Token meets Wallet:
Formalizing Privacy and Revocation for FIDO2

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

Julian Loss
CISPA Helmholtz Center
for Information Security
loss@cispa.de

Benedikt Wagner
CISPA Helmholtz Center
for Information Security
Saarland University
benedikt.wagner@cispa.de

Abstract—**The FIDO2 standard is a widely-used class of challenge-response type protocols that allows to authenticate to an online service using a hardware token. Barbosa et al. (CRYPTO '21) provided the first formal security model and analysis for the FIDO2 standard. However, their model has two shortcomings: (1) It does not include privacy, one of the key features claimed by FIDO2. (2) It only covers tokens that store *all secret keys locally*. In contrast, due to limited memory, most existing FIDO2 tokens either derive all secret keys from a common seed or store keys on the server (the latter approach is also known as *key wrapping*).**

In this paper, we revisit the security of the WebAuthn component of FIDO2 as implemented in practice. Our contributions are as follows. (1) We adapt the model of Barbosa et al. so as to capture authentication tokens using key derivation or key wrapping. (2) We provide the *first formal definition of privacy for the WebAuthn component of FIDO2*. We then prove the privacy of this component in common FIDO2 token implementations if the underlying building blocks are chosen appropriately. (3) We address the unsolved problem of *global key revocation* in FIDO2. To this end, we introduce and analyze a simple revocation procedure that builds on the popular BIP32 standard used in cryptocurrency wallets and can efficiently be implemented with existing FIDO2 servers.

Index Terms—**FIDO2; BIP32; unlinkability; revocation;**

1. Introduction

Online authentication is one of the most pressing challenges faced by security engineers and cryptographers today. Reliable authentication is an important concern for both the security of user's accounts as well as the reputation of service providers. A simple way to strengthen the security of an authentication process is to introduce additional authentication factors. Usually, the user has to just provide a login and password (*something she knows*). A popular way to introduce a second factor is to use a device (*something she has*) that is registered to the user's account. Universal Second Factor (U2F) (or CTAP1, as it is currently named) is a protocol to achieve two-factor authentication using a designated device that we refer to as a *token*. The token runs a simple piece of code and interacts with the user's interface, e.g., a web browser. We refer to this interface as the *client* or *agent*. The agent acts as a proxy device during the authentication process between a token and a server that we call the *relying party*. The main benefit of this solution over the login/password approach is the protection against phishing attacks and database breaches. CTAP1/U2F and its successor CTAP2 are part of the FIDO specification [1]. Together with the complement W3C web authentication [2] (WebAuthn) they form the state-of-the-art for online authentication using security tokens.

In recent work, Barbosa et al. [3] gave the first formal model for token-based authentication and provided a security proof for the FIDO2 standard. While their model provides an important starting point for further exploration of FIDO2's security properties, it does not accurately model several key aspects of FIDO2 as used in practice. In this work, we revisit the FIDO2 standard and give a more complete security analysis of its security features. We also augment the existing standard with a new feature that allows to easily revoke keys of compromised tokens. Before presenting our contributions in more detail, we first explain several disparities between the model of Barbosa et al. and how FIDO2 is commonly used in practice.

Attestation. Barbosa et al. consider a model where the token and the relying party are initially provided a pair of attestation secret key and attestation public key by the manufacturer. Thereby, servers can verify that they interact with tokens from trusted manufacturers. While attestation is common practice for critical interactions such as banking transactions, it is rarely used for more mundane scenarios such as logging in to an internet account [4]. Also, naive attestation violates privacy [5]. Therefore, to reason about the security and privacy of FIDO2 in its most widely used form, we consider a model *without attestation*.

Privacy. There are two ways to implement FIDO2 tokens. The first is to store all secret keys of the token locally, i.e. on the authentication token. In practice however, keys are almost exclusively stored *externally* on the server [6], [7]. External key storage can be implemented securely using so called *key wrapping*, where an encryption of the key is stored with the server. Another approach is to use a *key derivation*

function to derive keys on the fly. Throughout, we refer to both key wrapping and key derivation function approaches as *external key storage*.

Since tokens are limited with regards to storage space, externally stored keys provide a distinct storage advantage over locally stored keys. On the downside, externally stored keys may give the server additional information that it may use to link separate sessions of the token.

Let us illustrate this using a (pathological) implementation of key wrapping. Assume that the server stores ciphertexts of the form $(\mathsf{Enc}(k, \mathsf{sk}), H(k))$, where Enc is a symmetric encryption scheme, H is a hash function, k denotes the encryption key, and sk denotes the signing key used for authentication. Observe that the above scheme appends a hash of the secret encryption key to the ciphertext of Enc. Intuitively, in the random oracle model, the above scheme is as secure as Enc. However, ciphertexts produced by the same token (i.e. using the same key k) can trivially be linked using the second component.

The above discussion shows that the use of externally stored keys leads to unexpected subtleties with regards to privacy across different sessions of the same token. Unfortunately, as Barbosa et al. consider exclusively the setting of locally stored keys, their model does not provide any treatment of privacy. Thus, proving the privacy properties claimed by FIDO2 remains an open problem.

1.1. Our Contribution

In this work, we present a new security model for the WebAuthn component of FIDO2. We analyze existing FIDO2 implementations in this model. In a second step, we identify a lack of practical key revocation in these implementations that is compatible with our privacy goals. We now explain our contributions in more detail.

New Security Models for Privacy and Impersonation. Our first contribution is to augment the model of Barbosa et al. with a suitable notion of privacy, and analyzing existing token implementations using key wrapping and key derivation functions in this augmented model. Concretely, we define a natural security game with three different winning conditions, leading to *strong, medium and weak unlinkability*. Intuitively, the winning condition for strong unlinkability only disallows a generic linking attack that no token implementation can mitigate [6]. For medium and weak unlinkability, we slightly weaken the restrictions on this attack. The results of our analysis in these models are given in the top part of Table 1. Notably, while key wrapping provides strong unlinkability if we rely on *anonymous* authenticated encryption, our analysis shows that implementations using the key derivation function approach provide only medium unlinkability. In this sense, our privacy notions show exactly which linking attacks are possible for these implementation, and reveal a gap between the key wrapping and key derivation approaches for externally storing keys. On top of these new privacy definitions, we also define an appropriate notion of security against impersonation in the

context of externally stored keys and show that existing implementations satisfy it.

Adding Global Revocation. When using FIDO2, a crucial aspect is how to securely revoke keys in case access to the authentication token is ever lost. Indeed, the usability study by Lyastani et al. [8] shows that one of the top concerns of users is that an adversary can gain access to their account in case their FIDO2 token gets lost or is stolen. For all current token implementations, a user has to revoke all keys separately by logging in at each server, which is highly impracticable. To address this revocation problem in a usable way, we desire a procedure that allows to simultaneously revoke *all keys* associated with an token without communicating with each server individually. This includes also any future keys that the (compromised) token may attempt to register with the relying party. Ideally, revocation is done by globally publishing some short revocation information, which is then noticed by all servers. The revocation information can be stored externally (e.g. on a piece of paper) by the user in a secure location. We refer to such a procedure as *global revocation*. In terms of security, we desire three properties: First, as long as the revocation information is not published, all credentials should be unlinkable to guarantee privacy. Second, we demand security against impersonation even after publishing the revocation information. This is of great importance, as otherwise an attacker could actively look for such published revocation information and impersonate a user before a server that periodically checks for revocations may notice. Finally, it should be difficult for a malicious party to launch a denial of service (DoS) attack in which it revokes keys on behalf of an honest user. In Section 6 we formalize these security properties of global revocation.

As we observe, however, achieving these properties for global revocation is a challenging endeavour. To get an intuition why, consider a naive approach where every credential of a user is prefixed with some common tag t. To revoke, one can simply publish t. This approach clearly links all of the prefixed credentials and hence violates any reasonable notion of privacy. Additionally, a DoS attack as mentioned above is possible, because t is not kept secret. This shows the technical challenge introduced by global revocation: On the one hand, keys can not be generated independently, as the published revocation information has to identify all keys. On the other hand, generating keys using some correlated randomness may violate the users privacy and security.

Crypto Wallets to the Rescue. Fortunately, we can rely on a practical solution from the cryptocurrency space to solve this issue. A common approach to store a multitude of keys compactly is to use *deterministic key derivation*[1]. Considering the case of BIP32, the most widely implemented procedure for deterministic key derivation, ECDSA keys are derived from a pair of master keys $\mathsf{msk} \in \mathbb{Z}_p, \mathsf{mpk} = g^{\mathsf{msk}}$ and a *chaincode* chain. Here, g denotes the base point of an elliptic curve and chain can be thought of as a random seed. To derive a fresh key pair for an identity id,

---

[1]Such a mechanism is colloquially referred to as a *wallet*.

Scheme	Unlinkability	Authentication	Revocation
kdfPA	Medium	Secure	Local
kwrPA	Strong	Secure	Local
bip32PA	Weak	Secure	Global/Local

TABLE 1. OVERVIEW OF OUR RESULTS. WE GIVE THE FIRST ANALYSIS OF EXISTING SCHEMES kdfPA, kwrPA, AND PROPOSE AND ANALYZE SCHEME bip32PA. IN LOCAL REVOCATION, THE USER HAS TO REVOKE ALL KEYS SEPARATELY, WHILE GLOBAL REVOCATION MEANS THAT THE USER CAN PUBLISH ONE STRING TO REVOKE ALL KEYS SIMULTANEOUSLY.

BIP32 first computes $w := \mathsf{Hash}(\mathsf{id}, \mathsf{chain}, \mathsf{mpk})$, then sets $\mathsf{sk_{id}} := \mathsf{msk} + w, \mathsf{pk_{id}} := \mathsf{mpk} \cdot g^w$. We observe that this procedure provides a simple means of global revocation. To revoke all public keys associated with a pair $\mathsf{msk}, \mathsf{mpk}$, one can simply publish mpk and chain. Now, each server can find deterministically recompute the public keys from mpk and chain, and revoke them. After introducing this new token implementation based on BIP32, we show that it satisfies the security and privacy properties that we defined. We highlight that proving that ECDSA signatures remain unforgeable with respect to keys derived in this fashion once chain is leaked turns out be non-trivial. In recent work, Das et al. were the first to consider this stronger form of unforgeability and showed that it holds for key-prefixed ECDSA given that no message is ever signed twice [9]. As we want our new token implementation to be compatible with existing FIDO2 server implementations, we can not rely on key-prefixing as it is used by Das et al. [9], and have to find a way to allow signing the same message multiple times. This makes the analysis challenging. For further details, we refer the reader to Section 7. Finally, we give experimental results that showcase the efficiency of our protocol during authentication and revocation.

1.2. Related Work

The first formal model for the FIDO2 protocol was proposed by Barbosa et al. [3]. The authors introduced the notion of passwordless authentication that models the WebAuthn protocol in the scenario when the token uses keys stored locally. They also present a security model for PIN-based access control that tries to formally define the CTAP protocol executed between the token and the client. Barbosa et al. show that CTAP only provides a weaker access control notion and proposes an alternative protocol based on password-authenticated key agreement (PAKE).

Bindel, Cremers, and Zhao [4] (to appear at S&P 2023) recently extended the Barbosa et al. model to support the latest specification, introducing CTAP in version 2.1. Their model uses explicit user verification assurance and guarantees stronger token binding properties. Relevant to our contribution is that they consider the case where attestation of the token is not used (attestation mode `None`) and argue that this is the most widely used mode. Unfortunately, their model, similar to [3], only works if the token permanently

stores all keys. They also do not consider any privacy definition that would model the desired property of cross relying party unlinkability of the token-generated credentials.

Guirat et al. [5] analyzed the WebAuthn protocol using automated verification and modeled a simple privacy definition. They showed that if the same key is used for attestation, it allows a server to link the same token. The security and privacy issues faced by developers during the implementation of FIDO and WebAuthn components were discussed by Alam et al. [10]. Kepkowski et al. [6] report that existing relying parties are implemented based on existing open-source solutions and, by default, support only externally stored keys. The authors also show a timing attack allowing a malicious relying party to link users across different services, proving that a formal definition and analysis of privacy is needed. Another line of related work provides alternative protocols and usability studies. Chakraborty et al. [11] used a TPM implementation (simTPM) based on a sim card as a secure and convenient FIDO2 token. A usability and acceptability study of FIDO2 was done by Lyastani et al. [8]. The authors report that one of the main concerns of users is that an adversary can access their account if they lose their token or it gets stolen. An efficient and usable revocation mechanism would solve some of those concerns and increase the acceptance rate among users.

Revocation in anonymous authentication and privacy-preserving signatures is not a new problem. Camenisch and Lysyanskaya [12] showed how cryptographic accumulators can be used to revoke users in an anonymous setting,. The system's authority provides users membership certificates, allowing them to prove zero-knowledge membership against a publicly available accumulator. Boneh and Shacham [13] used a different approach in the context of group signatures, the so-called verifier-local revocation. The verifier is the only party involved in the revocation process. Camenisch et al. [14] discussed revocation techniques for pseudonymous systems. Their solution allows checking the revocation status in logarithmic time in the number of revoked users, independently of all valid users in the system. A simple alternative approach to VLR is blocklisting in pseudonymous systems. During the interaction, users generate a domain/application-specific unlinkable pseudonym that the verifier compares with elements on the blocklist leading to the check taking logarithmic time in the number of revoked users.

Dauterman et al. [15] discussed the problem of backdooring FIDO tokens. They introduced the notion of verifiable identity families (VIF) and used it to generate FIDO identities. The key feature is that the tokens can prove an identity (i.e. public key) to be correct (given a relying party) and pseudorandom to any party holding a VIF master public key. It allows the user to verify that the token deterministically creates identities and a manufacturer did not implement any backdoor in the key generation algorithm. Their solution does not support revocation without per-server interaction since identities are not publicly linkable. To link identities, the token must create proof of correct VIF evaluation.

Frymann et al. [16] discussed the token backup procedure proposed by Yubico, a recovery solution in case of

token loss/theft. The solution allows a primary token to generate an identity that the backup device will use with a specific relying party. Thus, instead of registering both devices, a user can register once using the primary device and use the backup token if the former is lost. The problem of a backup device is orthogonal to revocation, i.e. in case the primary device gets stolen/lost, it still has to be revoked to protect the user's account.

From a cryptographic point of view, FIDO2 is a simple authentication scheme, a primitive which has been extensively studied. One of the first and well-known protocols is the Schnorr identification scheme [17], which provides impersonation resistance under the discrete logarithm assumption. Authenticated key exchange (AKE) protocols provide means to agree on a secret key between the interacting parties and simultaneously authenticate both parties [18], [19], [20]. A different approach to authentication is the folklore challenge-response protocol based on signature schemes, on which FIDO2 is based. This technique provides a framework for creating authentication protocols. Instantiated with group signatures [21], [22] the protocol provides a way for group members to authenticate to a server anonymously. A direct solution introduced by Teranishi, Furukawa, and Sako is anonymous authentication [23]. The extended access control protocol (EAC) [24], [25] is an alternative approach to authentication using a hardware device for machine-readable travel documents (e.g., e-Passports).

