Cryptology ePrint Archive: Report 2011/596

Parallel Homomorphic Encryption

Seny Kamara and Mariana Raykova

Abstract: In the problem of private outsourced computation, a client wishes to delegate the evaluation of a function f on a private input x to an untrusted worker without the latter learning anything about x and f(x). This problem occurs in many applications and, most notably, in the setting of cloud computing.

In this work, we consider the problem of privately outsourcing computation to a cluster of machines, which typically happens when the computation is to be performed over massive datasets, e.g., to analyze large social networks or train machine learning algorithms on large corpora. At such scales, computation is beyond the capabilities of any single machine so it is performed by large-scale clusters of workers.

We address this problem by introducing the notion of parallel homomorphic encryption (PHE) schemes, which are encryption schemes that support computation over encrypted data via evaluation algorithms that can be efficiently executed in parallel. We also consider delegated PHE schemes which, in addition, can hide the function being evaluated. More concretely, we focus on the MapReduce model of parallel computation and show how to construct PHE schemes that can support various MapReduce operations on encrypted datasets including element testing and keyword search. More generally, we construct schemes that can support the evaluation of functions in NC0 with locality 1 and m = polylog(k) (where k is the security parameter).

Underlying our PHE schemes are two new constructions of (local) randomized reductions (Beaver and Feigenbaum, STACS '90) for univariate and multivariate polynomials. Unlike previous constructions, our reductions are not based on secret sharing and are fully-hiding in the sense that the privacy of the input is guaranteed even if the adversary sees all the client's queries.

Our randomized reduction for univariate polynomials is information-theoretically secure and is based on permutation polynomials, whereas our reduction for multivariate polynomials is computationally-secure under the multi-dimensional noisy curve reconstruction assumption (Ishai, Kushilevitz, Ostrovsky, Sahai, FOCS '06).

Category / Keywords: homomorphic encryption, local random reductions, parallel computing, MapReduce, Hadoop, cluster-computing, cloud computing

Date: received 3 Nov 2011

Contact author: senyk at microsoft com