On Tightly Secure Non-Interactive Key Exchange

Julia Hesse[1](✉), Dennis Hofheinz[2], and Lisa Kohl[2]

[1] Technische Universität Darmstadt, Darmstadt, Germany
julia.hesse@crisp-da.de
[2] Karlsruhe Institute of Technology, Karlsruhe, Germany
{dennis.hofheinz,lisa.kohl}@kit.edu

Abstract. We consider the reduction loss of security reductions for non-interactive key exchange (NIKE) schemes. Currently, no tightly secure NIKE schemes exist, and in fact Bader et al. (EUROCRYPT 2016) provide a lower bound (of $\Omega(n^2)$, where n is the number of parties an adversary interacts with) on the reduction loss for a large class of NIKE schemes.

We offer two results: the first NIKE scheme with a reduction loss of $n/2$ that circumvents the lower bound of Bader et al., but is of course still far from tightly secure. Second, we provide a generalization of Bader et al.'s lower bound to a larger class of NIKE schemes (that also covers our NIKE scheme), with an adapted lower bound of $n/2$ on the reduction loss. Hence, in that sense, the reduction for our NIKE scheme is optimal.

1 Introduction

Tight security reductions. A security reduction relates the security of a cryptographic construction to the difficulty to solve some assumed-to-be-hard problem. In other words, to base the security of a scheme S on the hardness of a problem P, one has to show how to solve P given an adversary that successfully attacks S. As one usually considers asymptotic security, both adversary and problem solver are required to have polynomial running time and non-negligible success probability.

Many security reductions now guess where in S to embed problem P. For example, in case of a signature scheme, the security reduction might guess in which generated signature (an instance of) P is embedded. Asymptotically, this is fine, as an S-attacker can only ask for a polynomial number of signatures.

J. Hesse—Parts of work done while at École Normale Supérieure, Paris, supported by ERC Project CryptoCloud FP7/2007-2013 Grant Agreement no. 339563.

D. Hofheinz—Supported by ERC Project PREP-CRYPTO (724307), and by DFG grants HO 4534/4-1 and HO 4534/2-2.

L. Kohl—Supported by ERC Project PREP-CRYPTO (724307), by DFG grant HO 4534/2-2 and by a DAAD scholarship. Parts of work done while visiting Interdisciplinary Center Herzliya.

© International Association for Cryptologic Research 2018
H. Shacham and A. Boldyreva (Eds.): CRYPTO 2018, LNCS 10992, pp. 65–94, 2018.
https://doi.org/10.1007/978-3-319-96881-0_3

But when instantiating the scheme with concrete parameters, this guessing step leads to the following paradox: Considering a number of, say, 2^{30} signature queries (which is realistic when thinking of servers) and a security parameter $\lambda = 100$, the concrete loss in success probability introduced by the reduction would actually be larger than a factor of $2^{\lambda/4}$. When aiming at concrete security guarantees (derived from the hardness of P), one thus has to account for the number of expected signatures at the time of set-up, when choosing keylengths.

This makes so called *tight* security reductions a desirable goal. A security reduction is regarded as tight, if (with comparable running times) the success probability of the problem solver is close to the success probability of the underlying attacker. More precisely, one usually requires the success probabilities to only differ up to a small constant factor (or, for a broader notion of tightness, up to a factor linear in the security parameter). Tight security reductions allow to choose the security parameter for concrete instantiation independently of the number of expected instantiations (or, say, generated signatures in case of a signature scheme).

Positive and negative results on tight security. Schemes with tight security reductions could already be constructed for a variety of cryptographic applications (such as public-key encryption [2,6,19,20,26,28,36,37], identity-based encryption [3,7,11,23,31], digital signature schemes [1,27,34,36,37], or zero-knowledge proofs [20,28]). For public-key encryption schemes, the price to pay for an (almost) tight reduction has been reduced to essentially only one additional group element in ciphertexts [19,20].

On the other hand, starting with the work of Coron [12], a number of works show that certain types of reductions are inherently non-tight (in the sense that a problem solver derived from a given adversary has a significantly reduced success probability). For instance, [4,12,29,32] prove that any "simple" reduction for a sufficiently "structured" signature scheme must lose a factor of $\Omega(q_{\mathsf{sig}})$, where q_{sig} is the number of adversarial signature queries. (Here, the definitions of "simple" and "structured" vary across these papers.) Similar lower bounds exist also for specific schemes and other primitives [4,16,18,35,39]. Particularly interesting to our case is the work of Bader et al. [4], which proves lower bounds on the reduction loss of signature, encryption, and non-interactive key exchange schemes in the standard model.

Our focus: non-interactive key exchange. In this work, we investigate tight reductions for non-interactive key exchange (NIKE) schemes in the two-party setting[1]. Intuitively, a NIKE scheme enables any two parties P_i and P_j to compute a common shared key K_{ij} using a public-key infrastructure only, but *without any interaction*. (That is, K_{ij} should be an efficiently computable

[1] We focus on the two-party setting assuming a public key infrastructure (PKI) since this setting allows for efficient standard-model constructions. Intuitively, stronger settings (multi-party, identity-based with/without setup) appear to require qualitatively stronger tools to give any construction at all, tightly secure or not. However, since any n-party NIKE can be viewed as a 2-party NIKE by fixing n-2 identities, our lower bound trivially generalizes to multi-user NIKE schemes.

Reference	$\|pk\|$	model	sec. loss	assumption	uses
Diffie–Hellman [14]	$1 \times \mathbb{G}$	HKR	n^2	DDH	-
Ours, Sec. 3	$3 \times \mathbb{G}$	HKR	$n/2$	DDH	-
CKS08 [10]	$2 \times \mathbb{G}$	DKR	2	CDH	ROM
FHKP13 [17]	$1 \times \mathbb{Z}_N$	DKR	n^2	factoring	ROM
FHKP13 [17]	$2 \times \mathbb{G} + 1 \times \mathbb{Z}_p$	DKR	n^2	DBDH	asymm. pairing
Ours, full version [25]	$12 \times \mathbb{G}$	DKR	$n/2$	DLIN	symm. pairing

Fig. 1. Comparison of existing NIKE schemes. $|pk|$ denotes the size of the public keys, measured in numbers of group elements and exponents. "DKR" or "HKR" denote the CKS-heavy security notion from [17] with dishonest, resp. honest key registrations. Regarding security loss, n denotes the number of honest parties the adversary interacts with and q is the total number of queries made by the adversary. The losses of the two constructions from [17] stems from applying a generic transformation (from the same paper) to level the security guarantees of all compared schemes. Our construction from Sect. 3 is instantiated with the HPS of Cramer–Shoup based on DDH. For more details we refer to the full version [25]. We omit the second scheme from [17] since we focus on non-interactive key registration procedures.

function of P_i's public and P_j's private key, and we require $K_{ij} = K_{ji}$.) Already the original Diffie-Hellman key exchange [14] forms a NIKE scheme (although one that only satisfies a weak form of security). However, the formal investigation of NIKE schemes started with the work of Cash et al. [10], with a more detailed investigation provided in [17].

While there exist highly secure and efficient NIKE schemes (e.g., [10,17]), currently there is no NIKE scheme with a tight security reduction to a standard assumption (and in the standard model). We believe that this is no coincidence: as we will detail below, the rich interdependencies among NIKE keys prevent existing techniques to achieve tight security. Also, it might be interesting to note that the already mentioned work of Bader et al. [4] presents a particularly strong (i.e., *quadratic*) lower bound of $\Omega(n^2)$ on the reduction loss of NIKE schemes, where n is the number of parties that the adversary interacts with. While the scheme of [10] is proven only in the random oracle model, this lower bound applies to the scheme of [17].

Our results. In this work, we provide two contributions. First, we construct an efficient and modular NIKE scheme with a reduction significantly tighter than previous reductions. Concretely, our reduction targets the ℓ-Linear assumption in pairing-friendly groups, and has a loss of $n/2$, where n is the number of users an adversary interacts with. Thus, our scheme is the first to break (or, rather, circumvent) the lower bound of Bader et al. [4]. As a technical tool, we also present a generic transformation that turns any mildly secure NIKE scheme (i.e., secure only against passive adversaries) into a strongly secure one (secure against active adversaries).

Second, we show that our security reduction is optimal, in the sense that we can generalize the result of Bader et al. [4] to our scheme, at the price of a smaller lower bound (of precisely $n/2$). Our generalization follows the high-level ideas of Bader et al. (who in turn follow Coron's work [12]). However, unlike their result, we even consider NIKE schemes and reductions that make nontrivial changes to the public-key infrastructure itself. We believe that our second result points out the inherent difference between the public-key or signature settings (in which we already have tightly secure schemes from standard assumptions), and the NIKE setting (in which a broader range of lower bounds holds, and, to our knowledge, no tight schemes exist).

We note that in line with previous works [4,24], our negative result does not consider schemes or reductions in the random oracle model.

1.1 Technical Overview

In order to describe our results, it will be helpful to first recall existing lower bounds results (and in particular the result of Bader et al. [4]). This way, we will be able to detail how we circumvent these lower bounds, and what other obstacles still block the way to a tight reduction.

A closer look on existing lower bound results. It might be interesting to see why these lower bounds do not contradict any of the constructions mentioned above. All mentioned lower bounds use a "meta-reduction" (cf. [9]) that turns any tight reduction into a successful problem solver (even *without* a given successful adversary). To describe how a meta-reduction works, assume a reduction R that interacts with an adversary $\mathcal{A}$. Assume further that R first solves a number of problem instances for $\mathcal{A}$, and then expects $\mathcal{A}$ to solve a new problem instance. (For instance, in the signature setting, R might first generate many signatures for $\mathcal{A}$ on messages of $\mathcal{A}$'s choice, and then expect $\mathcal{A}$ to forge a signature for a fresh message.) R will then try to solve its own input instance using the fresh solution provided by $\mathcal{A}$.

Now a meta-reduction M runs R, and takes the place of $\mathcal{A}$ in an interaction with R. Intuitively, M will try to feed R with R's own problem solutions, and hope that R can use one of those to solve its own input. Of course, security games generally require the adversary to generate a *fresh* problem solution to avoid trivial attacks. (For instance, the standard security game for signatures [22] requires the adversary to forge a signature for a message that has not been signed before.) Hence, M runs R *twice*: in the first run, M asks R for the solutions to, say, q randomly chosen problem instances $z_1, \ldots, z_q$. Then, M rewinds R, asks for solutions to *different* problem instances $\tilde{z}_i$, and submits the previously obtained solution to one z_i as fresh solution.

Of course, R may fail to convert a z_i-solution into a solution to its own input *sometimes* (depending on its reduction loss), and this leaves a "loophole" for R to escape the meta-reduction strategy of M. However, a combinatorial argument of [12] shows that R must have a reduction loss of $\Omega(q_{\mathsf{sig}})$ to use this loophole.

For this strategy of M, it is essential that the reduction R will "accept" a problem solution that it has generated itself. To this end, [12,32] require unique signatures (i.e., problem solutions), and [4,29] require re-randomizable signatures (so that any valid signature produced by R can be converted in a random signature by M). However, this property is violated (in a very strong sense) by many of the tightly secure signature schemes mentioned above (e.g., [1,27,36,37]). Specifically, the corresponding (tight) reductions find a way to produce special valid-looking signatures for an adversary that are however useless to solving a problem instance. (Of course, these signatures are not re-randomizable or unique.)

The argument of Bader et al. for NIKE schemes. Bader et al. [4] adapt the above argument to NIKE schemes. To describe their argument, we first recall the NIKE security experiment (according to [10]). A NIKE adversary may request an arbitrary number n of public keys pk_i, and may adaptively corrupt an arbitrary subset of them (in which case the adversary gets the corresponding secret keys sk_i).[2] Finally, the adversary selects two public keys $\mathsf{pk}_{i^*}, \mathsf{pk}_{j^*}$ that have not been corrupted, and then must distinguish between their shared key K_{i^*,j^*}, and an independently random value.[3]

Now assume a reduction R that turns any NIKE adversary into a successful problem solver. This reduction R has to be able to answer adversarial corruption queries, and come up with the corresponding secret keys sk_i. Intuitively, a meta-reduction M can take the role of an adversary, and first obtain some of these keys sk_i from R. Then, M can rewind R, and choose to be challenged on a shared key K_{i^*,j^*} that can be computed from one previously obtained sk_i.

The main difference to the signature case above is that n public keys pk_i give rise to $O(n^2)$ shared keys (or, problem instances/solutions) K_{ij}. In particular, $O(n)$ corruptions enable M to compute $O(n^2)$ shared keys (and thus to potentially solve a quadratic number of shared key challenges). If R turns any of those challenge solutions into a problem solution, then M succeeds. Hence, R must fail with probability $1 - O(1/n^2)$. (Another way to view this is that the reduction's success has to vanish with the failures of the simulation.)

How to circumvent the NIKE lower bound. However, similar to previous works, Bader et al. assume that any secret key (or, more generally, problem solution) output by R can be used to solve corresponding challenges posed by R. This assumption can in fact be violated easily, e.g., by allowing many different secret keys per public key. (That is, a secret key is not uniquely determined by a given public key and, e.g., R may hand out different secret keys upon a corruption query.) Furthermore, different secret keys (for a given public key) may behave differently in the computation of shared keys, and thus may not necessarily be useful in solving a given challenge. Similar ideas are at the core of known techniques for improving tightness, in particular in the context of corruptions [5].

While this first thought allows to circumvent the lower bound of Bader et al., its concrete implementation is not clear at all in the context of NIKE schemes.

[2] We omit additional capabilities of the adversary which are not relevant for this overview.

[3] Like [4], we consider only one challenge pair of public keys (and not an arbitrary number, like the "m-CKS-heavy" notion of [17].

In particular, there should be many secret keys (with different functionality) for a given public key, but the secret keys obtained through corruptions should still satisfy correctness (in the sense that pk_i and sk_j lead to the same shared key as sk_i and pk_j). (We note that this obstacle is specific to NIKE schemes, and in our opinion the main reason why obtaining tightly secure NIKE schemes appears to be particularly difficult.)

