

Constrained PRFs for Unbounded Inputs

Hamza Abusalah*

Georg Fuchsbauer*

Krzysztof Pietrzak*

Institute of Science and Technology Austria
{habusalah,gfuchsbauer,pietrzak}@ist.ac.at

Abstract

A constrained pseudorandom function $F: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ for a family $\mathcal{T} \subseteq 2^{\mathcal{X}}$ of subsets of \mathcal{X} is a function where for any key $k \in \mathcal{K}$ and set $S \in \mathcal{T}$ one can efficiently compute a constrained key k_S which allows to evaluate $F(k, \cdot)$ on all inputs $x \in S$, while even given this key, the outputs on all inputs $x \notin S$ look random. At Asiacrypt'13 Boneh and Waters gave a construction which supports the most general set family so far. Its keys k_C are defined for sets decided by boolean circuits C and enable evaluation of the PRF on any $x \in \mathcal{X}$ where $C(x) = 1$. In their construction the PRF input length and the size of the circuits C for which constrained keys can be computed must be fixed beforehand during key generation.

We construct a constrained PRF that has an unbounded input length and whose constrained keys can be defined for any set recognized by a Turing machine. The only a priori bound we make is on the description size of the machines. We prove our construction secure assuming public-coin differing-input obfuscation.

As applications of our constrained PRF we build a broadcast encryption scheme where the number of potential receivers need not be fixed at setup (in particular, the length of the keys is independent of the number of parties) and the first identity-based non-interactive key exchange protocol with no bound on the number of parties that can agree on a shared key.

Keywords: Constrained PRFs, broadcast encryption, identity-based non-interactive key exchange.

1 Introduction

Constrained PRFs. A pseudorandom function (PRF) [GGM86] is a keyed function $F: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ for which no efficient adversary, given access to an oracle $\mathcal{O}(\cdot)$, can distinguish the case where $\mathcal{O}(\cdot)$ is $F(k, \cdot)$ with a random key $k \in \mathcal{K}$ from the case where $\mathcal{O}(\cdot)$ is a uniformly random function $\mathcal{X} \rightarrow \mathcal{Y}$.

Three papers [BW13, BGI14, KPTZ13] independently introduce the concept of a *constrained* PRF. Consider a set \mathcal{P} , where each $v \in \mathcal{P}$ specifies some predicate $p_v: \mathcal{X} \rightarrow \{0, 1\}$ that defines a (potentially exponential-size) subset $S_v = \{x \in \mathcal{X} \mid p_v(x) = 1\}$. A constrained PRF for a predicate family \mathcal{P} is a PRF F with an additional constrain algorithm $k_v \leftarrow F.\text{Constr}(k, v)$ that on input a key $k \in \mathcal{K}$ and a predicate $v \in \mathcal{P}$ outputs a constrained key k_v that can be used to compute $F(k, x)$ on all $x \in S_v$, while, given this key, all values $F(k, x)$ for $x \notin S_v$ still look random.

Constrained PRFs (CPRF) have been constructed for several interesting predicates. All three papers [BW13, BGI14, KPTZ13] show that the classical GGM construction [GGM86] of a PRF with input domain $\{0, 1\}^n$ yields a *prefix-constrained* PRF. This means $\mathcal{P} = \{0, 1\}^{\leq n}$ and for any $v \in \mathcal{P}$ the derived key k_v allows to compute $F(k, x)$ for all x with prefix v , i.e., $x = v \| x' \in \{0, 1\}^n$ for some x' . Assuming (leveled) multilinear maps [GGH13a, CLT13, LSS14], Boneh and Waters [BW13] construct CPRFs for much more general set systems. They present a bit-fixing PRF, where $\mathcal{P} = \{0, 1, ?\}^n$ and for $v \in \mathcal{P}$ we have $p_v(x) = 1$ if x agrees with v on all indices different from '?',

*Supported by the European Research Council, ERC Starting Grant (259668-PSPC)

i.e., for all $i = 1, \dots, n$, either $v[i] = ?$ or $v[i] = x[i]$. They moreover construct a circuit-constrained PRF, where the predicates are arbitrary circuits $C: \{0, 1\}^n \rightarrow \{0, 1\}$ of some fixed depth.

CPRFs have already found many interesting applications. From a prefix CPRF, one can construct a puncturable PRF [SW14], which is a constrained PRF for predicates $\mathcal{P} = \{0, 1\}^n$ where for $v \in \mathcal{P}$, the key k_v lets one compute $F(k, x)$ on all $x \neq v$. The GGM construction yields a puncturable PRF with punctured keys whose length is linear in the PRF input length. Puncturable PRFs play a crucial role in the security proofs of most of the recent constructions based on indistinguishability obfuscation [BGI⁺12, GGH⁺13b], and we will also use them in this paper.

The more general bit-fixing and circuit-constrained PRFs can be used to construct a variety of sophisticated cryptographic tools including broadcast encryption (BE) and identity-based non-interactive key-exchange, as outlined next.

BROADCAST ENCRYPTION. In a BE scheme [FN94, YFDL04, BGW05, BH08, PPS11, BWZ14] there is a set of n users, and for any given subset $S \subseteq \{1, \dots, n\}$ of users, we want to be able to encrypt a message (as a short ciphertext) that can be decrypted only by the users included in S . This can be achieved using a bit-fixing PRF with domain $\{0, 1\}^n$: Sample a random key k , and give a constrained key k_{v_i} to user i where $v_i[i] = 1$ and $v_i[j] = ?$ for any $j \neq i$. Thus, k_{v_i} allows to evaluate the PRF on exactly those inputs with a ‘1’ in position i .

To broadcast a message m to a set S of users, we simply send a symmetric encryption of m under the key $F(k, x_S)$, where $x_S[i] = 1$ if $i \in S$ and $x_S[i] = 0$ otherwise. Note that user i can compute $F(k, x_S)$ (and thus decrypt) iff her key k_{v_i} satisfies $v_i[i] = x_S[i]$, which by construction holds iff $i \in S$.

NON-INTERACTIVE KEY EXCHANGE. In an identity-based non-interactive key exchange (ID-NIKE) [SOK00, FHPS13, BW13] scheme there are parties that each have some identity $id \in \{0, 1\}^\ell$. For any set S of at most n parties we want the parties in S to be able to locally compute a shared key K_S which is indistinguishable from random for all parties outside of S . Such a scheme can be constructed from a bit-fixing PRF F with domain $\{0, 1\}^{n \cdot \ell}$. At setup, sample a key k for F and give to party $id \in \{0, 1\}^\ell$ a set of n constrained keys $k_{id}^{(1)}, \dots, k_{id}^{(n)}$, where $k_{id}^{(i)}$ is a key for the set $?^{(i-1)\ell} \| id \| ?^{(n-i)\ell}$. Now, only parties id_1, \dots, id_n can compute the joint key $K_S := F(k, id_1 \| id_2 \| \dots \| id_n)$.

CPRFs with unbounded input length. The disadvantage of the BE and ID-NIKE constructions just outlined is that the number n of possible recipients (for BE) or parties agreeing on a key (for ID-NIKE) must be fixed when setting up the system. Moreover, the length of the constrained keys given to every user is at least linear in n .

In this paper we construct a constrained PRF for which there is no a priori bound on the input length. The constraints on keys are given by Turing machines (TM), that is, given a key k and a machine M , we can derive a constrained key k_M that allows to compute $F(k, x)$ for any input x for which $M(x) = 1$. The only thing that must be bounded beforehand is the *size* of the TMs that we want to support. In our construction a constrained key for a machine M will be an obfuscated circuit whose size only depends on the size of M .

ADAPTIVE VS. SELECTIVE SECURITY. Security of constrained PRFs is defined via a game in which a challenger picks a random key k and the adversary chooses $x^* \in \mathcal{X}$ and must distinguish $F(k, x^*)$ from random. The adversary has also access to oracles to query constrained keys for sets $S \not\ni x^*$. We prove *selective* security of our CPRF where the adversary must choose x^* before it can query constrained keys. From a selectively secure CPRF we can get an *adaptively* secure CPRF (where the adversary can decide on x^* after its key queries) via “complexity leveraging”—but this reduction loses a factor that is exponential in the input length. Proving adaptive security for CPRFs without an exponential security loss is generally hard and Fuchsbaauer et al. [FKPR14]

show that for the bit-fixing CPRF from [BW13] any “simple” security reduction must lose an exponential factor.

Adaptive security of CPRFs was proved for the GGM-based prefix-constrained PRF in [FKPR14] losing only a quasi-polynomial (rather than an exponential) factor. Moreover, Hohenberger, Koppula and Waters [HKW14] construct an adaptively secure puncturable PRF with polynomial security loss using indistinguishability obfuscation ($i\mathcal{O}$) [GGH⁺13b, SW14, PST14]. Hofheinz et al. [HKKW14] construct an adaptively secure bit-fixing PRF, also using $i\mathcal{O}$, and additionally relying on the random-oracle model. We leave the construction of adaptively secure constrained unbounded-length PRFs (for any interesting set of constraints) as a challenging open problem.

Applications. As two applications of our constrained PRFs we show that they directly yield broadcast encryption and ID-NIKE for an unbounded number of parties. In particular, all parameters (private/public key size and for BE also ciphertext overhead) are poly-logarithmic in the number of potential parties (or equivalently, polynomial in the length of the identities). For BE, this has only recently been achieved by Boneh, Waters and Zhandry [BWZ14], who construct a BE scheme supporting n parties directly from $O(\log(n))$ -way multilinear maps. For ID-NIKE, our construction is the first to achieve this; all previous schemes require the maximum size of the group of users agreeing on a key to be fixed at setup, and they have parameters that depend at least linearly on this size.

A circuit-constrained PRF. A first idea for constructing a constrained PRF is to start with a standard PRF F ; given a key k and a set S , we can define a constrained key as a program P_S which on input x checks whether $x \in S$, and if so, outputs $F(k, x)$. We cannot define the constrained key as the program P_S as such, since an adversary could extract the key k from P_S , and hence $F(k, \cdot)$ would not be pseudorandom for $x \notin S$ given P_S .

When S is decided by a circuit, the above issue can be avoided by *obfuscating* P_S before outputting it. The strong notion of *virtual black-box obfuscation*, which requires that an obfuscated program leaks nothing about the program apart from its input/output behavior, is not achievable for general functionalities [BGI⁺12]. We therefore use *indistinguishability obfuscation* ($i\mathcal{O}$), which only guarantees that obfuscations of two circuits (of the same size) that output the same on any input are indistinguishable. A candidate $i\mathcal{O}$ scheme was proposed by Garg et al. [GGH⁺13b]. Although the notion seems weak, it has proven to be surprisingly useful.

Consider a CPRF derived from a puncturable PRF F for which a constrained key k_C for a circuit C is defined as an $i\mathcal{O}$ obfuscation of the circuit P_C , which on input x returns $F(k, x)$ if $C(x) = 1$ and \perp otherwise. In the selective-security game an adversary \mathcal{A} chooses an input x^* , can then ask for constrained keys for circuits C with $C(x^*) = 0$ and must distinguish $F(k, x^*)$ from random. In the security proof we first define a modified game where \mathcal{A} , when asking for a constrained key for a circuit C , is given an $i\mathcal{O}$ obfuscation of a circuit P'_C that outputs $F(k_{x^*}, x)$ if $C(x) = 1$ and \perp otherwise. The difference between P_C and P'_C is that in the latter F is evaluated using a key k_{x^*} that is punctured at x^* .

