



HAL
open science

Certificates for triangular equivalence and rank profiles

Jean-Guillaume Dumas, David Lucas, Clément Pernet

► **To cite this version:**

Jean-Guillaume Dumas, David Lucas, Clément Pernet. Certificates for triangular equivalence and rank profiles. ACM International Symposium on Symbolic and Algebraic Computations - ISSAC'17, Jul 2017, Kaiserslautern, Germany. pp.133-140, 10.1145/3087604.3087609 . hal-01466093v2

HAL Id: hal-01466093

<https://hal.science/hal-01466093v2>

Submitted on 20 Oct 2019

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.

Certificates for triangular equivalence and rank profiles*

Jean-Guillaume Dumas

David Lucas

Clément Pernet

Université Grenoble Alpes, Laboratoire Jean Kuntzmann, CNRS, UMR 5224
700 avenue centrale, IMAG - CS 40700, 38058 Grenoble cedex 9, France
{firstname.lastname}@univ-grenoble-alpes.fr

ABSTRACT

In this paper, we give novel certificates for triangular equivalence and rank profiles. These certificates enable to verify the row or column rank profiles or the whole rank profile matrix faster than recomputing them, with a negligible overall overhead. We first provide quadratic time and space non-interactive certificates saving the logarithmic factors of previously known ones. Then we propose interactive certificates for the same problems whose Monte Carlo verification complexity requires a small constant number of matrix-vector multiplications, a linear space, and a linear number of extra field operations. As an application we also give an interactive protocol, certifying the determinant of dense matrices, faster than the best previously known one.

1. INTRODUCTION

Within the setting of verifiable computing, we propose in this paper *interactive certificates* with the taxonomy of [4]. Indeed, we consider a protocol where a *Prover* performs a computation and provides additional data structures or exchanges with a *Verifier* who will use these to check the validity of the result, faster than by just recomputing it. More precisely, in an interactive certificate, the Prover submits a *Commitment*, that is some result of a computation; the Verifier answers by a *Challenge*, usually some uniformly sampled random values; the Prover then answers with a *Response*, that the Verifier can use to convince himself of the validity of the commitment. Several *rounds* of challenge/response might be necessary for the Verifier to be fully convinced.

By Prover (resp. Verifier) *time*, we thus mean bounds on the number of arithmetic operations performed by the Prover (resp. Verifier) during the protocol, while by extra *space*, we mean bounds on the volume of data being exchanged, not counting the size of the input and output of the computation.

Such protocols are said to be *complete* if the probability that a true statement is rejected by the Verifier can be made arbitrarily small; and *sound* if the probability that a false statement is accepted by the Verifier can be made arbitrarily small. In practice it is suffi-

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

cient that those probabilities are < 1 , as the protocols can always be run several times. Some certificates will also be *perfectly complete*, that is a true statement is never rejected by the Verifier. All these certificates can be simulated non-interactively by Fiat-Shamir heuristic [10]: uniformly sampled random values produced by the Verifier are replaced by cryptographic hashes of the input and of previous messages in the protocol. Complexities are preserved.

We do not use generic approaches to verified computation (where protocols check circuits with polylogarithmic depth [12] or use amortized models and homomorphic encryption [2]). Rather, we use dedicated certificates as those designed for dense [11, 14] or sparse [4, 5] exact linear algebra. The obtained certificates are problem-specific, but try to reduce as much as possible the overhead for the Prover, while preserving a fast verification procedure.

We will consider an $m \times n$ matrix A of rank r over a field \mathbb{F} . The *row rank profile* of A is the lexicographically minimal sequence of r indices of independent rows of A . Matrix A has *generic row rank profile* if its row rank profile is $(1, \dots, r)$. The *column rank profile* is defined similarly on the columns of A . Matrix A has *generic rank profile* if its r first leading principal minors are nonzero. The *rank profile matrix* of A , denoted by \mathcal{R}_A is the unique $m \times n$ $\{0, 1\}$ -matrix with r nonzero entries, of which every leading sub-matrix has the same rank as the corresponding sub-matrix of A . It is possible to compute \mathcal{R}_A with a deterministic algorithm in $O(mnr^{\omega-2})$ or with a Monte-Carlo probabilistic algorithm in $(r^\omega + m + n + \mu(A))^{1+o(1)}$ field operations [8], where $\mu(A)$ is the arithmetic cost to multiply A by a vector.

We first propose quadratic, space and verification time, non-interactive practical certificates for the row or column rank profile and for the rank profile matrix that are rank-sensitive. Previously known certificates have additional logarithmic factors to the quadratic complexities: replacing matrix multiplications by quadratic verifications in recursive algorithms yields at least one $\log(n)$ factor [14], graph-based approaches cumulate this and other logarithmic factors, at least from a compression by magical graphs and from a dichotomic search [16].

We then propose two linear space interactive certificates: one certifying that two non-singular matrices are triangular equivalent, i.e. there is a triangular change of basis from one to the other; the other one, certifying that a matrix has a generic rank profile. These certificates are then applied to certify the row or column rank profile, the Q (permutation) and D (diagonal) factors of a LDUP factorization, the determinant and the rank profile matrix. These certificates require, for the Verifier, between 1 and 3 applications of A to a vector and a linear amount of field operations. They are still elimination-based for the Prover, but do not require to communicate the obtained triangular decomposition. For the Determinant, this new certificates require the computation of a PLUQ decompo-

sition for the Prover, linear communication and Verifier time, with no restriction on the field size.

Table 1 compares linear quadratic volumes of communication, as well as sub-cubic (PLUQ, CHARPOLY) or quadratic matrix operations (one matrix-vector multiplication with a dense matrix is denoted fgemv). The results shows first that it is interesting to use linear space certificates even when they have quadratic Verification time. The table also presents a practical constant factor of about 5 between PLUQ and CHARPOLY computations. Computations use the FFLAS-FFPACK library (<http://linbox-team.github.io/fflas-ffpack>) on a single Intel Skylake core @3.4GHz, while we measured some communications between two workstations over an Ethernet Cat. 6, @1Gb/s network cable.

Dimension	$2k$	$10k$	$50k$
PLUQ	0.28s	17.99s	1448.16s
CHARPOLY	1.96s	100.37s	8047.56s
Linear comm.	0.50s	0.50s	0.50s
Quadratic comm.	1.50s	7.50s	222.68s
fgemv	0.0013s	0.038s	1.03s

Table 1: Communication of 64 bit words versus computation modulo 131071

A summary of our contributions is given in Table 3, to be compared with the state of the art in Table 2. We identify the symmetric group with the group of permutation matrices, and write $P \in \mathcal{S}_n$ to denote that a matrix P is a permutation matrix. There, $P[i]$ is the row index of the nonzero element of its i -th column; $\mathcal{D}_n(\mathbb{F})$ is the group of invertible diagonal matrices over the field \mathbb{F} and $[A]_{I,J}^J$ is the (I, J) -minor of the matrix A (the determinant of the submatrix of A with row indices in I and column indices in J). Lastly, $x \stackrel{\$}{\leftarrow} S$ denotes that x is sampled uniformly at random from S .

2. NON INTERACTIVE AND QUADRATIC COMMUNICATION CERTIFICATES

In this section, we propose two certificates, first for the column (resp. row) rank profile, and, second, for the rank profile matrix. While the certificates have a quadratic space communication complexity, they have the advantage of being non-interactive.

