


# Building an End-to-End Toolchain for Fully Homomorphic Encryption with MLIR

Alexander Viand, Patrick Jattke, Miro Haller, Anwar Hithnawi

**ETH** zürich

# Cloud Computing



“Where the sensitive information is concentrated, that is where the spies will go. This is just a fact of life.”  
former NSA official Ken Silva.

Software Vulnerabilities

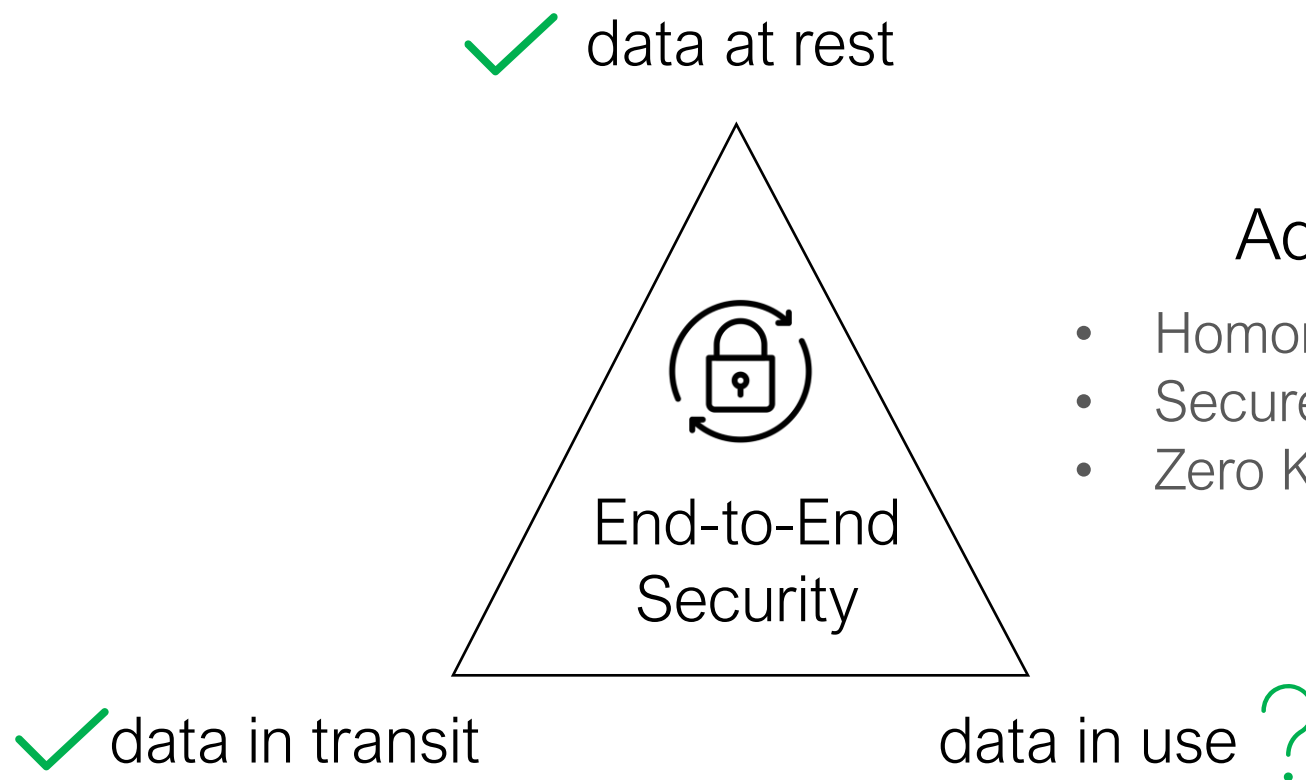
Insider Threats

Physical Attacks

# End-to-End Encrypted Systems



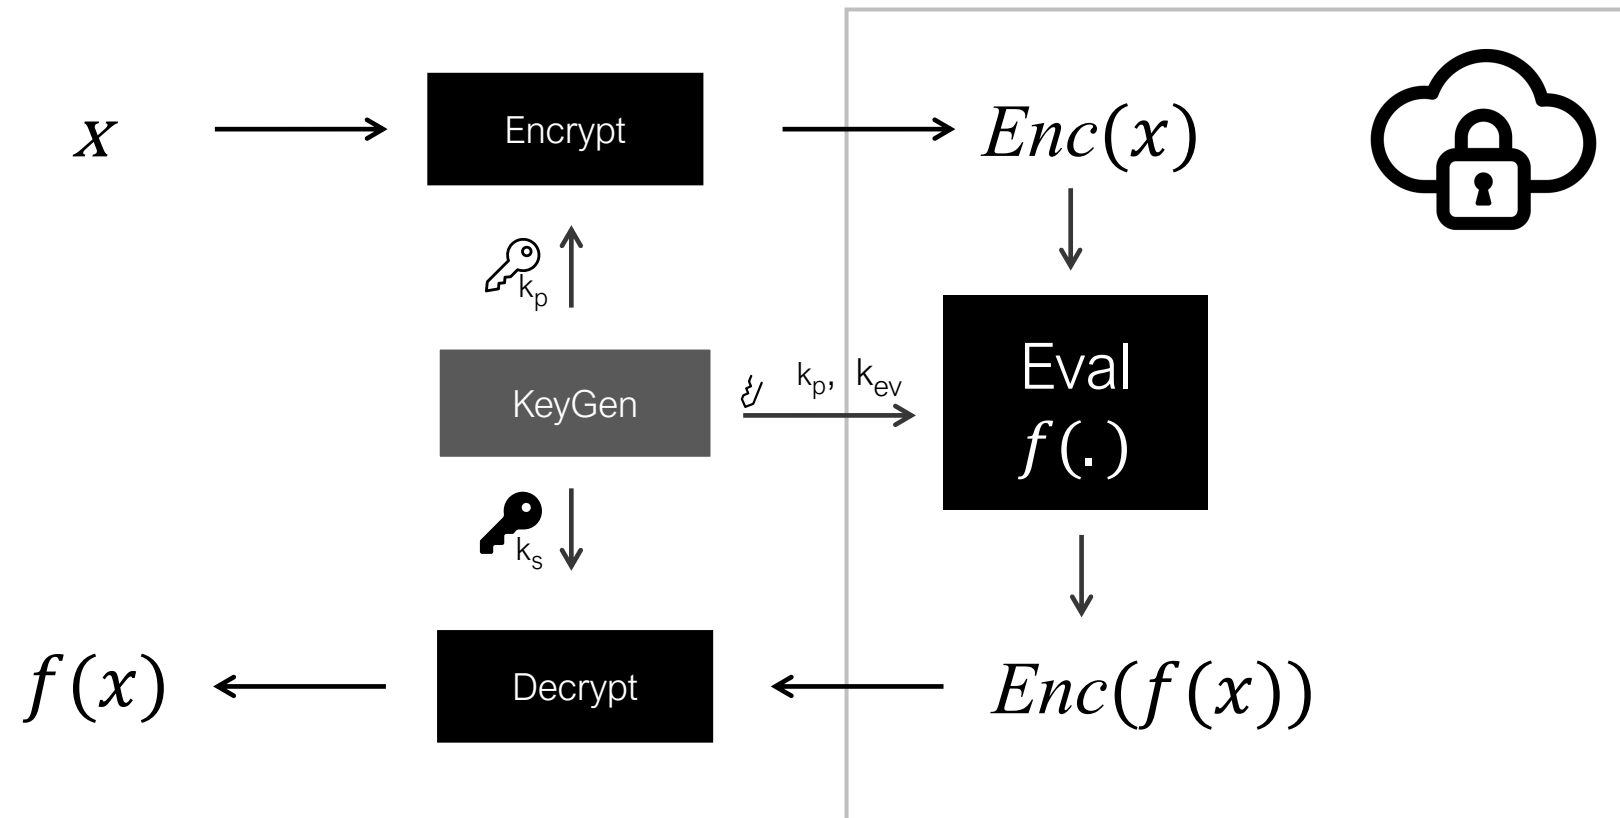
