Prover efficient public verification of dense or sparse/structured matrix-vector multiplication

Jean-Guillaume Dumas, Vincent Zucca

Prover efficient public verification of dense or sparse/structured matrix-vector multiplication*

Jean-Guillaume Dumas† Vincent Zucca‡

April 7, 2017

Abstract

With the emergence of cloud computing services, computationally weak devices (Clients) can delegate expensive tasks to more powerful entities (Servers). This raises the question of verifying a result at a lower cost than that of recomputing it. This verification can be private, between the Client and the Server, or public, when the result can be verified by any third party. We here present protocols for the verification of matrix-vector multiplications, that are secure against malicious Servers. The obtained algorithms are essentially optimal in the amortized model: the overhead for the Server is limited to a very small constant factor, even in the sparse or structured matrix case; and the computational time for the public Verifier is linear in the dimension. Our protocols combine probabilistic checks and cryptographic operations, but minimize the latter to preserve practical efficiency. Therefore our protocols are overall more than two orders of magnitude faster than existing ones.

1 Introduction

With the emergence of cloud computing services, computationally weak devices (Clients, such as smart phones or tablets) can delegate expensive tasks to more powerful entities (Servers). Such heavy tasks can, e.g., be cryptographic operations, image manipulation or statistical analysis of large data-sets. This raises the question of verifying a result at a lower cost than that of recomputing it. This verification can be private, between the Client and the Server, or public, when the result can be verified by any third party.

For instance within computer graphics (image compression and geometric transformation), graph theory (studying properties of large networks), big data analysis, one deals with linear transformations of large amount of data, often arranged in large matrices with large dimensions that are in the order of thousands or millions in some applications. Since a linear transformation on a vector x can be expressed by a matrix-vector multiplication (with a matrix of size $m \times n$), a weak client can use one of the protocols in the literature [6, 14, 4] to outsource and verify this computation in the optimal time $O(m + n)$, i.e., linear in the input and the output size. However as these protocols use expensive cryptographic operations, such as pairings, the constants hidden in the asymptotic complexity are usually extremely large [13].

In this paper, we propose an alternative protocol, achieving the same optimal behavior, but which is also practical: the overhead for the Prover is now very close to the time required to compute the matrix-vector multiplication, thus gaining two orders of magnitude with respect to the literature. Our protocol not only does this for dense matrices, but is also sensitive to any structure or sparsity of the linear transformation. For this, we first remove any quadratic operation that is not a matrix-vector multiplication (that is we use projections and rank-1 updates) and second we separate operations in the base field from cryptographic operations so as to minimize the latter.

*This work is partly funded by the OpenDreamKit Horizon 2020 European Research Infrastructures project (#676541).

†Université Grenoble Alpes, CNRS, LJK, 700 av. centrale, IMAG - CS 40700, 38058 Grenoble cedex 9, France, Jean-Guillaume.Dumas@imag.fr.

‡Sorbonne Universités, Univ. Pierre et Marie Curie, Laboratoire LIP6, CNRS, umr 7606, 4 place Jussieu, F75252 Paris, France, Vincent.Zucca@lip6.fr.

More precisely, we first combine rank-one updates of [6] and the projecting idea of [4] with Freivalds' probabilistic check [7]. Second, we use a novel strategy of vectorization. For instance, with a security parameter s (e.g., an $s = 128$-bits equivalent security), exponentiations or pairings operations usually cost about $O(s^3)$ arithmetic operations. To make the whole protocol work practical, we thus reduce its cost from $\mathcal{O}\left(s^3 mn\right)$ to $\mathcal{O}\left(\mu(A) + s^3(m + n^{4/3})\right)$, where $\mu(A) < 2mn$ is the cost of one, potentially structured, matrix-vector multiplication. We also similarly reduce the work of the Verifier. This allows us to gain two orders of magnitude on the Prover's work and therefore on the overall costs of outsourcing, while preserving and sometimes even improving the practical efficiency of the Verifier.

Thus, after some background in Section 2, our first improvement is given in a relaxed public verification setting in Section 3 via matrix projection and probabilistic checks. Our second improvement is given in Section 4 where the verification is bootstrapped efficiently by vectorization. We then show how to combine all improvements in Section 5 in order to obtain a complete and provably secure protocol. Finally, we show in Section 6 that our novel protocol indeed induces a global overhead factor lower than 3 with respect to non verified computations. This is gaining several orders of magnitude on the Prover side with respect to previously known protocols, while keeping the Verification step an order of magnitude faster.

2 Background and definitions

In this paper, we want to be able to prove fast that a vector is a solution to a linear system, or equivalently that a vector is the product of another vector by a matrix. This is useful, e.g., to perform some statistical analysis on some medical data. We distinguish the matrix, a static data, from the vectors which are potentially diverse. In the following, $\mathbb{F}_p$ will denote a prime field and we consider:

- Data: matrix $A \in \mathbb{F}_p^{m \times n}$.
- Input: one or several vectors $x_i \in \mathbb{F}_p^n$, for $i = 1..k$.
- Output: one or several vectors $y_i = Ax_i \in \mathbb{F}_p^m$, for $i = 1..k$.

Then, we denote by $\star$ an operation performed in the exponents (for instance, for $u \in \mathbb{G}^n$ and $v \in \mathbb{Z}^n$, the operation $u^T \star v$ actually denotes $\prod_{j=1}^n u[j]^{v[j]}$).

Publicly Verifiable Computation. A publicly verifiable computation scheme, in the formal setting of [11], is in fact four algorithms (*KeyGen*, *ProbGen*, *Compute*, *Verify*), where *KeyGen* is some (amortized) preparation of the data, *ProbGen* is the preparation of the input, *Compute* is the work of the *Prover* and *Verify* is the work of the *Verifier*. Usually the Verifier also executes *KeyGen* and *ProbGen* but in a more general setting these can be performed by different entities (respectively called a *Preparator* and a *Trustee*). More formally we define these algorithms as follow:

- $KeyGen(1^\lambda, f) \to (\texttt{param}, EK_f, VK_f)$: a randomized algorithm run by a *Preparator*, it takes as input a security parameters 1^λ and the function f to be outsourced. It outputs public parameters `param` which will be used by the three remaining algorithms, an evaluation key EK_f and a verification key VK_f.
- $ProbGen(x) \to (\sigma_x)$: a randomized algorithm run by a *Trustee* which takes as input an element x in the domain of the outsourced function f. It returns σ_x, an encoded version of the input x.
- $Compute(\sigma_x, EK_f) \to (\sigma_y)$: an algorithm run by the *Prover* to compute an encoded version σ_y of the output $y = f(x)$ given the encoded input σ_x and the evaluation key EK_x.
- $Verify(\sigma_y, VK_f) \to y$ or $\perp$: given the encoded output σ_y and the verification key VK_f, the *Verifier* runs this algorithm to determine whether $y = f(x)$ or not. If the verification passes it returns y otherwise it returns an error $\perp$.

Completeness. A publicly verifiable computation scheme for a family of function $\mathcal{F}$ is considered to be *perfectly complete* (or *correct*) if for every function belonging to $\mathcal{F}$ and for every input in the function domain, an honest *Prover* which runs faithfully the algorithm *Compute* will *always* (with probability 1) output an encoding σ_y which will pass *Verify*.

Soundness. A publicly verifiable computation scheme for a family of function $\mathcal{F}$ is called *sound* when a prover cannot convince a verifier to accept a wrong result $y' \neq y$ except with negligible probability. More formally we evaluate the capability of an adversary $\mathcal{A}$ to deceive the verifier through a *soundness experiment*. In this experiment, we assume that the adversary $\mathcal{A}$ accesses to the output of the algorithm *KeyGen* by calling an oracle $\mathcal{O}_{KeyGen}$ with inputs 1^λ and the function to evaluate f. This oracle $\mathcal{O}_{KeyGen}$ returns public parameters for the protocol `param`, an evaluation key EK_f and a verification key VK_f. Afterwards the adversary $\mathcal{A}$ sends its challenge input x to an oracle $\mathcal{O}_{ProbGen}$ which returns σ_x. Finally $\mathcal{A}$ outputs an encoding $\sigma_{y^*} \neq \sigma_y$ and runs the *Verify* algorithm on inputs σ_{y^*} and VK_f, whether it outputs y or $\perp$ the experiment has either succeeded or failed.

Definition 1. *A publicly verifiable computation scheme for a family of function $\mathcal{F}$ is sound if and only if for any polynomially bounded adversary $\mathcal{A}$ and for any f in $\mathcal{F}$ the probability that $\mathcal{A}$ succeeds in the soundness experiment is negligible in the security parameter.*

Adversary model. The protocol in [6] (recalled for the sake of completeness in Appendix A) is secure against a *malicious Server only*. That is the Client must trust both the Preparator and the Trustee. We will stick to this model of attacker in the remaining of this paper. Otherwise some attacks can be mounted:

- Attack with a *Malicious Preparator only*: send A' and a correctly associated W' to the Server, but pretend that A is used. Then, all verifications do pass, but for $y' = A'x$ and not $y = Ax$.
- Attack with a *Malicious Trustee only*: there, a malicious assistant can provide a wrong VK_x making the verification fail even if the Server and Preparator are honest and correct.
- Attack with *Malicious Server and Trustee*: the Server sends any y' and any z' to the Trustee, who computes $\text{VK'}_x[i] = e(z'[i]; g_2)/a^{y'[i]}$, that will match the verification with y' and z'.

Public delegatability. One can also further impose that there is not interaction between the Client and the Trustee after the Client has sent his input to the Server. Publicly verifiable protocols with this property are said to be publicly delegatable [4]. The protocol in [6] does not achieve this property, but some variants in [14, 4] already can.

Bilinear Pairings. The protocols we present in this paper use bilinear pairings and their security is based on the co-CDH assumption, for the sake of completeness we recall hereafter these definitions.

