Anonymous Tokens with Public Metadata and Applications to Private Contact Tracing

Tjerand Silde[1] and Martin Strand[2]

[1] Department of Mathematical Sciences,
Norwegian University of Science and Technology – NTNU,
tjerand.silde@ntnu.no
[2] Norwegian Defence Research Establishment – FFI,
martin.strand@ffi.no

Abstract. Anonymous single-use tokens have seen recent applications in private Internet browsing and anonymous statistics collection. We develop new schemes in order to include public metadata such as expiration dates for tokens. This inclusion enables planned mass revocation of tokens without distributing new keys, which for natural instantiations can give 77 % and 90 % amortized traffic savings compared to Privacy Pass (Davidson *et al.*, 2018) and DIT: De-Identified Authenticated Telemetry at Scale (Huang *et al.*, 2021), respectively. By transforming the public key, we are able to append public metadata to several existing protocols essentially without increasing computation or communication.
Additional contributions include expanded definitions, a more complete framework for anonymous single-use tokens and a description of how anonymous tokens can improve the privacy in dp^{3}t-like digital contact tracing applications. We also extend the protocol to create efficient and conceptually simple tokens with both public and private metadata, and tokens with public metadata and public verifiability from pairings.

Keywords: anonymous tokens · public metadata · contact tracing

1 Introduction

Anonymous credentials have been an active research area since the 1980's [23,24], involving schemes like blind signatures, partially blind signatures, anonymous tokens, attribute-based credentials, group signatures, ring signatures etc. This enables more complex systems for e.g., electronic cash or electronic voting, but also, to protect the privacy of the users in chat applications like Signal.

Recent work by Davidson *et al.* [30] presents a very practical protocol, named Privacy Pass [29], for anonymous single-use tokens. This protocol allows users to browse anonymously, e.g., using Tor, without having to solve a CAPTCHA every time they visit a website. Privacy Pass gives the user a set of randomized tokens whenever they solve a CAPTCHA, which they then later can redeem instead of solving a new CAPTCHA. This improves the usability of anonymous browsing. It also gives protection against spam, prevents DDoS attacks and provides fraud

resistance without the need for cross-site tracking or fingerprinting. However, the only way to expire or revoke batches of unspent tokens is by replacing the private-public key pair in a trusted way, which is impractical [28].

Privacy Pass has gained a lot of attention, and is currently being integrated to improve privacy in several applications, e.g., for private file storage[3] and for basic attention tokens (BATs) in the Brave browser[4]. It can also be used for private click measurement when making a purchase or signing up for a service[5].

Facebook uses partially blind signatures for combating fraud [41], and they have developed an extension of Privacy Pass called DIT: De-Identified Authenticated Telemetry at Scale [39], which is used for privately collecting client-side telemetry from WhatsApp. DIT requires daily key-rotation to prevent DoS attacks, which led to the development of an attribute-based verifiable oblivious pseudorandom function for transparent key-rotation.

The IETF is currently standardizing Privacy Pass [40], while Trust Token [51] is currently being standardized by the World Wide Web Consortium. Both standardization processes mention private and public metadata, in addition to public verifiability, as desirable extensions to the Privacy Pass protocol. Public metadata allows for more efficient key-rotation, and opens for applications using public labeling and public anonymity sets, while private metadata allows for allow/deny lists, rate-limiting, or trust-indication. Public verifiability allows for outsourcing signing or verification of tokens.

Kreuter *et al.* [44] gave the first construction of anonymous tokens with private metadata, while we give the first construction with public metadata. Our construction can also be combined with private metadata or public verifiability.

Privacy Pass guarantees anonymity for all tokens generated by the same key. The addition of any metadata reduces the anonymity set. We have designed the protocol in such a way that the user and the signer must agree on the metadata. Any application should restrict its use of metadata to generic, predefined values that would otherwise have triggered a change of keys, e.g., expiry dates. Client software should validate that the metadata is in accordance to the policy, and reject any malformed tokens. Furthermore, private metadata bits also reduces the anonymity set. Our protocol can easily be extended to include more than one private metadata bit, but this must be done with great care, as it opens for secretly tracking smaller sets of individual users.

Independently of this work, Tyagi *et al.* [50] have proposed a similar construction to include public metadata, along with a novel hardness assumption and a reduction to a more conventional problem, to be used in partially oblivious pseudo-random functions. We discuss their work further in Sections 1.4 and 3.

[3] PrivateStorage: https://medium.com/least-authority/the-path-from-s4-to-privatestorage-ae9d4a10b2ae.

[4] Brave: https://github.com/brave/brave-browser/wiki/Security-and-privacy-model-for-ad-confirmations.

[5] Private Click Measurement: https://privacycg.github.io/private-click-measurement.

1.1 Our Contribution

Our contribution in this work is threefold: First, we present new definitions and a new framework for anonymous tokens – extending the work by Kreuter *et al.* [44] – to also consider public metadata and/or public verifiability. Secondly, we present three efficient protocols for anonymous tokens with efficient batched revocation: 1) Privacy Pass [30] with public metadata, 2) Kreuter *et al.* [44] with public and private metadata, and 3) a Privacy Pass inspired protocol using pairings to satisfy public verifiability while including public metadata. Thirdly, we present contact tracing as a new and important application for anonymous single-use tokens, and discuss the implementation of Privacy Pass used in the Norwegian contact tracing app Smittestopp to improve users' privacy.

Updated Definitions and New Framework. Several works have asked for efficient batched revocation of anonymous tokens without key-rotation [28, 30]. Additionally, there is a need for anonymous tokens with public verifiability [51], so that token generation can be delegated, and verification can be performed locally for token redemption. We provide updated definitions for all of these cases: designated verifier anonymous tokens with or without public and/or private metadata and public verifier anonymous tokens with and without public metadata. Details can be found in Section 2.

Anonymous Tokens with Public Metadata. We present the first anonymous tokens protocols with efficient batched revocation, meaning that the protocol only requires one round of communication based on lightweight primitives and that we avoid key-rotation. The key insight in our protocol is conceptually very simple: all parties locally update the public key based on the hash of the public metadata, and then execute the protocols with respect to the new key pair. The main challenge is to sign tokens in a way that does not allow the user to forge tokens initially signed under metadata md to be valid under metadata md' instead. Let k be the secret key and let $d = \mathtt{H}(\mathsf{md})$ be the hash of the metadata. Our solution, inspired by Zhang *et al.* [53], is to use the inverse $e = (d + k)^{-1}$ as the new signing key. This allows us to replace the secret keys in the previous protocols in a modular way.

Furthermore, to avoid subliminal channels, the signer needs to prove that the signed token is computed correctly. This is easily solved for Privacy Pass [30]. In the original protocol they use a zero-knowledge protocol to prove, given generator G, public key $K = [k]G$, blinded token T' and signed token $W' = [k]T'$, the equality of discrete logarithms $\log_G K = k = \log_{T'} W'$ to ensure correctness. In our updated protocol, including metadata md, updated public key $U = [d]G + K$ and signed token $W' = [e]T'$, we prove the equality of discrete logarithms $\log_G U = d + k = \log_{W'} T'$ to ensure correctness.

However, it is not as easy to ensure correctness in the extended version of the protocol by Kreuter *et al.* [44] including both public and private metadata. We solve this by combining an $\mathtt{OR}$-proof with two $\mathtt{AND}$-proofs to make sure that the correct key is used. Further improvement is an open problem.

Next, we give a protocol based on pairings. The protocol is an adapted version of the partially blind signatures by Zhang *et al.* [53], where we tweak it into the same structure as Privacy Pass. We note that the communication in the protocol is the same, but in addition to get a more streamlined protocol structure, we also allow for more efficient instantiation in practice using the BLS12-381 pairing [6]. Ideally, we would like to avoid pairings altogether, but this seems necessary in practice. See more details about the protocols in Section 3.

Finally, we detail the communication efficiency of the protocols in Section 4, and compare our constructions with the current state of the art with respect to efficient batched revocation in Table 1. We show that our protocols are much more efficient in practice. We also make a concrete comparison with DIT [39] for collecting telemetry-data from WhatsApp, and show that our protocol in Figure 4 would decrease the size of the signed token in a natural setting by 90 %, saving the Facebook servers up to 1.7 TB of communication every day.

More Private Contact Tracing. Many countries have recently developed contact tracing apps as one of the measurements to battle the ongoing pandemic. These apps are inherently storing sensitive information about the user, e.g., the users' location graph and social graph. To avoid large, centralized databases with such sensitive information about a large portion of a country's adult population, most apps are based on the decentralized Google/Apple Exposure Notification System (ENS). However, there are still privacy issues with regards to uploading the randomized exposure keys to the central server, as the user would have to identify themselves to ensure that only people who have tested positive for COVID-19 are able to upload keys. We implemented Privacy Pass into the Norwegian contact tracing app to improve the users' privacy. We present more details about the contact tracing infrastructure and improvements in Section 5.

1.2 Comparison to Anonymous Credentials

There is a long line of research on more generalized anonymous credentials with properties like – in addition to unlinkability and unforgeability – multi-show, multi-attributes, and revocability, allowing expiration dates to be attributes.

However, generalized anonymous credentials often depends on stronger assumptions, e.g., strong RSA [14, 15, 17, 18, 20], strong DH [3] or DL assumptions in bilinear groups [19, 36]. Some schemes only depend on DDH [5, 21, 22, 47], but these schemes require larger messages in general. In conclusion, generalized anonymous credentials inherently impose larger parameters, more rounds of communication and less efficient protocols in practice, resulting in thousands of bits on communication over multiple rounds.

Finally, more general and complex anonymous credentials make these schemes less suited for use in simpler single-use systems with many users, which is the case in our setting. We want to minimize the rounds of communication and data being sent, in addition to minimizing the local computation and the local state. Hence, we only compare to one-round single-use efficiently revocable anonymous credentials with minimal communication in Section 4.

1.3 Related work

Our work achieving designated verification and public metadata extends a long line of publications. Freedman *et al.* [33] introduced oblivious pseudo-random functions, and Jarecki *et al.* [42,43] gave an efficient instantiation based on DDH in the random oracle model. Papadopoulos *et al.* [46] gave a verifiable PRF from elliptic curves, and Burns *et al.* [13] gave an oblivious PRF from elliptic curves. Privacy Pass combined these results with an extended version of the Chaum-Pedersen zero-knowledge protocol [25] given by Henry and Goldberg [37,38] to prove knowledge of batches of elements having the same discrete logarithm, and Kreuter *et al.* [44] added private metadata to Privacy Pass. In a concurrent work, Tyagi *et al.* [50] recently extended this line of works to partially oblivious PRFs.

To achieve public verifiability we use parings, inspired by the seminal work of Boneh *et al.* [12] for short and efficient signatures and a series of constructions of (partially) blind signatures based on pairings [9,10,16,26,27,34,35,53].

1.4 Chronology

As we report on both an implementation and new protocols, we believe it can be helpful to lay out the chronology of this work to separate the contributions.

Mid-October 2020, the authors were made aware of a potential privacy weakness in Norway's upcoming second COVID-19 contact tracing app Smittestopp. The first iteration had been stopped by the Norwegian Data Protection Agency in June, due to privacy concerns following from lack of data minimalization. The new app had a set launch date in December.

The issue at hand was that the verification service would collect IDs in order to automatically verify the infection status, and then send a token to the app which could then be used for uploading exposure keys. This token would create a hard link between an ID-based service and the rest of the system, in which the users are assumed to be anonymous.

Within few days, we suggested using Privacy Pass in order to remove this link. Due to lack of capacity, our proposal was acknowledged, but we were encouraged to submit a pull request. We teamed up with Henrik Walker Moe to implement Privacy Pass in C#, and our implementation was eventually accepted into Smittestopp along with an improvised solution to rotate keys every three days.

Motivated by this process and the last-minute improvisation, we expanded the original Privacy Pass protocol. Our initial manuscript was posted on ePrint February 24th, 2021. We were then made aware of a complication to the security proof. A correct proof was posted on ePrint by Tyagi *et al.* [50] June 24th, 2021. The primary separation between these two manuscripts are that they present a correct proof, while we were the first to present this protocol along with its variations. We also present the protocols in a way that is compatible to previous work. In this sense, these works complement each other.

The new protocol has not been implemented in Smittestopp. This is due to lack of further development of the app, and we do not expect any major changes to be accepted into the codebase at this stage.