# Modern Cryptography



## Advanced Crypto

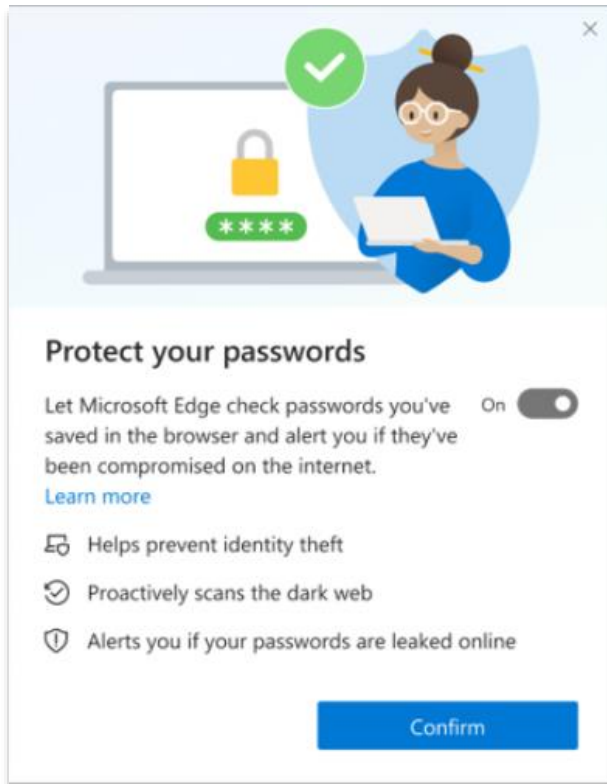
- Homomorphic Encryption
- Secure Multi-party Computation
- Zero Knowledge Proofs

# Fully Homomorphic Encryption



Computing on encrypted data

# Fully Homomorphic Encryption




- Alice91 | 0791 [lock icon] \*\*\*\*
- A.Sample | 1234 [lock icon] \*\*\*\*
- Bob | b4dp455 [lock icon] \*\*\*\*
- Alice | 12345 [lock icon] \*\*\*\*



- Alice | 12345
- Bob | wordpass
- Eve | pa55word
- Joe | correcth..
- ⋮
- Steve | hunter2

