

Anonymous Tokens with Hidden Metadata Bit from Algebraic MACs

Melissa Chase¹, F. Betül Durak¹, and Serge Vaudenay²

¹ Microsoft Research

² EPFL

Abstract. On the one hand, the web needs to be secured from malicious activities such as bots or DoS attacks; on the other hand, such needs ideally should not justify services tracking people’s activities on the web. Anonymous tokens provide a nice tradeoff between allowing an issuer to ensure that a user has been vetted and protecting the users’ privacy. However, in some cases, whether or not a token is issued reveals a lot of information to an adversary about the strategies used to distinguish honest users from bots or attackers.

In this work, we focus on designing an anonymous token protocol between a client and an issuer (also a verifier) that enables the issuer to support its fraud detection mechanisms while preserving users’ privacy. This is done by allowing the issuer to embed a hidden (from the client) metadata bit into the tokens. We first study an existing protocol from CRYPTO 2020 which is an extension of Privacy Pass from PoPETs 2018; that protocol aimed to provide support for a hidden metadata bit, but provided a somewhat restricted security notion. We demonstrate a new attack, showing that this is a weakness of the protocol, not just the definition. In particular, the metadata bit hiding is weak in the setting where the attacker can redeem some tokens and get feedback on what bit is extracted.

We then revisit the formalism of anonymous tokens with private metadata bit, consider the more natural notion, and design a scheme which achieves it. In order to design this new secure protocol, we base our construction on algebraic MACs instead of PRFs. Our security definitions capture a realistic threat model where adversaries could, through direct feedback or side channels, learn the embedded bit when the token is redeemed. Finally, we compare our protocol with one of the CRYPTO 2020 protocols which we obtain 20% more efficient implementation.

1 Introduction

There has been significant industry interest recently in anonymous tokens, including Google’s Trust Tokens (TT) [1] and Cloudflare’s Privacy Pass (PP) [12]. These protocols are used to transfer trust signals without compromising users’ privacy. Anonymous tokens define a protocol between three types of parties: a client, an issuer, and a redeemer. The client wishes to obtain tokens from an

issuer and then present them to a redeemer. The issuer determines the trustworthiness of clients and issues tokens; and the redeemer (a.k.a. the verifier) verifies the tokens. These systems are anonymous in that the token issuance and redemption are unlinkable, in the sense that the issuer and redeemer cannot tell which of the issued tokens was used in any given redemption. In the above systems, we consider the issuer and the redeemer which are controlled by the same entity, so these two parties are assumed to share a secret key.

In the case of PP, it is specifically designed in the context of CDNs to assess the trustworthiness of a client at the edge before the client is granted (or denied) access to a web server. In the typical use case, the client is required to solve a CAPTCHA before it accesses a web server for the first time. If the CAPTCHA is successfully solved, the client is given a set of tokens to redeem the next time it visits the web server. At subsequent visits the server checks that the user's token is valid and has not previously been used, and if so allows the client to bypass the CAPTCHA and directly access the content. This allows for a better experience for the user since they only have to complete the CAPTCHA once, even if they are accessing the webserver over Tor or a VPN. Because of the unlinkability, it does so without helping the web server to track the user across different visits.

Even though CAPTCHA is one way to detect bad actors, there are more advanced techniques to assess how trustworthy the client is, for example based on machine learning algorithms for fraud detection. Typically, such fraud detection algorithms are run on the issuer side; when these algorithms determine that a client is likely to be malicious, the issuer should refuse to issue it any tokens. However, such feedback allows a fraudulent client to improve their methods to bypass the fraud detection. Ideally the issuer would instead embed a bit (to indicate if the client is malicious or trusted) inside the token which would be hidden from the client and only recovered by the redeemer during the redemption. That way, the malicious client would not find out that its fraudulent activity has been detected until it tries to redeem the token. This would make this type of attack on the fraud detection algorithms significantly more cumbersome.

Given the idea behind anonymous tokens with private metadata bit, there are three desired security and privacy requirements: unforgeability (to prevent malicious clients from forging valid tokens), unlinkability (to prevent a malicious server and redeemer from linking the tokens they issued with those that are redeemed), and privacy of embedded metadata bit (to prevent the clients from learning immediately if they were identified as malicious actors). More specifically, the metadata bit is a covert channel between issuer and redeemer, who as described above will share a key. If this channel were allowed to convey unlimited information, unlinkability would be meaningless, so the protocols must ensure that the embedded signal is only a single bit and no more. Such covert channel with one bit is used to communicate whether or not the issued token should be accepted without revealing the decision to the client until it attempts to redeem.

The initial proposal for including a private metadata bit built directly on the PP protocol where the client picks an arbitrary message t and masks it to hide from the issuer (i.e., the issuer blindly "signs" the t), and then the client un.masks

the signature on t . This idea naturally extends to support private metadata bits with PP: the issuer would choose two PP issuing keys and generate a token under one key or the other depending on the bit it wished to encode.

Because PP is based on an oblivious pseudo random function, a token generated under either key was indistinguishable from random so the client was unable to determine the bit from looking at a correctly generated token. However, this protocol had significant weaknesses; a malicious client can easily make malformed token requests (keep using the same message t), and then tell from the responses whether the tokens issued encode the same metadata bit. This means that the attacker has to make one request for which he can predict the resulting bit (e.g. by using a genuine user device or by behaving badly enough that it will be guaranteed to be detected as fraudulent), and then it can make an incorrectly formed token request and directly tell from the (invalid) token it receives whether the attacker was attempting to issue a token with metadata bit 0 or 1. This problem essentially comes from the fact that the PP allows the client to pick the messages to be signed arbitrarily. In our design, messages generated by both the client and the issuer guaranteeing the randomness of it.

Two recent papers [14,13] and [16] aim to address this problem and to formally define and construct anonymous tokens with private metadata bit. The latter is more generic version of the former where the protocols can accept public metadata as well as the private metadata bit. The authors in these work identify the problem as being that the tokens are deterministic; they propose new randomized protocols. These proposals address the issues above and guarantee that an adversary who can maliciously interact with the token issuing server cannot learn anything about the private metadata bits encoded in the tokens it receives.

However, these schemes still have some counter intuitive properties. In particular, in their protocol, there are two notions for a token: “validity” where the adversary gets the feedback if the token is verified correctly and “embedded bit” where the adversary gets feedback if there is an embedded bit or not without revealing the bit (if it exists). Their definition of privacy for the metadata bit allows the adversary to learn at redemption whether a token is “valid”, but neither if there exists a bit nor the value of embedded bit. And this is not just a property of the definition: we will show that the proposed schemes have the property that if the redemption service reveals whether a bit is embedded during the token redemption, the adversary can use only a few malicious interactions with the issuing and redemption service to learn information about the hidden bits embedded in a large batch of tokens.

This separation may make sense in some contexts, where we can guarantee that the adversary gets no feedback at all from the redemption server on whether it’s token was accepted and what the included bit was. In other cases however, this seems to be a nontrivial weakness. For example, in the CDN application above, the adversary will clearly get feedback on the bit if it used to determine whether it will be allowed to access to the web content. A private metadata bit protocol with this type of weakness would allow the adversary to make many attempts to bypass the fraud detection and thus acquire many tokens, then

make a few redemption requests and identify exactly which of the remaining unredeemed tokens successfully avoided the fraud detection. Those tokens could then be collected and, for example, used to mount a DDOS attack.

Thus, it is clear that in at least some settings, we would like a token system which provides stronger guarantees. Moreover, identifying all sources of feedback is challenging. It seems likely that if a new primitive for anonymous tokens is released, it will at some point be used in settings where the adversary can get feedback on the encoded bit, whether or not the security definition allows for that. Thus, the best solution would be the one that provides the most natural security guarantees - an adversary interacting with an issuer and a redeemer may learn the bits encoded *in the tokens it redeems*, but nothing else.

That leaves us with the following questions: *Is such a definition efficiently realizable? What is the overhead as compared to the solutions described above?*

Our Contributions. In the rest of the paper, we begin by summarizing the private metadata bit proposals of [14,13], describing the known weakness in detail, and explaining an additional attack. Then we present our more natural security definitions, which are based on the definitions from [14,13] but do not allow for such weaknesses. Finally, we present a *new construction for anonymous tokens* with private metadata bit which we show satisfies *more natural security definitions*. We analyze the efficiency of our protocol in Section section 5 and show that, surprisingly, our protocol is faster than the weaker proposals of [14,13] (1.3 ms vs 1.6 ms). Finally, we show the *flexibility of our approach* by demonstrating that it extends easily to allow for tokens including *public metadata* visible to both issuer/redeemer and client.

Our Techniques. As mentioned above, the initial privacy pass protocol was based around oblivious pseudorandom functions (OPRFs). The proposals of [14,13,16], as described above, identify the problem as that OPRFs are deterministic, and attempt to address it by making the protocol randomized. However, once we recognize that we do not in fact need a deterministic function, we can ask whether it makes sense to base this primitive around OPRFs, whose defining characteristic is that they are deterministic. Moreover, the obliviousness property of OPRFs turns out to be *not* a very good fit for hiding the metadata bit.

This begs the question: *Is there a better primitive to start from if we want to encode hidden data in anonymous tokens?*

For this, we turn to authentication primitives like MACs. Since we need privacy, we look at the work on anonymous credentials, which allows issuers to certify attributes which can later be presented unlinkably [10,3]. Specifically, we will borrow from keyed verification anonymous credentials (KVAC)[6,7], where the credential issuer and verifier share a secret key, and on constructions based on algebraic MACs.

KVAC directly gives us a protocol for blindly issuing credentials, in which an issuer issues a MAC on a set of attributes, some of which are only known to the client. At a high level we can apply this as follows: in token issuance, the client chooses a random nonce, and the issuer uses blind issuance to give a MAC

on a pair of messages consisting of the nonce and the hidden bit. If we use the first construction from [6], MACGGM, that gives a construction in elliptic curve groups where tokens consist of the nonce and two additional group elements, which is roughly comparable with PMBT, and only twice as long as PP. But, if we consider the blind issuance protocol from [6], it would be roughly twice as expensive as the issuance for PMBT, and 4 times the cost of PP. It also does not directly provide metadata privacy as that is not a property generally considered in the credentials setting. Technically then, we are left with two questions: 1) *Can we optimize the MACGGM based issuance protocol to the point where it is competitive with PP or PMBT?* and 2) *If the client does not know one of the attributes in the credential, do the protocols still work? Can we prove that this attribute will not be leaked to the client?*

We address the first of these questions with a very careful optimization of the blind issuance protocol to get a result with comparable cost to PMBT. If we wanted to directly reduce to MAC security, we would need the request to include something from which we could extract the nonce in \mathbb{Z}_p which will be the message for the MAC. Extracting messages in \mathbb{Z}_p is extremely expensive.³ Moreover, the blind issuance protocol in [6] has the client form an ElGamal encryption of the message to be signed, which would result in a client-to-server message twice as long as in PP or PMBT, even before the proofs are added. Instead, we design an optimized blind issuance protocol and prove in the generic group model that the resulting token scheme is unforgeable.

Unlinkability follows in a straightforward way from the privacy of the KVAC scheme. Privacy of the metadata bit is more challenging, as there is no analog in the KVAC context. First, we note that the blind issuance protocol works even if the client does not know one of the messages, and the verifier (the issuance server) can simply verify with both possible bit messages, and output the bit for which the MAC verifies. Intuitively, we also might hope to get privacy for the metadata bit because MACGGM has some pseudorandomness properties: the basic MAC on a message pair (m_0, m_1) is $(U, (x + ym_0 + zm_1)U)$, where U is a random group element (x, y, z) are the secret key. DDH then guarantees that this will look like a random pair of messages. However, proving that this satisfies the metadata bit privacy property, where the adversary can interact maliciously with issuance and redemption oracles is significantly more challenging. Here, we again prove security in the generic group model.

Related Work. Beyond the works on anonymous tokens mentioned above, the most closely related work is in anonymous credentials. While there are also works based on RSA groups (beginning with [3,4]) and based on pairings, we focus here on works that can be implemented in prime order elliptic curve groups, since those provide the best efficiency. In that setting, besides MACGGM, there is one other proposal for a MAC based anonymous credential scheme [8] which is strictly more complex and more expensive than MACGGM.

³ [6] addresses this by making non-standard assumptions about the extraction properties of Fiat-Shamir.

In addition, there are several anonymous credential constructions in elliptic curve groups that take a blind-signature based approach [15,2,18].⁴ It is not clear how to add private issuer values (like the metadata bit) in these schemes, but even if we consider the simpler setting where the bit is known to the user, these schemes have several downsides: First, token redemption is significantly more expensive (4x for [15] or [18], and 8x for [2]). In a setting like the CDN application where we will be using tokens to decrease spam/prevent DDOS attacks, we want the cost to verify tokens to be as low as possible. Secondly, they all require a multi-round issuance protocol, with two round trips between the user and issuer. This requires that the token issuance server to be stateful, and in fact if the issuer can be tricked into completing a protocol in two different ways, then his secret key will be leaked. This means that implementing a token issuer requires careful state management, including storage per client session. In a setting where there are many many clients, many of which may be untrusted or on flaky connections, this can be quite expensive.

2 PMBT: a Case Study

We start working with [14] (and its full version [13]) and investigate the shortcomings as well as a new attack. There are two main constructions in these works called Private Metadata Bit Tokens (PMBT) and CMBT. For PMBT, the authors state that `Verify` returning always `true` is not meaningful and then announced their new protocol named CMBT in the full version of their paper [13]. Before the description of protocols, we borrow the interface of anonymous tokens (AT) with private metadata bit and its security and privacy requirements.

2.1 AT Interface and Security

- $(\text{crs}, \text{td}) \leftarrow \text{AT.Setup}(1^\lambda)$, the protocol requires a set up with security parameter λ , and returns a common reference string crs and a trapdoor td .
- $(\text{pp}, \text{sk}) \leftarrow \text{AT.KeyGen}(\text{crs})$, the public parameters pp and a secret key sk is generated with `KeyGen` algorithm which takes a crs as input.
- $\{\sigma, \perp\} \leftarrow \langle \text{AT.Client}(\text{pp}, t), \text{AT.IssueToken}(\text{sk}, b) \rangle$, the interactive token generation protocol runs between a *client* (also called a *user*) and the issuer. The client inputs are a string t along with the issuer's public parameters, and the issuer inputs are the secret key sk and a *metadata bit* b . The protocol outputs a token (also referred to as signature) σ for the client or \perp .
- $\text{bool} \leftarrow \text{AT.Verify}(\text{sk}, t, \sigma)$, the verification algorithm is run by the *redeemer* (which is the issuer, or a related party holding sk). The inputs are sk and a token (t, σ) . It outputs a boolean value to indicate if the token was valid.
- $\text{ind} \leftarrow \text{AT.ReadBit}(\text{sk}, t, \sigma)$, the metadata bit is extracted with the `ReadBit` algorithm which takes the private key and a token (t, σ) as input. It outputs an indicator $\text{ind} \in \{\perp, 0, 1\}$, with the value of b , or \perp if extraction failed.

⁴ All of these construction are multi-show in that many presentations/redemptions of the same credential are clearly linkable.

The issuance protocol is interactive between the user and the issuer. In [14], it is assumed to be a user-to-issuer-to-user protocol (2-move, user-initiated).

Security properties of an AT scheme are unforgeability, unlinkability, and privacy of the metadata bit. They are formally (re)defined in section 3.

1. **Unforgeability** implies that **only** the issuer can generate *valid* tokens. More precisely, if the issuer is invoked n times, the adversary cannot exhibit $n + 1$ valid tokens. The “validity” notion is crucial and discussed below. (This is called one-more unforgeability.)
2. **Unlinkability** implies that redeemed tokens cannot be linked to issuing sessions. More precisely, a malicious issuer who is engaging n signing sessions and then given a shuffled list of (t, σ) tokens to redeem cannot link a given token from the shuffled list to the corresponding issuing session with a probability slightly less than $2/n$. (This is called κ -unlinkability for $\kappa = 2$.)
3. **Privacy** of metadata means that no one but the redeemer learns the private bit b which was inserted by the issuer. Tokens issued with $b = 1$ should be indistinguishable from tokens with $b = 0$, even with access to an oracle telling whether a token is valid. Again, the “validity” notion is crucial. It may deviate from the validity notion in unforgeability.

Notice that there is a distinction between AT.Verify and AT.ReadBit algorithms. Essentially for (one-more) unforgeability to work, the attacker should pass two verification conditions: AT.Verify returns true and AT.ReadBit does not return \perp as indicator. However, in the security games of privacy, the attacker is given access to only AT.Verify algorithm not the AT.ReadBit algorithm. So, the authors [14,13] freely make AT.Verify always return true and sweep everything to AT.ReadBit algorithm, hence no real access for the attacker.

In [14], the unforgeability game declares a token *valid* if Verify is true and ReadBit returns an element of $\{0, 1\}$ (i.e. not \perp). The unforgeability game does not allow tokens in forgeries to share the same t value.

We recall the PMBT issuance protocol in Figure 1. As specified in [14], Verify always returns true and ReadBit returns $b \in \{0, 1\}$ such that $W = x_b H_t(t) + y_b S$ when it exists or \perp otherwise.

2.2 Potential Attack for PMBT

In this section, we describe two attacks on PMBT. Both attacks are possible in real-world but not considered in the formal security definitions. Attack I is acknowledged by the authors of PMBT [14], but Attack II is not considered and can be used without the issuance servers detecting them.

Attack I: Since Verify is always true, the validity is defined with the condition that ReadBit does not return \perp . However, what the user can do is to request tokens with $(T'_1, T'_2, \dots, T'_n)$ where $T'_i = r_i^{-1} H_t(t)$ for $i = 1, \dots, n$ and a fixed t in order to prepare his attack. As a return, he gets many tokens $(s_1, W'_1), (s_2, W'_2), \dots, (s_n, W'_n)$ where s_i 's are random and $W'_i = x_b T'_i + y_b S'_i$

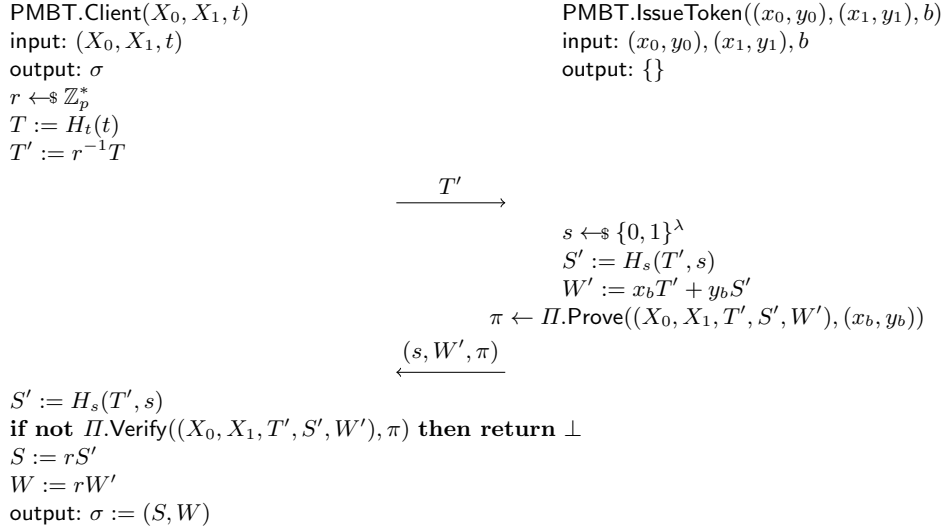


Fig. 1. PMBT token issuance protocol as given in [14, Fig 8, p 325]

with $S'_i = H_s(T'_i; s_i)$. The user gets these tokens and unblinds. As a result, it gets (t, S_i, W_i) verified with the secret by $W_i = x_b H_t(t) + y_b S_i$. From n linear equations $W_i = x_b H_t(t) + y_b S_i$ with the same bit b , the user can now take random coefficients α_i such that $\sum_i \alpha_i = 1$ and compute

$$\sum_i \alpha_i W_i = x_b H_t(t) + y_b \left(\sum_i \alpha_i S_i \right) \quad (1)$$

$$W = x_b H_t(t) + y_b S \quad (2)$$

which makes the triplet (t, S, W) pass verification. Thus, when $n \geq 2$, a malicious client can generate many valid tokens with same tag t and hiding the same bit b . The unforgeability game does not take it as a forgery because the tokens share the same t , so this attack is excluded by rule.

We can have an attack against the privacy of b by using the above way to make forgeries. If all collected tokens use the same b , the forged token is valid. However, if the equations do not use the same b , the above equation is not satisfied so the triplet is not valid. This enables the attacker to apply “cut-and-choose” to learn which tokens are valid and which are not without consuming them. The attack assumes access to a validity oracle which can be expected with a side-channel attack. This oracle would tell whether AT.ReadBit returns \perp . In the PMBT model, the adversary has only access to an AT.Verify oracle. As it always return **true**, the adversary does not have this side channel by rule.

Attack II: In another attack, we consider an adversary who collected n tokens with distinct t_i 's, $(t_1, S_1, W_1), \dots, (t_n, S_n, W_n)$, satisfying $W_i = x_{b_i} H(t_i) + y_{b_i} S_i$.

Assuming an oracle returning whether `AT.ReadBit` returns \perp or not, we describe an attack allowing the adversary to check whether a subset of the collected tokens hide the same bit or not. Let I be a subset of $\{1, \dots, n\}$. The adversary chooses a random t^* and runs the token request protocol one more time, but this time uses $T^* = H(t^*) - \sum_{i \in I} H(t_i)$ instead of $T^* = H(t^*)$, and obtains (S^*, W^*) such that $W^* = x_{b^*} T^* + y_{b^*} S^*$. Now, the adversary attempts to redeem token $(t^*, S^* + \sum S_i, W^* + \sum W_i)$. This will be a valid token iff all of the b_i for $i \in I$ are equal to b^* . If `AT.ReadBit` does not return \perp , the adversary deduces that all b_i for $i \in I$ are equal. If `AT.ReadBit` returns \perp , the adversary may try again, because it could still be the case that all b_i are equal but that b^* was not. Eventually, the adversary infers information and can apply a cut-and-choose strategy. Notice that this attack would be impossible to detect: the protocol ensures that malicious token to redeem is random. Furthermore, even if the redeem oracle checks for repeating tags, our attack ensures that all redeem attempts have different tags so no such countermeasure works.

As conclusion, the attack method does not violate the unforgeability game by rule (that t 's must be different) and does not violate the privacy game by rule (that there is no validity oracle). As acknowledged by the authors [14, Section 6.1], having `Verify` returning always `true` and pushing the validity to `ReadBit` not returning \perp is a bit cheating. They propose a way to enable token verification in the CMBT scheme which is presented in the full version of the paper [13, Appendix J].

2.3 The CMBT Fix

The CMBT protocol is presented in the eprint version of the paper [13]. We summarize this protocol in Figure 2.

