



Virtual Secure Platform: A Five-Stage Pipeline Processor over TFHE

Kotaro Matsuoka, Ryotaro Banno, Naoki Matsumoto,
Takashi Sato, and Song Bian, *Kyoto University*

<https://www.usenix.org/conference/usenixsecurity21/presentation/matsuoka>

**This paper is included in the Proceedings of the
30th USENIX Security Symposium.**

August 11-13, 2021

978-1-939133-24-3

**Open access to the Proceedings of the
30th USENIX Security Symposium
is sponsored by USENIX.**

Virtual Secure Platform: A Five-Stage Pipeline Processor over TFHE

Kotaro Matsuoka*, Ryotaro Banno†, Naoki Matsumoto†, Takashi Sato‡, Song Bian‡

*Undergraduate School of Electrical and Electronic Engineering, Kyoto University,

†Undergraduate School of Informatics and Mathematical Science, Kyoto University,

‡Department of Communications and Computer Engineering, Kyoto University

*matsuoka.kotaro@gmail.com, †{ryotaro.banno, m.naoki9911}@gmail.com,

‡{takashi, sbian}@easter.kyoto-u.ac.jp

Abstract

We present Virtual Secure Platform (VSP), the first comprehensive platform that implements a multi-opcode general-purpose sequential processor over Fully Homomorphic Encryption (FHE) for Secure Multi-Party Computation (SMPC). VSP protects both the data and functions on which the data are evaluated from the adversary in a secure computation offloading situation like cloud computing. We proposed a complete processor architecture with a five-stage pipeline, which improves the performance of the VSP by providing more parallelism in circuit evaluation. In addition, we also designed a custom Instruction Set Architecture (ISA) to reduce the gate count of our processor, along with an entire set of toolchains to ensure that arbitrary C programs can be compiled into our custom ISA. In order to speed up instruction evaluation over VSP, CMUX Memory based ROM and RAM constructions over FHE are also proposed. Our experiments show that both the pipelined architecture and the CMUX Memory technique are effective in improving the performance of the proposed processor. We provide an open-source implementation of VSP which achieves a per-instruction latency of less than 1 second. We demonstrate that compared to the best existing processor over FHE, our implementation runs nearly $1,600\times$ faster.

1 Introduction

In a typical cloud computing scheme, clients want to offload their computations, that are, the evaluations of some programs over their private data, to some cloud server. The problem we try to tackle in this paper is to protect the programs and data of the clients against the server per se, or some third-party intruder who has physical access to the server. Since current mainstream physical processors, like Intel Xeon, cannot directly run encrypted instructions (i.e., the program to be offloaded), encrypted functions and data must be decrypted at run-time. Therefore, current cloud computing schemes suffer from side channel attacks [1, 2]. In addition, processor vendors may also plant backdoors. As a result, the cloud service

vendors and those who can physically access the servers are, in theory, able to steal the program along with the input data from the clients.

The key idea to solve the problem mentioned above is to directly run encrypted instructions [3]. In other words, the client of the cloud service sends the encrypted instructions which represent the function to be evaluated and the encrypted inputs to the cloud sever. Meanwhile, the cloud server evaluates the function using the inputs, without decryption. After the evaluation, the cloud server sends back the encrypted results to the client. During the entire evaluation process, the cloud server does not have access to any plaintext, so the evaluated function and the data are protected. The above scheme can be established by representing the processor as Boolean circuits, and the evaluation of the circuits are conducted through the use of Secure Multi-Party Computation (SMPC) protocols. Because Boolean circuits can be represented by a graph containing different types of logic gates as graph nodes (e.g., in Figure 1b), if we can perform the logical operations over encrypted input bits, we can emulate the operation of a processor by evaluating the processor circuit with the associated encrypted inputs.

Currently, we have two well-known SMPC candidates for evaluating Boolean circuits directly over encrypted inputs, namely Garbled Circuit (GC) [4] and Fully Homomorphic Encryption (FHE) [5]. GC implements SMPC operations by providing a set of encrypted truth tables for the outputs of the corresponding logic gates. During GC evaluation, the truth tables are evaluated obliviously to carry out the encrypted function evaluation. On the other hand, FHE is intrinsically more of a Secure Computation Offloading (SCO) scheme, where inputs to some public function are encrypted. The evaluator directly evaluates the public function over the ciphertexts, and returns the results to the encryption party. There are two previous works which propose to run encrypted instructions using GC: TinyGarble [6] and GarbledCPU [7]. These works implement a processor with the MIPS Instruction Set Architecture (ISA). Since most modern compilers support MIPS, both TinyGarble and GarbledCPU support the evaluation of most

conventional programs, e.g., programs written in the C language. However, in theory, we cannot achieve SCO with GC, as the generation of the GC truth tables always take more computational resources than locally evaluating the program. In contrast, as mentioned, FHE is inherently an SCO scheme [5]. Unlike GC, there is no need for tables generation for the evaluation of logic gates in FHE. To the best of our knowledge, FURISC [8] is the only previous work which implements a processor over the Smart-Vercauteren FHE Cryptosystem [9]. The processor only accepts one Turing-complete instruction, Subtract Branch if Negative (SBN). This means that it is necessary to modify modern compilers like Clang or GCC to work with FURISC, which is a highly non-trivial task. In fact, FURISC does not have any compiler support.

We propose Virtual Secure Platform (VSP), a comprehensive platform that provides a full set of tools for a complete two-party SCO scheme. Our standalone platform includes open-sourced designs and implementations of HE libraries, processor architectures, custom ISA and compiler environments. Building upon the well-known Torus Fully Homomorphic Encryption (TFHE) scheme, VSP allows any user with an arbitrary C program to execute their codes in an SCO manner. To the best of our knowledge, VSP is the fastest and most complete (in terms of the set of tool sets we provide) open-source processor platform to date.

Contributions: In brief, our contributions are as follows:

- We present VSP, the first comprehensive platform that implements a multi-opcode general-purpose sequential processor over TFHE, which enables two-party SCO. We also provide an open-source Proof of Concept (PoC) implementation of VSP, including our pipelined processor.
- We implemented the entire toolchain including a C compiler based on LLVM in order to fully support C language in VSP. The toolchain is based on our custom ISA named CAHPv3.
- We propose CMUX Memory, an optimized memory structure over TFHE. We fully leverage the Leveled Homomorphic Encryption (LHE) mode of the TFHE to ensure fast memory access, which is one of the main performance bottlenecks of VSP.
- Our open-source PoC implementation can evaluate one clock cycle of the processor in less than 1 second. This translates to nearly $1,600\times$ per-instruction latency reduction compared to FURISC, the state-of-the-art FHE-based SCO scheme.

2 Preliminaries

In this section, we define and explain some basic concepts used throughout this work. We first review the properties and constructions of HE in Section 2.1. Then, we give an overview

on the security properties of the SMPC protocols focused in this work in Section 2.2. Finally, we briefly summarize the general terminologies involved in processor designs in Section 2.3.

2.1 Homomorphic Encryption

2.1.1 Overview of Homomorphic Encryption

Homomorphic Encryption (HE) is a form of encryption which permits encrypted data to be evaluated without decryption [10]. HE can be classified into several categories depending on the types of functions that are permitted to be evaluated. A Fully Homomorphic Encryption (FHE) scheme allows one to evaluate arbitrary functions. Some popular FHE candidates include Torus Fully Homomorphic Encryption (TFHE) [11], Smart-Vercauteren Cryptosystem [9] and Brakerski-Gentry-Vaikuntanathan (BGV) [12]. All of the above mentioned candidates can evaluate arbitrary Boolean circuit over encrypted ciphertexts. Beside FHE, another category of HE is called Leveled Homomorphic Encryption (LHE). LHE has limitations on the depth of function that can be expressed, but are much faster than FHE in general. Depth here means the number of consecutive multiplications to be performed on the same ciphertext. Lastly, we note that some FHE schemes like TFHE and BGV have LHE modes.

In VSP, we cannot know *a priori* how many times we have to evaluate the circuit of the processor. This is because a general solution to the problem of determining how many clock cycles a program will take written in a Turing-complete language solves the famous Halting Problem, which is known to be undecidable. Therefore, FHE is most suitable for constructing processor-like architectures as in VSP.

Bootstrapping: Bootstrapping is one of the most important idea in the construction of FHE. It is proposed in the seminal work of Gentry [5]. The bootstrapping can be thought as evaluating a decryption function over HE. Bootstrapping is needed for FHE schemes, as we can remove the noises from the ciphertexts generated during the evaluations. Bootstrapping needs additional keys for evaluation, including an encrypted secret key.

Key switching: Key switching is a function that maps a ciphertext $Enc_{s_1}(m)$ to $Enc_{s_2}(m)$ without decryption, where $Enc_{s_i}(m)$ means encrypted m with a secret key s_i . As with bootstrapping, this function also requires an encrypted secret key, but its format is different from that required for bootstrapping, because key switching needs $Enc_{s_2}(s_1)$.

In this paper, we call the set of the keys which are required to evaluate both the bootstrapping and the key switching as Bootstrapping Key.

2.1.2 TFHE

TFHE [11, 13] is one kind of FHE. TFHE natively supports one-operand logic operations like NOT, two-operand logical

operations like NAND, NOR, XNOR, AND, OR and XOR, and the three-operand MUX. There are two reasons for choosing TFHE as a foundation of VSP. First, bootstrapping of TFHE only takes 10 milliseconds order. This is the fastest one to our best knowledge. In contrast, bootstrapping of BGV takes order of minutes with HELib [14]. Second, TFHE supports LHE mode which we find to be efficient in constructing memory units, and a detailed construction of memory units over the LHE mode of TFHE is explained in Section 8.

In what follows, we describe TFHE briefly. We will strip away unnecessary generality in order to keep the explanations straightforward.

Notations: In this work, we adopt a similar notation style as in [11], which is listed below.

\mathbb{B} : The set $\{1, 0\}$ without any structure.