Definition 2 (bilinear pairing).
Let $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$ be three groups of prime order p, a bilinear pairing is a map $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ with the following properties:

1. bilinearity*: $\forall a, b \in \mathbb{F}_p,\ \forall (g_1, g_2) \in \mathbb{G}_1 \times \mathbb{G}_2,\ e(g_1^a, g_2^b) = e(g_1, g_2)^{ab}$;*
2. non-degeneracy*: if g_1 and g_2 are generators of $\mathbb{G}_1$ and $\mathbb{G}_2$ respectively then $e(g_1, g_2)$ is a generator of $\mathbb{G}_T$;*
3. computability*: $\forall (g_1, g_2) \in \mathbb{G}_1 \times \mathbb{G}_2$, there exist an efficient algorithm to compute $e(g_1, g_2)$.*

Definition 3 (co-CDH assumption).
Let $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$ be three groups of prime order p, such that there exist a bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$. Let $g_1 \in \mathbb{G}_1$, $g_2 \in \mathbb{G}_2$ be generators and a, $b \xleftarrow{\$} \mathbb{F}_p$ be chosen randomly. We say that the co-computational Diffie-Hellman assumption (co-CDH) holds in $\mathbb{G}_1$, if given g_1, g_2, g_1^a, g_2^b the probability to compute g_1^{ab} is negligible.

Related work. The work of [6] introduced the idea of performing twice the computations, once in the classical setting and once on encrypted values. This enables the Client to only have to check consistency of both results. The protocol is sound under the Decision Linear and co-CDH hypothesis. Then [14] extended the part on matrix-vector multiplication to matrix-matrix while adding public delegatability. Finally, [4] introduced the idea of projecting the random additional matrix and the extra-computations of the server, which allows to reduce the cost of the *Verify* algorithm. It also decreases the size of the verification key

by a factor m. For an $m \times n$ dense matrix, the protocol in [6] has a constant time overhead for the Prover, but this constant is on the order of cryptographic public-key operations like pairings. Similarly, the Verifier has $\mathcal{O}(mn)$ cryptographic public-key pre-computations and $\mathcal{O}(n)$ of these for the public verification. These cryptographic operations can then induce some 10^6 slow-down [13] and do no improve even if the initial matrix is sparse or structured (as the rank one updates, $s \cdot t^T$ and $\sigma \cdot \tau^T$, are always dense).

3 A first step towards public verifiability

Freivalds' probabilistic verification of matrix multiplications [7] allows for private verifiability of matrix-vector computations. This can be naturally extended in the random oracle model via Fiat-Shamir heuristic [5]. This however forces the vectors to be multiplied to be known in advance (the full details are given in Appendix B), whereas our goal is instead to obtain public verifiability with an *unbounded number of vector inputs*. As an upstart, we thus first present an improvement if the public verification model is slightly relaxed: in this Section, we allow the Trustee to perform some operations after the computations of the Server. We will see in next sections how to remove the need for the Trustee's intervention. For this, we combine Freivalds projection (to check that $Ax_i = y_i$, one can first precompute $w^T = u^T A$ and check that $w^T x_i = u^T y_i$) with Fiore & Gennaro's protocol, in order to improve the running time of both the Trustee and the Client: we let the Prover compute its projection in the group. That way most of the pairings computations of the Trustee and Client are transformed to classical operations: the improvement is from $\mathcal{O}(n)$ cryptographic operations to $\mathcal{O}(n)$ classical operations and a single cryptographic one. Further, the projection can be performed beforehand, during the precomputation phase. That way the preparation requires only one matrix-vector for the Freivalds projection and the dense part is reduced to a single vector. The cryptographic operations can still be delayed till the last check on pairings. This is shown in Figure 1.

- Preparator: secret random $u \in \mathbb{F}_p^m$, $t \in \mathbb{F}_p^n$, then $\omega^T = g_1^{u^T A + t^T} \in \mathbb{G}_1^n$.
- Preparator to Prover: $A \in \mathbb{F}_p^{m \times n}$, $\omega \in \mathbb{G}_1^n$
- Preparator to Trustee: u, t in a secure channel.
- Verifier to Prover: $x_i \in \mathbb{F}_p^n$
- Prover to Verifier: $y_i \in \mathbb{F}_p^m$, $\zeta_i \in \mathbb{G}_1$ such that $y_i = Ax_i$ and $\zeta_i = \omega^T \star x_i$.
- Verifier to Trustee: x_i, y_i
- Trustee to Verifier: $h_i = (u^T \cdot y_i) \in \mathbb{F}_p$ and $d_i = (t^T \cdot x_i) \in \mathbb{F}_p$, then send $\eta_i = e(g_1; g_2)^{h_i + d_i} \in \mathbb{G}_T$.
- Verifier public verification: $e(\zeta_i; g_2) \stackrel{?}{=\!=} \eta_i$ in $\mathbb{G}_T$.

Figure 1: Interactive protocol for Sparse-matrix vector multiplication verification under the co-CDH.

Theorem 4. *The protocol of Figure 1 is perfectly complete and sound under the co-Computational Diffie-Hellman assumption.*

The proof of Theorem 4 is given in Appendix C.

4 Verifying the dot-products by bootstrapping and vectorization

To obtain public verifiability and public delegatability, the Client should perform both dot-products, $u^T \cdot y$ and $t^T \cdot x$ (from now on, for the sake of simplicity, we drop the indices on x and y). But as u and t must remain secret, they will be encrypted beforehand. To speed-up the Client computation, the idea is then to let the Server perform the encrypted dot-products and to allow the Client to verify them mostly with classical operations.

For this trade-off, we use vectorization. That is, for the vectors u and y, we form another representation

as $\sqrt{m}\times\sqrt{m}$ matrices:

$$U = \begin{bmatrix} u_1 & \dots & u_{\sqrt{m}} \\ u_{1+\sqrt{m}} & \dots & u_{2\sqrt{m}} \\ \dots & \dots & \dots \\ u_{1+m-\sqrt{m}} & \dots & u_m \end{bmatrix} \text{ and } Y = \begin{bmatrix} y_1 & \dots & y_{1+m-\sqrt{m}} \\ y_2 & \dots & y_{2+m-\sqrt{m}} \\ \dots & \dots & \dots \\ y_{\sqrt{m}} & \dots & y_m \end{bmatrix}.$$

Then $u^T \cdot y = Trace(UY)$. Computing with this representation is in general slower than with the direct dot-product, $\mathcal{O}\left(\sqrt{m}^3\right)$ instead of $\mathcal{O}(m)$. As shown next, this can be circumvented with well-chosen left-hand sides and at least mitigated, with unbalanced dimensions.

4.1 dot-product with rank 1 left-hand side

The first case is if u is of rank 1, that is if in matrix form, u can be represented by a rank one update, $U = \mu \cdot \eta^T$ for $\mu, \eta \in \mathbb{F}_p^{\sqrt{m}}$. Then both representations require roughly the same number of operations to perform a dot-product since then:

$$Trace(\mu \cdot \eta^T \cdot Y) = \eta^T \cdot Y \cdot \mu \tag{1}$$

Therefore, we let the Prover compute $z^T = g_1^{\eta^T} \star Y$, where $z[i] = g_1^{\sum \eta[j]Y[j,i]}$, and then the Verifier can check this value via Freivalds with a random vector v: $g_1^{\eta^T} \star (Y \cdot v) \stackrel{?}{=\!=} z^T \star v$. The point is that the Verifier needs now $\mathcal{O}(m)$ operations to compute $(Y \cdot v)$, but these are just classical operations over the field. Then its remaining operations are cryptographic but there is only $\mathcal{O}(\sqrt{m})$ of these. Finally, the Verifier concludes the computation of the dotproduct, still with cryptographic operations, but once again with only $\mathcal{O}(\sqrt{m})$ of them. Indeed, the dot product $d = u^T y = Trace(UY) = Trace(\mu \cdot \eta^T \cdot Y) = \eta^T \cdot Y \cdot \mu$ is checked by $e(g_1^d; g_2) = e(g_1; g_2)^d = \prod_{i=1}^{\sqrt{m}} e(z[i]; g_2^{\mu[i]})$ and the latter is $e(g_1; g_2)^{\sum\sum \mu[i]\eta[j]Y[j,i]} = e(g_1; g_2)^{\eta^T \cdot Y \cdot \mu}$.

In practice, operations in a group can be slightly faster than pairings. Moreover $\lceil\sqrt{m}\rceil^2$ can be quite far off m. Therefore it might be interesting to use a non square vectorization $b_1 \times b_2$, as long as $b_1 b_2 \geq m$ and $b_1 + b_2 = \Theta(\sqrt{m})$ (and 0 padding if needed). Then we have $U \in \mathbb{F}_p^{b_1 \times b_2}$, $\mu \in \mathbb{F}_p^{b_1}$, $\eta \in \mathbb{F}_p^{b_2}$, $Y \in \mathbb{F}_p^{b_2 \times b_1}$ and $z \in \mathbb{G}_2^{b_1}$. The obtained protocol can compute $e(g_1; g_2)^{u^T \cdot y}$ with $\mathcal{O}(\sqrt{m})$ cryptographic operations on the Verifier side and is given in Figure 2.

Lemma 5. *The protocol of Figure 2 for publicly delegation of a size m external group dot-product verification with rank-1 left hand side is sound, perfectly complete and requires the following number of operations where $b_1 b_2 \geq m$ and $b_1 + b_2 = \Theta(\sqrt{m})$:*

- *Preparation: $\mathcal{O}(b_1 + b_2)$ in $\mathbb{G}_i$;*
- *Prover: $\mathcal{O}(m)$ in $\mathbb{G}_i$;*
- *Verifier: $\mathcal{O}(m)$ in $\mathbb{F}_p$, $\mathcal{O}(b_1 + b_2)$ in $\mathbb{G}_i$ and $\mathcal{O}(b_1)$ pairings.*

Proof. Correctness is ensured by Equation (1). Soundness is given by Freivalds check. Complexity is as given in the Lemma: indeed, for the Verifier, we have: for $Y \cdot v$: $\mathcal{O}(b_2 b_1) = \mathcal{O}(m)$ classic operations; for $g_1^{\eta^T} \star (Yv)$: $\mathcal{O}(b_2)$ cryptographic (group) operations; for $z^T \star v$: $\mathcal{O}(b_1)$ cryptographic (group) operations; and for $\prod_{i=1}^{b_1} e(z[i]; g_2^{\mu[i]})$: $\mathcal{O}(b_1)$ cryptographic (pairings) operations. Then the preparation requires to compute $g_1^{\eta} \in \mathbb{G}_1^{b_2}$ and $g_2^{\mu} \in \mathbb{G}_1^{b_1}$, while the Prover needs to compute $g_1^{\eta^T} \star Y$ for $Y \in \mathbb{F}_p^{b_2 \times b_1}$ and $b_1 b_2 = \mathcal{O}(m)$. □

4.2 Rectangular general dot-product

Now if u is not given by a rank 1 update, one can still verify a dot-product with only $\mathcal{O}(\sqrt{m})$ pairings operations but as the price of slightly more group operations as given in Figure 3.