2 Definitions for Anonymous Tokens

Anonymous tokens as used in Privacy Pass are conceptually simple: both issuance and verification require the private key, and the final token is uniquely determined by the token seed t and the private key. Kreuter *et al.* [44] extended this notion by adding a private bit in the token. We further extend the definition in two different directions: we want to add public metadata, and we want to make the token publicly verifiable. Now, private bits do not make immediate sense in the context of a publicly verifiable token scheme, but public metadata can be relevant in both settings.

The metadata can for instance be used to indicate an expiry date, replacing the need for frequent key rotation in certain applications [39] . We model it as a value that the user and issuer must agree upon, which should restrict the issuer from using arbitrary, identifiable values.

Lending terminology from programming, we would like the definition to provide backwards compatibility, and handle the notational incompatibility between private and public verifiability. To this end, we imitate the notion of [optional arguments] from programming. The notation $\mathsf{vk}|\mathsf{sk}$ is meant as "at least one of the public or the secret key". We align our definitions as close as possible to those by Kreuter *et al.* [44].

Definition 1 (Anonymous tokens). *An anonymous token scheme with zero or more of* ***private metadata bit****,* ***public metadata****, or* ***public verifiability*** *consists of the following algorithms:*

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{AT.Setup}(1^\lambda)$*, the setup algorithm that takes as input the security parameter* λ *in unary form, and returns a common reference string* crs *and trapdoor* td*. All the remaining algorithms take* crs *as their first input.*
- $(\mathsf{pp}, \mathsf{sk}, [\mathsf{vk}]) \leftarrow \mathsf{AT.KGen}(\mathsf{crs})$*, the key generation algorithm that generates a signing key* sk *and optionally a verification key* vk *along with public parameters* pp*. All the remaining algorithms take* pp *as their second input.*
- $\sigma \leftarrow \langle \mathsf{AT.User}(\mathsf{pp}, [\mathsf{vk}], t, [\mathsf{md}]), \mathsf{AT.Sign}(\mathsf{sk}, [\mathsf{md}], [b]) \rangle$*, the token issuance protocol, which involves interactive algorithms* $\mathsf{AT.User}$ *and* $\mathsf{AT.Sign}$*. The user algorithm takes as input values the public parameters and the token seed* $t \in \{0,1\}^\lambda$*, and potentially the verification key* vk *and the public metadata* md*. The signing algorithm takes the private key* sk *and potentially metadata* md *and the private bit* b*. At the end of the interaction, the issuer outputs nothing, while the user outputs* σ*, or* $\perp$ *to indicate error.*
- *bool* $\leftarrow \mathsf{AT.Vf}(\mathsf{vk}|\mathsf{sk}, t, [\mathsf{md}], \sigma)$*, the verification algorithm that takes as input either the public verification key* vk *or the private key* sk*, a token seed* t*, metadata* md *and the signature* σ*. It returns true if the token was valid.*
- $[\mathsf{ind} \leftarrow \mathsf{AT.ReadBit}(\mathsf{sk}, t, [\mathsf{md}], \sigma)]$*, the private bit extraction algorithm that takes as input the private key* sk *and token* $(t, [\mathsf{md}], \sigma)$*. It returns an indicator* $\mathsf{ind} \in \{\perp, 0, 1\}$ *which is either the private bit, or* $\perp$*.*

The notation of the above definition should be interpreted in a global sense. If one – for example – wants to use public metadata, it should be included everywhere it is mentioned. This listing then defines the following six notions:

1. With designated verification:
   (a) Anonymous single-use tokens
   (b) Anonymous single-use tokens with private metadata bit
   (c) Anonymous single-use tokens with public metadata
   (d) Anonymous single-use tokens with public and private metadata
2. With public verification:
   (a) Anonymous single-use tokens
   (b) Anonymous single-use tokens with public metadata

Examples of 1a and 1b are well known from previous work [30, 44]. A previous example of 2b is known as a partially blind signature scheme [2]. We will provide new examples of the last four (2a is implicit in 2b) in Section 3, Appendix C and Appendix D. We collectively refer to all of these as anonymous tokens.

We follow the convention of dividing the interactive protocol $\langle \mathsf{AT.User}, \mathsf{AT.Sign} \rangle$ into the non-interactive algorithms $\mathsf{AT.User}_0$, $\mathsf{AT.Sign}_0$ and $\mathsf{AT.User}_1$.

An anonymous token scheme must satisfy the following properties:

Definition 2 (Token correctness). *An anonymous token scheme* AT *is correct if any honestly generated token verifies. For any honestly generated* crs, $(\mathsf{pp}, \mathsf{sk}, [\mathsf{vk}])$, t *and* $[\mathsf{md}]$,

$$\Pr[\mathsf{AT.Vf}(\mathsf{vk}, t, [\mathsf{md}], \langle \mathsf{AT.User}(\mathsf{pp}, [\mathsf{vk}], t, \mathsf{md}), \mathsf{AT.Sign}(\mathsf{sk}, [\mathsf{md}], [b]) \rangle) = 1] = 1 - \mathsf{negl}(\lambda).$$

We split correctness of the private metadata bit into a separate definition in order to reduce notational clutter. This definition only applies in the private-key setting, and the parameters have been fixed accordingly.

Definition 3 (Correct private bit). *An anonymous token scheme* AT *is correct with respect to private metadata if the correct bit is retrieved successfully:*

$$\Pr[\mathsf{AT.ReadBit}(\mathsf{sk}, t, \langle \mathsf{AT.User}(\mathsf{pp}, t, [\mathsf{md}]), \mathsf{AT.Sign}(\mathsf{sk}, [\mathsf{md}], b) \rangle) = b] = 1 - \mathsf{negl}(\lambda).$$

No adversary should be able to redeem other tokens than those that have been correctly issued. The *one-more unforgeability* notion has become the common notion for anonymous credentials. It allows the adversary to claim ℓ tokens from the issuer, and the adversary should not be able to redeem $\ell + 1$ tokens. We require the tokens to be unique with respect to the value of the seed t.

Definition 4 (One-more unforgeability). *An anonymous token scheme* AT *is one-more unforgeable if for any* PPT *adversary* $\mathcal{A}$*, and any* $\ell \geq 0$*:*

$$\mathsf{Adv}^{\mathrm{omuf}}_{\mathsf{AT},\mathcal{A},\ell}(\lambda) := \Pr[\mathrm{OMUF}_{\mathsf{AT},\mathcal{A},\ell}(\lambda) = 1] = \mathsf{negl}(\lambda),$$

where $\mathrm{OMUF}_{\mathsf{AT},\mathcal{A},\ell}$ *is the game defined in Figure 1.*

Next, we want to provide user anonymity. The right notion for this is unlinkability, which guarantees that even colluding issuers and verifiers are unable to link tokens. Arbitrary metadata is a strong way of creating a link, and we omit this problem by only considering fixed public metadata for this notion. Notice that the adversary may query the user oracles for any public metadata md, but that we expect the post-processing to implicitly fail if $\mathsf{md} \neq \mathsf{md}'$. This is in line with for example expiry dates, which would otherwise have been solved in practice using key rotation, and the definition is (as usual) also using a fixed key. Private metadata is outside the control of the user, and gives one bit leakage. We fix it for this game. Note that the adversary controls the keys, and that we therefore do not need to provide access to signing and verification oracles.

Definition 5 (Unlinkability). *An anonymous token scheme* AT *is* κ*-unlinkable if for any* PPT *adversary* $\mathcal{A}$*, fixed* b*,* md*, and any* $m > 0$*,*

$$\mathsf{Adv}^{\mathrm{unlink}}_{\mathsf{AT},\mathcal{A},m,[b],[\mathsf{md}]}(\lambda) := \Pr\left[\mathrm{UNLINK}_{\mathsf{AT},\mathcal{A},m,[b],[\mathsf{md}]}(\lambda) = 1\right] \quad \leq \frac{\kappa}{m} + \mathsf{negl}(\lambda),$$

where $\mathrm{UNLINK}_{\mathsf{AT},\mathcal{A},m}$ *is the game defined in Figure 2.*

We finally consider the private metadata bit. We give the adversary access to two signing oracles: One uses the adversary's chosen private bit, the other is using a fixed bit for the game. The adversary can also query a verification oracle. At the end, the adversary outputs its guess for the fixed challenge bit.

Definition 6 (Private metadata bit). *An anonymous token scheme* AT *provides private metadata bit if for any* PPT *adversary* $\mathcal{A}$*,*

$$\mathsf{Adv}^{\mathrm{pmb}}_{\mathsf{AT},\mathcal{A}}(\lambda) := \left|\Pr[\mathrm{PMB}^0_{\mathsf{AT},\mathcal{A}}(\lambda)] - \Pr[\mathrm{PMB}^1_{\mathsf{AT},\mathcal{A}}(\lambda)]\right| = \mathsf{negl}(\lambda)$$

where $\mathrm{PMB}^\beta_{\mathsf{AT},\mathcal{A}}$ *is the game defined in Figure 3.*

Game $\mathrm{OMUF}_{\mathsf{AT},\mathcal{A},\ell}(\lambda)$	Oracle $\textsc{Sign}(\mathsf{msg}, [\mathsf{md}], [b])$
$(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{AT.Setup}(1^\lambda)$	$q_{b,\mathsf{md}} := q_{b,\mathsf{md}} + 1$
$(\mathsf{pp}, \mathsf{sk}, [\mathsf{vk}]) \leftarrow \mathsf{AT.KGen}(\mathsf{crs})$	**return** $\mathsf{AT.Sign}_0(\mathsf{sk}, \mathsf{msg}, [\mathsf{md}], [b])$
for $(b \in \{0,1\}, \mathsf{md}), q_{b,\mathsf{md}} := 0$	Oracle $\textsc{Verify}(t, [\mathsf{md}], \sigma)$
$(t_i, \mathsf{md}_i, \sigma_i)_{i \in [\ell+1]} \leftarrow \mathcal{A}^{\textsc{Sign},\textsc{Verify},\textsc{Read}}(\mathsf{crs}, \mathsf{pp})$	**return** $\mathsf{AT.Vf}(\mathsf{sk}\|\mathsf{vk}, t, [\mathsf{md}], \sigma)$
return $(\forall b \in \{0,1\}\ \forall \mathsf{md}, q_{b,\mathsf{md}} \leq \ell$ **and**	Oracle $\textsc{Read}(t, \sigma)$
$\forall i \neq j$ **in** $[\ell+1]$ $(t_i, \mathsf{md}_i, \sigma_i) \neq (t_j, \mathsf{md}_j, \sigma_j)$	**return** $\mathsf{AT.ReadBit}(\mathsf{sk}, t, [\mathsf{md}], \sigma)$
and $\exists (b, \mathsf{md}) \in \{0,1\} \times \{\mathsf{md}\} : \forall i \in [\ell+1],$	
$\mathsf{AT.ReadBit}(\mathsf{sk}, t_i, \sigma_i) = b$ **and**	
$\mathsf{AT.Vf}(\mathsf{sk}\|\mathsf{vk}, t_i, [\mathsf{md}], \sigma_i) = \mathbf{true})$	

Fig. 1. One-more unforgeability with metadata.

Game $\text{UNLINK}_{\mathsf{AT},\mathcal{A},m,[b],[\mathsf{md}]}(\lambda)$	Oracle $\text{USER}_0()$
$(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{AT.Setup}(1^\lambda)$	$q_0 := q_0 + 1$
$(\mathsf{st},\mathsf{pp},[\mathsf{vk}]) \leftarrow \mathcal{A}(\mathsf{crs},[b],[\mathsf{md}])$	$t_{q_0} \leftarrow\!\!\$\ \{0,1\}^\lambda$
$q_0 := 0; q_1 := 0, \mathcal{Q} := \emptyset$	$(\mathsf{msg}_{q_0},\mathsf{st}_{q_0}) \leftarrow \mathsf{AT.User}_0(\mathsf{pp},[\mathsf{vk}],t_{q_0},[\mathsf{md}'])$
$(\mathsf{st},(\mathsf{msg}_i)_{i\in\mathcal{Q}}) \leftarrow \mathcal{A}^{\text{USER}_0,\text{USER}_1}(\mathsf{st})$	$\mathcal{Q} := \mathcal{Q} \cup \{q_0\}$
if $\mathcal{Q} = \emptyset$ **then return** 0	**return** $(q_0, \mathsf{msg}_{q_0})$
$j \leftarrow\!\!\$\ \mathcal{Q}; \mathcal{Q} = \mathcal{Q} \setminus \{j\}$	Oracle $\text{USER}_1(j,\mathsf{msg})$
$\sigma_j \leftarrow \mathsf{AT.User}_1(\mathsf{st}_j,[\mathsf{vk}],\mathsf{msg}_j,[\mathsf{md}])$	**if** $j \notin \mathcal{Q}$ **then**
for $i \in \mathcal{Q}$	**return** $\perp$
$\sigma_i \leftarrow \mathsf{AT.User}_1(\mathsf{st}_i,[\mathsf{vk}],\mathsf{msg}_i,[\mathsf{md}])$	$\sigma \leftarrow \mathsf{AT.User}_1(\mathsf{st}_j,[\mathsf{vk}],\mathsf{msg},[\mathsf{md}'])$
$\phi \leftarrow\!\!\$\ \mathcal{S}_\mathcal{Q}$	**if** $\sigma \neq \perp$ **then**
$j' \leftarrow \mathcal{A}(\mathsf{st},(t_j,\sigma_j),(t_{\phi(i)},\sigma_{\phi(i)})_{i\in\mathcal{Q}})$	$\mathcal{Q} := \mathcal{Q} \setminus \{j\}$
return $q_0 - q_1 \geq m$ **and** $j' = j$	$q_1 := q_1 + 1$
	return σ

Fig. 2. Public-key unlinkability with fixed metadata. If X is a set, then $\mathcal{S}_X$ is the symmetric group of X.

