Privacy-Preserving Access Control and Computations of Encrypted Data in the Cloud

Robert Deng
AXA Chair Professor of Cybersecurity
Director, Secure Mobile Center
Singapore Management University

13 December 2017
SpaCCS 2017

Cybersecurity Faculty Members

Prof Robert DENG: Applied cryptography, cloud computing security, data & multimedia security, security protocols, IoT and network security

Associate Prof DING Xuhua: Cryptography, data security & privacy, network security, system security, trusted computing

Associate Prof GAO Debin: Software security, system security, intrusion detection

Associate Prof LI Yingjiu: Mobile app and platform security, IoT security & privacy, data security & privacy, data & media rights management, user authentication

Prof PANG Hwee Hwa: Privacy in search engine and DBMS, query answer authentication and data integrity, indexing and query processing for spatial & high-dimensional data, distributed systems

Assistant Prof WANG Qiuhong: Policy analytics of information security, economics of information systems

Projects at Secure Mobile Centre @ SMU

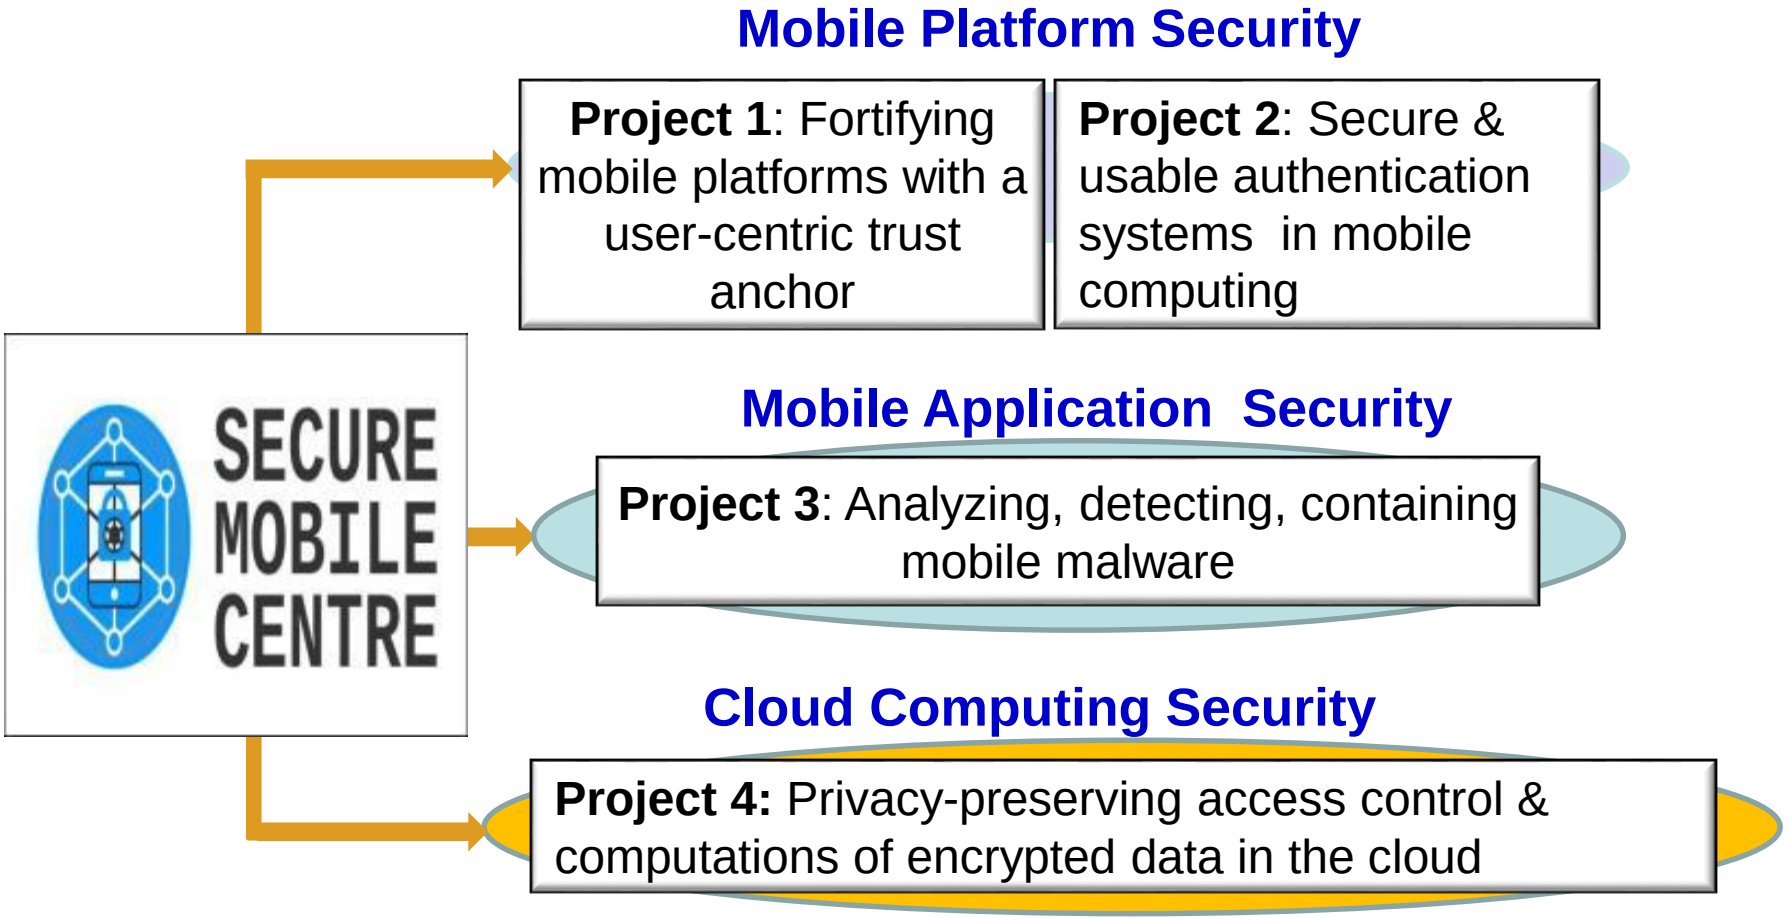

Access Control of Encrypted Data in the Cloud

Assumption and Objective

- Assumption

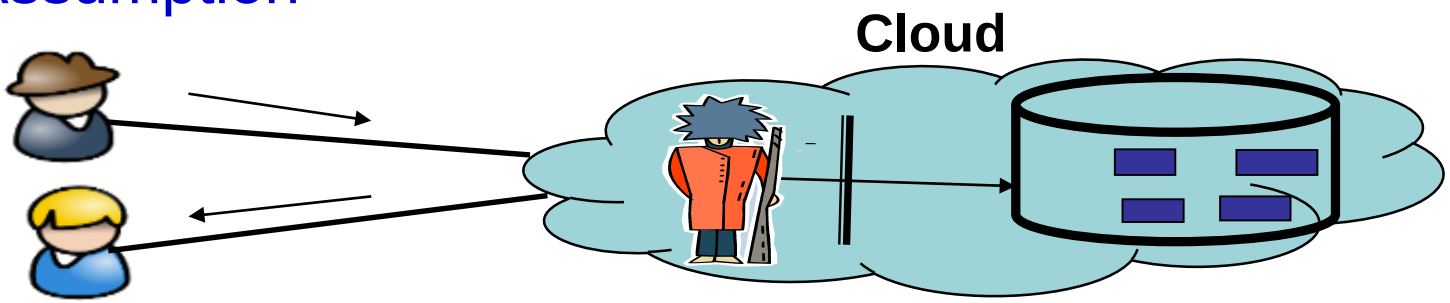

Honest but curious
- Not trusted to keep data Confidential
- Not trusted to enforce access control correctly

