

Cryptographic Security of Macaroon Authorization Credentials

Adriana López-Alt
New York University

December 6, 2013

Abstract

Macaroons, recently introduced by Birgisson et al. [BPUE⁺14], are authorization credentials that provide support for controlled sharing in decentralized systems. Macaroons are similar to cookies in that they are bearer credentials, but unlike cookies, macaroons include *caveats* that attenuate and contextually confine when, where, by who, and for what purpose authorization should be granted.

In this work, we formally study the cryptographic security of macaroons. We define *macaroon schemes*, introduce corresponding security definitions and provide several constructions. In particular, the MAC-based and certificate-based constructions outlined in [BPUE⁺14], can be seen as instantiations of our definitions. We also present a new construction that is privately-verifiable (similar to the MAC-based construction) but where the verifying party does not learn the intermediate keys of the macaroon, a problem already observed in [BPUE⁺14].

We also formalize the notion of a protocol for “discharging” third-party caveats and present a security definition for such a protocol. The encryption-based protocol outlined by Birgisson et al. [BPUE⁺14] can be seen as an instantiation of our definition, and we also present a new signature-based construction.

Finally, we formally prove the security of all constructions in the given security models.

1 Introduction

Macaroons, as introduced by Birgisson et al. [BPUE⁺14], are bearer tokens that act as authorization credentials. In its most basic form, a macaroon identifies a list of restrictions or *caveats* that specify when, where, by who, and for what purpose, access to a resource can be granted. For example, a caveat might specify that the macaroon is valid (and thus access should be granted) as long as the operation being performed is ‘read’. Another caveat might specify that the macaroon is valid as long as the object being accessed is ‘myimage.jpg’. These examples are what Birgisson et al. call *first-party caveats*. The conditions they impose can be verified by the target service that manages the resource when the access request is made and the macaroon is presented. More interesting are *third-party caveats*, caveats that require a third-party to assert that a certain condition holds true. For example, a third-party caveat might require proof of membership in a group G at a third-party service A . A recipient of a macaroon holding a third-party caveat must first obtain such a proof or *discharge* from the corresponding third-party, and pass along the discharge with the macaroon when it accesses the intended resource at the target service.

Of course, care must be taken to ensure that caveats are ‘bound together’ so that a principal holding a macaroon cannot remove or modify caveats at his or her convenience. Therefore, macaroons include a cryptographic signature that guarantees the integrity of the list of caveats and the caveats themselves.

Macaroons also allow *delegation* and *attenuation* by permitting additional caveats to be added to a macaroon, further restricting the validity of the token. More importantly, this can be done using only the macaroon itself (and perhaps an additional *extension key* created at the same time as the macaroon), and does not require communication with other principals or the target service. It is this flexibility that makes macaroons a leading candidate for authorization in decentralized distributed systems, and in ‘cloud’ services in particular.

Our Results. In this work, we formally study the cryptographic security of macaroons. We divide this task into three parts. First, we study the security of *macaroon schemes*, a new primitive that captures the requirements of integrity and extension (for delegation or attenuation) of a macaroon. Second, we study *discharge protocols*, a new primitive that specifies the syntax of third-party caveats and how they can be used to obtain discharges from the corresponding third-party. Finally, we study *macaroon-tree protocols*, a new primitive that captures the full functionality of macaroons: issuing, extending with a first-party caveat, extending with a third-party caveat, discharging, and verifying.

We follow the formalization of Birgisson et al. and model a macaroon as follows. A macaroon includes a list of caveats, where each caveat is a bit-string $m_i \in \{0,1\}^*$. Additionally, a macaroon includes a key identifier id that opaquely encodes a root key k , and a signature σ that guarantees the integrity of the macaroon and can be verified using the key k .

Macaroon Schemes. We define a new primitive called a *macaroon scheme* that defines how macaroons can be issued, extended (for delegation or attenuation), finalized, and verified. A key generation algorithm generates root keys k , with which a macaroon can be issued. When a macaroon is issued, an *extension key* is also given. Using this extension key, the new macaroon can be extended, essentially adding more caveats to it, and a new extension key is issued so as to allow further extensions. A verification algorithm is also defined such that given the corresponding root key k , the validity of a macaroon μ can be verified. Finally, an algorithm `Finalize` is defined, that modifies a given macaroon so that it can be verified by the verification algorithm. This last algorithm can be seen as optional – indeed, only one of our constructions requires it.

We formalize a security definition for macaroon schemes that captures the desired guarantee that a principal cannot forge a macaroon. In other words, a principal cannot create a valid macaroon μ^* with respect to a root key k unless it obtained the macaroon via “valid means”, i.e. μ^* was given to him, and μ^* was issued using k or is an extension of such a macaroon. Notice that this provides security against removing or modifying caveats from a macaroon obtained legally.

We provide three constructions of macaroon schemes. Two of these, a PRF-based construction and a certificate-based construction were already described in [BPUE⁺14]. We formalize these constructions and the properties required to prove their security. The third construction is new, and offers efficiency vs. security trade-offs when compared to the other two.

Finally, we make a distinction between a privately-verifiable macaroon scheme (as informally defined above), where a macaroon is issued and verified using the same root key k , and a publicly-verifiable macaroon scheme, where a macaroon is issued with a secret root key sk and verified with a corresponding (public) verification key vk . See Section 3.1 for a formal discussion of privately-verifiable macaroon schemes, and Section 3.2 for a formal discussion of publicly-verifiable macaroon schemes.

Discharge Protocols. We formally define *discharge protocols*, a new primitive that specifies the syntax of third-party caveats and their corresponding discharges, which are also macaroons. An initial setup creates a discharging key sk_D and an issuing key sk_I . We define a token-issuing algorithm, which takes as input the issuing key sk_I and outputs two tokens - a state token α_S and a discharge token α_D . The discharge token α_D , together with the discharging key sk_D , is used to obtain a discharge macaroon. The state token α_S is used to verify the validity of the discharge obtained.

A third-party caveat includes both the state token α_S , as well as the discharge token α_D . The discharge token α_D can be used by the holder of the macaroon to obtain the discharge from the third-party who holds the discharging key sk_D , and the state token α_S can be used by the target service to verify the validity of the discharge presented.

Note that in the use-case of discharge protocols, a principal will present a macaroon and the corresponding discharges to a target service. The target service will then use the state tokens that are part of the third-party caveats in the macaroon and use these to verify the discharges that are presented along with the macaroon. Therefore, our definition of security of a discharge protocol ensures that without knowing the issuing or discharge keys, a principal cannot forge a state token α_S^* and a discharge macaroon μ^* such that μ^* is a valid discharge macaroon for α_S^* , unless they were legally obtained.

We present two constructions of discharge protocols, one that uses encryption and a privately-verifiable macaroon scheme, and another that uses a message authentication code (MAC) and a publicly-verifiable

macaroon scheme. Though the first was already outlined in [BPUE⁺14], we formalize the properties required to prove its security. In particular, we show that the encryption used must be non-malleable.

Macaroon-Tree Protocols. Finally, we show how to combine macaroon schemes and discharge protocols to obtain *macaroon-tree protocols*, which have the full functionality of macaroons described in [BPUE⁺14]. In particular, macaroons can be issued, extended with a first-party caveat, extended with a third-party caveat, discharged, finalized, and verified. The security definition of macaroon-tree protocols follows directly from combining the security definitions of macaroon schemes and discharge protocols.

Our construction of a macaroon-tree protocol “glues” together a macaroon scheme and a discharge protocol using symmetric-key and public-key encryption, and a message authentication code. The symmetric-key encryption and the MAC are used to create the key identifier of a macaroon: the key identifier is a symmetric-key encryption of the root key k , together with a tag of this ciphertext. The public-key encryption is used to encrypt the token issuing key of the discharge protocol under the verifying party’s public key. We note that this encryption can be removed in certain cases and discuss this optimization in Section 5. We further note that the construction is *generic*, so that *any* macaroon scheme and *any* discharge protocol can be combined. In particular, a privately verifiable macaroon scheme can be used with the discharge protocol that uses publicly-verifiable discharges, and vice versa.

Concrete vs. Asymptotic Security. The proofs of security in this work use asymptotics (e.g. $\text{poly}(\cdot)$ and $\text{negl}(\cdot)$ functions). Though they establish the *theoretical* security of the constructions, a more concrete analysis of probabilities that takes into account the computational power of the adversary and possible upper bounds for the number of queries, must be performed in specific use cases to guarantee their concrete security.

Organization. We begin by presenting a few preliminary definitions in Section 2. In Section 3, we present our definitions and constructions of macaroon schemes. In Section 4, we discuss discharge protocols, and finally, in Section 5, we discuss macaroon-tree protocols. For clarity of presentation, all proofs are presented in the Appendix.

2 Preliminaries

We assume basic knowledge of cryptographic primitives such as PRFs, MACs, signatures, and symmetric-key and public-key encryption. For completeness, we review their security definitions below.

Notation. For a randomized function f , we write $f(x; r)$ to denote the unique output of f on input x with random coins r . We write $f(x)$ to denote a random variable for the output of $f(x; r)$ over the random coins r . For a distribution or random variable X , we write $x \leftarrow X$ to denote the operation of sampling a random x according to X . For a set S , we overload notation and use $s \leftarrow S$ to denote sampling s from the uniform distribution over S . We use $y := f(x)$ to denote the deterministic evaluation of f on input x with output y , and in general, as the assignment operator. Finally, we use $a||b$ to denote the concatenation of bit strings a and b , and λ to denote the empty string.

2.1 Pseudorandom Functions (PRFs)

A *pseudorandom function* or *PRF* is a function that is indistinguishable from a truly random function. A more formal definition follows.

Definition 2.1. A function $F : \{0, 1\}^n \times \{0, 1\}^B \rightarrow \{0, 1\}^n$ is pseudorandom if for all PPT distinguishers \mathcal{A} ,

$$\left| \Pr[\mathcal{A}^{F_k(\cdot)}(1^n) = 1] - \Pr[\mathcal{A}^{f(\cdot)}(1^n) = 1] \right| \leq \text{negl}(n),$$

where $k \leftarrow \{0, 1\}^n$ and f is chosen uniformly at random from the set of functions mapping $\{0, 1\}^B$ to $\{0, 1\}^n$.

Note that the oracle $f(\cdot)$ can be simulated efficiently by answering every query with an independent and uniformly random value in $\{0, 1\}^n$.

Definition 2.2. Let $F : \{0, 1\}^n \times \{0, 1\}^B \rightarrow \{0, 1\}^n$. Define the functions $F^{(1)}, F^{(2)}, \dots$ inductively as follows:

$$F^{(1)}(k, x_1) = F(k, x_1) \quad \text{and} \quad F^{(n)}(k, x_1, \dots, x_n) := F(F^{(n-1)}(k, x_1, \dots, x_{n-1}), x_n)$$

Define $F^*(k, x_1, \dots, x_n) := F^{(n)}(k, x_1, \dots, x_n)$.

The above construction is called the *cascade construction*. Bellare, Canetti, and Krawczyk [BCK96] showed that if F is a PRF, then F^* is pseudorandom against prefix-free adversaries, that is, against adversaries whose query set is a prefix-free set.

Theorem 2.3 ([BCK96]). *If F is a PRF, then F^* described above is pseudorandom against distinguishers whose query set is prefix-free.*

2.2 Message Authentication Codes (MAC)

Definition 2.4. A message authentication code (MAC) $\mathcal{M} = (\text{KeyGen}, \text{Mac}, \text{Verify})$ is existentially-unforgeable under a chosen message attack (EU-CMA-secure) if for all PPT adversaries \mathcal{A} , the probability that \mathcal{A} wins the following game is negligible in κ .

Setup: The challenger generates a key $k \leftarrow \text{KeyGen}(1^\kappa)$.

Tag Queries: The adversary \mathcal{A} is given access to a tag oracle. On its i th query, the adversary sends a message $m^{(i)}$, the challenger computes $t^{(i)} \leftarrow \text{Mac}(k, m^{(i)})$ and sends $t^{(i)}$ as response.

Winning Condition: The adversary outputs a message and forgery m^*, t^* , and he wins the game if t^* is a valid tag for m^* under key k , that is, if $\text{Verify}(k, m^*, t^*) = 1$.

2.3 Signatures

Definition 2.5. A signature scheme $\mathcal{S} = (\text{KeyGen}, \text{Sign}, \text{Verify})$ is existentially-unforgeable under a chosen message attack (EU-CMA-secure) if for all PPT adversaries \mathcal{A} , the probability that \mathcal{A} wins the following game is negligible in κ .

Setup: The challenger generates a key pair $(\text{vk}, \text{sk}) \leftarrow \text{KeyGen}(1^\kappa)$.

Tag Queries: The adversary \mathcal{A} is given access to a signature oracle. On its i th query, the adversary sends a message $m^{(i)}$, the challenger computes $\sigma^{(i)} \leftarrow \text{Sign}(\text{sk}, m^{(i)})$ and sends $\sigma^{(i)}$ as response.

Winning Condition: The adversary outputs a message and forgery m^*, σ^* , and he wins the game if σ^* is a valid signature for m^* under key vk , that is, if $\text{Verify}(\text{vk}, m^*, \sigma^*) = 1$.

2.4 CPA-Secure Symmetric-Key Encryption

Definition 2.6. A symmetric-key encryption scheme $\mathcal{E} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ has indistinguishable ciphertexts under a chosen plaintext attack (is CPA-secure) if for all PPT adversaries \mathcal{A} , the adversary's advantage in winning the following game is negligible in κ .

Setup: The challenger generates a key $\text{sk} \leftarrow \text{KeyGen}(1^\kappa)$.

Encryption Queries: The adversary \mathcal{A} is given access to an encryption oracle. On its i th query, the adversary sends a message $m^{(i)}$, the challenger computes $c^{(i)} \leftarrow \text{Enc}(\text{sk}, m^{(i)})$ and sends $c^{(i)}$ as response.

Challenge: The adversary picks two challenge plaintexts, m_0, m_1 , and sends them to the challenger. The challenger picks a bit at random, $b \leftarrow \{0, 1\}$, computes $c^* \leftarrow \text{Enc}(\text{sk}, m_b)$, and sends c^* to \mathcal{A} .

Encryption Queries: The adversary can make more encryption queries.

Winning Condition: The adversary outputs a bit b^* and wins if $b^* = b$. We define the advantage of \mathcal{A} in this game to be: $|\Pr[\mathcal{A} \text{ wins}] - 1/2|$.

2.5 CCA2-Secure Public-Key Encryption

Definition 2.7. A public-key encryption scheme $\mathcal{E} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ has indistinguishable ciphertexts under a chosen ciphertext attack (is CCA2-secure) if for all PPT adversaries \mathcal{A} , the adversary's advantage in winning the following game is negligible in κ .

Setup: The challenger generates a key pair $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\kappa)$.

Decryption Queries: The adversary \mathcal{A} is given access to a decryption oracle. On its i th query, the adversary sends a ciphertext $c^{(i)}$, the challenger computes $m^{(i)} \leftarrow \text{Dec}(\text{sk}, c^{(i)})$ and sends $m^{(i)}$ as response.

Challenge: The adversary picks two challenge plaintexts, m_0, m_1 , and sends them to the challenger. The challenger picks a bit at random, $b \leftarrow \{0, 1\}$, computes $c^* \leftarrow \text{Enc}(\text{sk}, m_b)$, and sends c^* to \mathcal{A} .

Decryption Queries: The adversary can make more decryption queries, as long as the query ciphertext $c^{(i)}$ is not the challenge ciphertext c^* , i.e. as long as $c^{(i)} \neq c^*$.

Winning Condition: The adversary outputs a bit b^* , and wins if $b^* = b$. We define the advantage of \mathcal{A} in this game to be: $|\Pr[\mathcal{A} \text{ wins}] - 1/2|$.

3 Macaroons

A *macaroon* is a bearer token that acts as an authorization credential and includes *caveats* that attenuate and contextually confine when, where, by who, and for what purpose authorization should be granted. We model each caveat by a bit-string $m_i \in \{0, 1\}^*$. Additionally, a macaroon includes a key identifier id that opaquely encodes a root key k , and a signature σ that guarantees the integrity of the macaroon and can be verified using the key k .

Definition 3.1. A macaroon consists of a key identifier id , an optional tuple of messages or caveats (m_1, \dots, m_n) for some $n \geq 1$, and a signature σ . We use the notation $\mu := (\text{id} \mid m_1, \dots, m_n \mid \sigma)$ to denote a macaroon μ , or $\mu := (\text{id} \mid \mid \sigma)$ if the macaroon does not have any caveats.

3.1 Macaroon Schemes

We define a new primitive called a *macaroon scheme* that defines how macaroons can be issued, extended (for delegation or attenuation), finalized, and verified. A key generation algorithm generates root keys k , with which a macaroon can be issued. When a macaroon is issued, an *extension key* is also given. Using this extension key, the new macaroon can be extended or delegated, essentially adding more caveats to it, and a new extension key is issued so as to allow further delegation. A verification algorithm is also defined such that given the corresponding root key k , the validity of a macaroon μ can be verified. Finally, an algorithm `Finalize` is defined, that modifies a given macaroon so that it can be verified by the verification algorithm. This last algorithm can be seen as optional – indeed, only one of our constructions requires it.

We make a distinction between a privately-verifiable macaroon scheme (as informally defined above), where a macaroon is issued and verified using the same root key k , and a publicly-verifiable macaroon scheme, where a macaroon is issued with a secret root key sk and verified with a corresponding verification key vk . We formally define privately-verifiable macaroon schemes below, present constructions, and discuss their security in this section. See Section 3.2 for a formal discussion of publicly-verifiable macaroon schemes.

