

Updatable Functional Encryption

Afonso Arriaga¹, Vincenzo Iovino¹, and Qiang Tang¹

SnT, University of Luxembourg, Luxembourg City, Luxembourg
afonso.delerue@uni.lu, vincenzo.iovino@uni.lu, tonyrhul@gmail.com

Abstract. Functional encryption (FE) allows an authority to issue tokens associated with various functions, allowing the holder of some token for function f to learn only $f(D)$ from a ciphertext that encrypts D . The standard approach is to model f as a circuit, which yields inefficient evaluations over large inputs. Here, we propose a new primitive that we call *updatable functional encryption* (UFE), where instead of circuits we deal with RAM programs, which are closer to how programs are expressed in von Neumann architecture. We impose strict efficiency constraints in that the run-time of a token P on ciphertext CT is proportional to the run-time of its clear-form counterpart (program P on memory D) up to a *polylogarithmic* factor in the size of D , and we envision tokens that are capable to *update* the ciphertext, over which other tokens can be subsequently executed. We define a security notion for our primitive and propose a candidate construction from obfuscation, which serves as a starting point towards the realization of other schemes and contributes to the study on how to compute RAM programs over public-key encrypted data.

Keywords: Updatable functional encryption, RAM model, Persistent memory.

1 Introduction

The concept of functional encryption (FE), a generalization of identity-based encryption, attribute-based encryption, inner-product encryption and other forms of public-key encryption, was independently formalized by Boneh, Sahai and Waters [7] and O’Neil [20]. In an FE scheme, the holder of a master secret key can issue tokens associated with functions of its choice. Possessing a token for f allows one to recover $f(D)$, given an encryption of D . Informally, security dictates that only $f(D)$ is revealed about D and nothing else.

Garg et al. [13] put forth the first candidate construction of an FE scheme supporting all polynomial-size circuits based on indistinguishability obfuscation (iO), which is now known as a central hub for the realization of many cryptographic primitives [22].

The most common approach is to model functions as a circuits. In some works, however, functions are modeled as Turing machines (TM) or random-access machines (RAM). Recently, Ananth and Sahai [3] constructed an adaptively secure functional encryption scheme for TM, based on indistinguishability

obfuscation. Nonetheless, their work does not tackle the problem of having the token update the encrypted message, over which other tokens can be subsequently executed.

In the symmetric setting, the notion of garbled RAM, introduced by Lu and Ostrovsky [18] and revisited by Gentry et al. [14], addresses this important use-case where garbled memory data can be reused across multiple program executions. Garbled RAM can be seen as an analogue of Yao’s garbled circuits [23] (see also [5] for an abstract generalization) that allows a user to garble a RAM program without having to compile it into a circuit first. As a result, the time it takes to evaluate a garbled program is only proportional to the running time of the program on a random-access machine. Several other candidate constructions were also proposed in [15,10,11,9].

Desmedt et al. [12] proposed an FE with controlled homomorphic properties. However, their scheme updates and re-encrypts the entire data, which carries a highly inefficient evaluation-time.

OUR CONTRIBUTION. We propose a new primitive that we call *updatable functional encryption* (UFE). It bears resemblance to functional encryption in that encryption is carried out in the public-key setting and the owner of the master secret key can issue tokens for functions—here, modeled as RAM programs—of its choice that allow learning the outcome of the function on the message underneath a ciphertext. We envision tokens that are also capable to *update* the ciphertext, over which other tokens can be subsequently executed. We impose strict efficiency constraints in that the run-time of a token \bar{P} on ciphertext CT is proportional to the run-time of its clear-form counterpart (program P on memory D) up to a *polylogarithmic* factor in the size of D. We define a security notion for our primitive and propose a candidate construction based on an instance of distributional indistinguishability (DI) obfuscation, a notion introduced by [4] in the context of point function obfuscation and later generalized by [2]. Recent results put differing-inputs obfuscation (diO) [1] with auxiliary information in contention with other assumptions [6]; one might question if similar attacks apply to the obfuscation notion we require in our reduction. As far as we can tell, the answer is negative. However, we view our construction as a starting point towards the realization of other updatable functional encryption schemes from milder forms of obfuscation.

2 Preliminaries

NOTATION. We denote the security parameter by $\lambda \in \mathbb{N}$ and assume it is implicitly given to all algorithms in unary representation 1^λ . We denote the set of all bit strings of length ℓ by $\{0, 1\}^\ell$ and the length of a string \mathbf{a} by $|\mathbf{a}|$. We write $\mathbf{a} \leftarrow \mathbf{b}$ to denote the algorithmic action of assigning the value of \mathbf{b} to the variable \mathbf{a} . We use $\perp \notin \{0, 1\}^*$ to denote a special failure symbol and ϵ for the empty string. A vector of strings \mathbf{x} is written in boldface, and $\mathbf{x}[i]$ denotes its i th entry. The number of entries of \mathbf{x} is denoted by $|\mathbf{x}|$. For a finite set X , we denote its

cardinality by $|\mathbf{X}|$ and the action of sampling a uniformly random element \mathbf{a} from \mathbf{X} by $\mathbf{a} \leftarrow_{\$} \mathbf{X}$. If \mathcal{A} is a probabilistic algorithm we write $\mathbf{a} \leftarrow_{\$} \mathcal{A}(i_1, i_2, \dots, i_n; r)$ for the action of running \mathcal{A} on inputs i_1, i_2, \dots, i_n with random coins r , and assigning the result to \mathbf{a} . For a circuit \mathbf{C} we denote its size by $|\mathbf{C}|$. We call a real-valued function $\mu(\lambda)$ negligible if $\mu(\lambda) \in \mathcal{O}(\lambda^{-\omega(1)})$ and denote the set of all negligible functions by NEGL. We adopt the code-based game-playing framework. As usual “ppt” stands for probabilistic polynomial time.

CIRCUIT FAMILIES. Let $\mathbf{MSp} := \{\mathbf{MSp}_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathbf{OSp} := \{\mathbf{OSp}_\lambda\}_{\lambda \in \mathbb{N}}$ be two families of finite sets parametrized by a security parameter $\lambda \in \mathbb{N}$. A circuit family $\mathbf{CSp} := \{\mathbf{CSp}_\lambda\}_{\lambda \in \mathbb{N}}$ is a sequence of circuit sets indexed by the security parameter. We assume that for all $\lambda \in \mathbb{N}$, all circuits in \mathbf{CSp}_λ share a common input domain \mathbf{MSp}_λ and output space \mathbf{OSp}_λ . We also assume that membership in sets can be efficiently decided. For a vector of circuits $\mathbf{C} = [C_1, \dots, C_n]$ and a message \mathbf{m} we define $\mathbf{C}(\mathbf{m})$ to be the vector whose i th entry is $C_i(\mathbf{m})$.

TREES. We associate a tree \mathbf{T} with the set of its nodes $\{\text{node}_{i,j}\}$. Each node is indexed by a pair of non-negative integers representing the position (level and branch) of the node on the tree. The root of the tree is indexed by $(0, 0)$, its children have indices $(1, 0)$, $(1, 1)$, etc. A binary tree is *perfectly balanced* if every leaf is at the same level.

2.1 Public-key encryption

SYNTAX. A public-key encryption scheme $\text{PKE} := (\text{PKE.Setup}, \text{PKE.Enc}, \text{PKE.Dec})$ with message space $\mathbf{MSp} := \{\mathbf{MSp}_\lambda\}_{\lambda \in \mathbb{N}}$ and randomness space $\mathbf{RSp} := \{\mathbf{RSp}_\lambda\}_{\lambda \in \mathbb{N}}$ is specified by three ppt algorithms as follows. (1) $\text{PKE.Setup}(1^\lambda)$ is the probabilistic key-generation algorithm, taking as input the security parameter and returning a secret key sk and a public key pk . (2) $\text{PKE.Enc}(\text{pk}, \mathbf{m}; r)$ is the probabilistic encryption algorithm. On input a public key pk , a message $\mathbf{m} \in \mathbf{MSp}_\lambda$ and possibly some random coins $r \in \mathbf{RSp}_\lambda$, this algorithm outputs a ciphertext c . (3) $\text{PKE.Dec}(\text{sk}, c)$ is the deterministic decryption algorithm. On input of a secret key sk and a ciphertext c , this algorithm outputs a message $\mathbf{m} \in \mathbf{MSp}_\lambda$ or failure symbol \perp .