- Objective
  - End-to-end data security and privacy
  - Scalable, efficient and flexible solution

- Approach
  - Attribute-based encryption (ABE) [Sahai & Waters'05]

System Architecture Based on CP-ABE

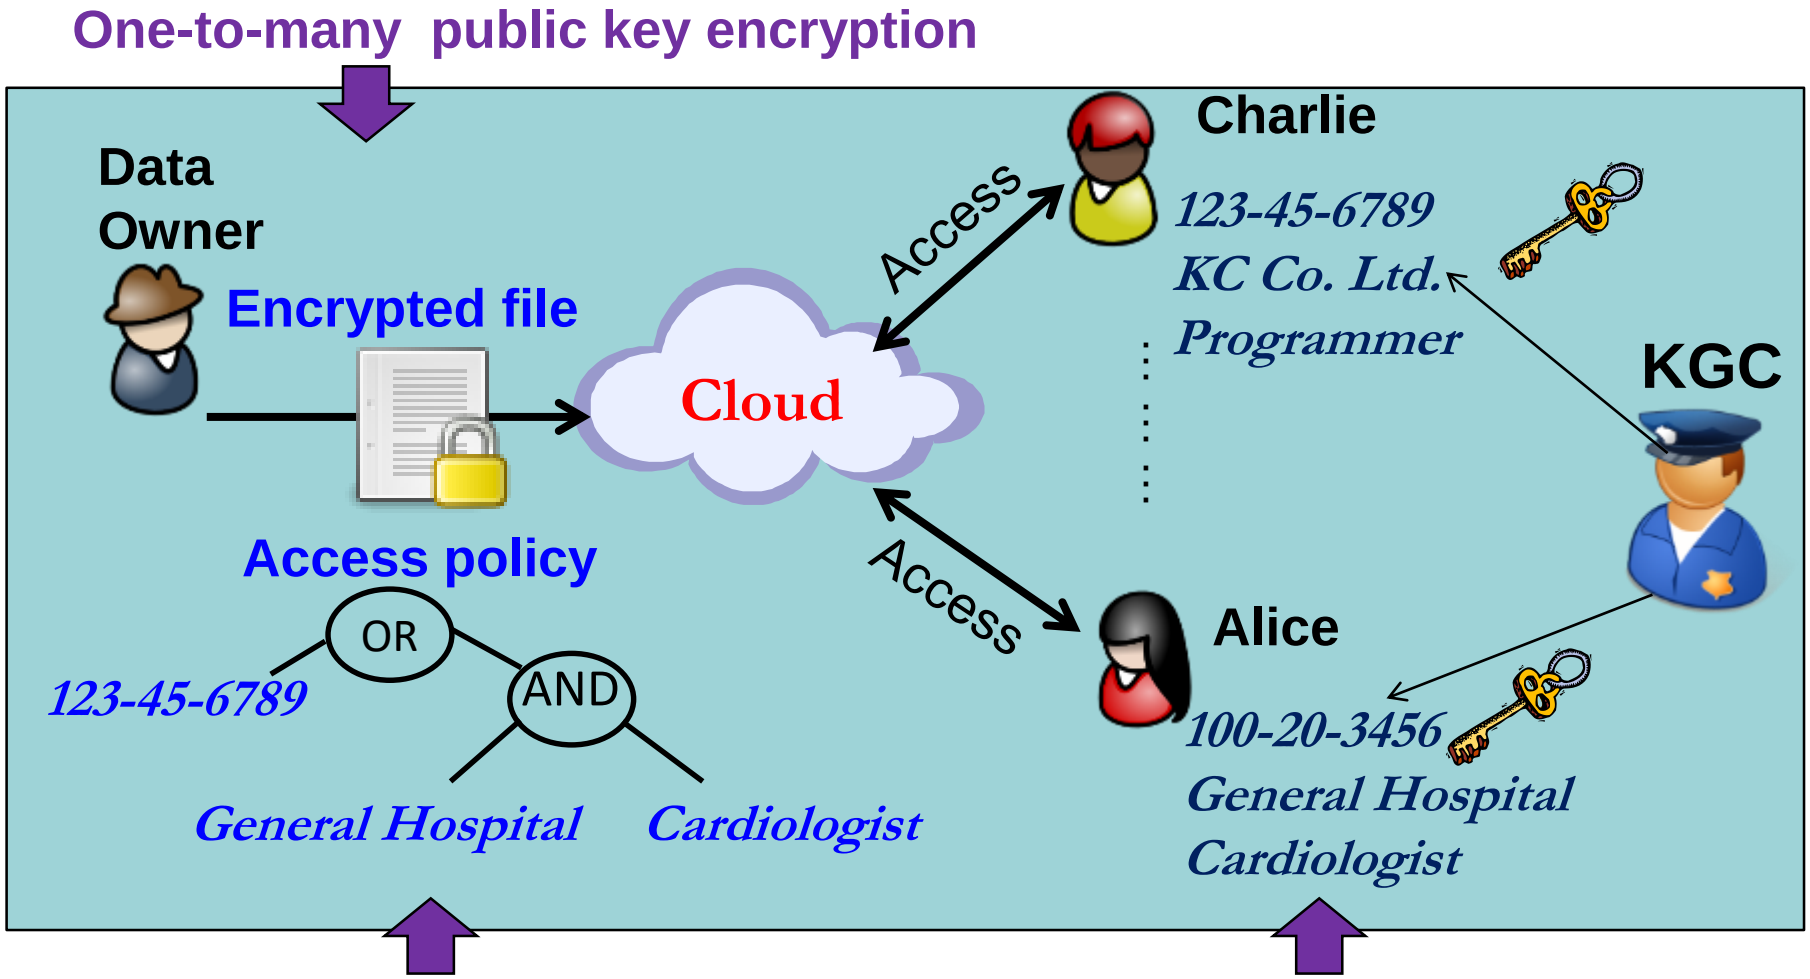

Our Contributions

- Verifiable outsourced decryption of ABE [TIFS 2013, TIFS 2015]
- Efficient user and attribute revocation [EOSRICS'15 & '16, SecureComm'17]
- CP-ABE with partial hidden access policy [AsiaCCS'12, ProvSec'16]
- Deduplication on encrypted data [TBD 2016], Best Paper Award
- Attribute-based secure messaging system in the cloud [SG-CRC'17]
- Lightweight sharable and traceable secure mobile health system [TDSC, accepted]

