl}(\lambda)$.

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 $(\mathsf{pk}_{\mathsf{EQ}}, \mathbf{m}, \sigma_{\mathsf{EQ}}, \mu)$ satisfying

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

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

In this work, we will use the pairing-based construction of SPS-EQ due to [18].

2.4 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 $x \in \mathscr{L}$ such that $\mathscr{R}(x, w) = 1$, and $\mathcal{V}$ gains no additional information. A ZK proof system consists of the following polynomial time algorithms:

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

Definition 2.8 (Zero-knowledge proof)**.** A zero-knowledge proof between $\mathcal{P}$ and $\mathcal{V}$ for an NP relation $\mathscr{R}$ must satisfy the following properties:

- **Completeness:** If $\mathscr{R}(x, w) = 1$ 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., $\forall w : \mathscr{R}(x, w) = 0$), 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

$$\Pr\left[\begin{array}{c|c} \mathsf{Verify}(\mathsf{crs}, x, \pi) = 1 & \mathsf{crs} \leftarrow\!\$\ \mathsf{Setup}(1^\lambda) \\ \wedge\ \mathscr{R}(x, \mathcal{E}(\mathsf{crs}, x, \pi)) = 0 & (x, \pi) \leftarrow \mathcal{A}(1^\lambda, \mathsf{crs}) \end{array}\right]$$

- **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 $|x|$

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

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 [13].

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 [17].

3 Non-interactive Anonymous Tokens

We now formally define a non-interactive anonymous tokens (NIAT). As discussed in the introduction, our definitions are inspired by existing definitions of anonymous tokens with private metadata bit (ATPM) [11, 26] properly adapted to capture the non-interactive properties of non-interactive blind signatures (NIBS) [3, 21].

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):

- $\mathsf{Setup}(1^\lambda) \rightarrow (\mathsf{pp}, \mathsf{aux})$. 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.
- $\mathsf{KeyGen}_I(\mathsf{pp}) \rightarrow (\mathsf{pk}_I, \mathsf{sk}_I, \mathsf{ek}_I)$. 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.
- $\mathsf{KeyGen}_C(\mathsf{pp}) \rightarrow (\mathsf{pk}_C, \mathsf{sk}_C)$. The client's key generation algorithm takes pp as input, and outputs a public-secret key pair $(\mathsf{pk}_C, \mathsf{sk}_C)$.
- $\mathsf{Issue}(\mathsf{sk}_I, \mathsf{ek}_I, \mathsf{pk}_C, b, \tau, \mathsf{aux}) \rightarrow (\mathsf{psig}, \mathsf{nonce}, \mathsf{aux})$. 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).
- $\mathsf{Obtain}(\mathsf{sk}_C, \mathsf{pk}_I, \mathsf{psig}, \mathsf{nonce}, \tau, \mathsf{id}) \rightarrow (t, \sigma)$ 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.
- $\mathsf{Verify}(\mathsf{pk}_I, t, \sigma, \tau) \rightarrow 1/0$. 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 1 (accept) or 0 (reject).
- $\mathsf{ReadBit}(\mathsf{ek}_I, t) \rightarrow b \in \{0, 1\}$ 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:

- $\mathsf{DSIden}(\mathsf{pk}_I, t, \sigma, \sigma', \mathsf{aux}) \rightarrow (\mathsf{pk}_C, \Pi)$ or $(\perp, \perp)$. The double-spend identification algorithm takes as input the issuer's public key pk_I, a tag t, two different presentations (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 $(\perp, \perp)$.
- $\mathsf{DSVer}(\mathsf{pk}_C, \Pi, \mathsf{aux}) \rightarrow b \in \{0, 1\}$. 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 $b \in \{0, 1\}$.

Our interface has two algorithms that could signal whether or not a token is valid: the Verify and ReadBit algorithms. However, as pointed out by [11], 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 applications where token verification is outsourced to a standalone server.

3.1 Security Definitions

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

$$\Pr\left[\mathsf{Verify}(\mathsf{pk}_I, t, \sigma, \tau) = 1 \wedge \mathsf{ReadBit}(\mathsf{pk}_I, \mathsf{ek}_I, t, \sigma, \tau) = b\right] = 1$$

Given that our scheme is non-interactive, we formalize the notion of *reusability* for NIAT based on a similar notion from the

[3] Note that ReadBit is a deterministic function of the issuer's secret extraction key and the tag. Additionally, as discussed in the introduction, ek_I could be potentially delegated to "premium" verifiers. This delegation process is orthogonal to the core scheme.

NIBS literature [3]. 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.

We note that this property is not needed in the ATPM definition given by Chase et al. [11] 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.

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 $\tau \in \mathcal{M}$ (for some message space $\mathcal{M}$) the following probability is negligible

$$\Pr\left[\begin{array}{l} \mathsf{nonce}_0 = \mathsf{nonce}_1 \ \vee \ t_0 = t_1 \ : \\ \qquad\qquad \mathsf{pp} \leftarrow\!\$\ \mathsf{Setup}(1^\lambda) \\ \qquad\qquad (\mathsf{pk}_I, \mathsf{sk}_I, \mathsf{ek}_I) \leftarrow\!\$\ \mathsf{KeyGen}_I(\mathsf{pp}) \\ \qquad\qquad (\mathsf{pk}_C, \mathsf{sk}_C) \leftarrow\!\$\ \mathsf{KeyGen}_C(\mathsf{pp}) \\ \quad \forall i \in \{0,1\} : (\mathsf{psig}_i, \mathsf{nonce}_i) \leftarrow\!\$\ \mathsf{Issue}\left(\begin{array}{l}\mathsf{sk}_I, \mathsf{ek}_I, \\ \mathsf{pk}_C, b, \tau\end{array}\right) \\ \quad \forall i \in \{0,1\} : (t_i, \sigma_i) \leftarrow\!\$\ \mathsf{Obtain}\left(\begin{array}{l}\mathsf{sk}_C, \mathsf{pk}_I, \\ \mathsf{psig}_i, \mathsf{nonce}_i, \tau\end{array}\right) \end{array}\right]$$

We now define one-more unforgeability for NIAT. Our adversary is given access to an issuance oracle $\mathcal{O}_{\mathsf{Issue}}$ that allows it to request pre-signatures (and the corresponding nonces) for any client public key pk_C, metadata bit $b \in \{0,1\}$, and public metadata τ of its choice. It is also given access to a ReadBit oracle.

Our definition captures the inability of an adversary to increase the number of usable tokens corresponding to any fixed metadata bit. Concretely, after interacting with the issuance oracle, the adversary outputs a target bit b^*, a public metadata value τ^*, and a collection of tokens. It wins if it produces more valid tokens extracting to b^* under τ^* than the number of issuance queries made for that same pair (b^*, τ^*).

This formulation also captures classical one-more forgery attacks. In particular, if an adversary produces more valid tokens than the total number of issuance queries, then by the pigeonhole principle it must produce more tokens for some bit b, which yields a winning strategy in the above game.

As with the existing definition of [11], we stress that we do not aim to rule out 'one-for-one' bit switching, where an adversary transforms a token issued under bit 0 into a token under bit 1 and vice-versa, preserving the *total* number of tokens per bit. Such transformations do not violate the above counting condition. However, in practice such a transformation does not give the adversary any additional power as it does not increase the number of usable tokens of any type, nor does it enable amplification of trust signals or bypass of system policies. In particular, any meaningful attack would require a *net increase* in tokens of a given bit, which is precisely what our definition rules out.

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

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

$\mathsf{pp}, \mathsf{aux} \leftarrow\!\$\ \mathsf{Setup}(1^\lambda)$
$(\mathsf{pk}_I, \mathsf{sk}_I, \mathsf{ek}_I) \leftarrow\!\$\ \mathsf{KeyGen}_I(\mathsf{pp})$
$b^*, \tau^*, \{(t_i, \sigma_i)\}_{i\in[N]} \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{Issue}}, \mathcal{O}_{\mathsf{Read}}}(\mathsf{pp}, \mathsf{pk}_I)$
if $\#\{t_i\}_{i\in[N]} \le \mathsf{ctr}_{b^*,\tau^*}$ **then reject**
foreach $i \in [N]$ **do**
  if $\mathsf{Verify}(\mathsf{pk}_I, t_i, \sigma_i, \tau^*) \neq 1$ **then reject**
  if $\mathsf{ReadBit}(\mathsf{ek}_I, t_i) \neq b^*$ **then reject**
accept

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

if $\mathsf{ctr}_{b,\tau}$ is uninitialized **then** $\mathsf{ctr}_{b,\tau} \coloneqq 0$
$(\mathsf{psig}, \mathsf{nonce}, \mathsf{aux}) \leftarrow\!\$\ \mathsf{Issue}(\mathsf{sk}_I, \mathsf{ek}_I, \mathsf{pk}_C, b, \tau, \mathsf{aux})$
$\mathsf{ctr}_{b,\tau} \leftarrow \mathsf{ctr}_{b,\tau} + 1$
return $(\mathsf{psig}, \mathsf{nonce})$

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

return $\mathsf{ReadBit}(\mathsf{ek}_I, t)$

Figure 3: The one-more unforgeability experiment.

We also define public-validity consistency which ensures that public verification and private bit extraction define compatible notions of validity: any token accepted by Verify must have a well-defined embedded metadata bit. This property is not needed in [11], but is important in publicly verifiable settings, where verification may be performed by parties that do not hold the extraction key.

Definition 3.4 (Public-validity consistency)**.** Let $\mathsf{Adv}^{\mathsf{PVC}}_{\mathcal{A}} = \Pr[\mathsf{PVC}_{\mathcal{A}}(\lambda) = \textbf{accept}]$ be the advantage of an adversary $\mathcal{A}$ in the experiment defined in Figure 4. We say a NIAT scheme satisfies public-validity consistency if for any PPT adversary $\mathcal{A}$ this advantage is negligible.

Remark 3.5. While designing our security definitions, we also identified an alternate conception of unforgeability that folds the conditions of public validity consistency into itself. While our proofs use the version of the definitions presented here in order to remain consistent with prior work, we also present our alternate definition in Appendix A for completeness. We conjecture that the two ways of defining unforgeability are equivalent, but the formal proof (or refutation) is left as future work.

Next, we formalize NIAT unlinkability. The overall intuition here remains the same as in the definitions of [11], except that the adversary is now equipped with a GenUser oracle $\mathcal{O}_{\mathsf{GenUser}}$ that issues new public keys, and an Obtain oracle $\mathcal{O}_{\mathsf{Obtain}}$ that returns the tokens on chosen $(\mathsf{psig}, \mathsf{nonce}, \tau)$ triples.

Game $\mathrm{PVC}_{\mathcal{A}}(\lambda)$

1 : $\mathsf{pp}, \mathsf{aux} \leftarrow\$ \mathsf{Setup}(1^\lambda)$
2 : $(\mathsf{pk}_I, \mathsf{sk}_I, \mathsf{ek}_I) \leftarrow\$ \mathsf{KeyGen}_I(\mathsf{pp})$
3 : $(\tau^*, t^*, \sigma^*) \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{Issue}}, \mathcal{O}_{\mathsf{Read}}}(\mathsf{pp}, \mathsf{pk}_I)$
4 : **if** $\mathsf{Verify}(\mathsf{pk}_I, t^*, \sigma^*, \tau^*) = 1$
$\quad \wedge\ \mathsf{ReadBit}(\mathsf{ek}_I, t^*) = \perp$ **then accept**
5 : **reject**

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

1 : $(\mathsf{psig}, \mathsf{nonce}, \mathsf{aux}) \leftarrow\$ \mathsf{Issue}(\mathsf{sk}_I, \mathsf{ek}_I, \mathsf{pk}_C, b, \tau, \mathsf{aux})$
2 : **return** $(\mathsf{psig}, \mathsf{nonce})$

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

1 : **return** $\mathsf{ReadBit}(\mathsf{ek}_I, t)$

Figure 4: The public-validity consistency experiment.

Definition 3.6 (κ-unlinkability)**.** Let $\mathsf{Adv}^{\mathsf{UNLINK}}_{\mathcal{A},n} = \Pr\left[\mathsf{UNLINK}_{\mathcal{A},n}(\lambda) = \textbf{accept}\right]$ 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 5. 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 $\kappa/n + \mathsf{negl}(\lambda)$.[4]

Concretely, the parameter n is the number of pairwise distinct $(\mathsf{pk}_C, \mathsf{psig}, \mathsf{nonce})$ triples in the challenge set Q^*. Notice that we do not enforce any other constraints on the values inside Q^*, so some public keys can repeat. In particular, when $n = 2$, $\kappa = 1$, and $\mathsf{pk}_{C_1} \neq \mathsf{pk}_{C_2}$ (resp. $\mathsf{pk}_{C_1} = \mathsf{pk}_{C_2}$), we can view this experiment as being analogous to the stronger versions of NIBS receiver (resp. nonce) blindness [3]. Thus we view our unlinkability definition as a generalization of the NIBS blindness definitions.