Our scheme. To explain our solution, it might be easiest to first outline our scheme (which, in its basic form, is a variation of the password-authenticated key exchange scheme of [21,33]). Let L be a language, and assume a hash proof system (HPS) for L with public keys hpk and secret keys hsk. We write $H_{\mathsf{hsk}}(x)$ for hash proof of an L-instance x under key hsk. Then, public and secret keys of our NIKE scheme are of the following form:

$$\mathsf{pk} = (\mathsf{hpk}, x) \qquad\qquad \mathsf{sk} = (\mathsf{hsk}, x, w),$$

where $x \in L$ with witness w, and a HPS keypair $(\mathsf{hpk}, \mathsf{hsk})$ are randomly chosen. Given $\mathsf{pk}_i = (\mathsf{hpk}_i, x_i)$ and $\mathsf{sk}_j = (\mathsf{hsk}_j, x_j, w_j)$, the corresponding NIKE shared key is computed as $K_{ij} = H_{\mathsf{hsk}_j}(x_i) \cdot H_{\mathsf{hsk}_i}(x_j)$, where the hash value $H_{\mathsf{hsk}_i}(x_j)$ is computed from (and uniquely determined by) hpk_i and w_j. We have correctness in the sense $K_{ji} = H_{\mathsf{hsk}_i}(x_j) \cdot H_{\mathsf{hsk}_j}(x_i) = H_{\mathsf{hsk}_j}(x_i) \cdot H_{\mathsf{hsk}_i}(x_j) = K_{ij}$.

Recall that there are many HPS secret keys hsk for any given public key hpk. However, all these secret keys act identically on any $x \in L$. Hence, in order to benefit from the non-uniqueness of hsk, a NIKE reduction will have to switch at least one $x \in L$ in a NIKE public key pk_i to a no-instance $x \notin L$. Let us call such a NIKE public key (with $x \notin L$) "invalid". For an invalid pk_i, no (full) secret key exists. This means that our reduction must hope that no invalid pk_i is ever corrupted. Since a NIKE adversary may corrupt all public keys except for the two selected challenge keys $\mathsf{pk}_{i^*}, \mathsf{pk}_{j^*}$, this means that our reduction may instead fail with probability $1 - 2/n$.

In other words, already with one invalid public key, our reduction has a loss of at least $n/2$. On the bright side, we will present a strategy that uses precisely one invalid public key to leverage a NIKE security reduction (with loss $n/2$). This reduction is of course far from tight, but it has a loss still considerably better than the $O(n^2)$ lower bound by Bader et al., and thus is significantly tighter than previous constructions. In a nutshell, our security proof proceeds in game hops:

1. We start with the NIKE security game.
2. We guess one index i^*, and hope that pk_{i^*} is one of the challenge public keys finally selected in the adversary's challenge. (If this is not the case, the reduction fails.) Since there are 2 challenge public keys, this step loses a factor of $n/2$.
3. We choose $x_{i^*} \notin L$. Since we may assume that pk_{i^*} is selected as challenge, this change will not be detectable (assuming L has a hard subset membership problem).

4. Finally, we observe that now, *all* keys K_{i^*j} (for arbitrary j) are randomized by the smoothness of the underlying HPS. In fact, HPS smoothness implies that K_{i^*j} is close to uniform, even given pk_j. In particular, this holds for $j = j^*$ and the final challenge $K_{i^*j^*}$.

Note that while [10] also crucially relies on HPSs, there are significant technical differences. Namely, [10] uses hash proof systems mainly as a tool to implement a "replacement decryption method" that allows to forget parts of the secret key. In other words, they use HPSs exclusively in "proof mode". In contrast, for our basic NIKE scheme we use the HPS only in "randomization mode", i.e. to randomize shared keys.

Instantiations and variants. Our basic scheme only requires a HPS for a language with hard subset membership problem, and thus can be implemented efficiently from various computational assumptions (such as the DDH [13], ℓ-Linear [30], DCR [13], or QR [13] assumptions). However, this basic scheme satisfies only a relatively mild form of security called "honest key registration" or "HKR" security in [17]. Hence, we also present a general transformation that turns *any* mildly secure NIKE scheme into one that satisfies a stronger form of security (dubbed "dishonest key registration" or "DKR" security in [17]). Our scheme requires a suitable non-interactive zero-knowledge proof system, and, very loosely speaking, adapts the Naor-Yung paradigm [38] to NIKE schemes. We finally give a concrete and optimized instance under the ℓ-MDDH assumption [15] (for any $\ell \geq 2$ in pairing-friendly groups). For details we refer to the full version [25].

We note that we view our construction as a "first" that demonstrates how to circumvent existing lower bounds for a particularly challenging application. We do not claim superior efficiency of our (fully secure) scheme over existing state-of-the-art NIKE schemes, not even when taking into account the reduction loss in the choice of group sizes. Still, Fig. 1 provides an overview over existing NIKE schemes, in particular in comparison to our scheme.

Our new lower bound. Even though it breaks the existing bound of Bader et al. [4], the reduction loss (of $O(n)$) of our scheme might be a bit disappointing. Our second result shows that we can extend the results from [4] to show that the reduction loss (at least for our scheme) is optimal. Specifically, we are able to give new lower bounds on the tightness of NIKE reductions even for schemes with invalid public keys.

In more detail, we show that a weak validity check (on public keys) is sufficient to prove a meaningful lower bound. Namely, we require that validity of a public key (in the sense that two valid public keys admit only one shared key) is verifiable given that public key *and one of its possible secret keys*. Hence, as long as a given public key is not corrupted, its validity may not be efficiently verifiable, and a reduction can hope to substitute it with an invalid key. (Note that this is precisely what happens in the proof of our NIKE scheme.)

On the other hand, this weak validity check allows us to again apply a rewinding argument as in [4]. Namely, as soon as the reduction returns a secret key on an extraction query, we can check whether the given public key was actually valid and in this case use the obtained secret key later to compute the unique shared key. The only case where we fail to do so is if the reduction does not return a valid secret key for a certain public key in all rewinding attempts. But then we can simply abort with high probability, namely in case this public key is part of the extraction queries (which happens with probability $1-2/n$). In other words, we prove that the best a reduction can do is to switch one public key to invalid and hope that this public key is not part of the extraction queries. We can thus conclude that a NIKE (such as ours) that admits a non-public validity check still suffers from a security reduction loss of at least $n/2$.

Roadmap. In Sect. 2 we provide the necessary preliminaries. In Sect. 3 we present our construction of a mildly secure NIKE with a security reduction whose tightness significantly improves upon existing NIKEs. In Sect. 4 we show how to transform a mildly secure NIKE into a strongly secure one. In Sect. 5 we prove a new lower bound for a broad class of NIKE schemes including ours. In the full version [25] we provide a concrete instantiation of our NIKE. Further, we show how to tweak efficiency of the transformation from mild to strong security when using our NIKE construction.

2 Preliminaries

Notation. Throughout the paper, λ denotes the security parameter. We say that a function is *negligible in* λ if its inverse vanishes asymptotically faster than any polynomial in λ. If a probabilistic algorithm $\mathcal{A}$ has running time polynomial in λ, we say that $\mathcal{A}$ is *probabilistic polynomial time* (PPT). We use $y \leftarrow \mathcal{A}(x)$ to denote that y is assigned the output of $\mathcal{A}$ running on input x, and we write $y \leftarrow \mathcal{A}(x; r)$ to make the randomness r used by a probabilistic algorithm explicit. We use $y \overset{\$}{\leftarrow} X$ to denote sampling from a set X uniformly at random. For $n \in \mathbb{N}$ by $[n]$ we denote the set $\{1, \ldots, n\}$. Let $\varepsilon \in [0, 1]$ and $\mathcal{X}, \mathcal{Y}$ distributions. To denote that $\mathcal{X}$ and $\mathcal{Y}$ have statistical distance at most ε, we write $\mathcal{X} \equiv_\varepsilon \mathcal{Y}$ and say $\mathcal{X}$ and $\mathcal{Y}$ are *ε-close.*

2.1 Hash Proof Systems

Definition 1 (Subset membership problem). *We call* $\mathtt{SMP} := \mathtt{Setup}$ *a* subset membership problem*, if* $\mathtt{Setup}$ *is a PPT algorithm with the following properties.*

$\mathtt{Setup}(1^\lambda)$ *outputs a compact (i.e. with length polynomial in* λ*) description* (X, L, R)*, where* $L \subset X$ *are sets and* R *is an efficiently computable relation with*

$$x \in L \iff \exists \text{ witness } w \text{ with } (x, w) \in R.$$

(We say a relation R *is efficiently computable if given a pair* (x, w) *it can be efficiently checked whether* $(x, w) \in R$*.)*

Further we require for all (X, L, R) in the image of $\texttt{Setup}$ *that it is possible to efficiently sample elements x uniformly at random from $X \backslash L$ (written $x \xleftarrow{\$} X \backslash L$) and to sample elements x uniformly random from L together with witness w (written $(x, w) \xleftarrow{\$} R$).*

Definition 2 (Subset membership assumption). *Let* $\mathtt{SMP}$ *be a subset membership problem. We say that the* subset membership assumption *holds for* $\mathtt{SMP}$*, if for all PPT algorithms $\mathcal{A}$ it holds that*

$$\begin{aligned}\mathrm{Adv}^{\mathsf{smp}}_{\mathcal{A},\mathtt{SMP}}(\lambda) :=& |\Pr[\mathcal{A}(1^\lambda, (X, L, R), x) = 1 | (x, w) \xleftarrow{\$} R] \\ &- \Pr[\mathcal{A}(1^\lambda, (X, L, R), x) = 1 | x \xleftarrow{\$} X \setminus L]|\end{aligned}$$

is negligible in λ, where $(X, L, R) \xleftarrow{\$} \mathtt{SMP.Setup}(1^\lambda)$.

We will employ the notion of a hash proof system based on [13].

Definition 3 (Hash Proof Systems (HPS)). *Let* $\mathtt{SMP}$ *be a subset membership problem. We call* $\mathtt{HPS} := \mathtt{Setup}$ *a hash proof system for* $\mathtt{SMP}$*, if it is a PPT algorithm of the following form.*

$\mathtt{Setup}(1^\lambda)$ *first samples public parameters $\mathcal{PP}_{\mathtt{SMP}} := (X, L, R) \leftarrow \mathtt{SMP.Setup}(1^\lambda)$ for the underlying subset membership problem. Further* $\mathtt{Setup}$ *chooses sets $\mathcal{HSK}, \Pi, \mathcal{HPK}$ such that elements can be efficiently sampled at random from $\mathcal{HSK}$ (denoted $\mathsf{hsk} \xleftarrow{\$} \mathcal{HSK}$). Further* $\mathtt{Setup}$ *chooses an efficiently computable map*

$$\alpha : \mathcal{HSK} \longrightarrow \mathcal{HPK},$$

a family of efficiently computable functions

$$\mathcal{H} := \{H_{\mathsf{hsk}} : X \longrightarrow \Pi \mid \mathsf{hsk} \in \mathcal{HSK}\}$$

and an efficiently computable map

$$F : R \times \mathcal{HPK} \longrightarrow \Pi$$

such that for all $\mathsf{hsk} \in \mathcal{HSK}, \mathsf{hpk} \in \mathcal{HPK}$ with $\alpha(\mathsf{hsk}) = \mathsf{hpk}$ and for all $(x, w) \in R$ we have

$$H_{\mathsf{hsk}}(x) = F(x, w, \mathsf{hpk}).$$

Finally, $\mathtt{Setup}$ *outputs $\mathcal{PP} := (\mathcal{PP}_{\mathtt{SMP}}, \mathcal{HSK}, \mathcal{H}, \alpha, F)$, which contains $\mathcal{PP}_{\mathtt{SMP}}$ together with the compact (i.e. with length polynomial in λ) description of $\mathcal{HSK}, \mathcal{H}, \alpha$ and F.*

We need a property of a HPS called smoothness, introduced in [13].

Definition 4 (Smoothness). *Let* `SMP` *be a subset membership problem and* `HPS` *be a hash proof system for* `SMP`*. We call* `HPS` ε-smooth *if for all* $\mathcal{PP} := ((X, L, R), \mathcal{HSK}, \mathcal{H}, \alpha, F)$ *in the image of* `HPS.Setup`*, the following distributions are* ε*-close:*

$$\left\{ (x, \mathsf{hpk}, H_{\mathsf{hsk}}(x)) \;\middle|\; \begin{array}{c} \mathsf{hsk} \xleftarrow{\$} \mathcal{K} \\ \mathsf{hpk} := \alpha(\mathsf{hsk}) \\ x \xleftarrow{\$} X \setminus L \end{array} \right\} \equiv_\varepsilon \left\{ (x, \mathsf{hpk}, \pi) \;\middle|\; \begin{array}{c} \mathsf{hsk} \xleftarrow{\$} \mathcal{K} \\ \mathsf{hpk} := \alpha(\mathsf{hsk}) \\ x \leftarrow X \setminus L, \pi \xleftarrow{\$} \Pi \end{array} \right\}.$$

(Recall that Π *is the image set of* H_{hsk}*.) In other words, on statements* x *outside the language* L*, the output of the private evaluation algorithms is* ε*-close to uniformly random even under knowledge of the public key. Note though that this statement only holds as long as no image of* H_{hsk} *on input* $x \in X \setminus L$ *is known.*

2.2 Non-Interactive Key Exchange (NIKE)

We formally define the notion of NIKE, following [10,17] and also adopting most of their notation. A NIKE scheme `NIKE` consists of three algorithms (`Setup`, `KeyGen`, `SharedKey`), an identity space $\mathcal{IDS}$ and a shared key space $\mathcal{K}$ which is the output space of `SharedKey`.

- `Setup`: On input 1^λ, this probabilistic algorithm outputs the system parameters $\mathcal{PP}$.
- `KeyGen`: On input $\mathcal{PP}$ and an ID ID, this probabilistic algorithm outputs a tuple $(\mathsf{pk}, \mathsf{sk}) \in \mathcal{PK} \times \mathcal{SK}$.
- `SharedKey`: On input of the public parameters $\mathcal{PP}$ and two identity, public key pairs $(\mathsf{ID}_1, \mathsf{pk}_1), (\mathsf{ID}_2, \mathsf{sk}_2)$, this deterministic algorithm outputs a shared key $K_{12} \in \mathcal{K}$. We assume that $\mathcal{K}$ contains a failure symbol $\bot$.

