

Multi-Party Key Exchange for Unbounded Parties from Indistinguishability Obfuscation

Dakshita Khurana^{1*}, Vanishree Rao^{2**}, and Amit Sahai^{1*}

¹ Dept of Computer Science, UCLA and Center for Encrypted Functionalities, USA.

{dakshita,sahai}@cs.ucla.edu

² PARC, a Xerox Company.

Vanishree.Rao@parc.com

Abstract. Existing protocols for non-interactive multi-party key exchange either (1) support a bounded number of users, (2) require a trusted setup, or (3) rely on knowledge-type assumptions.

We construct the first non-interactive key exchange protocols which support an unbounded number of parties and have a security proof that does not rely on knowledge assumptions. Our non-interactive key-exchange protocol does not require a trusted setup and extends easily to the identity-based setting. Our protocols suffer only a polynomial loss to the underlying hardness assumptions.

1 Introduction

Non-interactive key exchange (NIKE) enables a group of parties to derive a shared secret key without any interaction. In a NIKE protocol, all parties simultaneously broadcast a message to all other parties. After this broadcast phase, each party should be able to locally compute a shared secret key for any group of which he is a member. All members of a group should generate an identical shared key, and the shared key for a group should look random to a non-member.

This notion was introduced by Diffie and Hellman [16], who also gave a protocol for non-interactive key exchange in the two-party setting. More than two decades later, Joux [32] constructed the first non-interactive key exchange protocol for three parties. Given a set of N parties (where N is a polynomial in the security parameter), Boneh and Silverberg [4] obtained a multiparty NIKE protocol based on multilinear maps. The recent candidates for multilinear maps

* Research supported in part from a DARPA/ONR PROCEED award, NSF grants 1228984, 1136174, 1118096, and 1065276, a Xerox Faculty Research Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant. This material is based upon work supported by the Defense Advanced Research Projects Agency through the U.S. Office of Naval Research under Contract N00014-11-1-0389. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

** Work done while studying at UCLA.

given by [22,14,27,15] can be used to instantiate the scheme of [4], assuming a trusted setup. After the advent of candidate constructions for indistinguishability obfuscation (iO) starting with the result of Garg et. al. [23], Boneh and Zhandry [6] demonstrated how to obtain static secure NIKE and ID-NIKE based on indistinguishability obfuscation, without relying on a trusted setup.

However, all these constructions require an a-priori bound on the number of parties. The only known protocols which can possibly handle an unbounded number of parties are the ones by Ananth et. al. and Abusalah et. al. [2,1], but their solutions rely on *differing-inputs obfuscation* (diO)[3,2,7]. Unfortunately, diO is a knowledge-type assumption, and recent work [9,25,31] demonstrates that it may suffer from implausibility results. In this paper, we address the following question:

Can we obtain NIKE supporting an a-priori unbounded number of parties and not requiring any setup, based on indistinguishability obfuscation?

We give a positive answer to this question, by demonstrating non-interactive key exchange protocols that achieve static security and support an a-priori unbounded number of parties, based on indistinguishability obfuscation.

1.1 Our Contributions

We consider a setting where an a-priori unbounded number of parties may broadcast or publish messages, such that later any party can derive a shared secret key for a group of which it is a member. In our setting, parameters do not grow with the number of parties. Our results can be summarized as follows.

Theorem 1. *Assuming indistinguishability obfuscation and fully homomorphic encryption, it is possible to obtain static-secure non-interactive multi-party key exchange for an a-priori unbounded number of parties, without any setup.*

Theorem 2. *Assuming indistinguishability obfuscation and fully homomorphic encryption, it is possible to obtain static-secure identity-based non-interactive multi-party key exchange for an a-priori unbounded number of parties.*

Fully homomorphic encryption was first constructed by Gentry [26] and subsequently Brakerski and Vaikuntanathan [10] constructed it under the learning with errors assumption. Alternatively, it can be constructed based on sub-exponentially secure iO for circuits and sub-exponential one-way functions [12].

1.2 Technical Overview

Bottlenecks in known constructions. Our starting point the static-secure NIKE protocol of Boneh and Zhandry [6] based on indistinguishability obfuscation, in the simplest case where parties have access to a trusted setup. The adversary fixes a set of parties to corrupt, independent of the system parameters. Then the setup generates public parameters, and each party broadcasts its public values.

We require that the shared group key for all the honest parties should look indistinguishable from random, from the point of view of the adversary.

The basic Boneh-Zhandry construction uses a trusted setup to generate an obfuscated program with a secret PRF key. Parties pick a secret value uniformly at random, and publish the output of a length-doubling PRG applied to this value, as their public value. To derive the shared secret key for a group of users, a member of the group inputs public values of all users in the group (including himself) according to some fixed ordering, as well as his own secret value into this obfuscated program. The program checks if the PRG applied to the secret value corresponds to one of the public values in the group, and if the check passes, outputs a shared secret key by applying a PRF to the public value of all parties in the group.

To prove security, we note that the adversary never corrupts any party in the challenge set, so we never need to reveal the secret value for any party in this set. Thus, (by security of the PRG) we can set the public values of parties in the challenge set, to uniformly random values in the co-domain of the PRG. Then with overwhelming probability, there exists no pre-image for any of the public values in the challenge set, and therefore there exists no secret value for which the PRG check in the program would go through. This allows us to puncture the program at the challenge set, and then replace the shared secret key for this set with random. However, in this construction it is necessary to set an a-priori bound on the number of participants, since the size of inputs to the setup circuit must be bounded.

The construction of Ananth et. al. [2] works without this a-priori bound, by making use of differing-inputs obfuscation and collision resistant hashing. To obtain a shared group key, parties first hash down the group public values and then generate a short proof of membership in the group. The program takes the hashed value and proof, and outputs a shared secret key for the group if and only if the proof verifies. However, because the hash is compressing, there exist collisions and thus there exist false proofs of membership. The only guarantee is that these proofs are ‘hard’ to find in a computational sense. Unfortunately, proving security in such a situation requires the use of differing-inputs obfuscation, since we must argue that any non-member which distinguishes from random the shared key of a group, could actually have generated false proofs. Such an argument inherently involves an extractable assumption.

First attempt. Now, it seems that we may benefit from using an iO-friendly tool for hashing, such that the proof of membership is unconditionally sound for a select piece of the input to the hash. Moreover, the description of the hash should computationally hide which part of the input it is sound for. Then, like in the previous proof, we can set the public values of parties in the challenge set to uniformly random values in the co-domain of the PRG, and try to puncture the program at each position one by one. This resembles the “selective enforcement” techniques of Koppula et. al. [34] who construct accumulators (from iO) as objects enabling bounded commitments to an unbounded storage, which are unconditionally binding for a select piece of the storage.

At this point, it may seem that we can use the accumulator hash function and we should be done. Parties can hash down an unbounded number of public values to a bounded size input using the accumulator, and generate a short proof of correctness. On input the hashed value and proof, the program can be set to output a shared key if the proof verifies. To prove security, we could begin by generating the public values for parties in the challenge set, uniformly at random in the co-domain of the PRG. Then, it should be possible to make the accumulator binding at the first index, and consequently puncture out the first index from the program. Then we can continue across indices and finally puncture out the hash value at the challenge set for all indices.

Even though this proof seems straightforward, the afore-mentioned approach fails. This is because an accumulator can be made unconditionally sound at a particular index, *only conditioned on the previous indices being equal to an a-priori fixed sequence*. In the setting of obfuscating unbounded-storage Turing Machines, for which such accumulators were first introduced [34], there was indeed a well-defined “correct” path that was generated by the machine itself, and consequently there was a way to enforce correct behaviour on all previous indices. However, in our setting the adversary is allowed to hash completely arbitrary values for all indices. Very roughly, we require a tool that enables selective enforcing even when the adversary is allowed to behave arbitrarily on all other indices.

Our solution. At this point, we require a hash function which can enforce soundness at hidden indices, while allowing arbitrary behaviour on all other indices. Such a hash function was introduced recently in the beautiful work of Hubacek and Wichs [30], in the context of studying the communication complexity of secure function evaluation with long outputs. They call it a *somewhere statistically binding* (SSB) hash and give a construction based on fully homomorphic encryption. An SSB hash can be used like other other hash functions, to hash an a-priori unbounded number of chunks of input, each of bounded size, onto a bounded space. Moreover, this hash can operate in various *modes*, where each mode is statistically binding on some fixed index i determined at setup; yet, the description of the hash function computationally hides this index.

Equipped with this tool, it is possible to argue security via a selective enforcing hybrid argument. As usual, we begin by generating the public values of all parties in the challenge set, uniformly at random in the co-domain of the PRG. With overwhelming probability, this ensures that there exist no secret values that could generate the public values in the challenge set. Now, we zoom into each index (of the hash) one by one, and make the hash function statistically binding at that particular index. Specifically, we know that a specific output of the hash h^* can only be achieved via a single fixed public value at the enforcing index (say i). Moreover, with overwhelming probability, this public value lies outside the range of the PRG, and thus there exist no false proofs for hash value h^* at the enforcing index i .

This allows us to alter the obfuscated circuit to always ignore the value h^* at index i . Once we have changed the program, we generate the hash function

to be statistically binding at index $(i + 1)$, and repeat the argument. Note that once we are at this next index, there may exist false proofs for previous indices – however, at this point, we have already programmed the obfuscation to eliminate the value h^* for all previous indices.

In the identity-based NIKE setting, we generate secret keys for identities as PRF outputs on the identity (in a manner similar to [6]). In addition to using the enforce-and-move technique detailed above, we need to avoid simultaneously programming in an unbounded number of public values. We handle this using Sahai-Waters [37] punctured programming techniques (using PRGs) to puncture and then *un-puncture* the PRF keys before moving on to the next value.