Definition 3.2. A privately-verifiable macaroon scheme is a tuple of algorithms $\mathcal{M} = (\text{Setup}, \text{Issue}, \text{Extend}, \text{Finalize}, \text{Verify})$ with the following syntax:

- $k \leftarrow \text{KeyGen}(1^\kappa)$: Given a security parameter κ , outputs a root key k .
- $(\mu, \text{ek}) \leftarrow \text{Issue}(k, \text{id})$: Given a root key k and an identifier id , outputs a macaroon $\mu = (\text{id} \mid \mid \sigma)$ with no caveats, and an extension key ek .
- $(\mu', \text{ek}) \leftarrow \text{Extend}(\text{ek}, \mu, m)$: Given an extension key ek , a macaroon $\mu = (\text{id} \mid m_1, \dots, m_n \mid \sigma)$ and a message $m \in \{0, 1\}^*$, outputs a new macaroon $\mu' := (\text{id}' \mid m_1, \dots, m_n, m \mid \sigma')$ with key identifier id' such that id is a prefix of id' , and whose caveat list is the same as that of μ , with m appended at the end. In particular, it should be easily determined if a macaroon μ' is an extension of a macaroon μ .

- $\mu' \leftarrow \text{Finalize}(\text{ek}, \mu)$: Given an extension key ek and a macaroon $\mu = (\text{id} \mid m_1, \dots, m_n \mid \sigma)$, outputs a new macaroon $\mu' = (\text{id} \mid m_1, \dots, m_n \mid \sigma')$ with the same key identifier and the same caveat list as μ .
- $b := \text{Verify}(k, \mu)$: Given a root key k and a macaroon μ , outputs a bit b .

We require that these algorithms satisfy correctness: for all $\text{id} \in \{0, 1\}^*$, for all $n \geq 1$ and all $m_1, \dots, m_n \in \{0, 1\}^*$, if $k \leftarrow \text{KeyGen}(1^\kappa)$, $(\mu_0, \text{ek}_0) \leftarrow \text{Issue}(k, \text{id})$, $(\mu_i, \text{ek}_i) \leftarrow \text{Extend}(\text{ek}_{i-1}, \mu_{i-1}, m_i)$ for $i \in [n]$, $\mu_{\text{final}} \leftarrow \text{Finalize}(\text{ek}_n, \mu_n)$, then $\text{Verify}(k, \mu_{\text{final}}) = 1$ with probability 1.

The above definition outlines the syntax of a macaroon scheme; we now turn to defining its security. The desired security guarantee for a macaroon scheme is that a principal cannot *forge* a macaroon. In other words, a principal cannot create a valid macaroon μ^* with respect to a root key k unless it obtained the macaroon via “valid means”, i.e. μ^* was given to him, and μ^* was issued using k or is an extension of such a macaroon. Notice that this includes removing caveats from a macaroon obtained legally.

We wish to guarantee security even against colluding principals and as such define a global “adversary” that might obtain valid macaroons and whose goal is to forge a *new* macaroon. This has a similar flavor to the security definition of MACs or signatures, and so we follow the same structure, allowing the macaroon equivalent of “signature queries”. In our case, however, macaroons can be obtained from issuing a new macaroon or extending an existing one. Therefore, we allow the adversary to make both *issuing queries* and *extension queries*.

Finally, we must define a “valid forgery”. For signatures, a valid forgery means a signature of message that was not given as part of a signature query. For macaroons, we must extend this to mean a macaroon that was not a result of a query, or an extension of such a macaroon. This is because the extension key can be either implicit in the macaroon (as in our first construction) or empty (as in our second construction). Therefore, obtaining the macaroon itself is enough to extend or delegate it. To this end, we have the challenger keep track of the macaroons that have been issued and extended in queries, and have the adversary explicitly ask to receive these macaroons. This allows a valid forgery to include “parents” of macaroons that have been revealed to the adversary. In other words, this guarantees a breach of security if an adversary successfully removes caveats in macaroon that it obtained legally.

Definition 3.3. We say that a macaroon scheme $\mathcal{M} = (\text{KeyGen}, \text{Issue}, \text{Extend}, \text{Finalize}, \text{Verify})$ is secure if for all PPT adversaries \mathcal{A} , the probability that \mathcal{A} wins in the following game is negligible in κ .

Setup: The challenger generates a key $k \leftarrow \text{KeyGen}(1^\kappa)$. The challenger keeps a tree of macaroons \mathcal{T} , which has macaroon/extension key pairs as nodes and ids/caveats as edges. As a general rule, a macaroon in \mathcal{T} whose path from the root is (id, \vec{m}) has the form $(\tilde{\text{id}} \mid \vec{m} \mid \sigma)$ for some signature σ and key identifier $\tilde{\text{id}} = \text{id} \parallel \dots$ that has id as prefix. Note that its caveats are exactly \vec{m} . The challenger initializes \mathcal{T} to contain the empty string, λ , as root.

Issue Queries: On the i th query, the adversary chooses an identifier $\text{id}^{(i)}$. The challenger creates a macaroon $(\mu_0^{(i)}, \text{ek}_0^{(i)}) \leftarrow \text{Issue}(k, \text{id}^{(i)})$. It adds $(\mu_0^{(i)}, \text{ek}_0^{(i)})$ as a child of the root under an edge labeled with $\text{id}^{(i)}$.

Extend Queries: On the i th query, the adversary sends an identifier $\text{id}^{(i)}$, a tuple of caveats $\vec{m}^{(i)}$, and a message $m^{(i)} \in \{0, 1\}^*$ to the challenger. If $(\text{id}^{(i)}, \vec{m}^{(i)})$ is a valid path in \mathcal{T} leading to a node (μ, ek) , then the challenger computes $(\mu^{(i)}, \text{ek}^{(i)}) \leftarrow \text{Extend}(\text{ek}, \mu, m^{(i)})$, and adds $(\mu^{(i)}, \text{ek}^{(i)})$ to \mathcal{T} as a child of (μ, ek) under an edge labeled with $m^{(i)}$.

Macaroon Queries: On the i th query, the adversary sends an identifier $\text{id}^{(i)}$ and a tuple of caveats $\vec{m}^{(i)}$ to the challenger. If $(\text{id}^{(i)}, \vec{m}^{(i)})$ is a valid path in \mathcal{T} leading to a node (μ, ek) , then the challenger returns (μ, ek) to the adversary.

Winning Condition: The adversary \mathcal{A} outputs a macaroon $\mu^* := (\tilde{\text{id}}^* \parallel \vec{m}^* \parallel \sigma^*)$ with $\tilde{\text{id}}^* := \text{id}^* \parallel \dots$. He wins if $\text{Verify}(k, \mu^*) = 1$ and μ^* is not an extension of a macaroon given to \mathcal{A} in a macaroon query. This can be easily checked by checking if (id^*, \vec{m}^*) leads down a path that includes a macaroon that was given to \mathcal{A} in a macaroon query, or equivalently, if the caveat list of any of the macaroons given to \mathcal{A} in a macaroon query with identifier id^* , is a prefix of μ^* 's caveat list.

3.1.1 Construction 1

Below we show how to construct a macaroon scheme from a variable-input PRF. This construction is precisely the MAC-based construction described in [BPUE⁺14], where the key identifier and the caveats are MAC'ed in a cascade fashion, i.e. the root key k is used to create a signature, σ_0 , of the key identifier id ; σ_0 is then used as a key to create a signature, σ_1 , of the first caveat m_1 , and so on. Here, we formalize the properties required from the ‘‘MAC’’ function, and show that using a variable-input PRF suffices. Intuitively, the pseudorandomness of the output guarantees that it is computationally indistinguishable from a valid and fresh key, and the unpredictability of the output ensures unforgeability.

Let $F : \{0, 1\}^n \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a variable-input PRF. We define:

- $\text{KeyGen}(1^\kappa) : \text{Output } k \leftarrow \{0, 1\}^n$.
- $(\mu, \text{ek}) \leftarrow \text{Issue}(k, \text{id}) : \text{Output } \mu := (\text{id} \mid \mid \sigma)$ and $\text{ek} := \sigma$, where $\sigma := F(k, \text{id})$.
- $(\mu', \text{ek}') \leftarrow \text{Extend}(\text{ek}, \mu, m) : \text{Parse } \mu = (\text{id} \mid m_1, \dots, m_n \mid \sigma)$. Output $\mu' := (\text{id} \mid m_1, \dots, m_n, m \mid \sigma')$ and $\text{ek}' := \sigma'$, where $\sigma' := F(\sigma, m)$.
- $\mu' \leftarrow \text{Finalize}(\text{ek}, \mu) : \text{Output } \mu' := \mu$ (i.e. don't do anything)
- $b := \text{Verify}(k, \mu) : \text{Parse } \mu = (\text{id} \mid m_1, \dots, m_n \mid \sigma)$. Compute $\sigma_0 := F(k, \text{id})$. For $i = 1, \dots, n$, compute $\sigma_i := F(k, \sigma_i)$. Output 1 if $\sigma_n = \sigma$, and 0 otherwise.

Theorem 3.4. . *Let $F : \{0, 1\}^n \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a variable-input PRF. Then the above construction is a secure macaroon scheme.*

To prove Theorem 3.4, we must assume that if an adversary will ask a macaroon queries id, \vec{m}_1 and id, \vec{m}_2 such that \vec{m}_1 is a prefix of \vec{m}_2 , then it will ask the query id, \vec{m}_1 *before* the query id, \vec{m}_2 . Even though making this assumption weakens the theorem statement, it might not be an unreasonable assumption to make in the setting of macaroons. This is because the macaroon μ_2 corresponding to id, \vec{m}_2 is an extension of the macaroon μ_1 corresponding to id, \vec{m}_1 , and therefore μ_1 must be computed before μ_2 . In assuming that the adversary learns μ_1 before μ_2 , we are assuming *static* (vs. adaptive) corruptions of the parties, or in other words, we assume that if a party is corrupted at any point in time then it was corrupted from the beginning.

3.1.2 Construction 2

We now show a new construction of a privately-verifiable macaroon scheme. To issue a macaroon, one simply computes a public-key signature of the key identifier. Extending a macaroon requires sampling a new signature key pair, and signing the previous signature, together with the new verification key, and the new caveat. However, notice that this simple construction is insecure, as it allows an adversary to obtain a valid macaroon and remove its caveats, by simply removing the corresponding signatures. To safe-guard against this, we encrypt the intermediate signatures using a public-key encryption scheme. The macaroon can then be verified using the original verification key (corresponding to the secret key used to sign the key identifier), and the decryption key.

Let $\mathcal{S} = (\text{KeyGen}_{\mathcal{S}}, \text{Sign}, \text{Verify})$ be a signature scheme, and let $\mathcal{E} = (\text{KeyGen}_{\mathcal{E}}, \text{Enc}, \text{Dec})$ be a CPA-secure public-key encryption encryption scheme. We define:

- $\text{KeyGen}(1^\kappa) : \text{Output } k := (\text{vk}_{\mathcal{S}}, \text{pk}_{\mathcal{E}}, \text{sk}_{\mathcal{S}}, \text{sk}_{\mathcal{E}})$, where

$$(\text{vk}_{\mathcal{S}}, \text{sk}_{\mathcal{S}}) \leftarrow \text{KeyGen}_{\mathcal{S}}(1^\kappa) \quad , \quad (\text{pk}_{\mathcal{E}}, \text{sk}_{\mathcal{E}}) \leftarrow \text{KeyGen}_{\mathcal{E}}(1^\kappa)$$

- $(\mu, \text{ek}) \leftarrow \text{Issue}(k, \text{id}) : \text{Parse } k := (\text{vk}_{\mathcal{S}}, \text{pk}_{\mathcal{E}}, \text{sk}_{\mathcal{S}}, \text{sk}_{\mathcal{E}})$. Output $\mu = (\tilde{\text{id}} \mid \mid \sigma)$, where

$$\tilde{\text{id}} := \text{id} \parallel \text{pk}_{\mathcal{E}} \quad , \quad \sigma \leftarrow \text{Sign}(\text{sk}_{\mathcal{S}}, \tilde{\text{id}}) \quad , \quad \text{ek} := \lambda,$$

and λ is the empty string.

- $(\mu', \text{ek}') \leftarrow \text{Extend}(\text{ek}, \mu, m) : \text{Parse } \mu = (\tilde{\text{id}} \mid m_1, \dots, m_n \mid \sigma)$ and parse $\tilde{\text{id}} = \text{id} \parallel \text{pk}_{\mathcal{E}} \parallel \dots$. Output $\mu' := (\tilde{\text{id}}' \mid m_1, \dots, m_n, m \mid \sigma')$, where:

$(\text{vk}', \text{sk}') \leftarrow \text{KeyGen}_{\mathcal{S}}(1^\kappa)$, $c \leftarrow \text{Enc}(\text{pk}_{\mathcal{E}}, \sigma)$, $\tilde{\text{id}}' := \tilde{\text{id}} \parallel (\text{vk}', c)$, $\sigma' \leftarrow \text{Sign}(\text{sk}', \text{vk}' \parallel c \parallel m)$, $\text{ek} := \lambda$
and λ is the empty string.

- $\mu' \leftarrow \text{Finalize}(\text{ek}, \mu) : \text{Output } \mu' := \mu$ (i.e. don't do anything)
- $b := \text{Verify}(k, \mu) : \text{Parse } \mu = (\tilde{\text{id}} \mid m_1, \dots, m_n \mid \sigma)$ and parse $\tilde{\text{id}} = \text{id} \parallel \text{pk}_{\mathcal{E}} \parallel (\text{vk}_1, c_1) \parallel \dots \parallel (\text{vk}_n, c_n)$. For $i = 0, \dots, n-1$, compute $\sigma_i := \text{Dec}(\text{sk}_{\mathcal{E}}, c_{i+1})$, and let $\sigma_n := \sigma$. Check that $\text{Verify}(\text{vk}_{\mathcal{S}}, \text{id} \parallel \text{pk}_{\mathcal{E}}, \sigma_0) = 1$ and $\text{Verify}(\text{vk}_i, \text{vk}_i \parallel c_i \parallel m_i, \sigma_i) = 1$ for $i = 1, \dots, n$.

Theorem 3.5. *Let $\mathcal{S} = (\text{KeyGen}_{\mathcal{S}}, \text{Sign}_{\mathcal{S}}, \text{Verify}_{\mathcal{S}})$ be an EU-CMA-secure signature scheme, and let $\mathcal{E} = (\text{KeyGen}_{\mathcal{E}}, \text{Enc}, \text{Dec})$ be a CPA-secure public-key encryption scheme. Then the above construction is a secure macaroon scheme.*

TRADE-OFFS. Though Construction 2 is considerably less efficient than Construction 1 (as it uses public-key primitives), it has one advantage. Notice that in Construction 1, the process of verifying a macaroon reveals all intermediate signatures, which double as extension keys in that construction. As observed in [BPUE⁺14], this can be a problem when used within macaroon-tree protocols, as is our goal (see Section 5). Construction 2 does not use extension keys, and therefore does not suffer from this weakness.

3.2 Publicly-Verifiable Macaroon Schemes

Similarly to privately-verifiable macaroons schemes, we define publicly-verifiable macaroon schemes. The only difference is that a macaroon is issued with a secret root key sk and verified with a corresponding public verification key vk .

Definition 3.6. *A public-key macaroon scheme is a tuple of algorithms $\mathcal{M} = (\text{Setup}, \text{Issue}, \text{Extend}, \text{Verify})$ with the following syntax:*

- $(\text{vk}, \text{sk}) \leftarrow \text{KeyGen}(1^\kappa) : \text{Given a security parameter, outputs a verification key vk and a secret key sk.}$
- $(\mu, \text{ek}) \leftarrow \text{Issue}(\text{sk}, \text{id}) : \text{Given a secret key sk and an identifier id, outputs a macaroon } \mu = (\text{id} \mid \mid \sigma)$ with no caveats, and an extension key ek .
- $(\mu', \text{ek}) \leftarrow \text{Extend}(\text{ek}, \mu, m) : \text{Given an extension key ek, a macaroon } \mu = (\text{id} \mid m_1, \dots, m_n \mid \sigma)$ and a message $m \in \{0, 1\}^*$, outputs a new macaroon $\mu' := (\text{id}' \mid m_1, \dots, m_n, m \mid \sigma')$ with key identifier id' such that id is a prefix of id' , and whose caveat list is the same as that of μ , with m appended at the end. In particular, it should be easily determined if a macaroon μ' is an extension of a macaroon μ .
- $\mu' \leftarrow \text{Finalize}(\text{ek}, \mu) : \text{Given an extension key ek and a macaroon } \mu = (\text{id} \mid m_1, \dots, m_n \mid \sigma)$, outputs a new macaroon $\mu' = (\text{id} \mid m_1, \dots, m_n \mid \sigma')$ with the same key identifier and the same caveat list as μ .
- $b := \text{Verify}(\text{vk}, \mu) : \text{Given a verification key vk and a macaroon } \mu$, outputs a bit b .

We require that these algorithms satisfy correctness: for all $\text{id} \in \{0, 1\}^$, for all $n \geq 1$ and all $m_1, \dots, m_n \in \{0, 1\}^*$, if $(\text{vk}, \text{sk}) \leftarrow \text{KeyGen}(1^\kappa)$, $(\mu_0, \text{ek}_0) \leftarrow \text{Issue}(\text{sk}, \text{id})$, $(\mu_i, \text{ek}_i) \leftarrow \text{Extend}(\text{ek}_{i-1}, \mu_{i-1}, m_i)$ for $i \in [n]$, $\mu_{\text{final}} \leftarrow \text{Finalize}(\text{ek}_n, \mu_n)$, then $\text{Verify}(\text{vk}, \mu_{\text{final}}) = 1$ with probability 1.*

The definition of security of a public-key macaroon scheme is the analogous to that of a symmetric-key macaroon scheme. The only differences are that the adversary is also given the verification key vk , macaroons are issued with the secret key sk , and the adversary wins if μ^* verifies under the verification key vk .