Table 1. Types of queries for different security models, taken from [17], where q_x denotes the maximum number of allowed queries of the adversary to oracle $\mathcal{O}_x$. ✓, - and n mean that an adversary is allowed to make arbitrary, zero or n queries of this type, in an arbitrary order.

Model	q_{regH}	q_{regC}	q_{extr}	q_{revH}	q_{revC}	q_{test}
DKR CKS-light	2	✓	-	-	✓	1
DKR CKS	✓	✓	-	-	✓	✓
DKR CKS-heavy	✓	✓	✓	✓	✓	1
DKR m-CKS-heavy	✓	✓	✓	✓	✓	✓
HKR CKS-light	2	-	-	-	-	1
HKR CKS	✓	-	-	-	-	✓
HKR CKS-heavy	✓	-	✓	✓	-	1
HKR m-CKS-heavy	✓	-	✓	✓	-	✓

```
$\mathrm{Exp}^{[\mathsf{hkr}|\mathsf{dkr}]-\mathsf{cks}-\mathsf{heavy}}_{\mathcal{A},\mathtt{NIKE}}(\lambda)$:
  $\mathcal{PP} \xleftarrow{\$} \mathtt{NIKE.Setup}(1^{\lambda})$
  $Q_{\mathsf{regH}} := \emptyset$, $Q_{\mathsf{regC}} := \emptyset$, $Q_{\mathsf{extr}} := \emptyset$,
  $Q_{\mathsf{rev}} := \emptyset$
  $b^\star \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{H}},\ \mathcal{O}_{\mathsf{regC}}(\cdot),\ \mathcal{O}_{\mathsf{revC}}(\cdot,\cdot)}(\mathcal{PP})$
  if $b = b^\star \wedge \mathsf{ID}_1^\star, \mathsf{ID}_2^\star \notin Q_{\mathsf{extr}}$
   $\wedge \{\mathsf{ID}_1^\star, \mathsf{ID}_2^\star\} \notin Q_{\mathsf{rev}}$
    output $1$
  else
    $b' \xleftarrow{\$} \{0,1\}$
    output $b'$

$\mathcal{O}_{\mathsf{regH}}(\mathsf{ID})$:
  if $(\mathsf{ID},\cdot,\cdot) \notin \mathcal{O}_{\mathsf{regC}} \cup Q_{\mathsf{regH}}$
    $(\mathsf{pk},\mathsf{sk}) \xleftarrow{\$} \mathtt{NIKE.KeyGen}(\mathcal{PP},\mathsf{ID})$
    $Q_{\mathsf{regH}} := Q_{\mathsf{regH}} \cup \{(\mathsf{ID},\mathsf{pk},\mathsf{sk})\}$
    return $\mathsf{pk}$
  else return $\perp$

$\mathcal{O}_{\mathsf{regC}}(\mathsf{ID},\mathsf{pk})$:
  if $(\mathsf{ID},\cdot,\cdot) \notin \mathcal{O}_{\mathsf{regH}} \cup \mathcal{O}_{\mathsf{regC}}$
    $Q_{\mathsf{regC}} := Q_{\mathsf{regC}} \cup \{(\mathsf{ID},\mathsf{pk},\perp)\}$
  else return $\perp$

$\mathcal{O}_{\mathsf{extr}}(\mathsf{ID})$:
  if $\exists \mathsf{sk}\colon (\mathsf{ID},\mathsf{pk},\mathsf{sk}) \in Q_{\mathsf{regH}}$
    $Q_{\mathsf{extr}} := Q_{\mathsf{extr}} \cup \{\mathsf{ID}\}$
    return $\mathsf{sk}$
  else return $\perp$
```

```
$\mathcal{O}_{\mathsf{revH}}(\mathsf{ID}_1,\mathsf{ID}_2)$:
  if $\exists \mathsf{sk}_1,\mathsf{sk}_2\colon (\mathsf{ID}_1,\mathsf{pk}_1,\mathsf{sk}_1),$
           $(\mathsf{ID}_2,\mathsf{pk}_2,\mathsf{sk}_2) \in Q_{\mathsf{regH}}$
    $Q_{\mathsf{rev}} := Q_{\mathsf{rev}} \cup \{\{\mathsf{ID}_1,\mathsf{ID}_2\}\}$
    return $\mathtt{NIKE.SharedKey}(\mathsf{ID}_1,\mathsf{pk}_1,\mathsf{ID}_2,\mathsf{sk}_2)$
  else return $\perp$

$\mathcal{O}_{\mathsf{revC}}(\mathsf{ID}_1,\mathsf{ID}_2)$:
  if $\exists \mathsf{sk}_1\colon (\mathsf{ID}_1,\mathsf{pk}_1,\mathsf{sk}_1) \in Q_{\mathsf{regH}},$
           $(\mathsf{ID}_2,\mathsf{pk}_2,\cdot) \in Q_{\mathsf{regC}}$
    $Q_{\mathsf{rev}} := Q_{\mathsf{rev}} \cup \{\{\mathsf{ID}_1,\mathsf{ID}_2\}\}$
    return $\mathtt{NIKE.SharedKey}(\mathsf{ID}_2,\mathsf{pk}_2,\mathsf{ID}_1,\mathsf{sk}_1)$
  if $\exists \mathsf{sk}_2\colon (\mathsf{ID}_2,\mathsf{pk}_2,\mathsf{sk}_2) \in Q_{\mathsf{regH}},$
           $(\mathsf{ID}_1,\mathsf{pk}_1,\cdot) \in Q_{\mathsf{regC}}$
    $Q_{\mathsf{rev}} := Q_{\mathsf{rev}} \cup \{\{\mathsf{ID}_1,\mathsf{ID}_2\}\}$
    return $\mathtt{NIKE.SharedKey}(\mathsf{ID}_1,\mathsf{pk}_1,\mathsf{ID}_2,\mathsf{sk}_2)$
  else return $\perp$

$\mathcal{O}_{\mathsf{test}}(\mathsf{ID}_1^\star,\mathsf{ID}_2^\star)$:
  $b \xleftarrow{\$} \{0,1\}$
  if $\exists \mathsf{sk}_1^\star,\mathsf{sk}_2^\star\colon (\mathsf{ID}_1^\star,\mathsf{pk}_1^\star,\mathsf{sk}_1^\star),$
           $(\mathsf{ID}_2^\star,\mathsf{pk}_2^\star,\mathsf{sk}_2^\star) \in Q_{\mathsf{regH}}$
    $K_0 = \mathtt{NIKE.SharedKey}(\mathsf{ID}_1^\star,\mathsf{pk}_1^\star,\mathsf{ID}_2^\star,\mathsf{sk}_2^\star)$
    $K_1 \xleftarrow{\$} \mathcal{K}$
    return $K_b$
  else return $\perp$
```

Fig. 2. Experiment for HKR and DKR CKS-heavy security of a NIKE scheme NIKE with shared key space $\mathcal{K}$. The highlighted parts only occur in the setting of dishonest key registration. The oracle $\mathcal{O}_{\mathsf{test}}$ may only be queried once. $\mathcal{O}_{\mathsf{H}}$ comprises the oracles $\mathcal{O}_{\mathsf{regH}}, \mathcal{O}_{\mathsf{revH}}, \mathcal{O}_{\mathsf{extr}}$ and $\mathcal{O}_{\mathsf{test}}$. We use $\cdot$ to denote an arbitrary entry of a tuple. I.e., $\mathcal{O}_{\mathsf{regH}} \setminus \{(\mathsf{ID},\cdot,\cdot)\}$ denotes the set $\mathcal{O}_{\mathsf{regH}}$ without any tuple that contains ID in the first position.

We always require NIKE to be perfectly correct, meaning that for all corresponding key pairs $(\mathsf{ID}_1,\mathsf{pk}_1,\mathsf{sk}_1), (\mathsf{ID}_2,\mathsf{pk}_2,\mathsf{sk}_2)$ generated by KeyGen it holds

$$\mathtt{SharedKey}(\mathsf{ID}_1,\mathsf{pk}_1,\mathsf{ID}_2,\mathsf{sk}_2) = \mathtt{SharedKey}(\mathsf{ID}_2,\mathsf{pk}_2,\mathsf{ID}_1,\mathsf{sk}_1) \neq \perp$$

SECURITY. We quickly recall the game-based security notion from [10], called the *CKS model*, with its refinements from [17]. The model is defined via adversarial queries to oracles implemented by a challenger $\mathcal{C}$. The challenger $\mathcal{C}$ keeps track of all honest and corrupt registered identities and their keys. We informally describe the oracles provided to the adversary attacking a NIKE $\mathtt{NIKE}$ below.

- $\mathcal{O}_{\mathsf{regH}}$ for registering an honest user. $\mathcal{C}$ generates a key pair using $\mathtt{NIKE.KeyGen}$ and hands the public key to the adversary.
- $\mathcal{O}_{\mathsf{regC}}$ for registering a corrupt user. The adversary may introduce a public key without providing the corresponding secret key.
- $\mathcal{O}_{\mathsf{extr}}$ for extracting a secret key of an honest user.
- $\mathcal{O}_{\mathsf{revH}}$ for revealing a shared key of an honest pair of users.
- $\mathcal{O}_{\mathsf{revC}}$ for revealing a shared key between a corrupted and an honest user.
- $\mathcal{O}_{\mathsf{test}}$ for obtaining a challenge. $\mathcal{A}$ provides a pair of users it wishes to be challenged upon. $\mathcal{C}$ then flips a coin and replies either with their real shared key or a random one.

First, $\mathcal{C}$ runs $\mathcal{PP} \overset{\$}{\leftarrow} \mathtt{NIKE.Setup}(1^\lambda)$ and gives $\mathcal{PP}$ to $\mathcal{A}$. Then, the adversary may make an arbitrary number of the above queries, in an arbitrary order. Finally, the adversary outputs a bit $\hat{b}$ and wins if $\hat{b} = b$. Note that the adversary may register each ID only once[4].

To obtain different notions of CKS security, the adversary is restricted in the number of its queries. See Table 1 for a complete list. Notions that admit $\mathcal{O}_{\mathsf{regC}}$ and $\mathcal{O}_{\mathsf{revC}}$ queries are said to *allow dishonest key registrations*, dubbed *DKR*. Notions that do not allow such types of queries are called *with honest key registration*, or *HKR* for short.

In this paper, we are interested in *CKS - heavy* secure NIKE schemes. We provide the corresponding security experiment in Fig. 2.

Definition 5 (HKR- and DKR-CKS-heavy security). *Let NIKE be a NIKE. We say NIKE is* CKS-heavy secure with honest key registration*, or* HKR-CKS-heavy secure*, if for any PPT adversary $\mathcal{A}$ the advantage*

$$\mathrm{Adv}^{\mathsf{hkr-cks-heavy}}_{\mathcal{A},\mathit{NIKE}}(\lambda) = |\Pr[\mathrm{Exp}^{\mathsf{hkr-cks-heavy}}_{\mathcal{A},\mathit{NIKE}}(\lambda) \Rightarrow 1] - 1/2|$$

is negligible in λ, where $\mathrm{Exp}^{\mathsf{hkr-cks-heavy}}_{\mathcal{A},\mathit{NIKE}}$ is provided in Fig. 2. Similarly, we say that NIKE is CKS-heavy secure with dishonest key registration*, or* DKR-CKS-heavy secure*, if for any PPT adversary $\mathcal{A}$ the advantage*

$$\mathrm{Adv}^{\mathsf{dkr-cks-heavy}}_{\mathcal{A},\mathit{NIKE}}(\lambda) = |\Pr[\mathrm{Exp}^{\mathsf{dkr-cks-heavy}}_{\mathcal{A},\mathit{NIKE}}(\lambda) \Rightarrow 1] - 1/2|$$

is negligible in λ.

2.3 Public Key Encryption

Definition 6 (Public key encryption). *We call a tuple of PPT algorithms $\mathit{PKE} := (\mathit{KeyGen}, \mathit{Enc}, \mathit{Dec})$ a* public key encryption scheme *if the following holds.*

[4] In practice, this can be implemented by appending a counter to an identity string.

```
$\mathrm{Exp}^{\mathsf{ind-cpa}}_{\mathcal{A}=(\mathcal{A}_1,\mathcal{A}_2),\mathtt{PKE}}(\lambda)$:
    $(\mathsf{ppk}, \mathsf{psk}) \leftarrow \mathtt{PKE.KeyGen}(1^\lambda)$
    $(M_0, M_1, st) \leftarrow \mathcal{A}_1(1^\lambda, \mathsf{ppk})$
    $b \xleftarrow{\$} \{0,1\}$
    $C := \mathtt{Enc}(\mathsf{ppk}, M_b)$
    $b^\star \leftarrow \mathcal{A}_2(st, C)$
    if  $b = b^\star$  output  $1$
    else output  $0$
```

Fig. 3. IND-CPA experiment.

- *$\mathtt{KeyGen}(1^\lambda)$ returns a key pair $(\mathsf{ppk}, \mathsf{psk})$.*
- *$\mathtt{Enc}(\mathsf{ppk}, M)$ returns a ciphertext C.*
- *$\mathtt{Dec}(\mathsf{psk}, C)$ returns a message M or a special rejection symbol $\perp$.*

We further require Correctness*, that is for all $(\mathsf{ppk}, \mathsf{psk})$ in the range of $\mathtt{KeyGen}(1^\lambda)$, for all messages M and for all C in the range of $\mathtt{Enc}(\mathsf{pk}, M)$ we require*

$$\mathtt{Dec}(\mathsf{sk}, C) = 1.$$

Definition 7 (IND-CPA). *Let $\mathtt{PKE}$ be a public key encryption scheme. We say $\mathtt{PKE}$ is* IND-CPA secure *if for all PPT adversaries $\mathcal{A}$ we have that*

$$\mathrm{Adv}^{\mathsf{ind-cpa}}_{\mathcal{A},\mathtt{PKE}}(\lambda) := |\Pr[\mathrm{Exp}^{\mathsf{ind-cpa}}_{\mathcal{A},\mathtt{PKE}}(\lambda) \Rightarrow 1] - 1/2|$$

is negligible in λ, where $\mathrm{Exp}^{\mathsf{ind-cpa}}_{\mathcal{A},\mathtt{PKE}}(\lambda)$ is defined as in Fig. 3 and we require $|M_0| = |M_1|$.

