

Elimination-based certificates for triangular equivalence and rank profiles¹

Jean-Guillaume Dumas

*Université Grenoble Alpes, Laboratoire Jean Kuntzmann, CNRS, UMR 5224,
700 avenue centrale, IMAG - CS 40700, 38058 Grenoble cedex 9, France*

Erich Kaltofen

*North Carolina State University, Department of Mathematics,
Raleigh, North Carolina 27695-8205, USA*

David Lucas

*Université Grenoble Alpes, Laboratoire Jean Kuntzmann, CNRS, UMR 5224,
700 avenue centrale, IMAG - CS 40700, 38058 Grenoble cedex 9, France*

Clément Pernet

*Université Grenoble Alpes, Laboratoire Jean Kuntzmann, CNRS, UMR 5224,
700 avenue centrale, IMAG - CS 40700, 38058 Grenoble cedex 9, France*

Abstract

In this paper, we give novel certificates for triangular equivalence and rank profiles. These certificates enable somebody 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, with a linear number of interactions. As an application we also give an interactive protocol, certifying the determinant or the signature of dense matrices, faster for the Prover than the best previously known one. Finally we give linear space and constant round certificates for the row or column rank profiles.

¹This work is partly funded by the [OpenDreamKit Horizon 2020 European Research Infrastructures](#) project (#676541) and the French National Research Agency program ([ANR-15-IDEX-02](#)) (Dumas, Lucas, Pernet) and by NSF (USA), grants CCF-1421128 and CCF-1717100 (Kaltofen).

Email addresses: Jean-Guillaume.Dumas@univ-grenoble-alpes.fr (Jean-Guillaume Dumas), kaltofen@math.ncsu.edu (Erich Kaltofen), David.Lucas@univ-grenoble-alpes.fr (David Lucas), Clement.Pernet@univ-grenoble-alpes.fr (Clément Pernet)

URL: <http://www-ljk.imag.fr/membres/Jean-Guillaume.Dumas/> (Jean-Guillaume Dumas), <http://www.kaltofen.us> (Erich Kaltofen), <http://www-ljk.imag.fr/membres/David.Lucas/> (David Lucas), <http://www-ljk.imag.fr/membres/Clement.Pernet/> (Clément Pernet)

Contents

1	Introduction	3
2	Non interactive and quadratic communication certificates	5
2.1	Freivalds' certificate for matrix product	5
2.2	Column rank profile certificate	5
2.3	Rank profile matrix certificate	6
3	Linear communication certificate toolbox	7
3.1	Triangular one sided equivalence	7
3.2	Generic rank profile-ness	9
3.3	LDUP decomposition	12
4	Linear communication interactive certificates	14
4.1	Linear communication certificate for the determinant	15
4.2	Column or row rank profile certificate	16
4.3	Rank profile matrix certificate	19
5	Certificate for the signature of an integer matrix	22
6	Constant round certificates	23
6.1	Representative Laurent polynomial of a matrix	24
6.2	Constant round triangular equivalence certificate	24
6.3	Constant round certificates for the row and column rank profiles	26
7	Conclusion	26

1. Introduction

Within the setting of verifiable computing, we propose in this paper *interactive certificates* with the taxonomy of [3]. 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 sufficient 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 the Fiat-Shamir heuristic [10]: publicly and 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.

Our protocols follow the proof-of-work protocols of [12, 13] in that they verify that the Prover has performed some LU matrix factorization. However, they do so by verifying the factorization and the triangularity of the factors, which remain stored on the Prover side and are not communicated to the Verifier, rather than verifying the entire circuit that computes those factors by Lund-Fortnow-Karloff-Nisan polylog-compressive sumcheck protocols. In [6] we have applied [13] to matrices of exponential dimensions where the entries are computed from their indices by efficient circuits. Our version of the GKR proof-of-work protocol has a Verifier complexity that is, within a polylog factor, the depth of a parallel circuit whose local structure can be compute in polylog time, plus one linear scan of the input. The Prover complexity is within a polylog factor of the size of the circuit. The protocols here avoid those polylog factors.

It is possible to reduce the communication complexity in [13] to a constant number of rounds by when the space complexity is bounded [21] but it is not apparent to us how to asymptotically preserve the Prover's time complexity then (it remains polynomial-time).

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 worst case 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 [15], graph-based approaches cumulate this and other logarithmic factors, at least from a compression by magical graphs and from a dichotomic search [24].

We then propose two linear space interactive certificates. The first certificate is used to prove that two non-singular matrices are triangular equivalent (i.e. there is a triangular change of basis from one to the other). The second certificate is used to prove that a matrix has a generic rank profile. These two certificates are then applied to certify the row or column rank profile, the P (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 4 applications of A to a vector and a linear number of field operations. They are still elimination-based for the Prover, but do not require to communicate the obtained triangular decomposition.

An interesting setting would be for instance the case when the matrix A is sparse. Blackbox methods could then be used, when elimination-based method would suffer from some fill-in. Quite often though, elimination-based methods are then more limited by the available memory than by the number of computation. A Verifier could then outsource its computations to a server, for which fill-in would not be an issue, and use only still sparse matrix-vector multiplications to Verify the result.

For instance, for the Determinant, our new certificates require the computation of a PLUQ decomposition for the Prover, linear communication and Verifier time, with no restriction on the field size. The previously best known certificate for the determinant required instead some characteristic polynomial (CHARPOLY) computations.

With respect to [7] we propose a complete analysis of the rank profile matrix certificate 11 only sketched there; an application to computing the signature of a symmetric integral matrix; and a whole set of new certificates: for triangular equivalence, row and column rank profile, we are now able to propose protocols that preserve Prover and Verifier efficiency, while reducing the number of rounds from linear to constant. The constant round complexity is an important additional bonus in the delegation scenario, where network latency can make communication rounds more expensive. Note that the probabilistic analysis of [7, Theorem 4] omitted to account for several possibilities of failure, which is corrected here yielding a smaller probability of detecting a dishonest Prover.

