

Computing on Authenticated Data for Adjustable Predicates

Björn Deiseroth Victoria Fehr Marc Fischlin Manuel Maasz
Nils Fabian Reimers Richard Stein

Darmstadt University of Technology, Germany

Abstract. The notion of P -homomorphic signatures, introduced by Ahn et al. (TCC 2012), generalizes various approaches for public computations on authenticated data. For a given predicate P anyone can derive a signature for a message m' from the signatures of a set of messages M , as long as $P(M, m') = 1$. This definition hence comprises notions and constructions for concrete predicates P such as homomorphic signatures and redactable signatures.

In our work we address the question of how to combine P_i -homomorphic schemes for different predicates P_1, P_2, \dots to create a richer and more flexible class of supported predicates. One approach is to statically combine schemes for predicates into new schemes for logical formulas over the predicates, such as a scheme for AND ($P_1 \wedge P_2$). The other approach for more flexibility is to derive schemes which allow the signer to dynamically decide which predicate to use when signing a message, instead of supporting only a single, fixed predicate.

We present two main results. One is to show that one can indeed devise solutions for the static combination for AND, and for dynamically adjustable solutions for choosing the predicate on the fly. Moreover, our constructions are practical and add only a negligible overhead. The other main result is an impossibility result for static combinations. Namely, we prove that, in contrast to the case of AND, many other formulas like the logical OR ($P_1 \vee P_2$) and the NOT ($\neg P$) do not admit generic combinations through so-called canonical constructions. This implies that one cannot rely on general constructions in these cases, but must use other methods instead, like finding new predicate-specific solutions from scratch.

1 Introduction

The notion of P -homomorphic signatures has been put forward by Ahn et al. [ABC⁺12] as a generalization of several concurrent approaches to compute on authenticated data. The predicate P takes as input a set of messages M and determines the admissible messages m' which can be derived from M , and for which a signature can be publicly computed from the signatures for the messages in M . Examples covered by such signatures include homomorphic signatures [Des93, JMSW02, CJL09, BFKW09, GKRR10, JWL11, AL11, BF11b, BF11a, Fre12, CFW12] where m' is the sum of all messages in M , transitive signatures [MR02, BN02, SSM04, Yi07, WCZ⁺07, CH12] where m' describes a path in a graph given by M , and redactable signatures [JMSW02, SBZ01, MSI⁺03, ACdMT05, HHH⁺08, CLX09, NTKK09, BBD⁺10] where m' is a substring of the single message M .

Ahn et al. [ABC⁺12] proposed two general security notions for P -homomorphic signatures. The first one is unforgeability and says that one should not be able to forge signatures for fresh messages which have not been signed before, and which are not publicly derivable. The other

notion is called context hiding and provides strong privacy. It says that a derived signature for an admissible message m' and freshly created signatures for m' have statistically close distributions. This guarantees for instance that the original message in case of redactable signatures remains hidden. The context hiding notion has been subsequently refined in [ALP12].

P-HOMOMORPHIC SIGNATURES WITH ADJUSTABLE PREDICATES. While the abstract notion of P-homomorphic signatures is very handy for arguing about the security of solutions, any construction so far, even the ones in [ABC⁺12, ALP12], are for a specific fixed predicate P , such as quoting substrings of a message. What is currently unknown is how to adjust solutions for fixed predicates in the following sense:

- One desirable option may be the possibility to combine a set of given homomorphic schemes for predicates P_1, P_2, \dots into one for a new P-homomorphic signature scheme. Here, P may be a simple combination such as $P_1 \wedge P_2$ or $P_1 \vee P_2$, or describe even more complex functions. An example are two redactable schemes, one allowing for redaction only at the front of the message (P_1), and the other one enabling redaction only at the end (P_2). Then a $P_1 \vee P_2$ -homomorphic scheme would be a scheme for quoting substrings, by first pruning at the front and then truncating in another step at the end. Note that the problem here is to present a general transformation which supports a rich set of combinations from, say, basic predicates P_1, P_2, \dots , instead of having to build schemes for P from scratch.
- Another desirable feature, which is not offered by the previous ability to combine predicates, is that signer can decide “on the fly” for each signature which predicate P the signature should support. Here, the set of admissible predicates is only bound by the universe \mathcal{P} of predicates for which such signature schemes have been devised yet. This would allow to make the set of admissible message derivatives depend on the message itself, e.g., supporting selective redaction for different messages.

We call general constructions with the first property *statically adjustable* because the combined predicate P is fixed at the time of key generation. The latter schemes are called *dynamically adjustable*. Both approaches have their merits and display their full power only in combination. One can first derive (statically) adjustable schemes for a larger universe \mathcal{P} , and then use this universe for the dynamically adjustable scheme.

CONSTRUCTING SCHEMES WITH STATICALLY ADJUSTABLE PREDICATES. We first investigate simple static combinations such as $P_1 \wedge P_2$, $P_1 \vee P_2$, and $\neg P$. Having solutions for these cases would immediately allow arbitrarily complex combinations of predicates. Our first result is to confirm for the logical AND that the “componentwise” solution works: sign each message with the schemes for predicates P_1, P_2 individually, and derive signatures by applying the corresponding algorithms for each component.

Our main result is to show that the logical OR, $P_1 \vee P_2$, in general does not admit *canonical* constructions. Such canonical constructions can combine given signatures of the individual schemes into one for the $P_1 \vee P_2$ predicate, and can vice versa split any signature for the OR into parts for the individual schemes. Our AND construction is of this type. Our negative result for the OR holds for (almost) arbitrary predicates P_1, P_2 , essentially only excluding trivial examples like $P_1 \vee P_1$. Note that we cannot hope to show a similar result for *non-canonical* solutions, as for some cases we know constructions from scratch for $P_1 \vee P_2$ (e.g., for quotable substrings).

We actually present a more general result, saying that one cannot find canonical constructions for any predicate combination $f(P_1, P_2, \dots)$ if one is able to efficiently find a derivable message

m' under $f(P_1, P_2, \dots)$ and from a message set M , such that m' is not derivable under one of the predicates individually. This excludes the AND case, because any derivable message m' in $P_1 \wedge P_2$ must be also valid according to both in P_1 and P_2 . Yet, this notion includes the OR case if m' can be derived under one predicate, and therefore the OR, but not under the other predicate. It also covers the NOT case straightforwardly, because if m' is derivable under $f(P_1) = \neg P_1$, then it is clearly not derivable under P_1 . The impossibility result holds even if the canonical construction depends on f and the predicates. Put differently, it seems that the only general and non-trivial solutions for statically adjustable predicates are the ones for logical ANDs.

CONSTRUCTING SCHEMES WITH DYNAMICALLY ADJUSTABLE PREDICATES. Does the negative result for statically adjustable parameters also rule out solutions for the dynamic case? Not necessarily, because in this case we assume that the signer adaptively chooses the predicate P from the universe \mathcal{P} for which constructions are already known. Indeed we show that the “certify-then-sign” construction provides a solution in this case: use a regular signature scheme to certify a public key for the P -homomorphic scheme for the chosen predicate $P \in \mathcal{P}$ and sign the message under the secret key for P . Some care must be taken, though, because in order to preserve context hiding the key pair for the P -homomorphic scheme must remain fixed throughout the life time.

2 Preliminaries

We recall the definition and security notions of P -homomorphic signatures, as given in [ABC⁺12, ALP12], and adopt them slightly for our adjustable setting.

2.1 Adjustable \mathcal{P} -homomorphic Signature Schemes

We assume a fixed but public universe \mathcal{P} of predicates P_1, P_2, \dots , each predicate associated with a publicly known P_i -homomorphic signature scheme. A predicate $P_i : 2^{\mathcal{M}} \times \mathcal{M} \rightarrow \{0, 1\}$ indicates whether a set of messages M allows to derive another message m' from the message space \mathcal{M} or not. We give the signer and the verifier the predicate P in question as additional input. In case of a single fixed predicate P , as for the statically adjustable setting, where the universe \mathcal{P} is a singleton, this is an invariant for the scheme and could be ignored by both algorithms. In fact, in this case the notion basically coincides with the definition of a P -homomorphic scheme, the only difference being the predicate given to the signers and verifier as additional input. In this sense the definition of schemes with statically adjustable predicates is a rehash of the notion of P -homomorphic signatures. We stress that we do not suggest to change the terminology for P -homomorphic schemes. The reader should bear in mind, however, that schemes with statically adjustable predicates in this paper implicitly assume a *construction* from selected P -homomorphic schemes underneath. In light of this it matches the dynamic counterpart where predicates are chosen adaptively for each signature.