Remark 3.7. An important requirement for unlinkability is that all $(\mathsf{psig}, \mathsf{nonce})$ pairs in Q^* must be under the same public metadata, otherwise the property is trivially broken. This also holds for the definition given in [11].

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

Definition 3.8 (Privacy of metadata bit)**.** Let $\mathsf{Adv}^{\mathsf{PMB}}_{\mathcal{A}} = \left|\Pr\left[\mathsf{PMB}_{\mathcal{A},1}(\lambda) = \textbf{accept}\right] - \Pr\left[\mathsf{PMB}_{\mathcal{A},0}(\lambda) = \textbf{accept}\right]\right|$ be the advantage of an adversary $\mathcal{A}$ in the experiment defined in Figure 6. We say a NIAT scheme NIAT preserves privacy of the metadata bit if for any PPT adversary $\mathcal{A}$, this advantage is negligible.

NIAT with double-spending identification security. Recall from the syntax above that in a NIAT scheme with double-spend identification, the verifier identifier id appears as an input to the Obtain algorithm. In particular, presignatures are generated by the issuer with respect to (pk_C, b, τ) without knowledge of id, while the identifier is chosen later by the client at token derivation time. Thus, tokens are not issued toward specific verifier identifiers, but are instantiated for a particular identifier only at presentation time.

This asymmetry is fundamental to, both, the functionality and security of double-spend identification. The role of the verifier id is to enable double-spend *detection* when the same token is used multiple times, while also ensuring that different presentations of the same token produce unlinkable views. Thus, in the real protocol, the issuer does not observe id during issuance, but may learn it later when tokens are presented (e.g., via verifier-side forwarding).

This point is crucial in order to correctly model unlinkability in the double-spend identification context—allowing the adversary (which controls the issuer in the unlinkability game) to freely choose

[4] 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 $\mathsf{UNLINK}_{\mathcal{A},n}(\lambda)$

1 : $\mathsf{pp}, \mathsf{aux} \leftarrow\$ \mathsf{Setup}(1^\lambda)$
2 : $(\mathsf{pk}_I, \mathsf{st}_1) \leftarrow \mathcal{A}_1(\mathsf{pp}, \mathsf{aux})$
3 : $(Q^*, \tau, \mathsf{id}, \mathsf{st}_2) \leftarrow \mathcal{A}_2^{\mathcal{O}_{\mathsf{GenUser}}, \mathcal{O}_{\mathsf{Obtain}}}(\mathsf{st}_1)$
4 : **foreach** $(\mathsf{pk}_C, \mathsf{psig}, \mathsf{nonce}) \in Q^*$
5 : $\quad$ **if** $(\mathsf{pk}_C, \cdot) \notin Q_{\mathsf{usr}}$
$\quad\ \vee (\mathsf{pk}_C, \cdot, \mathsf{nonce}) \in Q_{\mathsf{obt}}$ **then reject**
6 : **if** $|Q^*| \neq n$ **then reject**
7 : **foreach** $i \in [n]$ **do**
8 : $\quad \mathsf{tok}_i \leftarrow\$ \mathsf{Obtain}(\mathsf{sk}_{C_i}, \mathsf{pk}_I, \mathsf{psig}_i, \mathsf{nonce}_i, \tau, \mathsf{id})$
9 : $\quad$ **if** $\mathsf{tok}_i = \perp$ **then reject**
10 : $\hat{i} \leftarrow\$ [n]$
11 : Pick a random permutation $\varphi : [n] \to [n]$
12 : $i^* \leftarrow \mathcal{A}_3\left(\mathsf{st}_2, \mathsf{tok}_{\hat{i}}, \left(\mathsf{tok}_{\varphi(i)}\right)_{i \in [n] \setminus \{\varphi^{-1}(\hat{i})\}}\right)$
13 : **if** $i^* = \hat{i}$ **then accept**

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

1 : **if** Q_{usr} is uninitialized **then** $Q_{\mathsf{usr}} \coloneqq \varnothing$
2 : $(\mathsf{pk}_C, \mathsf{sk}_C) \leftarrow\$ \mathsf{KeyGen}_C(\mathsf{pp})$
3 : $Q_{\mathsf{usr}} \coloneqq Q_{\mathsf{usr}} \cup \{(\mathsf{pk}_C, \mathsf{sk}_C)\}$
4 : **return** pk_C

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

1 : **if** Q_{obt} is uninitialized **then** $Q_{\mathsf{obt}} \coloneqq \varnothing$
2 : **if** $(\mathsf{pk}_C, \cdot) \notin Q_{\mathsf{usr}}$ **then return** $\perp$
3 : Retrieve sk_C from Q_{usr}
4 : $\mathsf{tok} \leftarrow\$ \mathsf{Obtain}(\mathsf{sk}_C, \mathsf{pk}_I, \mathsf{psig}, \mathsf{nonce}, \tau, \mathsf{id})$
5 : **if** $\mathsf{tok} \neq \perp$ **then**
6 : $\quad Q_{\mathsf{obt}} \coloneqq Q_{\mathsf{obt}} \cup \{(\mathsf{pk}_C, \mathsf{psig}, \mathsf{nonce})\}$
7 : **return** tok

Figure 5: The unlinkability experiment. With slight abuse of notation, we interpret $(\mathsf{pk}_{C_i}, \mathsf{psig}_i, \mathsf{nonce}_i)$ as the ith entry in Q^* (line 8).

Game $\mathsf{PMB}_{\mathcal{A},\hat{b}}(\lambda)$

$\mathsf{pp}, \mathsf{aux} \leftarrow\!\$\ \mathsf{Setup}(1^\lambda)$
$(\mathsf{pk}_I, \mathsf{sk}_I, \mathsf{ek}_I) \leftarrow\!\$\ \mathsf{KeyGen}_I(\mathsf{pp})$
$\mathsf{flag} \coloneqq \mathsf{False}$
$b^* \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{Issue}}, \mathcal{O}_{\mathsf{Read}}, \mathcal{O}_{\mathsf{Valid}}, \mathcal{O}_{\mathsf{Challenge}}}(\mathsf{pp}, \mathsf{pk}_I)$
if $b^* = \hat{b}$ **then accept**

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

$(\mathsf{psig}, \mathsf{nonce}, \mathsf{aux}) \leftarrow\!\$\ \mathsf{Issue}(\mathsf{sk}_I, \mathsf{ek}_I, \mathsf{pk}_C, b, \tau, \mathsf{aux})$
return $(\mathsf{psig}, \mathsf{nonce})$

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

if $\mathsf{flag} \wedge (\tau = \hat{\tau})$ **then return** $\perp$
if $\mathsf{Verify}(\mathsf{pk}_I, t, \sigma, \tau) = 1$ **then**
  return $\mathsf{ReadBit}(\mathsf{ek}_I, t)$
return $\perp$

Oracle $\mathcal{O}_{\mathsf{Valid}}(t, \sigma, \tau)$

$b \leftarrow \mathsf{ReadBit}(\mathsf{ek}_I, t)$
return $(b \neq \perp \wedge \mathsf{Verify}(\mathsf{pk}_I, t, \sigma, \tau) = 1)$

Oracle $\mathcal{O}_{\mathsf{Challenge}}(\mathsf{pk}_C, \tau)$

if flag **then return** $\perp$
$\mathsf{flag} \coloneqq \mathsf{True}, \hat{\tau} \coloneqq \tau$
$(\mathsf{psig}, \mathsf{nonce}, \mathsf{aux}) \leftarrow\!\$\ \mathsf{Issue}(\mathsf{sk}_I, \mathsf{ek}_I, \mathsf{pk}_C, \hat{b}, \hat{\tau}, \mathsf{aux})$
return $(\mathsf{psig}, \mathsf{nonce})$

Figure 6: The privacy of metadata bit experiment.

all identifiers $\{\mathsf{id}_i\}$ for challenge tokens would not faithfully reflect the protocol as described above. Rather, it would enable artificial attacks where the adversary can break unlinkability simply by comparing the identifier corresponding to tok_{i^*} from the challenge and mapping it to the position of the identifier in its list $\{\mathsf{id}_i\}$ of chosen identifiers. Instead, in our unlinkability experiment (Definition 3.6) we fix a common identifier id across all challenge tokens. This sidesteps the aforementioned artificial attack while also capturing the desired privacy guarantees. Indeed, linking tokens corresponding to the same client across different identifiers would translate into an attack in our game by allowing the adversary to recover the hidden permutation in the challenge.

For NIAT with double-spend identification, we define two additional properties: double-spending identification and exculpability. Roughly, 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 3.9 (Double-spending identification)**.** Let $\mathsf{Adv}^{\mathsf{Iden}}_{\mathcal{A}} = \Pr[\mathsf{Iden}_{\mathcal{A}}(\lambda) = \textbf{accept}]$ be the advantage of an adversary $\mathcal{A}$ in the experiment defined in Figure 7. We say a NIAT scheme NIAT achieves double-spending identification if for any PPT adversary $\mathcal{A}$, this advantage is negligible.

Game $\mathsf{Iden}_{\mathcal{A}}(\lambda)$

$(\mathsf{pp}, \mathsf{aux}) \leftarrow\!\$\ \mathsf{Setup}(1^\lambda)$
$(\mathsf{pk}_I, \mathsf{sk}_I, \mathsf{ek}_I) \leftarrow\!\$\ \mathsf{KeyGen}_I(\mathsf{pp})$
$(t, \sigma, \sigma') \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{Issue}}, \mathcal{O}_{\mathsf{Read}}}(\mathsf{pp}, \mathsf{pk}_I)$
if $\sigma = \sigma' \vee \mathsf{Verify}(\mathsf{pk}_I, t, \sigma, \tau) \neq 1$
  $\vee\ \mathsf{Verify}(\mathsf{pk}_I, t, \sigma', \tau) \neq 1$ **then reject**
$(\mathsf{pk}_C, \cdot) \leftarrow \mathsf{DSIden}(\mathsf{pk}_I, t, \sigma, \sigma', \mathsf{aux})$
if $\mathsf{pk}_C \in Q \setminus \{\perp\}$ **then reject**
accept

Figure 7: The double-spending identification experiment for NIAT. Oracles $\mathcal{O}_{\mathsf{Issue}}$ and $\mathcal{O}_{\mathsf{Read}}$ are as defined in Figure 4, except $\mathcal{O}_{\mathsf{Issue}}$ also stores all pk_C queried by the adversary in a list Q.

Finally, exculpability ensures that an issuer cannot wrongly accuse an honest client of double-spending. Formally,

Definition 3.10 (Double-spending exculpability)**.** Let $\mathsf{Adv}^{\mathsf{Excul}}_{\mathcal{A}} = \Pr[\mathsf{Excul}_{\mathcal{A}}(\lambda) = \textbf{accept}]$ be the advantage of an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in the experiment defined in Figure 8. We say a NIAT scheme NIAT achieves double-spending exculpability if for any PPT adversary $\mathcal{A}$, this advantage is negligible.

Game $\mathsf{Excul}_{\mathcal{A}}(\lambda)$

$(\mathsf{pp}, \mathsf{aux}) \leftarrow\!\$\ \mathsf{Setup}(1^\lambda)$
$(\mathsf{pk}_I, \mathsf{st}_1) \leftarrow \mathcal{A}_1(\mathsf{pp})$
$(\mathsf{pk}_C, \Pi) \leftarrow \mathcal{A}_2^{\mathcal{O}_{\mathsf{GenUser}}, \mathcal{O}_{\mathsf{Obtain}}}(\mathsf{st}_1)$
if $(\mathsf{pk}_C, \cdot) \in Q_{\mathsf{usr}} \wedge \mathsf{DSVer}(\mathsf{pk}_C, \Pi, \mathsf{aux}) = 1$
  then accept

Figure 8: The double-spend exculpability experiment. Oracles $\mathcal{O}_{\mathsf{GenUser}}$ and $\mathcal{O}_{\mathsf{Obtain}}$ are as defined in Figure 5, except $\mathcal{O}_{\mathsf{Obtain}}$ behaves as an honest user that accepts a given $(\mathsf{psig}, \mathsf{nonce})$ pair exactly once.

4 NIAT from Equivalence Class Signatures

We now present our NIAT construction from structure-preserving signatures on equivalence classes (SPS-EQ). At a high level, we extend the NIBS construction of [21] 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 ignore the public metadata τ for simplicity. In the full version of our article [4] we show how to extend all our constructions to include τ in a manner similar to [21].

Construction 1. Our construction requires a hash function $\mathsf{H} : \{0,1\}^\lambda \to \mathbb{G}_1$ modeled as a random oracle, an SPS-EQ scheme $\mathsf{EQ} = (\mathsf{EQ.Setup}, \mathsf{EQ.KeyGen}, \mathsf{EQ.Sign}, \mathsf{EQ.ChRep}, \mathsf{EQ.Verify})$, and a NIZK $\mathsf{NIZK} = (\mathsf{NIZK.Setup}, \mathsf{NIZK.Prove}, \mathsf{NIZK.Verify})$ for the following language:

Language $\mathscr{L}_{\mathsf{iss}}$

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 $\mathsf{ek}_I := (x_1, x_2)$, $\mathsf{sk}_I := (y_1, y_2, y_3)$ and a bit b.
Membership: w is a valid witness for x if the following are satisfied:

$$b \in \{0,1\} \ \wedge \ S = R^{(1-b)x_1 + bx_2} \ \wedge$$
$$\mathsf{pk}_I = (g_1^{x_1}, g_1^{x_2}, g_2^{y_1}, g_2^{y_2}, g_2^{y_3})$$

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 $\mathscr{L}_{\mathsf{iss}}$.