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} ; $\mathcal{D}_n^{(2)}(\mathbb{F})$ represents block diagonal matrices with diagonal or anti-diagonal blocks of size 1 or 2. For two subsets of row indices \mathcal{I} and of column indices \mathcal{J} , $A_{\mathcal{I},\mathcal{J}}$ denotes the submatrix extracted from A in these rows and columns. The set of prime numbers will be denoted by \mathbb{P} . Lastly, $x \xrightarrow{\text{u.i.d.}} S$ denotes that x is

uniformly independently randomly sampled from S . In what follows, while computing the communication space, we consider that field elements and indices have the same size.

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}	Choose $S \subset \mathbb{F}$
		$v \xleftarrow{\text{u.i.d.}} S^n$
		$A(Bv) - Cv \stackrel{?}{=} 0$

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 Protocol 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 s.t. UQ is in row echelon form	$\xrightarrow{P,L,U,Q}$	UQ row echelonized?
		$A \stackrel{?}{=} PLUQ$, by Protocol 1
		Extract $Q[1], \dots, Q[r]$

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

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

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 [14] or [8, § 6.4]. 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 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, 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 $p \geq 1 - \frac{1}{|S|}$ 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 [14], 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.

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 $P \begin{bmatrix} L & 0 \end{bmatrix} P^T$ 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 $P \begin{bmatrix} L & 0_{m \times (m-r)} \end{bmatrix} P^T$ 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 $\bar{U} = Q^T \begin{bmatrix} U \\ 0 \\ & I_{n-r} \end{bmatrix} Q$ 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 check that $A = PLUQ$, using Freivalds' certificate [11], and check that L is echelonized by P and U^T by Q^T . If successful, the Verifier can just compute the rank profile matrix of A from P and Q , as shown in Protocol 3.

Prover	Verifier
	$A \in \mathbb{F}^{m \times n}$
	$\xrightarrow{P, L, U, Q}$
a $PLUQ$ decomp. of A revealing \mathcal{R}_A .	<ol style="list-style-type: none"> 1. $A \stackrel{?}{=} PLUQ$ by Protocol 1 2. Is PLP^T lower triangular? 3. Is $Q^T UQ$ upper triangular? <p style="margin-left: 20px;">Extract $\mathcal{R}_A = P \begin{bmatrix} I_r & \\ & 0_{(m-r) \times (n-r)} \end{bmatrix} Q$</p>

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 $O(r(m+n))$, 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 $Q^T UQ$ 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 $p \geq 1 - \frac{1}{|S|}$ by Freivalds' certificate or the $PLUQ$ decomposition does not reveal the rank profile matrix of A . In that case, Lemma 2 implies that either $P \begin{bmatrix} L & 0 \\ & \end{bmatrix} P^T$ 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 $O(r(m+n))$. A rank profile matrix revealing $PLUQ$ decomposition can be computed in $O(mnr^{\omega-2})$ operations [4]. The Verifier has to check if $A = PLUQ$, which can be achieved in $O(r(m+n)) + \mu(A)$ 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, and A is supposed to be *regular* (full rank). 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 number 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 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 Prover's responses depend on the Verifier challenges.

Prover		Verifier
	$A, B \in \mathbb{F}^{m \times n}$ A regular, $m \geq n$	
T lower triangular matrix s.t. $AT = B$	$\xrightarrow{1: T \text{ exists}}$	
$y_1 = T_{1,*} \begin{bmatrix} x_1 \\ 0 \\ \vdots \end{bmatrix}$	$\xleftarrow{2: x_1}$ $\xrightarrow{3: y_1}$	$x_i \xleftarrow{\text{u.i.d.}} S \subset \mathbb{F}$
\vdots	\vdots	
$y_n = T_{n,*} \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$	$\xleftarrow{2n: x_n}$ $\xrightarrow{2n+1: y_n}$	$y = [y_1 \ \dots \ y_n]^T$ $Ay \stackrel{?}{=} Bx$

Protocol 4: Lower triang. right equivalence of regular matrices

Theorem 3. *Let $A, B \in \mathbb{F}^{m \times n}$, $m \geq n$, 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 and occupies $2n$ communication space. The Prover complexity is $O(mn^{\omega-1})$ field operations and the Verifier computation cost is $\mu(A) + \mu(B)$ field operations.*

Proof. If the Prover is honest, then $AT = B$ with T triangular and she just computes $y = Tx$, so that $Ay = ATx = Bx$. If the Prover is dishonest, then she must try to convince the Verifier even if the matrices are inequivalent. For the sake of the argument, replace the random values x_1, \dots, x_n by algebraically independent variables X_1, \dots, X_n . Then there are two cases, either $AT \neq B$ for any T or there exists at least one such matrix T but none of them are triangular.

In the former case, $AT \neq B$, there is thus at least one inconsistent column in B , say the j -th. Then, there exists a Farkas certificate of inconsistency for that column (a vector z such that $z^T A = 0$ and $z^T B_{*,j} \neq 0$). This means that $z^T Ay = 0$ for any y , but $z^T B[X_1, \dots, X_n]^T$ is a not identically zero polynomial (at least the coefficient of X_j is non zero) of degree 1. Therefore, by the DeMillo-Lipton/Schwartz/Zippel lemma [2, 25, 23], its evaluation will be zero with probability at most $1/|S|$.

In the latter case, $AT = B$ but T is not triangular. Since A is regular, there is thus a unique $n \times n$ matrix T (that is, $T = A_{\text{left}}^{-1} B$, for any A_{left}^{-1} left inverse of A) such that $AT = B$: indeed $T = A_{\text{left}}^{-1} AT = A_{\text{left}}^{-1} B$. For the same reason, the equality $Ay = Bx = ATx$ implies $y = Tx$. If T is not lower triangular, there is a row-index i such that the entry $t_{i,j_m} \neq 0$ for some $j_m > i$. The test $y = Tx$ only succeeds if $y_i = \sum_{j=0}^n t_{i,j} x_j$. Now the Prover selects y_i before x_{j_m} is revealed. Therefore, with probability no more than $1/|S|$ the Verifier selects the field element $x_{j_m} = 1/t_{i,j_m} (y_i - \sum_{j \neq j_m} t_{i,j} x_j)$, and the test succeeds for false T .

