UNIVERSITY OF DELAWARE

HE and ZKP: Privacy-Enhancing Technologies

Dept. of Electrical and Computer Engineering
University of Delaware

Charles Gouert & Dimitris Mouris
{cgouert, jimouris}@udel.edu

Trustworthy Computing Group

April 21, 2022

Trustworthy Computing (TwC) Group

Nektarios G. Tsoutsos
Assistant Professor

Charles Gouert
PhD Student

Dimitris Mouris
PhD Student

Lars Folkerts
PhD Student

- Privacy-Enhancing Technologies & Applied Cryptography
- Protecting the smart-grid & distributed energy resources
- Digital-Manufacturing & 3D-printing

https://github.com/TrustworthyComputing

Homomorphic Encryption

Zero-Knowledge Proofs

The “Holy Grail” of Cryptography

Enables computing on encrypted data!

Solves privacy issues in cloud computing:

- Protects *data in use*
- Completely outsource computation

Applications:

- Privacy-Preserving Statistics & Surveys
- Privacy-Preserving personalized advertising
- Private Machine Learning
- Secure e-voting

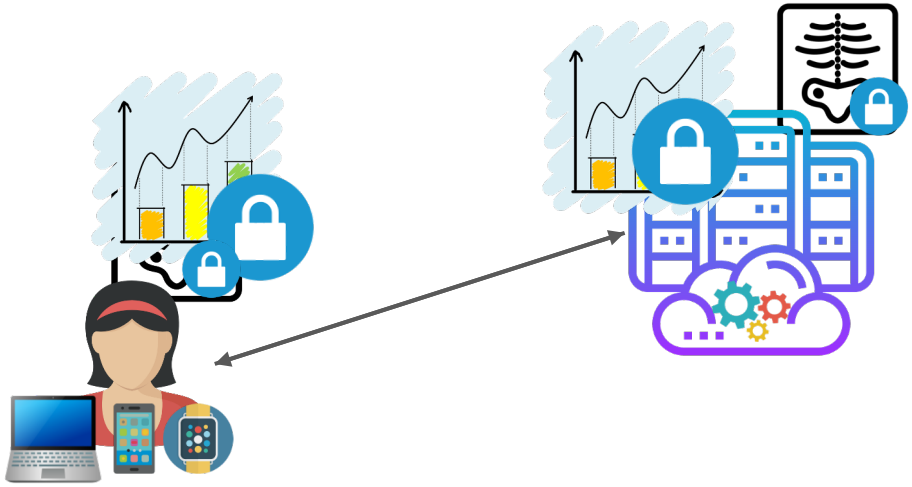

Partially Homomorphic Encryption (PHE)

Fastest and Oldest Form of HE

Allows for unlimited addition OR multiplication on ciphertexts

Famous schemes include:

- RSA
- ElGamal
- Paillier

Applications:

- Electronic voting
- Privacy preserving machine learning (PPML)

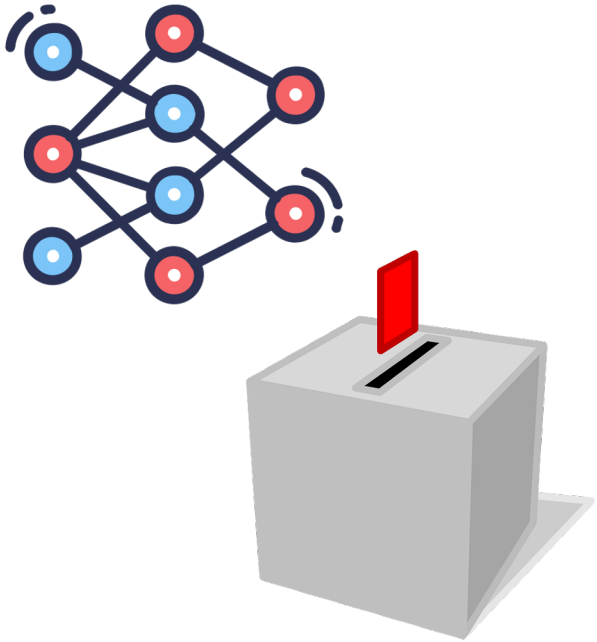

Leveled Homomorphic Encryption (LHE)

Bounded addition AND multiplication on ciphertexts

Ciphertexts are tuples of high-degree polynomials:

- Integer/Floating point plaintexts
- Poly degrees between 1024 and 32768
- Large coefficients (hundreds of bits)

Popular schemes:

- BGV (Integer)
- BFV (Integer)
- CKKS (Floating Point)

PALISADE

Why are the number of operations bounded?