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

Key generation. The issuer's key generation algorithm samples a random integer vector $\mathbf{x}$ and runs the SPS-EQ key generation algorithm. It sets its secret key as the SPS-EQ secret signing key, $\mathbf{x}$ as its secret extraction key, and $g_1^{\mathbf{x}}$ and $g_2^{\mathbf{y}}$ as its public key. The client's key generation algorithm samples a random value as its secret key, and sets the public key with respect to this secret.

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

$\mathsf{KeyGen}_C(\mathsf{pp})$
1: $\alpha \leftarrow\!\$\ \mathbb{Z}_p^*$
2: **return** $\mathsf{pk}_C := g_1^\alpha$, $\mathsf{sk}_C := \alpha$

Presignature issuance and token generation. We show the presignature issuance and token generation protocols in detail in Figure 9. To issue a presignature, the issuer creates an equivalence class signature on $(\mathsf{pk}_C, \mathsf{H}(r), S)$, where group element $S \in \mathbb{G}_1$ embeds the metadata bit b. The presignature psig includes the signature $\overline{\sigma}$ 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 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 $\overline{\sigma}$ to the equivalence class signature on $(g_1, \mathsf{H}(r)^{\alpha^{-1}}, S^{\alpha^{-1}})$. The final token is the tag $\mathbf{t} = \left(\mathsf{H}(r)^{\alpha^{-1}}, S^{\alpha^{-1}}\right)$ along with the transformed signature $\boldsymbol{\sigma}$.

Public verification (token redemption). The public verification algorithm consists of verifying the equivalence class signature $\boldsymbol{\sigma}$.

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

Bit extraction. The bit extraction proceeds by first calling the public verification algorithm. Then, the issuer simply checks which $b \in \{0,1\}$ satisfies $t_1^{x_{1+b}} = t_2$ (i.e., whether $t_1^{x_1}$ or $t_1^{x_2}$ equals t_2) and returns that bit if one is found, otherwise it returns an error.

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

Security. Below, we present our main theorems for NIAT security of Construction 1. Due to space constraints, we provide the proofs in the full version of our article [4] along with the proofs of correctness and reusability.

Theorem 4.1 (One-more unforgeability). *Construction 1 is one-more unforgeable (in the random oracle model) assuming* NIZK *satisfies zero knowledge and* EQ *is existentially unforgeable under adaptively chosen-message attacks.*

Theorem 4.2 (Public-validity Consistency). *Construction 1 is public-validity consistent (in the random oracle model) assuming* NIZK *satisfies zero knowledge and* EQ *is existentially unforgeable under adaptively chosen-message attacks.*

Theorem 4.3 (κ-unlinkability). *Construction 1 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.*

Remark 4.4 (Knowledge of secret keys)**.** For the proof of κ-unlinkability, we leverage the knowledge of secret keys (KOSK) model [8, 28] (also called the honest key model [21]). 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 5). A NIAT scheme that is secure in the KOSK model can be transformed into a fully maliciously secure scheme simply by having the issuer provide an one-time NIZK argument of knowledge (NIZKAoK) of the secret key sk_I as part of pk_I.

Theorem 4.5 (Privacy of metadata bit). *Construction 1 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.*

4.1 Instantiation

We use the SPS-EQ scheme given by [18] to instantiate our protocol from Figure 9. Due to space constraints, we provide the expanded token issuance and generation protocol in the full version of our article [4].

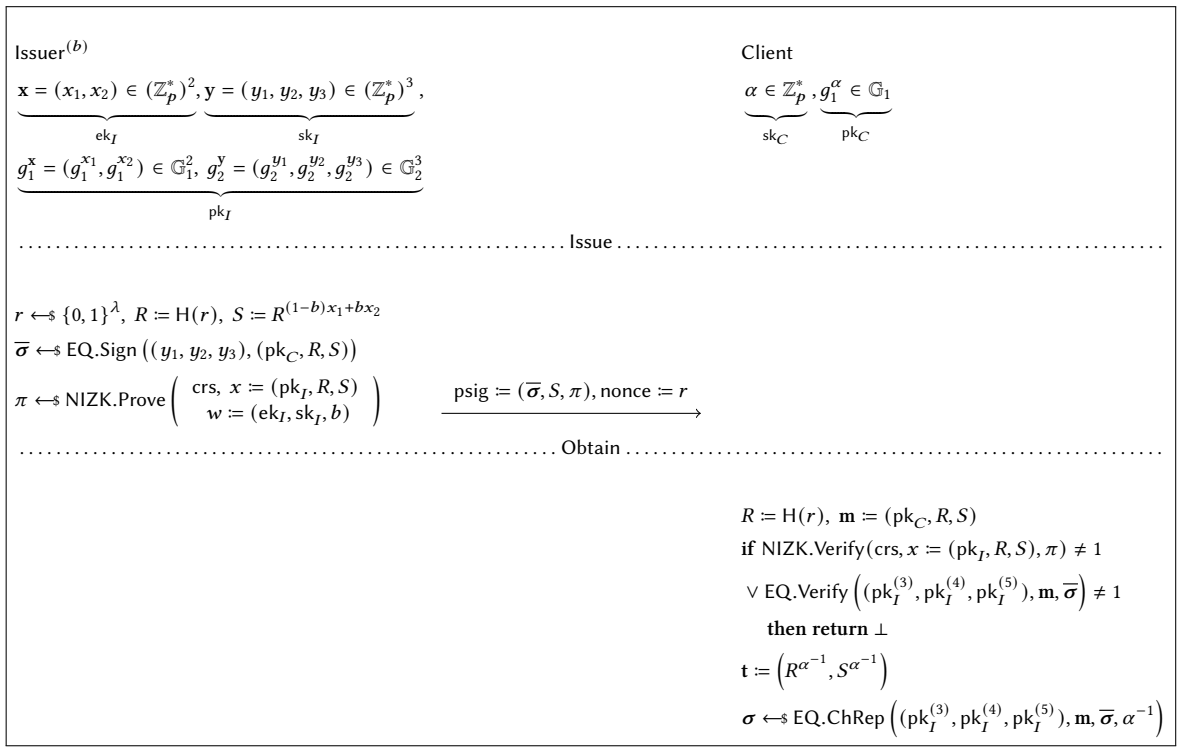

Figure 9: SPS-EQ construction for NIAT presignature issuance and token generation.

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 [13]. We detail the full ZK protocol for the language $\mathscr{L}_{\mathsf{iss}}$ in the full version of our article [4].

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 the full version of our article [4].

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 3.6.

We give a modification of Construction 1 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). To avoid repetition, we omit algorithms which are identical to (or trivially extendable from) their counterparts in Construction 1.

Construction 2. Our modified construction requires hash functions $\mathsf{H}_{\mathbb{G}} : \{0,1\}^{\lambda} \to \mathbb{G}_1$, $\mathsf{H}_{\mathbb{Z}} : \{0,1\}^{\lambda} \to \mathbb{Z}_p$, $\mathsf{H}_{\mathbb{G}\to\mathbb{G}} : \mathbb{G}_1 \to \mathbb{G}_1$, an SPS-EQ scheme $\mathsf{EQ} = (\mathsf{EQ.Setup}, \mathsf{EQ.Sign}, \mathsf{EQ.ChRep}, \mathsf{EQ.Verify})$, $\mathsf{NIZK}_{\mathsf{iss}} = (\mathsf{NIZK}_{\mathsf{iss}}.\mathsf{Setup}, \mathsf{NIZK}_{\mathsf{iss}}.\mathsf{Prove}, \mathsf{NIZK}_{\mathsf{iss}}.\mathsf{Verify})$ and $\mathsf{NIZK}_{\mathsf{obt}} = (\mathsf{NIZK}_{\mathsf{obt}}.\mathsf{Setup}, \mathsf{NIZK}_{\mathsf{obt}}.\mathsf{Prove}, \mathsf{NIZK}_{\mathsf{obt}}.\mathsf{Verify})$ where the former is a NIZK for $\mathscr{L}_{\mathsf{iss}}$ defined previously, and the latter is a NIZK for the following language:

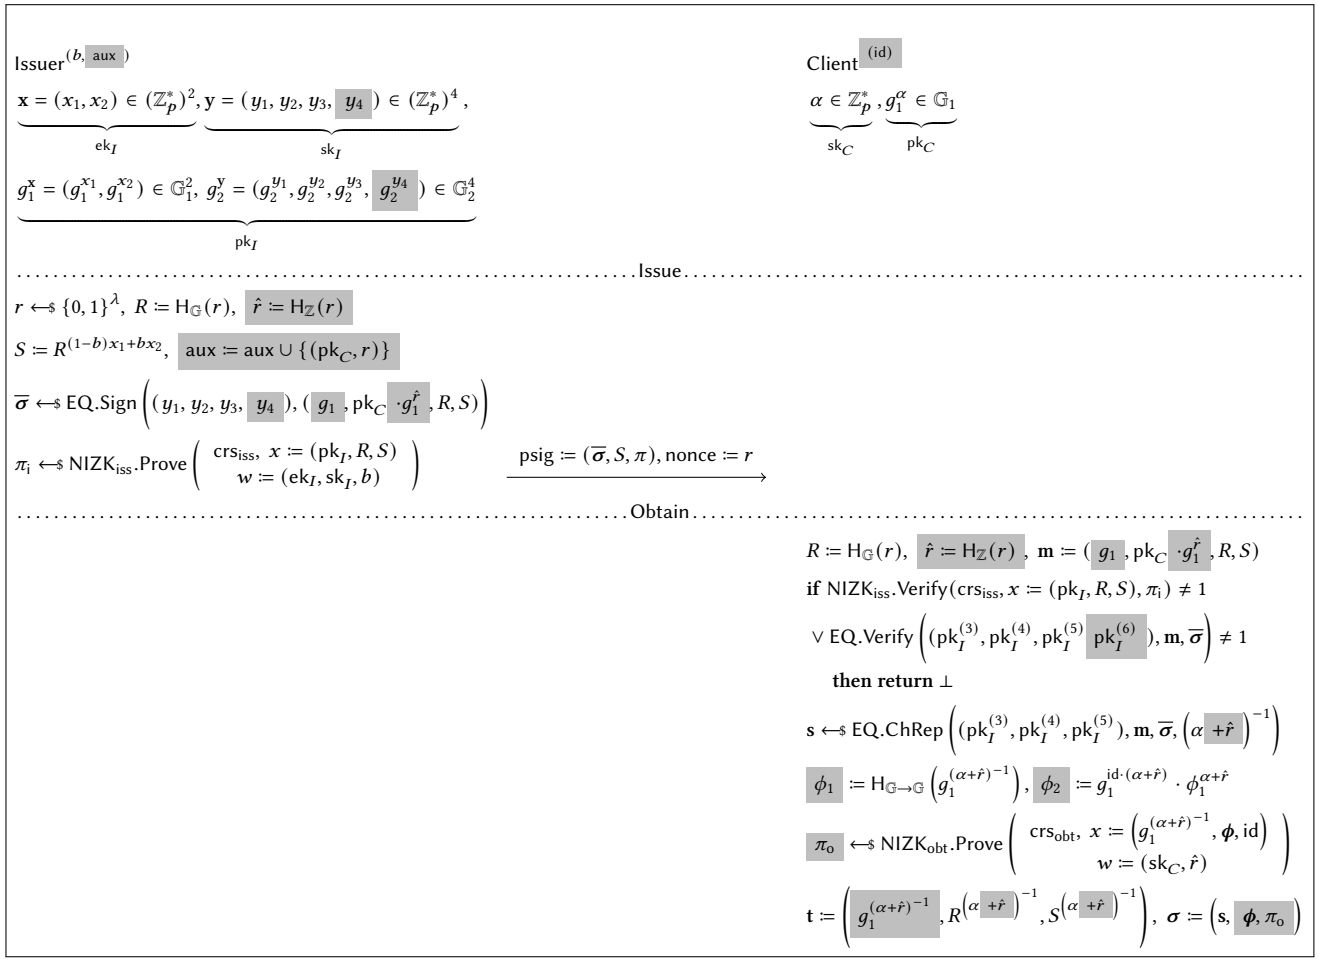

Figure 10: Modified SPS-EQ construction for NIAT presignature issuance and token generation.

