HElium: A Language and Compiler for Fully Homomorphic Encryption with Support for Proxy Re-Encryption

Mirko Günther
SAP SE
Dresden, Germany
mirko.guenther@sap.com

Lars Schütze
Technische Universität Dresden
Dresden, Germany
lars.schuetze@tu-dresden.de

Kilian Becher
Technische Universität Dresden
Dresden, Germany
kilian.becher@tu-dresden.de

Thorsten Strufe
Karlsruhe Institute of Technology
Karlsruhe, Germany
thorsten.strufe@kit.edu

Jeronimo Castrillon
Technische Universität Dresden
Dresden, Germany
jeronimo.castrillon@tu-dresden.de

ABSTRACT

Privacy-preserving analysis of confidential data can increase the value of such data and even improve peoples' lives. Fully homomorphic encryption (FHE) can enable privacy-preserving analysis. However, FHE adds a large amount of computational overhead and its efficient use requires a high level of expertise. Compilers can automate certain aspects such as parameterization and circuit optimizations. This in turn makes FHE accessible to non-cryptographers. Yet, multi-party scenarios remain complicated and exclude many promising use cases such as analyses of large amounts of health records for medical research. Proxy re-encryption (PRE), a technique that allows the conversion of data from multiple sources to a joint encryption key, can enable FHE for multi-party scenarios. Today, there are no optimizing compilers for FHE with PRE capabilities.

We propose HElium, the first optimizing FHE compiler with native support for proxy re-encryption. HElium features HEDSL, a domain-specific language (DSL) specifically designed for multi-party scenarios. By tracking encryption keys and transforming the computation circuit during compilation, HElium minimizes the number of expensive PRE operations. We evaluate the effectiveness of HElium's optimizations based on the real-world use case of the tumor recurrence rate, a well-known subject of medical research. Our empirical evaluation shows that HElium substantially reduces the overhead introduced through complex PRE operations, an effect that increases for larger amounts of input data.

KEYWORDS

fully homomorphic encryption, proxy re-encryption, healthcare, optimizing compiler, domain-specific language

1 INTRODUCTION

Today, healthcare institutions collect large amounts of data about their patients such as medical-treatment protocols and surgical reports. Analyses of combined data sets could increase the value of that data and provide more significant insights. For example, with larger data sets, more accurate information about correlations between genome mutations and the properties of tumors can be identified. This can enable researchers to develop more specific treatment methods against these tumors [38, 46]. However, patient data is highly confidential, and strict data-protection requirements
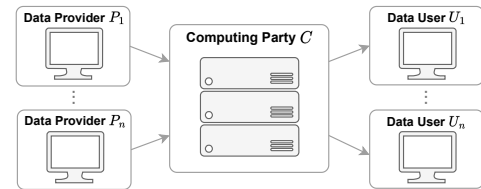

Figure 1: High-Level Multi-Party Scenario with n Data Providers, One Central Computing Party, and n Data Users

apply [23]. Furthermore, the pool of relevant medical data records could change over time. Consequently, sharing and processing patients' data with research institutions for scientific studies is anything but trivial. Our work aims to solve this and similar problems by simplifying the application of privacy-preserving computation for domain experts with limited to no experience in privacy-enhancing technologies. We propose HElium, an optimizing compiler for fully homomorphic encryption. It is particularly designed for multi-party scenarios as depicted in Figure 1. HElium paves the way for leveraging medical data for more effective medical research without sacrificing the privacy of individuals. Fully homomorphic encryption (FHE) schemes allow computation on encrypted data without intermediate decryption [27], thus, enabling privacy-preserving computation. However, FHE is limited in terms of computations with multiple parties. Since standard FHE requires ciphertexts to be encrypted under the same key, all parties have to agree on a common key pair before encryption. However, the generation of a joint key adds organizational overhead and requires some level of synchronicity. This might rule out many use cases in the first place, especially those that involve numerous data owners that provide confidential inputs, potentially over a longer period of time.

Approaches like multi-key FHE [36] and threshold FHE [3] address this problem. Multi-key FHE enables homomorphic computation with ciphertexts encrypted under different keys and outputs a result that is encrypted under a combination of the data providers' keys. In threshold FHE, data providers jointly run a key generation procedure that yields a common encryption key and one share of the secret decryption key per participant. However, both multi-key and threshold FHE require data providers' active participation in the decryption procedure, which limits flexibility and convenience.

Both hinder adoption in scenarios where data providers are natural persons like patients, especially if data is added to the pool of medical data records over a longer period.

In contrast, proxy re-encryption (PRE) [4, 32] enables the transformation of a ciphertext encrypted under one key into a ciphertext of the same plaintexts encrypted under a different key without intermediate decryption. Hence, it can be used to convert ciphertexts encrypted under different keys to a common key.

FHE schemes with PRE capabilities enable computations on distributed data sets that are encrypted under different encryption keys [40]. This renders them well-suited for multi-party scenarios. Therefore, among many other problems, they allow to combine patients' health records in a privacy-preserving way. This, in turn, allows for analyzing combined data sets to gain more significant insights, eventually leading to more effective medical research. However, FHE and PRE operations are computationally expensive. Therefore, they bear great potential for optimization, e.g., through parallelizing operations or rearranging their order. This is particularly the case when individuals contribute multiple data points, for instance across different phases of treatment.

Developing efficient FHE applications with PRE is challenging and requires cryptographic expertise. Compilers can help developers by automating tasks like parameterization and the efficient use of re-encryption operations. Even though a variety of compilers for FHE have been proposed in the past [5, 8, 12, 17, 19, 20, 22, 44, 45], none of them natively offers optimizations for the computational overhead caused by re-encryption operations. This renders them unsuitable for scenarios with large or frequently changing participant sets.

We propose HElium, a compiler for fully homomorphic encryption with native support for proxy re-encryption. HElium performs a plethora of optimizations such as the balancing of the computation graph as well as the minimization of the number of PRE operations. Furthermore, we propose HEDSL, HEliums's domain-specific language (DSL) that acts as an abstraction layer for complex FHE and PRE operations and captures the different keys in multi-party scenarios. Its syntax significantly improves the ease of use of FHE and PRE, especially for non-cryptographers.

We evaluate HElium using the real-world example of tumor recurrence rate, an active area of medical research. We show that HElium minimizes PRE operations compared to a baseline approach for scenarios where the number of inputs n is larger than the number of input keys k, i.e., $n/k > 1$. Most notably, for $n/k = 8$, the number of PRE operations is reduced by 87.5 % compared to the baseline approach, which results in a speedup of 1.34×. Consequently, our optimizations achieve high efficiency. They ensure that the execution time overhead introduced by PRE operations quickly becomes negligible, even for relatively small n/k.

The structure of this paper is as follows. We first provide preliminaries such as homomorphic encryption and proxy re-encryption in Section 2 followed by a detailed description of our scenario and the resulting adversary model in Section 3. We give an overview of the HElium compiler architecture in Section 4 and present the empirical evaluation of HElium in Section 5. Relevant related work is discussed in Section 6. We conclude the paper with a summary and approaches to future work in Section 7.

2 PRELIMINARIES

This section introduces homomorphic encryption and proxy re-encryption. While the former enables to outsource of computations to an untrusted party, i.e., a cloud service, the latter enables multi-party computations with data encrypted under different keys.

2.1 Homomorphic Encryption

Asymmetric cryptosystems are tuples $CS = (G, E, D)$ that consist of a probabilistic key-generation algorithm $G(\cdot)$, which generates pairs of a (public) encryption key pk and a (secret) decryption key sk, as well as a (probabilistic) encryption procedure $E(\cdot)$ and a decryption procedure $D(\cdot)$. We further denote plaintext and ciphertext spaces by $\mathcal{M}$ and C, respectively.

