Project PROMETHEUS
PRivacy preserving pOst-quantuM systEms
from advanced crypTograpHic mEchanisms
Using latticeS

D5.2

Intermediate results on privacy-preserving cryptographic protocols

Contractual submission date
Month 24

Actual submisision date
20/12/19

Deliverable version
1.0

Main author
Thijs Veugen (TNO)

http://www.h2020prometheus.eu/

h2020prometheus

PROMETHEUS project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 780701. The contents of this deliverable are the responsibility of the PROMETHEUS consortium, and do not necessarily reflect the official views of the European Union.

Document information

Grant agreement no.	780701
Project acronym	PROMETHEUS
Project full title	PRivacy preserving pOst-quantuM systEms from advanced crypTograpHic mEchanisms Using latticeS
Type of action	Research and Innovation Action (RIA)
Topic	H2020-DS-06-2017-Cybersecurity PPP: Cryptography
Project dates	1st January 2018 (Month 1) / 31st December 2021 (Month 48)
Duration	48 months
Project URL	http://www.h2020prometheus.eu/
EU Project Officer	Carmen Ifrim
Work package	WP5 – Privacy-preserving protocols
Deliverable title	Intermediate results on privacy-preserving cryptographic protocols
Deliverable no.	D5.2
Deliverable version	1.0
Deliverable filename	PROMETHEUS-780701-WP5-D5.2.pdf
Nature of deliverable	Report
Dissemination level	Public
Number of pages	19
Responsible partner	TNO (participant number 9)
Authors	Thijs Veugen (TNO), Thomas Ricosset (THA) Olivier Sanders (ORA) Javier Herranz (UPC)

Abstract. This deliverable describes the progress that partners have achieved in the half-life of the project, regarding privacy-preserving lattice-based protocols for the domains: anonymous credentials, e-cash and e-democracy. The deliverable also contains the related problems that remain open, and that will be the object of research in the 24 remaining months.

Keywords: Anonymous credentials, e-cash, e-democracy.

Signatures

Written by	Thijs Veugen	TNO	20/12/19
Reviewed by	Alon Rosen	IDC	17/12/19
Reviewed by	Adria Rodriguez and Enrique Larraia	SCYTL	17/12/19
Approved by	Benoît Libert as Project coordinator	ENSL	23/12/2019
Approved by	Sébastien Canard as Technical leader	ORA	23/12/2019

Partners

ENSL	ENS de Lyon
ORA	Orange SA
CWI	Centrum voor Wiskunde en Informatica
IDC	IDC Herzliya
RHUL	Royal Holloway, University of London
RUB	Ruhr-Universität Bochum
SCYTL	Scytl Secure Electronic Voting, S.A.
THA	Thales Communications & Security S.A.S.
TNO	Nederlandse organisatie voor Toegepast-Natuurwetenschappelijk Onderzoek
UPC	Universitat Politècnica de Catalunya · BarcelonaTech
UR1	Université de Rennes 1
WEI	Weizmann Institute of Science

Contents

1 Introduction

Within PROMETHEUS, WP5 deals with privacy-preserving cryptographic protocols for anonymous credentials, e-cash, and e-democracy. The protocols developed within WP5 are in relation with the building blocks produced in WP4. It aims at developing a specific cryptographic protocol related to anonymous credentials, electronic cash, and electronic voting. The design and implementation are then given on input to the use cases in WP6. The purpose of this document is to describe the intermediate results we have obtained so far within this work package, in order to prepare the use case specifications that are due in one year. As we will see, the way to proceed are different from one task to another, depending on the maturity of the cryptographic protocol we have studied.

Within anonymous credentials, we are working on the way to design and/or improve the cryptographic building blocks that form the basis of all existing anonymous credential systems: group signatures, blind signatures and zerk-knowledge proofs. our main purpose is to significantly reduce the efficiency gap between post-quantum candidates and those based on discrete logarithms. In particular, it aims to design solutions that rely on more efficient zero-knowledge proofs techniques, compatible with our needs.

In the context of e-cash, the maturity is less important and we have been obliged to put the things out to define a general framework for e-cash. This has given a significant paper in one of the major conference in cryptography, and has also permit us to better understand how to proceed and which primitives we need for our purpose.

Regarding e-voting and variants, we have worked on existing frameworks, taken from the literature, and based on several existing lattice-based cryptographic primitives. We have worked on the missing ones, and the way to put things together so as to obtain the best secure and efficient e-voting system based on lattices. We have in particular obtained a significant result on the way to improve the mix-net approach.

These three domains will now be presented independently, leading to three main sections. Each section will be organized differently, depending on the level of maturity of the work that has been done since the beginning of the project. We more or less first introduce the context, the concrete goals and the main way to design such systems. This first step gives, from deliverable D5.1, the main possible approaches. Next, we give some obtained results or ongoing works related to the primitives, as achieved during the first half of our project. This is also done in relation with the work done in WP4, so that we here recall (but not give details) of some results that are fully given in deliverable D4.1. Finally, we describe some remaining open problems that will be investigated in the second half of our project.