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 [3, 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 non-singular 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.

Prover		Verifier
	$A \in \mathbb{F}^{n \times n}$ non-singular	
$A = LU$	$\xrightarrow{A \text{ has g.r.p.}}$	
	for i from n downto 1	
$[x \ y] = U [\phi \ \psi]$	$\xleftarrow{\phi_i, \psi_i}$ $\xrightarrow{x_i, y_i}$	$(\phi_i, \psi_i) \xleftarrow{\text{u.i.d.}} S^2 \subset \mathbb{F}^2$
$z^T = \lambda^T L$	$\xleftarrow{\lambda_i}$ $\xrightarrow{z_i}$	$\lambda_i \xleftarrow{\text{u.i.d.}} S \subset \mathbb{F}$
		$z^T [x \ y] \stackrel{?}{=} (\lambda^T A) [\phi \ \psi]$

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 $\geq (1 - \frac{1}{|S|})^{2n}$, perfectly complete, communicates $6n$ field elements, and can be computed in $O(n^\omega)$ field operations for the Prover and $\mu(A) + 8n$ field operations for the Verifier.*

Proof of Theorem 4. The protocol is perfectly complete: if $A = LU$, then $z^T [x \ y] = \lambda^T LU [\phi \ \psi] = \lambda^T A [\phi \ \psi]$, and the answer of any honest Prover will pass the Verifier test.

For any i such that the $(i - 1) \times (i - 1)$ leading submatrix of A has generic rank profile, we can write a partial LU decomposition of A with the following notations:

$$A = \underbrace{\begin{bmatrix} L^{(i)} & 0 \\ B^{(i)} & I_{n-i+1} \end{bmatrix}}_B \underbrace{\begin{bmatrix} U^{(i)} & V^{(i)} \\ 0 & C^{(i)} \end{bmatrix}}_C, \quad (1)$$

where $L^{(i)} \in \mathbb{F}^{(i-1) \times (i-1)}$ is non-singular lower triangular, $U^{(i)} \in \mathbb{F}^{(i-1) \times (i-1)}$ is non-singular upper triangular, $B^{(i)} \in \mathbb{F}^{(n-i+1) \times (i-1)}$, $V^{(i)} \in \mathbb{F}^{(i-1) \times (n-i+1)}$, $C^{(i)} \in \mathbb{F}^{(n-i+1) \times (n-i+1)}$.

Let $v^{[i\dots n]} = [v_i, \dots, v_n]^T \in \mathbb{F}^{n-i+1}$ for a vector $v \in \mathbb{F}^n$, and let

$$\eta_i = (\lambda^{[i\dots n]})^T C^{(i)} \phi^{[i\dots n]}, \quad \xi_i = (\lambda^{[i\dots n]})^T C^{(i)} \psi^{[i\dots n]}. \quad (2)$$

Consider the following predicate:

$$H_i: \eta_i = (z^{[i\dots n]})^T x^{[i\dots n]} \text{ and } \xi_i = (z^{[i\dots n]})^T y^{[i\dots n]}. \quad (3)$$

Note that H_1 is what the Verifier checks because then $B = I_n$. Note also that when A is in generic rank profile with $A = LU$ and $z^T = \lambda^T L$ and $x = U\phi$ and $y = U\psi$ then H_i is true for all i . To see this consider an LU-factorization $C^{(i)} = \bar{L}^{(i)} \bar{U}^{(i)}$ and the identity

$$A = \begin{bmatrix} L^{(i)} & 0 \\ B^{(i)} & I_{n-i+1} \end{bmatrix} \begin{bmatrix} U^{(i)} & V^{(i)} \\ 0 & C^{(i)} \end{bmatrix} = \begin{bmatrix} L^{(i)} & 0 \\ B^{(i)} & \bar{L}^{(i)} \end{bmatrix} \begin{bmatrix} U^{(i)} & V^{(i)} \\ 0 & \bar{U}^{(i)} \end{bmatrix} = LU. \quad (4)$$

Then $(z^{[i\dots n]})^T = (\lambda^{[i\dots n]})^T \bar{L}^{(i)}$ and $x^{[i\dots n]} = \bar{U}^{(i)} \phi^{[i\dots n]}$ and $y^{[i\dots n]} = \bar{U}^{(i)} \psi^{[i\dots n]}$ verify H_i . Note that the conditions are only tested by the Verifier for $i = 1$.

At stage i , let Λ_i, Φ_i and Ψ_i be variables for the random choices for λ_i, ϕ_i and ψ_i and Z_i be a variable for the Prover's choice of z_i . Then H_i in (3) expands as:

$$\begin{cases} x_i Z_i = \left(d\Phi_i + \overbrace{\sum_{j=i+1}^n C_{1,j-i+1}^{(i)} \phi_j}^e \right) \Lambda_i + a\Phi_i + f, \\ y_i Z_i = \left(d\Psi_i + \underbrace{\sum_{j=i+1}^n C_{1,j-i+1}^{(i)} \psi_j}_g \right) \Lambda_i + a\Psi_i + h, \end{cases} \quad (5)$$

where $d = C_{1,1}^{(i)}$ and $a = \sum_{k=i+1}^n \lambda_k C_{k-i+1,1}^{(i)}$, or equivalently

$$\begin{bmatrix} -(d\Phi_i + e) & x_i \\ -(d\Psi_i + g) & y_i \end{bmatrix} \begin{bmatrix} \Lambda_i \\ Z_i \end{bmatrix} = \begin{bmatrix} a\Phi_i + f \\ a\Psi_i + h \end{bmatrix}. \quad (6)$$

Suppose now that A is not in generic rank profile, and let i_0 be minimal such that the leading $i_0 \times i_0$ minor of A is equal to 0. On any corresponding partial LU decomposition this means that $d = C_{1,1}^{(i_0)} = 0$. Furthermore, because A is assumed to be non-singular, there exist indices k_0 with $2 \leq k_0 \leq n - i_0 + 1$, and j_0 with $2 \leq j_0 \leq n - i_0 + 1$ such that $C_{k_0,1}^{(i_0)} \neq 0$ and $C_{1,j_0}^{(i_0)} \neq 0$.

