SEC 595: Special Topics in Information Assurance & Security

Encrypted Computing

(or Computing on Encrypted Data)

Lecture 1: Introduction

Instructor: Muhamad Felemban

Why A Course on Data Privacy?

Personal Information is Everywhere

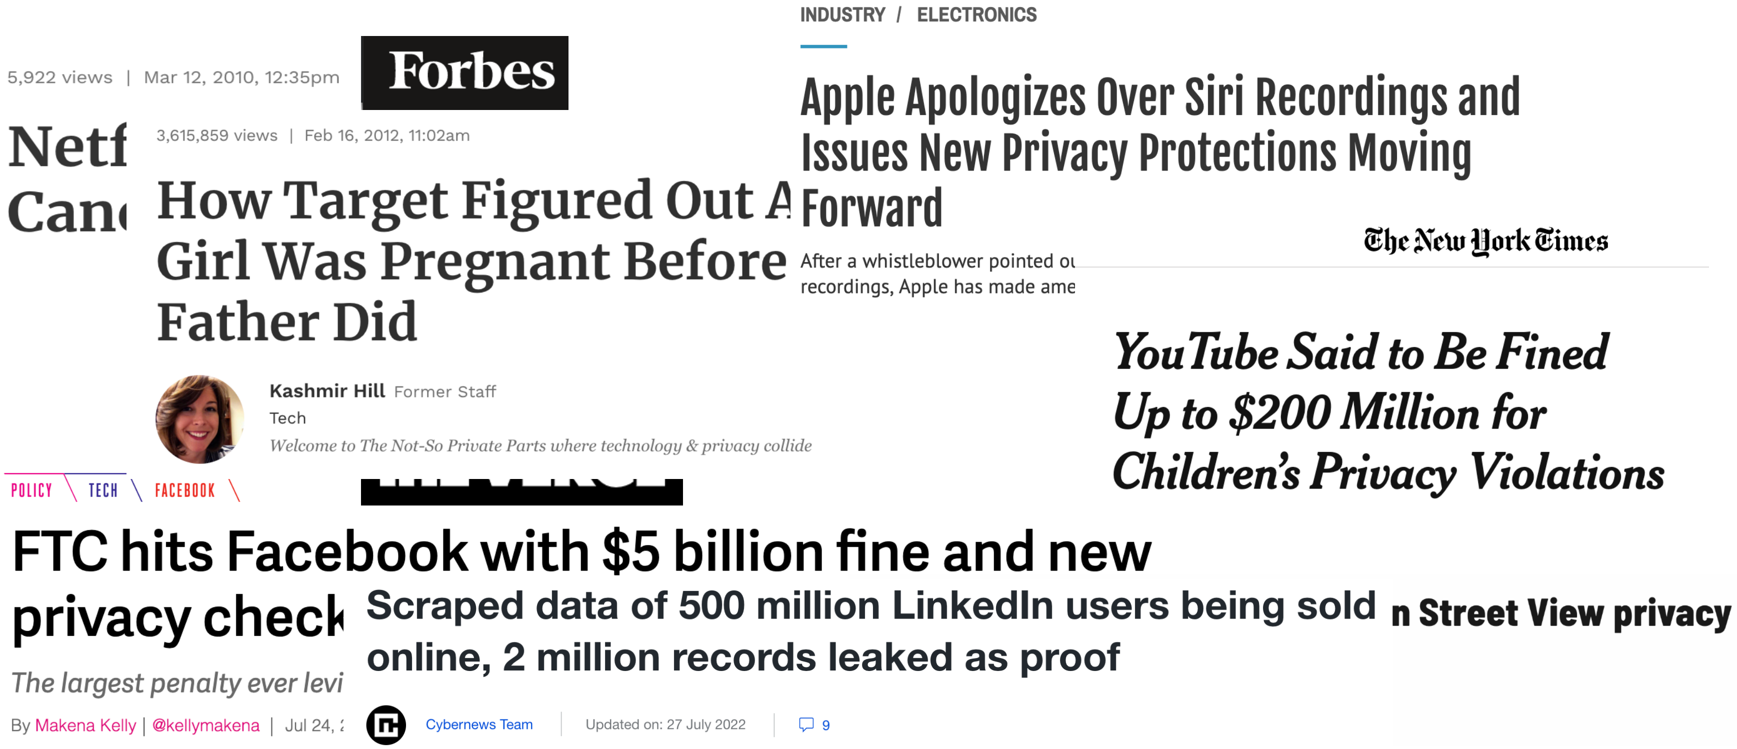

5,922 views | Mar 12, 2010, 12:35pm

Forbes

Netf
Can

3,615,859 views | Feb 16, 2012, 11:02am

How Target Figured Out A
Girl Was Pregnant Before
Father Did

Kashmir Hill Former Staff
Tech
Welcome to The Not-So Private Parts where technology & privacy collide

INDUSTRY / ELECTRONICS

Apple Apologizes Over Siri Recordings and Issues New Privacy Protections Moving Forward

After a whistleblower pointed ou
recordings, Apple has made ame

The New York Times

YouTube Said to Be Fined Up to $200 Million for Children's Privacy Violations

POLICY \ TECH \ FACEBOOK \

FTC hits Facebook with $5 billion fine and new
privacy check

The largest penalty ever levi

By Makena Kelly | @kellymakena | Jul 24, 2

Scraped data of 500 million LinkedIn users being sold
online, 2 million records leaked as proof

Cybernews Team | Updated on: 27 July 2022 | 9

n Street View privacy

Netflix Recommendation

THEY TOOK YOUR DATA.
THEN THEY TOOK CONTROL.
A NETFLIX ORIGINAL DOCUMENTARY
THE GREAT HACK
From the Academy Award® nominated filmmakers behind THE SQUARE
sundance
JULY 24 | NETFLIX

"Our personal data is being used against us in ways that we don't understand and it could bring down severe harm on us."

"So, as individuals, we can limit the flood of data that we're leaking all over the place. But there's no silver bullet. **There's no way to go off the grid. So, you have to understand how your data is affecting your life.** Our dignity as humans is at stake. But the hardest part in all of this is that all of these wreckage sites and crippling division begins with **the manipulation of one individual then the other then the other.** So I can't help but to ask myself: can I be manipulated? Can you?"
– David Carroll

What Is Security?

- Security is about protecting data/systems from external or internal forces by preserving the three "pillars" of security
  - Confidentiality
    - Access to systems or data is limited to authorized parties
  - Integrity
    - When you receive data, you get the "right" data
  - Availability
    - The system (or data) is "available" when you need it

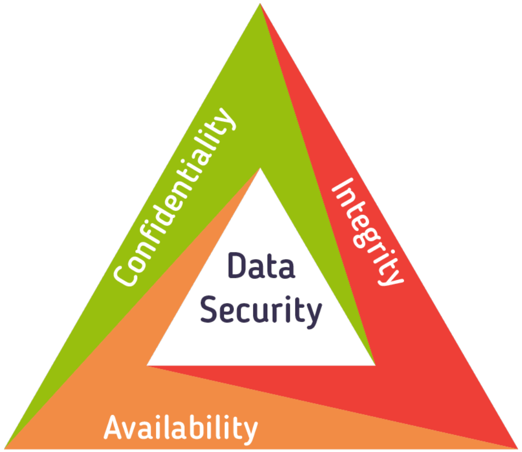

What Is Privacy?

- Before we define privacy, allow me to ask you two questions
  - Who enjoys using all the features provided by tech. companies in return for the data we share?
    - Example: Google maps, Fitbit programs, etc.
  - Who gets paranoid about his privacy and keeps the share data features off all the times?
    - Example: customized ads, shared location, incognito mode, etc.
- What Is Privacy?
  - Let's hear from you!

Security Vs. Privacy

- What Is Privacy?
- The terms are often used interchangeably
  - In this course, hopefully you will be able to distinguish between the two
- Two widely accepted definitions of privacy
  - "The right to be let alone"
  - "The fair use of information“
- Privacy is relative to individuals, families, societies, countries, etc.