Delegate the **processing** of data without giving away **access** to it

# Fully Homomorphic Encryption




Alice91 | 0791

A.Sample | 1234

Bob | b4dp455

Alice | 12345



🔒 🔒 ? \*

🔒 🔒 ? \*

🔒 🔒 ? \*

🔒 🔒 ? \*



Alice | 12345

Bob | wordpass

Eve | pa55word

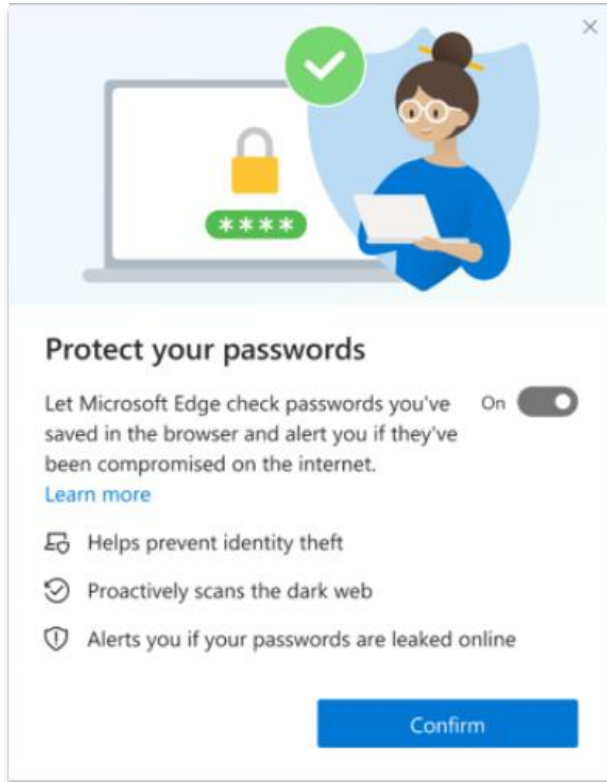
Joe | correcth..



Steve | hunter2

Delegate the **processing** of data without giving away **access** to it

# Fully Homomorphic Encryption



- Alice91 | 0791
- A.Sample | 1234
- Bob | b4dp455
- Alice | 12345



- Alice | 12345
- Bob | wordpass
- Eve | pa55word
- Joe | correcth..
- ⋮
- Steve | hunter2


Delegate the **processing** of data without giving away **access** to it



# FHE: Theory to Practice

$Enc(0) \odot Enc(1)$

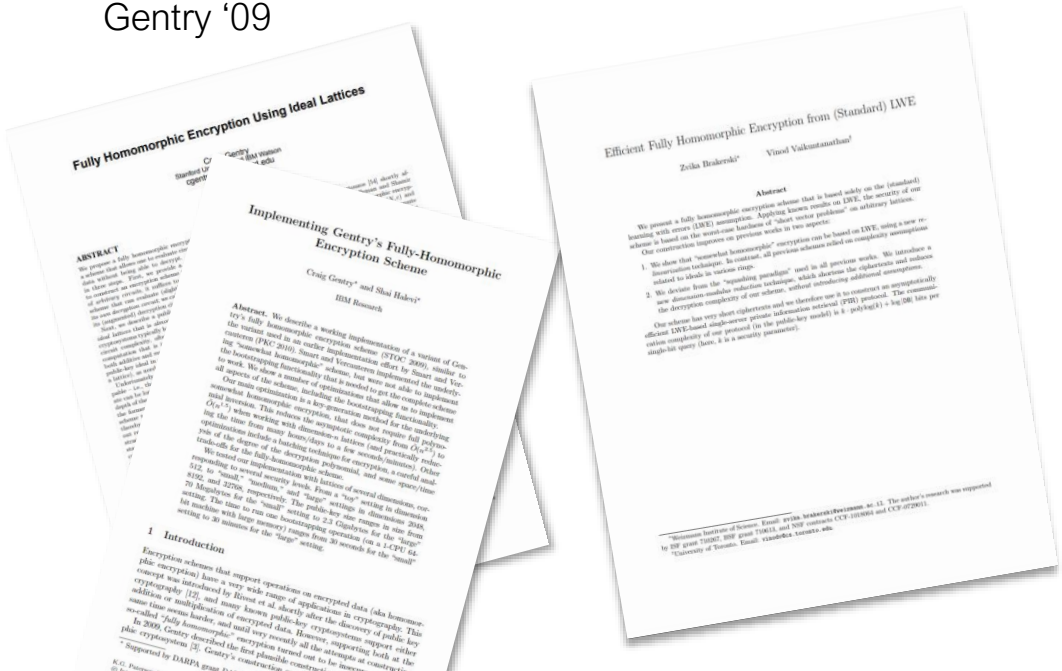
Learning **With Errors**



$$c = \sum_{i=1}^n a_i s_i + e$$


where  $a \xleftarrow{\$} \mathbb{Z}_q^n, e \xleftarrow{\chi} \mathbb{Z}_q, s \in \mathbb{Z}_q^n$

Gentry '09



# FHE: Theory to Practice

## $Enc(0) \odot Enc(1)$ (Ring-) Learning With Errors



$$c = a \cdot s + e$$


where  $a \xleftarrow{\$} R_q, e \xleftarrow{\chi} R_q, s \in R_q$   
 $R = \mathbb{Z}[X]/(X^n + 1)$

Gentry '09



# FHE: Theory to Practice

$Enc(0) \odot Enc(1)$  **(Ring-) Learning With Errors**




$$c = a \cdot s + e$$

where  $a \xleftarrow{\$} R_q, e \xleftarrow{\chi} R_q, s \in R_q$   
 $R = \mathbb{Z}[X]/(X^n + 1)$

Gentry '09


# FHE: Theory to Practice

## Binary FHE



FHEW'14, TFHE '16

## $Enc(0) \odot Enc(1)$ (Ring-) Learning With Errors





$$c = a \cdot s + e$$

where  $a \xleftarrow{\$} R_q, e \xleftarrow{\chi} R_q, s \in R_q$   
 $R = \mathbb{Z}[X]/(X^n + 1)$