CORRECTNESS. The correctness of a public-key encryption scheme requires that for any $\lambda \in \mathbb{N}$, any $(\text{sk}, \text{pk}) \in [\text{PKE.Setup}(1^\lambda)]$, any $\mathbf{m} \in \mathbf{MSp}_\lambda$ and any random coins $r \in \mathbf{RSp}_\lambda$, we have that $\text{PKE.Dec}(\text{sk}, \text{PKE.Enc}(\text{pk}, \mathbf{m}; r)) = \mathbf{m}$.

SECURITY. We recall the standard security notions of *indistinguishability under chosen ciphertext attacks* (IND-CCA) and its weaker variant known as *indistinguishability under chosen plaintext attacks* (IND-CPA). We say that a public-key encryption scheme PKE is IND-CCA secure if for every *legitimate* ppt adversary \mathcal{A}

$$\text{Adv}_{\text{PKE}, \mathcal{A}}^{\text{ind-cca}}(\lambda) := 2 \cdot \Pr[\text{IND-CCA}_{\text{PKE}}^{\mathcal{A}}(\lambda)] - 1,$$

where game $\text{IND-CCA}_{\text{PKE}}^{\mathcal{A}}$ described in Figure 1, in which the adversary has access to a left-or-right challenge oracle (LR) and a decryption oracle (Dec). We

$\text{IND-CCA}_{\text{PKE}}^{\mathcal{A}}(\lambda):$ $(\text{sk}, \text{pk}) \leftarrow_{\mathcal{S}} \text{PKE.Setup}(1^\lambda)$ $b \leftarrow_{\mathcal{S}} \{0, 1\}$ $b' \leftarrow_{\mathcal{S}} \mathcal{A}^{\text{LR}, \text{Dec}}(1^\lambda, \text{pk})$ return $(b = b')$	$\text{LR}(m_0, m_1):$ $c \leftarrow_{\mathcal{S}} \text{PKE.Enc}(\text{pk}, m_b)$ $\text{List} \leftarrow c : \text{List}$ return c	$\text{Dec}(c):$ $m \leftarrow \text{PKE.Dec}(\text{sk}, c)$ return m
---	---	---

Fig. 1. Game defining IND-CCA security of a public-key encryption scheme PKE.

say that \mathcal{A} is legitimate if: (1) $|m_0| = |m_1|$ whenever the left-or-right oracle is queried; and (2) the adversary does not call the decryption oracle with $c \in \text{List}$. We obtain the weaker IND-CPA notion if the adversary is not allowed to place any decryption query.

2.2 NIZK proof systems

SYNTAX. A non-interactive zero-knowledge proof system for an **NP** language \mathcal{L} with an efficiently computable binary relation \mathcal{R} consists of three ppt algorithms as follows. (1) $\text{NIZK.Setup}(1^\lambda)$ is the setup algorithm and on input a security parameter 1^λ it outputs a common reference string crs ; (2) $\text{NIZK.Prove}(\text{crs}, x, w)$ is the proving algorithm and on input a common reference string crs , a statement x and a witness w it outputs a proof π or a failure symbol \perp ; (3) $\text{NIZK.Verify}(\text{crs}, x, \pi)$ is the verification algorithm and on input a common reference string crs , a statement x and a proof π it outputs either **true** or **false**.

PERFECT COMPLETENESS. Completeness imposes that an honest prover can always convince an honest verifier that a statement belongs to \mathcal{L} , provided that it holds a witness testifying to this fact. We say a NIZK proof is *perfectly complete* if for every (possibly unbounded) adversary \mathcal{A}

$$\text{Adv}_{\text{NIZK}, \mathcal{A}}^{\text{complete}}(\lambda) := \Pr \left[\text{Complete}_{\text{NIZK}}^{\mathcal{A}}(\lambda) \right] = 0 ,$$

where game $\text{Complete}_{\text{NIZK}}^{\mathcal{A}}(\lambda)$ is shown in Fig. 2 on the left.

COMPUTATIONAL ZERO KNOWLEDGE. The zero-knowledge property guarantees that proofs do not leak information about the witnesses that originated them. Technically, this is formalized by requiring the existence of a ppt simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ where Sim_0 takes the security parameter 1^λ as input and outputs a simulated common reference string crs together with a trapdoor tp , and Sim_1 takes the trapdoor as input tp together with a statement x for which it must forge a proof π . We say a proof system is *computationally zero knowledge* if, for every ppt adversary \mathcal{A} , there exists a ppt simulator Sim such that

$$\text{Adv}_{\text{NIZK}, \mathcal{A}, \text{Sim}}^{\text{zk}}(\lambda) := \left| \Pr \left[\text{ZK-Real}_{\text{NIZK}}^{\mathcal{A}}(\lambda) \right] - \left[\text{ZK-Ideal}_{\text{NIZK}}^{\mathcal{A}, \text{Sim}}(\lambda) \right] \right| \in \text{NEGL} ,$$

where games $\text{ZK-Real}_{\text{NIZK}}^{\mathcal{A}}(\lambda)$ and $\text{ZK-Ideal}_{\text{NIZK}}^{\mathcal{A}, \text{Sim}}(\lambda)$ are shown in Fig. 3.

STATISTICAL SIMULATION SOUNDNESS. Soundness imposes that a malicious prover cannot convince an honest verifier of a false statement. This should be true even when the adversary itself is provided with simulated proofs. We say NIZK is *statistically simulation sound* with respect to simulator Sim if, for every (possibly unbounded) adversary \mathcal{A}

$$\mathbf{Adv}_{\text{NIZK}, \mathcal{A}}^{\text{sound}}(\lambda) := \Pr \left[\text{Sound}_{\text{NIZK}}^{\mathcal{A}, \text{Sim}}(\lambda) \right] \in \text{NEGL} ,$$

where game $\text{Sound}_{\text{NIZK}}^{\mathcal{A}}(\lambda)$ is shown in Fig. 2 on the right.

$\text{Complete}_{\text{NIZK}}^{\mathcal{A}}(1^\lambda):$ $\text{crs} \leftarrow_{\$} \text{NIZK.Setup}(1^\lambda)$ $(x, w) \leftarrow_{\$} \mathcal{A}(1^\lambda, \text{crs})$ if $(x, w) \notin \mathcal{R}$ return 0 $\pi \leftarrow_{\$} \text{NIZK.Prove}(\text{crs}, x, w)$ return $\neg(\text{NIZK.Verify}(\text{crs}, x, \pi))$	$\text{Sound}_{\text{NIZK}}^{\mathcal{A}}(1^\lambda):$ $\text{crs} \leftarrow_{\$} \text{NIZK.Setup}(1^\lambda)$ $(x, \pi) \leftarrow_{\$} \mathcal{A}(1^\lambda, \text{crs})$ return $(x \notin \mathcal{L} \wedge \text{NIZK.Verify}(\text{crs}, x, \pi))$
--	---

Fig. 2. Games defining the completeness and soundness properties of a non-interactive zero-knowledge proof system NIZK.