3 Anonymous Token Protocols

The Privacy Pass protocol [30] and its siblings [39, 44] are based on verifiable oblivious pseudo random functions (VOPRF). Here, a user holds some secret input x and the signer holds a secret key k and they evaluate the function F obliviously such that the user learns $F(x,k)$ but nothing about k, and the signer learns nothing about the input x nor the output $F(x,k)$. Additionally, the user is ensured that the function is evaluated by the correct secret key.

We give three protocols for anonymous tokens (AT) with 1) public metadata, 2) public and private metadata, and 3) public metadata and public verifiability, respectively, constructed from the same framework.

At the core of our protocols lies a verifiable key transformation. Let $d := \mathtt{H}_m(\mathsf{md})$ and the curve point $U := [d]G + K$, where G is a public generator and K is the public key with a corresponding private key k. Let $e = (d+k)^{-1}$ be the new signing key and $W' = [e]T'$. Notice the relation

$$\mathsf{KT} : \log_G([d]G + K) = (d+k) = \log_{W'} T'. \tag{1}$$

We give background on elliptic curves in Appendix A, and detail zero-knowledge proofs for equal discrete logarithms, AND-proofs and OR-proofs in Appendix B.

3.1 Secure Key Transformation

We argue that the key-transformation from k to e is secure against one-more unforgeability attacks. Several papers has been written using this transforma-

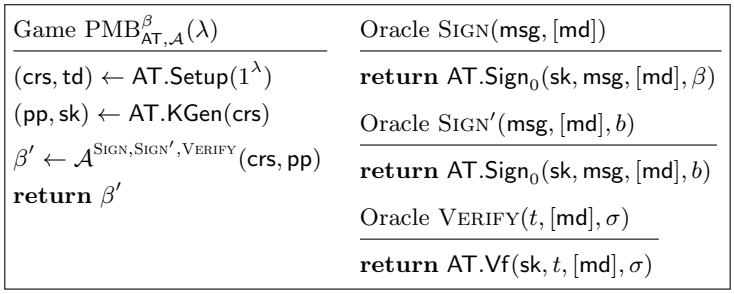

Fig. 3. Game for private metadata bit for anonymous tokens.

tion. Boneh and Boyen [11] shows that this transformation is secure against a non-adaptive attacker for arbitrary metadata md when used for signatures. Furthermore, Dodis and Yampolskiy [31] shows that this transformation is secure against active attackers when the set of possible metadata values is small, and give applications to PRFs. However, these works only prove security with respect to a fixed generators, while our construction signs arbitrary new generators in each execution of the protocol. Recently, Tyagi *et al.* [50] proved that this transformation is secure against an active attacker with respect to arbitrary generators and arbitrary set of metadata. They reduce the security of the transform to a new one-more gap strong inversion Diffie-Hellman problem (see Appendix A.1). They also show that this new problem is equivalent to the simpler q-DL assumption.

We summarize these results in a lemma:

Lemma 1. *Let* AT *be a scheme with keys* $(\mathsf{pk}, \mathsf{vk})$ *with security property* P *within adversarial advantage* $\mathsf{Adv}^{\mathrm{p}}_{\mathsf{AT},\mathcal{A}}(\lambda)$*, and assume we can prove the relation in Equation 1 within adversarial advantage* $\mathsf{Adv}^{\mathrm{rel}}_{\mathsf{KT},\mathcal{A}}(\lambda)$*. Then* $\mathcal{A}$ *has advantage* $\mathsf{Adv}^{\mathrm{p}}_{\mathsf{AT},\mathcal{A}}(\lambda) + \mathsf{Adv}^{\mathrm{rel}}_{\mathsf{KT},\mathcal{A}}(\lambda)$ *against property* P *in the scheme* AT *with transformed keys* $(\{e = (\mathsf{md} + \mathsf{sk})^{-1}, [e]G\})$.

3.2 Anonymous Tokens with Public Metadata

In Figure 4 we present an extension of Privacy Pass [30] with public metadata. The protocol is designated verifier, as the secret key is needed to verify tokens.

Setup and Key Generation. Let λ be the security parameter, let p be a prime and let E be an elliptic curve group of order p with generator G. Let $\mathtt{H}_t : \{0,1\}^* \to E$ and $\mathtt{H}_m : \{0,1\}^* \to \mathbb{Z}_p$ be hash functions, and assume that group elements and integers can be encoded uniquely as strings. Furthermore, let metadata md be an element of a public set of valid strings. Finally, let $\mathsf{sk} := k \leftarrow\!\!\$\ \mathbb{Z}_p^*$ be the signing key, and let $\mathsf{pk} := K := [k]G$ be the public key. We consider $G, E, p, \mathtt{H}_t, \mathtt{H}_m$ and K to be implicit knowledge in the protocol in Fig. 4.

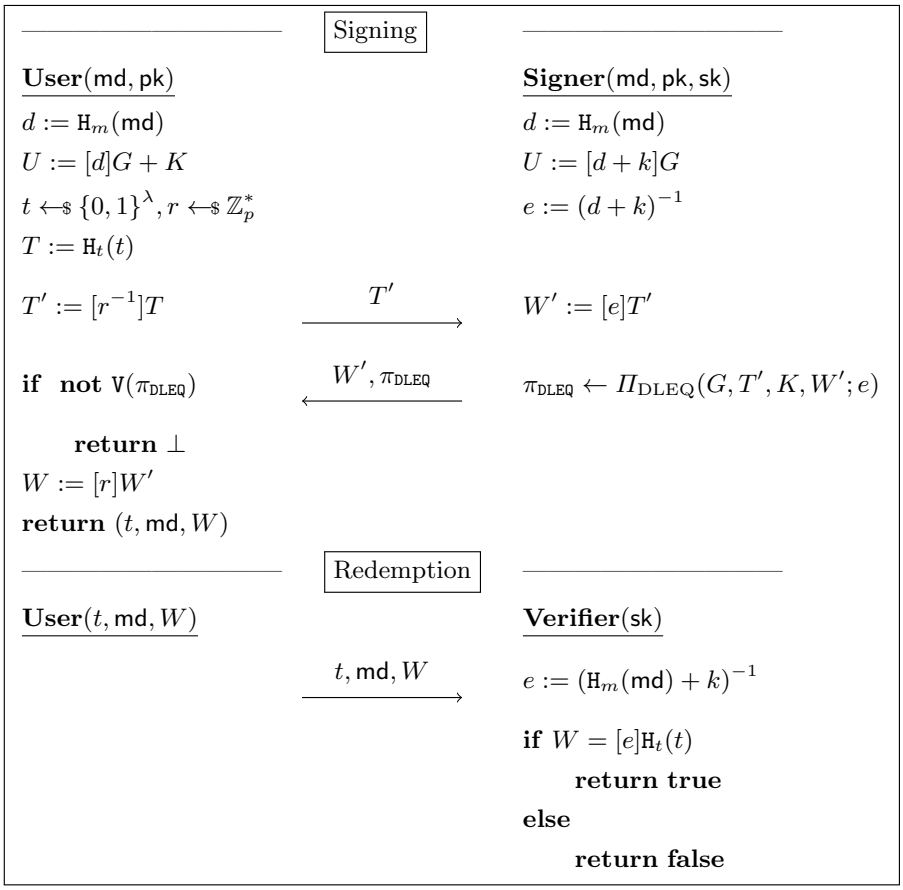

Fig. 4. Designated verifier anonymous tokens with public metadata. Our protocol is a direct extension of Privacy Pass [30].

Signing and Verification. The anonymous tokens protocol in Figure 4 uses the Π_{DLEQ}-protocol defined in Appendix B.1. The signer computes a proof $\pi_{\mathtt{DLEQ}} := (c, z)$ of equality of discrete logarithms by instantiating the protocol $\Pi_{\mathrm{DLEQ}}(G, T', K, W'; e)$. Given the public parameters G and K, and $U := [d]G + K$, this is a proof that $\log_G U = d + k = \log_{W'} T'$. This proves that $W' = [e]T'$, where $e := (d+k)^{-1}$, is computed correctly with respect to d and K. To verify, the user instantiates the verification algorithm, denoted by $\mathtt{V}(\pi_{\mathtt{DLEQ}})$.

Theorem 1 (Completeness). *The anonymous token protocol with public metadata in Figure 4 is complete according to Definition 2.*

Proof. The completeness follows from expanding W:

$$W = [r]W' = [r][e]T' = [r][e][r^{-1}]T = [e]\mathtt{H}_t(t).$$

Theorem 2 (Unforgeability). *The anonymous token protocol with public metadata in Figure 4 achieve one-more unforgeability with respect to Definition 4.*

Proof. Using the key transformation as described in Lemma 1, the security of the protocol reduces to the security of the one-more gap strong inversion Diffie-Hellman game as shown in Figure 6. The advantage of an attacker follows from Definition 7 and is proven secure by Tyagi *et al.* [50, Theorem 1].

Theorem 3 (Unlinkability). *Fix metadata* md*. Within the set defined by all tokens using* md*, the anonymous token protocol with public metadata in Figure 4 achieve unlinkability with respect to Definition 5.*

Proof. This proof is identical to [30, Theorem 1]: As we sample $r \leftarrow\!\!\$\ \mathbb{Z}_p$ uniformly at random, it follows that our protocol is unconditionally unlinkable. Since T is a generator of E, then $T' = [r^{-1}]T$ is uniformly random and contain no information about t nor T. As the signer only sees T', and the verifier only receive t, and they are independent, there is no link between the view of the signer and the view of the verifier.

3.3 Tokens with Private Metadata and Public Verification

Using the same framework, we present anonymous tokens with public and private metadata in Appendix C and anonymous tokens with public metadata and public verification in Appendix D. We provide security proofs for both protocols.

4 Performance and Comparison

In this section, we briefly describe the most efficient anonymous single-use token protocols with public metadata in the literature, for example, to enable batched revocation. We only consider protocols with one round of communication. We compare the protocols with our schemes in Table 1. To streamline the comparison, we assume that all parties know the public metadata, for example that md is the current date, and assume that this implicit knowledge is not sent. We instantiate the schemes with $\lambda = 128$ bits of security. Finally, we present a concrete example to show that we can replace DIT with our protocol in Figure 4 to improve both communication size and computational efficiency.

4.1 Anonymous single-use Tokens with Public Metadata

Privacy Pass. Our protocol in Figure 4 is inspired by Privacy Pass [30], and they have identical structure and communication. The main difference is the change of private key used for signing, and the updated zero-knowledge proof with respect to the new public key, both depending on the public metadata. The zero-knowledge proofs are of the same size, and it follows that the communication sizes are equal. However, Privacy Pass does not allow public metadata unless we have one public key for each valid string of metadata, and hence, to allow for 2^N possible messages md, Privacy Pass must publish 2^N public keys.

DIT: De-Identified Authenticated Telemetry at Scale. DIT [39] is also inspired by Privacy Pass [30], but uses an attribute-based VOPRF to generate new public keys on the fly. To allow for 2^N strings of public metadata, there are two main differences: 1) the public key consists of $N+2$ group elements, and 2) the token consists of an additional N group elements and zero-knowledge proofs to ensure that the correct public key is used in the signature.

