Functional Encryption and its Impact on Cryptography

Hoeteck Wee*

ENS, Paris, France

Abstract. Functional encryption is a novel paradigm for public-key encryption that enables both fine-grained access control and selective computation on encrypted data, as is necessary to protect big, complex data in the cloud. In this article, we provide a brief introduction to functional encryption, and an overview of its overarching impact on the field of cryptography.

1 Introduction

Recent computing and technological advances such as the ubiquity of high-speed network access and the proliferation of mobile devices have had a profound impact on our society, our lives and our behavior. In the past decade, we have seen a substantial shift towards a digital and paperless society, along with a migration of data and computation to the cloud. Storing data in the cloud offers tremendous benefits: easy and convenient access to the data and reliable data storage for individuals, as well as scalability and financial savings for organizations. On the flip side, storing data remotely poses an acute security threat as these data – government, financial, medical records as well as personal information exchanged over email and social networks – are outside our control and could potentially be accessed by untrusted parties. Without taking measures to protect our data, we are at risk of devastating privacy breaches and living under digital surveillance in an Orwellian future.

However, traditional public-key encryption lacks the expressiveness needed to protect big, complex data:

(i) First, traditional encryption only provides coarse-grained access to encrypted data, namely, only a single secret key can decrypt the data. Corporate entities want to share data with groups of users based on their credentials. Similarly, individuals want to selectively grant access to their personal data on social networks and Google Docs.

(ii) Second, access to encrypted data is "all or nothing": one either decrypts the entire plaintext or learns nothing about the plaintext. In applications such as data-mining on encrypted medical records or social networks, we want to provide only partial access and selective computation on the encrypted data, for instance, restricted classes of statistical or database queries.

* wee@di.ens.fr. CNRS (UMR 8548) and INRIA. Supported in part by NSF Awards CNS-1237429 and CNS-1319021 and a fellowship from the Alexander von Humboldt Foundation.

Ideally, we want to encrypt data while enabling fine-grained access control and selective computation; that is, we want control over *who* has access to the encrypted data and *what* they can compute. Such a mechanism would reconcile the conflict between our desire to outsource and compute on data and the need to protect the data.

2 Functional Encryption

Over the past decade, cryptographers have put forth a novel paradigm for public-key encryption [30, 23, 3, 28] that addresses the above goal: **(i)** *attribute-based encryption* (ABE), which enables fine-grain access control, and **(ii)** its generalization to *functional encryption*, which enables selective computation.

- In **attribute-based encryption (ABE)**, encrypted data are associated with a set of attributes and secret keys with policies that control which ciphertexts the key can decrypt. For instance, a digital content provider can issue keys that decrypt basic and premium channel contents on weekdays and only basic ones on weekends.
- In **functional encryption**, a secret key enables a user to learn a specific function of the encrypted data and nothing else. For example, decrypting an encrypted image with a cropping key will reveal a cropped version of the image and nothing else about the image.

A salient feature of both attribute-based and functional encryption is that there are many possible secret keys with different decryption capabilities. Moreover, the keys are resilient 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. Together, attribute-based and functional encryption constitute a crisp generalization of several advanced notions of encryption, such as broadcast and identity-based encryption as well as searching on encrypted data; indeed, many advances in public-key encryption over the past decade can be viewed as special cases of attribute-based and functional encryption.

State of the art. The fundamental goals in the study of attribute-based and functional encryption are two-fold: **(i)** to build expressive schemes that support a large class of policies and functions; and **(ii)** to obtain efficient instantiations based on widely-believed intractability of basic computational problems.

The simplest example of attribute-based encryption (ABE) is that of identity-based encryption (IBE), where both the ciphertext and secret key are associated with identities i.e. bit strings, and decryption is possible exactly when the identities are equal. Starting with identity-based encryption (IBE), substantial advances in ABE were made over the past decade showing how to support fairly expressive but nonetheless limited subset of policies, culminating most recently in schemes supporting any policy computable by general circuits [22, 4].

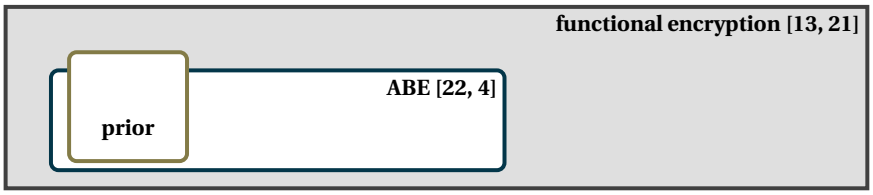

Fig. 1. Advances in attribute-based and functional encryption since 2012. The white region refers to ABE and functionalities for which we have efficient instantiations under standard assumptions; the grey region refers to functionalities beyond ABE for which our understanding is much more limited.