1.3 Other Related Work

Cash, Kiltz, and Shoup [13] and Freire, Hofheinz, Kiltz, and Paterson [19] formalized various security models in the two-party NIKE setting. Bones and Zhandry [6] first resolved NIKE for bounded $N > 3$ parties without relying on a trusted setup, assuming indistinguishability obfuscation. However, their security proofs worked only for the static and semi-static scenarios. Hofheinz et. al. [29] realized adaptive secure bounded N -party NIKE in the random oracle model without setup, and Rao [36] realized bounded N -party NIKE with adaptive security and without setup based on assumptions over multilinear maps. A recent independent work of Yamakawa et. al. [39] gives multilinear maps where the multilinearity levels need not be bounded during setup, and the size of the representations of elements is independent of the level of multi-linearity. In the same paper, these maps are used to construct multiparty NIKE with unbounded parties, however, requiring a trusted setup. Furthermore, their scheme does not seem to extend directly to the identity-based NIKE setting.

In the identity-based NIKE setting, there is a trusted master party that generates secret values for identities using a master secret key. This (seemingly weaker) setting has been extensively studied both in the standard and the random oracle models, and under various static and adaptive notions of security [38,18,35,20], but again for an a-priori bounded number of parties.

Zhandry [40] uses somewhere statistically binding hash along with obfuscation to obtain adaptively secure broadcast encryption with small parameters.

2 Preliminaries

2.1 Indistinguishability Obfuscation and PRFs

Definition 1 (Indistinguishability Obfuscator (iO)). *A uniform PPT machine iO is called an indistinguishability obfuscator for circuits if the following conditions are satisfied:*

- For all security parameters $\kappa \in \mathbb{N}$, for all circuits C , for all inputs x , we have that

$$\Pr[C'(x) = C(x) : C' \leftarrow iO(\kappa, C)] = 1$$

- For any (not necessarily uniform) PPT adversaries $Samp, D$, there exists a negligible function α such that the following holds: if $\Pr[|C_0| = |C_1|]$ and $\forall x, C_0(x) = C_1(x) : (C_0, C_1, \sigma) \leftarrow Samp(1^\kappa) > 1 - \alpha(\kappa)$, then we have:

$$\left| \Pr[D(\sigma, \text{iO}(\kappa, C_0)) = 1 : (C_0, C_1, \sigma) \leftarrow Samp(1^\kappa)] - \Pr[D(\sigma, \text{iO}(\kappa, C_1)) = 1 : (C_0, C_1, \sigma) \leftarrow Samp(1^\kappa)] \right| \leq \alpha(\kappa)$$

Such indistinguishability obfuscators for circuits were constructed under novel algebraic hardness assumptions in [24].

Definition 2. A puncturable family of PRFs F is given by a triple of Turing Machines $\text{Key}_F, \text{Puncture}_F$, and Eval_F , and a pair of computable functions $n(\cdot)$ and $m(\cdot)$, satisfying the following conditions:

- [**Functionality preserved under puncturing**] For every PPT adversary A such that $A(1^\kappa)$ outputs a set $S \subseteq \{0, 1\}^{n(\kappa)}$, then for all $x \in \{0, 1\}^{n(\kappa)}$ where $x \notin S$, we have that:

$$\Pr[\text{Eval}_F(K, x) = \text{Eval}_F(K_S, x) : K \leftarrow \text{Key}_F(1^\kappa), K_S = \text{Puncture}_F(K, S)] = 1$$

- [**Pseudorandom at punctured points**] For every PPT adversary (A_1, A_2) such that $A_1(1^\kappa)$ outputs a set $S \subseteq \{0, 1\}^{n(\kappa)}$ and state σ , consider an experiment where $K \leftarrow \text{Key}_F(1^\kappa)$ and $K_S = \text{Puncture}_F(K, S)$. Then we have

$$\left| \Pr[A_2(\sigma, K_S, S, \text{Eval}_F(K, S)) = 1] - \Pr[A_2(\sigma, K_S, S, U_{m(\kappa) \cdot |S|}) = 1] \right| = \text{negl}(\kappa)$$

where $\text{Eval}_F(K, S)$ denotes the concatenation of $\text{Eval}_F(K, x_1), \dots, \text{Eval}_F(K, x_k)$ where $S = \{x_1, \dots, x_k\}$ is the enumeration of the elements of S in lexicographic order, $\text{negl}(\cdot)$ is a negligible function, and U_ℓ denotes the uniform distribution over ℓ bits.

For ease of notation, we write $\text{PRF}(K, x)$ to represent $\text{Eval}_F(K, x)$. We also represent the punctured key $\text{Puncture}_F(K, S)$ by $K\{S\}$.

The GGM tree-based construction of PRFs [28] from one-way functions are easily seen to yield puncturable PRFs, as recently observed by [5,8,33]. Thus,

Imported Theorem 1. [28,5,8,33] If one-way functions exist, then for all efficiently computable functions $n(\kappa)$ and $m(\kappa)$, there exists a puncturable PRF family that maps $n(\kappa)$ bits to $m(\kappa)$ bits.

2.2 Somewhere Statistically Binding Hash

Definition 3. We use the primitive somewhere statistically binding hash (SSB hash), constructed by Hubacek and Wichs [30]. Intuitively, these are a special type of collision resistant hash function that is binding on a hidden index, and can be used with indistinguishability obfuscation. An SSB hash is a tripe of algorithms $(\text{Gen}, \text{Open}, \text{Ver})$ where:

- $\text{Gen}(s, i)$ takes as input two integers s and i , where $s \leq 2^\kappa$ denotes the number of blocks that will be hashed, and $i \in [s]$ indexes a particular block. The output is a function $H : \Sigma^s \rightarrow \mathcal{Z}$. The size of the description of H is independent of s and i (though it will depend on the security parameter).
- $\text{Open}(H, x = \{x_\ell\}_{\ell \in [s]}, j)$ for $x_\ell \in \Sigma$ and $j \in [s]$ produces an “opening” π that proves that the j^{th} element in x is x_j .
- $\text{Ver}(H, h \in \mathcal{Z}, j \in [s], u \in \Sigma, \pi)$ either accepts or rejects. The idea is that Ver should only accept when $h = H(x)$ where $x_j = u$.
- **Correctness:** $\text{Ver}(H, H(x), j, x_j, \text{Open}(H, x, j))$ accepts.
- **Index hiding:** $\text{Gen}(s, i_0)$ is computationally indistinguishable from $\text{Gen}(s, i_1)$ for any i_0, i_1 .
- **Somewhere Statistically Binding:** If $H \leftarrow \text{Gen}(s, i)$ then if $\text{Ver}(H, h, i, u, \pi)$ and $\text{Ver}(H, h, i, u', \pi')$ accept, it must be that $u = u'$.

Remark. Note that using SSB hash functions, one can efficiently hash down an a-priori unbounded (polynomial) number of values in the security parameter κ .

Imported Theorem 2. *Assuming the existence of FHE, there exists a somewhere statistically binding hash function family mapping unbounded polynomial size inputs to outputs of size κ bits (where κ denotes the security parameter), according to Definition 3.*

3 Definitions

Definition 4 (Multiparty Non-interactive Key Exchange). *An adaptive multiparty NIKE protocol has the following three algorithms:*

- $\text{Setup}(1^\kappa)$: The setup algorithm takes a security parameter κ , and outputs public parameters params .
- $\text{Publish}(1^\kappa, i)$: Each party executes the publishing algorithm, which takes as input the index of the party, and outputs two values: a secret key sv_i and a public value pv_i . Party P_i keeps sv_i as his secret value, and publishes pv_i to the other parties.
- $\text{KeyGen}(\text{params}, \mathcal{S}, (pv_i)_{i \in \mathcal{S}}, j, sv_j)$: To derive the common key $k_{\mathcal{S}}$ for a subset \mathcal{S} , each party in \mathcal{S} runs KeyGen with params , its secret value sv_j and the public values $(pv_i)_{i \in \mathcal{S}}$ of the parties in \mathcal{S} .

Then, these algorithms should satisfy the following properties:

- **Correctness:** For all $\mathcal{S}, i, i' \in \mathcal{S}$,

$$\text{KeyGen}(\text{params}, \mathcal{S}, (pv_j)_{j \in \mathcal{S}}, i, sv_i) = \text{KeyGen}(\text{params}, \mathcal{S}, (pv_j)_{j \in \mathcal{S}}, i', sv_{i'}).$$
- **Security:** The adversary is allowed to (statically) corrupt any subset of users of his choice. More formally, for $b \in \{0, 1\}$, we denote by $\text{expmt}(b)$ the following experiment, parameterized only by the security parameter κ and an adversary \mathcal{A} , and $\text{params} \xleftarrow{\$} \text{Setup}(1^\kappa)$ and $b' \leftarrow \mathcal{A}^{\text{Reg}(\cdot), \text{RegCor}(\cdot, \cdot), \text{Ext}(\cdot), \text{Rev}(\dots), \text{Test}(\dots)}(1^\kappa, \text{params})$ where:

- $\text{Reg}(i \in [2^\kappa])$ registers an honest party P_i . It takes an index i , and runs $(sv_i, pv_i) \leftarrow \text{Publish}(\text{params}, i)$. The challenger then records the tuple $(i, sk_i, pv_i, \text{honest})$ and sends pv_i to \mathcal{A} .
- $\text{RegCor}(i \in [2^\kappa], pk_i)$ registers a corrupt party P_i^* . It takes an index i and a public value pv_i . The challenger records $(i, \perp, pv_i, \text{corrupt})$. The adversary may make multiple queries for a particular identity, in which case the challenger only uses the most recent record.
- $\text{Ext}(i)$ extracts the secret key for an honest registered party. The challenger looks up the tuple $(i, sv_i, pv_i, \text{honest})$ and returns sv_i to \mathcal{A} .
- $\text{Rev}(\mathcal{S}, i)$ reveals the shared secret for a group \mathcal{S} of parties, as calculated by the i th party, where $i \in \mathcal{S}$. We require that party P_i was registered as honest. The challenger uses the secret key for party P_i to derive the shared secret key $k_{\mathcal{S}}$, which it returns to the adversary.
- $\text{Test}(\mathcal{S})$: Takes a set \mathcal{S} of users, all of which were registered as honest. Next, if $b = 0$ the challenger runs KeyGen to determine the shared secret key (arbitrarily choosing which user to calculate the key), which it returns to the adversary. Else if $b = 1$, the challenger generates a random key k to return to the adversary.