$\text{ZK-Real}_{\text{NIZK}}^{\mathcal{A}}(1^\lambda):$ $\text{crs} \leftarrow_{\$} \text{NIZK.Setup}(1^\lambda)$ $b \leftarrow_{\$} \mathcal{A}^{\text{Prove}}(1^\lambda, \text{crs})$ $\text{Prove}(x, w):$ if $(x, w) \notin \mathcal{R}$ return \perp $\pi \leftarrow_{\$} \text{NIZK.Prove}(\text{crs}, x, w)$ return π	$\text{ZK-Ideal}_{\text{NIZK}}^{\mathcal{A}, \text{Sim}}(1^\lambda):$ $(\text{crs}, \text{tp}) \leftarrow_{\$} \text{Sim}_1(1^\lambda)$ $b \leftarrow_{\$} \mathcal{A}^{\text{Prove}}(1^\lambda, \text{crs})$ $\text{Prove}(x, w):$ if $(x, w) \notin \mathcal{R}$ return \perp $\pi \leftarrow_{\$} \text{Sim}_2(\text{crs}, \text{tp}, x)$ return π
---	---

Fig. 3. Games defining the zero-knowledge property of a non-interactive zero-knowledge proof system NIZK.

2.3 Collision-resistant hash functions

A hash function family $\mathbf{H} := \{\mathbf{H}_\lambda\}_{\lambda \in \mathbb{N}}$ is a set parametrized by a security parameter $\lambda \in \mathbb{N}$, where each \mathbf{H}_λ is a collection of functions mapping $\{0, 1\}^m$ to $\{0, 1\}^n$ such that $m > n$. The hash function family \mathbf{H} is said to be collision-resistant if no ppt adversary \mathcal{A} can find a pair of colliding inputs, with noticeable probability, given a function picked uniformly from \mathbf{H}_λ . More precisely, we require that

$$\mathbf{Adv}_{\mathbf{H}, \mathcal{A}}^{\text{cr}}(\lambda) := \Pr[\text{CR}_{\mathbf{H}}^{\mathcal{A}}(\lambda)] \in \text{NEGL},$$

where game $\text{CR}_{\mathbf{H}}^{\mathcal{A}}(\lambda)$ is defined in Fig. 4.

$\begin{array}{l} \text{CR}_H^A(\lambda): \\ \bar{h} \leftarrow_{\$} H_\lambda \\ (x_0, x_1) \leftarrow_{\$} \mathcal{A}(1^\lambda, \bar{h}) \\ \text{return } (x_0 \neq x_1 \wedge h(x_0) = h(x_1)) \end{array}$

Fig. 4. Game defining collision-resistance of a hash function family H .

2.4 Puncturable pseudorandom functions

A puncturable pseudorandom function family $\text{PPRF} := (\text{PPRF.Gen}, \text{PPRF.Eval}, \text{PPRF.Punc})$ is a triple of ppt algorithms as follows. (1) PPRF.Gen on input the security parameter 1^λ outputs a uniform element in K_λ ; (2) PPRF.Eval is deterministic and on input a key $k \in K_\lambda$ and a point $x \in X_\lambda$ outputs a point $y \in Y_\lambda$; (3) PPRF.Punc is probabilistic and on input a $k \in K_\lambda$ and a polynomial-size set of points $S \subseteq X_\lambda$ outputs a punctured key k_S . As per [22], we require the PPRF to satisfy the following two properties:

FUNCTIONALITY PRESERVATION UNDER PUNCTURING : For every $\lambda \in \mathbb{N}$, every polynomial-size set $S \subseteq X_\lambda$ and every $x \in X_\lambda \setminus S$, it holds that

$$\Pr \left[\text{PPRF.Eval}(k, x) = \text{PPRF.Eval}(k_S, x) \mid \begin{array}{l} k \leftarrow_{\$} \text{PPRF.Gen}(1^\lambda) \\ k_S \leftarrow_{\$} \text{PPRF.Punc}(k, S) \end{array} \right] = 1.$$

PSEUDORANDOMNESS AT PUNCTURED POINTS : For every ppt adversary $\mathcal{A} := (\mathcal{A}_0, \mathcal{A}_1)$,

$$\text{Adv}_{\text{PPRF}, \mathcal{A}}^{\text{pprf}}(\lambda) := 2 \cdot \Pr[\text{PPRF}_{\text{PPRF}}^{\mathcal{A}}(\lambda)] - 1 \in \text{NEGL},$$

where game $\text{PPRF}_{\text{PPRF}}^{\mathcal{A}}(\lambda)$ is defined in Fig. 5.

$\begin{array}{l} \text{PPRF}_{\text{PPRF}}^{\mathcal{A}}(\lambda): \\ (\mathcal{S}, \text{st}) \leftarrow_{\$} \mathcal{A}_0(1^\lambda) \\ k \leftarrow_{\$} \text{PPRF.Gen}(1^\lambda) \\ k_S \leftarrow_{\$} \text{PPRF.Punc}(k, \mathcal{S}) \\ b \leftarrow_{\$} \{0, 1\} \\ b' \leftarrow_{\$} \mathcal{A}_1^{\text{Fn}}(1^\lambda, k_S, \text{st}) \\ \text{return } (b = b') \end{array}$	$\begin{array}{l} \text{Fn}(x): \\ \text{if } x \notin \mathcal{S} \text{ return } \text{PPRF.Eval}(k_S, x) \\ \text{if } T[x] = \perp \text{ then} \\ \quad T[x] \leftarrow_{\$} Y_\lambda \\ \text{if } b = 1 \text{ return } T[x] \\ \text{else return } \text{PPRF.Eval}(k, x) \end{array}$
---	---

Fig. 5. Game defining pseudorandomness at punctured points of $\text{PPRF} := (\text{PPRF.Gen}, \text{PPRF.Eval}, \text{PPRF.Punc})$.

2.5 Obfuscators

SYNTAX. An obfuscator for a circuit family CSp is a uniform ppt algorithm Obf that on input the security parameter 1^λ and the description of a circuit $C \in \text{CSp}_\lambda$

outputs the description of another circuit \bar{C} . We require any obfuscator to satisfy the following two requirements.

FUNCTIONALITY PRESERVATION : For any $\lambda \in \mathbb{N}$, any $C \in \text{CSp}_\lambda$ and any $m \in \text{MSp}_\lambda$, with overwhelming probability over the choice of $\bar{C} \leftarrow_s \text{Obf}(1^\lambda, C)$ we have that $C(m) = \bar{C}(m)$.

POLYNOMIAL SLOWDOWN : There is a polynomial poly such that for any $\lambda \in \mathbb{N}$, any $C \in \text{CSp}_\lambda$ and any $\bar{C} \leftarrow_s \text{Obf}(1^\lambda, C)$ we have that $|\bar{C}| \leq \text{poly}(|C|)$.

In this paper we rely on the security definitions of *indistinguishability obfuscation* (iO) [13] and *distributional indistinguishability* (DI). The latter definition was first introduced by [4] in the context of point function obfuscation and later generalized by [2] to cover samplers that output not only point circuits. We note that the work of [2] considers only *statistically* unpredictable samplers, which is a more restricted class of samplers, and therefore is a more amenable form of obfuscation. Unfortunately, for the purpose of proving the construction we present in Section 3.2 secure, we rely on a DI obfuscator against a computationally unpredictable sampler.

INDISTINGUISHABILITY OBFUSCATION (IO). This property requires that given any two functionally equivalent circuits C_0 and C_1 of equal size, the obfuscations of C_0 and C_1 should be computationally indistinguishable. More precisely, for any ppt adversary \mathcal{A} and for any sampler \mathcal{S} that outputs two circuits $C_0, C_1 \in \text{CSp}_\lambda$ such that $C_0(m) = C_1(m)$ for all inputs m and $|C_0| = |C_1|$, we have that

$$\text{Adv}_{\text{Obf}, \mathcal{S}, \mathcal{A}}^{\text{io}}(\lambda) := 2 \cdot \Pr[\text{iO}_{\text{Obf}}^{\mathcal{S}, \mathcal{A}}(\lambda)] - 1 \in \text{NEGL},$$

where game $\text{iO}_{\text{Obf}}^{\mathcal{S}, \mathcal{A}}(\lambda)$ is defined in Fig. 6 on the left.