Homomorphic encryption (HE) schemes additionally provide at least one operation "$\circ$" on C that corresponds to an operation "$\bullet$" on $\mathcal{M}$ such that $E_{pk}(m_1) \circ E_{pk}(m_2)$ results in encryption of $m_1 \bullet m_2$ for two plaintexts m_1, m_2. We formalize this as

$$D_{sk}(E_{pk}(m_1) \circ E_{pk}(m_2)) = m_1 \bullet m_2, \tag{1}$$

where $\bullet$ denotes Boolean or arithmetic operations like XOR, AND, addition, and multiplication as shown below for homomorphic addition and multiplication of two ciphertexts.

$$D_{sk}(E_{pk}(m_1) \oplus E_{pk}(m_2)) = m_1 + m_2 \tag{2}$$

$$D_{sk}(E_{pk}(m_1) \odot E_{pk}(m_2)) = m_1 \cdot m_2 \tag{3}$$

In addition to homomorphic operations with two encrypted operands, common HE schemes [11, 26, 39] also enable operations where only one operand is provided as ciphertext whereas the second operand is a constant. These are often referred to as addition/multiplication-by-constant or ciphertext-plaintext operations. They are typically more efficient than ciphertext-ciphertext operations. We formalize them as follows.

$$D_{sk}(E_{pk}(m_1) \circ m_2) = m_1 \bullet m_2 \tag{4}$$

Homomorphic operations can be used to construct more complex functions. Functions consisting of Boolean operations are typically interpreted as *circuits* consisting of gates and wires, while arithmetic operations are often interpreted as *Abstract Semantic Graphs*, also known as *Term Graphs*.

There are four different types of HE schemes [2]. Partially homomorphic encryption (PHE) schemes offer one homomorphic operation, e.g., addition [39] or multiplication [41] of underlying plaintexts. Somewhat homomorphic encryption (SHE) schemes enable the evaluation of two different operations, e.g., XOR and AND or addition and multiplication, but only for limited-depth circuits. Leveled fully homomorphic encryption (LFHE) schemes offer two homomorphic operations for circuits with pre-determined bounded depths. Fully homomorphic encryption (FHE) schemes enable the evaluation of arbitrary circuits of unlimited depth.

In 2009, Gentry presented the first FHE scheme [27], which acts as a blueprint for future FHE schemes. Public and secret key are elements that cancels out to zero during decryption [44]. To ensure security, a small error, referred to as *noise*, is incorporated into the public key and ciphertexts. Homomorphic operations cause the noise to grow. As long as the noise is small enough, the plaintext can be reconstructed correctly. If, however, the noise grows

too large, it affects the underlying plaintext such that correct decryption can no longer be guaranteed. Therefore, this construction alone is bounded to a limited number of consecutive operations. To overcome this problem, Gentry proposed a function referred to as *bootstrapping*. This computationally complex operation reduces the noise of a ciphertext. In the bootstrapping operation, the decryption function of the encryption scheme is translated into a circuit $\mathcal{D}$. This bootstrapping circuit is homomorphically evaluated given the public key pk, the double-encrypted secret m, and the encrypted secret key sk as follows.

$$E_{pk}(m) \leftarrow Eval(pk, \mathcal{D}, E_{pk}(E_{pk}(m)), E_{pk}(sk)) \qquad (5)$$

This yields a *fresh* encryption of the plaintext, where fresh means that the resulting ciphertext contains less noise than the original ciphertext, allowing for further homomorphic operations. Consequently, such a scheme can evaluate an arbitrary number of computation gates.

Due to the fact that homomorphic operations cause large computational overhead, FHE has been an active area of research since Gentry's seminal work. More recent FHE schemes offer higher efficiency. In [7], Brakerski, Gentry, and Vaikuntanathan presented a scheme that is based on the ring learning with errors (RLWE) problem [37]. This and following schemes like Brakerski/Fan-Vercauteren (BFV) [6, 26] and Cheon-Kim-Kim-Song (CKKS) [11] reduced the growth of noise during homomorphic operations. This allows evaluation of larger circuits without bootstrapping by using the scheme in its leveled form. *Batching*, a method for packing multiple plaintexts into a single ciphertext [42], allows computations in a SIMD fashion, where an operation is applied to multiple plaintext elements simultaneously. Therefore, this technique can vastly improve the efficiency of FHE-based programs.

2.2 Proxy Re-Encryption

Re-encryption transforms a ciphertext $c_1 = E_{pk_1}(m)$ encrypted under a key pk_1 into a ciphertext $c_2 = E_{pk_2}(m)$ of the same plaintext but encrypted under a different key pk_2. Proxy re-encryption (PRE) allows an untrusted party to perform this transformation without affecting confidentiality. PRE was first proposed by Blaze, Bleumer, and Strauss [4] for bidirectional settings and later by Ivan and Dodis [32] for unidirectional settings.

Following the notation of [40], we denote a PRE scheme as a tuple $\mathcal{PRE} = (PG, KG, ReKG, E, D, RE)$ of six procedures. Parameter generation $PG(\cdot)$ computes a set of public parameters related to the security parameter λ. The key generation algorithm $KG(\cdot)$ outputs a key pair (pk, sk). Re-encryption-key generation $ReKG(\cdot)$ takes a secret key sk_i and a public key $pk_{j \neq i}$ and computes a re-encryption key $rk_{i \rightarrow j}$. The re-encryption algorithm $RE(\cdot)$ transforms a ciphertext c_i of m encrypted under pk_i into a ciphertext c_j of m such that c_j encrypts m under $pk_{j \neq i}$. $E(\cdot)$ and $D(\cdot)$ denote encryption and decryption algorithms, respectively.

A generic construction to obtain a PRE scheme from an FHE scheme is described in Gentry's seminal work [27]. Gentry's bootstrapping construction (see Equation (5)) implies a one-way proxy re-encryption scheme where re-encryption is conducted as follows.

$$E_{pk_2}(m) \leftarrow Eval(pk_2, \mathcal{D}, E_{pk_2}(E_{pk_1}(m)), E_{pk_2}(sk_1)) \qquad (6)$$
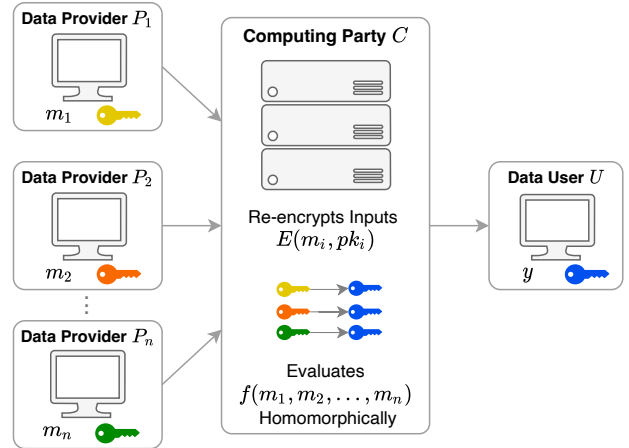

Figure 2: Abstract FHE Scenario with PRE

Here, the re-encryption key is $rk_{1 \rightarrow 2} = E_{pk_2}(sk_1)$, i.e., the first secret key encrypted under the second public key.

3 MULTI-PARTY COMPUTING SCENARIOS

This section gives an overview of the class of multi-party scenarios that our compiler enables through proxy re-encryption.