\mathbb{T} : The real Torus \mathbb{R}/\mathbb{Z} , the set of real number modulo 1. In this work, we define the interval of \mathbb{T} to be $[-0.5, 0.5)$.

$\mathbb{T}_N[X]$, $\mathbb{Z}_N[X]$: The rings of polynomials $\mathbb{R}[X]/(X^N + 1) \bmod 1$ and $\mathbb{Z}[X]/(X^N + 1)$.

$\mathbb{B}_N[X]$: The polynomials in $\mathbb{Z}_N[X]$ with binary coefficients.

$1[X]$: The polynomial in $\mathbb{Z}_N[X]$ whose coefficients are all 1.

$\text{sgn}(a[X])$: The polynomial whose i th coefficient is $\text{sgn}(i$ th coefficient of $a[X])$.

E^p : The set of vectors of dimension p with entries in E .

$M_{p,q}(E)$: The set of $p \times q$ -size matrices with elements in E .

$U(E)$: The uniform distribution over E .

\leftarrow : $x \leftarrow D$ means x itself or its entries or coefficients are drawn from the distribution D .

n, N, l, α, μ : $n, N, l \in \mathbb{N}$, $\alpha \in \mathbb{R}$ and $\mu = 1/8$.

$\mathbf{a}, a[X], b[X]$: $\mathbf{a} \in \mathbb{T}^n$ and $a[X], b[X] \in \mathbb{T}_N[X]$.

Modular Gaussian Distribution: Let $k \geq 1$ and $\sigma \in \mathbb{R}^+$. For all $\mathbf{x} \in \mathbb{R}^k$, we refer to the Gaussian function of center 0 and standard deviation σ as $\rho_{\mathbb{R}^k, \sigma}(\mathbf{x}) = \exp(-\|\mathbf{x}\|^2/2\sigma^2)$. Meanwhile, $\mathcal{D}_{\mathbb{T}^k, \sigma}(\mathbf{x})$ defines a (restricted) Gaussian Distribution of center $\mathbf{0}$ and standard deviation σ over \mathbb{T}^k , and is derived by $\mathcal{D}_{\mathbb{T}^k, \sigma}(\mathbf{x}) = \sum_{l \in \mathbb{Z}} \rho_{\mathbb{R}^k, \sigma}(\mathbf{x} + l \cdot \mathbf{1})$.

TLWE: TLWE is the Torus version of the learning with errors (LWE) problem [15]. TLWE can be represented as (\mathbf{a}, b) , an $n + 1$ dimensional Torus vector. $\mathbf{s} \in \mathbb{B}^n$ is the secret key and $\mathbf{s} \leftarrow \mathbb{B}$.

Encryption: Let $e \leftarrow \mathcal{D}_{\mathbb{T}, \sigma}(x)$ and $\mathbf{a} \leftarrow U(\mathbb{T}^n)$. $m \in \mathbb{B}$ is the plaintext message. Then, $b = \mathbf{a} \cdot \mathbf{s} + \mu(2m - 1) + e$.

Decryption: Return $(1 + \text{sgn}(b - \mathbf{a} \cdot \mathbf{s}))/2$.

TRLWE: TRLWE is the Torus version of ring-LWE. TRLWE can be represented as $(a[X], b[X])$, a two dimensional Torus polynomial vector. $s[X] \in \mathbb{T}_N[X]$ represents the secret key and $s[X] \leftarrow U(\mathbb{B})$.

Encryption: Let $e[X] \in \mathbb{T}_N[X] \leftarrow \mathcal{D}_{\mathbb{T}^N, \sigma}(\mathbf{x})$ and $a[X] \leftarrow U(\mathbb{T}^N)$. $m[X] \in \mathbb{B}_N[X]$ is the plaintext message. Then, $b[X] = a[X] \cdot s[X] + \mu(2m[X] - 1[X]) + e[X]$

Decryption: Return $(1[X] + \text{sgn}(b[X] - a[X] \cdot s[X]))/2$.

TRGSW: This is a Torus and ring version of GSW, which is represented as a vector of TRLWE ciphertexts, or equivalently, a matrix of polynomials. TRGSW ciphertexts are in $M_{2,2l}(\mathbb{T}_N[X])$.

Encryption: Let $l, Bg \in \mathbb{N}$, $i \in [1, 2l]$. $e[X] \in \mathbb{T}_N[X] \leftarrow \mathcal{D}_{\mathbb{T}^N, \sigma}(\mathbf{x})$ and $a[X] \leftarrow U(\mathbb{T}^N)$. $m \in \mathbb{B}$ is the plaintext message. Then, $b_i[X] = a_i[X] \cdot s[X] + e_i[X]$ and the ciphertext \mathbf{C} is defined as follows:

$$\mathbf{C} = \begin{pmatrix} a_1[X] + \frac{m}{Bg} & b_1[X] \\ a_2[X] + \frac{m}{Bg^2} & b_2[X] \\ \vdots & \vdots \\ a_l[X] + \frac{m}{Bg^l} & b_l[X] \\ a_{l+1}[X] & b_{l+1}[X] + \frac{m}{Bg} \\ \vdots & \vdots \\ a_{2l}[X] & b_{2l}[X] + \frac{m}{Bg^l} \end{pmatrix}$$

We omit the explanation on the decryption of TRGSW as it is not needed in this paper.

Sample Extraction and Identity Key Switching: This operation converts a TRLWE ciphertext into a TLWE ciphertext. Identity Key Switching (IKS) denotes the special case of Public Key Switching where the public function is the identity function [11]. The noise variance of the output TLWE ciphertext becomes larger than the input TRLWE ciphertext because IKS adds noises. The construction of Bootstrapping in TFHE uses this as a fundamental block.

Bootstrapping in TFHE: In TFHE, Bootstrapping can be defined for TRLWE and TLWE. This can be configured by Sample Extraction (SE) and IKS at the beginning or end of the Bootstrapping procedure. The type of output ciphertext is the same as that of the input but the noise is refreshed to the level of a freshly encrypted ciphertext.

CMUX: CMUX is short for Controlled MultipleXer. This is one of the LHE-mode operations of TFHE and is equivalent to a homomorphic multiplexer. CMUX takes two TRLWE ciphertexts as inputs and a TRGSW ciphertext as its selector input. CMUX outputs a TRLWE ciphertext. The noise variance of the output TRLWE ciphertext is bigger than that of the inputs because additional noise is induced by the CMUX operation.

Homomorphic Gates: These are FHE-mode operations of TFHE and they represent logic gates. Their inputs and outputs are TLWE ciphertexts. All Homomorphic Gates except for HomNOT perform bootstrapping in their evaluation. HomNOT only negates the coefficients of its input TLWE ciphertext, so the noise variances remain the same for its input and output ciphertexts.

HomMUX without SE and IKS: This is MUX of Homomorphic Gates (HomMUX) without SE and IKS in its Bootstrapping. By definition, HomMUX without SE and IKS maps three TLWE ciphertexts to a TRLWE ciphertext. HomMUX without SE and IKS is used in the construction of our CMUX Memory.

Circuit Bootstrapping: This is a function which converts a TLWE ciphertext into TRGSW ciphertext proposed in [11]. The noise variance is always the same between the input and output of Circuit Bootstrapping, as bootstrapping is performed during the process.

Parameters of TFHE: Parameters of TFHE are one of the most important things in security analysis of VSP since they determine the security level of TFHE. In our PoC implementation, we adopt parameters recommended in [13, 16]. The estimated security of the parameter set is 80-bit [11].

2.2 Terms for Security Analysis

2.2.1 Definitions for Protocols

The main protocol we treat in this paper is two-party Secure Computation Offloading (SCO). Two-party SCO is a special case of Private Function Evaluation (PFE) [17]. To clarify the difference between VSP and GarbledCPU or TinyGarble, we also explain Private Function Secure Function Evaluation (PF-SFE).

Definition of Alice and Bob: In this paper, Bob is someone who provides most of computational resource, like cloud vendors, and Alice is someone who is the user of the cloud service and possesses the secret key. Both of them are interested in learning as much private information as possible from the other party.

Two-party SCO: In this protocol, only Alice has private information, which is a function to be evaluated along with the inputs. Furthermore, only Alice learns the result of the evaluation.

Two-party PF-SFE: In this protocol, Bob has a function to be evaluated and Alice has its input, but only Alice learns the result of the function.

2.2.2 Security Assumptions

There are two assumptions in our security analysis: the 1-circular security defined in [18], which relates to security of the TFHE scheme, and the honest-but-curious model, which limits the behavior of the adversary.

1-circular security: Circular security is classified into Key-Dependent-Message (KDM) security [18, 19]. 1-circular security means that encryption of a secret key using the secret key itself is secure. This is assumed in [11] to simplify the implementation.

Honest-but-curious model: A honest-but-curious adversary is a legitimate participant in a communication protocol who will not deviate from the defined protocol but will attempt




Figure 1: The (a) circuit and (b) graph representation of a half adder.

to learn all possible information from legitimately received messages [20].