3.2.1 Construction

Below we show how to construct a publicly-verifiable macaroon scheme from a public-key signature scheme. This construction is a formalization of the certificate-based construction described in [BPUE⁺14]. Intuitively, in issuing and extending macaroons, one constructs two “strands” of signatures, denoted here with

ν and τ . The first is a certificate chain that authenticates the second, and the second is used to sign the key identifier and each of the caveats.

Note that the `Finalize` algorithm is crucial in this construction, as it ensures knowledge of the last signing key in the chain. This ensures that caveats cannot be removed from a legally-obtained macaroon, except by the principal who created the macaroon in the first place.

Let $\mathcal{S} = (\text{KeyGen}_{\mathcal{S}}, \text{Sign}_{\mathcal{S}}, \text{Verify}_{\mathcal{S}})$ be a signature scheme. We define:

- $\text{KeyGen}(1^\kappa) : \text{Output } (\text{vk}, \text{sk}) \leftarrow \text{KeyGen}_{\mathcal{S}}(1^\kappa)$.
- $(\mu, \text{ek}) \leftarrow \text{Issue}(\text{sk}, \text{id}) : \text{Output } \mu := (\tilde{\text{id}} \mid \mid \sigma)$ and $\text{ek} = \text{sk}'$, where:

$$(\text{vk}', \text{sk}') \leftarrow \text{KeyGen}_{\mathcal{S}}(1^\kappa) \quad , \quad \nu \leftarrow \text{Sign}_{\mathcal{S}}(\text{sk}, \text{vk}') \quad , \quad \tau \leftarrow \text{Sign}_{\mathcal{S}}(\text{sk}', \text{id}) \quad , \quad \sigma := (\nu, \tau) \quad , \quad \tilde{\text{id}} := \text{id} \parallel \text{vk}'$$
- $(\mu', \text{ek}') \leftarrow \text{Extend}(\text{ek}, \mu, m) : \text{Parse } \mu := (\tilde{\text{id}} \mid m_1, \dots, m_n \mid \sigma)$ and $\tilde{\text{id}} = \text{id} \parallel \dots$. Output \perp if $m = \text{id}$. Otherwise, output $\mu' := (\tilde{\text{id}}' \mid m_1, \dots, m_n, m \mid \sigma')$ and $\text{ek}' = \text{sk}'$, where:

$$(\text{vk}', \text{sk}') \leftarrow \text{KeyGen}_{\mathcal{S}}(1^\kappa) \quad , \quad \nu \leftarrow \text{Sign}_{\mathcal{S}}(\text{ek}, \text{vk}') \quad , \quad \tau \leftarrow \text{Sign}_{\mathcal{S}}(\text{sk}', m) \quad , \quad \sigma' := \sigma \parallel (\nu, \tau) \quad , \quad \tilde{\text{id}}' := \tilde{\text{id}} \parallel \text{vk}'$$
- $\mu' \leftarrow \text{Finalize}(\text{ek}, \mu) : \text{Parse } \mu := (\tilde{\text{id}} \mid m_1, \dots, m_n \mid \sigma)$. Output $\mu' := (\tilde{\text{id}} \mid m_1, \dots, m_n, m \mid \sigma')$, where $\sigma' := \sigma \parallel \sigma_{\text{final}}$ and $\sigma_{\text{final}} \leftarrow \text{Sign}_{\mathcal{S}}(\text{ek}, \text{id})$.
- $b := \text{Verify}(\text{vk}, \mu) : \text{Parse } \mu := (\tilde{\text{id}} \mid m_1, \dots, m_n \mid \sigma)$, parse $\tilde{\text{id}} = \text{id} \parallel \text{vk}_1 \parallel \dots \parallel \text{vk}_{n+1}$, and parse $\sigma = (\nu_0, \tau_0) \parallel \dots \parallel (\nu_n, \tau_n) \parallel \sigma_{\text{final}}$. For simplicity, let $\text{vk}_0 := \text{vk}$ and $m_0 := \text{id}$. For $i = 0, \dots, n$, check that $\text{Verify}_{\mathcal{S}}(\text{vk}_i, \text{vk}_{i+1}, \nu_i) = 1$ and $\text{Verify}_{\mathcal{S}}(\text{vk}_{i+1}, m_i, \tau_i) = 1$. Finally, check that $\text{Verify}_{\mathcal{S}}(\text{vk}_{n+1}, \text{id}, \sigma_{\text{final}}) = 1$.

Theorem 3.7. *Let $\mathcal{S} = (\text{KeyGen}_{\mathcal{S}}, \text{Sign}_{\mathcal{S}}, \text{Verify}_{\mathcal{S}})$ be an EU-CMA-secure signature scheme. Then the above construction is a secure macaroon scheme.*

4 Discharge Protocols

Having defined and constructed macaroon schemes, we now turn to studying the “discharge” of so-called “third-party” caveats. Third-party caveats require a third-party to assert that a given condition holds true. As an example (previously used in [BPUE⁺14]), a third-party caveat can require proof of membership in a group G at service A . Before presenting a macaroon to a target service for verification, a holder of a macaroon must obtain “discharges” from the appropriate third-parties, asserting that the conditions specified in the third-party caveats holds true. These discharges are also macaroons.

In what follows, we formally define *discharge protocols*, specifying the syntax of third-party caveats and their corresponding discharges. After an initial discharge setup that creates a discharging key sk_{D} and an initial issuing setup that creates an issuing secret key sk_{I} , tokens can be issued using the issuing key and discharged using the discharging key. Every invocation to the token-issuing algorithm outputs two tokens - a state token α_{S} and a discharge token α_{D} . The discharge token α_{D} is used to obtain a discharge macaroon and the state token α_{S} is used to verify the validity of the discharge obtained.

Looking ahead, a third-party caveat will include both the state token α_{S} , as well as the discharge token α_{D} . The discharge token α_{D} can be used by the holder of the macaroon to obtain the discharge from the discharging third-party who holds the discharging key sk_{D} , and the state token α_{S} can be used by the target service to verify the validity of the discharge presented. See Section 5 for more details.

Definition 4.1. *A macaroon discharge protocol is a tuple of algorithms $\mathcal{D} = (\text{Setup}, \text{Token}, \text{Discharge}, \text{Verify})$ with the following syntax:*

- $(\text{params}, \text{sk}_{\text{D}}) \leftarrow \text{SetupDischarge}(1^\kappa) : \text{Given a security parameter } \kappa, \text{ outputs public parameters } \text{params} \text{ and a discharging secret key } \text{sk}_{\text{D}}. \text{ All other algorithms take } \text{params} \text{ as argument, even if not explicitly written.}$
- $\text{sk}_{\text{T}} \leftarrow \text{SetupToken}(1^\kappa) : \text{Given a security parameter } \kappa, \text{ outputs a token-issuing secret key } \text{sk}_{\text{T}}.$

- $(\alpha_S, \alpha_D) \leftarrow \text{Token}(\text{sk}_T)$: Given the token-issuing secret key sk_T , outputs a state token α_S and a discharge token α_D .
- $\mu \leftarrow \text{Discharge}(\text{sk}_D, \alpha_D)$: Given a discharge token α_D and the discharging secret key sk_D , outputs a macaroon μ .
- $b := \text{Verify}(\text{sk}_T, \alpha_S, \mu)$: Given a macaroon μ , a state token α_S , and the token-issuing secret key sk_T , outputs a bit b .

We require that these algorithms satisfy correctness: if $(\text{params}, \text{sk}_T, \text{sk}_D) \leftarrow \text{Setup}(1^\kappa)$, $(\alpha_S, \alpha_D) \leftarrow \text{Token}(\text{sk}_T)$, $\mu \leftarrow \text{Discharge}(\text{sk}_D, \alpha_D)$, then $\text{Verify}(\text{sk}_T, \alpha_S, \mu) = 1$ with probability 1.

Note that in the use-case of discharge protocols, a principal will present a macaroon and the corresponding discharges to a target service. The target service will then use the state tokens that are part of the third-party caveats in the macaroon and use these to verify the discharges that are presented along with the macaroon. Therefore, for the security of a discharge protocol, we wish to ensure that without knowing the issuing or discharge keys, sk_I and sk_D , a principal cannot forge a state token α_S^* and a discharge macaroon μ^* such that μ^* is a valid discharge macaroon for α_S^* , unless they were legally obtained.

To ensure this, in our definition we allow the adversary to obtain tokens and discharges via *token queries* and *discharge queries*, and only consider a forgery (α_S^*, μ^*) to be valid if μ^* is a valid discharge macaroon for α_S^* , and α_S^* and μ^* were not *both* legally obtained from token and discharge queries. That is, we allow α_S^* to be the output of a token query (α_S^*, α_D^*) , as long as μ^* was not the output of a discharge query using the corresponding token α_D^* . Thus, security is broken if a principal is able to obtain a valid discharge for a third-party caveat that it did not create (i.e. for a state token that it did not issue), without the help of the third-party who owns the discharging key sk_D . Furthermore, we consider security broken if a principal is able to forge a state token sk_S^* .

Definition 4.2. We say that a macaroon discharge protocol $\mathcal{D} = (\text{Setup}, \text{Issue}, \text{Discharge}, \text{Verify})$ is secure if for all PPT adversaries \mathcal{A} , the probability that \mathcal{A} wins in the following game is negligible in κ .

Setup: The challenger chooses $(\text{params}, \text{sk}_D) \leftarrow \text{SetupDischarge}(1^\kappa)$ and $\text{sk}_T \leftarrow \text{SetupToken}(1^\kappa)$ and gives params to \mathcal{A} . The challenger keeps track of a set of tokens \mathcal{Q}_T and a set of discharges \mathcal{Q}_D , both of which start out empty.

Token Queries: The adversary \mathcal{A} asks the challenger to issue a token. On the i th query, the challenger runs $(\alpha_S^{(i)}, \alpha_D^{(i)}) \leftarrow \text{Token}(\text{sk}_T)$ and returns $(\alpha_S^{(i)}, \alpha_D^{(i)})$ to \mathcal{A} . It also adds $(\alpha_S^{(i)}, \alpha_D^{(i)})$ to \mathcal{Q}_T . Issuing queries and discharge queries can be interleaved.

Discharge Queries: On the i th query, the adversary \mathcal{A} sends a token $\alpha_D^{(i)*}$ to the challenger, who runs $\mu^{(i)} \leftarrow \text{Discharge}(\text{sk}_D, \alpha_D^{(i)*})$ and returns $\mu^{(i)}$ to \mathcal{A} . The challenger also adds $(\alpha_D^{(i)*}, \mu^{(i)})$ to \mathcal{Q}_D . Issuing queries and discharge queries can be interleaved.

Winning condition: The adversary wins if it produces a token α_S^* and a macaroon μ^* such that μ^* is not an extension of any macaroon μ such that $(\alpha_D^*, \mu) \in \mathcal{Q}_D$, where α_D^* is a corresponding discharge token to α_S^* in \mathcal{Q}_T , and $\text{Verify}(\text{sk}_I, \alpha_S^*, \mu^*) = 1$. In other words, \mathcal{A} wins if it outputs a valid macaroon that is not an extension of a macaroon that was the output of a discharge query with token α_D^* corresponding to α_S^* .

4.1 Construction 1

We first present the encryption-based construction that was described in [BPUE⁺14]. The idea is for the state and discharge tokens to encode a key k , which will be the root key for the discharge macaroon. The state token is a symmetric-key authenticated encryption of k , whereas the discharge token encrypts k using the third-party's public key. Given a discharge token, the third-party obtains k by decrypting the token, and issues a discharge macaroon using k as root key. A verifying party first checks the integrity of a state token, decrypts it to obtain the root key k , and uses this key to verify the validity of the discharge macaroon.

Let $\mathcal{E}_1 := (\text{KeyGen}_1, \text{Enc}_1, \text{Dec}_1)$ be a CPA-secure symmetric-key encryption scheme, let $\mathcal{E}_2 := (\text{KeyGen}_2, \text{Enc}_2, \text{Dec}_2)$ be a CCA2-secure public-key encryption scheme, let $\mathcal{T} := (\text{KeyGen}_T, \text{Mac}, \text{Verify}_T)$ be a message authentication code, and let $\mathcal{M} = (\text{KeyGen}_M, \text{Issue}, \text{Extend}, \text{Verify}_M)$ be a macaroon scheme. We define:

- $\text{SetupDischarge}(1^\kappa)$: Output $\text{params} := \text{pk}_2$, $\text{sk}_D := \text{sk}_2$, where $(\text{pk}_2, \text{sk}_2) \leftarrow \text{KeyGen}_2(1^\kappa)$.
- $\text{SetupToken}(1^\kappa)$: Output $\text{sk}_T := (\text{sk}_1, k)$, where

$$\text{sk}_1 \leftarrow \text{KeyGen}_1(1^\kappa) \quad , \quad k \leftarrow \text{KeyGen}_T(1^\kappa)$$

- $\text{Token}(\text{sk}_T)$: Parse $\text{params} = \text{pk}_2$, $\text{sk}_T = (\text{sk}_1, k)$ and output $\alpha_S := (c_1, \sigma)$ and $\alpha_D := c_2$, where

$$r \leftarrow \text{KeyGen}_{\mathcal{M}}(1^\kappa) \quad , \quad c_1 \leftarrow \text{Enc}_1(\text{sk}_1, r) \quad , \quad c_2 \leftarrow \text{Enc}_2(\text{pk}_2, r) \quad , \quad \sigma \leftarrow \text{Mac}(k, c_1)$$

- $\text{Discharge}(\text{sk}_D, \alpha_D)$: Parse $\text{sk}_D = \text{sk}_2$, $\alpha_D = c_2$ and output $\beta = (c_2, t)$ where

$$r \leftarrow \text{Dec}_2(\text{sk}_2, c_2) \quad , \quad \mu \leftarrow \text{Issue}(r, c_2)$$

- $\text{Verify}(\text{sk}_T, \alpha_S, \mu)$: Parse $\alpha_D = (c_1, \sigma)$ and $\text{sk}_T = (\text{sk}_1, k)$. Check that $\text{Verify}_T(k, c_1, \sigma) = 1$. Let $r = \text{Dec}_1(\text{sk}_1, c_1)$ and check that $\text{Verify}_{\mathcal{M}}(r, \mu) = 1$. Output 0 if either of the checks fail; otherwise output 1.

Theorem 4.3. *Let $\mathcal{E}_1 := (\text{KeyGen}_1, \text{Enc}_1, \text{Dec}_1)$ be a CPA-secure symmetric-key encryption scheme, let $\mathcal{E}_2 := (\text{KeyGen}_2, \text{Enc}_2, \text{Dec}_2)$ be a CCA2-secure public-key encryption scheme, let $\mathcal{T} := (\text{KeyGen}_T, \text{Mac}, \text{Verify}_T)$ be a message authentication code, and let $\mathcal{M} = (\text{KeyGen}_{\mathcal{M}}, \text{Issue}, \text{Extend}, \text{Verify}_{\mathcal{M}})$ be a macaroon scheme. Then the above construction is a secure macaroon discharge protocol.*

NEED FOR CCA2-SECURE PKE. We highlight the fact that the proof of Theorem 4.3 would not go through if \mathcal{E}_2 is not CCA2-secure but only CPA (or even CCA1) secure. The reason for this is that we need the encryption to be non-malleable. If an adversary were able to tamper with a ciphertext c_2 and obtain a ciphertext c'_2 that encrypts a related plaintext, then he could submit c'_2 as a discharge query and obtain a macaroon under a related key. If the macaroon is implemented using a signature or MAC (as proposed in Section 3), then the adversary obtains a signature or tag under a related key. In short, using an encryption scheme that permits malleability could allow the adversary to mount a *related-key attack* against the underlying signature or MAC scheme of the macaroon scheme. See, e.g [PSW12], for related-key attacks on HMAC.

USING SYMMETRIC-KEY ENCRYPTION FOR \mathcal{E}_2 . It is possible, and possibly preferable for efficiency concerns, to have \mathcal{E}_2 be a symmetric-key encryption scheme instead of a public-key one. In this case, it is possible to achieve the required CCA2-security by encrypting the plaintext using a CPA-secure symmetric-key encryption scheme and concatenating a MAC tag of the ciphertext to the ciphertext. The new shared key includes both the encryption key and the MAC key.

4.2 Construction 2

We now present a new construction of a discharge protocol based on MACs and a publicly-verifiable macaroon scheme. In this construction, tokens are issued by first sampling a nonce r . The state token is a MAC of r , and the discharge token is simply r . The third-party “authenticates” itself by presenting a valid macaroon with r as key identifier.

Let $\mathcal{T} := (\text{KeyGen}_T, \text{Mac}, \text{Verify}_T)$ be a message authentication code, and let $\mathcal{M} = (\text{KeyGen}_{\mathcal{M}}, \text{Issue}, \text{Extend}, \text{Verify}_{\mathcal{M}})$ be a publicly-verifiable macaroon scheme. We define:

- $\text{SetupDischarge}(1^\kappa)$: Output $\text{params} := \text{vk}$ and $\text{sk}_D := \text{sk}$, where $(\text{vk}, \text{sk}) \leftarrow \text{KeyGen}_{\mathcal{M}}(1^\kappa)$.
- $\text{SetupToken}(1^\kappa)$: Output $\text{sk}_T := k \leftarrow \text{KeyGen}_T(1^\kappa)$.
- $\text{Token}(\text{sk}_T)$: Output $\alpha_S := (r, \sigma)$ and $\alpha_D := r$, where

$$r \leftarrow \{0, 1\}^\kappa \quad , \quad \sigma \leftarrow \text{Mac}(\text{sk}_T, r)$$

- $\text{Discharge}(\text{sk}_D, \alpha_D)$: Output $\mu \leftarrow \text{Issue}(\text{sk}_D, r)$.