Tokens from RSA. Abe and Fujisaki [1] presents a partially blind signature scheme based on RSA. The public exponent e must be at least two bits longer than the public metadata, and we fix this to be of length 130 bits. The user updates the public key to $e_{\mathsf{md}} = e \cdot \tau(\mathsf{md})$, for a public formatting function τ, when they blind the message, and the signer updates the secret key $d_{\mathsf{md}} = (e \cdot \tau(\mathsf{md}))^{-1} \mod N$ when signing. Otherwise, the partially blind signature scheme [1] is similar to the blind signature by Chaum [23].

Tokens with Private Metadata. Kreuter *et al.* [44] presents an extension of Privacy Pass [30] to include private metadata. They publish two public keys, and the signer proves in zero-knowledge that the token is signed with one of the corresponding private keys. To ensure metadata privacy, each token is randomized based on a fresh seed s that is given to the user, and hence, the signature consists of a seed, a group element, and a proof. The token consists of the initial seed t in addition to two group elements. Like Privacy Pass, this protocol must publish a new pair of public keys for each valid string of metadata.

4.2 Comparison

We present a comparison of schemes in Table 1, where we focus on communication complexity. We note that both RSA and pairing based cryptography is usually slower than elliptic curve cryptography, in addition to requiring larger parameters. We also note that the updated keys in our protocols are only dependent on the secret key and the metadata, and can often be pre-computed. We conclude that when allowing for batched token-revocation, our protocols are more efficient than the state of the art in all categories.

While RSA and elliptic curve cryptography are primitives implemented in all mainstream cryptographic libraries, there are few trustworthy implementations of pairings. Even though there exists a few implementations[6], they are mostly for academic use, maybe except for the implementation in Rust used by Zcash[7]. We refer to [50, Table 1] for a comparison in computation between some protocols.

4.3 Telemetry Collection in WhatsApp

DIT [39] was designed to allow users of WhatsApp to anonymously report telemetry data to Facebook. We present a concrete comparison to our protocols in Table 2. Here, we assume that Facebook wants to update their public keys only once a year, rotate signing keys every day, and only sign one token per user each day. We fix a year and encode public metadata as strings "YYYY-MM-DD".

[6] Pairings: https://hackmd.io/@zkteam/eccbench
[7] Zcash: https://github.com/zkcrypto/bls12_381

Public Metadata (PM)	PubKey	Request	Signature	Token
Privacy Pass [30]	$257 \cdot 2^N$	257	769	385
DIT [39]	$257 \cdot (N+2)$	257	$769 \cdot (N+1)$	385
Our scheme (Figure 4)	257	257	769	385
PM + Private Metadata	PubKey	Request	Signature	Token
Kreuter *et al.* [44]	$514 \cdot 2^N$	257	1921	642
Our Scheme (Figure 8)	1028	257	3203	642
PM + Public Verifiability	PubKey	Request	Signature	Token
Abe and Fujisaki [1]	3202	3072	3072	3200
Our scheme (Figure 9)	763	382	382	510

Table 1. Size given in bits. We compare the schemes for 128 bits of security, allowing for 2^N strings md of metadata. Token seed t is of size 128 bits, and metadata md is implicit knowledge. Privacy Pass, DIT, Kreuter *et al.* and our protocols in Fig 4 and Fig 8 are instantiated with curve x25519 [8], Abe and Fujisaki is instantiated with RSA-3072 and our protocol in Fig 9 is instantiated with BLS12-381 [48].

Privacy Pass [30] is very efficient in terms of communication, but requires one public key per day. Hence, the public key is of size 93805 bits over a year of 365 days, that is, approximately 12 KB. An alternative method to download all keys and store them until usage is to use a Merkle-tree for key-transparency and give paths corresponding to the current public key as a part of each signature. Then, the public key consists of the root of size 256 bits, while each signature consists of $\lceil \log_2(365) \rceil = 9$ hashes of 256 bits in addition to the public key, the token, and the zero-knowledge proof. We give both instantiations in the table, and denote the alternative protocol as Privacy Pass+.

Our scheme in Figure 4 has the smallest overall communication complexity of all schemes. It offers much smaller keys than Privacy Pass, and much smaller signatures than Privacy Pass+ and DIT, saving up to 90 % in communication. If all 2 billion users of WhatsApp report their telemetry every day, our scheme in Figure 4 would save more than 1.7 TB of communication for the Facebook servers on a daily basis compared to the current implementation of DIT.

Our scheme in Figure 9 offers similar improvements to communication, in addition to public verifiability using pairings, but at the cost of less standardized cryptography and less efficient computation.

Protocol	PubKey	Request	Signature	Token
Privacy Pass [30]	93805	257	769	385
Privacy Pass+	256	257	3330	385
DIT [39]	2313	257	7690	385
Our scheme (Fig 4)	257	257	769	385
Our scheme (Fig 9)	763	382	382	510

Table 2. Size given in bits. We compare Privacy Pass, DIT, and the protocols in Fig 4 and Fig 9 with daily key-rotation in a year, signing one token at a time.

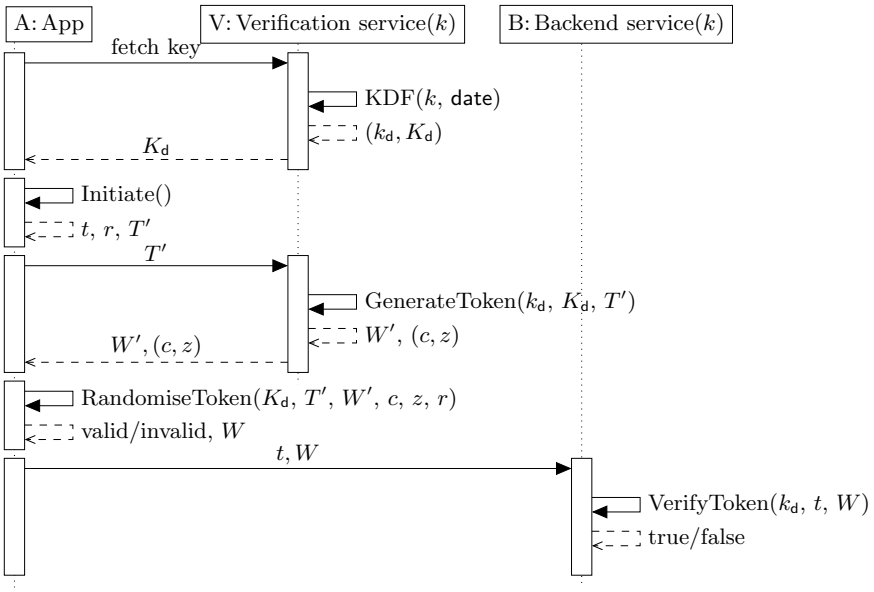

Fig. 5. A sequence diagram of anonymous tokens in the Norwegian app Smittestopp.

5 Application to Contact Tracing

As nations started adopting digital contact tracing during the COVID-19 pandemic, privacy experts warned that such systems could enable the collection of people's contact graphs. The dp^{3}t protocol [49] was eventually adopted as the *de facto* method for digital contact tracing through its implementation and deployment in iOS and Android as the Exposure Notification System (ENS).

We provide a brief overview of the basic dp^{3}t idea in order to put our contribution into context. The protocol is instantiated on each participating phone, which generates a random key (Temporary Exposure Key, TEK) every day. The TEK is used to generate new Rotating Proximity Identifiers (RPI) every 10–20 minutes, which is then broadcast from the phone using Bluetooth Low Energy (BLE). Other phones in the proximity store any RPI they hear.

If Alice tests positive for COVID-19 she can upload her TEKs (now renamed to diagnosis keys, DK) along with her BLE transmission strength to a health authority bulletin board. Bob's phone regularly checks the board to see if there is a sufficiently large overlap between published the DKs and the RPIs stored locally, and with sufficiently low difference between transmission strength and received strength. If this is the case, then Bob is given a suitable alert to let him know that he most likely has been in close vicinity of an infected individual, and should follow any advice given by the health authorities.

The process of uploading TEKs should depend on some sort of authorization. The dp^{3}t documentation describes a simplified model where a doctor receives the test results, and sends the patient an SMS with a short upload code. Now,

this process may take precious person-hours during a pandemic. Some countries have therefore opted to connect their exposure notification with already existing centralized registries of positive test results, e.g., Norway, Denmark, and Estonia.

When starting the upload process, the user is prompted to log in to some government service ("verification"). Once the user has identified herself, the service makes a query to the relevant health registry. The service returns an access token to the app if there exists a recent positive test, which is then used to upload the keys to "backend". Unfortunately, this token may create an identifiable link from the meant-to-be-anonymous database of DKs, and unique identities in the health registry. Using anonymous single-use tokens, one can break this link (up to traffic analysis, e.g., logging timings and network addresses).

The Norwegian Institute of Public Health (NIPH) wanted the tokens to be timestamped in order to avoid users posting severely delayed keys: this would have allowed an attacker to get well again, move back out among other people, and only then upload to the backend service. Notice that merely tying the token to keys – e.g., by using a hash of the TEKs as the token seed t – would not avoid this attack, as those could have been generated and stored until the time of the attack. As a result, it was decided that the keys should be rotated regularly.

The original Privacy Pass protocol was reimplemented as a reusable C# package, to ease the integration into the Norwegian contact tracing app Smittestopp. The verification and backend services keep a master secret key k, and generate daily keys from some $\mathsf{KDF}(k, \mathsf{date})$. The public key is posted from the verification service. The full integration of anonymous tokens is described in Figure 5.

We finally note that this key distribution method suffers from a potential attack by a dishonest verification service that could serve special public keys to track individuals. It is, however, detectable by the users if they share their view of the public keys with each other to ensure consistency. The current solution was accepted by all involved stakeholders due to limited time and a weighting of the practical risk against the potential reward. The challenges with respect to key-rotation and key-sharing strongly motivated the authors' work in Section 3.

6 Conclusion

In this work, we have updated the definitions for anonymous single-use tokens to also include public metadata, and we have constructed three protocols that satisfy these definitions. Additionally, we combine public metadata with either private metadata or public verifiability, and show that all instantiations are efficient in practice. For situations with frequent key-rotation, we show that our protocols can save up to 90 % in communication over the state of the art. Furthermore, our protocols fit nicely into the Privacy Pass framework, which makes it easy to incorporate our contributions in the ongoing standardization processes by IETF and W3C, solving an open problem.

We also provide a description of how anonymous one-time tokens can be used to improve the user's privacy in contact tracing applications, and implemented this into the solution used in Norway. The app has more than one million users

at the time of writing[8]. As the Norwegian app is built on top of the same code base as the Danish app, we consider it to be easy to extend the adaption of anonymous tokens to their app, and most likely others as well.

We would also like to suggest new use-cases for anonymous tokens. For example, anonymous tokens can improve the privacy of users traveling with public transport. Bus or train companies may require patrons to verify their period tickets for each journey, perhaps primarily to analyze traffic data. However, this can easily reveal the routes of single users while traveling in-between their home and workplace, but also to the abortion clinic, their church or to a public demonstration etc. If all travelers with valid tickets are given a series of tokens (e.g., with public metadata being the date or week or month the ticket is valid), then these can be redeemed when boarding. This way, the companies get the statistics they are interested in, without invading the user's privacy. In general, any systems with leveled authenticated login but anonymous actions can make use of our protocols, e.g., systems with electronic locks that only care if the user has certain privileges or not. We also note that Tyagi *et al.* [50] detail applications of a construction similar to ours to reduce key management complexity in the OPAQUE password authenticated key exchange protocol, and to ensure stronger security for password breach alerting services.

Finally, we would like to see improvements in three directions. Firstly, the zero-knowledge proofs used by the anonymous tokens protocol with public and private metadata in Figure 8 are much larger than the ones by Kreuter *et al.* [44], in contrast to our protocol with public metadata in Figure 4 achieving the exact same communication cost as Privacy Pass [30]. In particular, we would like to reduce the number of proofs and extra group elements in the protocol in Appendix C. Secondly, we would like to provide protocols free of zero-knowledge proofs, to reduce the communication and computational cost, as provided in [44, Section 7]. Finally, we would like to extend our protocols to achieve post-quantum security, continuing the work by Albrecht *et al.* [4] on lattice-based protocols.

Acknowledgments. The authors are very grateful to Henrik Walker Moe (Bekk Consulting AS) for stellar collaboration during the C# implementation phase. The final integration into Smittestopp was primarily a collaboration between Henrik, Johannes Brodwall (Sopra Steria) and Sindre Møgster Braaten (NIPH), with the authors and others as close consultants. We thank Nirvan Tyagi, Sofia Celi, Thomas Ristenpart and Christopher Wood for pointing out a flaw in the security game and unforgeability proof in an earlier version of this paper. The second author is grateful to the students Teodor Dahl Knutsen and Tallak Mannum for many useful comments to an earlier version of the manuscript. We would also like thank the anonymous reviewers at PETS 2021 and Financial Crypto 2022 for their feedback which greatly improved the presentation of this paper.