- *KeyGen*$(1^\lambda, \mu, \eta)$: given the security parameter 1^λ and vectors $\mu \in \mathbb{F}_p^{b_2}$ and $\eta \in \mathbb{F}_p^{b_1}$ such that $u = \mu \cdot \eta^t$, it selects two cyclic groups $\mathbb{G}_1$ and $\mathbb{G}_2$ of prime order p that admit a bilinear pairing $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ and generators g_1, g_2 and g_T of the three groups. Finally it outputs $\texttt{params} = \{b_1, b_2, p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2, g_T\}$ and $EK_f = (g_1^{\eta^T})$ and $VK_f = (g_2^{\mu})$.
- *ProbGen*(y): from $y \in \mathbb{F}_p^m$ it builds $Y \in \mathbb{F}_p^{b_1 \times b_2}$ and outputs $\sigma_x = Y$.
- *Compute*(σ_x, EK_f): compute $z^T = g_1^{\eta^T} \star Y$ and outputs $\sigma_y = (z^T)$.
- *Verify*(σ_y, VK_f): it starts by sampling randomly a vector $v \in \mathbb{F}_p^{b_2}$ then it checks whether $z^T \star v$ is equal to $g_1^{\eta^T} \star (Yv)$ or not. If the test passes it returns $\prod e(z[i]; g_2^{\mu[i]})$ and if it fails it returns $\perp$.

Figure 2: Publicly delegatable protocol for the dot-product in an external group with a rank-1 left hand side.

- *KeyGen*$(1^\lambda, u)$: given the security parameter 1^λ and vector $u \in \mathbb{F}_p^m$, it selects two cyclic groups $\mathbb{G}_1$ and $\mathbb{G}_2$ of prime order p that admit a bilinear pairing $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ and generators g_1, g_2 and g_T of the three groups and also integers b_1, b_2 such that $m = b_1 b_2$ and it outputs $\texttt{params} = \{b_1, b_2, p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2, g_T\}$. Then it samples a random $w \in \mathbb{F}_p^{b_1}$, creates $U \in \mathbb{F}_p^{b_1 \times b_2}$ from u and finally it outputs and $EK_f = (g_1^U)$ and $VK_f = (g_1^{w^T \cdot U}, g_2^{\omega^T})$.
- *ProbGen*(y): from $y \in \mathbb{F}_p^m$ it builds $Y \in \mathbb{F}_p^{b_2 \times b_1}$ and outputs $\sigma_x = Y$.
- *Compute*(σ_x, EK_f): compute $C = g_1^U \star Y$ and outputs $\sigma_y = (C)$.
- *Verify*(σ_y, VK_f): it starts by sampling randomly a vector $v \in \mathbb{F}_p^{b_1}$ then it computes $z = C \star v$ and checks whether $\prod e(z[i]; g_2^{\omega[i]})$ is equal to $e(g_1^{w^T \cdot U} \star (Yv); g_2)$ or not. If the test passes it returns $Trace(C)$ and if it fails it returns $\perp$.

Figure 3: Publicly delegatable protocol for the external dot-product.

Lemma 6. *The protocol of Figure 3 is sound, perfectly complete and requires the following number of operations with $b_1 b_2 \geq m$:*

- *Preparation: $\mathcal{O}(m)$ in $\mathbb{F}_p$ and $\mathcal{O}(m)$ in $\mathbb{G}_i$;*
- *Prover: $\mathcal{O}(mb_1)$ in $\mathbb{G}_1$;*
- *Verifier: $\mathcal{O}(m)$ in $\mathbb{F}_p$, $\mathcal{O}\left(b_1^2 + b_2\right)$ in $\mathbb{G}_i$ and $\mathcal{O}(b_1)$ pairings.*

Proof. Correctness is ensured by the vectorization in Equation (1). Soundness is given by the Freivalds check. Complexity is as given in the Lemma: indeed, for the Verifier, we have:

1. $Y \cdot v$: $\mathcal{O}(b_2 b_1) = \mathcal{O}(m)$ classic operations;
2. $g_1^{w^T U} \star (Yv)$: $\mathcal{O}(b_2)$ cryptographic (group) operations;
3. $z = C \star v$: $\mathcal{O}\left(b_1^2\right)$ cryptographic (group) operations;
4. $\prod_{i=1}^{b_1} e(z[i]; g_2^{w[i]})$: $\mathcal{O}(b_1)$ cryptographic (pairings) operations;

Then the preparation requires to compute $w^T \times U$. This is $\mathcal{O}(b_1 b_2 = m)$ operations. Finally, the Prover needs to compute the matrix multiplication $g_1^U \star Y$ for $U \in \mathbb{F}_p^{b_1 \times b_2}$ and $Y \in \mathbb{F}_p^{b_2 \times b_1}$, in $\mathcal{O}\left(b_1^2 b_2\right) = \mathcal{O}(mb_1)$. □

Therefore, one can take $b_1 = \mathcal{O}(\sqrt[3]{m})$ and $b_2 = \mathcal{O}\left(m^{2/3}\right)$ which gives only $\mathcal{O}\left(m^{2/3}\right)$ cryptographic operations for the Verifier, and $\mathcal{O}\left(m^{4/3}\right)$ cryptographic operations for the Prover. Now a dot-product can be cut in $\frac{n}{k}$ chunks of size k. Then Each chunk can be checked with the protocol of Figure 3 and the final dot-product obtained by adding the chunks. This gives the complexity of Corollary 7.

Corollary 7. *There exist a protocol for the dot-product using: $\mathcal{O}\left(n^{1-a/3}\right)$ cryptographic operations for the Verifier and $\mathcal{O}\left(n^{1+a/3}\right)$ cryptographic operations for the Prover, for any $0 < a < 1$.*

Proof. We let $k = n^a$ and use n/k times the protocol of the previous point with $b_1 = \sqrt[3]{k}$ and $b_2 = k^{\frac{2}{3}}$. Overall this gives $n/k(k^{4/3})$ for the Prover and $n/k(k^{2/3})$ for the Verifier. □

5 Public delegatability via bootstrapping

To recover the public delegatability model, we use the protocol of Figure 1 but we trade back some cryptographic operations using the protocol of Figure 2 to the Verifier. With an initial matrix $A \in \mathbb{F}_p^{m\times n}$ we however trade back only on the order of $\mathcal{O}(\sqrt{m}+\sqrt{n})$ cryptographic operations. This gives a slower verification in practice but interaction is not needed anymore. We present our full novel protocol for matrix vector product in Figure 4 (with the flow of exchanges shown in Figure 7, Appendix E). Apart from Freivalds's

- *Keygen*$(1^\lambda, A)$: given a security parameter 1^λ and a matrix $A \in \mathbb{F}_p^{m\times n}$, it selects two cyclic groups $\mathbb{G}_1$ and $\mathbb{G}_2$ of prime order p that admit a bilinear pairing $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$. Then it selects generators g_1, g_2 and g_T of the three groups, and integers b_1, b_2, c_1, c_2, d_1, d_2 such that $b_1b_2 = m$ and $c_1c_2 = d_1d_2 = n$ with $b_1 + b_2 \in \mathcal{O}\left(m^{1/2}\right)$, $c_1 + c_2 \in \mathcal{O}\left(n^{1/2}\right)$, $d_1 \in \mathcal{O}\left(n^{1/3}\right)$, $d_2 \in \mathcal{O}\left(n^{2/3}\right)$ and it outputs parameters:
  $\texttt{param} = (m, n, b_1, b_2, c_1, c_2, d_1, d_2, p, e, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, g_1, g_2, g_T)$.
  Then it samples randomly $\mu \in \mathbb{F}_p^{b_1}$, $\eta \in \mathbb{F}_p^{b_2}$, $(\rho_i)_{i=1}^2 \in \mathbb{F}_p^{c_1}$, $(\tau_i)_{i=1}^2 \in \mathbb{F}_p^{c_2}$, $\varpi \in \mathbb{F}_p^{d_1}$, $(\gamma, \delta) \in \mathbb{F}_p^2$ and $v \in \mathbb{F}_p^n$ with $v \simeq V \in \mathbb{F}_p^{d_1\times d_2}$.
  Finally it computes $U = \mu\cdot\eta^T \simeq u \in \mathbb{F}_p^n$, $T = \rho_1\cdot\tau_1^T + \rho_2\cdot\tau_2^T \simeq t^T \in \mathbb{F}_p^n$, $\omega^T = g_1^{u^T\cdot A + t^T + \gamma\delta v^T} \in \mathbb{F}_p^n$ and publishes:
  $EK_f = \{A, \omega^T, (g_1^{\tau_1^T}), (g_1^{\tau_2^T}), (g_1^{\eta^T}), (g_1^{\delta\cdot V})\}$
  $VK_f = \{(g_1^{\tau_1^T}), (g_1^{\tau_2^T}), (g_2^{\rho_1}), (g_2^{\rho_2}), (g_1^{\eta^T}), (g_2^{\mu}), (g_1^{\delta\varpi^T\cdot V}), (g_2^{\gamma\varpi}), g_2^{\gamma}\}$.
- *ProbGen*(x): given a vector $x \in \mathbb{F}_p^n$, output $\sigma_x = x$.
- *Compute*(σ_x, EK_f): given the encoded input σ_x and the evaluation key EK_f, it computes $y = Ax$, $\zeta = \omega^T \star x$, $s_i^T = g_1^{\tau_i^T} \star X$ for $i \in \{1, 2\}$ with $x \simeq X \in \mathbb{F}_p^{c_2\times c_1}$, $z^T = g_1^{\eta^T} \star Y$ with $y \simeq Y \in \mathbb{F}^{b_2\times b_1}$ and $C = g_1^{\delta\cdot V} \star X$ with $x \simeq X \in \mathbb{F}_p^{d_2\times d_1}$. Finally it outputs $\sigma_y = \{y, \zeta, (s_i^T)_{i=1}^2, z, C\}$.
- *Verify*(σ_y, VK_f): it starts by sampling randomly four vectors $(v_1, v_2) \in \mathbb{F}_p^{c_1}$, $v_3 \in \mathbb{F}_p^{b_1}$ and $v_4 \in \mathbb{F}_p^{d_1}$. Then for $i = 1, 2$ it checks whether $s_i^T \star v_i$ is equal to $g_1^{\tau_i^T} \star (X \cdot v_1)$ or not, then it computes $D_i = \prod e(s_i[j]; g_2^{\rho_1[j]})$ for both i. Then it checks whether $z^T \star v_3$ is equal to $g_1^{\eta^T} \star (Yv_3)$ or not, then it computes $H = \prod e(z[i]; g_2^{\mu[i]})$. From there it computes $\vartheta = C \star v_4$, and it checks whether $\prod e(\vartheta[i]; g_2^{\gamma\varpi[i]})$ is equal to $e(g_1^{\delta\varpi^T\cdot V} \star (Xv_4); g_2^{\gamma})$. If all the previous tests have passed it performs a final check which is to know if $e(\zeta; g_2)$ is equal to $H \cdot D_1 \cdot D_2 \cdot e(Trace(C); g_2^{\gamma})$ or not, if it passes it output y, but if any of the previous check has failed it outputs $\perp$.