Recall that the adversary can only submit circuits C with $C(x^*) = 0$ to its oracle. For every such C we have $P_C(x^*) = P'_C(x^*) = \perp$, and on any other input x , P_C and P'_C also return the same output (namely $F(k, x)$ if $C(x) = 1$ and \perp otherwise). By security of $i\mathcal{O}$, obfuscations of P_C and P'_C are thus indistinguishable, which means that the modified game is indistinguishable from the original game. From an adversary \mathcal{A} winning the modified game we obtain an adversary \mathcal{B} that breaks the puncturable PRF F : When \mathcal{A} commits to x^* , \mathcal{B} asks for a punctured key k_{x^*} , which allows \mathcal{B} to answer \mathcal{A} 's constrained-key queries in the modified game. If \mathcal{A} distinguishes $F(k, x^*)$ from random then so does \mathcal{B} .

A TM-constrained PRF. The above construction uses $i\mathcal{O}$ for circuits. Recently, Koppula, Lewko and Waters [KLW15] constructed $i\mathcal{O}$ for Turing machines. However, we cannot simply

replace circuits by TMs in the construction just sketched. In the security proof we need to upper-bound the size of the TM to be obfuscated when we switch from a TM using k to a TM using k_{x^*} . This is however impossible because the size of the underlying punctured key k_{x^*} cannot be a priori bounded for unbounded inputs x^* .

To overcome this problem, we use a collision-resistant hash function H to map long inputs to inputs of fixed length. Concretely, we define our CPRF as $F(k, x) := \text{PF}(k, H(x))$, where PF is a puncturable PRF. Consequently, a constrained key would be an obfuscation of the TM P_M that checks the input legitimacy of x , i.e., whether $M(x) = 1$, and evaluates PF on $H(x)$. In order to give a reduction of security to the puncturable PRF PF , we would, as before, replace the obfuscation of P_M by one of P'_M , which uses $k_{H(x^*)}$ instead of k . While this solves the size-bounding problem, it poses new challenges. Namely, $i\mathcal{O}$ is not sufficient as P_M and P'_M are in general not functionally equivalent: consider a machine M with $M(x^*) = 0$ and $M(x) = 1$ for some x with $H(x) = H(x^*)$; then $P_M(x) = F(k, x)$, whereas $P'_M(x) = \perp$.

DIFFERING-INPUT OBFUSCATION. Instead of $i\mathcal{O}$, we resort to a stronger form of obfuscation. Whereas $i\mathcal{O}$ yields indistinguishable obfuscations of programs that are functionally equivalent, differing-input obfuscation ($di\mathcal{O}$, a.k.a. extractability obfuscation) introduced by [BGI⁺12, BST14] for circuits and later by [BCP14, ABG⁺13] for TMs, guarantees that from an adversary that distinguishes obfuscations of two circuits (or TMs), one can extract an input on which they differ. $di\mathcal{O}$ is a strong assumption and in fact was shown to be implausible to exist [GGHW14]. We will use a weaker assumption suggested by Ishai, Pandey and Sahai [IPS15] and called public-coin $di\mathcal{O}$, for which no such implausibility results are known. Informally, this notion only implies indistinguishability for pairs of programs if it is hard to find an input on which the two programs differ *even when given the randomness used to sample this pair of circuits*.

We replace $i\mathcal{O}$ in our CPRF construction by public-coin $di\mathcal{O}$ for TMs with unbounded inputs from [IPS15] and define a constrained key for M as a $di\mathcal{O}$ obfuscation of P_M . This solution is not without limitations; a constrained key is now a $di\mathcal{O}$ -obfuscated TM and therefore keys are large and their size depends on the running time of the constraining TM M , which we show how to avoid next.

SNARKS. We “outsource” the check of input legitimacy (whether x satisfies $M(x) = 1$) to the party that evaluates the PRF. The latter first computes a succinct non-interactive argument of knowledge (SNARK) of a legitimate x and passes this SNARK to the obfuscated program. A SNARK system is a computationally sound non-interactive proof of knowledge for which proofs are universally succinct. That is, the length of a proof π for a statement η as well as its verification time are bounded by an a-priori-fixed polynomial in the length $|\eta|$ of the statement. In particular, we use a SNARK system for the language $L_{\text{legit}} := \{(H, M, h) \mid \exists x : M(x) = 1 \wedge H(x) = h\}$.

Instead of running M , the program P_M now only needs to verify a SNARK, which can be implemented by a *circuit*; we thus only require obfuscation of circuits. Let P_M be the circuit that on input (h, π) outputs $\text{PF}(k, h)$ iff π is a valid SNARK for (H, M, h) . A constrained key is then a public-coin $di\mathcal{O}$ obfuscation of P_M , whose size only depends on the size of M but not on its running time.

As we use public-coin $di\mathcal{O}$, we require the hash function H to be public-coin [HR04], that is, collision-resistance (CR) must hold when the adversary is given the randomness used to sample a hash function from the family; moreover, the SNARK must be in the common random string model. Such hash functions and SNARKs exist as discussed in [IPS15]. Assuming puncturable PRFs, public-coin CR hash functions, a SNARK for the language L_{legit} and public-coin $di\mathcal{O}$ for circuits, our construction is a TM-constrained PRF for inputs of unbounded length.

$\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{\mathcal{O}, b}(\lambda)$	Oracle $\text{CONSTR}(S)$	Oracle $\text{EVAL}(x)$
$k \leftarrow \text{F.Smp}(1^\lambda); C, E := \emptyset$	If $S \notin \mathcal{S}_\lambda \vee S \cap C \neq \emptyset$	If $x \notin \mathcal{X} \vee x \in C$
$(x^*, st) \leftarrow \mathcal{A}_1^{\mathcal{O}_1}(1^\lambda)$	Return \perp	Return \perp
If $x^* \in E$ then abort	$E := E \cup S$	$E := E \cup \{x\}$
If $b = 1$ then $y := \text{F.Eval}(k, x^*)$	$k_S \leftarrow \text{F.Constr}(k, S)$	$y = \text{F.Eval}(k, x)$
Else $y \leftarrow \mathcal{Y}$	Return k_S	Return y
$C := C \cup \{x^*\}$		
Return $b' \leftarrow \mathcal{A}_2^{\mathcal{O}_2}(st, y)$		

Figure 1: The security game for constrained PRFs.

2 Preliminaries

2.1 Constrained and Puncturable PRFs

Definition 1 (Constrained Functions). *A family of keyed functions $\mathcal{F}_\lambda = \{\text{F}: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}\}$ over a key space \mathcal{K} , a domain \mathcal{X} and a range \mathcal{Y} is efficiently computable if there exist a PPT sampler F.Smp and a deterministic polynomial-time (PT) evaluator F.Eval such that:*

- $k \leftarrow \text{F.Smp}(1^\lambda)$: On input a security parameter λ , F.Smp outputs a key $k \in \mathcal{K}$.
- $\text{F}(k, x) := \text{F.Eval}(k, x)$, for $k \in \mathcal{K}$ and $x \in \mathcal{X}$.

We say \mathcal{F}_λ is constrained w.r.t. a family \mathcal{S}_λ of subsets of \mathcal{X} , with constrained key space $\mathcal{K}_\mathcal{S}$ such that $\mathcal{K}_\mathcal{S} \cap \mathcal{K} = \emptyset$, if F.Eval accepts inputs from $(\mathcal{K} \cup \mathcal{K}_\mathcal{S}) \times \mathcal{X}$ and there exists the following PPT algorithm:

- $k_S \leftarrow \text{F.Constr}(k, S)$: On input a key $k \in \mathcal{K}$ and a description¹ of a set $S \in \mathcal{S}_\lambda$, F.Constr outputs a constrained key $k_S \in \mathcal{K}_\mathcal{S}$ such that

$$\text{F.Eval}(k_S, x) = \begin{cases} \text{F}(k, x) & \text{if } x \in S \\ \perp & \text{otherwise} \end{cases} .$$

Definition 2 (Security of Constrained PRFs). *A family of (efficiently computable) constrained functions $\mathcal{F}_\lambda = \{\text{F}: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}\}$ is selectively pseudorandom, if for every PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in $\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{\mathcal{O}, b}$, defined in Fig. 1, with $\mathcal{O}_1 = \emptyset$ and $\mathcal{O}_2 = \{\text{CONSTR}(\cdot), \text{EVAL}(\cdot)\}$, it holds that*

$$\text{Adv}_{\mathcal{F}, \mathcal{A}}^{\mathcal{O}}(\lambda) := \left| \Pr [\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{\mathcal{O}, 0}(\lambda) = 1] - \Pr [\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{\mathcal{O}, 1}(\lambda) = 1] \right| = \text{negl}(\lambda) . \quad (1)$$

\mathcal{F}_λ is adaptively pseudorandom if the same holds for $\mathcal{O}_1 = \mathcal{O}_2 = \{\text{CONSTR}(\cdot), \text{EVAL}(\cdot)\}$.

Puncturable PRFs [SW14] are a simple type of constrained PRFs whose domain is $\{0, 1\}^n$ for some n , and constrained keys can only be derived for the sets $\{\{0, 1\}^n \setminus \{x_1, \dots, x_m\} \mid x_1, \dots, x_m \in \{0, 1\}^n, m = \text{poly}(\lambda)\}$, i.e., a punctured key $k_{x_1 \dots x_m}$ can evaluate the PRF on all inputs *except* on polynomially many x_1, \dots, x_m . We only require pseudorandomness to hold against selective adversaries.

Definition 3 (Puncturable PRFs [SW14]). *A family of PRFs $\mathcal{F}_\lambda = \{\text{F}: \mathcal{K} \times \{0, 1\}^n \rightarrow \mathcal{Y}\}$ is puncturable if it is constrainable for sets $\{0, 1\}^n \setminus T$, where $T \subseteq \{0, 1\}^n$ is of polynomial size. \mathcal{F}_λ is (selectively) pseudorandom if for every PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in $\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{\text{PCT-}b}(\lambda)$, defined in Fig. 2, we have*

$$\text{Adv}_{\mathcal{F}, \mathcal{A}}^{\text{PCT}}(\lambda) := \left| \Pr [\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{\text{PCT-}0}(\lambda) = 1] - \Pr [\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{\text{PCT-}1}(\lambda) = 1] \right| = \text{negl}(\lambda) .$$

¹As outlined in the introduction, we assume that every set in \mathcal{S} can be specified by a short and efficiently computable predicate.

$\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{\text{PCT-}b}(\lambda)$ $(x^*, T, st) \leftarrow \mathcal{A}_1(1^\lambda); \text{ If } x^* \notin T, \text{ then abort}$ $k \leftarrow \text{F.Smp}(1^\lambda); k_{\overline{T}} \leftarrow \text{F.Constr}(k, \{0, 1\}^n \setminus T)$ $\text{If } b = 1 \text{ then } y := \text{F.Eval}(k, x^*); \text{ else } y \leftarrow \mathcal{Y}$ $\text{Return } b' \leftarrow \mathcal{A}_2^{\text{EVAL}(\cdot)}(st, k_{\overline{T}}, y)$	$\mathbf{Oracle EVAL}(x)$ $\text{If } x = x^*, \text{ return } \perp$ $\text{Return F.Eval}(k, x)$
--	--

Figure 2: The selective-security game for puncturable PRFs.