[8] Smittestopp: https://www.fhi.no/om/smittestopp/nokkeltall-fra-smittestopp, last accessed 2021-12-01.

References

1. Abe, M., Fujisaki, E.: How to date blind signatures. In: Kim, K., Matsumoto, T. (eds.) Advances in Cryptology – ASIACRYPT'96. Lecture Notes in Computer Science, vol. 1163, pp. 244–251. Springer, Heidelberg (Nov 1996). https://doi.org/10.1007/BFb0034851
2. Abe, M., Okamoto, T.: Provably secure partially blind signatures. In: Bellare, M. (ed.) Advances in Cryptology – CRYPTO 2000. Lecture Notes in Computer Science, vol. 1880, pp. 271–286. Springer, Heidelberg (Aug 2000). https://doi.org/10.1007/3-540-44598-6_17
3. Akagi, N., Manabe, Y., Okamoto, T.: An efficient anonymous credential system. In: Tsudik, G. (ed.) FC 2008: 12th International Conference on Financial Cryptography and Data Security. Lecture Notes in Computer Science, vol. 5143, pp. 272–286. Springer, Heidelberg (Jan 2008)
4. Albrecht, M.R., Davidson, A., Deo, A., Smart, N.P.: Round-optimal verifiable oblivious pseudorandom functions from ideal lattices. Cryptology ePrint Archive, Report 2019/1271 (2019), https://eprint.iacr.org/2019/1271
5. Baldimtsi, F., Lysyanskaya, A.: Anonymous credentials light. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 2013: 20th Conference on Computer and Communications Security. pp. 1087–1098. ACM Press (Nov 2013). https://doi.org/10.1145/2508859.2516687
6. Barreto, P.S.L.M., Lynn, B., Scott, M.: Constructing elliptic curves with prescribed embedding degrees. In: Cimato, S., Galdi, C., Persiano, G. (eds.) SCN 02: 3rd International Conference on Security in Communication Networks. Lecture Notes in Computer Science, vol. 2576, pp. 257–267. Springer, Heidelberg (Sep 2003). https://doi.org/10.1007/3-540-36413-7_19
7. Bellare, M., Namprempre, C., Pointcheval, D., Semanko, M.: The power of RSA inversion oracles and the security of Chaum's RSA-based blind signature scheme. In: Syverson, P.F. (ed.) FC 2001: 5th International Conference on Financial Cryptography. Lecture Notes in Computer Science, vol. 2339, pp. 319–338. Springer, Heidelberg (Feb 2002)
8. Bernstein, D.J.: Curve25519: high-speed elliptic curve cryptography (2005), https://cr.yp.to/ecdh.html
9. Blazy, O., Pointcheval, D., Vergnaud, D.: Compact round-optimal partially-blind signatures. In: Visconti, I., Prisco, R.D. (eds.) SCN 12: 8th International Conference on Security in Communication Networks. Lecture Notes in Computer Science, vol. 7485, pp. 95–112. Springer, Heidelberg (Sep 2012). https://doi.org/10.1007/978-3-642-32928-9_6
10. Boldyreva, A.: Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In: Desmedt, Y. (ed.) PKC 2003: 6th International Workshop on Theory and Practice in Public Key Cryptography. Lecture Notes in Computer Science, vol. 2567, pp. 31–46. Springer, Heidelberg (Jan 2003). https://doi.org/10.1007/3-540-36288-6_3
11. Boneh, D., Boyen, X.: Short signatures without random oracles. In: Cachin, C., Camenisch, J. (eds.) Advances in Cryptology – EUROCRYPT 2004. Lecture Notes in Computer Science, vol. 3027, pp. 56–73. Springer, Heidelberg (May 2004). https://doi.org/10.1007/978-3-540-24676-3_4
12. Boneh, D., Lynn, B., Shacham, H.: Short signatures from the Weil pairing. In: Boyd, C. (ed.) Advances in Cryptology – ASIACRYPT 2001. Lecture Notes in Computer Science, vol. 2248, pp. 514–532. Springer, Heidelberg (Dec 2001). https://doi.org/10.1007/3-540-45682-1_30

13. Burns, J., Moore, D., Ray, K., Speers, R., Vohaska, B.: EC-OPRF: Oblivious pseudorandom functions using elliptic curves. Cryptology ePrint Archive, Report 2017/111 (2017), http://eprint.iacr.org/2017/111
14. Camenisch, J., Groß, T.: Efficient attributes for anonymous credentials. In: Ning, P., Syverson, P.F., Jha, S. (eds.) ACM CCS 2008: 15th Conference on Computer and Communications Security. pp. 345–356. ACM Press (Oct 2008). https://doi.org/10.1145/1455770.1455814
15. Camenisch, J., Hohenberger, S., Lysyanskaya, A.: Compact e-cash. In: Cramer, R. (ed.) Advances in Cryptology – EUROCRYPT 2005. Lecture Notes in Computer Science, vol. 3494, pp. 302–321. Springer, Heidelberg (May 2005). https://doi.org/10.1007/11426639_18
16. Camenisch, J., Kohlweiss, M., Soriente, C.: An accumulator based on bilinear maps and efficient revocation for anonymous credentials. In: Jarecki, S., Tsudik, G. (eds.) PKC 2009: 12th International Conference on Theory and Practice of Public Key Cryptography. Lecture Notes in Computer Science, vol. 5443, pp. 481–500. Springer, Heidelberg (Mar 2009). https://doi.org/10.1007/978-3-642-00468-1_27
17. Camenisch, J., Lysyanskaya, A.: An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In: Pfitzmann, B. (ed.) Advances in Cryptology – EUROCRYPT 2001. Lecture Notes in Computer Science, vol. 2045, pp. 93–118. Springer, Heidelberg (May 2001). https://doi.org/10.1007/3-540-44987-6_7
18. Camenisch, J., Lysyanskaya, A.: A signature scheme with efficient protocols. In: Cimato, S., Galdi, C., Persiano, G. (eds.) SCN 02: 3rd International Conference on Security in Communication Networks. Lecture Notes in Computer Science, vol. 2576, pp. 268–289. Springer, Heidelberg (Sep 2003). https://doi.org/10.1007/3-540-36413-7_20
19. Camenisch, J., Lysyanskaya, A.: Signature schemes and anonymous credentials from bilinear maps. In: Franklin, M. (ed.) Advances in Cryptology – CRYPTO 2004. Lecture Notes in Computer Science, vol. 3152, pp. 56–72. Springer, Heidelberg (Aug 2004). https://doi.org/10.1007/978-3-540-28628-8_4
20. Camenisch, J., Van Herreweghen, E.: Design and implementation of the idemix anonymous credential system. In: Atluri, V. (ed.) ACM CCS 2002: 9th Conference on Computer and Communications Security. pp. 21–30. ACM Press (Nov 2002). https://doi.org/10.1145/586110.586114
21. Chase, M., Meiklejohn, S., Zaverucha, G.: Algebraic MACs and keyed-verification anonymous credentials. In: Ahn, G.J., Yung, M., Li, N. (eds.) ACM CCS 2014: 21st Conference on Computer and Communications Security. pp. 1205–1216. ACM Press (Nov 2014). https://doi.org/10.1145/2660267.2660328
22. Chase, M., Perrin, T., Zaverucha, G.: The signal private group system and anonymous credentials supporting efficient verifiable encryption. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 20: 27th Conference on Computer and Communications Security. pp. 1445–1459. ACM Press (Nov 2020). https://doi.org/10.1145/3372297.3417887
23. Chaum, D.: Blind signatures for untraceable payments. In: Chaum, D., Rivest, R.L., Sherman, A.T. (eds.) Advances in Cryptology – CRYPTO'82. pp. 199–203. Plenum Press, New York, USA (1982)
24. Chaum, D.: Blind signature system. In: Chaum, D. (ed.) Advances in Cryptology – CRYPTO'83. p. 153. Plenum Press, New York, USA (1983)
25. Chaum, D., Pedersen, T.P.: Wallet databases with observers. In: Brickell, E.F. (ed.) Advances in Cryptology – CRYPTO'92. Lecture Notes in Computer Science,

vol. 740, pp. 89–105. Springer, Heidelberg (Aug 1993). https://doi.org/10.1007/3-540-48071-4_7
26. Chen, X., Zhang, F., Mu, Y., Susilo, W.: Efficient provably secure restrictive partially blind signatures from bilinear pairings. In: Di Crescenzo, G., Rubin, A. (eds.) FC 2006: 10th International Conference on Financial Cryptography and Data Security. Lecture Notes in Computer Science, vol. 4107, pp. 251–265. Springer, Heidelberg (Feb / Mar 2006)
27. Chow, S.S.M., Hui, L.C.K., Yiu, S.M., Chow, K.P.: Two improved partially blind signature schemes from bilinear pairings. In: Boyd, C., Nieto, J.M.G. (eds.) ACISP 05: 10th Australasian Conference on Information Security and Privacy. Lecture Notes in Computer Science, vol. 3574, pp. 316–328. Springer, Heidelberg (Jul 2005)
28. Davidson, A.: Supporting the latest version of the privacy pass protocol. https://blog.cloudflare.com/supporting-the-latest-version-of-the-privacy-pass-protocol, (Accessed 01-December-2021)
29. Davidson, A., Goldberg, I., Sullivan, N., Tankersley, G., Valsorda, F.: Privacy pass: A privacy-enhancing protocol and browser extension. https://privacypass.github.io, (Accessed 01-December-2021)
30. Davidson, A., Goldberg, I., Sullivan, N., Tankersley, G., Valsorda, F.: Privacy pass: Bypassing internet challenges anonymously. Proceedings on Privacy Enhancing Technologies **2018**(3), 164–180 (Jul 2018). https://doi.org/10.1515/popets-2018-0026
31. Dodis, Y., Yampolskiy, A.: A verifiable random function with short proofs and keys. In: Vaudenay, S. (ed.) PKC 2005: 8th International Workshop on Theory and Practice in Public Key Cryptography. Lecture Notes in Computer Science, vol. 3386, pp. 416–431. Springer, Heidelberg (Jan 2005). https://doi.org/10.1007/978-3-540-30580-4_28
32. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) Advances in Cryptology – CRYPTO'86. Lecture Notes in Computer Science, vol. 263, pp. 186–194. Springer, Heidelberg (Aug 1987). https://doi.org/10.1007/3-540-47721-7_12
33. Freedman, M.J., Ishai, Y., Pinkas, B., Reingold, O.: Keyword search and oblivious pseudorandom functions. In: Kilian, J. (ed.) TCC 2005: 2nd Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 3378, pp. 303–324. Springer, Heidelberg (Feb 2005). https://doi.org/10.1007/978-3-540-30576-7_17
34. Fuchsbauer, G., Hanser, C., Kamath, C., Slamanig, D.: Practical round-optimal blind signatures in the standard model from weaker assumptions. In: Zikas, V., De Prisco, R. (eds.) SCN 16: 10th International Conference on Security in Communication Networks. Lecture Notes in Computer Science, vol. 9841, pp. 391–408. Springer, Heidelberg (Aug / Sep 2016). https://doi.org/10.1007/978-3-319-44618-9_21
35. Fuchsbauer, G., Hanser, C., Slamanig, D.: Practical round-optimal blind signatures in the standard model. In: Gennaro, R., Robshaw, M.J.B. (eds.) Advances in Cryptology – CRYPTO 2015, Part II. Lecture Notes in Computer Science, vol. 9216, pp. 233–253. Springer, Heidelberg (Aug 2015). https://doi.org/10.1007/978-3-662-48000-7_12
36. Hanzlik, L., Slamanig, D.: With a little help from my friends: Constructing practical anonymous credentials. In: Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security. p. 2004–2023. CCS '21, Association for Computing Machinery, New York, NY, USA (2021). https://doi.org/10.1145/3460120.3484582, https://doi.org/10.1145/3460120.3484582