2.4 Non-Interactive Zero Knowledge Proof of Knowledge

The notion of a quasi-adaptive non-interactive zero-knowledge proof was introduced in [8]. The following definition of non-interactive zero-knowledge is an adaptation of [20] with some differences. Note for instance, that we consider computational zero-knowledge instead of perfect zero-knowledge. We will employ such proofs to generically transform a NIKE which is secure in the HKR-CKS-heavy security model to a NIKE which is secure in the DKR-CKS-heavy security model.

Definition 8 (QANIZK). *Let $\mathtt{SMP}$ be a subset membership problem. Let $(X, L, R) \leftarrow \mathtt{SMP.Setup}(1^\lambda)$. A* quasi adaptive non-interactive zero-knowledge proof (QANIZK) *for $\mathtt{SMP}$ is a tuple of PPT algorithms $\mathtt{PS} := (\mathtt{Setup}, \mathtt{Gen}, \mathtt{Ver}, \mathtt{Sim})$ of the following form.*

- *$\mathtt{Setup}(1^\lambda, (X, L, R))$ generates a common reference string crs and a trapdoor trp. We assume (X, L, R) to be part of the crs.*

```
$\mathrm{Exp}^{\mathsf{extr}}_{\mathcal{A},\mathsf{PS}}(\lambda)$:
  $(X, L, R) \leftarrow$ SMP.Setup$(1^\lambda)$
  $(\mathsf{crs}, \mathsf{trp}, \mathsf{extr}) \overset{\$}{\leftarrow}$ PS.Setup$(1^\lambda, (X, L, R))$
  $Q_{\mathsf{sim}} := \emptyset$
  $(x^\star, \Pi^\star) \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{sim}}(\cdot), \mathcal{O}_{\mathsf{extract}}(\cdot,\cdot)}(1^\lambda, \mathsf{crs})$
  $w \leftarrow \mathcal{O}_{\mathsf{extract}}(x^\star, \Pi^\star)$
  if PS.Ver$(x^\star, \Pi^\star) = 1 \wedge (x^\star, w) \notin R$
   $\wedge x^\star \notin Q_{\mathsf{sim}}$
     output $1$
  else output $0$
```

```
$\mathcal{O}_{\mathsf{sim}}(x)$:
  $Q_{\mathsf{sim}} := Q_{\mathsf{sim}} \cup \{x\}$
  $\Pi \leftarrow$ PS.Sim$(\mathsf{crs}, \mathsf{trp}, x)$
  return $\Pi$

$\mathcal{O}_{\mathsf{extract}}(x, \Pi)$:
  if $x \notin Q_{\mathsf{sim}}$
     $w \leftarrow$ PS.Extract$(\mathsf{crs}, \mathsf{extr}, x, \Pi)$
     return $w$
  else return $\perp$
```

Fig. 4. Experiment for a extraction in the presence of simulated proofs. The adversary tries to set up a pair (x, Π) such that a witness w is not extractable from Π.

- *Prove*(crs, x, w) *given a word* $x \in L$ *and a witness* w *with* $R(x, w) = 1$*, outputs a proof* Π*.*
- *Ver*(crs, x, Π) *on input* crs*,* $x \in X$ *and* Π *outputs a verdict* $b \in \{0, 1\}$*.*
- *Sim*$(\mathsf{crs}, \mathsf{trp}, x)$ *given a* crs *with corresponding trapdoor* trp *and a word* $x \in X$*, outputs a proof* Π*.*

Further we require the following properties to hold.

Perfect completeness: *For all security parameters* λ*, all* (X, L, R) *in the image of SMP.Setup*(1^λ)*, all* $(\mathsf{crs}, \mathsf{trp})$ *in the range of Setup*$(1^\lambda, (X, L, R))$*, all words* $x \in L$*, all witnesses* w *such that* $R(x, w) = 1$ *and all* Π *in the range of Prove*(crs, x, w) *we have*

$$\textit{Ver}(\mathsf{crs}, x, \Pi) = 1.$$

Computational zero-knowledge: *For all security parameters* λ*, all* (X, L, R) *in the range of SMP.Setup*(1^λ)*, all tuples* $(\mathsf{crs}, \mathsf{trp})$ *in the range of Setup*$(1^\lambda, (X, L, R))$*, we have for all PPT adversaries* $\mathcal{A}$ *that*

$$\mathrm{Adv}^{\mathsf{zk}}_{\mathcal{A},PS}(\lambda) :=\mid \Pr[\mathcal{A}^{\mathcal{O}_{\mathsf{prv}}(\cdot,\cdot)}(1^\lambda, \mathsf{crs}) = 1] - \Pr[\mathcal{A}^{\mathcal{O}_{\mathsf{sim}}(\cdot,\cdot)}(1^\lambda, \mathsf{crs}) = 1 \mid$$

is negligible in λ*, where both oracles on input* (x, w) *first check whether* $(x, w) \in R$*. If this is the case,* $\mathcal{O}_{\mathsf{prv}}$ *returns Prove*(crs, x, w) *and* $\mathcal{O}_{\mathsf{sim}}$ *returns Sim*$(\mathsf{crs}, \mathsf{trp}, x)$ *(and* $\perp$ *otherwise).*

The following definition is tailored to our purposes. We require a strong notion of proof of knowledge in the sense that we need to be able to extract a witness while simulating proofs ourselves.

```
NIKE.Setup($1^\lambda$)
    $(\mathcal{PP}_{\mathtt{SMP}}, \mathcal{HSK}, \mathcal{H}, \alpha, F) \xleftarrow{\$}$ HPS.Setup($1^\lambda$)
    $\mathcal{PP} := (\mathcal{PP}_{\mathtt{SMP}}, \mathcal{HSK}, \mathcal{H}, \alpha, F)$
    return $\mathcal{PP}$

NIKE.SharedKey($\mathcal{PP}, \mathsf{ID}_1, \mathsf{pk}_1, \mathsf{ID}_2, \mathsf{sk}_2$)
    parse $\mathcal{PP} =: (\mathcal{PP}_{\mathtt{SMP}}, \mathcal{HSK}, \mathcal{H}, \alpha, F)$
    parse $\mathsf{pk}_1 =: (\mathsf{hpk}_1, x_1)$
    parse $\mathsf{sk}_2 =: (\mathsf{hsk}_2, x_2, w_2)$
    $K_{12} := H_{\mathsf{hsk}_2}(x_1) \cdot F(x_2, w_2, \mathsf{hpk}_1)$
    return $K_{12}$
```

```
NIKE.KeyGen($\mathcal{PP}, \mathsf{ID}$)
    parse $\mathcal{PP} =: (\mathcal{PP}_{\mathtt{SMP}}, \mathcal{HSK}, \mathcal{H}, \alpha, F)$
    parse $PP_{\mathtt{SMP}} =: (X, L, R)$
    $\mathsf{hsk} \xleftarrow{\$} \mathcal{HSK}$
    $\mathsf{hpk} := \alpha(\mathsf{hsk})$
    $(x, w) \xleftarrow{\$} R$
    $\mathsf{pk} := (\mathsf{hpk}, x)$
    $\mathsf{sk} := (\mathsf{hsk}, x, w)$
    return $(\mathsf{pk}, \mathsf{sk})$
```

Fig. 5. Our NIKE scheme. Recall that $\mathcal{H} = \{H_{\mathsf{hsk}} : X \to \Pi \mid \mathsf{hsk} \in \mathcal{K}\}$ is a family of functions and $F : R \times \mathcal{HPK} \to \Pi$ a function (where $\mathcal{HPK}$ is the image of α).

Definition 9 (QANIZK Proof of knowledge). *Let* PS' *be a QANIZK for a subset membership problem* SMP*, where* $\mathit{SMP}.\mathtt{Setup}$ *returns tuples* (X, L, R)*. Let* $\mathtt{Setup}$ *denote an algorithm that, on input* $(1^\lambda, (X, L, R))$ *runs* $(\mathsf{crs}, \mathsf{trp}) \xleftarrow{\$} \mathit{PS}'.\mathtt{Setup}(1^\lambda, (X, L, R))$ *and outputs* $(\mathsf{crs}, \mathsf{trp}, \mathsf{extr})$ *with an additional trapdoor* extr*. Let* $\mathtt{Gen} := \mathit{PS}'.\mathtt{Gen}, \mathtt{Prove} := \mathit{PS}'.\mathtt{Prove}, \mathtt{Ver} := \mathit{PS}'.\mathtt{Ver}, \mathtt{Sim} := \mathit{PS}'.\mathtt{Sim}$*. Let further* $\mathtt{Extract}$ *be an algorithm that on input* $(\mathsf{crs}, \mathsf{extr}, x, \Pi)$ *returns a witness* w*. We say* $\mathit{PS} = (\mathtt{Setup}, \mathtt{Gen}, \mathtt{Prove}, \mathtt{Ver}, \mathtt{Sim}, \mathtt{Extract})$ *is a* QANIZK Proof of Knowledge for SMP (QANIZKPoK)*, if for all PPT adversaries* $\mathcal{A}$ *the advantage*

$$\mathrm{Adv}^{\mathsf{extr}}_{\mathcal{A},\mathit{PS}}(\lambda) := \Pr[\mathrm{Exp}^{\mathsf{extr}}_{\mathcal{A},\mathit{PS}}(\lambda) \Rightarrow 1]$$

is negligible in λ*, where* $\mathrm{Exp}^{\mathsf{extr}}_{\mathcal{A},\mathit{PS}}(\lambda)$ *is as defined in Fig. 4.*

3 Our Construction

We now present a NIKE scheme that is secure in the HKR setting. Our reduction loses a factor of $q_{\mathsf{regH}}/2$, where q_{regH} is the number of honest users. Our scheme uses a hash proof system and its security relies on the hardness of the underlying subset membership problem as well as the smoothness of the HPS. It is presented in Fig. 3.

Let us first elaborate on why our NIKE scheme does not fall under the impossibility result of Bader et al. [4]. To enforce that the output of a successful NIKE attacker can always be used to solve the challenge given to the reduction, Bader et al. require that the NIKE scheme allows only public keys whose corresponding secret keys are uniquely determined. This way, the shared key between two public keys is uniquely determined and will be useful to solve the challenge. Moreover, the uniqueness condition has to be *efficiently checkable given only a public key.* This essentially prevents a reduction from switching public keys to "invalid" public keys that violate the uniqueness condition. Formally, Bader et al. require an efficient algorithm $\texttt{PKCheck}$ for testing uniqueness.

Our scheme does not provide such an algorithm, since essentially deciding uniqueness amounts to deciding a subset membership problem that we assume to be hard. This way, our reduction will have a way to indistinguishably switch one of the public keys to "invalid" by drawing it from outside the subgroup. Note that for such an invalid public key there exist no secret key, since secret keys contain a witness for the public key belonging to the subgroup. While this non-existence of a secret key helps us in arguing security, it also introduces an inherent loss in our reduction; namely, our reduction has to abort whenever the adversary wants to see the secret key corresponding to the invalid key, which occurs with probability $2/q_{\mathsf{regH}}$ and thus results in a loss of $q_{\mathsf{regH}}/2$. We now provide a proof of security that meets exactly this loss.

Theorem 1. *Let $\mathtt{SMP}$ be a subset membership problem, and let $\mathtt{HPS}$ be a hash proof system for $\mathtt{SMP}$, such that for all $\mathcal{PP} := (\mathcal{PP}_{\mathtt{SMP}}, \mathcal{HSK}, \mathcal{H}, \alpha, F)$ in the range of $\mathtt{HPS.Setup}$ the image Π of F and all $H_{\mathsf{hsk}} \in \mathcal{H}$ is a commutative multiplicative group. If the subset membership assumption holds for $\mathtt{SMP}$ and if $\mathtt{HPS}$ is ε-smooth with ε negligible in λ, then the NIKE scheme $\mathtt{NIKE}$ described in Fig. 5 is a perfectly correct, HKR-CKS-heavy secure NIKE. Further, the reduction to $\mathtt{SMP}$ loses a factor of $q_{\mathsf{regH}}/2$, where q_{regH} is the number of queries to $\mathcal{O}_{\mathsf{regH}}$ that $\mathcal{A}$ makes. More formally, if $\mathcal{A}$ is an adversary with running time $t_{\mathcal{A}}$ against the scheme in the HKR-CKS-heavy model, there exists an adversary $\mathcal{B}$ with running time $t_{\mathcal{B}} \approx t_{\mathcal{A}}$ breaking the subset membership problem $\mathtt{SMP}$ such that*

$$\mathrm{Adv}^{\mathsf{hkr-cks-heavy}}_{\mathcal{A},\mathtt{NIKE}}(\lambda) \;\le\; q_{\mathsf{regH}}/2 \cdot (\mathrm{Adv}^{\mathsf{smp}}_{\mathcal{B},\mathtt{SMP}}(\lambda) + \varepsilon)$$

Proof. Perfect correctness. Let the public parameters be $\mathcal{PP} := (\mathcal{PP}_{\mathtt{SMP}}, \mathcal{HSK}, \mathcal{H}, \alpha, F) \xleftarrow{\$} \mathtt{NIKE.Setup}(1^\lambda)$ and $(\mathsf{pk}_1, \mathsf{sk}_1) \leftarrow \mathtt{NIKE.KeyGen}(\mathcal{PP}, \mathsf{ID}_1), (\mathsf{pk}_2, \mathsf{sk}_2) \leftarrow \mathtt{NIKE.KeyGen}(\mathcal{PP}, \mathsf{ID}_2)$. Let further $\mathsf{pk}_1 =: (\mathsf{hpk}_1, x_1), \mathsf{pk}_2 =: (\mathsf{hpk}_2, x_2)$ and $\mathsf{sk}_1 =: (\mathsf{hsk}_1, x_1, w_1), \mathsf{sk}_2 =: (\mathsf{hsk}_2, x_2, w_2)$. As $\mathtt{HPS}$ is a hash proof system and as $x_1, x_2 \in L$, $\mathsf{hpk}_1 = \alpha(\mathsf{hsk}_1)), \mathsf{hpk}_2 = \alpha(\mathsf{hsk}_2))$ we have

$$H_{\mathsf{hsk}_2}(x_1) = F(x_1, w_1, \mathsf{hpk}_2) \text{ and } H_{\mathsf{hsk}_1}(x_2) = F(x_2, w_2, \mathsf{hpk}_1).$$

This yields

$$K_{12} = H_{\mathsf{hsk}_2}(x_1) \cdot F(x_2, w_2, \mathsf{hpk}_1) = H_{\mathsf{hsk}_1}(x_2) \cdot F(x_1, w_1, \mathsf{hpk}_2) = K_{21}$$

as required.