3.1 Multi-Party Computation with PRE

Figure 2 depicts an abstract scenario of a proxy re-encryption program. It comprises three types of participants: the data providers P_i with $i \in \{1, 2, \ldots, n\}$, data user(s) U, and a computing party C in between. Without loss of generality, we assume a single data user U for the remainder of this paper. The data providers have confidential data denoted as messages m_i. They provide their encrypted messages to the computing party C, which computes the function $f(\ldots, m_i, \ldots)$ specified by the data user U. The P_i could provide multiple data points m_i encrypted under the same key, which could for instance be treatment details of different phases of a medical treatment.

All data providers and the data user have their key pairs to encrypt and decrypt messages. The computing party in between enables the joint computation of the function and acts as a re-encryption proxy. It re-encrypts the ciphertexts $E_{pk_i}(m_i)$ given the corresponding, pre-computed re-encryption keys to bring them under U's public key, and homomorphically computes the function $f(\ldots, m_i, \ldots)$. The result $E_{pk_U}(y)$ is forwarded to the data user.

This construction enables use cases with frequently changing data providers since new data providers only have to provide a corresponding re-encryption key. Furthermore, it enables asynchronicity between encryption and usage of the data. Since data is encrypted under participant-specific keys, no party can read another party's confidential data. Neither the data user nor the target function $f(\cdot)$ with the required m_i's need to be determined at the time of encryption. The data user's identity and key pair need to be known only at the time of re-encryption-key generation. Therefore, PRE adds a high level of flexibility concerning the participant setup.
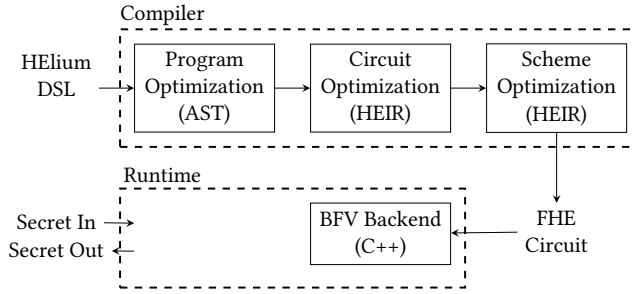
Figure 3: Compiler Pipeline of HElium

3.2 Threat Model

Similar to the related work [20, 21, 45], we assume a semi-honest threat model [35] where the computing party C faithfully executes the specified homomorphic evaluation but is curious about the data providers' confidential inputs. We assume that C does not collude with the data user U. This model is plausible in scenarios where C is a server entrusted with homomorphic evaluation and an adversary A has read access to C's internal state without the ability to corrupt it.

Attestation of correct execution of the specified homomorphic evaluation is beyond the scope of this paper. However, following [24] we note that the evaluation could be computed in a trusted execution environment (TEE), like Intel SGX. This model would leverage the guarantees of correct execution that are inherent in TEEs while still relying on the strong confidentiality guarantees of homomorphic encryption.

4 HELIUM COMPILER

This section introduces the architecture of the HElium compiler. The HElium compiler infrastructure consists of three components: a domain-specific language (DSL) to represent multi-party FHE computations, an intermediate representation (IR), and the backend that outputs scheme-specific circuits. This section specifically highlights the use of Proxy Re-Encryption (PRE) operations (see Section 2.2) to facilitate multi-party scenarios.

4.1 Architecture Overview

The compiler is designed to use staged compilation where the IR is passed through IR-specific optimizations, which are grouped into phases subsequently lowering the program representation. Figure 3 depicts the phases of the HElium compiler and the respective IR they work on. The front end parses the DSL and constructs the corresponding abstract syntax tree (AST). The AST is de-sugared, optimized, and transformed into HElium IR (HEIR), our high-level IR representing HE circuits. The high-level IR is subsequently transformed and optimized with attention to FHE-specific needs. Thereafter, scheme-specific optimizations are carried out on the IR such as the automatic insertion of PRE operations. Finally, the backend emits a circuit that can be evaluated by an FHE runtime or generates a scheme-specific IR which can be used as an input to other FHE tools. The compiler is modular, provided as a C++ library, such that passes, phases, and backends can be added.

⟨Program⟩	:=	⟨Stmt⟩⁺
⟨Stmt⟩	:=	⟨InputStmt⟩ | ⟨OutputStmt⟩
	|	⟨VarDecl⟩ | ⟨VarAssignment⟩
	|	⟨FuncDecl⟩ | ⟨ForStmt⟩ | ⟨ReturnStmt⟩
⟨TypeExpr⟩	:=	plain? (int | int[⟨Expr⟩])
⟨InputStmt⟩	:=	input ID : ⟨TypeExpr⟩ (@ ID <= ID)? ;
⟨OutputStmt⟩	:=	output ID => ID (@ ID : ⟨Expr⟩)? ;
⟨VarDecl⟩	:=	var ID : ⟨TypeExpr⟩ (= ⟨Expr⟩)? ;
⟨VarAssignment⟩	:=	ID = ⟨Expr⟩ ;
⟨ReturnStmt⟩	:=	return ⟨Expr⟩ ;
⟨FuncDecl⟩	:=	fun ID (⟨ArgDecl⟩) { ⟨Stmt⟩⁺ }
⟨ArgDecl⟩	:=	ID | ID , ⟨ArgDecl⟩
⟨ForStmt⟩	:=	for (ID : ⟨Expr⟩) { ⟨Stmt⟩⁺ }
⟨Expr⟩	:=	⟨Expr⟩ (+ | - | ∗) ⟨Expr⟩
	|	ID | LITERAL | ID (⟨ArgDecl⟩)

Figure 4: The grammar of the proposed HElium DSL in EBNF

4.2 HElium DSL

This section introduces the syntax of HEDSL and highlights key features. The HElium DSL (HEDSL) is a high-level representation of multi-party FHE programs. The DSL hides scheme-specific details to improve usability for non-cryptographers and automates complex tasks such as scheme parameter selection. We kept the DSL simple enough so that it can be easily embedded into mainstream programming languages. The grammar of the DSL presented in EBNF can be seen in Figure 4. A HElium program consists of a list of statements, which can be, for instance, inputs, outputs, variable declarations, or function declarations. The individual components are discussed in the following.

4.2.1 Data Types. HEDSL, depending on the used backend, provides primitive types for *integers* and *fixed-point numbers*. It also supports *fixed-size* arrays containing elements of these types.

The type system further differentiates between variables representing *constants*, *ciphertexts*, or *plaintexts*. Each variable is declared as ciphertext by default and captures a label of the key the ciphertext was encrypted with. This information is later in optimization stages such as the automatic insertion of PRE operations into the circuit (see Section 4.5).

To facilitate security by default, the usage of plaintext must be annotated explicitly to the qualifier of a variable using the `plain` keyword. HElium uses the type information to automatically introduce more efficient ciphertext-plaintext FHE instructions when possible.

4.2.2 Inputs and Outputs. To support multi-party FHE, HElium programs may have multiple input and output statements each declaring the data-providing and data-consuming party, respectively. The information is used to retrieve and assign data of different parties according to the inputs of the program. Furthermore, data can be provided under different encryption keys. These are captured by key labels annotated with the input statements. The labels are used in later optimization stages such as automatic insertion of PRE operations into the circuit.

Figure 5 shows the declaration of inputs. Similarly, outputs require a name to allow identification in programs with multiple

Table 1: Built-in Operations of HElium

Symbol	Operation	Argument / Result Type
+	Addition	$(int, int) \rightarrow (int)$
−	Subtraction	$(int[\,], int \mid int[\,]) \rightarrow (int[\,])$
∗	Multiplication	$(int \mid int[\,], int[\,]) \rightarrow (int[\,])$
**	Exponentiation	$(int, int) \rightarrow (int)$
		$(int[\,], int) \rightarrow (int[\,])$
>>	Right rotation	
<<	Left rotation	$(int[\,], int) \rightarrow (int[\,])$

outputs. They are encrypted under the target key whose public key is provided to the application. The secrets are sent back to the respective party and can be decrypted locally.

