

Domain-Specific Pseudonymous Signatures for the German Identity Card

Jens Bender¹, Özgür Dagdelen², Marc Fischlin², and Dennis Kügler¹

¹ Bundesamt für Sicherheit in der Informationstechnik (BSI), Germany

² Technische Universität Darmstadt, Germany

Abstract. The restricted identification protocol for the new German identity card basically provides a method to use pseudonyms such that they can be linked by individual service providers, but not across different service providers (even not malicious ones). The protocol can be augmented to allow also for signatures under the pseudonyms. In this paper, we thus view —and define— this idea more abstractly as a new cryptographic signature primitive with some form of anonymity, and use the term domain-specific pseudonymous signatures. We then analyze the restricted identification solutions in terms of the formal security requirements.

1 Introduction

The protocols designed for the new German identity card [7] secure the communication between the card and the reader resp. the terminal. See Figure 1 for an overview. Initially, the parties run the password-authenticated connection establishment protocol (PACE) to derive a secure channel between the card and the local reader, with the user’s consent who needs to enter the password. Then the card and the remote end point, called the terminal or service provider, run the extended access protocol (EAC) to authorize mutually and to establish another secure channel between these points. In the card authentication step anonymity is guaranteed by the fact that a large number of cards share the same secret. The security of these protocols has been analyzed in [4, 6, 12].

The overall design of the identity card includes another protocol, called restricted identification. In this optional protocol, card holders can use domain-specific pseudonyms to interact with service providers such that (a) a service provider can recognize pseudonyms of individual cards and use this information for the service (domain-specific linkability), and (b) different service providers cannot link interactions of one user in their respective domains (cross-domain anonymity). Although the concept of restricted identification in [7] —and the Diffie-Hellman based solution— currently only support recognition of pseudonyms, it can be easily




Fig. 1. Protocols for the German ID Card

extended for more functionality, by allowing users to create signatures under their pseudonyms.

1.1 Domain-Specific Pseudonymous Signatures

To analyze the protocols of the new German identity card we formally introduce the concept of *domain-specific pseudonymous signatures* which captures the scenario of the restricted identification protocols, augmented through signatures, but also allows to reason about the security of the basic restricted identification protocol.¹ In a sense, (domain-specific) pseudonymous signatures can be seen as a relaxed version of group signatures with a limited form of linkability: While group signature schemes, as formalized in [3], provide a very strong form of anonymity, preventing an adversary to identify signers even when knowing the secret keys, pseudonymous signatures are designed specifically to allow a well-defined verifier to link signatures.

The group-signature ancestry also lays the ground to model the security of pseudonymous signatures. As mentioned, according to [3], secure group signature schemes should satisfy full anonymity, the resistance to identify the origin of signatures, and full traceability, the ability of the group manager to trace the origin of a signature with the help of some trapdoor information. The latter property implies, for example, unforgeability (because a forgery could not be traced) and non-frameability, saying that a set of malicious parties, potentially including the group manager, could falsely blame an honest user (which would again contradict traceability).

Concerning anonymity, we note that the goal of restricted identification is exactly to allow a service provider to link previously seen pseudonyms. We thus relax the security requirement for pseudonymous signatures and only demand cross-domain anonymity, i.e., the inability to link signatures given for different service providers, even if both of them are malicious. Furthermore, since the restricted-identification scenario does not involve an authority for traceability, we revert to unforgeability directly for our setting, meaning that one cannot forge signatures on behalf of honest users.

We also introduce another property saying that it is infeasible to make the verifier accept a signature for an invalid domain-specific pseudonym. This assumes a white- or black-listing approach, in the former case making sure that valid pseudonyms are those in the white list, and in the other case demanding that the verifier would not accept a black-listed pseudonym. We observe that this property contrasts with unforgeability which merely protects honest signers from signature forgeries under their name. While this additional property, which we call *seclusiveness*, follows from full-traceability for group signatures, we need to state it explicitly in our setting without a tracing authority. Roughly, unforgeability together with seclusiveness provide a weaker, yet “best-we-can-hope-for” form of full-traceability for domain-specific pseudonymous signatures.

¹ Interestingly, the term “pseudonymous signatures” has occasionally been already mentioned in the literature, typically referring to regular signatures under pseudonyms [14], but not considered formally, from a cryptographic point of view.

1.2 Analysis of the Restricted-Identification Solutions

We then continue to analyze the current solution for restricted identification (and its signature-augmented version). The basic solution is roughly to give users a random value x_1 and the service provider a certified group element R , such that the user will derive the domain-specific pseudonym as the Diffie-Hellman key $I_R = R^{x_1}$ (or, to be precise, the hash value thereof). This value is then sent over a previously established secure channel to the service provider (but we ignore the channel part in our analysis).

To achieve the additional unforgeability and seclusiveness properties in the signature-augmented case the user is instead given a random representation (x_1, x_2) of the authority's public key y , i.e., $y = g_1^{x_1} g_2^{x_2}$ for generators $g_1, g_2 = g_1^z$. Here, the authority can create such representations easily with the knowledge of $z = \log_{g_1} g_2$. Each domain will again hold a certified group element R , and the user will derive once more the domain-specific pseudonym as the Diffie-Hellman key $I_R = R^{x_1}$ and transmits this value to the provider.

Obviously, if two malicious users could pool their distinct secrets (x_1, x_2) and (x'_1, x'_2) then they could recover the authority's secret key z . However, we note that these secrets are protected through the hardware of the identity card and are not available to users. In a sense, our security analysis relies on this fact, but at the same time even allows a single (malicious) user access to its secret key. In case of suspicion of leakage of the authority's secret key z a new key can be generated, the old key can be revoked, but white-listing can be used to mark the user keys under the old key as still valid.

In the signature-augmented version, the user additionally signs a message m by giving two intertwined non-interactive proofs of knowledge (where the message enters the hash evaluation to derive the challenge non-interactively in the random oracle model). This first proof shows that the user knows the discrete logarithm x_1 of I_R to base R , and the other proof shows that it additionally knows x_2 such that (x_1, x_2) forms a representation of $y = g_1^{x_1} g_2^{x_2}$. The first proof is basically a Schnorr proof of knowledge (PoK) [21], and the second one is an Okamoto PoK [17] but re-using the data from the Schnorr proof.