A static adversary \mathcal{A} must have the following restrictions:

- \mathcal{A} commits to a set S^* before seeing the public parameters, and,
- \mathcal{A} makes a single query to Test , and this query is on the set S^* .

We require that all register queries and register-corrupt queries are for distinct i , and that $pv_i \neq pv_j$ for any $i \neq j$. For $b = 0$, let W_b be the event that $b' = 1$ in $\text{expmt}(b)$ and we define $\text{Adv}_{\text{NIKE}}(\kappa) = |\Pr[W_0] - \Pr[W_1]|$.

Then, a multi-party key exchange protocol ($\text{Setup}, \text{Publish}, \text{KeyGen}$) is statistically secure if $\text{Adv}_{\text{NIKE}}(\kappa)$ is $\text{negl}(\kappa)$ for any static PPT adversary \mathcal{A} .

4 Static Secure NIKE for Unbounded Parties

4.1 Construction

Let PRF denote a puncturable PRF mapping κ bits to κ bits, and PRG denote a length-doubling pseudorandom generator with inputs of size κ bits. Let $(\text{Gen}, \text{Open}, \text{Ver})$ denote the algorithms of a somewhere statistically binding hash scheme as per Definition 3. Then the static secure NIKE algorithms are constructed as follows.

Setup(1^κ): Pick puncturable PRF key $K \xleftarrow{\$} \{0, 1\}^\kappa$. Run $\text{Gen}(1^\kappa, 2^\kappa, 0)$ to obtain H . Obfuscate using iO the program P_{KE} in Figure 1, padded to the appropriate length. Output $P_{\text{O}} = \text{iO}(P_{KE})$ as public parameters.

Publish: Party i chooses a random seed $s_i \in \{0, 1\}^\lambda$ as a secret value, and publishes $x_i = \text{PRG}(s_i)$.




Fig. 1: Static Secure NIKE Parameters P_{KE}

$\text{KeyGen}(P_{KE}, i, s_i, S, \{pv_j\}_{j \in S})$: Compute $h = H(\{pv_j\}_{j \in S})$. Compute $\pi = \text{Open}(H, \{pv_j\}_{j \in S}, i)$. Run program P_{KE} on input $(h, i, pv, sv, \pi, |S|)$ to obtain shared key K_S .

4.2 Security Game and Hybrids

Hybrid_0 : This is the real world attack game, where \mathcal{A} commits to a set \hat{S} . In response \mathcal{A} gets the public parameters from the setup, and then makes the following queries.

- Register honest user queries: \mathcal{A} submits an index i . The challenger chooses a random s_i , and sends $x_i = \text{PRG}(s_i)$ to \mathcal{A} .
- Register corrupt user queries: \mathcal{A} submits an index i such that $i \notin \hat{S}$, along with a string x_i as the public value for party i . We require that i was not, and will not be registered as honest.
- Extract queries: \mathcal{A} submits an $i \in [N] \setminus \hat{S}$ that was previously registered as honest. The challenger responds with s_i .
- Reveal shared key queries: The adversary submits a subset $\mathcal{S} \neq \hat{S}$ of users, of which at least one is honest. The challenger uses PRF to compute and send the group key.
- Finally, for set \hat{S} , the adversary receives either the correct group key (if $b = 0$) or a random key (if $b = 1$). The adversary outputs a bit b' and wins if $\Pr[b' = b] > \frac{1}{2} + 1/\text{poly}(\kappa)$ for some polynomial $\text{poly}(\cdot)$.

We now demonstrate a sequence of hybrids, via which we argue that the advantage of the adversary in guessing the bit b is $\text{negl}(\kappa)$, where $\text{negl}(\cdot)$ is a function that is asymptotically smaller than $1/\text{poly}(\kappa)$ for all polynomials $\text{poly}(\cdot)$. We give short overviews of indistinguishability between the hybrids, with full proofs in Appendix A. We use underline changes between subsequent hybrids.

Hybrid_1 : For each $i \in \hat{S}$, choose $x_i \xleftarrow{\$} \{0, 1\}^{2\lambda}$. When answering register honest user queries for $i \in \hat{S}$, use these x_i values instead of generating them from PRG.

Follow the rest of the game same as Hybrid_0 . This hybrid is indistinguishable from Hybrid_0 , by security of the PRG.

Let $t^* = |S|$. We start with $\text{Hybrid}_{2,1,a}$, and go across hybrids in the following sequence: $\text{Hybrid}_{2,1,a}, \text{Hybrid}_{2,1,b}, \text{Hybrid}_{2,2,a}, \text{Hybrid}_{2,2,b} \dots \text{Hybrid}_{2,t^*,a}, \text{Hybrid}_{2,t^*,b}$. The hybrids $\text{Hybrid}_{2,i^*,a}$ and $\text{Hybrid}_{2,i^*,b}$ are described below, for $i^* \in [t^*]$.

$\text{Hybrid}_{2,1,a}$: Generate hash function $H \xleftarrow{\$} \text{Gen}(1^\kappa, 2^\kappa, 1)$. Follow the rest of the game same as Hybrid_1 . This is indistinguishable from Hybrid_1 because of indistinguishability of statistical binding index.

$\text{Hybrid}_{2,i^*,a}$: Generate hash function $H \xleftarrow{\$} \text{Gen}(1^\kappa, 2^\kappa, i^*)$. Follow the rest of the game same as $\text{Hybrid}_{2,i^*-1,b}$. This is indistinguishable from $\text{Hybrid}_{2,i^*-1,b}$ because of indistinguishability of statistical binding index.

$\text{Hybrid}_{2,i^*,b}$: Generate hash function $H \xleftarrow{\$} \text{Gen}(1^\kappa, 2^\kappa, i^*)$. Let $h^* = H(\{pv_j\}_{j \in S})$. Set P_{KE} to C_{i^*} , an obfuscation of the circuit in Figure 2. This is indistinguishable from $\text{Hybrid}_{2,i^*,a}$ because of iO between functionally equivalent circuits P_{KE,i^*} and P_{KE,i^*-1} when the hash is statistically binding at index i^* .

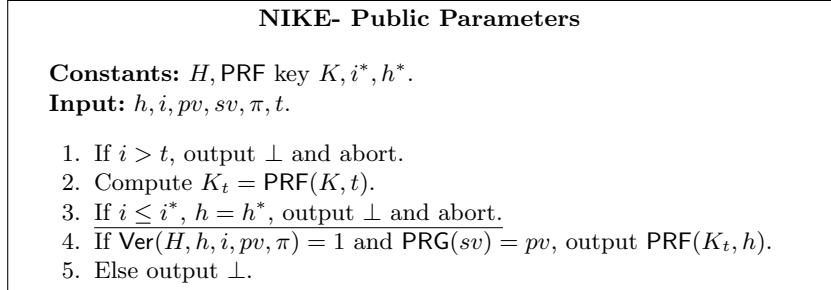


Fig. 2: Static Secure NIKE Parameters P_{KE,i^*}

$\text{Hybrid}_{2,t^*,b}$: Generate hash function $H \xleftarrow{\$} \text{Gen}(1^\kappa, 2^\kappa, t^*)$. Let $h^* = H(\{pv_j\}_{j \in S})$. Set P_{KE} to C_i which is an obfuscation of the circuit in Figure 3.

Hybrid_3 : Generate hash function $H \xleftarrow{\$} \text{Gen}(1^\kappa, 2^\kappa, t^*)$. Let $h^* = H(\{pv_j\}_{j \in S})$. Set $k^* = \text{PRF}(K, t^*)$. Set P_{KE} to C_i which is an obfuscation of the circuit in Figure 4, using punctured key $K\{t^*\}$. Follow the rest of the game the same as in $\text{Hybrid}_{2,t^*,b}$. This program is functionally equivalent to the program in Hybrid_3 , and thus the hybrid is indistinguishable by iO.

Hybrid_4 : Generate hash function $H \xleftarrow{\$} \text{Gen}(1^\kappa, 2^\kappa, t^*)$. Let $h^* = H(\{pv_j\}_{j \in S})$. Set $k^* \xleftarrow{\$} \{0, 1\}^\kappa$. Follow the rest of the game honestly according to Hybrid_3 .