2 Anonymous Credentials

2.1 Introduction

Digital signatures are a widely used cryptographic tool in our everyday life. We here focus on the use of digital signatures in credential systems where a user gets some certified credentials issued by some organization. These certified credentials enable the user to prove access rights, get some advantages, or anything else which requires authentication. Anonymous credential systems [?] allow users to prove possession of credentials in an anonymous way. When a credential is shown, a proof of possession is done without leaking (even for the issuing organization) any knowledge about neither the owner of the credential, nor the credential itself.

Several ways exist to design anonymous credential systems. The most common ones are based on either the use of group signatures (or some close variants)or on blind signatures.

- In a group signature scheme, any member of the group can sign messages on behalf of the group. Such signature remains anonymous and unlinkable for anyone except a designated authority (sometime called the Opener) who has the ability to identify the signer. The group is typically controlled by some Issuer that handles enrolment of members. It is then necessary for a user to interact with such Issuer to become a group member, and then be allowed to sign messages anonymously. An anonymous credential is a natural extension of group signatures, permitting a user to interact with an Issuer (the above organization) in order to sign (or prove the possession) on behalf of the group of users having the same credential. Having a group signature scheme is then a possible first step to obtain an anonymous credential system.

- A blind signature scheme permits a user to obtain a signature on a chosen message by interacting with a signing authority. The main difference with a classical signature is that at the end of the signature generation, the authority has never seen the message and is not able to link the signature outputted by the user to its corresponding view of the interactions. Thus, the user is anonymous among the set of users having requesting a signature to this authority. In an anonymous credential system, the organization plays the role of the signer and the user can then obtain a credential, as a signature, that can be used anonymously to prove possession of a certified credential.

But this is definitely not enough to have a group or a blind signature scheme to obtain an anonymous credential. As explained above (and also in D5.1), an anonymous credential system is much more expressive that those two building blocks, so that it is necessary to work on them to obtain what is needed. Moreover, in some situations (see e.g., D6.3 on Anonymous credential use case requirements), a credential is not directly given, but only used to prove that it satisfy some requested requirement, such as for example being under 25. This is typically done using zero-knwoledge proofs of knowledge that the non-revealed credential, that has been signed by some organization, verifies the requested property.

In order to provide a lattice-based anonymous credential system, one possible option is to work on lattice-based group and blind signature schemes, and also on lattice-based zero-knowledge proofs. Several options have been considered so far, and there are several ongoing works.

2.2 Ongoing Works and Related Results

As part of the PROMETHEUS work, several works are in progress that are given. Even if most of them have not yet given concrete publications, some are currently in submission and will certainly be published before the end of the project.

2.2.1 Towards practical lattice-based anonymous credentials

Thales (THA) is currently running an internal project (with a research internship) to study recent lattice-based privacy-preserving primitives and efficient group signatures to determine how they could be used in the construction of anonymous credentials protocols. The lattice-based anonymous credential protocol developed by THA is based on two specific cryptographic primitives: a signature scheme and a verifiable encryption scheme, which are both described in [dLS18]. The encryption scheme is said to be verifiable in the sense that its decryption step is coupled with a NIZK proof allowing to ensure that the decryption is known by the user without requiring to disclose the plaintext.

The group signature scheme presented in [dLS18] leverages new zero-knowledge proofs to provide an efficient lattice-based group signature scheme. It consists of a selectively secure signature scheme in which the adversary must declare the forgery message before seeing the public key. Such a scheme can be converted to a standard signature scheme provided that the size of the message space is limited. To overcome this constraint, the authors chose to strategically pick small subsets of their message space, where they are able to exhibit efficient zero-knowledge proofs of membership. Their scheme, although efficiently running on a laptop in the form of a C implementation, suffers from several drawbacks as a lack of tight security proofs and low flexibility.

In [YAZ$^+$19], the authors present new Zero-Knowledge Arguments of Knowledge (ZKAoKs) of quadratically constrained variables that describe the solution of a system of linear equations, allowing them to construct several privacy-preserving primitives including a range proof protocol. They are able to provide tight security proofs but are limited to the case of unstructured lattices.

In this study, primitives from [dLS18] and [YAZ$^+$19] have been integrated into a new anonymous credential protocol that enables to gradually choose how much of the authenticated value is revealed. Thanks to the addition of range proofs, the user can choose between disclosing the value of the proven attribute, revealing only an interval where this value lies, or hiding it completely. Moreover, several values can be jointly sent for verification. Lastly, users can be added to the group as long as it doesn't exceed the maximal size of 2^{40} users.

2.2.2 Improving Group and Blind Signature Schemes

Group signature without zero-knowledge proofs. A recent paper by Katsumata and Yamada [KY19] proposed to remove the need for zero-knowledge proofs in group signatures by using attribute-based signatures. This has a direct application in the design of post-quantum anonymous credential systems since group signatures and anonymous credentials are closely related. The main problem of this construction is that both the public key and the signature size are proportional to the number of users in the group, which makes the scheme quite inefficient in practice. Even if they propose a variant for which those parameters are independent of the number of users in the system, the construction is secure under non-standard lattice-based assumptions (namely the subexponential hardness of the SIS problem).