We simplify the notation below, and write $\text{Verify}(pk, M, \Sigma, P)$ as shorthand for $\bigwedge_{m \in M} \text{Verify}(pk, m, \sigma_m, P)$ with $\Sigma = \{\sigma_m\}_{m \in M}$. Similarly, we sometimes write $\Sigma \leftarrow \text{Sign}(sk, M, P)$ for $\Sigma = \{\text{Sign}(sk, m, P) \mid m \in M\}$.

Definition 2.1 (Adjustable \mathcal{P} -homomorphic Signature Scheme) *A (statically or dynamically) adjustable \mathcal{P} -homomorphic signature scheme is a tuple of PPT algorithms $(\text{KeyGen}, \text{Sign}, \text{SignDer}, \text{Verify})$ such that:*

- $(sk, pk) \leftarrow \text{KeyGen}(1^\lambda)$ maps the security parameter $\lambda \in \mathbb{N}$, given in unary, to a key pair.

- $\sigma \leftarrow \text{Sign}(sk, m, P)$ on input the secret key sk , a message $m \in \mathcal{M}$, and a predicate $P \in \mathcal{P}$ returns a signature σ to m and P .
- $\sigma' \leftarrow \text{SignDer}(pk, M, \Sigma, m', P)$ takes as input the public key pk , a set of messages $M \subseteq \mathcal{M}$ along with signatures $\Sigma = \{\sigma_m\}_{m \in M}$, a message $m' \in \mathcal{M}$, and the predicate $P \in \mathcal{P}$ to be applied, and outputs a signature σ' (or a special symbol \perp indicating failure).
- $b \leftarrow \text{Verify}(pk, m, \sigma, P)$, given the public key pk , a signature σ , a message $m \in \mathcal{M}$, and a predicate $P \in \mathcal{P}$, returns 1 if the signature is valid for the given message, and 0 if not.

We assume the usual correctness condition, namely, that for any $\lambda \in \mathbb{N}$, any $(sk, pk) \leftarrow \text{KeyGen}(1^\lambda)$, any $(m, M, m') \in \mathcal{M} \times 2^\mathcal{M} \times \mathcal{M}$ and any $P \in \mathcal{P}$ we have:

- if $\sigma \leftarrow \text{Sign}(sk, m, P)$, then $\text{Verify}(pk, m, \sigma, P) = 1$ with probability 1; and
- for any $\Sigma = \{\sigma_m\}_{m \in M}$, if $\text{Verify}(pk, M, \Sigma, P) = 1$ and $P(M, m') = 1$, then for any $\sigma' \leftarrow \text{SignDer}(pk, M, \Sigma, m', P)$ we have $\text{Verify}(pk, m', \sigma', P) = 1$ with probability 1.

2.2 Unforgeability

For any predicate P and set M of messages it is convenient to consider the set of messages which can be derived (recursively) from M through P . Hence, similar to [ABC⁺12], we define $P(M) = \{m' \in \mathcal{M} \mid P(M, m') = 1\}$ for any $M \subseteq \mathcal{M}$, as well as $P^0(M) = M$ and $P^i(M) = P(P^{i-1}(M))$ for $i > 0$. Let $P^*(M) = \bigcup_{i \in \mathbb{N}_0} P^i(M)$. We sometimes switch between the set $P^*(M)$ and its predicate analogue, with $P^*(M, m') = 1$ iff $m' \in P^*(M)$. Unless mentioned differently, we assume that any predicate can be evaluated efficiently.

We also presume, without further mentioning it, that predicates are *monotone*, that is, $P(M') \subseteq P(M)$ if $M' \subseteq M$. It follows inductively that $P^*(M') \subseteq P^*(M)$ in this case as well. This is necessary to ensure that, below in the unforgeability game, the set of messages for which a signature can be trivially derived from known signatures for M , does not shrink by asking for more signatures.¹ An alternative is to consider below all subsets $M' \subseteq M$ and declare that any message which is in $P(M')$ to be a message for which a signature is trivial to derive from the signatures for messages in M' .

We again consider both the static and the dynamic case simultaneously, with the understanding that the predicate is fixed in the static case via $\mathcal{P} = \{P\}$.

Definition 2.2 (Unforgeability) *A (statically or dynamically) adjustable \mathcal{P} -homomorphic signature scheme $(\text{KeyGen}, \text{Sign}, \text{SignDer}, \text{Verify})$ is called unforgeable, if any PPT adversary \mathcal{A} has a negligible advantage in the following game:*

1. The challenger \mathcal{C} generates the key pair $(sk, pk) \leftarrow \text{KeyGen}(1^\lambda)$ and gives pk to the adversary \mathcal{A} . The challenger initializes two empty sets T and Q .
2. \mathcal{A} interleaves adaptively the following queries:
 - *Signing queries:* \mathcal{A} chooses a message $m \in \mathcal{M}$ and a predicate $P \in \mathcal{P}$, upon which \mathcal{C} returns a unique handle h to \mathcal{A} , runs $\sigma \leftarrow \text{Sign}(sk, m, P)$, and stores (h, m, σ, P) in T .

¹Interestingly, this is not stipulated explicitly in previous works [ABC⁺12, ALP12]. Still, the predicates for the constructions there satisfy this property. It is, nonetheless, generally required for a reasonable definition in order to avoid trivial examples of schemes which are formally unforgeable, but intuitively insecure.

- *Derivation queries:* \mathcal{A} chooses a set of handles $\mathbf{h} = \{h_i\}_i$, a message $m' \in \mathcal{M}$ and a predicate P . The challenger \mathcal{C} retrieves the tuples $(h_i, m_i, \sigma_i, P_i)$ from T and returns \perp if one of these tuples does not exist, $P_i \neq P$ for some i , or $P(M, m') = 0$. Otherwise, the challenger returns a unique handle h' to \mathcal{A} , runs $\sigma' \leftarrow \text{SignDer}(pk, M, \{\sigma_m\}_{m \in M}, m', P)$ for $M = \{m_i\}_i$ and stores (h', m', σ', P) in T .
- *Reveal queries:* If \mathcal{A} chooses a handle h then \mathcal{C} returns \perp if there does not exist a tuple of the form (h, m, σ, P) in T . Otherwise, it returns σ to \mathcal{A} and adds (m, σ, P) to the set Q .

3. \mathcal{A} outputs a pair (m, σ, P) and wins if the following conditions hold:

- $\text{Verify}(pk, m, \sigma, P) = 1$, and
- $m \notin P^*(M_P)$, where $M_P = \{m \in \mathcal{M} \mid (m, *, P) \in Q\}$, the set of messages in the query set Q for the same predicate P .

Note that the condition on $m \notin P^*(M_P)$ can be relaxed by considering the set M of messages which have been signed under *some* predicate (and not only those which have been signed under the same predicate P as in the forgery attempt). In the static case both cases coincide, of course.

2.3 Context Hiding

The original definition of Ahn et al. [ABC⁺12] requires a strong privacy requirement, basically saying that a derived signature (from previously signed messages M), and a fresh signature for the new message m' are statistically close. It follows that a derived signature does not leak any information about the starting messages M , and thus implies other common privacy notions for, say, redactable signature schemes [BBD⁺10]. Still, the notion has been strengthened in [ALP12] to adaptive context hiding and complete context hiding, basically saying that derived signatures (for messages with any valid signatures) and fresh signatures are close. The generalization to valid signatures as input, instead of only signed messages, allows to cover previously excluded cases like rerandomizable signatures.

While the notion of adaptive context hiding is game-based, the notion of complete context hiding is defined through statistically close distributions of signatures. It is convenient for us here to present the latter definition also through a game, but considering *unbounded* adversaries (as opposed to *efficient* adversaries for adaptive context hiding). Otherwise the notions are identical. Our game-based definition of complete context hiding can be seen easily to be equivalent to the distributional approach in [ALP12].

Definition 2.3 ((Complete and Adaptive) Context Hiding) *A (statically or dynamically) adjustable \mathcal{P} -homomorphic signature scheme $(\text{KeyGen}, \text{Sign}, \text{SignDer}, \text{Verify})$ is called completely (resp. adaptively) context hiding, if any unbounded (resp. PPT) adversary \mathcal{A} has a negligible advantage in the following game:*

1. The challenger \mathcal{C} generates the key pair $(sk, pk) \leftarrow \text{KeyGen}(1^\lambda)$ and gives (sk, pk) to the adversary \mathcal{A} .
2. The adversary selects a set M of messages, and set $\{\sigma_m\}_{m \in M}$ of signatures, a predicate $P \in \mathcal{P}$, and a message m' and hands it to the challenger. If $P(M, m') = 0$ or if $\text{Verify}(pk, M, \{\sigma_m\}_{m \in M}, P) = 0$ then the challenger immediately returns \perp . Else it picks a random bit $b \leftarrow \{0, 1\}$ and computes a derived signature $\sigma' \leftarrow \text{SignDer}(pk, M, \{\sigma_m\}_{m \in M}, m', P)$ if $b = 0$, and a fresh signature $\sigma' \leftarrow \text{Sign}(sk, m', P)$ in case $b = 1$. It returns σ' to the adversary.