The Role of Noise in HE

LHE (and FHE) schemes derive their security from the **Learning With Errors** problem

Noise injected during encryption step

Noise increases during processing:

- Linearly for addition
- Exponentially for multiplication

Ciphertexts that grow too noisy are undecryptable!

Modulus switching scales down the noise

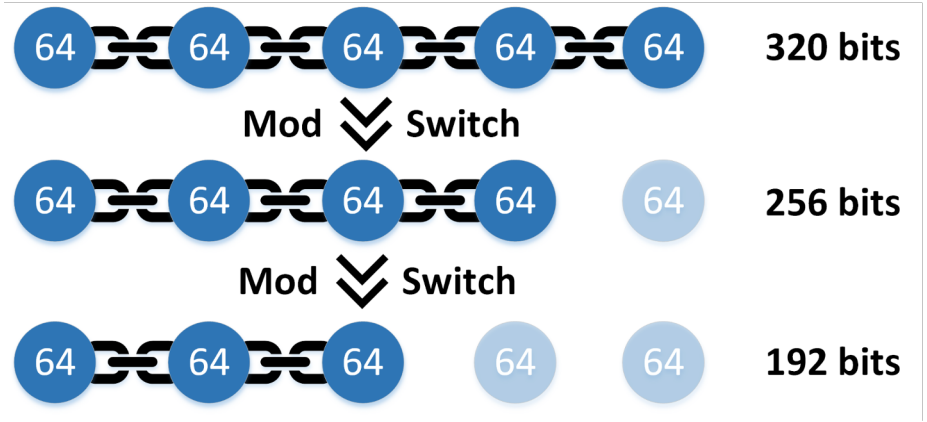

Fully Homomorphic Encryption (FHE)

Unlimited addition AND multiplication on ciphertexts

Any LHE scheme can become FHE by adding a bootstrapping procedure

Popular schemes:

- FHEW (Binary)
- TFHE/Zama (Binary/Floating Point)

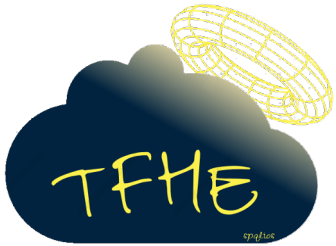

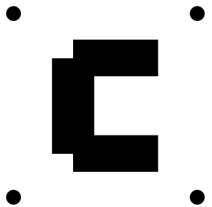

PALISADE

FHE Bootstrapping

Intuition: decryption eliminates noise

Solution: perform decryption in *HE domain* using an encryption of the secret key