On this subject, ORA and UR1 recently started to work on an improvement of this new group signature scheme. The main difference is the attribute based signature which will no more be based on a Bonsai Trees based signature [CHKP12] but rather on an merkle type scheme.

In parallel, UPC is currently working on the design of attribute-based signatures instantiated from ideal lattices, in order to improve the construction from Kaafarani and Katsumata (PKC 2018).

More efficient group-signature based anonymous credential systems. One possibility we are also exploring is to construct an efficient group signature scheme based on the ZKAoKs of [YAZ$^+$19]. This would allow tight security reductions and flexible parameters. The second step is then to depart from it to construct a new anonymous credential protocol. Thanks to a better tuning of the rejection sampling algorithm used in that protocol and by lowering the occurrences of this algorithm we expect a significant gain in the execution time of the showing protocol.

Lattice-based blind signatures. ORA and UR1 is also working on a new blind signature scheme, which follows the initial Rückert [Rüc10] scheme. In fact, there are two restarts during Rückert' signing process, leading to a signature generated after $\exp^{2/\phi}, \phi \in [1, 15] \cap \mathbb{Z}$ trials in average. Our main objective is then to find more efficient alternatives to the problems such restarts are solving. We then use several tricks to reach our goal and design a more efficient blind signature scheme:

- we first replace all the uniform sampling distributions by gaussian distributions, which permits us to benefit better parameters and a more efficient rejection sampling, compared to the Rückert scheme;
- we then make use of the ring version of the efficient trapdoor function due to [MP12, GM18], in order to sample elements on the kernel of a public matrix $\mathbf{A}$. Instead of generating a new challenge or an ephemeral vector in case of error, as it is done in the first *restart from scratch* of Rückert scheme, the signer can execute some rejection sampling to efficiently output a signature where the secret key is always sufficiently hidden;
- we add an oversized vector, compared to the signature sent by the signer, which is generated by the user, thanks to statistical distances between gaussian distribution centered on 0 and gaussian distribution centered on a vector $\mathbf{v}$. This naturally hides the information that can later be used by the signer to recognize the outputted blind signature, since the final signature distribution does not depend on the signature outputted by the signer. The consequence is that we do not need anymore Rückert' second restart from scratch, that has exactly the same objective;
- we finally remark that the removal of the trigger restarts leads to the uselessness of a commitment initially computed by the user during the challenge generation step.

We are also working on a partially blind variant of our scheme, which is necessary in a blind signature based anonymous credential as for the organization to verify some part of the issued credential. The key idea is that the signer generates a GPV signature [GPV08] of the known information. Concretely, the signer uses the pre-sampling technique on the value $\mathcal{F}(\mathsf{info}) = \mathbf{x}$, and computes the pre-sample $\mathbf{u}$ verifying $\mathbf{A} \cdot \mathbf{u} = \mathbf{x}$. This element $\mathbf{u}$ is added in the signature generated by the signer. Then it suffices to subtract this hash value $\mathcal{F}(\mathsf{info})$ in the verification step to get a signing and verification protocol very similar to the classical blind variant, but that now includes the common information.

2.2.3 Improving Zero-Knowledge Proofs

Part of the PROMETHEUS consortium, including ENS de Lyon (ENSL), Orange SA (ORA), THA and UR1, have been exploring the possibility of making the specific noninteractive zero-knowledge proofs faster, which is necessary to make such anonymous credential system practical. The first step in this challenge consists in instantiating those ZKAoKs on ideal lattices in order to make them more efficient.

ENSL wrote two papers related to anonymous credentials. The first one, "Lattice-Based Zero-Knowledge Arguments for Integer Relations" [LLNW18] (Crypto 2018), is on lattice-based protocols allowing to prove relations such as inequalities among commited integers. Range queries aim to verify that one or several secrets lie in a specific pre-determined interval. These queries are frequently encountered in anonymous credentials use cases, for example when one would wish to verify that an individual is of legal age without disclosing any exact value. The second one, "Zero-Knowledge Elementary Databases with More Expressive Queries" [LNTW19] (PKC 2019), is on some cryptographic schemes that allow a prover to commit to a set of D key-value pairs so as to be able to prove statements such as: x belongs to the support of D and $D(x) = y$, or x is not in the support of D. This can lead to new possibilities for anonymous credentials systems. Those results are further details in the WP4 deliverable D4.1.

2.3 Open problems

The main drawbacks of the current approaches to design a lattice-based anonymous credential protocol are related to the following points.

- Firstly, the slowness of the showing protocol, which is currently at least 100 times slower than existing quantum-unsafe anonymous credential protocols. This protocol is even expected to be much slower when using range queries, which are typically used to prove e.g., that a credential is less that a given public value ("I'm under 25");
- Secondly, security proofs (our current work is here adapted from [dLS18]) are far from being tight and does not allow for flexible choice of parameters, which is essential to obtain the best possible efficiency.

These drawbacks are directly related to the underlying Zero-Knowledge Arguments of Knowledge (ZKAoKs) involving the satisfiability of specific matrix-vector relations and integer relations such as inequalities. The different options we are considering (and which are described above) go in this direction, with several possible approaches: improving zero-knowledge proofs, designing systems without zero-knowledge proofs, making use of blind signatures.