- $\text{Verify}(\text{sk}_T, \alpha_S, \mu)$: Parse $\text{sk}_T = (\text{vk}, k)$ and $\alpha_S = (r, \sigma)$. Check that $\text{Verify}_T(k, r, \sigma) = 1$, μ 's key identifier is r , and $\text{Verify}_{\mathcal{M}}(\text{vk}, \mu) = 1$. Output 0 if either of the checks fail; otherwise output 1.

Theorem 4.4. *Let $\mathcal{S} := (\text{KeyGen}_{\mathcal{S}}, \text{Sign}, \text{Verify}_{\mathcal{S}})$ be an EU-CMA-secure signature scheme, let $\mathcal{T} := (\text{KeyGen}_{\mathcal{T}}, \text{Mac}, \text{Verify}_{\mathcal{T}})$ be a message authentication code, and let $\mathcal{M} = (\text{KeyGen}_{\mathcal{M}}, \text{Issue}, \text{Extend}, \text{Verify}_{\mathcal{M}})$ be a macaroon scheme. Then the above construction is a secure macaroon discharge protocol.*

5 Macaroon-Tree Protocols

In this section, we show how to combine macaroon schemes and discharge protocols to obtain *macaroon-tree protocols*, which have the full functionality of macaroons described in [BPUE⁺14]. In particular, macaroons can be issued, extended with a first-party caveat, extended with a third-party caveat, discharged, finalized, and verified. The security definition of macaroon-tree protocols also follows directly from combining the security definitions of macaroon schemes and discharge protocols.

Definition 5.1. *A macaroon-tree protocol is a tuple of algorithms $\mathcal{M} = (\text{Setup}, \text{Issue}, \text{Extend1stParty}, \text{Extend3rdParty}, \text{Verify})$ with the following syntax:*

- $(\text{params}, \text{sk}_I, \text{sk}_D) \leftarrow \text{Setup}(1^\kappa)$: Given a security parameter κ , outputs public parameters params , an issuing secret key sk_I and a discharge secret key sk_D . All other algorithms take params as argument, even if not explicitly written.
- $(k, \mu, \text{ek}) \leftarrow \text{Issue}(\text{sk}_I)$: Given an issuing secret key sk_I , outputs a key k and a macaroon μ .
- $(\mu', \text{ek}') \leftarrow \text{Extend1stParty}(\text{ek}, \mu, m)$: Given an extension key ek , a macaroon $\mu = (\text{id} \mid m_1, \dots, m_n \mid \sigma)$ and a message $m \in \{0, 1\}^*$, outputs a new macaroon $\mu' := (\text{id}' \mid m_1, \dots, m_n, m \mid \sigma')$ with key identifier id' such that id is a prefix of id' , and whose caveat list is the same as that of μ , with m appended at the end. In particular, it should be easily determined if a macaroon μ' is an extension of a macaroon μ .
- $(\mu', \text{ek}', \alpha_D) \leftarrow \text{Extend3rdParty}(\text{ek}, \mu)$: Given an extension key ek , a macaroon $\mu = (\text{id} \mid m_1, \dots, m_n \mid \sigma)$ and a message $m \in \{0, 1\}^*$, outputs a new macaroon $\mu' := (\text{id}' \mid m_1, \dots, m_n, m \mid \sigma')$ with key identifier id' such that id is a prefix of id' , and whose caveat list is the same as that of μ , with a new message m appended at the end. In particular, it should be easily determined if a macaroon μ' is an extension of a macaroon μ . This algorithm also outputs a discharge token α_D .
- $\mu' \leftarrow \text{Finalize}(\text{ek}, \mu)$: Given an extension key ek and a macaroon $\mu = (\text{id} \mid m_1, \dots, m_n \mid \sigma)$, outputs a new macaroon $\mu' = (\text{id} \mid m_1, \dots, m_n \mid \sigma')$ with the same key identifier and the same caveat list as μ .
- $\mu \leftarrow \text{Discharge}(\text{sk}_D, \alpha_D)$: Given a discharge token α_D , outputs a macaroon μ . The key identifier of μ must be α_D .
- $b := \text{Verify}(\text{sk}_I, \mu, M = \{\mu_j\})$: Given an issuing secret key sk_I , a macaroon μ , and a set of discharges $M = \{\mu_j\}$, outputs a bit b .

We require that these algorithms satisfy correctness. To state the correctness requirement, we first define the algorithm $(\mu', \text{ek}', \alpha) \leftarrow \text{Extend}(\text{ek}, \mu, m)$ that first checks if $m = \lambda$, where λ is the empty string. If it is, it outputs $\text{Extend3rdParty}(\text{ek}, \mu)$, otherwise it outputs $\text{Extend1stParty}(\text{ek}, \mu, m)$ and $\alpha = \lambda$. The correctness requirement is then stated as follows: if $(\text{params}, \text{sk}_I, \text{sk}_D) \leftarrow \text{Setup}(1^\kappa)$, $(k, \mu_0, \text{ek}_0) \leftarrow \text{Issue}(\text{sk}_I)$, $(\mu_i, \text{ek}_i, \alpha_i) \leftarrow \text{Extend}(\text{ek}_{i-1}, \mu_{i-1}, m_i)$ for $i \in [n]$, $\mu_{\text{final}} \leftarrow \text{Finalize}(\text{ek}_n, \mu_n)$, $\mu_j \leftarrow \text{Discharge}(\text{sk}_D, \alpha_j)$ for every j such that $\mu_j = \lambda$ and $M = \{\mu_j\}$, then $\text{Verify}(\text{sk}_I, \mu_{\text{final}}, M) = 1$ with probability 1.

Definition 5.2. *We say that a macaroon-tree protocol is secure if for all PPT adversaries \mathcal{A} , the probability that \mathcal{A} wins in the following game is negligible in κ .*

Setup: *The challenger chooses $(\text{params}, \text{sk}_I, \text{sk}_D) \leftarrow \text{Setup}(1^\kappa)$ and gives params to \mathcal{A} . The challenger keeps a tree of macaroons \mathcal{T} , which has macaroon/extension key pairs as nodes and ids/caveats as edges. As a general rule, a macaroon in \mathcal{T} whose path from the root is (id, \vec{m}) has the form $(\tilde{\text{id}} \mid \vec{m} \mid \sigma)$ for some signature σ and key identifier $\tilde{\text{id}} = \text{id} \parallel \dots$ that has id as prefix. Note that its caveats are exactly \vec{m} . The challenger initializes \mathcal{T} to contain the empty string, λ , as root. The challenger also keeps track of a set of discharges \mathcal{Q}_D , which starts out empty.*

Issue Queries: On the i th query, the challenger creates a macaroon $(\mu_0^{(i)}, \text{ek}_0^{(i)}) \leftarrow \text{Issue}(\text{sk}_1)$. It adds $(\mu_0^{(i)}, \text{ek}_0^{(i)})$ as a child of the root under an edge labeled with $\mu_0^{(i)}$'s key identifier. All types of queries can be interleaved.

Extend Queries: On the i th query, the adversary sends an identifier $\text{id}^{(i)}$, a tuple of caveats $\vec{m}^{(i)}$, and a message $m^{(i)} \in \{0, 1\}^* \cup \lambda$ to the challenger, where λ is the empty string. If $(\text{id}^{(i)}, \vec{m}^{(i)})$ is a valid path in \mathcal{T} leading to a node (μ, ek) , then the challenger computes $(\mu^{(i)}, \text{ek}^{(i)}, \alpha_D^{(i)}) \leftarrow \text{Extend}(\text{ek}, \mu, m^{(i)})$, and adds $(\mu^{(i)}, \text{ek}^{(i)}, \alpha_D^{(i)})$ to \mathcal{T} as a child of (μ, ek) under an edge labeled with $m^{(i)}$. Here, $\text{Extend}(\text{ek}, \mu, m)$ is defined as follows: If $m = \lambda$, it outputs $\text{Extend3rdParty}(\text{ek}, \mu)$, otherwise it outputs $\text{Extend1stParty}(\text{ek}, \mu, m)$ and $\alpha = \lambda$. All types of queries can be interleaved.

Discharge Queries: On the i th query, the adversary \mathcal{A} sends a token $\alpha_D^{(i)*}$ to the challenger, who runs $\mu^{(i)} \leftarrow \text{Discharge}(\text{sk}_D, \alpha_D^{(i)*})$ and returns $\mu^{(i)}$ to \mathcal{A} . The challenger also adds $(\alpha_D^{(i)*}, \mu^{(i)})$ to \mathcal{Q}_D . All types of queries can be interleaved.

Macaroon Queries: On the i th query, the adversary sends an identifier $\text{id}^{(i)}$ and a tuple of caveats $\vec{m}^{(i)}$ to the challenger. If $(\text{id}^{(i)}, \vec{m}^{(i)})$ is a valid path in \mathcal{T} leading to a node (μ, ek) , then the challenger returns (μ, ek) to the adversary. All types of queries can be interleaved.

Winning condition: The adversary \mathcal{A} outputs a macaroon μ^* and a set of discharges $M^* = \{\mu_j^*\}$. He wins if $\text{Verify}(\text{sk}_1, \mu^*, M^*) = 1$ and either (1). μ^* is not an extension of a macaroon given to \mathcal{A} in a macaroon query, or (2). μ^* is an extension of a macaroon μ given to \mathcal{A} in a macaroon query, and there exists a macaroon $m_j^* \in M^*$ corresponding to a third-party caveat $3||c_j||\alpha_S^{(j)}||\alpha_D^{(j)}$ in μ such that μ_j^* is not an extension of any macaroon μ' such that $(\alpha_D^{(j)}, \mu') \in \mathcal{Q}_D$. In other words, \mathcal{A} wins if it outputs a valid macaroon that is not an extension of a macaroon revealed to \mathcal{A} in a macaroon query, or if μ^* is an extension of a macaroon μ revealed to \mathcal{A} in a macaroon query and there is a discharge macaroon in M^* that corresponds to a third-party caveat in μ and that is not an extension of a macaroon that was the output of a discharge query for that caveat.

5.1 Construction

We describe a construction of a macaroon-tree protocol that “glues” together a macaroon scheme and a discharge protocol using symmetric-key and public-key encryption, and a message authentication code. The symmetric-key encryption and the MAC are used to create the key identifier of a macaroon; in other words, the key identifier of a macaroon is a symmetric-key encryption of the root key k , together with a tag of this ciphertext. Finally, the public-key encryption is used to encrypt the token issuing key of the discharge protocol under the verifying party’s public key. We note that this encryption can be removed in certain cases and will discuss this optimization later in this section.

Let $\mathcal{E}_1 = (\text{KeyGen}_1, \text{Enc}_1, \text{Dec}_1)$ be a CPA-secure symmetric-key encryption scheme, let $\mathcal{E}_2 = (\text{KeyGen}_2, \text{Enc}_2, \text{Dec}_2)$ be a CPA-secure public-key encryption scheme, let $\mathcal{T} = (\text{KeyGen}_{\mathcal{T}}, \text{Mac}_{\mathcal{T}}, \text{Verify}_{\mathcal{T}})$ be a message authentication code, let $\mathcal{M} = (\text{KeyGen}_{\mathcal{M}}, \text{Issue}_{\mathcal{M}}, \text{Extend}_{\mathcal{M}}, \text{Finalize}_{\mathcal{M}}, \text{Verify}_{\mathcal{M}})$ be a macaroon scheme, let $\mathcal{D} = (\text{SetupDischarge}_{\mathcal{D}}, \text{SetupToken}_{\mathcal{D}}, \text{Token}_{\mathcal{D}}, \text{Discharge}_{\mathcal{D}}, \text{Verify}_{\mathcal{D}})$ be a macaroon discharge protocol. We define:

- $\text{Setup}(1^\kappa)$: Output $(\text{params}, \text{sk}_1, \text{sk}_D)$, where:

$$(\text{params}_D, \text{sk}_D) \leftarrow \text{SetupDischarge}_D(1^\kappa) \quad , \quad \text{sk}_1 \leftarrow \text{KeyGen}_1(1^\kappa) \quad , \quad (\text{pk}_2, \text{sk}_2) \leftarrow \text{KeyGen}_2(1^\kappa)$$

$$k_{\mathcal{T}} \leftarrow \text{KeyGen}_{\mathcal{T}}(1^\kappa) \quad , \quad \text{sk}_1 := (\text{sk}_1, \text{sk}_2, k_{\mathcal{T}}) \quad , \quad \text{params} := (\text{params}_D, \text{pk}_2)$$

- $\text{Issue}(\text{sk}_1)$: Parse $\text{sk}_1 = (\text{sk}_1, \text{sk}_2)$. Output (k, μ, ek) , where:

$$k \leftarrow \text{KeyGen}_{\mathcal{M}}(1^\kappa) \quad , \quad c \leftarrow \text{Enc}(\text{sk}_1, k) \quad , \quad \sigma \leftarrow \text{Mac}(k_{\mathcal{T}}, c) \quad , \quad \text{id} := c||\sigma \quad , \quad (\mu, \text{ek}) \leftarrow \text{Issue}_{\mathcal{M}}(k, \text{id})$$

- $\text{Extend1stParty}(\text{ek}, \mu, m)$: Output $(\mu', \text{ek}') \leftarrow \text{Extend}_{\mathcal{M}}(\text{ek}, \mu, 1||m)$.

- $\text{Extend3rdParty}(\text{ek}, \mu)$: Output $(\mu', \text{ek}', \alpha_D)$, where:

$$\begin{aligned} \text{sk}_T &\leftarrow \text{SetupToken}_D(1^\kappa) \quad , \quad c \leftarrow \text{Enc}_2(\text{pk}_2, \text{sk}_T) \quad , \quad (\alpha_S, \alpha_D) \leftarrow \text{Token}_D(\text{sk}_T) \\ (\mu', \text{ek}') &\leftarrow \text{Extend}_{\mathcal{M}}(\text{ek}, \mu, 3\|c\|\alpha_S\|\alpha_D) \end{aligned}$$

- $\mu' \leftarrow \text{Finalize}(\text{ek}, \mu)$: Output $\text{Finalize}_{\mathcal{M}}(\text{ek}, \mu)$.
- $\text{Discharge}(\text{sk}_D, \alpha_D)$: Output $\mu \leftarrow \text{Discharge}_D(\text{sk}_D, \alpha_D)$.
- $b := \text{Verify}(\text{sk}_1, \mu, M = \{\mu_j\})$: Parse $\text{sk}_1 = (\text{sk}_1, \text{sk}_2)$ and $\mu := (\tilde{\text{id}} \mid m_1, \dots, m_n \mid \sigma)$ and $\tilde{\text{id}} = c\|\sigma\|\dots$. Check that $\text{Verify}(k_T, c, \sigma) = 1$ and if so, compute $k := \text{Dec}(\text{sk}_1, \text{id})$ and verify that $\text{Verify}_{\mathcal{M}}(k, \mu) = 1$. For every third-party caveat m_i in μ (ie. every caveat prefixed with 3), parse $m_i := 3\|c_i\|\alpha_S^{(i)}\|\alpha_D^{(i)}$ and compute $\text{sk}_T^{(i)} := \text{Dec}(\text{sk}_2, c_i)$. Finally, verify that there exists $\mu_j \in M$ with key identifier $\alpha_D^{(i)}$, and verify that $\text{Verify}_D(\text{sk}_T^{(i)}, \alpha_S^{(i)}, \mu_j) = 1$.

Theorem 5.3. *Let $\mathcal{E}_1 = (\text{KeyGen}_1, \text{Enc}_1, \text{Dec}_1)$ be a CPA-secure symmetric-key encryption scheme, let $\mathcal{E}_2 = (\text{KeyGen}_2, \text{Enc}_2, \text{Dec}_2)$ be a CPA-secure public-key encryption scheme, let $\mathcal{M} = (\text{KeyGen}_{\mathcal{M}}, \text{Issue}_{\mathcal{M}}, \text{Extend}_{\mathcal{M}}, \text{Verify}_{\mathcal{M}})$ be a macaroon scheme, let $\mathcal{D} = (\text{SetupDischarge}_D, \text{SetupToken}_D, \text{Token}_D, \text{Discharge}_D, \text{Verify}_D)$ be a macaroon discharge protocol. Then the above construction is a secure macaroon-tree protocol.*

OPTIMIZATION. If the macaroon scheme has *pseudorandom* extension keys, as is the case in the first macaroon-scheme construction (see Section 3.1.1), then in the algorithm Extend3rdParty , instead of sampling and encrypting sk_T , it is possible to use a (strong) computational extractor or a KDF to expand the extension key into randomness to run SetupToken_D , as well as a new pseudorandom extension key ek' to use when running $\text{Extend}_{\mathcal{M}}$. This allows Verify to recover sk_T and ek' when verifying the validity of a macaroon, and does not require the expensive encryption step.

Acknowledgements

We are very grateful to Úlfar Erlingsson for valuable discussions about the definitions and constructions. We also thank Tyler Close, Edward Knapp, and Yevgeniy Dodis for helpful discussions and feedback.

References

- [BCK96] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Pseudorandom functions revisited: The cascade construction and its concrete security. In *Proceedings of the 37th Symposium on Foundations of Computer Science, IEEE*, pages 514–523. IEEE, 1996.
- [BPUE⁺14] Arnar Birgisson, Joe Politz, Úlfar Erlingsson, Ankur Taly, Michael Vrbale, and Mark Lentczner. Macaroons: Cookies with contextual caveats for decentralized authorization in the cloud. In *21th Annual Network and Distributed System Security Symposium (NDSS)*, 2014.
- [PSW12] Thomas Peyrin, Yu Sasaki, and Lei Wang. Generic related-key attacks for hmac. In *ASIACRYPT*, volume 7658, pages 580–597. Springer, 2012.