2. Notation and Preliminaries

We denote by $z \leftarrow \mathcal{A}(x)$ the execution of algorithm $\mathcal{A}$ on input x and with output z. If we want to make the random coins ρ of algorithm $\mathcal{A}$ explicit, we write $z := \mathcal{A}(x; \rho)$ instead. We write $y \in \mathcal{A}(x)$ to indicate that y is a possible output of $\mathcal{A}$ on input x. By $r \xleftarrow{\$} S$ we mean that r is chosen uniformly at random over the set S. We will use $[n]$ to denote the set $\{1, \ldots, n\}$. By λ, we denote the security parameter. Throughout the paper, we assume public parameters par are given implicitly to all algorithms. We will use standard notions of digital signatures and one of the rerandomizable signature schemes discussed in [26], [27]. To formally model the key wrapping technique, we will use a definition of symmetric key encryption that is both authenticated and anonymous [28]. We present their syntax and security definitions in more detail in Appendix A.

3. WebAuthn and its Implementations

In this section, we give an overview of the WebAuthn protocol contained in the FIDO2 standard, and introduce the syntax we need for our formal analysis. FIDO2 can be used as means of passwordless authentication or as a second-factor for the standard login-via-password scenario. The WebAuthn protocol template consists of a registration and an authentication process executed between a token and a server. In both registration and authentication, the protocol consists of two messages sent between the server and the

token, relayed through a client which is implemented e.g. by the user's browser. The first message is sent from the server to the token (via the client) and contains a challenge rs. This challenge is then signed by the token, and the resulting signature σ is sent back to the server. In the registration protocol, this second message also contains a public key pk, and a so called *credential identifier* cid. The server stores this key and additional information in a *credential* cred. In subsequent authentication interactions, the server first finds cid (e.g. by looking up the username), and then includes it in the first message of the above flow.

Different token implementations may generate key material in a different way. To accommodate for memory-constraints, the most commonly used implementations do not store the key material locally, but instead use the value cid to securely outsource it to the server. In this paper we will focus on the following two variants:

- *Key Derivation Function.* The signing key is generated in a pseudorandom way using a master secret key, the server's identifier, and the randomly chosen credential identifier. This method is e.g. used by the open-source FIDO2 token SoloKey[2].
- *Key wrapping.* The signing keys are encrypted together with the server's identifier and the ciphertext is the credential identifier. This method is e.g. used by Yubico in their implementation of FIDO2 tokens [7].

We will now describe the above two schemes in more detail. An overview (in our syntax) can be found in Figure 1. We refer to the two variants as kdfPA and kwrPA, respectively. WebAuthn makes use of a hash function $H : \{0,1\}^* \rightarrow \{0,1\}^{2\lambda}$, modeled as a random oracle.

Key Generation. Initialization of a token is done using the Gen algorithm that outputs a master secret key msk. For kdfPA this key is generated by executing $\text{msk} \xleftarrow{\$} \{0,1\}^{\lambda}$. In the case of the key wrapping scheme, this key is the secret key for a symmetric key encryption scheme SKE with length function $\nu : \mathbb{N} \rightarrow \mathbb{N}$. Thus, we set $\text{msk} \leftarrow \text{SKE.Gen(par)}$.

Registration. We assume that every server S holds a registration context rcs_S, which stores users credentials. During registration, the server S generates a random nonce $rs \xleftarrow{\$} \{0,1\}^{\lambda}$ and sends the challenge $c = (\text{id}_S, rs)$ to the client, where id_S is a server identifier. The client verifies that id_S is correct and sends $M_r = H(rs)$ and id_S to the token. Here, the token implementations kwrPA and kdfPA differ slightly. In kdfPA, the token first chooses a random identifier $\text{cid} \xleftarrow{\$} \{0,1\}^{\lambda}$. It then uses a pseudorandom key derivation function PRF to generate a secret key $\text{sk} := \text{PRF(msk}, (\text{cid}, \text{id}_S))$. We assume that PRF outputs values in the secret key space of the signature scheme. The token computes the corresponding public key $\text{pk} := \text{ToPK(sk)}$. In kwrPA the key pair is computed using the signature scheme key generation algorithm $(\text{sk}, \text{pk}) \leftarrow \text{Gen(par)}$. The secret key sk and id_S is then encrypted by the token as $\text{cid} := \text{Enc(msk}, (\text{id}_S, \text{sk}))$. We assume that pairs (sk, id_S) of signing secret keys and server identifiers all

[2]see https://github.com/solokeys/solo

have the same length $l_0 = l_0(\lambda) \in \mathbb{N}$, which is achieved using appropriate padding. We set $\nu^* := \nu(l_0)$. For both schemes the secret key sk is used to create the signature $\sigma \leftarrow \mathsf{Sig}(\mathsf{sk}, m)$, where $m := (H(\mathsf{id}_S), \mathsf{cid}, \mathsf{pk}, M_r)$. Finally, the token sends a response message $R_r = (\mathsf{pk}, \sigma)$ together with cid to the client, which forwards it to the server. The server verifies the token's response as follows. It aborts and stores no credential in case $\mathsf{Ver}(\mathsf{pk}, \sigma, m) = 0$, where $m := (H(\mathsf{id}_S), \mathsf{cid}, \mathsf{pk}, M_r)$. Otherwise, it accepts. This means that it adds pk to its registration context, i.e. $\mathsf{rcs}_S[\mathsf{cid}] := \mathsf{pk}$.

Authentication. For authentication, we assume that the server knows the credential identifier cid of interest. Finding cid depends on the concrete use case of WebAuthn, e.g. using FIDO2 as a second factor, or in a standalone fashion. For example, a user could enter its username, and the server holds a mapping from usernames to credential identifiers. The authentication protocol begins with the server generating a random nonce $rs \xleftarrow{\$} \{0,1\}^\lambda$ and sending the challenge $c = (\mathsf{id}_S, rs)$ and cid to the client. The client verifies that id_S is correct and sends the message $M_a := H(rs)$, the server identifier id_S and identifier cid to the token. Now, in the first step the token recreates the signing key sk that it created during registration. For kdfPA this means that the token runs $\mathsf{sk} := \mathsf{PRF}(\mathsf{msk}, (\mathsf{cid}, \mathsf{id}_S))$. In case of the kwrPA the token first decrypts the cid to receive $(\mathsf{id}, \mathsf{sk}) := \mathsf{Dec}(\mathsf{msk}, \mathsf{cid})$. It then checks if this secret key correspond to the server, i.e. it returns an error if $\mathsf{id} \neq \mathsf{id}_S$. For both schemes the secret key sk is used to create the signature $\sigma \leftarrow \mathsf{Sig}(\mathsf{sk}, m)$, where $m := (H(\mathsf{id}_S), M_a)$. Finally, the token sends $R_a = \sigma$ to the client, which forwards it to the server. The server verifies the token's response as follows. First, the server uses the registration context rcs to get the token's public key $\mathsf{pk} \leftarrow \mathsf{rcs}_S[\mathsf{cid}]$. Then, it sets $m := (H(\mathsf{id}_S), H(rs))$ and accepts if $\mathsf{Ver}(\mathsf{pk}, \sigma, m) = 1$.

Formal Syntax. In our model, we consider parties $\mathcal{P} = \mathcal{T} \cup \mathcal{S}$, partitioned into the set of tokens $\mathcal{T}$ and the set of servers $\mathcal{S}$. Each server $S \in \mathcal{S}$ keeps as a registration context rcs_S as internal state, which is an initially empty key-value table. When new tokens register, an entry to rcs_S is added. Each token $T \in \mathcal{T}$ keeps a fixed state that is initialized once with a key msk_T and does not change. We do not model clients explicitly. Instead, we allow the adversary to do the computation that the clients do. As in [3], we assume that each server has a unique identifier id_S. In reality, id_S corresponds to a URL, which justifies this assumption. In all experiments that we define, we assume that these identifiers are given to all parties. We do not explicitly model the initial contact of user and server, as this may differ depending on the use case. That means that we assume that the server already knows the users account and therefore its cid.

We postpone the formal syntax and completeness definition to Appendix B, and give a short overview here. The reader may also consider Figure 1 to familiarize with the syntax. A *passwordless authentication scheme* (PlA) is a tuple $\mathsf{PLA} = (\mathsf{Gen}, \mathsf{Reg}, \mathsf{Auth})$, where $\mathsf{Gen}(\mathsf{par})$ generates a master secret key msk, stored on a token, and Reg and

Auth are the registration and authentication protocols. Registration is as follows: The server runs $(c, \mathsf{st}) \leftarrow \mathsf{rchall}(\mathsf{id}_S)$ and the client runs $M_r \leftarrow \mathsf{rcomm}(\mathsf{id}_S, c)$. Then, the token runs $(\mathsf{cid}, R_r) \leftarrow \mathsf{rresp}(\mathsf{msk}, \mathsf{id}_S, M_r)$ which generates a credential identifier cid. Finally, the server runs $(b_r, \mathsf{cred}) \leftarrow \mathsf{rcheck}(\mathsf{st}, \mathsf{cid}, R_r)$, where cred is the credential that will then be stored on the server as $\mathsf{rcs}[\mathsf{cid}] := \mathsf{cred}$. Authentication is done analogously, except that the token algorithm aresp takes cid as input instead of outputting it, and acheck only returns a bit and no credential cred.

4. Modeling Security and Privacy

In this section, we formally define security and privacy models for WebAuthn. While privacy is not considered in [3], our security model is closely related to their model. Before presenting our models, we first define server and token oracles an adversary may access. These oracles model the capability of an adversary to freely communicate with tokens and servers, and will be used in all of our security definitions. For server S, rcs_S is the registration context, i.e the set of credentials it stores, st_S is used to model the state transferred between algorithms rchall, achall and rcheck, acheck, respectively, and the map C_S is used for bookkeeping to bind registration interactions to authentication interactions.

Definition 1 (Server and Token Oracles). *Let $\mathcal{A}$ be an algorithm and $\mathsf{PLA} = (\mathsf{Gen}, \mathsf{Reg}, \mathsf{Auth})$ be a PlA. We associate each party $P \in \mathcal{T} \cup \mathcal{S}$ with a set of handles $\pi_P^{i,j}$ which model two types of instances corresponding to registration and authentication. Each party is represented by a number of these instances. Concretely, we refer to $\pi_P^{i,j}$ for $j = 0$ as the ith registration instance of party P and for $j \geq 1$ as the jth authentication instance of P corresponding to the ith registration.*

We assume that for each token $T \in \mathcal{T}$, a secret key $\mathsf{msk}_T \leftarrow \mathsf{Gen}(\mathsf{par})$ is given, and that for each server $S \in \mathcal{S}$, key-value tables rcs_S, C_S, and st_S are given. Per default, these are empty. Then, adversary $\mathcal{A}$ has access to oracles Start, Challenge, Complete *as follows.*

- Start$(\pi_S^{i,j})$*: This executes $(c, \mathsf{st}) \leftarrow \mathsf{rchall}(\mathsf{id}_S)$ in case $j = 0$ or $(c, \mathsf{st}) \leftarrow \mathsf{achall}(\mathsf{id}_S)$ in case $j > 0$. The oracle sets $\mathsf{st}_S[i,j] := \mathsf{st}$, and returns c to $\mathcal{A}$.*
- Challenge$(\pi_T^{i,j}, \mathsf{id}_S, \mathsf{cid}, M)$*: This runs algorithm $(\mathsf{cid}, R_r) \leftarrow \mathsf{rresp}(\mathsf{msk}_T, \mathsf{id}_S, M)$ if $j = 0$ or $R_a \leftarrow \mathsf{aresp}(\mathsf{msk}_T, \mathsf{id}_S, \mathsf{cid}, M)$ if $j > 0$. The result $((\mathsf{cid}, R_r)$ or $R_a)$ is returned to $\mathcal{A}$. Note that the input cid is ignored for $j = 0$.*
- Complete$(\pi_S^{i,j}, \mathsf{cid}, R)$*: This aborts if Start$(\pi_S^{i,j})$ has not been queried before. If $j = 0$, it runs $(b, \mathsf{cred}) \leftarrow \mathsf{rcheck}(\mathsf{st}_S[i,j], \mathsf{cid}, R)$, sets $C_S[i] := \mathsf{cid}$, and $\mathsf{rcs}[\mathsf{cid}] := \mathsf{cred}$. If $j > 0$, it aborts if $\mathsf{cid} \neq C_S[i]$. Otherwise, it runs $b \leftarrow \mathsf{acheck}(\mathsf{st}_S[i,j], \mathsf{rcs}_S, \mathsf{cid}, R)$. In both cases, b is returned to $\mathcal{A}$.*

We assume that for each $(i, j, T, S) \in \mathbb{N} \times \mathbb{N} \times \mathcal{T} \times \mathcal{S}$, the oracles Start$(\pi_S^{i,j})$, Challenge$(\pi_T^{i,j}, \cdot, \cdot, \cdot)$*, and* Complete$(\pi_S^{i,j}, \cdot, \cdot)$ *are executed only once.*

Token	Client	Server
$(\mathsf{cid}, R_r) \leftarrow \mathsf{rresp}(\mathsf{msk}, \mathsf{id}_S, M_r):$	$M_r \leftarrow \mathsf{rcomm}(\mathsf{id}_S, c):$	$(c, \mathsf{st}) \leftarrow \mathsf{rchall}(\mathsf{id}_S):$
$\mathsf{cid} \xleftarrow{\$} \{0,1\}^\lambda$	$(\mathsf{id}, rs) := c$	$rs \xleftarrow{\$} \{0,1\}^{\geq \lambda}$
$\mathsf{sk} := \mathsf{PRF}(\mathsf{msk}, (\mathsf{cid}, \mathsf{id}_S))$ **/** kdfPA $\quad \xleftarrow{\mathsf{id}_S, M_r}$	**if** $\mathsf{id} \neq \mathsf{id}_S$: **abort** $\quad \xleftarrow{\quad c \quad}$	$c := \mathsf{st} := (\mathsf{id}_S, rs)$
$\mathsf{pk} := \mathsf{ToPK}(\mathsf{sk})$ **/** kdfPA	$M_r := H(rs)$	
$(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{Gen}(\mathsf{par})$ **/** kwrPA		$(b, \mathsf{cred}) \leftarrow \mathsf{rcheck}(\mathsf{st}, \mathsf{cid}, R_r)$
$\mathsf{cid} := \mathsf{Enc}(\mathsf{msk}, (\mathsf{id}_S, \mathsf{sk}))$ **/** kwrPA		$m := (H(\mathsf{id}_S), \mathsf{cid}, \mathsf{pk}, H(rs))$
$\sigma \leftarrow \mathsf{Sig}(\mathsf{sk}, (H(\mathsf{id}_S), \mathsf{cid}, \mathsf{pk}, M_r)) \quad \xrightarrow{\mathsf{cid}, R_r}$	$\xrightarrow{\mathsf{cid}, R_r}$	$b := \mathsf{Ver}(\mathsf{pk}, \sigma, m)$
$R_r := (\mathsf{pk}, \sigma)$		**if** $b = 0$: $\mathsf{cred} := \perp$
		else : $\mathsf{cred} := \mathsf{pk}$
$R_a \leftarrow \mathsf{aresp}(\mathsf{msk}, \mathsf{id}_S, \mathsf{cid}, M_a):$	$M_a \leftarrow \mathsf{acomm}(\mathsf{id}_S, c):$	$(c, \mathsf{st}) \leftarrow \mathsf{achall}(\mathsf{id}_S):$
$\mathsf{sk} := \mathsf{PRF}(\mathsf{msk}, (\mathsf{cid}, \mathsf{id}_S))$ **/** kdfPA $\;\xleftarrow{\mathsf{id}_S, \mathsf{cid}, M_a}$	$(\mathsf{id}, rs) := c \quad \xleftarrow{\mathsf{cid}, c}$	$rs \xleftarrow{\$} \{0,1\}^{\geq \lambda}$
$(\mathsf{id}, \mathsf{sk}) := \mathsf{Dec}(\mathsf{msk}, \mathsf{cid})$ **/** kwrPA	**if** $\mathsf{id} \neq \mathsf{id}_S$: **abort**	$c := \mathsf{st} := (\mathsf{id}_S, rs)$
if $\mathsf{id} \neq \mathsf{id}_S$: **abort** **/** kwrPA	$M_a := H(rs)$	$b \leftarrow \mathsf{acheck}(\mathsf{st}, rcs, \mathsf{cid}, R_a)$
$\sigma \leftarrow \mathsf{Sig}(\mathsf{sk}, (H(\mathsf{id}_S), M_a))$		$\mathsf{pk} := rcs[\mathsf{cid}]$
$R_a := \sigma \qquad\qquad \xrightarrow{\quad R_a \quad}$	$\xrightarrow{\quad R_a \quad}$	$m := (H(\mathsf{id}_S), H(rs))$
		$b := \mathsf{Ver}(\mathsf{pk}, \sigma, m)$

Figure 1. The WebAuthn registration (top) and authentication protocol (bottom) protocol for the variations kdfPA and kwrPA. The highlighted statements are only executed in the variation that is given in the respective comment. Functions V_t, V_s (cf. Definition 16) are given as $(H(\mathsf{id}_S), H(rs), \mathsf{cid})$.