4.2.3 Operations. The homomorphic operations provided by the DSL work with *scalar* and *vector* operands. These operations are expressed by operators as shown in Table 1. If one operand is a vector while the other is a scalar, the scalar will be converted into a vector first. Operations on vectors are performed element-wise. For rotations, the compiler provides tools to generate the respective rotation keys for participating parties.

4.2.4 Function Definition. To increase code reuse and encapsulation HEDSL supports the definition of functions. Functions are defined using the `fun` keyword followed by a function name and the function signature. Optionally, the signature may be annotated with types. Figure 5 depicts the definition of a function `mul` in HEDSL. Omitting a signature increases reuse when used in different contexts with unrelated types allowing to write functions that are polymorphic in the types of their arguments and the return type. In such a case, the compiler creates specializations for each case where a function application applies arguments of different types not seen before. For each specialization types of intermediate expressions are inferred according to the specification shown in Table 1.

HElium also provides intrinsics, i.e., built-in functions, to work with built-in data types. For example, the function $\text{size}(\cdot)$ returns the size of a vector. Since the size of vectors is fixed and known at compile time, HElium can directly inline the numerical value of the vector's size.
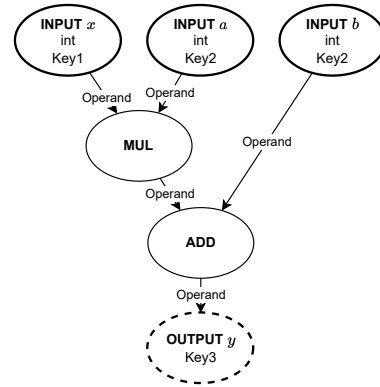
4.2.5 Loops. Loops help to improve readability by fostering code reuse. HEDSL provides `for` loops, which enable the iteration over the elements of a vector. As homomorphically encrypted programs are subject to several constraints, e.g., the absence of branching in arithmetic schemes, HElium focuses on loops with boundaries that can be determined statically. Due to the aforementioned constraints, HElium unrolls loops during compilation.

4.3 Intermediate Representation

A program in HElium represents operations on constants, ciphertext, or plaintext. The HElium IR (HEIR) in-memory representation uses *Term Graphs*, that is to interpret a term t as a graph $g(t)$, to represent HElium programs. In the remainder of this paper, we will use the term *computation graph*.

```
1   input a1: int @Key1 <= Party1;
2   input b1: int @Key1 <= Party1;
3   input a2: int @Key2 <= Party2;
4   input b2: int @Key2 <= Party2;
5   input a3: int @Key3 <= Party3;
6   input b3: int @Key3 <= Party3;
7
8   fun mul(a, b) { return a * b; }
9
10  output sum: mul(a1, b1) +
11              mul(a2, b2) +
12              mul(a3, b3);
13  output quantity: b1 + b2 + b3;
```

Figure 5: Example HElium DSL Program

Figure 6: HElium IR Graph Representation of $y = a \cdot x + b$

Each term has a particular number of operands depending on the type of operation. Unary terms take one operand as input while binary terms require two operands. The terms `input` and `output` are special unary terms that not only define the dataflow in the IR but also capture the types of incoming values. Constant values are represented by `const` terms without an operand. Figure 6 shows the computation graph of an example program that calculates the function $y = a \cdot x + b$. It has three inputs x, a, b and one output y.

The nodes represent operations either on ciphertexts or on plaintexts. There are special nodes that represent constant values, inputs, and outputs. For example, the *INPUT* nodes of the IR represent the inputs of the HElium program.

While the aforementioned properties are standard properties of contemporary IRs, HEDSL also captures HE-specific properties in its type system. It is required to annotate types of (encrypted) inputs that are materialized in the computation graph. The HElium compiler can infer types and key labels for the subsequent program. The IR allows the HElium compiler to reason about the program and automatically injects proxy re-encryption (PRE) operations to minimize PRE operations.

4.4 Transformations

The parameter selection of FHE programs depends on the noise growth of ciphertexts. As a ciphertext's noise grows with every operation performed, it is important to perform as few operations on a ciphertext as possible. The way HE-specific components are captured in HEIR simplifies many optimizations known from existing compiler frameworks. IRs used in contemporary compiler frameworks such as static single-assignment (SSA) forms require data-flow analyses to compute and extract information from the program not explicitly available. In this section, we present approaches to optimizing the computation graph to improve the execution runtime.

4.4.1 Constant Propagation and Constant Folding. A HElium program contains constant values, either added by the developer or introduced by language constructs, e.g., loop bounds and index accesses. While arithmetic operations with constant plaintext data are computationally less complex than ciphertext operations, precalculating operations with constants lowers the overall size of the computation graph. Therefore, HElium implements constant folding, i.e., the evaluation of constant expressions at compile time. Traditionally, this requires a reaching definition analysis to recompute *use-def* and *def-use* chains from the IR. Edges between nodes in the graph directly represent use-def and use-def chains. More precisely, our compiler evaluates if operands have use-def edges to constant values and replaces them with the constant value. If all operands of an operation are constant the operation is statically evaluated and replaced with the result. Due to the lack of branch instructions, this optimization only requires a single pass over the complete computation graph.

4.4.2 Unreachable and Dead-Code Elimination. While traditionally this requires a liveness analysis to find unused assignments, our IR uses edges between nodes to directly represent this information. A definition without a use edge is called dead and will be removed from the graph. This also applies to subgraphs that are not used. Therefore, HElium has a compiler pass that removes such subgraphs iteratively. To do so, it iterates over the drains of the computation graph and removes them, except for outputs and inputs. Then, HElium removes the operand-usage connections to their operands. In case their operands have no further usage, they are added to the list of drains of the computation graph. To remove all unused nodes, this procedure is repeated until the list of drains contains only output and input nodes. In case this list contains an input node, the compiler can notify the developer that the input is not necessary for program execution.

4.4.3 Rebalancing. Rebalancing can reduce the depth of the computation graph to $O(log_2 n)$ given n nodes in the computation graph. In the case of multiplications, it can reduce the multiplicative depth which may have a positive impact on the parameterization and the resulting runtime performance of an HE program. For multithreaded evaluation, this also improves performance as sub-graphs can be evaluated in parallel.

HElium uses an approach to re-balancing similar to [16]. The idea is to look for nodes that represent commutative operations like addition and multiplication. This is implemented as a compiler pass iterating over the computation graph. If an operation has a single usage that is of the same type of operation, both nodes are merged into a *super node* of that operation. After applying this procedure to the computation graph, all commutative chains of the same operation are transformed into super nodes with multiple operands.

The compiler then lowers the super nodes to return to a consistent IR state. Supernodes that contain more than two operands are transformed back into a balanced computation tree. This is applied recursively until all super nodes have been processed.

4.5 Placing Proxy Re-Encryption Operations

In HElium, inputs can be annotated with *key labels* that declare the key input is encrypted with. Homomorphic operations require that ciphertext operands are encrypted under the same key. Operands with different keys must be proxy re-encrypted under the same key before they can be used for homomorphic evaluation. The result of each homomorphic operation will be encrypted under the key of its operands. Finally, each output must be re-encrypted under the key of the receiving party.