Figure 4: Proven publicly delegatable protocol for matrix-vector product

checks and vectorization, we need to use a masking of the form $u^TA + t^T$ (see Figure 1) indistinguishable from a random distribution, but:

1. We have to add an extra component $\gamma\delta v^T$ to u^TA+t^T so that it is possible, when proving the reduction to co-CDH, to make up a random vector $\omega^T = g^{u^TA+t}$ where the components of u^TA+t^T are canceled out. This component cannot be revealed to the Prover nor the Verifier, otherwise its special structure could have been taken into account by the reduction. Also this component cannot have the rank-1 update structure as its has to be a multiple of $u^TA + t^T$. Therefore only the protocol of Figure 3 can be used to check the dotproduct with g^{v^T}.
2. To be able to apply the analysis of [6, Theorem 3] while allowing fast computations with t, we use a special form for t, namely: $t^T = \rho_1\tau_1^T + \rho_2\tau_2^T$.

With these modifications we are able to prove the soundness of the protocol in Figure 4.

Theorem 8. *Let $A \in \mathbb{F}_p^{m\times n}$ whose matrix-vector products costs $\mu(A)$ arithmetic operations. Protocol 4 is sound under the co-CDH assumption, perfectly complete and its number of performed operations is bounded as follows:*

	Preparation	*Prover*	*Verifier*
$\mathbb{F}_p$	$\mu(A)+\mathcal{O}(m+n)$	$\mu(A)$	$\mathcal{O}(m+n)$
$\mathbb{G}_i$	$\mathcal{O}(m+n)$	$\mathcal{O}\left(m+n^{4/3}\right)$	$\mathcal{O}\left(\sqrt{m}+n^{2/3}\right)$
Pairings	*0*	*0*	$\mathcal{O}(\sqrt{m}+\sqrt{n})$

The proof of Theorem 8 is given in Appendix E.

Remark 9. *Fast matrix multiplication can be used for the computation of C in the protocol of Figure 4. This decreases the $\mathcal{O}\left(n^{4/3}\right)$ factor of the Prover to $\mathcal{O}\left(n^{(1+\omega)/3}\right)$ where ω is the exponent of matrix-matrix multiplication. The currently best known exponent, given in [8], is $\omega \leq 2.3728639$. This immediately yields a reduced bound for the Prover of $\mu(A) + \mathcal{O}\left(m + n^{1.12428797}\right)$. This together with Corollary 7 can produce a protocol with Prover complexity bounded by $\mu(A) + \mathcal{O}\left(m + n^{1+0.12428797a}\right)$, for any $0 < a < 1$, while the Verifier complexity is $\mathcal{O}(n)$ classical operations and $o(n)$ cryptographic operations.*

6 Conclusion and experiments

We first recall in Table 1 the leading terms of the complexity bounds for our protocols and those of [6, 14, 4] (that is each value x in a cell is such that the actual cost is bounded by $x + o(x)$). There, we denote the base field operations by $\cdot\mathcal{F}$, the cryptographic group exponentiations or pairing operations by $\cdot\mathcal{G}$, and the cost of a product of the matrix $A \in \mathbb{F}_p^{m\times n}$ by a vector is $\mu(A)$. We see that our protocols are

Table 1: Leading terms for the time and memory complexity bounds (exchange of A, x and y excluded).

Scheme	[6]	[14]	[4]
Mode	Public verif.	Public deleg.	Public deleg.
Preparator (KeyGen)	$2mn\cdot\mathcal{F} + mn\cdot\mathcal{G}$	–	$2mn\cdot\mathcal{F} + 2mn\cdot\mathcal{G}$
Trustee (ProbGen)	$2(m+n)\cdot\mathcal{F} + 2m\cdot\mathcal{G}$	$mn\cdot\mathcal{F} + (2m+n)\cdot\mathcal{G}$	$n\cdot\mathcal{G}$
Prover (Compute)	$\mu(A)\cdot\mathcal{F} + 2mn\cdot\mathcal{G}$	$\mu(A)\cdot\mathcal{F} + 2mn\cdot\mathcal{G}$	$\mu(A)\cdot\mathcal{F} + 2mn\cdot\mathcal{G}$
Verifier	$2m\cdot\mathcal{G}$	$2m\cdot\mathcal{G}$	$m\cdot\mathcal{G}$
Extra storage	$\mathcal{O}(mn)$	$\mathcal{O}(mn)$	$\mathcal{O}(mn)$
Extra communications	$\mathcal{O}(m)$	$\mathcal{O}(m)$	$\mathcal{O}(1)$

Scheme	Figure 1	Figure 4
Mode	Public verif.	Public deleg.
Preparator (KeyGen)	$(\mu(A)+n)\cdot\mathcal{F} + n\cdot\mathcal{G}$	$(\mu(A)+m+5n)\cdot\mathcal{F} + 2n\cdot\mathcal{G}$
Trustee (ProbGen)	$2(m+n+1)\cdot\mathcal{F} + 1\cdot\mathcal{G}$	0
Prover (Compute)	$\mu(A)\cdot\mathcal{F} + 2n\cdot\mathcal{G}$	$\mu(A)\cdot\mathcal{F} + (2n^{4/3}+m)\cdot\mathcal{G}$
Verifier	$1\cdot\mathcal{G}$	$(2m+4n)\cdot\mathcal{F} + (6\sqrt{m}+2n^{2/3})\cdot\mathcal{G}$
Extra storage	$\mathcal{O}(m+n)$	$\mathcal{O}(n)$
Extra communications	$\mathcal{O}(1)$	$\mathcal{O}\left(n^{2/3}+\sqrt{m}\right)$

suitable to sparse or structured matrix-vector multiplication as they never require $\mathcal{O}(mn)$ operations but rather $\mu(A)$. Moreover, we see that most of the Verifier's work is now in base field operations were it was cryptographic operations for previously known protocols. As shown in Table 2 and in Figure 5, this is very useful in practice, even for dense matrices. For these experiments we compare with our own implementations of the protocols of [6, 14, 4] over the PBC library[1] [9] for the pairings and the FFLAS-FFPACK library[2] [3]

[1] https://crypto.stanford.edu/pbc, version 0.5.14
[2] http://linbox-team.github.io/fflas-ffpack, version 2.2.2

Table 2: Matrix-vector multiplication public verification over a 256-bit finite field with different protocols on a i7 @3.4GHz.

	1000×1000					2000×2000			
	[12]	[6]	[14]	[4]	Fig. 4	[6]	[14]	[4]	Fig. 4
KeyGen	141.68s	152.62s	-	154.27s	**0.80s**	615.81s	-	612.72s	**1.75s**
ProbGen	-	1.25s	2.28s	2.30s	-	2.13s	4.98s	4.56s	-
$Ax = y$	20.14s	**0.19s**	**0.19s**	**0.19s**	**0.19s**	**0.78s**	**0.78s**	**0.78s**	**0.78s**
Compute	188.60s	273.06s	433.88s	271.03s	**2.26s**	1097.96s	1715.46s	1079.71s	**5.37s**
Verify	2.06s	26.62s	27.56s	**0.33s**	0.90s	52.60s	55.79s	**0.62s**	1.19s

	4000×4000				8000×8000			
	[6]	[14]	[4]	Fig. 4	[6]	[14]	[4]	Fig. 4
KeyGen	2433.10s	-	2452.98s	**4.89s**	9800.42s	-	9839.26s	**15.64s**
ProbGen	3.81s	13.29s	9.24s	-	7.41s	43.44s	18.46s	-
$Ax = y$	**3.28s**	**3.28s**	**3.28s**	**3.28s**	**13.30s**	**13.30s**	**13.30s**	**13.30s**
Compute	4360.43s	6815.40s	4329.46s	**13.76s**	17688.69s	27850.90s	17416.38s	**37.00s**
Verify	103.14s	107.99s	**1.20s**	1.65s	211.07s	220.69s	2.37s	**2.25s**

for the exact linear algebra over finite fields (C++ source files are available on request via the PC and will be publicly posted on our web site if the paper is accepted). We used randomly generated dense matrices