DISTRIBUTIONAL INDISTINGUISHABILITY (DI). We define this property with respect to some class of unpredictable samplers \mathbb{S} . A sampler is an algorithm \mathcal{S} that on input the security parameter 1^λ and possibly some state information st outputs a pair of vectors of CSp_λ circuits (C_0, C_1) of equal dimension and possibly some auxiliary information z . We require the components of the two circuit vectors to be encoded as bit strings of equal length. \mathcal{S} is said to be *unpredictable* if no ppt predictor with oracle access to the circuits can find a differing input m such that $C_0(m) \neq C_1(m)$. An obfuscator Obf is DI secure with respect to a class of unpredictable samplers \mathbb{S} if for all $\mathcal{S} \in \mathbb{S}$ the obfuscations of C_0 and C_1 output by \mathcal{S} are computationally indistinguishable. More precisely, for every $\mathcal{S} \in \mathbb{S}$ and every ppt adversary \mathcal{A} we have that

$$\text{Adv}_{\text{Obf}, \mathcal{S}, \mathcal{A}}^{\text{di}}(\lambda) := 2 \cdot \Pr[\text{DI}_{\text{Obf}}^{\mathcal{S}, \mathcal{A}}(\lambda)] - 1 \in \text{NEGL},$$

where game $\text{DI}_{\text{Obf}}^{\mathcal{S}, \mathcal{A}}(1^\lambda)$ is defined in Fig. 6 (middle). Furthermore, we say sampler \mathcal{S} is *computationally unpredictable* if for any ppt predictor \mathcal{P}

$$\text{Adv}_{\mathcal{S}, \mathcal{P}}^{\text{pred}}(\lambda) := \Pr \left[\text{Pred}_{\mathcal{S}}^{\mathcal{P}}(1^\lambda) \right] \in \text{NEGL} ,$$

where game $\text{Pred}_{\mathcal{S}}^{\mathcal{P}}(1^\lambda)$ is shown in Fig. 6 on the right.

$\text{iO}_{\text{Obf}}^{\mathcal{S}, \mathcal{A}}(\lambda):$ $(\mathbf{C}_0, \mathbf{C}_1, z) \leftarrow_{\mathcal{S}} \mathcal{S}(1^\lambda)$ $b \leftarrow_{\mathcal{S}} \{0, 1\}$ $\overline{\mathbf{C}} \leftarrow_{\mathcal{S}} \text{Obf}(1^\lambda, \mathbf{C}_b)$ $b' \leftarrow_{\mathcal{A}} \mathcal{A}_1(1^\lambda, z, \overline{\mathbf{C}})$ return $(b = b')$	$\text{DI}_{\text{Obf}}^{\mathcal{S}, \mathcal{A}}(\lambda):$ $(\text{st}, \text{st}') \leftarrow_{\mathcal{A}_0} \mathcal{A}_0(1^\lambda)$ $(\mathbf{C}_0, \mathbf{C}_1, z) \leftarrow_{\mathcal{S}} \mathcal{S}(1^\lambda, \text{st})$ $b \leftarrow_{\mathcal{S}} \{0, 1\}$ $\overline{\mathbf{C}} \leftarrow_{\mathcal{S}} \text{Obf}(1^\lambda, \mathbf{C}_b)$ $b' \leftarrow_{\mathcal{A}_1} \mathcal{A}_1(1^\lambda, z, \text{st}', \overline{\mathbf{C}})$ return $(b = b')$	$\text{Pred}_{\mathcal{S}}^{\mathcal{P}}(\lambda):$ $(\text{st}, \text{st}') \leftarrow_{\mathcal{P}_0} \mathcal{P}_0(1^\lambda)$ $(\mathbf{C}_0, \mathbf{C}_1, z) \leftarrow_{\mathcal{S}} \mathcal{S}(\text{st})$ $\mathbf{m} \leftarrow_{\mathcal{P}_1} \mathcal{P}_1^{\text{Fn}}(1^\lambda, z, \text{st}')$ return $(\mathbf{C}_0(\mathbf{m}) \neq \mathbf{C}_1(\mathbf{m}))$ $\text{Fn}(\mathbf{m}):$ return $(\mathbf{C}_0(\mathbf{m}))$
--	--	--

Fig. 6. Games defining iO and DI security of an obfuscator Obf , and unpredictability of a sampler \mathcal{S} .

2.6 RAM programs

In the RAM model of computation, a program \mathbf{P} has random-access to some initial *memory data* \mathbf{D} , comprised of $|\mathbf{D}|$ *memory cells*. At each *CPU step* of its execution, \mathbf{P} reads from and writes to a single memory cell *address*, which is determined by the previous step, and updates its internal state. By convention, the address in the first step is set to the first memory cell of \mathbf{D} , and the initial internal state is empty. Only when \mathbf{P} reaches the final step of its execution, it outputs a result y and terminates. We use the notation $y \leftarrow \mathbf{P}^{\mathbf{D} \rightarrow \mathbf{D}^*}$ to indicate this process, where \mathbf{D}^* is the resulting memory data when \mathbf{P} terminates, or simply $y \leftarrow \mathbf{P}^{\mathbf{D}}$ if we don't care about the resulting memory data. We also consider the case where the memory data *persists* between a sequential execution of n programs, and use the notation $(y_1, \dots, y_n) \leftarrow (\mathbf{P}_1, \dots, \mathbf{P}_n)^{\mathbf{D} \rightarrow \mathbf{D}^*}$ as short for $(y_1 \leftarrow \mathbf{P}_1^{\mathbf{D} \rightarrow \mathbf{D}_1}; \dots; y_n \leftarrow \mathbf{P}_n^{\mathbf{D}_{n-1} \rightarrow \mathbf{D}^*})$. In more detail, a RAM program description is a 4-tuple $\mathbf{P} := (\mathcal{Q}, \mathcal{T}, \mathcal{Y}, \delta)$, where:

- \mathcal{Q} is the set of all possible states, which always includes the empty state ϵ .
- \mathcal{T} is the set of all possible contents of a memory cell. If each cell contains a single bit, $\mathcal{T} = \{0, 1\}$.
- \mathcal{Y} is the output space of \mathbf{P} , which always includes the empty output ϵ .
- δ is the transition function, modeled as a circuit, which maps $(\mathcal{Q} \times \mathcal{T})$ to $(\mathcal{T} \times \mathcal{Q} \times \mathbb{N} \times \mathcal{Y})$. On input an internal state $\text{st}_i \in \mathcal{Q}$ and a content of a memory cell $\text{read}_i \in \mathcal{T}$, it outputs a (possibly different) content of a memory cell $\text{write}_i \in \mathcal{T}$, an internal state $\text{st}_{i+1} \in \mathcal{Q}$, an address of a memory cell $\text{addr}_{i+1} \in \mathbb{N}$ and an output $y \in \mathcal{Y}$.

In Figure 7 we show how program \mathbf{P} is executed on a random-access machine with initial memory data \mathbf{D} .

To conveniently specify the *efficiency* and *security* properties of the primitive we propose in the following section, we define functions `runTime` and `accessPattern` that on input a program \mathbf{P} and some initial memory data \mathbf{D} return the number of steps required for \mathbf{P} to complete its execution on \mathbf{D} and the list of addresses accessed during the execution, respectively. In other words, as per description in Fig. 7, `runTime` returns the value i when \mathbf{P} terminates, whereas `accessPattern`

returns `List`. More generally, we also allow these functions to receive as input a *set* of programs (P_1, \dots, P_n) to be executed sequentially on persistent memory, initially set to `D`.

```

EXECUTE  $P^D$ :
 $i \leftarrow 0$ ;  $\text{addr}_i \leftarrow 0$ ;  $\text{st}_i \leftarrow \epsilon$ ;  $y \leftarrow \epsilon$ ;  $\text{List} \leftarrow []$ 
while ( $y = \epsilon$ )
  // step  $i$ 
   $\text{List} \leftarrow \text{addr}_i : \text{List}$  // record the access pattern
   $\text{read}_i \leftarrow D[\text{addr}_i]$  // read from memory
   $(\text{write}_i, \text{st}_{i+1}, \text{addr}_{i+1}, y) \leftarrow \delta(\text{st}_i, \text{read}_i)$ 
   $D[\text{addr}_i] \leftarrow \text{write}_i$  // write to memory
   $i \leftarrow i + 1$ 
return ( $y$ )

```

Fig. 7. Execution of program P on a RAM machine with memory D .