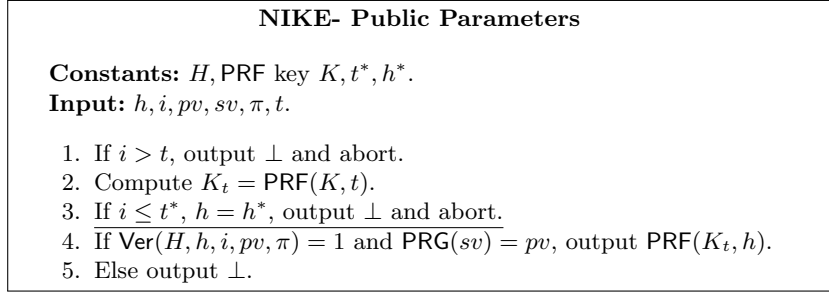


Fig. 3: Static Secure NIKE Parameters P_{KE, t^*}

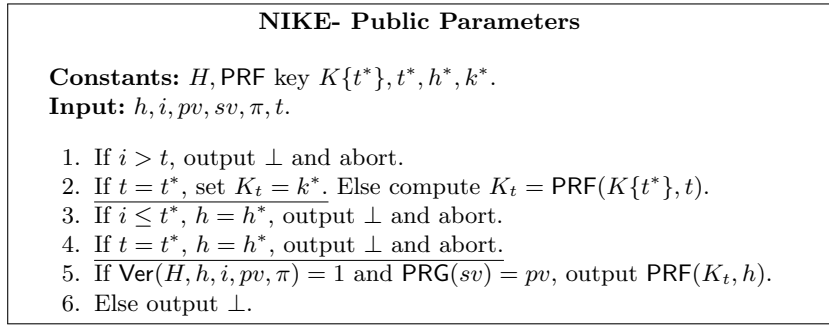


Fig. 4: Static Secure NIKE Parameters $P_{KE'}$

with this value of k^* . This hybrid is indistinguishable from Hybrid_3 because of security of the punctured PRF.

Hybrid_5 : Generate hash function $H \xleftarrow{\$} \text{Gen}(1^\kappa, 2^\kappa, t^*)$. Set $h^* = H(\{pv_j\}_{j \in S})$. Set $k^* \xleftarrow{\$} \{0, 1\}^\kappa$. Set P_{KE} to C_i which is an obfuscation of the circuit in Figure 5, using punctured key $k^*\{h^*\}$. This program is functionally equivalent to the program in Hybrid_4 , thus the hybrids are indistinguishable by iO .

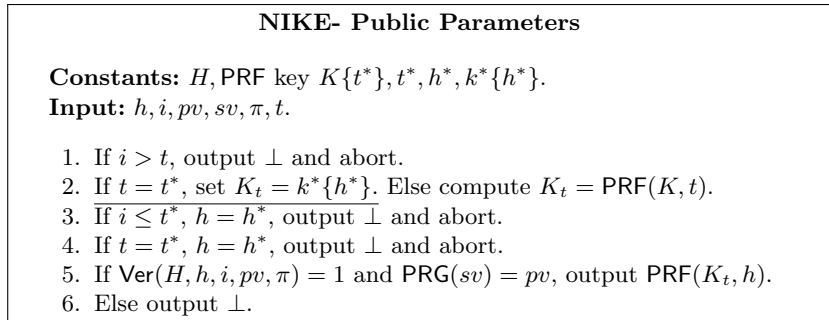


Fig. 5: Static Secure NIKE Parameters P_{KE}

Finally, by security of punctured PRF, \mathcal{A} 's advantage in Hybrid_5 is $\text{negl}(\kappa)$.

4.3 Removing the Setup

We note that in our protocol, the `Publish` algorithm is independent of the `Setup` algorithm. In such a scenario, [6] gave the following theorem, which can be used to remove setup from our scheme in the case of static corruptions.

Imported Theorem 3. [6] *Let $(\text{Setup}, \text{Publish}, \text{KeyGen})$ be a statically secure NIKE protocol where `Publish` does not depend on `params` output by `Setup`, but instead just takes as input (λ, i) . Then there is a statically secure NIKE protocol $(\text{Setup}', \text{Publish}', \text{KeyGen}')$ with no setup.*

5 ID-NIKE for Unbounded Parties

In the identity-based NIKE setting, there is a trusted setup that outputs public parameters, and generates secret keys for parties based on their identity. These parties then run another key-generating algorithm on their secret keys and setup, to compute shared group keys.

Our NIKE scheme can be extended to obtain identity-based NIKE for unbounded parties with a polynomial reduction to the security of indistinguishability obfuscation and fully homomorphic encryption. In this section, we describe our protocol for ID-NIKE for a-priori unbounded parties, and give an overview of the hybrid arguments.

5.1 Construction

Let κ denote the security parameter, `PRF` denote a puncturable pseudo-random function family mapping κ bits to κ bits and `PRG` denote a length-doubling pseudo-random generator with inputs of size κ bits. ID-NIKE consists of the following algorithms.

- `Setup`(1^κ): Sample random PRF keys K_1 and K_2 . Compute the program P_{IBKE} in Figure 6 padded to the appropriate length, and compute $P_{\text{IO}} = \text{iO}(P_{\text{IBKE}})$. Sample SSB hash $H \leftarrow \text{Gen}(1^\kappa)$. Publish the public parameters $\text{params} = P_{\text{IO}}, H$.
- `Extract`(K_2, id): Output $sk_{\text{id}} = \text{PRF}(K_2, \text{id})$.
- `KeyGen`($\text{params}, S, \text{id}, s_{\text{id}}$): To compute shared key k_S for lexicographically ordered set S , compute $h = H(S)$ and $\pi = \text{Open}(h, S, i)$; where i denotes the index of id in sorted S . Then obtain output $k_S = P_{\text{IO}}(h, i, \text{id}, \pi, s_{\text{id}}, |S|)$.

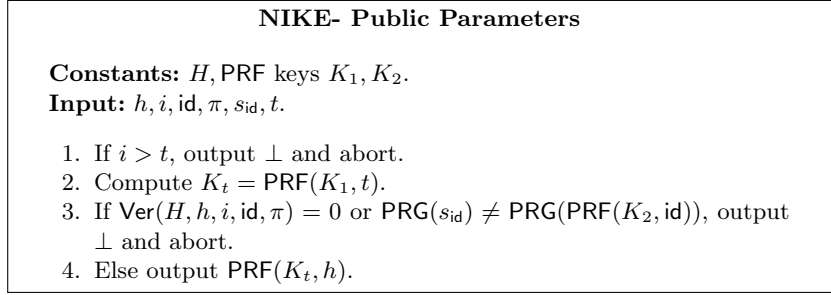


Fig. 6: Static Secure ID-NIKE Parameters P_{IBKE}

5.2 Security Game and Hybrids

Hybrid₀: This is the real world attack game, where \mathcal{A} commits to a set \hat{S} . In response \mathcal{A} gets the public parameters as hash function H and the obfuscation of P_{IBKE} and then makes the following queries.

- Obtain secret keys: \mathcal{A} submits an identity id such that $\text{id} \leq \text{poly}(\kappa)$ and $\text{id} \notin \hat{S}$. The challenger outputs $\text{Extract}(K_2, \text{id})$ to \mathcal{A} .
- Reveal shared keys: The adversary submits a subset $\mathcal{S} \neq \hat{S}$ of users, of which at least one is honest. The challenger uses the public parameters to compute and send the group key.
- Finally, for set \hat{S} , the adversary receives either the correct group key (if $b = 0$) or a random key (if $b = 1$). The adversary outputs a bit b' and wins if $\Pr[b' = b] > \frac{1}{2} + 1/\text{poly}(\kappa)$ for some polynomial $\text{poly}(\cdot)$.

We now demonstrate a sequence of hybrids, via which we argue that the advantage of the adversary in guessing the bit b is $\text{negl}(\kappa)$, where $\text{negl}(\cdot)$ is a function that is asymptotically smaller than $1/\text{poly}(\kappa)$ for all polynomials $\text{poly}(\cdot)$. We give short arguments for indistinguishability between the hybrids, with complete proofs in the full version.

Let $\hat{S} = \{\text{id}_1^*, \text{id}_2^*, \text{id}_3^*, \dots, \text{id}_{|\hat{S}|}^*\}$. Then, for $p \in [1, |\hat{S}|]$, we have the sequence of hybrids: $\text{Hybrid}_{p-1,j}, \text{Hybrid}_{p,a}, \text{Hybrid}_{p,b}, \dots, \text{Hybrid}_{p,j}, \text{Hybrid}_{p+1,a}, \text{Hybrid}_{p+1,b}, \dots$. Here, $\text{Hybrid}_0 \equiv \text{Hybrid}_{0,j}$. We now write out the experiments and demonstrate the sequence of changes between $\text{Hybrid}_{p-1,j}$ and $\text{Hybrid}_{p,j}$ for any $p \in [1, |\hat{S}|]$.

Hybrid_{p-1,j}: This is the same as Hybrid_0 except that the hash is generated as $H \leftarrow \text{Gen}(1^\kappa, 2^\kappa, p-1)$. The challenger computes $h^* = H(\hat{S})$ and outputs the program P_{IBKE} in Figure 7 padded to the appropriate length. He publishes $P_{\text{IO}} = \text{iO}(P_{\text{IBKE}})$.

Hybrid_{p,a}: This is the same as $\text{Hybrid}_{p-1,j}$ except that the challenger computes $r_p^* = \text{PRF}(K_2, \text{id}_p^*)$, $z_p^* = \text{PRG}(r_p^*)$. He computes punctured PRF key $K_2\{\text{id}_p^*\}$ and

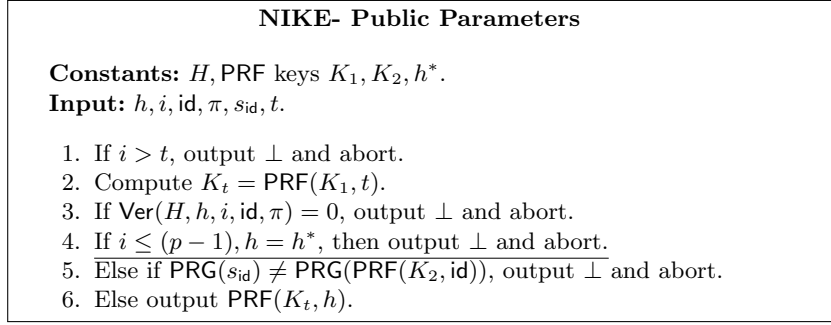


Fig. 7: Static Secure ID-NIKE Parameters P_{BKE}

using the program P_{BKE} in Figure 8 padded to the appropriate length, computes $P_{\text{IO}} = \text{iO}(P_{\text{BKE}})$. This is indistinguishable from $\text{Hybrid}_{p-1,j}$ because of iO between functionally equivalent circuits.