In a naïve approach, each input may be immediately proxy re-encrypted. While this approach fulfills the requirement of operands having the same key, it is not optimal w.r.t. the amount of PRE operations inserted into the computation graph. However, proxy re-encryption can be postponed until operands of an operation are encrypted under different keys. This allows to reduce the amount of PRE operations, thus reducing the overall execution time.

Algorithm 1 shows how we assign labels to nodes in the computation graph and automatically insert the least amount of PRE operations necessary. $SUCC(v)$ denotes the nodes that outgoing edges of v point to, while $PRED(v)$ denote the nodes that v has incoming edges from. The algorithm takes as input all nodes that represent the input to the computation graph. These nodes are already labeled with the key of the ciphertexts. A forward traversal computes whether the keys of operands differ. If they do not, the operand's key is propagated to that node. Otherwise, the key of the receiving party is assigned and a PRE operation for each operand is inserted into the graph.

4.6 Backend

The backend of HElium translates the IR to a low-level HE instruction set which is executable by the targeted FHE scheme implementation. During that lowering, HElium directly maps to or encodes operations of the IR in the instructions of the given scheme. The backend collects metrics of the generated computation graph such as the minimal necessary plaintext size or the multiplicative depth. These metrics are used to suggest suitable encryption parameters. However, the user is not forced to use the suggested parameters. In scenarios where the data is already encrypted before compilation, developers can provide their parameters.

HElium and its IR are designed to be scheme-agnostic. The staged compilation and lowering allow for the extension of HElium by registering new backends. Each backend must register capabilities and its HE instruction set. HElium automatically assists with scheme decisions based on the capabilities of the backends. For example, we implemented backends for the TFHE [13] and CKKS [10] schemes to integrate support for efficient Boolean operations and fixed-point

Algorithm 1: PRE Operation Insertion Algorithm

Input: The computation graph with annotated nodes V_{in}, and the output node V_{out}.

Result: The computation graph with PRE operations.

1 nodes ← SUCC (V_{in});
2 **foreach** $n \in nodes$ **do**
3   nodes ← nodes \ {n} ∪ SUCC (n);
4   keys ← ∅;
5   **foreach** *operand* ∈ PRED *(n)* **do**
6   keys ← keys ∪ KEY (operand);
7   **if** *|keys| = 1* **then**
8   KEY (n) ← KEY (operand);
9   **if** *|keys| > 1* **then**
10   KEY (n) ← KEY (V_{out});
11   **foreach** *operand* ∈ PRED *(n)* **do**
12   PRE (operand, n);

arithmetic. However, their discussion is beyond the scope of this work.

5 EVALUATION

This section presents the evaluation of the effectiveness of the optimizations provided by the HElium compiler. We demonstrate their effectiveness in a use-case taken from the aforementioned health-care scenario.

5.1 Use-Case: Tumor Recurrence Rate

Hospitals, universities, and private healthcare companies collect a large amount of data about their patients. This includes, for example, protocols from medical treatments, drug applications, and surgical reports. With the introduction of the electronic medical file, more information will be digitalized [18]. This information about patients is highly confidential and high data protection requirements apply [23]. Therefore, sharing and aggregation of patient data is very complicated due to regulatory reasons. FHE in conjunction with PRE can enable researchers to combine encrypted patient-data sets with other researchers without sacrificing patients' privacy. It allows performing analyses of the shared data sets without disclosing individual patient records.

For example, records of genome mutations can be compared with histology analysis or surgical reports of tumor patients to obtain information about correlations between specific genome mutations and the properties of the tumor [38, 46]. That enables researchers to develop more specific treatment methods for tumors. One measure that is analyzed in the context of cancer research is the recurrence rate of variants of tumors. It represents the percentage of patients for whom particular cancer reappears. This measure indicates the medical treatment of the cancer patient. The particular tumor type can often only be determined by surgery and the following histology analysis. A non-invasive determination of the tumor type and its properties can allow treatment of the tumor without surgery. Therefore, finding correlations between genome mutations and tumor properties is an active field of research.

```
1  input a0: int       @Key0 <= Party0;
2  input b0: int[1000] @Key0 <= Party0;
3  input a1: int       @Key1 <= Party1;
4  input b1: int[1000] @Key1 <= Party1;
5  input a2: int       @Key2 <= Party2;
6  input b2: int[1000] @Key2 <= Party2;
7  input a3: int       @Key3 <= Party3;
8  input b3: int[1000] @Key3 <= Party3;
9
10 output R => Party_Out @Key_Out:
11     a0*b0 + a1*b1 + a2*b2 + a3*b3;
12 output n => Party_Out @Key_Out:
13     b0 + b1 + b2 + b3;
```

Figure 7: HElium Program: The Recurrence Rate of Tumors. For Simplicity $n = 4$.

The tumor recurrence rate r of patients is calculated as an average of the recurrence of a tumor in relation to the absolute number of patients with the same tumor variant. Namely,

$$r = \frac{\text{Number of Tumor Patients with Recurrence}}{\text{Number of Tumor Patients}} \quad (7)$$

In order to find correlations between genome mutations and the tumor recurrence rate, r can be calculated with the presence of different mutations. The presence of mutations for the i-th patient P_i can be efficiently encoded as bits of a bit vector denoted by $\mathbf{b}_i$. Each element b_j of $\mathbf{b}_i$ represents the presence or absence of a mutation j. Similarly, the tumor recurrence of a patient i can be encoded as a single bit a_i. The recurrence rate r in relation to the presence of mutations is defined as shown in Equation (8). Each element of the result vector $\mathbf{r}$ represents the recurrence rate in relation to the j-th mutation. This measure can help to find correlations between the presence of mutations and the recurrence of tumors.

$$\mathbf{r} = \frac{\sum_{i=0}^{n} a_i \cdot \mathbf{b_i}}{\sum_{i=0}^{n} \mathbf{b_i}} \quad (8)$$

Figure 7, shows the HElium code of an example implementation of the presented use case. Each patient data set consists of two values a_i and b_i that are provided under a unique key Key_i. In lines 1 to 8, the inputs are defined for four data sets of four different participants. The outputs and the function are declared in lines 10 and 11. The division of two unknown ciphertexts is an operation with high computational complexity. Therefore, the example program computes only the two sums with FHE. The division is performed afterwards on client side as a plaintext operation.

5.2 Evaluation Setup

All experiments were conducted on a dedicated server with 8 Intel(R) Core(TM) i7-9700T CPU cores and 32 GB of RAM. Each experiment is performed 100 times. The runtimes of compilation and execution are measured, and the average, as well as the standard deviation, are calculated. After each run, the results are decrypted and compared with the expected result of the corresponding plaintext function to verify correctness of the compiled program. The experiments are conducted for set sizes $n \in \{128, 256, 512, 768, 1024\}$. This
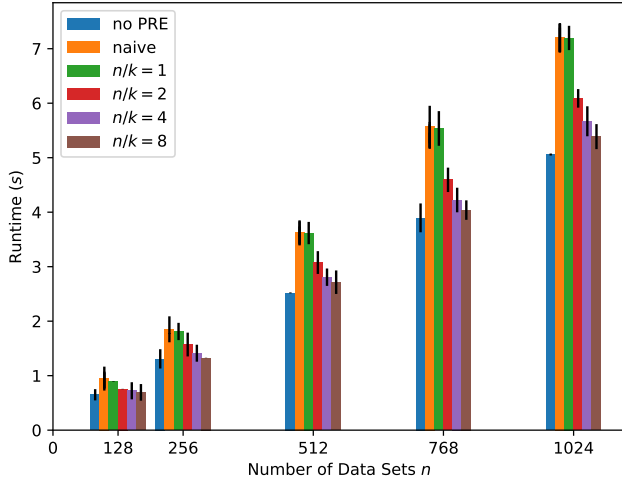
Figure 8: Evaluation Run-Time

relates to typical data-set sizes of medical research [46]. Often, only a few hundred patient records are available at a research institution depending on the type of tumor. The set sizes also allow to showcase the impact of HElium's optimizations such as the minimization of PRE operations and the balancing of computation graphs. Data-set sizes of power of two are well-suited to demonstrate the effectiveness of rebalancing. However, our compiler is not limited to those and supports arbitrary set sizes.