We show that the basic restricted-identification protocol ensures cross-domain anonymity according to our notion. Since pseudonyms $I_R = R^{x_1}$ are essentially Diffie-Hellman keys, this holds under the Decisional Diffie-Hellman assumption. For the signature-augmented version we show that unforgeability holds as well (in the random oracle model, and under the discrete log assumption). This follows from the Schnorr PoK for x_1 . Seclusiveness follows from the Okamoto proof showing that the pair (x_1, x_2) is a representation of y and has thus been issued by the authority. Here, however, due to the construction we require that no two malicious users collaborate. As explained above, this is prevented by hiding and protecting the secret of users within the secure hardware of the chip. Thereby, users can use their secrets, however, they cannot extract the keys to collude. Note that the basic version does not implement a signature functionality and thus cannot satisfy these two notions.

1.3 Related Work

As noted above, domain-specific pseudonymous signatures can be seen as descendants of group signatures, only with weaker anonymity requirements but with domain-specific verifiers. The related concept of ring signatures [20] can be viewed as “ad-hoc” group signatures without a central manager. For such schemes, rings of users can be formed at will, and one user can sign on behalf of the ring, but there is usually no mean to identify the actual signer later. As such, domain-specific pseudonymous signatures provide stronger notions of traceability than ring signatures, but a weaker form of anonymity—the domain holder can link signatures.

We note that credential systems [5, 8] are very similar to our pseudonymous signatures, but diverge in some important aspects. Most importantly, credential systems typically provide multi-show unlinkability, as opposed to our pseudonymous signatures. As such, solutions for multi-show credential systems are usually slightly more complex [9–11]. For one-show solutions, where the user can use a credential (under a pseudonym) only once, domain-specific linkability—and therefore cross-domain anonymity—has not been considered before. We also note that the question of turning cross-domain pseudonymous signatures into full-fledged multi-show credential systems is beyond the scope of the paper here: the requirement of recognizing pseudonyms for domain holders is inherent in the application requirement.

Finally, let us point out that our notion of domain-specific pseudonymous signatures is close to a recent proposal of Bernhard et al. [2] for defining direct anonymous attestation (DAA). Their security definition for DAA also resembles the one for group signature schemes but comes with a limited form of linkability: signatures of the same user in the same domain (called base there) should be publicly linkable. At the same time, one should be able to identify signatures given the user’s secret key, a requirement which we do not impose in our setting. Moreover, their scenario assumes that linkability must be enforced via cryptographic means, domain-specific pseudonymous signatures allow this by default through the pseudonyms. Besides minor technical differences concerning security in presence of compromised keys or incorporating black-listing, the main difference to their setting is that our model takes into account the extra layer of domain-specific pseudonyms and its unlinkability to the pseudonym layer.

Similarly, Wei [22] also uses DAA as a motivation for his work on tracing-by-linking group signatures, but he considers a weaker form of anonymity where a signer’s identity is only hidden if he signs at most a fixed number of times for $k \geq 1$. Any additional signature enables to trace back the identity of the signer by a public algorithm. The approach also includes a notion of k -linkability which corresponds to (public) traceability and essentially means that more than k signatures are linkable and allow to identify the origin. It also contains a notion of non-slanderness which means that no group of malicious users (including the issuing authority) can sign more than k times such that it points to an honest user outside of the group. This primitive does not help in our scenario, because we demand unlinkability only among different domain sectors. Furthermore, our

notion of seclusiveness, in contrast with non-slanderiness which addresses the linkability of more than k signatures, refers to the fact that one cannot produce any signature on behalf of honest users.

The idea of extending restricted identification to allow unlinkable sector signatures has appeared concurrently in [15]. The scheme does not cover the issue of seclusiveness, though, and is less explicit about the underlying security model, e.g., it remains unclear if unlinkability holds for multiple signatures. Exploring such questions and providing sound models, including issues related to black-listing or white-listing, and to prove security according to these models for the (augmented) restricted identification protocol of the new German identity cards is our contribution here.

2 Domain-Specific Pseudonymous Signatures

To simplify the notation we denote by $\{a\}_n$ a set of n elements a_1, \dots, a_n where the indices $i = 1, 2, \dots, n$ are implicit.

2.1 Preliminaries

Below we define our domain-specific pseudonymous signature scheme first for static groups. Since both the secret keys of users and the domain keys are chosen by the authority, we can imagine that a sufficiently large set is chosen at the outset, and individual entries only become active if required. We hence assume an algorithm NymKGen generating the group manager's key pair, as well as sufficiently many pseudonyms nym (and secret keys $\text{gsk}[\text{nym}]$) and public domain keys dpk . Each pseudonym and its secret key can now be combined via an algorithm NymDSGen to build a domain-specific pseudonym $\text{dsnym} = \text{nym}[\text{dpk}]$ which, together with $\text{gsk}[\text{nym}]$ can be used to sign messages.

In addition, we can also introduce the concept of black-lists to the model. For this, we assume that the list is represented by a (dynamically updated) set \mathbf{B} of domain-specific pseudonyms which have been revoked earlier. We note that, alternatively, one could employ white-listing where valid pseudonyms are those included in the list; to revoke such a pseudonym one simply removes the entry from the white-list.

Definition 1 (Domain-Specific Pseudonymous Signature). *A domain-specific pseudonymous signature scheme is a collection of the following efficient algorithms $\mathcal{NYS} = (\text{NymKGen}, \text{NymDSGen}, \text{NymSig}, \text{NymVf})$ defined as follows.*

$\text{NymKGen}(1^\kappa, 1^n, 1^d)$ is a probabilistic algorithm which, on input a security parameter 1^κ and parameters $1^n, 1^d$ (both polynomial in κ) outputs a pair $(\text{gpk}, \text{gmsk})$ where gpk is the group public key and gmsk the secret key of the group manager, and outputs n (unique) pseudonyms nym with their corresponding secret keys $\text{gsk}[\text{nym}]$, and d domain descriptions dpk .

$\text{NymDSGen}(\text{nym}, \text{gsk}[\text{nym}], \text{dpk})$ is a deterministic algorithm which maps a pseudonym nym (and its secret key $\text{gsk}[\text{nym}]$) and the domain dpk to a domain-specific pseudonym $\text{dsnym} = \text{nym}[\text{dpk}]$.

$\text{NymSig}(\text{dsnym}, \text{gsk}[\text{nym}], \text{dpk}, m)$ is a probabilistic algorithm which, on input a domain-specific pseudonym dsnym , a secret key $\text{gsk}[\text{nym}]$, a domain dpk , and message m , outputs the signature σ of m under $\text{gsk}[\text{nym}]$ for domain dpk .