37. Henry, R.: Efficient Zero-Knowledge Proofs and Applications. Ph.D. thesis, University of Waterloo (2014), http://hdl.handle.net/10012/8621
38. Henry, R., Goldberg, I.: Batch proofs of partial knowledge. In: Jacobson Jr., M.J., Locasto, M.E., Mohassel, P., Safavi-Naini, R. (eds.) ACNS 13: 11th International Conference on Applied Cryptography and Network Security. Lecture Notes in Computer Science, vol. 7954, pp. 502–517. Springer, Heidelberg (Jun 2013). https://doi.org/10.1007/978-3-642-38980-1_32
39. Huang, S., Iyengar, S., Jeyaraman, S., Kushwah, S., Lee, C.K., Luo, Z., Mohassel, P., Raghunathan, A., Shaikh, S., Sung, Y.C., Zhang, A.: Dit: De-identified authenticated telemetry at scale. Tech. rep., Facebook Inc., https://research.fb.com/wp-content/uploads/2021/04/DIT-De-Identified-Authenticated-Telemetry-at-Scale_final.pdf (April 2021)
40. Internet Engineering Task Force: Privacy pass datatracker. https://datatracker.ietf.org/wg/privacypass, (Accessed 01-December-2021)
41. Iyengar, S., Taubeneck, E.: Fraud resistant, privacy preserving reporting using blind signatures. https://github.com/siyengar/private-fraud-prevention, (Accessed 01-December-2021)
42. Jarecki, S., Kiayias, A., Krawczyk, H.: Round-optimal password-protected secret sharing and T-PAKE in the password-only model. In: Sarkar, P., Iwata, T. (eds.) Advances in Cryptology – ASIACRYPT 2014, Part II. Lecture Notes in Computer Science, vol. 8874, pp. 233–253. Springer, Heidelberg (Dec 2014). https://doi.org/10.1007/978-3-662-45608-8_13
43. Jarecki, S., Krawczyk, H., Xu, J.: OPAQUE: An asymmetric PAKE protocol secure against pre-computation attacks. In: Nielsen, J.B., Rijmen, V. (eds.) Advances in Cryptology – EUROCRYPT 2018, Part III. Lecture Notes in Computer Science, vol. 10822, pp. 456–486. Springer, Heidelberg (Apr / May 2018). https://doi.org/10.1007/978-3-319-78372-7_15
44. Kreuter, B., Lepoint, T., Orrù, M., Raykova, M.: Anonymous tokens with private metadata bit. In: Micciancio, D., Ristenpart, T. (eds.) Advances in Cryptology – CRYPTO 2020, Part I. Lecture Notes in Computer Science, vol. 12170, pp. 308–336. Springer, Heidelberg (Aug 2020). https://doi.org/10.1007/978-3-030-56784-2_11
45. Kreuter, B., Lepoint, T., Orru, M., Raykova, M.: Efficient anonymous tokens with private metadata bit. Cryptology ePrint Archive, Report 2020/072 (2020), https://eprint.iacr.org/2020/072
46. Papadopoulos, D., Wessels, D., Huque, S., Naor, M., Včelák, J., Reyzin, L., Goldberg, S.: Making NSEC5 practical for DNSSEC. Cryptology ePrint Archive, Report 2017/099 (2017), http://eprint.iacr.org/2017/099
47. Paquin, C., Zaverucha, G.: U-prove cryptographic specification v1.1 (2013), https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/U-Prove20Cryptographic20Specification20V1.1.pdf
48. S. Yonezawa, S. Chikara, T. Kobayashi and T. Saito: Pairing-Friendly Curves. https://tools.ietf.org/id/draft-yonezawa-pairing-friendly-curves-02.html, (Accessed 01-December-2021)
49. Troncoso, C., et al.: Decentralized privacy-preserving proximity tracing. https://arxiv.org/abs/2005.12273 (2020)
50. Tyagi, N., Celi, S., Ristenpart, T., Sullivan, N., Tessaro, S., Wood, C.A.: A fast and simple partially oblivious prf, with applications. Cryptology ePrint Archive, Report 2021/864 (2021), https://ia.cr/2021/864
51. World Wide Web Consortium: Trust Token API Explainer. https://github.com/WICG/trust-token-api, (Accessed 01-December-2021)

52. Wu, Q., Susilo, W., Mu, Y., Zhang, F.: Efficient partially blind signatures with provable security. In: Gavrilova, M., Gervasi, O., Kumar, V., Tan, C.J.K., Taniar, D., Laganá, A., Mun, Y., Choo, H. (eds.) Computational Science and Its Applications - ICCSA 2006. pp. 345–354. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
53. Zhang, F., Safavi-Naini, R., Susilo, W.: Efficient verifiably encrypted signature and partially blind signature from bilinear pairings. In: Johansson, T., Maitra, S. (eds.) Progress in Cryptology - INDOCRYPT 2003: 4th International Conference in Cryptology in India. Lecture Notes in Computer Science, vol. 2904, pp. 191–204. Springer, Heidelberg (Dec 2003)

A Background on Elliptic Curves

We assume that the reader is familiar with the basics of elliptic curve cryptography. To fix notation, let q be a prime and let $\mathbb{F}_{q^\ell}$ for some $\ell > 0$ be a field of characteristic q. Let E be all points (x, y) that satisfy the elliptic curve equation $y^2 = x^3 + ab + b$ in the algebraic closure of $\mathbb{F}_{q^\ell}$, and let $E(\mathbb{F}_{q^\ell})$ denote the set of all such points in $\mathbb{F}_{q^\ell} \times \mathbb{F}_{q^\ell}$ along with the point at infinity $\mathcal{O}$. By abuse of notation, we often let E be a group of order p inside $E(\mathbb{F}_{q^\ell})$. Define the group law in the usual additive way. In particular, let $[m] : E \to E$ be the multiplication-by-m map, which takes the same role as exponentiation in multiplicative groups. Now follows a brief discussion of the Chosen-Target Gap Diffie-Hellman problem and some zero-knowledge proofs we will need as primitives.

A.1 The One-More Gap Strong Inversion Diffie-Hellman Problem

The strong Diffie-Hellman problem was introduced by Boneh and Boyen [11]. Given a sequence $g, g^x, g^{x^2}, \ldots, g^{x^q}$ from a group $\mathbb{G}$ of prime order p, output a pair $(c, g^{(x+c)^{-1}})$ with $c \in \mathbb{Z}_p$. We now present a variant of this game: the adversary must commit to fixed set of candidates $\{c_i\}$, and may then query an oracle for $B^{(\mathsf{sk}+c_i)^{-1}}$ for arbitrary B, along with an oracle for the decision variant. The adversary wins if it can present $\ell + 1$ correct tuples for a chosen c_i, but only having queried ℓ or less. The details are presented in Figure 6. The definition and game is due to Tyagi *et al.* [50].

Definition 7 ((m, n)-One-More Gap Strong Inversion Diffie-Hellman). *Let m, n be natural numbers, and let $\mathbb{G}$ be a cyclic group of order p with generator g produced by the algorithm $\mathsf{Gen}(1^\lambda)$. Let (m, n)-OM-Gap-SDHI be the game defined in Figure 6. (m, n)-One-More Gap Strong Diffie-Hellman Inversion holds for $\mathbb{G}$ if for any PPT adversary $\mathcal{A}$ and any $\ell \geq 0$,*

$$\mathsf{Adv}^{\text{om-gap-sdhi}}_{\mathsf{Gen},\mathcal{A},\ell}(\lambda) := \Pr[(m, n)\text{-OM-Gap-SDHI}_{\mathsf{Gen},\mathcal{A},\ell}(\lambda) = 1] = \mathsf{negl}(\lambda).$$

Tyagi *et al.* [50] have proven that this assumption is implied by the much simpler q-DL assumption, which asks the adversary to return x, given $g, g^x, g^{x^2}, \ldots, g^{x^q}$.

Game (m,n)-OM-GAP-SDHI$_{\mathsf{Gen},\mathcal{A},\ell}(\lambda)$	Oracle SDH(B,i)
$(\mathbb{G},p,g) \leftarrow \mathsf{Gen}(1^\lambda)$	**if** $i \notin [n]$ **then**
$\mathsf{sk} \leftarrow\!\!\$\ \mathbb{Z}_p$	**return** $\bot$
$(y_i)_{i\in[m]} \leftarrow\!\!\$\ \mathbb{Z}_p^m$	$q_i := q_i + 1$
$(\mathsf{st}_\mathcal{A},(c_i)_{i\in[n]}) \leftarrow \mathcal{A}_1(p,\mathbb{G})$	$Z := B^{(\mathsf{sk}+c_i)^{-1}}$
$\left(\gamma,(Z_i,\alpha_i)_{i\in[\ell+1]}\right) \leftarrow \mathcal{A}_2^{\text{SDH,SDDH}}\left(g,g^{\mathsf{sk}},(g^{y_i})_{i\in[m]};\mathsf{st}_\mathcal{A}\right)$	**return** Z
if $q_\gamma \le \ell$ **and** $(i\neq j \to \alpha_i \neq \alpha_j)$ **then**	Oracle SDDH(Y,Z,i)
return $(Z_i)_{i\in[\ell+1]} = \left(g^{y_{\alpha_i}(\mathsf{sk}+c_\gamma)^{-1}}\right)_{i\in[\ell+1]}$	**return** $Z = Y^{(\mathsf{sk}+c_i)^{-1}}$

Fig. 6. The one-more gap strong inversion Diffie-Hellman security game.

A.2 DDH vs. CDH in Pairings

Let $\mathbb{G}_1$ and $\mathbb{G}_2$ be two cyclic groups of prime order, written additively, and let $\mathbb{G}_T$ be another cyclic group of same prime order, written multiplicatively. A bilinear pairing $\hat{e}$ is a map

$$\hat{e} : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$$

such that the following properties hold:

Bilinearity For all $P_1, P_2 \in \mathbb{G}_1$ and $Q_1, Q_2 \in \mathbb{G}_2$, it holds that $\hat{e}(P_1+P_2,Q_1) = \hat{e}(P_1,Q_1)\hat{e}(P_2,Q_1)$ and $\hat{e}(P_1,Q_1+Q_2) = \hat{e}(P_1,Q_1)e(P_1,Q_2)$.
Non-degeneracy For all $P \neq \mathcal{O}$, $\hat{e}(P,P) \neq 1$.
Computability $\hat{e}$ can be efficiently computed.

The bilinearity property implies that for any scalars a, b, we have $\hat{e}([a]P,[b]Q) = \hat{e}(P,Q)^{ab}$, which is the crucial property used for verification.

Bilinear maps lend themselves to a variant of the well-known Diffie-Hellman problem, the Chosen-Target Gap Diffie-Hellman problem [7]. Even if the adversary is given oracle access to ℓ instances of the Computational Diffie-Hellman (CDH) problem and arbitrary many queries to a Decision Diffie-Hellman (DDH) oracle, it should still not be able to compute the final Diffie-Hellman instance $\ell+1$. We repeat the game and definition by Kreuter *et al.* [44].

Definition 8 (Chosen-Target Gap Diffie-Hellman). *Let $\mathbb{G}$ be a cyclic group of order p with generator G produced by the algorithm $\mathsf{Gen}(1^\lambda)$. Let* CTGDH *be the game defined in Figure 7. Chosen-Target Gap Diffie-Hellman holds for $\mathbb{G}$ if for any* PPT *adversary $\mathcal{A}$ and any $\ell \ge 0$,*

$$\mathsf{Adv}^{\text{ctgdh}}_{\mathsf{Gen},\mathcal{A},\ell}(\lambda) := \Pr[\text{CTGDH}_{\mathsf{Gen},\mathcal{A},\ell}(\lambda) = 1] = \mathsf{negl}(\lambda).$$

Game $\mathrm{CTGDH}_{\mathsf{Gen},\mathcal{A},\ell}(\lambda)$	Oracle $\textsc{Target}(t)$	Oracle $\textsc{Help}(Y)$
$\Gamma = (\mathbb{G}, p, G) \leftarrow \mathsf{Gen}(1^\lambda)$	**if** $t \in \mathcal{Q}$ **then**	$q := q + 1$
$x \leftarrow\!\!\$\ \mathbb{Z}_p; X := [x]G$	$Y := \mathcal{Q}[t]$	**return** $[x]Y$
$q := 0; \mathcal{Q} := [\,]$	**else**	Oracle $\mathrm{DDH}(Y, Z)$
$(t_i, Z_i)_{i\in[\ell+1]} \leftarrow \mathcal{A}^{\textsc{Target},\textsc{Help},\mathrm{DDH}}(\Gamma, X)$	$Y \leftarrow\!\!\$\ \mathbb{G}$	**return** $(Z = [x]Y)$
for $i \in [\ell + 1]$	$\mathcal{Q}[t] := Y$	
if $t_i \notin \mathcal{Q}$ **then return** 0	**return** Y	
$Y_i := \mathcal{Q}[t_i]$		
return $(q \leq \ell$ **and**		
$\forall i \neq j \in [\ell + 1], t_i \neq t_j$ **and**		
$\forall i \in [\ell + 1], [x]Y_i = Z_i)$		

Fig. 7. The Chosen-target gap Diffie-Hellman security game.