COE 426/526: Data Privacy

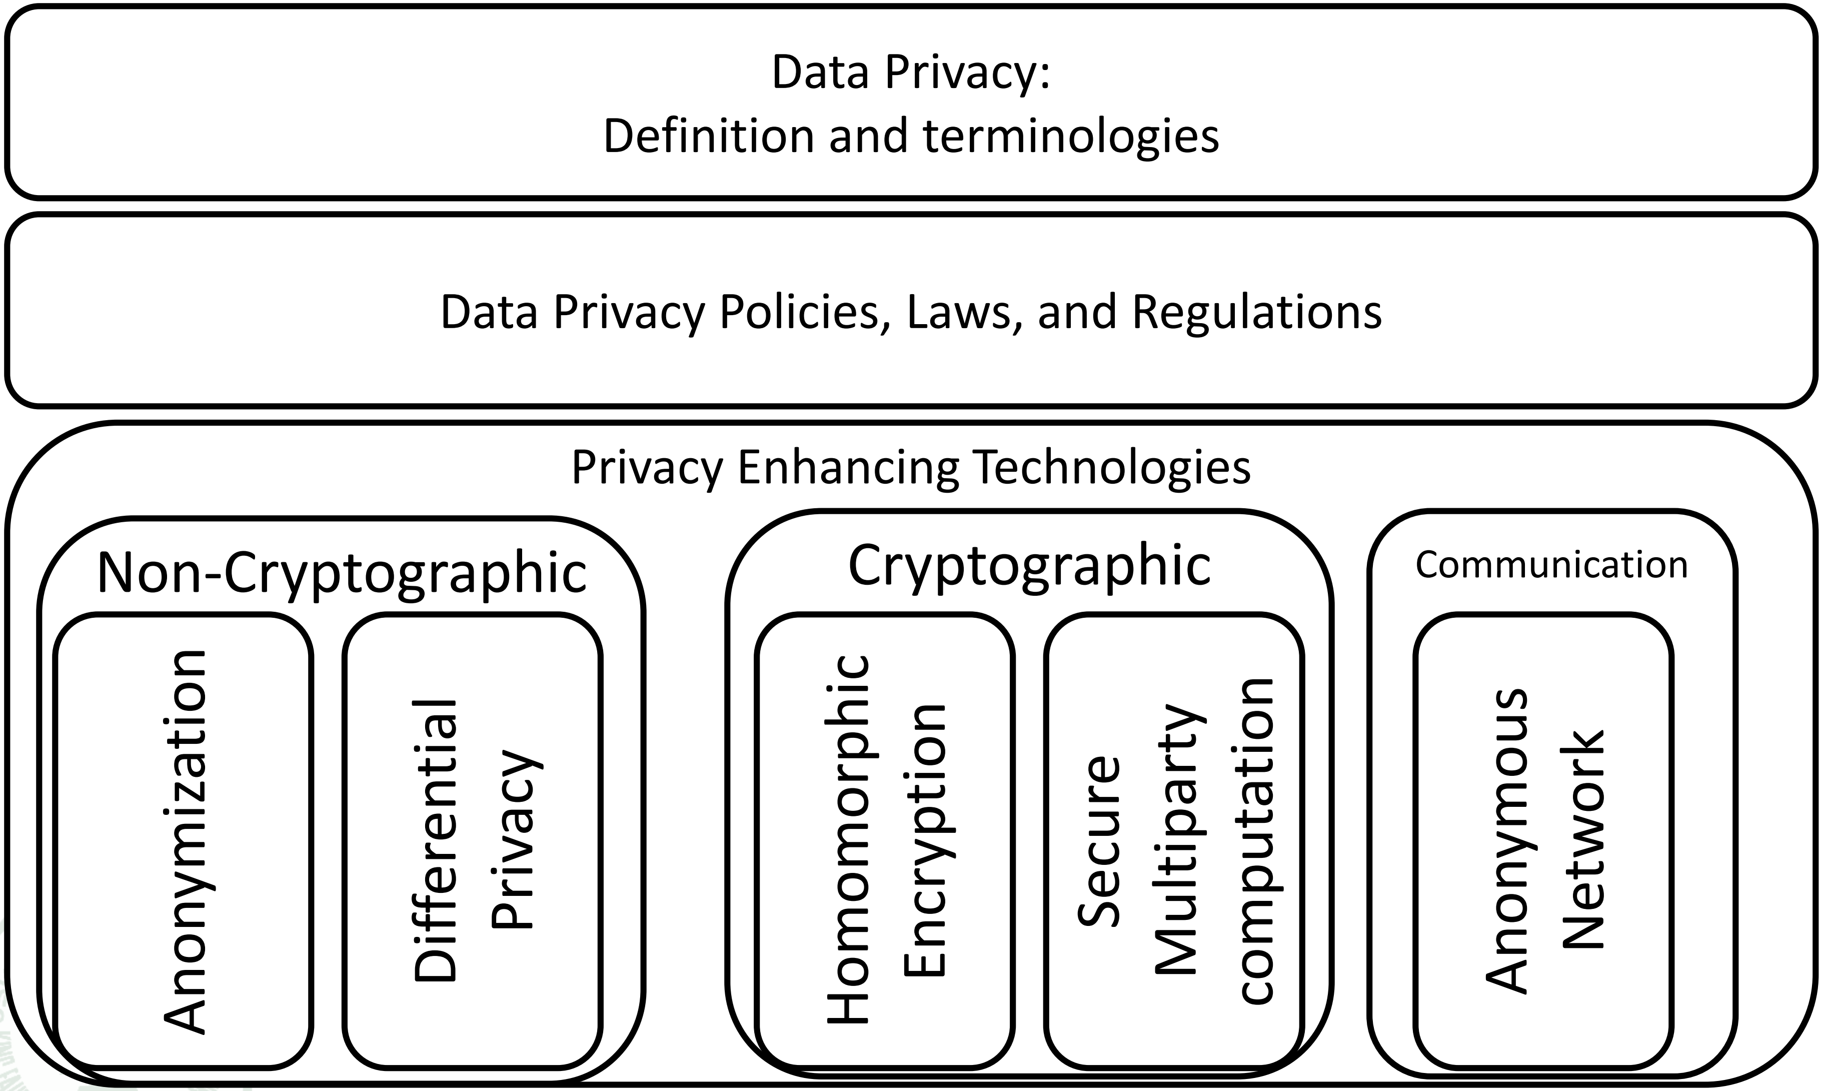

Data Privacy:
Definition and terminologies

Data Privacy Policies, Laws, and Regulations

Privacy Enhancing Technologies

Non-Cryptographic

Anonymization

Differential Privacy

Cryptographic

Homomorphic Encryption

Secure Multiparty computation

Communication

Anonymous Network

Privacy-Preserving Enhancing Technologies

Centre for the Fourth Industrial Revolution

WORLD ECONOMIC FORUM

In collaboration with Frontiers

Top 10 Emerging Technologies of 2024

FLAGSHIP REPORT

JUNE 2024

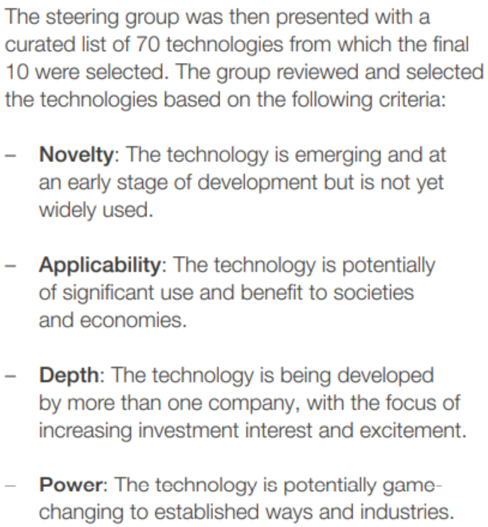

The steering group was then presented with a curated list of 70 technologies from which the final 10 were selected. The group reviewed and selected the technologies based on the following criteria:

- **Novelty**: The technology is emerging and at an early stage of development but is not yet widely used.
- **Applicability**: The technology is potentially of significant use and benefit to societies and economies.
- **Depth**: The technology is being developed by more than one company, with the focus of increasing investment interest and excitement.
- **Power**: The technology is potentially game-changing to established ways and industries.

