



HAL
open science

VESPo: Verified Evaluation of Secret Polynomials

Jean-Guillaume Dumas, Aude Maignan, Clément Pernet, Daniel S. Roche

► **To cite this version:**

Jean-Guillaume Dumas, Aude Maignan, Clément Pernet, Daniel S. Roche. VESPo: Verified Evaluation of Secret Polynomials. 2021. hal-03365854v1

HAL Id: hal-03365854

<https://hal.science/hal-03365854v1>

Preprint submitted on 5 Oct 2021 (v1), last revised 13 Mar 2023 (v5)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

VESPo: Verified Evaluation of Secret Polynomials

(with application to low-storage dynamic proofs of retrievability)

Jean-Guillaume Dumas* Aude Maignan* Clément Pernet* Daniel S. Roche †

October 5, 2021

Abstract

We consider the problem of efficiently evaluating a secret polynomial at a given public point, when the polynomial is stored on an untrusted server. The server performs the evaluation and returns a certificate, and the client can efficiently check that the evaluation is correct using some pre-computed keys. Our protocols support two important features: the polynomial itself can be encrypted on the server, and it can be dynamically updated by changing individual coefficients cheaply without redoing the entire setup. As an important application, we show how these new techniques can be used to instantiate a Dynamic Proof of Retrievability (DPoR) for arbitrary outsourced data storage that achieves low server storage size and audit complexity. Our methods rely only on linearly homomorphic encryption and pairings, and preliminary timing results indicate reasonable performance for polynomials with millions of coefficients, and efficient DPoR with for instance 1TB size databases.

1 Introduction

Verifiable computing.

Verifiable computing, first formalized by [23], consists in delegating the computation of some function to an untrusted server, who must return the result as well as a proof of its correctness. Generally, the verification of a correct result should be much less expensive than computing it directly and result in a provably low probability that the result is incorrect. While certified and verified computation protocols date back decades, the practical need for efficient methods is especially evident in cloud computing, wherein a low-powered device such as a mobile phone may wish to outsource expensive and critical computations to an untrusted, shared-resource commercial cloud provider.

The extensive literature on verifiable computation protocols can be divided into general-purpose computations — e.g., of an arbitrary algebraic circuit — and more limited and (hopefully) efficient special-purpose computations of certain functions. In the latter category, one of the most important problems is Verifiable Polynomial Evaluation (VPE), where a client wishes to outsource the evaluation of a univariate polynomial P at a given point x and efficiently verify the result.

Verifiable Polynomial Evaluation.

A VPE scheme is conventionally composed of three algorithms. First, a client runs $\text{Setup}(P)$ to compute some public representation of P (which may be stored on the server) as well as some private information which will be used to verify later evaluations. This step may be somewhat expensive, but only needs to be performed once.

The second algorithm, $\text{Eval}(x, \alpha)$, is run by the server using a public evaluation point x , as well as possibly some additional information α provided by the client. The server produces the evaluation $y = P(x)$ as well as some proof or certificate β that this evaluation is correct.

*Université Grenoble Alpes, Laboratoire Jean Kuntzmann, UMR CNRS 5224, Grenoble INP. 700 avenue centrale, IMAG — CS 40700, 38058 Grenoble, France. \{Jean-Guillaume.Dumas,Aude.Maignan,Clement.Pernet\}@univ-grenoble-alpes.fr.

†United States Naval Academy, Annapolis, Maryland, United States. roche@usna.edu.

Finally, the third algorithm, $\text{Verify}(y, \beta)$, is run by the client to check the correctness of the evaluation. This verification should be *always correct* and *probabilistically sound*, meaning that an honest server can always produce a result y and proof β that will pass the verification, whereas an incorrect evaluation y will always fail the verification with high probability. Furthermore, the Verify algorithm should be efficient, ideally much cheaper in time and/or space than the computation itself.

Additional protocol features.

In the simplest case, the considered polynomial P is static and stored in cleartext by both the server and the client.

But constraints can then be added to this framework.

- **Polynomial outsourcing.**

When the client device has limited storage, or to facilitate multiple clients, the polynomial and its computation must be externalized. This can always be trivially achieved by storing all client secrets on the server via symmetric encryption and a saved cryptographic hash digest; the challenge is to do so while minimizing the communication costs required for the client to verify an evaluation.

- **Secret polynomial.**

In some cases, to guarantee data privacy, the polynomial has to be hidden from the server, or the client, or both. Typically, the polynomial will be stored under a fully- or partially-homomorphic encryption scheme, in such a way that the server can still compute the (necessarily encrypted) evaluation and certificate for verification.

This setting has been extensively studied in the literature, with both general-purpose protocols as well as some specific for verified polynomial evaluation.

- **Public verification**

The verification protocol is said to be *private* when only a party which holds the secrets derived during Setup can verify evaluations. That is, any potential verifiers (sometimes called *readers*) must be trusted not to divulge secret information to the untrusted server.

In many applications, it is desirable also to have untrusted verifiers, who can check the result of an evaluation without knowing any secrets. In this *public verification* setting, the client at setup time publishes some additional information, distributed reliably but insecurely to any potential verifiers, which may be used to check evaluations and proofs issued by the server.

- **Dynamic updates.**

Recall that the initial Setup protocol is expected to be much more costly than each Verify for the client. This creates a challenge when the client wishes to update only a few of the coefficients of the polynomial and later compute evaluations of this new, modified polynomial P' .

A *dynamic* VPE protocol allows for such updates efficiently. Namely, the client and server storing polynomial P for verified evaluation can engage in an additional $\text{Update}(c, d)$ protocol, which effectively updates $P(x)$ to $P(x) + cx^d$ for future evaluations, along with any secret and/or public verification information.

To the best of our knowledge, no prior work in the literature discusses dynamic updates for verified polynomial evaluation, which is especially challenging when the polynomial (as well as any update) needs to be hidden from the server. The importance of allowing efficient updates is motivated by our application to verifiable data storage, which we explain next.

Proofs of Retrievability.

One important application of VC in general, and VPE in particular, is to *Proofs of Retrievability* (PoR), somewhat overlapping with the problem of *Provable Data Possession* (PDP) [25, 6]. In these settings, a client wishes to store her data on an untrusted server, then verify (without full retrieval) that the server still stores the data intact. The crucial protocol is an **Audit**, wherein the client issues some challenge to

the server, then verifies the response using some pre-computed information to prove that the original data is still recoverable in its entirety.

A variety of tools have been employed to develop efficient PoR and PDP protocols, and some of these are based on verifiable computing. Generally speaking, the PoR Audit can consist in performing some verified computation over the stored data. Retrieval is proven by arguing that any sequence of successful audits can be used to recover the original data, e.g., by polynomial interpolation.

Note that this recovery mechanism is not actually crucial except to *prove* the soundness of the audit protocol; the important feature is how cheaply the audits can be performed by a server and resource-constrained client.

1.1 Our contributions

We consider a client/server scheme where the client has a low computing and a low storage capacity and where the server is untrusted. The client wants to externalize the storage of a polynomial P of degree d . Then she needs to outsource the computation of an evaluation $P(x)$ at a public point x . Thus she sends the point x in clear to the server, which has to respond with a value y and a proof that y satisfies $y = P(x)$. Our contributions are the following:

1. first, we present a new *unencrypted* Verifiable Polynomial Evaluation scheme which supports efficient *dynamic updates*, meaning that updating only a few coefficients of P does not require performing the whole setup phase again (Section 4 and Table 3). The polynomial is stored in cleartext on the server, and the technique used to provide a correct and sound protocol uses both Merkle trees and pairings. A Horner-like evaluation scheme is used to optimize the evaluation of the difference polynomial for the proof, and no secrets are required to perform the verification.
2. Second, we combine a linearly homomorphic cryptosystem with techniques from the first scheme in order to obtain an *encrypted, dynamic and private protocol for verifiable polynomial evaluation* (Section 5 and Table 5). That is, the polynomial is stored encrypted on the server, and efficient updates to individual coefficients can be performed. Note however, this scheme does not support public verification as this verification now requires some secrets from the client.
3. Finally, we use our new encrypted, dynamic VPE protocol to develop a new Proof of Retrieval scheme that supports small server storage, dynamic updates, and efficient audits (Section 6 and Table 9). With high probability, a single audit will succeed only if the entire data is fully recoverable.

A complete security definition of verifiable polynomial evaluation can be found in Section 2. This definition follows previous results, with the novel inclusion of an **Update** protocol.

Section 3 introduces the tools for verification of polynomial evaluation. A motivating example is presented in the form of a direct extension of the bilinear pairing scheme proposed by [26] now supporting an encrypted input polynomial (Section 3 and Table 2). Since the privacy of this protocol is not proven and it does not support neither public verifiability nor dynamic updates, it motivates the more involved contributions of Section 4 (for public verifiability and dynamicity, but on an unciphered polynomial) and of Section 5 (for dynamicity on a ciphered polynomial, but without public verifiability).

The efficiency of our protocols is measured by the computational complexity of the server-side **Eval** algorithm, the volume of persistent client storage, and the amount of communication and client-side complexity to perform a **Verify** or **Audit**. Improving on previously-known results, our protocols all have $\mathcal{O}(d)$ server-side computation, $\mathcal{O}(\log d)$ communication and client-side computation time, and $\mathcal{O}(1)$ client-side persistent storage. We include some practical timings in Sections 5.2 and 6.3.

In addition, our new PoR scheme requires only $o(d)$ extra server space on the server. This improves on [36] in terms of server storage and on [5] in terms of communication and client computation complexity for **Audit**.

1.2 Related work

While ours is the first work we are aware of which considers verifiable polynomial computation while hiding the polynomial from the server and allowing efficient dynamic updates, there have been a number of prior works on different settings of the VPE problem.

One line of work considers *commitment schemes* for polynomial evaluation [16, 14, 30, 21, 37, 12, 33, 19]. Here, the polynomial P is known only to the server, who publishes a binding commitment without revealing P itself. The verifier then confirms that a given evaluation is consistent with the pre-published commitment. The protocol of Kate *et. al.* [26], which fits in this model, introduced some of the techniques that we employ for our private verification algorithm, namely the notion of the difference polynomial (see Section 3). By contrast, our protocols aim to hide the polynomial P from the server, while the client is fully trusted.

Another line of work considers polynomial evaluation as an encrypted function, which can be evaluated at any chosen point. Function-hiding inner product encryption (IPE) [10, 28, 2] can be used to perform polynomial evaluation without revealing the polynomial P , but this inherently requires linear-time for the client, who must compute the first d powers of the desired evaluation point x .

Similarly, protocols using a Private Polynomial Evaluation (PPE) scheme have been developed in [13]. This primitive, based on an ElGamal scheme, ensures that the polynomial is protected and that the user is able to verify the result given by the server. Here the aim of the protocol is not to outsource the polynomial evaluation, but to obtain $P(x)$ and a proof without knowing anything about the polynomial. To check the proof, as with IPE the client has to produce a computation which is linear in the degree of P .

A third and more general approach which can be applied to the VPE problem is that of secure evaluation of arithmetic circuits. These protocols make use of fully homomorphic encryption (FHE) to outsource the evaluation of an arbitrary arithmetic circuit without revealing the circuit itself to the server. The VC Scheme of [23] is based on a Yao’s label construction. During the preprocessing phase, P is transformed into an arithmetic circuit. The circuit is garbled once in a setup phase and sent to the server. To later perform a verified evaluation, the client sends an encryption of x , the server computes $P(x)$ through the garbled circuit, and the client can verify the result in time proportional to the circuit depth, which for us is $\mathcal{O}(\log d)$

Using similar techniques, Fiore *et al.* and Elkhyaoui *et al.* [8, 18, 17] propose high-degree polynomial evaluations with a fully secure public verification solution. Very recently, Fiore *et al.* [19] propose a new protocol for more general circuits, using SNARKs over a quotient polynomial ring. In contrast to our work, these protocols use more expensive cryptographic primitives, and they do not consider the possibility of efficiently updating the polynomial.

Then, Proof of retrievability (PoR) and Provable data possession (PDP) protocols also have an extensive literature. PDPs generally optimize server storage and efficiency at the cost of soundness; a PDP audit may succeed even when a constant fraction of the data is unrecoverable. PoRs have stronger soundness guarantees, but at the expense of larger and more complicated server storage, often based on erasure codes and/or ORAM techniques.