The index i in token handles $\pi_T^{i,j}$ may seem artificial at first, as tokens are stateless. However, this index will simplify the following definition significantly. Namely, we follow the work of Bellare et al. [29] to define partnering of handles, which will be used in the winning condition of security experiments. Two handles $\pi_S^{i,j}$ and $\pi_T^{i',j'}$ are partnered if they share the same session identifier. One should think of partnered handles as if the adversary just forwarded messages between the respective oracles. Formally, the notion of session identifiers must be defined by the protocol. A natural choice is that the session identifiers correspond to the view of a party. Note that the choice of session identifiers influences the notion of security that is achieved. We give a formal definition of session identifiers and partnering in Appendix B.

4.1. Impersonation Security

We now define what it means for a passwordless authentication protocol to be secure against impersonation. Informally, we say that if this property holds then the token that registered must be used to authenticate against a server and a single interaction cannot be used to authenticate multiple times. In the security game, the adversary can interact with tokens and servers in an arbitrary way, using the oracles from Definition 1.

Definition 2 (Impersonation Security). *For PlA* PLA $=$ $(\mathsf{Gen}, \mathsf{Reg}, \mathsf{Auth})$ *and an adversary* $\mathcal{A}$, *we define the experiment* $\mathbf{Imp}_{\mathsf{PLA}}^{\mathcal{A}}$ *as follows.*

- **Setup.** *For each token* $T \in \mathcal{T}$, *a key is generated by running* $\mathsf{msk}_T \leftarrow \mathsf{Gen}(\mathsf{par})$.
- **Online Phase.** *The adversary is allowed to interact with the oracles* Start, Challenge, Complete *as in Definition 1.*
- **Output Phase.** *Finally,* $\mathcal{A}$ *terminates, and the experiment outputs* 1 *if and only if there exists a server handle* $\pi_S^{i,j}$ *for* $j > 0$ *such that the following conditions hold:*

  1) $\pi_S^{i,0}$ *is partnered with a token handle* $\pi_T^{k,0}$.
  2) $\pi_S^{i,j}$ *accepted, i.e. in call* Complete$(\pi_S^{i,j}, \mathsf{cid}, R)$, *algorithm* $\mathsf{acheck}(\mathsf{st}_S[i,j], rcs_S, \mathsf{cid}, R)$ *returned* 1.
  3) $\pi_S^{i,j}$ *is not partnered with any token handle* $\pi_T^{i',j'}$, *or it is partnered with a token handle, which is partnered with a different server handle* $\pi_{S'}^{i'',j''}$.

We define the advantage of $\mathcal{A}$ *in winning the experiment as:*

$$\mathbf{Adv}_{\mathbf{Imp}, \mathsf{PLA}}^{\mathcal{A}} := \Pr[\mathbf{Imp}_{\mathsf{PLA}}^{\mathcal{A}} = 1].$$

Let us explain the three winning conditions that are defined. By the first condition, we know that the token T registered at server S. By the second and the third condition we know that the adversary was able to authenticate at server S for that particular registration. We highlight that this models a "trust-on-first-use" concept. Namely, servers consider the user communicating during registration as the honest user. In addition to the stateless nature of tokens, this is another difference to the model by Barbosa et al. [3], who use attestation keys as their trust anchor.

4.2. Unlinkability

We define unlinkability for passwordless authentication. Informally, we want that interactions of the same token are unlinkable, and different registrations of the same token can not be linked. Data that is exchanged outside of the protocol is out of the scope of our definition, e.g. metadata that could be used to link interactions of the token.

The Unlinkability Experiment. To model the above goal in a security experiment, we let the adversary $\mathcal{A}$ interact with all oracles in Definition 1 to get a global view of the system. Then, $\mathcal{A}$ targets two tokens T_0, T_1, and gets two additional oracles Left, Right, running T_b and T_{1-b} internally, for a random bit b. The goal of $\mathcal{A}$ is to determine which token is used in which oracle. That means that if $\mathcal{A}$ can link interactions with the same token, it can for example link the interaction with oracle Left to the interaction with token oracle Challenge for token T_b. This would allow him to win the game. Also, note that unlinkability can only make sense if there are at least two tokens used in the system, which is the reason why $\mathcal{A}$ has to output two tokens.

To prevent trivial linking, we have to introduce two winning conditions. The first one is a consequence of our definition of oracles and ensures that $\mathcal{A}$ can not violate the assumption in Definition 1. For the second condition, namely credential separation, we define three variants, leading to strong, medium, and weak unlinkability guarantees. Due to the stateless nature of tokens, the behavior of a token is determined by cid and id_S. This leads to the following trivial attack: Adversary $\mathcal{A}$ gets a cid from oracle Left during registration, and submits it to oracle Challenge for token T_0 and server identity id_{S_L}. It compares the resulting view (e.g. returned public keys) with the view that it gets from submitting cid to oracle Left in authentication. If the views are consistent, T_0 is used in Left. Our strongest notion of credential separation, leading to *strong unlinkability*, rules out exactly this behavior of $\mathcal{A}$.

In practice, servers can follow this strategy to determine if any credential cid in its database belongs to a user that aims to log in. The server just sends cid instead of the credential that is actually stored for this user and checks if the authentication goes through. It is worth noting that such an attack can be launched almost unnoticed by the user. In recent work, Kepkowski et al. [6] showed a timing attack against unlinkability where they exploit a small time difference on some hardware tokens when handling the cid. Their paper's key idea is to use the WebAuthn specification to amplify the attack. To support multiple tokens, the relying party sends a list of supported cid's, which are locally queried to the token by the browser to find the correct one, i.e. the one for which the token responds with a valid signature. Those queries do not require interaction and are oblivious to the user.

Going back to our example, the malicious relying party can add cid to the list containing the correct credential for the user account. If the response is a valid signature for cid, the relying party knows that cid corresponds to this account. If this distinguishing fails, the user is oblivious to the attack.

Weaker Notions. Not every implementation used in practice is strongly unlinkable. As we still aim for an analysis and a precise statement about which attacks allow to link, we define weaker variations of the above notion. To define *medium unlinkability*, we make the credential separation condition more restrictive. Namely, we also rule out that $\mathcal{A}$ queries the same cid at oracle Challenge (for the target $T_0, T_1, \mathrm{id}_{S_L}, \mathrm{id}_{S_R}$) and Left, Right, *even if the* cid *is made up by* $\mathcal{A}$ *and not returned in a registration of these oracles*. Looking ahead, this subtle difference will be required to prove unlinkability in case cid is not authenticated. For *weak unlinkability*, we additionally rule out the case that oracle Challenge (for the target $T_0, T_1, \mathrm{id}_{S_L}, \mathrm{id}_{S_R}$) and Left, Right *return the same* cid *during registration*. If cid has high enough entropy this implies medium unlinkability. However, there are schemes only satisfying weak unlinkability.

Definition 3 (Unlinkability). *For a PlA* PLA $=$ (Gen, Reg, Auth) *and an adversary* $\mathcal{A}$, *we define experiments* $\mathbf{wUnl}, \mathbf{mUnl}, \mathbf{sUnl}$ *as follows.*

- **Setup.** *For each token* $T \in \mathcal{T}$, *a key is generated by running* $\mathrm{msk}_T \leftarrow \mathsf{Gen}(\mathrm{par})$.
- **Phase 1.** *The adversary is allowed to interact with oracles* Start, Challenge, Complete *(see Definition 1).*
- **Phase 2.** *The adversary outputs two (not necessarily distinct) token identifiers* T_0, T_1, *and two (not necessarily distinct) server identifiers* $S_L, S_R \in \mathcal{S}$. *Let* i_0 *and* i_1 *be the smallest identifiers for which the token handles* $\pi_{T_0}^{i_0, 0}$ *and* $\pi_{T_1}^{i_1, 0}$ *were not queried to the* Challenge *oracle in Phase 1. The experiment chooses a bit* b *uniformly at random. It sets* $j_0 := 0, j_1 := 0$ *and initializes two oracles* Left, Right *as follows:*
  - Left(cid, M)*: Return* Challenge($\pi_{T_b}^{i_b, j_b}, \mathrm{id}_{S_L}, \mathrm{cid}, M$) *and set* $j_b = j_b + 1$.
  - Right(cid, M)*: Return* Challenge($\pi_{T_{1-b}}^{i_{1-b}, j_{1-b}}, \mathrm{id}_{S_R}, \mathrm{cid}, M$) *and set* $j_{1-b} = j_{1-b} + 1$.
- **Phase 3.** *The adversary is allowed to interact with all the oracles defined in Phase 1 and 2.*
- **Output Phase.** *Finally, the adversary outputs a bit* $\hat{b}$. *Consider the following lists of* cid*'s:*
  - $\mathcal{L}_{ch}^r$ *contains all* cid*'s returned by queries that are not issued via* Left, Right *and are of the form* Challenge($\pi_T^{i, 0}, \mathrm{id}_S, \cdot, \cdot$) *for any* $i, T \in \{T_0, T_1\}$ *and* $S \in \{S_L, S_R\}$.
  - $\mathcal{L}_{ch}^a$ *contains all* cid*'s that are part of the input of queries that are not issued via* Left, Right *and are of the form* Challenge($\pi_T^{i, j}, \mathrm{id}_S, \cdot, \cdot$) *for any* $j > 0, i, T \in \{T_0, T_1\}$ *and* $S \in \{S_L, S_R\}$.
  - $\mathcal{L}_{lr}^r$ *contains all* cid*'s returned by queries to* Left *or* Right *when* $j_b = 0$ *or* $j_{1-b} = 0$, *respectively.*
  - $\mathcal{L}_{lr}^a$ *contains all* cid*'s that are part of the queries to* Left *or* Right *when* $j_b > 0$ *or* $j_{1-b} > 0$, *respectively.*

  The experiment returns 1 if and only if:
  - *bit* $\hat{b}$ *is equal to bit* b, *and*
  - *(instance freshness) the adversary never made a query to oracle* Challenge *using handles* $\pi_{T_0}^{i_0, k_0}$ *and* $\pi_{T_1}^{i_1, k_1}$ *for any* k_0, k_1, *and*

– *(credential separation) The following set is empty:*

$$\mathbf{sUnl}: \quad S_{\mathbf{sUnl}} := (\mathcal{L}_{ch}^r \cap \mathcal{L}_{lr}^a) \cup (\mathcal{L}_{lr}^r \cap \mathcal{L}_{ch}^a),$$
$$\mathbf{mUnl}: \quad S_{\mathbf{mUnl}} := S_{\mathbf{sUnl}} \cup (\mathcal{L}_{ch}^a \cup \mathcal{L}_{lr}^a),$$
$$\mathbf{wUnl}: \quad S_{\mathbf{wUnl}} := S_{\mathbf{mUnl}} \cup (\mathcal{L}_{ch}^r \cup \mathcal{L}_{lr}^r).$$

For $x \in \{\mathbf{w}, \mathbf{m}, \mathbf{s}\}$, we define the advantage of $\mathcal{A}$ in winning the experiment as:

$$\mathbf{Adv}_{x\mathbf{Unl},\mathsf{PLA}}^{\mathcal{A}} := |\Pr[x\mathbf{Unl}_{\mathsf{PLA}}^{\mathcal{A}} = 1] - 1/2|.$$

The credential separation condition models the attacks a server can do when the same token is used twice at the same server. Even for schemes with strong unlinkability, there is a trivial active attack, as sketched above. For schemes with weak unlinkability, but no medium unlinkability, there is a passive attack that only involves the two registrations.

5. Analysis of Existing Implementations

In this section, we analyze the existing token implementations kwrPA and kdfPA of WebAuthn. Recall that these schemes are formally given in Figure 1. We analyze both impersonation security and unlinkability.

Analysis of Impersonation Security. As the first part of our analysis, we show that both implementations kwrPA and kdfPA satisfy impersonation security. We obtain the following statements.

Lemma 1. *Let $\mathcal{A}$ be an adversary in the impersonation game of* kdfPA. *Assume that $\mathcal{A}$ makes at most Q_H queries to random oracle H, at most Q_S queries to oracle* Start, *and at most Q_C queries to oracle* Challenge. *Then there exists algorithms $\mathcal{B}, \mathcal{B}'$ with the same running time as $\mathcal{A}$ such that*

$$\mathbf{Adv}_{\mathbf{Imp},\mathsf{kdfPA}}^{\mathcal{A}} \leq \frac{Q_S^2 + Q_C^2}{2^\lambda} + \frac{Q_H^2}{2^{2\lambda}} + |\mathcal{T}| \cdot \mathbf{Adv}_{\mathbf{prf},\mathsf{PRF}}^{\mathcal{B}'} + Q_C \cdot \mathbf{Adv}_{\mathbf{euf\text{-}cma},\mathsf{SIG}}^{\mathcal{B}}$$

Lemma 2. *Let $\mathcal{A}$ be an adversary in the impersonation game of* kwrPA. *Assume that $\mathcal{A}$ makes at most Q_H queries to random oracle H, and at most Q_S, Q_C queries to oracles* Start, Challenge, *respectively. Then there exist algorithms $\mathcal{B}$ and $\mathcal{C}$ with the same running time as $\mathcal{A}$ such that*

$$\mathbf{Adv}_{\mathbf{Imp},\mathsf{kwrPA}}^{\mathcal{A}} \leq \frac{Q_S^2}{2^\lambda} + \frac{Q_H^2}{2^{2\lambda}} + |\mathcal{T}| \cdot \mathbf{Adv}_{\mathbf{anon\text{-}auth},\mathsf{SKE}}^{\mathcal{B}} + Q_C \cdot \mathbf{Adv}_{\mathbf{euf\text{-}cma},\mathsf{SIG}}^{\mathcal{C}}.$$

We give a proof intuition here, and postpone the formal proofs to the full version [30] and Appendix C.2.

Proof Intuition. Our goal is to give a reduction from the **euf-cma** security of SIG. To do that, recall that the adversary wins the impersonation game, if a token first

registers a public key pk at a server, and then the adversary authenticates, i.e. it forges a valid signature for $m = (H(\mathsf{id}_S), H(rs))$ with respect to pk, without using the token. We call this interaction the forged authentication.