$\text{NymVf}(\text{gpk}, \text{dsnym}, \text{dpk}, m, \sigma, \mathbf{B})$ is a deterministic algorithm which, on input a message m and a signature σ together with the group public key gpk , a domain-specific pseudonym dsnym , the domain's key dpk , and a list \mathbf{B} , outputs either 1 (=valid) or 0 (=invalid).

We assume the usual completeness property, that for any honestly generated parameters, domain-specific pseudonyms, and signatures the verification algorithm accepts and outputs 1.

Note that we can assume that the group manager uses some standard way of certification for the public keys dpk given out to registered verifiers, and that the signing and verification algorithms check the validity of the keys. We thus often omit this step from the description of protocols.

2.2 Cross-Domain Anonymity

Cross-domain anonymity protects against linking pseudonyms across different domains —within one domain pseudonyms are meant to be linkable. We define cross-domain anonymity via a game between the adversary and a left-or-right oracle which, given two pseudonyms, a message, and a domain, generates a signature for either the left or the right pseudonym in the domain, according to a secret random bit b . We assume that the adversary can make adaptive calls to this left-or-right oracle; this is necessary since we cannot apply a hybrid argument to reduce such multiple queries to a single one (as opposed to the case of full-anonymity for group signatures as in [3], where this is possible since anonymity even holds if the adversary knows the users' secret keys). The adversary's goal is to predict b significantly beyond the pure guessing probability (condition (a) below).

In addition to challenge queries to the left-or-right oracle, the adversary may decide to blacklist domain-specific pseudonyms, to create additional signatures via NymSig , and to corrupt users. For simplicity, we define a version for *static corruptions* where all corruptions are made at the outset, before any other oracle calls are made, and discuss the adaptive version briefly below. To exclude trivial attacks, we must take into account that domain-specific pseudonyms are in principle linkable by the domain holder. Hence, in “transitivity” attacks where the adversary asks the left-or-right oracle first about a signature for domain-specific pseudonyms $\text{dsnym}_0, \text{dsnym}_1$ and then for $\text{dsnym}'_0, \text{dsnym}'_1$ for the same domain, but $\text{dsnym}_1 \neq \text{dsnym}'_1$, the signatures would point to the same domain-specific pseudonym if and only if the oracle signs under the left pseudonym. We thus exclude such queries in condition (b) below.

We require another case to be excluded. Namely, the additional signatures generated through NymSig cannot hide the pseudonyms behind the signatures; this would require further means like anonymous signatures and would only work if signatures are not publicly verifiable [13, 23]. Since such extra signatures for domain-specific pseudonyms would thus also allow to link the origin in the left-or-right queries to the pseudonyms, we must disallow the adversary from querying NymSig about domain-specific pseudonyms, which are also used in left-or-right queries (condition (c) below).

The model below assumes, to the advantage of the adversary, that all pseudonyms, all domain keys, and domain-specific pseudonyms are known at the outset; only the assignment of pseudonyms to domain-specific pseudonyms remains hidden. Nonetheless, we assume that the relation of domains and domain-specific pseudonyms is known (see below for the motivation). All this is captured by giving the adversary the corresponding data as sets, with $W \odot D$ being the set of all domain-specific pseudonyms dsnym and the corresponding domains dpk . The adversary will thus attack domain-specific pseudonyms from W . As common, we measure the adversary's running time including also all steps of honest parties, and covering both phases of the adversary.

In the definition, we presume that domain-specific pseudonyms are unique within a domain; global uniqueness can then be trivially achieved by attaching the domain key to the pseudonym. Indeed, domain-specific uniqueness will be later ensured by the unforgeability property anyway.

Definition 2 (Cross-Domain Anonymity). *A domain-specific pseudonymous signature scheme $\mathcal{NYS} = (\text{NymKGen}, \text{NymDSGen}, \text{NymSig}, \text{NymVf})$ is (n, d, t, Q, ϵ) cross-domain anonymous with $Q = (q_c, q_s, q_t)$ if for any algorithm \mathcal{A} running in time t , and making at most q_c queries to the corruption oracle, q_s queries to the signing oracle, and q_t queries to the left-or-right oracle, the probability that the following experiment returns 1 is at most ϵ :*

Experiment CD – Anon $_{\mathcal{A}}^{\mathcal{NYS}}(\kappa, n, d)$

$b \xleftarrow{\$} \{0, 1\}$

$T, S, W, B, C, N, D \leftarrow \emptyset$

$(\text{gpk}, \text{gmsk}, \{\text{gsk}[\text{nym}]\}_n, \{\text{nym}\}_n, \{\text{dpk}\}_d) \leftarrow \text{NymKGen}(1^\kappa, 1^n, 1^d)$

$N = \{\text{nym}\}_n$, and $D = \{\text{dpk}\}_d$

$W = \{\text{NymDSGen}(\text{nym}, \text{gsk}[\text{nym}], \text{dpk}) \mid \text{nym} \in N, \text{dpk} \in D\}$

$W \odot D := \{(\text{NymDSGen}(\text{nym}, \text{gsk}[\text{nym}], \text{dpk}), \text{dpk}) \mid \text{nym} \in N, \text{dpk} \in D\}$

$st \leftarrow \mathcal{A}^{\text{Corrupt}}(\text{gpk}, W, W \odot D, N)$

If \mathcal{A} queries $\text{Corrupt}(\text{nym})$ on input $\text{nym} \in N$

– set $C \leftarrow C \cup \{\text{nym}\}$

– return $\text{gsk}[\text{nym}]$

$d \leftarrow \mathcal{A}^{B', \text{NymSig}', \text{LoR}}(st)$

If \mathcal{A} queries $B'(\text{dsnym})$ on input $\text{dsnym} \in W$

– set $B \leftarrow B \cup \{\text{dsnym}\}$

If \mathcal{A} queries $\text{NymSig}'(\text{dsnym}, \text{dpk}, m)$ on input $\text{dsnym} \in W \setminus B$, $\text{dpk} \in D$, and message m