3 E-Cash

3.1 Introduction

Electronic Cash (e-cash) is the digital analogue of regular money. It allows users to withdraw coins from a bank and to spend them to merchants, in an anonymous way, thus perfectly emulating conventional cash transactions. Unfortunately, with e-cash, as any digital data, coins can easily be duplicated, and thus spent several times. It is therefore essential to be able to detect double-spending and even to identify the defrauders.

Designing an e-cash system which can handle any amount for a payment (as it is the case for regular cash) is not a trivial task and several kinds of solutions exist in the literature. One of them is to make use of coins of the smallest possible denomination (*e.g.* one cent), but this raises the problem of storing and spending the thousands of coins which become necessary to handle any amount, which problem was parially solved in [CHL05], in which this was possible to withdraw wallets of N coins at once and store them efficiently. Another solution is to manage several denominations but, in practice, a user can be unable to make a payment if his wallet does not contain the kind of denomination he needs, since giving change back is not easy. The last solution, commonly named *divisible e-cash*, enables users to withdraw a coin C of a large value V, and then to spend it in several transactions, but in such a way that the sum of the amount of these transactions v_i is at most the global amount: $V \geq \sum v_i$. This is currently the most relevant solution to solve the above problem and we have decided to focus on this type of e-cash.

It was known for decades that there are some close relations between e-cash and group and blind signatures (introduced in the previous section and in D4.1 and D5.1). It was also known that zero-knowledge proofs are also a very commonly used primitives to design an e-cash system. But, unlike anonymous credentials, there was no real common way to generically design an e-cash system using some basic primitives in a secure way.

In fact, during our work on the state-of-the-art (related in particular to the redaction of deliverable D5.1), we have pointed out a flaw in the security proofs of most previous e-cash papers (including lattice-based candidates). Our first work within PROMETHEUS was then to definitely put the things out to define a general framework for divisible e-cash. This was the work done by ORA and published at AsiaCrypt [BPS19], entitled "Divisible E-Cash from Constrained Pseudo-Random Functions". This paper explains how to fix the above flaw by describing a generic construction of an e-cash system that is proved secure under standard assumptions.

3.2 Main Result

To provide a truly practical solution to the problem of anonymous payment, e-cash systems must at least be compact (the wallet size must not be linear in the number of coins it contains) and, ideally, divisible (the system enables efficient batch spendings). Today, all such constructions roughly follow the framework implicitly defined in the seminal paper by Camenisch, Hohenberger and Lysyanskaya [CHL05]. This framework, based on pseudo-random functions, has in particular led to the first lattice-based e-cash system by Libert et al. [LLNW17], later improved by Yang et al. [YAZ$^+$19].

Flaw in existing schemes. However, we have shown in [BPS19] that this framework does not inherently ensure a very important security property called *exculpability*. Informally, this requires that only the coin's owner can spend it and that he cannot be falsely accused of overspending his coins. More specifically, we show in our paper that all security proofs for this notion were flawed and that only a small fraction of them can be fixed. Unfortunately, previous lattice-based constructions are not among the fixable ones, which concretely means that secure lattice-based e-cash systems no longer exist. In this paper, we also show that merchants' security was not considered by most papers, meaning that they were not always able to clear their transaction.

We therefore have several contributions in [BPS19]. First, we strengthen the security model of e-cash systems to ensure security for all actors. This is mostly done by defining a new notion, called clearing, which encompasses the natural security expectations from merchants.

Next, we identify the limits of the previous framework and propose some modifications to fix it. More specifically, we describe two frameworks in the spirit of [CHL05] that can be proven secure while using the same cryptographic tools as previous works. However, such a proof requires a series of properties from the pseudo-random functions that are not trivial. Before presenting them, we need to recall the notion of pseudo-random functions (PRF).

Informally, a PRF is a function F taking as input a seed s and an element $x \in \mathcal{S}$ and returning a value $F_s(x)$ that looks random, even with the knowledge of $F_s(x')$ for (adaptively chosen) $x' \neq x \in \mathcal{S}$. To support divisibility, we additionally need the ability to constrain the PRF [BGI14, KPTZ13a, BW13], that is, the ability to derive from s and a subset $\mathcal{S}' \subset \mathcal{S}$ a *constrained* key $k_{\mathcal{S}'}$ allowing to evaluate the PRF on all elements (and only them) of $\mathcal{S}'$.

A general framework. In [CHL05], a wallet of value N is associated with two secret seeds s and t and a subset $\mathcal{S}$ of size N. For each wallet, it is therefore possible to generate N pseudo-random pairs $(F_s(x), F_t(x))$. The first part, $F_s(x)$ is the coin serial number whereas the second part $F_t(x)$ roughly acts as a one-time pad on the spender's identity. If some user tries to double-spend some of its coins, then it must necessarily reuse some pair $(F_s(x), F_t(x))$ (suitable zero-knowledge proofs ensure this fact). In such a case, the same serial number $F_s(x)$ appears twice in the bank database, acting as a fraud alert. Moreover, the reuse of the same value $F_t(x)$ in different transactions enables the bank to remove this mask and so to recover the concealed identity by running an appropriate procedure.