Our strategy is to embed the public key that we get from the **euf-cma** game into one of the registration interactions between oracle Challenge and the adversary. Secret keys are not only used to sign, and we need to deal with that. Namely, in variant kdfPA, we first need to apply pseudorandomness of PRF to make secret keys independent and random. The variant kwrPA additionally outputs a ciphertext cid encrypting the secret key. To eliminate this dependency on the secret key, we perform a hybrid step to switch all cid's output by oracle Challenge to random, while internally storing a mapping from cid to the secret key to ensure consistency.

Now we can give a reduction that first guesses the registration interaction in which it embeds the given public key. To simulate the signatures for the embedded public key, the reduction can use the signing oracle provided by the **euf-cma** game. Finally, it can use the signature that the adversary sends in the forged authentication as a forgery for the **euf-cma** game. To see that the forgery is fresh, we rule out collisions for random oracle H and the challenges rs that servers send. □

Analysis of Unlinkability. Next, we give our results in terms of unlinkability. Notably, we obtain strong unlinkability for kwrPA, while kdfPA only satisfies medium unlinkability.

Lemma 3. *Let $\mathcal{A}$ be an adversary in the medium unlinkability game of* kdfPA. *Assume that $\mathcal{A}$ makes at most Q_C queries to oracle* Challenge. *Then, there is an algorithm $\mathcal{B}$, which has the same running time as $\mathcal{A}$, such that*

$$\mathbf{Adv}_{\mathbf{mUnl},\mathsf{kdfPA}}^{\mathcal{A}} \leq \frac{2Q_C}{2^\lambda} + |\mathcal{T}| \cdot \mathbf{Adv}_{\mathbf{prf},\mathsf{PRF}}^{\mathcal{B}}.$$

Lemma 4. *Let $\mathcal{A}$ be an adversary in the strong unlinkability game of* kwrPA. *Then there exists an algorithm $\mathcal{B}$, which has the same running time as $\mathcal{A}$, such that*

$$\mathbf{Adv}_{\mathbf{sUnl},\mathsf{kwrPA}}^{\mathcal{A}} \leq |\mathcal{T}|^2 \cdot \left(\mathbf{Adv}_{\mathbf{anon\text{-}auth},\mathsf{SKE}}^{\mathcal{B}} + \frac{2Q_C}{2^{\nu^*}}\right).$$

Again, we give a proof intuition here, and postpone formal proofs to Appendices C.1 and C.2.

Proof Intuition. To show unlinkability, our goal is to remove all information about tokens T_b, T_{1-b} from the oracles Left and Right. For the variant kdfPA and medium unlinkability, this can be done roughly as follows. We first use the entropy of msk_{T_b} to argue that the adversary can only access the pre-fixed random oracle $G := K(\mathsf{msk}_{T_b}, \cdot)$ via oracles Challenge and Left. Then, we rule out that two registrations sample the same cid. This essentially tells us that we can assume weak credential separation instead of medium credential separation. Using weak credential separation, we see that oracles Challenge and Left access oracle G on distinct

points, which allows us to split the oracle into two oracles. Using a similar argument for T_b and Right, we can conclude.

For kwrPA, our strategy is different. Namely, we first guess the target tokens and then apply the security of the encryption scheme, switching all cid's for these tokens to random. As in the proof of impersonation security, we hold a mapping to keep the keys that we use consistent. Now, for different cid's the tokens use independent keys. The only way to link between oracles Challenge and Left, Right, is to reuse values cid. This is forbidden by the strong credential separation condition. In contrast to the medium credential separation condition, the adversary is allowed to submit a fresh cid that is not output in an registration to Challenge and Left, Right. However, as the encryption scheme is authenticated, this will lead to aborting tokens. □

To see why kdfPA does not achieve strong unlinkability in general, consider an adversary that implements the following strategy. It first chooses two target tokens T_0, T_1 and server $S_L = S_R$ passes them to the experiment. Then, it runs honest registrations with oracle Left, oracle Right, and the oracle Challenge for token T_0. It then samples some random cid, and submits it during authentication interactions with these three oracles, using the same challenge M. All oracles will return valid signatures. Assuming that these are deterministically computed or reveal the public key, the adversary can simply compare the public keys to link. This attack would not be possible if cid's were authenticated, e.g. using a MAC. In this case, we are confident that one could prove medium unlinkability.

6. Defining Global Key Revocation

A useful feature that FIDO2 currently does not support is a means of revoking keys of a comprised token. Informally, we want to give the user the option to revoke its keys globally, without accessing all the servers which her token is registered to one by one. In this section, we focus on formally defining syntax and security for global key revocation. Then, in the next section we show a way to add this feature to FIDO2 using BIP32 key derivation.

In short, we define a global revocation mechanism as two algorithms that are associated with PLA. Intuitively, these should be understood as follows. First, when a user starts using its token $T \in \mathcal{T}$, it also obtains a revocation key rk, which is extracted from the long-term key msk using some algorithm Revoke, and should be stored externally, e.g. on a piece of paper. Recall from Definition 14 that when token T registers at a server $S \in \mathcal{S}$, the server stores a credential cred for this token in its state. If the user wants to revoke its key, it publishes rk. We do not further specify how the user publishes rk. However, we assume that all servers periodically scan for published revocation keys rk. Whenever a new revocation key is published, the server (with identity id_S) runs an algorithm $\text{CheckCred}(\text{id}_S, \text{cred}, \text{rk})$ for each credential $\text{cred} = \text{rcs}[\text{cid}]$ in its registration context. If this algorithm accepts then the credential is considered revoked.

Definition 4 (Global Key Revocation). *A PlA* PLA = (Gen, Reg, Auth) *satisfies global key revocation if there are* algorithms Revoke, CheckCred *with the following syntax:*

- Revoke(msk) *takes as input a master secret key* msk *and outputs a revocation key* rk.
- CheckCred(id_S, cred, rk) *takes as input a server identity* id_S, *a credential* cred *and a revocation key* rk *and outputs a bit* $b \in \{0, 1\}$.

Further, the algorithms should be complete in the following sense: For all msk $\in$ Gen(par)*, parties* T *and* S*, sets* R_{init}, R_{betw} *of tuples* (cid, cred) *the following experiment outputs 1 with probability 1:*

1) *Run steps (1)-(3) from the experiment in Definition 15.*
2) *Run* rk $\leftarrow$ Revoke(msk) *and* $b \leftarrow$ CheckCred(id_S, cred, rk)*. Return* b.

Clearly, the above definition is easily satisfied if we make algorithm CheckCred always output $b = 1$. This is not what we aim for. Instead, we need a security notion that ensures that only the owner of msk can revoke keys that are stored on an honest server. In the security experiment we define, the adversary gets arbitrary access to token and server oracles. Then, it has to choose two partnered oracles corresponding to a registration. It gets the corresponding credential and has to output a revocation key to revoke it. Intuitively, this means that it tries to revoke an honest registration, with which it can arbitrarily interfere.

Definition 5 (Revocation Soundness). *Let* PLA = (Gen, Reg, Auth) *be a PlA. Consider an algorithm* $\mathcal{A}$ *and the following experiment* $\text{rev-sound}_{\text{PLA}}^{\mathcal{A}}$:

- ***Setup.*** *For each token* $T \in \mathcal{T}$, *a key is generated by running* $\text{msk}_T \leftarrow$ Gen(par).
- ***Online Phase.*** $\mathcal{A}$ *gets access to oracles* Start, Challenge, Complete *as in Definition 1.*
- ***Output Phase.*** $\mathcal{A}$ *outputs tuples* (T^*, i_{T^*}) *and* (S^*, i_{S^*}). *If the oracle* $\text{Complete}(\pi_{S^*}^{i_{S^*},0}, \cdot, \cdot)$ *has never been queried, the experiment returns 0. Otherwise, the experiment returns* cred^* *to* $\mathcal{A}$, *where* cred^* *is the credential that oracle* $\text{Complete}(\pi_{S^*}^{i_{S^*},0}, \cdot, \cdot)$ *added to* rcs_{S^*}. *Then,* $\mathcal{A}$ *outputs* rk^*. *The experiment outputs 1 if and only if* $\pi_{T^*}^{i_{T^*},0}$ *and* $\pi_{S^*}^{i_{S^*},0}$ *are partnered and* $\text{CheckCred}(\text{id}_{S^*}, \text{cred}^*, \text{rk}^*) = 1$.

We define the advantage of $\mathcal{A}$ *in* $\text{rev-sound}_{\text{PLA}}^{\mathcal{A}}$ *as*

$$\mathbf{Adv}_{\text{rev-sound,PLA}}^{\mathcal{A}} := \Pr[\text{rev-sound}_{\text{PLA}}^{\mathcal{A}} = 1].$$

Even for globally revoked keys impersonation should be impossible, in case a server does not update its state in time. Therefore, we extend the notion of impersonation security.

Definition 6 (Impersonation Security - GR). *Let* PLA = (Gen, Reg, Auth) *be a PlA. We consider the experiment given in Definition 2 with the following modification, and call the resulting experiment* $\mathbf{Imp\text{-}GR}_{\text{PLA}}^{\mathcal{A}}$: *After generating* msk_T *for each* $T \in \mathcal{T}$, *the experiment also generates* $\text{rk}_T \leftarrow$ Revoke(msk_T) *for each* $T \in \mathcal{T}$. *Then,* $\{\text{rk}_T\}_{T \in \mathcal{T}}$

is given to algorithm $\mathcal{A}$ as an additional input. We define the advantage of an algorithm $\mathcal{A}$ in the experiment as

$$\mathbf{Adv}^{\mathcal{A}}_{\mathbf{Imp\text{-}GR},\mathsf{PLA}} := \Pr[\mathbf{Imp\text{-}GR}^{\mathcal{A}}_{\mathsf{PLA}} = 1].$$

For unlinkability, it is clear that all keys of a token can be linked once the revocation key is published. Thus, in our modified unlinkability experiment, the adversary gets all revocation keys except the ones for the challenge tokens. As our scheme only satisfies the weak version of unlinkability with global revocation, we only define this. The medium and strong version can be defined analogously.

Definition 7 (Unlinkability - GR). *Let* $\mathsf{PLA} = (\mathsf{Gen}, \mathsf{Reg}, \mathsf{Auth})$ *be a PlA. We consider the experiment* **wUnl** *given in Definition 3 with the following modification, and call the resulting experiment* $\mathbf{wUnl\text{-}GR}^{\mathcal{A}}_{\mathsf{PLA}}$: *After generating* msk_T *for each* $T \in \mathcal{T}$, *the experiment also generates* $\mathsf{rk}_T \leftarrow \mathsf{Revoke}(\mathsf{msk}_T)$ *for each* $T \in \mathcal{T}$. *Then, when* $\mathcal{A}$ *outputs* T_0, T_1 *and* S_L, S_R *in Phase 2, the experiment gives* $\{\mathsf{rk}_T\}_{T \in \mathcal{T} \setminus \{T_0, T_1\}}$ *to* $\mathcal{A}$. *The rest of the experiment is as in* **wUnl**. *We define the advantage of an algorithm* $\mathcal{A}$ *as*

$$\mathbf{Adv}^{\mathcal{A}}_{\mathbf{wUnl\text{-}GR},\mathsf{PLA}} := |\Pr[\mathbf{wUnl\text{-}GR}^{\mathcal{A}}_{\mathsf{PLA}} = 1] - 1/2|.$$

7. BIP32 Passwordless Authentication

In this section, we show how to instantiate FIDO2 using ECDSA with the BIP32 key derivation [31]. The resulting scheme, denoted by bip32PA supports global key revocation.

7.1. Scheme Description

We give a description of our scheme bip32PA. Then, we also explain how it can support global revocation.

ECDSA Signatures. We recall the ECDSA signature scheme and its key-prefixed variant. The system parameters par of the scheme contain a group $\mathbb{G}$ of prime order p with generator $g \in \mathbb{G}$. Let $H' : \{0,1\}^* \rightarrow \mathbb{Z}_p$ be a random oracle. We denote the scheme by $\mathsf{SIG}_{H'} = (\mathsf{Gen}, \mathsf{Sig}, \mathsf{SRerand}, \mathsf{PRerand}, \mathsf{Ver})$ and its key-prefixed variant as $\mathsf{SIG}^{kp}_{H'}$. We first describe key generation and randomization. Key generation and rerandomization are as follows:

- $\mathsf{Gen}(\mathsf{par})$: Sample $\mathsf{sk} \xleftarrow{\$} \mathbb{Z}_p$, $\mathsf{pk} := g^{\mathsf{sk}}$. Return $(\mathsf{sk}, \mathsf{pk})$.
- $\mathsf{SRerand}(\mathsf{sk}, \rho \in \mathbb{Z}_p)$: Return $\mathsf{sk}' := \mathsf{sk} + \rho$.
- $\mathsf{PRerand}(\mathsf{pk}, \rho \in \mathbb{Z}_p)$: Return $\mathsf{pk}' := \mathsf{pk} \cdot g^{\rho}$.

Next, let us explain signing and verification. Signing and verifcation makes use of algorithms $\mathsf{Sig}^{int}, \mathsf{Ver}^{int}$ that are used internally. For the purpose of this work, we can treat these as a black box. For details, see [9]. Using these algorithms, signing (i.e. algorithm $\mathsf{Sig}(\mathsf{sk}, m)$) is as follows:

1) Set $pm := m$ ((pk, m) for the key-prefixed variant).
2) Compute $z := H'(pm)$.
3) Return $\sigma \leftarrow \mathsf{Sig}^{int}(\mathsf{sk}, z)$.

Verification computes z in the same way and runs $\mathsf{Ver}^{int}(\mathsf{pk}, \sigma, z)$. Looking ahead, the fact that pm is hashed

will allow the reductions in our proof to map key-prefixed messages to messages prefixed with a server identifier.

Key Generation. Let us describe how master secret keys msk are generated for our scheme bip32PA. The key consists of an ECDSA key pair $(\mathsf{sk}_0, \mathsf{pk}_0)$, a so called chaincode ch and a seed $seed$. Looking ahead, the chaincode and the public key pk_0 can later be used to revoke keys, and the seed $seed$ will be used to derive randomness for signing. Concretely, the components of the key are generated as $\mathsf{sk}_0 \xleftarrow{\$} \mathbb{Z}_p, \mathsf{pk}_0 := g^{\mathsf{sk}_0}, ch \xleftarrow{\$} \{0,1\}^\lambda, seed \xleftarrow{\$} \{0,1\}^\lambda$ and we set $\mathsf{msk} := (\mathsf{sk}_0, \mathsf{pk}_0, ch, seed)$.

Registration and Authentication. Registration and authentication follow the WebAuthn specification. Thus, they are very similar to the protocols described in Section 5. Formally, we present protocols Reg and Auth in Figure 2. Let us describe the differences between bip32PA and the existing schemes. The most important difference is how keys are derived during registration and authentication. Namely, the token defines $\mathsf{cid} := \hat{H}(seed, \mathsf{id}_S)$ using a random oracle $\hat{H} : \{0,1\}^* \rightarrow \{0,1\}^\lambda$. Then, it derives a randomness $\rho := \hat{H}(\mathsf{pk}_0, ch, \mathsf{id}_S)$. This randomness is used to rerandomize the pair $(\mathsf{sk}_0, \mathsf{pk}_0)$ to a new keypair, i.e. $\mathsf{sk} := \mathsf{SRerand}(\mathsf{sk}_0, \rho)$ and $\mathsf{pk} := \mathsf{PRerand}(\mathsf{pk}_0, \rho)$. As in the scheme kdfPA, the server stores $\mathsf{cid}, \mathsf{pk}$ during registration and the key sk is used to sign challenges.