State-of-the-art PoR protocols either incur a constant-factor blowup in server storage with poly-logarithmic audit cost [15, 36], or use negligible extra server storage space but require polynomial-time for audits on the client and server [35, 5]. A lower bound argument from [5] proves that some time/space tradeoff is inherent, although the proof does not distinguish between server and client computation time during audits.

2 Security properties and assumptions

A verifiable dynamic polynomial evaluation (VDPE) scheme consists of three algorithms: **Setup**, **Update**, **VEval**, between a client \mathcal{C} with state $st_{\mathcal{C}}$, a server \mathcal{S} with state $st_{\mathcal{S}}$ and a verifier \mathcal{V} with (potentially public) state $st_{\mathcal{V}}$.

- $(st_{\mathcal{C}}, st_{\mathcal{V}}, st_{\mathcal{S}}) \leftarrow \mathbf{Setup}(1^\kappa, P)$: On input of the security parameters and the polynomial P of degree d , outputs the client state $st_{\mathcal{C}}$, the verifier $st_{\mathcal{V}}$ and the server state $st_{\mathcal{S}}$.
- $\{(st'_{\mathcal{C}}, st'_{\mathcal{V}}, st'_{\mathcal{S}}), \mathbf{reject}\} \leftarrow \mathbf{Update}(i, \delta, st_{\mathcal{C}}, st_{\mathcal{V}}, st_{\mathcal{S}})$: On input of an index $i \in 0..d$, data δ , the client/verifier/server states $st_{\mathcal{C}}/st_{\mathcal{V}}/st_{\mathcal{S}}$, outputs new client/verifier/server states $st'_{\mathcal{C}}/st'_{\mathcal{V}}/st'_{\mathcal{S}}$, representing the polynomial $P + \delta X^i$, or **reject**.
- $\{z, \mathbf{reject}\} \leftarrow \mathbf{VEval}(st_{\mathcal{V}}, st_{\mathcal{S}}, r)$: On input of the verifier state $st_{\mathcal{C}}$, the server state $st_{\mathcal{S}}$ and an evaluation point r , outputs a successful evaluation $z = P(r)$ or **reject**.

The client may use random coins for any algorithm. This is the general setting for *public verification*, the idea being that for a *private verification*, the client will play the role of the verifier too and their states will be identical: $st_V = st_C$.

Adapted from [26], in order to take into account dynamicity, we propose the following security properties:

Definition 1. ($\text{Setup, Update, VEval}$) is a secure publicly verifiable polynomial evaluation scheme if it satisfies the following properties:

Correctness. Let $d \in \mathbb{N}$, (a_0, \dots, a_d) in a ring \mathcal{R} and $P(X) = \sum_{i=0}^d a_i X^i$, then:

$$\text{VEval}(\text{Setup}(1^\kappa, P), r) = P(r)$$

and for any $\delta \in \mathcal{R}$ and $0 \leq i \leq d$:

$$\text{VEval}(\text{Update}(i, \delta, st_C, st_V, st_S), r) = \text{VEval}(st_V, st_S, r) + \delta r^i$$

or **reject** has been returned by one of the protocols.

Soundness. The soundness requirement stipulates that the client can always detect (except with negligible probability) if any message sent by the server deviates from honest behavior. We use the following game between two observers \mathcal{O}_1 & \mathcal{O}_2 (respectively playing the roles of the client and the verifier), a potentially malicious server \mathcal{A} and an honest server \mathcal{S} , with the game:

1. \mathcal{A} chooses an initial polynomial P . \mathcal{O}_1 runs **Setup** and sends the initial server part, st_S , of the memory layout to both \mathcal{A} and \mathcal{S} ; and the verifier part to \mathcal{O}_2 .
2. For a polynomial number of steps $t = 1, 2, \dots, \text{poly}(\kappa)$, \mathcal{A} picks an operation op_t where operation op_t is either **Update** or **VEval**. \mathcal{O}_1 executes the **Update** operations with both \mathcal{A} and \mathcal{S} , while \mathcal{O}_2 executes the **VEval** operations with also both \mathcal{A} and \mathcal{S} .
3. \mathcal{A} is said to win the game, if any message sent by \mathcal{A} differs from that of \mathcal{S} and neither \mathcal{O}_1 nor \mathcal{O}_2 did output **reject**.

A VDPE scheme is sound, if no polynomial-time adversary \mathcal{A} has more than negligible probability in winning the above security game.

Privacy. A VDPE scheme is private, if no polynomial-time adversary has more than negligible probability in obtaining any coefficient of P , given access to the transcript of all exchanged messages for any number of runs of **Setup**, **Update** or **VEval**, and the associated server parts st_S , of the memory layout.

Definition 2. ($\text{Setup, Update, VEval}$) is a secure privately verifiable polynomial evaluation scheme if it verifies the **Correctness**, **Soundness** and **Privacy** requirements of [Theorem 1](#), where the verifier state st_V is included in the client state st_C and no polynomial-time adversary \mathcal{A} has more than negligible probability in winning the soundness security game when \mathcal{O}_1 also plays the role of \mathcal{O}_2 .

In [Section 5](#) we apply our new verifiable polynomial evaluation protocols to the development of a new Proof of Retrievability (PoR) scheme, provably achieving correctness, soundness, and retrievability for PoR. We follow the exact same security definition for PoR as in [\[5\]](#), which we will not restate here for the sake of brevity.

To prove the security of our protocols we rely on classical discrete logarithm and Diffie-Hellman like assumptions, all related to polynomial computations.

The first assumption, a decisional one, is the linearly independent polynomial assumption: informally it states that linearly independent polynomial evaluations “in the exponents” are merely indistinguishable from randomness. The formal version is recalled in [Theorem 4](#).

Then we need computational assumptions, including the hardness to compute discrete logarithms, in [Theorem 3](#), and polynomial extensions of the hardness to produce Diffie-Hellman-like secrets, in [Theorem 5](#), even with bilinear pairings, in [Theorem 6](#).

Definition 3 (Discrete Logarithm, DLOG, hardness assumption [27, Def. 9.63]). A discrete-logarithm problem is hard relative a group \mathbb{G} of group order p , a generator g and a randomly sampled element h of the group, if for any probabilistic polynomial-time algorithms \mathcal{A} , there exists a negligible function negl such that

$$\Pr[\mathcal{A}_{DLOG}(\mathbb{G}, g, h) = x \text{ s.t. } h = g^x] \leq \text{negl}(\log(p)). \quad (1)$$

Definition 4 (Linearly Independent Polynomial, LIP, assumption [1]). Let $\mathbb{G} = \langle g \rangle$ be a group of prime order p . The advantage of an adversary \mathcal{A} against the (n, d) -LIP security of \mathbb{G} , denoted $\text{Adv}_{\mathbb{G}}^{(n, d)\text{-lip}}(\mathcal{A})$, is the probability of success in the game defined in Table 1, with \mathcal{A} being restricted to make queries $P \in \mathbb{Z}_p[T_1, \dots, T_n]$ such that for any query P , the maximum degree in one indeterminate in P is at most d , and for any sequence (P_1, \dots, P_q) of queries, the polynomials (P_1, \dots, P_q) are always linearly independent over \mathbb{Z}_p .

Table 1: Game defining the (n, d) -LIP security for a group \mathbb{G} [1]

Init	Challenge(P)	Response(b')
$\vec{r} \xleftarrow{\$} \mathbb{Z}_p^n$	If $b == 0$	Return $b' == b$
$\beta \xleftarrow{\$} \mathbb{Z}_p$	Then Return $y \leftarrow g^{\beta P(\vec{r})}$	
$b \xleftarrow{\$} \{0, 1\}$	Else Return $y \xleftarrow{\$} \mathbb{G}$	

In fact, the LIP security reduces to the Matrix Diffie-Hellman assumption (MDDH) [1, Theorem 1], a generalization of the widely used decision linear assumption [22, 32, 3, 4, 7].

Definition 5 (t-Strong Diffie-Hellman, t-SDH, assumption [26, 11]). Let $\alpha \in \mathbb{Z}_p^*$. Given as input a $(t + 1)$ -tuple $\langle g, g^\alpha, g^{\alpha^2}, \dots, g^{\alpha^t} \rangle \in \mathbb{G}^{t+1}$, then for every adversary $\mathcal{A}_{t\text{-SDH}}$, the probability

$$\Pr[\mathcal{A}_{t\text{-SDH}}(g, g^\alpha, g^{\alpha^2}, \dots, g^{\alpha^t}) = \langle c, g^{\frac{1}{\alpha+c}} \rangle] \leq \text{negl}(\log(p)) \quad (2)$$

for any value of $c \in \mathbb{Z}_p \setminus \{-\alpha\}$.

Definition 6 (t-Bilinear Strong Diffie-Hellman, t-BSDH, assumption, from [24, 26]). Let $\alpha \in \mathbb{Z}_p^*$. Given as input a $(t + 1)$ -tuple $\langle g, g^\alpha, g^{\alpha^2}, \dots, g^{\alpha^t} \rangle \in \mathbb{G}^{t+1}$, in a group \mathbb{G} with a symmetric bilinear pairing e , for every adversary $\mathcal{A}_{t\text{-BSDH}}$, the probability

$$\Pr[\mathcal{A}_{t\text{-BSDH}}(g, g^\alpha, g^{\alpha^2}, \dots, g^{\alpha^t}) = \langle c, e(g, g)^{\frac{1}{\alpha+c}} \rangle] \leq \text{negl}(\log(p)) \quad (3)$$

for any value of $c \in \mathbb{Z}_p \setminus \{-\alpha\}$.

We will also use a public-key partially homomorphic encryption scheme where both addition and multiplication are considered. We need the following properties on the linearly homomorphic encryption function E (according to the context, we use E_{pk} or just E to denote the encryption function, similarly for the decryption function, D or D_{sk}): computing several modular additions on ciphered messages and modular multiplications but only between a ciphered message and a cleartext. For instance, Paillier cryptosystems [34, 9, 20] can satisfy these requirements, *via* multiplication in the ground ring, for addition of enciphered messages, and *via* exponentiation for ciphered multiplication.

$$D(E(m_1)E(m_2)) = m_1 + m_2 \quad (4)$$

$$D(E(m_1)^{m_2}) = m_1 m_2 \quad (5)$$

In terms of security, any *IND-CPA* scheme will be sufficient. Since we consider the semantic security of the cryptosystem, we assume that adversaries are probabilistic polynomial time machines. More precisely we consider *Malicious adversaries*: a corrupted server that controls the network and stops, forges or listens to messages in order to gain information or fool the client.

Finally, we will use a Merkle hash tree to allow verifications of updates and therefore need to use a cryptographic hash function with *Second pre-image resistance*.

3 Tools for the verification of a polynomial evaluation

Our first step is to define a verification protocol for polynomial evaluation that supports a ciphered input polynomial over a finite ring \mathbb{Z}_p . For this we propose an adaptation of both [26, 18]. It seems not to be sufficient to cipher the polynomial, or to check consistency in the exponents, so we propose to use both. For this, as in the former paper, we first need to define a difference polynomial that we will use to check consistency.

Definition 7. For a polynomial $P(X) \in \mathbb{Z}_p[X] = \sum_{i=0}^d p_i X^i$ of degree d , let its subset polynomials be: $T_{k,P}(X) = \sum_{i=k+1}^d p_i X^{i-k-1} = \sum_{j=0}^{d-1-k} p_{j+k-1} X^j$.

Proposition 8. Let $Q_P(Y, X) = \frac{P(Y) - P(X)}{Y - X}$ be the difference polynomial of a polynomial P ; then

$$Q_P(Y, X) = \frac{P(Y) - P(X)}{Y - X} = \sum_{i=1}^d p_i \sum_{k=0}^{i-1} Y^{i-k-1} X^k = \sum_{k=0}^{d-1} T_{k,P}(Y) X^k \quad (6)$$

Proof. As $Y^i - X^i = (Y - X)(\sum_{k=0}^{i-1} Y^{i-k-1} X^k)$, we obtain that

$$Q_P(Y, X) = \sum_{i=1}^d p_i \sum_{k=0}^{i-1} Y^{i-k-1} X^k = \sum_{k=0}^{d-1} X^k \left(\sum_{i=k+1}^d p_i Y^{i-k-1} \right).$$

□

□

This identity relates two evaluations of P : $P(Y) = P(X) + (Y - X)Q_P(Y, X)$, which allow one to verify $z \stackrel{?}{=} P(r)$ by checking, for a secret s , that:

$$P(s) = z + (s - r)Q_P(s, r) \quad (7)$$

For this, let E, D be the encryption and decryption functions of a partially homomorphic cryptosystem, supporting addition of two ciphertexts and multiplication of ciphertext by a cleartext, as in Equations (4) and (5). Therefore it is possible to evaluate a ciphered polynomial at a clear evaluation point, using powers of the evaluation point: for $x = [1, r, r^2, \dots, r^d]$, denote by $E(P)^\top \square x = \prod_{i=0}^d E(p_i)^{r^i} = E(P(r))$, the homomorphic polynomial evaluation.