CKS-heavy security. We prove that the $\mathtt{NIKE}$ meets CKS-heavy security with honest key registration in a number of hybrid games. We provide an overview of the games in Fig. 6. By $\Pr[\mathbf{G}_i]$ we denote the probability that $\mathcal{A}$ wins game $\mathbf{G}_i$.

Game $\mathbf{G}_0$: The real experiment. Game $\mathbf{G}_0$ is the HKR-CKS-heavy experiment as presented in Fig. 2, where $\mathcal{A}$ plays with a challenger $\mathcal{C}$. We have thus

$$\mathrm{Adv}^{\mathsf{hkr-cks-heavy}}_{\mathcal{A},\mathtt{NIKE}}(\lambda) = |\Pr[\mathbf{G}_0] - 1/2|\,.$$

Game	$\mathcal{O}_{\mathsf{regH}}$ if $i = i^\star$	$\mathcal{O}_{\mathsf{extr}}(\mathsf{ID}_{i^\star})$	$\mathcal{O}_{\mathsf{revH}}(\{\mathsf{ID}, \mathsf{ID}_{i^\star}\})$	$\mathcal{O}_{\mathsf{test}}(\{\mathsf{ID}, \mathsf{ID}_{i^\star}\})$	**Explanation**
$\mathbf{G}_0$	$(x, w) \xleftarrow{\$} R$	$\mathsf{sk}_{i^\star}$	$\mathsf{sk}_{i^\star}/\mathsf{sk}$	$\mathsf{sk}_{i^\star}/\mathsf{sk}$	$= \mathrm{Exp}^{\mathsf{hkr-cks-heavy}}_{\mathcal{A},\mathtt{NIKE}}$
$\mathbf{G}_1$	$(x, w) \xleftarrow{\$} R$	$\mathsf{sk}_{i^\star}$	sk	sk	perfect correctness
$\mathbf{G}_2$	$(x, w) \xleftarrow{\$} R$	**abort**	sk	sk	$q_{\mathsf{regH}}/2$ loss
$\mathbf{G}_3$	$x \xleftarrow{\$} X \setminus L$	**abort**	sk	sk	`SMP` assumption
$\mathbf{G}_4$	$x \xleftarrow{\$} X \setminus L$	**abort**	sk	$K_0 \leftarrow \mathcal{K}$	smoothness `HPS`

Fig. 6. Games $\mathbf{G}_0$ to $\mathbf{G}_4$ we employ to prove the NIKE presented in Fig. 3 HKR-CKS-heavy secure. From game $\mathbf{G}_1$ on the index $i^\star \xleftarrow{\$} q_{\mathsf{regH}}$ is chosen ahead of time. By $\mathsf{ID}_{i^\star}$ we denote the $i^\star$-th registered honest user. The oracle $\mathcal{O}_{\mathsf{test}}$ may only be queried once. In Column 4 and 5, we give the secret key employed to compute `NIKE.SharedKey`. By denoting the input as a set $\{\cdot\}$ we want to indicate that we consider both inputs $\mathsf{pk}, \mathsf{pk}_{i^\star}$ and $\mathsf{pk}_{i^\star}, \mathsf{pk}$. In game $\mathbf{G}_0$ there is thus two possibility secret keys to be employed, depending on the order of the input.

Game $\mathbf{G}_1$: Guess the challenge. Recall that by q_{regH} we denote the number of $\mathcal{O}_{\mathsf{regH}}$ queries of $\mathcal{A}$. From game $\mathbf{G}_1$ on, an index $i^\star \leftarrow q_{\mathsf{regH}}$ is chosen ahead of time. The final goal will be to switch the $i^\star$-th registered honest user $\mathsf{ID}_{i^\star}$ to invalid and hope it is part of the test query. As a first step, from game $\mathbf{G}_1$ on we will make $\mathsf{sk}_{i^\star}$ redundant for $\mathcal{O}_{\mathsf{revH}}$ and $\mathcal{O}_{\mathsf{test}}$ queries. Namely, if $\mathcal{A}$ asks a query of this form with input $(\mathsf{ID}, \mathsf{ID}_{i^\star})$ (for an arbitrary identity ID) we will compute the shared key employing sk, where $(\mathsf{ID}, \mathsf{pk}, \mathsf{sk}) \in Q_{\mathsf{regH}}$, instead of $\mathsf{sk}_{i^\star}$. By perfect correctness of `NIKE` we have

$$\Pr[\mathbf{G}_1] = \Pr[\mathbf{G}_0].$$

Game $\mathbf{G}_2$: Abort upon wrong guess. We change the winning condition of the game as follows. If $\mathsf{ID}_{i^\star}$ is not included in the test query of $\mathcal{A}$, the experiment returns 1 with probability 1/2 and aborts. Then it holds

$$\begin{aligned}\Pr[\mathbf{G}_2] &= \Pr[\mathbf{G}_1] \cdot 2/q_{\mathsf{regH}} + 1/2 \cdot (1 - 2/q_{\mathsf{regH}}) \\ &= (\Pr[\mathbf{G}_1] - 1/2) \cdot 2/q_{\mathsf{regH}} + 1/2\end{aligned}$$

and thus

$$\Pr[\mathbf{G}_1] - 1/2 = q_{\mathsf{regH}}/2 \cdot (\Pr[\mathbf{G}_2] - 1/2).$$

Game $\mathbf{G}_3$: Remove the secret key. Upon receiving the $i^\star$-th register honest user query, $\mathcal{C}$ deviates from the `NIKE.KeyGen` procedure as follows: instead of drawing $(x_{i^\star}, w_{i^\star}) \xleftarrow{\$} R$, $\mathcal{C}$ draws $x_{i^\star} \xleftarrow{\$} X \setminus L$. Note that this way there is no $w_{i^\star}$ such that $R(x_{i^\star}, w_{i^\star}) = 1$ and thus $\mathcal{C}$ cannot compute a secret key $\mathsf{sk}_{i^\star}$. Instead, $\mathcal{C}$ adds $(\mathsf{ID}_{i^\star}, \mathsf{pk}_{i^\star}, \mathsf{sk}_{i^\star}) := (\mathsf{ID}_{i^\star}, (\mathsf{hpk}_{i^\star}, x_{i^\star}), (\mathsf{hsk}_{i^\star}, \perp))$ to Q_{regH}. A distinguisher between both games can be turned directly into a SMP attacker $\mathcal{B}$ putting his challenge in the place of $x_{i^\star}$. If the challenge was in L, Game $\mathbf{G}_2$ was simulated, else it was Game $\mathbf{G}_3$. Observe that it is crucial here that

$\mathcal{C}$ does not make use of $w_{i^\star}$ anymore due to the changes made in Game 1. This yields

$$|\Pr[\mathbf{G}_2] - \Pr[\mathbf{G}_3]| \leq \mathrm{Adv}^{\mathsf{smp}}_{\mathcal{B},\mathsf{SMP}}(\lambda).$$

Game $\mathbf{G}_4$: Randomize the test query. $\mathcal{C}$ changes the answer to the query $\mathcal{O}_{\mathsf{test}}(\mathsf{ID}_{i^\star}, \mathsf{ID})$[5] by drawing $K_0 \xleftarrow{\$} \mathcal{K}$, where $\mathcal{K} = \Pi$ is the image of the hash functions of the HPS. To analyze the distinguishing advantage, note that in the former game it holds that $K_0 = \mathtt{NIKE.SharedKey}(\mathsf{ID}_{i^\star}, \mathsf{pk}_{i^\star}, \mathsf{ID}, \mathsf{sk}) = H_{\mathsf{hsk}}(x_{i^\star}) \cdot F(x, w, \mathsf{hpk}_{i^\star})$ with $(\mathsf{ID}, \mathsf{pk}, \mathsf{sk}) = (\mathsf{ID}, (\mathsf{hpk}, x), (\mathsf{hsk}, w)) \in Q_{\mathsf{regH}}$ and $(\mathsf{ID}_{i^\star}, \mathsf{pk}_{i^\star}, \mathsf{sk}_{i^\star}) = (\mathsf{ID}_{i^\star}, (\mathsf{hpk}_{i^\star}, x_{i^\star}), (\mathsf{hsk}_{i^\star}, \perp)) \in Q_{\mathsf{regH}}$. The two distributions $(x_{i^\star}, \mathsf{hpk}, H_{\mathsf{hsk}}(x_{i^\star})), (x_{i^\star}, \mathsf{hpk}, r \xleftarrow{\$} \Pi)$ are ε-close by the ε-smoothness of the HPS, and thus K_0 was already statistically close to the uniform distribution over Π in Game $\mathbf{G}_3$. We thus have

$$|\Pr[\mathbf{G}_3] - \Pr[\mathbf{G}_4]| \leq \varepsilon.$$

We now show that the advantage of $\mathcal{A}$ playing the CKS-heavy game is negligible. We repeatedly use a folklore technique - add zero, then apply the triangle inequality - to go through all the above games until Game $\mathbf{G}_4$, for which the winning probability of $\mathcal{A}$ is 1/2 since its view does not depend on the challenge bit.

$$\begin{aligned}
\mathrm{Adv}^{\mathsf{hkr-cks-heavy}}_{\mathcal{A},\mathtt{NIKE}}(\lambda) &= |\Pr[\mathbf{G}_0] - 1/2| = |\Pr[\mathbf{G}_1] - 1/2| \\
&= q_{\mathsf{regH}}/2 \cdot |\Pr[\mathbf{G}_2] - \Pr[\mathbf{G}_3] + \Pr[\mathbf{G}_3] - 1/2| \\
&\leq q_{\mathsf{regH}}/2 \cdot |\Pr[\mathbf{G}_3] - \Pr[\mathbf{G}_4] + \Pr[\mathbf{G}_4] - 1/2| + q_{\mathsf{regH}}/2 \cdot \mathrm{Adv}^{\mathsf{smp}}_{\mathcal{B},\mathsf{SMP}}(\lambda) \\
&\leq q_{\mathsf{regH}}/2 \cdot |\Pr[\mathbf{G}_4] - 1/2| + q_{\mathsf{regH}}/2 \cdot (\mathrm{Adv}^{\mathsf{smp}}_{\mathcal{B},\mathsf{SMP}}(\lambda) + \varepsilon) \\
&= q_{\mathsf{regH}}/2 \cdot (\mathrm{Adv}^{\mathsf{smp}}_{\mathcal{B},\mathsf{SMP}}(\lambda) + \varepsilon)
\end{aligned}$$

Remark 1. A variant of our NIKE can be obtained if there is a total ordering $<$ on all identities. Then, the shared key of $\mathsf{ID}_1, \mathsf{ID}_2$ can be computed as the hash of the statement provided by the *smaller* identity. More formally, we modify `NIKE.SharedKey` as follows:

$$\begin{aligned}
\mathtt{NIKE.SharedKey}(\mathsf{ID}_1, \mathsf{pk}_1, \mathsf{ID}_2, \mathsf{sk}_2) :&= H_{hsk_2}(x_1) \\
&= F(x_1, w_1, hpk_2) =: \mathtt{NIKE.SharedKey}(\mathsf{ID}_2, \mathsf{pk}_2, \mathsf{ID}_1, \mathsf{sk}_1),
\end{aligned}$$

where $\mathsf{ID}_1 < \mathsf{ID}_2$. The only change in the proof of security is that in game $\mathbf{G}_2$ the challenger aborts if the guessed $i^\star$ is not the *smallest* identity contained in the test query. This yields a reduction loss of q_{regH}.

[5] Note that, starting with Game $\mathbf{G}_2$, $i^\star$ is always one of the inputs to $\mathcal{O}_{\mathsf{test}}$.

```
$\mathtt{NIKE}_{\mathsf{dkr}}.\mathtt{Setup}(1^\lambda)$
  $\mathcal{PP} \leftarrow \mathtt{NIKE.Setup}(1^\lambda)$
  $\mathcal{PP}_{\mathtt{PS}} \leftarrow \mathtt{PS.Setup}(1^\lambda, (X_{\mathtt{NIKE}}, L_{\mathtt{NIKE}}, R_{\mathtt{NIKE}}))$
  parse $\mathcal{PP}_{\mathtt{PS}} := (\mathsf{crs}, \mathsf{trp}, \mathsf{extr})$
  $\mathcal{PP}_{\mathsf{dkr}} := (\mathcal{PP}, \mathsf{crs})$
  return $\mathcal{PP}_{\mathsf{dkr}}$

$\mathtt{NIKE}_{\mathsf{dkr}}.\mathtt{SharedKey}(\mathcal{PP}_{\mathsf{dkr}}, \mathsf{ID}_1, \mathsf{pk}_1, \mathsf{ID}_2, \mathsf{sk}_2)$
  parse $\mathcal{PP}_{\mathsf{dkr}} =: (\mathcal{PP}, \mathsf{crs})$
  parse $\mathsf{pk}_1 =: (\mathsf{pk}_1', \Pi_1')$
  if $\mathtt{PS.Ver}(\mathsf{crs}, \mathsf{ID}_1, \mathsf{pk}_1', \Pi_1') = 1$
    return $\mathtt{NIKE.SharedKey}(\mathsf{ID}_1, \mathsf{pk}_1', \mathsf{ID}_2, \mathsf{sk}_2)$
  else return $\perp$

$\mathtt{NIKE}_{\mathsf{dkr}}.\mathtt{KeyGen}(\mathcal{PP}_{\mathsf{dkr}}, \mathsf{ID})$
  parse $\mathcal{PP}_{\mathsf{dkr}} =: (\mathcal{PP}, \mathsf{crs})$
  $r \leftarrow \mathcal{R}_{\mathsf{rand}}$
  $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathtt{NIKE.KeyGen}(\mathcal{PP}, \mathsf{ID}; r)$
  $\Pi \leftarrow \mathtt{PS.Prove}(\mathsf{crs}, \mathsf{ID}, \mathsf{pk}, \mathsf{sk}, r)$
  return $((\mathsf{pk}, \Pi), \mathsf{sk})$
```

Fig. 7. A generic transformation from HKR-CKS-heavy security to DKR-CKS-heavy security. $(X_{\mathtt{NIKE}}, L_{\mathtt{NIKE}}, R_{\mathtt{NIKE}})$ is defined as in Remark 2.

