Cryptology ePrint Archive: Report 2011/530

Key-Evolution Schemes Resilient to Space-Bounded Leakage

Stefan Dziembowski and Tomasz Kazana and Daniel Wichs

Abstract: Much recent work in cryptography attempts to build secure schemes in the presence of \emph{side-channel leakage} or leakage caused by malicious software, like computer viruses. In this setting, the adversary may obtain some additional information (beyond the control of the scheme designer) about the internal secret state of a cryptographic scheme. Here, we consider key-evolution schemes that allow a user to evolve a secret-key K_1 via a \emph{deterministic} function f, to get updated keys $K_2 = f(K_1), K_3 = f(K_2), \ldots$. Such a scheme is \emph{leakage-resilient} if an adversary that can leak on the first i steps of the evolution process does not get any useful information about any future keys. For such schemes, one must assume some restriction on the \emph{complexity} of the leakage to prevent \emph{pre-computation attacks}, where the leakage on a key K_i simply pre-computes a future key $K_{i+ t}$ and leaks even a single bit on it.

We notice that much of the prior work on this problem, and the restrictions made therein, can be divided into two types. Theoretical work offers rigor and provable security, but at the cost of having to make strong restrictions on the type of leakage and designing complicated schemes to make standard reduction-based proof techniques go through (an example of such an assumption is that only the data actually used in computation can leak to the adversary). On the other hand, practical work focuses on simple and efficient schemes, often at the cost of only achieving an intuitive notion of security without formal well-specified guarantees.

In this paper, we complement the two tracks via a middle-of-the-road approach. On one hand, we rely on the random-oracle model, acknowledging the usefulness of this methodology in practice despite its theoretical shortcomings. On the other hand, we show that even in the random-oracle model, designing secure leakage-resilient schemes with clear and meaningful guarantees requires great care and is susceptible to pitfalls. For example, just assuming that leakage ``cannot evaluate the random oracle" can be misleading. Instead, we define a new model in which we assume that the ``leakage" can be any arbitrary \emph{space bounded} computation that can make random oracle calls itself. We connect the space-complexity of a computation in the random-oracle modeling to the \emph{pebbling complexity} on graphs. Using this connection, we derive meaningful guarantees for relatively simple key-evolution constructions. Our security proofs do not rely on the assumption that only data used in the computation can leak.

Our scheme is secure also against a large and natural class of active attacks. This is especially important if the key evolution is performed on a PC that can be attacked by a virus. So far, all results that provided solutions against such attacks were secure under the assumption that the virus can download the data from the machine, but he cannot modify any information stored on it (that was called the {\em bounded retrieval model (BRM)}). This paper provides the first scheme were the adversary in the BRM can also modify the data stored on the machine.

Category / Keywords: secret-key cryptography / graph pebbling, leakage-resilient cryptography, bounded-retrieval model

Date: received 30 Sep 2011

Contact author: tkazana at mimuw edu pl