In addition, we have a wide spectrum of techniques for efficient IBE and ABE that yields various trade-offs between efficiency, expressiveness, security and intractability assumptions. The specific assumptions in use may be broadly classified into two categories: (i) pairing-based, such as variants of the Diffie-Hellman problem over bilinear groups, and (ii) lattice-based, notably the learning with errors (LWE) assumption.

Beyond ABE, our understanding of functional encryption is much more limited. The only efficient schemes we have are for very simple functionalities related to computing an inner product [24]. In a recent break-through work, Garg et al. [13] gave a beautiful construction of functional encryption for general circuits; however, the construction relies on "multi-linear maps", for which we have few candidates, along with complex intractability assumptions which are presently poorly understood. In contrast, if we consider collusions of a priori *bounded* size, a weaker guarantee that is still meaningful for many applications, then it is possible to obtain functional encryption for general circuits under a large class of standard assumptions.

Along with these cryptographic advances, the community has also made a greater push towards implementation, prototypes and deployment of attribute-based and functional encryption: several IBE schemes are now standardized in RFC 5091; the CHARM project provides a Python framework for rapidly prototyping cryptosystems and includes implementations of several IBE and ABE schemes; the SHARPS project explores the use of ABE for protecting health-care data; the Mylar project presents a web application platform that uses ABE to provide fine-grained access to encrypted data.

3 Impact on Cryptography

The study of functional encryption has significantly advanced the state of the art in the field of cryptography. In particular, it motivated the development of new and powerful tools and techniques, including trapdoor and delegation techniques in lattices [16, 9] and the first candidate construction of multi-linear maps [12]. These tools and techniques have in turn found numerous applications beyond functional

encryption, notably CCA-secure encryption [8], signatures schemes [5], leakage-resilient cryptography [26], delegating and verifying computation [29, 25, 11], and most recently, garbled circuits [19, 4] and program obfuscation [13, 20]. We highlight three examples, drawing upon recent developments closely related to our research in functional encryption.

Verifiable computation. In verifiable computation, a computationally weak client with input x wishes to delegate a complex computation f to an untrusted server, with the assurance that the server cannot convince the client to accept an incorrect computation [18, 14]. We focus on the online/offline setting, where the protocol proceeds in two phases. In the offline phase, the client sends to the server a possibly long message that may be expensive to compute. Later on, in the online phase (when the input x arrives), the client sends a short message to the server, and receives the result of the computation $f(x)$ together with a certificate for correctness. Applying an existing transformation [29] to our ABE for general circuits [22], we obtain a protocol for verifiable computation on general circuits f with a number of highly desirable properties: (i) the client's communication and computational complexity in the online phase depends only on the input/output lengths and depth of the circuit computing f but not the circuit size; (ii) anyone can check the server's work given a "verification" key published by the client; (iii) we may securely reuse the computation of the offline phase across multiple inputs in the online phase (in particular, our construction is immune to the "rejection problem" from [14]).

Short, scalable signatures. Many applications involving cloud computing and big data require cryptographic primitives that remain secure when used on huge data sets. In particular, we would like to design scalable signatures schemes that remain secure when used to sign a very large number of messages without any performance penalty. However, most known signature schemes are not scalable: their security guarantee degrades linearly in the number of signatures seen by an adversary; this implies a performance degradation as we need to increase key sizes to account for the security loss, which in turn increases the running time of the implementation. In a recent work [10], we presented the first scalable signature scheme in the standard model where each signature is a constant number of group elements. The signature scheme is derived from an IBE with a better security reduction that overcomes seemingly inherent limitations of prior proof techniques, via a delicate combination of techniques used for achieving full security in IBE/ABE [31, 32] and those for constructing efficient pseudo-random functions [27]. Our signature scheme has since been improved and extended to the multi-user setting in [2].

Fully homomorphic encryption. In 2009, Gentry [15] presented the first candidate fully homomorphic encryption (FHE) for all circuits, and substantial progress have since been made towards improving the efficiency and the underlying assumptions [6, 17]. We note that while both FHE and functional encryption support some form

of computation on encrypted data, it is not known how to construct functional encryption from FHE or vice versa. Nonetheless, our lattice-based ABE for branching programs [22] has recently inspired the first FHE schemes based on the LWE assumption with a polynomial modulus-to-noise ratio [7, 1]. Roughly speaking, we propagate LWE samples across computation during decryption in ABE, and during homomorphic evaluation in FHE. If we compute on circuits, the noise accumulated in the LWE samples grows exponentially with the depth D of the circuit (the noise grows as n^D where n is the length of the LWE secret). On the other hand, by exploiting an asymmetry in computation on branching programs, it is possible to achieve noise growth that is linear in the length of the branching program. The latest FHE schemes in [7, 1] then use a branching program instead of a log-depth circuit to compute the decryption function during bootstrapping, thus incurring a polynomial as opposed to a quasi-polynomial noise growth.

