

Privacy Preserving Computation in Cloud Using Reusable Garbled Oblivious RAMs

Yongge Wang¹ Qutaibah M. Malluhi²

UNC Charlotte, NC 28223, USA

Qatar University, Qatar

Bali, Indonesia, 18-22 December 2022

Outline

- 1 Privacy preserving data storage in cloud
 - Privacy preserving data storage
- 2 Reusable garbled ORAMs
 - Reusable garbled ORAMs
 - Construction of reusable garbled ORAMs

A motivating example

- encrypt gmail at Gmail server
- search gmail
 - download gmail to local machine, decrypt it and search
 - search without downloading: search encrypted data (pattern released, some more challenges)
 - how to search over encrypted data without leaking search pattern?

Privacy preserving data storage

- Store encrypted data in cloud
- Carry out computation on these encrypted data: download the encrypted data, process data, and upload
- move computation to the data?
 - computation over encrypted data in the cloud without downloading: does not hide the access pattern to data
 - solution: use oblivious RAM techniques by Goldreich and Ostrovsky for computation over encrypted data, which provably hides all access patterns.

ORAM in cloud

- ORAM schemes require trusted CPUs
- users may not trust the CPU powers at cloud environments
- users run the trusted CPU at client site and to treat the cloud as a large random access memory storage service
- disadvantage: heavy communication overhead between the client and the cloud
- e.g., the most efficient ORAM scheme requires at least $O(\log n)$ memory accesses¹ for each individual memory access, where the cloud database contains n unit blocks of data.

¹<https://web.cs.ucla.edu/~rafail/PUBLIC/128.pdf> or
<https://eprint.iacr.org/2021/1123.pdf>

ORAM without trusted CPUs

- Lu and Ostrovsky (2013) and Goldwasser et al (2013): garbled ORAM CPU instead of using trusted CPUs
- disadvantage: the garbled RAM CPU is not succinct and is for one-time use only
- if the ORAM CPU runs t -steps for input x , then the garbled ORAM CPU for the input x is at the size of $O(t)$
- Lu and Ostrovsky asked: Is that possible to use Goldwasser et al's reusable garbled circuits to design reusable garbled ORAMs?
- Goldwasser et al's reusable garbled circuits are based on fully homomorphic encryption (FHE) and Attribute Based Encryption (ABE) schemes, which are not practical

Our approach: Reusable garbled ORAMs

- Intuition: using indistinguishability obfuscation schemes
- our first design (2015): using multilinear map based Jigsaw puzzles:
 - S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: *Proc. IEEE 54th FOCS*
- Unfortunately, all proposed multilinear maps have been broken
- good news: Jain-Lin-Sahai's (2021) indistinguishability obfuscation schemes
- Our new design: using Jain-Lin-Sahai's technique

A sample ORAM

- the first oblivious RAM by Goldreich (1987): “square root” construction.
- for a RAM machine with n memory cells denoted by an array $R[1..n]$, design an oblivious RAM with a memory array $OR[1..n + 2\sqrt{n}]$
- the portion $OR[n + \sqrt{n} + 1..n + 2\sqrt{n}]$ of size \sqrt{n} is used by the ORAM as the cache space (or a shelter).
- For the first $n + \sqrt{n}$ cells, choose a random permutation

$$\pi : \{1, \dots, n + \sqrt{n}\} \rightarrow \{1, \dots, n + \sqrt{n}\}$$

and let $OR[\pi(i)] = R[i] = (v_i, x_i)$

A sample ORAM (continued)

- Each time when the ORAM accesses a data block (v_i, x_i) from $OR[\pi(i)] = R[l]$, it stores this value (v_i, x_i) in the cache $OR[n + \sqrt{n} + 1..n + 2\sqrt{n}]$
- For each new query of a data block (v_j, x_j) , ORAM checks all values in $OR[n + \sqrt{n} + 1..n + 2\sqrt{n}]$ to see whether (v_j, x_j) has been cached there already. If the data block is found, ORAM only needs to make a dummy access to another cell $OR[\pi(n + l)]$ where l is the counter. That is, this is the l -th dummy memory cell access. If the data block is not found, ORAM loads the data block (v_j, x_j) from $OR[\pi(j)]$ directly.
- After \sqrt{n} memory cell accesses, ORAM needs to re-shuffle data blocks in the memory cells using an oblivious sorting process.