3 Updatable Functional Encryption

We propose a new primitive that we call *updatable functional encryption*. It bears resemblance to functional encryption in that encryption is carried out in the public-key setting and the owner of the master secret key can issue tokens for functions of its choice that allows the holder of the token to learn the outcome of the function on the message underneath a ciphertext. Here, we model functions as RAM programs instead of circuits, which is closer to how programs are expressed in von Neumann architecture and avoids the RAM-to-circuit compilation. Not only that, we envision tokens that are capable to *update* the ciphertext, over which other tokens can be subsequently executed. Because the ciphertext evolves every time a token is executed and for better control over what information is revealed, each token is numbered sequentially so that it can only be executed *once* and *after all previous extracted tokens* have been executed on that ciphertext. Informally, the security requires that the ciphertext should not reveal more than what can be learned by applying the extracted tokens in order. As for efficiency, we want the run-time of a token to be proportional to the run-time of the program up to a *polylogarithmic* factor in the length of the encrypted message.

3.1 Definitions

SYNTAX. An updatable functional encryption scheme UFE for program family $\mathcal{P} := \{\mathcal{P}_\lambda\}_{\lambda \in \mathbb{N}}$ with message space $\text{MSP} := \{\text{MSP}_\lambda\}_{\lambda \in \mathbb{N}}$ is specified by three ppt algorithms as follows.

- $\text{UFE.Setup}(1^\lambda)$ is the setup algorithm and on input a security parameter 1^λ it outputs a master secret key `msk` and a master public key `mpk`;

- $\text{UFE.TokenGen}(\text{msk}, P, \text{tid})$ is the token-generation algorithm and on input a master secret key msk , a program description $P \in \mathcal{P}_\lambda$ and a token-id $\text{tid} \in \mathbb{N}$, outputs a token (i.e. another program description) \bar{P}_{tid} ;
- $\text{UFE.Enc}(\text{mpk}, D)$ is the encryption algorithm and on input a master public key mpk and memory data $D \in \text{MSp}_\lambda$ outputs a ciphertext CT .

We do not explicitly consider an evaluation algorithm. Instead, the RAM program \bar{P} output by UFE.TokenGen executes directly on memory data CT , a ciphertext resulting from the UFE.Enc algorithm. Note that this brings us close to the syntax of Garbled RAM, but in contrast encryption is carried out in the public-key setting.

CORRECTNESS. We say that UFE is correct if for every security parameter $\lambda \in \mathbb{N}$, for every memory data $D \in \text{MSp}_\lambda$ and for every sequence of polynomial length in λ of programs (P_1, \dots, P_n) , it holds that

$$\Pr \left[y_1 = y'_1 \wedge \dots \wedge y_n = y'_n \left| \begin{array}{l} (\text{msk}, \text{mpk}) \leftarrow_s \text{UFE.Setup}(1^\lambda) \\ \text{CT} \leftarrow_s \text{UFE.Enc}(\text{mpk}, D) \\ \text{for } i \in [n] \\ \bar{P}_i \leftarrow_s \text{UFE.TokenGen}(\text{msk}, P_i, i) \\ (y_1, \dots, y_n) \leftarrow (P_1, \dots, P_n)^D \\ (y'_1, \dots, y'_n) \leftarrow (\bar{P}_1, \dots, \bar{P}_n)^{\text{CT}} \end{array} \right. \right] = 1.$$

EFFICIENCY. Besides the obvious requirement that all algorithms run in polynomial-time in the length of their inputs, we also require that the run-time of token \bar{P} on ciphertext CT is proportional to the run-time of its clear-form counterpart (program P on memory D) up to a polynomial factor in λ and up to a polylogarithmic factor in the length of D . More precisely, we require that for every $\lambda \in \mathbb{N}$, for every sequence of polynomial length in λ of programs (P_1, \dots, P_n) and every memory data $D \in \text{MSp}_\lambda$, there exists a fixed polynomial function poly and a fixed polylogarithmic function polylog such that

$$\Pr \left[\begin{array}{l} \text{runTime}((\bar{P}_1, \dots, \bar{P}_n), \text{CT}) \leq \\ \text{runTime}((P_1, \dots, P_n), D) \cdot \\ \text{poly}(\lambda) \cdot \text{polylog}(|D|) \end{array} \left| \begin{array}{l} (\text{msk}, \text{mpk}) \leftarrow_s \text{UFE.Setup}(1^\lambda) \\ \text{CT} \leftarrow_s \text{UFE.Enc}(\text{mpk}, D) \\ \text{for } i \in [n] \\ \bar{P}_i \leftarrow_s \text{UFE.TokenGen}(\text{msk}, P_i) \end{array} \right. \right] = 1$$

over the random coins of all probabilistic algorithms. In particular, this means that for a program P running in sublinear-time in $|D|$, the run-time of \bar{P} over the encrypted data remains sublinear.

SECURITY. Let UFE be an updatable functional encryption scheme. We say UFE is *selectively secure* if for any legitimate ppt adversary \mathcal{A}

$$\mathbf{Adv}_{\text{UFE}, \mathcal{A}}^{\text{sel}}(\lambda) := 2 \cdot \Pr [\text{SEL}_{\text{UFE}}^{\mathcal{A}}(\lambda)] - 1 \in \text{NEGL},$$

where game $\text{SEL}_{\text{UFE}}^{\mathcal{A}}(\lambda)$ is defined in Fig. 8. We say \mathcal{A} is legitimate if the following two conditions are satisfied:

1. $(P_1, \dots, P_n)^{D_0} = (P_1, \dots, P_n)^{D_1}$
2. $\text{accessPattern}((P_1, \dots, P_n), D_0) = \text{accessPattern}((P_1, \dots, P_n), D_1)$

These conditions avoid that the adversary trivially wins the game by requesting tokens whose output differ on left and right challenge messages or have different access patterns.

```

SELUFEA(λ):
(D0, D1, (P1, ..., Pn), st) ←s A0(1λ)
(msk, mpk) ←s UFE.Setup(1λ)
b ←s {0, 1}
CT ←s UFE.Enc(mpk, Db)
for i ∈ [n]
    P̄i ←s UFE.TokenGen(msk, Pi)
b' ←s A1(CT, (P̄1, ..., P̄n), st)
return (b = b')
    
```

Fig. 8. Selective security of an Updatable FE scheme UFE.

3.2 Our construction

The idea of our construction is the following. Before encryption we append to the cleartext the token-id of the first token to be issued, the address of the first position to be read and the initial state of the program. These values are all pre-defined at the beginning. We then split the data into bits and label each of them with a common random tag, their position on the array and a counter that keeps track of how many times that bit was updated (initially 0). Then, we build a Merkle tree over the labeled bits. Later, this will allow us to check the consistency of the data without having to read through all of it. It also binds a token-id, a read-position and a state to the data at a particular stage. Finally, we encrypt each node of the tree, twice, and attach a NIZK proof attesting that they encrypt the same content. Tokens include the decryption key inside their transition circuit in order to perform the computation over the clear data and re-encrypt the nodes at the end of each CPU step. These circuits are obfuscated to protect the decryption key and the random coins necessary to re-encrypt come from a puncturable PRF. The proof then follows a mix of different strategies seen in [19,17,13,2,16].

- $\text{UFE.Setup}(1^\lambda)$ samples two public-key encryption key pairs $(\text{sk}_0, \text{pk}_0) \leftarrow_s \text{PKE.Setup}(1^\lambda)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow_s \text{PKE.Setup}(1^\lambda)$, a common reference string $\text{crs} \leftarrow_s \text{NIZK.Setup}(1^\lambda)$ and a collision-resistant hash function $\text{H} \leftarrow_s \text{H}_\lambda$. It then sets constants (l_1, l_2, l_3) as the maximum length of token-ids, addresses and possible states induced by the supported program set \mathcal{P}_λ , respectively, encoded as bit-strings. Finally, it sets $\text{msk} \leftarrow \text{sk}_0$ and $\text{mpk} \leftarrow (\text{pk}_0, \text{pk}_1, \text{crs}, \text{H}, (l_1, l_2, l_3))$ and outputs the key pair (msk, mpk) .