It is possible to prove that a defrauder will necessarily be identified by this procedure. Unfortunately, there is no equivalence here and an adversary could trick the latter into returning an identity that was not involved in the spendings. Intuitively, it stems from the fact that collisions could occur between the outputs of the PRF and that two transactions involving different users could be considered as double-spendings.

To solve the first problem, we explicitly require (and actually define) collision resistance from F, meaning that it should be hard to find two seeds s and s' and two elements x and x' such that $F_s(x) = F_{s'}(x')$. This way, we can ensure that, except with negligible probability, the same wallet is involved in transactions generating the same serial numbers.

Two different approaches. To solve the second problem, we need to modify the way the serial numbers are constructed. Concretely, we now want the latter to depend on the spender's identity. However, we identify two different approaches to achieve this, leading to two different frameworks.

Our first framework is based on key-homomorphic PRFs [BFP+15], *i.e.* PRFs such that $F_{s \cdot s'}(x) = F_s(x) \cdot F_{s'}(x)$. As in [CHL05], our serial number contains the element $F_s(x)$ but we now also add the element $F'_{s \cdot id}(x)$ where F' is another PRF and id is some element depending on the spender's identity. If two serial numbers are equal, the collision resistance of F ensures that they were generated using the same seed s and the same element x. Since $F'_{s \cdot id}(x) = F'_s(x) \cdot F'_{id}(x)$ we can conclude that the value $F'_{id}(x)$ is the same in both transactions and so is id thanks to the collision resistance of F'.

Our second framework is based on delegatable PRFs [KPTZ13b], *i.e.* PRFs such that a constrained key $k_{\mathcal{S}_1}$ can be derived from $k_{\mathcal{S}_0}$ for any subset $\mathcal{S}_1$ of $\mathcal{S}_0$. Here again we will force the serial numbers to depend on the spender's identity, this time by defining the second part of these elements as $ID \cdot F_s(x)$.

From theory to practice. Obviously, there are a lot of technical details to address, and the formalization of our frameworks is indeed rather complex, but the security essentially relies on this simple intuition. It then now essentially remains to find such suitable PRFs.

In [BPS19], we describe concrete instantiations of pseudo-random functions satisfying all these conditions, but all of them rely on number-theoretic assumptions and so cannot resist to quantum computers. In particular, it seems that existing lattice-based PRFs do not fulfil all these requirements. Concretely, this means that deriving a lattice-based e-cash system from our frameworks requires some additional work.

3.3 Open problems

Based on this work, ORA and ENSL have decided to work together on two different axes.

1. The first one is to construct a new lattice-based pseudo-random function that would achieve all the properties defined in [BPS19]. The advantage of this strategy is that we could leverage the results from [BPS19] to directly derive a divisible e-cash system from the resulting pseudo-random function. We plan to study both approaches so as to consider all possibilities for the design of a secure and efficient divisible e-cash system.

2. The second axe is to depart significantly from the frameworks of [CHL05, BPS19] and thus propose a new construction that would only use lattice-friendly tools. The challenge is much more important since it prevents us from using decades of work on e-cash but we believe that, in the end, it could lead to more efficient systems.

4 E-Democracy

4.1 Introduction

As for e-cash, e-voting is an emulation of traditional voting systems. e-voting allows more accurate and fast vote counts, reduces the logistic cost of organizing an election and also offers specific mechanisms for voters with disabilities to cast their votes independently. In particular, Internet voting systems provide voters with the chance to cast their votes from anywhere. Requirements for Internet voting systems include privacy (voters are given the opportunity to cast their vote privately and it should be impossible to link the content of a vote to the identity of the voter who has cast it) and verifiability. At the same time, it has to be ensured that only eligible voters can cast a vote, and that only one vote per voter is counted.

As shown in deliverable D5.1, the literature on e-voting is very important, and there are several ways to design such system. However, the most common system, as being at the same time secure, efficient and easy to deploy can be described as follows.

There is a public encryption key PK for the election, whose corresponding secret key SK is distributed in a (t, n)-threshold way among a set of n members of an election authority. The participation of at least t of these authorities is needed in order to decrypt any ciphertext $C = \mathsf{Enc}_{PK}(m)$. Each voter i of the election has a pair of signing / verification keys (sk_i, vk_i). In the voting phase, the voter i encrypts and signs his chosen option m_i, leading to a tuple (C_i, σ_i), where $C_i = \mathsf{Enc}_{PK}(m)$ and σ_i is a signature on C_i, computed using sk_i. To avoid replay attacks (and optionally to avoid processing incorrect votes) some zero-knowledge proofs of knowledge may be required before accepting (C_i, σ_i) in the ballot box, for instance proofs of knowledge of the m_i encrypted inside C_i and proofs of the fact that (the secret element) m_i is a valid voting option.