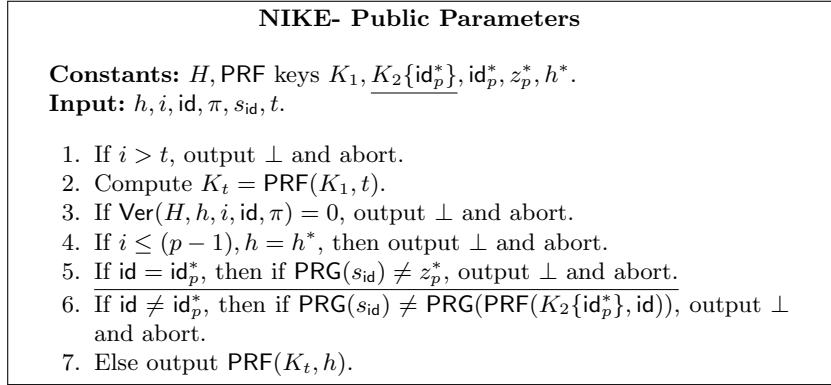


Fig. 8: Static Secure ID-NIKE Parameters P_{BKE}

$\text{Hybrid}_{p,b}$: This is the same as $\text{Hybrid}_{p,a}$ except that the challenger picks $r_p^* \xleftarrow{\$} \{0, 1\}^\kappa$ and sets $z_p^* = \text{PRG}(r_p^*)$. This is indistinguishable from $\text{Hybrid}_{p,a}$ by security of the puncturable PRF.

$\text{Hybrid}_{p,c}$: This is the same as $\text{Hybrid}_{p,b}$ except that the challenger sets $z_p^* \xleftarrow{\$} \{0, 1\}^{2\kappa}$. This is indistinguishable from $\text{Hybrid}_{p,b}$ by security of the PRG.

$\text{Hybrid}_{p,d}$: This is the same as $\text{Hybrid}_{p,c}$ except that the challenger computes the program P_{BKE} in Figure 9 padded to the appropriate length, and publishes $P_{\text{IO}} = \text{iO}(P_{\text{BKE}})$.

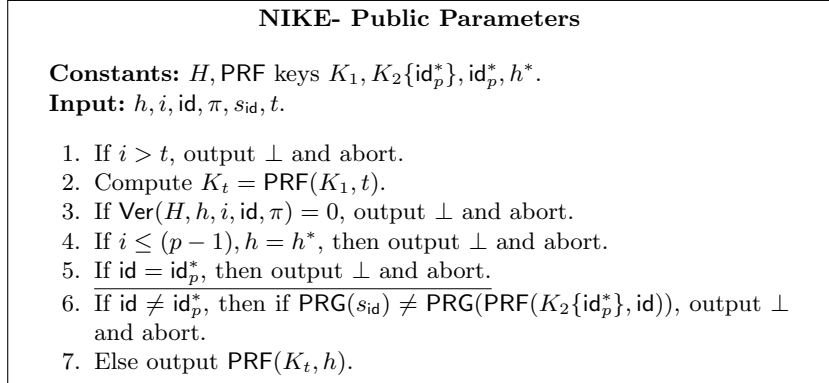


Fig. 9: Static Secure ID-NIKE Parameters P_{IBKE}

With probability $1/2^\kappa$ over random choice of z_1^* , the value z_1^* does not lie in the co-domain of the length-doubling PRG. Then this hybrid is indistinguishable from $\text{Hybrid}_{1,c}$ because of iO between functionally equivalent circuits.

$\text{Hybrid}_{p,e}$: In this hybrid, the challenger generates $H \xleftarrow{\$} \text{Gen}(1^\kappa, 2^\kappa, p)$ (such that it is statistically binding at index p). The rest of the game is same as $\text{Hybrid}_{p,d}$. This hybrid is indistinguishable from $\text{Hybrid}_{p,d}$ because of indistinguishability of statistical binding index.

$\text{Hybrid}_{p,f}$: This is the same as $\text{Hybrid}_{p,e}$ except that the challenger outputs the program P_{IBKE} in Figure 10 padded to the appropriate length. He publishes $P_{\text{iO}} = \text{iO}(P_{\text{IBKE}})$.

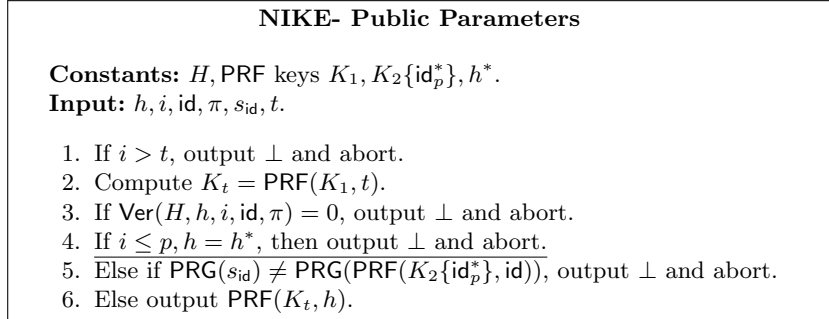


Fig. 10: Static Secure ID-NIKE Parameters P_{IBKE}

This is indistinguishable from $\text{Hybrid}_{p,e}$ because of iO between functionally equivalent circuits. The circuits are functionally equivalent because the hash is statistically binding at index p .

$\text{Hybrid}_{p,g}$: This is the same as $\text{Hybrid}_{p,f}$, except that the challenger picks $z_p^* \xleftarrow{\$} \{0, 1\}^{2\kappa}$, and then outputs the program P_{BKE} in Figure 11 padded to the appropriate length. He publishes $P_{\text{IO}} = \text{iO}(P_{\text{BKE}})$.

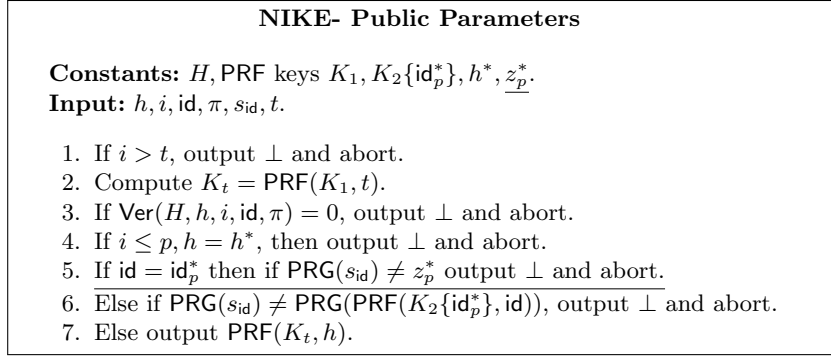


Fig. 11: Static Secure ID-NIKE Parameters P_{BKE}

With probability $1/2^\kappa$ over the randomness of choice of z_p^* , the value z_p^* lies outside the co-domain of the PRG. Thus, with over whelming probability, the extra statement is never activated and the circuit is functionally equivalent to the one in $\text{Hybrid}_{p,f}$. Then this is indistinguishable from $\text{Hybrid}_{p,f}$ because of iO between functionally equivalent circuits.

$\text{Hybrid}_{p,h}$: This is the same as $\text{Hybrid}_{1,g}$ except that the challenger picks $r_p^* \xleftarrow{\$} \{0, 1\}^\kappa$ and sets $z_p^* = \text{PRG}(r_p^*)$. It follows the rest of the game same as $\text{Hybrid}_{p,g}$ with this value of z_p^* . This hybrid is indistinguishable from $\text{Hybrid}_{p,g}$ because of security of length-doubling PRGs.

$\text{Hybrid}_{p,i}$: This is the same as $\text{Hybrid}_{p,h}$ except that the challenger sets $r_p^* = \text{PRF}(K_2, \text{id}_p^*)$ and $z_p^* = \text{PRG}(r_p^*)$. It follows the rest of the game same as $\text{Hybrid}_{p,h}$ with this value of z_p^* . This hybrid is indistinguishable from $\text{Hybrid}_{p,h}$ because of security of the puncturable PRF.

$\text{Hybrid}_{p,j}$: This is the same as $\text{Hybrid}_{p,i}$ except that the challenger outputs the program P_{BKE} in Figure 12 padded to the appropriate length. He publishes $P_{\text{IO}} = \text{iO}(P_{\text{BKE}})$.

This is indistinguishable from $\text{Hybrid}_{p,i}$ by iO between functionally equivalent circuits. Note that at this stage, we have un-punctured the PRF at value id_p^* . This is crucial for our hybrid arguments to go through, because we will eventually have to program in an a-priori un-bounded number of identities.

