ALFOMs and the Moirai: Quantifying the Performance/Security Tradeoff for ZK-friendly Hash Functions

Aurélien Boeuf*[1] and Léo Perrin[1]

[1]Inria de Paris – Université Paris-Sorbonne - Paris IV – France

Résumé

Zero-Knowledge (ZK) protocols rely internally on hash functions for their security arguments. However, the hash functions that are the most efficient in this context differ substantially from e.g. SHA-3: their round function R must enable an efficient arithmetization of its verification. In practice, it means that verifying if y = R(x) involves as little finite field multiplications as possible. In turn, this design requirement implies a greater vulnerability to algebraic attacks. In fact, improvement of those have proved devastating, and imply the need to completely rethink the methods used to ensure security against them. In this presentation, we show that it is possible to build a simple yet efficient security argument based on a precise estimate of the so-called "ideal degree" of a system of equations. Furthermore, we show that the increase of this quantity across rounds is tightly connected to the cost of the hash function in two different arithmetizations, namely AIR and R1CS. We precisely quantify this relation by introducing ALgebraic Figures Of Merit (ALFOMs) that capture how efficient a specific primitive (and in fact its round function) are at increasing the security per unit of cost. This new insight allows us to better understand sometimes puzzling performance differences between state-of-the-art hash functions in the R1CS and AIR cases, and to provide a fair and simple comparison of their round functions in this context. Furthermore, we present a new group of round functions we called the Moirai which allow us to explore what a round function providing optimal performance/security tradeoff could look like.

*Intervenant