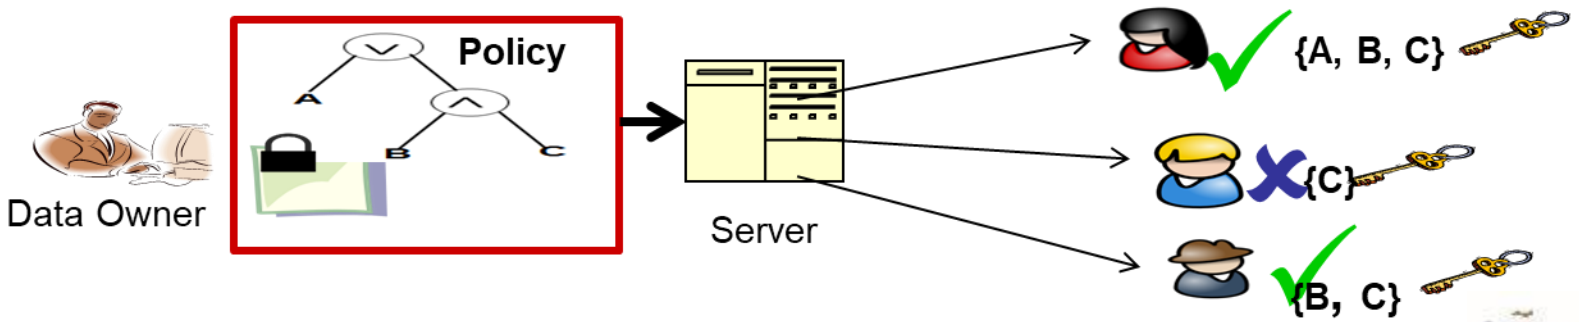

User Revocation

- User leaves the system, or user's private key is compromised

- User revocation in ABE is critical because of its one-to-many encryption nature

- However, efficient user revocation has been a very challenging problem in ABE

Existing Solutions

[Boneh & Franklin CRYPTO'01] **Basic approach**

- Associate encrypted message with a time stamp: CT(t)
- KGC periodically updates users keys over private channels

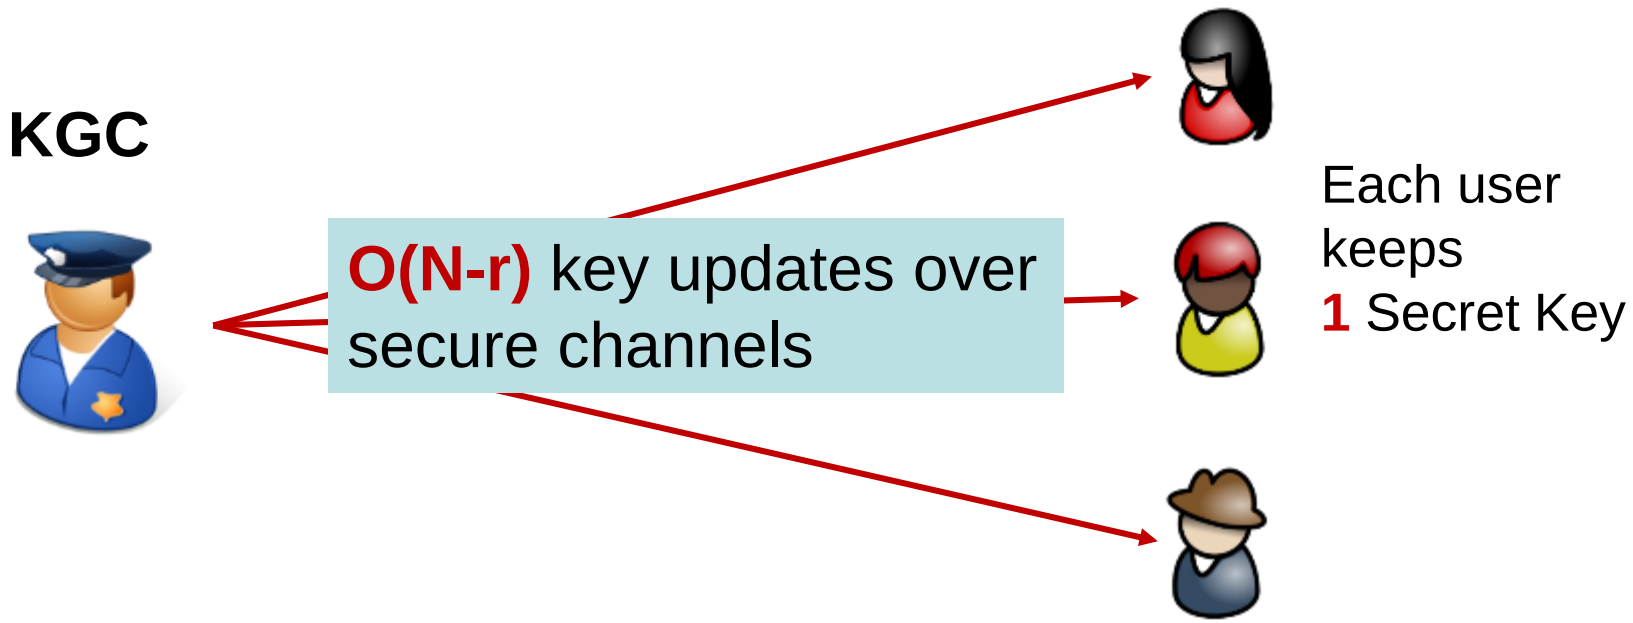

Existing Solutions

[Boldyreva, Goyal, Kumar CCS'08] [Seo & Emura PKC'13]

Tree-based approach

- Associate encrypted message with a time stamp: CT(t)
- KGC periodically broadcasts key updates to users over public channels
- Non-revoked user computes decryption key for current time t using her/his long-term secret keys and key updates

Existing Solutions

- [Attrapadung, Imai Pairing'09]: **Direct ABE revocation**
  - Data owners directly specify revocation list when encrypting

- Extension
  - [Yang, Ding, Lu, Wan, Zhou ISC'13] A semi-trusted server shares the decryption ability with data users, and terminates decryption for revoked users
  - [Attrapadung, Imai ICCC'09] A hybrid revocable ABE system allows data owner to select either direct or indirect revocation when encrypting a message

Limitations of Existing Solutions

- **Direct revocation**
  - Requires all data owners to keep a current revocation list. This makes the system impurely attribute-based, since data owners in ABE create a ciphertext based solely on attributes without caring each data user's status
- **Tree-based revocation solutions**, such as [BGK08], require
  - All non-revoked users to periodically update decryption keys themselves to decrypt newly encrypted data
  - Every user keeps O(log***N***) long-term private keys
- Is it possible to overcome the above limitations in tree-based revocations solutions?

Our ABE Scheme with Server-Aided Revocation [ESORICS'16]

- All user revocation related operations are delegated to a public server which keeps users' public keys
- KGC periodically sends **O(r log N/r)** key updates to public server which forms transformation keys for current time period for non-revoked users only
- Public server uses a user's transformation key to transform a ciphertext
- User uses his/her secret key to finally decrypt on transformed ciphertext