Remark 9. An implementation with Paillier cryptosystem of the evaluation $P(r)$ in a modular ring \mathbb{Z}_m , thus providing the functionalities of Equations (4) and (5), requires some care: indeed these equations are satisfied modulo an RSA composite number N not equal to m .

More precisely, Paillier cryptosystem will provide $D(E(P(r))) \equiv (\sum_{i=0}^d p_i r^i) \pmod{N}$. Thus a possibility to recover the correct value, is to precompute $r^i \pmod{m}$, thus use the following Algorithm 1, and require that:

$$(d + 1)(m - 1)^2 < N \quad (8)$$

Algorithm 1 Homomorphic modular polynomial evaluation with a different Paillier modulus

Input: An integer $r \in [0..m - 1]$;

Input: A Paillier cryptosystem (E, D) with modulus $N > (m - 1)^2$.

Input: $(E(p_0), \dots, E(p_d)) \in \mathbb{Z}_N^{d+1}$, such that $\forall i, p_i \in [0..m - 1]$ and $d < \frac{N}{(m-1)^2} - 1$.

Output: $c \in \mathbb{Z}_N$ such that $D(c) \pmod{m} \equiv P(r) \pmod{m} \equiv \sum_{i=0}^d p_i r^i \pmod{m}$.

1: let $x_0 = 1$ and $c = E(p_0)$;

2: **for** $i = 1$ **to** d **do**

3: let $x_i \equiv x_{i-1} \cdot r \pmod{m}$;

{Now $x_i \in [0..m - 1]$ }

4: let $c \leftarrow c \cdot E(p_i)^{x_i}$;

5: **end for**

6: **return** c .

Proof. If $0 \leq p_i \leq (m-1)$, then as x_i is considered as an integer between 0 and $m-1$, then $0 \leq \sum_{i=0}^d p_i x_i \leq (d+1)(m-1)^2 < N$ by the constraints on d and N . Therefore $\sum_{i=0}^d p_i x_i \bmod N = \sum_{i=0}^d p_i x_i \in \mathbb{Z}$ and now $D(c) \bmod m = \sum_{i=0}^d p_i x_i \bmod m \equiv P(r)$. \square

Similarly, if $H = [h_i] = [g^{a_i}]$, denote by $H \odot x = \prod_{i=0}^d h_i^{x_i} = g^{\sum a_i x_i}$ the dot-product in the exponents. Then [Table 2](#) shows how the server produces the evaluation via the partially homomorphic cipher and how this evaluation is bound to be correct by the consistency check in the exponents.

Table 2: Verifiable Ciphered Polynomial Evaluation

	Server	Communications	Client
Setup		\mathbb{G}, \mathbb{G}_T groups of order p symmetric pairing e generators $g, e(g; g)$ $\xleftarrow{W, H}$	$P \in \mathbb{Z}_p[X], 1 \leq d^\circ(P) \leq d$ $s \xleftarrow{\$} \mathbb{Z}_p$ $W \leftarrow E(P), \mathcal{K} \leftarrow g^{P(s)}$ $H \leftarrow [g^{T_{k,P}(s)}]_{k=0..d-1}$ Discard P, W, H
VEval	$x \leftarrow [1, r, r^2, \dots, r^d]^\top$ $\zeta = W^\top \square x$ $\xi = H^\top \odot x$	\xleftarrow{r} $\xrightarrow{\zeta, \xi}$	$r \xleftarrow{\$} \mathbb{Z}_p$ $e(\xi; g^{s-r})e(g; g)^{D(\zeta)} \stackrel{?}{=} e(\mathcal{K}; g)$

Proposition 10. *The protocol of [Table 2](#) is correct and sound (verifiable).*

Proof. Correctness. First, $\zeta = W^\top \square x = \prod_{i=0}^d E(p_i)^{r^i} = E(P(r))$. Then, second, $\xi = H^\top \odot x = \prod_{k=0}^{d-1} g^{T_{k,P}(s)r^k} = g^{Q_P(s,r)}$, by [Theorem 8](#). Therefore, the verification is that $e(g; g)^{Q_P(s,r)(s-r)+P(r)} \stackrel{?}{=} e(g; g)^{P(s)}$ and this is guaranteed by [Equation \(7\)](#).

Soundness. Let $\langle \mathbb{G}, g, g^s, g^{s^2}, \dots, g^{s^t} \rangle \in \mathbb{G}^{t+1}$ be a t-SDH instance. Let $[p_0, \dots, p_t] \xleftarrow{\$} \mathbb{Z}_p^{t+1}$ for a degree t polynomial and $d = t$. Then compute $W = E(P)$, $T_{k,P} = \sum_{i=k+1}^t p_i Y^{i-k-1} = \sum_{j=0}^{t-1-k} t_{k,j} Y^j$. Finally, homomorphically compute $\mathcal{K} = \langle g, g^s, g^{s^2}, \dots, g^{s^t} \rangle \odot [p_0, \dots, p_t]$ and $H = [h_k]$, where $h_k = \langle g, g^s, g^{s^2}, \dots, g^{s^{t-1-k}} \rangle \odot [t_{k,0}, \dots, t_{k,t-1-k}]$. These inputs are indistinguishable from random inputs to the protocol of [Table 2](#).

Finally, select a random evaluation point $r \neq s$, compute (ζ, ξ) and call an attacker of the **VEval** part of the protocol to get (ζ', ξ') such that $(D(\zeta'), \xi') \neq (D(\zeta), \xi)$, even though both are passing the verification. Now, if $D(\zeta') = D(\zeta)$, then $e(\xi; g^s/g^r)e(g; g)^{D(\zeta)} = e(\mathcal{K}; g)$ and $e(\xi'; g^s/g^r)e(g; g)^{D(\zeta)} = e(\mathcal{K}; g)$. Therefore, $\frac{\xi'}{\xi} = 1^{\frac{1}{s-r}} = 1$ and $\xi = \xi'$, contradicting the attacker result. Thus $D(\zeta') \neq D(\zeta)$ and we have:

$$\left(\frac{\xi}{\xi'} \right)^{\frac{1}{D(\zeta') - D(\zeta)}} = g^{\frac{1}{s-r}}. \quad (9)$$

This proves that the t-SDH $\langle -r, g^{\frac{1}{s-r}} \rangle$ is broken. \square

We see here that using a decipherable partially homomorphic function for the coefficients of P is required, otherwise one could not compute the exponentiation on ξ/ξ' in the soundness proof.

Several issues remain with this protocol: first it is not dynamic. Indeed, for a dynamic version, the problem is that updating only one coefficients of P requires to update up to $d-1$ coefficients of H . This work would be of the same order of magnitude as recomputing the whole setup. Second it is not fully hiding the coefficients of P as they are just put in the exponents without any masking, and we do not prove the privacy requirement. Third, the protocol is not fully publicly verifiable since the decryption key of the partially homomorphic system is required.

We incrementally solve the first two issues in the sequel of this paper and obtain a thus fully secure private protocol. We also are able to provide a dynamic protocol, publicly verifiable, but for an unciphered polynomial. Combining all three properties, that is, designing a similar dynamic protocol for ciphered polynomials, but publicly verifiable, remains an open question to us.

4 Outsourced verification of dynamic polynomial evaluation

4.1 Merkle trees for logarithmic client storage

To avoid storing the polynomial coefficients on the client side, we use a Merkle hash tree [31, 29]. Then it is sufficient to store the root of the Merkle tree. For our purpose, an implementation of such trees must just provide the following algorithms:

- $T \leftarrow \mathbf{MTCreatetree}(X)$ creates a Merkle hash tree from a database X .
- $r \leftarrow \mathbf{MTRootFromLeaves}(X)$ computes from scratch the root of the Merkle hash tree of the whole database X .
- $(a, L) \leftarrow \mathbf{MTElementAndPath}(i, X, T)$ is an algorithm providing the client with the requested leaf element a , together with the corresponding list L of Merkle tree uncles.
- $r \leftarrow \mathbf{MTRootFromPath}(i, a, L)$ computes the root of the Merkle hash tree from a leaf element a and the associated path of uncles L .
- $T' \leftarrow \mathbf{MTUpdateLeaf}(i, a, T)$ updates the whole Merkle tree T by changing the i -th leaf to be a .

The requirements are thus that:

$$\forall i, X, \mathbf{MTRootFromLeaves}(X) = \mathbf{MTRootFromPath}(i, \mathbf{MTElementAndPath}(i, X, \mathbf{MTCreatetree}(X))) \quad (10)$$

$$\forall i, a, X, \text{ Let } (b, L) \leftarrow \mathbf{MTElementAndPath}(i, X, \mathbf{MTCreatetree}(X)), \text{ and let } X' \leftarrow X \setminus \{(i, b)\} \cup \{(i, a)\}, \quad (11)$$

$$\text{then } \mathbf{MTUpdateLeaf}(i, a, \mathbf{MTCreatetree}(X)) = \mathbf{MTCreatetree}(X')$$

4.2 Public Dynamic unciphered Polynomial Evaluation

Here we propose a protocol capable of verifying the evaluation of a dynamic polynomial P which supports public verifiability. It consists in three algorithms (**Setup**, **Update**, **PVDUeval**) detailed Table 3. During the **Setup** algorithm, the Client sends the unciphered polynomial to the Server and deletes it to minimize the Server storage. The Client uses a random coin s to create some data to be published or to be sent to the server. We introduce a third part named the Verifier. The Verifier collects the published data and is authorized to run the **Read** and the **PVDUeval** algorithms. But she is not authorized to run the **Init** algorithm and s is not known by the Verifier.

Theorem 11. *The protocol of Table 3 is correct and sound.*

Proof. Correctness. First, $\xi = \prod_{i=1}^d \prod_{k=0}^{i-1} S_{i-k-1}^{p_i x_k} = g^{\sum \sum s^{i-k-1} p_i x_k} = g^{Q_P(r, s)}$ and, second, we have that:

$$\begin{aligned} e(\xi, \mathcal{K}_2/g^r) e(g, g)^\zeta &= e(\xi, g^{s-r}) e(g, g)^{P(r)} \\ &= e(g, g)^{Q_P(r, s)(s-r) + P(r)} = e(g, g)^{P(s)}. \end{aligned}$$

Hence we see that $e(\xi, \mathcal{K}_2/g^r) e(g, g)^\zeta = e(\mathcal{K}_1, g)$ and, therefore, the protocol is correct.

Soundness. Let $\langle \mathbb{G}, g, g^s, g^{s^2}, \dots, g^{s^t} \rangle \in \mathbb{G}^{t+1}$ be a t-BSDH instance. For the setup phase, set $d = t$ and randomly select $[p_0, \dots, p_t] \stackrel{\$}{\leftarrow} \mathbb{Z}_p^{t+1}$. Then set $S = \langle \mathbb{G}, g, g^s, g^{s^2}, \dots, g^{s^t} \rangle$. Finally compute $\mathcal{K}_1 = \langle g, g^s, g^{s^2}, \dots, g^{s^t} \rangle \odot [p_0, \dots, p_t]$. These inputs are indistinguishable from random inputs to the protocol of Table 3.

For any number of update phase, randomly select δ , receive p_i and L_i from the Server, compute $\mathcal{K}'_1 = \mathcal{K}_1 S_i^\delta$ and refresh r_p .

Finally, select a random evaluation point r , compute (ζ, ξ) and call an attacker of the **PVDUeval** part of the protocol to get (ζ', ξ') such that $(\zeta', \xi') \neq (\zeta, \xi)$, even though both are passing the verification.