If these proofs are accepted and the signature σ_i is verified as valid, the ciphertext C_i is added to the ballot box and the signature is removed. This leads to a list $L = \{C_i\}_{i \in \mathcal{I}}$ of valid encrypted votes, the final output of the voting phase. At this point, the first part of the tally phase, the mixing phase, starts: it consists of a sequential series of mixing nodes, each one applying a random and secret permutation, and randomization of each ciphertext in the list L, in a verifiable way: everybody should be able to verify the correctness of each mixing step, that is, that the output list of ciphertexts is really a randomization and permutation of the input list of ciphertexts. This is known as a verifiable shuffling process. As long as one of this mixing nodes behaves honestly, the mixing phase successfully breaks the link between the voters (or their verification keys vk_i), and the final list of ciphertexts $L' = \{C_j\}_{j \in \mathcal{I}}$ that will be actually decrypted, in the tally phase.

Finally, the last part of the tally phase consists of the mixing and decrypting of every ciphertext $C_j \in L'$. By using a threshold mechanism, one ensures that privacy and anonymity of the votes and voters is preserved even against an attacker who controls up to $t - 1$ of the (m members of the) election authority, and all-but-one of the mixing nodes. In the threshold decryption process, at least t election authorities must participate and jointly decrypt every $C_j \in L'$, proving somehow that they have correctly done their part of decryption (this may involve new zero-knowledge proofs). The list of resulting plaintext is returned as the final output of the election.

4.2 Ongoing Works and Related Results

4.2.1 Overall strategy

Our initial idea to come up with a lattice-based e-voting system is then to make use of this approach, taking advantage of decades of work on e-voting.

As discussed in the scientific group sessions of the project, specifically in the general scientific meeting (Amsterdam, May 2019) and the specific scientific meeting for electronic Voting (Barcelona, September 2019), the three main cryptographic challenges in the design of such a lattice-based electronic voting system are: (1) zero-knowledge systems for some specific languages related to encryption and decryption of a lattice-based scheme, (2) the verifiable shuffling of lattice-based ciphertexts, and (3) the secure and efficient threshold decryption of lattice-based ciphertexts.

For (1), there are several partners of PROMETHEUS working in the area of lattice-based zero-knowledge proofs of knowledge, inside WP4 of the project. Some of the results obtained therein will be directly applied to the electronic voting setting. For (2), researchers of UPC and Scytl have published a paper [CMM19], presented at workshop VOTING'19, proposing and analysing a specific protocol for verifiable shuffling of lattice-based ciphertexts, that we describe in more detail in next section. These same researchers are trying to obtain alternative protocols for verifiable shuffling, improving the size of the proof of correctness that each mixing node publishes; this is described in the section on ongoing work, below. Also members of IDC Herzliya are working on this same problem. Finally, for (3), researchers of different partners of PROMETHEUS are currently working on the design and analysis of specific mechanisms for threshold decryption of ring LWE ciphertexts, as we briefly describe in the section on ongoing work, below. Threshold decryption is a desirable functionality not only for electronic voting, but also for other use cases considered in this project, like Cyber Threat Intelligence.

4.2.2 Proof of a Shuffle

The verifiable protocol to shuffle N ciphertexts that is designed and analyzed in [CMM19] is the first proposal of a verifiable shuffle scheme with full post-quantum security. It can be used as a key component in the design of a post-quantum secure e-voting system.

The protocol works for ciphertexts produced by the encryption scheme by Lyubashevsky et al. [LPR10], whose security is based on the Ring Learning With Errors (RLWE) problem. The design of the protocol is inspired by the shuffling paradigm of Bayer and Groth [BG12], which was instantiated with ElGamal ciphertexts (in the classical, and so quantum-unsafe, discrete logarithm setting). When considering RLWE-based ciphertexts, some parts of the Bayer-Groth paradigm have to be carefully modified: now encrypted votes belong to a ring (instead of a group), and randomization of elements needs to satisfy several constraints, to avoid that the global noise leads to incorrect decryption, that have to be verified.

In a bit more of detail, each mixing node takes as input a list of RLWE ciphertexts, then re-randomizes each one independently and applies a global permutation π to the resulting N ciphertexts. The node has to add a proof that this operation has been done correctly. The first step of the proof, where the first difference with [BG12] lies, consists of committing the re-encryption parameters in order to demonstrate that they meet certain constraints. This is done using the commitment scheme and the zero-knowledge proofs of knowledge proposed by Benhamouda et al. [BKLP15], which fit perfectly with the scenario based on the hardness of the RLWE problem.

The next step consists in proving knowledge of the permutation π. The general idea here is to prove that two sets contain the same elements. Following the Bayer-Groth paradigm, this is done by computing two polynomials, each of them having as roots the elements of each set, and proving that both polynomials are equal.

The last step will prove knowledge of the re-encryption parameters, and this introduces another difference between Bayer and Groth's protocol and the lattice-based one. While they demonstrate that there exists a linear combination of the parameters such that an equality holds, here a different technique is needed, since the re-encryption parameters in a RLWE re-encryption scheme are taken from an error distribution and a linear combination of them would imply that the error grows uncontrollably, causing decryption errors. In particular, some compression techniques of Bayer-Groth cannot be applied here, which leads to a final proof of length $\mathcal{O}(N)$, where N, the cardinality of $\mathcal{I}$, is the number of ciphertexts that are shuffled.

The security of the whole mixing protocol is formally proved in detail, and is based on the hardness of the RLWE problem. The paper [CMM19] was presented in February 2019, in the workshop VOTING'19, a satellite workshop of the conference on Financial Cryptography. The final version will be published in an upcoming volume of the Lecture Notes in Computer Science series (Springer).