B Zero-Knowledge Proofs

B.1 Proof of Equal Discrete Logs

Chaum and Pedersen [25] introduced an elegant honest-verifier zero-knowledge protocol to prove that two group elements have the same discrete logarithm relative to their respective bases, $\log_G K = k = \log_T W$. We describe the protocol loosely to ensure the reader is familiar with the idea. Let $\mathbb{G}$ be a cyclic group of prime order p with independent generators G and T, and let $K := [k]G$, $W := [k]T$ where k is a scalar private to the prover $\mathcal{P}$.

$\mathcal{P}$.1 Choose a random scalar r in the underlying field, compute $A := [r]G$ and $B := [r]T$, and send (A, B) to $\mathcal{V}$.
$\mathcal{V}$.1 Choose a random challenge c modulo p and send it to $\mathcal{P}$.
$\mathcal{P}$.2 Compute the response $z := r - ck$ modulo p, and then send z to $\mathcal{V}$.
$\mathcal{V}$.2 Verify that $A = [z]G + [c]K$ and $B = [z]T + [c]W$.

This protocol satisfies unconditional special soundness and special honest-verifier zero-knowledge. One can make the protocol non-interactive by applying the Fiat-Shamir transformation [32]. The prover queries the oracle on the tuple $(\mathbb{G}, G, T, K, W, A, B)$. In addition, one can reduce communication by sending the oracle response c instead of (A, B), and modifying the final verification step to querying the oracle on $(\mathbb{G}, G, T, K, W, [z]G + [c]K, [z]T + [c]W)$, and then verify that it indeed returns c. We will use a shorthand notation to refer to this proof as $\Pi_{\mathrm{DLEQ}}(G, T, K, W; k)$, meaning that $\log_G([k]G) = \log_W T$.

The proof can be batched for many instances with respect to the same secret scalar using the techniques by Henry [37] as showed in [30, Section 3.2.1].

B.2 AND-Proof of Equal Discrete Logs

Let $\mathbb{G}$ be an additive group of prime order p with independent generators G, H, T, S, and let $K := [k_0]G + [k_1]H$ and $V := [k_0]T + [k_1]S$, where k_0, k_1 are scalars private to the prover $\mathcal{P}$. We want to prove that V is correctly computed with respect to T and S using the same secret scalars as K with respect to G and H. We present a simple protocol to prove this relation, by essentially computing two Chaum-Pedersen proofs in parallel.

$\mathcal{P}$.1 Choose two random scalars r_0, r_1 modulo p. Compute $A := [r_0]G + [r_1]H$, $B := [r_0]T + [r_1]S$, and send (A, B) to $\mathcal{V}$.
$\mathcal{V}$.1 Choose a random challenge c modulo p and send it to $\mathcal{P}$.
$\mathcal{P}$.2 Compute $z_0 := r_0 - ck_0$ and $z_1 := r_1 - ck_1$ modulo p and send them to $\mathcal{V}$.
$\mathcal{V}$.2 Verify that $A = [c]K + [z_0]G + [z_1]H$ and $B = [c]V + [z_0]T + [z_1]S$.

It is straightforward to verify that this is a sigma protocol with special soundness and special honest-verifier zero-knowledge. As above, we can apply the Fiat-Shamir [32] transformation to get a non-interactive protocol. We will refer to this proof as $\Pi_{\text{DLEQ2}}(G, H, T, S, K, V; k_0, k_1)$.

B.3 OR-Proof of Equal Discrete Logs

We present the honest-verifier zero-knowledge `OR`-proof of equal discrete logarithms instantiated by Kreuter et al. [45, Appendix B]. Let $\mathbb{G}$ be a cyclic group of prime order p with generators G, H, T, S, and let $V_0 := [e_{0,0}]G + [e_{0,1}]H$ and $V_1 := [e_{1,0}]G + [e_{1,1}]H$, where $e_{i,j}$ are distinct scalars private to the prover $\mathcal{P}$. Furthermore, let $W := [e_{b,0}]T + [e_{b,1}]S$ for $b \in \{0, 1\}$. We want to prove that W is computed using the same secret scalars as either V_0 or V_1.

$\mathcal{P}$.1 Choose random scalars $r_0, r_1, c_{b-1}, u_{b-1}, v_{b-1}$ in the underlying field and compute the following:

$$\begin{aligned} A_{b,0} &:= [r_0]G + [r_1]H, \\ A_{b,1} &:= [r_0]T + [r_1]S, \\ A_{1-b,0} &:= [u_{b-1}]G + [v_{b-1}]H - [c_{b-1}]V_{b-1}, \\ A_{1-b,1} &:= [u_{b-1}]T + [v_{b-1}]S - [c_{b-1}]W. \end{aligned}$$

Finally, send $(A_{0,0}, A_{0,1}A_{1,0}, A_{1,1})$ to $\mathcal{V}$.
$\mathcal{V}$.1 Choose a random challenge c modulo p and send it to $\mathcal{P}$.
$\mathcal{P}$.2 Compute the responses

$$c_b := c - c_{1-b}, \quad u_b := r_0 + c_b e_{b,0}, \quad v_b := r_1 + c_b e_{b,1},$$

modulo p. Send $(c_i, u_i, v_i)_{i=0,1}$ to $\mathcal{V}$.

$\mathcal{V}.2$ Verify that $c = c_0 + c_1$ and that

$$\begin{aligned} A_{0,0} &= [u_0]G + [v_0]H - [c_0]V_0, \\ A_{0,1} &= [u_0]T + [v_0]S - [c_0]W, \\ A_{1,0} &= [u_1]G + [v_1]H - [c_1]V_1, \\ A_{1,1} &= [u_1]T + [v_1]S - [c_1]W. \end{aligned}$$

We can make the proof non-interactive using Fiat-Shamir [32] like above. We will refer to this protocol as $\Pi_{\text{DLEQOR2}}(G, H, T, S, V_0, V_1, W; e_{b,0}, e_{b,1})$.

Π_{DLEQOR2} can be batched for many instances with respect to the same secret scalars using the techniques by Henry [37] as shown in [44, Appendix B.1].

C Anonymous Tokens with Public and Private Metadata

In Figure 8, we present an extension of the PMBTokens [44, Figure 8] with public metadata. This protocol is also designated verifier, requiring the secret key for verification.

Setup and Key Generation. Let λ be the security parameter, let p be a prime and let E be an elliptic curve group of order p with generators G_0, G_1. Let $\mathtt{H}_t : \{0,1\}^* \to E$, $\mathtt{H}_m : \{0,1\}^* \to \mathbb{Z}_p^*$ and $\mathtt{H}_s : \{0,1\}^* \to \mathbb{Z}_p^*$ be hash-functions, and assume that group elements and integers can be encoded uniquely as strings. Furthermore, let metadata md be an element of a public set of valid strings. Finally, let $\mathsf{sk} := (k_{0,0}, k_{0,1}, k_{1,0}, k_{1,1}) \leftarrow\!\!\$\ (\mathbb{Z}_p^*)^4$ (all $k_{i,j}$ being distinct) be the signing key, and let $\mathsf{pk} := \{K_{i,j}\} = \{[k_{i,j}]G_i\}$, for $i, j = 0, 1$, be the public key. This is implicit knowledge in the protocol description.

Signing and Verification. The anonymous tokens protocol in Figure 8 uses the Π_{DLEQ2}-protocol defined in Appendix B.2 twice as a subroutine to ensure that we afterwards can prove that the signed token W' was computed correctly. Given the generators G_0, G_1, T', S', the public keys $K_{i,j} := [k_{i,j}]G_i$ and the elements $V_i := [e_{i,0}]T' + [e_{i,1}]S'$, for $i, j = 0, 1$, we want to prove that the following relations hold:

$$\begin{bmatrix} G_0 + G_1 \\ V_i \end{bmatrix} = [e_{i,0}] \begin{bmatrix} [d]G_0 + K_{i,0} \\ T' \end{bmatrix} + [e_{i,1}] \begin{bmatrix} [d]G_1 + K_{i,1} \\ S' \end{bmatrix}.$$

We instantiate $\Pi_{\text{DLEQ2}}(G_0, G_1, V_0; e_{0,0}, e_{0,1})$ and $\Pi_{\text{DLEQ2}}(G_0, G_1, V_1; e_{1,0}, e_{1,1})$ in Figure 8 to get proofs $\pi_{\mathtt{AND}}$ and $\pi'_{\mathtt{AND}}$. We denote the verification by $\mathtt{V}(\pi_{\mathtt{AND}}, \pi'_{\mathtt{AND}})$.

We also use the Π_{DLEQOR2}-protocol defined in Appendix B.3 . The signer computes an $\mathtt{OR}$-proof of equality of discrete logs by instantiating the zero-knowledge protocol $\Pi_{\text{DLEQOR2}}(G_0, G_1, T', S', V_0, V_1, W'; e_{0,0}, e_{0,1}, e_{1,0}, e_{1,1})$. Consider the generators G_0, G_1, T', S', hashed metadata d and computed value W'. The signer then proves that W' is correctly computed, with respect to T' and

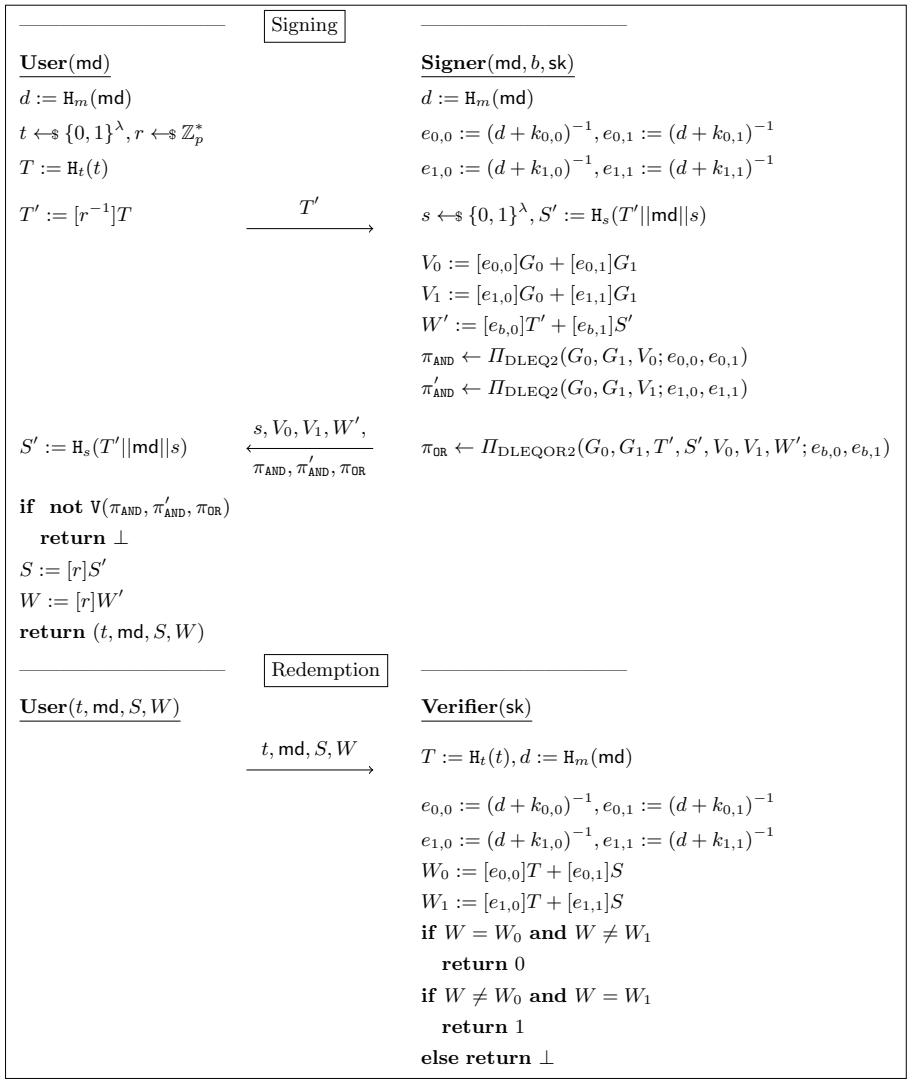

Fig. 8. Designated verifier anonymous tokens with public and private metadata, an adjusted extension of Kreuter et al. [44].

