Controlled Functional Encryption

Novel models can bring modern cryptographic technology to real applications.

Muhammad Naveed
UIUC
naveed.pro

U.S. Department of Health & Human Services reports that health records of more than 39 million individuals have been breached from the hospitals, including Carle hospital in Urbana, and other healthcare institutions. Therefore, a patient may worry about the privacy of his sensitive health data, and want healthcare providers to learn only limited information, e.g., result of a particular test. In general, a patient may want healthcare providers to learn information allowed by the patient specified policy. Such privacy concerns are not limited to healthcare domain. Existing cryptographic techniques do not provide realistic solution to this problem. For example, secure computation would require each patient, or an agent of the patient, to run a computationally intensive program on her computer for each computation (e.g., a test). Functional Encryption can solve the problem, but is extremely inefficient and requires untested cryptographic hardness assumptions. Fully homomorphic encryption is also currently infeasible.

In this work, we propose a new cryptographic model called "Controlled Functional Encryption (C-FE)" that allows us to construct realistic and efficient constructions. As in functional encryption, C-FE allows a user (client) to learn only certain functions of encrypted data, using keys obtained from an authority. However, we allow (and require) the client to send a fresh key request to the authority every time it wants to evaluate a function on a ciphertext. We obtain efficient solutions by carefully combining CCA2 secure public-key encryption with Yao's garbled circuit. Our main contributions in this work include developing and formally defining the notion of C-FE; designing efficient and practical constructions of C-FE schemes achieving these definitions for specific and general classes of functions; and evaluating the performance of our constructions on various application scenarios.

Our constructions are based on efficient cryptographic primitives and perform very well in practical applications. On a laptop (Intel Core i7, 8GB RAM), our construction takes 1.28s and consumes 132KB bandwidth for a 1,000[1] SNP disease marker in personalized medicine application. In genomic patient similarity application, comparing two 4-million[2] SNP profile costs \$0.014[3], takes 4 minutes and consumes 53.77MB bandwidth.

Paper

[1] M. Naveed, S. Agrawal, M. Prabhakaran, X. Wang, E. Ayday, J.-P. Hubaux, and C. A. Gunter. Controlled functional encryption. In *ACM Conference on Computer & communications security (CCS)*. Available online at https://web.engr.illinois.edu/~naveed2/pub/CCS2014CFE.pdf.

Talks at other places

- IBM T.J. Watson Research Center, Oct 2014
- ACM Conference on Computer and Communications Security (CCS), Nov 2014
- University of Washington, Dec 2014
- Microsoft Research Redmond, Dec 2014
- Massachusetts Institute of Technology, Feb 2015 *(Upcoming)*

[1] In practice disease marker size is less than 40 SNPs.
[2] Each human have at most 4-million SNPs.
[3] Based on Amazon EC2 pricing.