



Universiteit
Leiden
The Netherlands

Compressed Σ -protocol theory

Attema, T.

Citation

Attema, T. (2023, June 1). *Compressed Σ -protocol theory*. Retrieved from <https://hdl.handle.net/1887/3619596>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3619596>

Note: To cite this publication please use the final published version (if applicable).

SUMMARY

Summary

The field of (*probabilistic*) *proof systems* has developed into a flourishing subfield of cryptology and computer science. In analogy to mathematical proofs, the goal of a proof system is for a prover to convince a verifier of the correctness of a claim. However, by contrast, probabilistic proofs allow the verifier to make mistakes, i.e., to accept false claims (soundness error) or reject true claims (completeness error). In many occasions, the error probability can be made negligibly small by repetition, causing only a minor loss in efficiency, which is sufficient for most practical applications. Further, probabilistic proofs may have multiple rounds of interaction between the prover and the verifier, in which case they are also referred to as *interactive proofs*. These two relaxations, due to Babai, Goldwasser, Micali and Rackoff [Bab85; GMR85], revolutionized the theory of proofs. For instance, by trading absolute certainty for high probability and allowing interaction, it is possible to prove claims without revealing anything beyond their correctness, i.e., in *zero-knowledge*. Nowadays, zero-knowledge proofs are widely deployed; they are for instance essential in the public-key infrastructures (PKIs) that manage digital identities and secure communication channels on the internet.

Especially the theory of Σ -protocols [Cra96] now provides a well-understood basis for the modular design of zero-knowledge proof systems in a wide variety of application domains. A Σ -protocol is an interactive proof with three rounds; the prover first sends a message to the verifier, who replies with a challenge sampled uniformly at random from some finite set, and after receiving the prover's response the verifier decides whether to accept or reject the prover's claim. The theory of Σ -protocols stands out in its *modularity*; basic Σ -protocols are elegant and easy to analyze, and complex application scenarios are handled by appropriately combining these basic building blocks. This includes proving the satisfiability of an arithmetic circuit $C: \mathbb{Z}_q^n \rightarrow \mathbb{Z}_q$ [CD98], where $\mathbb{Z}_q$ denotes the ring of integers modulo q . More precisely, it includes proving that C admits a satisfiable input $\mathbf{x} \in \mathbb{Z}_q^n$ such that $C(\mathbf{x}) = 0$. In fact, Σ -protocols even offer a stronger functionality; they allow provers to not only prove that a circuit admits a satisfiable input, but also that they know one. This property is referred to as *knowledge soundness*, and knowledge sound interactive proofs are called *proofs of knowledge*. The circuit satisfiability problem is NP-complete, i.e., every problem for which solutions are efficiently verifiable can be written as a circuit satisfiability problem. Therefore, by means of a Σ -protocol, every efficiently verifiable claim can be proven in zero-knowledge. However, due to the modularity of Σ -protocol theory, there are often more direct and more efficient solutions that avoid the oftentimes cumbersome reduction to a circuit satisfiability problem.

Probabilistic proofs have various performance metrics, indicating for instance the (computational) complexity of generating or verifying a proof. The communi-

cation costs define another important performance metric, i.e., the number of bits communicated between the prover and the verifier. Unfortunately, for many application scenarios, the communication costs of standard Σ -protocols grow *linearly* in the size of the problem instance. For instance, the communication complexity of a Σ -protocol for the circuit satisfiability problem is linear in the size of the arithmetic circuit. More recently, a folding technique was introduced to reduce the communication complexity from linear down to logarithmic in the size of the problem instance [BCC+16; BBB+18]. The resulting protocols are referred to as Bulletproofs. Bulletproofs were introduced as a “drop-in replacement” for Σ -protocols in several applications, such as zero-knowledge proofs for arithmetic circuit satisfiability.

In this dissertation, we reconcile Bulletproofs’ folding technique with the established theory of Σ -protocols. We show that the folding technique can be cast as a significant *strengthening*, rather than a replacement, of Σ -protocols. Our starting point is a basic Σ -protocol for proving knowledge of a preimage of a group homomorphism $\Psi: \mathbb{G}^n \rightarrow \mathbb{H}$. More precisely, this Σ -protocol allows a prover to prove knowledge of a secret input vector $\mathbf{x} \in \mathbb{G}^n$ such that $\Psi(\mathbf{x}) = P$ for some public $P \in \mathbb{H}$, with communication complexity linear in $n \in \mathbb{N}$. Subsequently, we show that, by an appropriate adaptation of Bulletproofs’ folding technique, the communication complexity can be reduced down to logarithmic in n (or polylogarithmic depending on the concrete instantiation). In line with Bulletproofs, this reduction comes at the expense of a logarithmic number of rounds, instead of constant. Since the compression mechanism is cast as an extension of a basic Σ -protocol, many techniques well known from Σ -protocol theory directly carry over to this new *compressed* Σ -protocol theory.