2.1 Freivalds' certificate for matrix product

In this paper, we will use Freivalds' certificate [11] to verify matrix multiplication. Considering three matrices A, B and C in $\mathbb{F}^{n \times n}$, such that $A \times B = C$, a straightforward way of verifying the equality would be to perform the multiplication $A \times B$ and to compare its result coefficient by coefficient with C . While this method is deterministic, it has a time complexity of $O(n^\omega)$, which is the matrix multiplication complexity. As such, it cannot be a certificate, as there is no complexity difference between the computation and the verification.

Prover		Verifier
	$A, B \in \mathbb{F}^{n \times n}$	
$C = AB$	\xrightarrow{C}	$v \in \mathbb{F}^{n \times 1}$
		$A(Bv) - Cv \stackrel{?}{=} 0$

Protocol 1: Freivalds' certificate for matrix product

Freivalds' certificate proposes a probabilistic method to check this product in a time complexity of $\mu(A) + \mu(B) + \mu(C)$ using matrix/vector multiplication, as detailed in Figure 1.

2.2 Column rank profile certificate

We now propose a certificate for the column rank profile.

Prover		Verifier
	$A \in \mathbb{F}^{m \times n}$	
A PLUQ decomposition of A	$\xrightarrow{\text{PLUQ}}$	UQ row echelonized?
s.t. UQ is in row echelon form		$A \stackrel{?}{=} PLUQ$, by cert. 1
		Return $Q[1], \dots, Q[r]$

Protocol 2: Column rank profile, non-interactive

LEMMA 1. *Let $A = PLUQ$ be the PLUQ decomposition of an $m \times n$ matrix A of rank r . If UQ is in row echelon form then $(Q[1], \dots, Q[r])$ is the column rank profile of A .*

PROOF. Write $A = P \begin{bmatrix} L_1 \\ L_2 \end{bmatrix} \begin{bmatrix} U_1 & U_2 \end{bmatrix} Q$, where L_1 and U_1 are $r \times r$ lower and upper triangular respectively. If UQ is in echelon form, then $R = \begin{bmatrix} I_r & U_1^{-1}U_2 \\ 0_{(m-r) \times n} \end{bmatrix}$ is in reduced echelon form. Now

$$\begin{bmatrix} U_1^{-1} & \\ & I_{m-r} \end{bmatrix} \begin{bmatrix} L_1 & \\ & L_2 \end{bmatrix} \begin{bmatrix} I_r & \\ & I_{m-r} \end{bmatrix}^{-1} P^T A = \begin{bmatrix} U_1^{-1}UQ \\ 0_{(m-r) \times n} \end{bmatrix} = R$$

is left equivalent to A and is therefore the echelon form of A . Hence the sequence of column positions of the pivots in R , that is $(Q[1], \dots, Q[r])$, is the column rank profile of A . \square

Lemma 1 provides a criterion to verify a column rank profile from a PLUQ decomposition. Such decompositions can be computed in practice by several variants of Gaussian elimination, with no arithmetic overhead, as shown in [13] or [7, § 8]. Hence, we propose the certificate in Protocol 2.

THEOREM 1. *Let $A \in \mathbb{F}^{m \times n}$ with $r = \text{rank}(A)$. Certificate 2, verifying the column rank profile of A is sound, perfectly complete, with a communication bounded by $O(r(m+n))$, a Prover computation bounded by $O(mnr^{\omega-2})$ and a Verifier computation cost bounded by $O(r(m+n)) + \mu(A)$.*

PROOF. If the Prover is honest, then, UQ will be in row echelon form and $A = PLUQ$, thus, by Lemma 1, the Verifier will be able to read the column rank profile of A from Q . If the Prover is dishonest, either $A \neq PLUQ$, which will be caught by the Prover with probability $p \geq 1 - \frac{1}{q}$ using Freivalds' certificate [11] or UQ is not in row echelon form, which will be caught every time by the Verifier.

The Prover sends P, L, U and Q to the Verifier, hence the communication cost of $O(r(m+n))$, as P and Q are permutation matrices and L, U , are respectively $m \times r$ and $r \times n$ matrices, with $r = \text{rank}(A)$. Using algorithms provided in [13], one can compute the expected PLUQ decomposition in $O(mnr^{\omega-2})$. The Verifier has to check if $A = PLUQ$, and if UQ is in row echelon form, which can be done in $O(r(m+n))$. \square

Note that this holds for the row rank profile of A : in that case, the Verifier has to check if PL is in column echelon form.

	Algorithm	Inter.	Prover		Communication	Probabilistic Verifier Time	# \mathbb{F}
			Determ.	Time			
RANK	[14] over [1]	No	No	$\tilde{O}(r^\omega + \mu(A))$	$\tilde{O}(r^2 + m + n)$	$\tilde{O}(r^2 + \mu(A))$	≥ 2
	[4]	Yes	No	$O(n(\mu(A) + n))$	$O(m + n)$	$2\mu(A) + \tilde{O}(m + n)$	$\tilde{O}(\min\{m, n\})$
	[9]	Yes	Yes	$O(mnr^{\omega-2})$	$O(m + r)$	$O(r + \mu(A) + m + n)$	≥ 2
CRP/RRP	[14] over [16]	No	No	$\tilde{O}(r^\omega + m + n + \mu(A))$	$\tilde{O}(r^2 + m + n)$	$\tilde{O}(r^2 + m + n + \mu(A))$	$\tilde{O}(\min\{m, n\})$
	[14] over [13]	No	Yes	$O(mnr^{\omega-2})$	$\tilde{O}(mn)$	$\tilde{O}(mn)$	≥ 2
RPM	[14] over [8]	No	No	$\tilde{O}(r^\omega + m + n + \mu(A))$	$\tilde{O}(r^2 + m + n)$	$\tilde{O}(r^2 + m + n + \mu(A))$	$\tilde{O}(\min\{m, n\})$
	[14] over [6]	No	Yes	$O(mnr^{\omega-2})$	$\tilde{O}(mn)$	$\tilde{O}(mn)$	≥ 2
DET	[11] & PLUQ	No	Yes	$O(n^\omega)$	$O(n^2)$	$O(n^2) + \mu(A)$	≥ 2
	[5] & CHARPOLY	Yes	No	$O(n\mu(A))$ or $O(n^\omega)$	$O(n)$	$\mu(A) + O(n)$	$\geq n^2$

Table 2: State of the art certificates for the rank, the row and column rank profiles, the rank profile matrix and the determinant

	Algorithm	Interactive	Prover		Communication	Probabilistic Verifier Time	# \mathbb{F}
			Deterministic	Time			
CRP/RRP	§ 2.2	No	Yes	$O(mnr^{\omega-2})$	$O(r(m + n))$	$O(r(m + n)) + \mu(A)$	≥ 2
	§ 4.2	Yes	Yes	$O(mnr^{\omega-2})$	$O(m + n)$	$2\mu(A) + O(m + n)$	≥ 2
RPM	§ 2.3	No	Yes	$O(mnr^{\omega-2})$	$O(r(m + n))$	$O(r(m + n)) + \mu(A)$	≥ 2
	§ 4.3	Yes	Yes	$O(mnr^{\omega-2})$	$O(m + n)$	$4\mu(A) + O(m + n)$	≥ 4
DET	§ 4.1 & PLUQ	Yes	Yes	$O(n^\omega)$	$O(n)$	$\mu(A) + O(n)$	≥ 2

Table 3: This paper's contributions

2.3 Rank profile matrix certificate

LEMMA 2. A decomposition $A = PLUQ$ reveals the rank profile matrix, namely $\mathcal{R}_A = P \begin{bmatrix} I_r & \\ & 0 \end{bmatrix} Q$, if and only if $P \begin{bmatrix} L & 0 \\ & 0 \end{bmatrix} P^T$ is lower triangular and $Q^T \begin{bmatrix} U \\ & 0 \end{bmatrix} Q$ is upper triangular.