4.3 Open Problems and First Approaches

4.3.1 Verifiable Mixing

On the theoretical side, the above result gives a very long proof, which may be a problem for a practice instantiation. To improve this length, we need to consider alternative protocols for the verifiable proof of a correct shuffling in the lattice-based setting. Two teams are working on this: one with IDC and UR1 members and another one with UPC and Scytl members.

One idea is to use the results in [BBC$^+$18]: here ZKPK to prove knowledge of an input that satisfies a circuit C are provided, where the size of the proof is $\mathcal{O}\left(\sqrt{\lambda M \log^3(M)}\right)$, M being the number of gates of the circuit C and λ the security parameter.

The idea would be to apply this technique to the following circuit C: the public inputs are the two lists $L = \{C_i\}_{1 \leq i \leq N}$ and $L' = \{C'_j\}_{1 \leq j \leq N}$ of initial and shuffled ciphertexts, where $j = \pi(i)$ is some secret permutation chosen by the shuffling node; the private inputs are the re-randomization values r_i such that $C'_i = \texttt{Re-Randomize}(C_i, r_i)$; the output of the circuit on these inputs is one if and only if $\texttt{Re-Randomize}(L, \{r_i\}) = L'$, which is checked for instance by lexicographically ordering the two lists (via QuickSort) and comparing elements one by one. The number of gates of this circuit is $M = \mathcal{O}(N \log(N))$, so the results in [BBC$^+$18] should lead to a proof of correctness of a shuffle with size $\mathcal{O}\left(\sqrt{\lambda N \log^4(N)}\right)$, sub-linear in the number N of shuffled ciphertexts.

Our plan is then to implement these two alternative protocols (from both Scytl-UPC and IDC-UR1), and then compare them with the one in [CMM19], both theoretically and in practice, for specific (and realistic) values of N.

4.3.2 Threshold Decryption

For the task of threshold decryption of lattice-based ciphertexts, we are searching for a simple and efficient enough protocol which can be used in the specific setting of our whole lattice-based electronic voting system; in particular, the ciphertexts to be decrypted are produced by the encryption scheme by Lyubashevsky et al. [LPR10], with particularly chosen parameters so that the verifiable shuffling in [CMM19] (or the upcoming improvements) can be safely applied.

UPC and Scytl are currently considering different possibilities. The starting point is the paper [BGG$^+$18], where several solutions for the problem of threshold decryption are discussed and detailed, in the lattice-based setting. One of the solutions enjoys the feature that the efficiency (for instance, the length of the shares of the decryption key that the servers must hold) is independent of the threshold t and the number n of servers. However, this solution is conceptually quite complicated, in particular it involves ideas from fully homomorphic encryption, like bootstrapping. A simpler solution, also discussed in [BGG$^+$18], consists in considering secret sharing schemes where the reconstruction coefficients can be only 0 or 1. This solution produces shares that depend on t, but this may be not a serious problem in particular cases where this value of t is small; this seems to be the case in the specific e-voting prototype that will be designed in WP6 of the project, where the idea is to have small values (at most 5) for both n and t.

We plan to specify this protocol for the case of RLWE ciphertexts compatible with the shuffling protocol of [CMM19] and analyze its security in detail, as the first step, and then implement and test the protocol as the second step. We stress that, to provide public verifiability to the electronic voting system, each decryption server, which uses his secret share to compute a partial decryption of a ciphertext C'_j must also provide a zero-knowledge proof of the fact that this partial decryption has been correctly done. The security analysis of the threshold decryption protocol must take into account these zero-knowledge proofs, as well.

4.4 Other Approaches

Mix-nets based e-voting is not the only way to proceed. One can find several other ways to do in the literature. Within PROMETHEUS, we also plan to consider the following two possibilities.

1. Homomorphic encryption based. A homomorphic encryption scheme permits to perform operations over encrypted data. In the case of e-voting, the idea is for each voter to encrypt his/her vote using an additively homomorphic encryption scheme. Then, the talliers can use such homomorphic property to “add” all the votes and then obtain the encrypted result. We then use again a threshold mechanism during decryption to ensure privacy and anonymity of the votes and voters, and obtain the final decrypted result of the election. For obvious security reasons, it is also necessary to add some zero-knowledge proofs to prove the validity of a vote, or that the decryption has been done correctly. Having a lattice-based homomorphic encryption is not so hard, but it remains to design a compatible threshold mechanism and the additional zero-knowledge proofs to obtain a fully secure lattice-based system.

2. Anonymous signature based. As explain before (see the anonymous credential section), group and blind signature are cryptographic primitives which purpose is to provide the anonymity of users, and then break the link between e.g., a voter and his/her vote. Again, it is necessary to add a threshold encryption scheme (to prevent partial results) and zero-knowledge proofs (for vote’s validity and result verifiability) that are compatible with the used lattice-based group or blind signature scheme so as to obtain a fully secure lattice-based e-voting scheme.