3. Eventually the adversary outputs a bit $b^* \in \{0, 1\}$ and wins if $b^* = b$. The advantage of \mathcal{A} is defined to be $\mathbf{Adv}(\mathcal{A}) = |\text{Prob}[b^* = b] - \frac{1}{2}|$.

Some remarks are in place. First note that the adversary can ask the challenger only once. A standard hybrid argument shows that this remains true for multiple (polynomially many) queries for which the challenger re-uses the same bit b . For both cases, the static and the dynamic one, the advantage grows by a factor proportional to the number of queries.

Secondly, note that in the dynamically adjustable case we do not aim to hide the predicate P which has been used to compute the signature. In a stronger requirement one could demand that the actual predicate remains hidden, either among all predicates from the universe, or among the predicates for which the public derivation algorithm would succeed. The former would require a super-polynomial set \mathcal{P} (else the privacy attacker could probe the derivation algorithms for all predicates). The latter would mean a trade-off between privacy, usability, and the signer's intention for restricting the class of admissible public operations: if the signature would hide the corresponding predicate among multiple possibilities, then signatures for a different predicate than the original choice may be derivable. This would imply that the signer loses some control about the (in)ability to derive further signatures. Hence, we do not pursue such stronger requirements here.

3 Statically Adjustable Computations

In this section we investigate statically adjustable constructions for the basic operations AND, OR, and NOT. As explained in the introduction, we can give a general solution for AND, but cannot hope to give (general) transformations for the other two cases.

Below we consider combinations for arbitrary functions f over a fixed² number q of predicates P_1, P_2, \dots, P_q . We assume that such a function $f(P_1, P_2, \dots, P_q)$ over the predicates itself constitutes a predicate and defines a set of derivable messages from M in a straightforward way, by evaluating the predicates for (M, m') and plugging the results into the formula. If viewed as sets, our basic examples for OR, AND, and NOT can then be written as $f_{\vee}(P_1, P_2)(M) = P_1(M) \cup P_2(M)$, and $f_{\wedge}(P_1, P_2)(M) = P_1(M) \cap P_2(M)$, as well as $f_{-}(P_1)(M) = \mathcal{M} \setminus P_1(M)$.

Note that one could more generally also define $f(P_1, P_2, \dots, P_q)$ for divisible message sets $M = (M_1, M_2, \dots, M_q)$ by evaluating $f(M, m')$ as a logical formula over $P_1(M_1, m'), \dots, P_q(M_q, m')$, i.e., assigning only the i -th part M_i of M to the i -th predicate, instead of using the same set M for all predicates. This can be captured in our notion with a single M by having the predicates P_i first project M onto M_i and then evaluating the actual predicate on (M_i, m') . For sake of readability we use the simpler notion with identical M .

We also assume that the message spaces \mathcal{M}_i of all schemes are identical. This can always be achieved by setting $\mathcal{M} = \bigcap_{i=1}^q \mathcal{M}_i$. Note that, if message spaces are not identical this in principle allows to distinguish, say, in case of OR which predicate can be used to create a signature for some message. Since this would violate the idea of privacy immediately, we restrict ourselves to the case of identical message spaces.

²Note that, in general, the number of combined predicates is specific for the scheme and must not depend on the security parameter, i.e., the design of the scheme does not change with the security parameter. In this sense the number q of predicates is constant in the security parameter.

3.1 Statically Adjustable Computations for AND

We first confirm that the solution to sign each message component-wise under a set of public keys yields a secure solution for the AND. Instead of considering only two predicates we allow to combine any fixed number q of predicates.

Construction 3.1 (AND-Construction) *Let $(\text{KeyGen}_i, \text{Sign}_i, \text{SignDer}_i, \text{Verify}_i)$ be P_i -homomorphic signature schemes for predicates P_1, \dots, P_q . Then the following scheme $(\text{KeyGen}, \text{Sign}, \text{SignDer}, \text{Verify})$ is a P -homomorphic signature scheme for $P = P_1 \wedge \dots \wedge P_q$:*

- $\text{KeyGen}(1^\lambda)$ runs $(sk_i, pk_i) \leftarrow \text{KeyGen}_i(1^\lambda)$ for all $i = 1, 2, \dots, q$, and outputs $sk = (sk_1, \dots, sk_q)$ and $pk = (pk_1, \dots, pk_q)$.
- $\text{Sign}(sk, m, P)$ computes $\sigma_i \leftarrow \text{Sign}_i(sk_i, m, P_i)$ for all i and returns $\sigma = (\sigma_1, \dots, \sigma_q)$.
- $\text{SignDer}(pk, M, \Sigma, m', P)$ first checks that $P_i(M, m') = 1$ for all i , and then creates $\sigma'_i \leftarrow \text{SignDer}_i(pk_i, M, \Sigma_i, m', P_i)$ where Σ_i is the set of projections on the i -th component for each signature tuple in $\Sigma = \{\sigma_m\}_{m \in M}$. It returns $\sigma' = (\sigma'_1, \dots, \sigma'_q)$.
- $\text{Verify}(pk, M, \Sigma, P)$ returns 1 if and only if $\text{Verify}_i(pk_i, M, \Sigma_i, P_i) = 1$ for all i (where again Σ_i is the set of projections on the i -th component for each signature in Σ).

Correctness follows easily from the correctness of the underlying P_i -homomorphic schemes.

Proposition 3.2 *For any constant q and any unforgeable and completely (resp. adaptively) context-hiding P_i -homomorphic schemes, Construction 3.1 (AND-Construction) is unforgeable and completely (resp. adaptively) context-hiding.*

For concrete parameters our proof shows that the advantage of breaking unforgeability resp. context hiding for the AND scheme is bounded by the sum of the advantages for the corresponding property over all P_i -homomorphic schemes.

Proof. We first show unforgeability, then context hiding.

UNFORGEABILITY. Assume that there exists a successful adversary \mathcal{A} against unforgeability (Definition 2.2) for the P -homomorphic signature scheme where $P = P_1 \wedge \dots \wedge P_q$. For each $i \in \{1, 2, \dots, q\}$, we first construct an adversary \mathcal{A}_i against the unforgeability of the underlying P_i -homomorphic signature schemes:

- \mathcal{A}_i initially receives pk_i from the challenger \mathcal{C}_i for the game against the P_i -homomorphic signature schemes.
- \mathcal{A}_i creates an initially empty table T' and runs $(sk_j, pk_j) \leftarrow \text{KeyGen}_j(1^\lambda)$ for all $j = 1, 2, \dots, q$, $j \neq i$ to create the other keys.
- \mathcal{A}_i invokes adversary \mathcal{A} against the AND-scheme on $pk = (pk_1, \dots, pk_q)$.
- For every signing query (m, P) from \mathcal{A} , adversary \mathcal{A}_i creates a signing query for message m and the predicate P_i for its challenger and gets the handle h , then computes $\sigma_j \leftarrow \text{Sign}_j(sk_j, m, P_j)$ for all $j \neq i$, and stores (j, h, m, σ_j, P_j) in T' .

- For every derivation query $(\{h\}, m', P)$ of \mathcal{A} , adversary \mathcal{A}_i passes a derivation query for the corresponding handles $(\{h\}, m', P_i)$ to its challenger to receive a handle h' . If $h' \neq \perp$ adversary \mathcal{A}_i looks up all entries (j, h, m, σ_m, P_j) for $j \neq i$ in T' for the queried handles in $\{h\}$ to form $M = \{m\}$, internally checks $P_j(M, m') = 1$, and computes $\sigma'_j \leftarrow \text{SignDer}_j(pk_j, M, \{\sigma_m\}_{m \in M}, m', P_j)$. If no error occurs it returns h' to \mathcal{A} and stores $(j, h', m', \sigma'_j, P_j)$ in T' for all $j \neq i$; else it returns \perp .
- For every reveal request \mathcal{A}_i runs a reveal request for the corresponding handle h , combines the reply σ_i with the values σ_j from entries (j, h, m, σ_j, P_j) in T' to σ and sends it to \mathcal{A} ; in case of an error it simply returns \perp .
- When \mathcal{A} eventually outputs a tuple (m, σ, P) , then \mathcal{A}_i outputs the tuple (m, σ_i, P_i) for the i -th component σ_i in σ .