1- AI for Scientific Discovery

2- Privacy-Enhancing Technologies

3- Reconfigurable Intelligent Surfaces

4- High Altitude Platform Stations

5- Integrated Sensing and Communication

6- Immersive Technology for the Built World

7- Elastocalorics

8- Carbon-capturing Microbes

9- Alternative Livestock Feeds

10- Genomics and Transplants

COE 426/526: Data Privacy

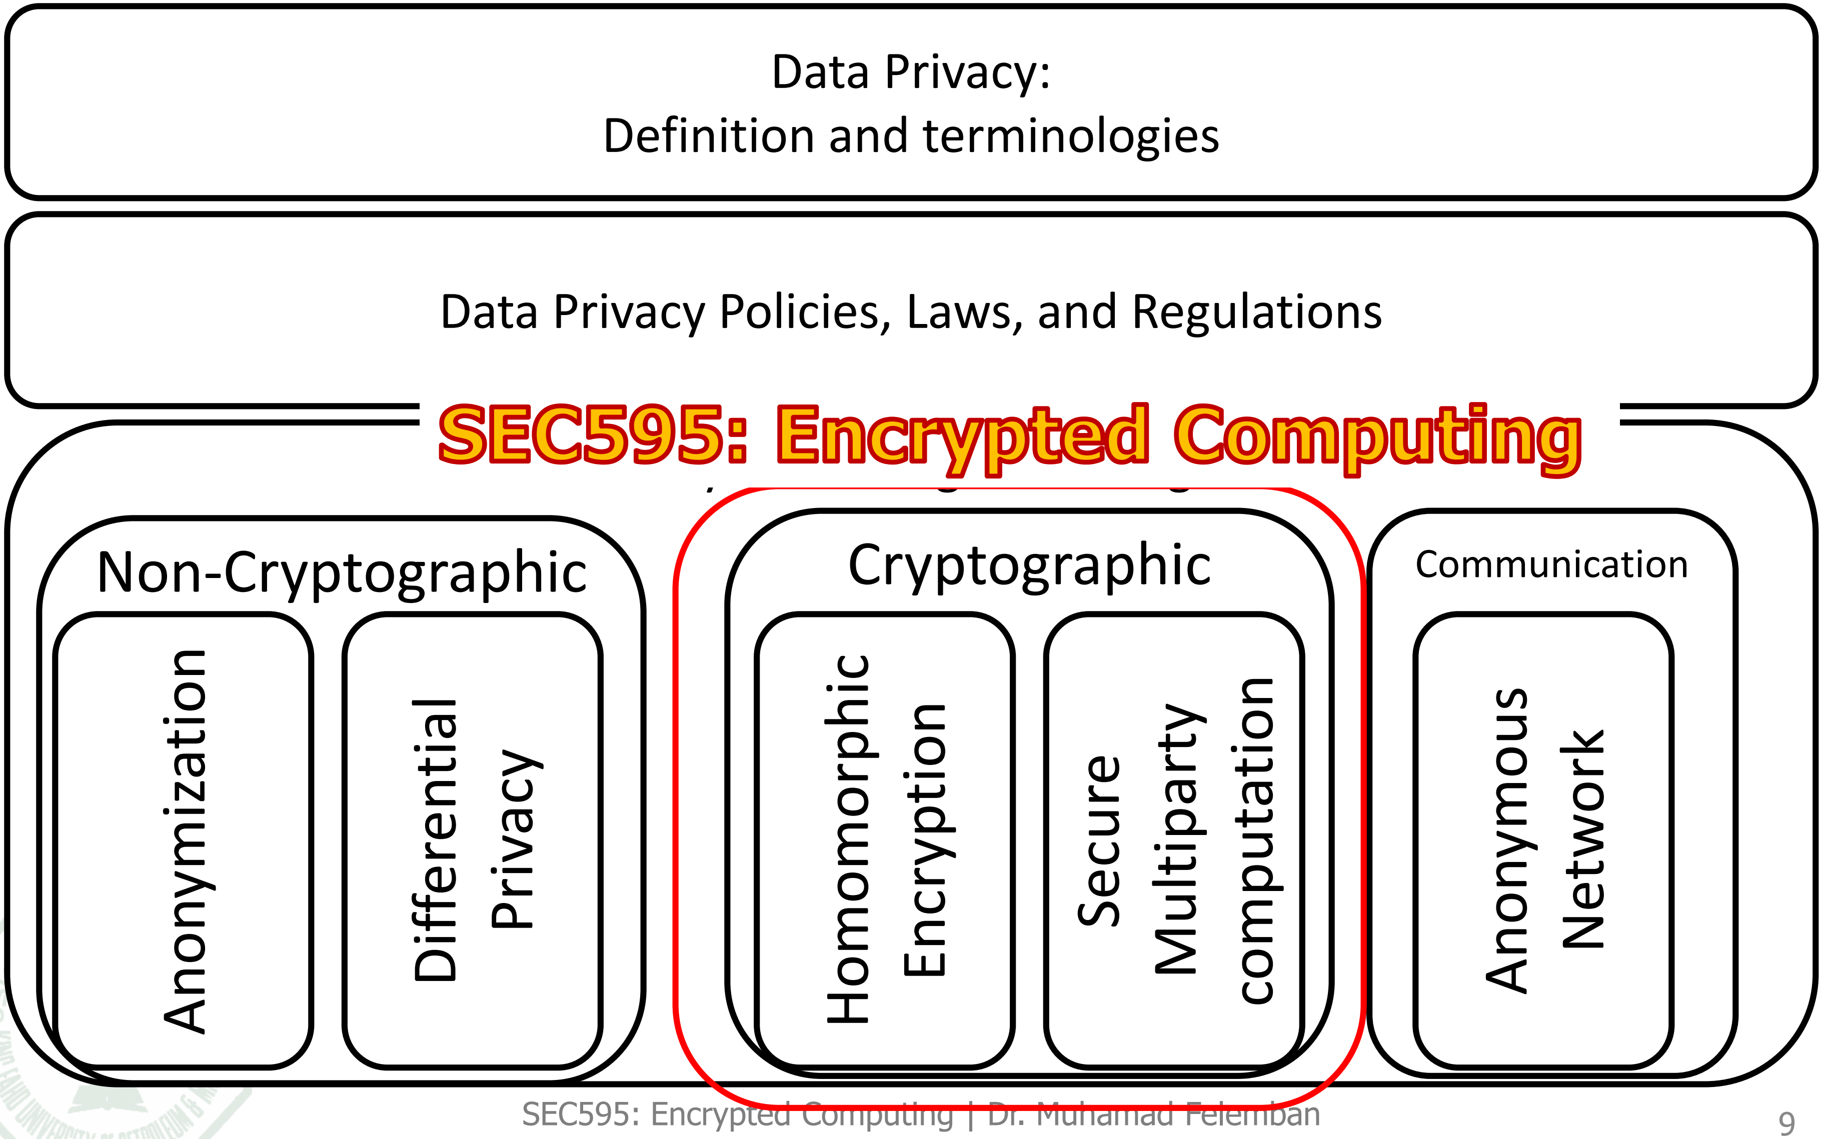

Data Privacy:
Definition and terminologies

Data Privacy Policies, Laws, and Regulations

SEC595: Encrypted Computing

Non-Cryptographic

Anonymization

Differential Privacy

Cryptographic

Homomorphic Encryption

Secure Multiparty computation

Communication

Anonymous Network

Computing on Encrypted Data

Nigel Smart | KU Leuven and Zama

"The ability to compute on encrypted data is fast becoming a practical reality. We discuss the progress in four technologies which enable this:"

- **Trusted Execution Environments**
- **Fully Homomorphic Encryption**
- **Multi-Party Computation**
- **Zero-Knowledge Proofs.**

https://www.computer.org/csdl/magazine/sp/2023/04/10194492/1P4BEqduz9C

Homomorphic Encryption

- Name inspired by ring-homomorphism

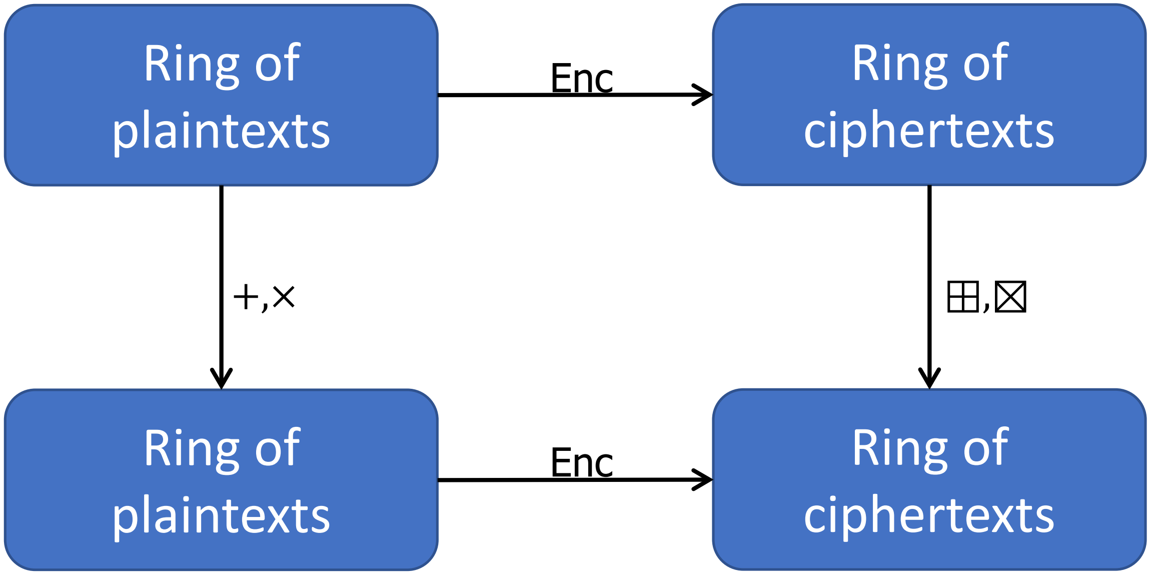
Ring of plaintexts

Enc

Ring of ciphertexts

+,×

⊞,⊠

Ring of plaintexts

Enc

Ring of ciphertexts

A Toy HE Scheme

- Encryption: Double the plaintext. x → 2x
- Decryption: Halve the ciphertext. x → x/2

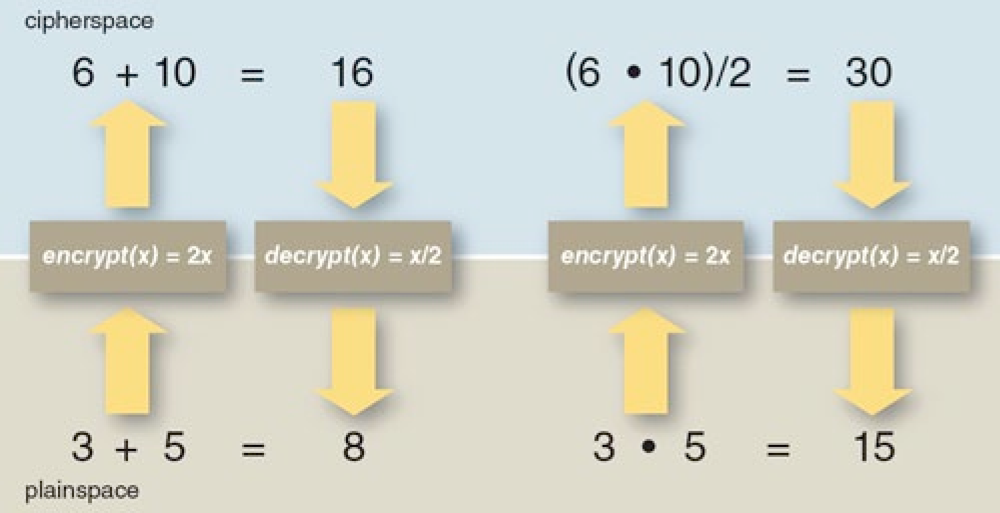

https://www.americanscientist.org/article/alice-and-bob-in-cipherspace

An Analogy

- Alice wants workers to assemble raw materials into jewelry
- But Alice is worried about theft:
  - She wants workers to process raw materials without having access.
- Alice puts raw materials in locked glovebox.
- Workers assemble jewelry inside glovebox, using the gloves.
- Alice unlocks box to get “results”.