5.3 Efficiency of the PRE Integration

This section discusses how efficiently the compiler chooses to insert PRE operations. The compiler aims to insert only as few PRE operations as necessary. A naive approach of inserting PRE operations is to immediately proxy re-encrypt all inputs to a common key. We refer to the number of introduced PRE operations in this naive approach as p_{naive}. It is determined by the number of inputs i of the program, i.e., $p_{naive} = i$.

From a theoretical perspective, the minimal number of PRE operations p_{min} is determined by the keys of the inputs and outputs of the program, as shown in Equation (9). K_I and K_O denote the set of input and output keys, respectively. In the best case, only one PRE operation is needed for all keys that are members of K_I but not of K_O, i.e., the difference of K_I and K_O.

$$p_{min} = |K_I \setminus K_O| \qquad (9)$$

However, p_{min} does not take the structure of a program into account. Often, this theoretic minimum cannot be reached. HElium aims to achieve a number of PRE operations p close to $p_{min} \leq p \ll p_{naive}$.

We analyze the efficiency of HEliums approach to optimizing the insertion of PRE operations by measuring the number of inserted PRE operations and the resulting execution time of the evaluation of the HE circuit. Furthermore, the ratio between data sets n and keys k, denoted as n/k, has a significant impact on the optimization possibilities. To explore the effects of n/k on p and the evaluation run-time, we conduct measurements for $n/k \in \{1, 2, 4, 8\}$.

Figure 8 depicts the measured execution times of data-set sizes n and ratios n/k. Furthermore, the computation of each data-set size n without PRE and the naive approach is shown to emphasize the impact of our proposed optimizations. The larger n/k, the higher is the potential for optimization, hence, the smaller is the overhead introduced by PRE. The four right-most bars for each n show that the overhead quickly becomes negligible, even for relatively small n/k.

For $n/k = 1$, each data set is provided encrypted under its own key. In that case each input needs to be proxy re-encrypted in both approaches, i.e., $p = p_{naive}$ and their evaluation run-time is on par with each other. With increasing n/k ratio, the execution time decreases due to a lower number of inserted PRE operations. For example, in $n/k = 2$, p is reduced by 50 % compared to p_{naive}. In the case of $n/k = 8$, p is reduced by 87.5 % which results in a speedup of 1.34×. Most notably, optimizing PRE operations reduces the run-time overhead from 42.4 %, for the naive variant, to 6.5 % for a ratio of $n/k = 8$.

In summary, our analysis shows that the HElium compiler effectively minimizes PRE operations compared to the naive approach for scenarios $n/k > 1$. Consequently, due to our optimizations, the execution time overhead introduced by PRE operations quickly becomes negligible, even for relatively small n/k.

6 RELATED WORK

There exist a variety of (open-source) software libraries that implement one or multiple FHE schemes. Compared to domain-specific compilers, these libraries often miss domain-specific analyses and optimizations. Often being implemented using low-level programming languages like C++ increases the cognitive load to write efficient and secure programs. However, their APIs and data structures often capture domain-specific aspects which are seen in both library implementations and external DSLs.

6.1 Software Libraries for Fully Homomorphic Encryption

One of the first efficient software libraries is HElib [30, 31], providing implementations of the BGV [7] and CKKS [11] schemes. It offers bootstrapping for BGV and focuses on effective use of ciphertext packing and so-called "GHS optimizations" [28].

The Simple Encrypted Arithmetic Library (SEAL) [9] software library provides the BGV, CKKS, and BFV [6, 26] schemes and also supports bootstrapping. It supports various platforms such as Android, iOS, and embedded systems. SEAL provides abstractions for batching to implement more efficient homomorphic computations. While offering an easy-to-use API, the difference between efficient and inefficient implementations can be several orders of magnitude.

PALISADE [29] supports the BGV, BFV, CKKS, and FHEW [25] schemes and a more secure variant of the TFHE scheme, including bootstrapping. Another important feature is the support for proxy re-encryption (PRE). However, developers need to remember the keys to be used for PRE operations. In our approach, the target key to re-encrypt is implicit from the input program.

Lattigo [1] is a software library written in Go implementing the CKKS and the BFV schemes. Similar to PALISADE, Lattigo provides threshold variants of the implemented schemes.

The TFHE-rs library [47] provides a Rust implementation of the TFHE scheme [14]. TFHE is optimized for Boolean and integer arithmetics and fast bootstrapping. Among others, TFHE-rs offers programmable bootstrapping [15]. This technique allows the application of unary functions during bootstrapping via lookup tables. It is a technology to enable efficient computation of non-linear functions like square roots or activation functions like ReLU. The ability to compute non-linear functions is a big advantage for machine learning use cases.

6.2 Compilers for Fully Homomorphic Encryption

To the best of our knowledge, HElium is the first compiler that optimizes proxy re-encryption. Existing compilers often simply assume that data providers agree on a common key before homomorphic evaluation. However, agreeing on a joint key among multiple parties adds organizational overhead and requires some level of synchronicity. Especially the latter might already exclude many use cases where data is added over a longer period and a joint key agreement might be hard to perform. Thus, in this section, we focus on approaches to improving the generated circuits that are going to be homomorphically evaluated. We do not compare to work on exploiting parallelism on hardware platforms such as GPUs and FPGAs to improve the performance of homomorphic computations as this is beyond the focus of this work. See [44] for a more detailed comparison of FHE compilers in general.

Some approaches use DSLs to capture FHE programs as computations over simple expressions, generate related circuits, and link them with FHE backends. This approach is also used in Alchemy [19], a language and compiler integrated into Haskell. The integrated compiler automatically inserts maintenance operations into the program and selects suitable encryption parameters. It uses its implementation of the BFV scheme. Encrypt-Everything-Everywhere (E3) [12] uses template meta-programming to generate circuits from C++ code. For each supported backend it provides specializations of the templated classes. It also provides a collection of auxiliary tools to assemble an executable from configuration files.

Another class of approaches has a strong focus on FHE for machine learning applications and tensor operations. The Compiler and Runtime for Homomorphic Evaluation of Tensor Programs (CHET) [21] comes with a high-level language for machine-learning operations. It focuses on optimizations for matrix operations using efficient batching and optimizes scaling factors automatically, given a set of test inputs. EVA [16, 20] is a low-level compiler for vector-arithmetic operations. It extends the idea of CHET and is a C++ library with a Python interface that targets the CKKS implementation of SEAL. EVA inserts maintenance operations into the computation graph using a cost model and automatically selects encryption parameters. HECATE [34] provides a DSL and Intermediate Representation (IR) to optimize scale management of circuits using the CKKS scheme. It presents a flexible approach to rescaling that explores multiple different rescaling approaches as well as a preconstructed estimation of the latency of each FHE operation at different levels and polynomial modulus N. This estimator is used to estimate the performance of a HECATE IR program during optimization. nGraph-HE [5] is a homomorphic encryption backend to the Intel nGraph compiler for machine learning. It focuses on the inference of machine learning models over encrypted data. nGraph-HE translates Tensorflow computations into FHE circuits for BFV or CKKS. It applies optimizations on the computation graph and supports SIMD-packing, i.e., the efficient packing of multi-dimensional tensors into batched ciphertexts. HECO [43] provides an end-to-end architecture from a high-level DSL down to hardware-accelerated FHE circuits. The compiler infrastructure is based on MLIR [33] and optimizations focus on the automatic generation of batched HE instructions.