2.3 Terms for Processor Design

We use a half adder as an example for explaining circuit-related vocabulary in this paper. A half-adder can be represented as Figure 1a. To simplify the explanation, we only use NAND and NOT gates here. We denote its input bits by A and B , and its output bits by S and C . A half adder computes a 1-bit addition. For example, if $A = B = 1$, then $S = 0, C = 1$, which calculates one plus one equals two. Let e, d, f denote intermediate outputs of the gates. If we represent a NAND gate by function $\text{NAND}(\cdot, \cdot)$ and NOT gate by $\text{NOT}(\cdot)$, we can interpret this circuit as a series of equations like the following:

$$\begin{cases} d = \text{NAND}(A, B) \\ e = \text{NAND}(A, d) \\ f = \text{NAND}(d, B) \\ S = \text{NAND}(e, f) \\ C = \text{NOT}(d) \end{cases} \quad (1)$$

Boolean Circuits over TFHE: The main idea for evaluating Boolean circuits over TFHE is replacing each logic gate in the Boolean circuit by a Homomorphic Gate from TFHE. In the half-adder circuit shown in Figure 1a, this means replacing $\text{NAND}(\cdot, \cdot)$ and $\text{NOT}(\cdot)$ in Equation (1) by equivalent TFHE operations, that is, $\text{HomNAND}(\cdot, \cdot)$ and $\text{HomNOT}(\cdot)$. Let $\text{Enc}(\cdot)$ denote encryption function of TFHE. Then, we can reinterpret Figure 1a by using the idea as follows:

$$\begin{cases} \text{Enc}(d) = \text{HomNAND}(\text{Enc}(A), \text{Enc}(B)) \\ \text{Enc}(e) = \text{HomNAND}(\text{Enc}(A), \text{Enc}(d)) \\ \text{Enc}(f) = \text{HomNAND}(\text{Enc}(d), \text{Enc}(B)) \\ \text{Enc}(S) = \text{HomNAND}(\text{Enc}(e), \text{Enc}(f)) \\ \text{Enc}(C) = \text{HomNOT}(\text{Enc}(d)) \end{cases}$$

This interpretation enables us to evaluate single-bit addition over TFHE with encrypted inputs and outputs. We can formulate, in a similar way, an entire processor circuit over TFHE.

Pipeline: Pipeline is a mechanism to increase the number of gates that can be evaluated in parallel (g) by dividing the




Figure 2: Examples of (a) unpipelined and (b) pipelined circuits.

circuit into several stages with registers. The registers hold the inputs to and outputs from the stages for synchronization.

Figure 2 shows the unpipelined and pipelined circuits. In the unpipelined circuit, $g = 2$ because only $\text{NAND}(A, B)$ and $\text{NAND}(C, D)$ can be evaluated simultaneously. Meanwhile, in the pipelined circuit, $g = 3$ because the register feeds the value to the NAND gate, such that $\text{NAND}(\text{Reg}(e), \text{Reg}(f))$, $\text{NAND}(A, B)$ and $\text{NAND}(C, D)$ can be evaluated in parallel. That is how the pipelining increases the parallelism of the processor. Lastly, we emphasize an important point that pipelining adds considerable costs to physical processor designs as physical registers need to be added to the processor circuit to enable pipelining. However, for FHE-based processors, we do not need to implement these pipeline registers using Homomorphic Gates. The intermediate ciphertexts can simply be stored into the physical memory, acting as a “pipeline register.” This reasoning holds true for all sequential elements (e.g., flip-flops) in the VSP processor architecture.

3 Related Works

There are some previous works which enable one to run encrypted programs by implementing a Boolean circuit of a processor over SMPC protocols. We only provide a brief summary on the most relevant works, and more works can be found in Appendix A.

3.1 Processor over HE

There have been a few works that have attempted to implement processors over HE to run encrypted instructions [21–24]. However, only FURISC [8, 25] represents the processor as a Boolean circuit. FURISC uses Smart-Vercauteren Cryptosystem [9, 26] to represent its processor. Smart-Vercauteren Cryptosystem is an FHE which supports XOR and AND over the ciphertexts. FURISC theoretically can be solutions for two-party SCO although it is not discussed in their paper [8]. FURISC implements an One Instruction Set Computer (OISC) processor which supports only one instruction, SBN. This means modifying modern compilers like Clang or GCC to work for it is not an easy task because it is far different from current mainstream instruction sets. In fact, there is no high-level language compiler available for FURISC. In

the experiments in Section 9, we show that VSP runs nearly $1,600\times$ faster than the estimated runtime of FURISC.

3.2 Garbled Processor

Garbled Processor is the name for the processor over Garbled Circuit (GC). There are three works, ARM2GC [27], TinyGarble [6], and GarbledCPU [7]. ARM2GC emulates an ARM processor, but it assumes the function to be evaluated as public. TinyGarble and GarbledCPU emulate a MIPS processor and enable to use conventional programming representation for two-party PF-SFE [6, 7]. The most critical weakness of Garbled Processors is that, in theory, such constructions cannot achieve two-party SCO. If Garbled Processor is used in SCO, Alice needs to generate a table of ciphertexts for all of the outputs of each gate for each clock cycle. This means Alice has to do more computationally intensive tasks than directly evaluating the function with the inputs.

4 Abstract Protocol Flow in Two-party SCO

In this section, we explain how VSP works in the two-party SCO protocol. Two-party PF-SFE can be theoretically achieved by modifying two-party SCO. See Appendix B.

Public/Private Data: The parameters of TFHE, Bootstrapping Key, the circuit of the processor, the upper-bound of the number of processor evaluation, the ciphertexts of ROM and RAM, and the sizes of ROM and RAM are public to all parties. The plaintext data of ROM and RAM data, the result of the evaluated function and the secret key are private for Alice.

4.1 Abstract Protocol Flow

The protocol flow of VSP can be divided into seven phases, and a visual depiction is shown in Figure 3. The phases are discussed as follows.

1. **Key Generation:** Alice generates a secret key.
2. **Registration:** Alice generates a Bootstrapping Key from the secret key and sends the Bootstrapping Key to Bob.
3. **Compilation:** Alice compiles the source code of the function to be evaluated into executable (instructions) for the processor using an ordinary compiler.
4. **Encryption:** Alice combines the executable with inputs, and encrypts them as ROM and RAM. The executable has a RAM part because of the initialization of global variables. In this phase, Alice also decides how many clock cycles Bob has to evaluate.
5. **Evaluation:** Bob evaluates the encrypted ROM and RAM by repeatedly evaluating the processor circuit using the TFHE ciphertexts from Alice for the designated




Figure 3: The proposed protocol flow of two-party SCO.

number of clock cycles. In this phase, what we refer to as the snapshot is also generated. A snapshot contains all necessary information for the Resumption phase, including ciphertexts of current register values, ROM and RAM.

6. **Decryption:** Alice decrypts the encrypted result using her secret key.
7. **Resumption:** Alice checks the termination flag which is included in the result. If the flag indicates that the evaluation of the function has finished, the protocol is terminated. If not, Alice re-generates the number of clock cycles Bob needs to additionally evaluate the processor circuits. Then, Bob executes the evaluation for the designated clock cycles using the information contained in the snapshot and returns to Decryption phase.

In the above procedures, 1. and 2. are needed only once. If Alice wants to evaluate multiple sets of functions and (or) inputs, the secret key and the Bootstrapping Key can be reused. Therefore, the computational and communication costs for them are negligible.

Client-Side Computation and Outsourcing: Here, we briefly show why VSP is able to provide a meaningful computation outsourcing scheme. To outsource a program in a meaningful way, the cost of client-side (i.e., Alice-side) computations for setting up the outsourcing protocol must be less than that of locally evaluating the program to be outsourced. In VSP, the client-side costs almost entirely depend on the security parameter and the size of the memory m , but not on the number of clock cycles n required to evaluate the compiled program. Therefore, for any program where $k \cdot m \leq o(n)$ for some constant k (k only depends on the security parameter), it holds that the client-side computation costs are a less than that of directly evaluating the program.

5 Security Analysis

In this section, we analyze security of VSP. We also describe the termination problem, which is one of the reasons why we assume honest-but-curious adversary model. In this paper, we also assume 1-circular security as assumed in TFHE.

5.1 Security Analysis in Two-party SCO

In this paper, we assume that Bob has physical access to the computational resource. More precisely, the assumption is that Bob can read even electric signals in the CPU dies between transistors. Therefore, any private information which is decrypted in the computational resource leaks to Bob.

Bob tries to guess Bootstrapping Key, ROM, RAM, registers, wires, etc. However, since we assume honest-but-curious adversary model, this can be reduced to the hardness of decryption of ciphertexts of TFHE in Chosen-plaintext Attack (CPA) setting. As LWE-based FHE schemes are generally based on well-established hardness assumptions, the security of VSP can be easily guaranteed.

5.2 The Termination Problem

In VSP, it is obvious that Bob cannot know if the evaluated program is halted or not, without run-time communication with Alice, as the state of the processor is entirely encrypted. The termination problem is also discussed in FURISC paper [8]. The protocol which is claimed to be a solution for the problem in the paper can be interpreted as the following procedures in VSP:

1. Bob sends to Alice a TLWE ciphertext of the termination flag. Here, the termination flag indicates if the function evaluation is finished or not.
2. Alice decrypts the termination flag and tells Bob to terminate or continue the evaluation.
3. If Alice decided to terminate the evaluation in step 2, Bob sends back the evaluation results of the function to

Alice. If Alice decided to continue, Alice re-generates the number of clock cycles and sends it to Bob. Then, Bob performs the evaluation and goes back to step 1.

This protocol is included in step 5 to 7 of the protocol flow of two-party SCO, since the ciphertext of the termination flag is included in the encrypted result. In our PoC implementation of VSP, the termination flag is (homomorphically) generated by the Instruction Decode stage of the processor.

Note that if the adversary model is not honest, Bob can try to send the Bootstrapping Key, which includes the encrypted secret key, or arbitrary ciphertexts to Alice for decryption, pretending that the TLWE ciphertext is encrypting the termination flag. As a result, to extend the threat model of VSP into a malicious setting, we need to ensure the existence of a decryption oracle and the malleability of the underlying FHE schemes are overcome. We point out that adopting IND-CCA1 FHE [28, 29] in combined with Verifiable Computation [30] can be a candidate solution for VSP in a malicious setting, and is one of our future works.

6 Design and Implementation of VSP

In this section, we explain how we designed and implemented VSP [31].

6.1 Design Goals

The following three design goals are prioritized during the design of VSP.

(i) C compatibility

Since it is obviously difficult to actually adopt a secure framework if the framework is inconvenient to use, we decided to support high-level program representations so that users can use VSP with ease. There are two reasons why we chose the C language as our high-level representation. First, C is one of the most widely used programming languages. Second, the C language is designed to be fast, where extensive optimizations have been devoted into the optimization of C-based programs, e.g., the LLVM framework [32]. Therefore, with C support, users of VSP can have easy access to efficient programs.