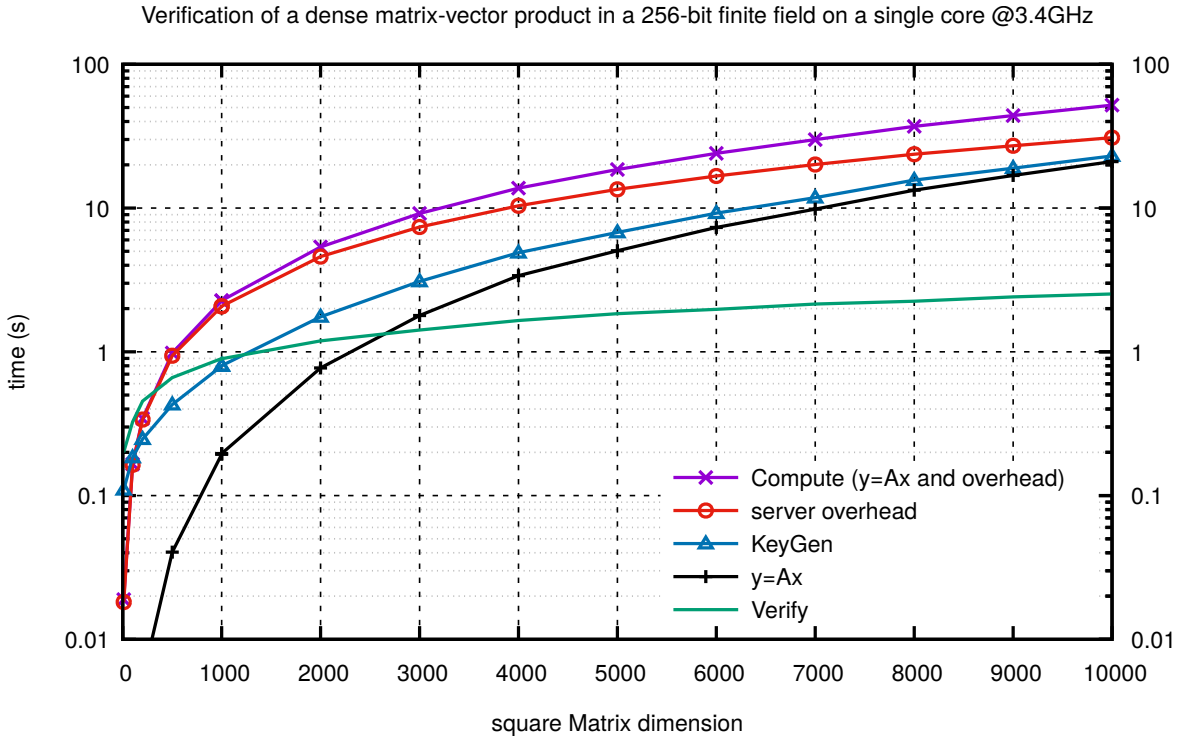

Figure 5: Protocol 4 performance.

and vectors and to optimize the costs (pairings are more expensive than exponentiations), we also chose the following parameters for the vectorizations: $b_1 = \lceil \sqrt{m}/10 \rceil$; $b_2 = \lceil 10\sqrt{m} \rceil$; $c_1 = \lceil \sqrt{n}/10 \rceil$; $c_2 = \lceil 10\sqrt{n} \rceil$;

$d_1 = \lceil n^{1/3}/3 \rceil$; $d_2 = \lceil 3n^{2/3} \rceil$. We indeed chose a type 3 pairing over a Barreto-Naehrig curve [1] based on a 256-bits prime field, which should guarantee 128 bits of security. First, with $\mathbb{F}_p$ the 256-bits prime field[3], $\mathbb{G}_1$ is the group of $\mathbb{F}_p$-rational points $E(\mathbb{F}_p)$ with parameters: $\mathbb{G}_1$ $(E) : y^2 = x^3 + 6$, modulo p. Second, $\mathbb{G}_2$ is a subgroup of a sextic twist of E defined over $\mathbb{F}_{p^2}$ denoted $E'(\mathbb{F}_{p^2})$ with parameters[4]: $\mathbb{G}_2$ $(E') : y^2 = x^3 + 6e$, $\mathbb{F}_{p^2} \cong \mathbb{F}_p[X]/(X^2 - 2)$, $e = a_0 + a_1 X \in \mathbb{F}_{p^2}$. The third group $\mathbb{G}_T$ is then a subgroup of the multiplicative group of the field $\mathbb{F}_{p^{12}}$. This curve is reasonably well-suited to our needs and is supported by the PBC library. We used it for our own implementation of the three protocols [6, 14, 4] as well as for ours, Section 5 and Figure 7.

In the first set of timings of Table 2, we also compare the latter protocols with a compiled verifiable version obtained via the Pepper software[5] [12]. This software uses a completely different strategy, namely that of compiling a C program into a verifiable one. We added the timings for $n = 1000$ as a comparison, but the Pepper compilation thrashed on our 64 GB machine for $n \geq 2000$.

In terms of Prover time, we see that our protocols are between two to three orders of magnitude faster than existing ones (further evidence is given in Table 4, Appendix E). Moreover, overall we see that with the new protocol, the data preparation (KeyGen) is now very close to a single non-verified computation and that the work of the Prover can be less than three times that of a non-verified computation (note first, that in both Table 2 and Figure 5, the "Compute" fields include the computation of $y = Ax$, and, second, that the Prover overhead being asymptotically faster than the compute time, this latter overhead is rapidly amortized). Finally, only the protocol of [4] did exhibit a verification step faster than the computation itself for size 2000×2000 whereas, as shown in Figure 5, our protocol achieves this only from size 3000×3000. However, we see that we are competitive for larger matrices. Moreover, as shown by the asymptotics of Theorem 8, our overall performance outperforms all previously known protocols also in practice, while keeping an order of magnitude faster Verification time.

References

[1] Paulo S. L. M. Barreto and Michael Naehrig. Pairing-friendly elliptic curves of prime order. In Bart Preneel and Stafford Tavares, editors, *Selected Areas in Cryptography: 12th International Workshop, SAC 2005, Kingston, ON, Canada, August 11-12, 2005, Revised Selected Papers*, pages 319–331, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. doi:10.1007/11693383_22.

[2] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Assche. Sponge-based pseudo-random number generators. In Stefan Mangard and François-Xavier Standaert, editors, *CHES 2010*, pages 33–47. Springer, 2010. doi:10.1007/978-3-642-15031-9_3.

[3] Jean-Guillaume Dumas, Pascal Giorgi, and Clément Pernet. Dense linear algebra over prime fields: the FFLAS and FFPACK packages. *ACM Transactions on Mathematical Software*, 35(3):1–42, November 2008. doi:10.1145/1391989.1391992.

[4] Kaoutar Elkhiyaoui, Melek Önen, Monir Azraoui, and Refik Molva. Efficient techniques for publicly verifiable delegation of computation. In *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security*, ASIA CCS '16, pages 119–128, New York, NY, USA, 2016. ACM. doi:10.1145/2897845.2897910.

[5] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In A. M. Odlyzko, editor, *Advances in Cryptology - CRYPTO'86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer-Verlag, 1987, 11–15 August 1986. URL: http://www.cs.rit.edu/~jjk8346/FiatShamir.pdf.

[3] $p = 57896044618658115533954196422662521694340972374557265300857239534749215487669$

[4] a_0=527250522724512898182991239521675688175482150373036387310978085617039101783 75, a_1=390302625865493553046028116363993748397589815144007427619200754037365

[5] https://github.com/pepper-project/pepper, git: fe3bf04

[6] Dario Fiore and Rosario Gennaro. Publicly verifiable delegation of large polynomials and matrix computations, with applications. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, CCS '12, pages 501–512, New York, NY, USA, 2012. ACM. doi:10.1145/2382196.2382250.

[7] Rūsiņš Freivalds. Fast probabilistic algorithms. In J. Bečvář, editor, *Mathematical Foundations of Computer Science 1979*, volume 74 of *Lecture Notes in Computer Science*, pages 57–69, Olomouc, Czechoslovakia, September 1979. Springer-Verlag. doi:10.1007/3-540-09526-8_5.

[8] François Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*, ISSAC '14, pages 296–303, New York, NY, USA, 2014. ACM. doi:10.1145/2608628.2608664.

[9] Ben Lynn. The pairing-based cryptography (PBC) library, 2010. URL: https://crypto.stanford.edu/pbc.

[10] NIST. *FIPS publication 202: SHA-3 standard: permutation-based hash and extendable-output functions*, August 2015. doi:10.6028/NIST.FIPS.202.

[11] Bryan Parno, Mariana Raykova, and Vinod Vaikuntanathan. How to delegate and verify in public: verifiable computation from attribute-based encryption. In Ronald Cramer, editor, *Theory of Cryptography: 9th Theory of Cryptography Conference, TCC 2012, Taormina, Sicily, Italy, March 19-21, 2012. Proceedings*, pages 422–439, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. doi:10.1007/978-3-642-28914-9_24.

[12] Srinath T. V. Setty, Richard McPherson, Andrew J. Blumberg, and Michael Walfish. Making argument systems for outsourced computation practical (sometimes). In *19th Annual Network and Distributed System Security Symposium, NDSS 2012, San Diego, California, USA, February 5-8, 2012*. The Internet Society, 2012. URL: http://www.internetsociety.org/sites/default/files/04_3.pdf.

[13] Michael Walfish and Andrew J. Blumberg. Verifying computations without reexecuting them. *Communications of the ACM*, 58(2):74–84, January 2015. doi:10.1145/2641562.

[14] Yihua Zhang and Marina Blanton. Efficient secure and verifiable outsourcing of matrix multiplications. In Sherman S.M. Chow, Jan Camenisch, Lucas C.K. Hui, and Siu Ming Yiu, editors, *Information Security*, pages 158–178. Springer International Publishing, 2014. doi:10.1007/978-3-319-13257-0_10.

A Fiore and Gennaro's protocol

For the sake of completeness, we present here the original protocol for matrix-vector verification in [6], but with our rank-one update view. It stems from the fact that if s, t, ρ, τ are randomly generated vectors then the function $g^{M[i,j]}$, where $M = s \cdot t^T + \rho \cdot \tau^T$, is a pseudorandom function [6, Theorem 3], provided that the *Decision Linear* assumption [6, Definition 3] holds (a generalization of the External Diffie-Hellman assumption for pairings).

- *KeyGen*: for $A \in \mathbb{F}_p^{m \times n}$, generate 3 multiplicative groups $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$ of prime order p, with g_1 generating $\mathbb{G}_1$ (resp. g_2 generating $\mathbb{G}_2$), and a bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$. Generate $2(m+n)$ secret random values $s \in \mathbb{F}_p^m$, $t \in \mathbb{F}_p^n$, $\rho \in \mathbb{F}_p^m$, $\tau \in \mathbb{F}_p^n$. Compute $W[i,j] = g_1^{\alpha A[i,j] + s[i]t[j] + \rho[i]\tau[j]} \in \mathbb{G}_1$, give W to the server and publish $a = e(g_1^\alpha, g_2) \in \mathbb{G}_T$.