- $\text{UFE.Enc}(\text{mpk}, D)$ parses mpk as $(\text{pk}_0, \text{pk}_1, \text{crs}, H, (l_1, l_2, l_3))$ and appends to the memory data D the token-id 1, address 0 and the empty state ϵ , encoded as bit-stings of length l_1 , l_2 and l_3 , respectively: $D \leftarrow (D, 1, 0, \epsilon)$. (We assume from now on that $|D|$ is a power of 2. This is without loss of generality since D can be padded.) UFE.Enc sets $z \leftarrow \log(|D|)$, samples a random string $\text{tag} \leftarrow_{\$} \{0, 1\}^\lambda$ and constructs a perfectly balanced binary tree $T := \{\text{node}^{(i,j)}\}$, where leafs are set as

$$\forall j \in \{0, \dots, (|D| - 1)\}, \text{node}^{(z,j)} \leftarrow (D[j], \text{tag}, (z, j), 0)$$

and intermediate nodes are computed as

$$\begin{aligned} \forall i \in \{(z-1), \dots, 0\}, \forall j \in \{0, \dots, (2^i - 1)\}, \\ \text{node}^{(i,j)} \leftarrow (H(\text{node}^{(i+1,2j)}, \text{node}^{(i+1,2j+1)})). \end{aligned}$$

UFE.Enc then encrypts each node independently under pk_0 and pk_1 , i.e.

$$\begin{aligned} \forall i \in \{0, \dots, z\}, \forall j \in \{0, \dots, (2^i - 1)\}, \\ r_0^{(i,j)} \leftarrow_{\$} \text{RSp}_\lambda; r_1^{(i,j)} \leftarrow_{\$} \text{RSp}_\lambda \\ \text{CT}_0^{(i,j)} \leftarrow \text{PKE.Enc}(\text{pk}_0, \text{node}^{(i,j)}; r_0^{(i,j)}) \\ \text{CT}_1^{(i,j)} \leftarrow \text{PKE.Enc}(\text{pk}_1, \text{node}^{(i,j)}; r_1^{(i,j)}) \end{aligned}$$

and computes NIZK proofs that $\text{CT}_0^{(i,j)}$ and $\text{CT}_1^{(i,j)}$ encrypt the same content. More precisely,

$$\begin{aligned} \forall i \in \{0, \dots, z\}, \forall j \in \{0, \dots, (2^i - 1)\}, \\ \pi^{(i,j)} \leftarrow_{\$} \text{NIZK.Prove}(\text{crs}, x^{(i,j)}, (\text{node}^{(i,j)}, r_0^{(i,j)}, r_1^{(i,j)})), \end{aligned}$$

where $x^{(i,j)}$ is the NP statement

$$\exists(m, r_0, r_1) : \text{CT}_0^{(i,j)} = \text{PKE.Enc}(\text{pk}_0, m; r_0) \wedge \text{CT}_1^{(i,j)} = \text{PKE.Enc}(\text{pk}_1, m; r_1).$$

Finally, UFE.Enc lets

$$\text{CT} := \{(\text{CT}_0^{(i,j)}, \text{CT}_1^{(i,j)}, \pi^{(i,j)})\},$$

which encodes a perfectly balanced tree, and outputs it as a ciphertext of memory data D under mpk .

- $\text{UFE.TokenGen}(\text{msk}, \text{mpk}, P, \text{tid})$ parses $(\text{pk}_0, \text{pk}_1, \text{crs}, H, (l_1, l_2, l_3)) \leftarrow \text{mpk}$, $(\mathcal{Q}, \mathcal{T}, \mathcal{Y}, \delta) \leftarrow P$ and $\text{sk}_0 \leftarrow \text{msk}$. It then samples a new puncturable PRF key $k \leftarrow_{\$} \text{PPRF.Gen}(1^\lambda)$. Next, it sets a circuit $\hat{\delta}$ as described in Fig. 9, using the parsed values as the appropriate hardcoded constants with the same naming. UFE.TokenGen then obfuscates this circuit by computing $\bar{\delta} \leftarrow_{\$} \text{Obf}(\hat{\delta})$. Finally, for simplicity in order to avoid having to explicitly deal with the

data structure in the ciphertext, and following a similar approach as in [8], we define token \bar{P} not by its transition function, but by pseudocode, as the RAM program that executes on CT the following:

1. Set initial state $\text{st} \leftarrow \epsilon$, initial address $\text{addr} \leftarrow 0$ and empty output $y \leftarrow \epsilon$.
2. While ($y = \epsilon$)
 - (a) Construct a tree \bar{T} by selecting from CT the leaf at address addr and the last $(l_1 + l_2 + l_3)$ leafs (that should encode tid , addr and st if CT is valid), as well as all the necessary nodes to compute the hash values of their path up to the root.
 - (b) Evaluate $(\bar{T}, \text{addr}, y) \leftarrow \bar{\delta}(\bar{T})$.
 - (c) Update CT by writing the resulting \bar{T} to it.
3. Output y .

Theorem 1. *Let PKE be an IND-CCA secure public-key encryption scheme, let NIZK be a non-interactive zero knowledge proof system with perfect completeness, computational zero knowledge and statistical simulation soundness, let H be a collision-resistant hash function family, let PPRF be a puncturable pseudorandom function and let Obf be an iO -secure obfuscator that is also DI-secure w.r.t. the class of samplers described in Game₄. Then, the updatable functional encryption scheme UFE[PKE, NIZK, H, PPRF, Obf] detailed in Section 3.2 is selectively secure (as per definition in Fig. 8).*

Proof (Outline). The proof proceeds via a sequence of games as follows.

Game₀ : This game is identical to the real SEL game when the challenge bit $b = 0$, i.e. the challenger encrypts D_0 in the challenge ciphertext.

Game₁ : In this game, the common reference string and NIZK proofs are simulated. More precisely, at the beginning of the game, the challenger executes $(\text{crs}, \text{tp}) \leftarrow_s \text{Sim}_0(1^\lambda)$ to produce the crs that is included in the mpk , and proofs in the challenge ciphertext are computed with Sim_1 and tp . The distance to the previous game can be bounded by the zero-knowledge property of NIZK.

Game₂ : Let $T_0 := \{\text{node}_0^{(i,j)}\}$ be the perfectly balanced tree resulting from the encoding of D_0 with tag_0 , and $T_1 := \{\text{node}_1^{(i,j)}\}$ the one resulting from the encoding of D_1 with tag_1 , where (D_0, D_1) are the challenge messages queried by the adversary and $(\text{tag}_0, \text{tag}_1)$ are independently sampled random tags. In this game, $\text{CT}_1^{(i,j)}$ in the challenge ciphertext encrypts $\text{node}_1^{(i,j)}$; the NIZK proofs are still simulated. This transition is negligible down to the IND-CPA security of PKE.

Game₃ : In this game we hardwire a pre-computed lookup table to each circuit $\hat{\delta}_l$, containing fixed inputs/outputs that allow to bypass the steps described in Fig. 9. If the input to the circuit is on the lookup table, it will immediately return the corresponding output. The lookup tables are computed such that executing the tokens in sequence starting on the challenge ciphertext will propagate the execution over D_0 in the left branch and D_1 in the

right branch. Because the challenge ciphertext evolves over time as tokens are executed, to argue this game hop we must proceed by hardwiring one input/output at the time, as follows: (1) We hardwire the input/output of the regular execution [iO property of Obf]; (2) we puncture the PPRF key of $\widehat{\delta}_l$ on the new hardwired input [functionality preservation under puncturing of PPRF + iO property of Obf]; (3) we replace the pseudorandom coins used to produce the hardwired output with real random coins [pseudorandomness at punctured points of PPRF]; (4) we use simulated NIZK proofs in the new hardwired output [zero-knowledge property of NIZK]; (5) we compute circuit δ_l independently on the right branch before encrypting the hardwired output [IND-CPA security of PKE].