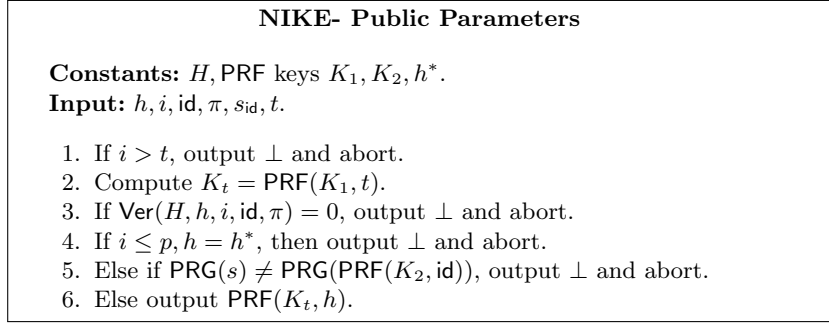


Fig. 12: Static Secure ID-NIKE Parameters P_{BKE}

$\text{Hybrid}_{|\hat{S}|,j}$: This is the final hybrid in the sequence, where the hash is generated as $H \leftarrow \text{Gen}(1^\kappa, 2^\kappa, |\hat{S}|)$. The challenger computes $h^* = H(\hat{S})$ and outputs the program P_{BKE} in Figure 7 padded to the appropriate length. He publishes $P_{\text{O}} = \text{iO}(P_{\text{BKE}})$.

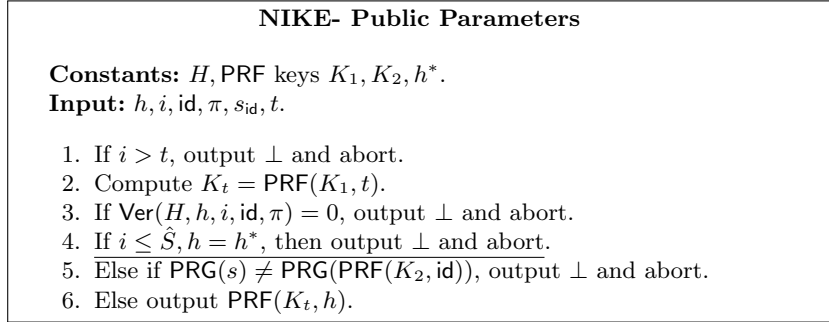


Fig. 13: Static Secure ID-NIKE Parameters P_{BKE}

Finally, we have the following two hybrids

$\text{Hybrid}_{\text{ante-penultimate}}$: In this hybrid, the challenger generates the hash function as $H \leftarrow \text{Gen}(1^\kappa, 2^\kappa, |\hat{S}|)$. The challenger computes $h^* = H(\hat{S})$ and sets $k^* = \text{PRF}(K_1, t^*)$. He punctures the PRF key K_1 on input t^* to obtain $\overline{K\{t^*\}}$. He outputs the program P_{BKE} in Figure 14 padded to the appropriate length. He publishes $P_{\text{O}} = \text{iO}(P_{\text{BKE}})$.

This is indistinguishable from $\text{Hybrid}_{|\hat{S}|,j}$ because of indistinguishability between functionally equivalent circuits.

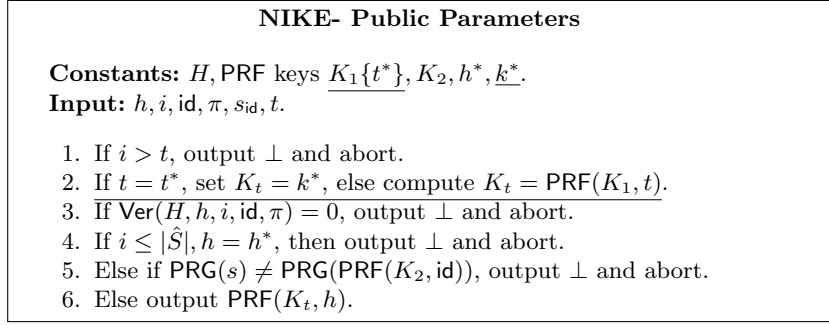


Fig. 14: Static Secure ID-NIKE Parameters P_{BKE}

$\text{Hybrid}_{\text{penultimate}}$: This is the same as $\text{Hybrid}_{\text{ante-penultimate}}$, except that the challenger sets $k^* \xleftarrow{\$} \{0, 1\}^{\kappa}$. This is indistinguishable from $\text{Hybrid}_{\text{ante-penultimate}}$ because of security of the puncturable PRF.

$\text{Hybrid}_{\text{ultimate}}$: This is the same as $\text{Hybrid}_{\text{penultimate}}$, except that the challenger punctures PRF key k^* on value h^* . Then, he sets the program P_{BKE} in Figure 15 using punctured key $\underline{k^*\{h^*\}}$ padded to the appropriate length. He publishes $P_{\text{IO}} = \text{iO}(P_{\text{BKE}})$.

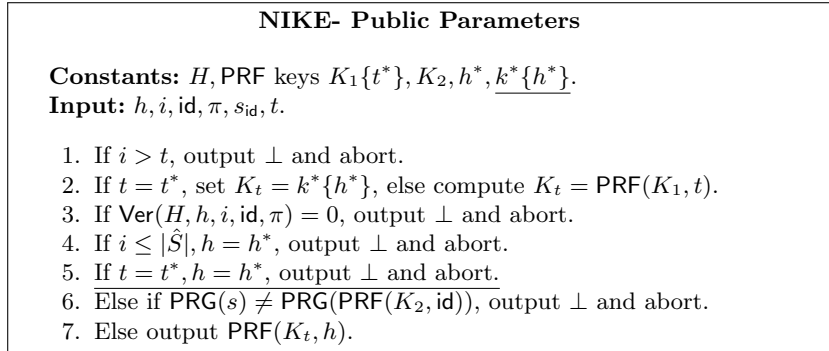


Fig. 15: Static Secure ID-NIKE Parameters P_{BKE}

This hybrid is indistinguishable from $\text{Hybrid}_{\text{penultimate}}$ because of iO between functionally equivalent programs. Finally, the distinguishing advantage of the adversary in this hybrid is at most $\text{negl}(\kappa)$, by security of the puncturable PRF.

6 Conclusion

We construct static-secure protocols that allow NIKE and ID-NIKE between an unbounded number of parties, relying on more feasible assumptions such as indistinguishability obfuscation and fully homomorphic encryption; as opposed to ‘knowledge-type’ assumptions such as differing-inputs obfuscation. It would be interesting to design protocols that tolerate more active attacks by adversaries, for an unbounded number of parties.

A NIKE: Proofs of Indistinguishability of the Hybrids

Lemma 1. *For all (non-uniform) PPT adversaries \mathcal{D} , $\mathcal{D}(\text{Hybrid}_0) \approx_c \mathcal{D}(\text{Hybrid}_1)$.*

Proof. We define a sub-sequence of polynomially many (concretely, $|\hat{S}|$) sub-hybrids $\text{Hybrid}_{0,j}$ for $j \in |\hat{S}|$, where $\text{Hybrid}_{0,j}$ is the same as Hybrid_0 except that the challenger samples the first j public values for parties in the set \hat{S} at random in $\{0,1\}^{2\kappa}$ instead of generating them as the output of the PRG. Note that $\text{Hybrid}_{0,0} \equiv \text{Hybrid}_0$ and $\text{Hybrid}_{0,|\hat{S}|} \equiv \text{Hybrid}_1$.

We show that hybrids $\text{Hybrid}_{0,j}$ and $\text{Hybrid}_{0,j+1}$ are computationally indistinguishable. We will prove this by contradiction.

Suppose there exists a distinguisher \mathcal{D} which distinguishes between $\text{Hybrid}_{0,j}$ and $\text{Hybrid}_{0,j+1}$ with advantage $1/\text{poly}(\kappa)$ for some polynomial $\text{poly}(\cdot)$. We construct a reduction that uses this distinguisher to break security of the underlying PRG. The reduction obtains a challenge value x , which may either be the output of the PRG on a uniform input, or may be chosen uniformly at random in $\{0,1\}^{2\kappa}$. It picks the first j public values for parties in the set \hat{S} at random in $\{0,1\}^{2\kappa}$ instead of generating them as the output of the PRG. It sets the $(j+1)^{\text{th}}$ public value to the challenge value x . It samples the remaining public values for parties in the set \hat{S} by picking a uniform $s_i \in \{0,1\}^\kappa$ and computing $x_i = \text{PRG}(s_i)$.

If x is the output of a PRG this is the experiment in $\text{Hybrid}_{0,j}$ else it is $\text{Hybrid}_{0,j+1}$. Therefore, the reduction can mimic the output of the distinguisher between $\text{Hybrid}_{0,j}$ and $\text{Hybrid}_{0,j+1}$, thereby breaking the security of the PRG with advantage $1/\text{poly}(\kappa)$.

Lemma 2. *For all (non-uniform) PPT adversaries \mathcal{D} , $\mathcal{D}(\text{Hybrid}_1) \approx_c \mathcal{D}(\text{Hybrid}_{2,1,a})$.*

Proof. Suppose there exists a distinguisher which distinguishes between Hybrid_1 and $\text{Hybrid}_{2,1,a}$ with advantage $1/\text{poly}(\kappa)$ for some polynomial $\text{poly}(\cdot)$. We construct a reduction that uses this distinguisher to break index-hiding security of the somewhere statistically hiding hash.

The reduction gives indices $\{0, 1\}$ to the hash challenger. The challenger then generates hash $H = \text{Gen}(s, b)$ for $b \xleftarrow{\$} \{0, 1\}$ and sends them to the reduction. The reduction uses this function H as the hash (instead of generating the hash itself), and continues the game with the distinguisher. If $b = 0$, this corresponds to Hybrid_1 , and if $b = 1$ the game corresponds to $\text{Hybrid}_{2,1,a}$. Therefore, the reduction can mimic the output of the distinguisher between $\text{Hybrid}_{2,i^*-1,b}$ and $\text{Hybrid}_{2,i^*,a}$, thereby breaking index-hiding security of the hash function with advantage $1/\text{poly}(\kappa)$.