Language $\mathscr{L}_{\text{obt}}$

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:

$$\phi_2 = g_1^{\mathsf{id}\cdot\left(\mathsf{sk}_C+\hat{r}\right)} \cdot \phi_1^{\mathsf{sk}_C+\hat{r}} \ \wedge\ t_1 = g_1^{\left(\mathsf{sk}_C+\hat{r}\right)^{-1}}$$

Presignature issuance and token generation. In Figure 10, 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 $\mathsf{H}_{\mathbb{Z}}(r)$ and the signature, which is now computed over $\mathbf{m} \coloneqq (g_1, \mathsf{pk}_C \cdot g_1^{\hat{r}}, R, S)$. 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 $\mathbf{t} \coloneqq \left(g_1^{(\alpha+\hat{r})^{-1}}, R^{(\alpha+\hat{r})^{-1}}, S^{(\alpha+\hat{r})^{-1}}\right)$.

Now, if a client intends to redeem its token with some verifier with public identifier id, as we show in Figure 10, the client's Obtain algorithm must additionally compute an ElGamal ciphertext $\boldsymbol{\phi}$ over id along with a proof π_{o} of the well-formedness of the ciphertext (see language $\mathscr{L}_{\text{obt}}$). Therefore, the final signature is the equivalence class signature $\mathbf{s}$, the ciphertext $\boldsymbol{\phi}$ and the proof π_{o}.

Public verification. The public verification algorithm verifies the equivalence class signature $\mathbf{s}$ and the ZK proof π_{o}.

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

1: **if** $\mathsf{Verify}(\mathsf{pk}_I, \mathbf{t}, \boldsymbol{\sigma}) \neq 1 \ \vee \ \mathsf{Verify}(\mathsf{pk}_I, \mathbf{t}, \boldsymbol{\sigma}') \neq 1$ $\vee\ \boldsymbol{\sigma} = \boldsymbol{\sigma}'$ **then return** $(\perp, \perp)$
2: $\boldsymbol{\sigma} \coloneqq (\mathbf{s}, \boldsymbol{\phi}, \pi_o, \mathsf{id})$, $\boldsymbol{\sigma}' \coloneqq (\mathbf{s}', \boldsymbol{\phi}', \pi_o', \mathsf{id}')$
3: $h \coloneqq \left(\phi_2^{-1} \cdot \phi_2'\right)^{(\mathsf{id}'-\mathsf{id})^{-1}}$
4: **foreach** $(\mathsf{pk}_C, \mathsf{nonce}) \in \mathsf{aux}$ **do**
5: **if** $h = \mathsf{pk}_C \cdot g_1^{\mathsf{H}_{\mathbb{Z}}(\mathsf{nonce})}$ **then**
6: **return** $\mathsf{pk}_C, \Pi \coloneqq (\mathbf{t}, \boldsymbol{\sigma}, \boldsymbol{\sigma}', \mathsf{nonce})$
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 $\boldsymbol{\phi}'$ over id' and the corresponding proof π_o' with the same $\mathbf{s}$ and $\mathbf{t}$. When a system detects that a double-spend of $(\mathbf{t}, \mathbf{s})$ has occurred, it can compute $\left(\phi_2^{-1} \cdot \phi_2'\right)^{(\mathsf{id}'-\mathsf{id})^{-1}} = g_1^{\alpha+\hat{r}}$ 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-spending.[5] A simple optimization that offers significant asymptotic advantage over the linear search on aux is to instead store the value $\mathsf{pk}_C \cdot g_1^{\mathsf{H}_{\mathbb{Z}}(\mathsf{nonce})}$ 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.

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

1: **if** $\mathsf{Verify}(\mathsf{pk}_I, \mathbf{t}, \boldsymbol{\sigma}) \wedge \mathsf{Verify}(\mathsf{pk}_I, \mathbf{t}, \boldsymbol{\sigma}') \neq 1$ **then**
2: **return** 0
3: $\boldsymbol{\sigma} \coloneqq (\mathbf{s}, \boldsymbol{\phi}, \pi_o, \mathsf{id})$, $\boldsymbol{\sigma}' \coloneqq (\mathbf{s}', \boldsymbol{\phi}', \pi_o', \mathsf{id}')$
4: **return** $\left(\phi_2^{-1} \cdot \phi_2'\right)^{(\mathsf{id}-\mathsf{id}')^{-1}} = \mathsf{pk}_C \cdot g_1^{\mathsf{H}_{\mathbb{Z}}(\mathsf{nonce})}$ $\wedge\ (\mathsf{pk}_C, \mathsf{nonce}) \in \mathsf{aux}$

Security. Due to space constraints, we present the theorems and proofs for NIAT security of Construction 2 in the full version of our article [4].

6 Implementation and Evaluation

We implemented our main NIAT Construction 1 in C++ using the `mcl` library [29].[6] For our implementation, we chose the pairing-friendly BLS12-381 curve [5, 10]. 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.

[5] One could 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.
[6] Source code : https://github.com/aayux/mcl.

Evaluation goals. Our evaluation focuses on understanding the practical impact of non-interactive issuance in high-volume anonymous token systems. In particular, we focus on three salient aspects: (i) the concrete cost of token issuance, generation, and redemption in our SPS-EQ-based NIAT construction; (ii) the extent to which pairing aggregation amortizes verification costs; and (iii) the operational trade-offs enabled by asynchronous presignature pre-computation. We stress that our implementation is a proof-of-concept prototype intended to evaluate cryptographic costs rather than a full production deployment and thus we evaluate only the core cryptographic operations on a single machine. While we provide some preliminary estimates for them, a full analysis of latency, memory consumption and optimized multi-server deployments are left as future, systems-oriented, work.

Intended deployment scenario. We envision NIAT as being deployed in a large-scale token issuance application such as a CDN attestation service, where the issuer serves millions of issuance requests per second and minimizing latency is critical. As we detailed in Section 1.1, the offline pre-computation enabled by NIAT is highly suitable for this setting.

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 11 (left) for batch size $n = 30, 60$ and 100. Concretely, under our scheme a protocol would require 1-2 TB of storage for 100 million clients, which is quite nominal for modern servers. 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.

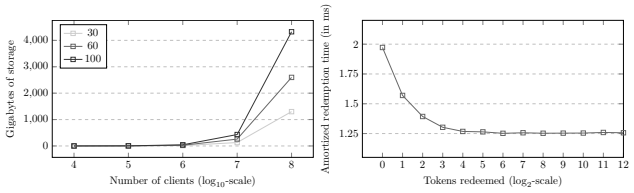

Figure 11: Storage estimates for an offline signing server storing batches of 30, 60 **and** 100 **presignatures per client (left); and Issuer's amortized redemption cost per token (right). Roughly, our protocol requires** 13.6 **KB to store** 30 **tokens per client.**

	# Moves	Public Verification	Private Metadata Bit	Issue	Obtain	Readbit	Token size $\lvert t\rvert + \lvert\sigma\rvert$
Const. 1*	1	✓	✓	16×	1e + 19×	3e + 1×	$5\mathbb{G}$
[6]	3	✓	✓	7×	11×	6×	$3\mathbb{G} + 4\mathbb{Z}$
[11]	2	✗	✓	11×	17×	1×	$2\mathbb{G} + 1\mathbb{Z}$
[26]	2	✗	✓	12×	15×	4×	$2\mathbb{G} + 1\mathbb{Z}$
[14]	2	✗	✗	7×	2×	1×	$1\mathbb{G}$
[33]	2	✓	✗	3×	2e + 1×	2e	$1\mathbb{G} + 1\mathbb{Z}$

Table 1: Comparison of anonymous token schemes. × represents group exponentiation and e represents pairing operation. $\mathbb{G}$ and $\mathbb{Z}$ stand for group elements and integers respectively. ∗ indicates the verification cost is amortized.

Computational costs. 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 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 $\mathsf{e}(\mathsf{pk}_C, \mathsf{pk}_I^{(3)})$) 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 $\mathsf{e}(g_1, \mathsf{pk}_I^{(3)})$) to verify the SPS-EQ signature, and then extracts the embedded bit with 1 exponentiation.

Comparison and trade-offs. Table 1 summarizes our results (after amortization) and presents them as part of a comparison with related works using standard operation counting. Our focus in this work is on structural improvement for the existing token issuance protocols rather than concrete computational gains of individual algorithms. While a full implementation and comparison of prior schemes would be valuable from a measurement standpoint, we believe that Table 1 sufficiently isolates the cryptographic trade-offs that differentiate the schemes for our purpose.

Experimental benchmarks. From our experiments we found that an exponentiation in $\mathbb{G}_1$ took approximately 0.047 ms, an exponentiation in $\mathbb{G}_2$ 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.

Offline pre-computation. The core premise of NIAT is that presignature generation can be decoupled from client request latency and executed opportunistically during periods of low load or on dedicated infrastructure. Based on the issuance times (Table 2), we estimate that a laptop computer can pre-compute slightly over 30 million tokens overnight at 1064 clients per second, thus serving more than a million clients without any latency. For comparison, an interactive system like Privacy Pass also handles about 1000 clients per second at peak (even without a private bit) but with a 32-core server [14]. Indeed, the throughput of our scheme will be significantly higher for a commercial server computer, virtually removing the online signing bottleneck during peak hours. This highlights an enormous advantage gained by our non-interactive system over interactive alternatives.

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

Table 2: Benchmarks for our implementation of Construction 1.

Batched verification. We also ran benchmark tests for a batched version of signature verification the details of which can be found in our full version [4]. 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 presignatures based on experimental evaluations averaged over 1000 runs[7]. 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 11 (right) we show the actual amortized costs over up to 2^{12} token redemptions.

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

Estimates for double-spending protection. We provide estimates of the *additional* costs incurred by the double-spending extension over the base scheme. 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 $\mathbb{G}_1$ (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.

Acknowledgments

We thank the reviewers and our shepherd (the 'anonymous Roman') for all the valuable comments and discussions that helped to improve the presentation and the clarity of our definitions. Baldimtsi, Yadav and Nguyen were supported by NSF #2143287 and #2247304. Hanzlik was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) #545066907.

References

[1] Ghous Amjad, Kevin Yeo, and Moti Yung. 2023. RSA Blind Signatures with Public Metadata. Cryptology ePrint Archive, Report 2023/1199. https://eprint.iacr.org/2023/1199

[2] Apple. 2021. iCloud Private Relay Overview. https://www.apple.com/icloud/docs/iCloud_Private_Relay_Overview_Dec2021.pdf. Accessed: 2024-04-29.

[3] Foteini Baldimtsi, Jiaqi Cheng, Rishab Goyal, and Aayush Yadav. 2024. Non-Interactive Blind Signatures: Post-Quantum and Stronger Security. In *Advances in Cryptology – ASIACRYPT 2024, Part II (Lecture Notes in Computer Science, Vol. 15485)*, Kai-Min Chung and Yu Sasaki (Eds.). Springer, Singapore, Singapore, Kolkata, India, 70–104. doi:10.1007/978-981-96-0888-1_3

[4] Foteini Baldimtsi, Lucjan Hanzlik, Quan Nguyen, and Aayush Yadav. 2025. Non-interactive Anonymous Tokens with Private Metadata Bit. Cryptology ePrint Archive, Report 2025/430. https://eprint.iacr.org/2025/430

[5] Paulo S. L. M. Barreto, Ben Lynn, and Michael Scott. 2003. Constructing Elliptic Curves with Prescribed Embedding Degrees. In *SCN 02: 3rd International Conference on Security in Communication Networks (Lecture Notes in Computer Science, Vol. 2576)*, Stelvio Cimato, Clemente Galdi, and Giuseppe Persiano (Eds.). Springer Berlin Heidelberg, Germany, Amalfi, Italy, 257–267. doi:10.1007/3-540-36413-7_19

[6] Fabrice Benhamouda, Tancrède Lepoint, Michele Orrù, and Mariana Raykova. 2022. Publicly verifiable anonymous tokens with private metadata bit. Cryptology ePrint Archive, Report 2022/004. https://eprint.iacr.org/2022/004

[7] Fabrice Benhamouda, Mariana Raykova, and Karn Seth. 2023. Anonymous Counting Tokens. In *Advances in Cryptology – ASIACRYPT 2023, Part II (Lecture Notes in Computer Science, Vol. 14439)*, Jian Guo and Ron Steinfeld (Eds.). Springer, Singapore, Singapore, Guangzhou, China, 245–278. doi:10.1007/978-981-99-8724-5_8

[8] Alexandra Boldyreva. 2003. 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 (Lecture Notes in Computer Science, Vol. 2567)*, Yvo Desmedt (Ed.). Springer Berlin Heidelberg, Germany, Miami, FL, USA, 31–46. doi:10.1007/3-540-36288-6_3

[9] Dan Boneh, Hart William Montgomery, and Ananth Raghunathan. 2010. Algebraic pseudorandom functions with improved efficiency from the augmented cascade. In *ACM CCS 2010: 17th Conference on Computer and Communications Security*, Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov (Eds.). ACM Press, Chicago, Illinois, USA, 131–140. doi:10.1145/1866307.1866323

[10] Sean Bowe. 2017. BLS12-381: New zk-SNARK Elliptic Curve Construction. https://electriccoin.co/blog/new-snark-curve/. Accessed: 2024-04-17.