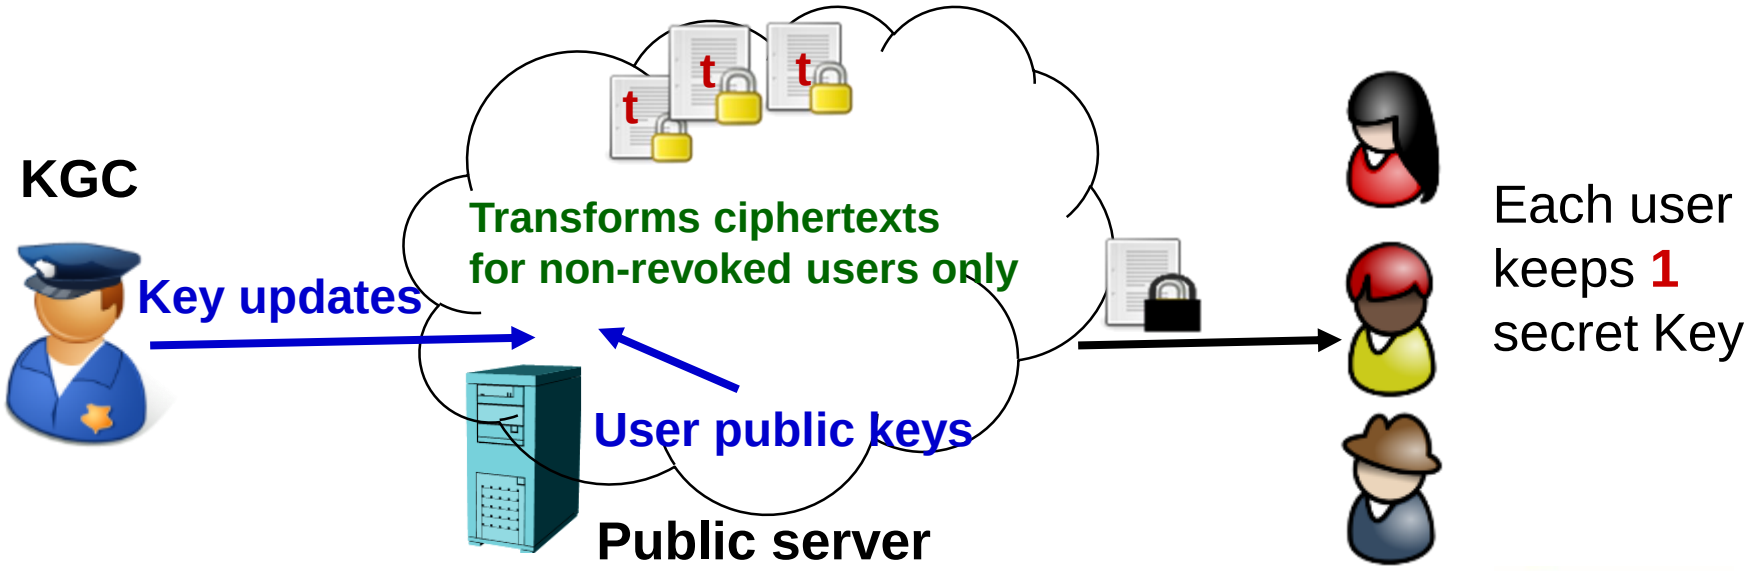

Security Properties

- Revoked user cannot decrypt ciphertexts generated under the current and future time periods
- Public server cannot obtain the message encrypted in a ciphertext
- Except the KGC, all other parties may collude
- No secure channels are required between the KGC, users and the public server

Basic Idea

X: function for user public key
Y: function for key update
Z: function for transformation key
X and Y have different parts of KGC's master key as input

- User $\mathbf{u_1}$'s public key: $PK_1 = (X(u_1,\mathbf{A}_1)_0, X(u_1,\mathbf{A}_1)_1, X(u_1,\mathbf{A}_1)_3, X(u_1,\mathbf{A}_1)_7)$
- User $\mathbf{u_3}$'s public key: $PK_3 = (X(u_3,\mathbf{A}_3)_0, X(u_3,\mathbf{A}_3)_1, X(u_3,\mathbf{A}_3)_4, X(u_3,\mathbf{A}_3)_9)$
- If no user is revoked at time t:
  Key update $KU_t = \{Y(t)_0\}$ and
  transformation key for user $\mathbf{u_i}$ at time t:
  $TK(u_i,t) = Z(X(u_i,\mathbf{A}_i)_0, Y(t)_0)$
- If user $\mathbf{u_3}$ is revoked at time t:
  $KU_t = \{Y(t)_3, Y(t)_{10}, Y(t)_2\}$ and
  transformation key for user $\mathbf{u_1}$ at time t:
  $TK(u_1,t) = Z(X(u_1,\mathbf{A}_1)_3, Y(t)_3)$

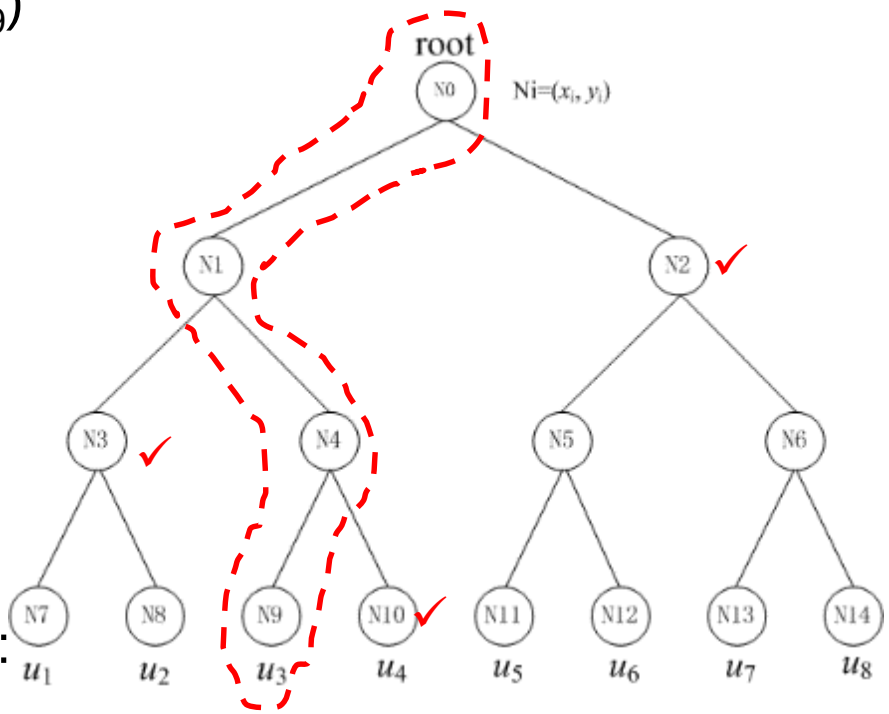

Comparison of Revocable ABE

	[BGK08]	[AI09]	[YDLWZ13]	[SSW12]	Ours
Revocation Mode	Indirect	Indirect & Direct	Direct	Indirect	Indirect
Server	–	–	Semi-trust	–	**Public**
Key Exposure Resistance	No	No	-	No	Yes
Security	Selective	Selective	Selective	Selective	Selective
Secure Channel	Yes	Yes	Yes	Yes	**No**
Size of Key Updates	$O(R \log(N/R))$	$O(R \log(N/R))$	_	$O(R \log(N/R))$	$O(R \log(N/R))$
No. of Key Stored by User	$O(l \log N)$	$O(l \log N)$	$O(1)$	$O(l \log N)$ & $O(k \log N)$	**$O(1)$**

N: the number of all data users; R: the number of revoked data users; l: the number of attributes presented in an access structure; k: the size of the attribute set associated with an attribute-key

Our IBE/ABE with Server-Aided Revocation

- [ESORICS'15] IBE with server-aided user revocation
- [ESPRICS'16] ABE with server-aided user revocation
- [SecureComm'16] ABE with server-aided granular revocation
  - Both user and attribute revocation