Note that for each i adversary \mathcal{A}_i perfectly simulates an attack of \mathcal{A} on the P -homomorphic scheme with the help of its challenger, such that \mathcal{A} would output a successful forgery with the same probability in the simulation as in the original attack. By construction, we also have that the message set M_P of queries $(m, *, P)$ in \mathcal{A} 's queries in the simulation is identical to the set M_{P_i} for queries $(m, *, P_i)$ of \mathcal{A}_i to its challenger for each i . Hence, from $m \notin P^*(M_P)$ it follows that $m \notin P_i^*(M_{P_i})$ for some $i \in \{1, 2, \dots, q\}$. Furthermore, since verification succeeds for all components, it also holds that $\text{Verify}_i(pk_i, m, \sigma, P_i) = 1$ for this i .

In other words, any successful forgery yields a successful forgery against (at least) one of the underlying schemes. It follows that the probability of breaking unforgeability for the AND scheme is bounded from above by the sum of the probabilities to break each underlying scheme.

CONTEXT HIDING. Assume next that there exists a successful adversary \mathcal{A} against context hiding (Definition 2.3) for our P -homomorphic signature scheme with $P = P_1 \wedge \dots \wedge P_q$. As in the case of unforgeability we construct, for each $i \in \{1, 2, \dots, q\}$, an adversary \mathcal{A}_i against context hiding of the i -th scheme. The advantage of \mathcal{A} will be bounded from above by the sum over all advantages of the \mathcal{A}_i 's via a standard hybrid argument. Furthermore, each \mathcal{A}_i will be efficient if \mathcal{A} is, such that the claim remains true for adaptive context hiding.

Adversary \mathcal{A}_i receives a pair (sk_i, pk_i) from its challenger and creates the other key pairs (sk_j, pk_j) for $j \neq i$ by running $\text{KeyGen}_j(1^\lambda)$. It hands $sk = (sk_1, \dots, sk_q)$ and $pk = (pk_1, \dots, pk_q)$ to adversary \mathcal{A} and waits for the adversary to create a challenge request M, Σ, m' . For each signature σ_m in Σ adversary \mathcal{A}_i extracts the i -th component and thereby forms the set Σ_i . It passes M, m' , and Σ_i to its own challenger to receive a signature σ'_i (or an error message). It creates the signatures σ'_j for $j < i$ by running the signing algorithm on m' ; for $j > i$ it runs the signature derivation algorithm on M, m', Σ_j to create the remaining signatures σ'_j . In all cases it checks the validity of the predicates and signatures. If there is an error it returns \perp to the adversary \mathcal{A} , and $(\sigma'_1, \dots, \sigma'_q)$ otherwise. If \mathcal{A} eventually outputs a bit b^* then \mathcal{A}_i , too, outputs this bit and stops.

For the analysis note that \mathcal{A}_1 , given that its challenger uses $b = 0$, describes the case that all signatures are derived via SignDer . It follows that the probability of \mathcal{A} correctly outputting 0 for derived signatures in the attack (and thus in the perfect simulation through \mathcal{A}_1) is exactly the probability that \mathcal{A}_1 returns 0, given $b = 0$ in its challenge. Analogously, given $b = 1$ adversary \mathcal{A}_q only creates fresh signatures via Sign in all components, hence given $b = 1$ the probability that \mathcal{A}_q returns 0 is exactly the same that \mathcal{A} outputs 0 in the case that all signatures are fresh. A standard

hybrid argument now yields:

$$\begin{aligned}
\text{Prob}[\mathcal{A} = b] - \frac{1}{2} &= \frac{1}{2} \cdot (\text{Prob}[\mathcal{A} = 0 \mid b = 0] - \text{Prob}[\mathcal{A} = 0 \mid b = 1]) \\
&= \frac{1}{2} \cdot (\text{Prob}[\mathcal{A}_1 = 0 \mid b = 0] - \text{Prob}[\mathcal{A}_q = 0 \mid b = 1]) \\
&= \frac{1}{2} \cdot (\text{Prob}[\mathcal{A}_1 = 0 \mid b = 0] - \text{Prob}[\mathcal{A}_q = 0 \mid b = 1]) \\
&\quad + \sum_{i=1}^{q-1} (\text{Prob}[\mathcal{A}_i = 0 \mid b = 1] - \text{Prob}[\mathcal{A}_i = 0 \mid b = 1])
\end{aligned}$$

and observing that $\text{Prob}[\mathcal{A}_i = a \mid b = 1] = \text{Prob}[\mathcal{A}_{i+1} = a \mid b = 0]$, because in both cases exactly the first i signatures are computed via **Sign**,

$$\begin{aligned}
&= \frac{1}{2} \cdot (\text{Prob}[\mathcal{A}_1 = 0 \mid b = 0] - \text{Prob}[\mathcal{A}_q = 0 \mid b = 1]) \\
&\quad + \sum_{i=1}^{q-1} (\text{Prob}[\mathcal{A}_{i+1} = 0 \mid b = 0] - \text{Prob}[\mathcal{A}_i = 0 \mid b = 1]) \\
&= \sum_{i=1}^q \frac{1}{2} \cdot (\text{Prob}[\mathcal{A}_i = 0 \mid b = 0] - \text{Prob}[\mathcal{A}_i = 0 \mid b = 1]) \\
&= \sum_{i=1}^q \text{Adv}(\mathcal{A}_i).
\end{aligned}$$

This proves context hiding. □

3.2 Statically Adjustable Computations for OR and NOT

Our impossibility result holds for canonical constructions which combine P_i -homomorphic schemes in a general way, ruling out specific constructions which ignore the underlying schemes and builds a new scheme from scratch. We require four algorithms, one for synthesizing public keys of the individual schemes into one for the combined scheme (**PKComb**), one for splitting keys (**PKSplit**), one for combining signatures (**SigComb**), and one to divide signatures for the combined scheme into signatures for the individual schemes (**SigSplit**). The latter is usually necessary to reduce the security to the security of the individual schemes.

For sake of readability we follow the statistical indistinguishability approach also used for (complete) context hiding, and require that the distributions of the algorithms above for combining and splitting keys and signatures have identical distributions as if running the actual algorithms of the combined scheme directly. As our proof below shows our impossibility result can be extended to cover computationally indistinguishable distributions.

Definition 3.3 (Canonical Construction) *Let f be a functional predicate over predicates P_1, \dots, P_q for a fixed number q of predicates. A statically adjustable $f(P_1, \dots, P_q)$ -homomorphic signature scheme $(\text{KeyGen}, \text{Sign}, \text{SignDer}, \text{Verify})$ is a canonical construction out of P_i -homomorphic signature schemes $(\text{KeyGen}_i, \text{Sign}_i, \text{SignDer}_i, \text{Verify}_i)$ if there exist PPT algorithms $(\text{PKComb}, \text{PKSplit}, \text{SigComb}, \text{SigSplit})$ such that:*

Identical distribution of combined keys: *The following random variables are identically distributed:*

- Let $(pk, sk) \leftarrow \text{KeyGen}(1^\lambda)$ and output pk ;

- Let $(pk_i, sk_i) \leftarrow \text{KeyGen}_i(1^\lambda)$ for all i , $pk \leftarrow \text{PKComb}(pk_1, \dots, pk_q)$, output pk ,

Identical distribution of split keys: *The following random variables are identically distributed:*

- Let $(pk, sk) \leftarrow \text{KeyGen}(1^\lambda)$ and output $(pk_1, \dots, pk_q) \leftarrow \text{PKSplit}(pk)$;
- Let $(pk_i, sk_i) \leftarrow \text{KeyGen}_i(1^\lambda)$ for all i , output (pk_1, \dots, pk_q) ,

Identical distribution of combined signatures: *For any PPT algorithm \mathcal{F} the following pairs of random variables are identically distributed:*

- Run $M \leftarrow \mathcal{F}(1^\lambda)$. Compute $(pk, sk) \leftarrow \text{KeyGen}(1^\lambda)$ and output $\Sigma \leftarrow \text{Sign}(sk, M, f(P_1, \dots, P_q))$;
- Run $M \leftarrow \mathcal{F}(1^\lambda)$. For all i , compute $(pk_i, sk_i) \leftarrow \text{KeyGen}_i(1^\lambda)$ along with $\Sigma_i \leftarrow \text{Sign}_i(sk_i, M, P_i)$. Synthesize the public key via $pk \leftarrow \text{PKComb}(pk_1, \dots, pk_q)$ and output $\Sigma \leftarrow \text{SigComb}(pk, pk_1, \dots, pk_q, \Sigma_1, \dots, \Sigma_q, M)$.