Game₄ : In all circuits $\widehat{\delta}_l$, we switch the decryption key sk_0 with sk_1 and perform the operations based on the right branch, i.e. we modify the circuits such that $\text{node}^{(i,j)} \leftarrow \text{PKE.Dec}(sk_1, CT_1^{(i,j)})$. This hop can be upper-bounded by the distributional indistinguishability of Obf. To show this, we construct an adversary $(\mathcal{S}, \mathcal{B})$ against the DI game that runs adversary \mathcal{A} as follows. Sampler \mathcal{S} runs \mathcal{A}_0 to get the challenge messages (D_0, D_1) and circuits δ_l . Then, it produces the challenge ciphertext (same rules apply on Game₃ and Game₄), and compute circuits $\widehat{\delta}_l$ according to rules of Game₃ (with decryption key sk_0) on one hand and according to rules of Game₄ (with decryption key sk_1) on the other. Finally, it outputs the two vectors of circuits and the challenge ciphertext as auxiliary information.

Adversary \mathcal{B} receives the obfuscated circuits $\widehat{\delta}_l$ either containing sk_0 or sk_1 and the challenge ciphertext. With those, it runs adversary \mathcal{A}_1 perfectly simulating Game₃ or Game₄. \mathcal{B} outputs whatever \mathcal{A}_1 outputs.

It remains to show that sampler \mathcal{S} is *computationally* unpredictable. Suppose there is a predictor Pred that finds a differing input for the circuits output by sampler \mathcal{S} . It must be because either the output contains a NIZK proof for a false statement (which contradicts the soundness property of NIZK), or there is a collision in the Merkle tree (which contradicts the collision-resistance of \mathbf{H}), or the predictor was able to guess the random tag in one of the ciphertexts (which contradicts the IND-CCA security of PKE). Note that (1) the random tag is high-entropy, so lucky guesses can be discarded; (2) we cannot rely only on IND-CPA security of PKE because we need the decryption oracle to check which random tag the predictor was able to guess to win the indistinguishability game against PKE. We also rely on the fact that adversary \mathcal{A}_0 is legitimate in its own game, so the outputs in clear of the tokens are the same in Game₃ and Game₄.

Game₅ : In this game, we remove the lookup tables introduced in Game₃. We remove one input/output at the time, from the last input/output pair added to the first, following the reverse strategy of that introduced in Game₃.

Game₆ : Here, the challenge ciphertext is computed exclusively from D_1 (with the same random tag on both branches). This transition is negligible down to the IND-CPA security of PKE.

Game₇ : In this game, we move back to regular (non-simulated) NIZK proofs in the challenge ciphertext. The distance to the previous game can be bounded by the zero-knowledge property of NIZK.

Game₈ : We now switch back the decryption key to sk_0 and perform the decryption operation on the left branch. Since NIZK is statistically sound, the circuits are functionally equivalent. We move from sk_1 to sk_0 one token at the time. This transition is down to the iO property of Obf. This game is identical to the real SEL game when the challenge bit $b = 1$, which concludes our proof. □

It is easy to check that the proposed scheme meets the correctness and efficiency properties as we defined in Section 3.1 for our primitive. The size of the ciphertext is proportional to the size of the cleartext. The size expansion of the token is however proportional to the number of steps of its execution, as the circuit $\bar{\delta}$ must be appropriately padded for the security proof.

4 Future Work

The problem at hand is quite challenging to realize even when taking strong cryptographic primitives as building blocks. Still, one might wish to strengthen the security model by allowing the adversary to obtain tokens adaptively, or by relaxing the legitimacy condition that imposes equal access patterns of extracted programs on left and right challenge messages using known results on Oblivious RAM. We view our construction as a starting point towards the realization of other updatable functional encryption schemes from milder forms of obfuscation.

Acknowledgements. The authors would like to thank Karol Zebrowski for his contribution to an earlier version of this work. Afonso Arriaga is supported by the National Research Fund, Luxembourg, under AFR Grant No. 5107187, and by the Doctoral School of Computer Science & Computer Engineering of the University of Luxembourg. Vincenzo Iovino is supported by the National Research Fund, Luxembourg. Qiang Tang is supported by a CORE (junior track) grant from the National Research Fund, Luxembourg.

References

1. P. Ananth, D. Boneh, S. Garg, A. Sahai, and M. Zhandry. Differing-inputs obfuscation and applications. IACR Cryptology ePrint Archive, Report 2013/689, 2013.
2. A. Arriaga, M. Barbosa, and P. Farshim. Private functional encryption: indistinguishability-based definitions and constructions from obfuscation. In *INDOCRYPT 2016*, vol. 10095 of LNCS, pp. 227–247. Springer, 2016. IACR Cryptology ePrint Archive, Report 2016/018, 2016.

3. P. Ananth, and A. Sahai. Functional encryption for Turing machines. In *TCC 2016*, vol. 9562 of LNCS, pp. 125–153. Springer, 2016.
4. N. Bitansky and R. Canetti. On strong simulation and composable point obfuscation. *Journal of Cryptology*, 27(2):317–357, 2014.
5. M. Bellare, V. Hoang, and P. Rogaway. Foundations of garbled circuits. In *CCS 2012*, pp. 784–796. ACM, 2012.
6. M. Bellare, I. Stepanovs and S. Tessaro. Contention in cryptoland: obfuscation, leakage and UCE. In *TCC 2015*, vol. 9563 of LNCS, pp. 542–564. Springer, 2015. IACR Cryptology ePrint Archive, Report 2015/487, 2015.
7. D. Boneh, A. Sahai, and B. Waters. Functional encryption: Definitions and challenges. In *TCC 2011*, vol. 6597 of LNCS, pp. 253–27. Springer, 2011.
8. R. Canetti, Y. Chen, J. Holmgren, and M. Raykova. Succinct Adaptive Garbled RAM. IACR Cryptology ePrint Archive, Report 2015/1074, 2015.
9. R. Canetti and J. Holmgren. Fully succinct garbled RAM. In *ITCS 2016*, pp. 169–178. ACM, 2016.
10. R. Canetti, J. Holmgren, A. Jain, and V. Vaikuntanathan. Indistinguishability obfuscation of iterated circuits and RAM programs. IACR Cryptology ePrint Archive, Report 2014/769, 2014.
11. R. Canetti, J. Holmgren, A. Jain, and V. Vaikuntanathan. Succinct garbling and indistinguishability obfuscation for RAM programs. In *STOC 2015*, pp. 429–437. ACM, 2015.
12. Y. Desmedt, V. Iovino, G. Persiano, and I. Visconti. Controlled homomorphic encryption: definition and construction. IACR Cryptology ePrint Archive, Report 2014/989, 2014.
13. S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai and B. Waters. Candidate indistinguishability obfuscation and functional encryption. for all circuits. In *FOCS 2013*, pp. 40–49. IEEE Computer Society, 2013.
14. C. Gentry, S. Halevi, S. Lu, R. Ostrovsky, M. Raykova, and D. Wichs. Garbled RAM revisited. In *EUROCRYPT 2014*, vol. 8441 of LNCS, pp. 405–422, Springer, 2014.
15. C. Gentry, S. Halevi, M. Raykova, and D. Wichs. Outsourcing private RAM computation. In *FOCS 2014*, pp. 404–4013. IEEE Computer Society, 2014.
16. V. Goyal, A. Jain, V. Koppula, A. Sahai. Functional encryption for randomized functionalities. In *TCC 2015*, vol. 9015 of LNCS, pp. 325–351. Springer, 2015.
17. Y. Ishai, O. Pandey, and A. Sahai. Public-coin differing-inputs obfuscation and its applications. In *TCC 2015*, vol. 9015 of LNCS, pp. 668–697. Springer, 2015.
18. S. Lu, and R. Ostrovsky. How to garble RAM programs. In *EUROCRYPT 2013*, vol. 7881 of LNCS, pp. 719–734, Springer, 2013.
19. M. Naor, and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC 1990*, pp. 427–437. ACM, 1990.
20. A. O’Neill. Definitional issues in functional encryption. IACR Cryptology ePrint Archive, Report 2010/556, 2010.
21. A. Sahai, and B. Waters. Fuzzy identity-based encryption. In *EUROCRYPT 2005*, vol. 3494 of LNCS, pp. 457–473, Springer, 2005.
22. A. Sahai, and B. Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *STOC 2014*, pp. 475–484. ACM, 2014.
23. A. Yao. How to generate and exchange secrets. In *FOCS 1986*, pp. 162–167. IEEE Computer Society, 1986.