Gentry '09

BGV '12, B/FV '12, CKKS '16

## Arithmetic FHE




(a) Approximation of ReLU using different methods

Figure from CryptoDL: Deep Neural Networks over Encrypted Data [HTG17]


# FHE: Theory to Practice

## Binary FHE



FHEW'14, TFHE '16

## $Enc(0) \odot Enc(1)$ (Ring-) Learning With Errors



$$c = a \cdot s + e$$

where  $a \xleftarrow{\$} R_q, e \xleftarrow{\chi} R_q, s \in R_q$   
 $R = \mathbb{Z}[X]/(X^n + 1)$



Gentry '09

## Hardware Acceleration



BGV '12, B/FV '12, CKKS '16

## Arithmetic FHE



(a) Approximation of ReLU using different methods

Figure from CryptoDL: Deep Neural Networks over Encrypted Data [HTG17]









Existing tools make important contributions,  
but are too **narrowly focussed**

# End-to-End FHE Toolchain

Computer Program


```
void hd(vector<bool>&u,  
        vector<bool>&v)  
{  
    int sum = 0;  
    for(int i = 0;  
        i < v.size();  
        ++i)  
    {  
        sum += (v[i]==u[i]);  
    }  
}
```



FHE Compiler



Arithmetic Circuit




# End-to-End FHE Toolchain


Computer Program

```
void hd(vector<bool>&u,  
        vector<bool>&v)  
{  
    int sum = 0;  
    for(int i = 0;  
        i < v.size();  
        ++i)  
    {  
        sum += (v[i]==u[i]);  
    }  
}
```