– set $S \leftarrow S \cup \{(\text{dsnym}, \text{dpk}, m)\}$

- find $\text{nym} \in \mathbf{N}$ such that $\text{dsnym} = \text{NymDSGen}(\text{nym}, \text{gsk}[\text{nym}], \text{dpk})$
 - return $\text{NymSig}(\text{dsnym}, \text{gsk}[\text{nym}], \text{dpk}, m)$
- If \mathcal{A} queries $\text{LoR}(\text{dsnym}_0, \text{dsnym}_1, \text{dpk}, m)$ on input $\text{dsnym}_0, \text{dsnym}_1 \in \mathbf{W} \setminus \mathbf{B}$ $\text{dpk} \in \mathbf{D}$, and message m ,
- set $\mathbf{T} \leftarrow \mathbf{T} \cup \{(\{\text{dsnym}_0, \text{dsnym}_1\}, \text{dpk}, m)\}$
 - find $\text{nym}_0, \text{nym}_1 \in \mathbf{N} \setminus \mathbf{C}$ such that
 - $\text{dsnym}_i = \text{NymDSGen}(\text{nym}_i, \text{gsk}[\text{nym}_i], \text{dpk})$ for $i = 0, 1$
 - return \perp if no such $\text{nym}_0, \text{nym}_1$ exist,
 else return $\text{NymSig}(\text{dsnym}_b, \text{gsk}[\text{nym}_b], \text{dpk}, m)$

Return 1 iff

- (a) $d = b$ and
- (b) for any $(\{\text{dsnym}_0, \text{dsnym}_1\}, \text{dpk}, m), (\{\text{dsnym}'_0, \text{dsnym}'_1\}, \text{dpk}, m') \in \mathbf{T}$ we have either $\{\text{dsnym}_0, \text{dsnym}_1\} = \{\text{dsnym}'_0, \text{dsnym}'_1\}$ or $\{\text{dsnym}_0, \text{dsnym}_1\} \cap \{\text{dsnym}'_0, \text{dsnym}'_1\} = \emptyset$, and
- (c) for any $(\text{dsnym}, \text{dpk}, m) \in \mathbf{S}$ there is no dsnym', m' such that $(\{\text{dsnym}, \text{dsnym}'\}, \text{dpk}, m') \in \mathbf{T}$.

The probability is taken over all coin tosses of NymKGen , NymSig , and \mathcal{A} , and the choice of b .

We note that the adversary in our game always accesses oracles through their domain-specific pseudonyms. This is possible since \mathcal{A} knows the set \mathbf{W} of such pseudonyms (but not the relation to the pseudonyms nym). Having this list at the outset is motivated by the fact that the adversary can potentially collect such domain-specific pseudonyms when acting as a domain holder, where it gets to learn the domain-specific pseudonyms and signatures under these pseudonyms. Note that this also allows to link domain-specific pseudonyms dsnym to domain keys dpk , hence we give $\mathbf{W} \odot \mathbf{D}$ as additional input. This also implies that we cannot grant the adversary access to another signature oracle which it can provide $\text{nym}, \text{dpk}, m$ to get a signature for m under the corresponding dsnym ; it would be easy to check for the validity of the signature under the domain-specific pseudonym with the help of \mathbf{W} and to link dsnym to nym . In other words, one can link pseudonyms to their domain-specific pseudonyms given a signature under the pseudonym for the respective domain.

For an adaptive version, the adversary may interleave **Corrupt** queries with the other oracle queries arbitrarily. Then we must ensure that no nym in a **Corrupt** query has appeared in an **LoR** query before, declaring the adversary to lose if this is the case.

2.3 Unforgeability

Unforgeability of domain-specific pseudonymous signatures follows the basic paradigm for regular signatures: It should be infeasible to create a valid signature on behalf of an honest pseudonym for a previously unsigned message. This should even hold if the adversary knows the group manager's secret key (but not the user's secret key, else trivial attacks would be possible). Since for

unforgeability we do not need to hide the link between pseudonyms and domain-specific pseudonyms, we assume that the adversary simply knows the tuples $(\text{dsnym}, \text{nym}, \text{dpk})$ of domain-specific pseudonyms, and corresponding pseudonyms and domain keys. Below we say that the adversary wins if it manages to forge a signature under a domain-specific pseudonym dsnym^* which is potentially derived from *some* pseudonym nym in *some* domain dpk ; but the adversary does not to specify these values.

Our notion of unforgeability is weaker than the non-frameability property of group signatures in the sense that, even though the adversary may know the group manager's secret key, it must not collaborate with the group manager during generation. We again consider only the version of static corruptions, although here it is straightforward to capture adaptive corruptions by giving the adversary simply the corruption oracle in the second phase, too.

Definition 3 (Unforgeability). *The domain-specific pseudonymous signature scheme $\mathcal{NYS} = (\text{NymKGen}, \text{NymDSGen}, \text{NymSig}, \text{NymVf})$ is (n, d, t, q, ϵ) -unforgeable if any algorithm \mathcal{A} , running in time t and making at most q signing queries, makes the following experiment output 1 with probability at most ϵ :*

Experiment $\text{Unforge}_{\mathcal{A}}^{\mathcal{NYS}}(\kappa, n, d)$

$S, B, C, N, D \leftarrow \emptyset$

$(\text{gpk}, \text{gmsk}, \{\text{gsk}[\text{nym}]\}_n, \{\text{nym}\}_n, \{\text{dpk}\}_d) \leftarrow \text{NymKGen}(1^\kappa, 1^n, 1^d)$

$N = \{\text{nym}\}_n$, and $D = \{\text{dpk}\}_d$

$W \odot N \odot D = \{(\text{NymDSGen}(\text{nym}, \text{gsk}[\text{nym}], \text{dpk}), \text{nym}, \text{dpk}) \mid \text{nym} \in N, \text{dpk} \in D\}$

$st \leftarrow \mathcal{A}^{\text{Corrupt}}(\text{gpk}, \text{gmsk}, W \odot N \odot D)$

If \mathcal{A} queries $\text{Corrupt}(\text{nym})$ on input $\text{nym} \in N$

– set $C \leftarrow C \cup \{\text{nym}\}$

– return $\text{gsk}[\text{nym}]$

$(m^*, \sigma^*, \text{dsnym}^*) \leftarrow \mathcal{A}^{B', \text{NymSig}', \text{Corrupt}}(st)$

If \mathcal{A} queries $B'(\text{dsnym})$ on input $\text{dsnym} \in W$

– set $B \leftarrow B \cup \{\text{dsnym}\}$

If \mathcal{A} queries $\text{NymSig}'(\text{dsnym}, \text{dpk}, m)$ on input $\text{dsnym} \in W \setminus B$, $\text{dpk} \in D$, and message m

