Abstract

The advent of cloud computing potentially allows individuals and organizations to outsource storage and processing of large volumes of data to third party servers. At the same time, this leads to privacy concerns. A simple and efficient solution for ensuring the privacy of user information is to encrypt the user data before offloading it to the cloud. However, this solution comes at a price - if the data is encrypted using some conventional encryption algorithm, then the cloud loses all ability to compute on the data without decrypting it first. There exist solutions such as fully homomorphic encryption, functional encryption and oblivious RAM, which allow evaluating rich classes of functionalities directly over encrypted data (without the need for decryption) and provide highly desirable security guarantees. Unfortunately, known implementations of these solutions incur considerable overhead in terms of performance and storage requirements that severely limit their practicality. This thesis studies the provisioning of *practically realizable* encryption schemes that allow evaluating certain *restricted* classes of functionalities on encrypted data, while maintaining a balance between security and efficiency.

We first study the problem of *searchable symmetric encryption* (SSE) for conjunctive keyword searches in a single-writer/single-reader framework. In a seminal work, Cash *et al.* [Crypto'13] proposed Oblivious Cross-Tags (OXT) - an SSE supporting conjunctive searches with high performance, but with significant *partial information leakages* to the untrusted server. We propose a novel leakage suppression technique for conjunctive SSE schemes based on symmetric-key hidden vector encryption and show how to instantiate the same efficiently, without resorting to relatively heavy cryptographic machinery such as bilinear maps. Based on this technique, we propose two new conjunctive SSE schemes for static datasets that improve significantly upon OXT in terms of leakage, albeit at the cost of an additional round of communication for searches. Our experiments show that the second of these two schemes requires $30 - 35$x lower turnaround time for conjunctive searches as compared to OXT. We then focus on designing *dynamic* that

support both updates as well as conjunctive searches over encrypted datasets. This requires satisfying two additional security notions - forward privacy and backward privacy. All previously known SSE schemes that are forward and backward private are restricted to single keyword searches. We address this issue by proposing the first forward and backward private conjunctive SSE with extremely fast updates and searches in practice.

We then study the problem of designing *function private predicate encryption* (PE) schemes in the public-key setting for predicate distributions with low min-entropy. Existing (statistically) function private PE schemes due to Boneh *et al.* [Crypto'13, Asiacrypt'13] *necessarily* require predicate distributions with min-entropy $(\lambda + \omega(\log \lambda))$, where λ is the security parameter. In this thesis, we relax this requirement to $\omega(\log \lambda)$ for certain widely used classes of predicate distributions, namely zero inner-product encryption (ZIPE) and non-zero inner-product encryption (NIPE). We believe that this is a more reasonable requirement than that imposed by existing constructions, and is likely to be satisfied by real-world predicates. As a trade-off, we only aim to achieve function privacy against *computationally bounded* adversaries. We note that for a vast majority of applications, such a relaxed computational notion of function privacy suffices. Technically, the function privacy guarantees of our constructions rely on a new variant of the well-known matrix decisional Diffie-Hellman (MDDH) family of assumptions, that considers matrices sampled from distributions with $\omega(\log \lambda)$ min-entropy. We believe that this is an independent contribution that may have other interesting applications.

We finally study *key-aggregate cryptosystems* (KAC), which were introduced by Chu *et al.* (IEEE TPDS'14) for secure data-sharing in cloud-based applications. In KAC, each plaintext message is encrypted with respect to an individual identity, and a single *aggregate* secret key of fixed overhead can be generated for any arbitrary subset of identities, such that only ciphertexts designated for identities in the set can be decrypted using this aggregate key. In this thesis, we formalize the security of KAC against chosen-ciphertext attacks (CCA), and propose the first concrete KAC construction that achieves this security notion. We also study *key-aggregate searchable encryption* (KASE) - an extension of KAC introduced by Cui *et al.* (IEEE TC'16) that allows keyword-search over restricted sub-parts of encrypted datasets. In this thesis, we propose a formal framework for proving the security of KASE schemes against chosen-database attacks. We then show a direct black-box construction of KASE from any CPA-secure KAC scheme.