CMBT was updated⁵ on April 21, 2022 [13] as shown on Figure 2. A token (S, W, \tilde{W}) is verified and reads bit b if $\tilde{W} = \tilde{x}H_t(t) + \tilde{y}S$ (verify part), $W = x_b H_t(t) + y_b S$ and $W \neq x_{1-b} H_t(t) + y_{1-b} S$ (read part). Unforgeability is enforced by the security of previous constructions. However, the privacy of metadata bit is still in question when we do not separate `Verify` and `ReadBit` as we will detail next.

Clearly, a linear combination attack with tokens sharing the same t will forge new tokens making `AT.Verify` true and `AT.ReadBit` returning the common bit to all tokens or \perp if there is no unanimous bit b . Hence, privacy is broken with access to an oracle telling whether a token is well formed.

As another attack, if a client engages in a protocol to issue three tokens $(t_i, S_i, W_i, \tilde{W}_i)$, $i = 1, 2, 3$ with $t_1 = t_2$ and hidden metadata bits b_i , respectively, from the verification equations, the client can compute $\Delta S = S_2 - S_1$, $\Delta \tilde{W} = \tilde{W}_2 - \tilde{W}_1 = \tilde{y} \Delta S$, and $\Delta W = W_2 - W_1$. Let us now forge $(t_4, S_4, W_4, \tilde{W}_4) =$

⁵ In the January 13, 2021 version (20210113:200918) of the paper, \tilde{W}' was returned by the issuer without the $\tilde{\pi}$ proof, so a malicious issuer could hide a marker in it and break unlinkability. The authors corrected the protocol on April 21, 2022 after our disclosure.

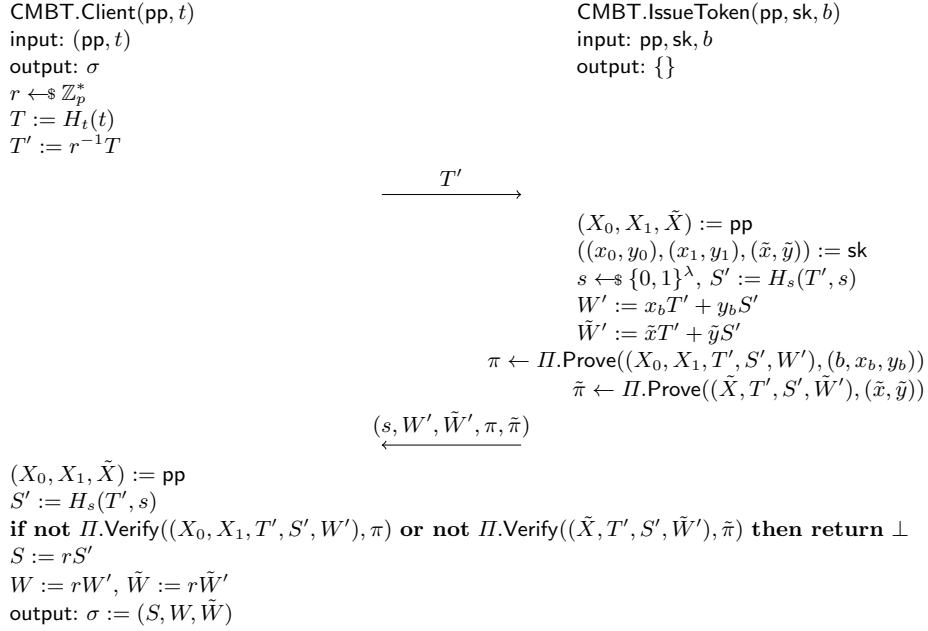


Fig. 2. CMBT token issuance protocol as given in [13, Fig 22, p 55] (current version 20220421:171853).

$(t_3, S_3 + \Delta S, W_3 + \Delta W, \tilde{W}_3 + \Delta \tilde{W})$. This is a token for which `Verify` returns `true`. Note that this does not violate the unforgeability result because the security game does not take as a valid forgery the creation of a new token with the same tag $t_3 = t_4$. However, `ReadBit` returns a bit if and only if $b_1 = b_2 = b_3$. In other words, either this new token is taken as valid but returns no bit, or it is fully valid and hides the same bit as all other tokens. Hence, having access to an oracle telling whether a bit is returned or not breaks the privacy of the metadata bit.

We can draw the following conclusion: if a side channel attack gives access to a *validity* oracle (i.e. whether `AT.ReadBit` returns \perp or not once it is known that `AT.Verify` is true), we have an attack against the privacy of the metadata bit b .

2.4 Anonymous Tokens with Public Metadata

Silde and Strand [16] design a protocol to add *public* metadata in Privacy Pass while also extending it for a private metadata bit and for public verifiability of the token. They consider a case study with the reporting phase in digital contact tracing when a user reports proximity keys with an anonymous token from the authority who has the positive test. To avoid changing credential secret keys every day, they add the date in public metadata md of the token.

Their scheme with private metadata is adapted from CMBT. It has similar problems as the ones we explained in the previous section. Moreover, the re-

demption scheme verifies validity and reads the hidden bit at the same time. However, given two tokens (S_i, W_i) ($i = 1, 2$) with the same (t, md, b) , a combination of the two tokens is a valid token for (t, md, b) as well while two tokens with the same (t, md) and different b , a combination is a token which reads no bit. The PMB game for the privacy of the metadata bit is specified with a $\mathcal{O}_{\text{verify}}$ oracle but it is not clear what this oracle answers for this scheme. If $\mathcal{O}_{\text{verify}}$ always returns `true`, this attack does not contradict the security result. However, if $\mathcal{O}_{\text{verify}}$ returns whether the token reads any bit (as it should do) then PMB security is broken.

3 Anonymous Tokens Revisited

3.1 AT Interface

We revisit the interface of AT and update the security notions as follows.

- $(\text{crs}, \text{td}) \leftarrow \text{AT.Setup}(1^\lambda)$ the protocol requires a set up with security parameter λ , and returns parameters common to all parties, denoted `crs`. Depending on the scheme, this will contain a group description, the security parameter, and could also contain a common reference string `crs` and a trapdoor `td`.
- $(\text{pp}, \text{sk}) \leftarrow \text{AT.KeyGen}(\text{crs})$, the public parameters `pp` and a secret key `sk` are generated by the token *issuer* with the `KeyGen` algorithm which takes `crs` as input.
- $\langle t, \sigma, \perp \rangle \leftarrow \langle \text{AT.Client}(\text{pp}, m), \text{AT.IssueToken}(\text{sk}, b, m) \rangle$, the interactive token generation protocol runs between a *client* (also called a *user*) and an issuer. The client input is the issuer's public parameters and the *public metadata* m . The issuer inputs are the secret key `sk`, a *metadata bit* $b \in \{0, 1\}$, and a public metadata m . (Both participants are assumed to agree on m .) The protocol outputs a token (t, σ) composed of a tag t , the public metadata m , and a token σ for the client and nothing (\perp) for the issuer. The elements m , b , and t are called *attributes* and can be optionally offered by the protocol. The attribute t is a *nonce*; b is the issuer's private metadata bit; m is a public metadata (on which both participants must agree). As we focus on a round-trip protocol which is initiated by the client (2-move, client-initiated), we can specify it by three algorithms:
 - $\text{AT.ClientQuery}(\text{pp}, m) \rightarrow (\text{query}, \text{st})$ //Client sends query to issuer
 - $\text{AT.IssueToken}(\text{sk}, b, m, \text{query}) \rightarrow \text{resp}$ //Server replies with resp
 - $\text{AT.ClientFinal}(\text{st}, \text{resp}) \rightarrow (t, \sigma)$ //Client locally computes token
- $\text{ind} \leftarrow \text{AT.ReadBit}(\text{sk}, m, t, \sigma)$, the verification algorithm is run with inputs of a secret key and a token σ with attributes (m, t) . It outputs either a bit, in which case we say the token is valid, or \perp in which case we say the token is invalid.

We deviate from the previous definitions in three ways: first of all, there is a unique `AT.ReadBit` algorithm (and no extra `AT.Verify`) which returns the hidden bit or \perp if invalid. Second, the client no longer chooses the t input in the issuing

protocol. Instead, a unique nonce t is returned to the client. Finally, we added the optional *public metadata* m attribute. For protocols not allowing it, input m is ignored in algorithms and games.

Security properties of an AT scheme are unforgeability, unlinkability, and privacy of metadata bit.

1. **Unforgeability** implies that an adversary cannot create valid tokens with modified attributes on an existing token. More precisely, if the issuer is invoked $n_{b,m}$ times for each attribute (b, m) , then, for no (b, t) the adversary can exhibit $n_{b,m} + 1$ valid tokens with pairwise different tags t . The adversary has access to a $\text{ReadBit}(\text{sk}, \cdot, \cdot)$ oracle and can choose the bit b to be hidden in the token by the issuer.
2. **Unlinkability** implies that a malicious issuer cannot link a redeemed token with one of the issuing sessions. The malicious issuer can maliciously set up the public parameters.
3. **Privacy** of metadata means that a malicious client cannot guess the metadata bit hidden during a issuing session, even with access to an oracle for checking if a token is valid (but without access to an oracle which extracts the bit).

3.2 Unforgeability

The one-more unforgeability game (OMUF) is defined on Figure 3.⁶ This is the same as in Kreuter et al. [14] except for a modification in the quantifiers⁷ and for the modification in the interface: having the token verification and the bit extraction in the same algorithm and oracle. Notably, (t, σ) making AT.Verify return true and AT.ReadBit return \perp would not exist any more as there is no AT.Verify .

Definition 1. *In the OMUF game on Figure 3, we define the advantage of an adversary \mathcal{A} by*

$$\text{Adv}_{\mathcal{A}}^{\text{OMUF}}(\lambda) = \Pr[\text{win}]$$

We say that AT is OMUF-secure if for any PPT adversary \mathcal{A} , the advantage is a negligible function.

⁶ For protocols with no public metadata m , the variables $n_{b,m}$ in the game shall be changed to n_b .

⁷ In their OMUF security definition, the first condition says that both $q_0 \leq \ell$ AND $q_1 \leq \ell$ where q_b is the number of oracle queries with bit b . Suppose, the adversary made 10 queries with $b = 0$ ($q_0 = 10$) and 1000 queries with $b = 1$ ($q_1 = 1000$). If the adversary forges 11 tokens with $b = 0$, for this to succeed as a forgery, ℓ must be at least 1000. If it is not, this does not succeed.

Game OMUF(1^λ):	Oracle $\mathcal{O}_{\text{sign}}(b, m, \text{query})$:
1: AT.Setup(1^λ) \rightarrow (crs, td)	7: increment $n_{b,m}$
2: AT.KeyGen(crs) \rightarrow (pp, sk)	8: AT.IssueToken(sk, b, m, query) \rightarrow resp
3: initialize $n_{b,m} \leftarrow 0$ for all (b, m)	9: return resp
4: $\mathcal{A}^{\mathcal{O}_{\text{sign}}, \mathcal{O}_{\text{read}}}(\text{crs}, \text{pp}) \rightarrow b, m, (t_i, \sigma_i)$	
5: if $\#\{t_i\} \leq n_{b,m}$ then abort	Oracle $\mathcal{O}_{\text{read}}(m, t, \sigma)$
6: win iff AT.ReadBit(sk, t_i , m, σ_i) = b for all i	10: return AT.ReadBit(sk, m, t, σ)

Fig. 3. One-More UnForgeability Game

3.3 Unlinkability

The unlinkability game (UNLINK) is defined on Figure 4. This is the same as in Kreuter et al. [14] except for the modification in the interface: the tag t is output instead of being a arbitrarily selected input by the client and the public metadata m must be the same for all challenge tokens.

Definition 2. *In the UNLINK game on Figure 4, we define the advantage of an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ for parameter n by*

$$\text{Adv}_{\mathcal{A}, n}^{\text{UNLINK}}(\lambda) = \Pr[\text{win}]$$

We say that AT is κ -UNLINK-secure if for any PPT adversary \mathcal{A} and any integer n , the advantage bounded by $\frac{\kappa}{n}$ plus a negligible function.

3.4 Privacy of the Metadata Bit

The privacy of the metadata bit game (PMB) is defined on Figure 5 with a challenge bit b^* . This is the same as in Kreuter et al. [14] except for the modification in the interface: the verify oracle is implemented by checking if AT.ReadBit does not return \perp . We also modified to have a single access to $\mathcal{O}_{\text{chal}}$ and to give access to $\mathcal{O}_{\text{read}}$ until the challenge is released. In the case of [14], the separation between AT.Verify and AT.ReadBit allowed $\mathcal{O}_{\text{verify}}$ to return true although AT.ReadBit would return \perp . Our interface does not allow it any more so the adversary has more information.

Definition 3. *In the PMB game on Figure 5, we define the advantage of an adversary \mathcal{A} by*

$$\text{Adv}_{\mathcal{A}}^{\text{PMB}}(\lambda) = \Pr[\text{PMB}_1 \rightarrow 1] - \Pr[\text{PMB}_0 \rightarrow 1]$$

We say that AT is PMB-secure if for any PPT adversary \mathcal{A} , the advantage is a negligible function.

Game UNLINK $_n(1^\lambda)$:

- 1: AT.Setup(1^λ) \rightarrow (crs, td)
- 2: initialize $\mathcal{Q}_{\text{query}}, \mathcal{Q}_{\text{final}} \leftarrow \emptyset$
- 3: $\mathcal{A}_1(\text{crs}) \rightarrow (\text{pp}, \text{state}_1)$
- 4: $\mathcal{A}_2^{\mathcal{O}_{\text{query}}, \mathcal{O}_{\text{final}}}(\text{state}_1) \rightarrow$
 $(\mathcal{Q}, (\text{resp}_i)_{i \in \mathcal{Q}}, \text{state}_2)$
- 5: if $\mathcal{Q} \not\subseteq \mathcal{Q}_{\text{query}} - \mathcal{Q}_{\text{final}}$ then abort
- 6: if $\#\mathcal{Q} < n$ then abort
- 7: **for** all $i \in \mathcal{Q}$ **do**
- 8: $\text{out}_i \leftarrow \text{AT.ClientFinal}(\text{st}_i, \text{resp}_i)$
- 9: if $\text{out}_i = \perp$ then abort
- 10: parse $(t_i, \sigma_i) \leftarrow \text{out}_i$
- 11: **end for**
- 12: if $\#\{m_i; i \in \mathcal{Q}\} > 1$ then abort
- 13: $i^* \leftarrow \$ \mathcal{Q}$
- 14: pick a random permutation σ of \mathcal{Q}
- 15: $\mathcal{A}_3(\text{state}_2, \text{out}_{i^*}, (\text{out}_{\sigma(i)})_{i \in \mathcal{Q}}) \rightarrow i$
- 16: win iff $i = i^*$

Oracle $\mathcal{O}_{\text{query}}(i, m)$:

- 17: if $i \in \mathcal{Q}_{\text{query}}$ then **return**
- 18: insert i in $\mathcal{Q}_{\text{query}}$
- 19: $m_i \leftarrow m$
- 20: AT.ClientQuery(pp, m_i) $\rightarrow \text{query}_i, \text{st}_i$
- 21: **return** query_i

Oracle $\mathcal{O}_{\text{final}}(i, \text{resp})$:

- 22: if $i \in \mathcal{Q}_{\text{final}}$ or $i \notin \mathcal{Q}_{\text{query}}$ then **return**
- 23: insert i in $\mathcal{Q}_{\text{final}}$
- 24: $\text{resp}_i \leftarrow \text{resp}$
- 25: **return** AT.ClientFinal($\text{st}_i, \text{resp}_i$)

Fig. 4. Unlinkability Game

Game PMB $_{b^*}(1^\lambda)$:

- 1: AT.Setup(1^λ) \rightarrow (crs, td)
- 2: AT.KeyGen(crs) \rightarrow (pp, sk)
- 3: flag \leftarrow false
- 4: **return** $\mathcal{A}^{\mathcal{O}_{\text{sign}}, \mathcal{O}_{\text{chal}}, \mathcal{O}_{\text{read}}, \mathcal{O}_{\text{valid}}}(\text{crs}, \text{pp})$

Oracle $\mathcal{O}_{\text{sign}}(b, m, \text{query})$:

- 5: AT.IssueToken(sk, b, m, query) \rightarrow resp
- 6: **return** resp

Oracle $\mathcal{O}_{\text{chal}}(m, \text{query})$:

- 7: if flag then **return** \perp
- 8: flag \leftarrow true
- 9: $m^* \leftarrow m$
- 10: AT.IssueToken(sk, b^*, m^*, query) \rightarrow
resp
- 11: **return** resp

Oracle $\mathcal{O}_{\text{read}}(m, t, \sigma)$:

- 12: if flag and $m = m^*$ then **return** \perp
- 13: **return** AT.ReadBit(sk, m, t, σ)

Oracle $\mathcal{O}_{\text{valid}}(m, t, \sigma)$:

- 14: **return** whether AT.ReadBit(sk, m, t, σ) \neq
 \perp

Fig. 5. Privacy of the Metadata Bit Game

4 ATHM: Anonymous Token with Hidden Metadata

Instead of relying on a deterministic PRF, we construct a protocol which is based on a randomized algebraic MAC. Our proposed protocol gets inspired from previous works on algebraic MACs used in anonymous credentials [6,7] and Signal’s private group management in group chats [8,9]. A valid MAC for an input b with a nonce t and a secret key (x, y, z) is a pair $\sigma = (P, Q)$ such that $Q = (x + by + tz)P$. The security assumptions for this MAC are detailed in the next section. In our scheme, a token with a hidden bit b will be a pair (t, σ) such that σ is a valid MAC for the attributes (b, t) with a secret key (x, y, z) . Several variants and extensions of ATHM are possible such as introducing public metadata m or using another MAC algorithm. In this section, we focus on the simplest options. The various options and the generalized protocol are presented in appendix.

4.1 The ATHM Components

Our scheme uses as a building block a simulatable non-interactive proof Π_2 .

Setup algorithm. Setup is composed of four phases. The Setup_1 algorithm generates an (additive) group, which is cyclic, of prime order p . The Setup_2 algorithm selects a group generator G . The Setup_4 algorithm selects common parameters for Π_2 (typically, a second group generator H and its logarithm). The crs string includes p , group parameters gp to be able to do operations in the group, the generator G , and crs_2 .

$\text{Setup}(1^\lambda)$:

- 1: $\text{Setup}_1(1^\lambda) \rightarrow (\text{gp}, p)$ ▷ group setup
- 2: $\text{Setup}_2(\text{gp}, p) \rightarrow G$ ▷ generator setup
- 3: $\text{Setup}_4(\text{gp}, p, G) \rightarrow (\text{crs}_2, \text{td}_2)$ ▷ Π_2 setup
- 4: $\text{crs} \leftarrow (\text{gp}, p, G, \text{crs}_2)$
- 5: $\text{td} \leftarrow \text{td}_2$

The KeyGen algorithm. Key generation is composed of several phases.

$\text{KeyGen}(\text{crs})$:

- 1: $\text{KeyGen}_0(\text{crs}) \rightarrow (\text{pp}_0, \text{sk}_0)$
- 2: $\text{KeyGen}_2(\text{crs}, \text{pp}_0, \text{sk}_0) \rightarrow (\text{pp}_2, \text{sk}_2)$ ▷ Π_2 key generation
- 3: $\text{pp} \leftarrow (\text{pp}_0, \text{pp}_2)$
- 4: $\text{sk} \leftarrow (\text{sk}_0, \text{sk}_2)$

In KeyGen_0 , the issuer selects three secrets $\text{sk}_0 = (x, y, z)$ with $y, z \neq 0$, and sets $\text{pp}_0 = Z = zG$.

Our proposed Π_2 scheme requires KeyGen_2 to add in pp_2 Pedersen commitments $C_x = xG + r_xH$, $C_y = yG + r_yH$, together with a proof of knowledge of (x, r_x, y, r_y) . Clients are assumed to verify this proof before starting the issuance protocol, but this is done only once for all.

The token issuance protocol. The user has public parameters. The server's input is the secret (x, y, z) and a bit b to hide inside the token. The protocol works as depicted on Figure 6: the client selects a random tag share t_C and a random mask $r \in \mathbb{Z}_p$ and sends $T = t_C Z + rG$ to the issuer. The issuer selects a random tag share t_S and generates a pair (U, V) such that $(U, V - rU)$ is a valid MAC for tag $t = t_C + t_S$ and metadata b with key (x, y, z) . For this, the issuer selects $U = dG$ for a random $d \in \mathbb{Z}_p^*$ and $V = d(xG + byG + t_S zG + T)$. A proof Π_2 must prove that the (U, V, t_S) triplet was correctly generated. Note that Π_2 must prove that b is a bit. Then, (U, V, t_S, π) is returned. The client computes $(U, V - rU)$. To make it unlinkable, the pair is multiplied by a random mask c to obtain another pair $\sigma = (P, Q)$.

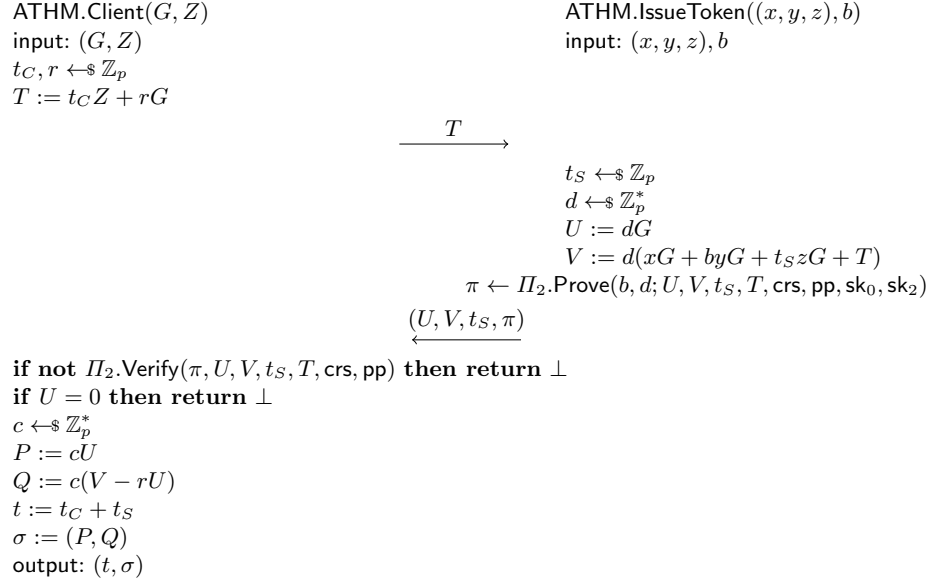


Fig. 6. ATHM token issuance protocol.

The proof Π_2 is a Fiat-Shamir transform of an OR proof of two Schnorr proofs for $b = 0$ and $b = 1$. It is specified in subsection 4.3.