– set $S \leftarrow S \cup \{(\text{dsnym}, \text{dpk}, m)\}$

– find $\text{nym} \in N$ such that $\text{dsnym} = \text{NymDSGen}(\text{nym}, \text{gsk}[\text{nym}], \text{dpk})$

– return $\text{NymSig}(\text{dsnym}, \text{gsk}[\text{nym}], \text{dpk}, m)$

Output 1 iff there are $\text{nym}^* \in N \setminus C$ and $\text{dpk}^* \in D$ such that

(a) $\text{NymDSGen}(\text{nym}^*, \text{gsk}[\text{nym}^*], \text{dpk}^*) = \text{dsnym}^*$, and

(b) $\text{NymVf}(\text{gpk}, \text{dsnym}^*, \text{dpk}^*, m^*, \sigma^*, B) = 1$, and

(c) $(\text{dsnym}^*, \text{dpk}^*, m^*) \notin S$.

The probability is taken over all coin tosses of NymKGen , NymSig , and \mathcal{A} .

Note that conditions (a) and (c) also imply that domain-specific pseudonyms of different users (within a domain) cannot collide, except with negligible probability. Otherwise, the adversary may corrupt one of the two parties, and any signature created under the domain-specific pseudonym of the one party would

immediately constitute a forgery under the other party's pseudonym. (In the above model it would then be more appropriate to let the adversary in calls to NymSig' also specify nym , instead of searching for it; since the adversary knows the list $W \odot N \odot D$ it can look this value up.)

2.4 Seclusiveness

Seclusiveness finally considers the case that the verifier would accept a signature under a domain-specific pseudonym which has not been created by the authority. Note that this assumes that only the black-listed domain-specific pseudonyms are available, but not the universe of all created pseudonyms. This is indeed a valid assumption, following the suggestion in [7] about revocation for pseudonyms through black-lists, with no intention for white-lists. It is also clear that, unlike in case of unforgeability, we thus cannot allow the adversary to know the manager's secret key; else generating keys for additional users would be easy.

As in unforgeability, we again consider only the version of static corruptions. One captures adaptive corruptions by giving the adversary simply the corruption oracle in the second phase, too.

Definition 4 (Seclusiveness). *The domain-specific pseudonymous signature scheme $\mathcal{N}\mathcal{Y}\mathcal{M}\mathcal{S} = (\text{NymKGen}, \text{NymDSGen}, \text{NymSig}, \text{NymVf})$ is (n, d, t, Q, ϵ) -secluding with $Q = (q_c, q_s)$ if any algorithm \mathcal{A} , running in time t and making at most q_s signing queries and q_c corruption queries, makes the following experiment output 1 with probability at most ϵ :*

Experiment $\text{Sec}_A^{\mathcal{N}\mathcal{Y}\mathcal{M}\mathcal{S}}(\kappa, n, d)$

$W, B, C, N, D \leftarrow \emptyset$

$(\text{gpk}, \text{gmsk}, \{\text{gsk}[\text{nym}]\}_n, \{\text{nym}\}_n, \{\text{dpk}\}_d) \leftarrow \text{NymKGen}(1^\kappa, 1^n, 1^d)$

$N = \{\text{nym}\}_n$ and $D = \{\text{dpk}\}_d$

$W = \{\text{NymDSGen}(\text{nym}, \text{gsk}[\text{nym}], \text{dpk}) \mid \text{nym} \in N, \text{dpk} \in D\}$

$W \odot N \odot D$

$:= \{(\text{NymDSGen}(\text{nym}, \text{gsk}[\text{nym}], \text{dpk}), \text{nym}, \text{dpk}) \mid \text{nym} \in N, \text{dpk} \in D\}$

$st \leftarrow \mathcal{A}^{\text{Corrupt}}(\text{gpk}, W \odot N \odot D)$

If \mathcal{A} queries $\text{Corrupt}(\text{nym})$ on input $\text{nym} \in N$

– set $C \leftarrow C \cup \{\text{nym}\}$

– return $\text{gsk}[\text{nym}]$

$(m^*, \sigma^*, \text{dsnym}^*) \leftarrow \mathcal{A}^{B', \text{NymSig}'}(st)$

If \mathcal{A} queries $B'(\text{dsnym})$ on input $\text{dsnym} \in W$

– set $B \leftarrow B \cup \{\text{dsnym}\}$

If \mathcal{A} queries $\text{NymSig}'(\text{dsnym}, \text{dpk}, m)$ on $\text{dsnym} \in W \setminus B$, $\text{dpk} \in D$, and message m

– find $\text{nym} \in N$ such that $\text{dsnym} = \text{NymDSGen}(\text{nym}, \text{gsk}[\text{nym}], \text{dpk})$

– return $\text{NymSig}(\text{dsnym}, \text{gsk}[\text{nym}], \text{dpk}, m)$

Output 1 iff there exists $\text{dpk}^* \in D$ such that

(a) $\text{NymVf}(\text{gpk}, \text{dsnym}^*, \text{dpk}^*, m^*, \sigma^*, B) = 1$

(b) $\text{dsnym}^* \notin W$

The probability is taken over all coin tosses of NymKGen , NymSig , and \mathcal{A} .

3 Construction

The idea of the discrete-log based construction is as follows: The group manager will hold two generators $g_1, g_2 = g_1^z$ with $z \in \text{gmsk}$ for which it knows the discrete log with respect to each other. In addition, it will hold a public key $y = g_1^x$, such that it can easily compute many pairs (x_1, x_2) such that $y = g_1^{x_1} g_2^{x_2}$ with the help of z ; this is the trapdoor property of such values [1, 18]. Each user pseudonym nym will receive one of these pairs as its secret key $\text{gsk}[\text{nym}]$. The domain parameters are given by values $\text{dpk} = g^r$. A user can then compute the domain-specific pseudonym as $\text{dsnym} = \text{dpk}^{x_1}$.

To sign a message the user can then use common discrete-log based protocols (in the random oracle model) to show that (a) it knows the discrete-log x_1 of dsnym with respect to dpk , and (b) it knows a matching value x_2 to this discrete logarithm x_1 such that the pair forms a representation of y . Essentially, this is accomplished by running the non-interactive version of the Okamoto proof of knowledge [17] for x_1, x_2 and y to base g_1, g_2 , where the x_1 -part can simultaneously be used to show knowledge of x_1 of dsnym with respect to base dpk . As usual, the message to be signed enters the hash computations.