Lemma 3. *For all (non-uniform) PPT adversaries \mathcal{D} , $\mathcal{D}(\text{Hybrid}_{2,i^*-1,b}) \approx_c \mathcal{D}(\text{Hybrid}_{2,i^*,a})$.*

Proof. Suppose there exists a distinguisher which distinguishes between $\text{Hybrid}_{2,i^*-1,b}$ and $\text{Hybrid}_{2,i^*,a}$ with advantage $1/\text{poly}(\kappa)$ for some polynomial $\text{poly}(\cdot)$. We construct a reduction that uses this distinguisher to break index-hiding security of the somewhere statistically hiding hash.

The reduction gives indices $\{i^* - 1, i^*\}$ to the hash challenger. The challenger then generates hash $H = \text{Gen}(s, b)$ for $b \xleftarrow{\$} \{i^* - 1, i^*\}$ and sends them to the reduction. The reduction uses this function H as the hash (instead of generating the hash itself), and continues the game with the distinguisher. If $b = i^* - 1$, this corresponds to $\text{Hybrid}_{2,i^*-1,b}$, and if $b = i^*$ the game corresponds to $\text{Hybrid}_{2,i^*,a}$. Therefore, the reduction can mimic the output of the distinguisher between Hybrid_1 and $\text{Hybrid}_{2,1,a}$, thereby breaking index-hiding security of the hash function with advantage $1/\text{poly}(\kappa)$.

Lemma 4. *For all (non-uniform) PPT adversaries \mathcal{D} , $\mathcal{D}(\text{Hybrid}_{2,i^*,a}) \approx_c \mathcal{D}(\text{Hybrid}_{2,i^*,b})$.*

Proof. Note that in the experiment of $\text{Hybrid}_{2,1,a/b}$, $H \leftarrow \text{Gen}(1^\kappa, 2^\kappa, i^*)$. Thus, by the statistical binding property of the hash function, if $\text{Ver}(H, h, i^*, u, \pi)$ and $\text{Ver}(H, h, i^*, u', \pi')$ accept, then it must be that $u = u'$.

Consider the programs P_{KE,i^*-1} and P_{KE,i^*} . Note that the only place where the two programs may differ is on inputs of the form $(h^*, i^*, pv, sv, \pi, t)$, where $h^* = H(\hat{S})$. Denote the public values $\{pv_j\}_{j \in S}$ by $pv_{x_1}, pv_{x_2}, \dots, pv_{x_{|S|}}$. In this case, P_{KE,i^*-1} (in $\text{Hybrid}_{2,i^*,a}$) checks if $\text{Ver}(H, h^*, i^*, pv, \pi) = 1$ and if $\text{PRG}(sv) = pv$, then outputs $\text{PRF}(K, h)$ else outputs \perp . On the other hand, P_{KE,i^*} (in $\text{Hybrid}_{2,i^*,b}$) always outputs \perp . Because of the statistical binding property of the hash at index i , if $\text{Ver}(H, h^*, i^*, pv, \pi)$ accepts for any value of (pv, π) , then $pv = pv_{x_{i^*}}$. Moreover, since $pv_{x_{i^*}}$ is uniformly chosen in the range of the PRG, then with overwhelming probability, there does not exist any value sv such that $\text{PRG}(sv) = pv_{x_{i^*}}$. Thus, the ‘if’ condition in P_{KE,i^*-1} in $\text{Hybrid}_{2,i^*,a}$ will never be activated, and the two programs are functionally equivalent.

Therefore, the obfuscated circuits $iO(P_{KE,i^*-1})$ and $iO(P_{KE,i^*})$ must be indistinguishable by security of the iO. Suppose they are not, then consider a distinguisher \mathcal{D} which distinguishes between these hybrids with non-negligible advantage. \mathcal{D} can be used to break selective security of the indistinguishability obfuscation (according to Definition 1) via the following reduction to iO. The reduction acts as challenger in the experiment of $\text{Hybrid}_{2,i^*,a}$.

The iO challenger $\text{Samp}(1^\kappa)$ first activates the reduction, which samples the two circuits P_{KE,i^*-1}, P_{KE,i^*} and gives them to $\text{Samp}(1^\kappa)$. The challenger then samples challenge circuit $C \xleftarrow{\$} \{P_{KE,i^*-1}, P_{KE,i^*}\}$, and sends $C' = iO(C)$ to the reduction. The reduction continues the game of $\text{Hybrid}_{2,i^*,a}$ with C' in place of the obfuscation of program P_{KE,i^*-1} .

If $C = P_{KE,i^*-1}$, this corresponds to $\text{Hybrid}_{2,i^*-1,b}$, and if $C = P_{KE,i^*}$ the game corresponds to $\text{Hybrid}_{2,i^*,a}$. Therefore, the reduction can mimic the output of the distinguisher between $\text{Hybrid}_{2,i^*,a}$ and $\text{Hybrid}_{2,i^*,b}$, thereby breaking security of the iO with advantage $1/\text{poly}(\kappa)$.

Lemma 5. *For all (non-uniform) PPT adversaries \mathcal{D} , $\mathcal{D}(\text{Hybrid}_{2,t^*,b}) \approx_c \mathcal{D}(\text{Hybrid}_3)$.*

Proof. Consider the programs P_{KE,t^*} ($\text{Hybrid}_{2,t^*,b}$) and $P_{KE'}$ (Hybrid_3). Note that the only place where the two programs may differ is on inputs where $t = t^*$.

In this case, if $i \leq t^*$, then for all $h = h^*$, both programs output \perp and abort. If $i > t^*$ and $t = t^*$, then $i > t$ and both programs output \perp and abort in Step 1. Moreover, for $t = t^*$, $k^* = \text{PRF}(K, t^*)$ and thus the programs are functionally equivalent.

Therefore, the obfuscated circuits $iO(P_{KE,t^*})$ and $iO(P_{KE'})$ must be indistinguishable by security of the iO. Suppose they are not, then consider a distinguisher \mathcal{D} which distinguishes between these hybrids with non-negligible advantage. \mathcal{D} can be used to break security of the indistinguishability obfuscation (according to Definition 1) via the following reduction to iO. The reduction acts as challenger in the experiment of Hybrid_3 .

The iO challenger $\text{Samp}(1^\kappa)$ first activates the reduction, which samples the two circuits $P_{KE,t^*}, P_{KE'}$ and gives them to $\text{Samp}(1^\kappa)$. The challenger then samples challenge circuit $C \xleftarrow{\$} \{P_{KE,t^*}, P_{KE'}\}$, and sends $C' = iO(C)$ to the reduction. The reduction continues the game of Hybrid_3 with C' in place of the obfuscation of program $P_{KE'}$.

If $C = P_{KE,t^*}$, this corresponds to $\text{Hybrid}_{2,t^*,b}$, and if $C = P_{KE'}$ the game corresponds to Hybrid_3 . Therefore, the reduction can mimic the output of the distinguisher between $\text{Hybrid}_{2,t^*,b}$ and Hybrid_3 , thereby breaking security of the iO with advantage $1/\text{poly}(\kappa)$.

Lemma 6. For all (non-uniform) PPT adversaries \mathcal{D} , $\mathcal{D}(\text{Hybrid}_3) \approx_c \mathcal{D}(\text{Hybrid}_4)$.

Proof. Suppose there exists a distinguisher \mathcal{D} which distinguishes between these hybrids with non-negligible advantage. \mathcal{D} can be used to break selective security of the puncturable PRF via the following reduction. The reduction acts as challenger in the experiment of Hybrid_3 .

It obtains challenge set \hat{S} from the distinguisher, and computes $t^* = |\hat{S}|$. Then, it gives t^* to the PRF challenger, and obtains punctured PRF key $K\{t^*\}$ and a challenge a , which is either chosen uniformly at random or is the output of the PRF at t^* . Then, the reduction continues the experiment of Hybrid_3 as challenger, except that he sets $r^* = a$.

If $a = \text{PRF}(K, t^*)$ then this is the experiment of Hybrid_3 , and if a is chosen uniformly at random, then this is the experiment of Hybrid_4 . Therefore, the reduction can mimic the output of the distinguisher between Hybrid_3 and Hybrid_4 , thereby breaking security of the puncturable PRF with advantage $1/\text{poly}(\kappa)$.

Lemma 7. For all (non-uniform) PPT adversaries \mathcal{D} , $\mathcal{D}(\text{Hybrid}_4) \approx_c \mathcal{D}(\text{Hybrid}_5)$.

Proof. Consider the programs $P_{KE'}$ (Hybrid_4) and $P_{KE''}$ (Hybrid_5). Note that the only place where the two programs may differ is on inputs where $t = t^*$. Then, for $t = t^*, h = h^*$, both programs output \perp . Moreover, because of functional equivalence of the punctured key $k^*\{h^*\}$ and k^* on all points where $h \neq h^*$, the programs are equivalent.

Suppose there exists a distinguisher \mathcal{D} which distinguishes between these hybrids with non-negligible advantage. \mathcal{D} can be used to break security of indistinguishability obfuscation via the following reduction. The reduction acts as challenger in the experiment of Hybrid_4 . It obtains challenge set \hat{S} from the distinguisher, and computes $h^* = H(\hat{S})$. Then, it constructs gives circuits $P_{KE'}, P_{KE''}$ as input to the Samp algorithm. Samp picks $C \xleftarrow{\$} \{P_{KE'}, P_{KE''}\}$ and sends $C' = iO(C)$ to the reduction. The reduction uses C' in place of $P_{KE'}$ in the experiment of Hybrid_4 .