Puncturable PRFs are easily obtained from prefix-constrained PRFs, which were constructed from the GGM PRF [GGM86] in [BW13, BGI14, KPTZ13].

2.2 Collision-Resistant Hash Functions

A family of hash functions is collision-resistant (CR) if given a uniformly sampled function, it is hard to find inputs on which it collides. It is called public-coin CR if this is hard even when given the coins used to sample the function.

Definition 4 (Public-Coin CR Hash Functions [HR04]). *A family of (efficiently computable) functions $\mathcal{H}_\lambda = \{H: \{0, 1\}^* \rightarrow \{0, 1\}^n\}$, for which Smp samples a random function, is public-coin collision-resistant if for every PPT adversary \mathcal{A} it holds that*

$$\Pr \left[\begin{array}{l} r \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; H := \text{Smp}(1^\lambda, r); \\ (x_1, x_2) \leftarrow \mathcal{A}(1^\lambda, r) \end{array} : \begin{array}{l} H(x_1) = H(x_2) \\ \wedge x_1 \neq x_2 \end{array} \right] = \text{negl}(\lambda) .$$

2.3 Indistinguishability and Differing-Input Obfuscation

As a consequence of the impossibility of virtual black-box obfuscation, Barak et al. [BGI⁺12], proposed two weaker notions: *indistinguishability obfuscation* ($i\mathcal{O}$) and *differing-input obfuscation* ($di\mathcal{O}$). The first, $i\mathcal{O}$, guarantees that obfuscations of equivalent functionalities are computationally indistinguishable.

Definition 5 ($i\mathcal{O}$ [GGH⁺13b]). *A uniform PPT algorithm $i\mathcal{O}$ is an indistinguishability obfuscator for a family of polynomial-size circuits \mathcal{C}_λ , if the following hold:*

- For all $\lambda \in \mathbb{N}$, $C \in \mathcal{C}_\lambda$, and x : $\Pr [\tilde{C} \leftarrow i\mathcal{O}(1^\lambda, C) : C(x) = \tilde{C}(x)] = 1$.
- For every PPT adversary \mathcal{A} and all $C_0, C_1 \in \mathcal{C}_\lambda$ s.t. $\forall x, C_0(x) = C_1(x)$:

$$|\Pr [\mathcal{A}(i\mathcal{O}(1^\lambda, C_0)) = 1] - \Pr [\mathcal{A}(i\mathcal{O}(1^\lambda, C_1)) = 1]| = \text{negl}(\lambda) .$$

Differing-input obfuscation is a stronger notion which requires that for any efficient adversary that distinguishes obfuscations of two functionalities, there exists an efficient *extractor* \mathcal{E} that extracts a point on which the functionalities differ. Ishai, Pandey and Sahai [IPS15] weaken this to *public-coin $di\mathcal{O}$* , where the extractor is given the coins used to sample the functionalities. We will use this concept for circuits, which is formalized by requiring that indistinguishability only holds for circuits output by a sampler Samp for which no differing-input extractor exists.

Definition 6 (Public-Coin Differing-Input Sampler [IPS15]). *A non-uniform PPT sampler Samp is a public-coin differing-input sampler for a polynomial-size family of circuits \mathcal{C}_λ if the output of Samp is in $\mathcal{C}_\lambda \times \mathcal{C}_\lambda$ and for every non-uniform PPT extractor \mathcal{E} it holds that*

$$\Pr [r \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; (C_0, C_1) := \text{Samp}(1^\lambda, r); x \leftarrow \mathcal{E}(1^\lambda, r) : C_0(x) \neq C_1(x)] = \text{negl}(\lambda) . \quad (2)$$

Definition 7 (Public-Coin diO [IPS15]). A uniform PPT algorithm diO is a public-coin differing-input obfuscator for a family of polynomial-size circuits \mathcal{C}_λ if the following hold:

- For all $\lambda \in \mathbb{N}$, $C \in \mathcal{C}_\lambda$ and x : $\Pr [\tilde{C} \leftarrow \text{diO}(1^\lambda, C) : C(x) = \tilde{C}(x)] = 1$.
- For every public-coin differing-input sampler Samp for a family of polynomial-size circuits \mathcal{C}_λ , every non-uniform PPT distinguisher \mathcal{D} and every $\lambda \in \mathbb{N}$:

$$\begin{aligned} & \left| \Pr [r \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; (C_0, C_1) := \text{Samp}(1^\lambda, r); \tilde{C} \leftarrow \text{diO}(1^\lambda, C_0) : 1 \leftarrow \mathcal{D}(r, \tilde{C})] - \right. \\ & \left. \Pr [r \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; (C_0, C_1) := \text{Samp}(1^\lambda, r); \tilde{C} \leftarrow \text{diO}(1^\lambda, C_1) : 1 \leftarrow \mathcal{D}(r, \tilde{C})] \right| = \text{negl}(\lambda) . \end{aligned} \quad (3)$$

A candidate $i\mathcal{O}$ was constructed based on a simplified variant of multilinear maps and proven secure in an idealized model [GGH⁺13b]. The candidate was conjectured to be a $di\mathcal{O}$ for NC^1 [BCP14]. Unfortunately, Garg et al. [GGHW14] present an implausibility result for $di\mathcal{O}$ for arbitrary distributions. However, Ishai et al. [IPS15] argue that current candidate constructions for $i\mathcal{O}$ satisfy their notion of public-coin $di\mathcal{O}$.

2.4 Succinct Non-interactive Arguments of Knowledge

A succinct non-interactive argument of knowledge (SNARK) is a computationally sound non-interactive proof-of-knowledge system for which proofs are universally succinct. A proof π of knowledge of a witness w to a statement η is succinct if the proof length and its verification time are bounded by an a priori fixed polynomial in the statement length $|\eta|$.

Definition 8 (The Universal Relation $\mathcal{R}_{\mathcal{U}}$ [BG08]). The universal relation $\mathcal{R}_{\mathcal{U}}$ is the set of instance/witness pairs of the form $((M, m, t), w)$ where M is a TM accepting an input/witness pair (m, w) within t steps. In particular $|w| \leq t$.

We define SNARK systems in the common-random-string model following Bitansky et al. [BCCT13, BCC⁺14, IPS15] as follows.

Definition 9 (SNARK). A pair of PPT algorithms (Prove, Verify) is a succinct non-interactive argument of knowledge (SNARK) in the common-random-string model for a language \mathcal{L} with witness relation $\mathcal{R} \subseteq \mathcal{R}_{\mathcal{U}}$ if there exist polynomials p, ℓ, q independent of \mathcal{R} such that the following hold:

1. *Completeness:* For every $(\eta = (M, m, t), w) \in \mathcal{R}$ the following holds:

$$\Pr [\text{crs} \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; \pi \leftarrow \text{Prove}(\text{crs}, \eta, w) : \text{Verify}(\text{crs}, \eta, \pi) = 1] = 1 .$$

Moreover, Prove runs in time $q(\lambda, |\eta|, t)$.

2. *(Adaptive) Soundness:* For every PPT adversary \mathcal{A} :

$$\Pr [\text{crs} \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; (\eta, \pi) \leftarrow \mathcal{A}(\text{crs}) : \text{Verify}(\text{crs}, \eta, \pi) = 1 \wedge \eta \notin \mathcal{L}] = \text{negl}(\lambda) .$$

3. *(Adaptive) Argument of knowledge:* For every PPT adversary \mathcal{A} there exists a PPT extractor $\mathcal{E}_{\mathcal{A}}$ such that

$$\Pr \left[\begin{array}{l} \text{crs} \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; r \leftarrow \{0, 1\}^{\text{poly}(\lambda)} \\ (\eta, \pi) := \mathcal{A}(\text{crs}; r); w \leftarrow \mathcal{E}_{\mathcal{A}}(1^\lambda, \text{crs}, r) \end{array} : \begin{array}{l} \text{Verify}(\text{crs}, \eta, \pi) = 1 \\ \wedge (\eta, w) \notin \mathcal{R} \end{array} \right] = \text{negl}(\lambda) .$$

4. *Succinctness:* For all $(\text{crs}, \eta, w) \in \{0, 1\}^{\text{poly}(\lambda)} \times \mathcal{R}$, the length of a proof $\pi \leftarrow \text{Prove}(\text{crs}, \eta, w)$ is bounded by $\ell(\lambda, \log t)$ and the running time of $\text{Verify}(\text{crs}, \eta, \pi)$ is bounded by $p(\lambda + |\eta|) = p(\lambda + |M| + |m| + \log t)$.

Bitansky et al. [BCC⁺14] construct SNARKs for $\mathcal{R}_c \subset \mathcal{R}_{\mathcal{U}}$ where $t = |m|^c$ and c is a constant, based on knowledge-of-exponent assumptions [BCCT13] and extractable collision-resistant hash functions (ECRHF) [BCC⁺14]. These are both non-falsifiable assumptions, but Gentry and Wichs [GW11] prove that SNARKs cannot be built from falsifiable assumptions via black-box reductions. Relying on exponentially hard one-way functions and ECRHF, [BCC⁺14] construct SNARKs for $\mathcal{R}_{\mathcal{U}}$.

3 Constrained PRFs for Unbounded Inputs

As a warm-up we first construct a constrained PRF for sets decided by polynomial-size circuits and show how to extend it to Turing machines in Sect. 3.2.

3.1 A Circuit-Constrained PRF

Our circuit-constrained PRF F uses a puncturable PRF PF with input space $\mathcal{X} = \{0, 1\}^n$. The output of $F(k, x)$ is simply $PF(k, x)$. To constrain F w.r.t. a circuit C , we construct a circuit $P_{k,C}$, which on input x runs C on x and outputs $PF(k, x)$ if $C(x) = 1$, and \perp otherwise. A constrained key k_C for C is then an indistinguishability obfuscation of $P_{k,C}$, i.e., $k_C \leftarrow \text{iO}(1^\lambda, P_{k,C})$.

Construction 1 (Circuit-Constrained PRF). *Let $\mathcal{C}_\lambda = \{C: \{0, 1\}^n \rightarrow \{0, 1\}\}$ be a family of polynomial-size circuits, $\mathcal{PF}_\lambda = \{PF: \mathcal{K} \times \{0, 1\}^n \rightarrow \mathcal{Y}\}$ a family of selectively secure puncturable PRFs, and iO an indistinguishability obfuscator for a family of poly-size circuits \mathcal{P}_λ that contains all circuits defined in (4) for all $C \in \mathcal{C}_\lambda$. We construct a family of PRFs $\mathcal{F}_\lambda = \{F: \mathcal{K} \times \{0, 1\}^n \rightarrow \mathcal{Y}\}$ constrained w.r.t. \mathcal{C}_λ with a constrained-key space \mathcal{K}_C such that $\mathcal{K}_C \cap \mathcal{K} = \emptyset$.²*

$k \leftarrow F.\text{Smp}(1^\lambda)$: Given security parameter λ , output $k \in \mathcal{K}$ as $k \leftarrow PF.\text{Smp}(1^\lambda)$.