The ReadBit algorithm. The redemption of (t, P, Q) with (x, y, z) checks for which $b \in \{0, 1\}$, the equality $Q = (x + by + tz)P$ is satisfied.

Rationales. We can first observe that it is necessary that the issuer has an influence on the final t . If the client could decide $t = t_C$, then PMB security could be broken by getting a challenge with tag t^* then issuing a token with same tag $t = t^*$ and taking the linear combination of both with coefficient $\frac{1}{2}$ and $\frac{1}{2}$. The obtained ticket is valid if and only if b^* is equal to the bit b put in the second token.

We can also observe that it is necessary to have t_C , r , and c . Without any of them, the malicious issuer can easily link an issued token with a redeemed token. The same goes with the proof π without which the issuer could change the secret and use it as a marker.

Beware of double spending. Note that for this protocol, tag t needs to be a nonce as in other protocols [13], [14] i.e. the redeemer should check against double-spending of a token with the same t . Otherwise, it is easy to transform a valid token with tag t into another valid token with the same tag t : let $\sigma = (P, Q)$ be a signature on t . Then the client can forge another signature $\sigma' = (P', Q')$ on t as follows: Client samples another c' and computes $P' := c'P$ and $Q' := c'Q$. The redeemer can check that $Q' = (x + by + tz)P'$ and take (t, σ') as another valid token.

4.2 The MAC Building Block

Our security results will be based on the security of an algebraic MAC. The simplest one is the MACGGM algorithm [6] defined as follows: given a secret $(x, y, z) \in \mathbb{Z}_p^3$, a valid authentication for $(b, t) \in \mathbb{Z}_p^2$ is a pair $\sigma = (P, Q)$ such that $Q = (x + by + tz)P$. For this MAC to be secure, it is important that no adversary can find any linear relation between the random values of P . Hence, P is selected at random by the issuer.

The security of MACGGM was proven in the generic group model (GGM) [6]. So, we use the same model to prove the security of ATHM. However, our construction generalizes to other MAC algorithms which can be proven in the standard model, and we use non-GGM security for this generalization, as shown in appendix.

4.3 The Simulatable Proof Building Block

We assume that crs_2 is a new generator H . The trapdoor td_2 is the discrete logarithm of H . We further assume that pp_2 includes some commit $C_x = xG + r_xH$, $C_y = yG + r_yH$ together with a proof of knowledge of (x, y, r_x, r_y) . The issuer's secret sk_2 must keep r_x and r_y .

The proof π consists of first releasing a commit $C = bC_y + \mu H$, proving (with OR proof) that C either commits to 0 or to C_y , and proving (with Schnorr proof) knowledge of (d', ρ, w) such that $-G = d'U$, $-(C_x + C + t_S Z + T) = d'V + \rho H$, and $-T = d'V + wG$. The link with d is that $d' = -\frac{1}{d}$. Hence

$$\exists(d', \rho, w) \quad d' \begin{pmatrix} U \\ V \\ V \end{pmatrix} + \rho \begin{pmatrix} 0 \\ H \\ 0 \end{pmatrix} + w \begin{pmatrix} 0 \\ 0 \\ G \end{pmatrix} = - \begin{pmatrix} G \\ C_x + C + t_S Z + T \\ T \end{pmatrix} \quad (3)$$

where $\rho = -(r_x + br_y + \mu)$ and $w = x + by + t_S z$. Note that the issuer already computed $t_S Z$ in the protocol. However, the client would have to compute it. This costs one multiplication by the client. In addition, the computation of C by

the issuer costs one multiplication for μH (the bC_y multiplication is free since $b = 0$ or $b = 1$).

This proof consists of one OR proof and an independent AND proofs. Typically, we can merge them, but, for simplicity, we start with OR proof first and merge them as one protocol later.

OR Proof. The server needs to prove knowledge of a μ such that $C = bC_y + \mu H$ with $b \in \{0, 1\}$. It first picks a random r_μ and commits to it with $C_b = r_\mu H$. Picks a random e_{1-b} and a simulated response a_{1-b} to get the alternate commit $C_{1-b} = a_{1-b}H - e_{1-b}(C - (1-b)C_y)$ for the flipped bit $1-b$. The commit message is (C_0, C_1) . Based on the challenge e , set $e_b = e - e_{1-b}$ and compute the response $a_b = r_\mu + e_b\mu$ for the real bit b . The answer is the message (e_0, a_0, a_1) .

The verifier receives $(C_0, C_1, e_0, a_0, a_1)$, computes $e_1 = e - e_0$, and checks

$$a_0H == C_0 + e_0C \text{ and } a_1H == C_1 + e_1(C - C_y)$$

To formally define **Prove**, we merge this OR proof with the AND proofs with statement given in Equation 3 and we transform into a non-interactive proof. We formally define the algorithms in Π_2 below.

Setup₄. The algorithm **Setup₄**(gp, p, G) first selects $\text{td}_2 \in \mathbb{Z}_p^*$ and sets $\text{crs}_2 = H = \text{td}_2.G$.

KeyGen₂. The algorithm **KeyGen₂**($\text{crs}, Z, \text{sk}_0$) parses $\text{crs} = (\text{gp}, p, G, \cdot, H)$ and $\text{sk}_0 = (x, y, z)$, picks $r_x, r_y \in \mathbb{Z}_p$, and computes $C_x = xG + r_xH$ and $C_y = yG + r_yH$. We have

$$\text{sk}_2 = (r_x, r_y)$$

In the generic group model, an issuer who wants to maliciously register a (C_x, C_y, Z) key must know how to express C_x , C_y , and Z in terms of the only group elements which are available at setup, i.e. G and H . Hence, an extractor for (x, y, r_x, r_y) is trivial. For Z , the only thing to prove is that the component in H is zero. This can be done with a Schnorr proof: **KeyGen₂** picks a random ρ_z , computes $\Gamma_z = \rho_z G$, $\varepsilon = \text{Hash}(G, H, Z, \Gamma_z)$, $a_z = \rho_z + \varepsilon z$, and adds (ε, a_z) in pp_2 . We have

$$\text{pp}_2 = (C_x, C_y, \varepsilon, a_z)$$

To verify the proof (this must be done at least once for all by **ClientQuery**), the client computes $\Gamma_z = a_z G - \varepsilon Z$ and checks $\varepsilon = \text{Hash}(G, H, Z, \Gamma_z)$.

The protocol imposes that Z and Γ_z must be known before hashing, hence before ε is randomly selected by the hash function (otherwise, hashing gives an incorrect ε , except with $\frac{1}{p}$ probability). In the generic group model, this implies that the the expression of Z and Γ_z in terms of provided group elements is known. Then, a solution a_z can be searched for the equation $\Gamma_z = a_z G - \varepsilon Z$. As explained in the next paragraph, the generic group model advocates that finding a non-trivial linear relation between provided random group elements can be done with probability bounded by $\text{Adv}^{\text{DLOG}} + \frac{1}{p}$, where Adv^{DLOG} is the

advantage of a discrete logarithm solver of similar complexity. Here, we have $k = 2$ with G and H . Let $Z = zG + uH$. Except with probability $\text{Adv}^{\text{DLOG}} + \frac{1}{p}$, the coefficients v of Γ_z in H must satisfy $u = -ev$, with u and v fixed before e is picked. So, either the malicious issuer made a lucky selection (this can happen with probability $\frac{1}{p}$), or $u = v = 0$, meaning $Z = zG$. Therefore, a straightline extraction succeeds, except with probability $\text{Adv}^{\text{EXTRACT}} = \text{Adv}^{\text{DLOG}} + \frac{3}{p}$.

In the generic group model, we may be interested in several linear combinations of the provided group elements. By using the \mathcal{O}_{cmp} oracle, we can see if one linear combination vanishes. If any such situation occurs, a non-trivial linear combination is found. Given some DLOG input (X, Y) , we can simulate each of the provided random elements by random linear combination of X and Y . A non-trivial linear relation between them gives a linear relation between X and Y . Except with probability $\frac{1}{p}$, this yield the discrete logarithm of Y in basis X . This is why we can have an advantage overhead of $\text{Adv}^{\text{EXTRACT}}$ and assume that no nonzero linear combination of the provided elements vanishes.

Prove. The algorithm $\text{Prove}(b, d; U, V, t_S, T, \text{crs}, \text{pp}, \text{sk}_0, \text{sk}_2)$ parses different elements, picks μ , sets $C = bC_y + \mu H$, $d' = -\frac{1}{d}$, $\rho = -(r_x + br_y + \mu)$, and $w = x + by + t_S z$.

The issuer (prover) picks $e_{1-b}, a_{1-b}, r_\mu, r_d, r_\rho, r_w$ at random and computes $C_b = r_\mu H$, $C_{1-b} = a_{1-b}H - e_{1-b}(C - (1-b)C_y)$,

$$\begin{pmatrix} C_d \\ C_\rho \\ C_w \end{pmatrix} = r_d \begin{pmatrix} U \\ V \\ V \end{pmatrix} + r_\rho \begin{pmatrix} 0 \\ H \\ 0 \end{pmatrix} + r_w \begin{pmatrix} 0 \\ 0 \\ G \end{pmatrix}$$

$e = \text{Hash}(G, H, C_x, C_y, Z, U, V, t_S, C, C_0, C_1, C_d, C_\rho, C_w)$, $e_b = e - e_{1-b}$, $a_b = r_\mu + e_b \mu$, and $(a_d, a_\rho, a_w) = (r_d, r_\rho, r_w) + e(d', \rho, w)$. Finally, the output is

$$\pi = (C, e_0, e_1, a_0, a_1, a_d, a_\rho, a_w)$$

Verify. The algorithm $\text{Verify}(\pi, U, V, t_S, T, \text{crs}, \text{pp})$ parses π , crs , and pp , computes $C_0 = a_0H - e_0C$, $C_1 = a_1H - e_1(C - C_y)$, $e = e_0 + e_1$,

$$\begin{pmatrix} C_d \\ C_\rho \\ C_w \end{pmatrix} = a_d \begin{pmatrix} U \\ V \\ V \end{pmatrix} + a_\rho \begin{pmatrix} 0 \\ H \\ 0 \end{pmatrix} + a_w \begin{pmatrix} 0 \\ 0 \\ G \end{pmatrix} + e \begin{pmatrix} G \\ C_x + C + t_S Z + T \\ T \end{pmatrix}$$

then verifies $e = \text{Hash}(G, H, C_x, C_y, Z, U, V, t_S, C, C_0, C_1, C_d, C_\rho, C_w)$.

SimKeyGen. The algorithm $\text{SimKeyGen}(\text{crs}, Z)$ parses crs to retrieve G and H . It picks a fully random pp_2 , computes $\Gamma_x, \Gamma_y, \Gamma_z$ like in verification, then programs the random oracle to have a consistent Hash . Since no hash is done before, the simulation is perfect.

Simulate. The algorithm $\text{Simulate}(U, V, t_S, T, \text{crs}, \text{pp}, \text{td}_2)$ parses inputs. Then, it picks a random proof π with same distribution, i.e. a group element C , $e_0, e_1, a_d, a_0, a_1, a_\rho, a_w \in \mathbb{Z}_p$ all uniform and independent. It can then proceed as in $\text{Verify}(\pi, U, V, t_S, T, \text{crs}, \text{pp})$ to compute $C_0, C_1, C_d, C_\rho, C_w$, then the hash. Then, the random oracle is programmed to output the correct hash $e = e_0 + e_1$. Simulation fails if the random oracle cannot be programmed on the selected hash input, which happens with probability bounded by $\frac{q_{\text{Hash}}}{p^6}$, where q_{Hash} is the number of previous oracle calls to Hash . Otherwise, the simulation is perfect.

Straightline extraction. Extraction in the generic group model is straightline: the malicious issuer must know how to express $U, V, C, C_0, C_1, C_d, C_\rho, C_w$ in terms of provided group elements (i.e. G, H, T , and prior elements T_i provided by clients) when they are presented to the random oracle, hence before knowing e , or fails except with probability $\frac{1}{p}$. This expression is extractable in the generic group model. After e is known, the malicious issuer can adjust $e_0, a_0, a_1, a_d, a_\rho, a_w$ to satisfy the equations. As mentioned before, finding a non-trivial linear combination of random provided group elements is limited to $\text{Adv}^{\text{DLOG}} + \frac{1}{p}$. (As explained before, the $\text{Adv}^{\text{DLOG}} + \frac{1}{p}$ advantage overhead does not need to be counted multiple times.) We consider it for the 5 equations in $C_0, C_1, C_d, C_\rho, C_w$ and reduce to equalities for each component.

We have $C_0 + C_1 = (a_0 + a_1)H - eC + e_1C_y$. For the sum of components in T and the T_i in C , we deduce that it is zero unless good luck (with probability $\frac{1}{p}$). Looking at the equation in C_0 and C_1 , this implies that the coefficients in T or the T_i in C, C_0, C_1 are zero. Let α, β , and b (we do not know if b is a bit yet) be the components of C_0, C_1 , and C in G . We obtain $\alpha = -e_0b$ and $\beta = -e_1(b-1)$. Hence, $(b-1)\alpha + b\beta = -eb(b-1)$. Since α, β, b are fixed before e is picked, either there is good luck (this happens with probability $\frac{1}{p}$), or $b(b-1) = 0$ which means that b is a bit. Hence, we can write $C = bC_y + \mu H$.

For the sum of components in H, T and the T_i in U , the equation in C_d shows that either there is a good luck (with probability $\frac{1}{p}$) or it is zero. Hence, both C_d and U are multiple of G . We write $U = dG$. Similarly, the equation in C_w shows no component in H or the T_i in V (unless good luck of $\frac{1}{p}$) and we can write $V = wG + dT$.

As V has no component in the T_i , this is the same for C_ρ . Finally, in the equation in C_ρ we focus on the components in G and T . We let C'_ρ be the sum of these two components for C_ρ . We obtain $C'_\rho = a_dV + e((x + by + t_Sz)G + T)$. The values C'_ρ and V are fixed before e is picked. Except good luck (with probability $\frac{1}{p}$), we deduce V proportional to $(x + by + t_Sz)G + T$. Due to the expression $V = wG + dT$, we deduce $w = d(x + by + t_Sz)$. Hence, $V = d(x + by + t_Sz)G + dT$.

All in all, extraction works, except with probability $\text{Adv}^{\text{SOUND}} = \text{Adv}^{\text{DLOG}} + \frac{7}{p}$.

5 Performance

Implementation. We implemented our construction as given in Figure 6 in Rust (version 1.66.0). We use the Ristretto group using curve25519-dalek library. We

use `RistrettoBasePointTable` struct (which is a precomputed table for multiplications with the group generator G) to accelerate the scalar multiplications (PMBT implementation does *not* use these tables). We use two of these tables: one for G and one for H which requires 60 KB of memory for constant time cryptography and up to 4 times speed up. For OR proofs and verification, we did *not* rely on any external library, meaning we implemented it in pure Rust. It is available at <https://github.com/Microsoft/MacTok>.

We benchmarked the implementation on a machine with Intel(R) i7-1185G7 3.00GHz CPU. Our benchmarks excludes the key generation (because it is generated only once for all). They include client blinded message generation, server’s computation of MACs (blindly, along with the proof π), client’s unblinding (along with the verification of π), and server’s redemption. It takes 1.3 ms whereas PMBT takes 1.6 ms for the same operations⁸. We note that we disabled SIMD optimizations (which allows `curve25519-dalek` to run faster curve operations) in both ATHM and PMBT due to the unstable version (1.66.0-nightly) of the Rust compiler that does not allow building PMBT. When we run our ATHM protocol with SIMD optimization, we get 0.9 ms of running time.

Theoretical Complexity. We also computed the number of scalar multiplications to compare ATHM with PMBT and observed that ATHM computes 29 scalar multiplication whereas PMBT computes 31 multiplications in total, for benchmarked operations (client and server side computations including redemption along with the proof and verification).

In ATHM, the issuer computes 11 scalar multiplications during issuance (one for C , 7 for the proof, and 3 for the ATHM protocol). The client computes 17 scalar multiplications (one for $t_S Z$, 11 for the proof, and 5 for ATHM). The total is 28 multiplications.⁹ Furthermore, the number of transmission is of 4 group elements (C in the proof and (T, U, V) in the proof) and 8 integers (t_S in ATHM and the 7 elements of π). The total is 12.¹⁰ For redeem, the number of multiplications is 1 and the token length is 3 (t , P , and Q). For key generation, the issuer computes 2 multiplications (1 for Z and 1 for pp_2) and the client computes 2 multiplications for the II_2 verification. The public key contains 5 elements (1 for Z and 4 for pp_2).

As a comparison, in PMBT, for issuance, the issuer computes 12 multiplication in total and the client computes 15 multiplications. In total, the number of multiplications is 27 for issuance. The number of transmissions is 2 group elements and 7 scalars. The total is 9. The redemption needs 4 multiplication with a token length 2 group elements. For key generation, the issuer computes 4 multiplication. The public key contains 4 elements.

⁸ PMBT code is available at <https://github.com/mmaker/anonymous-tokens>

⁹ In this count, we took the computation of $(1 - b)C_y$ as free. Furthermore, the computation of $r_d V$ and $a_d V$ are done twice but count for a single operation.

¹⁰ By setting $t_S = \text{Hash}(U)$, the issuer would not have to send t_S any longer and save the transmission of one \mathbb{Z}_p element.

6 Security Proof for ATHM in the Generic Group Model

We consider ATHM as specified in section 4: with the MACGGM algorithm, no T_{ext} , and with the nonce t and private metadata bit b attributes. Considering more attributes would work the same. We prove OMUF and PMB security in the generic group model. The options for Π_2 do not play any role.

In the generic group model, all group operations are outsourced to an external oracle which also keeps the group element values. Initially, the oracle has registers set to the group elements from the initialization (such as G , other crs values, Z , and other pp values). These registers are given an address that the adversary or the game can use. The adversary uses addresses as references when a group operation is requested but never sees the element value itself. The output of the oracle is either the address where the result is stored or a numeric value. Actually, the oracle only computes subtractions (from which we can do additions, scalar multiplications, inversion, and get the neutral element) and comparisons (whether or not the group elements referred to by two addresses are equal).

The property of the generic group model is that each address can be mapped to a tuple of \mathbb{Z}_p elements such that the linear combination of the initial registers with these coefficients is equal to the content of the referred register. The adversary can keep track of these linear coefficients. When the operations are done by the game with coefficients unknown by the adversary, the adversary can associate these coefficients with unknowns and express (typically) the coefficients as polynomials in those unknowns.

6.1 OMUF Security

We consider an adversary \mathcal{A} playing the OMUF game. Without loss of generality, we assume that \mathcal{A} either aborts, or returns (b_j, t_j, P_j, Q_j) tuples which are all valid, with pairwise different t_j , same $b_j = b$, and in number equal to $n_b + 1$. Following the generic group model, we further assume that all group operations (selection of G , addition, scalar multiplication, comparison) which are done by the adversary, the game or the oracles, are outsourced to the generic group oracles. Only a register index for each group element is visible, as well as the Boolean result of the comparison oracle.

We first reduce to a game Γ_1 where the π proof in Π_2 is generated by a simulator using $\Pi_2.\text{SimKeyGen}$ and $\Pi_2.\text{Simulate}$. For our choice of Π_2 (see subsection 4.3), this simulation is almost perfect, so we have:

$$\text{Adv}^{\text{OMUF}} \leq \text{Adv}^{\Gamma_1} + \frac{q_{\text{Hash}}}{p^6}$$

In the i th query to $\mathcal{O}_{\text{sign}}$, we let $(b, \text{query}) = (b_i, T_i)$ denote the input and $(U_i, V_i, t_{S,i})$ denote the output.

Thanks to the generic group model, we can assume that every time the adversary \mathcal{A} produces a group element (such as T_i, P_j, Q_j or an input to $\mathcal{O}_{\text{read}}$), it comes with a representation as a linear combination of already seen group

elements. These seen group elements consist of G , Z , H , and every U_i and V_i which is returned by a previous $\mathcal{O}_{\text{sign}}$ oracle call. (C_x and C_y are now simulated by the adversary in Γ_1 .) Hence, for each group element A which is produced by \mathcal{A} , there is an extractor based on the view of \mathcal{A} which outputs the scalars of the following equation:

$$A = a_G \cdot G + a_Z \cdot Z + a_H \cdot H + \sum_i (a_{U_i} \cdot U_i + a_{V_i} \cdot V_i)$$

We write U_i and V_i in the form of a linear combination of G , Z , T_i with coefficients written as formal polynomials in the values d_i, x, y, z . I.e., we write

$$U_i = \bar{d}_i \cdot G \quad V_i = \bar{d}_i(\bar{x} + b\bar{y} + t_{S,i}\bar{z}) \cdot G + \bar{d}_i \cdot T_i$$

where $\bar{d}_i, \bar{x}, \bar{y}, \bar{z}$ are the formal variables of the polynomial which express the unknowns d_i, x, y, z respectively. They correspond to the values which are not revealed. There is another variable td_2 corresponding to td_2 . (We recall that $H = \text{td}_2 \cdot G$.) The returned value $t_{S,i}$ is given in clear so considered as a scalar. Hence, by induction, every group element A is expressed by a polynomial $\text{Pol}_A(\text{Val})$ multiplied by G , with $\text{Val} = ((d_i)_i, x, y, z, \text{td}_2)$. We consider the formal polynomial $\text{Pol}_A(\text{Var})$ with $\text{Var} = ((\bar{d}_i)_i, \bar{x}, \bar{y}, \bar{z}, \text{td}_2)$.

We prove by induction the following fact.

Fact 1 *For each A produced by the adversary after q queries to $\mathcal{O}_{\text{sign}}$, the Pol_A polynomial has total degree bounded by $q + 1$. Furthermore, every partial degree is bounded by 1. Monomials are square-free.*

Indeed, after $q = 0$ queries, the largest degree is for $\text{Pol}_Z(\text{Var}) = \bar{z}$. Making a new query to $\mathcal{O}_{\text{sign}}$ multiplies $\text{Pol}_{T_i}(\text{Var})$ by a fresh \bar{d}_i and adds a degree-2 polynomial $\bar{d}_i \cdot (\bar{x} + b\bar{y} + t_{S,i}\bar{z})$.

The values in Val are uniformly distributed and independent. In the generic group model, the adversary knows every polynomial Pol_A but not the evaluations $\text{Pol}_A((d_i)_i, x, y, z)$. The adversary can only know if two group elements are equal by using the comparison oracle, thus deduce that two polynomials evaluate to the same result. By using the Schwartz-Zippel lemma and the bound on the degree of polynomials, we have the following fact.