Splitting Signatures: *For any PPT algorithm \mathcal{F}' we have that for $(pk_i, sk_i) \leftarrow \text{KeyGen}_i(1^\lambda)$ for all i , $pk \leftarrow \text{PKComb}(pk_1, \dots, pk_q)$, $(M, m') \leftarrow \mathcal{F}'(1^\lambda)$ where $m' \in f(P_1, \dots, P_q)(M)$, $\Sigma_i \leftarrow \text{Sign}_i(sk_i, M, P_i)$, $\Sigma \leftarrow \text{SigComb}(pk_1, \dots, pk_q, \Sigma_1, \dots, \Sigma_q, M)$, $\sigma' \leftarrow \text{SignDer}(pk, M, \Sigma, m', f(P_1, \dots, P_q))$, the probability that $(\sigma'_1, \dots, \sigma'_q) \leftarrow \text{SigSplit}(pk, pk_1, \dots, pk_q, m', \sigma')$ does not contain some valid component and thus $\text{Verify}_i(pk_i, m', \sigma'_i) = 0$ for all i , is negligible.*

In other words, **SigSplit** returns at least one valid signature for one of the underlying predicates with sufficiently high probability. Our AND-construction is canonical in the above sense: **PKComb** and **SigComb** both concatenate their inputs (and **PKSplit** divides the concatenated keys again), and **SigSplit** simply returns the signature itself. Note that the definition allows **PKComb**, **PKSplit**, **SigComb**, and **SigSplit** to depend on the given predicates P_i ; the construction only follows a canonical pattern.

In what follows, we need to exclude trivial examples like $P_1 \vee P_2 = P_1 \vee P_1$. Hence, for the OR we assume below the existence of a message m' which can be derived from a set of messages M under one predicate, but not the other predicate. This clearly prevents $P_1 = P_2$. More generally, and to include for instance also the NOT case, we assume that m' can be derived under $f(P_1, P_2, \dots)$ but not under one of the predicates; the excluded predicate P_i can be arbitrary, but the output distribution of m' does not depend on this choice. The latter is necessary to ensure that m' does not contain any information about the predicate's index i . Furthermore, we assume that such pairs (M, m') are efficiently computable. We discuss an illuminating example after the definition.

Definition 3.4 (Efficiently Distinguishable Predicates) *Let f be a functional predicate over predicates P_1, \dots, P_q . Consider a statically adjustable $f(P_1, \dots, P_q)$ -homomorphic signature scheme $(\text{KeyGen}, \text{Sign}, \text{SignDer}, \text{Verify})$. Then the predicates are called efficiently distinguishable with respect to f , if there exists a PPT algorithm \mathcal{F} such that for any $i \in \{1, 2, \dots, q\}$ and for any $(M, m') \leftarrow \mathcal{F}(1^\lambda, i)$, we have $m' \in (f(P_1, \dots, P_q)(M) \setminus P_i^*(M))$. Moreover, for any $i, j \in \{1, 2, \dots, q\}$ the distribution of m' (over the coin tosses of \mathcal{F}) in the output of $\mathcal{F}(1^\lambda, i)$ resp. $\mathcal{F}(1^\lambda, j)$ is identical.*

Let us demonstrate the property for the introductory example of two redactable signature schemes (with message space $\mathcal{M} = \{0, 1\}^*$), one allowing to drop message bits only at the front (predicate P_1), and the other one only at the end (P_2). Consider the OR predicate $P_1 \vee P_2$ describing a scheme for quotable substrings. Then \mathcal{F} can simply pick $m' = 0^\lambda$ and for $i = 1$ output $M = \{0^\lambda 1\}$,

and for $i = 2$ it returns $M = \{10^\lambda\}$ instead. Clearly, for $i = 1$ one can derive m' from M via P_2 and therefore for the OR, but not via P_1 , because the '1' at the end cannot be redacted through P_1 . The same argument holds vice versa for $i = 2$, and the (trivial) distributions on m' are identical for both $i = 1$ and $i = 2$. Hence, this examples has efficiently distinguishable predicates.

The case of NOT is even simpler. Algorithm \mathcal{F} simply needs to find some M and some m' which lies in $(\neg P(M)) \setminus P^*(M) = \mathcal{M} \setminus P^*(M)$, i.e., if m' is not derivable according to $P^*(M)$. Finally note that constructions based only on AND cannot be distinguishable, since $(P_1(M) \cap P_2(M)) \setminus P_i^*(M) = \emptyset$ for any i .

Theorem 3.5 *Let f be a functional predicate over predicates P_1, \dots, P_q for a fixed number q of predicates. Assume further that the predicates are efficiently distinguishable with respect to f . Then there is no adaptively context-hiding, statically-adjustable $f(P_1, \dots, P_q)$ -homomorphic signature scheme which is a canonical construction out of unforgeable P_i -homomorphic signature schemes.*

The proof idea is as follows. Essentially we show how to forge a signature for one of the underlying schemes. For this we use the distinguishability of the predicates to create a set of messages M and a message m' which is derivable by $f(P_1, \dots, P_q)(M)$ but does not lie in $P_i^*(M)$ for some i . Then we ask for signatures for the messages in M , and derive a signature for m' via the public operation SignDer for the combined scheme and for $f(P_1, \dots, P_q)$. Splitting up the signature into its components via SigSplit we obtain (with sufficiently large probability) a valid signature for m' under the i -th scheme. But since $m' \notin P_i^*(M)$ we thus create a valid forgery, contradicting the security of the underlying scheme. In the course of the proof we use the context hiding property to show that the “skewed” choice of M, m' (with $m' \notin P_i^*(M)$) does not bias the success probability of SigSplit for returning a valid signature component for the i -th scheme significantly.

Proof. Take an arbitrary canonical construction. By the efficient distinguishability of the predicates there exists a PPT algorithm \mathcal{F} which on input i returns (M, m') , such that $m' \in (f(P_1, \dots, P_q) \setminus P_i^*)(M)$. Given this algorithm we next construct, for each $i \in \{1, 2, \dots, q\}$, an adversary \mathcal{A}_i trying to break the unforgeability game of the P_i -homomorphic signature scheme ($\text{KeyGen}_i, \text{Sign}_i, \text{SignDer}_i, \text{Verify}_i$):

- The adversary \mathcal{A}_i is given pk_i of $(sk_i, pk_i) \leftarrow \text{KeyGen}_i(1^\lambda)$. It generates the other keys $(pk_j, sk_j) \leftarrow \text{KeyGen}_j(1^\lambda)$ for $j \neq i$ on its own. It also computes the combined public key $pk \leftarrow \text{PKComb}(pk_1, \dots, pk_q)$.
- \mathcal{A}_i runs \mathcal{F} on $(1^\lambda, i)$ to find a suitable tuple (M, m') .
- It obtains signatures for the set M with the help of signing queries, namely $\Sigma_i \leftarrow \text{Sign}_i(sk_i, M, P_i)$. It computes the other signatures $\Sigma_j \leftarrow \text{Sign}_j(sk_j, M, P_j)$ with the secret keys sk_j , and synthesizes them all via $\Sigma \leftarrow \text{SigComb}(pk, pk_1, \dots, pk_q, \Sigma_1, \dots, \Sigma_q, M)$.
- Adversary \mathcal{A}_i obtains the forged signature by the derivative $\sigma' \leftarrow \text{SignDer}(pk, M, \Sigma, m', f(P_1, \dots, P_q))$ and by running $(\sigma'_1, \dots, \sigma'_q) \leftarrow \text{SigSplit}(pk, pk_1, \dots, pk_q, m', \sigma')$. It returns σ'_i .

Obviously, adversary \mathcal{A}_i runs in polynomial time because all steps can be executed efficiently and the set M is of at most polynomial size and therefore the number of signature queries and creations.

It remains to argue that the adversary outputs a valid forgery with non-negligible probability. To this end we need to show that the splitting algorithm returns with a valid signature for our

P_i -homomorphic scheme with sufficiently large probability. In particular, since M, m' are such that $m' \in (f(P_1, \dots, P_q) \setminus P_i^*)(M)$ it may be that the signatures for M and the input m' for `SigSplit` convey some information about i , and `SigSplit` then “avoids” to output some valid signature for i , but always uses a different index j . In fact, we need to make sure that `SigSplit` for our specific input does not return a valid signature for the i -th scheme with negligible probability, despite succeeding with non-negligible probability in general. We use the context hiding property to show that this cannot happen.

By assumption, the canonical construction is adaptively context hiding. We claim that the probability that `SigSplit` does not return a valid signature for the index i given to \mathcal{F} is negligible. To this end assume that `SigSplit`, instead of receiving the derived signature σ' , is run on a fresh signature σ' generated through `Sig` on m' for a key sk which is generated according to the key generation for the canonical construction, $(sk, pk) \leftarrow \text{KeyGen}(1^\lambda)$. We claim that this cannot increase the probability of `SigSplit` returning a valid signature for the index i given to \mathcal{F} by more than a negligible amount. Otherwise it we can devise a successful (and efficient!) adversary \mathcal{B}_i against the context hiding property, as explained next.