[11] Melissa Chase, F. Betül Durak, and Serge Vaudenay. 2023. Anonymous Tokens with Stronger Metadata Bit Hiding from Algebraic MACs. In *Advances in Cryptology – CRYPTO 2023, Part II (Lecture Notes in Computer Science, Vol. 14082)*, Helena Handschuh and Anna Lysyanskaya (Eds.). Springer, Cham, Switzerland, Santa Barbara, CA, USA, 418–449. doi:10.1007/978-3-031-38545-2_14

[12] David Chaum. 1983. Blind Signature System. In *Advances in Cryptology – CRYPTO'83*, David Chaum (Ed.). Plenum Press, New York, USA, Santa Barbara, CA, USA, 153. doi:10.1007/978-1-4684-4730-9_14

[13] Ivan Damgård. 2010. On $\sum$-protocols. https://www.cs.au.dk/ ivan/Sigma.pdf. Accessed: 2025-03-20.

[14] Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, and Filippo Valsorda. 2018. Privacy Pass: Bypassing Internet Challenges Anonymously. *Proceedings on Privacy Enhancing Technologies* 2018, 3 (July 2018), 164–180. doi:10.1515/popets-2018-0026

[15] Alex Davidson, Jana Iyengar, and Christopher A. Wood. 2024. The Privacy Pass Architecture. RFC 9576. doi:10.17487/RFC9576

[16] Yevgeniy Dodis and Aleksandr Yampolskiy. 2005. A Verifiable Random Function with Short Proofs and Keys. In *PKC 2005: 8th International Workshop on Theory and Practice in Public Key Cryptography (Lecture Notes in Computer Science, Vol. 3386)*, Serge Vaudenay (Ed.). Springer Berlin Heidelberg, Germany, Les Diablerets, Switzerland, 416–431. doi:10.1007/978-3-540-30580-4_28

[17] Amos Fiat and Adi Shamir. 1987. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In *Advances in Cryptology – CRYPTO'86 (Lecture Notes in Computer Science, Vol. 263)*, Andrew M. Odlyzko (Ed.). Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA, 186–194. doi:10.1007/3-540-47721-7_12

[18] Georg Fuchsbauer, Christian Hanser, and Daniel Slamanig. 2019. Structure-Preserving Signatures on Equivalence Classes and Constant-Size Anonymous Credentials. *Journal of Cryptology* 32, 2 (April 2019), 498–546. doi:10.1007/s00145-018-9281-4

[19] Stephen D. Galbraith, Kenneth G. Paterson, and Nigel P. Smart. 2006. Pairings for Cryptographers. Cryptology ePrint Archive, Report 2006/165. https://eprint.iacr.org/2006/165

[20] Google. 2024. Private State Tokens. https://developers.google.com/privacy-sandbox/protections/private-state-tokens. Accessed: 2024-04-29.

[21] Lucjan Hanzlik. 2023. Non-interactive Blind Signatures for Random Messages. In *Advances in Cryptology – EUROCRYPT 2023, Part V (Lecture Notes in Computer Science, Vol. 14008)*, Carmit Hazay and Martijn Stam (Eds.). Springer, Cham, Switzerland, Lyon, France, 722–752. doi:10.1007/978-3-031-30589-4_25

[22] Lucjan Hanzlik, Eugenio Paracucchi, and Riccardo Zanotto. 2025. 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 (Lecture Notes in Computer Science, Vol. 15602)*, Serge Fehr and Pierre-Alain Fouque (Eds.). Springer, 365–394. doi:10.1007/978-3-031-91124-8_13

[23] Lucjan Hanzlik and Daniel Slamanig. 2021. With a Little Help from My Friends: Constructing Practical Anonymous Credentials. In *ACM CCS 2021: 28th Conference on Computer and Communications Security*, Giovanni Vigna and Elaine Shi (Eds.). ACM Press, Virtual Event, Republic of Korea, 2004–2023. doi:10.1145/3460120.3484582

[24] Shumon Huque, Christian Elmerot, and Olafur Gudmundsson. 2025. Compact Denial of Existence in DNSSEC. RFC 9824. doi:10.17487/RFC9824

[25] IETF. 2024. Privacy Pass (privacypass). https://datatracker.ietf.org/wg/privacypass/about/. Accessed: 2024-04-29.

[26] Ben Kreuter, Tancrède Lepoint, Michele Orrù, and Mariana Raykova. 2020. Anonymous Tokens with Private Metadata Bit. In *Advances in Cryptology – CRYPTO 2020, Part I (Lecture Notes in Computer Science, Vol. 12170)*, Daniele Micciancio and Thomas Ristenpart (Eds.). Springer, Cham, Switzerland, Santa Barbara, CA, USA, 308–336. doi:10.1007/978-3-030-56784-2_11

[27] Thibault Meunier, Cefan Daniel Rubin, and Armando Faz-Hernández. 2024. Privacy Pass: upgrading to the latest protocol version. https://blog.cloudflare.com/privacy-pass-standard. Accessed: 2024-04-29.

[28] Silvio Micali, Kazuo Ohta, and Leonid Reyzin. 2001. Accountable-Subgroup Multisignatures: Extended Abstract. In *ACM CCS 2001: 8th Conference on Computer and Communications Security*, Michael K. Reiter and Pierangela Samarati (Eds.). ACM Press, Philadelphia, PA, USA, 245–254. doi:10.1145/501983.502017

[29] Shigeo Mitsunari. 2024. MCL. https://github.com/herumi/mcl.

[30] Michele Orrù, Stefano Tessaro, Greg Zaverucha, and Chenzhi Zhu. 2024. Oblivious Issuance of Proofs. In *Advances in Cryptology – CRYPTO 2024, Part IX (Lecture Notes in Computer Science, Vol. 14928)*, Leonid Reyzin and Douglas Stebila (Eds.). Springer, Cham, Switzerland, Santa Barbara, CA, USA, 254–287. doi:10.1007/978-3-031-68400-5_8

[31] David Pointcheval and Jacques Stern. 1996. Provably Secure Blind Signature Schemes. In *Advances in Cryptology – ASIACRYPT'96 (Lecture Notes in Computer Science, Vol. 1163)*, Kwangjo Kim and Tsutomu Matsumoto (Eds.). Springer Berlin Heidelberg, Germany, Kyongju, Korea, 252–265. doi:10.1007/BFb0034852

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

$\mathsf{pp}, \mathsf{aux} \leftarrow\!\$\ \mathsf{Setup}(1^\lambda)$
$(\mathsf{pk}_I, \mathsf{sk}_I, \mathsf{ek}_I) \leftarrow\!\$\ \mathsf{KeyGen}_I(\mathsf{pp})$
$\tau^*, \{(t_i, \sigma_i)\}_{i\in[N]} \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{Issue}}, \mathcal{O}_{\mathsf{Read}}}(\mathsf{pp}, \mathsf{pk}_I)$
// Type 1: token forgery
if $\#\{t_i\}_{i\in[N]} > \ell \wedge \bigwedge_{i\in[N]} \mathsf{Verify}(\mathsf{pk}_I, t_i, \sigma_i, \tau^*) = 1$ **then accept**
$\mathsf{ctr}^*_{0,\tau^*} \coloneqq 0, \mathsf{ctr}^*_{1,\tau^*} \coloneqq 0$
foreach $i \in [N]$ **do**
  if $\mathsf{Verify}(\mathsf{pk}_I, t_i, \sigma_i, \tau^*) = 1$ **then**
    $b' \leftarrow \mathsf{ReadBit}(\mathsf{ek}_I, t_i)$
  // Type 2: public-validity consistency
    if $b' = \perp$ **then accept**
    $\mathsf{ctr}^*_{b',\tau^*} \leftarrow \mathsf{ctr}^*_{b',\tau^*} + 1$
  // Type 3: metadata bit integrity
if $\mathsf{ctr}^*_{0,\tau^*} > \mathsf{ctr}_{0,\tau^*} \vee \mathsf{ctr}^*_{1,\tau^*} > \mathsf{ctr}_{1,\tau^*}$ **then accept**
reject

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

if $\ell, \mathsf{ctr}_{0,\tau}, \mathsf{ctr}_{1,\tau}$ are uninitialized **then** $\ell \coloneqq 0, \mathsf{ctr}_{0,\tau} \coloneqq 0, \mathsf{ctr}_{1,\tau} \coloneqq 0$
$(\mathsf{psig}, \mathsf{nonce}, \mathsf{aux}) \leftarrow\!\$\ \mathsf{Issue}(\mathsf{sk}_I, \mathsf{ek}_I, \mathsf{pk}_C, b, \tau, \mathsf{aux})$
$\ell \leftarrow \ell + 1$
$\mathsf{ctr}_{b,\tau} \leftarrow \mathsf{ctr}_{b,\tau} + 1$
return $(\mathsf{psig}, \mathsf{nonce})$

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

return $\mathsf{ReadBit}(\mathsf{ek}_I, t)$

Figure 12: The one-more unforgeability experiment for NIAT.

[32] David Pointcheval and Jacques Stern. 2000. Security Arguments for Digital Signatures and Blind Signatures. *Journal of Cryptology* 13, 3 (June 2000), 361–396. doi:10.1007/s001450010003

[33] Tjerand Silde and Martin Strand. 2022. Anonymous Tokens with Public Metadata and Applications to Private Contact Tracing. In *FC 2022: 26th International Conference on Financial Cryptography and Data Security (Lecture Notes in Computer Science, Vol. 13411)*, Ittay Eyal and Juan A. Garay (Eds.). Springer, Cham, Switzerland, Grenada, 179–199. doi:10.1007/978-3-031-18283-9_9

[34] Nirvan Tyagi, Sofia Celi, Thomas Ristenpart, Nick Sullivan, Stefano Tessaro, and Christopher A. Wood. 2022. A Fast and Simple Partially Oblivious PRF, with Applications. In *Advances in Cryptology – EUROCRYPT 2022, Part II (Lecture Notes in Computer Science, Vol. 13276)*, Orr Dunkelman and Stefan Dziembowski (Eds.). Springer, Cham, Switzerland, Trondheim, Norway, 674–705. doi:10.1007/978-3-031-07085-3_23

[35] John Wilander. 2021. Introducing Private Click Measurement, PCM. https://webkit.org/blog/11529/introducing-private-click-measurement-pcm/. Accessed: 2024-04-29.

A A Note on Public Validity Consistency

As we suggested in Section 3, the public validity consistency property can be cleanly folded into the one-more unforgeability definition. In this section, we provide the alternate definition which we believe to be equivalent to Definitions 3.3 and 3.4 taken together. As above, our adversary is given access to an issuance oracle $\mathcal{O}_{\mathsf{Issue}}$ that allows it to request pre-signatures (and the corresponding nonces) for any client public key pk_C, metadata bit $b \in \{0, 1\}$, and public metadata τ of its choice. It is also given access to an oracle $\mathcal{O}_{\mathsf{ReadBit}}$ for the bit extraction algorithm. We now consider three types of NIAT forgeries explained as follows:

- **Type 1 (Token forgery).** The adversary breaks unforgeability if, after receiving ℓ pre-signatures from $\mathcal{O}_{\mathsf{Issue}}$, it outputs more than ℓ valid tokens (t, σ) under the same public metadata τ^*. This corresponds to the standard "one-more" forgery notion for blind-signature-based constructions.
- **Type 2 (Public-validity consistency).** The adversary also breaks unforgeability if it outputs a token (t, σ) such that $\mathsf{Verify}(\mathsf{pk}_I, t, \sigma, \tau^*) = 1$ but $\mathsf{ReadBit}(\mathsf{ek}_I, t) = \perp$ on any public metadata τ^* of its choice. This guarantees, that public verification and private bit extraction define compatible notions of validity, i.e., any token accepted by Verify must have a well-defined embedded metadata bit. This property is not needed in [11], but is important in publicly verifiable settings, where verification may be performed by parties that do not hold the extraction key.
- **Type 3 (Metadata bit integrity).** Finally, the adversary breaks unforgeability if it violates the integrity of the embedded metadata bit. Concretely, let ctr_0 and ctr_1 denote the number of issuance queries made by the adversary for bits 0 and 1, respectively, under public metadata τ^*. The adversary wins if it outputs a set of valid tokens for which the number of tokens extracting to some bit $b \in \{0, 1\}$ exceeds ctr_b. In other words, the adversary cannot produce more valid tokens carrying bit b than the number of times it queried that specific bit during issuance.

Definition A.1 (One-more unforgeability). Let $\mathsf{Adv}^{\text{OM-Unf}}_{\mathcal{A}} = \Pr[\text{OM-Unf}_{\mathcal{A}}(\lambda) = \textbf{accept}]$ be the advantage of an adversary $\mathcal{A}$ in the experiment defined in Figure 12. We say a NIAT scheme NIAT is one-more unforgeable if for any PPT adversary $\mathcal{A}$ this advantage is negligible.