Table 3: Public and Dynamic unciphered polynomial evaluation

	Server	Communications	Client
Setup	$T_P \leftarrow \text{MTCreatetree}(P)$ Store P, T_P, S	order p group \mathbb{G} pairing e to \mathbb{G}_T gen. $g, e(g; g)$ $\xleftarrow{P, S}$	$P \in \mathbb{Z}_p[X], 0 \leq d^\circ(P) \leq d$ Let $s \xleftarrow{\$} \mathbb{Z}_p$ $\mathcal{K}_1 \leftarrow g^{P(s)}, S \leftarrow [g^{s^k}]_{k=0..d-1}$ $r_p \leftarrow \text{MTRootFromLeaves}(P)$ Publish $\mathcal{K}_1, \mathcal{K}_2 \leftarrow g^s$ discard P, S .
Update	$(p_i, L_i) \leftarrow \text{MTElementAndPath}(i, P, T_P)$ $T_P \leftarrow \text{MTUpdateLeaf}(i, p_i + \delta, T_P)$	$\xleftarrow{i, \delta}$ $\xrightarrow{p_i, L_i}$	$r_P \stackrel{?}{=} \text{MTRootFromPath}(i, p_i, L_i)$ $\mathcal{K}_1 \leftarrow \mathcal{K}_1 \cdot g^{s^i \delta}$ $r_P \leftarrow \text{MTRootFromPath}(i, p_i + \delta, L_i)$ publish \mathcal{K}_1
PVDUeval	Form $x \leftarrow [1, r, r^2, \dots, r^d]^\top$ $\zeta \leftarrow P(r)$ $\xi \leftarrow \prod_{i=1}^d \prod_{k=0}^{i-1} S_{i-k-1}^{p_i x_k}$	\xleftarrow{r} $\xrightarrow{\zeta, \xi}$	$r \xleftarrow{\$} \mathbb{Z}_p$ $e(\xi, \mathcal{K}_2/g^r) e(g^\zeta, g) \stackrel{?}{=} e(\mathcal{K}_1, g)$

- if $\zeta' = \zeta$, then $\frac{\xi}{\xi'} = 1^{\frac{1}{s-r}} = 1$ and $\xi = \xi'$, contradicting the attacker result.
- otherwise, $\zeta' \neq \zeta$ and we have both $e(\xi', \mathcal{K}_2/g^r) e(g, g)^{\zeta'} = e(\mathcal{K}_1, g) = e(\xi, \mathcal{K}_2/g^r) e(g, g)^\zeta$. This gives $e(\frac{\xi'}{\xi}, g^{s-r}) = e(g^{\zeta-\zeta'}, g)$ and thus $e((\frac{\xi'}{\xi})^{s-r}, g) = e(g^{\zeta-\zeta'}, g)$. Finally:

$$\left(\frac{\xi}{\xi'} \right)^{\frac{1}{\zeta'-\zeta}} = g^{\frac{1}{s-r}}. \quad (12)$$

This proves that the t-BSDH

$\langle -r, g^{\frac{1}{s-r}} \rangle$ is broken. □ □

4.3 Efficient linear-time evaluation

As a first approach to evaluate our protocols, we consider that the cardinality of the coefficient domain is a constant. Therefore, we count as arithmetic operations in the field not only the usual addition, subtraction, multiplication and inversion, but also the exponentiations that are independent of the degree of the polynomial. We thus express our asymptotic complexity bounds in Table 4, only with respect to that degree d .

Table 4: Complexity bounds for the publicly verifiable dynamic and unciphered polynomial evaluation of Table 3 for a degree d polynomial.

		Server	Communication	Client
Storage		$\mathcal{O}(d)$		$\mathcal{O}(1)$
Comput.	Setup	$\mathcal{O}(d)$	$\mathcal{O}(d)$	$\mathcal{O}(d)$
	Update	$\mathcal{O}(\log(d))$	$\mathcal{O}(\log(d))$	$\mathcal{O}(\log(d))$
	PVDUeval	$\mathcal{O}(d)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$

Theorem 12. *The setup protocol of Table 3 requires $\mathcal{O}(d)$ arithmetic operations.*

The update protocol of Table 3 requires $\mathcal{O}(\log(d))$ arithmetic operations.

The verification protocol of Table 3 requires $\mathcal{O}(1)$ communications and arithmetic operations for the client, and $\mathcal{O}(d)$ arithmetic operations for the Server.

Proof. For the update phase, the client has to compute the root of the Merkle tree from the new value $p_i + \delta$ and the path L_i given by the server in $\mathcal{O}(\log(d))$. She also has to compute an exponentiation and a product in $\mathbb{Z}_p[X]$, this is in $\mathcal{O}(1)$.

For the verification phase, communications are just 3 group elements. The client work is only 2 pairing and 2 exponentiations and 1 product.

Now for the server. First, computing ζ is $d+1$ homomorphic multiplications and d additions. Second, the server has to compute $\xi = \prod_{i=1}^d \prod_{k=0}^{i-1} S_{i-k-1}^{p_i x_k} = \prod_{i=1}^d \left(\prod_{k=0}^{i-1} S_{i-k-1}^{r^k} \right)^{P_i}$. Therefore, one can use a Horner-like prefix computation: consider $t_0 = 1$, and $t_i = S_{i-1} \cdot t_{i-1}^r$, then $t_1 = S_0$, $t_2 = S_1 S_0^r$ and $t_i = S_{i-1} (S_{i-2} \dots (S_2 (S_1 S_0^r)^r) \dots)^r = \prod_{k=0}^{i-1} S_{i-k-1}^{r^k}$. Thus one can use the following Algorithm 2 to compute ξ .

Algorithm 2 Homomorphic prefix evaluation of the difference polynomial

```

1:  $\xi = 1; t = 1;$ 
2: for  $i = 1$  to  $d$  do
3:    $t \leftarrow S_{i-1} \cdot t^r;$ 
4:    $\xi \leftarrow \xi \cdot t^{P_i}.$ 
5: end for
6: return  $\xi.$ 

```

Thus, computing ξ requires at most $2d$ exponentiations and $2d$ multiplications. □ □

5 Fully private, dynamic and ciphered protocol for polynomial evaluation

So far we have a polynomial evaluation verification, that allows efficient updates of its coefficients. We now propose a scheme which combine the polynomial evaluation with the externalization of the polynomial itself. For this, two more ingredients are added: an efficient masking in the exponents in order to fulfill the hiding security property and an outsourcing of the (ciphered) polynomial itself in Section 5.1. This latter feature allows the client to not even store the polynomial and reduces her need for storage to a small constant number of field elements. For this we use Merkle hash trees presented in Section 4.1. They ensure the authenticity of the coefficient updates, with the storage of only one hash.

5.1 Private, dynamic, ciphered protocol

Now we add a masking of the polynomial coefficients in order to make the protocol hiding. For this we use Theorem 4: indeed, LIP security states that in a group \mathbb{G} of prime order, the values $(g^{P_1(s)}, \dots, g^{P_m(s)})$ are indistinguishable from a random tuple of the same size, when P_1, \dots, P_m are linearly independent multivariate polynomials of bounded degree and s is the secret. Therefore, in our modified protocol, the coefficients $g^{\beta \varphi^s}$ for a secret φ , are indistinguishable from a random tuple (g^{ρ_i}) of the same size since the polynomials X^j , $j = 1..m$ are independent distinct monomials.

The dynamic externalized polynomial evaluation scheme consist of the following algorithms Setup, Update and VEval between a client \mathcal{C} with state $st_{\mathcal{C}}$ and the server \mathcal{S} of state $st_{\mathcal{S}}$.

- $(st_{\mathcal{C}}, st_{\mathcal{S}}) \leftarrow \text{Setup}(1^k, P)$: on input of the security parameters and the polynomial P , outputs the client state $st_{\mathcal{C}}$ and the server state $st_{\mathcal{S}}$, as detailed in Algorithm 3.
- $\{(st'_{\mathcal{C}}, st'_{\mathcal{S}}), \text{reject}\} \leftarrow \text{Update}(i, \delta, st_{\mathcal{C}}, st_{\mathcal{S}})$: on input of an index $i \in 0..d$, the difference data δ , the client state $st_{\mathcal{C}}$ and the server state $st_{\mathcal{S}}$, outputs a new client state $st'_{\mathcal{C}}$ and a new server state $st'_{\mathcal{S}}$ (such that now the new i -th coefficient of the polynomial is $P'_i = P_i + \delta$, for P_i the previous i -th coefficient), or **reject**, as detailed in Algorithm 4.

Table 5: Private, Dynamic, Ciphered and logarithmic polynomial evaluation

	Server	Communications	Client
Setup	$T_W \leftarrow \text{MTCreatetree}(W)$ Store W, T_W, \bar{H}, S	order p group \mathbb{G} pairing e to \mathbb{G}_T , gen. $g, e(g; g)$	$P \in \mathbb{Z}_p[X], 1 \leq d^\circ(P) \leq d$ $s, \alpha, \beta, \varphi \xleftarrow{\$} \mathbb{Z}_p^*$ Let $\bar{P}(X) \leftarrow \sum_{i=0}^d (\alpha p_i + \beta \varphi^i) X^i$ $W \leftarrow E(P), \bar{H} \leftarrow [g^{\bar{P}^i}]_{i=1..d}$ $\bar{\mathcal{K}} \leftarrow g^{\bar{P}(s)}, S \leftarrow [g^{s^k}]_{k=0..d-1}$ $r_W \leftarrow \text{MTRootFromLeaves}(W)$ discard $P, \bar{P}, W, \bar{H}, S$.
Update	$\bar{h}'_i \leftarrow \bar{h}_i \cdot \Delta$ $(w_i, L_i) \leftarrow \text{MTElementAndPath}(i, W, T_W)$ $w'_i \leftarrow w_i \cdot e_\delta$ $T_W \leftarrow \text{MTUpdateLeaf}(i, w'_i, T_W)$	$\xleftarrow{i, e_\delta, \Delta}$ $\xrightarrow{w_i, L_i}$	$e_\delta \leftarrow E(\delta), \Delta \leftarrow g^{\alpha \delta}$ $r_W \stackrel{?}{=} \text{MTRootFromPath}(i, w_i, L_i)$ $w'_i \leftarrow w_i \cdot e_\delta, \bar{\mathcal{K}} \leftarrow \bar{\mathcal{K}} \cdot \Delta^{s^i}$ $r_W \leftarrow \text{MTRootFromPath}(i, w'_i, L_i)$
VEval	Form $x \leftarrow [1, r, r^2, \dots, r^d]^\top$ $\zeta \leftarrow W^\top \square x$ $\bar{\xi} \leftarrow \prod_{i=1}^d \prod_{k=0}^{i-1} e(\bar{H}_i; S_{i-k-1})^{x_k}$	\xleftarrow{r} $\xrightarrow{\zeta, \bar{\xi}}$	For $r \in \mathbb{Z}_p$ $c \leftarrow \beta \frac{(r\varphi)^{d+1} - 1}{r\varphi - 1}$ $\bar{\xi}^{s-r} e(g; g)^{\alpha D(\zeta) + c} \stackrel{?}{=} e(\bar{\mathcal{K}}; g)$

- $\{\xi, \text{reject}\} \leftarrow \text{VEval}(st_C, st_S, r)$: on input of the client state st_C , the server state st_S and an evaluation point r , outputs a successful evaluation $\xi = P(r)$ or **reject**, as detailed in [Algorithm 5](#).

Algorithm 3 Setup($1^\kappa, P$)

Input: $1^\kappa; p \in \mathbb{P}, P = \sum_{i=0}^d p_i X^i \in \mathbb{Z}_p[X]$;

Input: a partially homomorphic cryptosystem E/D satisfying [Equations \(4\)](#) and [\(5\)](#), for any dot-product of size $d+1$, modulo p .

Output: st_S, st_C .

- 1: Client: generates an order p group \mathbb{G} with generator g and non-degenerate pairing $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$;
 - 2: Client: generates a public/private key pair (pk, sk) for E/D ;
 - 3: Client: randomly selects $s, \alpha, \beta, \varphi \xleftarrow{\$} \mathbb{Z}_p^*$;
 - 4: Client: computes $\bar{P}(X) = \sum_{i=0}^d (\alpha p_i + \beta \varphi^i) X^i$, $W = E_{pk}(P)$, $\bar{H} = [g^{\bar{P}^i}]_{i=1..d}$, $\bar{\mathcal{K}} = g^{\bar{P}(s)}$ and $S = [g^{s^k}]_{k=0..d-1}$;
 - 5: Client: $r_W = \text{MTRootFromLeaves}(W)$; {root of the Merkle tree}
 - 6: Client: sends $pk, \mathbb{G}, g, \mathbb{G}_T, e, W, \bar{H}, S$ to the Server;
 - 7: Client: **return** $st_C \leftarrow (pk, sk, \mathbb{G}, g, \mathbb{G}_T, e, s, \alpha, \beta, \varphi, \bar{\mathcal{K}}, r_W)$;
 - 8: Server: $T_W \leftarrow \text{MTCreatetree}(W)$; {the Merkle tree}
 - 9: Server: **return** $st_S \leftarrow (pk, \mathbb{G}, g, \mathbb{G}_T, e, W, T_W, \bar{H}, S)$.
-

We have now in [Theorem 13](#), the complete result for the Dynamic Verified Evaluation of Secret Polynomials.