FHE Compiler



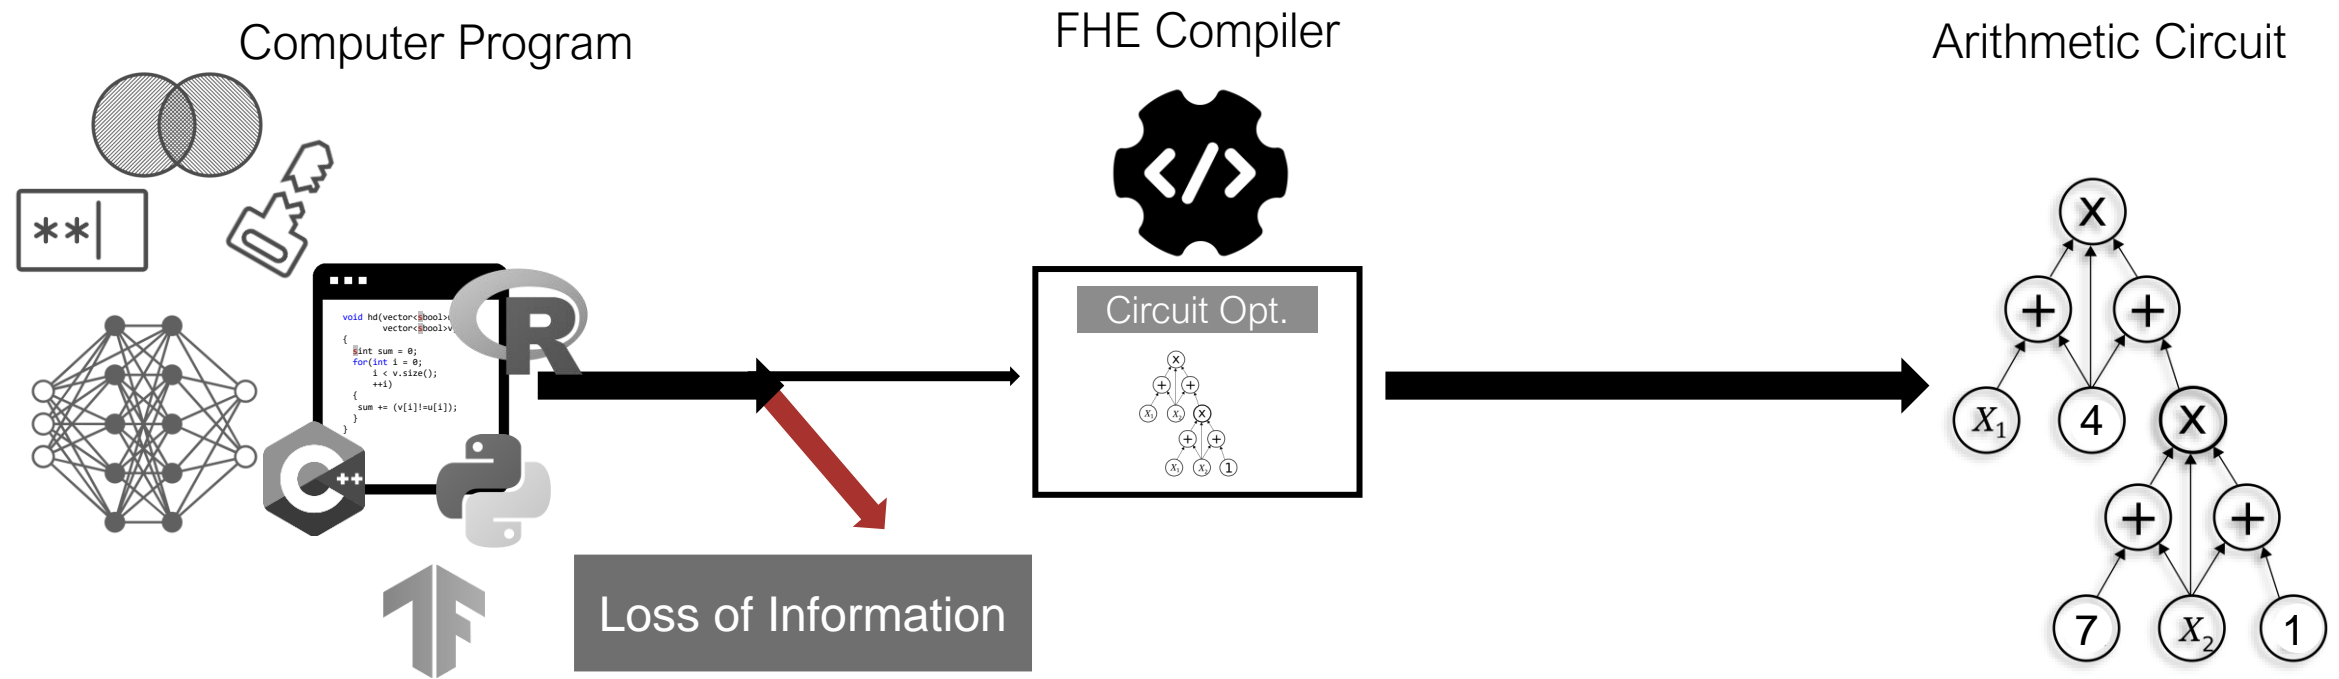
Arithmetic Circuit



# End-to-End FHE Toolchain



# End-to-End FHE Toolchain




# End-to-End FHE Toolchain

Computer Program


```
void hd(vector<bool>u,  
vector<bool>v)  
{  
  int sum = 0;  
  for(int i = 0;  
      i < v.size();  
      ++i)  
  {  
    sum += (v[i]==u[i]);  
  }  
}
```



FHE Compiler



Arithmetic Circuit




# End-to-End FHE Toolchain

Computer Program

```
void hd(vector<bool>&u,  
vector<bool>&v)  
{  
  int sum = 0;  
  for(int i = 0;  
      i < v.size();  
      ++i)  
  {  
    sum += (v[i]==u[i]);  
  }  
}
```



FHE Compiler



Secure and Efficient FHE Solutions




# End-to-End FHE Toolchain

Computer Program


```
void hd(vector<bool>&u,  
       vector<bool>&v)  
{  
    int sum = 0;  
    for(int i = 0;  
        i < v.size();  
        ++i)  
    {  
        sum += (v[i]==u[i]);  
    }  
}
```



FHE Compiler



Secure and Efficient FHE Solutions






# End-to-End FHE Toolchain

Computer Program


```
void hd(vector<bool>&u,  
        vector<bool>&v)  
{  
    int sum = 0;  
    for(int i = 0;  
        i < v.size();  
        ++i)  
    {  
        sum += (v[i]==u[i]);  
    }  
}
```



FHE Compiler



Secure and Efficient FHE Solutions




# End-to-End FHE Toolchain

Computer Program



```
void hd(vector<bool>&u,  
vector<bool>&v)  
{  
  int sum = 0;  
  for(int i = 0;  
      i < v.size();  
      ++i)  
  {  
    sum += (v[i]==u[i]);  
  }  
}
```



FHE Compiler



Secure and Efficient FHE Solutions




# End-to-End FHE Toolchain

Computer Program




```
void hd(vector<bool>&u,  
        vector<bool>&v)  
{  
    int sum = 0;  
    for(int i = 0;  
        i < v.size();  
        ++i)  
    {  
        sum += (v[i]==u[i]);  
    }  
}
```



FHE Compiler



Secure and Efficient FHE Solutions




# End-to-End FHE Toolchain

Computer Program




```
void hd(vector<bool>&u,  
        vector<bool>&v)  
{  
    int sum = 0;  
    for(int i = 0;  
        i < v.size();  
        ++i)  
    {  
        sum += (v[i]==u[i]);  
    }  
}
```



FHE Compiler



Secure and Efficient FHE Solutions




# End-to-End FHE Toolchain

Computer Program


```
void hd(vector<bool>&u,  
vector<bool>&v)  
{  
  int sum = 0;  
  for(int i = 0;  
      i < v.size();  
      ++i)  
  {  
    sum += (v[i]==u[i]);  
  }  
}
```