Benefit: more powerful than mod switching

Drawback: computationally expensive

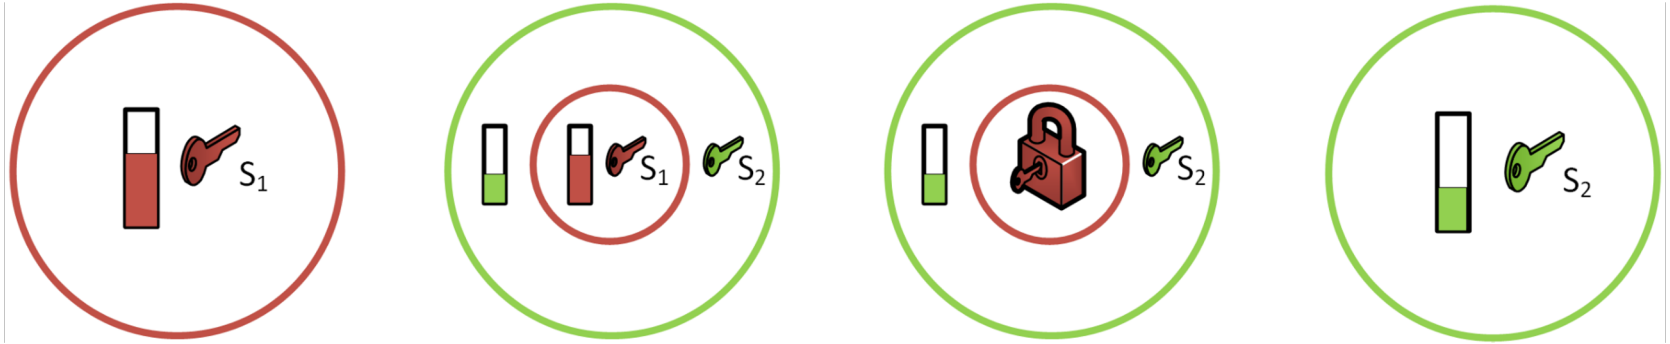

T2: Benchmarking FHE and LHE Libraries

How do you choose which HE library to use for a given application?

T2 is a framework for comparing HE libraries/schemes

- Standardized benchmark suite (in T2DSL)
- Compiler that maps T2DSL to popular open-source libraries

Benchmarks include:

- ML Inference
- Sorting
- Private Information Retrieval (PIR)

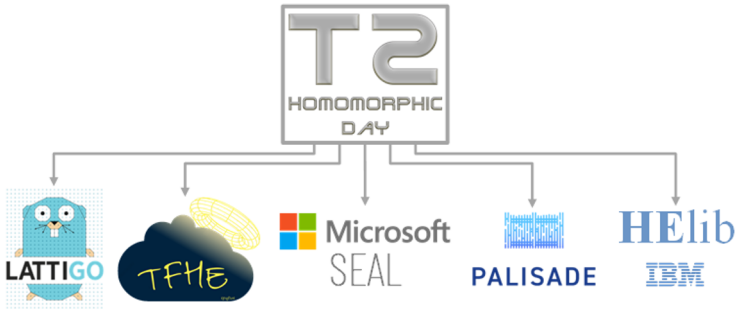

Gouert, C., Mouris, D., & Tsoutsos, N. G. (2022). New Insights into Fully Homomorphic Encryption Libraries via Standardized Benchmarks. Cryptology ePrint Archive.

ROMEO: Converting HDL Designs to FHE Programs

Synthesis with Yosys:

- Perform logic optimization
- Export netlist to EDIF

TFHE as a backend:

- Encrypted logic gates
- Ciphertexts encode bits

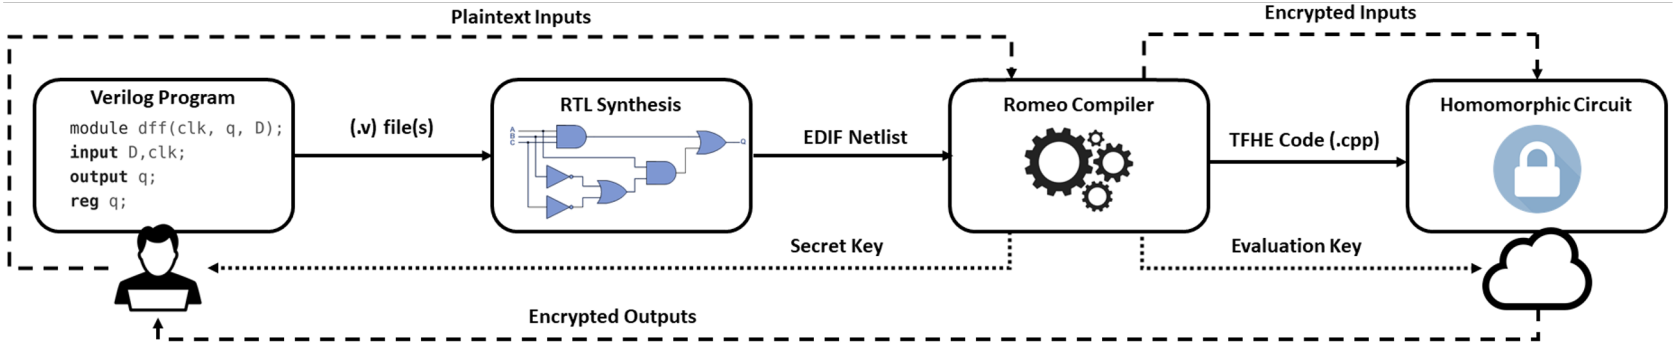

Gouert, C., & Tsoutsos, N. G. (2020, July). Romeo: conversion and evaluation of HDL designs in the encrypted domain. In 2020 57th ACM/IEEE Design Automation Conference (DAC) (pp. 1-6). IEEE.

REDsec: Running Encrypted DNNs in Seconds

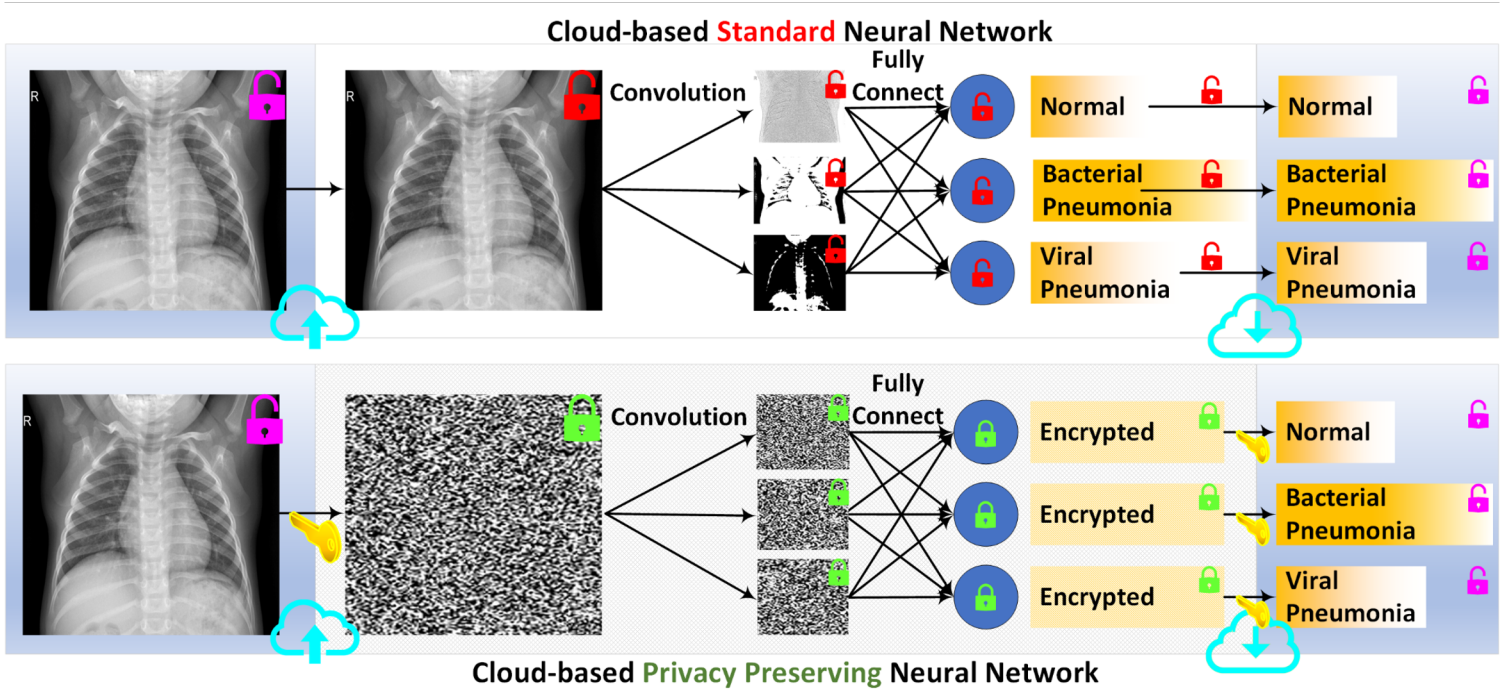

Folkerts, L., Gouert, C., & Tsoutsos, N. G. (2021). REDsec: Running Encrypted DNNs in Seconds. Cryptology ePrint Archive.