$k_C \leftarrow F.\text{Constr}(k, C)$: On input a secret key $k \in \mathcal{K}$ and a description of a circuit $C \in \mathcal{C}_\lambda$, output $k_C \in \mathcal{K}_C$ as $k_C \leftarrow \text{iO}(1^\lambda, P_{k,C})$, with $P_{k,C} \in \mathcal{P}_\lambda$ defined as:

$$P_{k,C}(x) := \begin{cases} PF(k, x) & \text{if } |x| = n \wedge C(x) = 1 \\ \perp & \text{otherwise} \end{cases} . \quad (4)$$

$y := F.\text{Eval}(\kappa, x)$: On input $\kappa \in \mathcal{K} \cup \mathcal{K}_C$ and $x \in \{0, 1\}^n$, do the following:

- If $\kappa \in \mathcal{K}$, output $PF.\text{Eval}(\kappa, x)$.
- If $\kappa \in \mathcal{K}_C$, interpret κ as a circuit and output $\kappa(x)$.

The proof of selective security of \mathcal{F} , as just constructed, is relatively straightforward. Recall that in the selective-security game the adversary \mathcal{A} outputs x^* , then the challenger chooses $k \leftarrow F.\text{Smp}$ and gives \mathcal{A} access to a constrained-key oracle CONSTR , which can be queried on any C with $C(x^*) = 0$. \mathcal{A} must then distinguish $F(k, x^*)$ from random. We modify this game by deriving from k a key k_{x^*} which is punctured at x^* and computing constrained keys as obfuscations of $P_{k_{x^*}, C}$ (defined like $P_{k,C}$ but using k_{x^*} instead of k). Since $PF(k, x) = PF(k_{x^*}, x)$ for all $x \neq x^*$, and since for any circuit C that the adversary can query we have $P_{k,C}(x^*) = P_{k_{x^*}, C}(x^*) = \perp$, the circuits $P_{k_{x^*}, C}$ and $P_{k,C}$ are functionally equivalent, and thus by Def. 5 the two games are indistinguishable. Note that we also need to ensure that these circuits are of the same size, which can be achieved by appropriate padding.

An adversary \mathcal{A} winning the modified game can be translated into an adversary \mathcal{B} against \mathcal{PF} . In the security game for \mathcal{PF} (Fig. 2), \mathcal{B} runs $(x^*, st) \leftarrow \mathcal{A}$ and outputs $(x^*, \{x^*\}, st)$. Given k_{x^*} and y , \mathcal{B} can simulate the modified game and output whatever \mathcal{A} outputs. \mathcal{B} 's probability of breaking the security of \mathcal{PF} is the same as that of \mathcal{A} winning the modified game.

²W.l.o.g. we assume from now on that $\mathcal{K} \cap \mathcal{K}_C = \emptyset$, as this can always be achieved by simply prepending a ‘0’ to elements from \mathcal{K} and a ‘1’ to elements from \mathcal{K}_C .

3.2 A TM-Constrained PRF

In this section we construct a family of constrained PRFs for unbounded inputs whose keys can be constrained to sets decided by Turing machines (TM). As a first attempt, in Construction 1 we could replace C in $P_{k,C}$ with a TM M , yielding a TM $P_{k,M}$. We would thus have to use obfuscation for Turing machines rather than just circuits. However, the problem with this construction is that in the proof we would have to replace the underlying PRF key k with a punctured key k_{x^*} for some x^* whose length is not a priori bounded. It is thus not clear how to pad the original key, which could be done in our previous construction.

To overcome this problem we compress the unbounded input to a fixed length by applying a collision-resistant hash function H to it, that is, we evaluate the PRF on hashed inputs. Moreover, we outsource the check of input legitimacy outside the program $P_{k,M}$ by using a SNARK. In particular, when evaluating the PRF, the user computes a SNARK proving that a given hash is the hash value of a legitimate input. The program $P_{k,M}$ is then only given the *hash* of the input to the PRF and a SNARK proof confirming the legitimacy of a preimage, and evaluates the PRF on the hash if the proof verifies.

Note that $P_{k,M}$ can now be implemented by a circuit, which means that we can avoid obfuscation of Turing machines altogether. In our construction a constrained key k_M for a TM M is a public-coin *diO* obfuscation of a circuit $P_{k,M}$ which is given (h, π) and checks whether π proves knowledge of an x such that $H(x) = h$ and $M(x) = 1$, and if so, evaluates PF on h .

Let us justify the use of (public-coin) *diO* and SNARKs. As for our circuit-constrained PRF, we want to reduce the selective security of the TM-constrained PRF F to the selective security of the underlying puncturable PRF PF . In a first game hop we replace $P_{k,M}$ with $P_{k_{h^*},M}$, which is identical to $P_{k,M}$ except that the key k is replaced with a key k_{h^*} that punctures out $h^* := H(x^*)$. Unfortunately, the two circuits $P_{k,M}$ and $P_{k_{h^*},M}$ are not equivalent: there exists $x \neq x^*$ such that $H(x) = H(x^*)$, and on input $H(x)$, $P_{k,M}$ outputs $PF(k, H(x)) = PF(k, h^*)$ and $P_{k_{h^*},M}$ outputs \perp . We thus cannot use *iO* and hence we use *diO* instead. This requires that it be hard to find an input (h, π) on which the two circuits differ, which means that either π proves a wrong statement or it proves knowledge of some x with $H(x) = H(x^*)$. That is, finding a differing input amounts to either breaking soundness of the SNARK or breaking collision-resistance of H . Since both are hard even for adversaries that know the coins used to sample the hash function or the common random string for the SNARK, it suffices to use public-coin *diO*.

Finally, hash-function collisions are also the reason we need to use SNARKs rather than SNARGs: if an adversary can distinguish obfuscations of $P_{k,M}$ and $P_{k_{h^*},M}$ by finding a collision for H then we need to extract this collision in the security proof, which SNARKs (arguments of *knowledge*) allow us to do.

Definition 10 (R_{legit}). We define the relation $R_{legit} \subset \mathcal{R}_{\mathcal{U}}$, with $\mathcal{R}_{\mathcal{U}}$ defined in Def. 8, to be the set of instance/witness pairs $((H, M), h, t, x)$ such that $M(x) = 1$ and $H(x) = h$ within t steps, and M is a TM and H is a hash function. We let L_{legit} be the language corresponding to R_{legit} . For notational convenience, we abuse the notation and write $((H, M, h), x) \in R_{legit}$ to mean $((H, M), h, t, x) \in R_{legit}$ while implicitly setting $t = 2^\lambda$.

Remark 1. Let $t = 2^\lambda$ in the definition of R_{legit} ; then by succinctness of SNARKs (Def. 9), the length of a SNARK proof is bounded by $\ell(\lambda)$ and its verification time is bounded by $p(\lambda + |M| + |H| + |h|)$, where p, ℓ are a priori fixed polynomials that do not depend on R_{legit} .

Construction 2 (TM-Constrained PRF). Let $\mathcal{PF}_\lambda = \{PF: \mathcal{K} \times \{0, 1\}^n \rightarrow \mathcal{Y}\}$ be a selectively secure puncturable PRF, $\mathcal{H}_\lambda = \{H: \{0, 1\}^* \rightarrow \{0, 1\}^n\}_\lambda$ a family of public-coin CR hash functions, *diO* a public-coin *diO* obfuscator for a family of polynomial-size circuits \mathcal{P}_λ , and SNARK a SNARK system for R_{legit} (cf. Def. 10). We construct a family of PRFs $\mathcal{F}_\lambda = \{F: \mathcal{K} \times \{0, 1\}^* \rightarrow \mathcal{Y}\}$ constrained w.r.t. to any polynomial-size family of TMs \mathcal{M}_λ as follows:

<p>Exp$_{\mathcal{F}, \mathcal{A}}^{(\emptyset, \{\text{CONSTR}, \text{EVAL}\}), b}(\lambda)$</p> <p>$(x^*, st) \leftarrow \mathcal{A}_1(1^\lambda)$ $K \leftarrow \text{F.Smp}(1^\lambda)$ If $b = 1$ $y^* := \text{F.Eval}(K, x^*)$ Else $y^* \leftarrow \mathcal{Y}$ $b' \leftarrow \mathcal{A}_2^{\text{CONSTR}(\cdot), \text{EVAL}(\cdot)}(st, y^*)$ Return b'</p> <p>Oracle $\text{CONSTR}(M)$</p> <p>If $M \notin \mathcal{M}_\lambda \vee M(x^*) = 1$ Return \perp $k_M \leftarrow \text{F.Constr}(K, M)$ Return k_M</p> <p>Oracle $\text{EVAL}(x)$</p> <p>If $x = x^*$ Return \perp $y = \text{F.Eval}(K, x)$ Return y</p>	<p>Exp$_{\mathcal{F}, \mathcal{A}}^{b, (c)}(\lambda) \quad // c \in \{0, 1, 2\}$</p> <p>$(x^*, st) \leftarrow \mathcal{A}_1(1^\lambda)$ $H \leftarrow \text{H.Smp}(1^\lambda)$ $\text{crs} \leftarrow \{0, 1\}^{\text{poly}(\lambda)}$ $k \leftarrow \text{PF.Smp}(1^\lambda)$ $k_{h^*} \leftarrow \text{PF.Constr}(k, \{0, 1\}^n \setminus \{H(x^*)\})$</p> <p>$pp := (H, \text{crs})$ If $b = 1$, $y^* := \text{PF.Eval}(k, H(x^*))$, else $y^* \leftarrow \mathcal{Y}$ $b' \leftarrow \mathcal{A}_2^{\text{CONSTR}(\cdot), \text{EVAL}(\cdot)}(st, y^*)$ Return b'</p> <p>Oracle $\text{CONSTR}(M)$</p> <p>If $M \notin \mathcal{M}_\lambda \vee M(x^*) = 1$ Return \perp If $c = 0$ $P := P_{M, H, \text{crs}, k}$ (as defined in (5)) Else $P := P_{M, H, \text{crs}, k_{h^*}}$ $\tilde{P} \leftarrow \text{diO}(1^\lambda, P)$ Return $k_M := (M, \tilde{P}, pp)$</p> <p>Oracle $\text{EVAL}(x)$</p> <p>If $x = x^*$ Return \perp If $c \leq 1$ $y := \text{PF.Eval}(k, H(x))$ Else If $H(x) = H(x^*)$, abort Else $y := \text{PF.Eval}(k_{h^*}, H(x))$ Return y</p>
---	---

Figure 3: The original security game and hybrids used in the proof of Theorem 1.

$K \leftarrow \text{F.Smp}(1^\lambda)$: On input a security parameter λ , sample $H \leftarrow \text{H.Smp}(1^\lambda)$, $\text{crs} \leftarrow \{0, 1\}^{\text{poly}(\lambda)}$, $k \leftarrow \text{PF.Smp}(1^\lambda)$, set $pp := (H, \text{crs})$ and return $K := (k, pp)$.

$k_M \leftarrow \text{F.Constr}(K, M)$: On input $K = (k, pp = (H, \text{crs}))$ and $M \in \mathcal{M}_\lambda$, set

$$P_{M, H, \text{crs}, k}(h, \pi) := \begin{cases} \text{PF.Eval}(k, h) & \text{if } \text{SNARK.Verify}(\text{crs}, (H, M, h), \pi) = 1 \\ \perp & \text{otherwise} \end{cases} \quad (5)$$

and compute $\tilde{P} \leftarrow \text{diO}(1^\lambda, P_{M, H, \text{crs}, k})$. Return $k_M := (M, \tilde{P}, (H, \text{crs}))$.