PROOF. The only if case is proven in [8, Th. 21]. Now suppose that $P \begin{bmatrix} L & 0_{m \times (m-r)} \\ & 0 \end{bmatrix} P^T$ is lower triangular. Then we must also have that $\bar{L} = P \begin{bmatrix} L & 0 \\ & I_{m-r} \end{bmatrix} P^T$ is lower triangular and non-singular. Similarly suppose that $Q^T \begin{bmatrix} U \\ & 0 \end{bmatrix} Q$ is upper triangular so that $\bar{U} = Q^T \begin{bmatrix} U \\ & 0 \\ & 0 & I_{n-r} \end{bmatrix} Q$ is non-singular upper triangular. We have $A = \bar{L}P \begin{bmatrix} I_r & \\ & 0 \end{bmatrix} Q\bar{U}$. Hence the rank of any (i, j) leading submatrix of A is that of the (i, j) leading submatrix of $P \begin{bmatrix} I_r & \\ & 0 \end{bmatrix} Q$, thus proving that $\mathcal{R}_A = P \begin{bmatrix} I_r & \\ & 0 \end{bmatrix} Q$. \square

We use this characterization to verify the computation of the rank profile matrix in the following protocol: Once the Verifier receives P, L, U and Q , he has to:

1. Check that $A = PLUQ$, using Freivalds' certificate [11]
2. Check that L is echelonized by P and U^T by Q^T .
3. If successful, compute the rank profile matrix of A as $\mathcal{R}_A =$

$$P \begin{bmatrix} I_r & \\ & 0_{(m-r) \times (n-r)} \end{bmatrix} Q$$

Prover	Verifier
	$A \in \mathbb{F}^{m \times n}$
a PLUQ decomp. of A revealing \mathcal{R}_A .	\xrightarrow{PLUQ} 1. $A \stackrel{?}{=} PLUQ$ by Protoc. 2.1 2. Is PLP^T lower triangular? 3. Is $Q^T U Q$ upper triangular?

Protocol 3: Rank profile matrix, non-interactive

THEOREM 2. Certificate 3 verifies the rank profile matrix of A , it is sound and perfectly complete, with a communication cost bounded by $O(r(n + m))$, a Prover computation cost bounded by

$O(mnr^{\omega-2})$ and a Verifier computation cost bounded by $O(r(m + n)) + \mu(A)$.

PROOF. If the Prover is honest, then, the provided $PLUQ$ decomposition is indeed a factorization of A , which means Freivalds' certificate will pass. It also means this $PLUQ$ decomposition reveals the rank profile matrix. According to Lemma 2, PLP^T will be lower triangular and $Q^T U Q$ upper triangular. Hence the verification will succeed and $\mathcal{R}_A = P \begin{bmatrix} I_r & \\ & 0_{(m-r) \times (n-r)} \end{bmatrix} Q$ is indeed the rank profile matrix of A . If the Prover is dishonest, either $A \neq PLUQ$, which will be caught with probability $p \geq 1 - \frac{1}{q}$ by Freivalds' certificate or the $PLUQ$ decomposition does not reveal the rank profile matrix of A . In that case, Lemma 2 implies that either $P \begin{bmatrix} L & 0 \\ & 0 \end{bmatrix} P^T$ is not lower triangular or $P \begin{bmatrix} U \\ & 0 \end{bmatrix} Q$ is not upper triangular which will be detected.


The Prover sends P, L, U and Q to the Verifier, hence the communication cost of $O((n + m)r)$. A rank profile matrix revealing $PLUQ$ decomposition can be computed in $O(mnr^{\omega-2})$ operations [6]. The Verifier has to check if $A = PLUQ$, which can be achieved in $O((m + n)r) + \mu(A)$ field operations. \square

3. LINEAR COMMUNICATION CERTIFICATE TOOLBOX

3.1 Triangular one sided equivalence

Two matrices $A, B \in \mathbb{F}^{m \times n}$ are right (resp. left) equivalent if there exist an invertible $n \times n$ matrix T such that $AT = B$ (resp. $TA = B$). If in addition T is a lower triangular matrix, we say that A and B are lower triangular right (resp. left) equivalent. The upper triangular right (resp. left) equivalence is defined similarly. We propose a certification protocol that two matrices are left or right triangular equivalent. Here, A and B are input, known by the Verifier and the Prover. A simple certificate would be the matrix T itself, in which case the Verifier would check the product $AT = B$ using Freivalds' certificate. This certificate is non-interactive and

requires a quadratic amount of communication. In what follows, we present a certificate which allows to verify the one sided triangular equivalence without communicating T , requiring only $2n$ communications. It is essentially a Freivalds' certificate with a more constrained interaction pattern in the way the challenge vector and the response vector are communicated. This pattern imposes a triangular structure in the way the Provers' responses depend on the Verifier challenges which match with the structure of the problem.



Protocol 4: Lower triang. right equivalence of regular matrices

THEOREM 3. Let $A, B \in \mathbb{F}^{m \times n}$, and assume A is regular. Certificate 4 proves that there exists a lower triangular matrix T such that $AT = B$. This certificate is sound, with probability larger than $1 - \frac{1}{|S|}$, perfectly complete, occupies $2n$ communication space, and can be computed in $O(mn^{\omega-1})$ field operations and verified in $\mu(A) + \mu(B)$ field operations.

PROOF. If the Prover is honest, then $AT = B$ and she just computes $y = Tx$, so that $Ay = ATx = Bx$. If the Prover is dishonest, replace the random values x_1, \dots, x_n by algebraically independent variables X_1, \dots, X_n . Since A is regular, there is a unique $n \times n$ matrix T (that is, $T = A^\dagger B$ with A^\dagger the Moore-Penrose inverse of A) such that $AT = B$. For the same reason, there is a unique vector $\hat{Y} = A^\dagger BX$ such that $A\hat{Y} = BX$. The vector \hat{Y} is then formed by n degree-1 polynomials in X_1, \dots, X_n . If T is not lower triangular, let i be the first row such that $T_{i,j} \neq 0$ for some $j > i$, and let j_m be the largest such j . Then \hat{Y}_i has degree 1 in X_{j_m} . Let Y be the vector output by the Prover. At step $2i + 1$, the value for X_{j_m} was still not released, hence Y_i is constant in X_{j_m} . As A is regular, the verification $AY = BX = A\hat{Y}$ is equivalent to $Y - \hat{Y} = 0$. The i -th component in this equation is $Y_i - \hat{Y}_i = 0$, whose left hand-side contains a non zero monomial in X_j . There is therefore a probability lower than $1/|S|$ that the random choice for x_j makes this polynomial vanish.


This certificate requires to transmit x and y , which costs $2n$ in communication. The Verifier has to compute Ay and Bx , whose computational cost is $\mu(A) + \mu(B)$. The Prover has to compute T , this can be done by a PLUQ elimination on A followed by a triangular system solve, both in $O(mn^{\omega-1})$. Then $y = Tx$ requires only $O(n^2)$ operations. \square

Note that the case where T is upper triangular works similarly: the Verifier needs to transmit x in reverse order, starting by x_n .