(ii) ISA Optimization

Due to the high computational demand, the number of logic gates that can be evaluated in parallel (g) over TFHE is limited by the number of parallel processing capacity of the physical machine. In VSP, the evaluation time of the circuit is proportional to the total number of gate count (t), as g of a processor generally exceeds the parallel processing capacity of an ordinary desktop computer. Since the ISA plays a key role in determining t of the processor, we decided to design our custom ISA in such a way that the circuit of the processor can be minimized, while retaining C compatibility.

(iii) Maximizing Parallelism

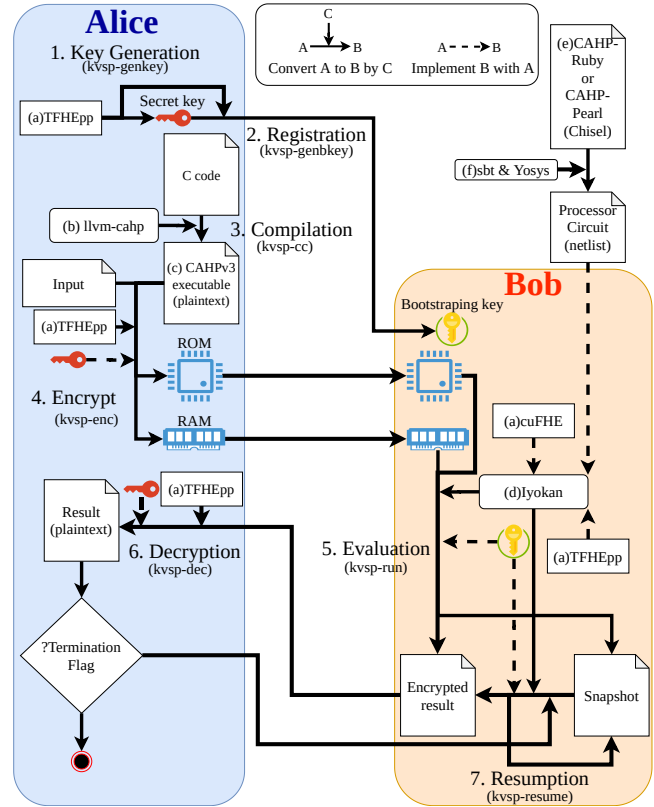


Figure 4: The VSP architecture and the main procedural flow.

As mentioned, the amount of parallelism in VSP (and generally in lattice-based cryptography) exceeds the parallel processing capability of conventional desktop computers. However, we point out that cloud vendors may have much more computational resources available than a home computer. To fully leverage the computational resources available in data centers, we designed the processor architecture in VSP to have a pipelined structure, where the processor circuit is divided into different pipeline stages that can be evaluated simultaneously. We assert that the pipeline technique does affect the execution time when the physical machine does not have much parallel processing capability. However, the runtime of VSP can be significantly reduced by the pipeline technique when there are enough physical processor cores.

6.2 The Architecture of VSP

Notation: In this paper, physical machine is the actual processing unit that runs VSP. CPU and GPU refers the physical CPU or GPU in the physical machine. In contrast, we use processor to refer to the virtual processor constructed over TFHE in VSP.

The visual overview of the implemented protocol flow of the proposed VSP framework is given in Figure 4, which details the abstract protocol flow in Figure 3. Table 1 shows

Table 1: The Phases of VSP and the Associated Subcommands of the Command-Line Interface kvsp

Phase	Key Generation	Registration	Compilation	Encryption	Evaluation	Decryption	Resumption
Subcommand	kvsp-gen	kvsp-genbkey	kvsp-cc	kvsp-enc	kvsp-run	kvsp-dec	kvsp-resume
Modules	(a)	(a)	(b), (c)	(c)	(a), (d)	(a)	(a), (d)

which phase each subcommand of kvsp [31] (a command-line user interface for VSP) corresponds to and by which module is called. Each subcommand of kvsp takes its inputs as a file, and outputs its results to a file. Therefore, the communication between the parties can be done via files transferring through public channels.

We first describe how the modules (a)-(f) are used here. Then, we explain each module. In this work, we name our proposed processor circuits as (e) CAHP-Ruby and CAHP-Pearl, and the details on the circuits are explained in Section 7. It is assumed that Alice and Bob agree on which processor architecture will be used in advance. (f) sbt [33] and Yosys [34] are used to convert the Chisel code for the processor into a JSON netlist. Here, the netlist is a graph of nodes, where each node corresponds to a logic gate. The netlist is provided to Bob before the start of the protocol. In Key Generation phase, Alice uses (a) TFHEpp, a C++ implementation of TFHE on CPU, to generate a secret key. In Registration phase, Alice uses (a) TFHEpp one more times to generate the Bootstrapping Key from the secret key and sends the Bootstrapping Key to Bob. In Compilation phase, Alice uses (b) llvm-cahp, our C compiler for our custom ISA called (c) CAHPv3, to generate executable binaries. In Encryption phase, Alice uses (a) TFHEpp to encrypt the executable binaries and the input into encrypted ROM and RAM. Then, Alice sends the ROM and RAM to Bob. In Evaluation phase, Bob uses (d) Iyokan to evaluate the processor circuit netlist over TFHE with the given encrypted ROM and RAM data. (a) TFHEpp and cuFHE, the CUDA implementation of TFHE on GPU, are used in (d) Iyokan to perform homomorphic computations. Then, Bob sends back the encrypted result to Alice. In Decryption phase, (a) TFHEpp is used to decrypt the encrypted result. In Resumption phase, Alice checks the termination flag in the result. If it is 1, she terminates the protocol. Otherwise, Alice tells Bob to resume evaluation. Bob again runs (d) Iyokan for the new round of homomorphic evaluation.

(a) TFHEpp and cuFHE: The TFHE libraries

TFHEpp is our fully-scratch C++17 implementation of TFHE on the CPU, while cuFHE is a TFHE library on the GPU (we optimized the original TFHE library from [35, 36]).

In general, cuFHE is faster than TFHEpp, especially when multiple logic gates are run in parallel, as the throughput of GPU is higher than that of CPU. We describe how we use these libraries in (d) Iyokan.

While TFHEpp supports Circuit Bootstrapping, which is a

necessary component of CMUX Memory, cuFHE does not. cuFHE uses Number Theoretic Transform (NTT) to perform fast polynomial multiplication, where the ciphertext modulus is kept to be $2^{64} - 2^{32} + 1$. This bit width constraint is to ensure that the operands involved in NTT fit into the multiply instructions on the GPUs. Unfortunately, due to this bit width constraint, cuFHE cannot directly perform Circuit Bootstrapping, as the moduli required by Circuit Bootstrapping needs to be larger than 64-bit. While we can simply increase the size of modulus to be compatible with Circuit Bootstrapping, the performance of cuFHE in practice will be significantly reduced as more multiplication instructions (on the GPUs) and memory accesses are required to perform a single polynomial multiplication operation. The efficient implementation of Circuit Bootstrapping on GPUs currently remain as an open field of study.

(b) llvm-cahp: The C Compiler

We implemented a new C compiler *llvm-cahp* for our ISA, CAHPv3, using LLVM9.

The LLVM compiler infrastructure project is an assemblage of compiler and toolchain technologies [32], which serves as a good foundation for our custom processor architecture and ISA. LLVM is widely used in both open and closed projects as well as used in academia [37]. In particular, LLVM surpasses GCC to win the ACM Software System Award in 2012 [38]. LLVM includes four parts. First, we have language-dependent frontends that compile the program source code into the intermediate representation named LLVM IR. Second, LLVM has a target-independent optimizer that operates on LLVM IR. Third, the LLVM target-specific backends are used to generate the object code of each target from LLVM IR. Finally, the LLVM linker turns multiple object codes into one executable. Since we defined a custom ISA, we implemented a new backend for CAHPv3. We also added support for CAHPv3 to the frontend of the C language (i.e., Clang), and to the LLVM linker (i.e., LLD). By putting them together, we can directly compile C program into a CAHPv3 executable binary file.

Our compiler supports almost all features of C such as basic arithmetic operations, control expressions, function calls including recursion, structures, and so on. Furthermore, since LLVM has the target-independent optimizer as mentioned above, llvm-cahp can output fast and small executables by using the `-O3` or `-Oz` compiler options.

Since the proposed processor is a virtual one, our modified compiler does not provide functions in standard libraries that

require physical processor components (e.g., the print function). There are also some minor limitations (e.g., jump over 1kiB) in our compiler.

(c) CAHPv3: Instruction Set Architecture

CAHPv3¹ is our RISC ISA based on RISC-V 32-bit integer and 16-bit compressed instructions (RV32IC). CAHPv3 has 16-bit datapath and sixteen 16-bit registers. However, the instruction bit width is a mixture of 24 bits and 16 bits, since we want to minimize the size of the machine code.

CAHPv3 has two important features from the perspectives of our design goals. First, it is relatively easy to implement the LLVM backend for CAHPv3, due to its similarities to mainstream ISAs such as x86 and RISC-V. We note that this is one of the main reasons why the OISC used in FURISC is considered impractical. Second, CAHPv3 reduces the complexity of the processor circuitry because it is a RISC ISA, and the datapath is only of 16-bit wide. Unlike RV32IC, CAHPv3 does not include instructions that are not necessary in VSP, such as privileged instructions and synchronization instructions, further reducing the total gate count. The specification is here [39].

(d) Iyokan: The Gate Evaluation Engine

Iyokan is our main software written in C++17 to run the processor over TFHE. The fundamental features of Iyokan are to receive an arbitrary Boolean circuit along with the encrypted input data, evaluate the circuits according to the inputs over TFHE, and return encrypted results of the evaluation. Therefore, we can execute encrypted programs without decryption by feeding Iyokan with the processor as a logical circuit and the associated inputs.