Fact 2 *Let q is the number of $\mathcal{O}_{\text{sign}}$ queries before an input (A, B) is presented for the first time to a comparison oracle. Except with a probability bounded by $\frac{q+2}{p}$, we have $A = B$ if and only if $\text{Pol}_A = \text{Pol}_B$.*

If the call to the comparison oracle is made by the adversary, the polynomials have degree bounded by $q + 1$ so we obtain the result. Otherwise, the call must come from the usage of AT.ReadBit in either the game of the $\mathcal{O}_{\text{read}}$ oracle, which multiply a degree- $(q + 1)$ -bounded polynomial by a degree-1 polynomial $\bar{x} + b\bar{y} + t\bar{z}$. So, the degree is bounded by $q + 2$.

The adversary, game, or $\mathcal{O}_{\text{read}}$ oracle can simulate that comparison oracle by checking equality between $\text{Pol}_A = \text{Pol}_B$. Hence, by induction, using hybrids,

we reduce to a game Γ_2 where no access to the comparison oracle is made, and for every final (b, t, P, Q) output, we have $\text{Pol}_Q(\text{Var}) = (\bar{x} + b\bar{y} + t\bar{z})\text{Pol}_P(\text{Var})$ (in winning cases). The total number of comparisons is bounded by $2(n + 1 + q_{\mathcal{O}_{\text{read}}}) + q_{\mathcal{O}_{\text{cmp}}}$, where $q_{\mathcal{O}_{\text{read}}}$ is the number of calls to $\mathcal{O}_{\text{read}}$ by the adversary and $q_{\mathcal{O}_{\text{cmp}}}$ is the number of calls to the comparison by the adversary. We obtain

$$|\text{Adv}^{\Gamma_1} - \text{Adv}^{\Gamma_2}| \leq (2(n + 1 + q_{\mathcal{O}_{\text{read}}}) + q_{\mathcal{O}_{\text{cmp}}}) \times \frac{n + 2}{p}$$

We now focus on one of the $n + 1$ final (b, t, P, Q) produced by \mathcal{A} in winning cases. We already assumed that $\text{Pol}_Q(\text{Var}) = (\bar{x} + b\bar{y} + t\bar{z})\text{Pol}_P(\text{Var})$. \mathcal{A} is only making scalar linear combinations of G, Z, H , and the U_i and V_i . We write

$$\begin{aligned} P &= a_G \cdot G + a_Z \cdot Z + a_H \cdot H + \sum_i (a_{U_i} \cdot U_i + a_{V_i} \cdot V_i) \\ Q &= b_G \cdot G + b_Z \cdot Z + b_H \cdot H + \sum_i (b_{U_i} \cdot U_i + b_{V_i} \cdot V_i) \\ &= (x + by + tz)P \end{aligned}$$

Given that the partial degree of Pol_Q in \bar{z} is bounded by 1, we can see from the last equation $\text{Pol}_Q = \text{Pol}_P \times (\bar{x} + b\bar{y} + t\bar{z})$ that $a_Z = 0$.

Pol_P has constant term a_G . From the last equation, Pol_Q has a_G as a coefficient of monomial \bar{x} . However, no monomial \bar{x} can appear in the linear expression of Pol_Q . (We have seen that $b_x = 0$ and \bar{x} never stands alone outside Pol_{C_x} . For instance, Pol_{V_i} is always a multiple of \bar{d}_i .) Hence, $a_G = 0$.

Similarly, $\bar{x} + b\bar{y} + t\bar{z}$ has no constant term so we must have $b_G = 0$.

By inspecting the monomial \bar{z} we now obtain that $b_Z = 0$.

Hence,

$$\begin{aligned} P &= \sum_i (a_{U_i} \cdot U_i + a_{V_i} \cdot V_i) \\ Q &= \sum_i (b_{U_i} \cdot U_i + b_{V_i} \cdot V_i) \\ &= (x + by + tz)P \end{aligned}$$

Given a polynomial and a variable \bar{u} , we say that \bar{u} appears if the partial degree in \bar{u} is at least 1. We have the following fact.

Fact 3 \bar{d}_i appears in Pol_P if and only if it appears in Pol_Q .

In the i th oracle call to $\mathcal{O}_{\text{sign}}$, we have

$$\text{Pol}_{V_i} = \bar{d}_i(\bar{x} + b_i\bar{y}) + t_{S,i}\bar{d}_i\bar{z} + \bar{d}_i \times \text{Pol}_{T_i}$$

with $t_{S,i}$ sampled as a fresh uniform scalar. Note that \bar{d}_i cannot appear in Pol_{T_i} because it is formed before sampling d_i . Pol_{T_i} may have the monomial \bar{z} (it is actually supposed to) but Pol_{T_i} is set before sampling $t_{S,i}$. Hence, except with probability $\frac{1}{p}$, the $\bar{d}_i\bar{z}$ monomial is present in Pol_{V_i} .

Fact 4 For any i , the monomial $\bar{d}_i \bar{z}$ has a nonzero coefficient in Pol_{V_i} , except with probability $\frac{1}{p}$.

We reduce to a game Γ_3 where $\bar{d}_i \bar{z}$ never has a zero coefficient in Pol_{V_i} for every i . We have

$$|\text{Adv}^{\Gamma_2} - \text{Adv}^{\Gamma_3}| \leq \frac{n}{p}$$

Fact 5 In Γ_3 , among all group elements given to \mathcal{A} , the monomial $\bar{d}_i \bar{z}$ has a nonzero coefficient in Pol_{V_i} and only in Pol_{V_i} .

Indeed, even though it could be put in a Pol_{T_j} for $j > i$, the monomial would be multiplied by \bar{d}_j and thus $\bar{d}_i \bar{z}$ would not appear as a standalone monomial in Pol_{V_j} . Coming back to the representation of a final P and Q , we deduce that $a_{V_i} = 0$ for every i (as we cannot have $\bar{d}_i \bar{z}^2$ in Pol_Q). By writing $\text{Pol}_{U_i} = \bar{d}_i$, Pol_P is linear in every \bar{d}_i . This implies that no monomial in a final Pol_Q can be divisible by any $\bar{d}_i \bar{d}_{i'}$.

We have $\text{Pol}_{V_i} = (\bar{x} + b_i \bar{y} + t_{S,i} \bar{z}) \bar{d}_i + \bar{d}_i \text{Pol}_{T_i}$. For every i such that $b_{V_i} \neq 0$, we deduce that no $\bar{d}_{i'}$ appear in Pol_{T_i} . The same holds for td_2 . Hence, for every i such that $b_{V_i} \neq 0$, we have that T_i must be a known linear combination of G and Z . We write $T_i = t_{C,i} Z + r_i G$. Hence, $\text{Pol}_{V_i} = \bar{d}_i (\bar{x} + b_i \bar{y} + (t_{C,i} + t_{S,i}) \bar{z})$. By writing $t_i = t_{C,i} + t_{S,i}$, we have

$$\begin{aligned} \text{Pol}_P &= \sum_i a_{U_i} \bar{d}_i \\ \text{Pol}_Q &= \sum_i (b_{U_i} + b_{V_i} (\bar{x} + b_i \bar{y} + t_i \bar{z} + r_i)) \bar{d}_i \\ &= \text{Pol}_P \times (\bar{x} + b \bar{y} + t \bar{z}) \end{aligned}$$

By inspecting \bar{d}_i we can further see that we must have $b_{U_i} = -b_{V_i} r_i$ and $a_{U_i} = b_{V_i}$. Hence, $P = \sum_i a_{U_i} U_i$ and $Q = \sum_i a_{U_i} (V_i - r_i U_i)$.

We say that the i th query is well formed if Pol_{T_i} is a linear combination of Pol_G and Pol_Z (i.e. that Pol_{T_i} is a polynomial in \bar{z} with degree bounded by 1: $\text{Pol}_{T_i} = t_{C,i} \bar{z} + r_i$). For each well formed query we can define $t_i = t_{C,i} + t_{S,i}$. It follows that for every i such that $a_{U_i} \neq 0$, we have that the i th query is well formed and that $b_i = b$ and $t_i = t$. This proves that for any valid (b, t, P, Q) , (P, Q) is a known linear combination of all $(U_i, V_i - r_i U_i)$ for well-formed queries satisfying $(b, t) = (b_i, t_i)$. Since P is nonzero, there exists i such that the i th query is well formed and $(b, t) = (b_i, t_i)$. Hence, the number of pairwise different t cannot exceed n_b . We deduce there is no winning case in Γ_3 .

We can wrap up by collecting all Adv overheads to get an upper bound for Adv^{OMUF} .

Theorem 1. For every \mathcal{A} playing OMUF in the generic group model and making n oracle calls to $\mathcal{O}_{\text{sign}}$, if H_2 is perfectly simulatable, we have

$$\text{Adv}^{\text{OMUF}} \leq (2(n+1) + q_{\mathcal{O}_{\text{read}}}) + q_{\mathcal{O}_{\text{cmp}}} \times \frac{n+2}{p} + \frac{n}{p} + \frac{q_{\text{Hash}}}{p^6}$$

where $q_{\mathcal{O}_{\text{read}}}$, $q_{\mathcal{O}_{\text{cmp}}}$, and q_{Hash} are the number of queries to $\mathcal{O}_{\text{read}}$, \mathcal{O}_{cmp} , and to the random oracle.

One attack strategy to match this bound would be, for the adversary, to do $n/|\text{Var}|$ $\mathcal{O}_{\text{sign}}$ queries with T_i result of previous queries, in order to make a large-degree polynomial, then to make variations of this large degree polynomial by adding small-degree ones, with available elements. This would generate many large-degree polynomials. Doing $q_{\mathcal{O}_{\text{cmp}}} \sim |\text{Var}| \times p/n$ comparisons would eventually yield a non-trivial large degree polynomial with the secrets as root. Iterating this attack $|\text{Var}|$ times and solving the equations would yield the secrets and allow to make forgeries. Hence, the result is tight.

6.2 Unlinkability

The role of Π_2 is important as it makes sure that the issuer must use either $b = 0$ or $b = 1$ and therefore hides only one bit. Without Π_2 , the issuer could use more than one bit with b and use this as a marker to link tokens to redeem to clients requesting a token. So, the client must verify the π proof for unlinkability. The client must also verify $U \neq 0$, because $U = 0$ could be used by the issuer to mark a token.

The issuer knows which bit is hidden during an issuing session and can extract the hidden bit during redeem. Hence, we should only consider unlinkability when the bits are the same.

Theorem 2. *ATHM is 2-UNLINK-secure. More precisely, given an UNLINK-adversary \mathcal{A} making oracle calls to $\mathcal{O}_{\text{query}}$ with index set $\mathcal{Q}_{\text{query}}$, there exist a DLOG-adversaries \mathcal{B} such that*

$$\text{Adv}_{\mathcal{A}}^{\text{UNLINK}} \leq \frac{2}{n} + \frac{3}{p} + \frac{6}{p} \times \#\mathcal{Q}_{\text{query}} + \text{Adv}_{\mathcal{B}}^{\text{DLOG}}$$

The proof uses the fact that Π_2 is sound. The dominant part of the bound is $2/n$, which is tight.

Proof. We start with an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ playing the UNLINK $_n$ game.

We assume that `AT.ClientQuery` verifies `pp`, at least at the first time it is run. Using the extractor of the proof of x, r_x, y, r_y, z in `pp2` (having an advantage overhead of $\frac{1}{p} + \text{Adv}^{\text{DLOG}}$ for the discrete logarithm and of $\frac{2}{p}$ for the rest) and the extractor of (b_i, d_i) for all π_i which are returned by an $\mathcal{O}_{\text{query}}$ call (having an extra advantage overhead of $\frac{6}{p}$ for each i , the discrete logarithm being already covered), we reduce to a game Γ in which this extraction succeeds. We have

$$\text{Adv}_{\mathcal{A}}^{\text{UNLINK}} \leq \text{Adv}_{\mathcal{A}}^{\Gamma} + \frac{3}{p} + \frac{6}{p} \times \#\mathcal{Q}_{\text{query}} + \text{Adv}_{\mathcal{B}}^{\text{DLOG}}$$

For every i , the $(\text{query}_i, \text{resp}_i)$ pair, parsed as $\text{query}_i = T_i$ and $\text{resp}_i = (U_i, V_i, t_{i,S}, \pi_i)$, we extract (b_i, d_i) such that $b_i \in \{0, 1\}$, $U_i = d_i G$, and $V_i =$

$d_i(x + b_i y + t_{i,S} z)G + dT_i$. Note that (b_i, d_i) is unique when it exists. We deduce that $Q_i = (x + b_i y + t_i z)P_i$ with $b_i \in \{0, 1\}$ in the game Γ .

The rest of the proof is an information theoretic argument for which complexities do not matter. Given $(x, y, z, T_i, U_i, V_i, t_{i,S}, \pi_i)$ we can uniquely determine b_i . We observe that $(t_i, P_i)|(x, y, z, T_i, U_i, V_i, t_{i,S}, \pi_i)$ is uniformly distributed as a pair composed of a scalar and a nonzero group element. Hence, whenever \mathcal{A}_2 returns \mathcal{Q} and the list of resp_i , it determines the values of the b_i but the (t_i, P_i) to be released are still uniform. After permutation, $(t_{\sigma(i)}, P_{\sigma(i)}, Q_{\sigma(i)})$ has a value of $Q_{\sigma(i)}$ which is imposed by $Q_{\sigma(i)} = (x + b_{\sigma(i)} y + t_{\sigma(i)} z)P_{\sigma(i)}$ so brings $b_{\sigma(i)}$ as only information.

This reduces to the following game: the adversary chooses a list of bits $(b_i)_{i \in \mathcal{Q}}$ with $\#\mathcal{Q} \geq n$, the game selects a random i^* and a random permutation σ then provides b_{i^*} and $(b_{\sigma(i)})_{i \in \mathcal{Q}}$ to the adversary, and the adversary finally makes a guess i and win if and only if $i = i^*$. If the adversary puts n_0 zeros and n_1 ones, the adversary can only win with probability $\frac{1}{n_0}$ when it is a zero (which happens with probability $\frac{n_0}{n_0+n_1}$), and with probability $\frac{1}{n_1}$ when it is a one (which happens with probability $\frac{n_1}{n_0+n_1}$). Overall, the adversary wins with probability $\frac{2}{n_0+n_1}$ which is at most $\frac{2}{n}$ since $n_0 + n_1 = \#\mathcal{Q} \geq n$. \square

6.3 PMB Security

We now consider an adversary \mathcal{A} playing the PMB_{b^*} game.

The only new element in the generic group model treatment is that there is a new variable \bar{b}^* appearing in polynomials. This variable appears as soon as \mathcal{A} queries $\mathcal{O}_{\text{chal}}$. It appears as a term of form $\text{Pol}_{V_{i^*}} = (\bar{x} + \bar{b}^* \bar{y} + t_{S,i^*} \bar{z} + \text{Pol}_{T_{i^*}}) \bar{d}_{i^*}$, where i^* is the index number of the \mathcal{O} query for the $\mathcal{O}_{\text{chal}}$ query.

We treat the variable \bar{b}^* differently than others from Var (as \bar{b}^* takes random values in $\{0, 1\}$ instead of random values in \mathbb{Z}_p like other variables.) For each polynomial $\text{Pol}_A(\bar{b}^*, \text{Var})$ in which \bar{b}^* appears, we can make two partial evaluations $\text{Pol}_A(0, \text{Var})$ and $\text{Pol}_A(1, \text{Var})$ corresponding to $\bar{b}^* = 0$ and $\bar{b}^* = 1$. We obtain two polynomials with no \bar{b}^* variable. Hence, this increase the number of polynomials by a factor at most 2.

We first proceed like for OMUF security with games Γ_1 to have no π in the return from the issuer. The transition to Γ_2 to get rid of the \mathcal{O}_{cmp} oracle (and thus of the $\mathcal{O}_{\text{read}}$ and $\mathcal{O}_{\text{valid}}$ oracles) is also using hybrid arguments but is more complicated. Oracles are simulated in the order they are called. For the simulation of \mathcal{O}_{cmp} and $\mathcal{O}_{\text{read}}$ until the challenge is made, it works like for OMUF security. After the challenge is made, the $\mathcal{O}_{\text{cmp}}(A, B)$ made by the adversary are simulated by answering 1 if and only if there exists $\beta \in \{0, 1\}$ such that $\text{Pol}_A - \text{Pol}_B$ vanishes in the partial evaluation $\bar{b}^* = \beta$. To simulate $\mathcal{O}_{\text{valid}}(t, P, Q)$, the answer is 1 if and only if there exists $b, \beta \in \{0, 1\}$ such that $\text{Pol}_Q - (\bar{x} + b\bar{y} + t\bar{z}) \times \text{Pol}_P$ vanishes in the partial evaluation $\bar{b}^* = \beta$. In the first case, we prove the following variant of Fact 2.

Fact 6 *Let q is the number of $\mathcal{O}_{\text{sign}}$ queries (including $\mathcal{O}_{\text{chal}}$) before an input (A, B) is presented for the first time to the \mathcal{O}_{cmp} oracle. We assume that the*

call is made by the adversary. Except with a probability bounded by $\frac{q+2}{p}$, we have $A = B$ if and only if $\text{Pol}_A(0, \text{Var}) = \text{Pol}_B(0, \text{Var})$ or $\text{Pol}_A(1, \text{Var}) = \text{Pol}_B(1, \text{Var})$.

The direct implication works like in Fact 2 as the match happens for $\bar{b}^* = b^*$. We now want to show the converse implication: if $\text{Pol}_A - \text{Pol}_B$ vanishes for $\bar{b}^* = \beta$ with $\beta \in \{0, 1\}$, then we have $A = B$.

Thanks to the generic group model, the adversary knows a linear combination of provided group elements to obtain $A - B$. Let λ_i and μ_i be the coefficients of U_i and V_i respectively. We let i be the largest index such that $(\lambda_i, \mu_i) \neq (0, 0)$. The partial derivative of Pol_{A-B} with respect to \bar{d}_i is $\lambda_i + \mu_i(\bar{x} + b_i\bar{y} + t_{S,i}\bar{z} + \text{Pol}_{T_i})$ (with b_i replaced by \bar{b}^* in the $i = i^*$ case). Since $\text{Pol}_{A-B}(\beta, \text{Var}) = 0$, the partial derivative vanishes for $\bar{b}^* = \beta$ too. The adversary knows how to express T_i as a linear combination of provided group elements. We notice that no group element has any monomial $\bar{b}^*\bar{x}$. No group element has the monomial \bar{x} , so there is no way to make a T_i such that $\text{Pol}_{T_i}(\beta, \text{Var})$ has a monomial \bar{x} . Hence, it cannot cancel \bar{x} . We deduce that $A - B$ has no U_i and V_i as a component in the linear combination. Hence, it does not depend on b^* and the result is trivial: if $\text{Pol}_A - \text{Pol}_B$ vanishes for $\bar{b}^* = \beta$, it vanishes with the partial evaluation $\bar{b}^* = b^*$ too. Hence, we can apply Fact 2 and conclude.

To simulate $\mathcal{O}_{\text{valid}}$, we use the following fact.

Fact 7 *Let q is the number of $\mathcal{O}_{\text{sign}}$ queries (including $\mathcal{O}_{\text{chal}}$) before a $\mathcal{O}_{\text{valid}}(t, P, Q)$ call is made for the first time. We assume there is no \mathcal{O}_{cmp} call before. Except with a probability bounded by $\frac{q+2}{p}$, the oracle returns 1 if and only if there exists $b, \beta \in \{0, 1\}$ such that $\text{Pol}_Q(\beta, \text{Var}) = (\bar{x} + b\bar{y} + t\bar{z})\text{Pol}_P(\beta, \text{Var})$.*

The direct implication uses Fact 2 with $\beta = b^*$ and the right b which makes (b, t, P, Q) valid. For the converse implication, we assume that the polynomial equality is verified for a given (b, β) pair of bits.

We let $a_{U_i}, a_{V_i}, b_{U_i}, b_{V_i}$ be the coefficients of U_i and V_i for P_i and of U_i and V_i for Q_i . Let $\lambda_i = b_{U_i} - (\bar{x} + b\bar{y} + t\bar{z})a_{U_i}$ and $\mu_i = b_{V_i} - (\bar{x} + b\bar{y} + t\bar{z})a_{V_i}$. Like in the previous case, let i be the largest index such that (λ_i, μ_i) is nonzero. The partial derivative of $\text{Pol}_Q - (\bar{x} + b\bar{y} + t\bar{z})\text{Pol}_P$ in terms of \bar{d}_i is again $\lambda_i + \mu_i(\bar{x} + b_i\bar{y} + t_{S,i}\bar{z} + \text{Pol}_{T_i})$ (with b_i replaced by \bar{b}^* in the $i = i^*$ case). We know it vanishes when evaluated on $\bar{b}^* = \beta$. Clearly, a_{V_i} must be zero (otherwise, \bar{x}^2 appears and cannot vanish). The same argument about the monomial \bar{x} and also about the monomial \bar{y} implies that $(\bar{x} + b\bar{y})a_{U_i} = (\bar{x} + b_i\bar{y})b_{V_i}$. Hence,

$$\frac{\partial \text{Pol}_{A-B}}{\partial \bar{d}_i}(\bar{b}^*, \text{Var}) = b_{U_i} + b_{V_i} ((b_i - b)\bar{y} + (t_{S,i} - t)\bar{z} + \text{Pol}_{T_i}(\bar{b}^*, \text{Var}))$$

The adversary also know a linear combination of T_i in terms of the provided group elements. Let λ'_j and μ'_j be the coefficients of U_j and V_j and let j be the largest index for which they are nonzero. The second partial derivative gives

$$\frac{\partial^2 \text{Pol}_{A-B}}{\partial \bar{d}_j \partial \bar{d}_i}(\bar{b}^*, \text{Var}) = b_{U_i} (\lambda'_j + \mu'_j(\bar{x} + b_j\bar{y} + t_j\bar{z} + \text{Pol}_{T_j}(\bar{b}^*, \text{Var})))$$

which should vanish for $\bar{b}^* = \beta$. The same argument than before shows that \bar{x} cannot disappear. Hence, T_i cannot have components in U_j or V_j and does not depend on b^* .