REDsec: Running Encrypted DNNs in Seconds

A framework for private neural network inference with support for both CPU and GPU cloud instances

Modules:

- Neural network compiler
- Secure inference library with TFHE (CPU)
- Secure inference library with (RED)cuFHE (GPU)

MNIST classification in 2 seconds (for 97.2% accuracy)

First FHE work to classify ImageNet (Real-life ML)

- Binary AlexNet: 1.7 hours

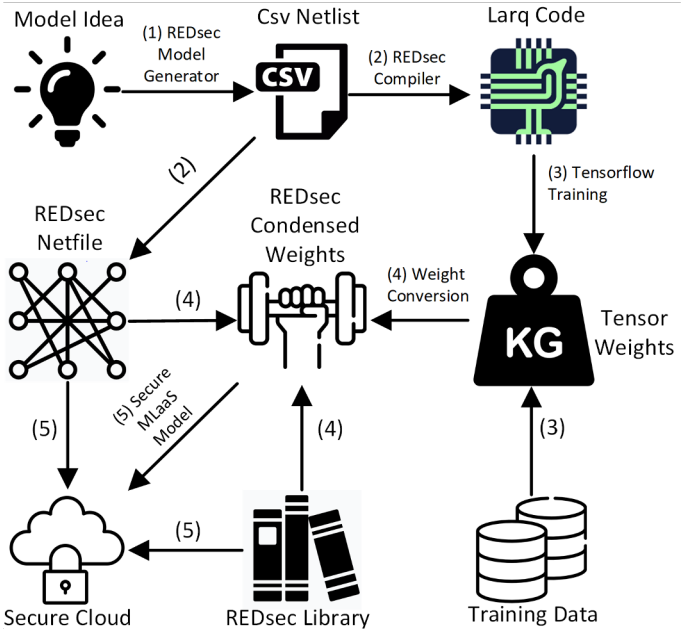

Folkerts, L., Gouert, C., & Tsoutsos, N. G. (2021). REDsec: Running Encrypted DNNs in Seconds. Cryptology ePrint Archive.

The Future of HE

FHE is getting faster every year!

Bootstrapping has improved from over 30 minutes to approximately 10 ms

Software:

- New schemes
- Optimizations (like RNS)

Hardware:

- Increasing GPU support
- Dedicated hardware for FHE

Homomorphic Encryption

Zero-Knowledge Proofs

Zero-Knowledge Proofs

Prove that a statement is true without conveying any additional information about the statement!

Solves privacy issues in cloud computing:

- Protects *data in use*
- Completely outsource computation

Applications:

- Privacy-preserving statistics & surveys
- Privacy-protecting digital currency
- Secure e-voting

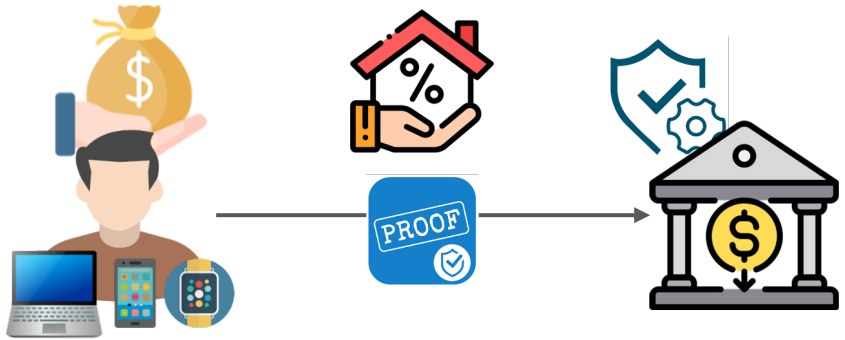

How can Bob convince the Bank?

Mortgage Application

How can a developer use these tools?

Zilch: Framework for usable ZKP development

How can Bob create a Zero-Knowledge Proof that he is eligible for a loan?