Let $x \in \mathbb{F}_p^n$ be a query vector

- *ProbGen*: compute $\mathrm{VK}_x \in \mathbb{G}_T^m$, such that $d = t^T \cdot x \in \mathbb{F}_p$, $\delta = \tau^T \cdot x \in \mathbb{F}_p$, and $\mathrm{VK}_x[i] = e(g_1^{s[i]d + \rho[i]\delta}; g_2) \in \mathbb{G}_T$.

- *Compute*: compute $y = Ax$ and $z = W \star x \in \mathbb{G}_1^m$ (that is $z[i] = \prod_{j=1}^n W[i,j]^{x[j]}$).
- *Verify*: check that $e(z[i]; g_2) = a^{y[i]}\mathrm{VK}_x[i]$, for all $i = 1, \ldots, m$.

In this protocol, the flow of communications is as follows:

1. Preparator: secret random $\alpha \in \mathbb{F}_p$, $s \in \mathbb{F}_p^m$, $t \in \mathbb{F}_p^n$, $\rho \in \mathbb{F}_p^m$, $\tau \in \mathbb{F}_p^n$. then $W = g_1^{\alpha A + s \cdot t^T + \rho \cdot \tau^T}$
2. Preparator to Prover: $A \in \mathbb{F}_p^{m \times n}$, $W \in \mathbb{G}_1^{m \times n}$.
3. Preparator to Trustee: s, ρ, t, τ, in a secure channel.
4. Preparator publishes and signs $a = e(g_1; g_2)^\alpha \in \mathbb{G}_T$.
5. Verifier to both Prover and Trustee: $x \in \mathbb{F}_p^n$.
6. Trustee publishes and signs $\mathrm{VK}_x \in \mathbb{G}_T^m$ such that : $\mathrm{VK}_x = e(g_1; g_2)^{s \cdot (t^T \cdot x) + \rho \cdot (\tau^T \cdot x)}$.
7. Prover to Verifier: $y \in \mathbb{F}_p^m$, $z \in \mathbb{G}_1^m$.
8. Verifier public verification: $e(z; g_2) \stackrel{?}{=\!=} a^y \mathrm{VK}_x$ (component-wise in $\mathbb{G}_T^m$).

This protocol is sound, complete and publicly verifiable. It however uses many costly exponentiations and pairings operations that renders it inefficient in practice: even though the Client and Trustee number of operations is linear in the vector size, it takes still way longer time that just computing the matrix-vector product in itself, as shown in the experiment Section 6.

In this paper, our aim was first to adapt this protocol to the sparse/structured case, and, second, to reduce the number of cryptographic operations in order to obtain a protocol efficient in practice.

B Probabilistic verification and the random oracle model

First we recall that private verification is very fast and does not require any cryptographic routines. Then we show that this allows to obtain a very efficient protocol in the random oracle model, but for a fixed number of inputs.

B.1 Private verification

Without any recourse to cryptography, it is well known how to privately verify a matrix-vector multiplication. The idea is to use Freivalds test [7], *on the left*, provided that multiplication by the transpose matrix is possible:

- Verifier to Prover: A, $x_i \in \mathbb{F}_p^n$, for $i = 1, \ldots, k$.
- Prover to Verifier: $y_i \in \mathbb{F}_p^m$, for $i = 1, \ldots, k$.
- Verifier verification: random $u \in \mathbb{F}_p^m$, then $w^T = u^T \cdot A$, and finally check, for $i = 1, \ldots, k$, that $w^T \cdot x_i \stackrel{?}{=\!=} u^T \cdot y_i$ in $\mathbb{F}_p$.

On the one hand, this protocol uses only classical arithmetic and is adaptable to sparse matrices, that is when a matrix vector product costs $\mu(A)$ operations with $\mu(A) < 2mn$ (this is the case for instance if the matrix is not structured but is sparse with $\mu(A)/2 < mn$ non-zero elements). Indeed, in the latter case, the cost for the Prover is $k\mu(A)$, where the cost for the Verifier is $\mu(A) + 4kn$.

On the other hand, the protocol has now Freivalds probability of revealing an error in any of the y_i: $1 - 1/p$, if $\mathbb{F}_p$ is of cardinality p (or $1 - 1/p^\ell$ if u is chosen in an extension of degree ℓ of $\mathbb{F}_p$).

B.2 Public verification in the random oracle model

Using Fiat-Shamir heuristic [5], the privately verifiable certificate of Section B.1 can be simulated non-interactively: uniformly sampled random values produced by the Verifier are replaced by cryptographic hashes (to prove security in the random oracle model) of the input and of previous messages in the protocol. Complexities are preserved, as producing cryptographically strong pseudo-random bits by a cryptographic hash function (e.g., like the extendable output functions of the SHA-3 family defined in [2, 10]), is linear in the size of both its input and output (with atomic operations often even faster than finite field ones):

- Preparator to Prover: $A \in \mathbb{F}_p^{m\times n}$.
- Verifier to Prover: $x_i \in \mathbb{F}_p^n$, for $i = 1, \ldots, k$.
- Prover to Verifier: $y_i \in \mathbb{F}_p^m$, for $i = 1, \ldots, k$.
- Verifier to Trustee: all the x_i and y_i.
- Trustee publishes and signs both $u \in \mathbb{F}_p^m$ and $w \in \mathbb{F}_p^n$ such that: $u = Hash(A, x_1, \ldots, x_k, y_1, \ldots, y_k) \in \mathbb{F}_p^m$, then $w^T = u^T \cdot A \in \mathbb{F}_p^n$.
- Verifier public verification: $w^T \cdot x_i \stackrel{?}{=\!=} u^T \cdot y_i$ in $\mathbb{F}_p$.

There is absolutely no overhead for the Prover; the cost for the Trustee is a single matrix-vector product for any k plus a cost linear in the input size; and the cost for the Verifier is $\mathcal{O}(nk)$. Using Fiat-Shamir heuristic this gives a possibility for an afterwards public verification (that is after the computations), but this not possible to test new vectors once u has been revealed.

C Proof of Theorem 4

We here give the proof of Theorem 4, page 4, recalled hereafter.

Theorem 4. *The protocol of Figure 1 is perfectly complete and sound under the co-Computational Diffie-Hellman Problem assumption.*

Proof. For the correctness, we have that: $\zeta_i = g_1^{(u^TA+t^T)\cdot x_i} = g_1^{u^T\cdot y_i + t^T\cdot x_i} = g_1^{h_i} g_1^{d_i} = g_1^{h_i+d_i}$. Then, by bilinearity, $e(\zeta_i; g_2) = e(g_1; g_2)^{h_i+d_i} = \eta_i$.

For the soundness, a malicious Prover can guess the correct output values, but this happens once in the number of elements of $\mathbb{G}_T$. Otherwise he could try to guess some matching h_i and d_i, but that happens less than one in the number of elements of $\mathbb{F}_p$. Finally, the Prover could produce directly ζ_i. Suppose then it is possible to pass our verification scheme for some A, x and $y' \neq y = Ax$. Then without loss of generality, we can suppose that the first coefficients of both vectors are different, $y'[1] \neq y[1]$ (via row permutations) and that $y'[1] - y[1] = 1$ (via a scaling).

Take a co-computational Diffie-Hellman problem (g_1^c, g_2^d), where g_1^{cd} is unknown. Then denote by $a = e(g_1^c; g_2^d) = e(g_1^{cd}; g_2)$ and consider the vector $z^T = [a, e(1;1), \ldots, e(1;1)]$. Compute $\chi^T = z^T \star A$. The latter correspond to $\chi^T = e(g_1^{u^TA}; g_2)$ for (a not computed) $u^T = [cd, 0, \ldots, 0]$. Now randomly choose $\psi^T = [\psi_1, \ldots, \psi_n]$ and compute $\omega^T = g_1^{\psi^T}$. Compute also the vector $\phi^T = e(\omega^T; g_2)/\chi^T$ coefficient-wise. The latter correspond to $\phi^T = e(g_1^{t^T}; g_2)$ for $t^T = \psi^T - u^TA$. Finally, compute $\zeta = g_1^{\psi^T\cdot x}$ (indeed, then $\mu = e(\zeta; g_2) = e(g_1^{\psi^T\cdot x}; g_2) = \eta = e(g_1^{u^T\cdot y}; g_2)e(g_1^{t^T\cdot x}; g_2)$, that is $\eta = (\chi^T \star x)(\phi^T \star x)$ is actually $\eta = (z^T \star y)(\phi^T \star x)$). Now, if it is possible to break the scheme, then it is possible to compute ζ' that will pass the verification for y' as Ax, that is $e(\zeta'; g_2) = (z^T \star y')(\phi^T \star x)$. Let $h = u^Ty$, $d = t^Tx$ and $h' = u^Ty'$. Then $e(\zeta; g_2) = e(g_1^h; g_2)e(g_1^d; g_2)$ and $e(\zeta'; g_2) = e(g_1^{h'}; g_2)e(g_1^d; g_2)$. But $h' - h = u^T(y' - y) = cd(y'[1] - y[1]) = cd$ by construction. Therefore $\zeta'/\zeta = g_1^{cd}$, as e is non-degenerate, and the co-CDH is solved. □

D Small fields

The protocol of Figure 1 is quite efficient. We have made experiments with randomly generated dense matrices and vectors with the PBC library[1] for the pairings and the FFLAS-FFPACK library[2] for the exact linear algebra over finite fields. For instance, it is shown in Table 3, that for a 8000×8000 matrix over a field of size 256 bits, the protocol is highly practical: first, if the base field and the group orders are of similar sizes, the verification phase is very efficient; second, the overhead of computing ζ for the server is quite negligible and third, the key generation is dominated by the computation of one matrix-vector product.

Table 3: Verification of a 8000×8000 matrix-vector multiplication with different field sizes via the protocol in Figure 1 on a single core @3.4GHz.