Iyokan works in the following way:

1. Split the input sequential logical circuit into two parts: combinational circuits and flip-flops to represent general Boolean circuits.
2. Convert the combinational circuits into a directed acyclic graph (DAG), where the logical gates are represented as graph nodes, and wires as directed edges. Figure 1b shows an example graph representation of the half adder circuit in Figure 1a.
3. Evaluate the DAG by using the converted circuit along with its inputs and the outputs of the flip-flops. Since every node in the DAG has to be evaluated, Iyokan uses the list scheduling algorithm to assign the tasks to workers which are physical CPU and GPU processing units. Note here that the scheduling algorithm also needs to resolve the dependency relations between nodes represented as

¹CAHP is short for “CAHP Ain’t for Hardware Processors,” and v3 means this is our third version ISA for VSP (the former two did not work well).

edges in the DAG. Almost all the tasks are executed on GPUs via cuFHE, and the rest of the tasks which cannot be run on GPUs, such as Circuit Bootstrapping, are executed on CPUs via TFHEpp.

This step gives us the output of the combinational circuit in the current cycle, which is used as the inputs to the flip-flops.

4. Save the inputs the previous step provides to the flip-flops (physical memories).
5. Output the stored values in the flip-flops.
6. Exit if the number of clock cycles exceeds the threshold which is specified by the user through command-line option. Otherwise, go to step 3.

Each evaluation from step 3 to 6 corresponds to one clock cycle. As mentioned in Section 4.1, Alice has to decide a threshold, that is, how many times the steps between step 3 and 6 should be repeated.

There are two important features of Iyokan. First, Iyokan can handle not only normal logic gates but also CMUX Memory. CMUX Memory can be represented as a scheduled graph, so it can also be embedded in the DAG. Second, Iyokan can run more than one worker on CPUs and GPUs in parallel.

(e) CAHP-Ruby, CAHP-Pearl: Processor

We developed two processors, CAHP-Ruby and CAHP-Pearl, for VSP:

CAHP-Ruby CAHP-Ruby is a 5-stage pipeline processor that implements CAHPv3 ISA. We will explain its details in Section 7.

CAHP-Pearl CAHP-Pearl is a single cycle processor that also implements CAHPv3 ISA. We made it by just removing pipeline registers from CHAP-Ruby.

(f) sbt and Yosys: Logic Synthesis

We chose Chisel [40], a particular Hardware Description Language (HDL), to instantiate our processors for VSP, as Chisel is widely adopted in the industry [41]. The sbt program compiles Chisel to the Verilog HDL. Then, Yosys [34] is utilized to compile Verilog codes into JSON netlists.

7 The Proposed Processor Architecture

Figure 5 conceptually illustrates CAHP-Ruby, the proposed custom processor architecture. CAHP-Ruby has a five-stage pipeline structure consisting of an *Instruction Fetch*, an *Instruction Decode*, an *Execution*, a *Memory Access*, and a *Write*




Figure 5: The architecture of the five-stage pipelined CAHP-Ruby processor.

Back stage. We chose a five-stage construction, as this structure is widely used in physical processor designs [42–45]. Determining the optimal number of pipeline is actually platform-dependent, i.e., it depends on the physical resources available to VSP. A framework that automatically optimizes the number of pipeline stage is one of our main future works.

CAHP-Ruby has two different memory areas: ROM and RAM, as shown in Figure 5. This structure greatly simplifies the processor circuitry and enables each memory area to have different and optimized implementations, further discussed in Section 8. Here, ROM is a read-only memory area, designated for the compiled instructions. RAM permits both read and write operations, and is mainly for program data handling. We note that CAHP-Ruby does not support any peripheral devices nor interruption because they are not needed in a virtual processor. Through such design decisions, we are able to reduce the complexity of the CAHP-Ruby circuitry.

In what follows, we detail the operational behavior of each of our custom processor stages.

Instruction Fetch (IF): IF is responsible for producing an instruction. First, IF fetches a 32-bit block from ROM. However, the block may not contain any complete instructions due to the fact that our custom ISA contains both 16-bit and 24-bit instructions. Therefore, IF includes a 32-bit instruction cache to resolve this 24/16-bit boundary alignment problem. The cache contains ROM output value of the previous clock cycle. If the currently fetched 32-bit ROM block does not contain a complete instruction, data from the instruction cache can be read, and it is guaranteed that there will always be a complete instruction in a 64-bit ROM block. Therefore, IF constructs a complete instruction with the assistance of the instruction cache and the current ROM output value.

Instruction Decode (ID): ID decodes the instruction to provide operands for the execution stage. This stage also reads the data from the registers specified by the instruction in the main register file. ID is also responsible for generating the termination flag. In this work, we indicate a program termination by inserting a jump instruction which jumps to the same its own memory address, creating an infinite loop. Once the ID stage detects such loop, the termination flag is set, and can be read from the dedicated port.

Execution (Ex): This stage consists of an arithmetic and logical unit (ALU) and a branch controller. ALU performs

(homomorphic) arithmetic operations such as addition and subtraction, and logical operations such as logical summation, and shift. In the case of a jump instruction or a branch instruction, the branch controller generates a flag indicating whether to jump or not according to the result of the ALU operation. We assert that all the computations and branches are over FHE ciphertexts, guaranteeing that the processor circuit evaluator does not observe any private information.

Memory Access (Mem): This stage consists of two parts: memory controller and RAM. We defer a detailed presentation of the RAM in Section 8. The memory controller takes write data from the execution stage as its input. When the write data is 8-bit wide, the controller converts the write data to be of 16-bit wide, for the RAM only accepts 16-bit data. The memory controller also reads the data from the read port of RAM and format when the output value to be of 8-bit wide. Finally, the memory controller passes the read data from the RAM to the write back stage.

Write Back (WB): This stage simply writes data into the main register files.

8 CMUX Memory

In this section, we present CMUX Memory, a new construction of memory unit over HE that leverages the LHE mode of TFHE for optimization. As mentioned, there are two types of memories: RAM and ROM.

8.1 Theoretical Speed Predictions

Informally, the reason why CMUX Memory is fast can be explained by the fact that the evaluation of Circuit Bootstrapping takes about 10 times as long as it takes to evaluate any two-input homomorphic gate. Let $v, w \in \mathbb{N}$ be the number of bits of the address and the data bus, respectively. Assuming that we ignore the time it takes to process CMUX, because CMUX is several hundred times faster than any two-input homomorphic gate, the time it requires to evaluate the ROM of the CMUX Memory is roughly equivalent to $10v + w$ Homomorphic Gates. Meanwhile, the time it takes to evaluate the RAM is roughly equivalent to $10v + w(2^v + 1)$ Homomorphic Gates. The w term comes from HomMUX without SE and IKS and the $w \cdot 2^v$ term comes from the noise refreshment. The construction of the ROM and RAM by logic gates takes at least $w \cdot 2(2^v - 1)$ two-input Homomorphic Gates each to construct the tree for data fetching. Therefore, in theory, CMUX Memory can be expected to be faster than constructing the memory by logic gates.

8.2 RAM

In this paper, we only treat one cycle, single port RAM since it requires the minimum amount of Bootstrapping to implement. The RAM has following characteristics: (i) Read and write




Figure 6: The architecture of the one-cycle single-port RAM.

are exclusive. (ii) Both read and write are done in one cycle. (iii) Read and write use the same address.

The visual overview of the RAM architecture is given in Figure 6. There are three inputs for RAM: address, write flag, and write data. The address is the memory address for write or read. The write flag is one bit data which selects the operation mode of RAM. The write data is the data which will be written in the address if the RAM is write mode. The RAM has one output port, where the data presented at the input address are retrieved. In VSP, the data bus in the processor uses TLWE as ciphertexts for memory elements, since TLWE ciphertexts are also used by the Homomorphic Gates in other parts of the processor circuit.

As shown in Figure 6, RAM consists of the read unit, the control unit, and the write unit. In what follows, we provide a comprehensive explanation on each of the unit. Note that addresses in the write and the read units are in TRGSW ciphertext, but the memory controller in Memory Access stage of the processor feeds the address as TLWE ciphertexts. Therefore, Circuit Bootstrapping is applied to TLWE ciphertexts to get TRGSW ciphertexts representations of the addresses.

Read Unit

The read unit reads the data at a given address. The visual overview of its architecture is given in Figure 7 and 8. The data of RAM are represented as $w \cdot 2^v$ TRLWE ciphertexts, where each TRLWE ciphertext contains one bit of plaintext information. The TRLWE ciphertexts are divided into w blocks, where the i th block contains the i th bit of each word. A CMUX tree is used to fetch the i th bit of the word from the i th block. We note that, although the message space of TRLWE is capable of holding a vector of N binary values, i.e., $\mathbb{B}_N[X]$, we only fill one entry with an actual plaintext value. If we pack multiple bits into a single TRLWE ciphertext for read, we also have to write in a packed manner. The problem for pack writing is that every instruction might have a chance to write only a small amount (e.g., a 16-bit register) of data to RAM,

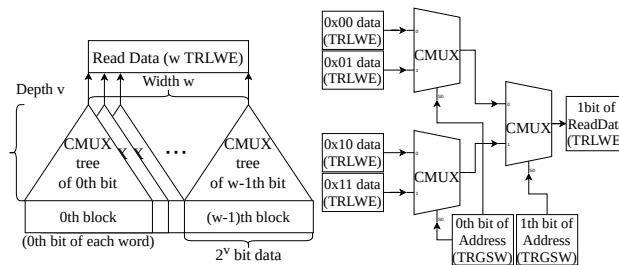


Figure 7: The architecture of the read unit. **Figure 8:** An example of the CMUX tree ($v = 2$).

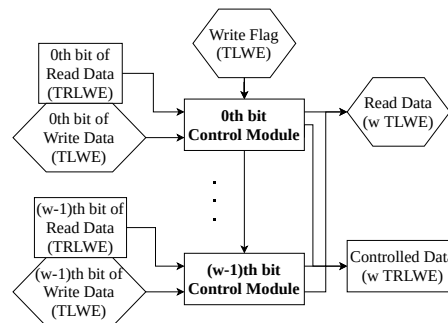


Figure 9: The architecture of the control unit.

and the amount of computations it takes to pack and unpack bits can be a significant overhead.