3.2 Generic rank profile-ness

The problem here is to verify whether a non-singular input matrix $A \in \mathbb{F}^{m \times n}$ has generic rank profile (to test non-singularity, one

can apply beforehand the linear communication certificate in [4, Fig. 2], see also Protocol 8 thereafter). A matrix A has generic rank profile if and only if it has an LU decomposition $A = LU$, with L unit lower triangular and U non-singular upper triangular. The protocol picks random vectors ϕ, ψ, λ and asks the Prover to provide the vectors $z^T = \lambda^T L$, $x = U\phi$, $y = U\psi$ on the fly, while receiving the coefficients of the vectors ϕ, ψ, λ one at a time. These vectors satisfy the fundamental equations $z^T x = \lambda^T A\phi$ and $z^T y = \lambda^T A\psi$ that will be checked by the Verifier.



Protocol 5: Generic rank profile with linear communication

THEOREM 4. Certificate 5 verifying that a non-singular matrix has generic rank profile is sound, with probability larger than $1 - \frac{1}{|S|}$, perfectly complete, communicates $3n$ field elements, and can be computed in $O(n^\omega)$ field operations for the Prover and $\mu(A) + 8n$ field operations for the Verifier.

We will need the following Lemma, used in Dodgson determinant condensation rule.

LEMMA 3 (DESNANOT-JACOBI, OR DODGSON RULE [3]).

$$[A]_{\{1..n\}}^{\{1..n\}} [A]_{\{2..n-1\}}^{\{2..n-1\}} = \begin{bmatrix} [A]_{\{1..n-1\}}^{\{1..n-1\}} & [A]_{\{1..n-1\}}^{\{2..n\}} \\ [A]_{\{2..n\}}^{\{1..n-1\}} & [A]_{\{2..n\}}^{\{2..n\}} \end{bmatrix}.$$

Applying the same permutation, the cyclic shift of order 1 to the left, on the rows and columns of A , yields the following formula with no change of sign:

$$[A]_{\{1..n\}}^{\{1..n\}} [A]_{\{1..n-2\}}^{\{1..n-2\}} = \begin{bmatrix} [A]_{\{1..n-2,n\}}^{\{1..n-2,n\}} & [A]_{\{1..n-1\}}^{\{1..n-1\}} \\ [A]_{\{1..n-2,n\}}^{\{1..n-1\}} & [A]_{\{1..n-1\}}^{\{1..n-1\}} \end{bmatrix}. \quad (1)$$

PROOF OF THEOREM 4. The protocol is perfectly complete: if $A = LU$, then $z^T [x \ y] = \lambda^T LU [\phi \ \psi] = \lambda^T A [\phi \ \psi]$.

Now, for the soundness, replace every ϕ, ψ, λ chosen at random by the Verifier by vectors of algebraically independent variables Φ, Ψ, Λ . Similarly, the responses of the Prover z, x, y are now vectors of algebraically independent variables Z, X, Y . Under the assumption of the success of the Verifier test,

$$\begin{cases} Z^T X = \Lambda^T A \Phi \\ Z^T Y = \Lambda^T A \Psi \end{cases}, \quad (2)$$

and that A is non-singular, we will prove the following induction hypothesis:

$$H_i : \begin{cases} Z_{i..n}^T X_{i..n} &= \frac{1}{d_{i-1}} \sum_{i \leq j, k \leq n} \Lambda_k [A]_{\{1..i-1, j\}}^{\{1..i-1, k\}} \Phi_j \\ Z_{i..n}^T Y_{i..n} &= \frac{1}{d_{i-1}} \sum_{i \leq j, k \leq n} \Lambda_k [A]_{\{1..i-1, j\}}^{\{1..i-1, k\}} \Psi_j \\ d_j \neq 0 \ \forall j < i \end{cases}$$

where $d_i = [A]_{\{1..i\}}^{\{1..i\}}$, $d_0 = 1$.

For $i = 1$, note that $[A]_{\{1..i-1,j\}}^{\{1..i-1,k\}} = A_{k,j}$, hence the right hand-sides of the first two equations of H_1 can be written as:

$$\begin{aligned} \sum_{1 \leq j, k \leq n} \Lambda_k A_{k,j} \Phi_j &= \Lambda^T A \Phi = Z^T X \\ \sum_{1 \leq j, k \leq n} \Lambda_k A_{k,j} \Psi_j &= \Lambda^T A \Psi = Z^T Y \end{aligned}$$

by (2). Finally $d_0 = 1$ is obviously nonzero.

Now suppose H_i is true for some $0 \leq i < n$. Then

$$\begin{cases} Z_i X_i + Z_{i+1..n}^T X_{i+1..n} = \frac{1}{d_{i-1}} \Lambda_i \sum_{j=i}^n [A]_{\{1..i-1,j\}}^{\{1..i\}} \Phi_j \\ \quad + \frac{1}{d_{i-1}} \sum_{j=i}^n \sum_{k=i+1}^n \Lambda_k [A]_{\{1..i-1,j\}}^{\{1..i-1,k\}} \Phi_j \\ Z_i Y_i + Z_{i+1..n}^T Y_{i+1..n} = \frac{1}{d_{i-1}} \Lambda_i \sum_{j=i}^n [A]_{\{1..i-1,j\}}^{\{1..i\}} \Psi_j \\ \quad + \frac{1}{d_{i-1}} \sum_{j=i}^n \sum_{k=i+1}^n \Lambda_k [A]_{\{1..i-1,j\}}^{\{1..i-1,k\}} \Psi_j \end{cases} \quad (3)$$

At the time of choosing the value for Λ_i , all variables are set, except Z_i . Hence for all value assigned to Λ_i , there is a value for Z_i that satisfies the above system of two linear equations in Z_i and Λ_i . Consequently this system is singular and the following two determinants vanish:

$$\begin{vmatrix} d_{i-1} X_i & \sum_{j=i}^n [A]_{\{1..i-1,j\}}^{\{1..i\}} \Phi_j \\ d_{i-1} Y_i & \sum_{j=i}^n [A]_{\{1..i-1,j\}}^{\{1..i\}} \Psi_j \end{vmatrix} = 0 \quad (4)$$

$$\begin{vmatrix} \sum_{j=i}^n [A]_{\{1..i-1,j\}}^{\{1..i\}} \Phi_j & d_{i-1} Z_{i+1..n}^T X_{i+1..n} - F_A(\Lambda, i, \Phi) \\ \sum_{j=i}^n [A]_{\{1..i-1,j\}}^{\{1..i\}} \Psi_j & d_{i-1} Z_{i+1..n}^T Y_{i+1..n} - F_A(\Lambda, i, \Psi) \end{vmatrix} = 0 \quad (5)$$

where $F_A(\Lambda, i, R) = \sum_{j=i}^n \sum_{k=i+1}^n \Lambda_k [A]_{\{1..i-1,j\}}^{\{1..i-1,k\}} R_j$, for $R = \Phi, \Psi$. Actually, Equation (5) is thus of the form $\begin{bmatrix} d_i \Phi_i + b & a \Phi_i + e \\ d_i \Psi_i + c & a \Psi_i + f \end{bmatrix} = U_1 = D^{-1}U$ is unitary. Thus, for any $\mathbb{F}^n \ni \psi = [\psi_1, \tilde{\psi}]^T$: $U\psi = DU_1\psi = D\left(\psi + \begin{bmatrix} \tilde{u}_1 \\ 0 \end{bmatrix} \tilde{\psi}\right)$.