Functional encryption

Definition

A functional encryption scheme FE for a class of functions $\{\mathcal{F}_n\}_{n \in \mathcal{N}}$ is a tuple of PPT algorithms

- $(\text{fmpk}, \text{fmsk}) = \text{FE.Setup}(1^\kappa)$ outputs a master public key fmpk and a master secret key fmsk
- $\text{fsk}_f = \text{FE.KeyGen}(\text{fmsk}, f)$ outputs a secret key for a function f .
- $c = \text{FE.Enc}(\text{fmpk}, x)$ outputs a ciphertext for x .
- $y = \text{FE.Dec}(\text{fsk}_f, c)$ outputs the value y which should equal $f(x)$.

The functional encryption scheme is correct if $y \neq f(x)$ with a negligible probability.

Functional encryption security

Definition

Let FE be a FE for $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$. For a pair of PPT algorithms $A = (A_0, A_1)$ and a PPT simulator S :

$\text{Exp}_{\text{FE}, A}^{\text{real}}(1^\kappa) :$

$(\text{fmpk}, \text{fmsk}) \leftarrow \text{FE.Setup}(1^\kappa)$

$(f, \text{state}_A) \leftarrow A_1(\text{fmpk})$

$\text{fsk}_f \leftarrow \text{FE.KeyGen}(\text{fmsk}, f)$

$(x, \text{state}'_A) \leftarrow A_2(\text{state}_A, \text{fsk}_f)$

$c \leftarrow \text{FE.Enc}(\text{fmpk}, x)$

output(state'_A, c)

$\text{Exp}_{\text{FE}, A, S}^{\text{ideal}}(1^\kappa) :$

$(\text{fmpk}, \text{fmsk}) \leftarrow \text{FE.Setup}(1^\kappa)$

$(f, \text{state}_A) \leftarrow A_1(\text{fmpk})$

$\text{fsk}_f \leftarrow \text{FE.KeyGen}(\text{fmsk}, f)$

$(x, \text{state}'_A) \leftarrow A_2(\text{state}_A, \text{fsk}_f)$

$\bar{c} \leftarrow S(\text{fmpk}, \text{fsk}_f, f, f(x), 1^{|x|})$

output(state'_A, \bar{c})

FE is secure if there exists a PPT simulator S such that for all pairs of PPT adversaries (A_0, A_1) , the outcomes of the two experiments are computationally indistinguishable.

Garbled circuits

Definition

A garbling scheme for a family of circuits $\mathcal{C} = \{C_n\}_{n \in N}$ is a tuple of PPT algorithms $GC = (GC.Garble, GC.Enc, GC.Eval)$ with

- $(\Gamma, sk) = GC.Garble(1^\kappa, C)$ outputs a garbled circuit Γ and a secret key sk .
- $c_x = GC.Enc(sk, x)$ outputs an encoding c_x for an input $x \in \{0, 1\}^n$.
- $y = GC.Eval(\Gamma, c_x)$ outputs $y = C(x)$.

The garbling scheme GC is *correct* if the probability that $GC.Eval(\Gamma, c_x) \neq C(x)$ is negligible. The garbling scheme GC is *efficient* if the size of Γ is bounded by a polynomial and the run-time of $c = GC.Enc(sk, x)$ is also bounded by a polynomial.