The task of the CMUX tree is to compare each bit of the address with that of RAM data by a tree of multiplexers implemented by CMUX, such that data at the designated memory address can be read.

Control Unit

The control unit is the interface between main processor circuit and CMUX Memory. We show an architectural illustration of the control unit and module in Figure 9 and 10, respectively. The control unit consists of w control modules, each of which processes a single bit of the write data. Since the processor only accepts TLWE ciphertexts, SE and IKS are inserted to convert the read data from the read unit into TLWE ciphertexts. The control module performs multiplexing between the read and the write data, depending on the write flag. The multiplexed result is sent to the write unit as the controlled data.

Write Unit

From the view of the main processor circuit, each word of the current cycle data is the multiplexed result between the word of the previous cycle data and the write data depending on the write flag and the address matching. Since the multiplexed result depends on the write flag that is fed as the controlled data, the write unit only needs to take care of




Figure 10: The architecture of a control module.




Figure 11: The architecture of the write unit.

the address matching part of the computation. The write unit also performs Bootstrapping over the entire contents of the RAM. An visual overview of the write unit is given in Figure 11, 12, and 13. The write unit consists of w write blocks, each handles a single word. Each write block is composed of 2^v write bars which handles a single bit. Therefore, the write unit consists of $w \cdot 2^v$ write bars arranged in parallel.

The working principle of the write bar is comparing each bit of the input address with the addresses in the RAM, through an array of CMUX gates. If all bits in the input address match a particular entry in the RAM, the controlled data is selected and becomes current cycle data. Here, when the write flag is false, the controlled data is same as the previous cycle data in the address, so current cycle data is same as previous one. On the other hand, if the write flag is true, the controlled data and current cycle data both become the write data, and the data are written into the memory. If the addresses do not match, previous cycle data is selected, and data in memory are not modified. The write bar refreshes the noise added by the CMUX array by bootstrapping the data at the end.

Remark: The implementation of the comparison between an input address bit and a constant address bit is, in fact, quite simple. More specifically, the comparison result between an input bit with a constant value of 1 is the bit itself. Meanwhile, the comparison with 0 can be implemented by a subtraction of a constant TRGSW ciphertext encrypting the constant 1 followed by a sign inversion of all coefficients in the resulting TRGSW ciphertext.




Figure 12: An example of the write block ($v = 2$).




Figure 13: An example of a write bar at address 0x01 ($v = 2$).

8.3 ROM

The construction of ROM with LHE mode of TFHE is trivial by using Look Up Table (LUT), which is described in [11]. We applied both optimization techniques mentioned in [11], namely Vertical Packing and Horizontal Packing.

9 Evaluation

In this section, we perform thorough experiments on VSP to demonstrate its performance. We will first characterize VSP over a set of benchmarks in Section 9.1, and then deliver the overall performance statistics in Section 9.2

9.1 Benchmarks

Benchmark environments

In our implementation, ROM and RAM are 512 bytes, that is, $v = 8$ and $w = 16$ when using the CMUX Memory for RAM. We also experimented 1 KiB ones. See Appendix C for the details.

The main benchmark program used in our evaluation is Hamming. Hamming takes two 8-digit hexadecimal numbers a and b as its arguments, and finds the Hamming distance between them. The programs are compiled into CAHPv3 executable by `llvm-cahp` with `-Oz` optimization flag, which minimizes the size of machine code. Then, the compiled programs are encrypted and executed on Iyokan with CAHP-Ruby (with pipeline) and CAHP-Pearl (without pipeline). The scripts to reproduce the runtime performance evaluation is available at [48].

We used four types of machines to evaluate VSP:

AWS c5.metal An HPC server hosted by Amazon Web Service equipped with Intel Xeon Platinum 8275CL CPU

Table 2: Processor Size Evaluation

Processor	MUX	NOT	Others
CAHP-Ruby	996	15	2422
CAHP-Pearl	877	22	2054
Lite MIPS [46]	1276	39	6241
PicoRV32 [47]	2732	11	5162

Table 3: Size of Keys and Ciphertexts

Type	Size[MiB]
Secret Key	0.023
Bootstrapping Key	2563.047
ROM	0.033
RAM	33.55

(96 vCPUs), 92GiB RAM, and no GPUs.

Sakura Koukaryoku An HPC server hosted by Sakura internet Inc. equipped with Intel Xeon CPU E5-2623 v3 (16 vCPUs), 128GB RAM, and single NVIDIA Tesla V100.

AWS p3.8xlarge An HPC server hosted by Amazon Web Service equipped with Intel Xeon CPU E5-2686 v4 (32 vCPUs), 244GB RAM, and 4 NVIDIA Tesla V100.

AWS p3.16xlarge An HPC server hosted by Amazon Web Service equipped with Intel Xeon CPU E5-2686 v4 (64 vCPUs), 488GB RAM, and 8 NVIDIA Tesla V100.

Runtime Performance Evaluation

Table 5 shows the run-time statistics required to evaluate the encrypted program of Hamming. Here, sec./cycle stands for seconds per clock cycle, which is the amount of program run-time divided by the number of required clock cycles.

While pipelining increases the number of gates of the processors, the technique enables more gates to be run in parallel. Therefore, when the physical machine has enough parallel processing units, pipelining reduces per-clock-cycle run-time of VSP, and eventually results in decreased total run-time (Compared between Cases #4 and 6, 7 and 8, 9 and 11, and 10 and 12). On the other hand, when the physical machine is not so powerful (Cases #1 and 2), the runtime ends up being slower due to the increased number of clock cycles. In addition, in Cases #3 and 5, the CMUX Memory is turned off and the machine does not have enough parallel processing units to fully parallelize the gates in ROM and RAM. Consequently, the physical processors do not have more machine resources for evaluating the pipelined core processor circuit.

Finally, we observe that while AWS p3.8xlarge (4 V100) is much faster than Sakura Koukaryoku (single V100), there

Table 4: Machine Code Size

Program	RV32IC [B]	CAHPv3 [B]
Fibonacci	36	31
Hamming	354	264
Brainf*ck	226	229

is almost no difference between AWS p3.16xlarge (8 V100) and p3.8xlarge. This is most likely caused by the fact that the parallel processing capabilities of both machines well exceed the number of logic gates that can be evaluated in parallel in our processor. Therefore, further pipelining may be conducted on such powerful computing platforms.

Besides pipelining, we also experiment on the performance impact of the proposed CMUX Memory. As shown in Table 5, CMUX Memory reduces runtime across all cases we tested. When CMUX Memory is not used, ROM and RAM need to be implemented by the Homomorphic Gates in the FHE mode of TFHE, which results in significant performance degradations.

The fastest instance we tested is Case #12, that is, AWS p3.16xlarge with pipelining and CMUX Memory applied, which is shown in bold in Table 5. We achieved a performance of about 0.8 sec./cycle, or equivalently, 1.25 Hz. From the results of the benchmark, we conclude that both pipelining and CMUX Memory are effective in improving the performance of VSP.

Processor Size Evaluation

In general, fewer logic gates means fewer computational complexity, so the total gate count of the processors is one of the most important factors which determine the performance of VSP. Table 2 shows the size of CAHP-Ruby and CAHP-Pearl. In the table MUX and NOT are counted separately because their performance characteristics are different from a normal homomorphic gate. In particular, MUX is twice as slow as other homomorphic gates, even with the cryptographic optimization proposed in [11]. Meanwhile, NOT can be evaluated much faster than other gates, as the only operations in a NOT gate are sign inversions. We compare the gate count of our processor to that of Lite MIPS [46] and PicoRV32 [47]. Lite MIPS is the processor which is implemented in TinyGarble [6]. PicoRV32 is one of open-source implementations of RISC V, where the design goal is to implement a small (in terms of gate count) processor. As shown in Table 2, our processors are smaller than both of the existing designs.

Data size Evaluation

We used two more programs except Hamming: Fibonacci and Brainf*ck here. Fibonacci takes a decimal digit n as its command-line argument, and calculates n th Fibonacci number. Here we used $n = 5$. Brainf*ck interprets code of brainf*ck,

Table 5: Performance Evaluation Using Hamming

Case #	Machine	# of V100	Pipelining?	CMUX Memory?	# of cycles	Runtime [s]	sec./cycle	# of tries
1	AWS c5.metal	0	No	Yes	936	2342.0 ± 13.3	2.502 ± 0.014	3
2			Yes	Yes	1216	2773.0 ± 2.8	2.280 ± 0.002	3
3	Sakura Koukaryoku	1	No	No	936	5919.0 ± 33.1	6.324 ± 0.035	5
4			No	Yes	936	2232.1 ± 1.7	2.385 ± 0.002	5
5			Yes	No	1216	7809.0 ± 45.8	6.422 ± 0.038	4
6			Yes	Yes	1216	2045.0 ± 4.6	1.682 ± 0.004	5
7	AWS p3.8xlarge	4	No	Yes	936	1455.7 ± 0.3	1.555 ± 0.000	3
8			Yes	Yes	1216	979.0 ± 12.5	0.805 ± 0.010	3
9	AWS p3.16xlarge	8	No	No	936	1627.0 ± 4.2	1.739 ± 0.004	3
10			No	Yes	936	1440.0 ± 2.5	1.538 ± 0.003	3
11			Yes	No	1216	1566.0 ± 9.7	1.288 ± 0.008	3
12			Yes	Yes	1216	965.9 ± 3.4	0.794 ± 0.003	3

which is a esoteric programming language, and returns the result. We inputted `++++[>+++++<-]>+++` to it, the result of which is 42.