0 where $d_i = [A]_{\{1..i\}}^{\{1..i\}}$, $a = -\sum_{k=i+1}^n \Lambda_k [A]_{\{1..i-1,j\}}^{\{1..i-1,k\}}$ and b, c, e, f are constants with respect to the variables Φ_i, Ψ_i . If $d_i = 0$, then, at least one $[A]_{\{1..i-1,j\}}^{\{1..i\}}$ for $j > i$ must be nonzero, otherwise A would be singular. Similarly, at least one $[A]_{\{1..i-1,k\}}^{\{1..i-1,j\}}$ for $k > i$ is nonzero, hence a is a nonzero polynomial in $\Lambda_{i+1}, \dots, \Lambda_n$ and b, c are nonzero polynomials in Φ_j, Ψ_j for $j > i$, but constant in Φ_i and Ψ_i . This is a contradiction, as the first column of the determinant, $\begin{bmatrix} b \\ c \end{bmatrix}$ can not be colinear with the second one. Hence $d_i \neq 0$.

Therefore $\begin{bmatrix} e \\ f \end{bmatrix} = \frac{a}{d_i} \begin{bmatrix} b \\ c \end{bmatrix}$ which is

$$\begin{cases} d_{i-1} Z_{i+1..n}^T X_{i+1..n} = \frac{1}{d_i} \sum_{j,k=i+1}^n \Lambda_k \left(d_i [A]_{\{1..i-1,j\}}^{\{1..i-1,k\}} - [A]_{\{1..i\}}^{\{1..i-1,k\}} [A]_{\{1..i-1,j\}}^{\{1..i\}} \right) \Phi_j \\ d_{i-1} Z_{i+1..n}^T Y_{i+1..n} = \frac{1}{d_i} \sum_{j,k=i+1}^n \Lambda_k \left(d_i [A]_{\{1..i-1,j\}}^{\{1..i-1,k\}} - [A]_{\{1..i\}}^{\{1..i-1,k\}} [A]_{\{1..i-1,j\}}^{\{1..i\}} \right) \Psi_j \end{cases}$$

Applying variant (1) of Lemma 3 to $[A]_{\{1..i,j\}}^{\{1..i,k\}}$, yields

$$\begin{cases} d_{i-1} Z_{i+1..n}^T X_{i+1..n} = \frac{1}{d_i} \sum_{j,k=i+1}^n \Lambda_k d_{i-1} [A]_{\{1..i,j\}}^{\{1..i,k\}} \Phi_j \\ d_{i-1} Z_{i+1..n}^T Y_{i+1..n} = \frac{1}{d_i} \sum_{j,k=i+1}^n \Lambda_k d_{i-1} [A]_{\{1..i,j\}}^{\{1..i,k\}} \Psi_j \end{cases}$$

and H_{i+1} is verified.

We have proven that if H_i is true, then either H_{i+1} is also true or the system (3) has a single solution and the Verifier randomly

chose precisely that λ_i . Therefore, suppose that A has not generic rank profile, it means that some $d_j = 0$ and H_j is false. But the Verifier checks that H_1 is true. If this is the case, then at least once, did the Verifier choose the value expected by the dishonest Prover. This happens with probability lower than $1/|S|$.

Finally, for the complexity, the Prover needs one Gaussian elimination to compute LU in time $O(n^\omega)$, then her extra work is just three triangular solve in $O(n^2)$. The extra communication is three vectors, ϕ, ψ, λ , and the Verifier's work is four dot-products and one multiplication by the initial matrix A . \square

3.3 LDUP decomposition

With Protocol 5, when the matrix A does not have generic rank profile, any attempt to prove that it has generic rank profile will be detected w.h.p. (soundness). However when it is the case, the verification will accept many possible vectors x, y, z : any scaling of z_i by α_i and x_i, y_i by $1/\alpha_i$ would be equally accepted for any non zero constants α_i . This slack correspond to our lack of specification of the diagonals' shape in the used LU decomposition. Indeed, for any diagonal matrix with non zero elements, $LD \times D^{-1}U$ is also a valid LU decomposition and yields x, y and z scaled as above. Specifying these diagonals is not necessary to prove generic rank profileness, so we left it as is for this task.

However, for the determinant or the rank profile matrix certificates of Sections 4.1 and 4.3, we will need to ensure that this scaling is independent from the choice of the vectors ϕ, ψ, λ . Hence we propose an updated protocol, where L has to be unit diagonal, and the prover has to first commit the main diagonal D of U .

For an $n \times n$ triangular matrix T , its strictly triangular part is denoted $\tilde{T} \in \mathbb{F}^{(n-1) \times (n-1)}$: for instance if T is upper triangular, then $\tilde{t}_{i,j} = t_{i,j+1}$ for $j \geq i$ and 0 otherwise.

For U an invertible upper triangular matrix we have for its diagonal (d_1, \dots, d_n) and the associated diagonal matrix D , that $U_1 = D^{-1}U$ is unitary. Thus, for any $\mathbb{F}^n \ni \psi = [\psi_1, \tilde{\psi}]^T$: $U\psi = DU_1\psi = D\left(\psi + \begin{bmatrix} \tilde{u}_1 \\ 0 \end{bmatrix} \tilde{\psi}\right)$.

So the idea is that the Prover will commit D beforehand, and that within a generic rank profile certificate, the Verifier will only communicate $\tilde{\phi}, \tilde{\psi}$ and $\tilde{\lambda}$ to obtain $\tilde{z} = \tilde{\lambda}^T \tilde{L}, \tilde{x} = \tilde{U}_1 \tilde{\phi}$ and $\tilde{y} = \tilde{U}_1 \tilde{\psi}$. Then the Verifier will compute by herself the complete vectors. This ensures that L is unitary and that $U = DU_1$ with U_1 unitary.

Finally, if an invertible matrix does not have generic rank profile, we note that it is also possible to incorporate the permutations, by committing them in the beginning and reapplying them to the matrix during the checks. The full certificate is given in Figure 6.

THEOREM 5. *The Protocol of Figure 6, committing a permutation matrix P and a diagonal matrix D for an invertible matrix A , such that there exists unitary triangular matrices L and U with $A = LDUP$, is sound, with probability larger than $1 - \frac{1}{|S|}$, and perfectly complete. For an $n \times n$ matrix, it requires less than $8n$ extra communications and the computational cost for the Verifier is bounded by $\mu(A) + 12n + o(n)$.*

PROOF. If the Prover is honest, then $A = LUP = LDU_1P$, so that for any choice of λ and ψ we have: $\lambda^T AP^T \psi = \lambda^T LDU_1 \psi$, that is $\begin{bmatrix} \tilde{\lambda}^T & \lambda_n \end{bmatrix} \left(I + \begin{bmatrix} 0 & 0 \\ \tilde{L} & 0 \end{bmatrix} \right) D \left(\begin{bmatrix} 0 & \tilde{\psi} \\ 0 & 0 \end{bmatrix} + I \right) \begin{bmatrix} \tilde{\psi} \\ \psi_n \end{bmatrix} = z^T y$ and the same is true for λ and ϕ , so that the protocol is perfectly complete.