5 Conclusion

One of the final goal of PROMETHEUS is to design and implement several demonstrators to show how mature is lattice-based cryptography in the context of privacy-preserving protocols. Most of those demonstrators will be based on the work done within WP5. In the first half of the project, we have listed the requirements for those use cases, study and improve the basic cryptographic primitives useful for those use cases (see in particular deliverable D4.2) and study the way we can design such privacy-preserving protocols. More precisely, this has been done within three different domains: anonymous credentials, e-cash, and e-democracy.

This deliverable has presented the main results that were produced by the PROMETHEUS consortium during the first two years within WP5. Some results have been published at highly respected cryptographic conferences and a lot a other works have been initiated, which is quite natural as WP5 highly depends on the results of the other technical work packages.

The challenges that WP5 faces in the remainder of the project have been given in the different sections above and the result will be describd at first in the use case specifications in one year (D6.5, D.6.6, D6.7 and D6.8) but also in the final version of this document (D5.4 on Final results on privacy-preserving cryptographic protocols, due to M48).

References

[BBC⁺18] Carsten Baum, Jonathan Bootle, Andrea Cerulli, Rafaël del Pino, Jens Groth, and Vadim Lyubashevsky. Sub-linear lattice-based zero-knowledge arguments for arithmetic circuits. In *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part II*, pages 669–699, 2018.

[BFP⁺15] Abhishek Banerjee, Georg Fuchsbauer, Chris Peikert, Krzysztof Pietrzak, and Sophie Stevens. Key-homomorphic constrained pseudorandom functions. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 31–60. Springer, Heidelberg, March 2015.

[BG12] Stephanie Bayer and Jens Groth. Efficient zero-knowledge argument for correctness of a shuffle. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 263–280. Springer, Heidelberg, April 2012.

[BGG⁺18] Dan Boneh, Rosario Gennaro, Steven Goldfeder, Aayush Jain, Sam Kim, Peter M. R. Rasmussen, and Amit Sahai. Threshold cryptosystems from threshold fully homomorphic encryption. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 565–596. Springer, 2018.

[BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In *PKC*, pages 501–519. Springer, 2014.

[BKLP15] Fabrice Benhamouda, Stephan Krenn, Vadim Lyubashevsky, and Krzysztof Pietrzak. Efficient zero-knowledge proofs for commitments from learning with errors over rings. In *Computer Security - ESORICS 2015 - 20th European Symposium on Research in Computer Security, Vienna, Austria, September 21-25, 2015, Proceedings, Part I*, pages 305–325, 2015.

[BPS19] Florian Bourse, David Pointcheval, and Olivier Sanders. Divisible e-cash from constrained pseudo-random functions. In *Asiacrypt 2019*, 2019.

[BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Heidelberg, December 2013.

[CHKP12] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. *Journal of Cryptology*, 25(4):601–639, October 2012.

[CHL05] Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. Compact e-cash. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 302–321. Springer, Heidelberg, May 2005.

[CMM19] Núria Costa, Ramiro Martínez, and Paz Morillo. Lattice-based proof of a shuffle. *IACR Cryptology ePrint Archive*, 2019:357, 2019.

[dLS18] Rafaël del Pino, Vadim Lyubashevsky, and Gregor Seiler. Lattice-based group signatures and zero-knowledge proofs of automorphism stability. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 574–591. ACM Press, October 2018.

[GM18] Nicholas Genise and Daniele Micciancio. Faster gaussian sampling for trapdoor lattices with arbitrary modulus. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 174–203. Springer, Heidelberg, April / May 2018.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008.

[KPTZ13a] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer &#38; Communications Security*, CCS '13, pages 669–684, New York, NY, USA, 2013. ACM.

[KPTZ13b] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 669–684. ACM Press, November 2013.

[KY19] Shuichi Katsumata and Shota Yamada. Group signatures without NIZK: From lattices in the standard model. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 312–344. Springer, Heidelberg, May 2019.

[LLNW17] Benoît Libert, San Ling, Khoa Nguyen, and Huaxiong Wang. Zero-knowledge arguments for lattice-based PRFs and applications to E-cash. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part III*, volume 10626 of *LNCS*, pages 304–335. Springer, Heidelberg, December 2017.

[LLNW18] Benoît Libert, San Ling, Khoa Nguyen, and Huaxiong Wang. Lattice-based zero-knowledge arguments for integer relations. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 700–732. Springer, Heidelberg, August 2018.

[LNTW19] Benoît Libert, Khoa Nguyen, Benjamin Hong Meng Tan, and Huaxiong Wang. Zero-knowledge elementary databases with more expressive queries. In Dongdai Lin and Kazue Sako, editors, *PKC 2019, Part I*, volume 11442 of *LNCS*, pages 255–285. Springer, Heidelberg, April 2019.

[LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Heidelberg, May / June 2010.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, April 2012.

[Rüc10] Markus Rückert. Lattice-based blind signatures. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 413–430. Springer, Heidelberg, December 2010.

[YAZ$^+$19] Rupeng Yang, Man Ho Au, Zhenfei Zhang, Qiuliang Xu, Zuoxia Yu, and William Whyte. Efficient lattice-based zero-knowledge arguments with standard soundness: Construction and applications. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 147–175. Springer, Heidelberg, August 2019.