4 Security Against Dishonest Key Generation

In this section we want to show how to achieve CKS-heavy security for our scheme allowing dishonest key registrations. That is the adversary is allowed to dishonestly register keys and ask for shared keys where one of the public keys is registered dishonestly.

Due to space limitations we only provide the generic transformation from a HKR-CKS-heavy secure NIKE to a DKR-CKS-heavy secure NIKE. For the proof of security and for a more efficient transformation of an instantiation of our NIKE from Sect. 3 we refer to the full version [25].

Remark 2. Every NIKE induces a SMP as follows. Let $\mathtt{NIKE}$ be a NIKE with public key space $\mathcal{PK}$ and secret key space $\mathcal{SK}$ and randomness space $\mathcal{R}_{\mathsf{rand}}$. Then we define an SMP $\mathtt{SMP}_{\mathtt{NIKE}}$ as follows. On input 1^λ, $\mathtt{SMP}_{\mathtt{NIKE}}.\mathtt{Setup}$ generates $\mathcal{PP} \leftarrow \mathtt{NIKE.Setup}(1^\lambda)$ and sets

$$\begin{aligned} X_{\mathtt{NIKE}} &:= \mathcal{IDS} \times \mathcal{PK}, \\ L_{\mathtt{NIKE}} &:= \{(\mathsf{ID}, \mathsf{pk}) \in X \mid \exists \mathsf{sk}, r \colon (\mathsf{pk}, \mathsf{sk}) = \mathtt{NIKE.KeyGen}(\mathcal{PP}, \mathsf{ID}; r)\} \text{ and} \\ R_{\mathtt{NIKE}} &:= \{(\mathsf{ID}, \mathsf{pk}, \mathsf{sk}, r) \mid (\mathsf{pk}, \mathsf{sk}) = \mathtt{NIKE.KeyGen}(\mathcal{PP}, \mathsf{ID}; r)\}. \end{aligned}$$

Theorem 2. *If $\mathtt{NIKE}$ is a perfectly correct, HKR-CKS-heavy secure NIKE and $\mathtt{PS}$ is an QANIZKPoK for the SMP $\mathtt{SMP}_{\mathtt{NIKE}}$, then the $\mathtt{NIKE}_{\mathsf{dkr}}$ presented in Fig. 7 with algorithms $\mathtt{NIKE}_{\mathsf{dkr}}.\mathtt{Setup}$, $\mathtt{NIKE}_{\mathsf{dkr}}.\mathtt{KeyGen}$, $\mathtt{NIKE}_{\mathsf{dkr}}.\mathtt{SharedKey}$ is perfectly correct and secure in the DKR-CKS-heavy model. More precisely, if $\mathcal{A}$ is an adversary on $\mathtt{NIKE}_{\mathsf{dkr}}$ with running time $t_{\mathcal{A}}$, there exists adversaries $\mathcal{B}, \mathcal{B}_1, \mathcal{B}_2$ with running times $t_{\mathcal{B}} \approx t_{\mathcal{B}_1} \approx t_{\mathcal{B}_2} \approx t_{\mathcal{A}}$ such that*

$$\mathrm{Adv}_{\mathcal{A},\mathtt{NIKE}_{\mathsf{dkr}}}^{\mathsf{dkr-cks-heavy}}(\lambda) \leq \mathrm{Adv}_{\mathcal{B},\mathtt{PS}}^{\mathsf{zk}}(\lambda) + \mathrm{Adv}_{\mathcal{B}_1,\mathtt{PS}}^{\mathsf{extr}}(\lambda) + \mathrm{Adv}_{\mathcal{B}_2,\mathtt{NIKE}}^{\mathsf{hkr-cks-heavy}}(\lambda).$$

```
$\mathrm{Exp}^{\mathsf{uf-cks-heavy}}_{\mathcal{A}=(\mathcal{A}_1,\mathcal{A}_2),n,\mathtt{NIKE}}(\lambda)$:
  $\mathcal{PP} \xleftarrow{\$} \mathtt{NIKE.Setup}(1^\lambda)$
  $\mathsf{ID}_1, ..., \mathsf{ID}_n \xleftarrow{\$} \mathcal{IDS}$ (all disjoint)
  $(\mathsf{pk}_i, \mathsf{sk}_i) \xleftarrow{\$} \mathtt{NIKE.KeyGen}(\mathcal{PP}, \mathsf{ID}_i), i = 1, ..., n$
  $(st, \{i^\star, j^\star\}) \leftarrow \mathcal{A}_1(\mathcal{PP}, \mathsf{ID}_1, \mathsf{pk}_1, ..., \mathsf{ID}_n, \mathsf{pk}_n)$
  $K^\star \leftarrow \mathcal{A}_2(st, (\mathsf{sk}_i)_{i \in [n] \setminus \{i^\star, j^\star\}})$
  if $K^\star = \mathtt{NIKE.SharedKey}(\mathsf{ID}_{i^\star}, \mathsf{pk}_{i^\star}, \mathsf{ID}_{j^\star}, \mathsf{sk}_{j^\star})$
    then output 1
  else output 0
```

Fig. 8. Experiment for *UF - CKS - heavy*$_n$ security of a NIKE scheme NIKE with shared key space $\mathcal{K}$, for any $n \in \mathbb{N}$. The set $C := \{i^\star, j^\star\}$ contains the indices of the two public keys $\mathcal{A}$ wishes to be challenged upon. The set $[n] \setminus C$ contains all indices of the $n-2$ public keys for which $\mathcal{A}$ learns a secret key from the experiment.

5 Optimality of Our Construction

Our NIKE scheme in Sect. 3 does not meet the lower bound regarding tightness proven in [4]. We can circumvent their result since our scheme does not offer a public and efficient algorithm for checking validity of public keys (called PKCheck in [4]): the reduction introduces invalid public keys where the statement is not from the language. It follows from the hardness of the subset membership problem that this is not detectable given *just the public key.*

This immediately raises the question whether, in this new setting without efficient and public PKCheck, we can still obtain a lower bound for the tightness of *HKR - CKS - heavy*-secure NIKE schemes. We answer this question in the affirmative and prove a new lower bound that meets the loss of our reduction in Sect. 3. To present our result, we first give some definitions.

Since *HKR - CKS - heavy* security provides several oracles to the adversary which can be queried in an arbitrary order, a reduction to *HKR - CKS - heavy*-security cannot be formalized as an algorithm in a short and easy way. As done in previous impossibility results before, we thus prove our result w.r.t a weaker security notion that is easier to present. Afterwards, we show that our result carries over to *HKR - CKS - heavy*-security. Our weaker notion is called *UF - CKS - heavy*$_n$[6]. The security experiment is depicted in Fig. 8. Observe that the experiment provides the adversary with all but two secret keys, and thus implicitly with all but one shared key. The adversary chooses which keys he wants to see after obtaining all public keys in the system. The notion is further weakened by letting the number of users in the system be a fixed $n \in \mathbb{N}$ instead of letting the adversary determine it on-the-fly (i.e., via $\mathcal{O}_{\mathsf{regH}}$ queries).

[6] We work with an even weaker notion that [4]. The main difference is that our adversary only has a secret key oracle (from which it can compute shared keys itself), while the adversary in [4] is provided with a shared key oracle.

The next lemma allows us to prove a lower bound w.r.t UF-CKS-heavy_n instead of HKR-CKS-heavy. It will become crucial that the reduction is tight.

Lemma 1 **(HKR-CKS-$\mathit{heavy} \Rightarrow \mathit{UF}$-$\mathit{CKS}$-$\mathit{heavy}_n$).** *For every adversary $\mathcal{A}$ attacking UF-CKS-heavy_n in running time $t_{\mathcal{A}}$ with success probability $\varepsilon_{\mathcal{A}}$, there exists an adversary $\mathcal{B}$ attacking CKS-heavy in running time $t_{\mathcal{B}} \approx t_{\mathcal{A}}$ and success probability $\varepsilon_{\mathcal{B}} = \varepsilon_{\mathcal{A}}$.*

Proof. Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a UF-CKS-heavy_n adversary. We show how to construct a HKR-CKS-heavy adversary $\mathcal{B}$.

On input $\mathcal{PP}$ by the challenger, the adversary $\mathcal{B}$ first generates random, disjoint identities $\mathsf{ID}_1, ..., \mathsf{ID}_n$ and calls the oracle $\mathcal{O}_{\mathsf{regH}}(\mathsf{ID}_i)$ for all $i \in [n]$. $\mathcal{B}$ thus obtains $\mathsf{pk}_1 ..., \mathsf{pk}_n$. Now, $\mathcal{B}$ runs $\mathcal{A}_1(\mathcal{PP}, \mathsf{pk}_1, ..., \mathsf{pk}_n)$ and obtains a state $st_{\mathcal{A}}$ and a set $C := \{i^\star, j^\star\}$. Now, for every $i \in [n] \setminus C$, $\mathcal{B}_1$ queries its oracle $\mathcal{O}_{\mathsf{extr}}(\mathsf{ID}_i)$ which returns a secret key sk_i. Next, $\mathcal{B}_1$ runs $\mathcal{A}_2(st_{\mathcal{A}}, (\mathsf{sk}_i)_{i \in [n] \setminus C})$ and obtains a key $K^\star$. The adversary $\mathcal{B}$ finally queries its test oracle on $(\mathsf{ID}_{i^\star}, \mathsf{ID}_{j^\star})$ which returns a key K. It outputs 0 if $K^\star = K$ and 1 otherwise. As we assume the shared key to be uniquely determined and as further $\mathcal{B}$ only queries $\mathcal{O}_{\mathsf{extr}}$ on identities ID_i with $i \notin C$ we obtain $\varepsilon_{\mathcal{B}} = \varepsilon_{\mathcal{A}}$.

We recall the definition of a non-interactive complexity assumption, taken verbatim from [4], Definitions 4 and 5.

Definition 10 (Non-interactive complexity assumption). *A non-interactive complexity assumption (NICA) $N = (T, V, U)$ consists of three turing machines. The instance generation machine $(c, w) \xleftarrow{\$} T(1^\lambda)$ takes the security parameter as input, and outputs a problem instance c and a witness w. U is a PPT machine, which takes as input c and outputs a candidate solution s. The verification TM V takes as input (c, w) and a candidate solution s. If $V(c, w, s) = 1$, then we say that s is a correct solution to the challenge c.*

Definition 11. *We say that $\mathcal{B}$ (t, ε)-breaks a NICA $N = (T, U, V)$ if $\mathcal{B}$ runs in time $t(\lambda)$ and it holds that*

$$|\Pr[\mathrm{Exp}^{\mathsf{nica}}_{\mathcal{B},N}(1^\lambda) \Rightarrow 1] - \Pr[\mathrm{Exp}^{\mathsf{nica}}_{U,N}(1^\lambda) \Rightarrow 1]| \geq \varepsilon(\lambda),$$

where $\mathrm{Exp}^{\mathsf{nica}}_{\mathcal{B},N}$ is the experiment defined in Fig. 9 and the probability is taken over the random coins consumed by T and the uniformly random choices in the experiment.

Now we are ready to formalize what we mean by a reduction Λ from a NICA to the UF-CKS-heavy_n security of `NIKE`. We closely follow the structure of [4] and similar to [4,12,29,32,35] only consider a certain class of reductions.

Definition 12 (Simple reduction). *We call a TM Λ a $(t_\Lambda, n, \varepsilon_\Lambda, \varepsilon_{\mathcal{A}})$*-reduction *from breaking a NICA $N = (T, U, V)$ to breaking the UF-CKS-heavy_n security of `NIKE`, if Λ turns an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ that runs in time $t_{\mathcal{A}}$ and has advantage $\varepsilon_{\mathcal{A}}$ to break $\mathrm{Exp}^{\mathsf{uf-cks-heavy}}_{\mathcal{A},n,\mathtt{NIKE}}$ (as provided in Fig. 8) into a TM $\mathcal{B}$ that runs in time*

$t_\Lambda + t_\mathcal{A}$ and has advantage ε_Λ to break N (see Definition 11). We call Λ simple*, if Λ has only black-box access to $\mathcal{A}$ and executes $\mathcal{A}$ only once (and in particular without rewinding).*

```
$\mathrm{Exp}^{\mathsf{nica}}_{\mathcal{B},N=(T,U,V)}(\lambda)$:
    $(c, w) \xleftarrow{\$} T(1^\lambda)$
    $s \leftarrow \mathcal{B}(c)$
    return $V(c, w, s)$
```

Fig. 9. Security experiment for a non-interactive complexity assumption (NICA).

In the following we will only consider *simple* reductions. Note that even though this seems to restrict the class of reductions heavily, actually most reductions (including reductions performing hybrid steps) are simple. The security proofs of all existing NIKE schemes [10,14,17] we are aware of[7] are simple reductions.

Since our notion of *UF - CKS - heavy*$_n$-security requires only two rounds of interaction between the adversary and the challenger, we are able to give a very compact formal description of the algorithm $\Lambda := (\Lambda_1, \Lambda_2, \Lambda_3)$ as follows:

- Λ_1 is a probabilistic algorithm that gets as input a (set of) NICA challenge(s) c and outputs public parameters $\mathcal{PP}$, a set of identities and public keys $\mathsf{ID}_1, \mathsf{pk}_1, ..., \mathsf{ID}_n, \mathsf{pk}_n$ and a state st_1.
- Λ_2 is a deterministic algorithmn that receives as input $C \subseteq [n]$ with $|C| = 2$ (else aborts) and st_1 and outputs $(st_2, (sk_i)_{i \in [n] \setminus C})$.
- Λ_3 is a deterministic algorithm that receives as input st_2 and $\tilde{K}$ and outputs an s.