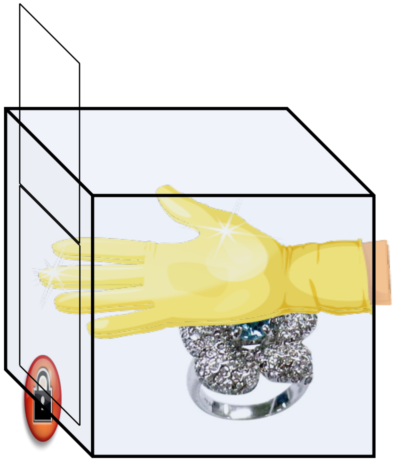

Evolution of FHE Schemes

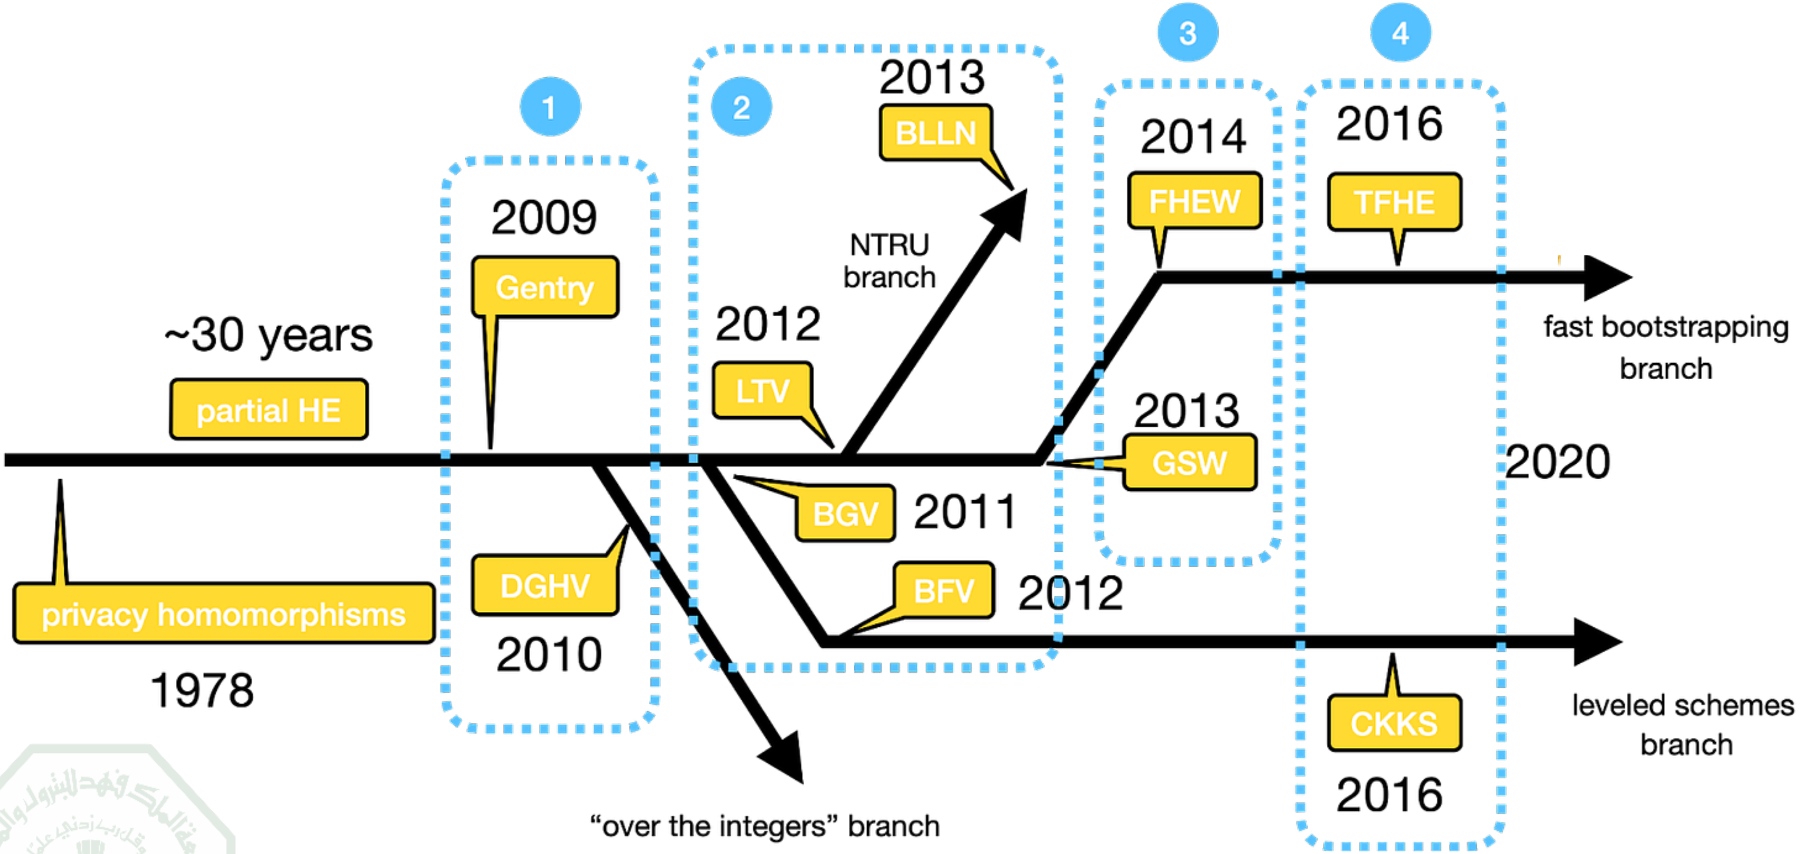

Examples of Homomorphic Encryption

- "Unpadded" RSA is a multiplicative homomorphic encryption scheme
  - Recall $E_{PK}(x) = x^e \, mod \, n$
  - $E_{PK}(x) \cdot E_{PK}(y) = x^e \, mod \, n \cdot y^e \, mod \, n$
    $= (x \cdot y)^e \; mod \; n = E_{PK}(x \cdot \mathrm{y})$
  - How about $E_{PK}(x + y)$?

Examples of Homomorphic Encryption

- Paillier cryptosystem is an additive homomorphic encryption scheme
  - $E_{PK}(x) = g^x \cdot r^n \, mod \, n^2$
  - $E_{PK}(x) \cdot E_{PK}(y) = (g^x r_1^n) \, mod \, n^2 \; (g^y r_1^n) \, mod \, n^2$
    $= g^{x+y}(r_1 r_2)^n \, mod \, n^2 = E_{PK}(x+y)$
  - How about $E_{PK}(x \;\cdot y)$?

Homomorphic Encryption Types

Partially Homomorphic Encryption (PHE)

- Can evaluate functions with very limited operations (+ or *)

Fully Homomorphic Encryption (FHE)

- Arbitrary functions with any operations

Towards FHE

- Genrty09: A bootstrapping technique

  Somewhat homomorphic → Fully homomorphic

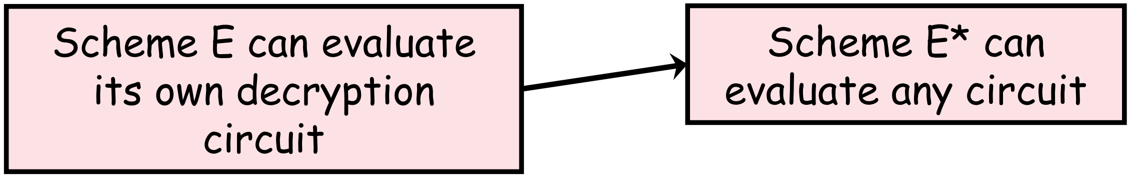

Scheme E can evaluate its own decryption circuit

Scheme E* can evaluate any circuit

- Gentry also described a candidate “bootstrappable” scheme
  - Based on ideal lattices
  - Then extended to integer values

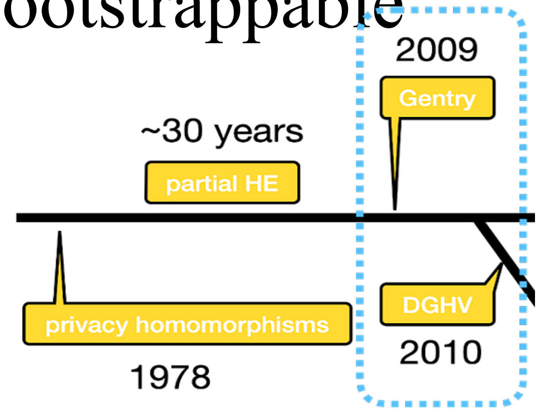

A homomorphic symmetric encryption

- Shared secret key: odd number p
- To encrypt a bit m:
  - Choose at random large q

    2r+m much smaller than p
  - Output $c = pq + 2r + m$
    - Ciphertext is close to a multiple of p
    - m = LSB of distance to nearest multiple of p
- To decrypt c:
  - Output $m = (c \bmod p) \bmod 2$

$$
\begin{aligned}
m &= c - p \cdot [c/p] \bmod 2 \\
&= c - [c/p] \bmod 2 \\
&= \mathrm{LSB}(c)\ \mathrm{XOR}\ \mathrm{LSB}([c/p])
\end{aligned}
$$

Why is this homomorphic?

- $c_1=q_1p+2r_1+m_1, \quad c_2=q_2p+2r_2+m_2$

Distance to nearest multiple of p

- $c_1+c_2 = (q_1+q_2)p + 2(r_1+r_2) + (m_1+m_2)$
  - $2(r_1+r_2)+(m_1+m_2)$ still much smaller than p

  ➔$c_1+c_2 \bmod p = 2(r_1+r_2) + (m_1+m_2)$

- $c_1 \text{ x } c_2 = (c_1q_2+q_1c_2-q_1q_2)p + 2(2r_1r_2+r_1m_2+m_1r_2) + m_1m_2$
  - $2(2r_1r_2+\ldots)$ still much smaller than p

  ➔$c_1\text{x}c_2 \bmod p = 2(2r_1r_2+\ldots) + m_1m_2$

How homomorphic is this?