Knowing that T_i has no component in any previous U_i of V_i , we can go to the second largest i such that (λ_i, μ_i) is nonzero and look at the partial derivative with respect to \bar{d}_i (which we know does not appear in a subsequent V_i). We obtain the same result that T_i does not have any component in U_j or V_j . We analyze like this all components in every (U_i, V_i) of $\text{Pol}_Q - (\bar{x} + b\bar{y} + t\bar{z})\text{Pol}_P$. Since it vanished on $\bar{b}^* = \beta$, we cannot have any other component. Hence

$$\text{Pol}_Q - (\bar{x} + b\bar{y} + t\bar{z})\text{Pol}_P = \sum_i (b_{U_i} + b_{V_i} ((b_i - b)\bar{y} + (t_{S,i} - t)\bar{z} + \text{Pol}_{T_i}(\bar{b}^*, \text{Var}))) \bar{d}_i$$

with b_{i^*} to be replaced by \bar{b}^* . For all indices i of nonzero terms, we deduce that $b = b_i$ and T_i is of form $T_i = t_{C,i}Z + r_iG$. This boils down to

$$\text{Pol}_Q - (\bar{x} + b\bar{y} + t\bar{z})\text{Pol}_P = b_{V_{i^*}}(\bar{b}^* - b)\bar{y}\bar{d}_{i^*} + \sum_{i \neq i^*} b_{V_i}(b_i - b)\bar{y}\bar{d}_i$$

If $b_{V_{i^*}} = 0$, this vanishing for $\bar{b}^* = \beta$ implies vanishing for $\bar{b}^* = b^*$ too, so we can apply Fact 2 to deduce that $Q = (x + by + tz)P$ most of the cases. If $b_{V_{i^*}} \neq 0$, this vanishing for $\bar{b}^* = \beta$ implies $b = \beta$. We can then see that it vanishes for $\bar{b}^* = b^*$ and $b = b^*$ too. We apply Fact 2 to deduce that $Q = (x + \beta y + tz)P$ most of the cases.

Using hybrids, we obtain

$$|\Pr[\text{PMB}_{b^*} \rightarrow 1] - \Pr[\Gamma_{2,b^*} \rightarrow 1]| \leq (2(n + 1 + q_{\mathcal{O}_{\text{read}}}) + q_{\mathcal{O}_{\text{cmp}}}) \times \frac{n + 2}{p} + \frac{q_{\text{Hash}}}{p^6}$$

In the game Γ_2 , the oracles $\mathcal{O}_{\text{verify}}$, $\mathcal{O}_{\text{read}}$, and \mathcal{O}_{cmp} are not used any more.

After getting rid of $\mathcal{O}_{\text{verify}}$ and $\mathcal{O}_{\text{read}}$, we obtain a game in which no information about b^* is given to the adversary. Thus,

$$\Pr[\Gamma_{2,0} \rightarrow 1] = \Pr[\Gamma_{2,1} \rightarrow 1]$$

Theorem 3. *For every \mathcal{A} playing PMB in the generic group model and making n oracle calls to $\mathcal{O}_{\text{sign}}$, if Π_2 is perfectly simulatable, we have*

$$\text{Adv}^{\text{PMB}} \leq 2(2(n + 1 + q_{\mathcal{O}_{\text{read}}}) + q_{\mathcal{O}_{\text{cmp}}}) \times \frac{n + 2}{p} + \frac{q_{\text{Hash}}}{p^6}$$

where $q_{\mathcal{O}_{\text{read}}}$, $q_{\mathcal{O}_{\text{cmp}}}$, and q_{Hash} are the number of queries to $\mathcal{O}_{\text{read}}$, \mathcal{O}_{cmp} , and to the random oracle.

Again, the bound is tight.

7 Conclusion

In our work, we studied the anonymous tokens with hidden metadata bit from issuer to the verifier. We started our studies with a security weakness in a protocol from CRYPTO 2020 which is an extension of Privacy Pass from PoPETs 2018. The protocol is based on oblivious PRFs. We discussed the real-world implications of the security notions defined in their protocol and showed that such problems can be overcome with a new token protocol from algebraic MACs. We defined the security with more realistic threat model. The security of our protocol relies on an extractable proof system which can be initialized in three ways: (1) verifiable encryption relying on standard notions (2) knowledge-of-exponents where its proof relies on generic group model (GGM) (3) with no extractable proof where the security left as an open problem.

We believe algebraic MACs suit better in anonymous tokens with hidden bit as a primitive. However, it is an open question to understand if there are other OPRFs or another new primitive that would help designing anonymous tokens with strong security and privacy guarantees.

Acknowledgements. We heartily thank Greg Zaverucha, Micheler Orrù, and Nirvan Tyagi for the insightful discussions and fruitful comments on the early version of the paper. We are also very thankful to Kim Laine, Radames Cruz Moreno, and Wei Dai for their valuable time and help with the implementation of the protocols and benchmarks in Rust.

References

1. Trust Tokens API. <https://developer.chrome.com/docs/privacy-sandbox/trust-tokens/>.
2. Foteini Baldimtsi and Anna Lysyanskaya. Anonymous Credentials Light. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13*, pages 1087–1098. ACM, 2013.
3. Jan Camenisch and Anna Lysyanskaya. An Efficient System for Non-Transferable Anonymous Credentials with Optional Anonymity Revocation. In *EUROCRYPT*. Springer International Publishing, 2001.
4. Jan Camenisch and Anna Lysyanskaya. A Signature Scheme with Efficient Protocols. In Stelvio Cimato, Giuseppe Persiano, and Clemente Galdi, editors, *Security in Communication Networks*, pages 268–289. Springer Berlin Heidelberg, 2003.
5. Pyrros Chaidos and Jens Groth. Making sigma-protocols non-interactive without random oracles. In *Public-Key Cryptography – PKC*, pages 650–670. Springer International Publishing, 2015.
6. Melissa Chase, Sarah Meiklejohn, and Greg Zaverucha. Algebraic MACs and Keyed-Verification Anonymous Credentials. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, CCS '14*, page 1205–1216. Association for Computing Machinery, 2014.
7. Melissa Chase, Sarah Meiklejohn, and Greg Zaverucha. Algebraic MACs and Keyed-Verification Anonymous Credentials. <https://eprint.iacr.org/2013/516>, 2014.

8. Melissa Chase, Trevor Perrin, and Greg Zaverucha. *The Signal Private Group System and Anonymous Credentials Supporting Efficient Verifiable Encryption*, page 1445–1459. Association for Computing Machinery, 2020.
9. Melissa Chase, Trevor Perrin, and Greg Zaverucha. The Signal Private Group System and Anonymous Credentials Supporting Efficient Verifiable Encryption. <https://eprint.iacr.org/2019/1416>, 2020.
10. David Chaum. Blind Signatures for Untraceable Payments. In *Advances in Cryptology – CRYPTO*. Springer International Publishing, 1982.
11. Ivan Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In Joan Feigenbaum, editor, *Advances in Cryptology — CRYPTO ’91*, pages 445–456, Berlin, Heidelberg, 1992. Springer Berlin Heidelberg.
12. Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, and Filippo Val-sorda. Privacy Pass: Bypassing Internet Challenges Anonymously. In *PoPETs*, pages 164–180, 2018.
13. Ben Kreuter, Tancrede Lepoint, Michele Orrù, and Mariana Raykova. Anonymous Tokens with Private Metadata Bit. <https://eprint.iacr.org/2020/072>.
14. Ben Kreuter, Tancrede Lepoint, Michele Orrù, and Mariana Raykova. Anonymous Tokens with Private Metadata Bit. In *Advances in Cryptology – CRYPTO*, pages 308–336. Springer International Publishing, 2020.
15. Christian Paquin and Greg Zaverucha. U-Prove Cryptographic Specification v1.1 (Revision 3), December 2013. Released under the Open Specification Promise (<http://www.microsoft.com/openspecifications/en/us/programs/osp/default.aspx>).
16. Tjerand Silde and Martin Strand. Anonymous Tokens with Public Metadata and Applications to Private Contact Tracing. <https://fc22.ifca.ai/preproceedings/40.pdf>.
17. Akira Takahashi and Greg Zaverucha. Verifiable encryption from MPC-in-the-head. <https://eprint.iacr.org/2021/1704>, 2021.
18. Stefano Tessaro and Chenzhi Zhu. Short Pairing-Free Blind Signatures with Exponential Security. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology - EUROCRYPT 2022*, volume 13276 of *Lecture Notes in Computer Science*, pages 782–811. Springer, 2022.

Additional Material

A ATHM: Anonymous Token with Hidden Metadata

The full version of ATHM includes several options. First of all, we add the public metadata attribute m . It can, for instance, include a coarse expiration data. It should be used carefully as it degrades unlinkability. A token with a hidden bit b will be a tuple (t, m, σ) such that σ is a valid MAC for the attributes (b, m, t) .

The objective of other options is to get rid of the dependency to the generic group model and possibly have provable security in the standard model. For that, one option is the choice of the MAC. A valid MAC for an input b , a public metadata m , a nonce t , and a secret key $(\vec{x}, \vec{y}, \vec{y}', \vec{z})$ is a pair $\sigma = (P, \vec{Q})$ such that $\vec{Q} = (\vec{x} + b\vec{y} + m\vec{y}' + t\vec{z})P$. As detailed below, several options can be made for the MAC algorithm in which case the secrets \vec{x} , \vec{y} , \vec{y}' , and \vec{z} will be vectors of various dimension and structure with elements in \mathbb{Z}_p .¹¹

Finally, we also have options for the choice of proof schemes Π_1 and Π_2 . The proof Π_1 is an extractable proof made by the client and Π_2 is the simulatable proof made by the issuer.

Options are summarized in Table 1.

Table 1. Options for the ATHM Protocol. Non-standard models for provable security are indicated by either an ad-hoc assumption, or the random-oracle model (ROM), or the generic group model (GGM). In the Attributes menu, several options can be selected.

Attributes	MAC	$\Pi_1 (T_{\text{ext}})$	$\Pi_2 (\pi)$
<ul style="list-style-type: none"> – Issuer private metadata bit b – Public metadata m – Nonce t 	<ul style="list-style-type: none"> – MAC-GGM – MAC-DDH – MAC3 	<ul style="list-style-type: none"> – no T_{ext} (GGM) – Knowledge-of-Exponent (Assumption 3 or GGM) – MPC-in-the-Head (ROM) 	<ul style="list-style-type: none"> – Fiat-Shamir (ROM) – Homomorphic encryption

A.1 The ATHM Components

Our scheme uses as a building block an extractable non-interactive proof Π_1 and a simulatable non-interactive proof Π_2 . Several options for Π_1 and Π_2 exist but we focus on one in this section. They will be discussed in separate sections.

¹¹ If we multiply a vector with elements in \mathbb{Z}_p by a group element, we obtain a vector with elements in the group.

Setup algorithm. Setup is composed of four phases. The Setup_1 algorithm generates an (additive) group, which is cyclic, of prime order p . The Setup_2 algorithm selects a group generator G . The crs string includes p , group parameters gp to be able to do operations in the group, and the generator G . It also includes parameters for a proof Π_1 which are generated by Setup_3 together with a trapdoor td_1 for extraction and parameters for Π_2 which are generated by Setup_4 together with a trapdoor td_2 .

$\text{Setup}(1^\lambda)$:

- 1: $\text{Setup}_1(1^\lambda) \rightarrow (\text{gp}, p)$ ▷ group setup
- 2: $\text{Setup}_2(\text{gp}, p) \rightarrow G$ ▷ generator setup
- 3: $\text{Setup}_3(\text{gp}, p, G) \rightarrow (\text{crs}_1, \text{td}_1)$ ▷ Π_1 setup
- 4: $\text{Setup}_4(\text{gp}, p, G) \rightarrow (\text{crs}_2, \text{td}_2)$ ▷ Π_2 setup
- 5: $\text{crs} \leftarrow (\text{gp}, p, G, \text{crs}_1, \text{crs}_2)$
- 6: $\text{td} \leftarrow (\text{td}_1, \text{td}_2)$

The KeyGen algorithm. Key generation is composed of several phases.

$\text{KeyGen}(\text{crs})$:

- 1: $\text{KeyGen}_0(\text{crs}) \rightarrow (\text{pp}_0, \text{sk}_0)$
- 2: $\text{KeyGen}_1(\text{crs}, \text{pp}_0) \rightarrow (\text{pp}_1, \text{sk}_1)$ ▷ Π_1 key generation
- 3: $\text{KeyGen}_2(\text{crs}, \text{pp}_0, \text{sk}_0) \rightarrow (\text{pp}_2, \text{sk}_2)$ ▷ Π_2 key generation
- 4: $\text{pp} \leftarrow (\text{pp}_0, \text{pp}_1, \text{pp}_2)$
- 5: $\text{sk} \leftarrow (\text{sk}_0, \text{sk}_1, \text{sk}_2)$

In KeyGen_0 , the issuer selects four secrets $\text{sk}_0 = (\vec{x}, \vec{y}, \vec{y}', \vec{z})$ with $\vec{y}, \vec{y}', \vec{z} \neq 0$, and sets $\text{pp}_0 = \vec{Z} = \vec{z}G$. The structure of those secrets depends on the choice of the MAC algorithm. We denote by $\mathcal{E}_x, \mathcal{E}_y, \mathcal{E}_{y'}$, and \mathcal{E}_z the respective domains of $\vec{x}, \vec{y}, \vec{y}'$, and z , depending on the choice of the MAC algorithm. We will also need the span $\bar{\mathcal{E}}_z$ of \mathcal{E}_z .

The token issuance protocol. The user has public parameters. The server's input is the secret $(\vec{x}, \vec{y}, \vec{y}', \vec{z})$, a bit b to hide inside the token, and an agreed public metadata m to embed. The protocol works as depicted on Figure 7: the client selects a random tag share t_C and a random mask $\vec{r} \in \bar{\mathcal{E}}_z$ and sends $\vec{T} = t_C \vec{Z} + \vec{r}G$ to the issuer. The client must also send an extractable proof T_{ext} for \vec{T} . The issuer selects a random tag share t_S and generates a pair (U, \vec{V}) such that $(U, \vec{V} - \vec{r}U)$ is a valid MAC for tag $t = t_C + t_S$ and metadata (b, m) with key $(\vec{x}, \vec{y}, \vec{y}', \vec{z})$. For this, the issuer selects $U = dG$ for a random $d \in \mathbb{Z}_p^*$ and $\vec{V} = d(\vec{x}G + b\vec{y}G + m\vec{y}'G + t_S\vec{z}G + \vec{T})$. A NIZK proof Π_2 must prove that the (U, \vec{V}, t_S) triplet was correctly generated. Note that Π_2 must prove that b is a bit and that m was embedded. Both participants are assumed to have agreed on m (alternately, one participant decides and sends m to the other). Then, (U, \vec{V}, t_S, π) is returned. The client computes $(U, \vec{V} - \vec{r}U)$. To make it unlinkable, the pair is multiplied by a random mask c to obtain another pair $\sigma = (P, \vec{Q})$.

<p>ATHM.Client(G, \vec{Z}, m) input: (G, \vec{Z}) $t_C \leftarrow \mathbb{Z}_p, \quad \vec{r} \leftarrow \mathcal{E}_z$ $\vec{T} := t_C \vec{Z} + \vec{r} G$ $\left[\vec{T}_{\text{ext}} := \Pi_1.\text{Prove}(t_C, \vec{r}; \text{crs}, \text{pp}, \vec{T}) \right]$</p>	<p>ATHM.IssueToken($(\vec{x}, \vec{y}, \vec{y}', \vec{z}), b, m$) input: $(\vec{x}, \vec{y}, \vec{y}', \vec{z}), b, m$</p>
$\xrightarrow{\vec{T} \left[\vec{T}_{\text{ext}} \right]}$	
<p>[if not $\Pi_1.\text{Verify}(\vec{T}, \vec{T}_{\text{ext}}, \text{crs}, \text{pp}, \text{sk}_1)$ then return \perp] $t_S \leftarrow \mathbb{Z}_p, \quad d \leftarrow \mathbb{Z}_p^*$ $U := dG$ $\vec{V} := d(\vec{x}G + b\vec{y}G + m\vec{y}'G + t_S \vec{z}G + \vec{T})$ $\pi \leftarrow \Pi_2.\text{Prove}(b, d; U, \vec{V}, t_S, m, \vec{T}, \text{crs}, \text{pp}, \text{sk}_0, \text{sk}_2)$ $\leftarrow (U, \vec{V}, t_S, \pi)$</p>	
<p>if not $\Pi_2.\text{Verify}(\pi, U, \vec{V}, t_S, m, \vec{T}, \text{crs}, \text{pp})$ then return \perp if $U = 0$ then return \perp $c \leftarrow \mathbb{Z}_p^*$ $P := cU$ $\vec{Q} := c(\vec{V} - \vec{r}U)$ $t := t_C + t_S$ $\sigma := (P, \vec{Q})$ output: (t, σ)</p>	

Fig. 7. ATHM token issuance protocol.

About the extractable proof \vec{T}_{ext} following Π_1 , there are several options which are discussed in a subsequent section. The role of \vec{T}_{ext} is to allow to formally prove the security of ATHM without relying on the generic group model.

The NIZK proof Π_2 is a Fiat-Shamir transform of an OR proof of two Schnorr proofs for $b = 0$ and $b = 1$. Variants without using a random oracle are possible too.

The ReadBit algorithm. The redemption of (m, t, P, \vec{Q}) with $(\vec{x}, \vec{y}, \vec{y}', \vec{z})$ checks for which $b \in \{0, 1\}$, the equality $\vec{Q} = (\vec{x} + b\vec{y} + m\vec{y}' + t\vec{z})P$ is satisfied. Assuming $0 \notin \mathcal{E}_y$ ensures that b is unique. If b exists, it is returned. Otherwise, \perp is returned.

Note that for this protocol, tag t needs to be a nonce as in other protocols [13], [14] i.e. the redeemer should check against double-spending of a token with the same t . Otherwise, it is easy to transform a valid token with attributes (b, m, t) into another valid token with the same attributes: let $\sigma = (P, \vec{Q})$. Then the client can forge another token $\sigma' = (P', \vec{Q}')$ by taking $\sigma' = c' \cdot \sigma$ for any $c' \in \mathbb{Z}_p^*$.

A.2 The MAC Building Block

Our security results will be based on the security of an algebraic MAC. The simplest one is the MACGGM algorithm [6] defined as follows: given a secret $(x, y, y', z) \in \mathbb{Z}_p^4$, a valid authentication for $(b, m, t) \in \mathbb{Z}_p^3$ is a pair $\sigma = (P, Q)$ such that $Q = (x + by + my' + tz)P$. For this MAC to be secure, it is important that no adversary can find any linear relation between the random values of P . Hence, P is selected at random by the issuer.

In MACDDH [6], we consider \vec{x} , \vec{y} , \vec{y}' , and \vec{z} as vectors of a particular form: $\vec{x} \in \mathcal{E}_x = \mathbb{Z}_p^3$, $\vec{z} = \{(\alpha, \beta, 0); \alpha, \beta \in \mathbb{Z}_p\}$, and $\vec{y}, \vec{y}', \vec{z} \in \mathcal{E}_y = \mathcal{E}_{y'} = \mathcal{E}_z = \bar{\mathcal{E}}_z - \{(0, 0, 0)\}$.

In general, we let $\vec{x} \in \mathcal{E}_x$, $\vec{y} \in \mathcal{E}_y$, $\vec{y}' \in \mathcal{E}_{y'}$, $\vec{z} \in \mathcal{E}_z$. A MAC $\sigma = (P, \vec{Q})$ for attributes (b, m, t) is valid if it satisfies $\vec{Q} = (\vec{x} + b\vec{y} + m\vec{y}' + t\vec{z})P$.

We will assume that no adversary can forge a valid σ for a new (b, t, m) tuple when given access to a MAC oracle and also that no adversary can distinguish $b \in \{b_0, b_1\}$ from a random valid tuple (t, m, σ) for this unknown b . The two respective games are depicted on Figure 8 and Figure 9.

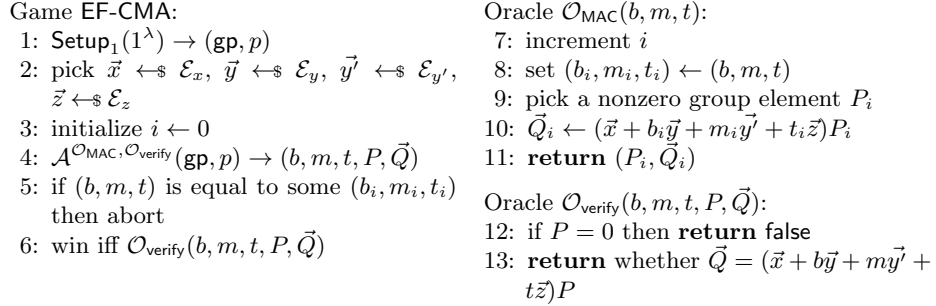


Fig. 8. Forgeability Game for the MAC

We will make the two following assumptions. For MACGGM, they are proven in the generic group model [6]. For MACDDH, they are proven under the DDH assumption [6].

Assumption 1 *In the EF-CMA game on Figure 8, we define the advantage of an adversary \mathcal{A} by*

$$\text{Adv}_{\mathcal{A}}^{\text{EF-CMA}}(\lambda) = \Pr[\text{win}]$$

We assume that for any PPT adversary \mathcal{A} , the advantage is a negligible function.

Assumption 2 *In the IND-CMA game on Figure 9, we define the advantage of an adversary \mathcal{A} by*

$$\text{Adv}_{\mathcal{A}}^{\text{IND-CMA}}(\lambda) = \Pr[\text{IND-CMA}_1 \rightarrow 1] - \Pr[\text{IND-CMA}_0 \rightarrow 1]$$

We assume that for any PPT adversary \mathcal{A} , the advantage is a negligible function.