We will now prove the two following statements:

1. H_{i_0} is false with probability $\geq (1 - 1/|S|)^4$;
2. If H_{i+1} is false then H_i is false with probability $\geq (1 - 1/|S|)^2$ for $1 \leq i < i_0$.

Informally, this means that the Prover cannot achieve H_{i_0} with any choice of returned values x_1, \dots, z_{i_0} , with high probability and then this failure propagates with high probability to H_1 which is checked by the Verifier. By induction, this leads to a probability of $\geq (1 - 1/|S|)^{4+2(i_0-1)}$ that the Verifier check will fail when the matrix A is not in

generic rank profile. Since A is non-singular, $i_0 \leq n - 1$, and therefore this probability is $\geq (1 - 1/|S|)^{2n}$.

First, we prove Statement 1, that is the case when $d = 0$. The Verifier selects a random λ_{i_0} , and then the Prover a z_{i_0} . If the coefficient matrix in (6) is non-singular, there is a unique solution for Λ_{i_0} , which the Verifier will choose with probability $\leq 1/|S|$. Otherwise, the coefficient matrix is singular and the only way for the system to have a solution is that the determinant

$$\Delta = \begin{vmatrix} -e & a\Phi_{i_0} + f \\ -g & a\Psi_{i_0} + h \end{vmatrix} = -e(a\Psi_{i_0} + h) + g(a\Phi_{i_0} + f)$$

is equal to 0, which exactly happens in the three following cases:

- $[e \ g] = [0 \ 0]$, which happens with probability $\leq 1/|S|^2$ as $C_{1,j_0}^{\langle i_0 \rangle} \neq 0$;
- $a = 0$ (which happens with probability $\leq 1/|S|$ as $C_{k_0,1}^{\langle i_0 \rangle} \neq 0$) and $\begin{vmatrix} -e & f \\ -g & h \end{vmatrix} = 0$;
- otherwise, $ea \neq 0$ or $ga \neq 0$ and Δ is a nonzero polynomial of degree 1 in Φ_{i_0}, Ψ_{i_0} and evaluates to 0 for the random choices ϕ_{i_0}, ψ_{i_0} with probability $\leq 1/|S|$.

Overall, H_{i_0} is false with probability

$$\geq \left(1 - \frac{1}{|S|^2}\right) \left(1 - \frac{1}{|S|}\right)^3 \geq \left(1 - \frac{1}{|S|}\right)^4 \geq 1 - \frac{4}{|S|}$$

based on the random choices of the Verifier: ϕ_{j_0}, ψ_{j_0} yield $[e \ g] \neq [0 \ 0]$; λ_{k_0} yields $a \neq 0$; ϕ_{i_0}, ψ_{i_0} yield $\Delta \neq 0$; λ_{i_0} avoids the unique solution to (6).

For Statement 2, consider the predicate H_i (3) at $i < i_0$, that is $d \neq 0$. Similarly, if the coefficient matrix in (6) is non-singular, there is a unique solution for Λ_i , which the Verifier will choose with probability $\leq 1/|S|$. Otherwise, the coefficient matrix is singular and the only way for the system to have a solution is that the following determinant is equal 0:

$$0 = \Delta = \begin{vmatrix} -(d\Phi_i + e) & a\Phi_i + f \\ -(d\Psi_i + g) & a\Psi_i + h \end{vmatrix} = (df - ae)\Psi_i - (dh - ag)\Phi_i - eh + gf.$$

We block decompose the bottom right block in the incomplete right factor in (1) $C^{\langle i \rangle} = \begin{bmatrix} d & r^T \\ s & W \end{bmatrix}$, where $d = C_{1,1}^{\langle i \rangle} \neq 0$. We have $C^{\langle i+1 \rangle} = W - \frac{1}{d}sr^T$. Now since $a = (\lambda^{[i+1\dots n]})^T s$, $e = r^T \phi^{[i+1\dots n]}$, we have $ae = (\lambda^{[i+1\dots n]})^T sr^T \phi^{[i+1\dots n]}$ and

$$\begin{aligned} f - \frac{ae}{d} &= (\lambda^{[i+1\dots n]})^T C^{\langle i+1 \rangle} \phi^{[i+1\dots n]} - (z^{[i+1\dots n]})^T x^{[i+1\dots n]} \\ &= \eta_{i+1} - (z^{[i+1\dots n]})^T x^{[i+1\dots n]}. \end{aligned}$$

Similarly, $h - \frac{ag}{d} = \xi_{i+1} - (z^{[i+1\dots n]})^T x^{[i+1\dots n]}$, and these two quantities are not equal to 0 simultaneously, for otherwise H_{i+1} would be true. Therefore Δ is a nonzero polynomial of degree 1 in Φ and Ψ . It is equal to 0 with probability $\leq 1/|S|$. Overall, H_i is false with probability $\geq (1 - 1/|S|)^2$ based on the random choices for λ_i, ϕ_i and ψ_i made by the Verifier.

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 six vectors, $\phi, \psi, \lambda, x, y, z$, and the Verifier's work is four dot-products and one multiplication by the initial matrix A (certifying the transposed to have a single matrix times λ -vector product). \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 corresponds to our lack of specification of the diagonals 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 triangular, and the prover has to first commit the main diagonal D of U .