- Can keep adding and multiplying until the “noise term” grows larger than q/2
  - Noise doubles on addition, squares on multiplication

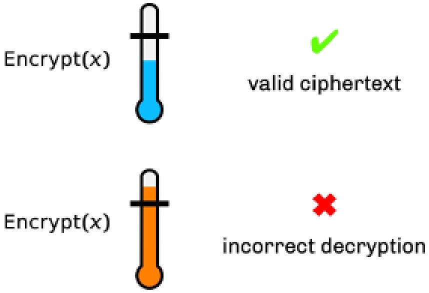

- We choose r ~ 2^n, p ~ 2^{n^2} (and q ~ 2^{n^5})
  - Can compute polynomials of degree ~n before the noise grows too large

Homomorphic Public-Key Encryption

- Secret key is an odd p as before
- Public key is many “encryptions of 0”
  - $[x_i = q_i p + 2r_i]_{x0}$ for i=1,2,…,n
- Enc$_{pk}$(m): Choose a random subset $S \subseteq \{1,2,\ldots,\tau\}$ and a random integer r and output

$$c \leftarrow \left[m + 2r + 2\sum_{i \in S} x_i \right]_{x_0}$$

- Dec$_{sk}$(c) = (c mod p) mod 2
- Eval as before

Bootstrapping

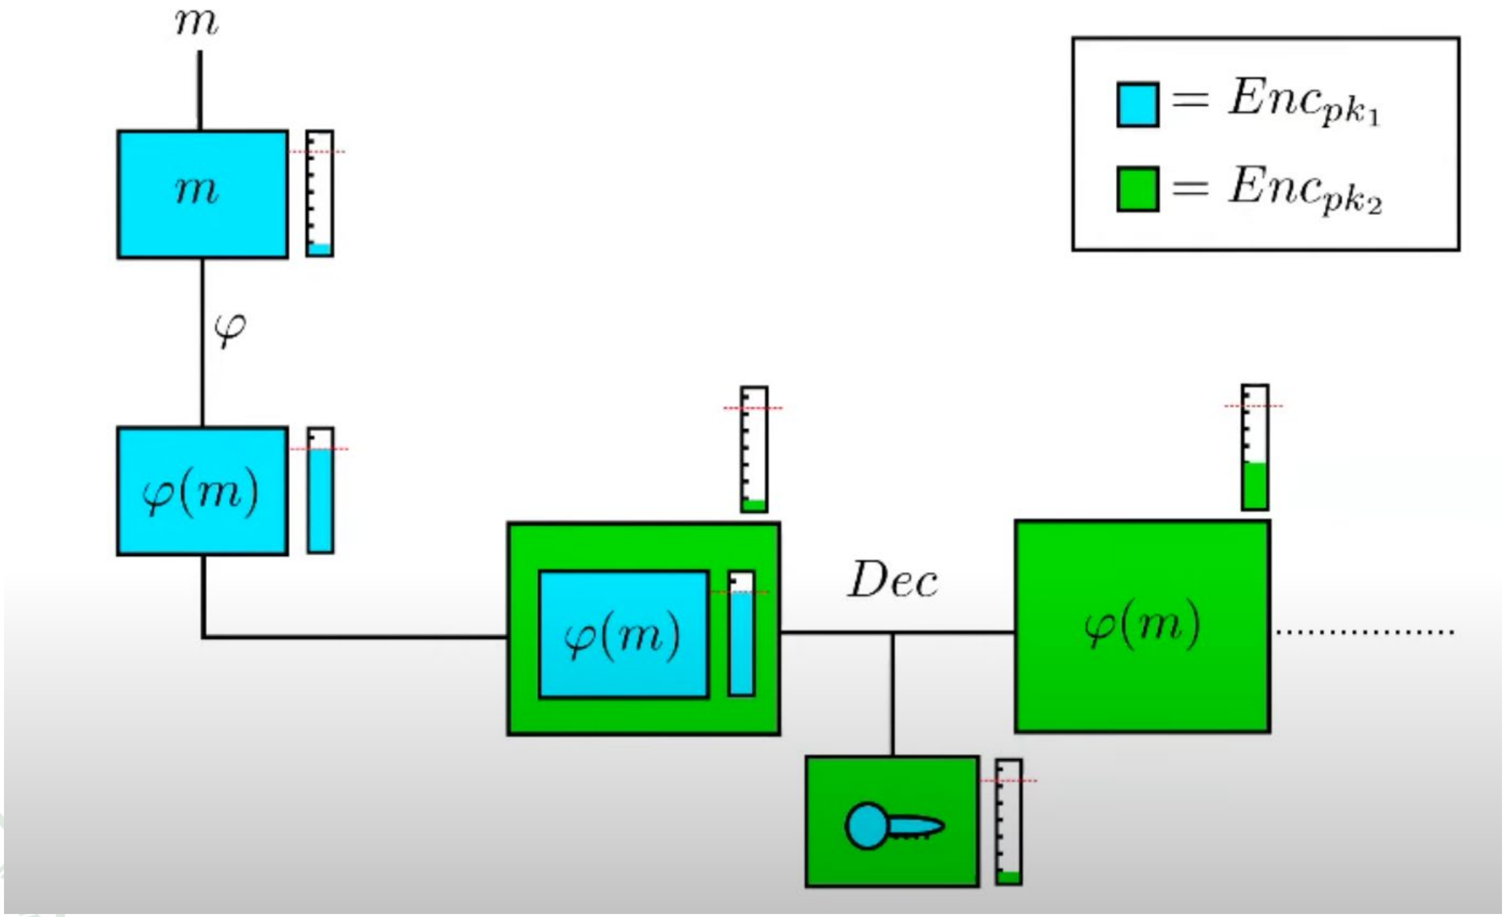

Use case

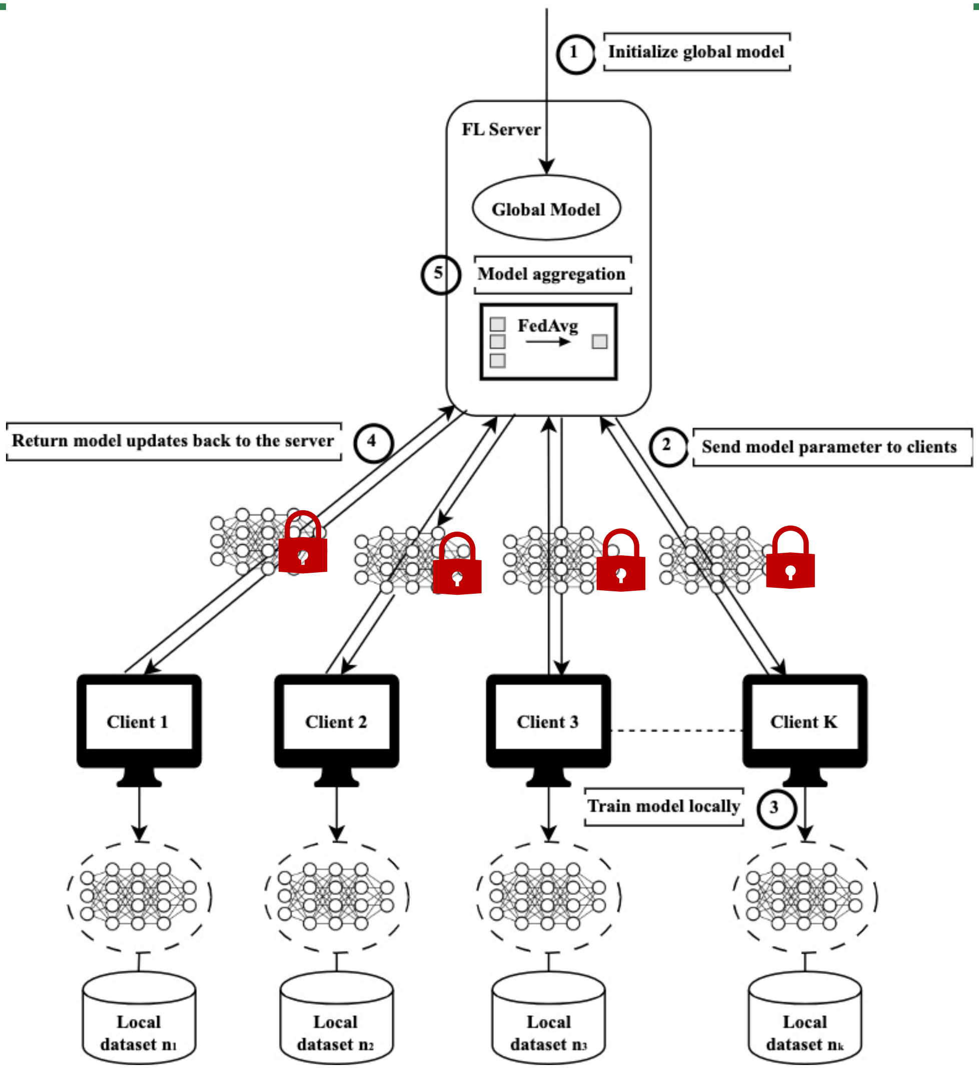
1
Initialize global model
FL Server
Global Model
5
Model aggregation
FedAvg
Return model updates back to the server
4
2
Send model parameter to clients
Client 1
Client 2
Client 3
Client K
Train model locally
3
Local dataset n1
Local dataset n2
Local dataset n3
Local dataset nk

]

Speeding Up Bootstrapping

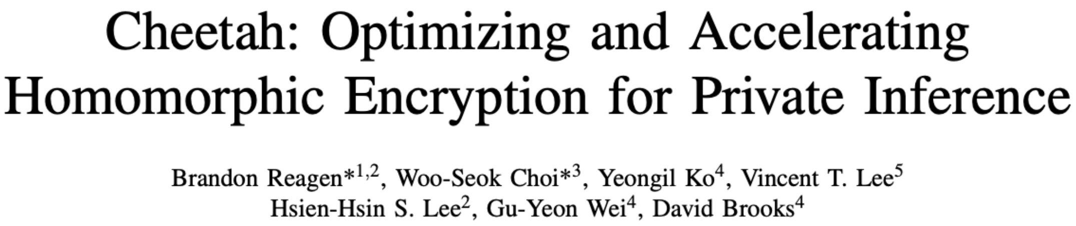

Cheetah: Optimizing and Accelerating Homomorphic Encryption for Private Inference

Brandon Reagen*[1,2], Woo-Seok Choi*[3], Yeongil Ko[4], Vincent T. Lee[5]
Hsien-Hsin S. Lee[2], Gu-Yeon Wei[4], David Brooks[4]