A Proof of Theorem 3.4: Security of SK Macaroons, Construction 1

Proof. We use a series-of-games argument to show the security of the public-key macaroon scheme construction. In all games, the setup phase is the same: the challenger runs $k \leftarrow \text{KeyGen}(1^\kappa)$.

We let p_i be the probability that \mathcal{A} wins in Game i . For simplicity, we do not write extension keys since these are always the same as the corresponding signature.

Game 0: Let Game 0 be the original security game from Definition 3.3. In this game, we have:

Issue Queries: On the i th query, \mathcal{A} gives $\text{id}^{(i)}$ to the challenger who adds $\mu_0^{(i)}$ to \mathcal{T} , where:

$$\sigma^{(i)} \leftarrow F(k, \text{id}^{(i)}) \quad , \quad \mu_0^{(i)} = (\text{id}^{(i)} \mid \mid \sigma^{(i)})$$

Extend Queries: On the i th query, \mathcal{A} gives $\text{id}^{(i)}, \vec{m}^{(i)}, m^{(i)}$ to the challenger, who checks that $(\text{id}^{(i)}, \vec{m}^{(i)})$ is a valid path in \mathcal{T} leading to a macaroon $\mu := (\text{id}^{(i)} \mid \vec{m}^{(i)} \mid \sigma)$. If it is, it adds $\mu^{(i)}$ to \mathcal{T} as a child of μ with edge $m^{(i)}$, where:

$$\sigma^{(i)} \leftarrow F(\sigma, m^{(i)}) \quad , \quad \mu^{(i)} := (\text{id}^{(i)} \mid \vec{m}^{(i)}, m^{(i)} \mid \sigma^{(i)})$$

Macaroon Queries: On the i th query, \mathcal{A} gives $\text{id}^{(i)}, \vec{m}^{(i)}$ to the challenger, who checks that $(\text{id}^{(i)}, \vec{m}^{(i)})$ is a valid path in \mathcal{T} leading to a macaroon $\mu := (\text{id}^{(i)} \mid \vec{m}^{(i)} \mid \sigma)$. If it is, it sends μ to the adversary.

Winning Condition: The adversary \mathcal{A} outputs a macaroon $\mu^* = (\text{id}^* \mid m_1^*, \dots, m_n^* \mid \sigma^*)$. He wins if $\sigma^* = F^*(k, \text{id}^*, m_1^*, \dots, m_n^*)$ and μ^* is not an extension of a macaroon given to \mathcal{A} in a macaroon query. Such a μ^* we call a valid macaroon forgery.

Game 1: In Game 1, we change how the challenger answers macaroon queries. At the beginning of the game, the challenger chooses a random function f . It receives issue and extend queries from \mathcal{A} and builds \mathcal{T} but does not compute the signatures immediately. When it receives a macaroon query id, \vec{m} , it checks if there exists a revealed macaroon $\mu' \in \mathcal{T}$ whose path from the root is id, \vec{m}' and \vec{m}' is a prefix of \vec{m} . Or in other words, it checks if any of the ancestors of the macaroon corresponding to id, \vec{m} in \mathcal{T} have been revealed to \mathcal{A} in a macaroon query. If no ancestor has been revealed, then the challenger computes $\sigma = f(\text{id}, \vec{m})$. If an ancestor $\mu := (\text{id} \mid \vec{m}' \mid \sigma')$ has already been revealed, then the challenger computes the new signature as an extension of σ' , that is, it outputs $\mu := (\text{id}, \vec{m}, \sigma)$ where $\sigma = F^*(\sigma', m_{i+1}, \dots, m_n)$ and $\vec{m} := (m_1, \dots, m_n)$ and $\vec{m}' := (m_1, \dots, m_i)$. The adversary wins if $\sigma^* = f(\text{id}^*, \vec{m}^*)$. Note that f can be simulated efficiently by independently sampling a uniformly random value in $\{0, 1\}^n$ every time the function is evaluated on a new point.

We claim that $|p_0 - p_1| = \text{negl}(\kappa)$ by the pseudorandomness of F^* (see Theorem 2.3). The reduction behaves the same as the challenger in Game 1, except that when it receives a macaroon query and no ancestor has been revealed, it queries its oracle at point (id, \vec{m}) to obtain σ . In the case that the oracle function is F^* , the reduction perfectly simulates Game 0, and if the oracle function is random, then the reduction perfectly simulates Game 1. Therefore, $|p_0 - p_1| = \text{negl}(\kappa)$ by the pseudorandomness of F^* . Notice that this is only true since we assume that an adversary will never ask to reveal a macaroon that is an ancestor of a macaroon that has already been revealed. This guarantees that the reduction's query set is prefix-free.

A Can't Win: We argue that $p_1 = \text{negl}(\kappa)$. Since f is a truly random function, then $f(\text{id}^*, \vec{m}^*)$ is uniformly random in $\{0, 1\}^n$, and therefore $p_1 = 2^{-n} = \text{negl}(\kappa)$.

We conclude that $p_0 \leq |p_0 - p_1| + p_1 = \text{negl}(\kappa)$. □

B Proof of Theorem 3.5: Security of SK Macaroons, Construction 2

Proof. We use a series-of-games argument to show the security of the public-key macaroon scheme construction. In all games, the setup phase is the same: the challenger runs $k \leftarrow \text{KeyGen}(1^\kappa)$ and obtains $k := (\text{vk}_S, \text{pk}_E, \text{sk}_S, \text{sk}_E)$ where $(\text{vk}_S, \text{sk}_S) \leftarrow \text{KeyGen}_S(1^\kappa)$ and $(\text{pk}_E, \text{sk}_E) \leftarrow \text{KeyGen}_E(1^\kappa)$.

We let p_i be the probability that \mathcal{A} wins in Game i . For simplicity, we do not write extension keys since these are always λ , the empty string.

Game 0: Let Game 0 be the original security game from Definition 3.3. In this game, we have:

Issue Queries: On the i th query, \mathcal{A} gives $\text{id}^{(i)}$ to the challenger who adds $\mu_0^{(i)}$ to \mathcal{T} , where:

$$\tilde{\text{id}}^{(i)} := \text{id}^{(i)} \parallel \text{pk}_E \quad , \quad \sigma^{(i)} \leftarrow \text{Sign}(\text{sk}_S, \tilde{\text{id}}^{(i)}) \quad , \quad \mu_0^{(i)} = (\tilde{\text{id}}^{(i)} \mid \mid \sigma^{(i)})$$

Extend Queries: On the i th query, \mathcal{A} gives $\text{id}^{(i)}, \vec{m}^{(i)}, m^{(i)}$ to the challenger, who checks that $(\text{id}^{(i)}, \vec{m}^{(i)})$ is a valid path in \mathcal{T} leading to a macaroon $\mu := (\tilde{\text{id}} \mid \vec{m}^{(i)} \mid \sigma)$. If it is, it adds $\mu^{(i)}$ to \mathcal{T} as a child of μ with edge $m^{(i)}$, where:

$$\begin{aligned} (\text{vk}^{(i)}, \text{sk}^{(i)}) &\leftarrow \text{KeyGen}_{\mathcal{S}}(1^\kappa) \quad , \quad c^{(i)} := \text{Enc}(\text{pk}_{\mathcal{E}}, \sigma) \quad , \quad \tilde{\text{id}}^{(i)} := \tilde{\text{id}} \parallel (\text{vk}^{(i)}, c^{(i)}) \\ \sigma^{(i)} &\leftarrow \text{Sign}(\text{sk}^{(i)}, \text{vk}^{(i)} \parallel c^{(i)} \parallel m^{(i)}) \quad , \quad \mu^{(i)} := (\tilde{\text{id}}^{(i)} \mid \vec{m}^{(i)}, m^{(i)} \mid \sigma^{(i)}) \end{aligned}$$

Macaroon Queries: On the i th query, \mathcal{A} gives $\text{id}^{(i)}, \vec{m}^{(i)}$ to the challenger, who checks that $(\text{id}^{(i)}, \vec{m}^{(i)})$ is a valid path in \mathcal{T} leading to a macaroon $\mu := (\tilde{\text{id}} \mid \vec{m}^{(i)} \mid \sigma)$. If it is, it sends μ to the adversary.

Winning Condition: The adversary \mathcal{A} outputs a macaroon $\mu^* = (\tilde{\text{id}}^* \mid m_1^*, \dots, m_n^* \mid \sigma^*)$. He wins if $\text{Verify}(k, \mu^*) = 1$ and μ^* is not an extension of a macaroon given to \mathcal{A} in a macaroon query. Such a μ^* we call a valid macaroon forgery. In more detail, it must be the case that, $\tilde{\text{id}}^* = \text{id}^* \parallel \text{pk}_{\mathcal{E}} \parallel (\text{vk}_1^*, c_1^*) \parallel \dots \parallel (\text{vk}_n^*, c_n^*)$ and $\text{Verify}(\text{vk}_{\mathcal{S}}, \text{id} \parallel \text{pk}_{\mathcal{E}}, \sigma_0) = 1$ and $\text{Verify}(\text{vk}_i^*, \text{vk}_i^* \parallel c_i^* \parallel m_i^*, \sigma_i^*) = 1$ for $i = 1, \dots, n$, where: $\sigma_i^* := \text{Dec}(\text{sk}_{\mathcal{E}}, c_{i+1}^*)$ for $i = 0, \dots, n-1$ and $\sigma_n^* := \sigma^*$.

Game 1: In Game 1, we change the winning condition so that \mathcal{A} wins only if $\text{id}^* = \text{id}^{(i)}$ for some $\text{id}^{(i)}$ given by \mathcal{A} in an issuing query. This guarantees that μ^* is in \mathcal{T} or is an extension of some macaroon in \mathcal{T} .

We argue that $|p_0 - p_1| = \text{negl}(\kappa)$ by the security of the signature scheme \mathcal{S} under keys $(\text{vk}_{\mathcal{S}}, \text{sk}_{\mathcal{S}})$. Let ε be the probability that \mathcal{A} outputs a valid macaroon forgery and $\text{id}^* \neq \text{id}^{(i)}$ for all $\text{id}^{(i)}$ given by \mathcal{A} in an issuing query. Then $|p_0 - p_1| = \varepsilon$ since the games are equivalent otherwise. We argue that $\varepsilon = \text{negl}(\kappa)$ by the security of \mathcal{S} . The reduction receives the verification key from the signature challenger and uses it as $\text{vk}_{\mathcal{S}}$. It computes $(\text{pk}_{\mathcal{E}}, \text{sk}_{\mathcal{E}}) \leftarrow \text{KeyGen}_{\mathcal{E}}(1^\kappa)$ on its own. When \mathcal{A} asks for an issuing query $\text{id}^{(i)}$, the reduction answers it by querying the signature oracle on $\text{id}^{(i)} \parallel \text{pk}_{\mathcal{E}}$. The reduction answers extend and macaroon queries by computing everything as specified. When \mathcal{A} outputs a forgery $\mu^* := (\tilde{\text{id}}^* \mid m_1^*, \dots, m_n^* \mid \sigma^*)$, it parses $\tilde{\text{id}}^* = \text{id}^* \parallel \text{pk}_{\mathcal{E}} \parallel (\text{vk}_1^*, c_1^*) \parallel \dots$, decrypts $\sigma := \text{Dec}(\text{sk}_{\mathcal{E}}, c_1^*)$ and submits σ as a signature forgery for message $\text{id}^* \parallel \text{pk}_{\mathcal{E}}$. We argue that the probability that the reduction outputs a valid forgery is at least ε . The reduction's forgery is valid only if \mathcal{A} outputs a macaroon μ^* such that $\sigma := \text{Dec}(\text{sk}_{\mathcal{E}}, c_1^*)$ is a valid signature for $\text{id}^* \parallel \text{pk}_{\mathcal{E}}$, and $\text{id}^* \parallel \text{pk}_{\mathcal{E}}$ was not queried to the signing oracle. But note that μ^* is valid only if σ is a valid signature for $\text{id}^* \parallel \text{pk}_{\mathcal{E}}$ since this is part of the macaroon's verification. Furthermore, if $\text{id}^* \neq \text{id}^{(i)}$ for all issuing queries, then $\text{id}^* \parallel \text{pk}_{\mathcal{E}}$ is never queried to the signing oracle. Therefore, by the security of \mathcal{S} , we must have $\varepsilon = \text{negl}(\kappa)$.

Game 2 : Let q_I be an upper bound on the number of issuing queries that \mathcal{A} makes, and let q_E be an upper bound on the number of extension queries that it makes. In Game 2, the challenger picks a uniformly random index $t \in \{1, \dots, q_I + q_E\}$. This number identifies a unique macaroon $\mu_{\text{ch}} := (\tilde{\text{id}}_{\text{ch}} \mid \vec{m}_{\text{ch}} \mid \sigma_{\text{ch}}) \in \mathcal{T}$: if $t \leq q_I$, then t corresponds to the node created in the t th issue query. If $t > q_I$, then t corresponds to the node created in the $(t - q_I)$ th extension query. We change the winning condition so that \mathcal{A} only wins if $\tilde{\text{id}}_{\text{ch}} = \text{id}^* \parallel \dots$ and μ_{ch} is the last macaroon in \mathcal{T} whose caveat list, \vec{m}_{ch} , is a prefix of \vec{m}^* . In other words, \mathcal{A} only wins if \vec{m}_{ch} is a prefix of \vec{m}^* and there does not exist another macaroon $\mu := (\tilde{\text{id}} \mid \vec{m} \mid \sigma) \in \mathcal{T}$ such that $\tilde{\text{id}} = \text{id}^* \parallel \dots$ and \vec{m} is a longer prefix of \vec{m}^* .

We argue that $p_2 = p_1 / (q_E + q_I)$. If \mathcal{A} outputs a valid forgery $\mu^* = (\text{id}^* \parallel \dots \mid \vec{m}^* \mid \sigma^*)$, then there is a unique macaroon $\mu := (\tilde{\text{id}} \mid \vec{m} \mid \sigma) \in \mathcal{T}$ such that $\tilde{\text{id}} = \text{id}^* \parallel \dots$ and \vec{m} is the longest prefix of \vec{m}^* in any macaroon in \mathcal{T} . Since t was sampled uniformly at random, the probability that $\mu_{\text{ch}} = \mu$ is $1 / (q_E + q_I)$.

Game 3.i for $i \in 1, \dots, q_E$: Let q_E and μ_{ch} be as in Game 2. In Game 3.i we change how the challenger answers the i th extension query that extends μ_{ch} . Since q_E is an upper bound on the number of extension queries, in particular it is an upper bound on the number of extensions of μ_{ch} . On the i th extension to μ_{ch} , instead of encrypting σ_{ch} , the challenger encrypts 0.

Extend Queries: Upon receiving $\tilde{\text{id}}_{\text{ch}}, \vec{m}_{\text{ch}}, m^{(i)}$, the challenger adds $\mu^{(i)}$ to \mathcal{T} as a child of μ_{ch} with edge $m^{(i)}$, where:

$$(\text{vk}^{(i)}, \text{sk}^{(i)}) \leftarrow \text{KeyGen}_{\mathcal{S}}(1^\kappa) \quad , \quad c^{(i)} \leftarrow \text{Enc}(\text{pk}_{\mathcal{E}}, 0) \quad , \quad \tilde{\text{id}}^{(i)} := \tilde{\text{id}}_{\text{ch}} \parallel (\text{vk}^{(i)}, c^{(i)})$$

$$\sigma^{(i)} \leftarrow \text{Sign}(\text{sk}^{(i)}, \text{vk}^{(i)} \| c^{(i)} \| m^{(i)}) \quad , \quad \mu^{(i)} := (\tilde{\text{id}}^{(i)} \mid \vec{m}_{\text{ch}}, m^{(i)} \mid \sigma^{(i)})$$

For simplicity let Game 3.0 be Game 2. We argue that $|p_{3.(i-1)} - p_{3.i}| = \text{negl}(\kappa)$ for all $i \in 1, \dots, q_E$ by the CPA-security of \mathcal{E} . The reduction receives pk from the CPA-challenger and uses it as $\text{pk}_{\mathcal{E}}$. It samples $(\text{vk}_{\mathcal{S}}, \text{sk}_{\mathcal{S}}) \leftarrow \text{KeyGen}_{\mathcal{S}}(1^\kappa)$ on its own. It answers issue queries using $\text{sk}_{\mathcal{S}}$ to sign. It answers the j th extension query as follows: if $j < i$, it lets $c^{(j)} \leftarrow \text{Enc}(\text{pk}_{\mathcal{E}}, \sigma)$. For the i th query, it sends $m_0 = \sigma, m_1 = 0$ to the CPA challenger, and receives a challenge ciphertext $c^* \leftarrow \text{Enc}(\text{pk}_{\mathcal{E}}, m_b)$. It lets $c^{(i)} = c^*$. Finally, if $j > i$, it lets $c^{(j)} \leftarrow \text{Enc}(\text{pk}_{\mathcal{E}}, 0)$. The challenger answers macaroon queries as specified. If $b = 0$, the reduction correctly simulates Game 3.($i - 1$), whereas if $b = 1$, the reduction correctly simulates Game 3. i . Therefore, by the CPA-security of \mathcal{E} , we must have $|p_{3.(i-1)} - p_{3.i}| = \text{negl}(\kappa)$.