Hardcoded: Transition circuit δ , token-id tid^* , secret key sk_0 , puncturable PRF key k , public keys pk_0 and pk_1 , common reference string crs , hash function H and bit-length constants (l_1, l_2, l_3) . **Input:** Tree $\bar{\text{T}}$.

1. Verify the NIZK proof in each node of tree $\bar{\text{T}}$, and decrypt the first ciphertext of each node with sk_0 . Let T be the resulting decrypted tree.

$$\begin{aligned} &\forall (i, j) \in \mathbb{N}^2 : \overline{\text{node}}^{(i, j)} \in \bar{\text{T}}, \\ &\quad \text{parse } \overline{\text{node}}^{(i, j)} \text{ as } (\text{CT}_0^{(i, j)}, \text{CT}_1^{(i, j)}, \pi^{(i, j)}) \text{ or return } \perp \\ &\quad \text{if } \text{NIZK.Verify}(\text{crs}, \text{x}^{(i, j)}, \pi^{(i, j)}) = \text{false return } \perp \\ &\quad \text{node}^{(i, j)} \leftarrow \text{PKE.Dec}(\text{sk}_0, \text{CT}_0^{(i, j)}) \\ &\text{let } \text{T} := \{\text{node}^{(i, j)}\} \end{aligned}$$
2. On the decrypted tree T , verify the path of each leaf up to the root (i.e. intermediate nodes must be equal to the hash of their children) and check that all leaves are marked with the same random tag.

$$\begin{aligned} &\mathbf{z} \leftarrow \max\{i \in \mathbb{N} : \text{node}^{(i, j)} \in \text{T}, \exists j \in \mathbb{N}\} \\ &\forall j \in \mathbb{N} : \text{node}^{(\mathbf{z}, j)} \in \text{T}, \\ &\quad \forall i \in \{\mathbf{z} - 1, \dots, 0\} \\ &\quad \quad \text{if } \text{node}^{(i, \lfloor \frac{j}{2^{\mathbf{z}-i}} \rfloor)} \neq \text{H}(\text{node}^{((i+1), 2\lfloor \frac{j}{2^{\mathbf{z}-i}} \rfloor)}, \text{node}^{((i+1), (2\lfloor \frac{j}{2^{\mathbf{z}-i}} \rfloor + 1))}) \text{ return } \perp \\ &\quad \text{parse } \text{node}^{(\mathbf{z}, j)} \text{ as } (\text{value}^{(\mathbf{z}, j)}, \text{tag}^{(\mathbf{z}, j)}, \text{position}^{(\mathbf{z}, j)}, \text{counter}^{(\mathbf{z}, j)}) \text{ or return } \perp \\ &\quad \text{if } \exists (j, j') \in \mathbb{N}^2 : \text{node}^{(\mathbf{z}, j)} \in \text{T} \wedge \text{node}^{(\mathbf{z}, j')} \in \text{T} \wedge \text{tag}^{(\mathbf{z}, j)} \neq \text{tag}^{(\mathbf{z}, j')} \text{ return } \perp \end{aligned}$$
3. Read the token-id, address and state of the current step encoded in tree T . Check that the token-id matches the one hardcoded in this token. Then, evaluate the transition circuit δ .

$$\begin{aligned} &\text{read } (\text{tid}, \text{addr}, \text{st}) \text{ with fixed bit-length } (l_1, l_2, l_3) \text{ from } \text{T} \text{ or return } \perp \\ &\text{if } \text{tid} \neq \text{tid}^* \text{ return } \perp \\ &(\text{value}^{(\mathbf{z}, \text{addr})}, \text{st}, \text{addr}, \mathbf{y}) \leftarrow \delta(\text{st}, \text{value}^{(\mathbf{z}, \text{addr})}) \end{aligned}$$
4. If the transition circuit δ outputs some result \mathbf{y} then increase the token-id and reset the internal state and address.

$$\text{if } \mathbf{y} \neq \epsilon \text{ then } \text{tid} \leftarrow \text{tid} + 1 \ ; \ \text{st} \leftarrow 0 \ ; \ \text{addr} \leftarrow 0$$
5. Write the (possibly new) token-id, address and state to tree T , update the counters of leaf nodes and recompute the path of each leaf up to the root.

$$\begin{aligned} &\text{write } (\text{tid}, \text{addr}, \text{st}) \text{ with fixed bit-length } (l_1, l_2, l_3) \text{ to } \text{T} \\ &\forall j \in \mathbb{N} : \text{node}^{(\mathbf{z}, j)} \in \text{T}, \text{counter}^{(\mathbf{z}, j)} \leftarrow \text{counter}^{(\mathbf{z}, j)} + 1 \\ &\forall j \in \mathbb{N} : \text{node}^{(\mathbf{z}, j)} \in \text{T}, \forall i \in \{\mathbf{z} - 1, \dots, 0\}, \\ &\quad \text{node}^{(i, \lfloor \frac{j}{2^{\mathbf{z}-i}} \rfloor)} \leftarrow \text{H}(\text{node}^{((i+1), 2\lfloor \frac{j}{2^{\mathbf{z}-i}} \rfloor)}, \text{node}^{((i+1), (2\lfloor \frac{j}{2^{\mathbf{z}-i}} \rfloor + 1))}) \end{aligned}$$
6. Re-encrypt all nodes of T (as before, encrypt under pk_0 and pk_1 and add NIZK proofs under crs). To extract the necessary random coins, we use the puncturable PRF under key k , providing as input the input of this circuit, i.e. $\bar{\text{T}}$.

$$\begin{aligned} &\forall (i, j) \in \mathbb{N}^2 : \text{node}^{(i, j)} \in \text{T}, (\text{r}_0^{(i, j)}, \text{r}_1^{(i, j)}, \text{r}_\pi^{(i, j)}) \leftarrow \text{PPRF.Eval}(\text{k}, (\bar{\text{T}}, (i, j))) \\ &\forall (i, j) \in \mathbb{N}^2 : \text{node}^{(i, j)} \in \text{T}, \\ &\quad \text{CT}_0^{(i, j)} \leftarrow \text{PKE.Enc}(\text{pk}_0, \text{node}^{(i, j)}; \text{r}_0^{(i, j)}); \text{CT}_1^{(i, j)} \leftarrow \text{PKE.Enc}(\text{pk}_1, \text{node}^{(i, j)}; \text{r}_1^{(i, j)}) \\ &\quad \pi^{(i, j)} \leftarrow \text{NIZK.Prove}(\text{crs}, \text{x}^{(i, j)}, (\text{node}^{(i, j)}, \text{r}_0^{(i, j)}, \text{r}_1^{(i, j)}); \text{r}_\pi^{(i, j)}) \end{aligned}$$
7. Finally, output the updated (encrypted) tree $\bar{\text{T}}$, the address for next iteration and possibly the outcome of the token.

$$\text{return } (\bar{\text{T}}, \text{addr}, \mathbf{y})$$

Fig. 9. Specification of circuit $\hat{\delta}$, as part of our updatable functional encryption scheme.