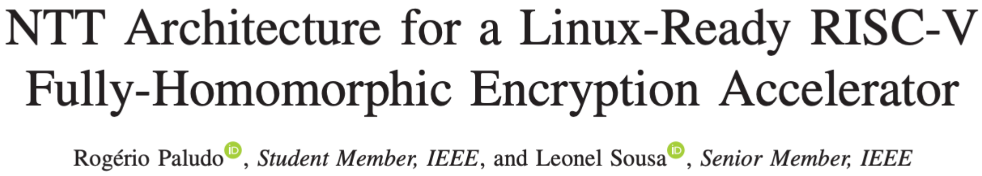

NTT Architecture for a Linux-Ready RISC-V Fully-Homomorphic Encryption Accelerator

Rogério Paludo, *Student Member, IEEE*, and Leonel Sousa, *Senior Member, IEEE*

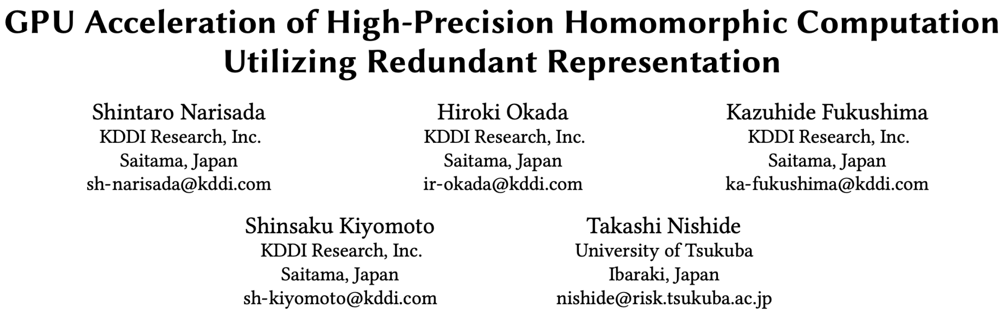

GPU Acceleration of High-Precision Homomorphic Computation Utilizing Redundant Representation

Shintaro Narisada
KDDI Research, Inc.
Saitama, Japan
sh-narisada@kddi.com

Hiroki Okada
KDDI Research, Inc.
Saitama, Japan
ir-okada@kddi.com

Kazuhide Fukushima
KDDI Research, Inc.
Saitama, Japan
ka-fukushima@kddi.com

Shinsaku Kiyomoto
KDDI Research, Inc.
Saitama, Japan
sh-kiyomoto@kddi.com

Takashi Nishide
University of Tsukuba
Ibaraki, Japan
nishide@risk.tsukuba.ac.jp

FHE Application Space

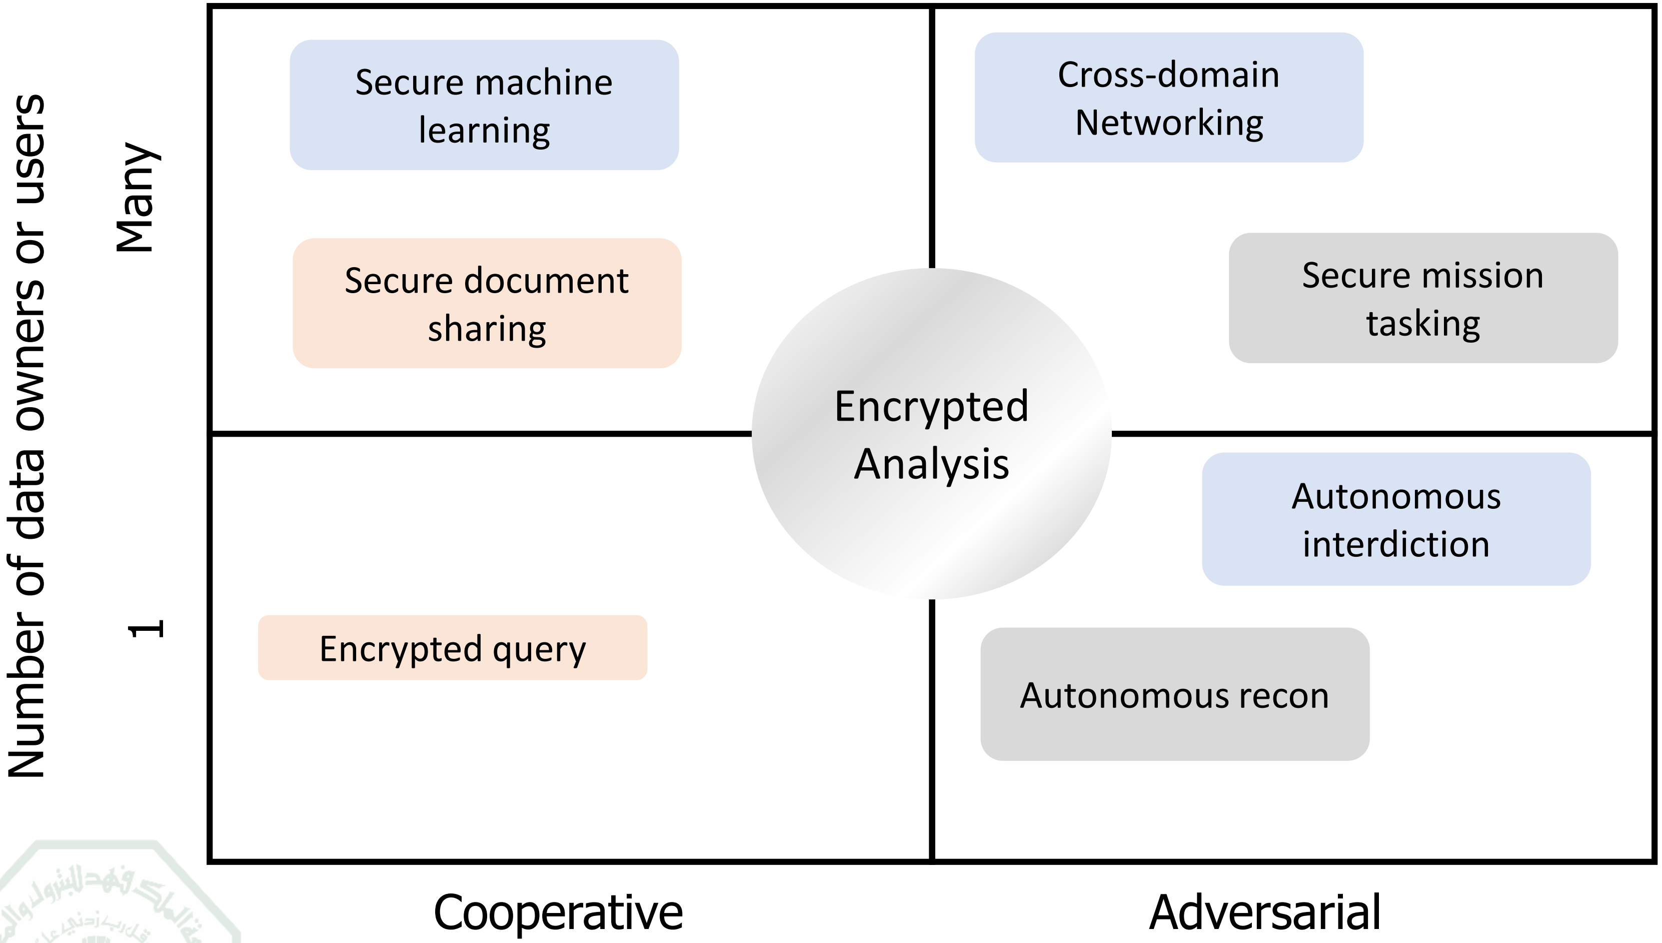

*Source: DARPA DPRIVE presentati

Computing on Encrypted Data

Nigel Smart | KU Leuven and Zama

"The ability to compute on encrypted data is fast becoming a practical reality. We discuss the progress in four technologies which enable this:"

- Trusted Execution Environments
- **Fully Homomorphic Encryption**
- Multi-Party Computation
- Zero-Knowledge Proofs.

https://www.computer.org/csdl/magazine/sp/2023/04/10194492/1P4BEqduz9C

Secure Multiparty Computation (SMPC or MPC)

- Secure multiparty computation is a subfield of encrypted computation
- MPC involves multiple parties (who do not trust each other) agreeing to compute a joint function of their inputs but only if the data is encrypted
- Parties wish to preserve
  - **Privacy:** parties can’t observe each others’ inputs
  - **Correctness:** the function gives the correct results
- Privacy must be preserved in the face of adversarial behavior by
  - One (or some) of the participants, or
  - An external party

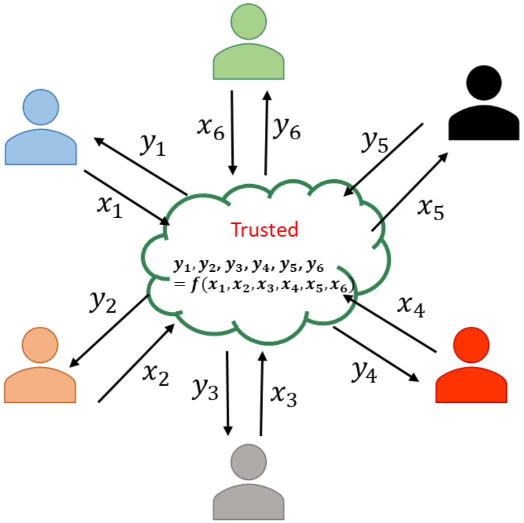

A Couple of Observations

- In all cases, we are dealing with distributed multi-party protocols
  - A protocol describes how parties are supposed to exchange messages on the network
- All these tasks can be easily computed by a trusted third party
- The goal of secure multi-party computation is to achieve the same result **without involving a trusted third party**