A different approach is taken by Porcupine [17]. It employs program synthesis to automatically generate vectorized homomorphic encryption programs from a high-level specification of a kernel and a sketch of that kernel. A sketch consists of operands which may be defined as holes and possible arithmetic instructions used to implement the specification. Holes in the sketch define the space that is searched. To reduce the search space, rotation is not an instruction living on its own but is considered as input to arithmetic instructions in the sketch. An SMT solver verifies that a solution of a sketch fulfills the specification.

7 CONCLUSION

Efficient privacy-preserving analysis of confidential data, e.g., for medical research, is a pressing problem of our time. Fully homomorphic encryption (FHE) schemes allow computations over encrypted data. FHE has proven to be a valuable means to leverage the value of data without sacrificing the privacy of individuals. However, FHE operations are computationally expensive and optimizations require a high level of expertise. Compilers for FHE can automate many optimization aspects and improve usability of FHE, especially for non-cryptographers. However, FHE operations require data to be encrypted under the same key. Yet, many real-world use cases imply multi-party scenarios where data is encrypted under different keys, thus, excluding standard FHE for numerous use cases.

FHE schemes with proxy re-encryption (PRE) capabilities allow the conversion of ciphertexts encrypted under different keys into ciphertexts encrypted under a common key, thus, filling the above-mentioned gap. The automated optimization of FHE programs with PRE operations, however, remains yet to be solved. We target this open question and propose HElium, the first compiler for fully homomorphic encryption with native proxy re-encryption support. HEliums's domain-specific language acts as an abstraction layer for complex FHE and PRE operations and, therefore, significantly improves the ease of use of FHE and PRE, especially for non-cryptographers. Among a variety of optimizations, our HElium compiler minimizes of the number of PRE operations and conducts balancing of the computation graph. This substantially improves the ease of use of fully homomorphic encryption in multi-party scenarios, especially for non-cryptographers.

Based on the real-world example of tumor recurrence rate, an active area of medical research, we show that HElium minimizes PRE operations compared to a baseline approach for scenarios where the number of inputs n is larger than the number of input keys k, i.e., $n/k > 1$. Most notably, for $n/k = 8$, the number of PRE operations is reduced by 87.5 % compared to the baseline approach, which results in a speedup of 1.34×. Consequently, our optimizations ensure that

the execution time overhead introduced by PRE operations quickly becomes negligible, even for relatively small n/k.

In the future we plan to further increase performance by extending the runtime of HElium to evaluate homomorphic circuits in parallel. We see a promising future research direction in adding a cost-model based optimization taking both performance and noise accumulation into account. The compiler cannot reason about these independently, as performance cost often depends on noise growth because of the underlying mathematical homomorphic encryption instruction implementations.

REFERENCES

[1] Lattigo v4. Online: https://github.com/tuneinsight/lattigo, Aug. 2022. EPFL-LDS, Tune Insight SA.

[2] ACAR, A., AKSU, H., ULUAGAC, A. S., AND CONTI, M. A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys 51*, 4 (jul 2018).

[3] ASHAROV, G., JAIN, A., LÓPEZ-ALT, A., TROMER, E., VAIKUNTANATHAN, V., AND WICHS, D. Multiparty computation with low communication, computation and interaction via threshold fhe. In *Advances in Cryptology – EUROCRYPT 2012* (Berlin, Heidelberg, 2012), D. Pointcheval and T. Johansson, Eds., Springer Berlin Heidelberg, pp. 483–501.

[4] BLAZE, M., BLEUMER, G., AND STRAUSS, M. Divertible protocols and atomic proxy cryptography. In *Advances in Cryptology — EUROCRYPT'98* (Berlin, Heidelberg, 1998), K. Nyberg, Ed., Springer Berlin Heidelberg, pp. 127–144.

[5] BOEMER, F., LAO, Y., CAMMAROTA, R., AND WIERZYNSKI, C. ngraph-he. In *Proceedings of the 16th ACM International Conference on Computing Frontiers* (New York, NY, USA, 2019), F. Palumbo, M. Becchi, M. Schulz, and K. Sato, Eds., ACM, pp. 3–13.

[6] BRAKERSKI, Z. Fully homomorphic encryption without modulus switching from classical gapsvp. In *Advances in Cryptology – CRYPTO 2012* (2012).

[7] BRAKERSKI, Z., GENTRY, C., AND VAIKUNTANATHAN, V. (leveled) fully homomorphic encryption without bootstrapping. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference on - ITCS '12* (New York, New York, USA, 2012), S. Goldwasser, Ed., ACM Press, pp. 309–325.

[8] CARPOV, S., DUBRULLE, P., AND SIRDEY, R. Armadillo. In *Proceedings of the 3rd International Workshop on Security in Cloud Computing* (New York, NY, 2015), ACM, pp. 13–19.

[9] CHEN, H., LAINE, K., AND PLAYER, R. Simple Encrypted Arithmetic Library - SEAL v2.1. In *Financial Cryptography and Data Security*, vol. 10323. Springer International Publishing, Cham, 2017, pp. 3–18.

[10] CHEN, H., LAINE, K., AND RINDAL, P. Fast private set intersection from homomorphic encryption. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (New York, NY, USA, 2017), ACM.

[11] CHEON, J. H., KIM, A., KIM, M., AND SONG, Y. Homomorphic encryption for arithmetic of approximate numbers. In *Advances in Cryptology – ASIACRYPT 2017*, T. Takagi and T. Peyrin, Eds., vol. 10624 of *Lecture notes in computer science*. Springer International Publishing, Cham, 2017, pp. 409–437.

[12] CHIELLE, E., MAZONKA, O., GAMIL, H., TSOUTSOS, N. G., AND MANIATAKOS, M. E3: A framework for compiling c++ programs with encrypted operands. Cryptology ePrint Archive, Report 2018/1013, 2018. https://ia.cr/2018/1013.

[13] CHILLOTTI, I., GAMA, N., GEORGIEVA, M., AND IZABACHÈNE, M. Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds. In *Advances in cryptology - ASIACRYPT 2016* (Berlin, 2016), Lecture notes in computer science, Springer, pp. 3–33.

[14] CHILLOTTI, I., GAMA, N., GEORGIEVA, M., AND IZABACHÈNE, M. Tfhe: Fast fully homomorphic encryption over the torus. Cryptology ePrint Archive, Paper 2018/421, 2018.

[15] CHILLOTTI, I., JOYE, M., AND PAILLIER, P. Programmable bootstrapping enables efficient homomorphic inference of deep neural networks. In *Cyber Security Cryptography and Machine Learning: 5th International Symposium (CSCML 2021)* (2021), Springer, pp. 1–19.

[16] CHOWDHARY, S., DAI, W., LAINE, K., AND SAARIKIVI, O. EVA Improved: Compiler and Extension Library for CKKS. In *Proceedings of the 9th on Workshop on Encrypted Computing & Applied Homomorphic Cryptography* (Virtual Event Republic of Korea, Nov. 2021), ACM, pp. 43–55.