Let us give a brief explanation of the design choices made here. Implementing the signing process in a provably secure way that is compatible with existing server implementations is non-trivial. This is because provable security with respect to randomized keys for ECDSA is only known for the key-prefixed version and for one signature per message [9]. To support multiple signatures per message, one idea is to let the token choose a random nonce and append it to messages. As we do not want to change the protocol on the server side, we can not rely on such random nonces. Instead, we derandomize the signing process by deriving the random coins used for signing a message m as $\hat{H}(seed, m)$. It remains to avoid key-prefixing, as this would also require changing verification on the server side. Here, we makes use of the fact that with high probability, the mapping from server identities to public keys is injective. Therefore, prefixing with server identities (which is already done in WebAuthn) is as good as prefixing with public keys. To make this idea work, we need to ensure that for each server, there is a fixed public key, which explains the definition of cid. As cid is deterministically derived from id_S, each registration on the same server will be associated with the same public key. This is also the reason why the scheme only achieves weak unlinkability.

Global Revocation. The advantage of the BIP32 key derivation compared to existing schemes is global key revocation, as defined in the previous section. We present algorithms Revoke and CheckCred for our scheme bip32PA. Algorithm $\mathsf{Revoke}(\mathsf{msk})$ is given as follows:

1) Parse $(\mathsf{sk}_0, \mathsf{pk}_0, ch, seed) := \mathsf{msk}$.
2) Return $\mathsf{rk} := (\mathsf{pk}_0, ch)$.

Algorithm $\mathsf{CheckCred}(\mathsf{id}_S, \mathsf{cred}, \mathsf{rk})$ is as follows:

$(\text{cid}, R_r) \leftarrow \text{rresp}(\text{msk} = (\text{sk}_0, \text{pk}_0, ch, seed), \text{id}_S, M_r):$	$R_a \leftarrow \text{aresp}(\text{msk} = (\text{sk}_0, \text{pk}_0, ch, seed), \text{id}_S, \text{cid}, M_a):$
$\text{cid} := \hat{H}(seed, \text{id}_S)$	**if** $\text{cid} \neq \hat{H}(seed, \text{id}_S): $ **abort**
$\text{sk} := \text{SRerand}(\text{sk}_0, \hat{H}(\text{pk}_0, ch, \text{id}_S))$	$\text{sk} := \text{SRerand}(\text{sk}_0, \hat{H}(\text{pk}_0, ch, \text{id}_S))$
$\text{pk} := \text{PRerand}(\text{pk}_0, \hat{H}(\text{pk}_0, ch, \text{id}_S))$	$m := (H(\text{id}_S), M_a), \; coins := \hat{H}(seed, m)$
$m := (H(\text{id}_S), \text{cid}, \text{pk}, M_r), \; coins := \hat{H}(seed, m)$	$\sigma := \text{Sig}(\text{sk}, m; coins), \; R_a := \sigma$
$\sigma := \text{Sig}(\text{sk}, m; coins), \; R_r := (\text{pk}, \sigma)$	

Figure 2. The WebAuthn registration (left) and authentication protocol (right) for our new variation bip32PA. Functions V_t, V_s, and algorithms rcomm, acomm, rchall, achall, rcheck, acheck are given as in Figure 1.

1) Parse $(\text{pk}_0, ch) := \text{rk}$ and $\text{pk} := \text{cred}$.
2) Run $\text{pk}' := \text{PRerand}(\text{pk}_0, \hat{H}(\text{pk}_0, ch, \text{id}_S))$.
3) Return 1 if $\text{pk} = \text{pk}'$. Otherwise, return 0.

7.2. Revocation Soundness and Unlinkability

Next, we show revocation soundness and unlinkability of bip32PA. We postpone the formal proofs to the full version [30], and give proof intuitions here.

Revocation Soundness. We show revocation soundness. Recall that this means that only the owner of a token can revoke its keys. Our result is summarized in the following statement.

Lemma 5. *Let $\mathcal{A}$ be an adversary in the revocation soundness game of* bip32PA. *Assume that $\mathcal{A}$ makes at most $Q_{\hat{H}}$ queries to oracle $\hat{H}$ and at most Q_C queries to oracle* Challenge. *Then we have*

$$\mathbf{Adv}^{\mathcal{A}}_{\mathbf{rev\text{-}sound}, \mathsf{bip32PA}} \leq \frac{Q_{\hat{H}} \cdot |\mathcal{T}| + Q_{\hat{H}}^2 + 1 + Q_C \cdot Q_{\hat{H}}}{2^\lambda}.$$

Proof Intuition. In the experiment, the adversary first instructs a token T^* with key $\text{msk}_T = (\text{sk}_{T^*,0}, \text{pk}_{T^*,0}, ch_{T^*}, seed_{T^*})$ and a server S^* to interact in a registration. Then, it gets the resulting credential $\text{cred}^* = \text{pk}$, and has to output a revocation key $\text{rk}^* = (\text{pk}_0^*, ch^*)$. The adversary breaks revocation soundness, if $\text{CheckCred}(\text{id}_{S^*}, \text{cred}^*, \text{rk}^*) = 1$, i.e. if $\text{pk} = \text{PRerand}(\text{pk}_0^*, \hat{H}(\text{pk}_0^*, ch^*, \text{id}_{S^*}))$. To show that this is infeasible, we first argue using the entropy of ch_{T^*} that the adversary will not output $\text{rk}^* = (\text{pk}_{T^*,0}, ch_{T^*})$. Then, we guess the registration interaction of interest and embed an independent public key pk into this registration. We can do this by programming random oracle $\hat{H}$ accordingly. With this, winning the game reduces to solving the following isolated problem: Given a key pair (sk, pk) and access to a random oracle $\tilde{H}$, find pk' and ch such that $\text{pk} = \text{PRerand}(\text{pk}', \tilde{H}(\text{pk}', ch))$. This problem is statistically hard to solve, as we show in the full version [32]. $\square$

Unlinkability. We show unlinkability in presence of global revocation. We obtain the following result.

Lemma 6. *Let $\mathcal{A}$ be an adversary in the weak unlinkability with global revocation game of* bip32PA. *Assume that $\mathcal{A}$*

makes at most $Q_{\hat{H}}, Q_H$ *queries to random oracles $\hat{H}, H$, respectively. Then we have*

$$\mathbf{Adv}^{\mathcal{A}}_{\mathbf{wUnl\text{-}GR}, \mathsf{bip32PA}} \leq \frac{4 \cdot Q_{\hat{H}} + 2 \cdot |\mathcal{T}|^2}{2^\lambda} + \frac{Q_{\hat{H}}^2}{2^{2\lambda}}.$$

Proof Intuition. The proof is similar to the proof of variant kdfPA. Namely, we first argue that the adversary can access the random oracles $G_1 := \hat{H}(\text{pk}_{T_j,0}, ch_{T_j}, \cdot, \cdot)$ or $G_2 := \hat{H}(seed_{T_j}, \cdot)$ for the challenge tokens $j \in \{0, 1\}$ only indirectly via the oracles Left, Right and Challenge. Then, we claim that the oracles Left and Challenge access these prefixed random oracles on distinct inputs. As we have perfect rerandomization of keys, this means that the keys $\text{pk} = \text{PRerand}(\text{pk}_{T_j,0}, \hat{H}(\text{pk}_{T_j,0}, ch_{T_j}, \cdot, \cdot))$ that tokens use are independent, and the claim follows. To show the claimed domain separation, we can not rely on entropy of the values cid as we did for the variant kdfPA. Indeed, variant bip32PA derives cid deterministically from the server identity id_S. However, as we consider weak unlinkability, this argument is not necessary. Namely, for this specific scheme, credential separation implies that the adversary never uses the same server identity in both Challenge and Left, Right. Then, the domain separation follows from the fact that Left, Right and Challenge access the oracles G_1 and G_2 only on inputs that contain the server identity. $\square$

7.3. Impersonation Security

As showing impersonation security with global revocation is technically the most interesting part of our analysis, we dedicate this entire section to it. We will give an intuitive overview of our analysis, and postpone the formal definitions and proofs to the full version of the paper [32].

Unforgeability of Key-Prefixed ECDSA. The reader may wonder why we can not follow the proof idea of variant kdfPA and reduce to the **euf-cma** security of plain ECDSA. To understand this, we have to recall that in the security experiment, the adversary gets the revocation keys $\text{rk}_T = (\text{pk}_{T,0}, ch_T)$ for all tokens T, and a token computes the keys pk that it uses via $\text{pk} := \text{PRerand}(\text{pk}_{T,0}, \hat{H}(\text{pk}_{T,0}, ch, \text{id}_S))$. Therefore, the adversary knows a non-trivial correlation between these public keys, which may allow to run some related key attack. To rule out this attack, the security notion of unforgeability under honestly rerandomized keys (**uf-hrk**1)

has been introduced [27], [9]. This notion is similar to standard **euf-cma**, but in addition, the adversary gets access to an oracle RandO that outputs uniform randomness ρ. An adversary can also ask the signing oracle to sign a message m using rerandomized keys $\mathsf{sk}_\rho := \mathsf{SRerand}(\mathsf{sk}, \rho)$. In the end, a forgery is also allowed to be for a rerandomization of the public key that the adversary obtained. We have the additional restriction that the signing oracle can be queried at most once per pair (m, ρ). Das et al. [9] showed that the key-prefixed variant of ECDSA satisfies this security notion.

Lemma 7 (informal, [9])**.** *Let* $H : \{0,1\}^* \to \mathbb{Z}_p$ *be a random oracle and* SIG_H^{kp} *be the key-prefixed ECDSA signature scheme. Then, under suitable assumption, for each efficient algorithm* $\mathcal{A}$*, the advantage* $\mathbf{Adv}_{\mathbf{uf\text{-}hrk1}, \mathsf{SIG}_H^{kp}}^{\mathcal{A}}$ *is negligible.*

We could just use this result and obtain impersonation security for the key-prefixed variant of bip32PA. As we aim to avoid key-prefixing, we need to translate the above result into a result for plain ECDSA.

Translation to Plain ECDSA. Our main insight is that the proof by Das et al. still works if we prefix all signed messages with some index idx, as long as there is a mapping from idx to ρ, meaning that for each idx, we always use the same $\rho = \rho_{idx}$. Then, in our scheme, we can use the hash of the server identifier id_S as idx. This works because in WebAuthn, messages are always prefixed with the hash of id_S. However, to have this mapping also means that we can not use different cid for the same id_S, which explains why we deterministically derive cid from id_S.

We make this approach more formal by introducing a variant **uf-hrk-idx**1 of the above security notion. In this variant, oracle RandO takes as input an index idx and always outputs the same ρ_{idx} for the same input idx. Then, all signed messages are prefixed by idx in the signing oracle. The final forgery also has to be prefixed with the correct index. Then, we show that ECDSA without key-prefixing satisfies this notion.

Lemma 8. *Let* $H_0 : \{0,1\}^* \to \mathbb{Z}_p$ *and* $H_1 : \{0,1\}^* \to \mathbb{Z}_p$ *be random oracles. Let* SIG_{H_1} *be the ECDSA signature scheme and* $\mathsf{SIG}_{H_0}^{kp}$ *be its key-prefixed variant as above. Let* $\mathcal{A}$ *be an adversary in the game* **uf-hrk-idx**1 *for* SIG_{H_1}*. Assume that* $\mathcal{A}$ *makes at most* Q_{H_1} *queries to random oracle* H_1*, and at most* Q_R *queries to oracle* RandO*. Then there exists an algorithm* $\mathcal{B}$ *with the same running time as* $\mathcal{A}$ *such that*

$$\mathbf{Adv}_{\mathbf{uf\text{-}hrk\text{-}idx}1, \mathsf{SIG}_{H_1}}^{\mathcal{A}} \leq \frac{(Q_R + Q_{H_1})^2}{2^\lambda} + \mathbf{Adv}_{\mathbf{uf\text{-}hrk}1, \mathsf{SIG}_{H_0}^{kp}}^{\mathcal{B}}.$$

The idea of the proof is that the reduction can internally simulate random oracle $H_1(idx, \cdot)$ by random oracle $H_0(\mathsf{pk}_{idx}, \cdot)$, where pk_{idx} is given as $\mathsf{PRerand}(\mathsf{pk}, \rho_{idx})$.

Using this result, we can then show impersonation security with global revocation of bip32PA. This is done by guessing the token T^* for which the adversary will forge

Scheme	Registration Time [ms]	Authentication Time [ms]
kdfPA	369.42	176.31
bip32PA	436.70	213.40

TABLE 2. EXECUTION TIME FOR bip32PA AND kdfPA AVERAGED ACROSS 100 RUNS. FOR ALL MEASUREMENTS, WE INSTRUMENT THE TOKEN TO NOT WAIT FOR USER TOUCH. THE STANDARD DEVIATION FOR ALL MEASUREMENTS IS LESS THAN 6MS.

a signature, and then embedding the key given by the **uf-hrk-idx**1 experiment as $\mathsf{pk}_{T^*, 0}$. To see that we only need to query the signing oracle once per pair (m, ρ), we highlight that we derive the random coins for signing deterministically in bip32PA. We obtain the following statement.

Lemma 9. *Let* $\mathcal{A}$ *be an adversary in the impersonation with global revocation game of* bip32PA*. Assume that* $\mathcal{A}$ *makes at most* $Q_H, Q_{\hat{H}}$ *queries to random oracles* $H, \hat{H}$*, respectively, at most* Q_S *queries to oracle* Start*, and at most* Q_C *queries to oracle* Challenge*. Then there exists an algorithm* $\mathcal{B}$ *with the same running time as* $\mathcal{A}$ *such that* $\mathcal{A}$*'s advantage in the impersonation with global revocation game can be upper bounded by*

$$\frac{Q_H^2 + |\mathcal{S}|^2}{2^{2\lambda}} + \frac{Q_S^2 + Q_{\hat{H}} \cdot |\mathcal{T}|}{2^\lambda} + |\mathcal{T}| \cdot \mathbf{Adv}_{\mathbf{uf\text{-}hrk\text{-}idx}1, \mathsf{SIG}_{H'}}^{\mathcal{B}}.$$

7.4. Evaluation

To show the practicality of bip32PA, we created a prototype implementation and evaluated it on the FIDO2 compliant SoloKey token. The SoloKey firmware uses the secp256r1 curve, is open source and can be uploaded to a special version called Solo Hacker[3]. Additionally, we evaluate the efficiency of the revocation process introduced for bip32PA on a standard Macbook Pro with an Intel i7 processor @2,3 GHz with 4 cores and 16 GB of RAM. We show that revoking even 2^{20} (around one million) tokens takes no more than 3 minutes on this personal computer. It is worth noting that the revocation process is highly parallelizable, and executing it on a server-based platform with multiple cores will yield significantly better execution times. We make the source code for our prototype implementation publicly available[4].

Token Implementation. We evaluate the on-token execution time for registration and authentication of bip32PA and compare it to the execution time of the kdfPA variant that is implemented in the original firmware of SoloKey. The results are presented in Table 2. Our prototype implementation of bip32PA is only around $1.2\times$ as slow as the original firmware implementation of kdfPA. The difference in execution time is around 40 ms in the case of authentication, which will only negligibly influence the user experience.

