Efficient Fully Homomorphic Encryption from (Standard) LWE

Abstract

We present a fully homomorphic encryption scheme that is based solely on the (standard) learning with errors (LWE) assumption. Applying known results on LWE, the security of our scheme is based on the worst-case hardness of short vector problems on arbitrary lattices. As icing on the cake, our scheme is quite efficient, and has very short ciphertexts.

Our construction improves on previous works in two aspects:

1. We show that "somewhat homomorphic" encryption can be based on LWE, using a new *relinearization* technique. In contrast, all previous schemes relied on complexity assumptions related to ideals in various rings.
2. More importantly, we deviate from the "squashing paradigm" used in all previous works. We introduce a new *dimension reduction* technique, which shortens the ciphertexts and reduces the decryption complexity of our scheme, *without introducing additional assumptions*. In contrast, all previous works required an additional, very strong assumption (namely, the sparse subset sum assumption).

Since our scheme has very short ciphertexts, we use it to construct an asymptotically-efficient LWE-based single-server private information retrieval (PIR) protocol. The communication complexity of our protocol (in the public-key model) is $k \cdot \text{polylog}\, k + \log |\mathtt{DB}|$ bits per single-bit query, which is better than any known scheme (here, k is a security parameter). Previously, it was not known how to achieve a communication complexity of even $\text{poly}(k, \log |\mathtt{DB}|)$ based on LWE.

Joint Work with Zvika Brakerski (Weizmann).