Construction 1. *The construction of the domain-specific pseudonymous signature scheme $\mathcal{NYSMS} = (\text{NymKGen}, \text{NymDSGen}, \text{NymSig}, \text{NymVf})$ is as follows:*

NymKGen($1^\kappa, 1^n, 1^d$): *Let $\mathcal{G} = \langle g \rangle$ be a (public) cyclic group of prime order q .*

We also assume a public hash function H , modeled as a random oracle in the security proofs. Choose $z \in_{\mathbb{R}} \mathbb{Z}_q$ randomly and calculate $g_1 := g$ and $g_2 := g^z$. Define $\text{gpk} := g_1^x$ for random $x \in_{\mathbb{R}} \mathbb{Z}_q$. To generate the secrets for the pseudonyms choose n random elements $x_{2,1}, \dots, x_{2,n} \in_{\mathbb{R}} \mathbb{Z}_q^$ and calculate $x_{1,i} = x - z \cdot x_{2,i}$ for $i = 1, 2, \dots, n$. Define $\text{gsk}[i] := (x_{1,i}, x_{2,i})$. By x_j we denote the $x_{j,i}$ when pseudonym i is clear from context. For the domain-parameters pick random $r_1, \dots, r_d \in_{\mathbb{R}} \mathbb{Z}_q^*$ and define $\text{dpk}_i := g^{r_i}$ for $i = 1, \dots, d$. Store z in gmsk . (Note that once the values $\text{gsk}[\cdot]$ have been output resp. given to the users, the group manager deletes them.)*

NymDSGen($\text{nym}, \text{gsk}[\text{nym}], \text{dpk}$): *Compute and output the domain-specific pseudonym $\text{nym}[\text{dpk}] := \text{dpk}^{x_1}$, which is also sometimes denoted as dsnym when nym and dpk are known from context.*

NymSig($\text{dsnym}, \text{gsk}[\text{nym}], \text{dpk}, m$): *Let $a_1 = g_1^{t_1} \cdot g_2^{t_2}$ and $a_2 = \text{dpk}^{t_1}$, for random $t_1, t_2 \in_{\mathbb{R}} \mathbb{Z}_q$. Compute $c = H(\text{dpk}, \text{dsnym}, a_1, a_2, m)$. Let $s_1 = t_1 - c \cdot x_1$ and $s_2 = t_2 - c \cdot x_2$. Then, output $\sigma = (c, s_1, s_2)$. (Note that in the Restricted-Identification protocol the user also sends dsnym which we can include here in the signature, in order to match the protocol description.)*

NymVf($\text{gpk}, \text{dsnym}, \text{dpk}, m, \sigma, \mathbb{B}$): *To verify a signature perform the following steps:*

1. *Parse $(c, s_1, s_2) \leftarrow \sigma$.*
2. *Let $a_1 = y^c \cdot g_1^{s_1} \cdot g_2^{s_2}$ and $a_2 = \text{dsnym}^c \cdot \text{dpk}^{s_1}$.*
3. *Output 1 iff $c = H(\text{dpk}, \text{dsnym}, a_1, a_2, m)$ and $\text{dsnym} \notin \mathbb{B}$.*

Revocation Mechanisms. We presented our construction above in terms of black-listing, revoking fraudulent domain-specific pseudonyms by listing them explicitly. Alternatively, and our definitions and constructions are robust in this regard, one can use a white-listing approach to list valid entries only. To represent the whitelisting approach in our framework any delisted domain-specific pseudonym from W will be put in B , such that $W \setminus B$ corresponds to the set of currently whitelisted entries. Checking for whitelisting thus corresponds to verifying that the entry is in $W \setminus B$ in our framework.

Blacklisting and whitelisting is performed by calculating the domain-specific pseudonym dsnym for domain dpk , but without knowledge of the private keys x_1 and x_2 . One important difference between black- and whitelisting is that whitelisting allows to retain security even if authority's secret key z is compromised. If whitelisting is used it is not strictly required to keep z secret. While an attacker would be able to construct valid private keys (x_1, x_2) corresponding to the group public key y , the corresponding pseudonyms would not be listed on the whitelist and thus, the signatures would be rejected. Therefore, the attacker would have to find corresponding private keys for given pseudonyms on the whitelist, which in turn would require to calculate discrete logarithms.

4 Security Analysis

Our proofs work in the random oracle model, and thus, an adversary may also query a random hash function oracle. By q_h , we denote the maximum number of queries to hash function oracle made by the adversary.

4.1 Number-Theoretic Assumptions

Our proof for the anonymity follows by reduction to the Decisional Diffie-Hellman (DDH) assumption whereas the hardness of unforgeability and seclusiveness is reduced to the Discrete Logarithm (DL) problem. Both are considered as standard assumptions, and are believed to be hard for decades. Roughly, the DDH assumption says that distinguishing the tuples (g^a, g^b, g^{ab}) and (g^a, g^b, g') where g' is picked randomly from all group elements, is hard. The DL assumption says it is hard to find a given group \mathcal{G} generated by g and element $A = g^a$. In our security analysis we use concrete bounds for the underlying problems, i.e., we assume that the Discrete-Logarithm problem is (t, ϵ) -hard if any algorithm \mathcal{A} given $A \in \mathcal{G}$ (sampled uniformly in \mathcal{G}), running in time t , finds a with $g^{a'} = A$ with probability at most ϵ . Second, the Decisional Diffie-Hellman problem is (t, ϵ) -hard if any algorithm \mathcal{A} can differentiate tuples (g^a, g^b, g^c) and (g^a, g^b, g^{ab}) with probability at most $\frac{1}{2} + \epsilon$ in runtime t where $a, b, c \in \mathbb{Z}_\varphi(\mathcal{G})$ are sampled uniformly.

4.2 Anonymity

Informally, the protocol is cross-domain anonymous (with respect to static corruptions) because the domain-specific pseudonyms appear to be random to the

adversary under the decisional Diffie-Hellman assumption. Below we write $t' \approx t$ to denote the fact that t' is essentially the same running time as t , except for minor administrative overhead.