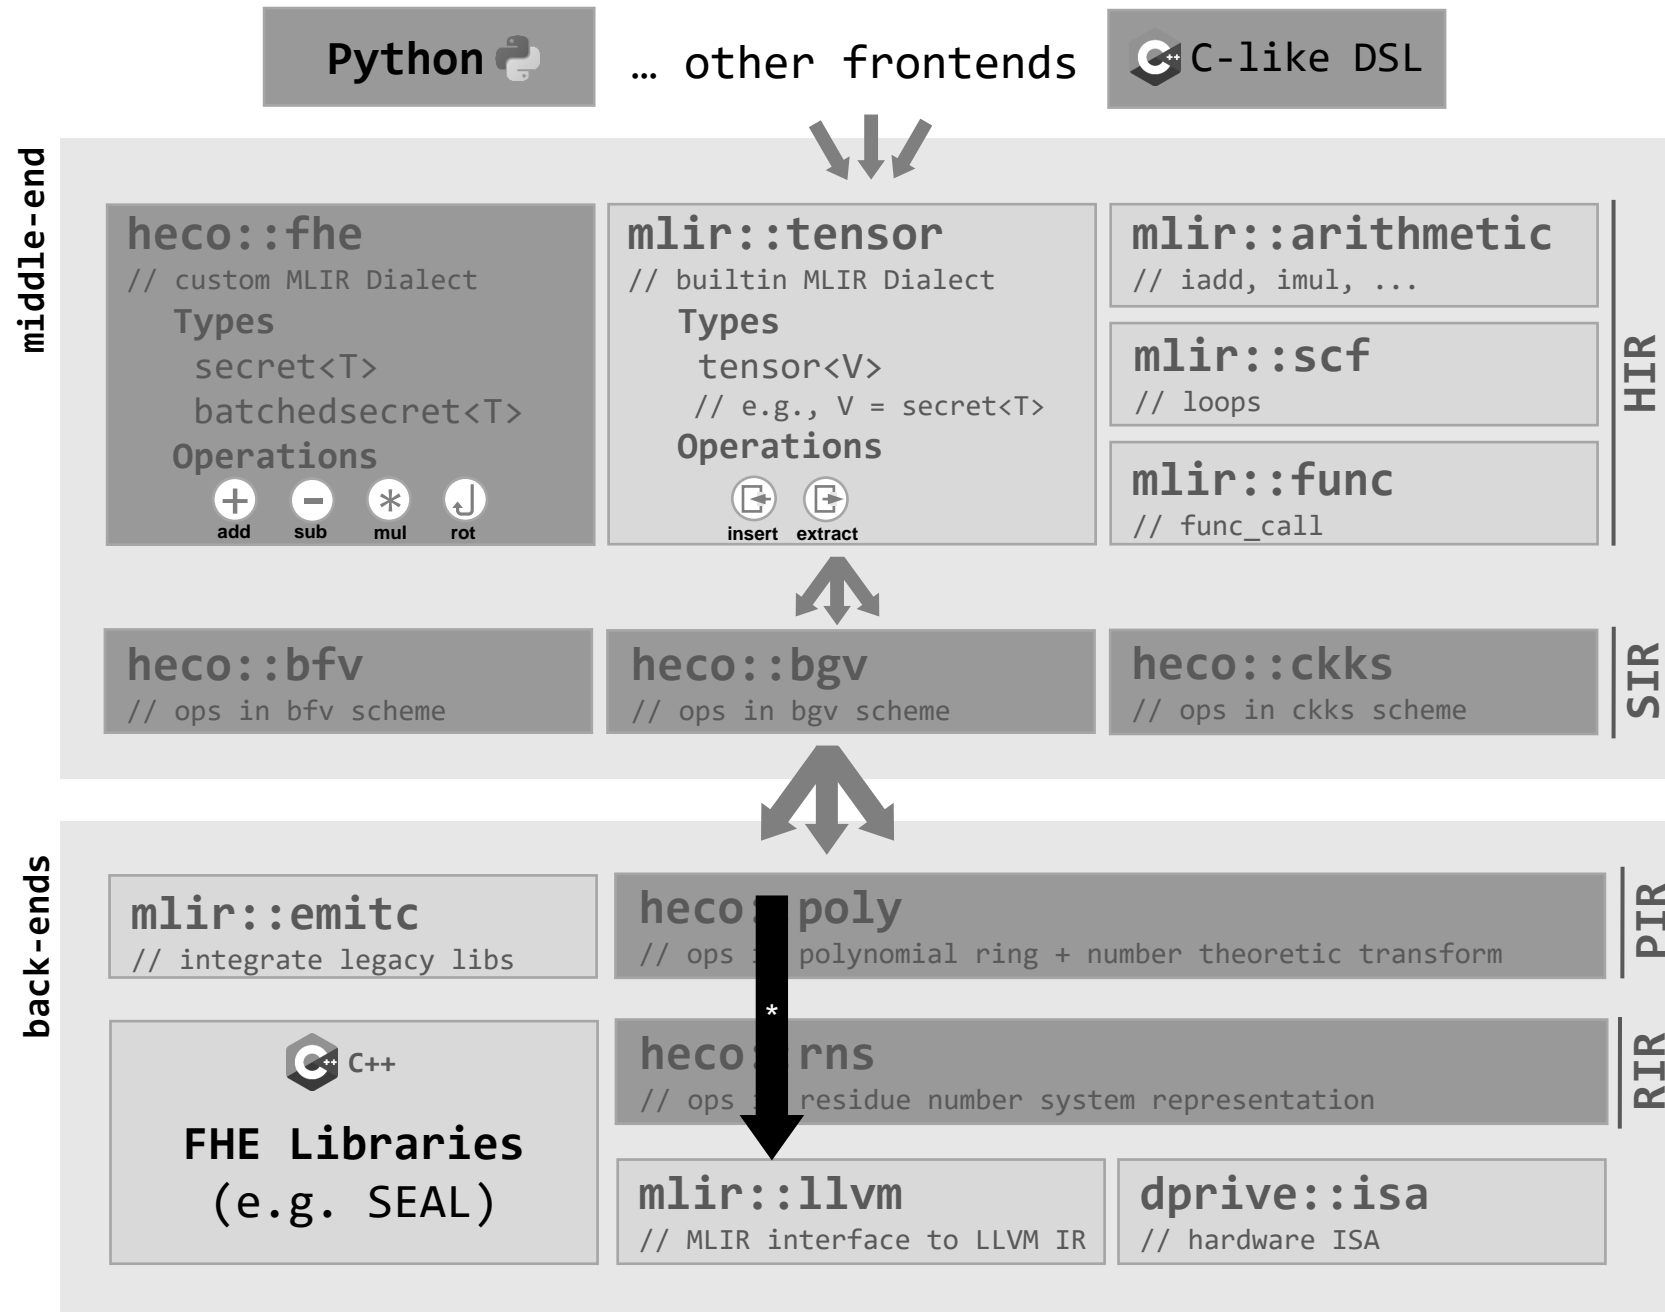
FHE Compiler



Secure and Efficient FHE Solutions



# HECO Architecture



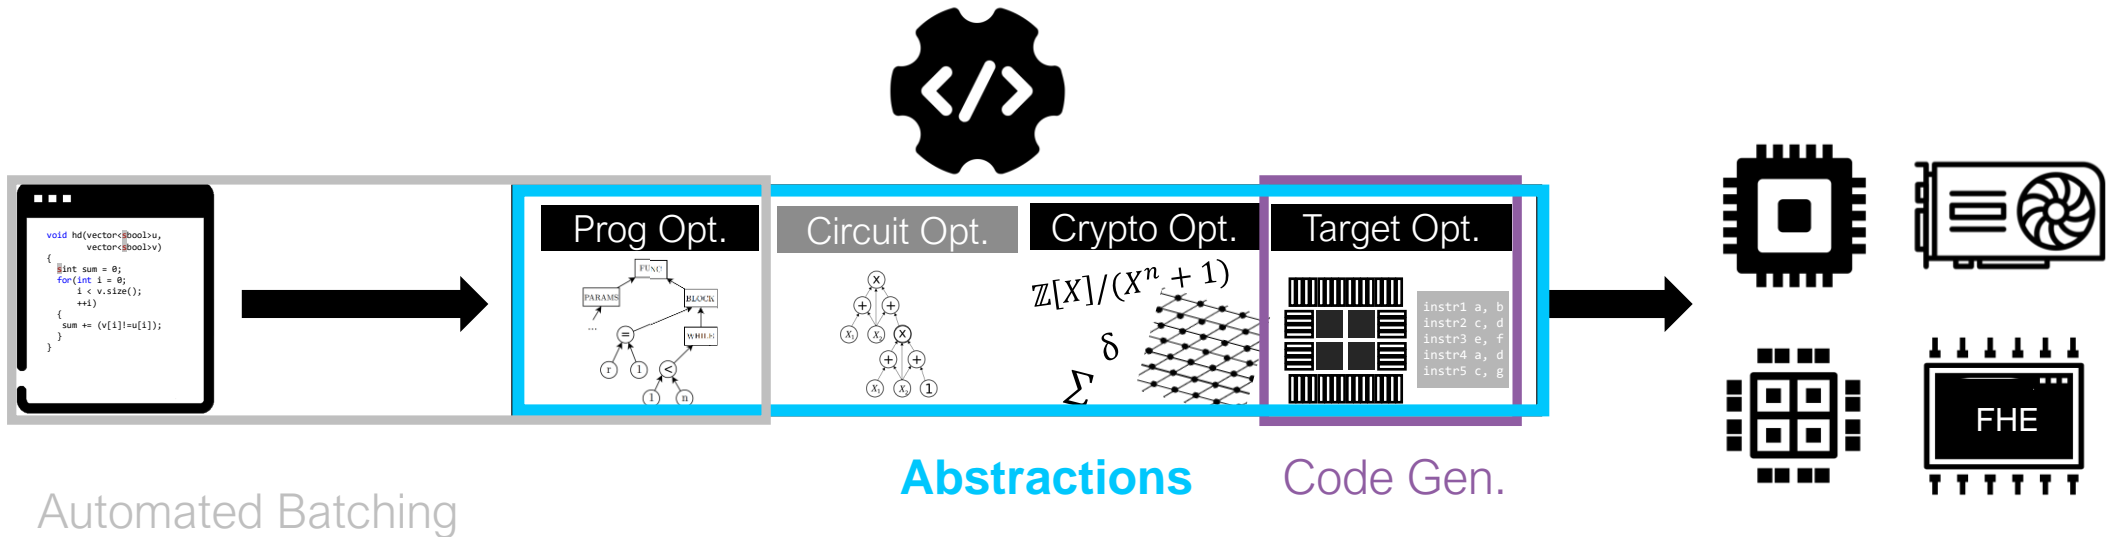
\* Juneyoung Lee, Prototyping a compiler for homomorphic encryption in MLIR, EuroLLVM'22

# HECO: Modular End-to-End Design

Computer Program

FHE Compiler

Secure and Efficient FHE Solutions



FHE Scheme Standardization

- T
- N



Microsoft



ZAMA


FHE Intermediate Representation Standardization

- 2018 draft API standard not adopted/implemented
- Significant FHE compiler efforts are accelerating
  - Need to re-visit standardization of abstractions
  - Expand beyond "FHE API" abstraction

# SIMD-like Parallelism

```
Standard C++  
int foo(int[] x,int[] y){  
    int[] r;  
    for(i = 0; i < 6; ++i){  
        r[i] = x[i] * y[i]  
    }  
    return r;  
}
```

```
Batched FHE  
int foo(int[] a,int[] b){  
    return a * b;  
}
```