Table 3 shows the size of keys and ciphertexts. We can see Bootstrapping Key is significantly bigger than other parts, so reusing Bootstrapping Key can reduce communication cost greatly. The reason why RAM is about 1024 times bigger than ROM is that Vertical Packing [11] is not applied to RAM. Table 4 shows the machine code size of the programs in CAHPv3. We also show RV32IC version for reference. RV32IC has more registers than CAHPv3 does, so register spills more often occurs in CAHPv3, which made code of Brainf*ck larger.

Client-side Cost Evaluation

It is noted that, on p3.8xlarge, it takes Alice (i.e., the client) about 57 seconds to complete the generation of the Bootstrapping Key, encryption of the memory, compilation of the program and the decryption of the evaluation results. Among these, the generation of Bootstrapping Key is the most time consuming procedure, where it takes about 55 seconds to finish. For simple programs like Hamming, evaluating the program locally by the client only takes around 0.5 microseconds on a conventional CPU, and program outsourcing in such case provides no practical merit. However, as discussed in Section 4.1, for programs that potentially contain infinite loops, VSP can obviously reduce the amount of client-side computations. Therefore, exploring practical applications of VSP is one of our main future works.

9.2 Overall Performance and Comparison to Existing Works

Because FURISC [8, 25] is the only previous work which represents the processor as a Boolean circuit and evaluates it over FHE, we compare FURISC as the-state-of-the-art to VSP

Table 6: Comparison between VSP and FURISC

Name	sec./cycle	# of instructions	Implementation
VSP	0.8	Small	Public [31]
FURISC	1278 (est.)	Large	Private

in Table 6. FURISC gives the FPGA-accelerated evaluation time for Subtraction, in Table 6.5 in [8]. Because SBN is the only instruction FURISC supports, the evaluation time of Subtraction corresponds to one clock cycle in the FURISC processor. Therefore, the estimated time for evaluating one clock cycle of FURISC is 21.3 minutes, over 1000 seconds. In contrast, our VSP implementation can evaluate one clock cycle in 0.8 seconds, as shown in Case #12 in Table 5.

The number of instructions for representing (almost all) programs in FURISC are larger than that of VSP, for that FURISC has an OISC ISA. Therefore, we can see that compiling the same program on VSP results in a smaller number of instructions, and that each instruction runs nearly 1600× faster than FURISC. Hence, we are confident that the open-source VSP is the fastest FHE-based SCO platform to date.

10 Conclusion

In this work, we presented VSP, the first comprehensive platform that implements a multi-opcode general-purpose sequential processor over TFHE for two-party Secure Computation Offloading (SCO). We proposed a complete SCO scheme and designed a custom five-stage pipelined processor along with a custom ISA CAHPv3. We also proposed CMUX Memory, the optimized structure of ROM and RAM over TFHE to speed up instruction evaluation. We thoroughly evaluated VSP on benchmarks to show that both pipelining and CMUX Memory are effective in speeding up VSP. Our open-source implementation is nearly 1600× faster than the-state-of-the-

art implementation while accepting conventional C language programs.

11 Acknowledgement

We thank all the people including our shepherd Thomas Ristenpart and anonymous reviewers for their insightful comments. This study was supported by Information-technology Promotion Agency (IPA), Japan, The MITOU Program in fiscal year 2019, and SAKURA internet Inc. We are grateful to Kazuyuki Shudo for his support as our project manager in The MITOU Program. This work was partially supported by JSPS KAKENHI Grant No. 20K19799, and 20H04156.

References

- [1] S. van Schaik, M. Minkin, A. Kwong, D. Genkin, and Y. Yarom, “CacheOut: Leaking data on Intel CPUs via cache evictions.” <https://cacheoutattack.com/>, 2020.
- [2] S. van Schaik, A. Kwong, D. Genkin, and Y. Yarom, “SGAxe: How SGX fails in practice.” <https://sgaxeattack.com/>, 2020.
- [3] S. Rass and P. Schartner, “On the security of a universal cryptocomputer: the chosen instruction attack,” *IEEE Access*, vol. 4, pp. 7874–7882, 2016.
- [4] A. C. Yao, “How to generate and exchange secrets,” in *27th Annual Symposium on Foundations of Computer Science (sfcs 1986)*, pp. 162–167, 1986.
- [5] C. Gentry, *A Fully Homomorphic Encryption Scheme*. PhD thesis, Stanford, CA, USA, 2009.
- [6] E. M. Songhori, S. U. Hussain, A. Sadeghi, T. Schneider, and F. Koushanfar, “Tinygarble: Highly compressed and scalable sequential garbled circuits,” in *2015 IEEE Symposium on Security and Privacy*, pp. 411–428, 2015.
- [7] E. M. Songhori, T. Schneider, S. Zeitouni, A. Sadeghi, G. Dessouky, and F. Koushanfar, “Garbledcpu: A mips processor for secure computation in hardware,” in *2016 53rd ACM/EDAC/IEEE Design Automation Conference (DAC)*, pp. 1–6, June 2016.
- [8] A. Chatterjee and K. M. M. Aung, *FURISC: FHE Encrypted URISC Design*, pp. 87–115. Singapore: Springer Singapore, 2019.
- [9] N. P. Smart and F. Vercauteren, “Fully homomorphic encryption with relatively small key and ciphertext sizes,” in *Proceedings of the 13th International Conference on Practice and Theory in Public Key Cryptography, PKC’10*, (Berlin, Heidelberg), p. 420–443, Springer-Verlag, 2010.
- [10] R. Rivest, L. Adleman, and M. Dertouzos, “On data banks and privacy homomorphisms,” in *Foundations on Secure Computation*, Academia Press, pp. 169–179, 1978.
- [11] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, “Tfhe: Fast fully homomorphic encryption over the torus,” *Journal of Cryptology*, vol. 33, no. 1, pp. 34–91, 2020.
- [12] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, “(leveled) fully homomorphic encryption without bootstrapping,” in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS ’12*, (New York, NY, USA), p. 309–325, Association for Computing Machinery, 2012.
- [13] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, “Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds,” in *Advances in Cryptology – ASIACRYPT 2016* (J. H. Cheon and T. Takagi, eds.), (Berlin, Heidelberg), pp. 3–33, Springer Berlin Heidelberg, 2016.
- [14] S. Halevi and V. Shoup, “Bootstrapping for helib.” Cryptology ePrint Archive, Report 2014/873, 2014. <https://eprint.iacr.org/2014/873>.
- [15] O. Regev, “On lattices, learning with errors, random linear codes, and cryptography,” in *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, STOC ’05*, (New York, NY, USA), p. 84–93, Association for Computing Machinery, 2005.
- [16] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, “TFHE: Fast fully homomorphic encryption library.” <https://tfhe.github.io/tfhe/>, August 2016.
- [17] P. Mohassel and S. Sadeghian, “How to hide circuits in mpc an efficient framework for private function evaluation,” in *Advances in Cryptology – EUROCRYPT 2013* (T. Johansson and P. Q. Nguyen, eds.), (Berlin, Heidelberg), pp. 557–574, Springer Berlin Heidelberg, 2013.
- [18] D. Cash, M. Green, and S. Hohenberger, “New definitions and separations for circular security,” in *Public Key Cryptography – PKC 2012* (M. Fischlin, J. Buchmann, and M. Manulis, eds.), (Berlin, Heidelberg), pp. 540–557, Springer Berlin Heidelberg, 2012.
- [19] J. Black, P. Rogaway, and T. Shrimpton, “Encryption-scheme security in the presence of key-dependent messages,” in *Revised Papers from the 9th Annual International Workshop on Selected Areas in Cryptography, SAC ’02*, (Berlin, Heidelberg), p. 62–75, Springer-Verlag, 2002.

