Cryptology ePrint Archive: Report 2011/250

A Parallel Repetition Theorem for Leakage Resilience

Zvika Brakerski and Yael Tauman Kalai

Abstract: A leakage resilient encryption scheme is one which stays secure even against an attacker that obtains a bounded amount of side information on the secret key (say λ bits of ``leakage''). A fundamental question is whether parallel repetition amplifies leakage resilience. Namely, if we secret share our message, and encrypt the shares under two independent keys, will the resulting scheme be resilient to 2λ bits of leakage?

Surprisingly, Lewko and Waters (FOCS 2010) showed that this is false. They gave an example of a public-key encryption scheme that is (CPA) resilient to λ bits of leakage, and yet its 2-repetition is not resilient to even $(1+\epsilon)\lambda$ bits of leakage. In their counter-example, the repeated schemes share secretly generated public parameters.

In this work, we show that under a reasonable strengthening of the definition of leakage resilience (one that captures known proof techniques for achieving non-trivial leakage resilience), parallel repetition \emph{does} in fact amplify leakage (for CPA security). In particular, if fresh public parameters are used for each copy of the Lewko-Waters scheme, then their negative result does not hold, and leakage is amplified by parallel repetition.

More generally, we show that given t schemes that are resilient to $\lambda_1, \ldots, \lambda_t$ bits of leakage, respectfully, their direct product is resilient to $\sum (\lambda_i-1)$ bits. We present our amplification theorem in a general framework that applies other cryptographic primitives as well.

[Cryptology ePrint archive]