For a non-singular upper triangular matrix U with diagonal $D = \text{Diag}(d_1, \dots, d_n)$, the matrix $U_1 = D^{-1}U$ is unit triangular. Thus, for any $\psi = \begin{bmatrix} \psi_1 \\ \tilde{\psi} \end{bmatrix} \in \mathbb{F}^n$: $U\psi = DU_1\psi = D\left(\psi + \begin{bmatrix} \tilde{U}_1\tilde{\psi} \\ 0 \end{bmatrix}\right)$, where $\tilde{U}_1 = (U_1 - I_n)_{\{1, \dots, n-1\}, \{2, \dots, n\}}$ upper triangular in $\mathbb{F}^{(n-1) \times (n-1)}$. 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 $\tilde{y} = \tilde{U}_1 \tilde{\psi}$, where $\tilde{L} = (L - I_n)_{\{2, \dots, n\}, \{1, \dots, n-1\}}$ lower triangular in $\mathbb{F}^{(n-1) \times (n-1)}$. Then the Verifier will compute by himself the complete vectors. This ensures that L is unit triangular and that $U = DU_1$ with U_1 unit triangular.

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 Protocol 6.

Theorem 5. *The Protocol 6 requires less than $8n$ extra communications. The computational cost for the Prover is $O(n^\omega)$ and the Verifier cost is bounded by $\mu(A) + 12n + o(n)$. The protocol is perfectly complete and fails the verification for a non generic rank profile matrix $AP^{-1} = AP^T$ with probability $\geq (1 - \frac{1}{|S|})^{2n}$.*

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:

$$\begin{aligned} z^T Dy &= (\lambda^T + [\tilde{z}^T \ 0])D \left(\psi + \begin{bmatrix} \tilde{y} \\ 0 \end{bmatrix} \right) \\ &= \begin{bmatrix} \lambda_1 & \tilde{\lambda}^T \end{bmatrix} \left(I + \begin{bmatrix} 0 & 0 \\ \tilde{L} & 0 \end{bmatrix} \right) D \left(\begin{bmatrix} 0 & \tilde{U}_1 \\ 0 & 0 \end{bmatrix} + I \right) \begin{bmatrix} \psi_1 \\ \tilde{\psi} \end{bmatrix}. \end{aligned}$$

The same is true for λ and ϕ , so that the protocol is perfectly complete.

Now, the last part of the Protocol 6 is actually a verification that AP^T has generic rank profile, in other words that there exists lower and upper triangular matrices L^*

Prover	$A \in \mathbb{F}^{n \times n}$ non-singular	Verifier
$A = LDU_1P$	$\xrightarrow{P, D}$	$P \stackrel{?}{\in} \mathcal{S}_n, D \in \mathcal{D}_n(\mathbb{F}^*)$
$\widetilde{U}_1 = (U_1 - I_n)_{\{1, \dots, n-1\}, \{2, \dots, n\}}$		
$\widetilde{L} = (L - I_n)_{\{2, \dots, n\}, \{1, \dots, n-1\}}$		
Choose $S \subset \mathbb{F}$ for i from n downto 2:		
$[\bar{x} \ \bar{y}] = \widetilde{U}_1 [\tilde{\phi} \ \tilde{\psi}]$	$\xleftarrow{\phi_i, \psi_i}$	$(\phi_i, \psi_i) \xleftarrow{\text{u.i.d.}} S^2$
	$\xrightarrow{\bar{x}_{i-1}, \bar{y}_{i-1}}$	
$\bar{z} = \tilde{\lambda}^T \widetilde{L}$	$\xleftarrow{\lambda_i}$	$\lambda_i \xleftarrow{\text{u.i.d.}} S$
	$\xrightarrow{\bar{z}_{i-1}}$	
	\vdots	
	\vdots	$\phi_1, \psi_1, \lambda_1 \xleftarrow{\text{u.i.d.}} S^3$
		$[x \ y] = [\phi \ \psi] + \begin{bmatrix} \bar{x} & \bar{y} \\ 0 & 0 \end{bmatrix}$
		$z^T = (\lambda^T + [\bar{z}^T \ 0])$
		$z^T D [x \ y] \stackrel{?}{=} (\lambda^T A) P^T [\phi \ \psi]$

Protocol 6: LDUP decomposition (linear communication)

and U^* such that $AP^T = L^*U^*$. This verification is sound by Theorem 4. Next, the multiplication by the diagonal D is performed by the Verifier, in order to be 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 $a + \begin{bmatrix} \tilde{b} \\ 0 \end{bmatrix}$ is also done by the Verifier, in order to have in fact a guarantee that L^* and U_1^* are unit triangular.

Overall, if the matrix AP^T does not have generic rank profile, 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 \bar{z} , \bar{x} and \bar{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 vector-matrix multiplication by A (for $(\lambda^T A)$), that is $n + 3(n-1) + 4(2n-1)$. \square

We do not need the following fact to show that Protocol 6 correctly verifies generic rank profileness, but furthermore, this protocol actually gives some guarantees on the actual values of D and x, y, z :

Proposition 1. *Let S be a finite subset of \mathbb{F} in Protocol 6, if AP^T is not in generic rank profile, or else if the committed D does not correspond to the unique decomposition $AP^T = LDU_1$ or $[x \ y] \neq U_1 [\phi \ \psi]$ or $z^T \neq \lambda^T L$, then the verification will fail with probability $\geq (1 - \frac{1}{|S|})^{2n}$, and therefore Protocol 6 is sound.*

Proof. For a dishonest Prover, either

(i) AP^T is not in generic rank profile, then Protocol 6 will detect it with the probability of Theorem 5;

(ii) or she could still try, to send modified vectors $\bar{x}, \bar{y}, \bar{z}$ or diagonal D .