5.1 A Weaker Validity Check

We expand the results from [4] by relaxing the assumptions on the publicly checkable validity of public keys. Recall that [4] requires a method $\mathtt{PKCheck}$ allowing to efficiently verify whether a public key pk was generated by $\mathtt{NIKE.KeyGen}(\mathcal{PP}, \mathsf{ID})$, e.g., whether there exists a secret key sk and random coins r such that $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathtt{NIKE.KeyGen}(\mathcal{PP}, \mathsf{ID}; r)$. We will only require the following notion of weak checkability of public keys. In particular, we only require it to be checkable whether a public key is valid *given a corresponding secret key.*

Definition 13. *Let $\mathtt{NIKE}$ be a NIKE with secret key space $\mathcal{SK}$, identity space $\mathcal{IDS}$ and public key space $\mathcal{PK}$. We say that $\mathtt{NIKE}$ satisfies* weak checkability of public keys*, if there exists a efficiently computable function*

$$\mathtt{wPKCheck}\colon \mathcal{IDS} \times \mathcal{PK} \times \mathcal{SK} \rightarrow \{0, 1\}$$

[7] Remember that we restrict to 2-party key exchange protocols in the setting where a PKI is available.

with the following properties:

$$\textit{For all } (\mathsf{pk},\mathsf{sk}) \leftarrow \textit{NIKE.KeyGen}(\mathcal{PP},\mathsf{ID}) \textit{ we have } \textit{wPKCheck}(\mathsf{ID},\mathsf{pk},\mathsf{sk}) = 1. \quad (1)$$

$$\begin{aligned}\textit{For all } (\mathsf{ID}_1,\mathsf{pk}_1,\mathsf{sk}_1),(\mathsf{ID}_1,\mathsf{pk}_1,\mathsf{sk}_1'),(\mathsf{ID}_2,\mathsf{pk}_2,\mathsf{sk}_2) \textit{ with } \textit{wPKCheck}(\mathsf{ID}_1,\mathsf{pk}_1,\mathsf{sk}_1)\\ = \textit{wPKCheck}(\mathsf{ID}_1,\mathsf{pk}_1,\mathsf{sk}_1') = \textit{wPKCheck}(\mathsf{ID}_2,\mathsf{pk}_2,\mathsf{sk}_2) = 1 \textit{ it holds}\\ \textit{NIKE.SharedKey}(\mathsf{ID}_2,\mathsf{pk}_2,\mathsf{ID}_1,\mathsf{sk}_1) = \textit{NIKE.SharedKey}(\mathsf{ID}_2,\mathsf{pk}_2,\mathsf{ID}_1,\mathsf{sk}_1').\end{aligned} \quad (2)$$

We call a secret key sk valid *for* $(\mathsf{ID},\mathsf{pk})$ *if* $\textit{wPKCheck}(\mathsf{ID},\mathsf{pk},\mathsf{sk}) = 1$*. We further define the* language of valid public keys

$$L^{\text{valid}} := \{(\mathsf{ID},\mathsf{pk}) \mid \exists \mathsf{sk}\colon \textit{wPKCheck}(\mathsf{ID},\mathsf{pk},\mathsf{sk}) = 1\}.$$

Property 2 now implies that any two tuples $(\mathsf{ID}_1,\mathsf{pk}_1),(\mathsf{ID}_2,\mathsf{pk}_2) \in L^{\text{valid}}$ *lead to a unique shared key independently of which valid secret key is employed to compute the shared key.*

Remark 3. Note that a NIKE for which it can be efficiently verified whether a pair $(\mathsf{pk},\mathsf{sk})$ lies in the image of $\mathtt{NIKE.KeyGen}(\mathcal{PP},\mathsf{ID})$ in particular satisfies weak checkability of public keys with

$$\mathtt{wPKCheck}(\mathsf{ID},\mathsf{pk},\mathsf{sk}) = \begin{cases}1 & \text{if } \exists r : (\mathsf{pk},\mathsf{sk}) = \mathtt{NIKE.KeyGen}(\mathcal{PP},\mathsf{ID};r)\\ 0 & \text{else}\end{cases}.$$

5.2 A Lower Bound on Tightness

In this section we show that if a NIKE $\mathtt{NIKE}$ satisfies weak checkable uniqueness, then any simple reduction from a NICA to the $\mathit{UF}\text{-}\mathit{CKS}\text{-}\mathit{heavy}_n$-security of $\mathtt{NIKE}$ it has to inherently lose a factor of $n/2$ in reduction, where n is the number of public keys. Further, we show that the NIKE $\mathtt{NIKE}$ presented in Fig. 5 satisfies weak checkability of public keys. Note that by definition any NIKE supporting weak checkability of public keys is *perfectly correct*, that is for all $(\mathsf{ID}_i,\mathsf{pk}_i,\mathsf{sk}_i) \xleftarrow{\$} \mathtt{NIKE.KeyGen}(\mathcal{PP},\mathsf{ID}_i), i \in \{1,2\}$, we have

$$\mathtt{NIKE.SharedKey}(\mathsf{ID}_1,\mathsf{pk}_1,\mathsf{ID}_2,\mathsf{sk}_2) = \mathtt{NIKE.SharedKey}(\mathsf{ID}_2,\mathsf{pk}_2,\mathsf{ID}_1,\mathsf{sk}_1).$$

Theorem 3. *Let $N = (T,U,V)$ be a non-interactive complexity assumption and $n \in \mathrm{poly}(\lambda)$. Let* $\mathtt{NIKE}$ *be a $\mathit{UF}\text{-}\mathit{CKS}\text{-}\mathit{heavy}_n$ secure NIKE with shared key space $\mathcal{K}$, public key space $\mathcal{PK}$ and secret key space $\mathcal{SK}$ which satisfies weak checkability of public keys via algorithm* $\mathtt{wPKCheck}$*. Let further evaluating* $\mathtt{wPKCheck}$ *require time t_{wPKCheck}. Then any reduction $\Lambda = (\Lambda_1,\Lambda_2,\Lambda_3)$ from N to* $\mathtt{NIKE}$ *has to lose a factor $n/2$ assuming N is hard. More formally, for any simple $(t_\Lambda, n, \varepsilon_\Lambda, 1)$-reduction from breaking the assumption N to breaking the $\mathit{UF}\text{-}\mathit{CKS}\text{-}\mathit{heavy}_n$-security of* $\mathtt{NIKE}$*, there exists an adversary $\mathcal{B}$ breaking N in running time*

$$t_{\mathcal{B}} \le \frac{n(n-1)}{2} t_\Lambda + \frac{n(n-1)(n-2)}{2} t_{\text{wPKCheck}}$$

with success probability

$$\varepsilon_{\mathcal{B}} \geq \varepsilon_{\Lambda} - \frac{2}{n}.$$

Remark 4. We have $\varepsilon_{\mathcal{A}} = 1$ and $\varepsilon_{\mathcal{B}} = \eta(\lambda)$ for a negligible function η (as N is assumed to be hard). We can thus transform the last equation into $\varepsilon_{\Lambda} \leq \frac{2}{n}\varepsilon_{\mathcal{A}} + \eta(\lambda)$. This implies the claimed reduction loss of $n/2$.

Proof. We follow the proof structure of [4,29,35].

The hypothetical adversary. In the following we describe a hypothetical adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. Note that this adversary might not be efficient, but in order to prove the reduction loss of $n/2$ we show how to simulate it efficiently.

$\mathcal{A}_1(\mathcal{PP}, \mathsf{ID}_1, \mathsf{pk}_1, \ldots, \mathsf{ID}_n, \mathsf{pk}_n)$ chooses $C := \{i^\star, j^\star\} \subseteq [n]$ with $|C| = 2$ uniformly at random. It outputs (st, C), where $st = (\mathcal{PP}, \mathsf{ID}_1, \mathsf{pk}_1, \ldots, \mathsf{ID}_n, \mathsf{pk}_n, C)$.

$\mathcal{A}_2(st, (\mathsf{sk}_i)_{i \in [n] \setminus C})$ checks whether $\texttt{wPKCheck}(\mathsf{ID}_i, \mathsf{pk}_i, \mathsf{sk}_i) = 1$ for all $i \in [n] \setminus C$ and whether $(\mathsf{ID}_i, \mathsf{pk}_i) \in L^{\text{valid}}$ for both $i \in C$. If this is the case $\mathcal{A}_2$ computes a secret key $\mathsf{sk}_{j^\star}$ s.t. $\texttt{wPKCheck}(\mathsf{ID}_{j^\star}, \mathsf{pk}_{j^\star}, \mathsf{sk}_{j^\star}) = 1$ and outputs $K^\star = \texttt{NIKE.SharedKey}(\mathsf{ID}_{i^\star}, \mathsf{pk}_{i^\star}, \mathsf{ID}_{j^\star}, \mathsf{sk}_{j^\star})$. Otherwise $\mathcal{A}_2$ outputs $\perp$.

As we have $(\mathsf{ID}, \mathsf{pk}, \mathsf{sk}) \in \mathcal{R}^{\text{unique}}$ for all $(\mathsf{pk}, \mathsf{sk}) \leftarrow \texttt{NIKE.KeyGen}(\mathcal{PP}, \mathsf{ID})$ and further $\texttt{NIKE.SharedKey}$ returns a unique key for all tuples passing $\texttt{wPKCheck}$, due to property 2 of Definition 13 the hypothetical adversary always wins in the *UF - CKS - heavy*$_n$ experiment.

We now describe an adversary $\mathcal{B}$ attempting to break $N = (T, U, V)$. The strategy is to run the reduction $\Lambda = (\Lambda_1, \Lambda_2, \Lambda_3)$ simulating $\mathcal{A}$ efficiently. Let c be the input of $\mathcal{B}$, where $(c, w) \leftarrow T(1^\lambda)$. Let $SK[\,], SK^\star[\,]$ be arrays of n entries initialized by $\emptyset$ and maintained throughout the reduction by $\mathcal{B}$.

1. The adversary $\mathcal{B}$ runs $(st_1, \mathcal{PP}, \mathsf{ID}_1, \mathsf{pk}_1, \ldots, \mathsf{ID}_n, \mathsf{pk}_n) \leftarrow \Lambda_1(c)$.
2. The adversary $\mathcal{B}$ samples $\{i^\star, j^\star\} = C^\star \subset [n]$ with $|C^\star| = 2$ uniformly at random.
3. For each $C \subset [n]$ with $|C| = 2$ the adversary $\mathcal{B}$ runs the reduction $\Lambda_2(st_1, C)$. Let $(st_2^C, (\mathsf{sk}_i^C)_{i \in [n] \setminus C})$ denote the output of the respective execution. Whenever $\texttt{wPKCheck}(\mathsf{ID}_i, \mathsf{pk}_i, \mathsf{sk}_i^C) = 1$ for an $i \in [n] \setminus C$ the adversary sets $SK[i] = \mathsf{sk}_i^C$. If $C = C^\star$, $\mathcal{B}$ additionally sets $SK^\star[i] = \mathsf{sk}_i^{C^\star}$
4. If there exists an $i \in [n] \setminus C^\star$ with $SK^\star[i] = \emptyset$ (i.e. $\texttt{wPKCheck}(\mathsf{ID}_i, \mathsf{pk}_i, \mathsf{sk}_i^{C^\star}) = 0$) or there exists a $i \in C^\star$ such that $SK[i] = \emptyset$ (i.e. $\texttt{wPKCheck}(\mathsf{ID}_i, \mathsf{pk}_i, \mathsf{sk}_i^C) = 0$ for all $C \subseteq [n]$ with $|C| = 2$) then $\mathcal{B}$ sets $K^\star = \perp$. Otherwise $\mathcal{B}$ computes $K^\star = \texttt{NIKE.SharedKey}(\mathsf{ID}_{i^\star}, \mathsf{pk}_{i^\star}, \mathsf{ID}_{j^\star}, SK[j^\star])$.
5. Finally, the adversary $\mathcal{B}$ outputs $s \overset{\$}{\leftarrow} \Lambda_3(st_2^{C^\star}, C^\star, K^\star)$.

Efficiency of $\mathcal{B}$. In the third step Λ_2 has to be executed $\binom{n}{2} = \frac{n(n-1)}{2}$ times. Each time the validity check has to be performed $n - 2$ times. For the running time of $\mathcal{B}$ it thus holds

$$t_{\mathcal{B}} \leq \frac{n(n-1)}{2} t_{\Lambda} + \frac{n(n-1)(n-2)}{2} t_{\text{wPKCheck}}.$$

Success probability of $\mathcal{B}$. Let $C^\star = \{i^\star, j^\star\}$ as before. Consider the following two events:

$$\begin{aligned} \mathsf{check\text{-}fails}: &\quad \exists i \in [n] \setminus C^\star \text{ such that } SK^\star[i] = \emptyset \\ \mathsf{pk\text{-}valid}: &\quad \forall i \in \mathcal{C}^\star \text{ it holds that } SK[i] \neq \emptyset \end{aligned}$$

We first want to show that in the case of $\mathsf{check\text{-}fails} \vee \mathsf{pk\text{-}valid}$, $\mathcal{B}$ simulates the hypothetical adversary $\mathcal{A}$ perfectly. If $\mathsf{check\text{-}fails}$ occurs, then $\mathcal{B}$ returns $\perp$. The hypothetical adversary would have returned $\perp$ as well because in this case it holds $\mathtt{wPKCheck}(\mathsf{ID}_i, \mathsf{pk}_i, \mathsf{sk}_i^{C^\star}) = 0$ for an $i \in [n] \setminus C^\star$. If $\mathsf{pk\text{-}valid}$ occurs, we have $(\mathsf{ID}_i, \mathsf{pk}_i) \in L^{\text{valid}}$ for all $i \in [n]$ (as in this case for each $i \in [n]$ there exists a set $C \subset [n]$ such that the reduction Λ_2 provided an sk_i^C with $\mathtt{wPKCheck}(\mathsf{ID}_i, \mathsf{pk}_i, \mathsf{sk}_i^C) = 1$ at some point). In this case the shared key $K^\star$ is unique by property 1 in Definition 13 and can be computed by $\mathcal{B}$ with the secret key $\mathcal{SK}[j^\star]$.

We summarize all other possible cases in the event

$$\mathsf{bad} = \neg\mathsf{check\text{-}fails} \wedge \neg\mathsf{pk\text{-}valid},$$

which is well-defined, as Λ_2 is deterministic.