Privacy-Preserving Outsourced Computations in the Cloud

Homomorphic Encryption

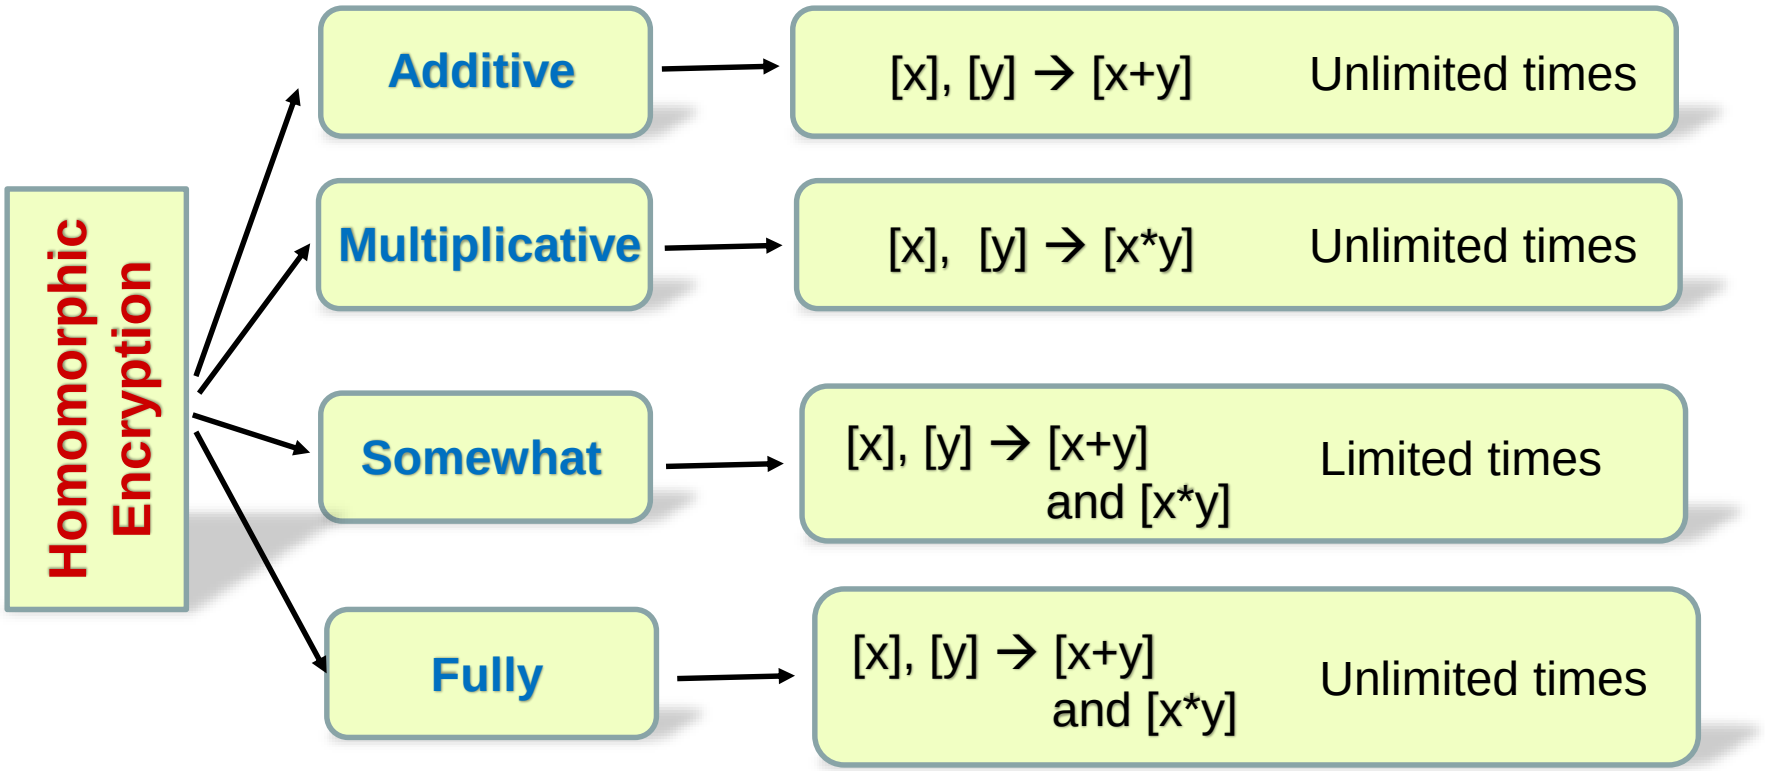

Homomorphic Encryption

Fully Homomorphic encryption can achieve secure arbitrary computation, but up to now, it's still too costly in computation and storage

Our approach: **semi-homomorphic encryption + system architecture**

Our Contributions

- Privacy-preserving outsourced calculation of integers and rational numbers [TDSC, accepted]
- Privacy-preserving outsourced calculation of floating point numbers [TIFS, 2016]
- Privacy-preserving outsourced calculation toolkits with multiple keys [TIFS, 2016]
- Encrypted data processing with homomorphic re-encryption [Information Sciences 2017]
- Privacy-preserving data processing with flexible access control [TDSC, accepted]
- Privacy-preserving outsourced clinical decision support system in the cloud [TSC, accepted]

Paillier Encryption

Notation: Let x $\in Z_N$, and $[x]$ denotes the encryption of x

Additive homomorphic: given $[x]$ and $[y]$, we have

$$[x] \cdot [y] \bmod N^2 = [x + y]$$

Scalar-product homomorphic: given $[x]$, $b \in Z_N$, we have

$$[x]^b \bmod N^2 = [bx]$$

Note: $[x + y] = [x + y \ mod \ N]$ and $[bx] = [bx \ mod \ N]$

Use Paillier Cryptosystem to Achieve Secure Integer Operations [1, 2]

Cloud Server (CP)

Computational Service Provider (CSP)

$[x], [y]$

sk

Two non-colluding servers

[1] Samanthula B K, Elmehdwi Y, Jiang W. K-nearest neighbor classification over semantically secure encrypted relational data. IEEE transactions on Knowledge and data engineering, 2015, 27(5): 1261-1273.
[2] Bost R, Popa R A, Tu R, and Goldwasser S. Machine learning classification over enecypted data, 2015 NDSS.

Use Paillier Cryptosystem to Achieve Secure Multiplication [1]

Given $[x]$ and $[y]$, output $[xy]$

Step-1@CP: randomly get $r_x, r_y \in Z_N$, compute

$$X = [x] \cdot [r_x] = [x + r_x]$$

$$Y = [y] \cdot [r_y] = [y + r_y]$$

send X and Y to CSP

Step-2@CSP:

use sk to decrypt X, Y to get X', Y'

compute $h = X' \cdot Y' = (x + r_x)(y + r_y)$

send $[h]$ to CP

[1] Samanthula B K, Elmehdwi Y, Jiang W. K-nearest neighbor classification over semantically secure encrypted relational data. IEEE transactions on Knowledge and data engineering, 2015, 27(5): 1261-1273.

Use Paillier Cryptosystem to Design Secure Multiplication [1]

Step-3@CP: compute

$$S_1 = [r_x \cdot r_y]^{N-1} = [-r_x \cdot r_y]$$

$$S_2 = [x]^{N-r_y} = [-r_y \cdot \mathrm{x}]$$

$$S_3 = [y]^{N-r_x} = [-r_x \cdot y]$$

$$[h] \cdot S_1 \cdot S_2 \cdot S_3 = [h - r_x y - r_y x - r_x r_y] = [x \cdot y]$$

Limitations of Existing Solutions

- Limitations of [1]
  - Private key is directly stored in CPS
  - Only support integer addition, multiplication, comparison, squared Euclidean distance etc

- Limitations of [2]
  - Private key is directly stored in CPS
  - Only support integer addition, comparison, argmax (the index of the largest value) and dot product

[1] Samanthula B K, Elmehdwi Y, Jiang W. K-nearest neighbor classification over semantically secure encrypted relational data. IEEE transactions on Knowledge and data engineering, 2015, 27(5): 1261-1273.
[2] Bost R, Popa R A, Tu R, and Goldwasser S. Machine learning classification over enecypted data, 2015 NDSS.