Theorem 1. *Assume the DDH problem is (t, ϵ) -hard. Then the domain-specific pseudonymous signature scheme \mathcal{NYS} of Section 3 is (n, d, t', Q, ϵ') cross-domain anonymous with $Q = (q_c, q_s, q_t, q_h)$, where $\epsilon' \leq nd\epsilon + \frac{(q_s + q_t + q_h)(q_s + q_t)}{q^2}$ and $t' \approx t$. This holds in the random oracle model.*

Proof. We combine the ideas of pseudorandom synthesizers of Naor and Reingold [16] with the simulation soundness of the non-interactive zero-knowledge proofs (aka. signatures) in the random oracle model. That is, assume the original attack of the adversary on our protocol. In a first game hop we replace the actual signature computations in LoR and NymSig' queries by simulated signatures (via programming the random oracle). See [19] for details. This strategy is valid if programming the hash values in such computations have not appeared in previous hash queries of the adversary, nor in previous signature queries. However, since the random group elements a_1, a_2 enter the hash evaluations, the probability that an input collision occurs in any of the at most $q_s + q_t$ signature generations, is at most $(q_s + q_t + q_h)(q_s + q_t)/q^2$. Given that no such collision occurs the simulation is perfectly indistinguishable such that the adversary's success probability cannot increase by more than this term.

Note that after the game hop, we can create valid signatures on behalf of users without knowing the secret keys. In the next game hop, we replace the domain-specific pseudonyms dsnym in LoR queries by random group elements (but in a consistent way). That is, whenever we are supposed to use dsnym we instead use a new random element $\text{dsnym}' \leftarrow \mathcal{G}$, unless nym in combination with dpk has been used before, either in a signature request or an LoR query, in which case we use again the previously generated random value dsnym' .

We claim that this hop cannot increase the adversary's success probability noticeable by the DDH assumption. To this end, we briefly recall the notion of a pseudorandom synthesizer in [16]. The pseudorandom synthesizers for the DH pairs is an $a \times b$ matrix, with the rows labeled by values g^{x_i} and the columns labeled by g^{r_j} , and entries at position i, j set to $g^{x_i r_j}$, such that this matrix is indistinguishable from an $a \times b$ matrix of independent and random group elements, even if the row and column labels g^{x_i} and g^{r_j} are given. In a sense, the matrix entries are correlated but still look random. As discussed in [16] this holds in our case under the DDH assumption. In fact, it allows for a reduction to the DDH problem with a loss of a factor ab where, in our case, after the initial corruption, there are at most $ab \leq nd$ entries of honest users.

In the next game hop, we always use the left domain-specific pseudonym dsnym_0 in LoR queries, independently of the value of b . We stress that, in case of $b = 1$, this does not change the adversary's success probability at all. Assume from now on $b = 1$. Note that each LoR query about $(\text{dsnym}_0, \text{dsnym}_1, \text{dpk}, m)$ is answered by a random element, just as it would be for $b = 0$. All other LoR queries involving dpk can only be about the same pair $(\text{dsnym}_0, \text{dsnym}_1)$ in this

order, in reverse order $(\text{dsnym}_1, \text{dsnym}_0)$, or for distinct entries. In the first case, we would answer again consistently with the (wrong) random element, in the second case, we would switch to the other random element, and in the third case use an independent random value. This behavior, however, is identical to the case $b = 0$ from the adversary's point of view. Similarly, the adversary cannot make any signature request for $(\text{dsnym}_0, \text{dpk})$ nor $(\text{dsnym}_1, \text{dpk})$ without losing. It follows that such signature requests do not depend on the bit b . Hence, the probability of the experiment returning 1 does not change.

In the final game, the adversary's success probability is independent of b , and the adversary cannot win with probability more than $\frac{1}{2}$. Collecting all probabilities from the game hops yields the claimed bound. \square

Anonymity of Restricted Identification. Recall that in the basic version of the restricted identification protocol, the user merely shows the domain-specific pseudonym dsnym to the service provider (who checks that this value has not been revoked yet). Anonymity of this solution follows from the proof above under the DDH assumption alone, noting that we do not need to simulate the additional signatures in the random oracle model.²

4.3 Unforgeability

Theorem 2. *Assume the DL problem is (t, ϵ) -hard on \mathcal{G} , then the domain-specific pseudonymous signature scheme $\mathcal{N}\mathcal{Y}\mathcal{M}\mathcal{S}$ of Section 3 is (n, d, t', Q, ϵ') -unforgeable with $Q = (q_s, q_h)$, where*

$$\epsilon' \approx (2q)^{-1} (q_h - \sqrt{q_h} \sqrt{\delta\epsilon + (2\delta/q)(q_s + q_h)^2 + q_h})$$

and $t' \approx t$ with $\delta = 4ndq^2$. This holds in the random oracle model.

Due to space limitations, the proof of Theorem 2 appears in the full version of this work; however, we provide a sketch here.

Proof (Sketch). We reduce the hardness of the unforgeability game (cf. Definition 3) to the discrete-logarithm problem. We embed a DL challenge $A = g^a$ into a domain-specific pseudonym dsnym^* of an randomly chosen pseudonym $\text{nym}^* \in \mathbb{N}$ and domain $\text{dpk}^* \in \mathbb{D}$. Similarly to the proof of Schnorr signatures [19], we leverage the Forking Lemma, in order to obtain two related forgeries $(c, s_1, s_2), (c', s'_1, s'_2)$ on a message m under dsnym^* from which we can extract the witness (resp. discrete logarithm) from the challenge as follows. Given both signatures σ, σ' we extract the discrete-log a of $\text{dsnym}^* = A = g^a$ as follows. Given $s_1 = t_1 - c \cdot a$ and $s'_1 = t_1 - c' \cdot a$, we have $a = (s_1 - s'_1)/(c' - c)$. Hence, we found the solution a for the DL instance A .

We require that \mathcal{A} succeeds to forge on behalf of pseudonym nym^* under domain public key dpk^* in the first signature. In addition, we loose a tightness

² The specification actually lets the user send a hash value of dsnym . This does not affect the discussion, though.

factor due to Forking Lemma, which yields the probability to find the discrete logarithm of A at most $\epsilon = \epsilon'/nd \cdot (\epsilon'/q_h + 1/q) - (2q_s + 2q_h)^2/q$ where ϵ' is the success probability of \mathcal{A} .