B Security Proofs for Section 4

Theorem 4.1 (One-more unforgeability). Construction 1 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 OM-Unf$_{\mathcal{A}}$ (see Definition 3.3) defined with respect to Construction 1.
- 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 $q_{\mathsf{H}} = \mathrm{poly}(\lambda)$ 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 $q_{\mathsf{H}}/2^\lambda$.
- 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 $\mathsf{Adv}^{j}_{\mathcal{A}}$ be the advantage of a PPT adversary $\mathcal{A}$ in hybrid $\mathbf{H}_j$. Then, we claim that $\mathsf{Adv}^{2}_{\mathcal{A}}$ 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 existentially unforgeability (under adaptively chosen-message attack) of the underlying SPS-EQ scheme EQ. In particular, the reduction does the following:

- It sets $(\mathsf{pk}_I^{(3)}, \mathsf{pk}_I^{(4)}, \mathsf{pk}_I^{(4)}) \coloneqq \mathsf{pk}_{\mathsf{EQ}}$ 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 $\overline{\sigma}$ on $(\mathsf{pk}_C, \mathsf{H}(r), S)$ and returns $\mathsf{psig} \coloneqq (\overline{\sigma}, S, \pi)$ and $\mathsf{nonce} \coloneqq r$, where π is the simulated NIZK proof.
- On receiving a Read query $(\mathbf{t}, \boldsymbol{\sigma})$ it first checks if $\mathsf{Verify}(\mathsf{pk}_I, \mathbf{t}, \boldsymbol{\sigma}) = 1$ 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 $\{(\mathbf{t}^{(i)}, \boldsymbol{\sigma}^{(i)})\}_{i\in[N]}$, $\mathcal{B}$ uniformly chooses a pair $(\mathbf{t}^*, \boldsymbol{\sigma}^*)$ and outputs it as its forgery.

For a successful adversary $\mathcal{A}$, each $\boldsymbol{\sigma}^{(i)}$ must be a signature on a unique class represented by $(g_1, t_1^{(i)}, t_2^{(i)})$.

To see this, first observe that the messages signed by the SPS-EQ signature scheme fall into equivalence classes determined by the discrete-logarithm relations among their components. Consider a message of the form (g_1, A, B). The discrete logarithm of B to base A partitions the message space according to the value x satisfying $B = A^x$. In our setting, there are only two possible values, namely x_1 and x_2. Thus, the relation between A and B, which encodes the metadata bit, induces two distinct sets of messages.

It follows that every query to the Issue oracle with bit $(1 - b^*)$ produces an SPS-EQ signature on an equivalence class different from that of $(g_1, t_1^{(i)}, t_2^{(i)})$. Moreover, within each of these two sets, the discrete logarithm of A to base g_1 further refines the partition.

It remains to argue that the class of $(g_1, t_1^{(i)}, t_2^{(i)})$ is unique among the messages signed in Issue queries with bit b^*. By definition, the N tags under consideration are new and distinct. Hence, for every $j \in [N]$ with $j \neq i$, we have $t_1^{(j)} \neq t_1^{(i)}$. Furthermore, once $t_1^{(j)}$ and the bit b^* are fixed, the third component is uniquely determined as $t_2^{(j)} = \left(t_1^{(j)}\right)^{x_{b^*}}$. This follows, since we know that for all tokens output by the adversary we have $\mathsf{ReadBit}(\mathsf{ek}_I, t_i) = b^*$. Therefore, no other signed message with bit b^* lies in the same SPS-EQ equivalence class as $(g_1, t_1^{(i)}, t_2^{(i)})$. Consequently, this equivalence class is unique among all messages signed by Issue.

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, $\mathsf{Adv}^{2}_{\mathcal{A}} = \frac{\ell}{N-\ell} \cdot \mathsf{Adv}^{\mathsf{EUnf\text{-}CMA}}_{\mathcal{B}}$, which is negligible for any PPT adversary $\mathcal{A}$. □

Theorem 4.2 (Public-validity Consistency). Construction 1 is public-validity consistent (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 public-validity consistency experiment $\mathsf{PVC}_{\mathcal{A}}$ (see Definition 3.4) defined with respect to Construction 1.
- 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 $q_{\mathsf{H}} = \mathsf{poly}(\lambda)$ 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 $q_{\mathsf{H}}/2^{\lambda}$.
- 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 $\mathsf{Adv}^{j}_{\mathcal{A}}$ be the advantage of a PPT adversary $\mathcal{A}$ in hybrid $\mathbf{H}_j$. Then, we claim that $\mathsf{Adv}^{2}_{\mathcal{A}}$ 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 existentially unforgeability (under adaptively chosen-message attack) of the underlying SPS-EQ scheme EQ. In particular, the reduction does the following:

- It sets $(\mathsf{pk}_I^{(3)}, \mathsf{pk}_I^{(4)}, \mathsf{pk}_I^{(4)}) \coloneqq \mathsf{pk}_{\mathsf{EQ}}$ 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 $\overline{\sigma}$ on $(\mathsf{pk}_C, \mathsf{H}(r), S)$ and returns $\mathsf{psig} \coloneqq (\overline{\sigma}, S, \pi)$ and $\mathsf{nonce} \coloneqq r$, where π is the simulated NIZK proof.
- On receiving a Read query $(\mathbf{t}, \boldsymbol{\sigma})$ it first checks if $\mathsf{Verify}(\mathsf{pk}_I, \mathbf{t}, \boldsymbol{\sigma}) = 1$ 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 $(t^* = (t_1^*, t_2^*), \sigma^*)$, $\mathcal{B}$ outputs $((g_1, t_1^*, t_2^*), \sigma^*)$ as its forgery.

For a successful adversary $\mathcal{A}$, the message (g_1, t_1^*, t_2^*) must be unique among the messages signed through the Issue oracle. To see this, observe, as in the proof above, that the messages signed by the SPS-EQ signature scheme are partitioned into equivalence classes determined by the discrete-logarithm relations among their components.

Consider a message of the form (g_1, A, B). The discrete logarithm of B to base A partitions the message space according to the value x satisfying $B = A^x$. In our setting, there are only two possible such values, namely x_1 and x_2. However, for the adversary's output, we have $(t_1^*)^{x_1} \neq t_2^*$ and $(t_1^*)^{x_2} \neq t_2^*$; otherwise, the ReadBit algorithm would output the corresponding bit rather than $\perp$.

It follows that (g_1, t_1^*, t_2^*) does not belong to any equivalence class previously signed through the Issue oracle. Hence, if σ^* is a valid SPS-EQ signature on this message, then $((g_1, t_1^*, t_2^*), \sigma^*)$ constitutes a valid SPS-EQ forgery. □

Theorem 4.3 (κ-unlinkability). Construction 1 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 $\mathsf{UNLINK}_{\mathcal{A},n}$ (see Definition 3.6) defined with respect to Construction 1.

- 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 $(\mathsf{pk}_C, \mathsf{psig} := (\overline{\boldsymbol{\sigma}}, S, \pi), \mathsf{nonce} := r)$, if $\mathsf{EQ.Verify}\Big((\mathsf{pk}_I^{(3)}, \mathsf{pk}_I^{(4)}, \mathsf{pk}_I^{(5)}), (\mathsf{pk}_C, \mathsf{H}(r), S)\Big) \neq 1$ it sets $\mathsf{tok} := (\perp, \perp)$. Otherwise, it sets $\mathbf{t}$ appropriately, and uses sk_I to compute $\boldsymbol{\sigma} \leftarrow\$ \mathsf{EQ.Sign}\Big((\mathsf{sk}_I^{(1)}, \mathsf{sk}_I^{(2)}, \mathsf{sk}_I^{(3)}, (g_1, t_1, t_2)\Big)$. 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 $(\mathsf{pk}_C^{(i)}, \mathsf{psig}^{(i)}, \mathsf{nonce}^{(i)})$ triple in Q^* such that for each $i \in [n]$, $\mathsf{psig}^{(i)} = (\overline{\boldsymbol{\sigma}}^{(i)}, S^{(i)}, \pi^{(i)})$. Then for any i, if $\mathsf{EQ.Verify}\Big((\mathsf{pk}_I^{(3)}, \mathsf{pk}_I^{(4)}, \mathsf{pk}_I^{(5)}), (\mathsf{pk}_C^{(i)}, \mathsf{H}(r^{(i)}), S^{(i)})\Big) \neq 1$ it sets $\mathsf{tok}_i := (\perp, \perp)$ for all i; otherwise, it sets each $\mathbf{t}^{(i)}$ appropriately, and uses sk_I to compute $\boldsymbol{\sigma}^{(i)} \leftarrow\$ \mathsf{EQ.Sign}\Big((\mathsf{sk}_I^{(1)}, \mathsf{sk}_I^{(2)}\mathsf{sk}_I^{(3)}), \Big(g_1, t_1^{(i)}, t_2^{(i)}\Big)\Big)$. 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 $(\mathsf{pk}_C, \mathsf{psig} := (\overline{\boldsymbol{\sigma}}, S, \pi), \mathsf{nonce} := r)$, if $\mathsf{EQ.Verify}\Big((\mathsf{pk}_I^{(3)}, \mathsf{pk}_I^{(4)}, \mathsf{pk}_I^{(5)}), (\mathsf{pk}_C, \mathsf{H}(r), S)\Big) \neq 1$ the challenger sets $\mathsf{tok} := (\perp, \perp)$. Otherwise, it uses its knowledge of sk_I to extract b from S, and then sets $t_1 := \mathsf{H}(r)^{\mathsf{sk}_C^{-1}}$ as usual, and $t_2 := t_1^{(1-b)\mathsf{ek}_I^{(1)} + b\mathsf{ek}_I^{(2)}}$. 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 $(\mathsf{pk}_C^{(i)}, \mathsf{psig}^{(i)}, \mathsf{nonce}^{(i)})$ triple in Q^* such that for each $i \in [n]$, $\mathsf{psig}^{(i)} = (\overline{\boldsymbol{\sigma}}^{(i)}, S^{(i)}, \pi^{(i)})$. Then, if $\mathsf{EQ.Verify}\Big((\mathsf{pk}_I^{(3)}, \mathsf{pk}_I^{(4)}, \mathsf{pk}_I^{(5)}), (pk_C^{(i)}, \mathsf{H}(r^{(i)}), S^{(i)})\Big) \neq 1$ it sets $\mathsf{tok}^{(i)} := (\perp, \perp)$ for every i; otherwise, it uses its knowledge of sk_I to extract $b^{(i)}$ from $S^{(i)}$, and then sets $t_1^{(i)} := \mathsf{H}(r^{(i)})^{1/\mathsf{sk}_C^{(i)}}$ as usual, and $t_2^{(i)} := \Big(t_1^{(i)}\Big)^{(1-b^{(i)})\mathsf{ek}_I^{(1)} + b^{(i)}\mathsf{ek}_I^{(2)}}$. 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 $(\mathsf{pk}_C, \mathsf{psig} := (\overline{\boldsymbol{\sigma}}, S, \pi), \mathsf{nonce} := r)$, instead of setting t_1 as $\mathsf{H}(r)^{sk_C^{-1}}$, 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 $\mathsf{pk}_C := g_1^\alpha$ and $\mathsf{sk}_C := \perp$. 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 $\mathsf{Adv}^{\mathsf{invDDH}}_{\mathcal{B}}$.
- Hybrid $\mathbf{H}_7$: This is the same as $\mathbf{H}_6$ except, after parsing each $(\mathsf{pk}_C^{(i)}, \mathsf{psig}^{(i)}, \mathsf{nonce}^{(i)})$ triple in Q^* such that for each $i \in [n]$, $\mathsf{psig}^{(i)} = (\overline{\boldsymbol{\sigma}}^{(i)}, S^{(i)}, \pi^{(i)})$, instead of setting $t_1^{(i)}$ as $\mathsf{H}(r^{(i)})^{sk_C^{(i)^{-1}}}$ 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 $(\mathbf{t}^{(i)}, \boldsymbol{\sigma}^{(i)})$ 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 $b_{\hat{i}}$ of $\mathsf{tok}_{\hat{i}}$ and output a random value i^* from the set of all $n_{b_{\hat{i}}}$ indices where the embedded bit was equal to $b_{\hat{i}}$. The probability that the adversary wins is

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

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

□

Theorem 4.5 (Privacy of metadata bit). Construction 1 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 $\mathsf{PMB}_{\mathcal{A},\hat{b}}$ for $\hat{b} \in \{0,1\}$ (see Definition 3.8) defined with respect to Construction 1.
- 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 $\mathsf{pk}_I^{(1)} := g_1^\alpha$ and $\mathsf{ek}_I^{(1)} := \perp$ otherwise if $\hat{b} = 1$ it sets $\mathsf{pk}_I^{(2)} := g_1^\alpha$ and $\mathsf{ek}_I^{(2)} := \perp$. 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 $t_1^{\mathsf{ek}_I^{(2-\hat{b})}} = t_2$ (note that $\mathcal{B}$ knows $\mathsf{ek}_I^{(2-\hat{b})}$) 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 $\mathsf{H}(r) \coloneqq g_1^{\beta}$, sets $S \coloneqq g_1^{\gamma}$ 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 $\boldsymbol{\sigma}$ on some tag $\mathbf{t}$ such that $\mathcal{B}$ answers the Read query incorrectly on $(\mathbf{t}, \boldsymbol{\sigma})$. 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 $\mathsf{Adv}^{\mathsf{DDH}}_{\mathcal{B}} + \mathsf{Adv}^{\mathsf{EUnf\text{-}CMA}}_{\mathcal{B}}$.