Garbled circuits privacy

Definition

A garbling scheme GC for a family of circuits \mathcal{C} is said to be *input and circuit private* if there exists a PPT simulator $\text{Sim}_{\text{Garble}}$ such that for all PPT adversaries A and D and all large κ , we have

$$\left| \Pr[D(\alpha, x, C, \Upsilon, c) = 1 | \text{REAL}] - \Pr[D(\alpha, x, C, \tilde{\Upsilon}, \tilde{c}) = 1 | \text{SIM}] \right| = \text{negl}(\kappa)$$

where REAL and SIM are the following events

$$\begin{aligned} \text{REAL} : & (x, C, \alpha) = A(1^\kappa); (\Upsilon, \text{sk}) = \text{GS.Garble}(1^\kappa, C); c = \text{GS.Enc}(x, \text{sk}) \\ \text{SIM} : & (x, C, \alpha) = A(1^\kappa); (\tilde{\Upsilon}, \tilde{c}) = \text{Sim}_{\text{Garble}}(1^\kappa, C(x), 1^{|\mathcal{C}|}, 1^{|\mathcal{X}|}). \end{aligned}$$

Reusable garbled circuits

Definition

RGC: a reusable garbling scheme for $\mathcal{C} = \{\mathcal{C}_n\}_{n \in \mathbb{N}}$ and $C \in \mathcal{C}_n$.
 PPT algorithms $A = (A_0, A_1)$ and PPT simulator $S = (S_0, S_1)$:

$$\overline{\text{Exp}_{\text{RGC}, A}^{\text{real}}(1^\kappa)} :$$

$$(C, \text{state}_A) \leftarrow A_0(1^\kappa)$$

$$(\text{sk}, \Upsilon) \leftarrow \text{RGC.Garble}(1^\kappa, C)$$

$$\alpha \leftarrow A_1^{\text{RGC.Enc}(\text{sk}, \cdot)}(M, \Upsilon, \text{state}_A)$$

$$\overline{\text{Exp}_{\text{RGC}, A, S}^{\text{ideal}}(1^\kappa)} :$$

$$(C, \text{state}_A) \leftarrow A_0(1^\kappa)$$

$$(\tilde{\Upsilon}, \text{state}_S) \leftarrow S_0(1^\kappa, C)$$

$$\alpha \leftarrow A_1^{O(\cdot, C)[[\text{state}_S]]}(M, \tilde{\Upsilon}, \text{state}_A)$$

RGC is *private* if $\exists S, \forall A = (A_0, A_1)$, we have

$$\left\{ \text{Exp}_{\text{RGC}, A}^{\text{real}}(1^\kappa) \right\}_{\kappa \in \mathbb{N}} =_c \left\{ \text{Exp}_{\text{RGC}, A, S}^{\text{ideal}}(1^\kappa) \right\}_{\kappa \in \mathbb{N}} \quad (1)$$

Indistinguishability obfuscation

Theorem

(Jain, Lin, and Sahai 2021) Let $\tau > 0$, and $\delta, \varepsilon \in (0, 1)$. Assume the following assumptions with security parameter κ where p is a κ -bit prime, and l, k, n below are large polynomials in κ :

- *LWE assumption over Z_p with subexponential modulus-to-noise ratio 2^{k^ε} , k is dimension of LWE secret*
- *LPN assumption over Z_p with polynomially many LPN samples and error rate $1/l^\delta$, l is dimension of LPN secret*
- *the existence of a Boolean PRG in NC^0 with stretch $n^{1+\tau}$,*
- *SXDH assumption on asymmetric bilinear groups of a order p .*

Then, (subexponentially secure) indistinguishability obfuscation for all polynomial-size circuits exists.

Functional encryption scheme for circuits $C \in NC^1$


- Choose two standard public-key encryption key pairs (pk_1, sk_1) and (pk_2, sk_2) in the key generation process of the Functional Encryption scheme.
- The encryption of an input x consists of two ciphertexts of x under the two public keys pk_1 and pk_2 together with a statistically simulation sound non-interactive zero knowledge (NIZK) proof that both ciphertexts encrypt the same message.
- The secret key sk_C for the circuit C is an indistinguishability obfuscation of a program that first checks the NIZK proof and, if the proof is valid, it uses one of the two secret keys sk_1 and sk_2 to decrypt x and then computes and outputs $C(x)$.