- [20] A. Paverd, A. Martin, and I. Brown, “Modelling and automatically analysing privacy properties for honest-but-curious adversaries,” tech. rep., University of Oxford, 2014.
- [21] M. Brenner, H. Perl, and M. Smith, “How practical is homomorphically encrypted program execution? an implementation and performance evaluation,” in *2012 IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications*, pp. 375–382, June 2012.
- [22] N. G. Tsoutsos and M. Maniatakos, “The heroic framework: Encrypted computation without shared keys,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 6, pp. 875–888, 2015.
- [23] P. T. Breuer and J. P. Bowen, “Fully encrypted high-speed microprocessor architecture: the secret computer in simulation,” *IJCCBS*, vol. 9, no. 1/2, pp. 26–55, 2019.
- [24] O. Mazonka, N. G. Tsoutsos, and M. Maniatakos, “Cryptoleq: A heterogeneous abstract machine for encrypted and unencrypted computation,” *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 9, pp. 2123–2138, 2016.
- [25] A. Chatterjee and I. Sengupta, “Furisc: Fhe encrypted urisc design.” Cryptology ePrint Archive, Report 2015/699, 2015. <https://eprint.iacr.org/2015/699>.
- [26] “libScarab.” <https://github.com/hcrypt-project/libScarab>, 2013. Accessed 06/19/2020.
- [27] E. M. Songhori, M. S. Riazi, S. U. Hussain, A.-R. Sadeghi, and F. Koushanfar, “Arm2gc: Succinct garbled processor for secure computation,” *2019 56th ACM/IEEE Design Automation Conference (DAC)*, pp. 1–6, 2019.
- [28] S. Yasuda, F. Kitagawa, and K. Tanaka, *Constructions for the IND-CCA1 Secure Fully Homomorphic Encryption*, pp. 331–347. Singapore: Springer Singapore, 2018.
- [29] J. Loftus, A. May, N. P. Smart, and F. Vercauteren, “On cca-secure somewhat homomorphic encryption,” in *Selected Areas in Cryptography* (A. Miri and S. Vaudenay, eds.), (Berlin, Heidelberg), pp. 55–72, Springer Berlin Heidelberg, 2012.
- [30] E. Ben-Sasson, A. Chiesa, D. Genkin, E. Tromer, and M. Virza, “Snarks for c: Verifying program executions succinctly and in zero knowledge,” in *Advances in Cryptology – CRYPTO 2013* (R. Canetti and J. A. Garay, eds.), (Berlin, Heidelberg), pp. 90–108, Springer Berlin Heidelberg, 2013.
- [31] K. Matsuoka, R. Banno, and N. Matsumoto, “Source codes of our implementation,” 2020. <https://github.com/virtualsecureplatform/kvsp>.
- [32] C. Lattner and V. Adve, “LLVM: A compilation framework for lifelong program analysis and transformation,” in *CGO*, (San Jose, CA, USA), pp. 75–88, Mar 2004.
- [33] “sbt.” <https://www.scala-sbt.org/index.html>. Accessed 06/19/2020.
- [34] C. Wolf, “Yosys open synthesis suite.” <http://www.clifford.at/yosys/>. Accessed 06/19/2020.
- [35] “Original implementation of cuFHE.” <https://github.com/vernamlab/cuFHE>, 2018. Accessed 06/19/2020.
- [36] W. Dai and B. Sunar, “cuhe: A homomorphic encryption accelerator library,” in *Cryptography and Information Security in the Balkans* (E. Pasalic and L. R. Knudsen, eds.), (Cham), pp. 169–186, Springer International Publishing, 2016.
- [37] “Users page in official LLVM website.” <http://llvm.org/Users.html>. Accessed 06/19/2020.
- [38] “ACM Software System Award 2012 (LLVM).” https://awards.acm.org/award_winners/lattner_5074762. Accessed 06/19/2020.
- [39] N. Matsumoto, R. Banno, and K. Matsuoka, “Specification for CAHPv3,” 2020. <https://github.com/virtualsecureplatform/cahvp3-spec>.
- [40] J. Bachrach, H. Vo, B. Richards, Y. Lee, A. Waterman, R. Avižienis, J. Wawrzynek, and K. Asanović, “Chisel: Constructing hardware in a scala embedded language,” in *DAC Design Automation Conference 2012*, pp. 1212–1221, 2012.
- [41] “Community page in official chisel website.” <https://www.chisel-lang.org/community.html>.
- [42] D. A. Patterson and J. L. Hennessy, *Computer Organization and Design, Fourth Edition, Fourth Edition: The Hardware/Software Interface (The Morgan Kaufmann Series in Computer Architecture and Design)*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 4th ed., 2008.
- [43] “Rocket core (risc-v).” <https://github.com/chipsalliance/rocket-chip>. Accessed 06/19/2020.
- [44] P. M. Sailer, P. M. Sailer, and D. R. Kaeli, *The DLX Instruction Set Architecture Handbook*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1st ed., 1996.

- [45] X. Inc., “Microblaze processor reference guide embedded development kitedk 11.4.” https://www.xilinx.com/support/documentation/sw_manuals/xilinx11/mb_ref_guide.pdf. Accessed 06/19/2020.
- [46] E. M. Songhori, S. U. Hussain, A. Sadeghi, T. Schneider, and F. Koushanfar. https://github.com/esonghori/TinyGarble/tree/d40454dd365a943c364c3e7de05039fe94728c7a/circuit_synthesis/mips. Accessed 06/19/2020.
- [47] C. Wolf, “PicoRV32.” <https://github.com/cliffordwolf/picorv32/tree/f9b1beb4cfd6b382157b54bc8f38c61d5ae7d785>. Accessed 06/19/2020.
- [48] R. Banno, K. Matsuoka, and N. Matsumoto, “Benchmark scripts for VSP,” 2020. <https://github.com/virtualsecureplatform/kvsp-benchmark>.
- [49] P. Paillier, “Public-key cryptosystems based on composite degree residuosity classes,” in *Advances in Cryptology — EUROCRYPT ’99* (J. Stern, ed.), (Berlin, Heidelberg), pp. 223–238, Springer Berlin Heidelberg, 1999.
- [50] N. G. Tsoutsos and M. Maniatakos, “Heroic: Homomorphically encrypted one instruction computer,” in *2014 Design, Automation Test in Europe Conference Exhibition (DATE)*, pp. 1–6, 2014.
- [51] O. Mazonka, “Higher subleq.” <http://mazonka.com/subleq/hsq.html>, March 2011. Accessed 06/19/2020.
- [52] P. T. Breuer, J. P. Bowen, E. Palomar, and Z. Liu, “A practical encrypted microprocessor,” in *Proceedings of the 13th International Joint Conference on E-Business and Telecommunications, ICETE 2016*, (Setubal, PRT), p. 239–250, SCITEPRESS - Science and Technology Publications, Lda, 2016.

A Related Work

Although FURISC is the most similar existing study to our proposed method, we will discuss other studies here. There are some previous works which uses Paillier Cryptosystem [49] to evaluate encrypted binaries. HEROIC [50] is the first one. Paillier cryptosystem is one kind of Partial Homomorphic Encryption (PHE). Paillier Cryptosystem only permits addition of integers. Therefore, HEROIC uses some tables to provide enough functionality to implement a processor as Arithmetic circuit. HEROIC implements an OISC processor which only supports SUBtract and branch if Less-than or Equal to zero (SUBLEQ) instruction. Unlike to FURSIC, there is a C like language compiler for SUBLEQ, HIGHER SUBLEQ, though its last update is in March 2011 [51]. The use of the tables




Figure 14: Protocol flow of PF-SFE

leads the ciphertexts becomes deterministic. This means the public key cannot be public. Therefore HEROIC theoretically cannot achieve two-party PF-SFE. In addition, the method of using the table has not been proven to be secure. The authors of HEROIC also proposed Cryptoleq in 2016 [24]. It also uses Paillier Cryptosystem with some tables and OISC. They also proposed assembly-like Domain Specific Language (DSL) in it. Cryptoleq depends on the random number generation of the server. This is not suitable characteristic for SMPC since it needs the verification of the random number generator. Cryptoleq also depends on heuristic code-based obfuscation. There is a Open RISC implementation based on idea of HEROIC [52], but this is suffered from too much memory consumption because of big tables. The authors estimated it to be between hundreds of gigabytes to terabytes.

B Abstract Protocol Flow in two-party PF-SFE

In this section, we explain how the protocol flow of VSP can be theoretically modified to do two-party PF-SFE.

Public/Private Data: In this protocol, the function to be evaluated is provided by Bob and the input data is provided by Alice. The most important fact for understanding how VSP works in this protocol is that TFHE supports “trivial” ciphertexts. “Trivial” here means the generation of them does not need any secret key nor random number generation. For example, trivial TLWE of 1 is $(0, \mu)$. In such a way, Bob can provide ROM and RAM data without the input of Alice.

ROM data and RAM data except for the input of Alice are private of Bob and the input is private of Alice.

B.1 Abstract protocol workflow

The visual image is shown in Figure 14. This shows only difference from two-party SCO case.

- 3. Compilation:** Bob compiles the source code of the desired function into the executable for the processor.

Table 7: Runtime performance evaluation in 1 KiB ROM and RAM setting

Machine	# of V100	Pipelining?	CMUX Memory?	# of cycles	Runtime [s]	sec./cycle	# of tries
AWS p3.8xlarge	4	No	No	937	3306.0 ± 12.9	3.528 ± 0.014	2
AWS p3.8xlarge	4	No	Yes	937	1733.4 ± 1.1	1.850 ± 0.001	3
AWS p3.8xlarge	4	Yes	No	1217	3804.0 ± 18.6	3.126 ± 0.015	3
AWS p3.8xlarge	4	Yes	Yes	1217	1217.0 ± 2.4	1.000 ± 0.002	5
AWS c5.metal	0	No	Yes	937	3230.0 ± 40.6	3.443 ± 0.043	3
AWS c5.metal	0	Yes	Yes	1217	3940.0 ± 68.4	3.238 ± 0.056	3

Table 8: Resource requirements of CAHP-Ruby.

Gate	IF	ID+WB	Ex	Mem	Total
AND	193	270	301	92	651
ANDNOT	59	110	44	0	223
MUX	54	683	256	116	996
NAND	300	336	356	10	1025
NOR	4	77	31	0	90
NOT	3	4	11	1	15
OR	77	56	95	8	215
ORNOT	39	142	40	8	195
XNOR	3	9	40	0	51
XOR	4	5	21	0	36

Table 9: The number of CMUXs in CMUX Memory

Component	# of CMUXs
RAM Read Unit	4,080
RAM Write Unit	32,768
ROM	8
Total	36,856

- Encryption:** Alice encrypts the input and sends it to Bob. Bob encrypts the executable using trivial ciphertexts and combines it with the encrypted input to generate the encrypted ROM and RAM.

B.2 Security Analysis in Two-party PF-SFE

Bob tries to reveal plaintexts of Bootstrapping Key, RAM, registers and each ciphertext of each wire, etc. ROM (the function to be evaluated) is not a target since it is provided by Bob. However, like two-party SCO case, this can be reduced to the hardness of decryption of ciphertexts of TFHE.

Alice tries to reveal ROM, RAM, registers and ciphertexts of each wire, etc. Though they will not be provided to Alice, Alice knows the result of the function, so if Bob uses always the same function and input, Alice can try to get the results for all possible inputs. Therefore, the protection of private information of Bob from Alice needs another method like indistinguishable obfuscation. This is beyond the scope of our proposed method.

We note that PF-SFE protocol does have the FHE malleability problem, since the program is provided by Bob. However, PF-SFE is still vulnerable to the termination problem mentioned in Section 4.1.

C Additional Evaluations

Runtime Performance Evaluation

Table 7 shows additional results in 1 KiB ROM and RAM setting.

Gate count evaluation

Table 8 shows the gate requirements of each stage of CAHP-Ruby. Note that these values are calculated by synthesising the components separately. Due to global optimizations in the synthesis software, the numbers do not add up to the size of the entire processor circuit (the column “Total”).

Also, Table 9 shows the number of CMUXs in CMUX Memory components.

Memory consumption evaluation

On p3.8xlarge, running our implementation consumes about 3.7 GB of main memory and about 0.6 GB per GPU. The most of memory consumption seems to be caused by holding Bootstrapping Key.