We now bound the probability that bad happens. For this, we observe that $\neg\mathsf{pk\text{-}valid}$ can only occur if the event $E := (\exists i \in [n] \text{ s.t. } \mathcal{SK}[i] = \emptyset)$ occurs. As $C^\star$ is chosen uniformly at random and the view of Λ_2 is independent of $C^\star$, we have $i \in [n] \setminus C^\star$ with probability $1 - 2/n$. In this case $\mathsf{check\text{-}fails}$ occurs and thus $\Pr[\mathsf{check\text{-}fails}|\ E] \geq 1 - 2/n$. Now since $\neg\mathsf{pk\text{-}valid} \Rightarrow E$ it holds that $\Pr[\neg\mathsf{check\text{-}fails} \wedge \neg\mathsf{pk\text{-}valid}] \leq \Pr[\neg\mathsf{check\text{-}fails} \wedge E] = \Pr[\neg\mathsf{check\text{-}fails}|E] \cdot \Pr[E] \leq \Pr[\neg\mathsf{check\text{-}fails}|E] = 1 - \Pr[\mathsf{check\text{-}fails}|E] \leq 2/n$. We thus obtain

$$\Pr[\mathsf{bad}] \leq 2/n.$$

Let $\varepsilon_{\mathcal{B}}\big|_{\neg\mathsf{bad}}$ denote the probability of $\mathcal{B}$ to win under the condition that bad does not occur and $\varepsilon_{\Lambda}\big|_{\neg\mathsf{bad}}$ accordingly. We have

$$|\varepsilon_{\mathcal{B}} - \varepsilon_{\Lambda}| \leq \left|\varepsilon_{\mathcal{B}}\big|_{\neg\mathsf{bad}} - \varepsilon_{\Lambda}\big|_{\neg\mathsf{bad}}\right| + \Pr[\mathsf{bad}] = \Pr[\mathsf{bad}] \leq \frac{2}{n}.$$

Remark 5. As shown in [4] it is straightforward to generalize Theorem 3 to simple $(t_\Lambda, n, \varepsilon_\Lambda, \varepsilon_{\mathcal{A}})$-reductions for general $\varepsilon_{\mathcal{A}}$ by letting the hypothetical adversary (and $\mathcal{B}$ respectively) toss a coin and only return $K^\star$ with probability $\varepsilon_{\mathcal{A}}$.

Remark 6. While Theorem 3 establishes the impossibility of tight security reductions for a large class of NIKE schemes, it thereby also gives a hint about how a tight NIKE scheme has to be constructed. Namely, such a scheme has to violate the assumptions made in the theorem such as the existence of an efficient $\mathtt{PKCheck}$ that, given the secret key, decides uniqueness of shared keys. More detailed, a tight NIKE scheme needs to allow a reduction to indistinguishably switch public keys to invalid (in fact, even *tightly* switch many of them in one step), such that invalid public keys admit *many* secret keys that lead do *different* shared keys. It is an interesting open question how to construct such a scheme.

5.3 Weak Checkable Uniqueness of Our NIKE

Lemma 2. *If instantiated with a hash proof system* $\mathtt{HPS}$ *where membership in* $\mathcal{HSK}$ *is efficient checkable for all sets of secret keys in the image of* $\mathtt{HPS.Setup}$*, the NIKE* $\mathtt{NIKE}$ *presented in Fig. 5 complies with weak checkability of public keys.*

Proof. Let $\mathcal{PP} := ((X, L, R), \mathcal{HSK}, \mathcal{H}, \alpha, F) \xleftarrow{\$} \mathtt{NIKE.Setup}(1^\lambda)$. We define

$$\mathtt{wPKCheck}(\mathsf{ID}, (\mathsf{hpk}, x), (\mathsf{hsk}, x, w)) := \begin{cases} 1 & \text{if } \mathsf{hsk} \in \mathcal{HSK} \wedge \alpha(\mathsf{hsk}) = \mathsf{hpk} \\ & \quad \wedge (x, w) \in R \\ 0 & \text{else} \end{cases} .$$

We have to show that $\mathtt{wPKCheck}$ is efficiently computable and further that $\mathtt{wPKCheck}$ meets properties 1 and 2 in Definition 13. By prerequisites we have that membership in $\mathcal{HSK}$ is efficiently checkable. Further, by definition of a hash proof system the map α and the relation R are efficiently computable. Property 1 follows straightforward from the definition of $\mathtt{wPKCheck}$. Note that actually we have equality, that is

$$\mathtt{wPKCheck}(\mathsf{ID}, \mathsf{pk}, \mathsf{sk}) = 1 \Leftrightarrow \exists r\colon (\mathsf{pk}, \mathsf{sk}) = \mathtt{NIKE.KeyGen}(\mathcal{PP}, \mathsf{ID}; r).$$

It remains to prove property 2: for all $(\mathsf{ID}_1, \mathsf{pk}_1, \mathsf{sk}_1)$, $(\mathsf{ID}_1, \mathsf{pk}_1, \mathsf{sk}_1')$, $(\mathsf{ID}_2, \mathsf{pk}_2, \mathsf{sk}_2)$ that all pass $\mathtt{wPKCheck}$ we have

$$\mathtt{NIKE.SharedKey}(\mathsf{ID}_2, \mathsf{pk}_2, \mathsf{ID}_1, \mathsf{sk}_1) = \mathtt{NIKE.SharedKey}(\mathsf{ID}_2, \mathsf{pk}_2, \mathsf{ID}_1, \mathsf{sk}_1').$$

Let in the following $\mathsf{pk}_1 =: (\mathsf{hpk}_1, x_1), \mathsf{pk}_2 =: (\mathsf{hpk}_2, x_2), \mathsf{sk}_1 =: (\mathsf{hsk}_1, w_1)$, $\mathsf{sk}_1' =: (\mathsf{hsk}_1', w_1')$ and $\mathsf{sk}_2 =: (\mathsf{hsk}_2, w_2)$. By the properties of the hash proof system we have that for $\mathsf{hsk}_1, \mathsf{hsk}_1' \in \mathcal{HSK}$ with $\alpha(\mathsf{hsk}_1) = \alpha(\mathsf{hsk}_1') = \mathsf{hpk}_1$ and $x_2 \in L$ it holds

$$H_{\mathsf{hsk}_1}(x_2) = F(x_2, w_2, \mathsf{hpk}) = H_{\mathsf{hsk}_1'}(x_2)$$

and for w_1' with $(x_1, w_1') \in \mathcal{R}$ it holds

$$F(x_1, w_1, \mathsf{hpk}_2) = H_{\mathsf{hsk}_2}(x_1) = F(x_1, w_1', \mathsf{hpk}_2).$$

This yields

$$\begin{aligned} \mathtt{NIKE.SharedKey}(\mathsf{ID}_2, \mathsf{pk}_2, \mathsf{ID}_1, \mathsf{sk}_1) &= H_{\mathsf{hsk}_1}(x_2) \oplus F(x_1, w_1, \mathsf{hpk}_2) \\ &= H_{\mathsf{hsk}_1'}(x_2) \oplus F(x_1', w_1', \mathsf{hpk}_2) \\ &= \mathtt{NIKE.SharedKey}(\mathsf{ID}_2, \mathsf{pk}_2, \mathsf{ID}_1, \mathsf{sk}_1'). \end{aligned}$$

Corollary 1 (Informal). *The security reduction in the proof of Theorem 1 is optimal regarding tightness among all simple reductions.*

Proof. Theorem 3 shows that simple security reductions for a NIKE admitting a weak PKCheck encounter a loss of at least $n/2$. Lemma 2 proves that our NIKE admits such a weak PKCheck and thus from Theorem 3 it follows that

$\mathit{UF}\text{-}\mathit{CKS}\text{-}\mathit{heavy}_n$-security of our NIKE can only be shown by a simple reduction if the reduction loses at least a factor of $n/2$. Now Lemma 1 shows that a $\mathit{UF}\text{-}\mathit{CKS}\text{-}\mathit{heavy}_n$ adversary tightly implies a $\mathit{HKR}\text{-}\mathit{CKS}\text{-}\mathit{heavy}$ adversary. Thus, any reduction with loss M from a NICA to $\mathit{HKR}\text{-}\mathit{CKS}\text{-}\mathit{heavy}$ security would imply a reduction with loss M to $\mathit{UF}\text{-}\mathit{CKS}\text{-}\mathit{heavy}_n$ security. It follows that $M \geq n/2$.

Remark 7. Since DKR-CKS-heavy security also tightly implies $\mathit{UF}\text{-}\mathit{CKS}\text{-}\mathit{heavy}_n$ security, our result carries over to DKR-CKS-heavy secure NIKE schemes that comply with weak checkable uniqueness.

Acknowledgements. We would like to thank Kenny Paterson for collaboration at early stages of this work. Also, we would like to thank the anonymous reviewers of Crypto 2018 for their helpful comments.

References

1. Abe, M., Hofheinz, D., Nishimaki, R., Ohkubo, M., Pan, J.: Compact structure-preserving signatures with almost tight security. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 548–580. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63715-0_19
2. Abe, M., David, B., Kohlweiss, M., Nishimaki, R., Ohkubo, M.: Tagged one-time signatures: tight security and optimal tag size. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 312–331. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36362-7_20
3. Attrapadung, N., Hanaoka, G., Yamada, S.: A framework for identity-based encryption with almost tight security. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 521–549. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48797-6_22
4. Bader, C., Jager, T., Li, Y., Schäge, S.: On the impossibility of tight cryptographic reductions. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 273–304. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_10
5. Bader, C., Hofheinz, D., Jager, T., Kiltz, E., Li, Y.: Tightly-secure authenticated key exchange. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part I. LNCS, vol. 9014, pp. 629–658. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46494-6_26
6. Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: security proofs and improvements. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 259–274. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-45539-6_18
7. Blazy, O., Kiltz, E., Pan, J.: (Hierarchical) identity-based encryption from affine message authentication. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 408–425. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44371-2_23
8. Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications (Extended Abstract). In: Proceedings of 20th ACM STOC, pp. 103–112. ACM Press, May 1988

9. Boneh, D., Venkatesan, R.: Breaking RSA may not be equivalent to factoring. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 59–71. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0054117
10. Cash, D., Kiltz, E., Shoup, V.: The Twin Diffie-Hellman problem and applications. In: Smart, N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 127–145. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78967-3_8
11. Chen, J., Wee, H.: Fully, (almost) tightly secure IBE and dual system groups. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 435–460. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_25
12. Coron, J.-S.: Optimal security proofs for PSS and other signature schemes. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 272–287. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_18
13. Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_4
14. Diffie, W., Hellman, M.E.: New directions in cryptography. IEEE Trans. Inf. Theory **22**(6), 644–654 (1976)
15. Escala, A., Herold, G., Kiltz, E., Ràfols, C., Villar, J.: An algebraic framework for Diffie-Hellman assumptions. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 129–147. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_8
16. Fleischhacker, N., Jager, T., Schröder, D.: On tight security proofs for Schnorr signatures. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014, Part I. LNCS, vol. 8873, pp. 512–531. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45611-8_27
17. Freire, E.S.V., Hofheinz, D., Kiltz, E., Paterson, K.G.: Non-interactive key exchange. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 254–271. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36362-7_17
18. Garg, S., Bhaskar, R., Lokam, S.V.: Improved bounds on security reductions for discrete log based signatures. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 93–107. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_6
19. Gay, R., Hofheinz, D., Kohl, L.: Kurosawa-Desmedt meets tight security. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 133–160. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63697-9_5
20. Gay, R., Hofheinz, D., Kiltz, E., Wee, H.: Tightly CCA-secure encryption without pairings. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016, Part I. LNCS, vol. 9665, pp. 1–27. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49890-3_1
21. Gennaro, R., Lindell, Y.: A framework for password-based authenticated key exchange. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 524–543. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-39200-9_33
22. Goldwasser, S., Micali, S., Rivest, R.L.: A Digital signature scheme secure against adaptive chosen-message attacks. SIAM J. Comput. **17**(2), 281–308 (1988)
23. Gong, J., Chen, J., Dong, X., Cao, Z., Tang, S.: Extended nested dual system groups, revisited. In: Cheng, C.-M., Chung, K.-M., Persiano, G., Yang, B.-Y. (eds.) PKC 2016, Part I. LNCS, vol. 9614, pp. 133–163. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49384-7_6

24. Guo, F., Chen, R., Susilo, W., Lai, J., Yang, G., Mu, Y.: Optimal security reductions for unique signatures: bypassing impossibilities with a counterexample. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 517–547. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63715-0_18
25. Hesse, J., Hofheinz, D., Kohl, L.: On tightly secure non-interactive key exchange. In: IACR Cryptology ePrint Archive 2018, p. 237 (2018). http://eprint.iacr.org/2018/237
26. Hofheinz, D.: Adaptive partitioning. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part II. LNCS, vol. 10212, pp. 489–518. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56617-7_17
27. Hofheinz, D.: Algebraic partitioning: fully compact and (almost) tightly secure cryptography. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016, Part I. LNCS, vol. 9562, pp. 251–281. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49096-9_11
28. Hofheinz, D., Jager, T.: Tightly secure signatures and public-key encryption. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 590–607. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_35
29. Hofheinz, D., Jager, T., Knapp, E.: Waters signatures with optimal security reduction. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 66–83. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-30057-8_5
30. Hofheinz, D., Kiltz, E.: Secure hybrid encryption from weakened key encapsulation. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 553–571. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74143-5_31
31. Hofheinz, D., Koch, J., Striecks, C.: Identity-based encryption with (almost) tight security in the multi-instance, multi-ciphertext setting. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 799–822. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46447-2_36
32. Kakvi, S.A., Kiltz, E.: Optimal security proofs for full domain hash, revisited. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 537–553. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_32
33. Katz, J., Ostrovsky, R., Yung, M.: Efficient password-authenticated key exchange using human-memorable passwords. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 475–494. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44987-6_29
34. Katz, J., Wang, N.: Efficiency improvements for signature schemes with tight security reductions. In: Jajodia, S., Atluri, V., Jaeger, T. (eds.) ACM CCS 03, pp. 155–164. ACM Press, October 2003
35. Lewko, A., Waters, B.: Why proving HIBE systems secure is difficult. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 58–76. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_4
36. Libert, B., Peters, T., Joye, M., Yung, M.: Compactly hiding linear spans. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 681–707. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48797-6_28
37. Libert, B., Joye, M., Yung, M., Peters, T.: Concise multi-challenge CCA-secure encryption and signatures with almost tight security. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014, Part II. LNCS, vol. 8874, pp. 1–21. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45608-8_1

38. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: Proceedings of 22nd ACM STOC, pp. 427–437. ACM Press, May 1990
39. Seurin, Y.: On the exact security of Schnorr-type signatures in the random oracle model. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 554–571. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_33