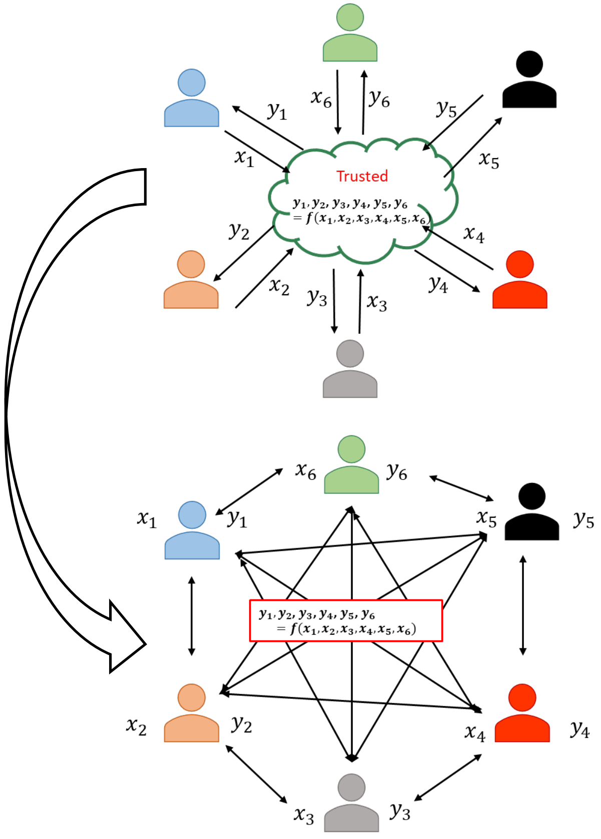

y_1
x_6
y_6
y_5
x_1
x_5
Trusted
$y_1, y_2, y_3, y_4, y_5, y_6 = f(x_1, x_2, x_3, x_4, x_5, x_6)$
y_2
x_4
x_2
y_3
x_3
y_4
x_6
y_6
x_1
y_1
x_5
y_5
$y_1, y_2, y_3, y_4, y_5, y_6 = f(x_1, x_2, x_3, x_4, x_5, x_6)$
x_2
y_2
x_4
y_4
x_3
y_3

Yao’s Millionaire Problem

- Yao’s Millionaire Problem is another example o

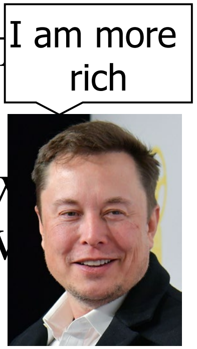

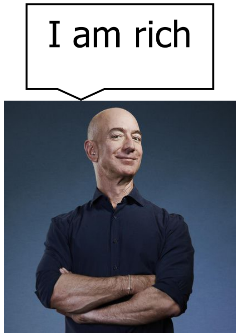

- Two millionaires, Alice and Bob want to know w of them is richer without revealing their actual w
- This problem is analogous to a more general problem where there are two numbers *A* and *B* and the goal is to solve the inequality without revealing the actual values of *A* and *B*
- We will see a solution to Yao’s Millionaire problem using encrypted computing

Garbled Circuit

- A “**garbled**” version of a Boolean circuit
  - Also known as **encrypted** circuit, or **scrambled** circuit

- Overview

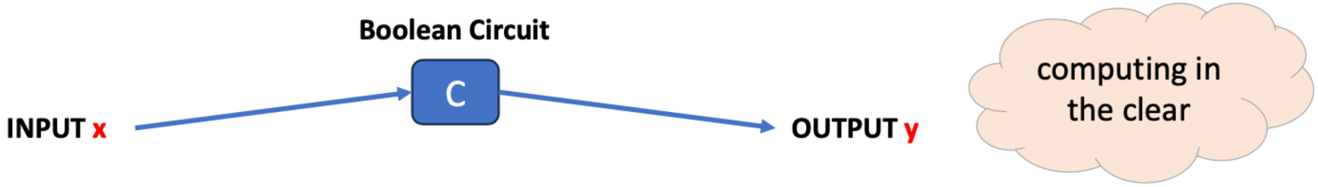

Bellare, Mihir, Viet Tung Hoang, and Phillip Rogaway. "Foundations of garbled circuits." *Proceedings of the 2012 ACM conference on Computer and communications security*. 2012.

Garbled Circuit

- A “**garbled**” version of a Boolean circuit
  - Also known as **encrypted** circuit, and **scrambled** circuit

- Overview

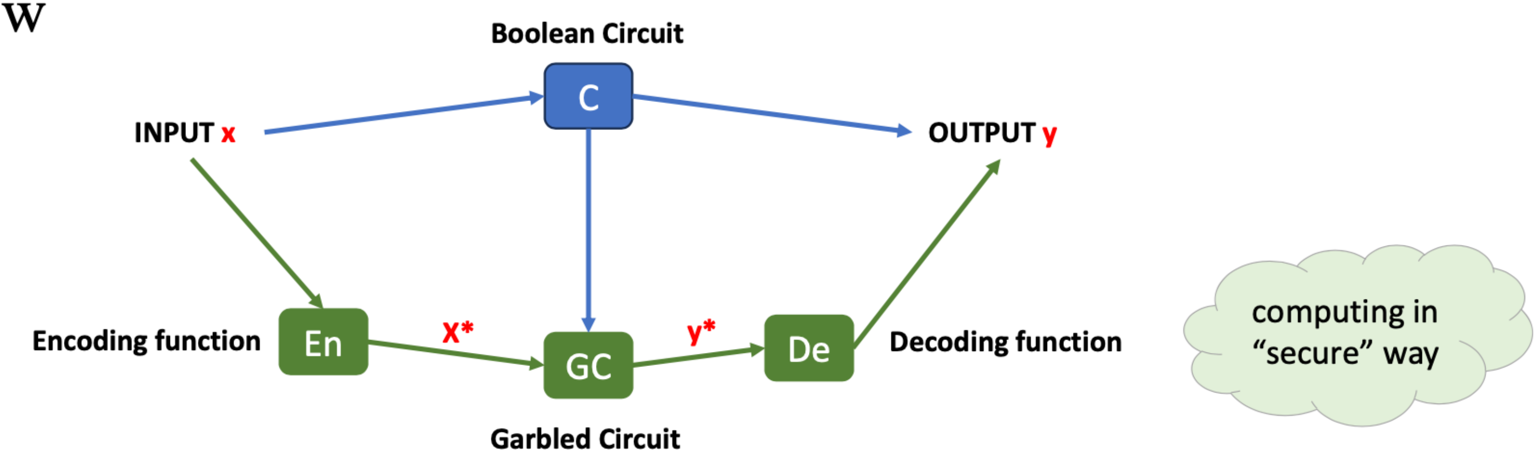

Bellare, Mihir, Viet Tung Hoang, and Phillip Rogaway. "Foundations of garbled circuits." *Proceedings of the 2012 ACM conference on Computer and communications security*. 2012.

