Encryption for Fine-Grained Access Control

Supervisors.
Michel Abdalla mabdalla@di.ens.fr http://www.di.ens.fr/~mabdalla
Hoeteck Wee wee@di.ens.fr http://www.di.ens.fr/~wee

Location. Ecole Normale Supérieure, Paris.

Project

We live in an era of "Big Data", wherein a deluge of data is being generated, collected, and stored all around us. These data include not only government, financial and medical records, but also personal information exchanged over email, social networks and other data-sharing sites. Without taking measures to protect these data, we risk living under digital surveillance in an Orwellian future. Unfortunately, traditional encryption systems lack the expressiveness needed in applications involving big, complex data.

Functional encryption. Functional encryption [2] is an emerging paradigm for public-key encryption that enables more fine-grained access control to encrypted data. For instance, it provides the ability to specify a decryption policy in the ciphertext so that only individuals who satisfy the policy can decrypt, and the ability to associate keywords to a secret key so that it can only decrypt documents containing the keyword. This stands in contrast to traditional encryption, where access to the encrypted data is all or nothing one can either decrypt and read the entire plaintext or one learns nothing at all about the plaintext (other than its length). Furthermore, functional encryption enforces resilience to collusion attacks, namely any group of users holding different secret keys learns nothing about the plaintext beyond what each of them could individually learn.

Conditional disclosure of secrets. Let us begin with a much simpler information-theoretic primitive, namely that of conditional disclosure of secrets (CDS) [5] (Fig 1). CDS allows two parties Alice and Bob to disclose a secret $\alpha \in \mathbb{Z}_q$ to an external party Carol, subject to a given condition on their respective inputs x and y. Carol knows x, y but not α, so she knows whether the condition holds and whether she will obtain the secret. Alice and Bob on the other hand only sees their portion of the input and does not necessarily know whether Carol will obtain the secret. We also allow Alice and Bob to share randomness w which is independent of their inputs.

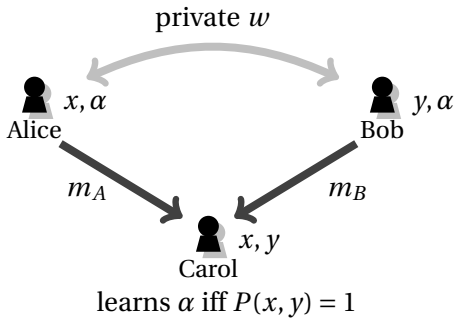

Figure 1: Conditional disclosure of secrets.

For example, consider the following two predicates:

- **equality.** $P : \mathbb{Z}_q \times \mathbb{Z}_q \to \{0,1\}$ and $P(x,y) = 1$ iff $x = y$.

  Here, we have $w = (w_0, w_1) \in \mathbb{Z}_q^2$. Alice sends $m_A = w_0 x + w_1$ and Bob sends $m_B = \alpha + w_0 y + w_1$. Carol outputs $m_B - m_A$.

  It is straight-forward to verify that if $P(x,y) = 1$, then Carol outputs α. On the other hand, if $P(x,y) = 0$, then Carol just sees two uniformly and independently random values in $\mathbb{Z}_q$ which reveal no information about α.

- **index.** $P : \{0,1\}^n \times \{1,2,\ldots,n\} \to \{0,1\}$ and $P(x,y) = x_y$, i.e., the predicate is true iff the y'th bit of the string x is 1.

  Here, we have $w = (w_1, \ldots, w_n) \in \{0,1\}^n$. Alice sends $m_A = (m_A^1, \ldots, m_A^n) \in \{0,1\}^n$, where

  $$m_A^i = \begin{cases} w_i & \text{if } x_i = 1 \\ 0 & \text{otherwise} \end{cases}$$

  Bob sends $\alpha \oplus w_y \in \{0,1\}$. Carol outputs $m_A^y \oplus m_B$. In particular, if $x_y = 0$, then Alice's message leaks no information about w_y, which serves as a one-time pad that completely hides α.

  It is not hard to see that there is another protocol where Alice sends 1 bit and Bob sends n bits, and also one where both Alice and Bob sends $O(\sqrt{n})$ bits.

Project goals. We will study protocols for conditional disclosure of secrets, either to investigate new constructions and lower bounds [4, 6], or to explore the connection to functional encryption [7, 3, 1]. The student is strongly encouraged to attend seminars and to interact with Ph.D. students, post-docs, and permanent researchers in the group.

Requirements. Being comfortable with mathematical proofs and the notion of efficient reductions between computational problems. Knowledge of probability and basic computational complexity is strongly recommended. It is also very strongly recommended to have taken at least one course in cryptography and having been exposed to the notion of "provable security" (as covered in 2.12.1 and 2.30 of the MPRI). Fluency in English is a necessity.

References.

[1] M. Abdalla, R. Gay, M. Raykova, and H. Wee. Multi-input inner-product functional encryption from pairings. In *EUROCRYPT*, 2017.

[2] D. Boneh, A. Sahai, and B. Waters. Functional encryption: a new vision for public-key cryptography. *Commun. ACM*, 55(11):56–64, 2012.

[3] J. Chen, R. Gay, and H. Wee. Improved dual system ABE in prime-order groups via predicate encodings. In *Eurocrypt*, pages 595–624, 2015.

[4] R. Gay, I. Kerenidis, and H. Wee. Communication complexity of conditional disclosure and attribute-based encryption. In *CRYPTO*, pages 485–502, 2015.

[5] Y. Gertner, Y. Ishai, E. Kushilevitz, and T. Malkin. Protecting data privacy in private information retrieval schemes. *J. Comput. Syst. Sci.*, 60(3):592–629, 2000.

[6] T. Liu, V. Vaikuntanathan, and H. Wee. Conditional disclosure of secrets via non-linear reconstruction. In *CRYPTO*, 2017.

[7] H. Wee. Dual system encryption via predicate encodings. In *TCC*, pages 616–637, 2014.