$y := \text{F.Eval}(\kappa, x)$: On input $\kappa \in \mathcal{K} \cup \mathcal{K}_\mathcal{M}$ and $x \in \{0, 1\}^*$, do the following:

- If $\kappa \in \mathcal{K}$, $\kappa = (k, (H, \text{crs}))$: return $\text{PF.Eval}(k, H(x))$.
- If $\kappa \in \mathcal{K}_\mathcal{M}$, $\kappa = (M, \tilde{P}, (H, \text{crs}))$: if $M(x) = 1$, let $h := H(x)$ (thus $((H, M, h), x) \in R_{\text{legit}}$), compute $\pi \leftarrow \text{SNARK.Prove}(\text{crs}, (H, M, h), x)$, interpret \tilde{P} as a circuit and return $\tilde{P}(h, \pi)$.

Remark 2. Note that \mathcal{P}_λ is in fact a family of circuits with an input length $n + |\pi|$ where $|\pi|$ is upper bounded by $\ell(\lambda)$ even for an exponentially long x (cf. Remark 1).

Theorem 1. \mathcal{F}_λ of Construction 2 is a selectively secure family of constrained PRFs with input space $\{0, 1\}^*$ for which constrained keys can be derived for any set that can be decided by a polynomial-size Turing machine.

Proof. Let \mathcal{A} be an arbitrary PPT adversary for game $\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{(\emptyset, \{\text{CONSTR}, \text{EVAL}\}), b}(\lambda)$, as defined in Fig. 3, which we abbreviate as \mathbf{Exp}^b for simplicity. We need to show that \mathbf{Exp}^0 and \mathbf{Exp}^1 are indistinguishable. Our proof will be by game hopping and we define a series of hybrid games

$\mathbf{Exp}^{b,(0)} := \mathbf{Exp}^b$, $\mathbf{Exp}^{b,(1)}$ and $\mathbf{Exp}^{b,(2)}$, which are all defined in Fig. 3. We show that for $b = 0, 1$ and $c = 0, 1$ the games $\mathbf{Exp}^{b,(c)}$ and $\mathbf{Exp}^{b,(c+1)}$ are indistinguishable and that $\mathbf{Exp}^{0,(2)}$ and $\mathbf{Exp}^{1,(2)}$ are also indistinguishable, which concludes the proof.

$\mathbf{Exp}^{b,(0)}$ is the original game $\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{b, (\emptyset, \{\text{CONSTR}, \text{EVAL}\})}(\lambda)$ for Construction 2.

$\mathbf{Exp}^{b,(1)}$ differs from $\mathbf{Exp}^{b,(0)}$ by replacing the full key of the puncturable PRF PF, with one that is punctured at $H(x^*)$ in the definition of P .

$\mathbf{Exp}^{b,(2)}$ differs from $\mathbf{Exp}^{b,(1)}$ by answering EVAL queries using the punctured key k_{h^*} and aborting whenever the query is a collision with x^* for H .

The only difference between $\mathbf{Exp}^{b,(0)}$ and $\mathbf{Exp}^{b,(1)}$ is the definition of the circuits P that are obfuscated when the CONSTR oracle is queried. In $\mathbf{Exp}^{b,(0)}$ the circuit P is defined as in (5), with $k \leftarrow \text{PF.Smp}(1^\lambda)$. In $\mathbf{Exp}^{b,(1)}$, the key k is replaced by $k_{h^*} \leftarrow \text{PF.Constr}(k, \{0, 1\}^n \setminus \{H(x^*)\})$, a key that punctures out $H(x^*)$. By a hybrid argument there must exist some query (say the i th for M_i) where the adversary distinguishes a $di\mathcal{O}$ obfuscation of $P_{M_i, H, \text{crs}, k}$ from one of $P_{M_i, H, \text{crs}, k_{h^*}}$. Thus, there exists a $di\mathcal{O}$ extractor that extracts an input $(\hat{h}, \hat{\pi})$ on which $P_{M_i, H, \text{crs}, k}$ and $P_{M_i, H, \text{crs}, k_{h^*}}$ differ.

By correctness of PF, the circuits only differ on inputs $(\hat{h}, \hat{\pi})$, where

$$\hat{h} = H(x^*) \text{ ,} \quad (6)$$

as that is where the punctured key behaves differently. Moreover, the extracted proof $\hat{\pi}$ must be valid for (H, M_i, \hat{h}) , as otherwise both circuits output \perp . By SNARK extractability, we can extract a witness \hat{x} for $(H, M_i, \hat{h}) \in L_{\text{legit}}$, that is, (i) $M_i(\hat{x}) = 1$ and (ii) $H(\hat{x}) = \hat{h}$. Since M_i is a legitimate query, we have $M_i(x^*) = 0$, which together with (i) implies $\hat{x} \neq x^*$. On the other hand, (ii) and (6) imply $H(\hat{x}) = H(x^*)$. Together, this means (\hat{x}, x^*) is a collision for H .

Proposition 1. *For $b = 0, 1$, $\mathbf{Exp}^{b,(0)}$ and $\mathbf{Exp}^{b,(1)}$ are computationally indistinguishable if $di\mathcal{O}$ is a public-coin differing-input obfuscator and \mathcal{H} is public-coin collision-resistant.*

For the game hop from games $\mathbf{Exp}^{b,(1)}$ to $\mathbf{Exp}^{b,(2)}$, indistinguishability follows directly from collision resistance of \mathcal{H} , as the only difference is that $\mathbf{Exp}^{b,(2)}$ aborts when \mathcal{A} finds a collision.

Proposition 2. *For $b = 0, 1$, $\mathbf{Exp}^{b,(1)}$ and $\mathbf{Exp}^{b,(2)}$ are computationally indistinguishable for if \mathcal{H} is collision-resistant.*

We have now reached a game, $\mathbf{Exp}^{b,(2)}$, in which the key k is only used to create a punctured key k_{h^*} . The experiment can thus be simulated by an adversary \mathcal{B} against selective security of \mathcal{PF} , which first asks for a key for the set $\{0, 1\}^n \setminus \{H(x^*)\}$ and then uses \mathcal{A} to distinguish $y^* = \text{PF.Eval}(k, H(x^*))$ from random.

Proposition 3. *$\mathbf{Exp}^{0,(2)}$ and $\mathbf{Exp}^{1,(2)}$ are indistinguishable if \mathcal{PF} is a selectively secure family of puncturable PRFs.*

Theorem 1 now follows from Propositions 1, 2 and 3, whose proofs can be found in Appendices A.1, A.2 and A.3. \square

4 Applications

As a first application of TM-constrained PRFs we construct the first identity-based non-interactive key-agreement protocol with no a priori bound on the number of parties agreeing on a shared key. Our second application is a broadcast encryption scheme where during setup the number of potential receivers need not be known.

4.1 ID-Based Non-interactive Key Exchange for Unbounded Groups

Identity-based non-interactive key exchange (ID-NIKE) [SOK00] allows users to compute shared keys without any interaction—it suffices to know the identity of the users one wants to share a key with. In our construction a user can compute a shared key for any group of users and there is no a priori bound on the size of these groups. We generalize the construction of [BW13, HKKW14], in which identities are elements from $\{0, 1\}^\ell$ and the system is set up by creating a secret key msk for a constrained PRF. A key for a group of users $\{id_1, \dots, id_n\}$ is defined as $F.Eval(msk, x)$, where $x = id_1 \parallel \dots \parallel id_n$ and we assume identities are always ordered lexicographically. Boneh and Zhandry [BZ14] suggest a similar scheme (assuming indistinguishability obfuscation) which does not require any setup, but where each party publishes a value, and a key for a set of parties can be computed when knowing this values for each party in the set.

Since in the previous constructions the CPRF is set up for a fixed input length m there is an a-priori-fixed maximum number of users that can share a key, namely m/ℓ . As a user's id could appear in any position of the string x , the owner of id is given constrained keys for the sets $id \parallel ?^{(n-1)\ell} := \{id \parallel z \mid z \in \{0, 1\}^{(n-1)\ell}\}, ?^\ell \parallel id \parallel ?^{(n-2)\ell}, \dots, ?^{(n-1)\ell} \parallel id$. These keys then allow the user to compute the key for any set which she is part of.

We generalize this to sets of users of unbounded size. Again, a key for a (lexicographically ordered) set $\{id_1, \dots, id_n\}$ is defined as $F.Eval(msk, id_1 \parallel \dots \parallel id_n)$, but now n can be arbitrary and is not fixed in advance. In order to let a user with identity id compute the keys of the sets which she is part of—but not anything else—, she is given a constrained key for the following Turing machine M_{id} : on input $x \in \{0, 1\}^*$, machine M_{id} outputs 1 if and only if id is a substring of x , which starts at position $i \cdot \ell + 1$, for some $i \geq 0$, that is, at position 1 or $\ell + 1$ or $2\ell + 1$, etc.

ID-NIKE. An (unbounded) ID-NIKE scheme consists of three algorithms:

- $(pp, msk) \leftarrow \text{Setup}(1^\lambda)$: On input λ , output public parameters pp and a master secret key msk .
- $sk_{id} \leftarrow \text{Extract}(msk, id)$: On input msk and $id \in \{0, 1\}^\ell$, output a secret key sk_{id} .
- $k_{\mathcal{I}} \leftarrow \text{KeyGen}(pp, sk_{id}, \mathcal{I})$: On input pp , a key sk_{id} for id and a list $\mathcal{I} \subseteq \{0, 1\}^\ell$ of n (for arbitrary n) users with $id \in \mathcal{I}$, output a shared key $k_{\mathcal{I}}$.

Correctness is defined as follows: for all $id, id' \in \{0, 1\}^\ell$, all $\mathcal{I} \subseteq \{0, 1\}^\ell$ with $id, id' \in \mathcal{I}$, all $(pp, msk) \leftarrow \text{Setup}(1^\lambda)$, $sk_{id} \leftarrow \text{Extract}(msk, id)$ and $sk_{id'} \leftarrow \text{Extract}(msk, id')$, it holds that $\text{KeyGen}(pp, sk_{id}, \mathcal{I}) = \text{KeyGen}(pp, sk_{id'}, \mathcal{I})$.

Following [PS09], we define security via a game where an adversary can obtain secret keys sk_{id} for identities of his choice and can query secret keys $k_{\mathcal{I}}$ for sets \mathcal{I} of his choice. The scheme is secure if the adversary cannot distinguish a key $k_{\mathcal{I}^*}$ for a set \mathcal{I}^* of his choice from random, where we must have $id \notin \mathcal{I}^*$ for all id for which the adversary queried key extraction, and $\mathcal{I}^* \neq \mathcal{I}$ for all \mathcal{I} for which the adversary queried a shared key. We prove that our scheme satisfies the selective variant of this definition, where the adversary must output \mathcal{I}^* before getting access to its oracles.

ID-NIKE from constrained PRFs for unbounded inputs. Our unbounded ID-NIKE is obtained from a TM-constrained PRF ($F.Smp, F.Constr, F.Eval$) as follows.

- $\text{Setup}(1^\lambda)$: Return $msk \leftarrow F.Smp(1^\lambda)$. (There are no parameters pp .)
- $\text{Extract}(msk, id)$: On input $id \in \{0, 1\}^\ell$ do the following: define a Turing machine M_{id} that on input a string $x \in \{0, 1\}^*$ outputs 1 iff x is of the form $x' \parallel id \parallel x''$ with $x' \in \{0, 1\}^{n' \cdot \ell}$ and $x'' \in \{0, 1\}^{n'' \cdot \ell}$ for some $n', n'' \in \mathbb{N}$; return $sk_{id} \leftarrow F.Constr(msk, M_{id})$.
- $\text{KeyGen}(sk_{id}, \mathcal{I})$: If $\mathcal{I} = \{id_1, \dots, id_n\} \subseteq \{0, 1\}^\ell$ for some n and $id \in \mathcal{I}$ then define $x := id_{i_1} \parallel \dots \parallel id_{i_n}$, with $id_{i_j} < id_{i_{j+1}}$ for all j , and output $k_{\mathcal{I}} := F.Eval(sk_{id}, x)$; else output \perp .