Theorem 13. *Under the security assumptions of [Section 2](#), the protocol composed of [Algorithms 3](#) to [5](#) (summarized in [Table 5](#)) is a fully secure verifiable polynomial evaluation scheme, as defined in [Theorem 1](#) and the complexity bounds of its algorithms are given in [Table 6](#).*

For the complexity bounds we still consider the cardinality of the coefficient domain to be a constant (so that, again, even exponentiations not involving the degree are considered constant) and we also consider that one encryption/decryption with the linearly homomorphic cryptosystem requires a number of arithmetic operations constant with respect to the degree.

Algorithm 4 Update(i, δ, st_C, st_S)

Input: $i \in [0..d]$, $a \in \mathbb{Z}_p$, st_S, st_C .**Output:** st'_S, st'_C or **reject**.

- 1: Client: computes $e_\delta = E_{pk}(\delta)$, $\Delta = g^{\alpha\delta}$;
 - 2: Client: sends i, e_δ, Δ to the Server;
 - 3: Server: $(w_i, L_i) \leftarrow \mathbf{MTElementAndPath}(i, W, T_W)$; *{extracts w_i and its uncles in the Merkle tree}*
 - 4: Server: $T'_W \leftarrow \mathbf{MTUpdateLeaf}(i, w_i \cdot e_\delta, T_W)$; *{updates the Merkle tree}*
 - 5: Server: sends w_i, L_i to the Client;
 - 6: Server: **return** $st'_S \leftarrow st_S \setminus \{T_W, w_i, \bar{h}_i\} \cup \{T'_W, w_i \cdot e_\delta, \bar{h}_i \cdot \Delta\}$;
 - 7: **if** $r_W = \mathbf{MTRootFromPath}(i, w_i, L_i)$ **then**
 - 8: Client: computes $\bar{K}' \leftarrow \bar{K} \cdot \Delta^{s^i}$;
 - 9: Client: computes $r'_W = \mathbf{MTRootFromPath}(i, w_i \cdot e_\delta, L_i)$;
 - 10: Client: **return** $st'_C \leftarrow st_C \setminus \{\bar{K}, r_W\} \cup \{\bar{K}', r'_W\}$.
 - 11: **else** *{else the stored root does not match the received element and uncles}*
 - 12: Client: **return reject**.
 - 13: **end if**
-

Algorithm 5 VEval(st_C, st_S, r)

Input: st_C, st_S and $r \in \mathbb{Z}_p$;**Output:** $z = P(r)$ or **reject**.

- 1: Client: computes $c \leftarrow \beta \frac{(r\varphi)^{d+1} - 1}{r\varphi - 1}$;
 - 2: Client: sends r to the Server;
 - 3: Server: homomorphically computes $\zeta = W^\top \square x = \prod_{i=0}^d w_i^{(r^i \bmod p)}$ *{via Equations (4) and (5), see also, e.g., Theorem 9 and Algorithm 1}*
 - 4: Server: $\bar{\xi} = 1_{\mathbb{G}_T}$; $t = 1_{\mathbb{G}}$;
 - 5: **for** $i = 1$ **to** d **do**
 - 6: Server: $t \leftarrow S_{i-1} \cdot t^r$;
 - 7: Server: $\bar{\xi} \leftarrow \bar{\xi} \cdot e(\bar{H}_i; t)$;
 - 8: **end for**
 - 9: Server: sends $\zeta, \bar{\xi}$ to the Client;
 - 10: Client: computes $z = D_{sk}(\zeta) \bmod p$;
 - 11: **if** $\bar{\xi}^{s-r} e(g; g)^{\alpha z + c} = e(\bar{K}; g)$ **then**
 - 12: Client: **return** z .
 - 13: **else**
 - 14: Client: **return reject**.
 - 15: **end if**
-

Proof. Correctness. We use the left hand side of Theorem 8 and Equation (7). Applying this to \bar{P} , we obtain that:

$$\bar{\xi} = \prod_{i=1}^d \prod_{k=0}^{i-1} e(\bar{H}_i; S_{i-k-1})^{x_k} = \prod_{i=1}^d \prod_{k=0}^{i-1} e(g^{\bar{P}^i}; g^{s^{i-k-1}})^{r^k} = e(g; g)^{Q_{\bar{P}}(s, r)}.$$

Denote by $G(Z) = \frac{Z^{d+1} - 1}{Z - 1}$. Now $\bar{P}(X) = \alpha P(X) + \beta G(X\varphi)$, then $c = \beta G(r\varphi)$ and thus $\bar{P}(r) = \alpha D(\zeta) + c$. Therefore the verification is indeed that

$$e(g; g)^{Q_{\bar{P}}(s, r)(s-r) + \bar{P}(r)} \stackrel{?}{=} e(g; g)^{\bar{P}(s)}.$$

Complexity bounds. In terms of storage, the client just has to store four elements mod p , that is s , α , β , and φ , together with one group element, \bar{K} ; the server has to store the polynomial ciphered twice, the ciphered powers of s and the Merkle tree for the ciphered polynomial: all this is $O(d)$. In terms of communications, during the Update phase the client sends one index and two group elements, while receiving one group element and the list of its $\log(d)$ uncles. During the VEval phase, only three elements are exchanged. Finally, in terms of computations, the server performs $O(d)$ operations for the Merkle tree

Table 6: Complexity bounds for verifiable dynamic and ciphered polynomial evaluation: for groups and prime fields of supposed constant order/cardinality, the asymptotics are here function of the degree d of the evaluated polynomial: storage units are given in number of group/field elements, computational operations are given in number of group/prime field arithmetic operations.

		Server	Communication	Client
Storage		$\mathcal{O}(d)$		$\mathcal{O}(1)$
Comput.	Setup	$\mathcal{O}(d)$	$\mathcal{O}(d)$	$\mathcal{O}(d)$
	Update	$\mathcal{O}(\log(d))$	$\mathcal{O}(\log(d))$	$\mathcal{O}(\log(d))$
	VEval	$\mathcal{O}(d)$	$\mathcal{O}(1)$	$\mathcal{O}(\log(d))$

generation at **Setup**; fetches $O(\log(d))$ uncles at **Update**; and $O(d)$ (homomorphic) operations at **VEval**, thanks to [Algorithm 2](#). For the client, **Update** requires $O(\log(d))$ arithmetic operations to compute the exponentiations s^i and a constant number of other arithmetic operations, independent of the degree. Similarly, computing $(r\varphi)^{d+1}$ also requires $O(\log(d))$ classical arithmetic operations and the rest is a constant number of operations that are independent of the degree.

Soundness. Let $\langle \mathbb{G}, g, g^s, g^{s^2}, \dots, g^{s^t} \rangle \in \mathbb{G}^{t+1}$ be a t-BSDH instance. For the setup phase, randomly select $\alpha, \beta, \varphi \neq 0$ and $[p_0, \dots, p_t] \stackrel{\$}{\leftarrow} \mathbb{Z}_p^{t+1}$. Then compute $W = E(P)$, $\bar{H} = g^{\bar{P}}$, and let $S = \langle \mathbb{G}, g, g^s, g^{s^2}, \dots, g^{s^t} \rangle$. Finally homomorphically compute $\bar{\mathcal{K}} = \langle g, g^s, g^{s^2}, \dots, g^{s^t} \rangle \odot [\bar{p}_0, \dots, \bar{p}_t]$. These inputs are indistinguishable from random inputs to the protocol of [Table 5](#). For any number of update phase, randomly select δ and compute $e_\delta = E(\delta)$ and $\Delta = g^{\alpha\delta}$. Also compute $\mathcal{K}' = \mathcal{K}S_i^{\alpha\delta}$. Finally, select a random evaluation point r , compute $(\zeta, \bar{\xi})$ and call an attacker of the **VEval** part of the protocol to get $(\zeta', \bar{\xi}')$ such that $(D(\zeta'), \bar{\xi}') \neq (D(\zeta), \bar{\xi})$, even though both are passing the verification. This means, again, that if $D(\zeta') = D(\zeta)$, then $\frac{\bar{\xi}}{\bar{\xi}'} = 1^{s-\frac{1}{r}} = 1$ and $\bar{\xi} = \bar{\xi}'$, contradicting the attacker result. Therefore it must be that $D(\zeta') \neq D(\zeta)$ and then, as $\alpha \neq 0$, we have again:

$$\left(\frac{\bar{\xi}}{\bar{\xi}'} \right)^{\frac{1}{\alpha(D(\zeta') - D(\zeta))}} = e(g; g)^{\frac{1}{s-r}}. \quad (13)$$

This proves that the t-BSDH $\langle -r, e(g; g)^{\frac{1}{s-r}} \rangle$ is broken.

Privacy. We show that the protocol is hiding both p_i and \bar{p}_i .

For \bar{p}_i first. Let $B = g^b$ be a DLOG instance. For the setup phase, randomly select $s, \alpha, \varphi \neq 0, d$ and $[p_0, \dots, p_d] \stackrel{\$}{\leftarrow} \mathbb{Z}_p^{d+1}$. Then compute $W = E(P)$, $\bar{h}_i = g^{\alpha p_i} B^{\varphi^i}$, $S = \langle \mathbb{G}, g, g^s, g^{s^2}, \dots, g^{s^t} \rangle$, and $\bar{K} = g^{\alpha P(s)} B^{G(s\varphi)}$. These inputs are indistinguishable from random inputs to the protocol of [Table 5](#). For any update phase, randomly select δ and compute $e_\delta = E(\delta)$ and $\Delta = g^{\alpha\delta}$. Also compute $\mathcal{K}' = \mathcal{K}S_i^{\alpha\delta}$. Such updates are indistinguishable from random updates to the protocol of [Table 5](#). Randomly select any number of evaluation points r and run the associated **VEval** phases, randomly alternated with update phases. Now suppose that an attacker can find from this transcript one coefficient \bar{p}_i : from this, then compute $b = (\bar{p}_i - \alpha p_i)/\varphi^i$ and the DLOG is broken.

For p_i , we proceed with a sequence of indistinguishable games:

1. Under LIP security [[1](#), [Theorem 3.1](#)], the parameter \bar{h}_i , or more precisely, $(E(p_i), g^{\alpha p_i + \beta \varphi^i})$, is indistinguishable from $(E(p_i), g^{\alpha p_i + \rho_i})$ for a random ρ_i . Therefore the protocol of [Table 5](#) is indistinguishable, as a whole, from the same protocol where $\beta \varphi^i$ is everywhere replaced by ρ_i , and c is (now inefficiently) computed as $\sum \rho_i r^i$.
2. Now we prove the hiding property of the latter. Let $Z = E(\omega)$ be the cipher of a secret ω . Randomly select d and $[u_0, \dots, u_d] \stackrel{\$}{\leftarrow} \mathbb{Z}_p^{d+1}$. Compute $W_i = E(\omega)E(u_i) = E(\omega + u_i)$. Randomly select α and randomly select h_i (so that $\rho_i = \log(h_i) - \alpha(\omega + u_i)$ exists, but remains unknown) for $i = 1..d$. Randomly select s and compute $\bar{\mathcal{K}} = H \odot [1, s, \dots, s^d]$. For any number of updates, randomly select δ , compute $e_\delta = E(\delta)$, so that $\delta = p'_i - p_i = (\omega + u'_i) - (\omega + u_i) = u'_i - u_i = \delta$. Therefore, compute $\Delta = g^{\alpha\delta}$ and $\bar{\mathcal{K}}' = \bar{\mathcal{K}}\Delta^{s^i}$. Alternatively run such updates with random **VEval** phases; all this is

indistinguishable from a normal transcript of the protocol. Now if from this transcript an attacker can find one p_j , then compute $\omega = p_j - u_j$ and the homomorphic cryptosystem is broken. \square

\square

5.2 Experiments

To assess the efficiency of our protocol, we implemented Table 5 using the following libraries¹: `gmp-6.2.1` for modular operations, `libpaillier-0.8` for Paillier’s cryptosystem and `pbcc-0.5.14` for pairings. We used a “type A” symmetric pairing over 256-bits group size².