S', and in the same way as one of the committed values V_0 or V_1, with respect to G and H. That is, for either $b = 0$ or $b = 1$:

$$V_b = [e_{b,0}]G_0 + [e_{b,1}]G_1 \quad \wedge \quad W' = [e_{b,0}]T' + [e_{b,1}]S'.$$

We denote the verification of the proof $\pi_{\mathtt{OR}}$ by $\mathtt{V}(\pi_{\mathtt{OR}})$.

Theorem 4 (Completeness). *The anonymous token protocol with public and private metadata in Figure 8 is complete according to Definition 2 and will, according to Definition 3, return the correct metadata bit except with negligible probability.*

Proof. If the user submits (t, md, S, W), completeness follows from expanding W_b:

$$\begin{aligned} W_b &= [r]([e_{b,0}]T' + [e_{b,1}]S') = [r]([e_{b,0}][r^{-1}]T + [e_{b,1}]S') \\ &= [e_{b,0}]T + [r][e_{b,1}]S' = [e_{b,0}]\mathtt{H}_t(t||\mathsf{md}) + [e_{b,1}]S. \end{aligned}$$

Furthermore, the probability that this equation holds for both $b = 0$ and $b = 1$ is negligible. If that was the case, then

$$[e_{0,0}]T + [e_{0,1}]S = [e_{1,0}]T + [e_{1,1}]S.$$

As we require all keys $k_{i,j}$ to be distinct, it follows that all $e_{i,j}$ are distinct. Then, we have that

$$T = \left[\frac{e_{1,1} - e_{0,1}}{e_{0,0} - e_{1,0}}\right] S.$$

Since $T = \mathtt{H}_t(t)$ is sampled independently and uniformly at random, the probability that this equation holds is $1/p$, which is negligible.

Theorem 5 (Unforgeability). *The anonymous token protocol with public and private metadata in Figure 8 achieves one-more unforgeability with respect to Definition 4.*

Proof. For fixed metadata md we let the adversary query the signing oracle ℓ times for both $b = 0$ and $b = 1$. Using the key transformation as described in Lemma 1, the security of the protocol reduces to the security of the one-more gap strong inversion Diffie-Hellman game as shown in Figure 6. The advantage of an attacker follows from Definition 7 and is proven secure by Tyagi *et al.* [50, Theorem 1].

Theorem 6 (Unlinkability). *Fix private metadata b and public metadata md. Within the set defined by all tokens using b and md, the anonymous token protocol with public and private metadata in Figure 8 achieves unlinkability with respect to Definition 5.*

Proof. We note that it is easy to create many different anonymity sets to distinguish users based on private metadata being $b = 0$ or $b = 1$, and in combination with different values of public metadata md. We restrict the unlinkability to hold for users within the same anonymity sets based on b and md, both sampled according to the real distribution of private and public metadata. Let $\mathfrak{U}_{b,\mathsf{md}}$ be this set, and select two sessions from $\mathfrak{U}_{b,\mathsf{md}}$. Then it follows directly from [44, Theorem 9] that the probability of success of the adversary will be upper bounded by $2/m + \mathsf{negl}(\lambda)$

Theorem 7 (Private metadata bit). *The anonymous token protocol with public and private metadata in Figure 8 provides private metadata bit with respect to Definition 6.*

Proof. This statement follows directly from the proof of [44, Theorem 10], which describes a hybrid argument to prove that instances with private bit 0 are indistinguishable from instances with private bit 1. Notice in particular that the extra `OR`-proofs in our protocol are independent of the private bit b, and therefore need no additional simulation.

D Public Verifiability from Pairings

The authors of Privacy Pass [30] described an application where the issuer and the recipient of a token would be the same entity, possibly separated by time. For the application we present in Section 5, those two roles are in fact separate, and one should therefore have a scheme that supports public verifiability. It remains an open problem to achieve this without pairings, unless we allow for two rounds of communication [2, 52].

We move on to provide a new variant of a partially blinded signature by Zhang, Safavi-Naini and Susilo [53]. The protocol allows a user and a signer to generate a signature on a user-private message m and agreed-upon metadata md. Both the issuance protocol and the signature consists of a single curve point.

We show that the idea underlying this scheme can be viewed as a combination of Boneh-Lynn-Shacham signatures [12] and Privacy Pass, inheriting its attractive properties from both.

Setup and Key Generation. Let λ be the security parameter, let $\hat{e} : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ be a pairing, where G_1, G_2 and g_T are generators for their respective prime p order groups. Furthermore, let $\mathtt{H}_1 : \{0,1\}^* \to \mathbb{G}_1$ and $\mathtt{H}_m : \{0,1\}^* \to \mathbb{Z}_p^*$ be hash functions, and assume that group elements and integers can be encoded uniquely as strings. Also, let md be an element of a public set of valid metadata strings. Finally, let $\mathsf{sk} := k \leftarrow\!\$\ \mathbb{Z}_p^*$ be the signing key, and let $\mathsf{pk} := K = [k]G_2$ be the public key. This is implicit knowledge in the protocol description.

Signing and Verification. Recall that the BLS-scheme signs a message m by hashing it to the group generated by G_1 and multiplying it with the secret key

k; $W := [k]\mathtt{H}_1(m)$. The signature can then be verified by checking that

$$\hat{e}(\mathtt{H}_1(m), K) = \hat{e}(W, G_2).$$

Correctness follows from the linearity of the pairing.

We replace m by a token seed t, and use the same trick as earlier to concurrently update the key-pair based on metadata. Then we get the following anonymous token scheme:

Signing The user sends $T' := [r^{-1}]\mathtt{H}_1(t)$ to the issuer, who returns $W' := [e]T'$, for $e = (d+k)^{-1}$. The user can verify that the signature is correct by checking $\hat{e}(W', U) = \hat{e}(T', G_2)$, for $U := [d+k]G$, and then storing $(t, \mathsf{md}, W = [r]W')$.

Verification The user sends (t, md, W), and the recipient can verify the token by checking if $\hat{e}(W, U) = \hat{e}(T, G_2)$.

This scheme hides the token similarly to Privacy Pass, it can be verified without using the private key, and its unforgeability follows directly from BLS. We note that the check $\hat{e}(W', U) = \hat{e}(T', G_2)$ ensures that the tokens are signed correctly with respect to the public key. The complete protocol is listed in Figure 9. Finally, we note that we can batch-verify n tokens under the same key and metadata and check for equality by computing

$$\hat{e}\left(\sum_i [c_i]W_i, U\right) = \hat{e}\left(\sum_i [c_i]T_i, G_2\right)$$

where $c_1, \ldots, c_n$ are random coefficients. This saves the verifier of $2(n-1)$ expensive pairing-computations, which is especially useful in systems with large anonymity sets. Note that the verifier computes T_i from the received pre-tokens t_i, making sure that (T_i, W_i) is not just a scaling of a different valid token.

Theorem 8 (Completeness). *The anonymous token protocol with public metadata and public verifiability in Figure 9 is complete according to Definition 2.*

Proof. Completeness follows from expanding $\hat{e}(W, U)$:

$$\begin{aligned}\hat{e}(W, U) &= \hat{e}([r]W', [d+k]G_2) = \hat{e}([r][e]T', [d+k]G_2) \\ &= \hat{e}([r][e][r^{-1}]T, [d+k]G_2) = \hat{e}([e]T, [d+k]G_2) \\ &= \hat{e}(T, G_2)^{e\cdot(d+k)} = \hat{e}(T, G_2).\end{aligned}$$

Theorem 9 (Unforgeability). *The anonymous token protocol with public metadata and public verifiability in Figure 9 achieve one-more unforgeability with respect to Definition 4.*

Proof. Assume that we have an adversary who breaks unforgeability. In particular, this means that they can produce $\ell+1$ distinct and valid tuples $(t_i, \mathsf{md}, \sigma_i)$ but only query the signing oracle at most ℓ times. We use this adversary to construct an adversary against the (m, n)-OM-GAP-SDHI problem in the Random Oracle Model.

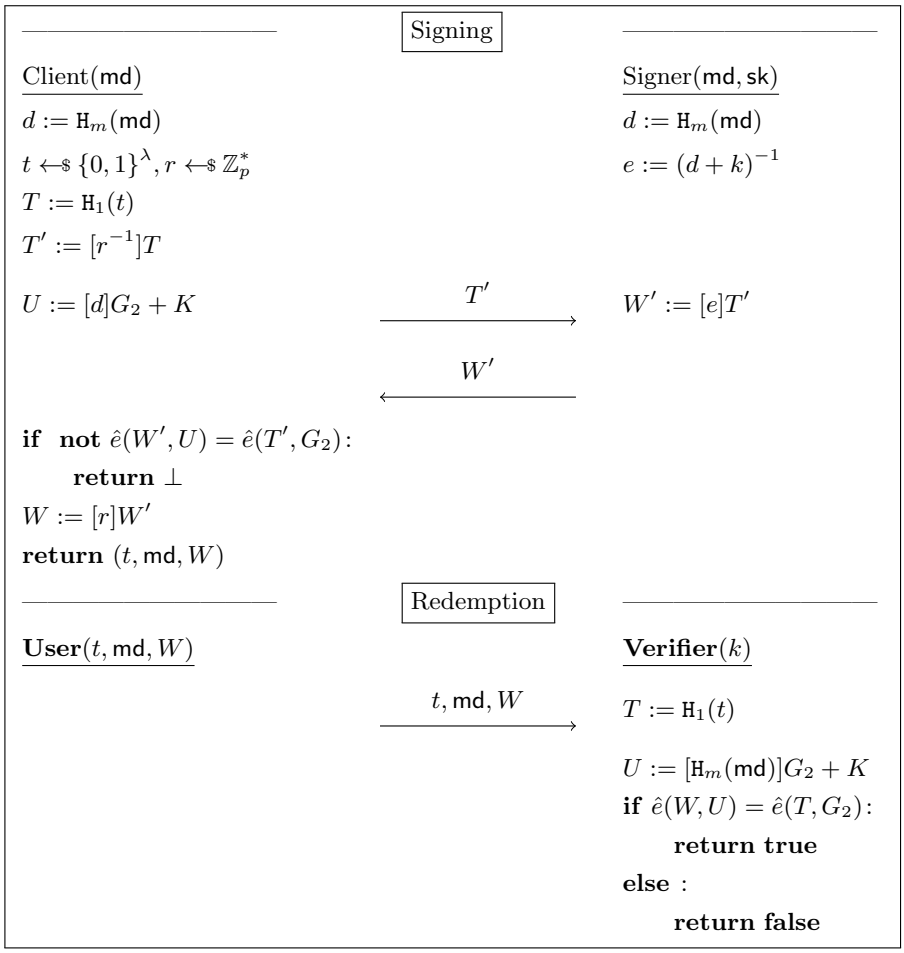

Fig. 9. Anonymous tokens with public metadata and public verifiability by adjusting Zhang et al. [53] for asymmetric pairings.

$\mathcal{A}_1^{\text{om-gap-sdhi}}$ Recall that we assume that the user and the signer agrees on the metadata. Run $\mathcal{O}_m$ on all acceptable metadata values to get the list $\{c_i\}$, and return it.

$\mathcal{A}_2^{\text{om-gap-sdhi}}$ Receive the input $G, K = [k]G, [y_i]G$. Set $\mathsf{vk} = K$ and the other parameters appropriatly. Reprogram $\mathcal{O}_1$ such that it on input t returns $[t][y_j]G$ for the next j, which looks like a random group element. Whenver the adversary queries the SIGN oracle, forward the query to the SDH oracle. If the adversary wins the OMUF game for some metadata value md corresponding to an index γ, we have $\ell+1$ signatures $\sigma_i = [e_{\mathsf{md}}]\mathcal{O}_1(t_i) = [e_{\mathsf{md}}][t_i][y_1]G$. Use the programming of $\mathcal{O}_1$ to return $(\gamma, ([t^{-1}]\sigma_i, \alpha_i))$

The result from $\mathcal{A}_2^{\text{om-gap-sdhi}}$ satisfies the (m, n)-OM-GAP-SDHI conditions.

One can construct a more detailed proof along the lines of [50, Appendix B] in order to get concrete bounds.

Theorem 10 (Unlinkability). *Fix metadata* md*. Within the set defined by all tokens using* md*, the anonymous token protocol with public metadata and public verifiability in Figure 9 achieve unlinkability with respect to Definition 5.*

Proof. Observe that given any valid token (t, md, W) and any honestly generated view (T', W') there exists a unique value r' such that both $W - [r']W'$ and $T - [r']T'$ holds, and hence, T is independent of any W. It follows that the anonymous token is unlinkable.