Game IND-CMA_{b^*} :

- 1: $\text{flag} \leftarrow \text{false}$
- 2: initialize $i \leftarrow 0$
- 3: $m^* \leftarrow \perp$
- 4: $\text{Setup}_1(1^\lambda) \rightarrow (\text{gp}, p)$
- 5: pick $\vec{x} \leftarrow_{\$} \mathcal{E}_x, \vec{y} \leftarrow_{\$} \mathcal{E}_y, \vec{y}' \leftarrow_{\$} \mathcal{E}_{y'}, \vec{z} \leftarrow_{\$} \mathcal{E}_z$
- 6: **return** $\mathcal{A}^{\mathcal{O}_{\text{MAC}}, \mathcal{O}_{\text{chal}}, \mathcal{O}_{\text{verify}}, \mathcal{O}_{\text{valid}}}(\text{gp}, p)$

Oracle $\mathcal{O}_{\text{MAC}}(b, m, t)$:

- 7: if **flag** and $(b, m, t) \in \{(b_0^*, m^*, t^*), (b_1^*, m^*, t^*)\}$ then **return** \perp
- 8: increment i
- 9: $(b_i, m_i, t_i) \leftarrow (b, m, t)$
- 10: pick a nonzero group element P
- 11: $\vec{Q} \leftarrow (\vec{x} + b\vec{y} + m\vec{y}' + t\vec{z})P$
- 12: **return** (P, \vec{Q})

Oracle $\mathcal{O}_{\text{chal}}(b, b', m, t)$:

- 13: if **flag** or $\exists i (b_i, m_i, t_i) \in \{(b, m, t), (b', m, t)\}$ then **return** \perp
- 14: $(b_0^*, b_1^*, m^*, t^*) \leftarrow (b, b', m, t)$
- 15: $\text{flag} \leftarrow \text{true}$
- 16: pick a nonzero group element P^*
- 17: $\vec{Q}^* \leftarrow (\vec{x} + b_0^*\vec{y} + m^*\vec{y}' + t^*\vec{z})P^*$
- 18: **return** (P^*, \vec{Q}^*)

Oracle $\mathcal{O}_{\text{verify}}(b, m, t, P, \vec{Q})$:

- 19: if **flag** and $b \in \{b_0^*, b_1^*\}$ and $(t, m) = (t^*, m^*)$ then **return** \perp
- 20: if $P = 0$ then **return** **false**
- 21: **return** whether $\vec{Q} = (\vec{x} + b\vec{y} + m\vec{y}' + t\vec{z})P$

Oracle $\mathcal{O}_{\text{valid}}(b_0, b_1, m, t, P, \vec{Q})$:

- 22: if $\neg \text{flag}$ then **return** \perp
- 23: if $(t, m) = (t^*, m^*)$ and $\{b_0, b_1\} \neq \{b_0^*, b_1^*\}$ then **return** \perp
- 24: if $P = 0$ then **return** **false**
- 25: if $\vec{Q} = (\vec{x} + b_0\vec{y} + m\vec{y}' + t\vec{z})P$ then **return** **true**
- 26: if $\vec{Q} = (\vec{x} + b_1\vec{y} + m\vec{y}' + t\vec{z})P$ then **return** **true**
- 27: **return** **false**

Fig. 9. Pseudorandomness Games for the MAC

A.3 The Extractable Proof Building Block

The Π_1 extractable proof is used in our analysis to formally prove security. We require a *straightline extraction* of t_C and \vec{r} by using a trapdoor and a verification algorithm which does not require the trapdoor. Namely, a dedicated algorithm should be able to extract t_C and \vec{r} from a valid proof and the view of the client and the issuer should be able to verify the proof. Furthermore, we need the proof to be simulatable in a computationally indistinguishable manner by using a trapdoor.

We propose several options for Π_1 .

- We use a construction from techniques of verifiable encryption, which has a complexity cost but a security proven under standard assumptions. Namely, we should encrypt the values t_C and \vec{r} with a public key, which can be decrypted using a trapdoor only, and in a way that we can publicly verify that it encrypts values satisfying $\vec{T} = t_C \vec{Z} + \vec{r}G$.
- We use a simple Π_1 scheme (i.e. the proof is computed and verified with a very low complexity overhead) but an ad-hoc (non-standard) extractability assumption.
- Use a trivial Π_1 scheme (i.e. with a void T_{ext} and an always true verification) and the generic group model.

Using verifiable encryption. Regarding the first approach, we can use the construction by Takahashi and Zaverucha [17]. Namely, t_C and \vec{r} are shares among several participants who run a multiparty computation protocol to compute $[t_C]\vec{Z} + [\vec{r}]G$, reveal it, and compare it to \vec{T} . The execution is simulated. The views are committed with an extractable commitment (for instance, a public-key encryption scheme with a public key set up by Setup_1 in the common reference string crs_1 , extractable from the secret key in td_1). Based on the commit values, a challenge is computed using a random oracle. This challenge gives the index of a participant whose view remains hidden but other commitments are opened. The \vec{T}_{ext} record consists of the commitments and the openings. Verification consists of checking that the openings match their respective commit values, that the opened views match and make their participant output `true`, and that the missing view is the one corresponding to the hash of the commit values through the random oracle.

To extract from \vec{T}_{ext} , we use the trapdoor which allows to extract all views from the commit values. The views include all shares of t_C and \vec{r} which can thus be reconstructed.

To forge \vec{T}_{ext} , we can program the random oracle to predict the view not to be revealed and make the corresponding participant malicious so that others would output `true`.

The proposed verifiable encryption construction from [17] is secure from any MPC-in-the-head zero-knowledge proof and an IND-CPA cryptosystem which is binding. This can be instantiated with quite standard assumptions.

Using an ad-hoc knowledge-of exponent assumption. We consider a key generation KeyGen_1 which sets an additional secret $\text{sk}_1 = \theta$ for the issuer and reveals $\text{pp}_1 = (G_{\text{ext}}, \vec{Z}_{\text{ext}})$ to the client, with $G_{\text{ext}} = \theta G$ and $\vec{Z}_{\text{ext}} = \theta \vec{Z}$. Then, we set $\vec{T}_{\text{ext}} = t_C \vec{Z}_{\text{ext}} + \vec{r} G_{\text{ext}}$ to be verified by $\vec{T}_{\text{ext}} = \theta \vec{T}$.

Simulating \vec{T}_{ext} is trivial with the help of the trapdoor $\text{sk}_1 = \theta$.

For extraction, we need a knowledge-of-exponent-like assumption [11]: that being able to make a valid $(\vec{T}, \vec{T}_{\text{ext}})$ pair with the knowledge of the client would imply knowing a linear combination $\vec{T} = t_C \vec{Z} + \vec{r} G$. In the standard knowledge-of-exponent assumption, we do not have any oracle whereas in our case, we need them. However, some oracles may break the knowledge-of-exponent assumption as shown below. Thus, the difficulty is to show that our oracles do not break this assumption. The oracle calls may help to provide other $(\vec{T}, \vec{T}_{\text{ext}})$ pairs to the client, without the client knowing the linear combination. Our assumption captures that the oracles somehow do not help for that. We formalize two types of oracles.

An oracle is of Type I if there exists vectors of low-degree polynomials $\vec{f}_1, \dots, \vec{f}_n$ and scalars $\lambda_1, \dots, \lambda_n$, for some integer n , such that the oracle is defined as follows:

Oracle $O(\vec{T}, \vec{T}_{\text{ext}})$:

- 1: if $\vec{T}_{\text{ext}} \neq \theta \vec{T}$ then **return** \perp
- 2: $\vec{A}_i \leftarrow \vec{f}_i(\text{aux})G + \lambda_i \vec{T}$ for $i = 1, \dots, n$
- 3: pick $d \leftarrow_{\$} \mathbb{Z}_p$ at random
- 4: **return** $(d\vec{A}_1, \dots, d\vec{A}_n)$

The values θ and aux as well as the group parameters are used by the oracle. What is crucial is that O is stateless, does not use θ or \vec{T}_{ext} after Step 1, and that it randomizes the output with a fresh d . Once the random coins for t_S are fixed, the issuer producing (U, \vec{V}) (without π) can be seen as a Type I oracle which would represent $\mathcal{O}_{\text{sign}}$.

An oracle is of Type II if there exists a deterministic polynomially bounded algorithm \vec{f} returning a vector of scalars, such that the oracle is defined as follows:

Oracle $O(\text{inp}, P, \vec{Q})$:

- 1: compute $\vec{s} \leftarrow \vec{f}(\text{inp}, \text{aux})$
- 2: **return** whether $\vec{Q} = \vec{s}P$

The value aux and the group parameters are used by the oracle. One crucial property is that the output of the oracle is independent from the group element representation. The $\mathcal{O}_{\text{verify}}$ oracle can be seen as a Type II oracle.

We can now define the following game with an adversary \mathcal{A} and an extractor **Extract**:

- 1: $\text{Setup}_1(1^\lambda) \rightarrow (\text{gp}, p)$ ▷ set up the group parameters
- 2: $\text{Setup}_2(\text{gp}, p) \rightarrow G$ ▷ select a group generator
- 3: pick $\text{aux}_1, \text{aux}_2, \dots \leftarrow_{\$} \mathbb{Z}_p$
- 4: $\vec{Z} \leftarrow G.\vec{L}(\text{aux})$
- 5: pick $\theta \leftarrow_{\$} \mathbb{Z}_p^*$

- 6: $G_{\text{ext}} \leftarrow \theta G, \vec{Z}_{\text{ext}} \leftarrow \theta \vec{Z}$
- 7: run $\mathcal{A}(\text{gp}, p, G, G_{\text{ext}}, \vec{Z}, \vec{Z}_{\text{ext}}) \rightarrow (\vec{T}, \vec{T}_{\text{ext}})$ with access to oracles of Type I and Type II
- 8: set **View** to the view of \mathcal{A} \triangleright including inputs, coins, and oracle answers
- 9: run $\text{Extract}(\text{View}) \rightarrow (t_C, \vec{r})$
- 10: **return** whether $\vec{T}_{\text{ext}} = \theta \vec{T}$ and $t_C \vec{Z} + \vec{r}G \neq \vec{T}$

The game includes a uniform vector **aux** of secrets from \mathbb{Z}_p and a linear function \vec{L} to define \vec{Z} .

Our assumption for the extraction which we will prove shortly is as follows:

Assumption 3 *For every PPT \mathcal{A} , every linear \vec{L} , and every set of oracles of Type I and Type II, there exists a polynomially bounded extractor Extract such that the above game returns 1 with negligible probability.*

Hence, Extract can extract t_C and \vec{r} on the fly from the view of \mathcal{A} whenever $(\vec{T}, \vec{T}_{\text{ext}})$ are valid pairs. With such an assumption, we can extract without rewinding and without the coins of the oracles. Note that the syntax for the extractor $\Pi_1.\text{Extract}$ is a bit different from the CRS model here: instead of extracting with the help of a trapdoor td_1 , we extract from the view **View** of the adversary.

To illustrate the difficulty, we stress that the assumption cannot generalize to *any* oracle. For instance, the oracle which uses more \vec{T}_{ext} by returning a random d multiplied by $(\vec{T}, \vec{T}_{\text{ext}})$ would give to the client a valid pair for which he cannot know the scalar d . The variant of this counterexample which returns a secret x (from **aux**) multiplied by the $(\vec{T}, \vec{T}_{\text{ext}})$ pair follows the same argument (unless x leaks). Consequently, the Boolean oracle defined by $O(i, P, \vec{Q})$ returning the i th bit of $(xP, x\vec{Q})$ with a secret x from **aux** makes the assumption invalid too.

So far, it is unknown how to prove Assumption 3 in the standard model. To convince ourselves that it is a reasonable assumption, we prove it in the generic model. However, our later security proofs will be in the standard model with the extra assumption that Assumption 3 is true.

Proof of Assumption 3 in the generic group model. We prove our assumption in the generic group model. We assume that all group operations that an adversary can make are only addition, subtractions, and comparison of previously obtained group elements. They are externalized with the values of group elements hidden.

In this model, an algorithm who knows the view of the adversary and who runs the adversary algorithm step by step to pay attention to every group operation is able to express any group element issued by the adversary as a linear combination of group elements which are in the view. These elements are $G, G_{\text{ext}}, \vec{Z}, \vec{Z}_{\text{ext}}$, and the outputs from the Type I oracles.

We let **Var** be a set of formal variables which represent the values in **aux**. We represent each group element A produced by the adversary as a formal polynomial $\text{Pol}_A(\theta, \text{Var}, \vec{d}_1, \dots, \vec{d}_q)$ in terms of secrets θ, aux , and each d_i which is selected by the Type I oracles. The polynomial is such that $A = \text{Pol}_A(\theta, \text{aux}, d_1, \dots, d_q)G$. Hence, $\text{Pol}_G = 1, \text{Pol}_{G_{\text{ext}}} = \theta$, for the i th components Z_i and $Z_{\text{ext},i}$ of \vec{Z} and \vec{Z}_{ext} ,

we have $\text{Pol}_{Z_{\text{ext},i}} = \bar{\theta} \times \text{Pol}_{Z_i}$. In the Type I oracles, we assume that the degree of each f_j is bounded by some δ . Each output A_j from the i -th oracle call can be expressed as a polynomial which is multiple of \bar{d}_i . For this, we take the f_j as polynomials in Var (of degree bounded by δ).

By induction, every appearing group element is represented by a polynomial of degree bounded by $\delta + q_I$, where q_I is the number of queries to a Type I oracle. Since aux is uniform, we notice that the probability that two group elements which are represented by two different polynomials will match with a probability bounded by $\frac{\delta + q_I}{p}$, due to the Schwartz-Zippel Lemma. As this is negligible, we assume it does not occur. Hence two appearing group elements are equal if and only if their polynomial representations are the same.

The first crucial observation is that the partial degree in θ of every polynomial representation of any produced group element is bounded by 1. This is because the known group elements are such that no f_j uses θ . Since the polynomial representation of \vec{T}_{ext} has no θ^2 , to have the equation $\vec{T}_{\text{ext}} = \theta \vec{T}$ satisfied, the polynomial representation of \vec{T} cannot have any θ . We thus assume that neither the input \vec{T} to oracle calls nor the final output \vec{T} has any θ .

The second observation is that the outputs of the oracle calls have no θ (since neither \vec{T} nor f_j has any) but are multiple of some d_i . The final \vec{T}_{ext} must be a linear combination of group elements in the view. Since none has any $d_i \theta$ monomial, this means \vec{T} must have no d_i inside. We know it must have no θ as well. Hence, it is a linear combination of G and \vec{Z} only. This linear combination is $\vec{T} = t_C \vec{Z} + \vec{r}G$.

By following step by step the group operations, and also doing the corresponding polynomial operation, we obtain t_C and \vec{r} . To lower the complexity of extraction, we can even do operations modulo the monomials d_i that we know will not appear in the end.

A.4 The Simulatable Proof Building Block

We adapt the Π_2 proof here.

Setup₄. First of all, $\text{Setup}_4(\text{gp}, p, G)$ first selects $\theta' \in \mathbb{Z}_p^*$ and sets $H = \theta'G$, $\text{crs}_2 = xH$, and $\text{td}_2 = \theta'$. The role of H is to do a Pedersen commitment.

KeyGen₂. The algorithm $\text{KeyGen}_2(\text{crs}, \vec{Z}, \text{sk}_0)$ parses $\text{crs} = (\text{gp}, p, G, \cdot, H)$ and $\text{sk}_0 = (\vec{x}, \vec{y}, \vec{y}', \vec{z})$, picks $\vec{r}_x \in \mathcal{E}_x$, $\vec{r}_y \in \mathcal{E}_y$, $\vec{r}_{y'} \in \mathcal{E}_{y'}$, and computes $\vec{C}_x = \vec{x}G + \vec{r}_x H$, $\vec{C}_y = \vec{y}G + \vec{r}_y H$, and $\vec{C}_{y'} = \vec{y}'G + \vec{r}_{y'} H$.

Then, KeyGen_2 computes a proof of knowledge for $(\vec{x}, \vec{y}, \vec{z}, \vec{r}_x, \vec{r}_y)$. For this, it picks random $\vec{\rho}_x, \vec{\rho}_{r_x} \in \mathcal{E}_x$, $\vec{\rho}_y, \vec{\rho}_{r_y} \in \mathcal{E}_y$, $\vec{\rho}_{y'}, \vec{\rho}_{r_{y'}} \in \mathcal{E}_{y'}$, $\vec{\rho}_z \in \mathcal{E}_z$, computes $\vec{I}_x = \vec{\rho}_x G + \vec{\rho}_{r_x} H$, $\vec{I}_y = \vec{\rho}_y G + \vec{\rho}_{r_y} H$, $\vec{I}_{y'} = \vec{\rho}_{y'} G + \vec{\rho}_{r_{y'}} H$, $\vec{I}_z = \vec{\rho}_z G$,

$$\varepsilon = \text{Hash}(G, H, \vec{C}_x, \vec{C}_y, \vec{C}_{y'}, \vec{Z}, \vec{I}_x, \vec{I}_y, \vec{I}_{y'}, \vec{I}_z)$$

$(\vec{a}_x, \vec{a}_{r_x}, \vec{a}_y, \vec{a}_{r_y}, \vec{a}_{y'}, \vec{a}_{r_{y'}}, \vec{a}_z) = (\vec{\rho}_x, \vec{\rho}_{r_x}, \vec{\rho}_y, \vec{\rho}_{r_y}, \vec{\rho}_{y'}, \vec{\rho}_{r_{y'}}, \vec{\rho}_z) + \varepsilon(\vec{x}, \vec{r}_x, \vec{y}, \vec{r}_y, \vec{y}', \vec{r}_{y'}, \vec{z})$, and the proof is $(\varepsilon, \vec{a}_x, \vec{a}_{r_x}, \vec{a}_y, \vec{a}_{r_y}, \vec{a}_{y'}, \vec{a}_{r_{y'}}, \vec{a}_z)$. Finally,

$$\text{pp}_2 = (\vec{C}_x, \vec{C}_y, \vec{C}_{y'}, \varepsilon, \vec{a}_x, \vec{a}_{r_x}, \vec{a}_y, \vec{a}_{r_y}, \vec{a}_{y'}, \vec{a}_{r_{y'}}, \vec{a}_z)$$

and $\text{sk}_2 = (\vec{r}_x, \vec{r}_y, \vec{r}_{y'})$ are the output of KeyGen_2 .

To verify the proof (this must be done at least once for all), the client computes $\vec{I}_x = \vec{a}_x G + \vec{a}_{r_x} H - \varepsilon \vec{C}_x$, $\vec{I}_y = \vec{a}_y G + \vec{a}_{r_y} H - \varepsilon \vec{C}_y$, $\vec{I}_{y'} = \vec{a}_{y'} G + \vec{a}_{r_{y'}} H - \varepsilon \vec{C}_{y'}$, $\vec{I}_z = \vec{a}_z G - \varepsilon \vec{Z}$, and checks $\varepsilon = \text{Hash}(G, H, \vec{C}_x, \vec{C}_y, \vec{C}_{y'}, \vec{Z}, \vec{I}_x, \vec{I}_y, \vec{I}_{y'}, \vec{I}_z)$.

Prove. The algorithm $\text{Prove}(b, d; U, \vec{V}, t_S, m, \vec{T}, \text{crs}, \text{pp}, \text{sk}_0, \text{sk}_2)$ parses different elements and picks $\vec{\mu}, \vec{a}_{1-b}, \vec{r}_\mu \in \bar{\mathcal{E}}_y$, $e_{1-b}, r_d \in \mathbb{Z}_p$, $\vec{r}_\rho, \vec{r}_w \in \mathcal{E}$ at random. Then, it sets $d' = -\frac{1}{d}$, $\vec{\rho} = -(\vec{r}_x + b\vec{r}_y + m\vec{r}_{y'} + \vec{\mu})$, $\vec{w} = \vec{x} + b\vec{y} + m\vec{y}' + t_S \vec{z}$ and computes $\vec{C} = b.\vec{C}_y + H.\vec{\mu}$, $\vec{C}_b = H.\vec{r}_\mu$, $\vec{C}_{1-b} = H.\vec{a}_{1-b} - e_{1-b}(\vec{C} - (1-b)\vec{C}_y)$,

$$\begin{pmatrix} C_d \\ \vec{C}_\rho \\ \vec{C}_w \end{pmatrix} = r_d \begin{pmatrix} U \\ \vec{V} \\ \vec{V} \end{pmatrix} + \begin{pmatrix} 0 \dots 0 \\ H.\mathcal{I} \\ 0.\mathcal{I} \end{pmatrix} \vec{r}_\rho + \begin{pmatrix} 0 \dots 0 \\ 0.\mathcal{I} \\ G.\mathcal{I} \end{pmatrix} \vec{r}_w$$

where \mathcal{I} is the identity matrix,

$$e = \text{Hash}(G, H, \vec{C}_x, \vec{C}_y, \vec{Z}, U, \vec{V}, t_S, m, \vec{C}, \vec{C}_0, \vec{C}_1, C_d, \vec{C}_\rho, \vec{C}_w)$$

$e_b = e - e_{1-b}$, $\vec{a}_b = \vec{r}_\mu + e_b \vec{\mu}$, and $(a_d, \vec{a}_\rho, \vec{a}_w) = (r_d, \vec{r}_\rho, \vec{r}_w) + e(d', \vec{\rho}, \vec{w})$. Finally,

$$\pi = (\vec{C}, e_0, e_1, \vec{a}_0, \vec{a}_1, a_d, \vec{a}_\rho, \vec{a}_w)$$

Verify. The algorithm $\text{Verify}(\pi, U, \vec{V}, t_S, m, \vec{T}, \text{crs}, \text{pp})$ computes $\vec{C}_0 = H.\vec{a}_0 - e_0.\vec{C}$, $\vec{C}_1 = H.\vec{a}_1 - e_1.(\vec{C} - \vec{C}_y)$, $e = e_0 + e_1$,

$$\begin{pmatrix} C_d \\ \vec{C}_\rho \\ \vec{C}_w \end{pmatrix} = a_d \begin{pmatrix} U \\ \vec{V} \\ \vec{V} \end{pmatrix} + \begin{pmatrix} 0 \dots 0 \\ H.\mathcal{I} \\ 0.\mathcal{I} \end{pmatrix} \vec{a}_\rho + \begin{pmatrix} 0 \dots 0 \\ 0.\mathcal{I} \\ G.\mathcal{I} \end{pmatrix} \vec{a}_w + e \begin{pmatrix} G \\ \vec{C}_x + \vec{C} + m\vec{C}_{y'} + t_S.\vec{Z} + \vec{T} \\ \vec{T} \end{pmatrix}$$

then verifies the $e = \text{Hash}(G, H, \vec{C}_x, \vec{C}_y, \vec{Z}, U, \vec{V}, t_S, m, \vec{C}, \vec{C}_0, \vec{C}_1, C_d, \vec{C}_\rho, \vec{C}_w)$.

SimKeyGen. The algorithm $\text{SimKeyGen}(\text{crs}, \vec{Z})$ parses crs to retrieve G and H . It picks a fully random pp_2 , computes $\vec{I}_x, \vec{I}_y, \vec{I}_{y'}, \vec{I}_z$ like in verification, then programs the random oracle to have a consistent Hash . The simulation is perfect.

Simulate. The algorithm $\text{Simulate}(U, \vec{V}, t_S, \vec{T}, \text{crs}, \text{pp}, \text{td}_2)$ parses inputs. Then, it picks a random proof π with same distribution, i.e. $\vec{C} \in G.\bar{\mathcal{E}}_y$, $e_0, e_1, a_d \in \mathbb{Z}_p$, $\vec{a}_0, \vec{a}_1 \in \bar{\mathcal{E}}_y$, $\vec{a}_\rho, \vec{a}_w \in \mathcal{E}$ all uniform and independent. It can then proceed as in $\text{Verify}(\pi, U, \vec{V}, t_S, \vec{T}, \text{crs}, \text{pp})$ to compute $\vec{C}_0, \vec{C}_1, C_d, \vec{C}_\rho, \vec{C}_w$, then the hash. Then, the random oracle is programmed to output the correct hash $e = e_0 + e_1$.m Programming may fail with probability at most $\frac{q_{\text{Hash}}}{p^6}$.