If $C = P_{KE'}$ then this is the experiment of Hybrid_4 , and if $C = P_{KE''}$, then this is the experiment of Hybrid_5 . Therefore, the reduction can mimic the output of the distinguisher between Hybrid_4 and Hybrid_5 , thereby breaking security of the indistinguishability obfuscation with advantage $1/\text{poly}(\kappa)$.

Lemma 8. For all (non-uniform) PPT adversaries \mathcal{D} , $\text{Adv}_{\mathcal{D}}(\text{Hybrid}_5) = \text{negl}(\kappa)$.

Proof. Suppose there exists a distinguisher \mathcal{D} which has non-negligible advantage in Hybrid_5 . \mathcal{D} can be used to break selective security of the puncturable PRF via

the following reduction. The reduction acts as challenger in the experiment of Hybrid_5 .

It obtains challenge set \hat{S} from the distinguisher, and computes $h^* = H(\hat{S})$. Then, it gives h^* to the PRF challenger, and obtains punctured PRF key $k^*\{t^*\}$ and a challenge a , which is either chosen uniformly at random or is the output of the PRF at h^* . Then, the reduction continues the experiment of Hybrid_3 as challenger, except that he sets the shared key to a .

If $a = \text{PRF}(k^*, h^*)$ then this corresponds to the correct shared key for group \hat{S} , whereas if a is chosen uniformly at random, this corresponds to a random key. Therefore, the reduction can mimic the output of the distinguisher, thereby breaking security of the puncturable PRF with advantage $1/\text{poly}(\kappa)$.

References

1. Abusalah, H., Fuchsbauer, G., Pietrzak, K.: Constrained prfs for unbounded inputs. IACR Cryptology ePrint Archive 2014, 840 (2014), <http://eprint.iacr.org/2014/840>
2. Ananth, P., Boneh, D., Garg, S., Sahai, A., Zhandry, M.: Differing-inputs obfuscation and applications. IACR Cryptology ePrint Archive 2013, 689 (2013), <http://eprint.iacr.org/2013/689>
3. Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. In: Kilian, J. (ed.) Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings. Lecture Notes in Computer Science, vol. 2139, pp. 1–18. Springer (2001), http://dx.doi.org/10.1007/3-540-44647-8_1
4. Boneh, D., Silverberg, A.: Applications of multilinear forms to cryptography. IACR Cryptology ePrint Archive 2002, 80 (2002), <http://eprint.iacr.org/2002/080>
5. Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. IACR Cryptology ePrint Archive 2013, 352 (2013)
6. Boneh, D., Zhandry, M.: Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In: Garay and Gennaro [21], pp. 480–499, http://dx.doi.org/10.1007/978-3-662-44371-2_27
7. Boyle, E., Chung, K., Pass, R.: On extractability obfuscation. In: Lindell, Y. (ed.) Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8349, pp. 52–73. Springer (2014), http://dx.doi.org/10.1007/978-3-642-54242-8_3
8. Boyle, E., Goldwasser, S., Ivan, I.: Functional signatures and pseudorandom functions. IACR Cryptology ePrint Archive 2013, 401 (2013)
9. Boyle, E., Pass, R.: Limits of extractability assumptions with distributional auxiliary input. IACR Cryptology ePrint Archive 2013, 703 (2013), <http://eprint.iacr.org/2013/703>
10. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE . SIAM J. Comput. 43(2), 831–871 (2014), <http://dx.doi.org/10.1137/120868669>

11. Canetti, R., Garay, J.A. (eds.): Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I, Lecture Notes in Computer Science, vol. 8042. Springer (2013), <http://dx.doi.org/10.1007/978-3-642-40041-4>
12. Canetti, R., Lin, H., Tessaro, S., Vaikuntanathan, V.: Obfuscation of probabilistic circuits and applications. In: Dodis and Nielsen [17], pp. 468–497, http://dx.doi.org/10.1007/978-3-662-46497-7_19
13. Cash, D., Kiltz, E., Shoup, V.: The twin diffie-hellman problem and applications. *J. Cryptology* 22(4), 470–504 (2009), <http://dx.doi.org/10.1007/s00145-009-9041-6>
14. Coron, J., Lepoint, T., Tibouchi, M.: Practical multilinear maps over the integers. In: Canetti and Garay [11], pp. 476–493, http://dx.doi.org/10.1007/978-3-642-40041-4_26
15. Coron, J., Lepoint, T., Tibouchi, M.: New multilinear maps over the integers. *IACR Cryptology ePrint Archive* 2015, 162 (2015), <http://eprint.iacr.org/2015/162>
16. Diffie, W., Hellman, M.E.: New directions in cryptography (1976)
17. Dodis, Y., Nielsen, J.B. (eds.): Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II, Lecture Notes in Computer Science, vol. 9015. Springer (2015), <http://dx.doi.org/10.1007/978-3-662-46497-7>
18. Dupont, R., Enge, A.: Provably secure non-interactive key distribution based on pairings. *Discrete Applied Mathematics* 154(2), 270 – 276 (2006), <http://www.sciencedirect.com/science/article/pii/S0166218X05002337>, coding and Cryptography
19. Freire, E.S.V., Hofheinz, D., Kiltz, E., Paterson, K.G.: Non-interactive key exchange. *IACR Cryptology ePrint Archive* 2012, 732 (2012), <http://eprint.iacr.org/2012/732>
20. Freire, E.S.V., Hofheinz, D., Paterson, K.G., Striecks, C.: Programmable hash functions in the multilinear setting. In: Canetti and Garay [11], pp. 513–530, http://dx.doi.org/10.1007/978-3-642-40041-4_28
21. Garay, J.A., Gennaro, R. (eds.): Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I, Lecture Notes in Computer Science, vol. 8616. Springer (2014), <http://dx.doi.org/10.1007/978-3-662-44371-2>
22. Garg, S., Gentry, C., Halevi, S.: Candidate multilinear maps from ideal lattices. In: EUROCRYPT (2013)
23. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: 54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA. pp. 40–49. IEEE Computer Society (2013), <http://dx.doi.org/10.1109/FOCS.2013.13>
24. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: FOCS (2013)
25. Garg, S., Gentry, C., Halevi, S., Wichs, D.: On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In: Garay and Gennaro [21], pp. 518–535, http://dx.doi.org/10.1007/978-3-662-44371-2_29
26. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: Mitzenmacher, M. (ed.) Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009. pp. 169–178. ACM (2009), <http://doi.acm.org/10.1145/1536414.1536440>

27. Gentry, C., Gorbunov, S., Halevi, S.: Graph-induced multilinear maps from lattices. In: Dodis and Nielsen [17], pp. 498–527, http://dx.doi.org/10.1007/978-3-662-46497-7_20
28. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions (extended abstract). In: FOCS. pp. 464–479 (1984)
29. Hofheinz, D., Jager, T., Khurana, D., Sahai, A., Waters, B., Zhandry, M.: How to generate and use universal parameters. IACR Cryptology ePrint Archive 2014, 507 (2014), <http://eprint.iacr.org/2014/507>
30. Hubacek, P., Wichs, D.: On the communication complexity of secure function evaluation with long output. In: Roughgarden, T. (ed.) Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11–13, 2015. pp. 163–172. ACM (2015), <http://doi.acm.org/10.1145/2688073.2688105>
31. Ishai, Y., Pandey, O., Sahai, A.: Public-coin differing-inputs obfuscation and its applications. In: Dodis and Nielsen [17], pp. 668–697, http://dx.doi.org/10.1007/978-3-662-46497-7_26
32. Joux, A.: A one round protocol for tripartite diffie-hellman. In: Bosma, W. (ed.) Algorithmic Number Theory, 4th International Symposium, ANTS-IV, Leiden, The Netherlands, July 2–7, 2000, Proceedings. Lecture Notes in Computer Science, vol. 1838, pp. 385–394. Springer (2000), http://dx.doi.org/10.1007/10722028_23
33. Kiayias, A., Papadopoulos, S., Triandopoulos, N., Zacharias, T.: Delegatable pseudorandom functions and applications. IACR Cryptology ePrint Archive 2013, 379 (2013)
34. Koppula, V., Lewko, A.B., Waters, B.: Indistinguishability obfuscation for turing machines with unbounded memory. IACR Cryptology ePrint Archive 2014, 925 (2014), <http://eprint.iacr.org/2014/925>
35. Paterson, K.G., Srinivasan, S.: On the relations between non-interactive key distribution, identity-based encryption and trapdoor discrete log groups. Des. Codes Cryptography 52(2), 219–241 (2009), <http://dx.doi.org/10.1007/s10623-009-9278-y>
36. Rao, V.: Adaptive multiparty non-interactive key exchange without setup in the standard model. IACR Cryptology ePrint Archive 2014, 910 (2014), <http://eprint.iacr.org/2014/910>
37. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Shmoys, D.B. (ed.) Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014. pp. 475–484. ACM (2014), <http://doi.acm.org/10.1145/2591796.2591825>
38. Sakai, R., Ohgishi, K., Kasahara, M.: Cryptosystems based on pairing. Symposium on Cryptography and Information Security SCIS (2000)
39. Yamakawa, T., Yamada, S., Hanaoka, G., Kunihiro, N.: Self-bilinear map on unknown order groups from indistinguishability obfuscation and its applications. Cryptology ePrint Archive, Report 2015/128 (2015), <http://eprint.iacr.org/>
40. Zhandry, M.: Adaptively secure broadcast encryption with small system parameters. IACR Cryptology ePrint Archive 2014, 757 (2014), <http://eprint.iacr.org/2014/757>