A Can't Win: We argue that $p_3 = \text{negl}(\kappa)$ by the security of signature scheme \mathcal{S} . Let $\tilde{\text{id}}^* = \text{id}^* \| \text{pk}_{\mathcal{E}} \| (\text{vk}_1, c_1) \| \dots \| (\text{vk}_n, c_n)$, and let j be such that $\text{id}_{\text{ch}} = \text{id}^* \| \text{pk}_{\mathcal{E}} \| (\text{vk}_1, c_1) \| \dots \| (\text{vk}_j, c_j)$, or let $j = 0$ if $\text{id}_{\text{ch}} = \text{id}^* \| \text{pk}_{\mathcal{E}}$ (in this case, define $\text{vk}_0 = \text{vk}_{\mathcal{S}}$). We know such a j exists by the definition of μ_{ch} . In the reduction, the challenger will receive vk from the signing challenger, and use $\text{vk}_j = \text{vk}$. It computes $(\text{pk}_{\mathcal{E}}, \text{sk}_{\mathcal{E}}) \leftarrow \text{KeyGen}_{\mathcal{E}}(1^\kappa)$ on its own. When \mathcal{A} asks for an issuing query or extension query that does not involve signing with sk_j (the corresponding secret key to vk_j), the challenger answers them accordingly. If a query involves signing a message with sk_j , then the challenger uses the signing oracle to compute the signature. However, the challenger never computes σ_{ch} . This means that the reduction never queries the signing oracle on $\text{id} \| \text{pk}_{\mathcal{E}}$ if $j = 0$ or on $\text{vk}_j \| c_j \| m_j$ if $j > 0$. This is possible since all extensions of μ_{ch} encrypt 0 instead of σ_{ch} , and if the adversary wins the game, then σ_{ch} is never revealed in a macaroon query. When \mathcal{A} outputs its forgery, the reduction computes $\sigma_j = \text{Dec}(\text{sk}_{\mathcal{E}}, c_{j+1})$, and outputs σ_j as its forgery. If \mathcal{A} wins, then σ_j must be a valid signature under vk_j for $\text{id} \| \text{pk}_{\mathcal{E}}$ if $j = 0$ or for $\text{vk}_j \| c_j \| m_j$ if $j > 0$ since this check is part of the verification algorithm. Since the signing oracle was never queried by the reduction on that message, σ_j is also a valid signature forgery. Therefore, by the security of \mathcal{S} , we must have $p_3 = \text{negl}(\kappa)$.

We wish to prove that $p_0 = \text{negl}(\kappa)$. We know that $|p_0 - p_1| = \text{negl}(\kappa)$, $p_1 = (q_1 + q_E) \cdot p_2$, $|p_2 - p_3| = \text{negl}(\kappa)$, and $p_3 = \text{negl}(\kappa)$. Putting it all together, we have:

$$|p_0 - p_1| + (q_1 + q_E) \cdot |p_2 - p_3| \geq (p_0 - p_1) + (q_1 + q_E) \cdot (p_2 - p_3) = p_0 - (q_1 + q_E) \cdot p_3$$

Since q_1 and q_E are number of queries made by \mathcal{A} and \mathcal{A} runs in polynomial time, we must have that $(q_1 + q_E) = \text{poly}(\kappa)$. We can therefore conclude that

$$p_0 \leq |p_0 - p_1| + (q_1 + q_E) \cdot |p_2 - p_3| + (q_1 + q_E) \cdot p_3 = \text{negl}(\kappa),$$

as required. □

C Proof of Theorem 3.7: Security of PK Macaroons

Proof. We use a series-of-games argument to show the security of the macaroon discharge protocol construction. In all games, the setup phase is the same: the challenger runs $(\text{params}, \text{sk}_I, \text{sk}_D) \leftarrow \text{Setup}(1^\kappa)$ and gives params to \mathcal{A} . We let p_i be the probability that \mathcal{A} wins in Game i .

Game 0: Let Game 0 be the original game from the definition of security of public-key macaroons (analogous to Definition 4.2).

Game 1 : Let q_1 be an upper bound on the number of issuing queries that \mathcal{A} makes, and let q_E be an upper bound on the number of extension queries that it makes. In Game 1, the challenger picks a uniformly random index $t \in \{1, \dots, q_1 + q_E\}$. This number identifies a unique macaroon $\mu_{\text{ch}} := (\tilde{\text{id}}_{\text{ch}} \mid \vec{m}_{\text{ch}} \mid \sigma_{\text{ch}}) \in \mathcal{T}$: if $t \leq q_1$, then t corresponds to the node created in the t th issue query. If $t > q_1$, then t corresponds to the node created in the $(t - q_1)$ th extension query. We change the winning condition so that in addition, \mathcal{A} only wins if $\tilde{\text{id}}_{\text{ch}} = \text{id} \| \text{vk}_1, \dots, \text{vk}_i$ is the longest prefix in \mathcal{T} of $\tilde{\text{id}}^*$.

This means $\tilde{\text{id}}^* = \text{id} \parallel \text{vk}_1 \parallel \dots \parallel \text{vk}_i \parallel \text{vk}_{i+1}^* \parallel \dots \parallel \text{vk}_n^* \parallel \text{vk}_{n+1}^*$, and there does not exist another macaroon $\mu := (\tilde{\text{id}} \mid \tilde{m} \mid \sigma) \in \mathcal{T}$ such that $\tilde{\text{id}}$ is a longer prefix of $\tilde{\text{id}}^*$.

We argue that $p_1 = p_0 / (q_E + q_I)$. If \mathcal{A} outputs a valid forgery $\mu^* = (\text{id}^* \parallel \dots \parallel \tilde{m}^* \mid \sigma^*)$, then there is a unique macaroon $\mu := (\tilde{\text{id}} \mid \tilde{m} \mid \sigma) \in \mathcal{T}$ such that $\tilde{\text{id}}$ the longest prefix of $\tilde{\text{id}}^*$ in any macaroon in \mathcal{T} . Since t was sampled uniformly at random, the probability that $\mu_{\text{ch}} = \mu$ is $1 / (q_E + q_I)$.

Game 2: We change the winning condition so that \mathcal{A} only wins if $\tilde{\text{id}}^* = \tilde{\text{id}}_{\text{ch}}$. Let ε be the probability that $i < n$. Then $|p_1 - p_2| = \varepsilon$ since the games are equivalent otherwise. We argue that $\varepsilon = \text{negl}(\kappa)$ by the security of \mathcal{S} under verification key vk_i . The reduction receives the verification key from the signature challenger and uses it as vk_i . The reduction answers issue, extension, and macaroon queries as specified, except when extending μ_{ch} . In this case, it samples a key pair $\text{vk}^{(j)}, \text{sk}^{(j)}$ as usual, and uses the signing oracle to compute the signature $\nu^{(j)}$ on $\text{vk}^{(j)}$. When the adversary outputs a macaroon forgery $\mu^* = (\tilde{\text{id}}^* \mid \tilde{m}^* \mid \sigma^*)$, it parses $\tilde{\text{id}}^* = \text{id} \parallel \text{vk}_1 \parallel \dots \parallel \text{vk}_i \parallel \text{vk}_{i+1}^* \parallel \dots \parallel \text{vk}_n^*$ and $\sigma^* = (\nu_0^*, \tau_0^*) \parallel \dots \parallel (\nu_i^*, \tau_i^*) \parallel \dots$ and outputs ν_i^* as a signature forgery of vk_{i+1}^* . If μ^* is a valid macaroon forgery, then ν_i^* is a valid signature since the macaroon verification checks this. Moreover, since μ_{ch} is the macaroon with the largest prefix, then $\text{vk}_{i+1}^* \neq \text{vk}^{(j)}$ for all j and therefore vk_{i+1}^* was never queried to the signature oracle. We thus conclude that $\varepsilon = \text{negl}(\kappa)$.

Game 3: Let $\mu_{\text{ch}} = (\text{id} \parallel \dots \parallel \tilde{m} \mid \sigma)$. In Game 3, we change the winning condition so that when \mathcal{A} submits a forgery $\mu^* := (\tilde{\text{id}}^* \mid \tilde{m}^* \mid \sigma^*)$, it wins only if $\tilde{\text{id}}^* = \text{id}^* \parallel \dots$ and $(\text{id}^*, \tilde{m}^*) = (\text{id}, \tilde{m})$. In other words, the key identifier and the caveats of the adversary's forgery must match those of μ_{ch} .

Let i be the last index at which $(\text{id}^*, \tilde{m}^*)$ and (id, \tilde{m}) are the same (if $\text{id}^* \neq \text{id}$ then $i = 0$). Let ε be the probability that $i < n$. Then $|p_1 - p_2| = \varepsilon$ since the games are equivalent otherwise. We argue that $\varepsilon = \text{negl}(\kappa)$ by the security of the signature scheme \mathcal{S} under key vk_{i+1} . The reduction answers issue, extend, and macaroon queries as usual, except when extending μ_{ch} . In this case, instead of sampling a new key pair, it queries the signature oracle. When \mathcal{A} outputs a forgery $\mu^* := (\tilde{\text{id}}^* \mid m_1^*, \dots, m_n^* \mid \sigma^*)$, the reduction parses $\sigma^* = (\nu_0^*, \tau_0^*) \parallel \dots \parallel (\nu_{i+1}^*, \tau_{i+1}^*) \parallel \dots$, and outputs τ_{i+1}^* as a signature forgery for message m_{i+1}^* .

The reduction's forgery is valid only if \mathcal{A} outputs a macaroon μ^* such that τ_{i+1}^* is a valid signature for m_{i+1}^* , and m_{i+1}^* was not queried to the signing oracle. Note that if μ^* is valid then the first condition is satisfied since this is part of the macaroon's verification. Furthermore, if $m_{i+1}^* \neq m_{i+1}$ then m_{i+1}^* is never queried to the signing oracle. Therefore, by the security of \mathcal{S} , we must have $\varepsilon = \text{negl}(\kappa)$.

\mathcal{A} Can't Win: We argue that \mathcal{A} can't win by the security of \mathcal{S} under verification key vk_{n+1} . The reduction uses the verification key from the signature challenger as vk_{n+1} , and answers issue, extend, and macaroon queries as usual, using the signing oracle when needed. When \mathcal{A} outputs a forgery $\mu^* := (\tilde{\text{id}}^* \mid m_1^*, \dots, m_n^* \mid \sigma^*)$, the reduction parses $\tilde{\text{id}}^* = \text{id}^* \parallel \dots$ and $\sigma^* = (\nu_0^*, \tau_0^*) \parallel \dots \parallel \sigma_{\text{final}}$, and outputs σ_{final} as a signature forgery for message id^* .

The reduction's forgery is valid only if \mathcal{A} outputs a macaroon μ^* such that σ_{final} is a valid signature for id^* , and id^* was not queried to the signing oracle. Note that if μ^* is valid then the first condition is satisfied since this is part of the macaroon's verification. Furthermore, the challenger never extends a macaroon with a caveat equal to its key identifier, so the reduction never queried id^* to the signing oracle. Therefore, by the security of \mathcal{S} , we must have $\varepsilon = \text{negl}(\kappa)$. □

D Proof of Theorem 4.3: Security of Discharge Protocol, Construction 1

Proof. We use a series-of-games argument to show the security of the macaroon discharge protocol construction. In all games, the setup phase is the same: the challenger runs $(\text{params}, \text{sk}_D) \leftarrow \text{SetupDischarge}(1^\kappa)$ and $\text{sk}_T \leftarrow \text{SetupToken}(1^\kappa)$ and gives params to \mathcal{A} . We let p_i be the probability that \mathcal{A} wins in Game i .

Game 0: Let Game 0 be the original game from Definition 4.2. In this game, we have:

Token Queries: On the i th query, \mathcal{A} receives $\alpha_S^{(i)} := (c_1^{(i)}, \sigma^{(i)})$, and $\alpha_D^{(i)} := c_2^{(i)}$, where

$$r^{(i)} \leftarrow \text{KeyGen}_{\mathcal{M}}(1^\kappa) \quad , \quad c_1^{(i)} \leftarrow \text{Enc}_1(\text{sk}_1, r^{(i)}) \quad , \quad \sigma^{(i)} \leftarrow \text{Mac}(k, c_1^{(i)}) \quad , \quad c_2^{(i)} \leftarrow \text{Enc}_2(\text{pk}_2, r^{(i)})$$

Discharge Queries: On the i th query $\alpha_D^{(i)*} = c_2^{(i)*}$, \mathcal{A} receives $\mu^{(i)}$, where

$$r^{(i)} \leftarrow \text{Dec}_2(\text{sk}_2, c_2^{(i)*}) \quad , \quad \mu^{(i)} \leftarrow \text{Issue}_{\mathcal{M}}(r^{(i)}, c_2^{(i)*})$$

Winning Condition: \mathcal{A} wins if it produces a token $\alpha_S^* = (c_1^*, \sigma^*)$ and a macaroon μ^* such that $\text{Verify}(r^*, \mu^*) = 1$ where $r^* = \text{Dec}_1(\text{sk}_1, c_1^*)$, and μ^* is not an extension of any macaroon μ such that $(\alpha_D^*, \mu) \in \mathcal{Q}_D$, where α_D^* corresponds to α_S^* in \mathcal{Q}_T . In other words, \mathcal{A} wins if it outputs a valid macaroon that is not an extension of a macaroon that was the output of a discharge query with a discharge token corresponding to α_S^* .

Game 1: Let q_T be an upper bound on the number of issuing queries that \mathcal{A} makes. In Game 1, the challenger remembers each $r^{(i)}$ that it samples in an issuing query, and we change the winning condition so that \mathcal{A} only wins if μ^* is a valid macaroon under key $r^{(i)}$ for some $i \in \{1, \dots, q_T\}$. In other words, \mathcal{A} only wins if it produces a valid macaroon under a key that the challenger sampled in a token query.

Winning Condition: \mathcal{A} wins if it outputs (α_S^*, μ^*) such that $\text{Verify}(r^{(i)}, \mu^*) = 1$ for some $i \in \{1, \dots, q_T\}$ and μ^* is not an extension of any macaroon μ such that $(\alpha_D^*, \mu) \in \mathcal{Q}_D$, where α_D^* corresponds to α_S^* in \mathcal{Q}_T .

We argue that $|p_0 - p_1| = \text{negl}(\kappa)$ by the security of the MAC \mathcal{T} . Consider the probability ε that \mathcal{A} outputs a valid forgery that contains $c_1^* \neq c_1^{(i)}$ for all $i \in \{1, \dots, q_T\}$, or in other words, that c_1^* is a ciphertext that was not given to \mathcal{A} as a response to a token query. It is easy to see that $|p_0 - p_1| = \varepsilon$ since by correctness of decryption of \mathcal{E}_1 , the games are equivalent if $c_1^* = c_1^{(i)}$ for some $i \in \{1, \dots, q_T\}$. By the security of the MAC \mathcal{T} , we must have $\varepsilon = \text{negl}(\kappa)$. The reduction samples $\text{sk}_1, \text{pk}_2, \text{sk}_2$, answers token queries by computing $c_1^{(i)}$ and $c_2^{(i)}$ on its own and using the MAC tagging oracle to obtain $\sigma^{(i)}$. It answers discharge queries using sk_2 , and given \mathcal{A} 's forgery (c_1^*, σ^*, μ^*) , it outputs c_1^*, σ^* . Since this simulates \mathcal{A} 's view perfectly, the reduction produces a tag on a new message with probability ε .

Game 2: In Game 2, the challenger chooses an index $i_{\text{ch}} \in \{1, \dots, q_T\}$ uniformly at random. We change the winning condition so that \mathcal{A} only wins if μ^* is a valid macaroon under key $r^{(i_{\text{ch}})}$.

Winning Condition: The challenger picks $i_{\text{ch}} \in \{1, \dots, q_T\}$ uniformly at random. \mathcal{A} wins if it outputs (α_S^*, μ^*) such that $\text{Verify}(r^{(i_{\text{ch}})}, \mu^*) = 1$ and μ^* is not an extension of any macaroon μ such that $(\alpha_D^*, \mu) \in \mathcal{Q}_D$. In other words, \mathcal{A} wins if it outputs a valid macaroon under $r^{i_{\text{ch}}}$ that is not an extension of a macaroon that was the output of a discharge query with token α_S^* .

Since i_{ch} is chosen uniformly at random, we have that $p_2 = p_1/q_T$.

Game 3: In Game 3, we change how the i_{ch} th token query is answered by the challenger. In particular, $c_1^{(i_{\text{ch}})}$ will now be an encryption of 0 instead of an encryption of $r^{(i_{\text{ch}})}$.

i_{ch} th Token Query: On the i_{ch} th query, \mathcal{A} receives $\alpha_S^{(i_{\text{ch}})} := (c_1^{(i_{\text{ch}})}, \sigma^{(i_{\text{ch}})})$, and $\alpha_D^{(i_{\text{ch}})} := c_2^{(i_{\text{ch}})}$, where

$$r^{(i_{\text{ch}})} \leftarrow \text{KeyGen}_{\mathcal{M}}(1^\kappa) \quad , \quad c_1^{(i_{\text{ch}})} \leftarrow \text{Enc}_1(\text{sk}_1, 0) \quad , \quad \sigma \leftarrow \text{Mac}(k, c_1^{(i_{\text{ch}})}) \quad , \quad c_2^{(i_{\text{ch}})} \leftarrow \text{Enc}_2(\text{pk}_2, r^{(i_{\text{ch}})})$$

We argue that $|p_2 - p_3| = \text{negl}(\kappa)$ by the CPA-security of \mathcal{E}_1 . The reduction samples $\text{pk}_2, \text{sk}_2, k$ and chooses a random $i_{\text{ch}} \in \{1, \dots, q_T\}$. It answers the i th token query as follows: if $i \neq i_{\text{ch}}$, it generates $r^{(i)}$ and obtains $c_1^{(i)}$ from the encryption oracle. It computes $\sigma^{(i)}, c_2^{(i)}$ on its own. On the i_{ch} th token query, the reduction gives $m_0 = r^{i_{\text{ch}}}$ and $m_1 = 0$ to the CPA-challenger as the challenge plaintexts. It obtains a challenge ciphertext $\text{Enc}_1(\text{sk}_1, m_b)$ in return, which it uses as $c_1^{(i_{\text{ch}})}$. It then computes $\sigma^{(i_{\text{ch}})}, c_2^{(i_{\text{ch}})}$ on its own. The reduction answers discharge queries as usual, decrypting with sk_2 .