Field size	$\lvert\mathbb{G}\rvert$	Security	KeyGen			Compute			Verify
			Total	u^TA	overhead	Total	$y=Ax$	overhead	
256	256	128	13.65s	12.34s	1.22s	15.72s	13.46s	2.26s	0.03s
10	322	128	1.96s	0.05s	1.81s	0.22s	0.09s	0.13s	0.04s

Differently, if the base field is small, say machine word-size, then having to use cryptographic sizes for the group orders can be penalizing for the Key Generation: multiplying a small field matrix A with a large field vector u^T is much slower than $y = Ax$ with x and A small. First of all, the computations must be compatible. For this, one possibility is to ask and verify instead for $y = Ax$ over $\mathbb{Z}$ and then to let the Verifier compute $y \mod p$ for himself. There, to reduce the overhead of computing u^TA, one can instead select the m values of the vector u as $u_\ell = \alpha r_i s_j$ with $\ell = i\lceil\sqrt{m}\rceil + j$ for α a randomly chosen large value and r_i, s_j some randomly chosen small values. Indeed then u^TA can be computed by first performing $(rs^T)A$ via $\mathcal{O}(\sqrt{m})$ matrix-vector computations with s (or a $\sqrt{m}\times n\sqrt{m}$ matrix-vector multiplication) followed by $\mathcal{O}(n\sqrt{m})$ multiplications by r (or a $n\times\sqrt{m}$ matrix-vector multiplication) where s_j and r_i are small values. Then it remains only to multiply a vector of small values by α. We have traded $\mathcal{O}(mn)$ operations with large values for $\mathcal{O}(\sqrt{m}n\sqrt{m} + n\sqrt{m})$ operations with small values and $\mathcal{O}(n)$ with large values.

Now, in order for the values to remain correct over $\mathbb{Z}$, the value of $(u^TA + t^T)x$ must not overflow. For this, one must choose a group order larger than mnp^4 (for $(rs^T)Ax$). Now the security is not anymore half the size of the group order but potentially half the size of the set from which t^T is selected, that is at most the group order size minus that of np (for t^Tx). To be conservative we even propose, as an estimated security of the obtained protocol, to consider only half the size of α (that is the size of the group order minus that of mnp^4). In terms of efficiency, the improvement is shown in Table 1, last row. On the one hand, the key generation is now dominant and can be amortized only after about 10 matrix-vector multiplications. On the other hand, the verification time starts to be faster than the computation time. This is also shown in Figure 6 where the equivalent of the last row in Table 3 is shown for different matrix dimensions.

E Proven publicly delegatable protocol with negligible cryptographic operations

We first give the proof of Theorem 8, page 8, recalled hereafter, for the correctness, soundness and complexity of the protocol in Figures 4 and 7. The flow of exchanges within our protocol is also illustrated in Figure 7.

Theorem 8. *Let $A \in \mathbb{F}_p^{m\times n}$ whose matrix-vector products costs $\mu(A)$ arithmetic operations. Protocol 4 is sound under the co-CDH assumption. It is also perfectly complete and its number of performed operations is bounded as follows:*

	Preparation	*Prover*	*Verifier*
$\mathbb{F}_p$	$\mu(A)+\mathcal{O}(m+n)$	$\mu(A)$	$\mathcal{O}(m+n)$
$\mathbb{G}_i$	$\mathcal{O}(m+n)$	$\mathcal{O}(m+n^{4/3})$	$\mathcal{O}(\sqrt{m}+n^{2/3})$
Pairings	*0*	*0*	$\mathcal{O}(\sqrt{m}+\sqrt{n})$

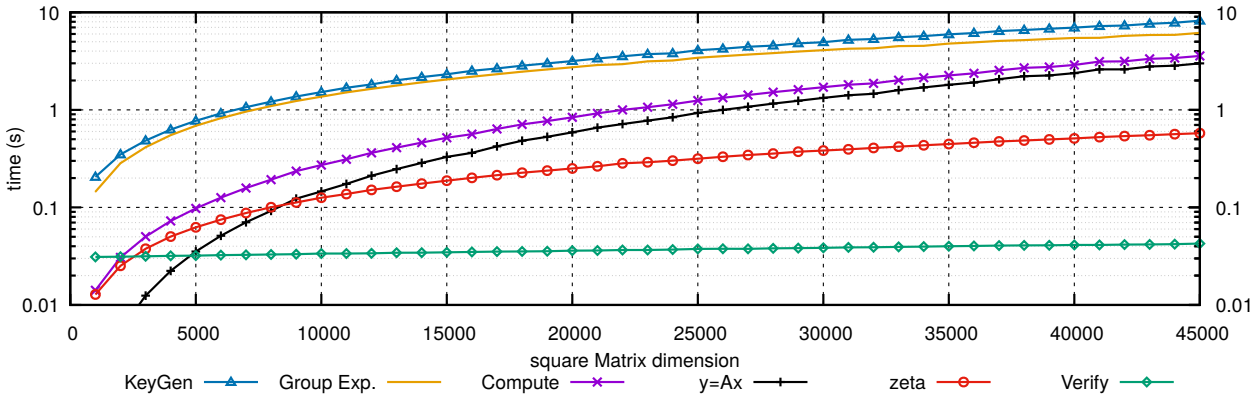

Figure 6: Trustee-helped Verification of a dense matrix-vector product in a 10-bits finite field on a single core @3.4GHz.

Proof. Completeness stems again directly from Equation (1).

For the complexity bounds, we fix $b_1 b_2 \geq m$ and $b_1 + b_2 = \Theta(\sqrt{m})$ (usually, pairing operations are costlier than group operations, therefore a good practice could be to take $b_1 < b_2$ and we, for instance, often have used $b_2 = 100 b_1$ with $b_1 b_2 \approx m$ which gave us a speed-up by a factor of 5), $c_1 c_2 \geq n$ and $c_1 + c_2 = \Theta(\sqrt{n})$, and finally $d_1 = \mathcal{O}\left(m^{1/3}\right)$ and $d_2 = \mathcal{O}\left(m^{2/3}\right)$. For the Prover, we then have that y obtained in $\mu(A)$ operations; ζ in $\mathcal{O}(n)$; s_1, s_2 and z computations are bounded by $\mathcal{O}(n+m)$ where C thus requires $\mathcal{O}\left(n^{4/3}\right)$ operations. The cost for the preparation is $\mathcal{O}(m+n)$ for U, T and $\varpi^T V$. ω requires $\mu(A) + 2m$ classical operations and $\mathcal{O}(m)$ group operations. $(g_1^{\delta})^{V^T}$ requires $\mathcal{O}(n)$ group operations while $g_1^{\tau_i}$, $g_1^{\eta_i}$, $g_2^{\rho_i}$, and $g_2^{\mu_i}$, require $\Theta(\sqrt{m} + \sqrt{n})$ operations, more than for $(g_2^{\gamma})^{\varpi^T}$ and $(g_1^{\delta})^{\varpi^T V}$. The complexity for the Verifier is then dominated by $\mathcal{O}\left(n^{2/3}\right)$ operations to check C, $\mathcal{O}(n)$ classical operations for $Y \cdot v_3$ and $\mathcal{O}(\sqrt{n})$ pairing operations.

Finally for the soundness, assume that there is an adversary $\mathcal{A}$ that breaks the soundness of our protocol with non-negligible advantage ϵ for a matrix $A \in \mathbb{F}_p^{m \times n}$. In the following we will prove how an adversary $\mathcal{B}$ can use adversary $\mathcal{A}$ to break the co-CDH assumption with non-negligible advantage $\epsilon' \simeq \epsilon$. Let assume that $\mathcal{B}$ was given a co-CDH sample $(L = g_1^a, R = g_2^b)$. First $\mathcal{B}$ simulates the soundness experiment to adversary $\mathcal{A}$ in the following manner: when $\mathcal{A}$ calls the oracle $\mathcal{O}_{KeyGen}$, adversary $\mathcal{B}$ first chooses integers, b_1, b_2, c_1, c_2, d_1, and d_2 such that $m = b_1 b_2$ and $n = b_1 b_2 = d_1 d_2$. Then it generates random vectors $\mu_0 \in \mathbb{F}_p^{b_1}, \eta_0 \in \mathbb{F}_p^{b_2}, \rho_{01} \in \mathbb{F}_p^{c_1}, \tau_{01} \in \mathbb{F}_p^{c_2}, \rho_{02} \in \mathbb{F}_p^{c_1}, \tau_{02} \in \mathbb{F}_p^{c_2}$, $\varpi \in \mathbb{F}_p^{d_1}$ and a value $r \in \mathbb{F}_p$. We let u_0 be the vector representation of $\mu_0 \cdot \eta_0^T$ and t_0 that of $\rho_{01} \cdot \tau_{01}^T + \rho_{02} \cdot \tau_{02}^T$. We also let $v = -(A^T u_0 + t_0) \in \mathbb{F}_p^n$. Finally, $\mathcal{B}$ forms $\omega^T = L^{r \cdot v^T}$; $g_1^{\eta} = L^{\eta_0}$, $g_2^{\mu} = R^{\mu_0}$; $g_1^{\tau_1} = L^{\tau_{01}}$, $g_2^{\rho_1} = R^{\rho_{01}}$; $g_1^{\tau_2} = L^{\tau_{02}}$, $g_2^{\rho_2} = R^{\rho_{02}}$; $g_1^{\delta} = L$, $g_2^{\gamma} = g_2^r$; $g_1^{\delta V} = L^V$, $(g_2^r)^{\varpi^T} = (g_2^{\gamma})^{\varpi^T}$ and $(g_1^{\delta})^{\varpi^T * V} = (g_1^{\gamma})^{\varpi^T V}$ and outputs:

$\texttt{param} = (m, n, b_1, b_2, c_1, c_2, d_1, d_2, p, e, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, g_1, g_2, g_T)$.

$EK_f = \{A, \omega^T, (g_1^{\tau_1^T}), (g_1^{\tau_2^T}), (g_1^{\eta^T}), (g_1^{\delta \cdot V})\}$

$VK_f = \{(g_1^{\tau_1^T}), (g_1^{\tau_2^T}), (g_2^{\rho_1}), (g_2^{\rho_2}), (g_1^{\eta^T}), (g_2^{\mu}), (g_1^{\delta \varpi^T \cdot V}), (g_2^{\gamma \varpi}), g_2^{\gamma}\}$.