Special soundness of pp₂. We can extract from a process issuing a valid pp₂ a single $(\vec{C}_x, \vec{C}_y, \vec{C}_{y'}, \vec{I}_x, \vec{I}_y, \vec{I}_{y'})$ and the answers $(\vec{a}_x, \vec{a}_{r_x}, \vec{a}_y, \vec{a}_{r_y}, \vec{a}_{y'}, \vec{a}_{r_{y'}})$ to two different challenges ε and $\varepsilon + \Delta\varepsilon$. From the equations in verification, we can then set $\vec{x} = \frac{1}{\Delta\varepsilon} \Delta\vec{a}_x$, $\vec{r}_x = \frac{1}{\Delta\varepsilon} \Delta\vec{a}_{r_x}$, $\vec{y} = \frac{1}{\Delta\varepsilon} \Delta\vec{a}_y$, $\vec{r}_y = \frac{1}{\Delta\varepsilon} \Delta\vec{a}_{r_y}$, $\vec{y}' = \frac{1}{\Delta\varepsilon} \Delta\vec{a}_{y'}$, $\vec{r}_{y'} = \frac{1}{\Delta\varepsilon} \Delta\vec{a}_{r_{y'}}$, and obtain $\vec{C}_x = G.\vec{x} + H.\vec{r}_x$, $\vec{C}_y = G.\vec{y} + H.\vec{r}_y$, and $\vec{C}_{y'} = G.\vec{y}' + H.\vec{r}_{y'}$.

Ideally, we should have an adaptive extractable proof in pp₂ to extract the secrets. Since we only need this in unlinkability and we should extract from the initialization adversary \mathcal{A}_1 who is only setting up pp₂, we keep the proof as it is for simplicity. An adversary \mathcal{A}_1 succeeding to build a valid pp₂ without being extractable would have a probability of success that we denote by $\text{Adv}_{\mathcal{A}_1}^{\text{EXTRACT}}$.

Special soundness of π . We assume that the secrets $\vec{x}, \vec{y}, \vec{y}', \vec{z}, \vec{r}_x, \vec{r}_y, \vec{r}_{y'}$ have been correctly extracted. We consider an issuer who responds with (U, \vec{V}, t_S, π) to a request \vec{T} with a proof π which verifies, with probability larger than $\frac{1}{p}$. Hence, there exists a single $(U, \vec{V}, t_S, \vec{C}, \vec{C}_0, \vec{C}_1, C_d, \vec{C}_\rho, \vec{C}_w)$ and the valid answer $(e_0, e_1, \vec{a}_0, \vec{a}_1, a_d, \vec{a}_\rho, \vec{a}_w)$ to two different challenges e and $e + \Delta e$.

Thanks to the equations in verification, we obtain $H.\Delta\vec{a}_0 = \Delta e_0.\vec{C}$, $H.\Delta\vec{a}_1 = \Delta e_1.(\vec{C} - \vec{C}_y)$, so $\vec{C} = H.\frac{\Delta(\vec{a}_0 + \vec{a}_1)}{\Delta e} + \frac{\Delta e_1}{\Delta e} \vec{C}_y$. We also obtain $G = -\frac{\Delta a_d}{\Delta e} U$, $-(\vec{C}_x + \vec{C} + m\vec{C}_{y'} + t_S \vec{Z} + \vec{T}) = \frac{\Delta a_d}{\Delta e} \vec{V} + \frac{1}{\Delta e} H.\Delta\vec{a}_\rho$, $-\vec{T} = \frac{\Delta a_d}{\Delta e} \vec{V} + \frac{1}{\Delta e} G.\Delta\vec{a}_w$. As $G \neq 0$, we can invert and multiply the last three equations by $d = -\frac{\Delta e}{\Delta a_d}$.

If $\frac{\Delta e_1}{\Delta e} \neq 1$, we have $\Delta e_0 \neq 0$ so we obtain $\vec{C} = \frac{1}{\Delta e_0} H.\Delta\vec{a}_0$. Otherwise, we obtain $\vec{C} = \vec{C}_y + \frac{1}{\Delta e} H.\Delta(\vec{a}_0 + \vec{a}_1)$. In any case, we obtain $b, \vec{\mu}, d, \vec{\rho}, \vec{w}$ such that $\vec{C} = b.\vec{C}_y + H.\vec{\mu}$, $U = d.G$, $\vec{V} = d(\vec{C}_x + \vec{C} + m\vec{C}_{y'} + t_S \vec{Z} + \vec{T}) + H.\vec{\rho} = G.\vec{w} + d.\vec{T}$.

Given that $\vec{C}_x = G.\vec{x} + H.\vec{r}_x$, $\vec{C}_y = G.\vec{y} + H.\vec{r}_y$, $\vec{C}_{y'} = G.\vec{y}' + H.\vec{r}_{y'}$, and $\vec{Z} = G.\vec{z}$, the equations in \vec{V} give $(\vec{x} + b\vec{y} + m\vec{y}' + t_S \vec{z} - \frac{1}{d}.\vec{w})G + (\vec{r}_x + b\vec{r}_y + m\vec{r}_{y'} + \vec{\mu} + \frac{1}{d}.\vec{\rho})H = 0$. Unless the issuer can compute the discrete logarithm of H , this implies that we have $\frac{1}{d}\vec{w} = \vec{x} + b\vec{y} + m\vec{y}' + t_S \vec{z}$ and $-\vec{\mu} - \frac{1}{d}\vec{\rho} = \vec{r}_x + b\vec{r}_y + m\vec{r}_{y'} + t_S \vec{z}$. These equations imply $U = d.G$ and $\vec{V} = d(G\vec{x} + bG\vec{y} + mG\vec{y}' + t_S G\vec{z} + \vec{T})$. Thus, if no such (b, d) exist with $b \in \{0, 1\}$, the proof succeeds with probability bounded by $\frac{1}{p}$ plus the advantage of a discrete logarithm adversary.

The solution (b, d) such that $b \in \{0, 1\}$, $U = d.G$, and $\vec{V} = d(G\vec{x} + bG\vec{y} + mG\vec{y}' + t_S G\vec{z} + \vec{T})$ is unique when it exists. We let $\text{Adv}_{\mathcal{A}}^{\text{SOUND}}$ be the probability that an adversary \mathcal{A} succeeds to build a valid proof π when no solution exists. Using the forking lemma, we obtain the answer to an alternate challenge with probability $\frac{(\text{Adv}_{\mathcal{A}}^{\text{SOUND}})^2}{4q_{\text{Hash}}} - \frac{1}{p}$. Using the above extraction method, we deduce the discrete logarithm of H . This defines an adversary \mathcal{B} such that

$$\text{Adv}_{\mathcal{A}}^{\text{SOUND}} \leq 2\sqrt{q_{\text{Hash}}}\sqrt{\text{Adv}_{\mathcal{B}}^{\text{DLOG}} + \frac{1}{p}}$$

Variant without ROM. The Fiat-Shamir relies on the random oracle model (ROM). We can get rid of this assumption by following usual techniques [5]. For instance, we can use an additive homomorphic encryption scheme. The client

would generate a key pair and sends the encryption key pk together with the encryption $\text{Enc}_{\text{pk}}(e)$ of a random challenge e . Then, the server would do all operations homomorphically to return and encrypted π . The client would decrypt it.

A.5 Unforgeability

Theorem 4. *Under Assumption 1, ATHM is OMUF-secure. More precisely, given an OMUF-adversary \mathcal{A} making q oracle calls to $\mathcal{O}_{\text{sign}}$ and q_{Hash} calls to the random oracle, there exists an EF-CMA-adversary \mathcal{B} making $q + 3$ oracle calls to \mathcal{O}_{MAC} such that*

$$\text{Adv}_{\mathcal{A}}^{\text{OMUF}} \leq \text{Adv}_{\mathcal{B}}^{\text{EF-CMA}} + \frac{q + 1}{p} + \frac{q_{\text{Hash}}}{p^6}$$

This result needs $\Pi_1.\text{Extract}$ to succeed if and only if $\Pi_1.\text{Verify}$ returns true and Π_2 to be perfectly simulatable with a trapdoor using $\Pi_2.\text{SimKeyGen}$ and $\Pi_2.\text{Simulate}$.

Proof. We start with an adversary \mathcal{A} playing the one-more unforgeability game against ATHM and we transform it into a sequence of adversaries playing other games with the same success.

First of all, we change the game so that the issuer no longer sends a proof π but rather make the adversary to set up Π_2 , generate the keys, and simulate π perfectly. The simulation uses $\Pi_2.\text{SimKeyGen}$ and $\Pi_2.\text{Simulate}$. We obtain an adversary Adv_0 playing a game Γ_0 such that

$$\text{Adv}_{\mathcal{A}}^{\text{OMUF}} \leq \text{Adv}_{\mathcal{A}_0}^{\Gamma_0} + \frac{q_{\text{Hash}}}{p^6}$$

Then, we define a new adversary \mathcal{A}_1 who can use the extractable proof \vec{T}_{ext} in our scheme. Giving consistent \vec{T} and \vec{T}_{ext} proves the knowledge of t_C and \vec{r} with a way to extract them. With a trapdoor td_1 (in the CRS model), \mathcal{A}_1 can extract $(t_C, \vec{r}) = \Pi_1.\text{Extract}(\text{td}_1, \vec{T}_{\text{ext}})$ (or $(t_C, \vec{r}) = \Pi_1.\text{Extract}(\text{View})$ with View being the view of \mathcal{A} in the model based on Assumption 3). Then, \mathcal{A}_1 can call a MAC oracle to build (t, P, \vec{Q}) such that t and P are random and $\vec{Q} = (\vec{x} + b\vec{y} + m\vec{y}' + t\vec{z})P$. Then, \mathcal{A}_1 sets $t_S = t - t_C$, $U = P$, and $\vec{V} = \vec{Q} + \vec{r}U$ to return to the client as written on Figure 10. Note that this requires to simulate the proof π based on the trapdoor θ' from td_2 (or programming the random oracle). This way, an attacker for our scheme transforms into a attacker in the Γ_1 game on Figure 10.

The extraction procedure $\Pi_1.\text{Extract}(\text{td}_1, \vec{T}_{\text{ext}})$ fails iff $\Pi_1.\text{Verify}(\vec{T}, \vec{T}_{\text{ext}}, G, \vec{Z})$ is false. As forging π is perfectly simulatable, \mathcal{A}_1 playing Γ_1 perfectly simulates \mathcal{A} playing OMUF. If \mathcal{A} wins, there exists a bit b such that there are more t_j with $b_j = b$ than oracle calls to $\mathcal{O}_{\text{sign}}$ with bit b . Hence, there must exist one for which t_j does not match any \mathcal{O} call with the bit b . Therefore, $(b_j, m_j, t_j, \sigma_j)$ is a valid forgery. Therefore, we have

$$\text{Adv}_{\mathcal{A}_0}^{\Gamma_0} \leq \text{Adv}_{\mathcal{A}_1}^{\Gamma_1}$$

<p>Adversary $\mathcal{A}_1^{\mathcal{O}_{\text{verify}}}$ (gp, p, G, \vec{Z}):</p> <ol style="list-style-type: none"> 1: $\text{Setup}_3(\text{gp}, p, G) \rightarrow (\text{crs}_1, \text{td}_1)$ 2: $\text{Setup}_4(\text{gp}, p, G) \rightarrow (\text{crs}_2, \text{td}_2)$ 3: $\text{crs} \leftarrow (\text{gp}, p, G, \text{crs}_1, \text{crs}_2)$ 4: $\text{KeyGen}_1(\text{crs}, \vec{Z}) \rightarrow (\text{pp}_1, \text{sk}_1)$ 5: $\text{SimKeyGen}(\text{crs}, \vec{Z}) \rightarrow \text{pp}_2$ 6: $\text{pp} \leftarrow (\text{crs}, \vec{Z}, \text{pp}_1, \text{pp}_2)$ 7: run $\mathcal{A}^{\mathcal{O}_{\text{sign}}, \mathcal{O}_{\text{read}}}(\text{crs}, \text{pp}) \rightarrow$ $(b, m, (t'_j, \sigma_j)_j)$ 8: find j such that for all i, $(b, m, t'_j) \neq$ (b_i, m_i, t_i) (if none, return \perp) 9: return (b, m, t'_j, σ_j) 	<p>Subroutine $\mathcal{O}_{\text{sign}}(b, m, \text{query})$:</p> <ol style="list-style-type: none"> 10: parse $\text{query} \rightarrow (\vec{T}, \vec{T}_{\text{ext}})$ 11: $\Pi_1.\text{Extract}(\text{td}_1, \text{sk}_1, \vec{T}_{\text{ext}}) \rightarrow (t_C, \vec{r})$ 12: if extraction failed or $\vec{T} \neq t_C \vec{Z} + \vec{r}G$ then return \perp 13: pick $t \leftarrow_{\\$} \mathbb{Z}_p$ 14: $\mathcal{O}(b, m, t) \rightarrow (P, \vec{Q})$ 15: $t_S \leftarrow t - t_C$ 16: $U \leftarrow P$ 17: $\vec{V} \leftarrow \vec{Q} + \vec{r}U$ 18: $\pi \leftarrow \Pi_2.\text{Simulate}(U, \vec{V}, t_S, \text{crs}, \text{pp}, \text{td}_2)$ 19: return (U, \vec{V}, t_S, π) <p>Subroutine $\mathcal{O}_{\text{read}}(m, t, \sigma)$:</p> <ol style="list-style-type: none"> 20: if $\mathcal{O}_{\text{verify}}(0, m, t, \sigma)$ then return 0 21: if $\mathcal{O}_{\text{verify}}(1, m, t, \sigma)$ then return 1 22: return \perp
<p>Game Γ_1:</p> <ol style="list-style-type: none"> 1: $\text{Setup}_1(1^\lambda) \rightarrow (\text{gp}, p)$ 2: $\text{Setup}_2(\text{gp}, p) \rightarrow G$ 3: pick $\vec{x} \leftarrow_{\\$} \mathcal{E}_x$, $y \leftarrow_{\\$} \mathcal{E}_y$, $y' \leftarrow_{\\$} \mathcal{E}_{y'}$, $z \leftarrow_{\\$} \mathcal{E}_z$ 4: set $\vec{Z} \leftarrow zG$ 5: initialize $i \leftarrow 0$ 6: $\mathcal{A}_1^{\mathcal{O}_{\text{verify}}}(\text{gp}, p, G, \vec{Z}) \rightarrow (b, m, t, P, \vec{Q})$ 7: if (b, m, t) is equal to some (b_i, m_i, t_i) then abort 8: win iff $\mathcal{O}_{\text{verify}}(b, m, t, P, \vec{Q})$ 	<p>Oracle $\mathcal{O}(b, m, t)$:</p> <ol style="list-style-type: none"> 9: increment i 10: $(b_i, m_i, t_i) \leftarrow (b, m, t)$ 11: pick a nonzero group element P_i 12: $\vec{Q}_i \leftarrow (\vec{x} + b_i \vec{y} + m_i \vec{y}' + t_i \vec{z})P_i$ 13: return (P_i, \vec{Q}_i) <p>Oracle $\mathcal{O}_{\text{verify}}(b, m, t, P, \vec{Q})$:</p> <ol style="list-style-type: none"> 14: if $P = 0$ then return false 15: return whether $\vec{Q} = (\vec{x} + b\vec{y} + m\vec{y}' + t\vec{z})P$

Fig. 10. Unforgeability Step 1: \mathcal{A}_1 playing Γ_1

We now get rid of the selection of G and of the input \vec{Z} to the adversary by a change of variable $G = PP_0$, $\vec{Z} = \vec{Q}\vec{Q}_0$, $\vec{x} = \vec{x}\vec{x}$, $\vec{y} = \vec{y}\vec{y}$, $\vec{y}' = \vec{y}\vec{y}'$, $\vec{z} = \vec{x}\vec{x} + \text{tt}_0 \cdot \vec{z}\vec{z}$, which uses an initial MAC $(\text{bb}_0, \text{mm}_0, \text{tt}_0, \text{PP}_0, \vec{Q}\vec{Q}_0)$ with secret $(\vec{x}\vec{x}, \vec{y}\vec{y}, \vec{y}\vec{y}', \vec{z}\vec{z})$ and input $\text{bb}_0 = 0$, $\text{mm}_0 = 0$, and $\text{tt}_0 = 1$. We set $\text{bb} = \frac{b}{1+t}$, $\text{mm} = \frac{m}{1+t}$, $\text{tt} = \frac{t \cdot \text{tt}_0}{1+t}$, $\text{PP} = P$ and $\vec{Q}\vec{Q} = (1+t)^{-1}\vec{Q}$ so that

$$\vec{x}\vec{x} + \text{bb} \cdot \vec{y}\vec{y} + \text{mm} \cdot \vec{y}\vec{y}' + \text{tt} \cdot \vec{z}\vec{z} = (1+t)^{-1}(\vec{x}\vec{x} + b\vec{y}\vec{y} + m\vec{y}\vec{y}' + t\vec{z}\vec{z})$$

Thus, an adversary \mathcal{A}_1 for Γ_1 transforms into another adversary \mathcal{B} in the EF-CMA game as depicted on Figure 11.

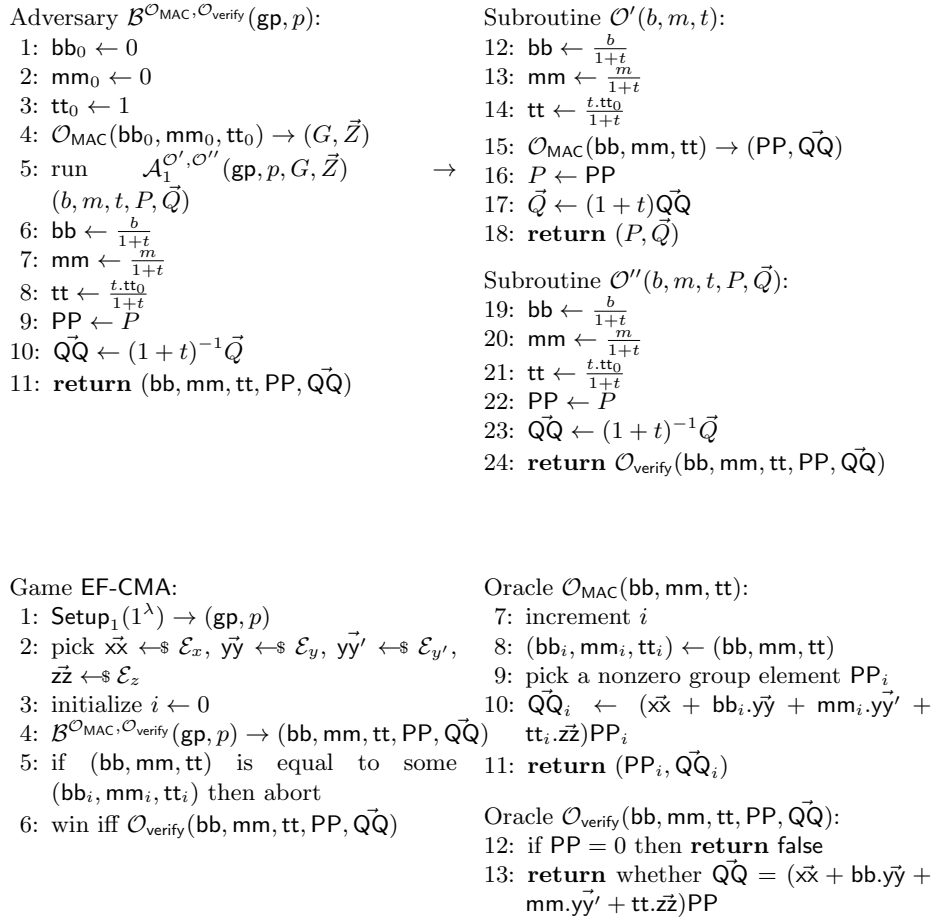


Fig. 11. Unforgeability Step 2: \mathcal{B} playing EF-CMA

Some rare failure event may happen: during the selection of \mathbf{tt}_0 , we may obtain $\vec{z} = \mathbf{xx} + \mathbf{tt}_0 \cdot \vec{\mathbf{z}} = 0$ making an invalid \vec{z} . Furthermore, during the selection of any t , we may obtain $t = -1$ making the change of variable fail. However, by the difference Lemma, we have

$$\text{Adv}_{\mathcal{A}_1}^{\Gamma_1} \leq \text{Adv}_{\mathcal{B}}^{\text{EF-CMA}} + \frac{q+1}{p}$$

□

A.6 Unlinkability

The role of the NIZK is important as it makes sure that the issuer must use either $b = 0$ or $b = 1$ and therefore hides only one bit. Without the NIZK, the issuer could use more than one bit with b and use this as a marker to link tokens to redeem to clients requesting a token. So, the client must verify the NIZK proof for unlinkability. The client must also verify $U \neq 0$, because $U = 0$ could be used by the issuer to mark a token.

The issuer knows which bit is hidden during an issuing session and can extract the hidden bit during redeem. Hence, we should only consider unlinkability when the bits are the same.

Theorem 5. *ATHM is 2-UNLINK-secure. More precisely, given an UNLINK-adversary \mathcal{A} making oracle calls to $\mathcal{O}_{\text{query}}$ with index set $\mathcal{Q}_{\text{query}}$, there exist an adversary \mathcal{A}' making a valid pp_2 without being extractable, $\#\mathcal{Q}_{\text{query}}$ SOUND-adversaries \mathcal{B}_i against the NIZK and $\#\mathcal{Q}_{\text{query}}$ distinguishers \mathcal{D}_i on T_{ext} such that*

$$\text{Adv}_{\mathcal{A}}^{\text{UNLINK}} \leq \frac{2}{n} + \text{Adv}_{\mathcal{A}'}^{\text{EXTRACT}} + \sum_{i \in \mathcal{Q}_{\text{query}}} \text{Adv}_{\mathcal{B}_i}^{\text{SOUND}} + \sum_{i \in \mathcal{Q}_{\text{query}}} \text{Adv}_{\mathcal{D}_i}^{\text{IND}}$$

The proof uses the fact that Π_2 is sound and that T_{ext} can be simulated using a trapdoor in a computationally indistinguishable manner.

Proof. We start with an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ playing the UNLINK_n game.