Reusable garbled circuit \overline{C} for a circuit $C \in NC^1$

- Based on Goldwasser et al (2013): “from FE to reusable garbled circuits”
- chooses a secret key sk to encrypt C as $E.Enc_{sk}(C)$.
- $U_E(sk, x) \in NC^1$ be a UC that decrypts $E.Enc_{sk}(C)$ and runs C on x
- sk_{U_E} be the secret key of the FE scheme for U_E
- The reusable garbled circuit \overline{C} is FE secret key sk_{U_E} .
- The secret key is (sk, puk_1, puk_2)
- input to \overline{C} consists of the two cipher texts of (sk, x) under the two public keys puk_1, puk_2 and a NIZK proof.
- By combining the results in Theorem 6, De Caro et al (2013), and Goldwasser et al (2013) we have the following result: With assumptions of Theorem 6, there exists a reusable garbling scheme RGC for circuits in NC^1 that is secure according to the Definition 5 in the random oracle model.

Construction of reusable garbled ORAMs

Figure: Garbled ORAM CPU



C_{CPU} outputs an encoded $(op, v, x, ctr, flag)$

Inputs: state $(\Sigma, v, x, sctr_l_r)$ and $(session, sctr_l_s)$.

Output: Encoded (Σ, v, x) , $session$, and interface command \overline{com}

- ⓧp1 if $sctr_l_r$ and $sctr_l_s$ are inconsistent then exit.
- ⓧp2 simulate ORAM CPU for one step with state Σ and input (v, x) , compute new state Σ and next interface command (op, v, x) .
- ⓧp3 update $session, sctr_l_r,$ and $sctr_l_s$.
- ⓧp4 encode $sk || (\Sigma, v, x) || sctr_l_r$ to obtain $\overline{(\Sigma, v, x)} = (e_1^\Sigma, e_2^\Sigma, \pi^\Sigma)$ as an input to a Goldwasser's reusable garbled circuit
- ⓧp5 encode $sk || session || sctr_l_s$ to obtain $\overline{session}$ as an input to a Goldwasser's reusable garbled circuit
- ⓧp6 encode $sk || (op, v, x, ctr, flag, str_l_c)$ to obtain \overline{com} as an input to a Goldwasser's reusable garbled circuit
- ⓧp7 output $\overline{(\Sigma, v, x)}, \overline{session},$ and \overline{com} .

Inputs: $(\Sigma, v, x, \text{sctrl}_r)$, $(op, v, x, \text{ctr}, \text{flag}, \text{sctrl}_c)$,
 $(\text{session}, \text{sctrl}_s)$

Output: Memory access command (op, i, z)

- ⓧp1 if sctrl_c , sctrl_r , and sctrl_s are consistent, use session to update sctrl_c , sctrl_r , and sctrl_s and go to next step 2. Otherwise, exit
- ⓧp2 if $\text{ctr} < t$, go to step 5.
- ⓧp3 if $\text{ctr} = t$ and $op = \text{READ}$, extract (v, x) from $(op, v, x, \text{ctr}, \text{flag}, \text{sctrl}_c)$ and put it in (Σ, v, x) .
- ⓧp4 output $sk || (\Sigma, v, x, \text{sctrl}_r)$ in the format of input to an Goldwasser's reusable garbled circuit and exit.
- ⓧp5 use ctr , flag , and oblivious memory access to output (op', i, z) , where op' and i are in plain text and $z = \overline{(v, x)}$.

C_{ORAM2}

Inputs: $(op', i, z), (op, v, x, ctr, flag, sctrl_c),$
 $(session, sctrl_s)$

Output: encoded $\overline{sk || (op, v, x, ctr, flag, sctrl_c, z)}$ using
 puk_1, puk_2

- ⓧ1 if $sctrl_c$ and $sctrl_s$ are inconsistent then exit.
- ⓧ2 use $session$ to update $sctrl_c$ and $sctrl_s$.
- ⓧ3 decode z to $sk || (v', x')$ using the key prk_1 .
- ⓧ4 if $v = v'$, then the required data block has been found. Set $flag = yes$. Furthermore, if $op = READ$, insert the value of x' to the x -field of $(op, v, x, ctr, flag, sctrl_c)$.
- ⓧ5 output encoded $\overline{sk || (op, v, x, ctr, flag, sctrl_c)}$ in the format of input to an Goldwasser's reusable garbled circuit using public keys puk_1, puk_2 .

Questions

Questions?