Thanks to the randomness and the decisional Diffie-Hellman assumption (DDH) in each group G_i, as well as [6, Theorem 3] for $\omega^T = (L^r)^{v^T}$, these public values are indistinguishable from randomly generated inputs. Further, we have $\omega^T = g_1^{arv^T} = g_1^{ab(u_0^T A + t_0^T + v^T) + arv^T} = g_1^{abu_0^T A + abt_0^T + a(b+r)v^T}$.

When adversary $\mathcal{A}$ calls the oracle $\mathcal{O}_{ProbGen}$ on input x, adversary $\mathcal{B}$ returns $\sigma_x = x$. Therefore, if $y = Ax$ and $\zeta = \omega^T \star x$, then the verification will pass: indeed the first two checks will ensure that $s_1^T = g_1^{\tau_1^T} \star X$ and $s_2^T = g_1^{\tau_2^T} \star X$ when the third check ensures that $z^T = g_1^{\eta^T} \star Y$. This shows that:

$$H = \left(\prod e(z[i]; g_1^{\mu[i]})\right) = e(g_1; g_2)^{abu_0^T y},$$

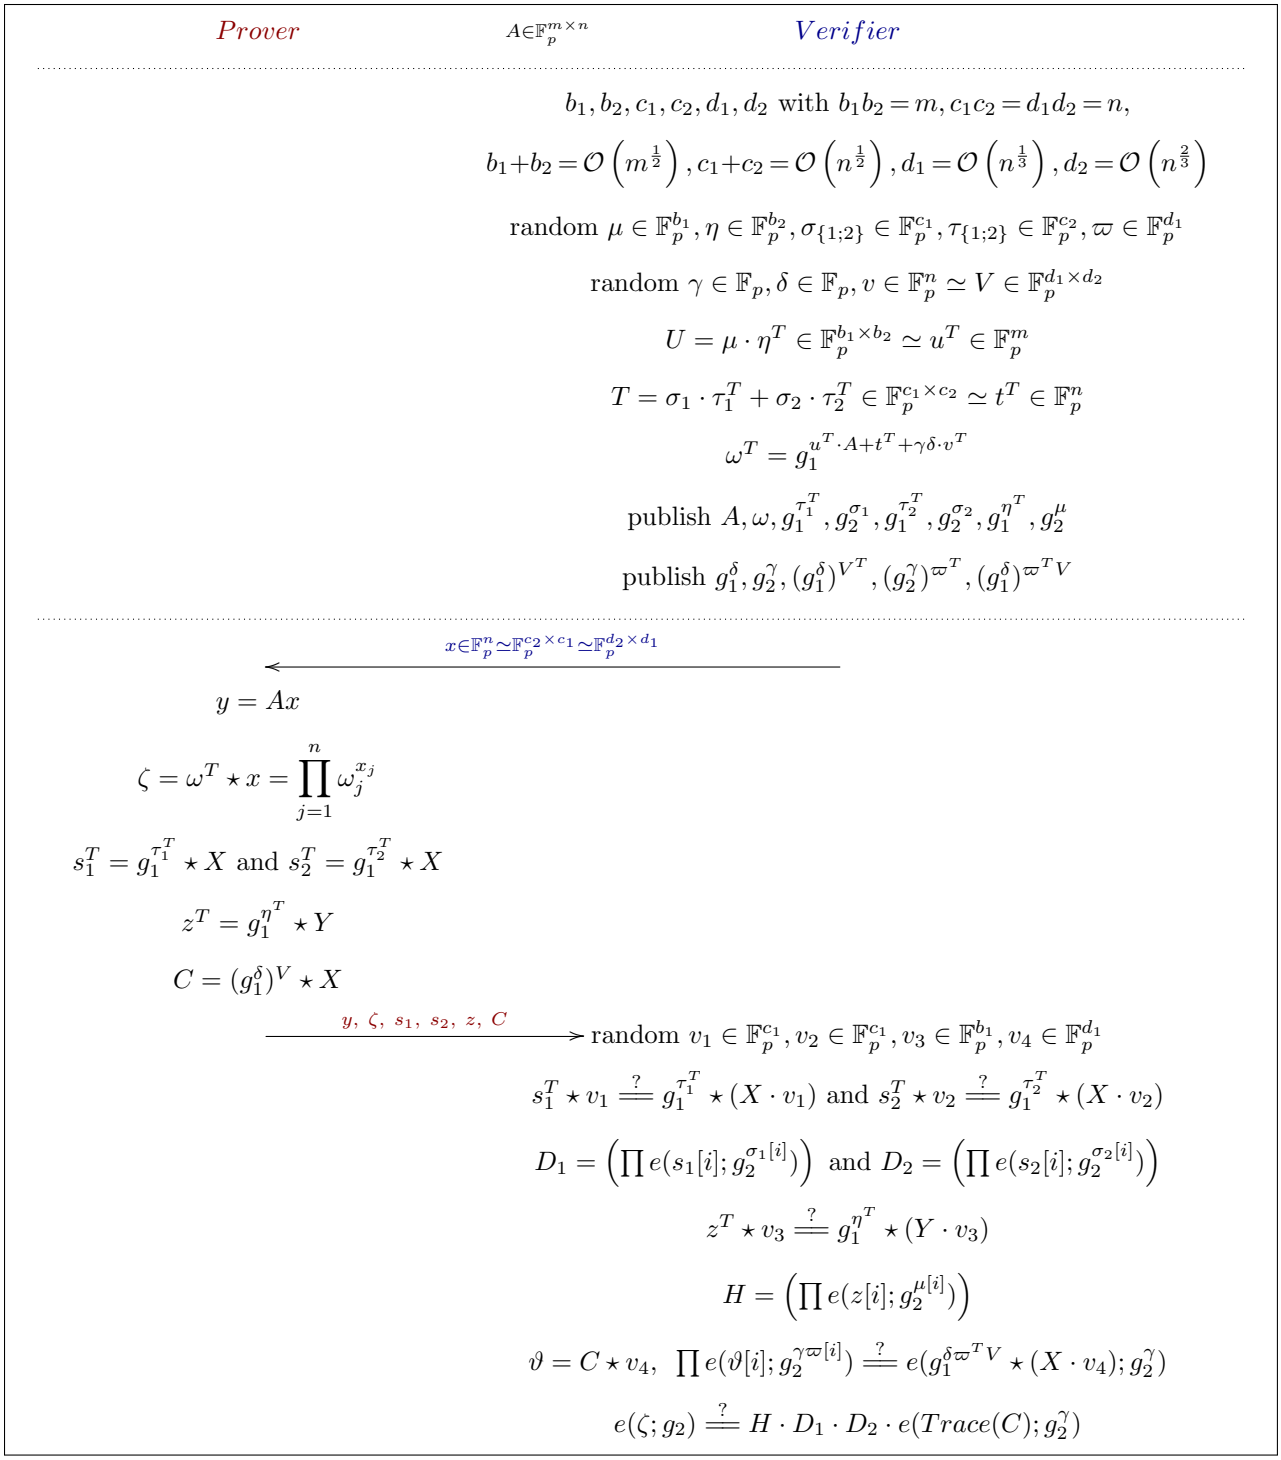

Figure 7: Exchanges in the proven publicly delegatable protocol with negligible cryptographic operations of Figure 4.

and that:

$$D_j = \left(\prod e(s_j[i]; g_2^{\rho_j[i]}) \right) \text{ for } j = 1, 2.$$

Finally, the last check is that these two parts, as well as the last one, which is $e(g_1^{\delta}; g_2^{\gamma})^{v^T \cdot x} = e(g_1; g_2)^{a(b+r)v^T \cdot x} = e(Trace(g_1^{\delta \cdot VX}); g_2^{\gamma})$, are coherent with the definitions of ω and ζ above. Now, with a non-negligible probability ϵ, adversary $\mathcal{A}$ can pass the check for another $y' \neq y$, by providing an adequate ζ'. First, z^T, s_1^T, s_2^T and C must be correct, as they are checked directly and independently by the Freivalds first four checks. Second, we have that $e(\zeta'; g_2) = e(g_1; g_2)^{abu_0^T y' + abt_0^T x} e(g_1^{\delta}; g_2^{\gamma})^{v^T \cdot x}$ and therefore, we must also have $e(\zeta(\zeta')^{-1}; g_2) = e(g_1; g_2)^{abu_0^T (y - y')}$. As u_0 is a secret unknown to adversary $\mathcal{A}$, for a random y' the probability that $u_0^T(y - y') = 0$ is bounded by $1/|\mathbb{G}_1|$ and thus negligible. Thus adversary $\mathcal{B}$ can compute $c \equiv \left(u_0^T(y' - y)\right)^{-1} \mod |\mathbb{G}_1|$ and $(\zeta/\zeta')^c = g_1^{ab}$. Therefore it breaks the co-CDH assumption with non-negligible probability $\epsilon' \simeq \epsilon$. The only other possibility is that adversary $\mathcal{A}$ was able to recover u_0^T. But that would directly implies that it has an advantage in the co-CDH: $g_1^{\eta} = L^{\eta_0}$, $g_2^{\mu} = R^{\mu_0}$. □

In Table 4, we present more timings for the comparison between our protocol and, to our knowledge and according to Table 2, the best previously known from [4]. The associated speed-ups supports our claim of a *Prover efficient* protocol with a gain of two orders of magnitude.

Table 4: Speed-up of our novel Protocol over a 256-bit finite field on a i7 @3.4GHz.

Size	100	200	500	1000	2000	3000	4000
[4]	2.77s	10.93s	67.93s	271.03s	1079.71s	2430.05s	4329.46s
Fig. 7	0.17s	0.34s	0.98s	2.26s	5.37s	9.16s	13.76s
Speed-up	17	32	69	120	201	265	315

Size	5000	6000	7000	8000	9000	10000
[4]	6790.15s	9780.24s	13309.61s	17416.38s	22002.51s	27175.12s
Fig. 7	18.55s	24.03s	29.93s	37.00s	44.00s	51.97s
Speed-up	366	407	445	471	500	523