Let then $D^*, x^* = \phi + \begin{bmatrix} \bar{x}^* \\ 0 \end{bmatrix} = U_1 \phi$, $y^* = \psi + \begin{bmatrix} \bar{y}^* \\ 0 \end{bmatrix} = U_1 \psi$ and $z^* = \begin{bmatrix} \bar{z}^* \\ 0 \end{bmatrix} + \lambda = L^T \lambda$ be the correct expected diagonal and vectors. Let also $i_0 \leq n$ be the largest index such that there is at least one discrepancy in $d_{i_0}, \bar{x}_{i_0}, \bar{y}_{i_0}$ or \bar{z}_{i_0} that makes at least one of them respectively different from $d_{i_0}^*, \bar{x}_{i_0}^*, \bar{y}_{i_0}^*$ or $\bar{z}_{i_0}^*$ ($\bar{x}_n = \bar{x}_n^* = 0, \bar{y}_n = \bar{y}_n^* = 0, \bar{z}_n = \bar{z}_n^* = 0$ by default). Then H_i of (3) is true for all i such that $n \geq i > i_0$, and thus in particular H_{i_0+1} is true (H_{n+1} is true by default). Now, H_{i_0} is also true if and only if we have both:

$$\begin{cases} z_{i_0} d_{i_0} x_{i_0} = z_{i_0}^* d_{i_0}^* x_{i_0}^*, \\ z_{i_0} d_{i_0} y_{i_0} = z_{i_0}^* d_{i_0}^* y_{i_0}^*. \end{cases} \quad (7)$$

Indeed, H_{i_0} is $(z^{[i_0 \dots n]})^T D^{[i_0 \dots n]} x^{[i_0 \dots n]} = (z^{*[i_0 \dots n]})^T D^{[i_0 \dots n]} x^{*[i_0 \dots n]}$ and similarly H_{i_0+1} is $(z^{[i_0+1 \dots n]})^T D^{[i_0+1 \dots n]} x^{[i_0+1 \dots n]} = (z^{*[i_0+1 \dots n]})^T D^{[i_0+1 \dots n]} x^{*[i_0+1 \dots n]}$. Further, Equations (7), with $a = \bar{z}_{i_0}^* d_{i_0}^* \bar{x}_{i_0}^* - \bar{z}_{i_0} d_{i_0} \bar{x}_{i_0}$, and $b = \bar{z}_{i_0}^* d_{i_0}^* \bar{y}_{i_0}^* - \bar{z}_{i_0} d_{i_0} \bar{y}_{i_0}$, is equivalent to:

$$\begin{cases} \lambda_{i_0} \phi_{i_0} (d_{i_0} - d_{i_0}^*) + \lambda_{i_0} (d_{i_0} \bar{x}_{i_0} - d_{i_0}^* \bar{x}_{i_0}^*) + \phi_{i_0} (d_{i_0} \bar{z}_{i_0} - d_{i_0}^* \bar{z}_{i_0}^*) - a = 0, \\ \lambda_{i_0} \psi_{i_0} (d_{i_0} - d_{i_0}^*) + \lambda_{i_0} (d_{i_0} \bar{y}_{i_0} - d_{i_0}^* \bar{y}_{i_0}^*) + \psi_{i_0} (d_{i_0} \bar{z}_{i_0} - d_{i_0}^* \bar{z}_{i_0}^*) - b = 0. \end{cases} \quad (8)$$

However, $\lambda_{i_0}, \phi_{i_0}, \psi_{i_0}$ are chosen by the Verifier after $d_{i_0}, \bar{x}_{i_0}, \bar{y}_{i_0}$ and \bar{z}_{i_0} have been committed. Hence, on the one hand, if $d_{i_0} \neq d_{i_0}^*$ then the coefficient of λ_{i_0} in one of the two polynomials is not equal to 0 for a random ϕ_{i_0} with probability $\geq 1 - 1/|S|^2$ and then that polynomial does not vanish for a random λ_{i_0} with probability $\geq (1 - 1/|S|^2)(1 - 1/|S|)$, based on the random choices made by the Verifier, and H_{i_0} is violated.

On the other hand, if $d_{i_0} = d_{i_0}^* \neq 0$, they can be removed from Equations (8) which then simplifies (for $i_0 < n$) as:

$$\begin{cases} \lambda_{i_0} (\bar{x}_{i_0} - \bar{x}_{i_0}^*) + \phi_{i_0} (\bar{z}_{i_0} - \bar{z}_{i_0}^*) - (\bar{z}_{i_0}^* \bar{x}_{i_0}^* - \bar{z}_{i_0} \bar{x}_{i_0}) = 0, \\ \lambda_{i_0} (\bar{y}_{i_0} - \bar{y}_{i_0}^*) + \psi_{i_0} (\bar{z}_{i_0} - \bar{z}_{i_0}^*) - (\bar{z}_{i_0}^* \bar{y}_{i_0}^* - \bar{z}_{i_0} \bar{y}_{i_0}) = 0. \end{cases} \quad (9)$$

When there is at least one discrepancy with the expected vector coefficients, then Equations (9) can be considered as 2 polynomials that are not simultaneously identically zero. Thus they both vanish with probability $\leq 1/|S|$ based on the random choices made by the Verifier. H_{i_0} is thus false with probability $\geq (1 - 1/|S|)$. As in the proof of Theorem 4, this propagates with high probability, to H_1 and the dishonest Prover is detected with probability $\geq (1 - 1/|S|)^{2(n-1)} (1 - 1/|S|^2) (1 - 1/|S|) \geq (1 - 1/|S|)^{2n}$.

Overall, both (i), AP^T is not GRP, or (ii), AP^T is GRP but some diagonal or vector elements is wrong, are detected with probability $\geq (1 - 1/|S|)^{2n}$. \square

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 [15, 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 [20]. 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 a field larger than $2n$, where that of [5, Theorem 14] requires 3 random elements but a field larger than n^2 . Finally the new protocol requires $O(n)$ rounds when 2 are sufficient in [5, Theorem 14].

For instance, using the routines shown in Table 1 (one matrix-vector multiplication with a dense matrix is denoted `fgemv`), the determinant of an $50k \times 50k$ random dense matrix can be computed in about 24 minutes, where with the certificate of Protocol 6, the overhead of the Prover is less than 5s and the Verifier time is about 1s.

Computations use the FFLAS-FFPACK library [16] 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. We see that a linear communication cost can be masked by a quadratic number of computations, when a quadratic communication cost could be up to two orders of magnitude worse.

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
<code>fgemv</code>	0.0013s	0.038s	1.03s

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

4.2. Column or row rank profile certificate

In Protocols 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 [3]. 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 ($\|\gamma\|_0$ denotes the Hamming weight of the vector γ). 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 . $LINSYS(r)$ denotes a complexity bound for solving a linear system of rank r by the Prover.