Now, the last part of the Protocol of Figure 6 is actually a verification that AP^T has generic rank profile, in other words that there exists lower and upper triangular matrices L^* and U^* such that $AP^T = L^*U^*$. This verification is sound by Theorem 4. Next,

Prover	Verifier
$A \in \mathbb{F}^{n \times n}$ non-singular	
$A = LDUP$	$P \stackrel{?}{\in} S_n, D \stackrel{?}{\in} \mathcal{D}_n(\mathbb{F})$
	Choose $S \subset \mathbb{F}$
	for i from n downto 2:
$[\tilde{x} \ \tilde{y}] \leftarrow \tilde{U}_1 [\tilde{\phi} \ \tilde{\psi}]$	$\phi_i, \psi_i \stackrel{\$}{\leftarrow} S^2$
	\vdots
	$\phi_i, \psi_i \stackrel{\$}{\leftarrow} S^2$
	\vdots
$\tilde{z} \leftarrow \tilde{\lambda}^T \tilde{L}$	$\lambda_i \stackrel{\$}{\leftarrow} S$
	\vdots
	$\phi_1, \psi_1, \lambda_1 \stackrel{\$}{\leftarrow} S^3$
	$[x \ y] \leftarrow [\phi \ \psi] + \begin{bmatrix} \tilde{x} & \tilde{y} \\ 0 & 0 \end{bmatrix}$
	$z^T \leftarrow (\lambda^T + [z^T \ 0])$
	$z^T D [x \ y] \stackrel{?}{=} (\lambda^T A) P^T [\phi \ \psi]$

Protocol 6: LDUP decomposition (linear communication)

the multiplication by the diagonal D is performed by the Verifier, so he is actually convinced that there exists lower and upper triangular matrices L^* and U_1^* such that $AP^T = L^*DU_1^*$. Finally, the construction of the vectors with the form $a + \begin{bmatrix} b \\ 0 \end{bmatrix}$ is also done by the Verifier, so he in fact has a guaranty that L^* and U_1^* are unitary.

Overall, if the Prover is dishonest, the Verifier will catch him with the probability of Theorem 4.

Finally, for the complexity bounds, the extra communications are: one permutation matrix P , a diagonal matrix D and 6 vectors $\tilde{\lambda}, \tilde{\phi}, \tilde{\psi}$ and \tilde{z}, \tilde{x} and \tilde{y} . That is n non-negative integers lower than n and $6(n-1) + n$ field elements. The arithmetic computations of the Verifier are one multiplication by a diagonal matrix, 3 vector sums, 4 dot-products and one matrix-vector multiplication by A (for $(\lambda^T A)$), that is $n + 3(n-1) + 4(2n-1)$. \square

We, furthermore, have some guaranties on the actual values of x, y, z :

PROPOSITION 1. *Let S be a finite subset of \mathbb{F} in Protocol 6, if $[x \ y] \neq U_1 [\phi \ \psi]$ then the verification will pass with probability at most $\frac{2}{|S|}$.*

PROOF. Equation (4) implies that, if the verification check passes, with (z, x, y) , then the vector $[x_i \ y_i]^T$ must be co-linear with the right column of this determinant, that can be written in the form $[d_i \phi_i + b \ d_i \psi_i + c]^T$ with $d_i \neq 0$ and b and c depending only on $\phi_k, \psi_k, x_k, y_k, \lambda_k, z_k$ with $k > i$. Hence, any value \tilde{x}_i, \tilde{y}_i , supplied by the Prover, must satisfy

$$\begin{vmatrix} \phi_i + \tilde{x}_i & d_i \phi_i + b \\ \psi_i + \tilde{y}_i & d_i \psi_i + c \end{vmatrix} = 0, \quad (6)$$

when ϕ_i and ψ_i are still unknown. This condition is ensured for any ϕ_i and ψ_i if and only if $[\tilde{x}_i \ \tilde{y}_i] = \frac{1}{d_i} [b \ c]$. If the Prover is dishonest and if $[x \ y] \neq U_1 [\phi \ \psi]$ then at least one couple $(\tilde{x}_i, \tilde{y}_i)$ is incorrect. Then, either the Verifier has chosen a couple of values (ϕ_i, ψ_i) making the degree 1 determinant (6) vanish, this happens with probability at most $1/|S|$, or System (3) has a unique solution (z_i, λ_i) . But if the latter is true and the final check succeeds then, as for Theorem 4, at least once the Prover chose to have $1/|S|$ chances that the Verifier picked the unique possibility for λ_j , $i \geq j \geq 1$. Overall, the Verification thus fails with probability at most $1 - \frac{2}{|S|}$. \square

REMARK 1. *Correctness of the vector z can also be ensured with the same probability: for the singular System (3), with respect to the unknowns Λ_i and Z_i , to have rank at least one, it is sufficient that one of X_i or Y_i is non zero. The Verifier, knowing \tilde{x}_i , can ensure this by restricting the set of choices for $\phi_i \in S \setminus \{-\tilde{x}_i\}$. Thus if x_i and y_i are correct, the Prover will have to provide a correct associated z_i or increase the probability of being caught.*

4. LINEAR COMMUNICATION INTERACTIVE CERTIFICATES

In this section, we give linear space communication certificates for the determinant, the column/row rank profile of a matrix, and for the rank profile matrix.

4.1 Linear communication certificate for the determinant

Existing certificates for the determinant are either optimal for the Prover in the dense case, using the strategy of [14, Theorem 5] over a PLUQ decomposition, but quadratic in communication; or linear in communication, using [5, Theorem 14], but using a reduction to the characteristic polynomial. In the sparse case the determinant and the characteristic polynomial both reduce to the same minimal polynomial computations and therefore the latter certificate is currently optimal for the Prover. Now in the dense case, while the determinant and characteristic polynomial both reduce to matrix multiplication, the determinant, via a single PLUQ decomposition is more efficient in practice [15]. Therefore, we propose here an alternative in the dense case: use only one PLUQ decomposition for the Prover while keeping linear extra communications and $O(n) + \mu(A)$ operations for the Verifier. The idea is to extract the information of a LDUP decomposition without communicating it: one uses Protocol 6 for $A = LDUP$ with L and U unitary, but kept on the Prover side, and then the Verifier only has to compute $Det(A) = Det(D)Det(P)$, with $n-1$ additional field operations.

COROLLARY 1. *For an $n \times n$ matrix, there exists a sound and perfectly complete protocol for the determinant over a field using less than $8n$ extra communications and with computational cost for the Verifier bounded by $\mu(A) + 13n + o(n)$.*

As a comparison, the protocol of [5, Theorem 14] reduces to CHARPOLY instead of PLUQ for the Prover, requires $5n$ extra communications and $\mu(A) + 13n + o(n)$ operations for the Verifier as well. Also the new protocol requires $3n$ random field elements for any field, where that of [5, Theorem 14] requires 3 random elements but a field larger than n^2 .

For instance, using the routines shown in Table 1, the determinant of an $50k \times 50k$ random dense matrix can be computed in about 24 minutes, where with the certificate of Figure 6, the overhead of the Prover is less than 5s and the Verifier time is about 1s.

4.2 Column or row rank profile certificate

In Figure 7 and 8, we first recall the two linear time and space certificates for an upper and a lower bound to the rank that constitute a rank certificate. We present here the variant sketched in [9, § 2] of the certificates of [4]. An upper bound r on the rank is certified by the capacity for the Prover to generate any vector sampled from the image of A by a linear combination of r column of A . A lower bound r is certified by the capacity for the Prover to recover the unique coefficients of a linear combination of r linearly independent columns of A .