Global Revocation in Practice. We discuss how services can practically implement revocation. Recall that given the revocation key rk (computed using Revoke), the relying party with identity id_S can verify using CheckCred if the key corresponds to a credential cred. The way we modeled global revocation, the knowledge of rk only allows an adversary to link credentials created by the same token and does not allow for impersonation attacks. Thus, users can generate the revocation key during a setup phase and securely extract it from the token to store it, e.g., on a piece of paper as a QR code. A straightaway way to implement the revocation mechanism is to use an approach similar to certificate revocation lists (CRLs) [33]. The relying party downloads a size B blocklist with revocation keys of revoked tokens and uses the CheckCred algorithm against an internal database containing N unique credentials cred. The complexity of this algorithm is $B \cdot N$, making it somewhat inefficient for services with many users (big N). Like in the case of CRLs, the relying party can download the list via delta updates and only periodically run this check. However, in the case of freshly registered tokens, the relying party must check the revocation status for all elements on the blocklist. For a more in-depth analysis of how to efficiently implement the update, storage, and operation of such a revocation list, we refer to the literature on CRLs.

We will now look at how to optimize the revocation process in the case of bip32PA. The revocation key contains (pk_0, ch) and the CheckCred algorithm computes the public key $pk' := PRerand(pk_0, \hat{H}(pk_0, ch, id_S))$ and compares pk' against the credential, which in this case is also a public key pk. In other words, the relying party recomputes the public key the revoked token used for registration, and checks it against the provided public key. There are two observations here that will provide an improvement to the direct approach to revocation described above. Firstly, for each entry on the blocklist, the corresponding public key pk' can be computed once and used against the whole database of the relying party. Secondly, since we compare public keys, we can use a binary search instead of comparing the recomputed public key pk' with every element in the database. The mechanism will look as follows, assuming a presorted database. For each revocation key $rk := (pk_0, ch)$ on the blocklist, first recompute $pk' := PRerand(pk_0, \hat{H}(pk_0, ch, id_S))$ using the identity id_S and then run a binary search on internal database. The complexity of this mechanism is B times the cost of PRerand and $O(B \cdot \log(N))$ times the cost of comparing public keys (done via BigInteger comparison). It is worth noting that the relying party can store the precomputed public keys and use them in case of new registrations. It can then check the status of a newly registered token with $O(\log(B))$ public key comparisons.

To evaluate its efficiency, we created a prototype implementation of the revocation mechanism described above. We designed a basic function that simulates just one revocation key checking. It generates the public key pk' by reusing the PRerand function from our bip32PA token implementation.

Size B	2^{18}	2^{20}	2^{22}	2^{24}
Time [s]	44	175	699	2761

TABLE 3. EXECUTION TIME FOR bip32PA REVOCATION AVERAGED ACROSS 100 RUNS. WE ASSUME THE SIZE OF THE RELYING PARTY'S DATABASE TO BE $N = 2^{40}$ AND THE BLOCKLIST SIZE TO BE B. THE STANDARD DEVIATION FOR ALL MEASUREMENTS IS LESS THAN 5S.

Then it simulates a binary search by performing $\log(N)$ comparisons of the public key pk' with a random public key. During each test, we execute this simple function B times. It is easy to see that this process is highly parallelizable, and we used this in our prototype implementation. We divided the work across $T = 16$ threads. Further increasing this number did not significantly improve the mechanism's efficiency on our test platform. However, executing this mechanism on a server-oriented platform can make use of this parallelization even more. We executed and computed the average of 100 test for parameters $N \in \{2^{30}, 2^{32}, 2^{34}, 2^{36}, 2^{38}, 2^{40}\}$ and $B \in \{2^{18}, 2^{20}, 2^{22}, 2^{24}\}$. It is worth noting that the binary search only constituted around 1 second in total for all cases of N. Thus, we can practically use our revocation mechanism with even bigger database sizes. Due to this reason, in Table 3 we present the results only for $N = 2^{40}$. Checking the revocation status of around 1 million ($\sim 2^{20}$) revocation keys takes only about 3 minutes on this personal computer platform. The relying party can check smaller lists of size $2^{18} = 262144$ in less than a minute. It is easy to see that this revocation mechanism is practical. As we mentioned above, blocklists can be provided periodically, and we can realistically assume that they will probably not exceed 2^{20} entries.

8. Conclusion

We analyzed the WebAuthn protocol in FIDO2 with a focus on its real-world use cases and adapted the existing security models accordingly. We showed that privacy (unlinkability) of the protocol is not immediately guaranteed by the specification if keys are stored externally as in common implementations. To solve this issue we introduced the first formal security definition to capture privacy. Our results can be used as a guideline for token vendors. As an important example, we observed that in the case of key-wrapping the underlying encryption scheme must provide an anonymity property, i.e. ciphertexts created using the same key must be unlinkable to each other. We also introduced the notion of global key revocation and gave the first formalization of this property. Finally, we have shown that BIP32 key derivation can be used to obtain an efficient token implementation that supports global key revocation and is fully compatible with existing server implementations.

References

[1] "Client to Authenticator Protocol (CTAP)," https://fidoalliance.org/specs/fido-v2.0-ps-20190130/fido-client-to-authenticator-protocol-v2.0-ps-20190130.pdf, 2019, [Online; accessed 14-January-2022].

[2] "Web Authentication: An API for accessing Public Key Credentials," https://www.w3.org/TR/webauthn/, 2020, [Online; accessed 14-January-2022].

[3] M. Barbosa, A. Boldyreva, S. Chen, and B. Warinschi, "Provable security analysis of FIDO2," in *CRYPTO 2021, Part III*, ser. LNCS, T. Malkin and C. Peikert, Eds., vol. 12827. Virtual Event: Springer, Heidelberg, Aug. 2021, pp. 125–156.

[4] N. Bindel, C. Cremers, and M. Zhao, "FIDO2, CTAP 2.1, and WebAuthn 2: Provable security and post-quantum instantiation," Cryptology ePrint Archive, Report 2022/1029, 2022, https://eprint.iacr.org/2022/1029.

[5] I. B. Guirat and H. Halpin, "Formal verification of the w3c web authentication protocol," in *Proceedings of the 5th Annual Symposium and Bootcamp on Hot Topics in the Science of Security*, ser. HoTSoS '18. New York, NY, USA: Association for Computing Machinery, 2018. [Online]. Available: https://doi-org.simsrad.net.ocs.mq.edu.au/10.1145/3190619.3190640

[6] M. Kepkowski, L. Hanzlik, I. Wood, and M. A. Kâafar, "How not to handle keys: Timing attacks on FIDO authenticator privacy," *CoRR*, vol. abs/2205.08071, 2022. [Online]. Available: https://doi.org/10.48550/arXiv.2205.08071

[7] Yubico, "Yubikey U2F Key Generation," 2020, [Online; accessed 14-January-2022]. [Online]. Available: {https://developers.yubico.com/U2F/Protocol_details/Key_generation.html}

[8] S. G. Lyastani, M. Schilling, M. Neumayr, M. Backes, and S. Bugiel, "Is FIDO2 the kingslayer of user authentication? A comparative usability study of FIDO2 passwordless authentication," in *2020 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2020, pp. 268–285.

[9] P. Das, A. Erwig, S. Faust, J. Loss, and S. Riahi, "The exact security of BIP32 wallets," in *CCS '21: 2021 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, Republic of Korea, November 15 - 19, 2021*, Y. Kim, J. Kim, G. Vigna, and E. Shi, Eds. ACM, 2021, pp. 1020–1042. [Online]. Available: https://doi.org/10.1145/3460120.3484807

[10] A. Alam, K. Krombholz, and S. Bugiel, "Poster: Let history not repeat itself (this time) - tackling WebAuthn developer issues early on," in *ACM CCS 2019*, L. Cavallaro, J. Kinder, X. Wang, and J. Katz, Eds. ACM Press, Nov. 2019, pp. 2669–2671.

[11] D. Chakraborty and S. Bugiel, "simFIDO: FIDO2 user authentication with simTPM," in *ACM CCS 2019*, L. Cavallaro, J. Kinder, X. Wang, and J. Katz, Eds. ACM Press, Nov. 2019, pp. 2569–2571.

[12] J. Camenisch and A. Lysyanskaya, "Dynamic accumulators and application to efficient revocation of anonymous credentials," in *CRYPTO 2002*, ser. LNCS, M. Yung, Ed., vol. 2442. Springer, Heidelberg, Aug. 2002, pp. 61–76.

[13] D. Boneh and H. Shacham, "Group signatures with verifier-local revocation," in *ACM CCS 2004*, V. Atluri, B. Pfitzmann, and P. McDaniel, Eds. ACM Press, Oct. 2004, pp. 168–177.

[14] J. Camenisch, M. Drijvers, and J. Hajny, "Scalable revocation scheme for anonymous credentials based on n-times unlinkable proofs," in *Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society, WPES@CCS 2016, Vienna, Austria, October 24 - 28, 2016*, E. R. Weippl, S. Katzenbeisser, and S. D. C. di Vimercati, Eds. ACM, 2016, pp. 123–133. [Online]. Available: http://dl.acm.org/citation.cfm?id=2994625

[15] E. Dauterman, H. Corrigan-Gibbs, D. Mazières, D. Boneh, and D. Rizzo, "True2F: Backdoor-resistant authentication tokens," in *2019 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2019, pp. 398–416.

[16] N. Frymann, D. Gardham, F. Kiefer, E. Lundberg, M. Manulis, and D. Nilsson, "Asynchronous remote key generation: An analysis of yubico's proposal for W3C WebAuthn," in *ACM CCS 2020*, J. Ligatti, X. Ou, J. Katz, and G. Vigna, Eds. ACM Press, Nov. 2020, pp. 939–954.

[17] C.-P. Schnorr, "Efficient signature generation by smart cards," *Journal of Cryptology*, vol. 4, no. 3, pp. 161–174, Jan. 1991.

[18] W. Diffie, P. C. van Oorschot, and M. J. Wiener, "Authentication and authenticated key exchanges," *Designs, Codes and Cryptography*, vol. 2, no. 2, pp. 107–125, Jun. 1992.

[19] L. Law, A. Menezes, M. Qu, J. A. Solinas, and S. A. Vanstone, "An efficient protocol for authenticated key agreement," *Des. Codes Cryptogr.*, vol. 28, no. 2, pp. 119–134, 2003.

[20] B. A. LaMacchia, K. Lauter, and A. Mityagin, "Stronger security of authenticated key exchange," in *ProvSec 2007*, ser. LNCS, W. Susilo, J. K. Liu, and Y. Mu, Eds., vol. 4784. Springer, Heidelberg, Nov. 2007, pp. 1–16.

[21] D. Chaum and E. van Heyst, "Group signatures," in *EUROCRYPT'91*, ser. LNCS, D. W. Davies, Ed., vol. 547. Springer, Heidelberg, Apr. 1991, pp. 257–265.

[22] M. Bellare, D. Micciancio, and B. Warinschi, "Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions," in *EUROCRYPT 2003*, ser. LNCS, E. Biham, Ed., vol. 2656. Springer, Heidelberg, May 2003, pp. 614–629.

[23] I. Teranishi, J. Furukawa, and K. Sako, "k-Times anonymous authentication (extended abstract)," in *ASIACRYPT 2004*, ser. LNCS, P. J. Lee, Ed., vol. 3329. Springer, Heidelberg, Dec. 2004, pp. 308–322.

[24] Ö. Dagdelen and M. Fischlin, "Security analysis of the extended access control protocol for machine readable travel documents," in *ISC 2010*, ser. LNCS, M. Burmester, G. Tsudik, S. S. Magliveras, and I. Ilic, Eds., vol. 6531. Springer, Heidelberg, Oct. 2011, pp. 54–68.

[25] J. Brendel and M. Fischlin, "Zero round-trip time for the extended access control protocol," in *ESORICS 2017, Part I*, ser. LNCS, S. N. Foley, D. Gollmann, and E. Snekkenes, Eds., vol. 10492. Springer, Heidelberg, Sep. 2017, pp. 297–314.

[26] N. Fleischhacker, J. Krupp, G. Malavolta, J. Schneider, D. Schröder, and M. Simkin, "Efficient unlinkable sanitizable signatures from signatures with re-randomizable keys," in *PKC 2016, Part I*, ser. LNCS, C.-M. Cheng, K.-M. Chung, G. Persiano, and B.-Y. Yang, Eds., vol. 9614. Springer, Heidelberg, Mar. 2016, pp. 301–330.

[27] P. Das, S. Faust, and J. Loss, "A formal treatment of deterministic wallets," in *ACM CCS 2019*, L. Cavallaro, J. Kinder, X. Wang, and J. Katz, Eds. ACM Press, Nov. 2019, pp. 651–668.

[28] P. Rogaway and T. Shrimpton, "A provable-security treatment of the key-wrap problem," in *EUROCRYPT 2006*, ser. LNCS, S. Vaudenay, Ed., vol. 4004. Springer, Heidelberg, May / Jun. 2006, pp. 373–390.

[29] M. Bellare, D. Pointcheval, and P. Rogaway, "Authenticated key exchange secure against dictionary attacks," in *EUROCRYPT 2000*, ser. LNCS, B. Preneel, Ed., vol. 1807. Springer, Heidelberg, May 2000, pp. 139–155.

[30] L. Hanzlik, J. Loss, and B. Wagner, "Token meets wallet: Formalizing privacy and revocation for FIDO2," Cryptology ePrint Archive, Report 2022/084, 2022, https://eprint.iacr.org/2022/084.

[31] B. Wiki, "Bip32 proposal," 2018. [Online]. Available: {https://en.bitcoin.it/wiki/BIP_0032}

[32] Anonymous, "Token meets Wallet - Full Version." [Online]. Available: https://anonymous.4open.science/r/bip32PA/FullVersion.pdf

[33] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," Internet Requests for Comments, RFC Editor, RFC 5280, May 2008. [Online]. Available: http://www.rfc-editor.org/rfc/rfc5280.txt

[34] P. Rogaway, M. Bellare, J. Black, and T. Krovetz, "OCB: A block-cipher mode of operation for efficient authenticated encryption," in *ACM CCS 2001*, M. K. Reiter and P. Samarati, Eds. ACM Press, Nov. 2001, pp. 196–205.

[35] J. Jonsson, "On the security of CTR + CBC-MAC," in *SAC 2002*, ser. LNCS, K. Nyberg and H. M. Heys, Eds., vol. 2595. Springer, Heidelberg, Aug. 2003, pp. 76–93.

Appendix A. Preliminaries

Notation. We denote by $z \leftarrow \mathcal{A}(x)$ the execution of algorithm $\mathcal{A}$ on input x and with output z. If we want to make the random coins ρ of algorithm $\mathcal{A}$ explicit, we write $z := \mathcal{A}(x; \rho)$ instead. We write $y \in \mathcal{A}(x)$ to indicate that y is a possible output of $\mathcal{A}$ on input x. By $r \xleftarrow{\$} S$ we mean that r is chosen uniformly at random over the set S. We will use $[n]$ to denote the set $\{1, \ldots, n\}$. By λ, we denote the security parameter. Throughout the paper, we assume public parameters par are given implicitly to all algorithms.

Signatures. We recall the standard notion of digital signatures. We also introduce rerandomizable signature schemes [26], [27] and the corresponding security notion.