If $b = 0$, then the reduction perfectly simulates Game 2 and if $b = 1$, the reduction perfectly simulates Game 1. Therefore, we conclude that $|p_2 - p_3| = \text{negl}(\kappa)$ by the CPA-security of \mathcal{E}_1 .

Game 4: In Game 4, we change how we answer discharge queries. More specifically, the challenger checks if $c_2^{(i)*} = c_2^{(i_{\text{ch}})}$, or in other words, if the query ciphertext is the same as the c_2 ciphertext given as a response to the i_{ch} th token query. If $c_2^{(i)*} \neq c_2^{(i_{\text{ch}})}$, then the discharge query is answered as usual, but if $c_2^{(i)*} = c_2^{(i_{\text{ch}})}$, then the challenger issues a macaroon under key $r^{(i_{\text{ch}})}$.

Games 3 and 4 are identically distributed by the correctness of decryption of \mathcal{E}_2 . Thus, $p_3 = p_4$.

Game 5: In Game 5, we again change how the i_{ch} th token query is answered by the challenger. In particular, $c_2^{(i_{\text{ch}})}$ will now be an encryption of 0 instead of an encryption of $r^{(i_{\text{ch}})}$.

i_{ch} th Token Query: On the i_{ch} th query, \mathcal{A} receives $\alpha_S^{(i_{\text{ch}})} := (c_1^{(i_{\text{ch}})}, \sigma^{(i_{\text{ch}})})$, and $\alpha_D^{(i_{\text{ch}})} := c_2^{(i_{\text{ch}})}$, where

$$r^{(i_{\text{ch}})} \leftarrow \text{KeyGen}_{\mathcal{M}}(1^\kappa) \quad , \quad c_1^{(i_{\text{ch}})} \leftarrow \text{Enc}_1(\text{sk}_1, 0) \quad , \quad \sigma \leftarrow \text{Mac}(k, c_1^{(i_{\text{ch}})}) \quad , \quad c_2^{(i_{\text{ch}})} \leftarrow \text{Enc}_2(\text{pk}_2, 0)$$

We argue that $|p_4 - p_5| = \text{negl}(\kappa)$ by the CCA2-security of \mathcal{E}_2 . The reduction obtains pk_2 from the CCA2-challenger, and samples sk_1 and k . To answer the i th token query, it computes $c_1^{(i)}, \sigma$ on its own. If $i \neq i_{\text{ch}}$, it also computes $c_2^{(i)}$. On the other hand, if $i = i_{\text{ch}}$, then it sends $m_0 = r^{(i_{\text{ch}})}, m_1 = 0$ to the CCA-challenger as the challenge plaintexts. It obtains a challenge ciphertext $\text{Enc}(\text{pk}_2, m_b)$, which it uses as $c_2^{(i)}$. It answers discharge queries as follows: If $c_2^{(i)*} \neq c_2^{(i_{\text{ch}})}$, then it gives $c_2^{(i)*}$ to the decryption oracle and uses the plaintext it receives back as the key to issue the macaroon. If $c_2^{(i)*} = c_2^{(i_{\text{ch}})}$, then the challenger simply issues the macaroon under key $r^{(i_{\text{ch}})}$. Note that this guarantees that we never query the decryption oracle with the challenge ciphertext.

If $b = 0$, then the reduction perfectly simulates Game 4 and if $b = 1$, the reduction perfectly simulates Game 5. Therefore, we conclude that $|p_4 - p_5| = \text{negl}(\kappa)$ by the CCA2-security of \mathcal{E}_2 .

A Can't Win: We argue that $p_5 = \text{negl}(\kappa)$ by the security of the macaroon scheme \mathcal{M} . The reduction samples $\text{sk}_1, \text{sk}_2, \text{pk}_2, k$. It answers token queries by computing $c_1^{(i)}, c_2^{(i)}, \sigma^{(i)}$ on its own. Note that this does not require encrypting $r^{(i_{\text{ch}})}$. It answers discharge queries as follows: If $c_2^{(i)*} \neq c_2^{(i_{\text{ch}})}$, it answers the query as usual: decrypts $c_2^{(i)*}$ and creates a macaroon using the plaintext. But if $c_2^{(i)*} = c_2^{(i_{\text{ch}})}$, it calls the issuing oracle with $c_2^{(i_{\text{ch}})}$ and the asks for it with a macaroon query. Finally, when \mathcal{A} outputs α_S^*, μ^* , it outputs μ^* as its forgery. If \mathcal{A} wins, then μ^* is a valid macaroon under $r^{(i_{\text{ch}})}$ and is not an extension of a macaroon μ such that $(\alpha_S^*, \mu) \in \mathcal{Q}_D$. In particular, this means that it is not an extension of a macaroon given to the reduction as a macaroon query. Therefore, the probability that the reduction gives a valid forgery is exactly p_5 since it simulated Game 5 perfectly for \mathcal{A} . By the security of the macaroon scheme \mathcal{M} , we conclude $p_5 = \text{negl}(\kappa)$.

We wish to prove that $p_0 = \text{negl}(\kappa)$. We know that $|p_0 - p_1| = \text{negl}(\kappa)$, $p_1 = q_T \cdot p_2$, $|p_2 - p_3| = \text{negl}(\kappa)$, $p_3 = p_4$, $|p_4 - p_5| = \text{negl}(\kappa)$, and $p_5 = \text{negl}(\kappa)$. Putting it all together, we have:

$$|p_0 - p_1| + q_T \cdot |p_2 - p_3| + q_T \cdot |p_4 - p_5| \geq (p_0 - p_1) + q_T \cdot (p_2 - p_3) + q_T \cdot (p_4 - p_5) = p_0 - q_T \cdot p_5$$

Since q_T is the number of queries that \mathcal{A} makes and \mathcal{A} runs in polynomial time, we must have that $q_T = \text{poly}(\kappa)$. We can therefore conclude that

$$p_0 \leq |p_0 - p_1| + q_T \cdot |p_2 - p_3| + q_T \cdot |p_4 - p_5| + q_T \cdot p_5 = \text{negl}(\kappa),$$

as required. \square

E Proof of Theorem 4.4: Security of Discharge Protocol, Construction 2

Proof. We use a series-of-games argument to show the security of the macaroon discharge protocol construction. In all games, the setup phase is the same: the challenger runs $(\text{params}, \text{sk}_D) \leftarrow \text{SetupDischarge}(1^\kappa)$ and $\text{sk}_T \leftarrow \text{SetupToken}(1^\kappa)$ and gives params to \mathcal{A} . We let p_i be the probability that \mathcal{A} wins in Game i .

Game 0: Let Game 0 be the original game from Definition 4.2. In this game, we have:

Token Queries: On the i th query, \mathcal{A} receives $\alpha_{\mathcal{S}}^{(i)} := (r^{(i)}, \sigma^{(i)})$, and $\alpha_{\mathcal{D}}^{(i)} := r^{(i)}$, where

$$r^{(i)} \leftarrow \{0, 1\}^{\kappa} \quad , \quad \sigma^{(i)} \leftarrow \text{Mac}(k, r^{(i)})$$

Discharge Queries: On the i th query $\alpha_{\mathcal{D}}^{(i)*} = r^{(i)*}$, \mathcal{A} receives $\mu^{(i)}$, where

$$\mu^{(i)} \leftarrow \text{Issue}_{\mathcal{M}}(\text{sk}_{\mathcal{D}}, r^{(i)*})$$

Winning Condition: \mathcal{A} *wins* if it produces a token $\alpha_{\mathcal{S}}^* = (r^*, \sigma^*)$ and a macaroon μ^* such that $\text{Verify}(\text{vk}, \mu^*) = 1$, μ^* 's key identifier is r^* , and μ^* is not an extension of any macaroon μ that was given to \mathcal{A} as a discharge query with token r^* .

Game 1: Let $q_{\mathcal{T}}$ be an upper bound on the number of issuing queries that \mathcal{A} makes. In Game 1, the challenger remembers each $r^{(i)}$ that it samples in a token query, and we change the winning condition so that \mathcal{A} only wins if $r^* = r^{(i)}$ for some $i \in \{1, \dots, q_{\mathcal{T}}\}$. In other words, \mathcal{A} only wins if it produces a valid discharge macaroon for a token it received from a token query.

Winning Condition: \mathcal{A} *wins* if it outputs $(\alpha_{\mathcal{S}}^*, \mu^*)$ such that $\text{Verify}(\text{vk}, \mu^*) = 1$, μ^* key identifier is r^* , $r^* = r^{(i)}$ for some $i \in \{1, \dots, q_{\mathcal{T}}\}$, and μ^* is not an extension of any macaroon μ that was given to \mathcal{A} as a discharge query with token r^* .

We argue that $|p_0 - p_1| = \text{negl}(\kappa)$ by the security of the MAC \mathcal{T} . Consider the probability ε that \mathcal{A} outputs a valid forgery that contains $r^* \neq r^{(i)}$ for all $i \in \{1, \dots, q_{\mathcal{T}}\}$, or in other words, that r^* was not given to \mathcal{A} as a response to an token query. It is easy to see that $|p_0 - p_1| = \varepsilon$ since the games are equivalent otherwise. By the security of the MAC \mathcal{T} , we must have $\varepsilon = \text{negl}(\kappa)$. The reduction samples (vk, sk) on its own, answers token queries by computing sampling random $r^{(i)}$ and using the MAC tagging oracle to compute $\sigma^{(i)}$. It answers discharge queries using sk . Given \mathcal{A} 's forgery (r^*, σ^*, μ^*) , it outputs σ^* and its forgery on message r^* . Since this simulates \mathcal{A} 's view perfectly, the reduction produces a tag on a new message with probability ε .

Game 2: In Game 2, the challenger chooses an index $i_{\text{ch}} \in \{1, \dots, q_{\mathcal{T}}\}$ uniformly at random. We change the winning condition so that \mathcal{A} only wins if $r^* = r^{(i_{\text{ch}})}$.

Winning Condition: \mathcal{A} *wins* if it outputs $(\alpha_{\mathcal{S}}^*, \mu^*)$ such that $\text{Verify}(\text{vk}, \mu^*) = 1$, μ^* key identifier is r^* , $r^* = r^{(i_{\text{ch}})}$, and μ^* is not an extension of any macaroon μ that was given to \mathcal{A} as a discharge query with token r^* .

Since i_{ch} is chosen uniformly at random, we have that $p_2 = p_1/q_{\mathcal{T}}$.

\mathcal{A} Can't Win: We argue that $p_2 = \text{negl}(\kappa)$ by the security of the macaroon scheme \mathcal{M} . The reduction samples k and receives vk from the macaroon scheme challenger. It answers token queries by computing $r^{(i)}, \sigma^{(i)}$ on its own. It answers discharge queries by calling its issuing oracle with $r^{(i)*}$ and immediately asking for the resulting macaroon with a macaroon query. Finally, when \mathcal{A} outputs $\alpha_{\mathcal{S}}^*, \mu^*$, it outputs μ^* as its forgery. If \mathcal{A} wins, then μ^* is a valid macaroon with $r^{(i_{\text{ch}})}$ as its key identifier, and it is not an extension of a macaroon μ given to \mathcal{A} as a discharge query for token $r^{(i_{\text{ch}})}$. In particular, this means that it is not an extension of a macaroon given to the reduction as a macaroon query. Therefore, the probability that the reduction gives a valid forgery is exactly p_2 since it simulated Game 2 perfectly for \mathcal{A} . By the security of the macaroon scheme \mathcal{M} , we conclude $p_2 = \text{negl}(\kappa)$.

We wish to prove that $p_0 = \text{negl}(\kappa)$. We know that $|p_0 - p_1| = \text{negl}(\kappa)$, $p_1 = q_{\mathcal{T}} \cdot p_2$, and $p_2 = \text{negl}(\kappa)$. Putting it all together, we have:

$$|p_0 - p_1| + q_{\mathcal{T}} \cdot p_2 \geq (p_0 - p_1) + q_{\mathcal{T}} \cdot p_2 = p_0$$

Since $q_{\mathcal{T}}$ is the number of queries that \mathcal{A} makes and \mathcal{A} runs in polynomial time, we must have that $q_{\mathcal{T}} = \text{poly}(\kappa)$. We can therefore conclude that

$$p_0 \leq |p_0 - p_1| + q_{\mathcal{T}} \cdot p_2 = \text{negl}(\kappa),$$

as required. □

F Proof of Theorem 5.3: Security of our Macaroon-Tree Protocol

Proof. We use a series-of-games argument to show the security of the macaroon discharge protocol construction. In all games, the setup phase is the same: the challenger runs $(\text{params}, \text{sk}_1, \text{sk}_D) \leftarrow \text{Setup}(1^\kappa)$ and gives params to \mathcal{A} . We let p_i be the probability that \mathcal{A} wins in Game i .

Game 0: Let Game 0 be the original game from Definition 5.2. In this game, we have:

Issue Queries: On the i th query, the challenger creates a macaroon $(\mu_0^{(i)}, \text{ek}_0^{(i)})$ and adds it as a child of the root under an edge labeled with $c^{(i)} \parallel \sigma^{(i)}$, where:

$$k^{(i)} \leftarrow \text{KeyGen}_{\mathcal{M}}(1^\kappa) \quad , \quad c^{(i)} \leftarrow \text{Enc}(\text{sk}_1, k^{(i)}) \quad , \quad \sigma^{(i)} \leftarrow \text{Mac}(k_{\mathcal{T}}, c^{(i)}) \quad , \quad \text{id}^{(i)} := c^{(i)} \parallel \sigma^{(i)}$$

$$(\mu_0^{(i)}, \text{ek}_0^{(i)}) \leftarrow \text{Issue}_{\mathcal{M}}(k^{(i)}, \text{id}^{(i)})$$

Extend Queries: On the i th query, the adversary sends an identifier $\text{id}^{(i)}$, a tuple of caveats $\vec{m}^{(i)}$, and a message $m^{(i)} \in \{0, 1\}^* \cup \lambda$ to the challenger, where λ is the empty string. If $(\text{id}^{(i)}, \vec{m}^{(i)})$ is a valid path in \mathcal{T} leading to a node (μ, ek) , then the challenger computes $(\mu^{(i)}, \text{ek}^{(i)}, \alpha_D^{(i)}) \leftarrow \text{Extend}(\text{ek}, \mu, m^{(i)})$, and adds $(\mu^{(i)}, \text{ek}^{(i)}, \alpha_D^{(i)})$ to \mathcal{T} as a child of (μ, ek) under an edge labeled with $m^{(i)}$. Here, $\text{Extend}(\text{ek}, \mu, m)$ is defined as follows: If $m = \lambda$, it outputs $\text{Extend3rdParty}(\text{ek}, \mu)$, otherwise it outputs $\text{Extend}_{\mathcal{M}}(\text{ek}, \mu, 3 \parallel c^{(i)} \parallel \alpha_S^{(i)} \parallel \alpha_D^{(i)})$ and $\alpha = \lambda$, where:

$$\text{sk}_{\mathcal{T}}^{(i)} \leftarrow \text{SetupToken}_{\mathcal{D}}(1^\kappa) \quad , \quad c^{(i)} \leftarrow \text{Enc}_2(\text{pk}_2, \text{sk}_{\mathcal{T}}^{(i)}) \quad , \quad (\alpha_S^{(i)}, \alpha_D^{(i)}) \leftarrow \text{Token}_{\mathcal{D}}(\text{sk}_{\mathcal{T}})$$

Discharge Queries: On the i th query, the adversary \mathcal{A} sends a token $\alpha_D^{(i)*}$ to the challenger, who runs $\mu^{(i)} \leftarrow \text{Discharge}_{\mathcal{D}}(\text{sk}_D, \alpha_D^{(i)*})$ and returns $\mu^{(i)}$ to \mathcal{A} . The challenger also adds $(\alpha_D^{(i)*}, \mu^{(i)})$ to \mathcal{Q}_D .

Macaroon Queries: On the i th query, the adversary sends an identifier $\text{id}^{(i)}$ and a tuple of caveats $\vec{m}^{(i)}$ to the challenger. If $(\text{id}^{(i)}, \vec{m}^{(i)})$ is a valid path in \mathcal{T} leading to a node (μ, ek) , then the challenger returns (μ, ek) to the adversary.

Winning condition: The adversary \mathcal{A} outputs a macaroon μ^* and a set of discharges $M^* = \{\mu_j^*\}$. He wins if:

1. $\text{Verify}(k_{\mathcal{T}}, c, \sigma) = 1$ and $\text{Verify}_{\mathcal{M}}(k, \mu) = 1$ where $k := \text{Dec}(\text{sk}_1, \text{id})$, and
2. For every third-party caveat $3 \parallel c_i \parallel \alpha_S^{(i)} \parallel \alpha_D^{(i)}$ in μ , there exists $\mu_j \in M$ with key identifier $\alpha_D^{(i)}$ such that $\text{Verify}_{\mathcal{D}}(\text{sk}_{\mathcal{T}}^{(i)}, \alpha_S^{(i)}, \mu_j) = 1$ where $\text{sk}_{\mathcal{T}}^{(i)} := \text{Dec}(\text{sk}_2, c_i)$, and
3. Either
 - (a) μ^* is not an extension of a macaroon given to \mathcal{A} in a macaroon query, or
 - (b) There exists a macaroon $m_j^* \in M^*$ corresponding to a third-party caveat $3 \parallel c_j \parallel \alpha_S^{(j)} \parallel \alpha_D^{(j)}$ in μ^* such that μ_j^* is not an extension of any macaroon μ such that $(\alpha_D^{(j)}, \mu) \in \mathcal{Q}_D$.