To observe the effect of the chosen homomorphic systems (Paillier and the pairing), we ran the experiments, on a single core of a i7-6700 3.4GHz, in two sets: the first one with a low RSA modulus size of 1024 bits, and the second one with a RSA modulus size of 2048 bits.

Table 7: Comparative behaviors of pairings and Paillier system on the Server and Client sides with a 256-bits group size for the protocol of Table 5: on the client side, column ‘pows’ is the time to perform the left hand-side exponentiations (by $s - r$ and by $\ell = \alpha D(\zeta) + c$); column ‘c’ is the time to perform the geometric sum (the only part non constant in d); column ‘D’ is the time to perform the single Paillier’s deciphering; and column ‘e’ is the time for the right-hand side (one application of the pairing to \bar{K}); column ‘Horner’ is a witness simple polynomial evaluation in that group. Each experiment was performed 11 times and we report the median value, with a maximum variance lower than 5% between runs.

Degree	Paillier	Server Certif.		Client Verif.				Horner
		ζ	ξ	pows	c	D	e	
256	1024	0.10s	0.39s	0.2ms	<0.1ms	1.4ms	1.3ms	<0.1ms
512	1024	0.20s	0.78s	0.2ms	<0.1ms	1.5ms	1.3ms	0.1ms
1024	1024	0.39s	1.56s	0.2ms	<0.1ms	1.5ms	1.3ms	0.1ms
2048	1024	0.77s	3.12s	0.2ms	<0.1ms	1.4ms	1.3ms	0.3ms
4096	1024	1.54s	6.33s	0.2ms	<0.1ms	1.4ms	1.3ms	0.5ms
8192	1024	3.10s	12.41s	0.2ms	<0.1ms	1.4ms	1.3ms	1.0ms
16384	1024	6.19s	25.26s	0.2ms	<0.1ms	1.4ms	1.4ms	2.0ms
32768	1024	12.43s	51.04s	0.2ms	<0.1ms	1.4ms	1.3ms	4.0ms
65536	1024	24.78s	101.53s	0.2ms	<0.1ms	1.5ms	1.3ms	8.0ms
131072	1024	49.42s	202.50s	0.3ms	<0.1ms	1.7ms	1.5ms	16.2ms
256	2048	0.34s	0.39s	0.2ms	<0.1ms	10.2ms	1.3ms	<0.1ms
512	2048	0.68s	0.79s	0.2ms	<0.1ms	10.3ms	1.4ms	0.1ms
1024	2048	1.40s	1.55s	0.2ms	<0.1ms	10.5ms	1.3ms	0.1ms
2048	2048	2.74s	3.12s	0.2ms	<0.1ms	10.3ms	1.3ms	0.3ms
4096	2048	5.49s	6.20s	0.2ms	<0.1ms	10.1ms	1.3ms	0.5ms
8192	2048	10.96s	12.42s	0.2ms	<0.1ms	10.1ms	1.3ms	1.0ms
16384	2048	21.96s	25.09s	0.2ms	<0.1ms	10.1ms	1.3ms	2.0ms
32768	2048	44.13s	50.73s	0.2ms	<0.1ms	10.2ms	1.4ms	4.0ms
65536	2048	88.04s	101.61s	0.2ms	<0.1ms	10.1ms	1.3ms	8.1ms
131072	2048	175.70s	202.84s	0.2ms	<0.1ms	10.6ms	1.3ms	16.1ms

In Table 7, we thus compare the Server time to the Client time of our protocol, to that of a simple (witness) polynomial evaluation in this group. First of all, of course, the Server time, using homomorphic arithmetic, can be several orders of magnitude slower than the simple polynomial evaluation, while indeed being clearly linear. Still for a large enough degree, we can observe the logarithmic Client time to win over the linear time polynomial evaluation.

Second, for the protocol itself, we see that both homomorphic evaluations of the Server are quite similar. Then, on the Client side and for the considered degrees, the dominant computation is that

¹<https://gmplib.org>, <http://hms.isi.jhu.edu/acsc/libpaillier>, <https://github.com/blynn/pbc>.

²Type A symmetric pairings are constructed on the curve $y^2 = x^3 + x$ over \mathbb{F}_q for some prime $q = 3 \pmod{4}$, with a subgroup of points of order r dividing $q + 1$, and with embedding degree 2.

of a single Paillier’s deciphering (and that the only part non-constant in the degree is by far the most negligible).

Future work involves trying the RELIC framework (<https://github.com/relic-toolkit/relic>), that might have more recent and more efficient implementations of Paillier’s system and pairings, thus showing the practicality of our protocol for even smaller degrees.

6 Low server storage dynamic proof of retrievability

As we discussed in the introduction, one application area of verifiable computing protocols is Proofs of Retrievability (PoR). These schemes allow a client with limited storage, who has outsourced her data to an untrusted server, to confirm via an efficient Audit protocol that the data is still being stored in its entirety.

The lower bound of [5, Theorem 4] proves that a tradeoff is inevitable between low/high audit cost and high/low storage overhead. Roughly speaking, for any PoR on an N -bit database, the product of persistent storage overhead times audit computational complexity must be at least N .

The dynamic PoR schemes of [15, 36] optimize for fast audits. They incur a large $\mathcal{O}(N)$ storage overhead on the server, but can perform audits with only $(\log N)^{\mathcal{O}(1)}$ communication and computation for the client and server.

The authors of [5] instead optimized for small storage overhead; their scheme (which we outline next) has only sub-linear storage overhead of $\mathcal{O}(N/\log N)$, but a higher audit cost of $\mathcal{O}(N)$ on the server, and $\mathcal{O}(\sqrt{N})$ client time and communication overhead. The authors demonstrate that, for reasonable deployment scenarios on commercial cloud platforms, the higher audit cost is more than offset by the greatly reduced costs of extra persistent storage, especially if audits are only performed a few times per day.

We here further improve on the low storage overhead approach of [5], by showing how to maintain a small $o(N)$ storage overhead, but achieving only $\mathcal{O}(\log N)$ communication and client computation cost for audits. That is, our new protocol still benefits from small storage overhead, while effectively pushing the higher computational cost of audits (which is inevitable from the lower bound) entirely off the client and onto the server. These savings are highlighted in Table 8.

Table 8: Attributes of some selected Proof of Retrievability schemes

Protocol	Server		Audit Comm.	Client	
	Extra Storage	Audit Comput.		Storage	Audit Comput.
Shi et al. [36]	$\mathcal{O}(N)$	$\mathcal{O}(\log N)$	$\mathcal{O}(\log N)$	$\mathcal{O}(1)$	$\mathcal{O}(\log N)$
Anthoine et al. [5]	$o(N)$	$N + o(N)$	$\mathcal{O}(\sqrt{N})$	$\mathcal{O}(1)$	$\mathcal{O}(\sqrt{N})$
Here	$o(N)$	$N + o(N)$	$\mathcal{O}(\log N)$	$\mathcal{O}(1)$	$\mathcal{O}(\log N)$

6.1 Matrix based approach for audits

Here we summarize the PoR presented in [5] upon which our new scheme is based.

The basic premise is to treat the data, consisting of N bits organized in machine words, as a matrix $M \in \mathbb{Z}_p^{m \times n}$, where \mathbb{Z}_p is a suitable finite field of size p . Crucially, the choice of ring \mathbb{Z}_p does not require any modification to the raw data itself; that is, any element of the matrix M can be retrieved in $\mathcal{O}(1)$ time from the underlying raw data storage.

The scheme is based on the commutativity of matrix-vector products. During the Setup phase, the client chooses a secret vector u of dimension m and computes $v^T = u^T M$; both vectors u and v are then stored by the client for later use, while the server stores the original data and hence the matrix M in the clear.

Reading or updating individual entries in M can be performed efficiently with the use of Merkle hash trees and from the observation that changing one element of M only requires changing one entry in the client’s secret control vector v .

To perform an audit, the client and server engage in a 1-round protocol:

1. Client chooses a random vector x of dimension n , and sends x to Server.
2. Server computes $y = Mx$ and sends the dimension- m vector y back to Client.
3. Client computes two dot products $u^\top y$ and $v^\top x$, and checks that they are equal.

The proof of retrievability relies on the fact that observing several successful audits allows, with high probability, recovery of the correct matrix M , and therefore of the entire database.

The communication costs are $\mathcal{O}(n)$ and $\mathcal{O}(m)$ in steps 1 and 2 respectively, and the client computation in step 3 is $\mathcal{O}(m+n)$, resulting in $\mathcal{O}(\sqrt{N})$ total communication and client computation when optimizing the matrix dimensions to roughly $m = n = \sqrt{N}$.

While this square-matrix setup is the basic protocol presented by [5], the authors also discuss a potential improvement in communication complexity. Instead of x being uniformly random over \mathbb{Z}_p^n , it can instead be a *structured* vector formed from a single random element $r \in \mathbb{Z}_p$ as $x = [r^i]_{i=1..n}$. Then the communication on step 1 is reduced to constant, and hence the total communication depends only on the row dimension $\mathcal{O}(m)$. By choosing a rectangular matrix M with few rows and many columns, the communication can be made arbitrarily small.

The tradeoff for this reduction in communication complexity is higher client storage of the control vector v as well as higher client computation cost for the n -dimensional dot product $v^\top x$.

In [5], the authors found that the savings in communication were not worth the higher client storage and computation, and their experimental evaluation was based on the square matrix version with overhead $\mathcal{O}(\sqrt{N})$.

6.2 Bootstrapping part of the client computation via ciphered and dynamic polynomial evaluation

Now we show how to modify the reduced communication version of the PoR protocol of [5] just presented in order to eliminate the costly client storage of $v \in \mathbb{Z}_p^n$ and computation of $v^\top x$ during audits.

Our improved protocol is based on the observation that, when the audit challenge vector x is structured as $x = [r^i]$, then the expensive client dot product computation of $v^\top x$ is actually a polynomial evaluation: if the entries of v are the coefficients of a polynomial P , then $v^\top x$ is simply $P(r)$.

We therefore eliminate the $\mathcal{O}(n)$ client persistent storage and computation cost during audits by outsourcing the (encrypted) storage of vector v and computation of $v^\top x = P(r)$ with our novel protocol for dynamic, encrypted, verifiable polynomial evaluation scheme of Table 5. The obtained private-verification PoR protocol, combining the PoR of [5] with our ciphered polynomial evaluation in Section 5, is presented in Table 9.

Theorem 14. *The protocol of Table 9 is correct and sound.*

Proof. For the sake of simplicity, we first consider the case $t = 1$, that is a single control vector.