Definition 8 (Digital Signature Scheme). *A digital signature scheme is a tuple of algorithms* $\mathsf{SIG} = (\mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$ *with the following properties.*

- *The randomized key generation algorithm* Gen *takes as input parameters* par. *It outputs a keypair* $(\mathsf{sk}, \mathsf{pk})$.
- *The randomized signing algorithm* Sig *takes as input a secret key* sk *and a message* m. *It outputs a signature* σ.
- *The deterministic verification algorithm* Ver *takes as input a public key* pk, *a signature* σ, *and a message* m. *It outputs* 0 *(reject) or* 1 *(accept).*

We say that a digital signature scheme is correct *if for all* $(\mathsf{sk}, \mathsf{pk}) \in \mathsf{Gen}(\mathsf{par})$ *and all messages* $m \in \{0, 1\}^*$ *we have* $\Pr_{\sigma \leftarrow \mathsf{Sig}(\mathsf{sk}, m)}[\mathsf{Ver}(\mathsf{pk}, \sigma, m) = 1] = 1$. *We assume that secret keys are chosen uniformly at random and there is an algorithm* ToPK *that maps secret keys to public keys.*

Definition 9 (Unforgeability under Chosen Message Attacks). *Let* $\mathsf{SIG} = (\mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$ *be a signature scheme and consider the experiment* $\mathbf{euf\text{-}cma}_{\mathsf{SIG}}^{\mathcal{A}}$ *defined as follows:*

- *Setup. The experiment generates* $(\mathsf{sk}, \mathsf{pk})$ *via* $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{Gen}(\mathsf{par})$. *It runs the adversary* $\mathcal{A}$ *on input* pk.
- *Online Phase. In this phase,* $\mathcal{A}$ *is given access to oracle* SignO, *which takes as input a message* m *and returns the signature* $\sigma \leftarrow \mathsf{Sig}(\mathsf{sk}, m)$.
- *Output Phase. When* $\mathcal{A}$ *returns* (m^*, σ^*), *the experiment returns* 1 *if* $\mathsf{Ver}(\mathsf{pk}, \sigma^*, m^*) = 1$ *and* m^* *was not queried to* SignO. *Otherwise, it returns* 0.

We define the advantage of $\mathcal{A}$ *in* $\mathbf{euf\text{-}cma}_{\mathsf{SIG}}^{\mathcal{A}}$ *as* $\mathbf{Adv}_{\mathbf{euf\text{-}cma}, \mathsf{SIG}}^{\mathcal{A}} := \Pr[\mathbf{euf\text{-}cma}_{\mathsf{SIG}}^{\mathcal{A}} = 1]$.

Definition 10 (Rerandomizable Signature [26], [27]). *A rerandomizable signature scheme is a tuple of algorithms* $\mathsf{SIG} = (\mathsf{Gen}, \mathsf{Sig}, \mathsf{SRerand}, \mathsf{PRerand}, \mathsf{Ver})$, *where* $(\mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$ *is a digital signature scheme, and algorithms* $\mathsf{SRerand}, \mathsf{PRerand}$ *have the following properties.*

- *The deterministic secret key rerandomization algorithm* $\mathsf{SRerand}$ *takes as input a secret key* sk *and a string* ρ. *It outputs a rerandomized key* sk_ρ.
- *The deterministic public key rerandomization algorithm* $\mathsf{PRerand}$ *takes as input a public key* pk *and a string* ρ. *It outputs a rerandomized key* pk_ρ.

Moreover, we require that for all $(\mathsf{sk}, \mathsf{pk}) \in \mathsf{Gen}(\mathsf{par})$, *key pairs* $(\mathsf{sk}_\rho, \mathsf{pk}_\rho)$ *generated as* $\rho \xleftarrow{\$} \{0, 1\}^\lambda, \mathsf{sk}_\rho := \mathsf{SRerand}(\mathsf{sk}, \rho)$, *and* $\mathsf{pk}_\rho := \mathsf{PRerand}(\mathsf{pk}, \rho)$ *are identically distributed to key pairs generated via* $(\mathsf{sk}', \mathsf{pk}') \leftarrow \mathsf{Gen}(\mathsf{par})$.

Symmetric Primitives. Here, we recall the definition of symmetric key encryption that is both authenticated and anonymous. For the definition of security, we follow [28]. For simplicity of exposition, we choose to leave nonces implicit, leading to a potentially weaker notion. We note that this notion is satisfied using the OCB mode of operation [34] or the CCM mode of operation [35].

Definition 11 (Symmetric Key Encryption Scheme). *A symmetric key encryption scheme (SKE) with length function* $\nu \colon \mathbb{N} \to \mathbb{N}$ *is a tuple of algorithms* $\mathsf{SKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *with the following properties.*

- *The randomized key generation algorithm* Gen *takes as input parameters* par. *It outputs a secret key* sk.
- *The randomized encryption algorithm* Enc *takes as input a secret* sk *and a message* m. *It outputs a ciphertext* $c \in \{0, 1\}^{\nu(|m|)}$.
- *The deterministic decryption algorithm* Dec *takes as input a secret key* sk *and a ciphertext* c. *It outputs either* $\bot$ *or a message* m.

Definition 12 (Authenticated Anonymous Security for SKE). *Let* $\mathsf{SKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *be a symmetric key encryption scheme with length function* ν. *For an algorithm* $\mathcal{A}$ *and a bit* $b \in \{0, 1\}$, *consider the following experiment* $\mathbf{anon\text{-}auth}_{\mathsf{SKE}, b}^{\mathcal{A}}$:

- *Setup. The experiment generates a key* $\mathsf{sk} \leftarrow \mathsf{Gen}(\mathsf{par})$ *and initializes a map* $L[\cdot]$.
- *Online Phase. The adversary* $\mathcal{A}$ *is run on input* par *with oracle access to oracles* $\mathsf{EncO}_b, \mathsf{DecO}_b$, *which are defined as follows.*
  - $\mathsf{EncO}_0(m)$: *Return* $c \leftarrow \mathsf{Enc}(\mathsf{sk}, m)$.
  - $\mathsf{DecO}_0(c)$: *Return* $m := \mathsf{Dec}(\mathsf{sk}, c)$.
  - $\mathsf{EncO}_1(m)$: *Sample a ciphertext* $c \xleftarrow{\$} \{0, 1\}^{\nu(|m|)}$ *uniformly at random, define* $L[c] := m$ *and return* c.
  - $\mathsf{DecO}_1(c)$: *If* $L[c] \neq \bot$, *return* $L[c]$. *Else, return* $\bot$.
- *Output Phase. The adversary outputs a bit* b'. *The experiment outputs* b'.

We define the advantage of $\mathcal{A}$ *against authenticated anonymous security of* SKE *as* $\mathbf{Adv}_{\mathbf{anon\text{-}auth}, \mathsf{SKE}}^{\mathcal{A}} := \left| \Pr[\mathbf{anon\text{-}auth}_{\mathsf{SKE}, 0}^{\mathcal{A}} = 1] - \Pr[\mathbf{anon\text{-}auth}_{\mathsf{SKE}, 1}^{\mathcal{A}} = 1] \right|$.

Definition 13 (Pseudorandom Function). *Let* $\ell_1 = \ell_1(\lambda), \ell_2 = \ell_2(\lambda) \in \mathbb{N}$. *Consider a efficiently computable function* $\mathsf{PRF} \colon \{0, 1\}^\lambda \times \{0, 1\}^{\ell_1} \to \{0, 1\}^{\ell_2}$, *an algorithm* $\mathcal{A}$, *a bit* $b \in \{0, 1\}$, *and the following experiment* $\mathbf{prf}_{\mathsf{PRF}, b}^{\mathcal{A}}$:

- *Setup. The experiment samples* $K \xleftarrow{\$} \{0, 1\}^\lambda$ *and initializes a map* $F[\cdot]$.
- *Online Phase. The adversary* $\mathcal{A}$ *is run on input* par *with oracle access to oracle* EvalO_b, *which is defined as follows.*
  - $\mathsf{EvalO}_0(x \in \{0, 1\}^{\ell_1})$: *If* $F[x] \neq \bot$, *return* $F[x]$. *Else, sample* $F[x] \xleftarrow{\$} \{0, 1\}^{\ell_2}$ *and return* $F[x]$.

- EvalO$_1$($x \in \{0,1\}^{\ell_1}$): Return $\mathsf{PRF}(K, x)$.
- **Output Phase.** *The adversary outputs a bit b'. The experiment outputs b'.*

We define the advantage of $\mathcal{A}$ against the pseudorandomness of PRF as $\mathbf{Adv}^{\mathcal{A}}_{\mathsf{prf},\mathsf{PRF}} :=$ $\left| \Pr[\mathbf{prf}^{\mathcal{A}}_{\mathsf{PRF},0} = 1] - \Pr[\mathbf{prf}^{\mathcal{A}}_{\mathsf{PRF},1} = 1] \right|$.

Appendix B.
Formal Syntax, Completeness, and Partnering

Definition 14 (PlA)**.** *A passwordless authentication scheme (PlA) is a tuple* $\mathsf{PLA} = (\mathsf{Gen}, \mathsf{Reg}, \mathsf{Auth})$ *with the following properties:*

- *The randomized key generation algorithm* Gen *takes as input parameters* par. *It outputs a master secret key* msk.
- *The registration protocol* Reg *given as a tuple of algorithms* $(\mathsf{rchall}, \mathsf{rcomm}, \mathsf{rresp}, \mathsf{rcheck})$.
  - *The randomized registration challenge generation algorithm* rchall *takes as input a server identity* id_S *and outputs a challenge value* c *and a state* st.
  - *The deterministic registration command creation algorithm* rcomm *takes as input a server identity* id_S *and a challenge value* c *and outputs a message* M_r.
  - *The randomized registration response algorithm* rresp *takes as input a master secret key* msk, *a server identity* id_S *and a message* M_r *and outputs and credential identifier* cid *and a response* R_r.
  - *The deterministic registration check algorithm* rcheck *takes as input a state* st, *a credential identifier* cid *and a response* R_r *and outputs a bit* $b \in \{0,1\}$ *and a credential* cred.
- *The authentication protocol* Auth *is given as a tuple of algorithms* $(\mathsf{achall}, \mathsf{acomm}, \mathsf{aresp}, \mathsf{acheck})$.
  - *The randomized authentication challenge generation algorithm* achall *takes as input a server identity* id_S *and outputs a challenge value* c *and a state* st.
  - *The deterministic authentication command creation algorithm* acomm *takes as input a server identity* id_S *and a challenge value* c *and outputs a message* M_a.
  - *The randomized authentication response algorithm* aresp *takes as input a master secret key* msk, *a server identity* id_S, *a credential identifier* cid, *and a message* M_a *and outputs a response* R_a.
  - *The deterministic authentication check algorithm* acheck *takes as input a state* st, *a registration context* rcs, *a credential identifier* cid *and a response* R_a *and outputs a bit* $b \in \{0,1\}$

Algorithms $\mathsf{rchall}, \mathsf{rcheck}, \mathsf{achall}, \mathsf{acheck}$ *are executed by servers,* $\mathsf{rcomm}, \mathsf{acomm}$ *are executed by clients, and* $\mathsf{rresp}, \mathsf{aresp}$ *are executed by tokens.*

Definition 15 (Completeness of PlA)**.** *We say that a PlA* $\mathsf{PLA} = (\mathsf{Gen}, \mathsf{Reg}, \mathsf{Auth})$ *with* $\mathsf{Reg} = (\mathsf{rchall}, \mathsf{rcomm}, \mathsf{rresp}, \mathsf{rcheck})$ *and* $\mathsf{Auth} = (\mathsf{achall}, \mathsf{acomm}, \mathsf{aresp}, \mathsf{acheck})$ *is complete, if for all* $\mathsf{msk} \in \mathsf{Gen}(\mathsf{par})$, *parties* T *and* S, *sets* R_{init}, R_{betw} *of tuples* $(\mathsf{cid}, \mathsf{cred})$, *the probability that the following experiment outputs 0 is 0:*

1) *Let* rcs *be a key-value table. For each* $(\mathsf{cid}, \mathsf{cred}) \in R_{init}$, *set* $\mathsf{rcs}[\mathsf{cid}] := \mathsf{cred}$.
2) *Run the registration protocol* Reg *of* T *at* S, *as follows:*

$$(c, \mathsf{st}) \leftarrow \mathsf{rchall}(\mathsf{id}_S), M_r \leftarrow \mathsf{rcomm}(\mathsf{id}_S, c),$$
$$(\mathsf{cid}, R_r) \leftarrow \mathsf{rresp}(\mathsf{msk}, \mathsf{id}_S, M_r),$$
$$(b_r, \mathsf{cred}) \leftarrow \mathsf{rcheck}(\mathsf{st}, \mathsf{cid}, R_r).$$

If $b_r = 0$, *output 0. Otherwise, set* $\mathsf{rcs}[\mathsf{cid}] := \mathsf{cred}$.
3) *For each* $(\mathsf{cid}, \mathsf{cred}) \in R_{betw}$, *set* $\mathsf{rcs}[\mathsf{cid}] := \mathsf{cred}$.
4) *Run the authentication protocol* Auth *of* T *at* S, *which is as follows:*

$$(c, \mathsf{st}) \leftarrow \mathsf{achall}(\mathsf{id}_S), M_a \leftarrow \mathsf{acomm}(\mathsf{id}_S, c),$$
$$R_a \leftarrow \mathsf{aresp}(\mathsf{msk}, \mathsf{id}_S, \mathsf{cid}, M_a),$$
$$b_a \leftarrow \mathsf{acheck}(\mathsf{st}, \mathsf{rcs}, \mathsf{cid}, R_a).$$

5) *Return* b_a.

We assume that cid *derived in Step 2 is not in the list* R_{betw}.

Definition 16 (Session Identifiers and Partnering)**.** *Let* PLA *be a PlA and consider the oracles from Definition 1. Let* V_t *be a function that takes as input the transcript* $tr_T^{i,j} = (\mathsf{id}_S, \mathsf{cid}, M, R)$ *that a token* $T \in \mathcal{T}$ *observes in an oracle call to* $\mathsf{Challenge}(\pi_T^{i,j}, \cdot, \cdot, \cdot)$, *and outputs a bitstring* $V_t(tr_T^{i,j})$. *Similarly, let* V_s *be a function that takes as input the transcript* $tr_S^{i,j} = (c, \mathsf{cid}, R)$ *that a server* $S \in \mathcal{S}$ *observes in oracle calls to* $\mathsf{Start}(\pi_S^{i,j}), \mathsf{Complete}(\pi_S^{i,j}, \cdot, \cdot)$, *and outputs a bitstring* $V_s(tr_S^{i,j})$. *We assume that these functions are specified by* PLA.

We say that handles $\pi_T^{i,j}$ *and* $\pi_S^{i',j'}$ *are partnered if the following hold:*

$$(j = 0 \iff j' = 0) \wedge V_t(tr_T^{i,j}) = V_s(tr_S^{i',j'}).$$

Appendix C.
Omitted Proofs for Existing Implementations

C.1. Omitted Proofs For Key Derivation Function

Proof of Lemma 3. We show the statement via a sequence of games. For each game $\mathbf{G}_i$, we denote the probability that $\mathbf{G}_i$ outputs 1 by pr_i. We note that games $\mathbf{G}_1, \mathbf{G}_2$ are taken verbatim from the proof of Lemma 1.