No efficient free permutation or scatter/gather



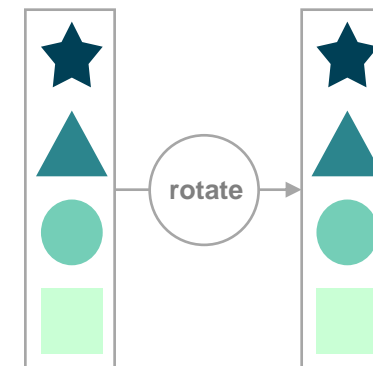
# SIMD-like Parallelism

```
Standard C++  
int foo(int[] x,int[] y){  
    int[] r;  
    for(i = 0; i < 6; ++i){  
        r[i] = x[i] * y[i]  
    }  
    return r;  
}
```

```
Batched FHE  
int foo(int[] a,int[] b){  
    return a * b;  
}
```



No efficient free permutation or scatter/gather




Only cyclical rotations


# SIMD-like Parallelism

```
Standard C++  
int foo(int[] x,int[] y){  
    int[] r;  
    for(i = 0; i < 6; ++i){  
        r[i] = x[i] * y[i]  
    }  
    return r;  
}
```

```
Batched FHE  
int foo(int[] a,int[] b){  
    return a * b;  
}
```



No efficient free permutation or scatter/gather



Only cyclical rotations

# HECO: Automatic Code Optimizations for Efficient Fully Homomorphic Encryption

Alexander Viand, Patrick Jattke, Miro Haller, Anwar Hühnawi  
ETH Zürich

## Abstract

In recent years, Fully Homomorphic Encryption (FHE) has undergone several breakthroughs and advancements leading to a leap in performance. Today, performance is no longer a major barrier to adoption. Instead, it is the complexity of deploying FHE in practice and at scale. Several FHE compilers have emerged recently to ease FHE development. However, none of these answer how to automatically transform imperative programs to secure and efficient FHE implementations. This is a fundamental issue that needs to be addressed before we can realistically expect broader use of FHE. We present HECO, a new FHE compiler that takes high-level code and automatically transforms it into FHE code. HECO is designed to be efficient and secure. It takes high-level code and automatically transforms it into FHE code. HECO is designed to be efficient and secure. It takes high-level code and automatically transforms it into FHE code.

## 1 Introduction

Privacy and security are gaining tremendous importance across all organizations, as public perception of these issues has increased. This has led to a surge in demand for secure and confidential computing solutions that protect data confidentiality while in transit, rest, and in-use. Fully Homomorphic Encryption (FHE) is a key secure computation technology that enables systems to preserve the confidentiality of data end-to-end, including while in use. Hence, allowing construction of computations without having to grant access to the data. In the last decade, advances in FHE schemes have

propelled FHE from a primarily theoretical breakthrough to a practical solution for a wide range of applications [1–3]. In recent years, we have seen FHE emerging in real-world applications, e.g., Microsoft's Edge browser uses FHE to utilize its privacy-preserving password monitor [4]. With upcoming dedicated hardware accelerators for FHE, promising further speedup [4, 5], FHE will soon be competitive for an even wider set of applications. Today, FHE is promising to become a major driver in transforming privacy. For example, a recent report by Gartner, Inc. predicts [6, 7] that FHE "will be a core technology for many future SaaS offerings."

Though promising in its potential, developing FHE applications remains a complex and tedious task. Writing efficient FHE programs that deliver on their potential performance is a challenge because of the need to map applications to their non-intuitive, highly-optimized FHE code. Code that is properly optimized for this paradigm is often overwhelmingly more efficient than poorly adapted code. Therefore, optimizing paradigms in general are being used to ease FHE development. Today, this remains an unsolved issue posing a major barrier to wider adoption. In order to facilitate easier development of FHE, general-purpose preserving technologies, the current FHE development ecosystem must better address the needs of non-expert developers. Existing tools fail to provide an easy-to-use, high-level purpose, and flexible toolchain to transform high-level code into FHE code. Existing tools are those created by experts, and they are not designed to adapt to FHE's unique paradigm, making only minimal changes to their code such as marking inputs as secret. Instead, tools should bridge the gap between traditional programming patterns and the restrictions of FHE development by automatically rewriting and optimizing programs to match FHE's unique programming paradigm.

the knowledge gap of-the-art of FHE objectives. First, to build applications to provide the successes and

ive survey of characteristics, ence by expert, nity applica- performance.

of a selection analysis of implement benchmarks of FHE for strengths



and approach.

urity

schemes rely on post-quantum hardness as- are widely believed to be secure for the fore- The community has developed a relatively of the concrete hardness of many of these parameter choices for a number of FHE standardized [28]. Nevertheless, some is worth noting that FHE does not the client requested or even no cal- techniques to address integrity, FHE does not by default provide might be able to learn informa- applied. Different techniques, of protection, can be used to the homomorphic encryption (HE). Since these arise only a third party gains access to ciphertext decryptions tion performed on the ence and refer to the

design

of HECO and

high-level language),

represent program

through a series of passes

and the code for this

model definition API, making it easy to use. In addition, the Graph FHE function, the programmatic interface, the concrete library has these are currently no

remain complex to use, in terms of usability of interpretability. For example, even efforts are being made to improve the usability of the API, e.g., by providing a high-level language

the complexity of the implementation. However, there are some subtle issues that have emerged in the context of FHE. For example, the complexity of the implementation is still a challenge. For example, the complexity of the implementation is still a challenge. For example, the complexity of the implementation is still a challenge.



github.com/MarbleHE/HECO

arxiv.org/abs/2202.01649