Prover		Verifier
	$A \in \mathbb{F}^{m \times n}$	
R s.t. $\text{rank}(A) \leq R$	\xrightarrow{R}	
		Choose $S \subset \mathbb{F}$
	\xleftarrow{w}	$v \xleftarrow{\text{u.i.d.}} S^n, w = Av$
$A\gamma = w$	$\xrightarrow{\gamma}$	$\ \gamma\ _0 \stackrel{?}{=} R$
		$A\gamma \stackrel{?}{=} w$

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 γ of an upper bound for the rank of A is sound, with probability larger than $1 - \frac{1}{|S|}$, perfectly complete, occupies $m + n$ communication space, can be computed in $LINSYS(r)$ and verified in $2\mu(A) + n$ time.*

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

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 β of a lower bound for the rank of A is sound, with probability larger than $1 - \frac{1}{|S|}$, perfectly complete and occupies $m + 2r$ communication space, can be computed in $LINSYS(r)$ and verified in $\mu(A) + r$ operations.*

Note that the communication in Protocol 8 involve sending r indices for \mathcal{J} , then m field elements for vector v , and only r field elements for vector β , as it has only r non-zero coefficients which positions are already indicated by \mathcal{J} . Hence the total communication cost is $m + 2r$.

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, one can reduce the 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 an upper triangular matrix Γ such that:

$$[A_{*,c_1} \ A_{*,c_2} \ \dots \ A_{*,c_{r-1}}] \Gamma = A \underbrace{\begin{bmatrix} v_1 & v_1 & \dots & \dots & v_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ v_{c_1-1} & \vdots & \vdots & \vdots & \vdots \\ 0 & v_{c_2-1} & \vdots & \vdots & \vdots \\ 0 & 0 & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & v_{c_{r-1}-1} & \vdots \\ 0 & 0 & 0 & 0 & v_n \end{bmatrix}}_V. \quad (10)$$

Note that $V = \text{Diag}(v_1, \dots, v_n)W$ where $W = [\mathbf{1}_{i < c_{j+1}}]_{i,j}$ (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 Protocol 9.

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 $m + n + 4r$ and a Verifier cost bounded by $2\mu(A) + n + 3r$.*

Prover		Verifier
	$A \in \mathbb{F}^{m \times n}$	
$\mathcal{J} = (c_1, \dots, c_r)$ CRP of A	Protocol 8 on A $\overline{\mathcal{J} = \text{CRP}(A)}$ $\xrightarrow{\quad}$ $\text{rank } A \stackrel{?}{\geq} r$	
	\xleftarrow{v}	Choose $S \subset \mathbb{F}$ $v \xleftarrow{\text{u.i.d.}} S^n$ $W = [\mathbf{1}_{i < c_{j+1}}] \in \mathbb{F}^{n \times r}$
$V = \text{Diag}(v_1, \dots, v_n)W$ (see (10)) Γ upper tri. s.t. $A_{*,\{c_1, \dots, c_r\}}\Gamma = AV$		
$y = \Gamma x$	$\xleftrightarrow{x \text{ (Cert. 4)} y}$	$x \xleftarrow{\text{u.i.d.}} S^r$ $z = \text{Diag}(v_1, \dots, v_n)(Wx)$ $z_{c_j} = z_{c_j} - y_j$ for $j = 1..r$ $Az \stackrel{?}{=} 0$

Protocol 9: Certificate for the column rank profile

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 cost of $n + 2r$, which adds up to the $m + 2r$ communication cost of Protocol 8. 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.

Prover	Verifier
$A \in \mathbb{F}^{n \times n}$ invertible	
$A = LDUP$, with $P = \mathcal{R}_A$	$P \stackrel{?}{\in} \mathcal{S}_n, D \stackrel{?}{\in} \mathcal{D}_n(\mathbb{F})$
Protocol 4 : P^T and $P^T U$ are left up. tri. equiv. with random proj.	
$\bar{U} = P^T U P$	\bar{U} is upper tri.
$f^T = e^T \bar{U}$	Choose $S \subset \mathbb{F}$
	<div style="display: flex; justify-content: space-between; align-items: center;"> <div style="text-align: center;"> $\xleftarrow{e_1, \dots, e_n}$ $\xrightarrow{f_1, \dots, f_n}$ </div> <div style="text-align: right;"> for $i = 1, \dots, n, e_i \xleftarrow{\text{u.i.d.}} S$ </div> </div>
	<div style="display: flex; align-items: center;"> <div style="border: 1px solid black; padding: 5px; margin-right: 10px;"> Protocol 6 on A $\xleftarrow{[\tilde{\phi} \ \tilde{\psi}]}$ $\xrightarrow{[\tilde{x} \ \tilde{y}]}$ </div> <div style="text-align: right;"> $\phi, \psi \xleftarrow{\text{u.i.d.}} S^n$ </div> </div>
	Now $\begin{bmatrix} x & y \end{bmatrix}$ is $U \begin{bmatrix} \phi & \psi \end{bmatrix}$
	$e^T P^T x \stackrel{?}{=} f^T P^T \phi$

Protocol 10: Rank profile matrix of an invertible matrix

Theorem 9. *Protocol 10 is sound, with probability $\geq (1 - \frac{1}{|S|})^{2n}$, 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 $\mu(A) + 16n + o(n)$.*

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 $(1 - \frac{1}{|S|})^{2n}$, 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

$$e^T P^T U P - f^T = 0. \quad (11)$$

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 (11) 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^TUP is not upper triangular, the verification will detect it with probability $\geq (1 - \frac{1}{|S|})^{2n}$.