We assume that AT.ClientQuery verifies pp , at least at the first time it is run. Using the soundness of the proof of $\vec{x}, \vec{r}_x, \vec{y}, \vec{r}_y, \vec{y}', \vec{r}_{y'}$ in pp_2 (based on that \mathcal{A}_1 is only running this unique proof), there is an extractor which extracts from \mathcal{A}' those secrets, except with negligible probability $\text{Adv}_{\mathcal{A}'}^{\text{EXTRACT}}$. We reduce to a game Γ_0 in which this extraction succeeds. We have

$$\text{Adv}_{\mathcal{A}}^{\text{UNLINK}} \leq \text{Adv}_{\mathcal{A}}^{\Gamma_0} + \text{Adv}_{\mathcal{A}'}^{\text{EXTRACT}}$$

Second, we use the simulator of Π_1 to simulate all \vec{T}_{ext} and reduce to a protocol using no \vec{T}_{ext} . For that, we define distinguishers \mathcal{D}_i between a valid \vec{T}_{ext} and a simulated one. We obtain a game Γ_1 with no Π_1 and an adversary \mathcal{A}_0 such that

$$\text{Adv}_{\mathcal{A}}^{\Gamma_0} \leq \text{Adv}_{\mathcal{A}_0}^{\Gamma_1} + \sum_{i \in \mathcal{Q}_{\text{query}}} \text{Adv}_{\mathcal{D}_i}^{\text{IND}}$$

Then, we reduce to a game Γ in which for every i , the $(m_i, \text{query}_i, \text{resp}_i)$ triplets, parsed as $\text{query}_i = \vec{T}_i$ and $\text{resp}_i = (U_i, \vec{V}_i, t_{i,S}, \pi_i)$, is such that there exists (b_i, d_i) such that $b_i \in \{0, 1\}$, $U_i = d_i G$, and $\vec{V}_i = d_i(G.\vec{x} + b_i G.\vec{y} + m_i.\vec{y}' + t_{i,S} G.\vec{z} + \vec{T}_i)$. Note that (b_i, d_i) is unique when it exists. Thanks to the soundness of Π_2 , we have $\#\mathcal{Q}_{\text{query}}$ adversaries \mathcal{B}_i to build a valid proof when no (b_i, d_i) exists, such that

$$\text{Adv}_{\mathcal{A}_0}^{\Gamma_1} \leq \text{Adv}_{\mathcal{A}_0}^{\Gamma} + \sum_{i \in \mathcal{Q}_{\text{query}}} \text{Adv}_{\mathcal{B}_i}^{\text{SOUND}}$$

We deduce that $\vec{Q}_i = (\vec{x} + b_i \vec{y} + m_i \vec{y}' + t_i \vec{z}) P_i$ with $b_i \in \{0, 1\}$ in the game Γ . We also note that all m_i ($i \in \mathcal{Q}$) must be equal to a common m as required in the UNLINK game.

The rest of the proof is an information theoretic argument for which complexities do not matter. Given $(\vec{x}, \vec{y}, \vec{y}', \vec{z}, \vec{T}_i, U_i, \vec{V}_i, t_{i,S}, m, \pi_i)$ we can uniquely determine b_i . We observe that $(t_i, P_i) | (\vec{x}, \vec{y}, \vec{y}', \vec{z}, \vec{T}_i, U_i, \vec{V}_i, t_{i,S}, m, \pi_i)$ is uniformly distributed as a pair composed of a scalar and a nonzero group element. Hence, whenever \mathcal{A}_2 returns \mathcal{Q} and the list of resp_i , it determines the values of the b_i but the (t_i, P_i) to be released are still uniform. After permutation, $(t_{\sigma(i)}, P_{\sigma(i)}, \vec{Q}_{\sigma(i)})$ has a value of $\vec{Q}_{\sigma(i)}$ which is imposed by $\vec{Q}_{\sigma(i)} = (\vec{x} + b_{\sigma(i)} \vec{y} + m \vec{y}' + t_{\sigma(i)} \vec{z}) P_{\sigma(i)}$ so brings $b_{\sigma(i)}$ as only information.

This reduces to the following game: the adversary chooses a list of bits $(b_i)_{i \in \mathcal{Q}}$ with $\#\mathcal{Q} \geq n$, the game selects a random i^* and a random permutation σ then provides b_{i^*} and $(b_{\sigma(i)})_{i \in \mathcal{Q}}$ to the adversary, and the adversary finally makes a guess i and win if and only if $i = i^*$. If the adversary puts n_0 zeros and n_1 ones, the adversary can only win with probability $\frac{1}{n_0}$ when it is a zero (which happens with probability $\frac{n_0}{n_0+n_1}$), and with probability $\frac{1}{n_1}$ when it is a one (which happens with probability $\frac{n_1}{n_0+n_1}$). Overall, the adversary wins with probability $\frac{2}{n_0+n_1}$ which is at most $\frac{2}{n}$ since $n_0 + n_1 = \#\mathcal{Q} \geq n$. \square

A.7 Privacy of Metadata

Theorem 6. *Under Assumption 2, ATHM is PMB-secure. More precisely, given an PMB-adversary \mathcal{A} making q oracle calls to $\mathcal{O}_{\text{sign}}$ and q_{Hash} calls to the random oracle, there exists an IND-CMA-adversary \mathcal{B} making $q + 3$ oracle calls to \mathcal{O}_{MAC} such that*

$$\text{Adv}_{\mathcal{A}}^{\text{PMB}} \leq \text{Adv}_{\mathcal{B}}^{\text{IND-CMA}} + 4 \frac{q+1}{p} + 2 \frac{q_{\text{Hash}}}{p^6}$$

This result needs $\Pi_1.\text{Extract}$ to be succeed if and only if $\Pi_1.\text{Verify}$ returns true and Π_2 to be perfectly simulatable with a trapdoor using $\Pi_2.\text{SimKeyGen}$ and $\Pi_2.\text{Simulate}$.

Proof. We start with an adversary \mathcal{A} playing the PMB game against ATHM and we transform it into a sequence of adversaries playing other games with same success.

We start by the same first step as in the proof of unforgeability. We define an adversary \mathcal{A}_1 who runs \mathcal{A} but simulates π in Π_2 . With a trapdoor td_1 (in the CRS model), \mathcal{A}_1 can also extract $(t_C, \vec{r}) = \Pi_1.\text{Extract}(\text{td}_1, \vec{T}_{\text{ext}})$. Then, \mathcal{A}_1 can call a MAC oracle to build (t, P, \vec{Q}) such that t and P are random and $\vec{Q} = (\vec{x} + b\vec{y} + m\vec{y}' + t\vec{z})P$. Then, \mathcal{A}_1 can set $t_S = t - t_C$, $U = P$, and $\vec{V} = \vec{Q} + \vec{r}U$ to return to \mathcal{A} . More precisely, the adversary \mathcal{A}_1 is fully written on Figure 12. This way, an adversary \mathcal{A} for our scheme transforms into an adversary \mathcal{A}_1 in the Γ_{b^*} game on Figure 12.

The extraction procedure is such that $\Pi_1.\text{Extract}(\text{td}, \vec{T}_{\text{ext}})$ fails is equivalent to $\Pi_1.\text{Verify}(\vec{T}, \vec{T}_{\text{ext}}, G, \vec{Z})$ is false. As forging π is perfect, \mathcal{A}_1 playing Γ_{b^*} perfectly simulates \mathcal{A} playing PMB_{b^*} .

$$\text{Adv}_{\mathcal{A}}^{\text{PMB}} \leq \text{Adv}_{\mathcal{A}_1}^{\Gamma} + \frac{q_{\text{Hash}}}{p^6}$$

In the next step, we remove G , \vec{Z} , and pp_2 as in the proof of unforgeability, with the same change of variable. We define an adversary \mathcal{B} playing the IND-CMA game on Figure 13.

For \mathcal{B} in IND-CMA_{b^*} to simulate \mathcal{A}_1 in Γ_{b^*} , there are a few things which can go wrong. Like in the previous proof, we can select the secrets such that $\vec{x}\vec{x} + \text{tt}_0.\vec{z}\vec{z} = 0$ or any t equal to -1 which make the change of variable fail. This happens with probability at most $\frac{q+1}{p}$. We can also select $t^* = \text{tt}_0$ (with probability $\frac{1}{p}$). Finally, any t in \mathcal{O}_{MAC} in Γ_{b^*} can be selected equal to t^* which makes $\mathcal{O}'_{\text{MAC}}$ abort. This happens with probability at most $\frac{q}{p}$. By the difference Lemma, we obtain

$$|\Pr[\Gamma_{b^*}(\mathcal{A}_1) \rightarrow 1] - \Pr[\text{IND-CMA}_{b^*}(\mathcal{B}) \rightarrow 1]| \leq 2\frac{q+1}{p}$$

We deduce

$$\left| \text{Adv}_{\mathcal{A}_1}^{\Gamma} - \text{Adv}_{\mathcal{B}}^{\text{IND-CMA}} \right| \leq 4\frac{q+1}{p}$$

□

A.8 Extensions of ATHM

We could add public verifiability of tokens by using a pairing. This can only be done with no private metadata (otherwise privacy is broken). We could also consider batching too like in PMBT, so that the issuer could make a single proof for a batch of responses. Clearly, algebraic MACs offer lots of flexibility for extensions. This is left as future work.

Adversary $\mathcal{A}_1^{\mathcal{O}_{\text{MAC}}, \mathcal{O}_{\text{chal}}, \mathcal{O}_{\text{verify}}, \mathcal{O}_{\text{valid}}}$
 (gp, p, G, Z) :

- 1: $\text{Setup}_3(\text{gp}, p, G) \rightarrow (\text{crs}_1, \text{td}_1)$
- 2: $\text{Setup}_4(\text{gp}, p, G) \rightarrow (\text{crs}_2, \text{td}_2)$
- 3: $\text{crs} \leftarrow (\text{gp}, p, G, \text{crs}_1, \text{crs}_2)$
- 4: $\text{KeyGen}_1(\text{crs}, \vec{Z}) \rightarrow (\text{pp}_1, \text{sk}_1)$
- 5: $\Pi_2.\text{SimKeyGen}(\text{crs}, \vec{Z}) \rightarrow \text{pp}_2$
- 6: $\text{pp} \leftarrow (\text{crs}, Z, \text{pp}_1, \text{pp}_2)$
- 7: **return** $\mathcal{A}^{\mathcal{O}_{\text{sign}}, \mathcal{O}'_{\text{sign}}, \mathcal{O}_{\text{read}}, \mathcal{O}_{\text{valid}}}(\text{crs}, \text{pp})$

Subroutine $\mathcal{O}_{\text{sign}}(b, m, \text{query})$:

- 8: parse $\text{query} = (\vec{T}, \vec{T}_{\text{ext}})$
- 9: $\Pi_1.\text{Extract}(\text{td}_1, \text{sk}_1, \vec{T}_{\text{ext}}) \rightarrow (t_C, \vec{r})$
- 10: if extraction failed or $\vec{T} \neq t_C \vec{Z} + \vec{r}G$
then **return** \perp
- 11: pick $t \leftarrow_{\$} \mathbb{Z}_p$
- 12: $\mathcal{O}_{\text{MAC}}(b, m, t) \rightarrow (P, \vec{Q})$
- 13: $t_S \leftarrow t - t_C$
- 14: $U \leftarrow P$
- 15: $\vec{V} \leftarrow \vec{Q} + \vec{r}U$
- 16: $\pi \leftarrow \Pi_2.\text{Simulate}(U, \vec{V}, t_S, \text{crs}, \text{pp}, \text{td}_2)$
- 17: **return** (U, \vec{V}, t_S, π)

Subroutine $\mathcal{O}'_{\text{sign}}(m, \text{query})$:

- 18: parse $\text{query} = (\vec{T}, \vec{T}_{\text{ext}})$
- 19: $\text{Extract}(\text{td}_1, \text{sk}_1, \vec{T}_{\text{ext}}) \rightarrow (t_C, \vec{r})$
- 20: if extraction failed or $\vec{T} \neq t_C \vec{Z} + \vec{r}G$
then **return** \perp
- 21: pick $t \leftarrow_{\$} \mathbb{Z}_p$
- 22: $\mathcal{O}_{\text{chal}}(m, t) \rightarrow (P, \vec{Q})$
- 23: $t_S \leftarrow t - t_C$
- 24: $U \leftarrow P$
- 25: $\vec{V} \leftarrow \vec{Q} + \vec{r}U$
- 26: $\pi \leftarrow \Pi_2.\text{Simulate}(U, \vec{V}, t_S, \text{crs}, \text{pp}, \text{td}_2)$
- 27: **return** (U, \vec{V}, t_S, π)

Subroutine $\mathcal{O}_{\text{read}}(m, t, \sigma)$:

- 28: if **flag** then **return** \perp
- 29: if $\mathcal{O}_{\text{verify}}(0, m, t, \sigma)$ then **return** 0
- 30: if $\mathcal{O}_{\text{verify}}(1, m, t, \sigma)$ then **return** 1
- 31: **return** \perp

Game Γ_{b^*} :

- 1: $\text{Setup}_1(1^\lambda) \rightarrow (\text{gp}, p)$
- 2: $\text{Setup}_2(\text{gp}, p) \rightarrow G$
- 3: pick $\vec{x} \leftarrow_{\$} \mathcal{E}_x, \vec{y} \leftarrow_{\$} \mathcal{E}_y, \vec{y}' \leftarrow_{\$} \mathcal{E}_{y'}, \vec{z} \leftarrow_{\$} \mathcal{E}_z$
- 4: set $\vec{Z} \leftarrow \vec{z}G$
- 5: initialize $i \leftarrow 0, \text{flag} \leftarrow \text{false}$
- 6: **return** $\mathcal{A}_1^{\mathcal{O}_{\text{MAC}}, \mathcal{O}_{\text{chal}}, \mathcal{O}_{\text{verify}}, \mathcal{O}_{\text{valid}}}(\text{gp}, p, G, \vec{Z})$

Oracle $\mathcal{O}_{\text{chal}}(m, t)$:

- 10: if **flag** then **return** \perp
- 11: **flag** \leftarrow **true**
- 12: $(m^*, t^*) \leftarrow (m, t)$
- 13: $\mathcal{O}_{\text{MAC}}(b^*, m^*, t^*) \rightarrow (P^*, \vec{Q}^*)$
- 14: **return** (P^*, \vec{Q}^*)

Oracle $\mathcal{O}_{\text{verify}}(b, m, t, P, \vec{Q})$:

- 15: if **flag** and $b \in \{0, 1\}$ and $(t, m) = (t^*, m^*)$ then **return** \perp
- 16: if $P = 0$ then **return** **false**
- 17: **return** whether $\vec{Q} = (\vec{x} + b\vec{y} + m\vec{y}' + t\vec{z})P$

Oracle $\mathcal{O}_{\text{MAC}}(b, m, t)$:

- 7: pick a nonzero group element P
- 8: $\vec{Q} \leftarrow (\vec{x} + b\vec{y} + m\vec{y}' + t\vec{z})P$
- 9: **return** (P, \vec{Q})

Oracle $\mathcal{O}_{\text{valid}}(t, m, P, \vec{Q})$:

- 18: if $P = 0$ then **return** **false**
- 19: **return** whether $\exists b \in \{0, 1\} \vec{Q} = (\vec{x} + b\vec{y} + m\vec{y}' + t\vec{z})P$

Fig. 12. Privacy Step 1: \mathcal{A}_1 playing Γ

Adversary $\mathcal{B}^{\mathcal{O}'_{\text{MAC}}, \mathcal{O}'_{\text{chal}}, \mathcal{O}'_{\text{verify}}, \mathcal{O}'_{\text{valid}}}(\text{gp}, p)$:

- 1: $\text{bb}_0 \leftarrow 0$
- 2: $\text{mm}_0 \leftarrow 0$
- 3: $\text{tt}_0 \leftarrow 1$
- 4: $\mathcal{O}'_{\text{MAC}}(\text{bb}_0, \text{mm}_0, \text{tt}_0) \rightarrow (G, \vec{Z})$
- 5: **return** $\mathcal{A}_1^{\mathcal{O}'_{\text{MAC}}, \mathcal{O}'_{\text{chal}}, \mathcal{O}'_{\text{verify}}, \mathcal{O}'_{\text{valid}}}(\text{gp}, p, G, \vec{Z})$

Subroutine $\mathcal{O}'_{\text{verify}}(b, m, t, P, \vec{Q})$:

- 6: $\text{bb} \leftarrow \frac{b}{1+t}, \text{mm} \leftarrow \frac{m}{1+t}, \text{tt} \leftarrow \frac{t \cdot \text{tt}_0}{1+t}$
- 7: $\text{PP} \leftarrow P, \vec{\text{QQ}} \leftarrow (1+t)^{-1} \vec{Q}$
- 8: **return** $\mathcal{O}'_{\text{verify}}(\text{bb}, \text{mm}, \text{tt}, \text{PP}, \vec{\text{QQ}})$

Subroutine $\mathcal{O}'_{\text{valid}}(m, t, P, \vec{Q})$:

- 9: $\text{bb} \leftarrow 0, \text{bb}' \leftarrow \frac{1}{1+t}$
- 10: $\text{mm} \leftarrow \frac{m}{1+t}, \text{tt} \leftarrow \frac{t \cdot \text{tt}_0}{1+t}$
- 11: $\text{PP} \leftarrow P, \vec{\text{QQ}} \leftarrow (1+t)^{-1} \vec{Q}$
- 12: **return** $\mathcal{O}'_{\text{valid}}(\text{bb}, \text{bb}', \text{mm}, \text{tt}, \text{PP}, \vec{\text{QQ}})$

Subroutine $\mathcal{O}'_{\text{MAC}}(b, m, t)$:

- 13: $\text{bb} \leftarrow \frac{b}{1+t}, \text{mm} \leftarrow \frac{m}{1+t}, \text{tt} \leftarrow \frac{t \cdot \text{tt}_0}{1+t}$
- 14: $\mathcal{O}'_{\text{MAC}}(\text{bb}, \text{mm}, \text{tt}) \rightarrow (\text{PP}, \vec{\text{QQ}})$
- 15: $P \leftarrow \text{PP}, \vec{Q} \leftarrow (1+t) \vec{\text{QQ}}$
- 16: **return** (P, \vec{Q})

Subroutine $\mathcal{O}'_{\text{chal}}(m, t)$:

- 17: $\text{bb} \leftarrow 0, \text{bb}' \leftarrow \frac{1}{1+t}$
- 18: $\text{mm} \leftarrow \frac{m}{1+t}, \text{tt} \leftarrow \frac{t \cdot \text{tt}_0}{1+t}$
- 19: $\mathcal{O}'_{\text{chal}}(\text{bb}, \text{bb}', \text{mm}, \text{tt}) \rightarrow (\text{PP}, \vec{\text{QQ}})$
- 20: $P \leftarrow \text{PP}$
- 21: $\vec{Q} \leftarrow (1+t) \vec{\text{QQ}}$
- 22: **return** (P, \vec{Q})

Game IND-CMA_{b^*} :

- 1: $\text{flag} \leftarrow \text{false}, \text{mm}^* \leftarrow \perp$
- 2: $\text{Setup}_1(1^\lambda) \rightarrow (\text{gp}, p)$
- 3: pick $\vec{x}\vec{x} \leftarrow \mathcal{E}_x, \vec{y}\vec{y} \leftarrow \mathcal{E}_y, \vec{y}\vec{y}' \leftarrow \mathcal{E}_{y'}, \vec{z}\vec{z} \leftarrow \mathcal{E}_z$
- 4: **return** $\mathcal{B}^{\mathcal{O}'_{\text{MAC}}, \mathcal{O}'_{\text{chal}}, \mathcal{O}'_{\text{verify}}, \mathcal{O}'_{\text{valid}}}(\text{gp}, p)$

Oracle $\mathcal{O}'_{\text{verify}}(\text{bb}, \text{mm}, \text{tt}, \text{PP}, \vec{\text{QQ}})$:

- 5: if flag and $\text{bb} \in \{\text{bb}_0^*, \text{bb}_1^*\}$ and $(\text{mm}, \text{tt}) = (\text{mm}^*, \text{tt}^*)$ then **return** \perp
- 6: if $\text{PP} = 0$ then **return** false
- 7: **return** whether $\vec{\text{QQ}} = (\vec{x}\vec{x} + \text{bb} \cdot \vec{y}\vec{y} + \text{mm} \cdot \vec{y}\vec{y}' + \text{tt} \cdot \vec{z}\vec{z}) \text{PP}$

Oracle $\mathcal{O}'_{\text{valid}}(\text{bb}_0, \text{bb}_1, \text{mm}, \text{tt}, \text{PP}, \vec{\text{QQ}})$:

- 8: if $\neg \text{flag}$ then **return** \perp
- 9: if $(\text{tt}, \text{mm}) = (\text{tt}^*, \text{mm}^*)$ and $\{\text{bb}_0, \text{bb}_1\} \neq \{\text{bb}_0^*, \text{bb}_1^*\}$ then **return** \perp
- 10: if $\text{PP} = 0$ then **return** false
- 11: if $\vec{\text{QQ}} = (\vec{x}\vec{x} + \text{bb}_0 \cdot \vec{y}\vec{y} + \text{mm} \cdot \vec{y}\vec{y}' + \text{tt} \cdot \vec{z}\vec{z}) \text{PP}$ then **return** true
- 12: if $\vec{\text{QQ}} = (\vec{x}\vec{x} + \text{bb}_1 \cdot \vec{y}\vec{y} + \text{mm} \cdot \vec{y}\vec{y}' + \text{tt} \cdot \vec{z}\vec{z}) \text{PP}$ then **return** true
- 13: **return** false

Oracle $\mathcal{O}'_{\text{MAC}}(\text{bb}, \text{mm}, \text{tt})$:

- 14: if flag and $(\text{bb}, \text{mm}, \text{tt}) \in \{(\text{bb}_0, \text{mm}^*, \text{tt}^*), (\text{bb}_1, \text{mm}^*, \text{tt}^*)\}$ then **return** \perp
- 15: increment i
- 16: $(\text{bb}_i, \text{mm}_i, \text{tt}_i) \leftarrow (\text{bb}, \text{mm}, \text{tt})$
- 17: pick a nonzero group element PP_i
- 18: $\vec{\text{QQ}}_i \leftarrow (\vec{x}\vec{x} + \text{bb}_i \cdot \vec{y}\vec{y} + \text{mm}_i \cdot \vec{y}\vec{y}' + \text{tt}_i \cdot \vec{z}\vec{z}) \text{PP}_i$
- 19: **return** $(\text{PP}_i, \vec{\text{QQ}}_i)$

Oracle $\mathcal{O}'_{\text{chal}}(\text{bb}, \text{bb}', \text{mm}, \text{tt})$:

- 20: if flag or $\exists i (\text{bb}_i, \text{mm}_i, \text{tt}_i) \in \{(\text{bb}, \text{mm}, \text{tt}), (\text{bb}', \text{mm}, \text{tt})\}$ then **return** \perp
- 21: $(\text{bb}_0^*, \text{bb}_1^*, \text{mm}^*, \text{tt}^*) \leftarrow (\text{bb}, \text{bb}', \text{mm}, \text{tt})$
- 22: $\text{flag} \leftarrow \text{true}$
- 23: pick a nonzero group element PP^*
- 24: $\vec{\text{QQ}}^* \leftarrow (\vec{x}\vec{x} + \text{bb}_0^* \cdot \vec{y}\vec{y} + \text{mm}^* \cdot \vec{y}\vec{y}' + \text{tt}^* \cdot \vec{z}\vec{z}) \text{PP}^*$
- 25: **return** $(\text{PP}^*, \vec{\text{QQ}}^*)$

Fig. 13. Privacy Step 2: \mathcal{B} playing IND-CMA