The adversary \mathcal{B}_i would receive a key pair $(sk, pk) \leftarrow \text{KeyGen}(1^\lambda)$ of the canonical construction, and mimic \mathcal{A}_i 's behavior by running \mathcal{F} on i to receive (M, m') . It creates all signatures Σ for the messages in M with the secret key sk , and then calls its challenge oracle about Σ, M, m' to obtain a signature σ' , either derived via `SignDer` or via `Sign`. It then runs `SigSplit` on $(pk, pk_1, \dots, pk_q, m', \sigma')$ to obtain signatures for the individual schemes, where it derives the keys pk_1, \dots, pk_q via the key splitting algorithm `PKSplit`. It returns 0 if this yields a valid signature σ'_i for the index i , which it can decide by running the verification algorithm `Verifyi` on $(pk_i, m', \sigma'_i, P_i)$; else it returns 1.

Note that from `SigSplit`'s perspective the simulations through \mathcal{A}_i and through \mathcal{B}_i (given $b = 0$ and thus a derived signature via `SignDer`) are identical, because the distribution of combined keys and signatures and genuine keys are identical. The same is true for \mathcal{A}_i 's run if giving `SigSplit` a fresh signature for m' instead, and \mathcal{B}_i 's simulation for $b = 1$. Hence, if the probability in \mathcal{A}_i 's simulation, when giving a fresh signature for m' instead, to find a suitable signature for the i -th scheme would increase from negligible to non-negligible, then this would also be true for the two cases in \mathcal{B}_i 's attack, contradicting the context hiding property. Hence, since m' is distributed independently of i according to the efficient distinguishability requirement on \mathcal{F} , the index i is information-theoretically hidden from algorithm `SigSplit` when receiving $pk, pk_1, \dots, pk_q, \sigma', m'$ for a fresh signature σ' . In this case, the output by `SigSplit` does not depend on i , and by context hiding this remains true up to a negligible error for the “skewed” parameters M, m' and the derived signature via `SignDer` as chosen by \mathcal{A}_i (for any i).

Now we can complete the argument of `SigSplit` outputting a valid signature for some i sufficiently often. Recall that, for infinitely many security parameters λ , algorithm `SigSplit` returns some valid signature with a success probability exceeding a polynomial fraction in the sense of the definition of a canonical construction. Hence, among the (constant number of) q predicates there must be a *fixed* index $i \in \{1, 2, \dots, q\}$ such that infinitely often the success probability exceeds a polynomial fraction for this index i . With the argument above, up to a negligible error, this remains true when run by \mathcal{A}_i for the parameters M, m' which depend on i now. Hence, this adversary still successfully forges infinitely often with sufficiently high probability. This, however, would contradict unforgeability of this scheme, concluding the proof. \square

We stress that the impossibility result holds for the computational notion of *adaptive* context hiding (with efficient distinguishers), which even strengthens our result. As mentioned before, a slightly more involved argument allows to extend the result also to algorithms `PKComb`, `PKSplit`, `SigComb` whose output is only computationally indistinguishable from the one

of the original algorithms (instead of being identical). This requires some additional steps to prove that gradually replacing the algorithms does not change the behavior of `SigSplit` in the above proof significantly.

4 Dynamically Adjustable Computations

In the dynamic case we assume a polynomial universe \mathcal{P} of predicates such that there exists a P_i -homomorphic scheme for each $P_i \in \mathcal{P}$. We furthermore assume that given (a description of) P_i one can efficiently recover the corresponding scheme, e.g., if the universe consists only of a fixed number of predicates. Vice versa, we assume that P_i is identifiable from the scheme's public key pk_i . This in particular implies that the public keys for predicates must be unique. For simplicity we assume an ordering on predicates in \mathcal{P} and often identify the predicate P_i and the scheme with its number i according to this order. We simply call sets \mathcal{P} as above *efficient*.

In the construction we need to assume that for a given predicate identifier i there is a fixed yet (pseudo)random key pair $(sk_i, pk_i) \leftarrow \text{KeyGen}_i(1^\lambda)$, generated according to the key generation algorithm for the scheme for predicate P_i . This key pair remains identical for all signature requests for P_i . For a polynomial universe \mathcal{P} this can be in principle implemented by generating the keys (sk_i, pk_i) when creating the scheme's keys (sk, pk) , and storing them in sk . In practice this may indeed be admissible for a small number of predicates, a more applicable approach may be to generate the keys on the fly via a pseudorandom function. Namely, store a key κ of a pseudorandom function in sk , and to create the key pair for predicate P_i , recover the (pseudo)random output $\omega_i = \text{PRF}(\kappa, P_i)$ and re-run $\text{KeyGen}_i(1^\lambda; \omega_i)$ for ω_i to derive the same pair (sk_i, pk_i) as before. For unforgeability it can be formally shown via standard techniques that this solution is (quasi) as secure as generating fresh key pairs and maintaining a table to look up previous keys; for context hiding, however, one requires an additional assumption on the security of the underlying scheme to preserve privacy, as discussed below.

Similarly, the public keys pk_i and their (fixed) certificates $cert_i$ may be published at once, or may be attached to each signature upon creation. Below we adopt the latter solution as it rather complies with our notion of (stateless) \mathcal{P} -homomorphic signatures. Hence, below we assume for simplicity that the efficient universe \mathcal{P} stores all pairs (sk_i, pk_i) with once-created certificates $cert_i$ at the beginning in sk . For certification we use a regular signature scheme which we can subsume as a special case under P -homomorphic schemes, without considering a `SignDer` algorithm nor context hiding. If we define $P(M) = M$ for this scheme, unforgeability for this “homomorphic” scheme corresponds to the common notion of unforgeability for regular schemes.

Construction 4.1 (Certify-Then-Sign Construction) *Let \mathcal{P} be an efficient set of predicates P_1, P_2, \dots, P_q . Let $(\text{KeyGen}_0, \text{Sign}_0, \text{Verify}_0)$ be a regular signature scheme. Define the following dynamically adjustable \mathcal{P} -homomorphic signature scheme $(\text{KeyGen}, \text{Sign}, \text{SignDer}, \text{Verify})$:*

- *$\text{KeyGen}(1^\lambda)$ generates $(sk_0, pk_0) \leftarrow \text{KeyGen}_0(1^\lambda)$, generates key pairs $(sk_i, pk_i) \leftarrow \text{KeyGen}_i(1^\lambda)$ for all predicates P_i , and certificates $cert_i \leftarrow \text{Sign}_0(sk_0, pk_i)$ for all i . It returns $sk = (sk_0, \{(sk_i, pk_i, cert_i)\}_i)$ and $pk = pk_0$.*
- *$\text{Sign}(sk, m, P_i)$ looks up $(sk_i, pk_i, cert_i)$ for P_i in sk and computes $\sigma_i \leftarrow \text{Sign}_i(sk, m)$ and returns $\sigma = (\sigma_i, pk_i, cert_i)$.*
- *$\text{SignDer}(pk, M, \Sigma, m', P')$ checks that all signatures carry the same pk_i and $cert_i$ for predicate P_i , that $P' = P_i$, that $P_i(M, m') = 1$, that $\text{Verify}_i(pk_i, M, \Sigma) = 1$, and, if all checks succeed, computes $\sigma'_i \leftarrow \text{SignDer}_i(pk_i, M, \Sigma, m')$ and returns $\sigma' = (\sigma'_i, pk_i, cert_i)$.*

- $\text{Verify}(pk, m, \sigma, P)$ checks that P corresponds to the public key in $(\sigma_i, pk_i, cert_i)$, that $\text{Verify}_0(pk, pk_i, cert_i) = 1$, and that $\text{Verify}_i(pk_i, m, \sigma_i) = 1$. Only if all checks succeed, it returns 1.

It is straightforward to verify that the above construction is correct in the sense that genuine (fresh and derived) signatures are accepted by Verify . This follows from the correctness properties of the regular scheme and of the P_i -homomorphic ones.

Proposition 4.2 *Assume that the signature scheme $(\text{KeyGen}_0, \text{Sign}_0, \text{Verify}_0)$ and all P_i -homomorphic schemes are unforgeable according to Definition 2.2. Then the Certify-then-Sign Construction 4.1 is also unforgeable for the efficient universe $\mathcal{P} = \{P_1, \dots, P_q\}$ for the fixed number q of predicates.*

In terms of concrete security, the success probability of any adversary against the construction is (for similar running time) bounded from above by the probability of forging certificates, plus q times the maximal advantage against any of the schemes from \mathcal{P} .