4.4 Seclusiveness

As remarked before, seclusiveness only holds as long as the adversary does not get a hold of the group manager's secret key. By construction, this means that the adversary can thus only corrupt one user, else z becomes known. When considering blacklisting for our construction, we stipulate this below by requiring that the secrets are stored securely in hardware, or, respectively, that the number of corrupt requests q_c is at most 1. If whitelisting is used instead, then we do not require any bound on the number of corruptions the adversary can make, since learning z does not help the adversary to compute a domain-specific `dsnym` which is listed in the (still trustworthy) whitelist.

Theorem 3. *Assume the DL problem is (t, ϵ) -hard on \mathcal{G} , then the domain-specific pseudonymous signature scheme of Section 3 is (n, d, t', Q, ϵ') -secluding with $Q = (q_c, q_s, q_h)$, where $q_c = 1$, $\epsilon' \approx (2q)^{-1}(q_h - \sqrt{q_h} \sqrt{q_h + 4q^2\epsilon + 4q^2\delta})$ and $t' \approx t$ with $\delta = 2(q_s + q_h)/q$. This holds in the random oracle model.*

Due to space limitations, the proof of Theorem 2 appears in the full version of this work; however, we provide a sketch here.

Proof (Sketch). We are given an adversary \mathcal{A} which wins in the seclusiveness game of \mathcal{NYMS} . Here, \mathcal{A} outputs a signature under a domain-specific pseudonym `dsnym` to no corresponding identity `nym` $\in \{\text{nym}\}_d$. Intuitively, the proof works as follows. We are asked for the discrete-log a of an element $A := g_{dl}^a$ from a presumably DL-hard group \mathcal{G}_{dl} . We embed A in generator g_2 such that the group's master key $z \in \text{gmsk}$ equals a . We are able to generate one secret key pair (x_1, x_2) satisfying $x = x_1 + zx_2$ for unknown x , `gmsk`. Using the signature given by \mathcal{A} we can extract a second pair (x_1^*, x_2^*) . Those two key pairs suffice to disclose z (resp. a) and thus, we solve the DL problem. \square

References

1. Brassard, G., Chaum, D., Crépeau, C.: Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.* 37(2), 156–189 (1988)
2. Bernhard, D., Fuchsbauer, G., Ghadafi, E., Smart, N.P., Warinschi, B.: Anonymous attestation with user-controlled linkability. *Cryptology ePrint Archive, Report 2011/658* (2011), <http://eprint.iacr.org/>
3. Bellare, M., Micciancio, D., Warinschi, B.: Foundations of Group Signatures: Formal Definitions, Simplified Requirements, and A Construction Based on General Assumptions. In: Biham, E. (ed.) *EUROCRYPT 2003*. LNCS, vol. 2656, pp. 614–629. Springer, Heidelberg (2003)
4. Bender, J., Fischlin, M., Kügler, D.: Security Analysis of the PACE Key-Agreement Protocol. In: Samarati, P., Yung, M., Martinelli, F., Ardagna, C.A. (eds.) *ISC 2009*. LNCS, vol. 5735, pp. 33–48. Springer, Heidelberg (2009)

5. Brands, S.: Rethinking Public Key Infrastructures and Digital Certificates; Building in Privacy. The MIT Press (2000)
6. Brzuska, C., Dagdelen, Ö., Fischlin, M.: TLS, PACE, and EAC: A Cryptographic View at Modern Key Exchange Protocols. In: GI-Sicherheit 2012. GI-LNI, pp. 71–82 (2012)
7. Advanced security mechanism for machine readable travel documents extended access control (eac). Technical Report (BSI-TR-03110) Version 2.05 Release Candidate, BSI (2010)
8. Chaum, D.: Security without identification: transaction systems to make big brother obsolete. *Commun. ACM* 28 (October 1985)
9. Camenisch, J., Lysyanskaya, A.: An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 93–118. Springer, Heidelberg (2001)
10. Camenisch, J., Lysyanskaya, A.: Dynamic Accumulators and Application to Efficient Revocation of Anonymous Credentials. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 61–76. Springer, Heidelberg (2002)
11. Camenisch, J., Lysyanskaya, A.: Signature Schemes and Anonymous Credentials from Bilinear Maps. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 56–72. Springer, Heidelberg (2004)
12. Dagdelen, Ö., Fischlin, M.: Security Analysis of the Extended Access Control Protocol for Machine Readable Travel Documents. In: Burmester, M., Tsudik, G., Magliveras, S., Ilić, I. (eds.) ISC 2010. LNCS, vol. 6531, pp. 54–68. Springer, Heidelberg (2011)
13. Fischlin, M.: Anonymous Signatures Made Easy. In: Okamoto, T., Wang, X. (eds.) PKC 2007. LNCS, vol. 4450, pp. 31–42. Springer, Heidelberg (2007)
14. Koops, C.f.B.J., Buitelaar, H., Lips, M. (eds.): D5.4: Anonymity in electronic government: a case-study analysis of governments? identity knowledge. FIDIS report (February 2012)
15. Kutylowski, M., Shao, J.: Signing with multiple ID's and a single key. In: 38th CCNC, pp. 519–520. IEEE (2011)
16. Naor, M., Reingold, O.: Number-theoretic constructions of efficient pseudo-random functions. In: 38th FOCS, pp. 458–467. IEEE (1997)
17. Okamoto, T.: Provably Secure and Practical Identification Schemes and Corresponding Signature Schemes. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 31–53. Springer, Heidelberg (1993)
18. Pedersen, T.P.: Non-interactive and Information-Theoretic Secure Verifiable Secret Sharing. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 129–140. Springer, Heidelberg (1992)
19. Pointcheval, D., Stern, J.: Security arguments for digital signatures and blind signatures. *Journal of Cryptology* 13(3), 361–396 (2000)
20. Rivest, R.L., Shamir, A., Tauman, Y.: How to Leak a Secret. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 552–565. Springer, Heidelberg (2001)
21. Schnorr, C.-P.: Efficient signature generation by smart cards. *Journal of Cryptology* 4(3), 161–174 (1991)
22. Wei, V.K.: Tracing-by-Linking Group Signatures. Cryptology ePrint Archive, Report 2004/370 (2004), <http://eprint.iacr.org/>
23. Yang, G., Wong, D.S., Deng, X., Wang, H.: Anonymous Signature Schemes. In: Yung, M., Dodis, Y., Kiayias, A., Malkin, T. (eds.) PKC 2006. LNCS, vol. 3958, pp. 347–363. Springer, Heidelberg (2006)