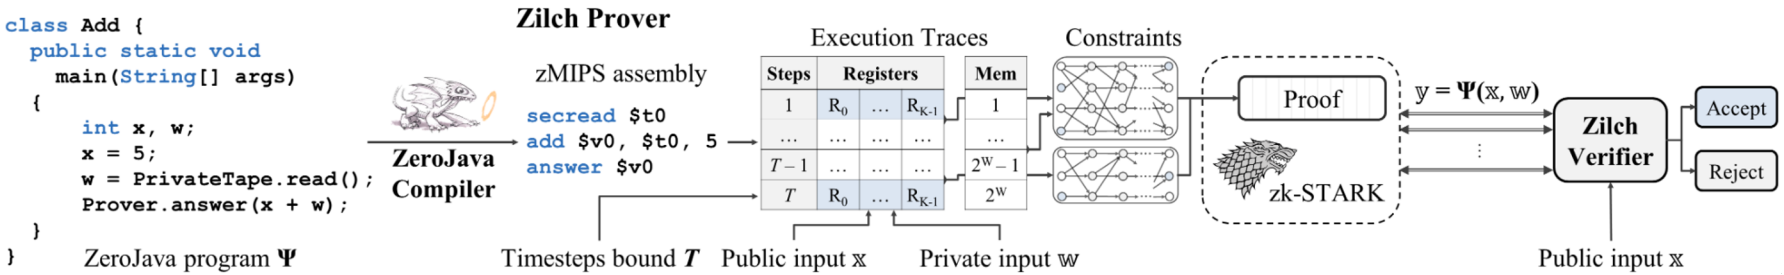

- **Range proofs:** Prove that a secret number is within a public range (e.g., loan application)
- **Private e-voting:** Prove that a vote is correct without revealing the vote!

- Mouris, D., & Tsoutsos, N. G. (2021). Zilch: A Framework for Deploying Transparent Zero-Knowledge Proofs. *IEEE Transactions on Information Forensics and Security*, *16*, 3269-3284.

Privacy-Preserving IP Verification

How to solve this deadlock without a trusted 3rd party?

A big company buys circuit modules from other vendors. Are these modules trustworthy?

[1]: Methodology for **proving functionality** of a private Integrated Circuit using **ZKPs**.

[2]: Methodology for **proving functional properties (area, performance, power)**.

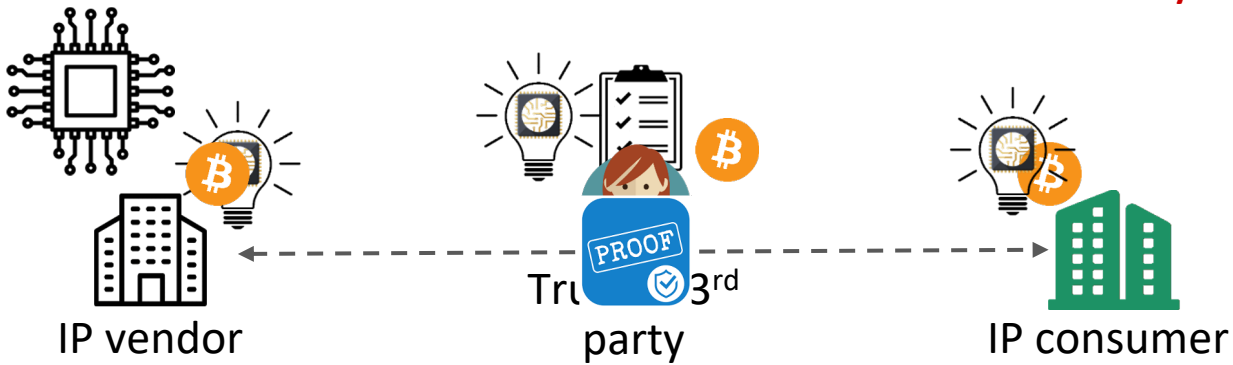

[1] Mouris, D., Gouert, C., & Tsoutsos, N. G. (2021). Privacy-preserving ip verification. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems.

