ETH Zurich, Department of Computer Science
SS 2021

Prof. Ueli Maurer
Dr. Martin Hirt
Konstantin Gegier
Chen-Da Liu Zhang

Cryptographic Protocols
Notes 4

Scribe: Sandro Coretti (modified by Chen-Da Liu Zhang and Konstantin Gegier)

About the notes: These notes serve as written reference for the topics not covered by the papers that are handed out during the lecture. The material contained therein is thus a *strict* subset of what is relevant for the final exam.

This week, the notes complement the proof shown in [Mau15, Theorem 1] to show that (most of) the protocols we have seen are proofs of knowledge.

4.1 Proofs of Knowledge

Proofs of knowledge (POKs) are defined relative to a (efficiently computable) predicate $Q : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}$ (corresponding some **NP**-language L). For some $z \in \{0,1\}^*$, x with $Q(z,x) = 1$ is called a *witness* for z (or, more precisely, for z's membership in L).

To formally define PoKs, one considers a *knowledge extractor*, which is an efficient algorithm K that, by interacting with a prover algorithm P' on some input z, tries to extract a witness x for z. Algorithm K may invoke P' arbitrarily many times and control its random tape.

Zero-Knowledge Proof-of-Knowledge: The definition of zero-knowledge was defined with respective to an instance set L. For proofs of statement, the instance set is the language. For proofs of knowledge, the instance set corresponds to the set of statements which have a witness $L_Q = \{z \in \{0,1\}^* | \ \exists x \ Q(z,x) = 1\}$. Moreover, the prover P receives an additional private input x. We say that a proof of knowledge is zero-knowledge if, for any $z \in L_Q$ and any x such that $Q(z,x) = 1$, the simulator (on input z), is able to produce a transcript U' distributed identically to the transcript U in the actual interaction between P on input (z,x), and V' on input z.

4.1.1 Proving the Proof-of-Knowledge Property

A convenient way of proving that an interactive proof is a proof of knowledge is via the following notion of 2-*extractability*, which we have already encountered (informally) in both the lecture and the exercises.

Definition 4.1. *A three-move round with challenge space $\mathcal{C}$ is* 2-extractable[1] *for a predicate Q if from any two accepting triples (t,c,r) and (t,c',r') with $c \neq c'$ for some input z, one can efficiently compute a x with $Q(z,x) = 1$.*

[1] This is also called *special soundness* in the literature.

Theorem 4.1. *An interactive protocol (P, V) consisting of s independent 2-extractable three-move rounds in which the challenge is chosen uniformly from some challenge space $\mathcal{C}$ is a proof of knowledge if $1/|\mathcal{C}|^s$ is negligible.*

Proof. Consider an arbitrary P' and fix $z \in \{0,1\}^*$. Denote by p the probability that V accepts an interaction with P' on input z.

The knowledge extractor K, which interacts with P' and controls its randomness ℓ, works as follows:

1. Choose ℓ uniformly at random.
2. Generate two independent protocol executions between P' with randomness ℓ and V.
3. If V accepts both executions and they have different challenge sequences, identify the first round in which the challenges differ and use 2-extractability to compute a witness x. Otherwise, return to step 1.

First note that since P''s randomness is fixed, the executions generated in step 2 are identical up to the point where V asks a different challenge for the first time. In particular, the first message in that round is the same. Thus, if such a round exists, 2-extractability implies that K indeed recovers x with $Q(z, x) = 1$.

It remains to bound the running time of K. Denote by $f(\ell)$ the probability that V accepts an interaction with P' when the randomness of P' is set to ℓ. Thus, if L denotes the random variable corresponding to the uniform choice of ℓ by K,

$$\mathbf{E}[f(L)] = p.$$

Moreover, the probability that both executions generated in step 2 are accepting is $f(\ell)^2$, and, therefore, the success probability of a single iteration of K is

$$\mathbf{E}[f(L)^2] \geq \mathbf{E}[f(L)]^2 = p^2,$$

where the first step uses Jensen's inequality. (This ignores that with negligible probability $1/|\mathcal{C}|^s$, the two executions are identical.) Hence, K runs in $\mathcal{O}(1/p^2)$ expected time, which is polynomial if p is non-negligible. □

References

[Mau15] Ueli Maurer. Zero-knowledge proofs of knowledge for group homomorphisms. In *Des. Codes Cryptogr.*, pages 663–676, 2015.