Further, we enhance compressed Σ -protocol theory with two higher level functionalities. First, by an arithmetic secret-sharing based technique, we show how to prove the correctness of m multiplication triples $(\alpha_i, \beta_i, \gamma_i = \alpha_i \cdot \beta_i) \in \mathbb{Z}_q^3$ for $1 \leq i \leq m$. More precisely, proving correctness of multiplication triples is reduced to proving knowledge of a homomorphism preimage, i.e., the nonlinear multiplication triple relation is linearized. This approach is known from Σ -protocol theory [CDM00; CDP12] and inspired by secure multiparty computation [CDN15], however, some adaptations are required to make it amenable for compression. By an appropriate and efficient reduction, we show that this functionality enhancement is sufficient for proving the satisfiability of an arithmetic circuit in (poly)logarithmic communication. As a second functionality enhancement, we construct a novel k -out-of- n proof of partial knowledge, allowing to prove knowledge of k -out-of- n homomorphism preimages without revealing which preimages the prover knows. Proofs of partial knowledge, especially 1-out-of- n , have seen myriad applications during the last decades, e.g., in electronic voting, ring signatures, and confidential transaction systems. Our construction shows how to reduce their communication complexity from linear down to (poly)logarithmic in k and n . We avoid the use of generic circuit satisfiability machinery and identify regimes of practical relevance where our approach achieves asymptotic and concrete performance improvements.

Compressed Σ -protocol theory is presented in a simple and abstract language, allowing for instantiations in a variety of cryptographic platforms. In particular,

we show how to instantiate compressed Σ -protocols from the discrete logarithm assumption, resulting in a logarithmic communication complexity. Moreover, we show how to extend this instantiation to bilinear pairing based platforms. Based on the knowledge of exponent assumption, the communication complexity can be reduced further down to constant. Finally, we present strong-RSA and lattice-based instantiations, the latter plausibly providing post-quantum security. Strong-RSA and lattice-based instantiations are subject to a so-called *soundness slack*. This warrants larger protocol parameters and causes the resulting communication complexity to be polylogarithmic rather than logarithmic or constant.

Additionally, we identify and close three gaps in the general theory of multi-round interactive proofs, with particular relevance to Bulletproofs and compressed Σ -protocols. More precisely, it is generally nontrivial to show that an interactive proof is knowledge sound and to find a tight bound on the knowledge error, i.e., the success probability of a dishonest prover. Therefore, in the context of Σ -protocols, the more convenient notion *special-soundness* was introduced [Cra96]. It is well known that special-soundness, or more precisely 2-out-of- N special-soundness, implies knowledge soundness with knowledge error $1/N$, where N is the size of the verifier's challenge set. More generally, k -out-of- N special-soundness implies knowledge soundness with knowledge error $(k - 1)/N$. Bulletproofs and compressed Σ -protocols have rendered natural *multi-round* generalizations of special-soundness relevant.

The first open problem that we address is the lack of a *tight* knowledge soundness analysis for special-sound multi-round interactive proofs. Non-tight bounds on the knowledge error warrant the use of overly conservative protocol parameters, possibly rendering concrete instantiations inefficient. We provide the first tight knowledge soundness analysis for the broad class of special-sound multi-round interactive proofs.

The second open problem questions the effect of parallel repetition on the knowledge error. In many occasions, the knowledge error κ is not small enough, and thus needs to be reduced. This can be done generically by repeating the interactive proof in parallel. The effect of parallel repetition on 2-out-of- N special-sound Σ -protocols is well known, but the situation becomes significantly more complicated when considering k -out-of- N special-soundness for $k > 2$, let alone its multi-round generalizations. More precisely, the t -fold parallel repetition of a 2-out-of- N special-sound interactive proof is easily seen to be 2-out-of- N^t special-sound, and thus has knowledge error $1/N^t$. A similar result does not hold for the (multi-round) generalizations of special-soundness. We solve the state-of-affairs by proving that, for all special-sound interactive proofs, t -fold parallel repetition optimally reduces the knowledge error from κ down to κ^t .

Third, we analyze the Fiat-Shamir transformation of special-sound multi-round interactive proofs. The Fiat-Shamir transformation is a commonly used heuristic that renders a public-coin¹ interactive proof non-interactive by replacing the verifier's messages by certain hash function evaluations. Unfortunately, the Fiat-Shamir transformation comes with a security loss; in general, the security loss is *exponential* in the number of rounds of the interactive proof. For multi-round

¹An interactive proof is said to be public-coin if the verifier publishes all its randomness during a protocol execution.

interactive proofs, this is a very unfortunate situation when it comes to choosing concrete security parameters. If one wants to rely on the proven security reduction, one needs to choose a large security parameter for the interactive proof, in order to compensate for the exponential security loss, affecting its efficiency. Alternatively, one has to give up on proven security and simply assume that the security loss is much milder than what the general bound suggests – indeed, for many interactive proofs, the known attacks do not feature such a large security loss. The latter, of simply assuming the loss to be milder, has become common practice. In this dissertation, we show that for special-sound interactive proofs the security loss is *independent* of the number of rounds. One can now rely on proven security without choosing overly conservative, and hence inefficient, protocol parameters.

Finally, as an application of compressed Σ -protocol theory, we construct a novel k -out-of- N threshold signature scheme (TSS). The TSS is succinct since a threshold signature has size sublinear in k and n , and in contrast to other succinct TSSs, our TSS does not require a trusted setup and is therefore transparent. We believe that, by the modular nature of compressed Σ -protocol theory, many more application scenarios can be handled in an intuitive and efficient manner.