[17] COWAN, M., DANGWAL, D., ALAGHI, A., TRIPPEL, C., LEE, V. T., AND REAGEN, B. Porcupine: a synthesizing compiler for vectorized homomorphic encryption. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation* (New York, NY, USA, June 2021), PLDI 2021, ACM, pp. 375–389.

[18] COWIE, M. R., BLOMSTER, J. I., CURTIS, L. H., DUCLAUX, S., FORD, I., FRITZ, F., GOLDMAN, S., JANMOHAMED, S., KREUZER, J., LEENAY, M., MICHEL, A., ONG, S.,

[18] PELL, J. P., SOUTHWORTH, M. R., STOUGH, W. G., THOENES, M., ZANNAD, F., AND ZALEWSKI, A. Electronic health records to facilitate clinical research. *Clinical research in cardiology : official journal of the German Cardiac Society 106*, 1 (2017), 1–9.

[19] CROCKETT, E., PEIKERT, C., AND SHARP, C. Alchemy: A language and compiler for homomorphic encryption made easy. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security* (New York, NY, USA, 2018), CCS '18, ACM, pp. 1020–1037.

[20] DATHATHRI, R., KOSTOVA, B., SAARIKIVI, O., DAI, W., LAINE, K., AND MUSUVATHI, M. Eva: an encrypted vector arithmetic language and compiler for efficient homomorphic computation. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (New York,NY,United States, 2020), ACM Digital Library, Association for Computing Machinery, pp. 546–561.

[21] DATHATHRI, R., SAARIKIVI, O., CHEN, H., LAINE, K., LAUTER, K., MALEKI, S., MUSUVATHI, M., AND MYTKOWICZ, T. Chet: An optimizing compiler for fully-homomorphic neural-network inferencing. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation* (New York, NY, USA, 2019), PLDI 2019, ACM, p. 142–156.

[22] DATHATHRI, R., SAARIKIVI, O., CHEN, H., LAINE, K., LAUTER, K. E., MALEKI, S., MUSUVATHI, M., AND MYTKOWICZ, T. Chet: Compiler and runtime for homomorphic evaluation of tensor programs. *CoRR abs/1810.00845* (2018).

[23] DOVE, E. S., AND PHILLIPS, M. Privacy law, data sharing policies, and medical data: A comparative perspective. In *Medical Data Privacy Handbook*. Springer International Publishing, Cham, 2015, pp. 639–678.

[24] DRUCKER, N., AND GUERON, S. Combining homomorphic encryption with trusted execution environment: A demonstration with paillier encryption and sgx. In *Proceedings of the 2017 International Workshop on Managing Insider Security Threats* (New York, NY, USA, 2017), MIST '17, Association for Computing Machinery, p. 85–88.

[25] DUCAS, L., AND MICCIANCIO, D. Fhew: Bootstrapping homomorphic encryption in less than a second. In *Advances in Cryptology - EUROCRYPT 2015* (Berlin, Heidelberg, 2015), Security and Cryptology, Springer Berlin Heidelberg and Imprint: Springer, pp. 617–640.

[26] FAN, J., AND VERCAUTEREN, F. Somewhat practical fully homomorphic encryption. In *Cryptology ePrint Archive, Report 2012/144* (2012).

[27] GENTRY, C. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009.

[28] GENTRY, C., HALEVI, S., AND SMART, N. P. Fully Homomorphic Encryption with Polylog Overhead. In *Advances in Cryptology – EUROCRYPT 2012*, vol. 7237. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 465–482.

[29] HALEVI, S., POLYAKOV, Y., AND SHOUP, V. An improved rns variant of the bfv homomorphic encryption scheme. In *Topics in Cryptology – CT-RSA 2019* (Cham, 2019), Springer International Publishing, pp. 83–105.

[30] HALEVI, S., AND SHOUP, V. Algorithms in HElib. In *Advances in Cryptology – CRYPTO 2014*, vol. 8616. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014, pp. 554–571.

[31] HALEVI, S., AND SHOUP, V. Design and implementation of helib: a homomorphic encryption library. Cryptology ePrint Archive, Paper 2020/1481, 2020.

[32] IVAN, A., AND DODIS, Y. Proxy cryptography revisited. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)* (2003), The Internet Society.

[33] LATTNER, C., AMINI, M., BONDHUGULA, U., COHEN, A., DAVIS, A., PIENAAR, J., RIDDLE, R., SHPEISMAN, T., VASILACHE, N., AND ZINENKO, O. Mlir: A compiler infrastructure for the end of moore's law. *arXiv preprint arXiv:2002.11054* (2020).

[34] LEE, Y., HEO, S., CHEON, S., JEONG, S., KIM, C., KIM, E., LEE, D., AND KIM, H. HECATE: Performance-Aware Scale Optimization for Homomorphic Encryption Compiler. In *2022 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)* (Seoul, Korea, Republic of, Apr. 2022), IEEE, pp. 193–204.

[35] LINDELL, Y. *Tutorials on the Foundations of Cryptography: Dedicated to Oded Goldreich*, 1st ed. Springer Publishing Company, Incorporated, 2017.

[36] LÓPEZ-ALT, A., TROMER, E., AND VAIKUNTANATHAN, V. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *Proceedings of the 44th Symposium on Theory of Computing* (New York, New York, USA, 2012), STOC '12, ACM Press, pp. 1219–1234.

[37] LYUBASHEVSKY, V., PEIKERT, C., AND REGEV, O. On ideal lattices and learning with errors over rings. In *Annual international conference on the theory and applications of cryptographic techniques* (2010), Springer, pp. 1–23.

[38] M BRADEN, A., V STANKOWSKI, R., M ENGEL, J., AND A ONITILO, A. Breast cancer biomarkers: risk assessment, diagnosis, prognosis, prediction of treatment efficacy and toxicity, and recurrence. *Current pharmaceutical design 20*, 30 (2014), 4879–4898.

[39] PAILLIER, P. Public-key cryptosystems based on composite degree residuosity classes. In *Advances in Cryptology — EUROCRYPT '99* (Berlin, Heidelberg, 1999), J. Stern, Ed., Springer Berlin Heidelberg, pp. 223–238.

[40] POLYAKOV, Y., ROHLOFF, K., SAHU, G., AND VAIKUNTANATHAN, V. Fast proxy re-encryption for publish/subscribe systems. *ACM Transactions on Privacy and Security 20*, 4 (2017), 1–31.

[41] Rivest, R. L., Shamir, A., and Adleman, L. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM 21*, 2 (1978), 120–126.

[42] Smart, N. P., and Vercauteren, F. Fully homomorphic simd operations. *Designs, Codes and Cryptography 71*, 1 (2014), 57–81.

[43] Viand, A., Jattke, P., Haller, M., and Hithnawi, A. HECO: Fully homomorphic encryption compiler. *usenix* (2023).

[44] Viand, A., Jattke, P., and Hithnawi, A. Sok: Fully homomorphic encryption compilers. In *2021 IEEE Symposium on Security and Privacy (SP)* (2021), IEEE, pp. 1092–1108.

[45] Viand, A., and Shafagh, H. Marble: Making fully homomorphic encryption accessible to all. In *Proceedings of the 6th Workshop on Encrypted Computing & Applied Homomorphic Cryptography* (New York, NY, 2018), M. Brenner, Ed., ACM Conferences, ACM, pp. 49–60.

[46] Yuzawa, S., Nishihara, H., and Tanaka, S. Genetic landscape of meningioma. *Brain tumor pathology 33*, 4 (2016), 237–247.

[47] Zama. TFHE-rs: A Pure Rust Implementation of the TFHE Scheme for Boolean and Integer Arithmetics Over Encrypted Data, 2022. https://github.com/zama-ai/tfhe-rs.