Our Approach

Use (n, k) threshold Paillier cryptosystem in which private key is split into n shares, such that any k shares can successfully decrypt

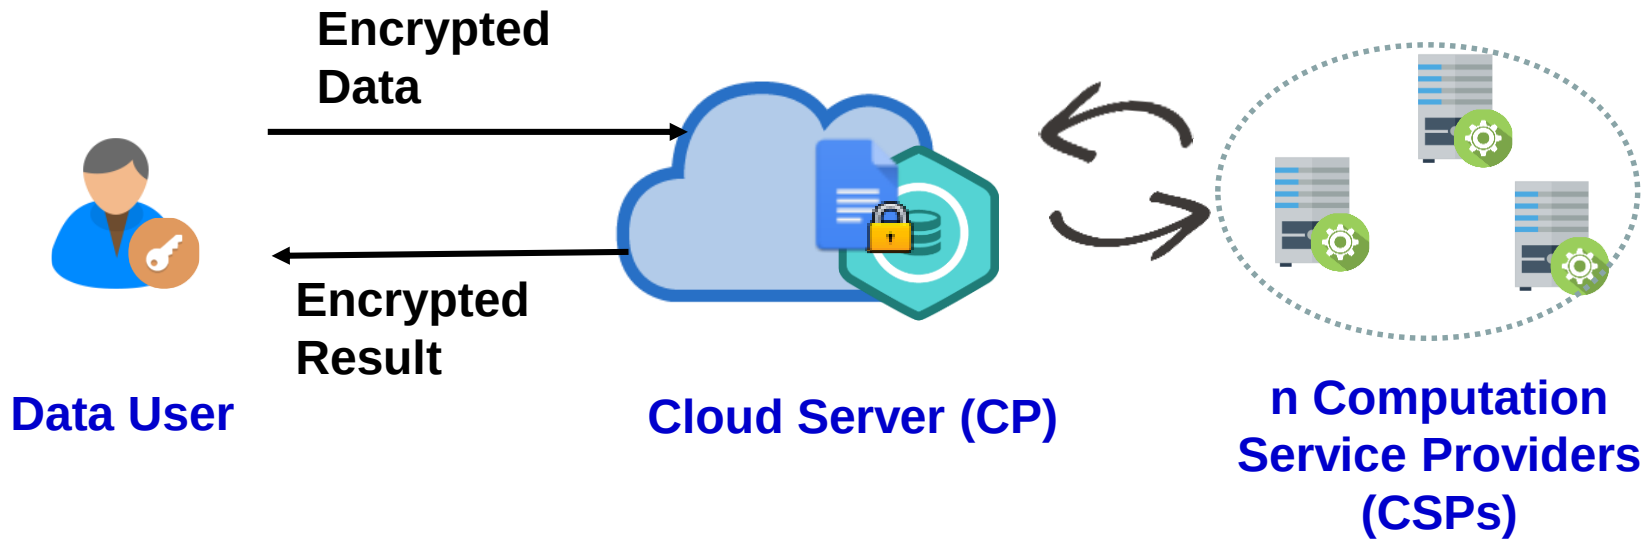

[3] Liu X, Choo R, Deng R, et al. Efficient and privacy-preserving outsourced calculation of rational numbers. IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), 2016, accepted.

Privacy-preserving Integer Calculation Toolkit

Secure Multiplication (SM): Given [x] and [y], output [xy]

Secure Less Than (SLT): Given [x] and [y], output [u], where u = 0 when $x \geq y$ and u = 1 when $x < y$

Secure Maximum and Minimum Sorting (SMMS): Given [x] and [y], output ([A], [B]), where $A \geq B$

[3] Liu X, Choo R, Deng R, et al. Efficient and privacy-preserving outsourced calculation of rational numbers. IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), 2016, accepted.

Privacy-preserving Integer Calculation Toolkit

Secure Equality Testing (SEQ): Given [x] and [y], output [f], where $f = 0\ if\ x = y, otherwise\ x \neq y$

Secure Division (SDIV): Given [x] and [y], output [q] and [r], where $y = q \cdot x + r$

Secure Greatest Common Divisor (SGCD): Given [x] and [y], output [c], where $c = GCD(x, y)$

[3] Liu X, Choo R, Deng R, et al. Efficient and privacy-preserving outsourced calculation of rational numbers. IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), 2016, accepted.

Privacy-preserving Integer Calculation Toolkit

Performance of secure calculations of integer numbers
PC with 3.6 GHz 6-core processor, 12G RAM, |N|=1024

Protocol	CP Comp	CSP Comp	Commu
SM	*82.7 ms*	*51.7 ms*	*1.25 KB*
SLT	*37.6 ms*	*29.9 ms*	*0.75 KB*
SEQ	*266.7 ms*	*165.6 ms*	3.99 *KB*
SMMS	*80.8 ms*	*45.8 ms*	2.74 *KB*
SDIV (10 bits)	*6.21 s*	*4.72 s*	*127.59 KB*
SGCD (10 bits)	*156.0 s*	*116.1 s*	*1.58 KB*

[3] Liu X, Choo R, Deng R, et al. Efficient and privacy-preserving outsourced calculation of rational numbers. IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), 2016, accepted.

Privacy-preserving Rational Numbers Calculations

Encryption of Rational Number

A rational number x can be expressed as $\frac{x^+}{x^-}$, and encrypted as $([x^+], [x^-])$

For example, -0.25 can be expressed as $\frac{x^+}{x^-} = -\frac{1}{4}$, and encrypted as $([1]^{N-1}, [4])$

[3] Liu X, Choo R, Deng R, et al. Efficient and privacy-preserving outsourced calculation of rational numbers. IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), 2016, accepted.

Privacy-preserving Rational Numbers Calculations

Secure Rational Number Processing

Similar to that of the plaintext rational numbers

For example, $\frac{1}{4} \cdot \frac{3}{4} = \frac{1 \cdot 3}{4 \cdot 4}$ and given ([1],[4]) and ([3],[4]), we have

$$(SM([1],[3]), SM([4],[4]))$$

[3] Liu X, Choo R, Deng R, et al. Efficient and privacy-preserving outsourced calculation of rational numbers. IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), 2016, accepted.

Privacy-preserving Rational Numbers Calculation Toolkit

Performance

Protocol	CP Comp	CSP Comp	Commu
ADD(R)	*280.7 ms*	*155.6 ms*	*3.74 KB*
MIN(R)	*283.7 ms*	*154.0 ms*	*3.74 KB*
MUL(R)	*190.3 ms*	*105.7 ms*	2.49 *KB*
DIV(R)	*195.3 ms*	*108.1 ms*	2.49*KB*
CMP(R)	*216.6 ms*	*125.5 ms*	*3.24KB*
EQ(R)	*495.1 ms*	*273.8 ms*	*6.49KB*

[3] Liu X, Choo R, Deng R, et al. Efficient and privacy-preserving outsourced calculation of rational numbers. IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), 2016, accepted.