Correctness. Assume that all the parties are honest. After each update phase, thanks to the correctness of the Merkle hash tree algorithms $w = E(uM)$ and $\bar{K} = g^{\bar{v}\sigma}$. To see this, suppose a modification of the database at indices i and j , and let $M' = M + (M'_{ij} - M_{ij})\mathcal{E}_{ij}$ where \mathcal{E}_{ij} is the single entry matrix with 1 at position (i, j) . We have $uM' = uM + u(M'_{ij} - M_{ij})\mathcal{E}_{ij} = uM + \gamma^i e_j (M'_{ij} - M_{ij})$ where e_j is the j -th canonical vector. Thus, $v' = v + \gamma^i (M'_{ij} - M_{ij})e_j = v + \delta e_j$ satisfies $uM' = v'$. Only the j -th coefficients are different in v and v' , and in w and w' as well. For the latter, $w'_j = E(v'_j) = E(v_j + \delta) = E(v_j)E(\delta) = w_j E(\delta)$. The server thus computes w' such that $w' = E(uM')$. Moreover, $\bar{v}' = \bar{v} + \alpha \delta e_j$, so that, similarly, $\bar{H}'_j = \bar{H}_j \Delta$ with $\Delta = g^{\alpha \delta}$, and $\bar{K}' = g^{\bar{v}'\sigma} = g^{\bar{v}\sigma} g^{\alpha \delta e_j \sigma} = \bar{K} g^{\alpha \delta s^j} = \bar{K} \Delta^{s^j}$. Now, concerning the audit phase: Since we consider the polynomial evaluation as a dotproduct, the application of Proposition 8 to our notations gives: $(s-r)(\sum_{i=1}^{n-1} \sum_{j=0}^{i-1} \bar{v}_i s^{i-j-1} r^j) + \sum_{i=0}^{n-1} \bar{v}_i r^i = \sum_{i=0}^{n-1} \bar{v}_i s^i$. Thus the audit phase is:

$$\begin{aligned} \bar{\xi} &= \prod_{i=1}^{n-1} \prod_{j=0}^{i-1} e(\bar{H}_i; S_{i-j-1})^{x_j} = \prod_{i=1}^{n-1} \prod_{j=0}^{i-1} e(g^{\bar{v}_i}; g^{s^{i-j-1}})^{r^j} \\ &= e(g; g)^{\sum_{i=1}^{n-1} \sum_{j=0}^{i-1} \bar{v}_i s^{i-j-1} r^j}. \end{aligned}$$

Table 9: Private verifiable Client/server PoR protocol with low storage server

	Server	Communications	Client
Setup	$T_M \leftarrow \mathbf{MTCreatetree}(M)$ $T_w \leftarrow \mathbf{MTCreatetree}(w)$ Store $M, T_M, w, T_w, \bar{H}, S$	$N = mn \log_2 p$ order p group \mathbb{G} pairing e to \mathbb{G}_T gen. $g, e(g; g)$	$\gamma \xleftarrow{\$} \mathbb{Z}_p^t$ form $u \leftarrow [\gamma_k^i]_{k=1..t, i=0..m-1} \in \mathbb{Z}_p^{t \times m}$ $v \leftarrow uM \in \mathbb{Z}_p^{t \times n}$ $\alpha, \beta \xleftarrow{\$} (\mathbb{Z}_p^*)^t, s, \varphi \xleftarrow{\$} \mathbb{Z}_p^*$ Let $\bar{v} \leftarrow [(\alpha v_{kj} + \beta \varphi^j)]_{k=1..t, j=0..n-1}$ $w \leftarrow E(v), \bar{H} \leftarrow [g^{\bar{v}_{kj}}]_{k=1..t, j=0..n-1}$ $\sigma \leftarrow [s^j]_{j=0..n-1}, \bar{K} = g^{\bar{v}\sigma}, S \leftarrow g^\sigma$ $r_M \leftarrow \mathbf{MTRootFromLeaves}(M)$ $r_w \leftarrow \mathbf{MTRootFromLeaves}(w)$ Discard $M, v, \bar{v}, w, \bar{H}, \sigma, S$. Store $\gamma, r_M, \alpha, \beta, s, \varphi, \bar{K}, r_w$
Update	$(M_{ij}, L_{M_{ij}}) \leftarrow \mathbf{MTElementAndPath}(j + i \cdot n, M, T_M)$ $(w_{\cdot j}, L_{w_{\cdot j}}) \leftarrow \mathbf{MTElementAndPath}(j, w, T_w)$ $\bar{H}_j \leftarrow \bar{H}_j \Delta$ $w_{\cdot j} \leftarrow w_{\cdot j} e_\delta$ $T_w \leftarrow \mathbf{MTUpdateLeaf}(j, w_{\cdot j}, T_w)$ $T_M \leftarrow \mathbf{MTUpdateLeaf}(j + i \cdot n, M'_{ij}, T_M)$ Store updated M	$\xrightarrow{i, j}$ $M_{ij}, L_{M_{ij}}, w_{\cdot j}, L_{w_{\cdot j}}$ $\xrightarrow{M'_{ij}, e_\delta, \Delta}$	$r_w \stackrel{?}{=} \mathbf{MTRootFromPath}(j, w_{\cdot j}, L_{w_{\cdot j}})$ $r_M \stackrel{?}{=} \mathbf{MTRootFromPath}(j + i \cdot n, M_{ij}, L_{M_{ij}})$ $\delta \leftarrow \gamma^i (M'_{ij} - M_{ij}), e_\delta \leftarrow E(\delta), \Delta = g^{\alpha \delta}$ $\bar{K} \leftarrow \bar{K} \cdot \Delta^{s^j}$ $r_w \leftarrow \mathbf{MTRootFromPath}(j, w_{\cdot j} e_\delta, L_{w_{\cdot j}})$ $r_M \leftarrow \mathbf{MTRootFromPath}(j + i \cdot n, M'_{ij}, L_{M_{ij}})$
Audit	form $x \leftarrow [r^j]_{j=0..n-1} \in \mathbb{Z}_p^n$ $y \leftarrow Mx$ $\zeta = w^\top \square x$ $\bar{\xi} = \prod_{i=1}^{n-1} \prod_{j=0}^{i-1} e(\bar{H}_{\cdot i}; S_{i-j-1})^{x_j}$	\xleftarrow{r} $\xrightarrow{y, \zeta, \bar{\xi}}$	$r \xleftarrow{\$} \mathbb{Z}_p^*$ $c \leftarrow \beta \frac{(r\varphi)^{n-1}}{r\varphi-1}$ $\bar{\xi}^{s-r} e(g; g)^{\alpha D(\zeta) + c} \stackrel{?}{=} e(\bar{K}; g)$ $u^\top y \stackrel{?}{=} D(\zeta)$

Moreover $\alpha D(\zeta) + c = \alpha vx + \beta \frac{(r\varphi)^{n-1}}{r\varphi-1} = \alpha vx + \beta \sum_{j=0}^{n-1} \varphi^j r^j = \bar{v}x$, thus we have that

$$\bar{\xi}^{s-r} e(g; g)^{\alpha D(\zeta) + c} = e(g, g)^{(s-r)(\sum_{i=1}^{n-1} \sum_{j=0}^{i-1} \bar{v}_i s^{i-j-1} r^j) + \bar{v}x} = e(g, g)^{\bar{v}\sigma} = e(\bar{K}; g).$$

Finally, we get $uy = UMx = vx$.

Soundness. An attacker to the protocol must provide (y', ζ') such that $(y', \zeta') \neq (y, \zeta)$, but still $u^\top y' = D_{sk}(\zeta')$, with a non negligible advantage ϵ . There are two cases: if $D_{sk}(\zeta') \neq D_{sk}(\zeta)$ then the attacker had to break the secure polynomial evaluation; otherwise, it must be that $u^\top y' = u^\top y$ with $y' \neq y$. For the first case, [Theorem 13](#) assesses the security of the polynomial evaluation.

For the second case, we consider $X = E_{pk}(x)$ the cipher of a secret x by the homomorphic scheme. Here, as in [Theorem 13](#), we use the fact that the protocol of [Table 9](#) is indistinguishable as a whole from the same protocol where, within the polynomial evaluation of, $\beta \varphi^i$ is everywhere replaced by a random ρ_i . Further, this is indistinguishable from a third protocol where, at each update of index i , a new ρ'_i is also randomly redrawn and replaces ρ_i in the client state. We thus continue the proof with this third game setting. Now, using e_k the k -th canonical vector of \mathbb{Z}_p^m , we can consider $\bar{u} = u + x e_k$ and $\bar{v}^\top = \bar{u}^\top M = (u^\top + x e_k^\top) M = v^\top + x M_{k,*}$. Then, for the **Setup** phase, we can randomly select m, n and $k \leq m$. Then also $M \in \mathbb{Z}_p^{m \times n}, u \in \mathbb{Z}_p^m$, and compute $v^\top = u^\top M$. From this, set $w_j = X^{M_{kj}} E(v_j) = E(v_j + x M_{kj}) = E(\bar{v}_j)$, for $j = 1..n$. We also randomly select s, α and h_j (so that $\rho_j = \log(h_j) - \alpha \bar{v}_j$ exists, but remains unknown). For any **Update** phases, compute $w'_j = w_j X^{M_{kj} - M_{kj}}$ and select randomly a Δ (so that $h'_j = h_j \Delta$ now corresponds to a new $\rho'_j = \log(h'_j) - \alpha \bar{v}'_j$ still unknown).

Finally, the attacker provides a vector y' such that both $\bar{u}^\top (y' - y) = 0$ and $y' \neq y \pmod p$. Since k is randomly chosen from $1..m$, the probability that the vectors are distinct at index k , in other words that $\bar{y}_k \neq y_k \pmod p$, is at least $1/m$. If this is the case, then, denoting $z = \bar{y} - y$, we have $z_k \neq 0 \pmod p$ and

$\bar{u}^\top z = 0$ implies that $u^\top z + xz_k = 0$ so that $x \equiv -z_k^{-1} \cdot (u^\top z) \pmod p$ and the homomorphic cryptosystem is broken with advantage ϵ/m .

The same reasoning works mutatis mutandis with $t > 1$, but selecting t at random, and then selecting t random locations to put the secret x , once in each one of the t columns of u . The advantage to break the homomorphic scheme would become $t\epsilon/m$. \square \square

6.3 Experiments

We now compare our modification of the PoR protocol with the one in [5], publicly available there: <https://github.com/dsroche/la-por>.

Table 10 has three blocks of experiments, each for four database sizes ranging from 1GB to 1TB. The first block of experiments is a run of the original statistically secure PoR protocol with two dotproducts for the verification, considering the matrix as 56 bits elements of modulo a 57-bits prime. The second block of experiments is our new modification, but still using close to square matrices. Subject now to computational security, we have to use a larger coefficient domain, namely here a 256-bits prime. We separate the timings of the **Update** phase in two phases, the remaining linear algebra phase and the new polynomial evaluation phase. In the third block of experiments we use a more rectangular matrix, trying to reduce communications while not increasing too much the computational effort.

Table 10: Modification of the PoR protocol, with 256-bits groups, 1024-bits Paillier, on 1 core i7-6700 3.40GHz & 64 GB ram

Database	1GB	10GB	100GB	1TB
Private-verified audit using 57-bits prime [5, Table 1]				
Matrix view	12339×12432	39131×39200	123831×123872	396281×396368
Client Audit	1.7ms	1.3ms	21.8ms	78.8ms
Server Audit	0.25s	2.38s	153.04s	1748.69s
Client Storage	194KB	612KB	1 935KB	49 540KB
Communications	194KB	612KB	1 935KB	49 540KB
Dynamic-ciphered delegated polynomial evaluation with 256-bits groups of Table 9				
Matrix view	5793×5793	18318×18318	57927×57927	185364×185364
Client Audit (dotproduct part)	1.7ms	20.3ms	56.8ms	79.9ms
Client Audit (polynomial part)	3.0ms	3.0ms	3.1ms	3.8ms
Server Audit (matrix-vector part)	4.3s	43.2s	433.2s	4 441.7s
Server Audit (polynomial part)	11.0s	35.2s	111.9s	357.8s
Client storage	0.38KB	0.38KB	0.38KB	0.38KB
Communications	181KB	572KB	1 810KB	5 792KB
Dynamic-ciphered delegated polynomial evaluation with 256-bits groups of Table 9				
Matrix view	6600×5085	7265×46187	7929×423187	8600×3995319
Client Audit (dotproduct part)	1.9ms	2.1ms	2.3ms	2.5ms
Client Audit (polynomial part)	3.0ms	3.1ms	3.1ms	3.1ms
Server Audit (matrix-vector part)	4.3s	42.1s	434.0s	4 534.0s
Server Audit (polynomial part)	9.7s	90.4s	830.4s	7 781.7s
Client storage	0.38KB	0.38KB	0.38KB	0.38KB
Communications	206KB	227KB	248KB	269KB

Overall, we see first in Table 10, that changing the coefficient domain size increases the computational effort of the server in the linear algebra phase. Still, reducing the dimension of the dotproduct for the client, as shown in the third block, allows the client to be faster for databases larger than 100GB. Still the client audit computational effort is never larger than a few milliseconds and thus the dominant part is most certainly communications. On this aspect, we see that our modification allows for large reductions in both the Client storage (even with square matrices) and the overall communications. Indeed, the client private state is the vector dimension, the Paillier’s private key, five integers modulo p , one group element and two Merkle tree roots; while the communications are mostly one vector of modular integers in the smallest dimension.

The price to pay is from about a factor of four to an order of magnitude for the server computations.

But the persistent client storage is going from dozens of *MB* to about 400 bytes, and the communication volume can be decreased by more than two orders of magnitude.

Moreover, we mention as possible future work the potential to parallelize the polynomial part of the server’s computation. The matrix-vector product part is already trivially parallelized, but as the dimensions become more rectangular, as we can see in [Table 10](#) the polynomial part starts to dominate. For this, a standard “baby steps / giant steps” approach could be employed, and all pairings computed in parallel as well. This could be used to further reduce the server latency for large databases and make the scheme more practically viable.