Acknowledgments. I am extremely grateful to Jie Chen, Sergey Gorbunov and Vinod Vaikuntanathan for many fruitful collaborations and to Dan Boneh, Yuval Ishai, Allison Lewko and Brent Waters for many illuminating discussions. I would also like to thank Michel Abdalla and the SCN 2014 PC for inviting me as a speaker.

References

[1] J. Alperin-Sheriff and C. Peikert. Faster bootstrapping with polynomial error. In *CRYPTO*, 2014. To appear.

[2] O. Blazy, E. Kiltz, and J. Pan. (hierarchical) identity-based encryption from affine message authentication. In *CRYPTO*, 2014. To appear.

[3] D. Boneh, A. Sahai, and B. Waters. Functional encryption: Definitions and challenges. In *TCC*, pages 253–273, 2011.

[4] D. Boneh, C. Gentry, S. Gorbunov, S. Halevi, V. Nikolaenko, G. Segev, V. Vaikuntanathan, and D. Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *EUROCRYPT*, pages 533–556, 2014.

[5] X. Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In *Public Key Cryptography*, pages 499–517, 2010.

[6] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *FOCS*, 2011. Cryptology ePrint Archive, Report 2011/344.

[7] Z. Brakerski and V. Vaikuntanathan. Lattice-based FHE as secure as PKE. In *ITCS*, pages 1–12, 2014.

[8] R. Canetti, S. Halevi, and J. Katz. Chosen-ciphertext security from identity-based encryption. In *EUROCRYPT*, pages 207–222, 2004.

[9] D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, pages 523–552, 2010.

[10] J. Chen and H. Wee. Fully, (almost) tightly secure IBE and dual system groups. In *CRYPTO (2)*, pages 435–460, 2013.

[11] J. Chen and H. Wee. Semi-adaptive attribute-based encryption and improved delegation for boolean formula. In *SCN*, 2014. To appear.

[12] S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT*, pages 1–17, 2013. Also, Cryptology ePrint Archive, Report 2012/610.

[13] S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, pages 40–49, 2013. Also, Cryptology ePrint Archive, Report 2013/451.

[14] R. Gennaro, C. Gentry, and B. Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *CRYPTO*, pages 465–482, 2010.

[15] C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.

[16] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.

[17] C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO (1)*, pages 75–92, 2013.

[18] S. Goldwasser, Y. T. Kalai, and G. N. Rothblum. Delegating computation: interactive proofs for muggles. In *STOC*, pages 113–122, 2008.

[19] S. Goldwasser, Y. T. Kalai, R. A. Popa, V. Vaikuntanathan, and N. Zeldovich. Reusable garbled circuits and succinct functional encryption. In *STOC*, pages 555–564, 2013.

[20] S. Goldwasser, S. D. Gordon, V. Goyal, A. Jain, J. Katz, F.-H. Liu, A. Sahai, E. Shi, and H.-S. Zhou. Multi-input functional encryption. In *EUROCRYPT*, 2014. To appear.

[21] S. Gorbunov, V. Vaikuntanathan, and H. Wee. Functional encryption with bounded collusions via multi-party computation. In *CRYPTO*, pages 162–179, 2012. Also Cryptology ePrint Archive, Report 2012/521.

[22] S. Gorbunov, V. Vaikuntanathan, and H. Wee. Attribute-based encryption for circuits. In *STOC*, pages 545–554, 2013. Also, Cryptology ePrint Archive, Report 2013/337.

[23] V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *ACM Conference on Computer and Communications Security*, pages 89–98, 2006.

[24] J. Katz, A. Sahai, and B. Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *EUROCRYPT*, pages 146–162, 2008.

[25] A. B. Lewko and B. Waters. New proof methods for attribute-based encryption: Achieving full security through selective techniques. In *CRYPTO*, pages 180–198, 2012.

[26] A. B. Lewko, Y. Rouselakis, and B. Waters. Achieving leakage resilience through dual system encryption. In *TCC*, pages 70–88, 2011. Cryptology ePrint Archive, Report 2010/438.

[27] M. Naor and O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. *J. ACM*, 51(2):231–262, 2004.

[28] A. O'Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010.

[29] B. Parno, M. Raykova, and V. Vaikuntanathan. How to delegate and verify in public: Verifiable computation from attribute-based encryption. In *TCC*, pages 422–439, 2012.

[30] A. Sahai and B. Waters. Fuzzy identity-based encryption. In *EUROCRYPT*, pages 457–473, 2005.

[31] B. Waters. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In *CRYPTO*, pages 619–636, 2009.

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