Privacy-Preserving Outsourced Calculations Across Domains

Data Providers, each has a public/private key pair

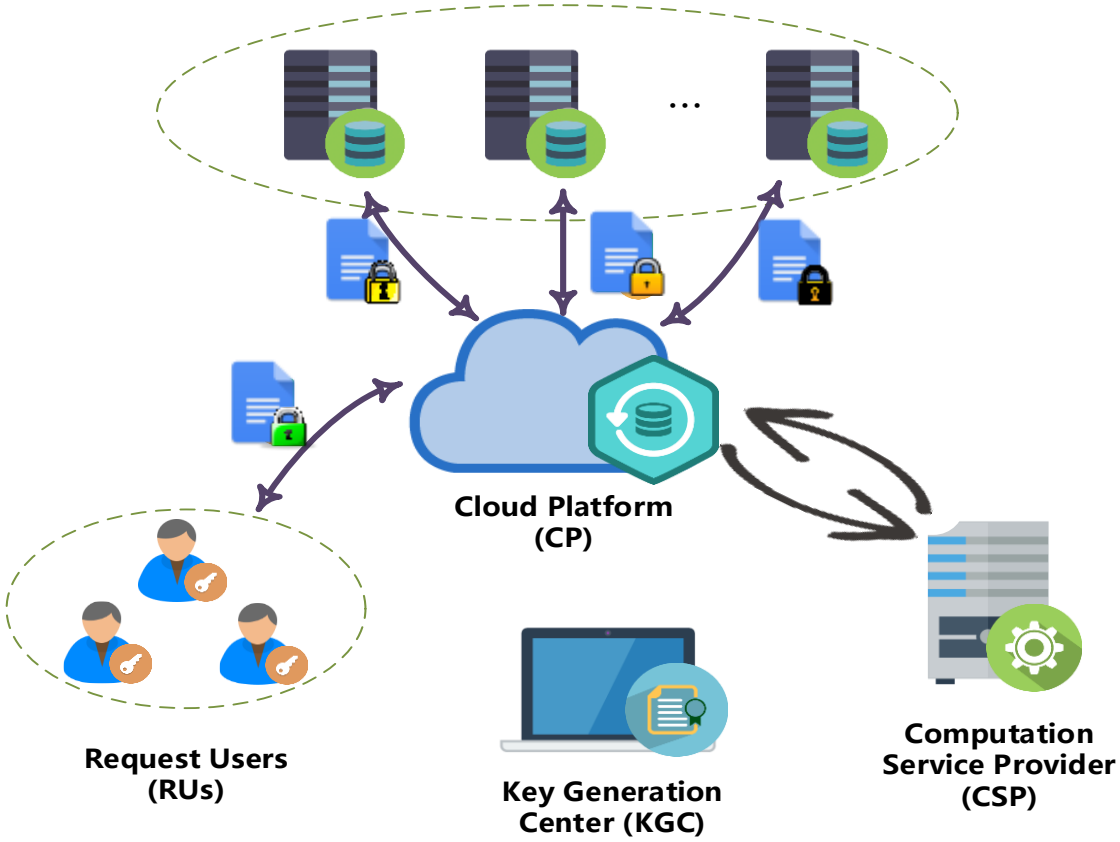

[4] Liu X, Deng R H, Choo K K R, et al. An Efficient Privacy-Preserving Outsourced Calculation Toolkit With Multiple Keys. IEEE Transactions on Information Forensics and Security, 2016, 11(11): 2401-2414.

Two Trapdoor Paillier Cryptosystem

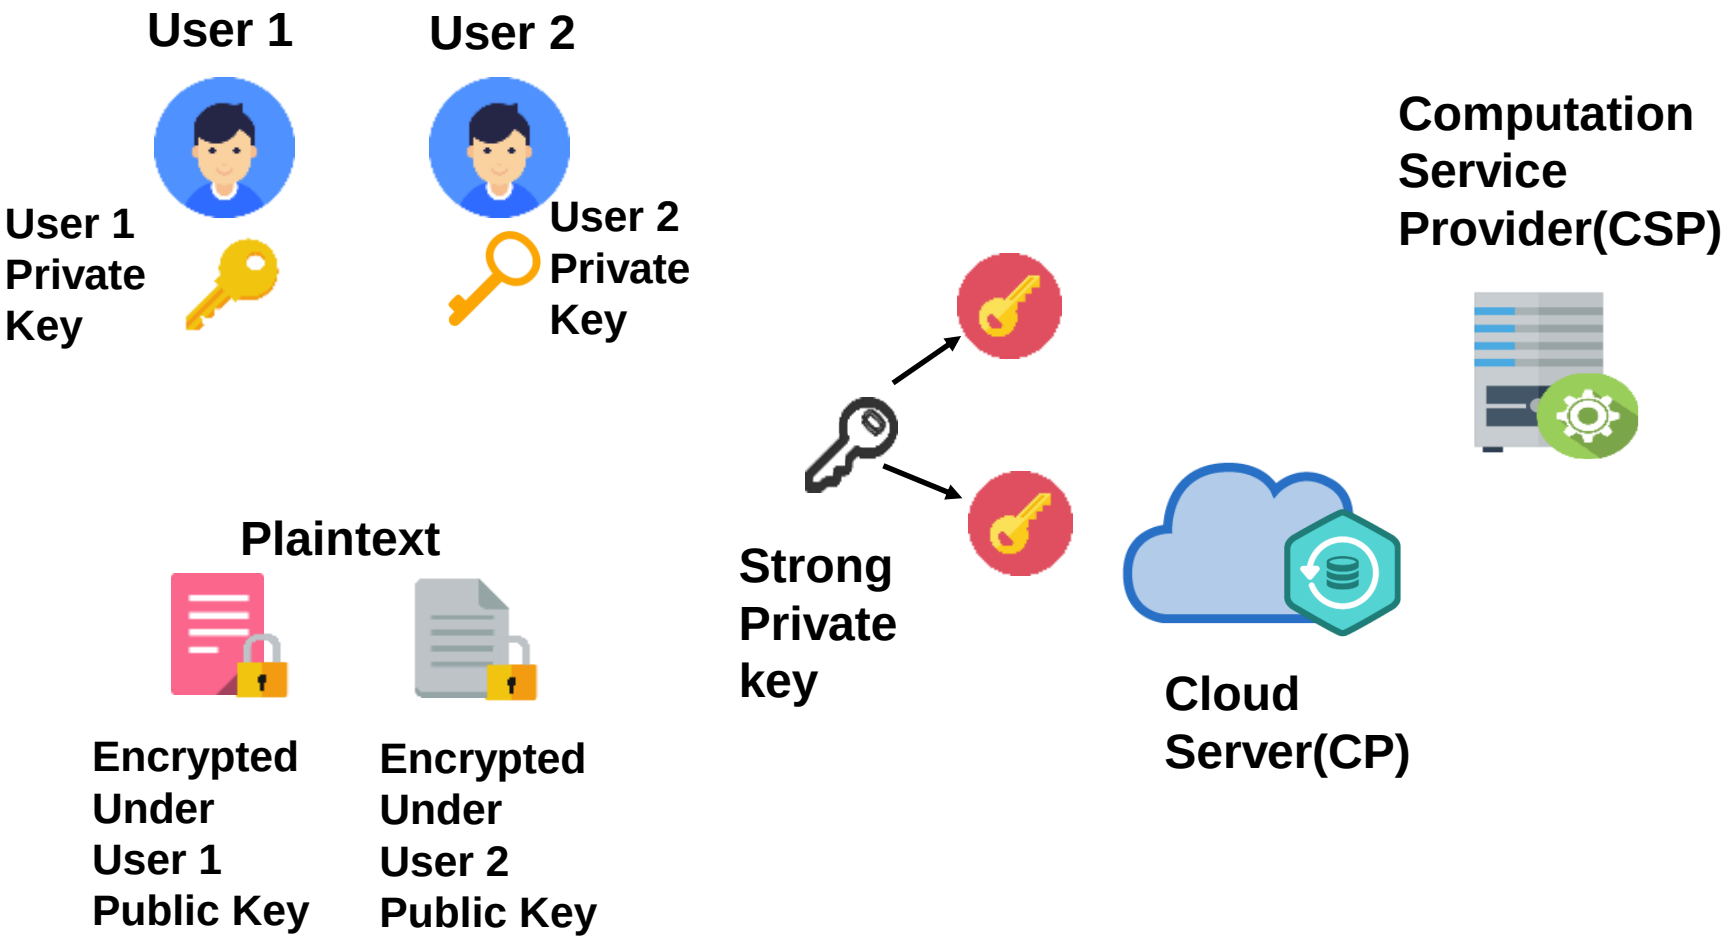

[5] Liu X, Deng R H, Choo K K R, et al. An Efficient Privacy-Preserving Outsourced Calculation Toolkit With Multiple Keys. IEEE Transactions on Information Forensics and Security, 2016, 11(11): 2401-2414.