Game 1: Let q_1 be an upper bound on the number of issuing queries that \mathcal{A} makes. In Game 1, the challenger remembers each $k^{(i)}$ that it samples in an issuing query, and we change the winning condition so that \mathcal{A} only wins if μ^* is a valid macaroon under key $k^{(i)}$ for some $i \in \{1, \dots, q_1\}$. In other words, \mathcal{A} only wins if it produces a valid macaroon under a key that the challenger sampled in an issuing query.

Winning condition: The adversary \mathcal{A} outputs a macaroon μ^* and a set of discharges $M^* = \{\mu_j^*\}$. He wins if:

1. $\text{Verify}(k_{\mathcal{T}}, c, \sigma) = 1$ and $\text{Verify}_{\mathcal{M}}(k^{(i)}, \mu) = 1$ for some $i \in \{1, \dots, q_1\}$, and
2. For every third-party caveat $3 \parallel c_i \parallel \alpha_S^{(i)} \parallel \alpha_D^{(i)}$ in μ , there exists $\mu_j \in M$ with key identifier $\alpha_D^{(i)}$ such that $\text{Verify}_{\mathcal{D}}(\text{sk}_{\mathcal{T}}^{(i)}, \alpha_S^{(i)}, \mu_j) = 1$ where $\text{sk}_{\mathcal{T}}^{(i)} := \text{Dec}(\text{sk}_2, c_i)$, and
3. Either
 - (a) μ^* is not an extension of a macaroon given to \mathcal{A} in a macaroon query, or

- (b) There exists a macaroon $m_j^* \in M^*$ corresponding to a third-party caveat $3||c_j||\alpha_S^{(j)}||\alpha_D^{(j)}$ in μ^* such that μ_j^* is not an extension of any macaroon μ such that $(\alpha_D^{(j)}, \mu) \in \mathcal{Q}_D$.

We argue that $|p_0 - p_1| = \text{negl}(\kappa)$ by the security of the MAC \mathcal{T} . Consider the probability ε that \mathcal{A} outputs a valid forgery μ^* with key identifier $\tilde{\text{id}}^* = c^*||\sigma^*||\dots$ such that $c^* \neq c^{(i)}$ for all $i \in \{1, \dots, q_1\}$, or in other words, that c^* is a ciphertext that was not given to \mathcal{A} as a response to an issue query. It is easy to see that $|p_0 - p_1| = \varepsilon$ since by correctness of decryption of \mathcal{E}_1 , the games are equivalent if $c^* = c^{(i)}$ for some $i \in \{1, \dots, q_1\}$. By the security of the MAC \mathcal{T} , we must have $\varepsilon = \text{negl}(\kappa)$. The reduction samples $\text{sk}_1, \text{pk}_2, \text{sk}_2$, answers token queries by sampling $k^{(i)}$ and computing $c^{(i)}$ on its own and using the MAC tagging oracle to obtain $\sigma^{(i)}$. It answers extend, macaroon, and discharge queries on its own. Given \mathcal{A} 's forgery μ^* with key identifier $\tilde{\text{id}}^* = c^*||\sigma^*||\dots$, it outputs σ^* as a signature forgery on message c^* . Since this simulates \mathcal{A} 's view perfectly, the reduction produces a tag on a new message with probability ε .

Game 2: In Game 2, the challenger chooses an index $i_{\text{ch}} \in \{1, \dots, q_1\}$ uniformly at random. We change the winning condition so that \mathcal{A} only wins if μ^* is a valid macaroon under key $r^{(i_{\text{ch}})}$.

Winning condition: The adversary \mathcal{A} outputs a macaroon μ^* and a set of discharges $M^* = \{\mu_j^*\}$. He wins if:

1. $\text{Verify}(k_{\mathcal{T}}, c, \sigma) = 1$ and $\text{Verify}_{\mathcal{M}}(k^{(i_{\text{ch}})}, \mu) = 1$, and
2. For every third-party caveat $3||c_i||\alpha_S^{(i)}||\alpha_D^{(i)}$ in μ , there exists $\mu_j \in M$ with key identifier $\alpha_D^{(i)}$ such that $\text{Verify}_{\mathcal{D}}(\text{sk}_{\mathcal{T}}^{(i)}, \alpha_S^{(i)}, \mu_j) = 1$ where $\text{sk}_{\mathcal{T}}^{(i)} := \text{Dec}(\text{sk}_2, c_i)$, and
3. Either
 - (a) μ^* is not an extension of a macaroon given to \mathcal{A} in a macaroon query, or
 - (b) There exists a macaroon $m_j^* \in M^*$ corresponding to a third-party caveat $3||c_j||\alpha_S^{(j)}||\alpha_D^{(j)}$ in μ^* such that μ_j^* is not an extension of any macaroon μ such that $(\alpha_D^{(j)}, \mu) \in \mathcal{Q}_D$.

Since i_{ch} is chosen uniformly at random, we have that $p_2 = p_1/q_1$.

Game 3: In Game 3, we change how the challenger answers the i_{ch} th issuing query. Instead of encrypting $k^{(i_{\text{ch}})}$, it encrypts 0.

Issue Queries: On the i_{ch} th query, the challenger creates a macaroon $(\mu_0^{(i_{\text{ch}})}, \text{ek}_0^{(i_{\text{ch}})})$ and adds it as a child of the root under an edge labeled with $c^{(i_{\text{ch}})}||\sigma^{(i_{\text{ch}})}$, where:

$$k^{(i_{\text{ch}})} \leftarrow \text{KeyGen}_{\mathcal{M}}(1^\kappa) \quad , \quad c^{(i_{\text{ch}})} \leftarrow \text{Enc}(\text{sk}_1, 0) \quad , \quad \sigma^{(i_{\text{ch}})} \leftarrow \text{Mac}(k_{\mathcal{T}}, c^{(i_{\text{ch}})}) \quad , \quad \text{id}^{(i)} := c^{(i)}||\sigma^{(i)}$$

$$(\mu_0^{(i_{\text{ch}})}, \text{ek}_0^{(i_{\text{ch}})}) \leftarrow \text{Issue}_{\mathcal{M}}(k^{(i_{\text{ch}})}, \text{id}^{(i_{\text{ch}})})$$

We argue that $|p_2 - p_3| = \text{negl}(\kappa)$ by the CPA-security of \mathcal{E}_1 . The reduction samples $\text{pk}_2, \text{sk}_2, k_{\mathcal{T}}$ on its own. For $i \neq i_{\text{ch}}$, it answers the i th issuing query by sampling $k^{(i)}$ and asking its encryption oracle for $c^{(i)}$. It then computes $\sigma^{(i)}$ and $(\mu_0^{(i)}, \text{ek}_0^{(i)})$ on its own. On the i_{ch} th issuing query, it samples $k^{(i_{\text{ch}})}$ and sends $m_0 := k^{(i_{\text{ch}})}, m_1 := 0$ to the encryption challenger as the challenge plaintexts, and receives a challenge ciphertext $\text{Enc}_1(\text{sk}_1, m_b)$. It uses the challenge ciphertext as $c^{(i_{\text{ch}})}$ and computes $\sigma^{(i)}$ on its own. The reduction answers extend, discharge, and macaroon queries on its own. Notice that if $b = 0$, the reduction perfectly simulates Game 2, whereas if $b = 1$, the reduction perfectly simulates Game 3. Therefore, by the CPA-security of \mathcal{E}_1 we must have $|p_2 - p_3| = \text{negl}(\kappa)$.

Game 4: Let q_E be an upper on the number of extension queries made by \mathcal{A} that add a third-party caveat. In Game 4, the challenger chooses an index $j_{\text{ch}} \in \{1, \dots, q_E\}$ uniformly at random. We change the winning condition so that \mathcal{A} only wins if either μ^* is not an extension of any macaroon μ revealed in a macaroon query, or it is such an extension, the j_{ch} th extension query added a third-party caveat $3||c_{j_{\text{ch}}}\alpha_S^{(j_{\text{ch}})}||\alpha_D^{(j_{\text{ch}})}$ caveat to μ , and there exists a macaroon $\mu^{(j)} \in M^*$ that is a valid discharge for this caveat and $\mu^{(j)}$ was not the output of a discharge query for token $\alpha_D^{(j_{\text{ch}})}$.

Winning condition: The adversary \mathcal{A} outputs a macaroon μ^* and a set of discharges $M^* = \{\mu_j^*\}$. He wins if:

1. $\text{Verify}(k_{\mathcal{T}}, c, \sigma) = 1$ and $\text{Verify}_{\mathcal{M}}(k^{(i_{\text{ch}})}, \mu) = 1$, and
2. For the third-party caveat $3||c_{j_{\text{ch}}}||\alpha_{\mathcal{S}}^{(j_{\text{ch}})}||\alpha_{\mathcal{D}}^{(j_{\text{ch}})}$ in μ , there exists $\mu_j \in M$ with key identifier $\alpha_{\mathcal{D}}^{(j_{\text{ch}})}$ such that $\text{Verify}_{\mathcal{D}}(\text{sk}_{\mathcal{T}}^{(j_{\text{ch}})}, \alpha_{\mathcal{S}}^{(j_{\text{ch}})}, \mu_j) = 1$ where $\text{sk}_{\mathcal{T}}^{(j_{\text{ch}})}$ is the key sampled by the challenger in the j_{ch} th extension query, and
3. Either
 - (a) μ^* is not an extension of a macaroon given to \mathcal{A} in a macaroon query, or
 - (b) $\mu^{(j)}$ was not the output of a discharge query for token $\alpha_{\mathcal{D}}^{(j_{\text{ch}})}$.

Since j_{ch} is chosen uniformly at random, we have that $p_4 \geq p_3/q_{\mathcal{E}}$.

Game 5 : In Game 5, we change how the j_{ch} th extension query is answered. If is a query to add a third-party caveat, instead of encrypting $\text{sk}_{\mathcal{T}}^{(j_{\text{ch}})}$, the reduction encrypts 0.

Extend Queries: On the j_{ch} th query that adds a third-party caveat, the adversary sends an identifier $\text{id}^{(j_{\text{ch}})}$, a tuple of caveats $\vec{m}^{(j_{\text{ch}})}$ to the challenger. If $(\text{id}^{(j_{\text{ch}})}, \vec{m}^{(j_{\text{ch}})})$ is a valid path in \mathcal{T} leading to a node (μ, ek) , then the challenger computes $(\mu', \text{ek}') \text{Extend}_{\mathcal{M}}(\text{ek}, \mu, 3||c^{(j_{\text{ch}})}||\alpha_{\mathcal{S}}^{(j_{\text{ch}})}||\alpha_{\mathcal{D}}^{(j_{\text{ch}})})$ and adds (μ', ek') to \mathcal{T} as a child of (μ, ek) , where:

$$\text{sk}_{\mathcal{T}}^{(j_{\text{ch}})} \leftarrow \text{SetupToken}_{\mathcal{D}}(1^{\kappa}) \quad , \quad c^{(i)} \leftarrow \text{Enc}_2(\text{pk}_2, 0) \quad , \quad (\alpha_{\mathcal{S}}^{(j_{\text{ch}})}, \alpha_{\mathcal{D}}^{(j_{\text{ch}})}) \leftarrow \text{Token}_{\mathcal{D}}(\text{sk}_{\mathcal{T}})$$

We argue that $|p_4 - p_5| = \text{negl}(\kappa)$ by the CPA-security of \mathcal{E}_2 .

Game 6 : In Game 6, we change the winning condition so that the adversary only wins if it outputs a macaroon μ^* that is valid under key $k^{(i_{\text{ch}})}$, and is not an extension of a macaroon revealed to \mathcal{A} in a macaroon query.

Winning condition: The adversary \mathcal{A} outputs a macaroon μ^* and a set of discharges $M^* = \{\mu_j^*\}$. He wins if:

1. $\text{Verify}(k_{\mathcal{T}}, c, \sigma) = 1$ and $\text{Verify}_{\mathcal{M}}(k^{(i_{\text{ch}})}, \mu) = 1$, and
2. μ^* is not an extension of a macaroon given to \mathcal{A} in a macaroon query.

We argue that $|p_5 - p_6| = \text{negl}(\kappa)$ by the security of the macaroon discharge protocol \mathcal{D} . Let E be the event that \mathcal{A} outputs m^*, M^* such that for the third-party caveat $3||c_{j_{\text{ch}}}||\alpha_{\mathcal{S}}^{(j_{\text{ch}})}||\alpha_{\mathcal{D}}^{(j_{\text{ch}})}$ in μ , there exists $\mu_j \in M$ with key identifier $\alpha_{\mathcal{D}}^{(j_{\text{ch}})}$ such that $\text{Verify}_{\mathcal{D}}(\text{sk}_{\mathcal{T}}^{(j_{\text{ch}})}, \alpha_{\mathcal{S}}^{(j_{\text{ch}})}, \mu_j) = 1$ where $\text{sk}_{\mathcal{T}}^{(j_{\text{ch}})}$ is the key sampled by the challenger in the j_{ch} th extension query, and $\mu^{(j)}$ was not the output of a discharge query for token $\alpha_{\mathcal{D}}^{(j_{\text{ch}})}$. Then $|p_5 - p_6| \leq \Pr[E]$, and we claim $\Pr[E] = \text{negl}(\kappa)$ by the security of \mathcal{D} .

The reduction samples $\text{sk}_1, \text{sk}_2, \text{pk}_2, k_{\mathcal{T}}$ on its own and answers issuing and macaroon queries on its own. It also answers extend queries on its own if the query is adding a first-party caveat. It answers all extension queries on its own, except the j_{ch} th extension query that adds a third-party caveat. It answers the j_{ch} th extension query that adds a third-party caveat by calling the token oracle and receiving $\alpha_{\mathcal{S}}^{(j_{\text{ch}})}, \alpha_{\mathcal{D}}^{(j_{\text{ch}})}$. It answers all discharge queries by calling its discharge oracle. When \mathcal{A} outputs μ^*, M^* , then the reduction looks for a discharge $\mu^{(j)} \in M^*$ with key identifier $\alpha_{\mathcal{D}}^{(j_{\text{ch}})}$ and submits $\alpha_{\mathcal{S}}^{(j_{\text{ch}})}, \alpha_{\mathcal{D}}^{(j_{\text{ch}})}, \mu^{(j)}$ as its discharge forgery. If there is more than one, it picks one at random. If the event E occurs, then there is at least one valid forgery macaroon in M^* and the reduction picks it with probability $1/|M^*|$. Therefore, the probability that the reduction outputs a valid forgery is at least $\Pr[E]/|M^*|$. Since $|M^*| = \text{poly}(\kappa)$, by the security of the discharge protocol \mathcal{D} , we must have $\Pr[E] = \text{negl}(\kappa)$, as desired.

A Can't Win : We argue that $p_6 = \text{negl}(\kappa)$ by the security of the macaroon scheme \mathcal{M} . The reduction samples $\text{sk}_1, \text{sk}_2, \text{pk}_2, k_{\mathcal{T}}, \text{params}_{\mathcal{D}}, \text{sk}_{\mathcal{D}}$ on its own. On the i th issue query, if $i \neq i_{\text{ch}}$, the reduction answers it on its own as well. On the i_{ch} th query, the reduction computes $c_{i_{\text{ch}}}$ and $\sigma_{i_{\text{ch}}}$ on its own (note that this does not require it to know $k_{i_{\text{ch}}}$), and calls its issuing oracle with $\text{id}_{i_{\text{ch}}} = c_{i_{\text{ch}}||\sigma_{i_{\text{ch}}}$. The reduction answers extend queries by first checking if $\text{id}^* = \text{id}_{i_{\text{ch}}}$. If it is, then it uses its extension oracle to complete the query, otherwise it computes everything on its own. It answers discharge queries on its own. When \mathcal{A} outputs a macaroon forgery μ^* , then the reduction outputs μ^* as well. It is easy to see that if \mathcal{A} wins, then the reduction outputs a valid forgery. Therefore, by the security of the macaroon scheme, we must have $p_6 = \text{negl}(\kappa)$.

□

We wish to prove that $p_0 = \text{negl}(\kappa)$. We know that $|p_0 - p_1| = \text{negl}(\kappa)$, $p_1 = q_1 \cdot p_2$, $|p_2 - p_3| = \text{negl}(\kappa)$, $p_3 \leq p_4 \cdot q_E$, $|p_4 - p_5| = \text{negl}(\kappa)$, $|p_5 - p_6| = \text{negl}(\kappa)$, and $p_6 = \text{negl}(\kappa)$. Putting it all together, we have:

$$|p_0 - p_1| + q_T \cdot |p_2 - p_3| + q_T \cdot q_E \cdot |p_4 - p_5| + q_T \cdot q_E \cdot |p_5 - p_6| \geq$$

$$(p_0 - p_1) + q_T \cdot (p_2 - p_3) + q_T \cdot q_E \cdot (p_4 - p_5) + q_T \cdot q_E \cdot (p_5 - p_6) = p_0 - q_T \cdot q_E \cdot p_6$$

Since q_T and q_E are numbers of queries that \mathcal{A} makes and \mathcal{A} runs in polynomial time, we must have that $q_T = \text{poly}(\kappa)$ and $q_E = \text{poly}(\kappa)$. We can therefore conclude that

$$p_0 \leq |p_0 - p_1| + q_T \cdot |p_2 - p_3| + q_T \cdot q_E \cdot |p_4 - p_5| + q_T \cdot q_E \cdot |p_5 - p_6| + q_T \cdot q_E \cdot p_6 = \text{negl}(\kappa),$$

as required.