Correctness of our scheme follows from correctness of the underlying constrained PRF. Selective security of the ID-NIKE follows from selective security of the CPRF (Def. 2). Given an adversary \mathcal{A} against the ID-NIKE, we construct an adversary \mathcal{B} against the CPRF. First \mathcal{B} runs \mathcal{A} to obtain \mathcal{I}^* and sends x^* , the concatenation of the lexicographically ordered elements of \mathcal{I}^* , to its challenger.

\mathcal{B} answers \mathcal{A} 's queries as follows: When \mathcal{A} queries a secret key for $id \in \mathcal{I}^*$ or the shared key for \mathcal{I}^* then reply with \perp . On a legal secret-key query for id , construct a Turing machine M_{id} as in the definition of Extract, query the CONSTR oracle on M_{id} and forward the reply to \mathcal{A} . When \mathcal{A} queries a shared key for a set $\mathcal{I} \neq \mathcal{I}^*$, construct x as in KeyGen, query EVAL on x and forward the reply.

Note that \mathcal{B} makes no illegal queries (any queried M evaluates x^* to 0 and x^* is never queried to EVAL) and perfectly simulates the game for \mathcal{A} . When \mathcal{B} receives a value y which is either $\text{F.Eval}(msk, x^*)$ or random, it forwards y as the challenge key $k_{\mathcal{I}^*}$ to \mathcal{A} and outputs whatever \mathcal{A} returns. \mathcal{B} thus breaks the CPRF with the same probability as \mathcal{A} breaking the ID-NIKE, which concludes the proof.

4.2 Broadcast Encryption to Unbounded Number of Users

As a second application we use a constrained PRF for unbounded inputs to construct a broadcast encryption (BE) [FN94] where there is no limit on the number of receivers. We start with defining *dynamic* BE, which allows users to join the system after it is set up. Each user is identified by a consecutive number i .

A broadcast encryption scheme \mathcal{BE} for a symmetric encryption scheme (enc, dec) with key space \mathcal{K}_{sym} consists of the following four PPT algorithms:

- $(bk, msk) \leftarrow \text{Setup}(1^\lambda)$: On input a security parameter λ , output a broadcast key bk and a master secret key msk , used to enroll new members in the system.
- $sk_i \leftarrow \text{KeyGen}(msk, i)$: On input a master key msk and a member identity i , output sk_i , a secret key for member i .
- $(hdr, K) \leftarrow \text{Encrypt}(bk, S)$: On input a set $S \subseteq \mathbb{N}$ and a broadcast key bk , output a header hdr and a key $K \in \mathcal{K}_{\text{sym}}$. (A message m is then broadcast as $(S, hdr, \text{enc}(K, m))$.)
- $K \leftarrow \text{Decrypt}(i, sk_i, S, hdr)$: On input a member identity i , an associated secret key sk_i , a set $S \subseteq \mathbb{N}$ and a header hdr , if $i \in S$ then output a symmetric key $K \in \mathcal{K}_{\text{sym}}$. (A broadcast (S, hdr, C) can then be decrypted via $m \leftarrow \text{dec}(K, C)$.)

Like Boneh and Waters [BW13], whose construction we build on, we will construct a *secret-key* BE scheme, where bk must only be known to the broadcaster. Correctness of a BE scheme is defined as follows: for all $S \subseteq \mathbb{N}$, $i \in S$, all $(bk, msk) \leftarrow \text{Setup}(1^\lambda)$, $sk_i \leftarrow \text{KeyGen}(msk, i)$ and $(hdr, K) \leftarrow \text{Encrypt}(bk, S)$, we have $K \leftarrow \text{Decrypt}(i, sk_i, S, hdr)$.

Selective security is defined via the following game $\text{Exp}^{\text{BE-}b}$ for an adversary \mathcal{A} :

$\text{Exp}_{\mathcal{BE}, \mathcal{A}}^{\text{BE-}b}(\lambda)$ $(bk, msk) \leftarrow \text{Setup}(1^\lambda)$ $(S^*, st) \leftarrow \mathcal{A}_1(1^\lambda)$ $(hdr^*, K^*) \leftarrow \text{Encrypt}(bk, S^*)$ <p>If $b = 0$ then $K^* \leftarrow \mathcal{K}_{\text{sym}}$</p> $b' \leftarrow \mathcal{A}_2^{\text{KEY}(\cdot), \text{ENCRYPT}(\cdot)}(st, (hdr^*, K^*))$ $\text{Return } b'$	$\text{Oracle KEY}(i)$ <p>If $i \in S^*$</p> $\text{Return } \perp$ $sk_i \leftarrow \text{KeyGen}(msk, i)$ $\text{Return } sk_i$	$\text{Oracle ENCRYPT}(S)$ <p>If $S = S^*$</p> $\text{Return } \perp$ $(hdr, K) \leftarrow \text{Encrypt}(bk, S)$ $\text{Return } (hdr, K)$
--	--	--

We say \mathcal{BE} is secure if

$$\text{Adv}_{\mathcal{BE}, \mathcal{A}}^{\text{BE}}(\lambda) := |\Pr[\text{Exp}_{\mathcal{BE}, \mathcal{A}}^{\text{BE-}0}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{BE}, \mathcal{A}}^{\text{BE-}1}(\lambda) = 1]| = \text{negl}(\lambda).$$

BE from constrained PRFs for unbounded inputs. For a finite set $S \subseteq \mathbb{N}$, we define the characteristic vector χ_S as the binary vector whose length equals the largest element in S and whose entry at position i is 1 iff $i \in S$. Let (enc, dec) be a symmetric encryption scheme with key space \mathcal{K}_{sym} . Let $\mathcal{F} = \{F: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}\}$ be a constrained PRF with input space $\mathcal{X} = \{0, 1\}^*$ and range $\mathcal{Y} = \mathcal{K}_{\text{sym}}$ for which constrained keys k_i for the following set can be computed:

$$S_i := \{x \in \{0, 1\}^* \mid x_i = 1\} . \quad (7)$$

(As S_i can be decided by a polynomial-time Turing machine, our construction from Sect. 3.2 can be used.) We define a broadcast encryption scheme \mathcal{BE} with *optimal ciphertext length* (that is, the header is empty: $\text{hdr} = \emptyset$) as follows:

- **Setup**(1^λ): Generate $k \leftarrow \text{F.Smp}(1^\lambda)$ and return $bk := k$, $msk := k$.
- **KeyGen**(msk, i): Return $k_i \leftarrow \text{F.Constr}(msk, S_i)$ with S_i as in (7).
- **Encrypt**(bk, S): Let $\chi_S \in \{0, 1\}^*$ be the characteristic vector of S ; compute $K \leftarrow \text{F.Eval}(bk, \chi_S)$ and output (\emptyset, K) .
- **Decrypt**(i, sk_i, S, hdr): With χ_S as above, output $K \leftarrow \text{F.Eval}(sk_i, \chi_S)$.

Correctness of \mathcal{BE} follows from correctness of \mathcal{F} ; security follows by reduction to selective pseudorandomness of \mathcal{F} . Let \mathcal{A} be a PPT adversary that breaks security of \mathcal{BE} ; then we construct a PPT algorithm $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ that breaks \mathcal{F} with the same probability:

$\mathcal{B}_1(1^\lambda)$ <ul style="list-style-type: none"> – $(S^*, st_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$. – Let x^* be the characteristic string of S^*. – Return $(x^*, st_{\mathcal{A}})$. 	$\mathcal{B}_2^{\text{CONSTR}(\cdot), \text{EVAL}(\cdot)}(st, K^*)$ <ul style="list-style-type: none"> – $b' \leftarrow \mathcal{A}_2^{\text{KEY}(\cdot), \text{ENCRYPT}(\cdot)}(st, (\emptyset, K^*))$; – simulate KEY(i): define S_i as in (7); query $k_i \leftarrow \text{CONSTR}(S_i)$; reply k_i; – simulate ENCRYPT(S): define the characteristic vector $\chi_S \in \{0, 1\}^*$ of S; query $K \leftarrow \text{EVAL}(\chi_S)$; reply (\emptyset, K). – Return b'.
--	---

By construction, we have $\mathbf{Exp}_{\mathcal{F}, \mathcal{B}}^{(\emptyset, \{\text{CONSTR}, \text{EVAL}\}), b} = \mathbf{Exp}_{\mathcal{BE}, \mathcal{A}}^{\text{BE-}b}$, which proves the claim.

References

- [ABG⁺13] Prabhajan Ananth, Dan Boneh, Sanjam Garg, Amit Sahai, and Mark Zhandry. Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689, 2013. <http://eprint.iacr.org/2013/689>.
- [BCC⁺14] Nir Bitansky, Ran Canetti, Alessandro Chiesa, Shafi Goldwasser, Huijia Lin, Aviad Rubinfeld, and Eran Tromer. The hunting of the SNARK. *IACR Cryptology ePrint Archive*, 2014:580, 2014.
- [BCCT13] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for SNARKS and proof-carrying data. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 111–120. ACM Press, June 2013.
- [BCP14] Elette Boyle, Kai-Min Chung, and Rafael Pass. On extractability obfuscation. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 52–73. Springer, Heidelberg, February 2014.
- [BG08] Boaz Barak and Oded Goldreich. Universal arguments and their applications. *SIAM J. Comput.*, 38(5):1661–1694, 2008.

- [BGI⁺12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6:1–6:48, May 2012.
- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, Heidelberg, March 2014.
- [BGW05] Dan Boneh, Craig Gentry, and Brent Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 258–275. Springer, Heidelberg, August 2005.
- [BH08] Dan Boneh and Michael Hamburg. Generalized identity based and broadcast encryption schemes. In Josef Pieprzyk, editor, *ASIACRYPT 2008*, volume 5350 of *LNCS*, pages 455–470. Springer, Heidelberg, December 2008.
- [BST14] Mihir Bellare, Igors Stepanovs, and Stefano Tessaro. Poly-many hardcore bits for any one-way function and a framework for differing-inputs obfuscation. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 102–121. Springer, Heidelberg, December 2014.
- [BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Heidelberg, December 2013.
- [BWZ14] Dan Boneh, Brent Waters, and Mark Zhandry. Low overhead broadcast encryption from multilinear maps. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 206–223. Springer, Heidelberg, August 2014.
- [BZ14] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 480–499. Springer, Heidelberg, August 2014.
- [CLT13] Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 476–493. Springer, Heidelberg, August 2013.
- [FHPS13] Eduarda S. V. Freire, Dennis Hofheinz, Kenneth G. Paterson, and Christoph Striecks. Programmable hash functions in the multilinear setting. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 513–530. Springer, Heidelberg, August 2013.
- [FKPR14] Georg Fuchsbauer, Momchil Konstantinov, Krzysztof Pietrzak, and Vanishree Rao. Adaptive security of constrained PRFs. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 82–101. Springer, Heidelberg, December 2014.
- [FN94] Amos Fiat and Moni Naor. Broadcast encryption. In Douglas R. Stinson, editor, *CRYPTO’93*, volume 773 of *LNCS*, pages 480–491. Springer, Heidelberg, August 1994.
- [GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 1–17. Springer, Heidelberg, May 2013.
- [GGH⁺13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.
- [GGHW14] Sanjam Garg, Craig Gentry, Shai Halevi, and Daniel Wichs. On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 518–535. Springer, Heidelberg, August 2014.

- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, October 1986.
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108. ACM Press, June 2011.
- [HKKW14] Dennis Hofheinz, Akshay Kamath, Venkata Koppula, and Brent Waters. Adaptively secure constrained pseudorandom functions. Cryptology ePrint Archive, Report 2014/720, 2014. <http://eprint.iacr.org/>.
- [HKW14] Susan Hohenberger, Venkata Koppula, and Brent Waters. Adaptively secure puncturable pseudorandom functions in the standard model. Cryptology ePrint Archive, Report 2014/521, 2014. <http://eprint.iacr.org/2014/521>.
- [HR04] Chun-Yuan Hsiao and Leonid Reyzin. Finding collisions on a public road, or do secure hash functions need secret coins? In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 92–105. Springer, Heidelberg, August 2004.
- [IPS15] Yuval Ishai, Omkant Pandey, and Amit Sahai. Public-coin differing-inputs obfuscation and its applications. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 668–697. Springer, Heidelberg, March 2015.
- [KLW15] Venkata Koppula, Allison Bishop Lewko, and Brent Waters. Indistinguishability obfuscation for turing machines with unbounded memory. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 419–428. ACM Press, June 2015.
- [KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 669–684. ACM Press, November 2013.
- [LSS14] Adeline Langlois, Damien Stehlé, and Ron Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 239–256. Springer, Heidelberg, May 2014.
- [PPS11] Duong Hieu Phan, David Pointcheval, and Mario Strefler. Security notions for broadcast encryption. In Javier Lopez and Gene Tsudik, editors, *ACNS 11*, volume 6715 of *LNCS*, pages 377–394. Springer, Heidelberg, June 2011.
- [PS09] Kenneth G. Paterson and Sriramkrishnan Srinivasan. On the relations between non-interactive key distribution, identity-based encryption and trapdoor discrete log groups. *Des. Codes Cryptography*, 52(2):219–241, 2009.
- [PST14] Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation from semantically-secure multilinear encodings. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 500–517. Springer, Heidelberg, August 2014.
- [SOK00] Ryuichi Sakai, Kiyoshi Ohgishi, and Masao Kasahara. Cryptosystems based on pairing. In *SCIS 2000*, Okinawa, Japan, January 2000.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014.
- [YFDL04] Danfeng Yao, Nelly Fazio, Yevgeniy Dodis, and Anna Lysyanskaya. ID-based encryption for complex hierarchies with applications to forward security and broadcast encryption. In Vijayalakshmi Atluri, Birgit Pfitzmann, and Patrick McDaniel, editors, *ACM CCS 04*, pages 354–363. ACM Press, October 2004.

A Proofs

A.1 Proof of Proposition 1

Let $b \in \{0, 1\}$ be arbitrarily fixed. Assume towards contradiction that there exists a PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ that distinguishes $\mathbf{Exp}^{b,(0)}$ and $\mathbf{Exp}^{b,(1)}$ with non-negligible probability, i.e., there exists a polynomial $p(\cdot)$ such that for infinitely many λ

$$|\Pr[\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{b,(0)}(\lambda) = 1] - \Pr[\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{b,(1)}(\lambda) = 1]| \geq \frac{1}{p(\lambda)}. \quad (8)$$

We use \mathcal{A} to construct a series of PPT adversaries $\mathcal{D}^{(1)}, \mathcal{D}^{(2)}, \dots$, one of which non-negligibly distinguishes public-coin $di\mathcal{O}$ -obfuscations.

Concretely, we construct a series of hybrids between $\mathbf{Exp}^{b,(0)}$ and $\mathbf{Exp}^{b,(1)}$ as follows. Let $q = q(\lambda)$ be a polynomial upper bound on the total number of constraining queries \mathcal{A} makes. Define the i -th hybrid $\mathbf{Exp}^{b,(0,i)}$ like $\mathbf{Exp}^{b,(0)}$, except that the first i constraining queries are answered by using the punctured key k_{h^*} of the puncturable PRF, and all remaining queries are answered by using private key k . By construction, we have $\mathbf{Exp}^{b,(0,0)} = \mathbf{Exp}^{b,(0)}$ and $\mathbf{Exp}^{b,(0,q)} = \mathbf{Exp}^{b,(1)}$.

We use \mathcal{A} to construct a PPT adversary $\mathcal{D}^{(i)}$ which distinguishes public-coin $di\mathcal{O}$ -obfuscations. We define the series of public-coin differing-input samplers $\mathbf{Samp}^{(i)}$ as follows:

$$\begin{aligned} & \underline{\mathbf{Samp}^{(i)}}(1^\lambda, r := r_{\mathcal{A}} \| r_H \| r_s \| r_k \| r_c \| r_y \| r_1 \| \dots \| r_{i-1}) \quad // \quad r \leftarrow \{0, 1\}^{\text{poly}(\lambda)}. \\ & - (x^*, st_{\mathcal{A}}) := \mathcal{A}_1(1^\lambda; r_{\mathcal{A}}). \\ & - H := \text{H.Smp}(1^\lambda; r_H) \text{ and } h^* := H(x^*). \\ & - crs := r_s. \\ & - k := \text{PF.Smp}(1^\lambda; r_k). \\ & - k_{h^*} := \text{PF.Constr}(k, \{0, 1\}^n \setminus \{h^*\}; r_c). \\ & - \text{If } b = 1 \text{ then } y^* := \text{PF.Eval}(k, H(x^*)) \\ & - \text{Else } y^* \leftarrow \mathcal{Y}, \text{ using randomness } r_y. \\ & - b' \leftarrow \mathcal{A}_2^{\text{CONSTR}(\cdot), \text{EVAL}(\cdot)}(st_{\mathcal{A}}, y^*; r_{\mathcal{A}}); \\ & \quad - \text{simulate the } j\text{-th query } \text{CONSTR}(M_j): \\ & \quad \quad - \text{if } M_j \notin \mathcal{M}_\lambda \vee M_j(x^*) = 1, \text{ reply } \perp; \\ & \quad \quad - \text{if } j < i, \text{ compute } P = P_{M_j, H, crs, k_{h^*}} \text{ as defined in (5)} \\ & \quad \quad \quad \text{and } \tilde{P} := \text{diO}(1^\lambda, P; r_j) \text{ and return } k_M := (M_j, \tilde{P}, (H, crs)). \\ & \quad \quad - \text{if } j = i, \text{ stop and output } (P_0 := P_{M_i, H, crs, k}, P_1 := P_{M_i, H, crs, k_{h^*}}). \\ & \quad - \text{simulate } \text{EVAL}(x): \\ & \quad \quad \text{if } x = x^*, \text{ reply } \perp; \text{ else reply } y := \text{PF.Eval}(k, H(x)). \end{aligned}$$

We now define a public-coin $di\mathcal{O}$ distinguisher $\mathcal{D}^{(i)}$ as follows:

$$\begin{aligned} & \underline{\mathcal{D}^{(i)}}(r, \tilde{P}) \quad // \quad \tilde{P} \text{ is either } \tilde{P}_0 \leftarrow \text{diO}(1^\lambda, P_0) \text{ or } \tilde{P}_1 \leftarrow \text{diO}(1^\lambda, P_1). \\ & - \text{Using randomness } r \text{ run all steps of } \mathbf{Samp}^{(i)}. \\ & - \text{Answer } \mathcal{A}'\text{s } i\text{-th } \text{CONSTR} \text{ query for } M_i \text{ with } (M_i, \tilde{P}). \\ & - \text{Simulate the remaining } \text{CONSTR}(M_j) \text{ as follows:} \\ & \quad - \text{if } M_j \notin \mathcal{M}_\lambda \vee M_j(x^*) = 1, \text{ reply } \perp; \\ & \quad - \text{compute } P_j = P_{M_j, H, crs, k} \text{ as defined in (5) and } \tilde{P}_j \leftarrow \text{diO}(1^\lambda, P_j); \\ & \quad - \text{return } k_M := (M_j, \tilde{P}_j, (H, crs)). \\ & - \text{When } \mathcal{A}_2 \text{ outputs } b', \text{ output } b'. \end{aligned}$$

If \tilde{P} is a public-coin $di\mathcal{O}$ obfuscation of P_0 then $\mathcal{D}^{(i)}$ simulates $\mathbf{Exp}^{b,(0,i-1)}$ for \mathcal{A} , while if \tilde{P} is an obfuscation of P_1 then it simulates $\mathbf{Exp}^{b,(0,i)}$ for \mathcal{A} . If \mathcal{A} distinguishes $\mathbf{Exp}^{b,(0)} = \mathbf{Exp}^{b,(0,0)}$ from $\mathbf{Exp}^{b,(1)} = \mathbf{Exp}^{b,(0,q)}$ with probability $1/p(\lambda)$, then by a hybrid argument for some i , \mathcal{A}

distinguishes $\mathbf{Exp}^{b,(0,i-1)}$ and $\mathbf{Exp}^{b,(0,i)}$ with probability $1/(p(\lambda)q(\lambda))$. Thus infinitely many λ , it holds that

$$\begin{aligned} \frac{1}{p(\lambda)q(\lambda)} &\leq \left| \Pr[\mathbf{Exp}_{\mathcal{A}}^{b,(0,i)}(\lambda) = 1] - \Pr[\mathbf{Exp}_{\mathcal{A}}^{b,(0,i-1)}(\lambda) = 1] \right| = \\ &\left| \Pr \left[r \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; (P_0, P_1) := \text{Samp}(1^\lambda, r); \tilde{P}_0 \leftarrow \text{diO}(1^\lambda, P_0) : 1 \leftarrow \mathcal{D}(r, \tilde{P}_0) \right] - \right. \\ &\left. \Pr \left[r \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; (P_0, P_1) := \text{Samp}(1^\lambda, r); \tilde{P}_1 \leftarrow \text{diO}(1^\lambda, P_1) : 1 \leftarrow \mathcal{D}(r, \tilde{P}_1) \right] \right|. \end{aligned}$$

By security of diO (Def. 7), this means that $\text{Samp}^{(i)}$ cannot satisfy Def. 6. Thus, there exists a non-uniform PPT extractor $\mathcal{E}^{(i)}$ that given r finds an input $\chi = (h, \pi)$ on which P_0 and P_1 , output by $\text{Samp}^{(i)}$ on coins r , differ. In particular, for some polynomial $\ell(\cdot)$ and infinitely many λ it holds that

$$\Pr \left[r \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; (P_0, P_1) \leftarrow \text{Samp}^{(i)}(1^\lambda, r); \chi \leftarrow \mathcal{E}^{(i)}(r) : P_0(\chi) \neq P_1(\chi) \right] \geq \frac{1}{\ell(\lambda)}. \quad (9)$$

Let $\hat{\chi} = (\hat{h}, \hat{\pi})$ be a differing input output by $\mathcal{E}^{(i)}$. Recall that $\hat{\pi}$ is a short proof of $\hat{\eta} = (H, M_i, \hat{h}) \in \text{Legit}$, i.e., a short proof of knowledge of a witness x such that $M_i(x) = 1$ and $H(x) = \hat{h}$. By the definition of $\text{Samp}^{(i)}$ we have $P_0 := P_{M_i, H, \text{crs}, k}$ and $P_1 := P_{M_i, H, \text{crs}, k_{h^*}}$, as defined in (5). Since $(\hat{h}, \hat{\pi})$ is a differing input, the following conditions must hold.

condition(1): $\text{SNARK.Verify}(\text{crs}, (H, M_i, \hat{h}), \hat{\pi}) = 1$, for otherwise both P_0 and P_1 output \perp .

condition(2): $\hat{h} = h^* = H(x^*)$, for otherwise P_0 outputs $\text{PF.Eval}(k, \hat{h})$ and P_1 outputs $\text{PF.Eval}(k_{h^*}, \hat{h})$, which are equal by the correctness of puncturing.

condition(3): $M_i(x^*) = 0$; otherwise the CONSTR query would have returned \perp .