Applications

Applications

How to securely process complex (non-liner) curves??

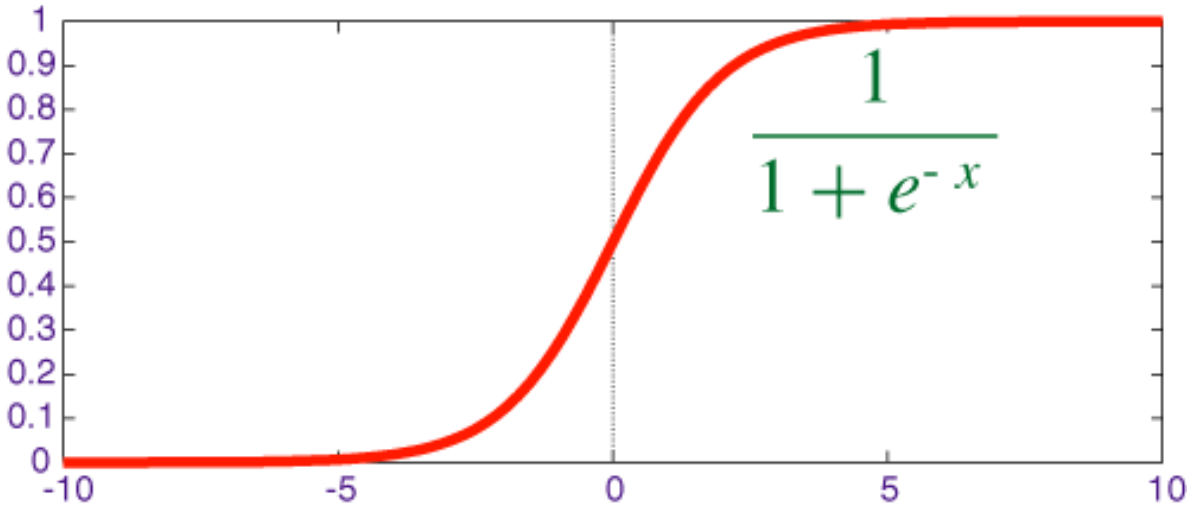

Piecewise vs Maclaurin/Taylor Series Approximations

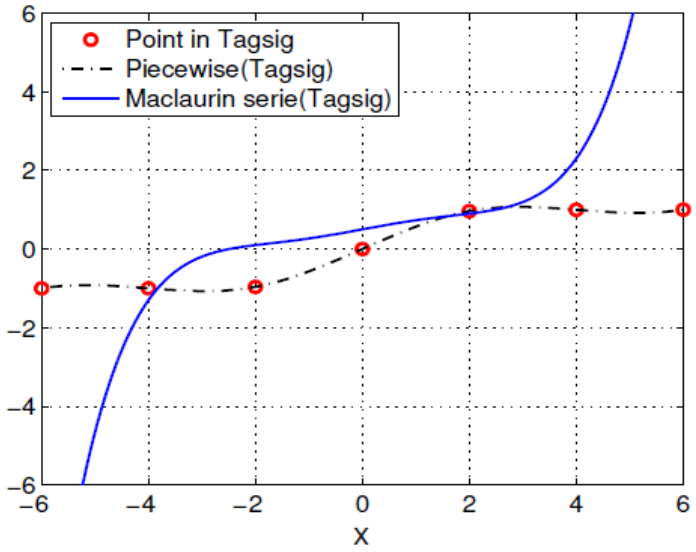

(a) Comparison

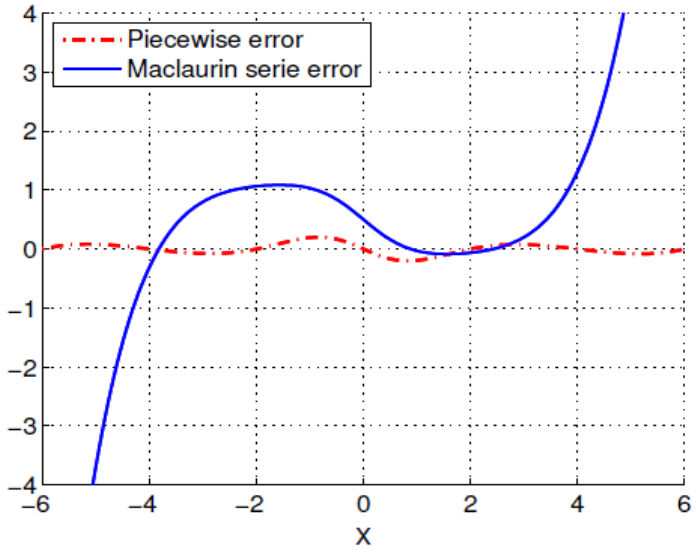

(b) Error Curves

[6] Liu X, Deng R H, Yang Y, et al. Hybrid Privacy-Preserving Clinical Decision Support System in Fog-Cloud Computing, Future Generation Computer Systems, 2017.

Privacy-Preserving Piecewise Function Calculation

Piecewise approximation

$$f(x) = \begin{cases} f_1(x) = \alpha_{1,k}x^k + \cdots + \alpha_{1,1}x + \alpha_{1,0}, & x \geq p_1 \\ f_2(x) = \alpha_{2,k}x^k + \cdots + \alpha_{2,1}x + \alpha_{2,0}, & p_2 \leq x < p_1 \\ \vdots & \\ f_z(x) = \alpha_{z,k}x^k + \cdots + \alpha_{z,1}x + \alpha_{z,0}, & x < q_{z-1} \end{cases}$$

Secure computation of piecewise approximation

$$[f(x)] = [u_1][f_1(x)]+[u_2][f_2(x)] +\ldots + [u_z][f_z(x)]$$

[6] Liu X, Deng R H, Yang Y, et al. Hybrid Privacy-Preserving Clinical Decision Support System in Fog-Cloud Computing, Future Generation Computer Systems, 2017.

Summary

- Scalable Access Control of Encrypted Data

- ABE is an one-to-many public key encryption and allows scalable access control of encrypted data in the cloud
- Verifiable outsourced decryption of ABE
- Efficient user and attribute revocation
- CP-ABE with partial hidden access police
- Deduplication on encrypted data
- Attribute-based secure messaging system in the cloud
- Usability study
- Integrated design and implementations

Summary
- Secure outsourced computations

- Combining semi-homomorphic encryption and system approaches to realize secure computation over encrypted data
- Secure integer computations, secure rational number computations, secure floating point number computations
- Secure computation across multiple domains
- Applications such as secure processing of complex curves
- Limitations
  - Efficiency, multiple servers, overflow/underflow problems which are common to all homomorphic encryption schemes

Thank You!

For more information please contact
robertdeng@smu.edu.sg