Prover	Verifier
$A \in \mathbb{F}^{m \times n}$	
r s.t. $\text{rank}(A) \leq r$	\xrightarrow{r}
	Choose $S \subset \mathbb{F}$
$A\gamma = w$	$\xrightarrow{\gamma}$
	$v \xleftarrow{\$} S^n, w = Av$ $ \gamma _H \stackrel{?}{=} r$ $A\gamma \stackrel{?}{=} w$

Protocol 7: Upper bound on the rank of a matrix

THEOREM 6. Let $A \in \mathbb{F}^{m \times n}$, and let S be a finite subset of \mathbb{F} . The interactive certificate 7 of an upper bound for the rank of A is sound, with probability larger than $1 - \frac{1}{|S|}$, perfectly complete, occupies $2n$ communication space, can be computed in $LINSYS(r)$ and verified in $2\mu(A) + n$ time.

Prover	Verifier
$A \in \mathbb{F}^{m \times n}$	
c_1, \dots, c_r indep. cols of A	$\xrightarrow{c_1, \dots, c_r}$
	Choose $S \subset \mathbb{F}$
Solve $A\beta = v$	$\xrightarrow{\beta}$
	$\alpha = \begin{cases} \alpha_{c_j} \xleftarrow{\$} S^* \\ 0 \text{ otherwise} \end{cases}$ $v = A\alpha$ $\beta \stackrel{?}{=} \alpha$

Protocol 8: Lower bound on the rank of a matrix

THEOREM 7. Let $A \in \mathbb{F}^{m \times n}$, and let S be a finite subset of \mathbb{F} . The interactive certificate 8 of a lower bound for the rank of A is sound, with probability larger than $1 - \frac{1}{|S|}$, perfectly complete and occupies $n + 2r$ communication space, can be computed in $LINSYS(r)$ and verified in $\mu(A) + r$ operations.

We now consider a column rank profile certificate: the Prover is given a matrix A , and answers the column rank profile of A , $\mathcal{J} = (c_1, \dots, c_r)$. In order to certify this column rank profile, we need to certify two properties:

1. the columns given by \mathcal{J} are linearly independent;
2. the columns given by \mathcal{J} form the lexicographically smallest set of independent columns of A .

Property 1 is verified by Certificate 8, as it checks whether a set of columns are indeed linearly independent. Property 2 could be certified by successive applications of Certificate 7: at step i , checking that the rank of $A_{*,(0, \dots, c_{i-1})}$ is at most $i - 1$ would certify that there is no column located between c_{i-1} and c_i in A which increases the rank of A . Hence, it would prove the minimality of \mathcal{J} . However, this method requires $O(nr)$ communication space.

Instead, we reduce these communication by seeding all challenges from a single n dimensional vector, and by compressing the responses with a random projection. The right triangular equivalence certificate plays here a central role, ensuring the lexicographic minimality of \mathcal{S} . More precisely, the Verifier chooses a vector $v \in \mathbb{F}^n$ uniformly at random and sends it to the Prover. Then, for each index $c_k \in \mathcal{S}$ the Prover computes the linear combination of the first $c_k - 1$ columns of A using the first $c_k - 1$ coefficients of v and has to prove that it can be generated from the $k - 1$ columns c_1, \dots, c_{k-1} . This means, find a vector $\gamma^{(k)}$ solution

to the system:

$$[A_{*,c_1} \ A_{*,c_2} \ \dots \ A_{*,c_{k-1}}] \gamma^{(k)} = A \begin{bmatrix} v_1 \\ \vdots \\ v_{c_k-1} \\ 0 \\ \vdots \end{bmatrix}.$$

Equivalently, find a strictly upper triangular matrix Γ such that:

$$[A_{*,c_1} \ A_{*,c_2} \ \dots \ A_{*,c_{r-1}}] \Gamma = A \underbrace{\begin{bmatrix} v_1 & v_1 & \dots & \dots & v_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ v_{c_1-1} & \vdots & \vdots & \vdots & \vdots \\ 0 & v_{c_2-1} & \vdots & \vdots & \vdots \\ 0 & 0 & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & v_{c_{r-1}-1} & \vdots \\ 0 & 0 & 0 & 0 & v_n \end{bmatrix}}_V.$$

Note that $V = \text{Diag}(v_1, \dots, v_n)W$ where $W = [\mathbb{1}_{i < c_{j+1}}]_{i,j}$ (with $c_{r+1} = n + 1$ by convention) In order to avoid having to transmit the whole $r \times r$ upper triangular matrix Γ , the Verifier only checks a random projection x of it, using the triangular equivalence Certificate 4. We then propose the certificate in Figure 9.

Prover	Verifier
$A \in \mathbb{F}^{m \times n}$	
(c_1, \dots, c_r) CRP of A	$\xrightarrow{(c_1, \dots, c_r)}$
	rank $A \geq r$ by Cert. 8
	Choose $S \subset \mathbb{F}$
Solve $A\beta = v$	$\xrightarrow{\beta}$
	$v \xleftarrow{\$} S^n$ $W = [\mathbb{1}_{i < c_{j+1}}]$ $D \leftarrow \text{Diag}(v_i)$
$y = \Gamma x$	$\xrightarrow{x \text{ (Cert. 4)} y}$
	$x \xleftarrow{\$} S^r$ $z \leftarrow D(Wx)$ $z_{c_j} \leftarrow z_{c_j} - y_j, j = 1..r$ $Az \stackrel{?}{=} 0$

Protocol 9: Certificate for the column rank profile

THEOREM 8. For $A \in \mathbb{F}^{m \times n}$ and $S \subset \mathbb{F}$, certificate 9 is sound, with probability larger than $1 - \frac{1}{|S|}$, perfectly complete, with a Prover computational cost bounded by $O(mnr^{\omega-2})$, a communication space complexity bounded by $2n + 4r$ and a Verifier cost bounded by $2\mu(A) + n + 3r$.

PROOF. If the Prover is honest, the protocol corresponds first to an application of Theorem 7 to certify that \mathcal{J} is a set of independent columns. This certificate is perfectly complete. Second the protocol also uses challenges from Certificate 7, which is perfectly complete, together with Certificate 4, which is perfectly complete as well. The latter certificate is used on $A_{*,\mathcal{J}}$, a regular submatrix, as \mathcal{J} is a set of independent columns of A . The final check then corresponds to $A(D(Wx)) - A_{*,\{c_1, \dots, c_r\}} y \stackrel{?}{=} 0$ and, overall, Certificate 9 is perfectly complete.

If the Prover is dishonest, then either the set of columns in \mathcal{J} are not linearly independent, which will be caught by the Verifier with probability at least $1 - \frac{1}{|S|}$, from Theorem 7, or \mathcal{J} is not lexicographically minimal, or the rank of A is not r . If the rank is wrong, it will not be possible for the prover to find a suitable Γ . This will be caught by the verifier with probability $1 - \frac{1}{|S|}$, from Theorem 3. Finally, if \mathcal{J} is not lexicographically minimal, there exists at least one column $c_k \notin \mathcal{J}, c_i < c_k < c_{i+1}$ for some fixed


i such that $\{c_1, \dots, c_i\} \cup \{c_k\}$ form a set of linearly independent columns of A . This means that $\text{rank}(A_{*,1,\dots,c_{i+1}-1}) = i + 1$, whereas it was expected to be i . Thus, the prover cannot reconstruct a suitable triangular Γ and this will be detected by the verifier also with probability $1 - \frac{1}{|S|}$, as shown in Theorem 3).