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

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

which is zero. This proves the theorem. □

C Full Instantiation of Construction 1

In this section, we explain in detail the individual components of our NIAT protocol, along with the full zero knowledge proof for the language $\mathscr{L}_{\mathsf{iss}}$.

C.1 The Concrete SPS-EQ Component

We use the SPS-EQ scheme due to Fuchsbauer, Hanser and Slamanig [18] to instantiate our protocol. Concretely, a signature on a message $\mathbf{m} = (m_1, m_2, \ldots, m_n) \in (\mathbb{G}_1^*)^{\ell}$ 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 ℓ elements in $\mathbb{Z}_p^*$, and the public key is an an ℓ-tuple from $\mathbb{G}_2^*$. Importantly, the signature can be adapted to another signature on message $\mathbf{m}^{\mu}$ without the secret key of the signer. A complete description is provided in Figure 13.

We also show the expanded token issuance and generation algorithms by instantiating with the concrete SPS-EQ of [18] in Figure 14.

C.2 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. As previously mentioned, we can apply the Fiat-Shamir heuristic [17] 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 $T_0 = \mathsf{pk}_I^{(1)} = g_1^{x_1}, T_1 = \mathsf{pk}_I^{(2)} = g_1^{x_2}, U = \mathsf{pk}_I^{(3)} = g_2^{y_1}, V = \mathsf{pk}_I^{(4)} = g_2^{y_2}$ and and $W = \mathsf{pk}_I^{(5)} = g_2^{y_3}$, and the witness $w = (\mathsf{sk}_I \coloneqq (x_1, x_2, y_1, y_2, y_3), b)$. 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 $\tilde{S}_0 \coloneqq R^{z_0}, \tilde{S}_1 \coloneqq R^{a_1} \cdot S^{-c_1}, \tilde{T}_0 \coloneqq g_1^{z_0}, \tilde{T}_1 \coloneqq g_1^{a_1} \cdot T_1^{-c_1}, \tilde{U} \coloneqq g_2^{z_u}, \tilde{V} \coloneqq g_2^{z_v}$ and $\tilde{W} \coloneqq g_2^{z_w}$. It then computes the challenge $c \coloneqq \mathsf{H}(g_1, g_2, \mathsf{pk}_C, S, T_0, T_1, U, V, , W, \tilde{S}_0, \tilde{S}_1, \tilde{T}_0, \tilde{T}_1, \tilde{U}, \tilde{V}, \tilde{W})$, followed by $a_u \coloneqq z_u + cy_1, a_v \coloneqq z_v + cy_2, a_w \coloneqq z_w + cy_3, c_0 \coloneqq c - c_1, a_0 \coloneqq z_0 + c_0 x_1$.

  (2) If $b = 1$, the issuer (prover) samples integers $z_1, z_u, z_v, z_w, c_0, a_0$, and computes commitments $\tilde{S}_0 \coloneqq R^{a_0} \cdot S^{-c_0}, \tilde{S}_1 \coloneqq R^{z_1}, \tilde{T}_0 \coloneqq g_1^{a_0} \cdot T_0^{-c_0}, \tilde{T}_1 \coloneqq g_1^{z_1}, \tilde{U} \coloneqq g_2^{z_u}, \tilde{V} \coloneqq g_2^{z_v}$ and $\tilde{W} \coloneqq g_2^{z_w}$. It then computes the challenge $c \coloneqq \mathsf{H}(g_1, g_2, \mathsf{pk}_C, S, T_0, T_1, U, V, W, \tilde{S}_0, \tilde{S}_1, \tilde{T}_0, \tilde{T}_1, \tilde{U}, \tilde{V}, \tilde{W})$, followed by $a_u \coloneqq z_u + cy_1, a_v \coloneqq z_v + cy_2, a_w \coloneqq z_w + cy_3, c_1 \coloneqq c - c_0, a_1 \coloneqq z_1 + c_1 x_2$.

  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 [11] 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 $c = c_0 + c_1, \tilde{S}_0 \coloneqq R^{a_0} \cdot S^{-c_0}, \tilde{S}_1 \coloneqq R^{a_1} \cdot S^{-c_1}, \tilde{T}_0 \coloneqq g_1^{a_0} \cdot T_0^{-c_0}, \tilde{T}_1 \coloneqq g_1^{a_1} \cdot T_1^{-c_1}, \tilde{U} \coloneqq g_2^{a_u} \cdot U^{-c}, \tilde{V} \coloneqq g_2^{a_v} \cdot V^{-c}$ and $\tilde{W} \coloneqq g_2^{a_w} \cdot W^{-c}$. Finally, the client checks that $c \coloneqq \mathsf{H}(g_1, g_2, \mathsf{pk}_C, S, T_0, T_1, U, V, W, \tilde{S}_0, \tilde{S}_1, \tilde{T}_0, \tilde{T}_1, \tilde{U}, \tilde{V}, \tilde{W})$.

C.3 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 $\mathsf{e}(\mathsf{pk}_C, \mathsf{pk}_I^{(3)})$ and $\mathsf{e}(g_1, \mathsf{pk}_I^{(3)})$ respectively, can be precomputed. This already reduces the number of pairing computations per presignature/token by one. For an issuing authority with key pair $(\mathsf{sk}_I, \mathsf{pk}_I)$, 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

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

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

Similarly, the batch verification check for the verifier is given by replacing n by n' in (2) 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.[8]

D Security Proofs for Section 5

Correctness and reusability of Construction 2 are easily extendable from that of the previous protocol (see Appendix B); and privacy of metadata bit is also proven identically. One-more unforgeability and public-validity consistency of Construction 2 can similarly be extended from that of the previous protocol with the additional requirement of soundness of $\mathsf{NIZK}_{\mathsf{obt}}$. 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 D.1 (κ-unlinkability). *Construction 2 is κ-unlinkable (in the random oracle and KOSK models) for $\kappa = 2$, assuming $\mathsf{NIZK}_{\mathsf{iss}}$ is an argument of knowledge, that DDH and k-DDH assumptions hold in $\mathbb{G}_1$ and EQ perfectly adapts signatures under a malicious signer $\mathsf{NIZK}_{\mathsf{obt}}$ satisfies zero knowledge.*

[8] 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.

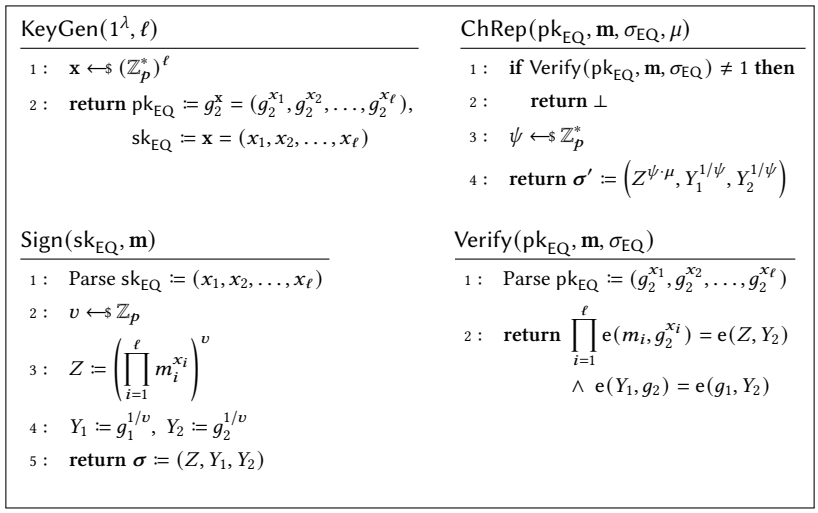

Figure 13: SPS-EQ scheme from [18].

Pf sketch. Recall the proof of Theorem 4.3 (Appendix B). The idea there was to make the final token effectively independent of $\mathsf{sk}_C =: \alpha$, $\mathsf{psig} =: (\overline{\sigma}, R, S, \pi)$ and $\mathsf{nonce} =: r$, 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 $\mathsf{H}(r)^{\alpha^{-1}}$ 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 $\mathsf{NIZK}_{\mathsf{obt}}$, the proof π_{o} is simulatable. Next, we can replace $\mathbf{t} = \left(g_1^{(\alpha+\hat{r})^{-1}}, R^{(\alpha+\hat{r})^{-1}}, S^{(\alpha+\hat{r})^{-1}}\right)$ by $\left(g_1^{\rho}, (g_1^{\varepsilon})^{\rho}, \left(g_1^{\varepsilon\cdot((1-b)x_1+bx_2)}\right)^{\rho}\right)$ where we 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 [16] which is adaptively secure under the k-DDH assumption [9] for $\mathsf{poly}(\lambda)$ sized domains. This, in particular means that we must ensure that the range of $\mathsf{H}_{\mathbb{Z}}(r)$ is of size $\mathsf{poly}(\lambda)$. 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 $\boldsymbol{\phi}$ as ElGamal encryption of $g_1^{\mathsf{id}\cdot(\alpha+\hat{r})}$ (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 $\mathsf{H}_{\mathbb{G}\to\mathbb{G}}(r) := g_1^{\nu_r}$. Then, for a query of the form $(\cdot, r, \mathsf{id})$, we create the encryption $\boldsymbol{\phi}$ as $\phi_1 := g_1^{\nu_r}$, and $\phi_2 := \mu \cdot \phi_1^{\alpha}$ for message $\mu := \mathsf{pk}_C^{\mathsf{id}} \cdot g_1^{(\mathsf{id}+\nu_r)\cdot\hat{r}}$. 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. □

Theorem D.2 (Double-spend identification). *Construction 2 satisfies double-spend identification (in the random oracle model) assuming* $\mathsf{NIZK}_{\mathsf{obt}}$ *is sound.*

Pf sketch. If $\mathsf{NIZK}_{\mathsf{obt}}$ is sound, then it follows that the adversary knows $(\mathsf{sk}_C, \hat{r})$ (resp. $(\mathsf{sk}'_C, \hat{r}')$) such that $t_1 = g_1^{\left(\mathsf{sk}_C+\hat{r}\right)^{-1}}$ (resp. $t_1 = g_1^{\left(\mathsf{sk}'_C+\hat{r}'\right)^{-1}}$) and $\phi_2 = g_1^{\mathsf{id}\cdot\left(\mathsf{sk}_C+\hat{r}\right)} \cdot \mathsf{H}_{\mathbb{G}\to\mathbb{G}}(t_1)^{\mathsf{sk}_C+\hat{r}}$ (resp. $\phi'_2 = g_1^{\mathsf{id}'\cdot\left(\mathsf{sk}'_C+\hat{r}'\right)} \cdot \mathsf{H}_{\mathbb{G}\to\mathbb{G}}(t_1)^{\mathsf{sk}'_C+\hat{r}}$) with respect to $(\mathbf{t}, \boldsymbol{\sigma})$ (resp. $(\mathbf{t}, \boldsymbol{\sigma}')$). Now since $\mathbf{t}$ is common between both tokens, we have $\left(\mathsf{sk}_C + \hat{r}\right) = \left(\mathsf{sk}'_C + \hat{r}'\right)$ (all computations are $\mod p$) so that $\left({\phi_2}^{-1} \cdot \phi'_2\right)^{\left(\mathsf{id}-\mathsf{id}'\right)^{-1}}$ simplifies to $g_1^{\mathsf{sk}_C+\hat{r}} = g_1^{\mathsf{sk}'_C+\hat{r}'}$, such that at least one of $(g_1^{\mathsf{sk}_C}, \mathsf{nonce})$ and $(g_1^{\mathsf{sk}'_C}, \mathsf{nonce}')$ is in aux with $\hat{r} := \mathsf{H}_{\mathbb{Z}}(\mathsf{nonce})$ (similarly, $\hat{r}'$). Thus, with all but negligible probability, $\mathsf{DSIden}(\mathsf{pk}_I, \mathbf{t}, \boldsymbol{\sigma}, \boldsymbol{\sigma}') \neq \perp$. □

Theorem D.3 (Double-spend exculpability). *Construction 2 satisfies double-spend exculpability assuming* $\mathsf{NIZK}_{\mathsf{obt}}$ *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 $\Pi := (\mathbf{t}, \boldsymbol{\sigma}, \boldsymbol{\sigma}', \mathsf{nonce})$ where $\boldsymbol{\sigma} := (\mathbf{s}, \boldsymbol{\phi}, \pi_{\mathsf{o}}, \mathsf{id})$ (similarly $\boldsymbol{\sigma}'$) such that $\left(\phi_2^{-1} \cdot \phi'_2\right)^{\left(\mathsf{id}-\mathsf{id}'\right)^{-1}} = \mathsf{pk}_C \cdot g_1^{\mathsf{H}_{\mathbb{Z}}(\mathsf{nonce})}$ and, both, the proof and signature verify. Having queried $\mathcal{O}_{\mathsf{Issue}}$ at most once per (psig, nonce) pair (w.l.o.g. suppose the query corresponds to the token $(\mathbf{t}, \boldsymbol{\sigma})$) 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 proof without sk_C, and we can reduce to the soundness of $\mathsf{NIZK}_{\mathsf{obt}}$. □

E 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 1. The only notable change is that the issuer now uses a tag-based equivalence class signature scheme instead of a structure-preserving signature.