References

- [1] Michel Abdalla, Fabrice Benhamouda, and Alain Passelègue. An algebraic framework for pseudorandom functions and applications to related-key security. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology – CRYPTO 2015*, pages 388–409, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg. doi:[10.1007/978-3-662-47989-6_19](https://doi.org/10.1007/978-3-662-47989-6_19).
- [2] Michel Abdalla, Florian Bourse, Hugo Marival, David Pointcheval, Azam Soleimani, and Hendrik Waldner. Multi-client inner-product functional encryption in the random-oracle model. In Clemente Galdi and Vladimir Kolesnikov, editors, *Security and Cryptography for Networks - 12th International Conference, SCN 2020, Amalfi, Italy, September 14-16, 2020, Proceedings*, volume 12238 of *Lecture Notes in Computer Science*, pages 525–545. Springer, 2020. doi:[10.1007/978-3-030-57990-6_26](https://doi.org/10.1007/978-3-030-57990-6_26).
- [3] Shweta Agrawal, Benoît Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016*, pages 333–362, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg. doi:[10.1007/978-3-662-53015-3_12](https://doi.org/10.1007/978-3-662-53015-3_12).
- [4] Miguel Ambrona, Gilles Barthe, and Benedikt Schmidt. Generic transformations of predicate encodings: Constructions and applications. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017*, pages 36–66, Cham, 2017. Springer International Publishing. doi:[10.1007/978-3-319-63688-7_2](https://doi.org/10.1007/978-3-319-63688-7_2).
- [5] Gaspard Anthoine, Jean-Guillaume Dumas, Michael Hanling, Mélanie de Jonghe, Aude Maignan, Clément Pernet, and Daniel S. Roche. Dynamic proofs of retrievability with low server storage. In *30th USENIX Security Symposium, August 11-13*, pages 1–18. USENIX Association, August 2021. URL: <https://www.usenix.org/conference/usenixsecurity21/presentation/anthoine>.
- [6] Giuseppe Ateniese, Randal Burns, Reza Curtmola, Joseph Herring, Lea Kissner, Zachary Peterson, and Dawn Song. Provable data possession at untrusted stores. In *14th ACM CCS*, pages 598–609. ACM, 2007. doi:[10.1145/1315245.1315318](https://doi.org/10.1145/1315245.1315318).
- [7] Nuttapon Attrapadung and Junichi Tomida. Unbounded dynamic predicate compositions in abe from standard assumptions. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2020*, pages 405–436, Cham, 2020. Springer International Publishing. doi:[10.1007/978-3-030-64840-4_14](https://doi.org/10.1007/978-3-030-64840-4_14).
- [8] Siavosh Benabbas, Rosario Gennaro, and Yevgeniy Vahlis. Verifiable delegation of computation over large datasets. In Phillip Rogaway, editor, *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 111–131. Springer, 2011. doi:[10.1007/978-3-642-22792-9_7](https://doi.org/10.1007/978-3-642-22792-9_7).
- [9] Josh Benaloh. Dense probabilistic encryption. In *First Annual Workshop on Selected Areas in Cryptography*, pages 120–128, Kingston, ON, May 1994. URL: http://sacworkshop.org/proc/SAC_94_006.pdf.

- [10] Allison Bishop, Abhishek Jain, and Lucas Kowalczyk. Function-hiding inner product encryption. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part I*, volume 9452 of *Lecture Notes in Computer Science*, pages 470–491. Springer, 2015. doi:[10.1007/978-3-662-48797-6_20](https://doi.org/10.1007/978-3-662-48797-6_20).
- [11] Dan Boneh and Xavier Boyen. Short signatures without random oracles. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology - EUROCRYPT 2004, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 56–73. Springer, 2004. doi:[10.1007/978-3-540-24676-3_4](https://doi.org/10.1007/978-3-540-24676-3_4).
- [12] Dan Boneh, Justin Drake, Ben Fisch, and Ariel Gabizon. Efficient polynomial commitment schemes for multiple points and polynomials. *IACR Cryptol. ePrint Arch.*, 2020:81, 2020. URL: <https://eprint.iacr.org/2020/081>.
- [13] Xavier Bultel, Manik Lal Das, Hardik Gajera, David Gérard, Matthieu Giraud, and Pascal Lafourcade. Verifiable private polynomial evaluation. In Tatsuaki Okamoto, Yong Yu, Man Ho Au, and Yannan Li, editors, *Provable Security*, pages 487–506, Cham, 2017. Springer International Publishing. doi:[10.1007/978-3-030-41702-4_4](https://doi.org/10.1007/978-3-030-41702-4_4).
- [14] Jan Camenisch, Maria Dubovitskaya, Kristiyan Haralambiev, and Markulf Kohlweiss. Composable and modular anonymous credentials: Definitions and practical constructions. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part II*, volume 9453 of *Lecture Notes in Computer Science*, pages 262–288. Springer, 2015. doi:[10.1007/978-3-662-48800-3_11](https://doi.org/10.1007/978-3-662-48800-3_11).
- [15] David Cash, Alptekin Küpçü, and Daniel Wichs. Dynamic proofs of retrievability via oblivious RAM. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 279–295. Springer, 2013. doi:[10.1007/978-3-642-38348-9_17](https://doi.org/10.1007/978-3-642-38348-9_17).
- [16] Dario Catalano and Dario Fiore. Vector commitments and their applications. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *Public-Key Cryptography - PKC 2013 - 16th International Conference on Practice and Theory in Public-Key Cryptography, Nara, Japan, February 26 - March 1, 2013. Proceedings*, volume 7778 of *Lecture Notes in Computer Science*, pages 55–72. Springer, 2013. doi:[10.1007/978-3-642-36362-7_5](https://doi.org/10.1007/978-3-642-36362-7_5).
- [17] Kaoutar Elkhiyaoui, Melek Önen, Monir Azraoui, and Refik Molva. Efficient techniques for publicly verifiable delegation of computation. In Xiaofeng Chen, XiaoFeng Wang, and Xinyi Huang, editors, *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security, AsiaCCS 2016, Xi'an, China, May 30 - June 3, 2016*, pages 119–128. ACM, 2016. doi:[10.1145/2897845.2897910](https://doi.org/10.1145/2897845.2897910).
- [18] Dario Fiore and Rosario Gennaro. Publicly verifiable delegation of large polynomials and matrix computations, with applications. In *ACM CCS*, pages 501–512, New York, NY, USA, 2012. ACM. doi:[10.1145/2382196.2382250](https://doi.org/10.1145/2382196.2382250).
- [19] Dario Fiore, Anca Nitulescu, and David Pointcheval. Boosting verifiable computation on encrypted data. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *Public-Key Cryptography - PKC 2020*, pages 124–154, Cham, 2020. Springer International Publishing. doi:[10.1007/978-3-030-45388-6_5](https://doi.org/10.1007/978-3-030-45388-6_5).
- [20] Laurent Fousse, Pascal Lafourcade, and Mohamed Alnuaimi. Benaloh’s dense probabilistic encryption revisited. In Abderrahmane Nitaj and David Pointcheval, editors, *Progress in Cryptology - AFRICACRYPT 2011 - 4th International Conference on Cryptology in Africa, Dakar, Senegal, July 5-7, 2011. Proceedings*, volume 6737 of *Lecture Notes in Computer Science*, pages 348–362. Springer, 2011. doi:[10.1007/978-3-642-21969-6_22](https://doi.org/10.1007/978-3-642-21969-6_22).

- [21] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. *IACR Cryptol. ePrint Arch.*, 2019:953, 2019. URL: <https://eprint.iacr.org/2019/953>.
- [22] Romain Gay, Dennis Hofheinz, Eike Kiltz, and Hoeteck Wee. Tightly CCA-secure encryption without pairings. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016*, pages 1–27, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg. doi:10.1007/978-3-662-49890-3_1.
- [23] Rosario Gennaro, Craig Gentry, and Bryan Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In Tal Rabin, editor, *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, volume 6223 of *Lecture Notes in Computer Science*, pages 465–482. Springer, 2010. doi:10.1007/978-3-642-14623-7_25.
- [24] Vipul Goyal. Reducing trust in the PKG in identity based cryptosystems. In Alfred Menezes, editor, *Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings*, volume 4622 of *Lecture Notes in Computer Science*, pages 430–447. Springer, 2007. doi:10.1007/978-3-540-74143-5_24.
- [25] Ari Juels and Burton S Kaliski Jr. Pors: Proofs of retrievability for large files. In *14th ACM CCS*, pages 584–597. ACM, 2007. doi:10.1145/1315245.1315317.
- [26] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings*, volume 6477 of *Lecture Notes in Computer Science*, pages 177–194. Springer, 2010. doi:10.1007/978-3-642-17373-8_11.
- [27] J. Katz and Y. Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC Cryptography and Network Security Series. CRC Press, 2020. URL: <https://books.google.fr/books?id=RsoOEAAAQBAJ>.
- [28] Sam Kim, Kevin Lewi, Avradip Mandal, Hart Montgomery, Arnab Roy, and David J. Wu. Function-hiding inner product encryption is practical. In Dario Catalano and Roberto De Prisco, editors, *Security and Cryptography for Networks - 11th International Conference, SCN 2018, Amalfi, Italy, September 5-7, 2018, Proceedings*, volume 11035 of *Lecture Notes in Computer Science*, pages 544–562. Springer, 2018. doi:10.1007/978-3-319-98113-0_29.
- [29] B. Laurie, A. Langley, E. Kasper, and Google. Certificate Transparency. RFC 6962, IETF, June 2013. URL: <https://tools.ietf.org/html/rfc6962>.
- [30] Benoît Libert, Somindu C. Ramanna, and Moti Yung. Functional commitment schemes: From polynomial commitments to pairing-based accumulators from simple assumptions. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPICs*, pages 30:1–30:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICs.ICALP.2016.30.
- [31] Ralph C. Merkle. A digital signature based on a conventional encryption function. In C. Pomerance, editor, *CRYPTO '87*, pages 369–378, 1988. doi:10.1007/3-540-48184-2_32.
- [32] Paz Morillo, Carla Ràfols, and Jorge L. Villar. The kernel matrix Diffie-Hellman assumption. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *Advances in Cryptology – ASIACRYPT 2016*, pages 729–758, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg. doi:10.1007/978-3-662-53887-6_27.
- [33] Alex Ozdemir, Riad S. Wahby, Barry Whitehat, and Dan Boneh. Scaling verifiable computation using efficient set accumulators. In Srdjan Capkun and Franziska Roesner, editors, *29th USENIX*

- Security Symposium, USENIX Security 2020, August 12-14, 2020*, pages 2075–2092. USENIX Association, 2020. URL: <https://www.usenix.org/conference/usenixsecurity20/presentation/ozdemir>.
- [34] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In Jacques Stern, editor, *Advances in Cryptology - EUROCRYPT '99, International Conference on the Theory and Application of Cryptographic Techniques, Prague, Czech Republic, May 2-6, 1999, Proceeding*, volume 1592 of *Lecture Notes in Computer Science*, pages 223–238. Springer, 1999. doi:10.1007/3-540-48910-X_16.
- [35] Hovav Shacham and Brent Waters. Compact proofs of retrievability. In Josef Pieprzyk, editor, *Advances in Cryptology - ASIACRYPT 2008, 14th International Conference on the Theory and Application of Cryptology and Information Security, Melbourne, Australia, December 7-11, 2008. Proceedings*, volume 5350 of *Lecture Notes in Computer Science*, pages 90–107. Springer, 2008. doi:10.1007/978-3-540-89255-7_7.
- [36] Elaine Shi, Emil Stefanov, and Charalampos Papamanthou. Practical dynamic proofs of retrievability. In *ACM CCS*, pages 325–336, New York, NY, USA, 2013. ACM. URL: <http://elaineshi.com/docs/por.pdf>, doi:10.1145/2508859.2516669.
- [37] Alin Tomescu, Ittai Abraham, Vitalik Buterin, Justin Drake, Dankrad Feist, and Dmitry Khovratovich. Aggregatable subvector commitments for stateless cryptocurrencies. In Clemente Galdi and Vladimir Kolesnikov, editors, *Security and Cryptography for Networks - 12th International Conference, SCN 2020, Amalfi, Italy, September 14-16, 2020, Proceedings*, volume 12238 of *Lecture Notes in Computer Science*, pages 45–64. Springer, 2020. doi:10.1007/978-3-030-57990-6_3.