Proof. Assume that there exists a successful forger \mathcal{A} . Then this adversary is able to forge with non-negligible probability a signature $\sigma^* = (\sigma, pk, cert)$ for a message m such that, in particular, $\text{Verify}_0(pk_0, pk, cert) = 1$. Note that if the probability that \mathcal{A} succeeds and that pk does not match any of the keys pk_i created by the signer for the predicates P_i , was non-negligible, then this would straightforwardly contradict the unforgeability of the certification scheme. Namely, construct an algorithm \mathcal{A}_0 against the certification scheme which, on input pk_0 , creates the polynomial number of key pairs $(sk_i, pk_i) \leftarrow \text{KeyGen}_i(1^n)$ and asks for signatures $cert_i$ for all pk_i from the signing oracle, and then emulates the attack of \mathcal{A} with the help of the secret keys. If \mathcal{A} eventually outputs $\sigma^* = (\sigma, pk, cert)$, then \mathcal{A}_0 returns $pk, cert$ as the forgery attempt.

If the probability that \mathcal{A} would succeed for a fresh pk with non-negligible probability as defined above, then our efficient algorithm \mathcal{A}_0 , which perfectly simulates the actual attack, would then successfully forge a signature $cert$ for a new “message” pk with non-negligible probability. Since this would contradict the unforgeability of the certification scheme, we can assume that this case happens with negligible probability only. It follows that \mathcal{A} must succeed with non-negligible probability for a key $pk = pk_i$ for some (unique) i , such that $\text{Verify}_i(pk_i, m, \sigma, P_i) = 1$, and the message is not trivially derivable under the corresponding predicate P_i from the signing queries for P_i .

Note that the specific choice pk_i may depend on the adversary’s randomness. However, there must exist at least one predicate P_i (among the q schemes) such that \mathcal{A} succeeds for this key fixed pk_i with non-negligible probability. We can now derive an adversary \mathcal{A}_i successfully forging signatures for this P_i -homomorphic scheme. Adversary \mathcal{A}_i receives from the challenger the public key pk_i and gets access to a Sign_i -oracle. It generates (sk_0, pk_0) and all other key pairs (sk_j, pk_j) and signs all of them, including pk_i . The adversary \mathcal{A}_i then runs \mathcal{A} on pk_0 , supplying all signatures requests for $P_j \neq P_i$ with the help of the secret keys, and using the external signing oracle for P_i . If \mathcal{A} finally returns m and $(\sigma, pk, cert)$ then \mathcal{A}_i returns m and σ_i .

Note that, if \mathcal{A} has a non-negligible success probability for forging under the key pk_i , then \mathcal{A}_i has the same success probability. This follows as the signature verifies under pk_i , and if the message m is not derivable from \mathcal{A} ’s queries for P_i , then this is also true for \mathcal{A}_i . This, however, would contradict the unforgeability assumption about the P_i -homomorphic scheme. \square

Proposition 4.3 *Assume that all P_i -homomorphic schemes are completely (resp. adaptively) context-hiding according to Definition 2.3. Then the Certify-then-Sign Construction 4.1 is also completely (resp. adaptively) context-hiding for an efficient universe $\mathcal{P} = \{P_1, P_2, \dots, P_q\}$ of a fixed number q of predicates.*

Proof. Assume that there exists a successful adversary \mathcal{A} with non-negligible advantage. We will show that at least one predicate will not be completely (resp. adaptively) context-hiding. For any $i \in \{1, \dots, q\}$ construct an adversary \mathcal{A}_i that plays against the predicate P_i . This adversary \mathcal{A}_i receives from the challenger the private key and the public key (sk_i, pk_i) for the i -th scheme. It then generates $(sk_0, pk_0) \leftarrow \text{KeyGen}_0(1^\lambda)$, the key pairs (sk_j, pk_j) for all $j \neq i$, and creates the certificates $cert_i$ for all i according to our construction. The private key $sk = (sk_0, \{(sk_i, pk_i, cert_i)\}_i)$ and the public key $pk = pk_0$ are then passed to adversary \mathcal{A} .³ Adversary \mathcal{A} then selects (M, Σ, m', P_j) . In the case of $i = j$ the adversary \mathcal{A}_i forwards it to the challenger. The challenger then returns either a derived signature, or a fresh signature σ' . The adversary \mathcal{A}_i then returns the final output bit b^* of \mathcal{A} . For the case $i \neq j$, adversary \mathcal{A}_i simply returns a random bit $b^* \leftarrow \{0, 1\}$.

Note that, if \mathcal{A} has a non-negligible advantage for the combined scheme, this means that it also has a non-negligible advantage when conditioning on \mathcal{A} choosing one of the fixed predicates P_i for the challenge (involving P_j). More precisely, there must be some i such that the probability of \mathcal{A} predicting correctly the bit b (in the actual attack and in the perfectly indistinguishable simulation through \mathcal{A}_i) and picking $j = i$, is at least $\frac{1}{q}$ times the general prediction probability. Since the number q of predicates is fixed this probability must be non-negligible as well. Fix such an i from now on.

The advantage of \mathcal{A}_i (against the underlying scheme) in terms of the advantage of \mathcal{A} (against the derived scheme) is then given by

$$\begin{aligned} \mathbf{Adv}(\mathcal{A}_i) &= \text{Prob}[\mathcal{A}_i = b] - \frac{1}{2} \\ &= \text{Prob}[\mathcal{A} = b \wedge i = j] + \text{Prob}[b^* = b \wedge i \neq j] - \frac{1}{2} \\ &\geq \frac{1}{q} \cdot \text{Prob}[\mathcal{A} = b] + \frac{1}{2} \cdot (1 - \frac{1}{q}) - \frac{1}{2} \\ &\geq \frac{1}{q} \cdot \mathbf{Adv}(\mathcal{A}). \end{aligned}$$

Hence, if \mathcal{A} breaks context hiding, i.e., has some non-negligible success probability, at least one of the adversary \mathcal{A}_i also breaks context hiding for P_i (losing a factor $\frac{1}{q}$ in the advantage compared to \mathcal{A}). This predicate would thus not be completely (resp. adaptively) context-hiding, which contradicts our assumption. \square

If we use a pseudorandom function with key κ to create $\omega_i = \text{PRF}(\kappa, P_i)$, then to inherit the context hiding property we need to assume that giving ω_i (in addition to sk_i) to the distinguisher does not violate context hiding of the underlying P_i -homomorphic scheme. If this is the case then we can formally reduce security as in the proof above. For this assume that the pseudorandom function has two keys, κ_0, κ_1 , and its output for P_i is defined by $\text{PRF}(\kappa_0, P_i) \oplus \kappa_1$. Then, receiving sk_i, ω_i, pk_i from the challenger we can forward the secret key $sk = (sk_0, (\kappa_0, \kappa_1))$ for random κ_0 and $\kappa_1 = \omega_i \oplus \text{PRF}(\kappa_0, P_i)$ to the simulated adversary. This choice of (κ_0, κ_1) lets the pseudorandom function map P_i to the desired value ω_i and is identically distributed to a truly random key pair. Hence, if we have an adversary against the certify-then-sign construction for the pseudorandom function, we obtain an adversary against context hiding of one of the underlying schemes as above.

Acknowledgments

We thank the anonymous reviewers for comments. Marc Fischlin was supported by a Heisenberg grant Fi 940/3-1 of the German Research Foundation (DFG).

³If we use a pseudorandom function to create the keys on the fly some care must be taken in this step, as discusses after completion of the proof here.