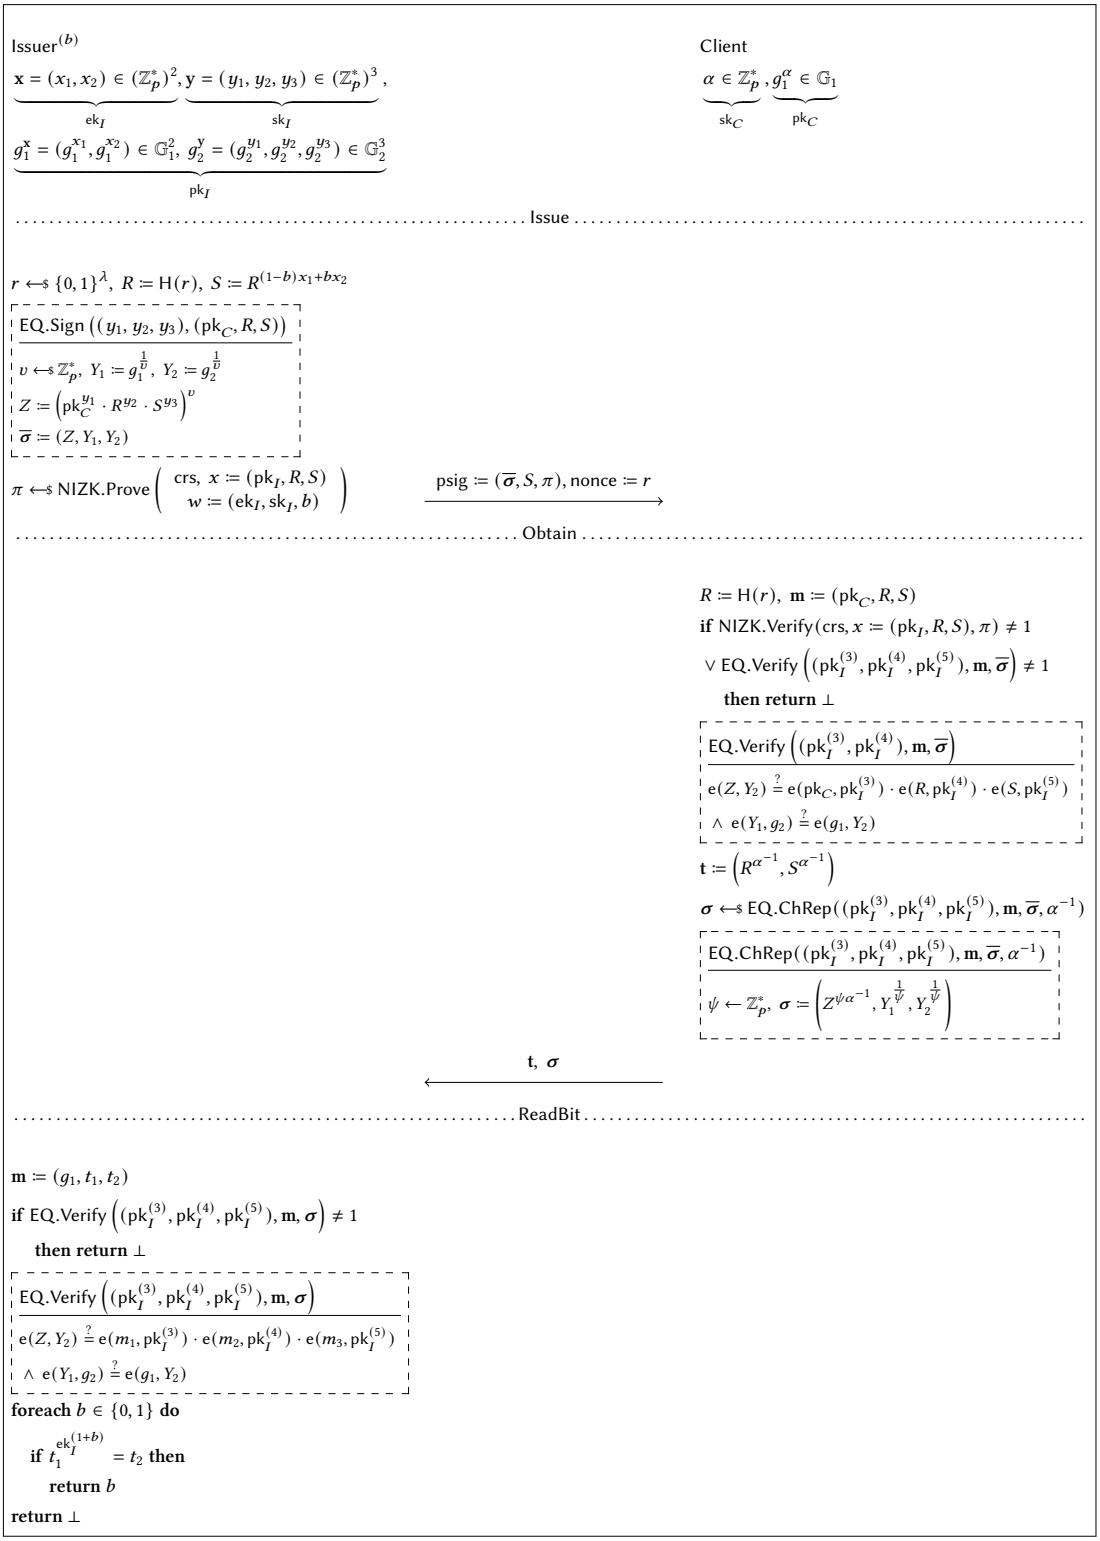

Figure 14: Expanded SPS-EQ construction for NIAT presignature issuance and token generation using [18].

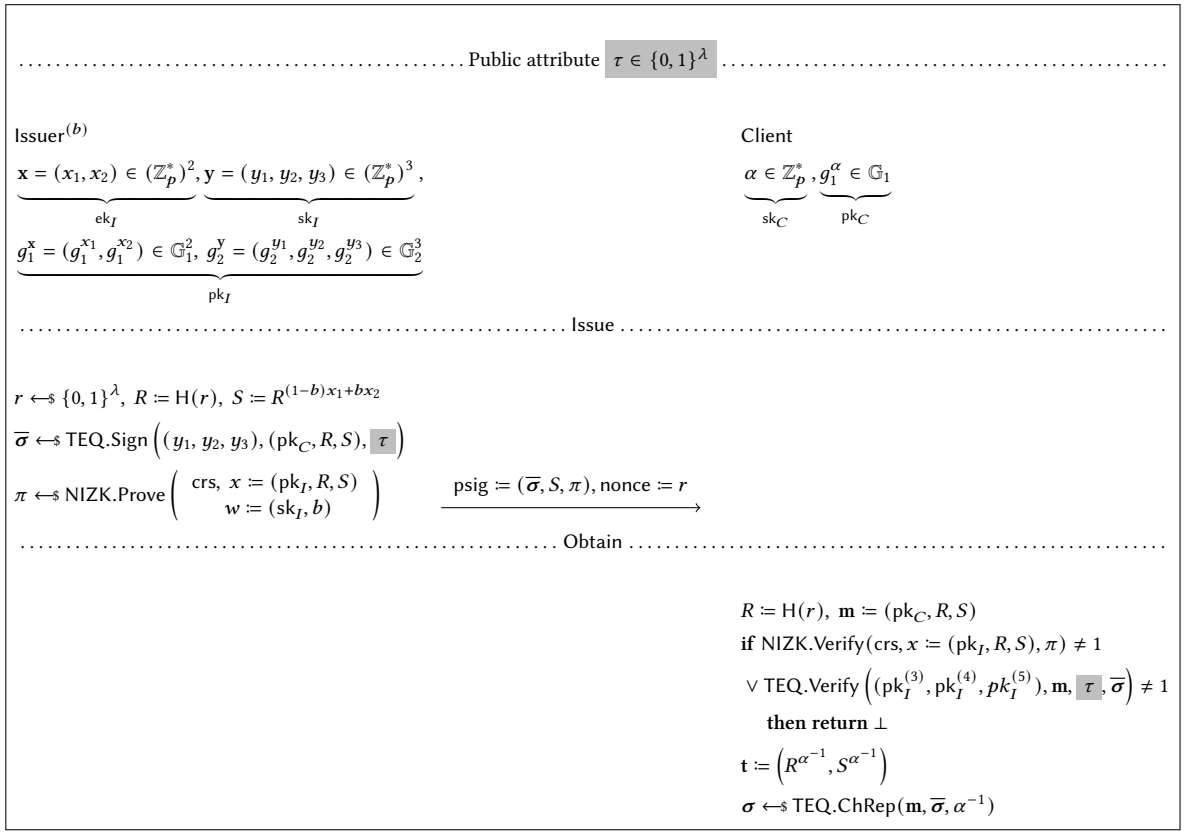

Figure 15: TB-EQS construction for NIAT (with public attributes) presignature issuance and token generation.

E.1 Tag-based Equivalence Class Signatures

We now recall the definition of Tag-based Equivalence Class Signatures (TB-EQS) [23] 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 E.1 (TB-EQS)**.** A TB-EQS scheme consists of the following PPT algorithms:

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

Definition E.2 (Existential unforgeability under chosen message attack)**.** Let $\mathsf{Adv}^{\mathsf{EUnf\text{-}CMA}}_{\mathcal{A},\ell} = \Pr\left[\mathsf{EUnf\text{-}CMA}_{\mathcal{A},\ell}(\lambda) = \mathbf{accept}\right]$ be the advantage of an PPT adversary $\mathcal{A}$ in the EUnf-CMA experiment defined in Figure 16. A TB-EQS scheme over $(\mathbb{G}_i^*)^\ell$ is existentially unforgeable under adaptive chosen-message attacks if for all $\ell > 1$, the advantage $\mathsf{Adv}^{\mathsf{EUnf\text{-}CMA}}_{\mathcal{A}} = \mathsf{negl}(\lambda)$.

Definition E.3 (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 $(\mathsf{pk}_{\mathsf{TEQ}}, \mathbf{m}, \sigma_{\mathsf{TEQ}}, \mu, \tau)$ satisfying

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

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

E.2 Construction

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

Presignature issuance and token generation. The presignature issuance and token generation protocols is presented in detail in Figure 15. The difference here from Construction 1 is the use of the

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

$Q \coloneqq \varnothing$
$(\mathsf{pp}_{\mathsf{TEQ}}) \leftarrow\!\$\ \mathsf{TEQ.Setup}(1^\lambda, \ell)$
$(\mathsf{pk}_{\mathsf{TEQ}}, \mathsf{sk}_{\mathsf{TEQ}}) \leftarrow\!\$\ \mathsf{TEQ.KeyGen}(\mathsf{pp}_{\mathsf{TEQ}})$
$(\mathbf{m}^*, \sigma^*_{\mathsf{TEQ}}, \tau^*) \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{Sign}}}(\mathsf{pk}_{\mathsf{TEQ}})$
return $([\mathbf{m}^*], \tau^*) \neq ([\mathbf{m}], \tau)\ \forall (\mathbf{m}, \tau) \in Q \wedge$
$\mathsf{TEQ.Verify}(\mathsf{pk}_{\mathsf{TEQ}}, \mathbf{m}^*, \tau^*, \sigma^*_{\mathsf{TEQ}}) = 1$

Oracle $\mathcal{O}_{\mathsf{Sign}}(\mathbf{m}, \tau)$

$\sigma_{\mathsf{TEQ}} \leftarrow\!\$\ \mathsf{TEQ.Sign}(\mathsf{sk}_{\mathsf{EQ}}, \mathbf{m}, \tau)$
$Q \coloneqq Q \cup \{(\mathbf{m}, \tau)\}$
return σ_{TEQ}

Figure 16: The existential unforgeability (under chosen message attack) experiment for TB-EQS.

public attribute τ as input to the TB-EQS scheme.

Public verification (token redemption). The public verification algorithm simply verifies the tag-based equivalence class signature $\boldsymbol{\sigma}$ for the public attribute τ.

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

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

Correctness, reusability and all security properties follow similarly from that of Construction 1.

Security. The main security theorems are presented below.

Theorem E.4 (One-more unforgeability). *Construction E 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. □

Theorem E.5 (Public-validity consistency). *Construction E is public-validity consistent (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 4.2 but with respect to EUnf-CMA security of the TB-EQS scheme. □

Theorem E.6 (κ-unlinkability). *Construction E 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. □

Theorem E.7 (Privacy of metadata bit). *Construction E 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.*

Proof. Proof is omitted as it is nearly identical to that of Theorem 4.5. □