Game $\mathbf{G}_0$: This game is the real medium unlinkability game. Recall that at the beginning of this game, a master secret key $\mathsf{msk}_T \xleftarrow{\$} \{0,1\}^{\lambda}$ is generated for each token $T \in \mathcal{T}$. The adversary $\mathcal{A}$ gets access to oracles $\mathsf{Start}, \mathsf{Challenge}, \mathsf{Complete}$. Then, it outputs two tokens T_0, T_1 and servers S_L, S_R. Afterwards, it also gets access to oracles $\mathsf{Left}, \mathsf{Right}$, which internally run $\mathsf{Challenge}(\pi_{T_b}^{i_b, j_b}, \mathsf{id}_{S_L}, \cdot, \cdot)$ and $\mathsf{Challenge}(\pi_{T_{1-b}}^{i_{1-b}, j_{1-b}}, \mathsf{id}_{S_R}, \cdot, \cdot)$, respectively. By definition, we have

$$\mathbf{Adv}^{\mathcal{A}}_{\mathbf{mUnl}, \mathsf{kdfPA}} = \left| \mathsf{pr}_0 - \frac{1}{2} \right|.$$

Game G_1: This game is as G_0, but we introduce a bad event and let the game abort if it occurs. Consider the lists $\mathcal{L}_{lr}^r$ and $\mathcal{L}_{ch}^r$ as in the definition of the unlinkability game. Game G_1 aborts if there is a collision, i.e. $\mathcal{L}_{lr}^r \cap \mathcal{L}_{ch}^r \neq \emptyset$. Recall that the values cid is sampled uniformly at random, and $|\mathcal{L}_{lr}^r| \leq 2$. Therefore, we have

$$|\mathsf{pr}_0 - \mathsf{pr}_1| \leq \frac{2Q_C}{2^\lambda}.$$

Game G_2: This game is as G_1, but we change the way the secret signing keys sk used in queries of the form $\mathsf{Challenge}(\pi_T^{i,j}, \mathsf{id}_S, \mathsf{cid}, \cdot)$ or queries to Left and Right are derived. This is similar to the proof of impersonation security. Recall that before they were derived as $\mathsf{sk} := \mathsf{PRF}(\mathsf{msk}_T, (\mathsf{cid}, \mathsf{id}_S)$ for each token $T \in \mathcal{T}$. Now, we introduce an initially empty map $K_T[\cdot]$ for each token $T \in \mathcal{T}$. Then, whenever G_1 would compute sk as above, game G_2 first checks if $K_T[\mathsf{cid}, \mathsf{id}_S]$ is already defined. If it is not, it samples a secret key sk uniformly at random, and sets $K_T[\mathsf{cid}, \mathsf{id}_S] := \mathsf{sk}$. In any case, it uses $\mathsf{sk} := K_T[\mathsf{cid}, \mathsf{id}_S]$ to continue. As in the proof of impersonation security, we can use the pseudorandomness of PRF in $|\mathcal{T}|$ hybrid steps to bound the distinguishing advantage between G_1 and G_2. To subsequent hybrids are shown to be close via a straight-forward reduction $\mathcal{B}$, and we get

$$|\mathsf{pr}_1 - \mathsf{pr}_2| \leq |\mathcal{T}| \cdot \mathbf{Adv}_{\mathbf{prf}, \mathsf{PRF}}^{\mathcal{B}}.$$

Game G_3: This game is as G_2, but we split the map K_{T_b} into two maps. To recall, T_b is the challenge token used in oracle Left. In game G_3, instead of having one map K_{T_b} that is used in $\mathsf{Challenge}(\pi_{T_b}^{i,j}, \cdot)$ and Left, we now only use K_{T_b} in $\mathsf{Challenge}(\pi_{T_b}^{i,j}, \cdot)$, and a separate independent map $K_L[\cdot]$ in oracle Left. It follows from credential separation and the bad event that we ruled out in G_1 that this is only a conceptual change. Namely, the change can only be noticed, if the game accesses both $K_{T_b}[\mathsf{cid}, \mathsf{id}_S]$ and $K_L[\mathsf{cid}, \mathsf{id}_S]$ for some $(\mathsf{cid}, \mathsf{id}_S)$. By definition of oracle Left, this means that $S = S_L$, and therefore credential separation applies. We get

$$\mathsf{pr}_2 = \mathsf{pr}_3.$$

Game G_4: This game is as game G_3, with a similar change as the previous one but for oracle Right and token T_{1-b}. Namely, we split the map $K_{T_{1-b}}$ into two maps. The map $K_{T_{1-b}}$ is used in $\mathsf{Challenge}(\pi_{T_{1-b}}^{i,j}, \cdot)$ and a map $K_R[\cdot]$ is used in oracle Right. Similarly as above, we get

$$\mathsf{pr}_3 = \mathsf{pr}_4.$$

We highlight that the steps from G_2 to G_4 would not be correct if we tried to prove strong unlinkability. This is because in this case, $\mathcal{A}$ would be allowed to come up with some fresh cid and submit it to oracle Left and Challenge, which react consistently in G_2 and inconsistently in G_3.

Finally, we see that the behavior of oracles Left and Right in game G_4 and thus $\mathcal{A}$'s view is independent of the bit b. Thus, we have $\mathsf{pr}_4 = 1/2$. This shows the claim. $\square$

C.2. Omitted Proofs For Key Wrapping

Proof of Lemma 2. We show the statement by presenting a sequence of games G_i, where for each such game G_i the probability that the game outputs 1 is denoted by $\mathbf{Adv}_i$.

Game G_0: This is the real impersonation game. At the beginning of this game, every token $T \in \mathcal{T}$ is initialized with a master secret key $\mathsf{msk}_T \leftarrow \mathsf{Gen}(\mathsf{par})$. Then, the adversary gets access to oracles Start, Challenge, Complete. By definition, we have

$$\mathbf{Adv}_0 = \mathbf{Adv}_{\mathbf{Imp}, \mathsf{kwrPA}}^{\mathcal{A}}.$$

Game G_1: We change the game as follows. The game is as G_0, but it aborts if for $x \neq x'$ we have $H(x) = H(x')$. The images of H are sampled uniformly at random from $\{0,1\}^{2\lambda}$, which implies that

$$|\mathbf{Adv}_0 - \mathbf{Adv}_1| \leq \frac{Q_H^2}{2^{2\lambda}}.$$

Game G_2: This game is as G_2, but we introduce another abort. To this end, consider the server-side oracles Start. Recall that during the execution of such an oracle, a random $rs \xleftarrow{\$} \{0,1\}^{\geq \lambda}$ is sampled. Game G_2 aborts if the same rs is sampled in two different invocations of the oracle Start. Clearly, we have the bound

$$|\mathbf{Adv}_1 - \mathbf{Adv}_2| \leq \frac{Q_S^2}{2^\lambda}.$$

Game G_3: In this game, we will no longer generate master secret keys msk_T for each token. Recall that these keys are used in the previous games to encrypt signing keys via $\mathsf{cid} \leftarrow \mathsf{Enc}(\mathsf{msk}_T, (\mathsf{id}_S, \mathsf{sk}))$ in queries of the form $\mathsf{Challenge}(\pi_T^{i,0}, \mathsf{id}_S, \cdot)$ (i.e. in registration) and to decrypt such cid in queries of the form $\mathsf{Challenge}(\pi_T^{i,j}, \mathsf{id}_S, \mathsf{cid}, \cdot)$, $j > 0$ (i.e. in authentication). In game G_2, we instead hold an initially empty map $L[\cdot, \cdot]$. In each registration query $\mathsf{Challenge}(\pi_T^{i,0}, \mathsf{id}_S, \cdot)$ the value cid is now sampled uniformly at random from $\{0,1\}^{\nu^*}$, and an entry $L[T, \mathsf{cid}] := (\mathsf{id}_S, \mathsf{sk})$ is added. In each authentication query $\mathsf{Challenge}(\pi_T^{i,j}, \mathsf{id}_S, \mathsf{cid}, \cdot), j > 0$, the entry $(\mathsf{id}, \mathsf{cid}) := L[T, \mathsf{cid}]$ is retrieved from the map instead of decrypting cid, and it is used if it is defined. If it is undefined, the oracle aborts its execution.

We can bound the distinguishing advantage between G_2 and G_3 using $|\mathcal{T}|$ intermediate hybrids. Namely, in hybrid i, we apply the change to the first i tokens $T \in \mathcal{T}$. For each hybrid step we can give a straight-forward reduction $\mathcal{B}$ from the anonymous authentication security of SKE. Thus, we have

$$|\mathbf{Adv}_1 - \mathbf{Adv}_2| \leq |\mathcal{T}| \cdot \mathbf{Adv}_{\mathbf{anon\text{-}auth}, \mathsf{SKE}}^{\mathcal{B}}.$$

Now, if we take a look at G_3, we see that each challenge that the adversary gets via oracle Start and has to sign is unique. Also, signing keys sk are only needed to sign, and not in the plain anymore. Thus, similar to the proof of Lemma 1 we can build a reduction $\mathcal{C}$ from the **euf-cma** security of SIG to bound $\mathbf{Adv}_3$. At a high level, the reduction guesses

in which query of the form $\mathsf{Challenge}(\pi_T^{i,0}, \cdot)$ the adversary obtained the public key, for which it forges a signature. We have

$$\mathbf{Adv}_3 \leq Q_C \cdot \mathbf{Adv}^{\mathcal{C}}_{\mathbf{euf\text{-}cma},\mathsf{SIG}}.$$

□

Proof of Lemma 4. We show the claim by presenting a sequence of games $\mathbf{G}_i$. For each game $\mathbf{G}_i$, we denote the probability that it outputs 1 by pr_i.

Game $\mathbf{G}_0$: This game is the real strong unlinkability game. In this game, a key $\mathsf{sk}_T \leftarrow \mathsf{Gen}(\mathsf{par})$ is generated for each token $T \in \mathcal{T}$. Then, the adversary $\mathcal{A}$ gets access to oracles $\mathsf{Start}, \mathsf{Challenge}, \mathsf{Complete}$ and outputs tokens T_0, T_1 and servers S_L, S_R. Afterwards, it also gets access to oracles $\mathsf{Left}, \mathsf{Right}$, which run $\mathsf{Challenge}(\pi_{T_b}^{i_b, j_b}, \mathsf{id}_{S_L}, \cdot, \cdot)$ and $\mathsf{Challenge}(\pi_{T_{1-b}}^{i_{1-b}, j_{1-b}}, \mathsf{id}_{S_R}, \cdot, \cdot)$, respectively. By definition, we have

$$\mathbf{Adv}^{\mathcal{A}}_{\mathbf{sUnl},\mathsf{kwrPA}} = \left| \mathsf{pr}_0 - \frac{1}{2} \right|.$$

Game $\mathbf{G}_1$: We change game $\mathbf{G}_0$ in the following way. In the beginning, $\mathbf{G}_1$ samples $T_0^*, T_1^* \xleftarrow{\$} \mathcal{T}$. Later, if $T_0 \neq T_0^*$ or $T_1 \neq T_1^*$, the game returns a random bit. Otherwise, it continues as $\mathbf{G}_0$ does. As $\mathcal{A}$ obtains no information about T_0^*, T_1^* until the potential abort, we have

$$\left| \mathsf{pr}_1 - \frac{1}{2} \right| = \frac{1}{|\mathcal{T}|^2} \left| \mathsf{pr}_0 - \frac{1}{2} \right|.$$

Game $\mathbf{G}_2$: We change $\mathbf{G}_1$ in the following way. In $\mathbf{G}_1$, whenever $\mathcal{A}$ starts a registration with token $T_r, r \in \{0, 1\}$ via oracles $\mathsf{Challenge}$ or $\mathsf{Left}, \mathsf{Right}$, a ciphertext cid is created as $\mathsf{cid} := \mathsf{Enc}(\mathsf{msk}_{T_r}, (\mathsf{id}_S, \mathsf{sk}))$. Furthermore, when $\mathcal{A}$ starts an authentication interaction with token $T_r, r \in \{0, 1\}$ via oracles $\mathsf{Challenge}$ or $\mathsf{Left}, \mathsf{Right}$, it provides a ciphertext cid, which is then decrypted as $(\mathsf{id}, \mathsf{sk}) := \mathsf{Dec}(\mathsf{msk}, \mathsf{cid})$. If Dec returns $\bot$ or id does not match with the server identity id_S provided, the oracles abort. Otherwise, they continue their execution using secret key sk.

Now, in game $\mathbf{G}_2$, we change this encryption and decryption for the tokens T_0^*, T_1^* that we guessed in $\mathbf{G}_1$. Note that if $\mathbf{G}_1$ does not abort, we know that $(T_0^*, T_1^*) = (T_0, T_1)$ and these tokens are also used in oracles Left and Right. Concretely, the game works as follows: Initially, two empty maps $L_0[\cdot], L_1[\cdot]$ are initialized. Then, in each registration interaction with token $T_i^*, i \in \{0, 1\}$ (including the ones in oracle Left or Right) the value cid is sampled randomly from $\{0, 1\}^{\nu^*}$. Then, an entry $L_i[\mathsf{cid}] := (\mathsf{id}_S, \mathsf{sk})$ is added. Furthermore, in each authentication interaction with token $T_i^*, i \in \{0, 1\}$ (including the ones in oracle Left or Right), where the adversary provides cid, we check if $L_i[\mathsf{cid}]$ is defined. If it is, we use it instead of decrypting cid. If it is not defined, we abort this interaction. A direct reduction $\mathcal{B}$ from the anonymous authentication security of SKE shows that

$$|\mathsf{pr}_1 - \mathsf{pr}_2| \leq \mathbf{Adv}^{\mathcal{B}}_{\mathbf{anon\text{-}auth},\mathsf{SKE}}.$$

We note that now, the only dependence of $\mathcal{A}$'s view on bit b is the shared use of the maps L_0, L_1. Namely, the table L_b

is used by challenge oracles for token T_b and by the oracle Left.

Game $\mathbf{G}_3$: Similar to game $\mathbf{G}_2$, but we partition map L_b into two tables: The first map, L_b', is used in oracle $\mathsf{Challenge}$ for token T_b. The second map, L_L is used in oracle Left. The maps are used as before, and the difference is that oracle $\mathsf{Challenge}$ never accesses L_L and oracle Left never accesses L_b'.

We claim that the view of $\mathcal{A}$ does not change from $\mathbf{G}_2$ to $\mathbf{G}_3$. To see this, note that $\mathcal{A}$ can only observe the change, if the same cid is given out in registration twice (once in $\mathsf{Left}, \mathsf{Right}$, and once in $\mathsf{Challenge}$), or it sends a value cid to one oracle in authentication (e.g. $\mathsf{Challenge}$), which was given out by the other oracle (e.g. Left) in registration. The former only occurs with probability at most $2Q_C / |\{0, 1\}^{\nu^*}|$, and the latter is forbidden due to (strong) credential separation. It follows that

$$|\mathsf{pr}_2 - \mathsf{pr}_3| \leq \frac{2Q_C}{2^{\nu^*}}.$$

Game $\mathbf{G}_4$: This game is as $\mathbf{G}_3$, but we partition the map L_{1-b} into two tables L_{1-b}' and L_R. The change is similar as above, and the same argument shows

$$\mathsf{pr}_3 = \mathsf{pr}_4.$$

We note that in $\mathbf{G}_4$, the view of $\mathcal{A}$ is independent of bit b, which implies that $\mathsf{pr}_4 = 1/2$. □