References

- [ABC⁺12] Jae Hyun Ahn, Dan Boneh, Jan Camenisch, Susan Hohenberger, Abhi Shelat, and Brent Waters. Computing on authenticated data. In Ronald Cramer, editor, *TCC 2012: 9th Theory of Cryptography Conference*, volume 7194 of *Lecture Notes in Computer Science*, pages 1–20, Taormina, Sicily, Italy, March 19–21, 2012. Springer, Berlin, Germany. (Cited on pages 1, 2, 3, 4, and 5.)
- [ACdMT05] Giuseppe Ateniese, Daniel H. Chou, Breno de Medeiros, and Gene Tsudik. Sanitizable signatures. In *ESORICS*, volume 3679 of *Lecture Notes in Computer Science*, pages 159–177. Springer, 2005. (Cited on page 1.)
- [AL11] Nuttapon Attrapadung and Benoît Libert. Homomorphic network coding signatures in the standard model. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011: 14th International Workshop on Theory and Practice in Public Key Cryptography*, volume 6571 of *Lecture Notes in Computer Science*, pages 17–34, Taormina, Italy, March 6–9, 2011. Springer, Berlin, Germany. (Cited on page 1.)
- [ALP12] Nuttapon Attrapadung, Benoit Libert, and Thomas Peters. Computing on authenticated data: New privacy definitions and constructions. In *Advances in Cryptology – ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 367–385. Springer, Berlin, Germany, 2012. (Cited on pages 2, 3, 4, and 5.)
- [BBD⁺10] Christina Brzuska, Heike Busch, Özgür Dagdelen, Marc Fischlin, Martin Franz, Stefan Katzenbeisser, Mark Manulis, Cristina Onete, Andreas Peter, Bertram Poettering, and Dominique Schröder. Redactable signatures for tree-structured data: Definitions and constructions. In Jianying Zhou and Moti Yung, editors, *ACNS 10: 8th International Conference on Applied Cryptography and Network Security*, volume 6123 of *Lecture Notes in Computer Science*, pages 87–104, Beijing, China, June 22–25, 2010. Springer, Berlin, Germany. (Cited on pages 1 and 5.)
- [BF11a] Dan Boneh and David Mandell Freeman. Homomorphic signatures for polynomial functions. In Kenneth G. Paterson, editor, *Advances in Cryptology – EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 149–168, Tallinn, Estonia, May 15–19, 2011. Springer, Berlin, Germany. (Cited on page 1.)
- [BF11b] Dan Boneh and David Mandell Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011: 14th International Workshop on Theory and Practice in Public Key Cryptography*, volume 6571 of *Lecture Notes in Computer Science*, pages 1–16, Taormina, Italy, March 6–9, 2011. Springer, Berlin, Germany. (Cited on page 1.)
- [BFKW09] Dan Boneh, David Freeman, Jonathan Katz, and Brent Waters. Signing a linear subspace: Signature schemes for network coding. In Stanislaw Jarecki and Gene Tsudik, editors, *PKC 2009: 12th International Conference on Theory and Practice of Public Key Cryptography*, volume 5443 of *Lecture Notes in Computer Science*, pages 68–87, Irvine, CA, USA, March 18–20, 2009. Springer, Berlin, Germany. (Cited on page 1.)

- [BN02] Mihir Bellare and Gregory Neven. Transitive signatures based on factoring and RSA. In Yuliang Zheng, editor, *Advances in Cryptology – ASIACRYPT 2002*, volume 2501 of *Lecture Notes in Computer Science*, pages 397–414, Queenstown, New Zealand, December 1–5, 2002. Springer, Berlin, Germany. (Cited on page 1.)
- [CFW12] Dario Catalano, Dario Fiore, and Bogdan Warinschi. Efficient network coding signatures in the standard model. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012: 15th International Workshop on Theory and Practice in Public Key Cryptography*, volume 7293 of *Lecture Notes in Computer Science*, pages 680–696, Darmstadt, Germany, May 21–23, 2012. Springer, Berlin, Germany. (Cited on page 1.)
- [CH12] Philippe Camacho and Alejandro Hevia. Short transitive signatures for directed trees. In Orr Dunkelman, editor, *Topics in Cryptology – CT-RSA 2012*, volume 7178 of *Lecture Notes in Computer Science*, pages 35–50, San Francisco, CA, USA, February 27 – March 2, 2012. Springer, Berlin, Germany. (Cited on page 1.)
- [CJL09] Denis Charles, Kamal Jain, and Kristin Lauter. Signatures for network coding. *Int. J. Inf. Coding Theory*, 1(1):3–14, March 2009. (Cited on page 1.)
- [CLX09] Ee-Chien Chang, Chee Liang Lim, and Jia Xu. Short redactable signatures using random trees. In Marc Fischlin, editor, *Topics in Cryptology – CT-RSA 2009*, volume 5473 of *Lecture Notes in Computer Science*, pages 133–147, San Francisco, CA, USA, April 20–24, 2009. Springer, Berlin, Germany. (Cited on page 1.)
- [Des93] Yvo Desmedt. Computer security by redefining what a computer is. In *Proceedings on the 1992-1993 workshop on New security paradigms*, NSPW ’92-93, pages 160–166. ACM, 1993. (Cited on page 1.)
- [Fre12] David Mandell Freeman. Improved security for linearly homomorphic signatures: A generic framework. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012: 15th International Workshop on Theory and Practice in Public Key Cryptography*, volume 7293 of *Lecture Notes in Computer Science*, pages 697–714, Darmstadt, Germany, May 21–23, 2012. Springer, Berlin, Germany. (Cited on page 1.)
- [GKKR10] Rosario Gennaro, Jonathan Katz, Hugo Krawczyk, and Tal Rabin. Secure network coding over the integers. In Phong Q. Nguyen and David Pointcheval, editors, *PKC 2010: 13th International Conference on Theory and Practice of Public Key Cryptography*, volume 6056 of *Lecture Notes in Computer Science*, pages 142–160, Paris, France, May 26–28, 2010. Springer, Berlin, Germany. (Cited on page 1.)
- [HHH⁺08] Stuart Haber, Yasuo Hatano, Yoshinori Honda, William Horne, Kunihiro Miyazaki, Tomas Sander, Satoru Tezoku, and Danfeng Yao. Efficient signature schemes supporting redaction, pseudonymization, and data deidentification. In Masayuki Abe and Virgil Gligor, editors, *ASIACCS 08: 3rd Conference on Computer and Communications Security*, pages 353–362, Tokyo, Japan, March 18–20, 2008. ACM Press. (Cited on page 1.)
- [JMSW02] Robert Johnson, David Molnar, Dawn Xiaodong Song, and David Wagner. Homomorphic signature schemes. In Bart Preneel, editor, *Topics in Cryptology – CT-RSA 2002*, volume 2271 of *Lecture Notes in Computer Science*, pages 244–262, San Jose, CA, USA, February 18–22, 2002. Springer, Berlin, Germany. (Cited on page 1.)

- [JWL11] Rob Johnson, Leif Walsh, and Michael Lamb. Homomorphic signatures for digital photographs. In George Danezis, editor, *FC 2011: 15th International Conference on Financial Cryptography and Data Security*, volume 7035 of *Lecture Notes in Computer Science*, pages 141–157, Gros Islet, St. Lucia, February 28 – March 4, 2011. Springer, Berlin, Germany. (Cited on page 1.)
- [MR02] Silvio Micali and Ronald L. Rivest. Transitive signature schemes. In Bart Preneel, editor, *Topics in Cryptology – CT-RSA 2002*, volume 2271 of *Lecture Notes in Computer Science*, pages 236–243, San Jose, CA, USA, February 18–22, 2002. Springer, Berlin, Germany. (Cited on page 1.)
- [MSI⁺03] K. Miyazaki, S. Susaki, M. Iwamura, T. Matsumoto, R. Sasaki, and H. Yoshiura. Digital documents sanitizing problem. In *Technical Report ISEC2003-20*. IEICE, 2003. (Cited on page 1.)
- [NTKK09] Ryo Nojima, Jin Tamura, Youki Kadobayashi, and Hiroaki Kikuchi. A storage efficient redactable signature in the standard model. In Pierangela Samarati, Moti Yung, Fabio Martinelli, and Claudio Agostino Ardagna, editors, *ISC 2009: 12th International Conference on Information Security*, volume 5735 of *Lecture Notes in Computer Science*, pages 326–337, Pisa, Italy, September 7–9, 2009. Springer, Berlin, Germany. (Cited on page 1.)
- [SBZ01] Ron Steinfeld, Laurence Bull, and Yuliang Zheng. Content extraction signatures. In *ICISC*, volume 2288 of *Lecture Notes in Computer Science*, pages 285–304. Springer, 2001. (Cited on page 1.)
- [SSM04] Siamak Fayyaz Shahandashti, Mahmoud Salmasizadeh, and Javad Mohajeri. A provably secure short transitive signature scheme from bilinear group pairs. In Carlo Blundo and Stelvio Cimato, editors, *SCN 04: 4th International Conference on Security in Communication Networks*, volume 3352 of *Lecture Notes in Computer Science*, pages 60–76, Amalfi, Italy, September 8–10, 2004. Springer, Berlin, Germany. (Cited on page 1.)
- [WCZ⁺07] Licheng Wang, Zhenfu Cao, Shihui Zheng, Xiaofang Huang, and Yixian Yang. Transitive signatures from braid groups. In K. Srinathan, C. Pandu Rangan, and Moti Yung, editors, *Progress in Cryptology - INDOCRYPT 2007: 8th International Conference in Cryptology in India*, volume 4859 of *Lecture Notes in Computer Science*, pages 183–196, Chennai, India, December 9–13, 2007. Springer, Berlin, Germany. (Cited on page 1.)
- [Yi07] Xun Yi. Directed transitive signature scheme. In Masayuki Abe, editor, *Topics in Cryptology – CT-RSA 2007*, volume 4377 of *Lecture Notes in Computer Science*, pages 129–144, San Francisco, CA, USA, February 5–9, 2007. Springer, Berlin, Germany. (Cited on page 1.)