Use case: Private Set Intersection

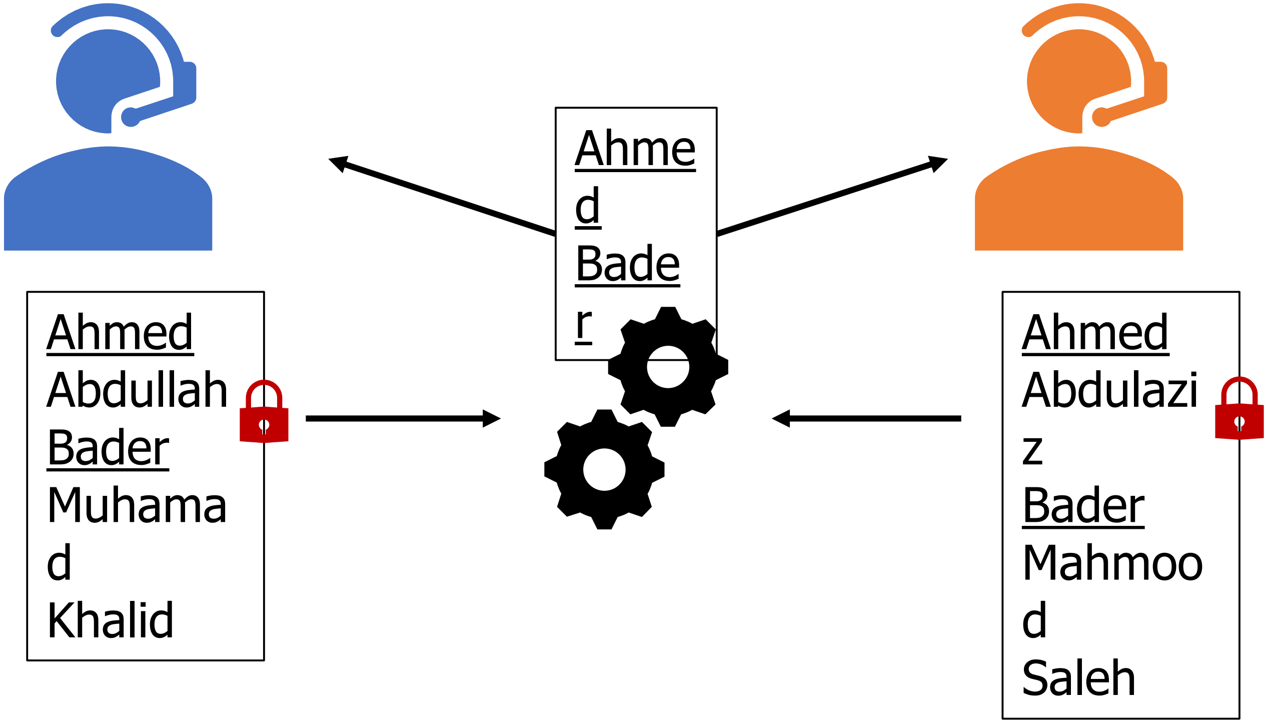

Ahmed
Abdullah
Bader
Muhamad
Khalid

Ahmed
Bader

Ahmed
Abdulaziz
Bader
Mahmood
Saleh

Practical Multi-Party Private Set Intersection Protocols

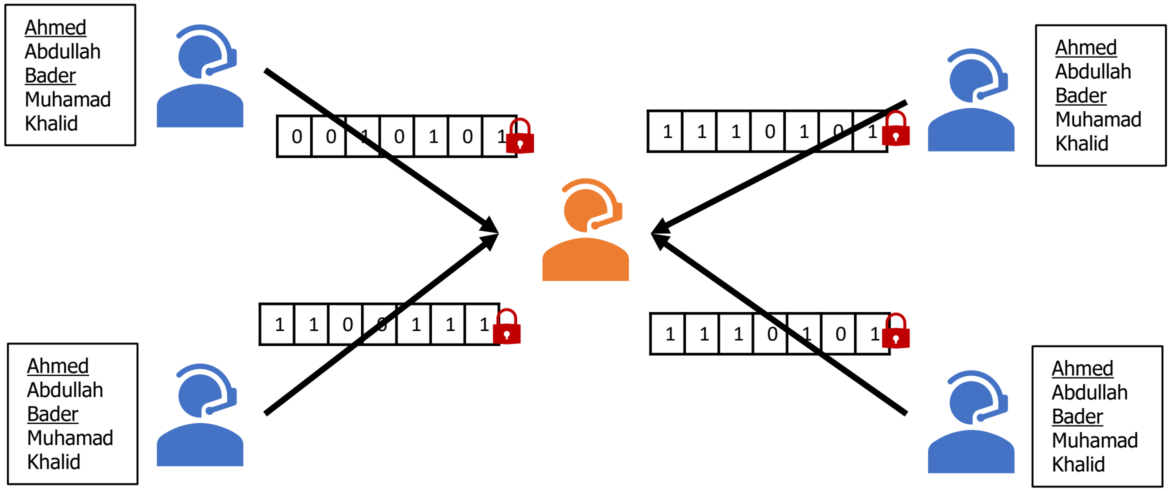

Ahmed
Abdullah
Bader
Muhamad
Khalid

0	0	1	0	1	0	1

Ahmed
Abdullah
Bader
Muhamad
Khalid

1	1	1	0	1	0	1

1	1	0	0	1	1	1

Ahmed
Abdullah
Bader
Muhamad
Khalid

1	1	1	0	1	0	1

Ahmed
Abdullah
Bader
Muhamad
Khalid

Bay, Aslı, et al. "Practical multi-party private set intersection protocols." *IEEE Transactions on Information Forensics and Security* 17 (2021): 1-15

Practical Multi-Party Private Set Intersection Protocols

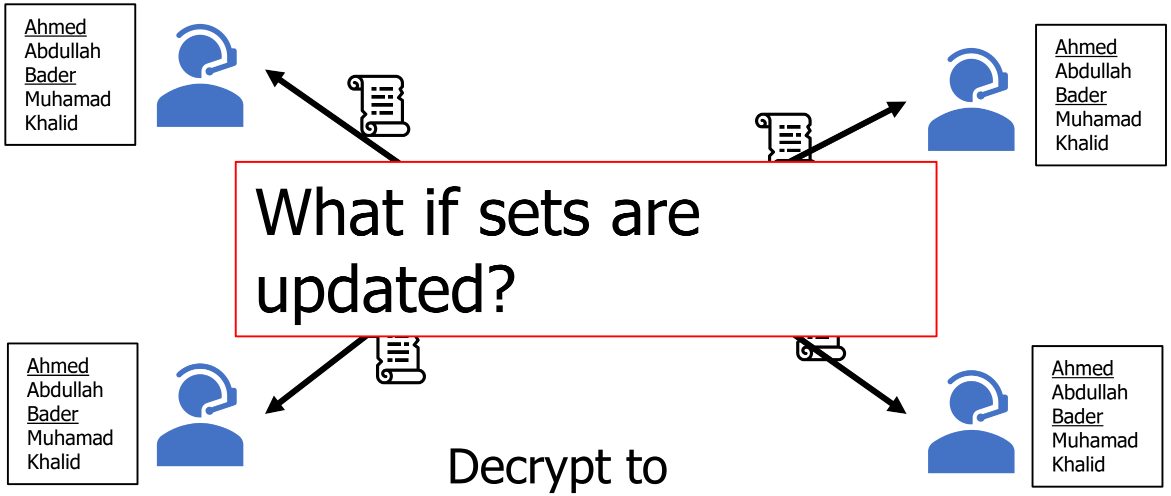

What if sets are updated?

Decrypt to find set intersection

Bay, Aslı, et al. "Practical multi-party private set intersection protocols." *IEEE Transactions on Information Forensics and Security* 17 (2021): 1-15

"**Updatable**" Multi-Party Private Set Intersection Protocols

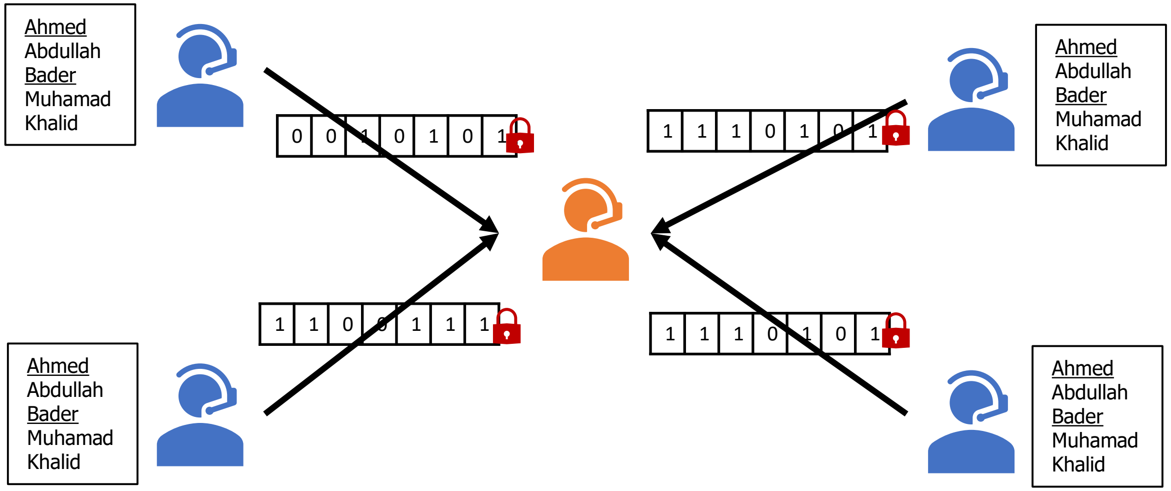

Ahmed
Abdullah
Bader
Muhamad
Khalid

0	0	1	0	1	0	1

Ahmed
Abdullah
Bader
Muhamad
Khalid

1	1	1	0	1	0	1

Ahmed
Abdullah
Bader
Muhamad
Khalid

1	1	0	0	1	1	1

Ahmed
Abdullah
Bader
Muhamad
Khalid

1	1	1	0	1	0	1

Exciting News!!

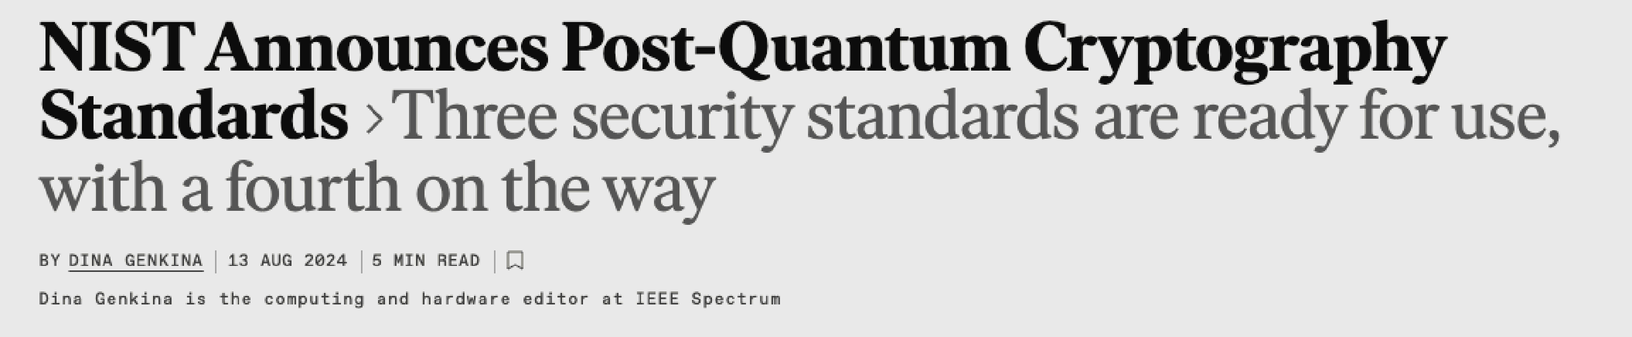

NIST Announces Post-Quantum Cryptography Standards › Three security standards are ready for use, with a fourth on the way

BY DINA GENKINA | 13 AUG 2024 | 5 MIN READ |

Dina Genkina is the computing and hardware editor at IEEE Spectrum

Course Objectives

- The objective of this course is to

i. Introduce the students to the emerging field of encrypted computing.

ii. Expose students to the state-of-the-art algorithms in homomorphic encryption and secure multiparty computation

Course Learning Outcomes

After taking this course, you will be able to

1. Explain various encrypted computing techniques and algorithms.
2. Identify the advantages and challenges of different encrypted computing algorithms.
3. Apply appropriate encrypted computing techniques based on the application's requirements.
4. Design new encrypted computing techniques and protocols.

Logistics

- Lectures: Sunday & Tuesday, 6:45-8:00 PM
- Office hours: UT 8:00-9:00 PM
  - Office# 22-214
  - Teams (or send me on Teams and I will try to accommodate your Qs online)
- Web page:
  - Blackboard page

Logistics

- Evaluation

Paper Presentation	10 %
Term Project	30%
Major Exam	30%
Final Exam	30%

Logistics

- Course Project:
  - Teams of 2-3 (if you want to work alone, please let me know ahead of time)
  - Project proposal: 0.5-1 page (TBA)
  - Deliverable Part I: 1-2 page Progress report (TBA)
  - Deliverable Part II: Written report: 4-8 pages (TBA)
  - In-class presentations (Week 15)