The Prover's time complexity is that of computing a $PLUQ$ decomposition of A . The transmission of v , x and y yields a communication space of $n + 2r$. Finally, in addition to Protocol 8, the Verifier computes Wx as a prefix sum with $r - 1$ additions, multiplies it by D , then subtracts y_i at the r correct positions and finally multiplies by A for a total cost bounded by $2\mu(A) + n + 3r - 1$. \square

4.3 Rank profile matrix certificate

We propose an interactive certificate for the rank profile matrix based on [8, Algorithm 4]: first computing the row and column support of the rank profile matrix, using Certificate 9 twice for the row and column rank profiles, then computing the rank profile matrix of the invertible submatrix of A lying on this grid.

In the following we then only focus on a certificate for the rank profile matrix of an invertible matrix. It relies on an LUP decomposition that reveals the rank profile matrix. From Theorem 2, this is the case if and only if $P^T U P$ is upper triangular. Protocol 10 thus gives an interactive certificate that combines Certificate 6 for a LDUP decomposition with a certificate that $P^T U P$ is upper triangular. The latter is achieved by Certificate 4 showing that P^T and $P^T U$ are left upper triangular equivalent, but since U is unknown to the Verifier, the verification is done on a random right projection with the vector ϕ used in Certificate 6.



Protocol 10: Rank profile matrix of an invertible matrix

THEOREM 9. *Protocol 10 is sound, with probability greater than $1 - \frac{2}{|S|}$, and perfectly complete. The Prover cost is $O(n^\omega)$ field operations, the communication space is bounded by $10n$ and the Verifier cost is bounded by $\mu(A) + 16n$.*

PROOF. If the Prover is dishonest and $\bar{U} = P^T U P$ is not upper triangular, then let (i, j) be the lexicographically minimal coordinates such that $i > j$ and $\bar{U}_{i,j} \neq 0$. Now either $\begin{bmatrix} x & y \end{bmatrix} \neq U \begin{bmatrix} \phi & \psi \end{bmatrix}$, and the verification will then fail to detect it with probability less than $\frac{2}{|S|}$, from Proposition 1. Or one can write $e^T P^T x - f^T P^T \phi = (e^T \bar{U} - f^T) P \phi = 0$. If

$$e^T P^T U P - f^T = 0. \quad (7)$$

is not satisfied, then a random ϕ will fail to detect it with probability less than $\frac{1}{|S|}$, since e, \bar{U} and f are set before choosing for ϕ . At the time of committing f_j , the value of e_i is still unknown, hence f_j is constant in the symbolic variable E_i . Thus the j -th coordinate in (7) is a nonzero polynomial in E_j and therefore vanishes with probability $1/|S|$ when sampling the values of e uniformly. Hence, overall if $P^T U P$ is not upper triangular, the verification will fail to detect it with probability at most $2/|S|$. \square

Finally, the rank profile matrix of any matrix, even a singular one, can thus be verified with two applications of Certificate 9 (one for the row rank profile and one for the column rank profile, themselves calling Certificate 8 only once), followed by Certificate 10 on the $r \times r$ selection of lexicographically minimal independent rows and columns. Overall this is $4\mu(A) + 2n + 21r$ operations for the Verifier, and $3n + 16r$ communications.

5. REFERENCES

- [1] H. Y. Cheung, T. C. Kwok, and L. C. Lau. Fast Matrix Rank Algorithms and Applications. *Journal of the ACM*, 60(5):31:1–31:25, Oct. 2013. doi:10.1145/2528404.
- [2] C. Costello, C. Fournet, J. Howell, M. Kohlweiss, B. Kreuter, M. Naehrig, B. Parno, and S. Zahur. Geppetto: Versatile verifiable computation. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 253–270, 2015. doi:10.1109/SP.2015.23.
- [3] C. L. Dodgson. Condensation of Determinants, Being a New and Brief Method for Computing their Arithmetical Values. *Proceedings of the Royal Society of London*, 15:150–155, 1866. URL: <http://www.jstor.org/stable/112607>.
- [4] J.-G. Dumas and E. Kaltofen. Essentially optimal interactive certificates in linear algebra. In K. Nabeshima, editor, *ISSAC'2014*, pages 146–153. ACM Press, New York, July 2014. doi:10.1145/2608628.2608644.
- [5] J.-G. Dumas, E. Kaltofen, E. Thomé, and G. Villard. Linear time interactive certificates for the minimal polynomial and the determinant of a sparse matrix. In X.-S. Gao, editor, *ISSAC'2016*, pages 199–206. ACM Press, New York, July 2016. doi:10.1145/2930889.2930908.
- [6] J.-G. Dumas, C. Pernet, and Z. Sultan. Simultaneous computation of the row and column rank profiles. In M. Kauers, editor, *ISSAC'2013*, pages 181–188. ACM Press, New York, June 2013. doi:10.1145/2465506.2465517.
- [7] J.-G. Dumas, C. Pernet, and Z. Sultan. Computing the rank profile matrix. In Yokoyama [17], pages 149–156. doi:10.1145/2755996.2756682.
- [8] J.-G. Dumas, C. Pernet, and Z. Sultan. Fast computation of the rank profile matrix and the generalized Bruhat decomposition. *Journal of Symbolic Computation*, 2016. in press. doi:10.1016/j.jsc.2016.11.011.
- [9] W. Eberly. A new interactive certificate for matrix rank. Technical Report 2015-1078-11, University of Calgary, June 2015. URL: <http://prism.ucalgary.ca/bitstream/1880/50543/1/2015-1078-11.pdf>.
- [10] A. Fiat and A. 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 LNCS, pages 186–194. Springer-Verlag, 1987, 11–15 Aug. 1986. URL: <http://www.cs.rit.edu/~jkk8346/FiatShamir.pdf>.
- [11] R. Freivalds. Fast probabilistic algorithms. *Mathematical Foundations of Computer Science, LNCS*, 74:57–69, Sept. 1979. doi:10.1007/3-540-09526-8_5.
- [12] S. Goldwasser, Y. T. Kalai, and G. N. Rothblum. Delegating computation: interactive proofs for muggles. In C. Dwork, editor, *STOC'2008*, pages 113–122. ACM Press, May 2008. doi:10.1145/1374376.1374396.
- [13] C.-P. Jeannerod, C. Pernet, and A. Storjohann. Rank-profile revealing gaussian elimination and the CUP matrix decomposition. *Journal of Symbolic Computation*, 56:pages 46–68, 2013. doi:10.1016/j.jsc.2013.04.004.

- [14] E. L. Kaltofen, M. Nehring, and B. D. Saunders. Quadratic-time certificates in linear algebra. In A. Leykin, editor, *ISSAC'2011*, pages 171–176. ACM Press, New York, June 2011.
[doi:10.1145/1993886.1993915](https://doi.org/10.1145/1993886.1993915).
- [15] C. Pernet and A. Storjohann. Faster algorithms for the characteristic polynomial. In C. W. Brown, editor, *ISSAC'2007*, pages 307–314. ACM Press, New York, July 29 – August 1 2007.
[doi:10.1145/1277548.1277590](https://doi.org/10.1145/1277548.1277590).
- [16] A. Storjohann and S. Yang. A Relaxed Algorithm for Online Matrix Inversion. In Yokoyama [17], pages 339–346.
[doi:10.1145/2755996.2756672](https://doi.org/10.1145/2755996.2756672).
- [17] K. Yokoyama, editor. *ISSAC'2015*. ACM Press, New York, July 2015.