Since the SNARK $\hat{\pi}$ extracted by $\mathcal{E}^{(i)}$ is a proof of knowledge, we can extract a witness \hat{x} for it. In order to formally apply item 3. of Def. 9, we first construct a machine $\mathcal{A}_{\text{snrk}}$ that outputs $\hat{\pi}$ together with the statement. $\mathcal{A}_{\text{snrk}}$ simply runs $\mathcal{E}^{(i)}$ except that it embeds the crs from its input into the randomness.

- $\mathcal{A}_{\text{snrk}}(\text{crs})$
- Sample randomness $r_{\mathcal{A}}, r_H, r_k, r_c, r_y, r_1, \dots, r_{i-1}$
 - Set $r := r_{\mathcal{A}} \| r_H \| \text{crs} \| r_k \| r_c \| r_y \| r_1 \| \dots \| r_{i-1}$.
 - Simulate $\text{Samp}^{(i)}(r)$;
 - let H be the sampled hash function.
 - let M_i be the i -th CONSTR query.
 - $(\hat{h}, \hat{\pi}) \leftarrow \mathcal{E}^{(i)}(\text{crs}, r)$.
 - Output $(\eta := (H, M_i, \hat{h}), \hat{\pi})$.

By the construction of $\mathcal{A}_{\text{snrk}}$, Eq. (9) and condition(1) we have that

$$\Pr \left[\text{crs} \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; (\eta, \hat{\pi}) \leftarrow \mathcal{A}_{\text{snrk}}(\text{crs}) : \text{Verify}(\text{crs}, \eta, \hat{\pi}) = 1 \right] \geq \frac{1}{\ell(\lambda)}. \quad (10)$$

Further, since SNARK is an adaptive argument of knowledge, there exists $\mathcal{E}_{\mathcal{A}_{\text{snrk}}}$ which extracts a witness, that is:

$$\Pr \left[\begin{array}{l} \text{crs} \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; r \leftarrow \{0, 1\}^{\text{poly}(\lambda)} \\ (\eta, \hat{\pi}) := \mathcal{A}_{\text{snrk}}(\text{crs}; r); \hat{x} \leftarrow \mathcal{E}_{\mathcal{A}_{\text{snrk}}}(\text{crs}, r) \end{array} : \begin{array}{l} \text{Verify}(\text{crs}, \eta, \hat{\pi}) = 1 \\ \wedge (\eta, \hat{x}) \notin R_{\text{legit}} \end{array} \right] = \text{negl}(\lambda),$$

which together with (10) yields:

$$\Pr \left[\begin{array}{l} \text{crs} \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; r \leftarrow \{0, 1\}^{\text{poly}(\lambda)} \\ (\eta, \hat{\pi}) := \mathcal{A}_{\text{snrk}}(\text{crs}; r); \hat{x} \leftarrow \mathcal{E}_{\mathcal{A}_{\text{snrk}}}(\text{crs}, r) \end{array} : (\eta, \hat{x}) \in R_{\text{legit}} \right] \geq \frac{1}{\ell(\lambda)} - \text{negl}(\lambda) . \quad (11)$$

We now construct an adversary $\mathcal{A}_{\text{cll-fnd}}$ against \mathcal{H} that on input λ and coins r_H , used to sample a function H , outputs a collision for H : Given r_H , $\mathcal{A}_{\text{cll-fnd}}$ generates a CRS for SNARKs, then runs $\mathcal{A}_{\text{snrk}}(\text{crs})$, but using the randomness r_H from its input, and then runs $\mathcal{E}_{\mathcal{A}_{\text{snrk}}}$ to extract a collision:

$$\begin{array}{l} \mathcal{A}_{\text{cll-fnd}}(1^\lambda, r_H^*) \\ - \text{crs} \leftarrow \{0, 1\}^{\text{poly}(\lambda)}. \\ - \text{Sample randomness } r_{\mathcal{A}}, r_H, r_k, r_c, r_y, r_1, \dots, r_{i-1} \\ - \text{Set } r := r_{\mathcal{A}} \| r_H^* \| \text{crs} \| r_k \| r_c \| r_y \| r_1 \| \dots \| r_{i-1}. \\ - (\hat{h}, \hat{\pi}) \leftarrow \mathcal{E}^{(i)}(r). \\ - \text{Set } r' := r_{\mathcal{A}} \| r_H^* \| r_k \| r_c \| r_y \| r_1 \| \dots \| r_{i-1} \text{ // } \mathcal{A}_{\text{snrk}} \text{'s coins} \\ - \hat{x} \leftarrow \mathcal{E}_{\mathcal{A}_{\text{snrk}}}(\text{crs}, r'). \\ - \text{Output } (\hat{x}, x^*) \text{ as a collision pair for } H. \end{array}$$

By (11), with non-negligible probability, the values $M_i, \hat{h}, \hat{\pi}$ computed during the execution of $\mathcal{A}_{\text{cll-fnd}}$ satisfy $((H, M_i, \hat{h}), \hat{x}) \in R_{\text{legit}}$, that is, $M_i(\hat{x}) = 1$ and $\hat{h} = H(\hat{x})$. By condition(3), $M_i(x^*) = 0$, and hence $\hat{x} \neq x^*$. By condition(2), $\hat{h} = H(x^*)$, and hence (x, x^*) is a collision. In particular, the following is non-negligible:

$$\Pr \left[\begin{array}{l} r_H \leftarrow \{0, 1\}^{\text{poly}(\lambda)}; H := \text{HSmp}(1^\lambda, r_H); \\ (x_1, x_2) \leftarrow \mathcal{A}_{\text{cll-fnd}}(1^\lambda, r_H) \end{array} : \begin{array}{l} H(x_1) = H(x_2) \wedge \\ x_1 \neq x_2 \end{array} \right] .$$

Therefore we have reached a contradiction to collision resistance of \mathcal{H} , and it must be that $\mathbf{Exp}^{b,(0)}$ and $\mathbf{Exp}^{b,(1)}$ are computationally indistinguishable, i.e.,

$$|\Pr[\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{b,(0)}(\lambda) = 1] - \Pr[\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{b,(1)}(\lambda) = 1]| = \text{negl}(\lambda) .$$

A.2 Proof of Proposition 2

Let $b \in \{0, 1\}$ be arbitrarily fixed. The only difference between games $\mathbf{Exp}^{b,(1)}$ and $\mathbf{Exp}^{b,(2)}$ is when \mathcal{A} queries $\text{EVAL}(\cdot)$ on x such that $H(x) = H(x^*)$. Then $\mathbf{Exp}^{b,(2)}$ aborts, while on any other query the oracle EVAL behaves equivalently in both games since $H(x) \neq H(x^*)$ implies $\text{PF.Eval}(k_{h^*}, H(x)) = \text{PF.Eval}(k, H(x))$.

We can therefore build an adversary $\mathcal{A}_{\text{cll-fnd}}$ against the hash-function family \mathcal{H} that on input $(1^\lambda, r_H)$ simulates $\mathbf{Exp}^{b,(2)}$ (except that instead of sampling H , it uses $H := \text{H.Smp}(1^\lambda; r_H)$) until in an oracle query $\text{EVAL}(x)$ the game would abort. $\mathcal{A}_{\text{cll-fnd}}$ then outputs (x^*, x) , which is a collision precisely when the game would have aborted.

A.3 Proof of Proposition 3

Assume towards contradiction that there exists a PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ and a polynomial $p(\cdot)$ such that for infinitely many λ

$$|\Pr[\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{0,(2)}(\lambda) = 1] - \Pr[\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{1,(2)}(\lambda) = 1]| \geq \frac{1}{p(\lambda)} .$$

Then we construct a PPT adversary $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ running in game $\mathbf{Exp}_{\mathcal{P}\mathcal{F}, \mathcal{B}}^{\text{PCT-}b}(\lambda)$, the selective-security game of $\mathcal{P}\mathcal{F}$ (cf. Fig. 2) as follows. (Note that \mathcal{B}_2 does not use its $\text{EVAL}(\cdot)$ oracle.)

$\mathcal{B}_1(1^\lambda)$

- $(x^*, st_A) \leftarrow \mathcal{A}_1(1^\lambda)$.
- $H \leftarrow \text{H.Smp}(1^\lambda)$ and $h^* := H(x^*)$.
- Return $(h^*, T := \{h^*\}, st := (H, x^*, st_A))$.

$\mathcal{B}_2^{\text{EVAL}(\cdot)}(st, k_{h^*}, y^*)$ // y^* is either $\text{PF.Eval}(k, H(x^*))$ or random.

- $crs \leftarrow \{0, 1\}^{\text{poly}(\lambda)}$.
- $pp := (H, crs)$.
- $b' \leftarrow \mathcal{A}_2^{\text{CONSTR}(\cdot), \text{EVAL}(\cdot)}(st_A, y^*)$.
 - simulate $\text{CONSTR}(M)$:
 - if $M \notin \mathcal{M}_\lambda \vee M(x^*) = 1$, reply \perp ;
 - else $P := P_{M, H, crs, k_{h^*}}$ as defined in (5); $\tilde{P} \leftarrow \text{diO}(1^\lambda, P_{M, H, crs, k_{h^*}})$.
 - Return $k_M := (M, \tilde{P}, (H, crs))$.
 - simulate $\text{EVAL}(x)$:
 - if $x = x^*$, reply \perp ; if $H(x) = H(x^*)$ then abort;
 - else reply $y := \text{PF.Eval}(k_{h^*}, H(x))$.
- Output b' .

By construction $\Pr[\mathbf{Exp}_{\mathcal{P}\mathcal{F}, \mathcal{B}}^{\text{PCT-}b}(\lambda) = 1] = \Pr[\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{b, (2)}(\lambda) = 1]$, and therefore for infinitely many λ , it holds that

$$\left| \Pr[\mathbf{Exp}_{\mathcal{P}\mathcal{F}, \mathcal{B}}^{\text{PCT-}0}(\lambda) = 1] - \Pr[\mathbf{Exp}_{\mathcal{P}\mathcal{F}, \mathcal{B}}^{\text{PCT-}1}(\lambda) = 1] \right| \geq \frac{1}{p(\lambda)} .$$

This contradicts the selective security of PF, and we conclude that

$$\left| \Pr[\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{0, (2)}(\lambda) = 1] - \Pr[\mathbf{Exp}_{\mathcal{F}, \mathcal{A}}^{1, (2)}(\lambda) = 1] \right| = \text{negl}(\lambda) .$$