[2] Mouris, D., & Tsoutsos, N. G. (2020, July). Pythia: Intellectual property verification in zero-knowledge. In 2020 57th ACM/IEEE Design Automation Conference (DAC) (pp. 1-6). IEEE.

Homomorphic Encryption

+

Zero-Knowledge Proofs

Masquerade: Verifiable Multi-Party Aggregation

Private Survey: Securely compute statistics & verify participants' inputs without seeing them.

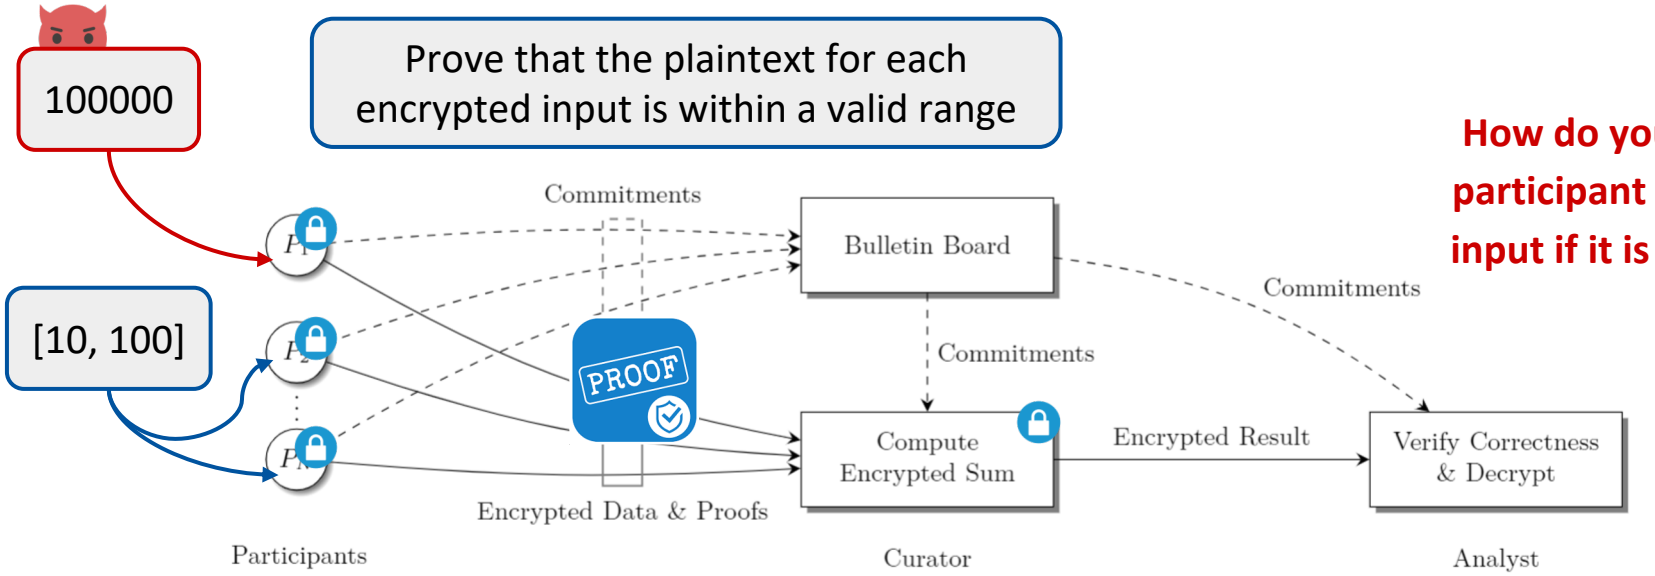

How do you know if a participant sent a valid input if it is encrypted?

- Mouris, D., & Tsoutsos, N. G. (2021). Masquerade: Verifiable Multi-Party Aggregation with Secure Multiplicative Commitments. Cryptology ePrint Archive.

Questions?

Thank you!

Charles Gouert
cgouert@udel.edu

github.com/cgouert

Dimitris Mouris
jimouris@udel.edu

github.com/jimouris

@jimouris

linkedin.com/in/jimouris