The Verifier's cost is that of Protocol 6 with two additional dot products for the last step, which is $\mu(A) + 16n + o(n)$. Similarly, the communication cost is that of Protocol 6 plus the size of e and f for a total of $10n$. The Prover remains unchanged. \square

Finally, we use [8, Algorithm 4] to certify the rank profile matrix of any matrix, even a singular one. To do so, we need to verify the row rank profile and the column rank profile of the input matrix, which can be done with two applications of Certificate 9. Then, we certify the rank profile matrix of the $r \times r$ selection of lexicographically minimal independent rows and columns we obtained before. This is done by an application of Certificate 10. We now define $\mathcal{E}_{m, \{i_1, \dots, i_n\}}$ as the $m \times n$ matrix whose j -th column is the i_j -th vector of the m -dimensional canonical basis. This certificate is detailed in Protocol 11, in the case where $m \leq n$. If $n < m$, one should first apply Protocol 9 on A to compute its column rank profile, and then apply the verification steps of the same protocol for the row rank profile of A . The application of Protocol 10 remains unchanged.

Theorem 10. *Protocol 11 is sound, with probability $\geq (1 - \frac{1}{|S|})^{2n}$, and perfectly complete. The Prover cost is $O(mnr^{\omega-2})$ field operations, the communication space is bounded by $m + n + \min(m, n) + 17r$ and the Verifier cost is bounded by $4\mu(A) + m + n + 21r$.*

Proof. If the Prover is honest, \mathcal{I} is the row rank profile of A and \mathcal{J} is the column rank profile of A . Then, the application of Protocol 10 will output the correct rank profile matrix of $A_{\mathcal{I}, \mathcal{J}}$ which will lead the Verifier to the correct rank profile matrix of A , as described in [8, Theorem 37]. Note that one only needs to verify the lower bound on the rank of A once, which is why Certificate 9 is fully executed once, while the second run only verifies that the committed rank profile is a rank profile indeed.

Now, for the soundness, Prover has a probability $\geq 1 - 1/|S|$ to be caught when cheating while running Certificate 9, and a probability $\geq (1 - \frac{1}{|S|})^{2n}$ to be caught when cheating while running Certificate 10. Overall, this makes a probability $\geq (1 - \frac{1}{|S|})^{2n}$ for the Verifier to catch a cheating Prover during the execution of Certificate 11.

For the complexity, Prover time complexity is bounded by the complexity of performing a PLUQ decomposition of the input matrix, $O(mnr^{\omega-2})$. The Verifier complexity is the one of one full application of Protocol 9 and one application of Protocol 9 without applying Protocol 8, which makes $3\mu(A) + n + m + 5r$, plus one application of Protocol 10 over an $r \times r$ matrix for a cost of $\mu(A) + 16r + o(r)$, the computation of \mathcal{R}_A only consists of memory operations, hence a total cost of $4\mu(A) + m + n + 21r + o(r)$ field operations. Communication space is computed as follows: a full application of Protocol 9 on A if $m \geq n$, on A^T otherwise, an application of the same Protocol without the underlying Protocol 8 which makes $n + m + \min(m, n) + 7r$ and the same application of Protocol 10 as above, for a cost of $10r$, hence a total communication space of $m + n + \min(m, n) + 17r$. \square

Prover	Verifier
$A \in \mathbb{F}^{m \times n}$ (assuming $m \leq n$ w.l.o.g.)	
<div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 0 auto;"> Protocol 9 on A^T $\bar{\mathcal{I}} = \bar{R}\bar{R}\bar{P}(\bar{A}) \in \llbracket \bar{1}, \bar{m} \rrbracket^{\bar{r}}$ </div>	
	Now, $\mathcal{I} = RRP(A)$ and $\text{rank}(A) \geq r$
$\mathcal{J} = CRP(A) = (c_1, \dots, c_r)$	$\xrightarrow{\mathcal{J}}$
	\xleftarrow{v}
$V = \text{Diag}(v_1, \dots, v_n)W$ (see (10))	Choose $S \subset \mathbb{F}$
Γ upper tri. s.t. $A_{*,\{c_1, \dots, c_r\}}\Gamma = AV$	$v \xleftarrow{\text{u.i.d.}} S^n$
$y = \Gamma x$	$W = [\mathbf{1}_{i < c_{j+1}}] \in \mathbb{F}^{n \times r}$
	$\xleftarrow{x \text{ (Cert. 4)} y}$
	$x \xleftarrow{\text{u.i.d.}} S^r$
	$z = \text{Diag}(v_1, \dots, v_n)(Wx)$
	$z_{c_j} = z_{c_j} - y_j$ for $j = 1..r$
	$Az \stackrel{?}{=} 0$
<div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 0 auto;"> Protocol 10 on $A_{\mathcal{I}, \mathcal{J}}$ $\bar{\mathcal{R}}_r = \bar{R}\bar{P}\bar{M}(A_{\mathcal{I}, \mathcal{J}})$ </div>	
	$\mathcal{R}_A = \mathcal{E}_{m, \mathcal{I}} \mathcal{R}_r \mathcal{E}_{n, \mathcal{J}}^T$

Protocol 11: Rank profile matrix

5. Certificate for the signature of an integer matrix

The signature of a symmetric matrix is the triple (n_+, n_-, n_0) indicating the number of positive, negative, and zero eigenvalues, respectively. Just like [3, Theorem 5], the idea is that the Prover commits the signature, and then certifies it modulo a Verifier chosen prime. This works directly for the signature algorithm in [15, Corollary 1] together with the CHARPOLY protocol of [5, Theorem 14]. As in § 4.1, in the dense case we propose here to replace the CHARPOLY computation with a symmetric Gaussian elimination.

Over the rationals, an algorithm for the Prover could be to first compute and certify the rank of A , and to compute a permutation matrix P such that $P^T A P$ has generic rank profile: for instance compute a $PL_p \Delta_p L_p^T P^T$ factorization modulo a sufficiently large prime p . Then $B = [I_r | 0] P^T A P \begin{bmatrix} I_r \\ 0 \end{bmatrix}$ is symmetric and non-singular. It is then sufficient to lift or reconstruct only the block diagonal matrix Δ over \mathbb{Q} of a non-pivoting symmetric factorization of B (the unit triangular matrix over \mathbb{Q} need not be computed). Compared to an integer characteristic polynomial computation this gains in practice an order of magnitude in efficiency for the Prover as shown on the logscale Figure 1, using LinBox-1.5.1 [17].


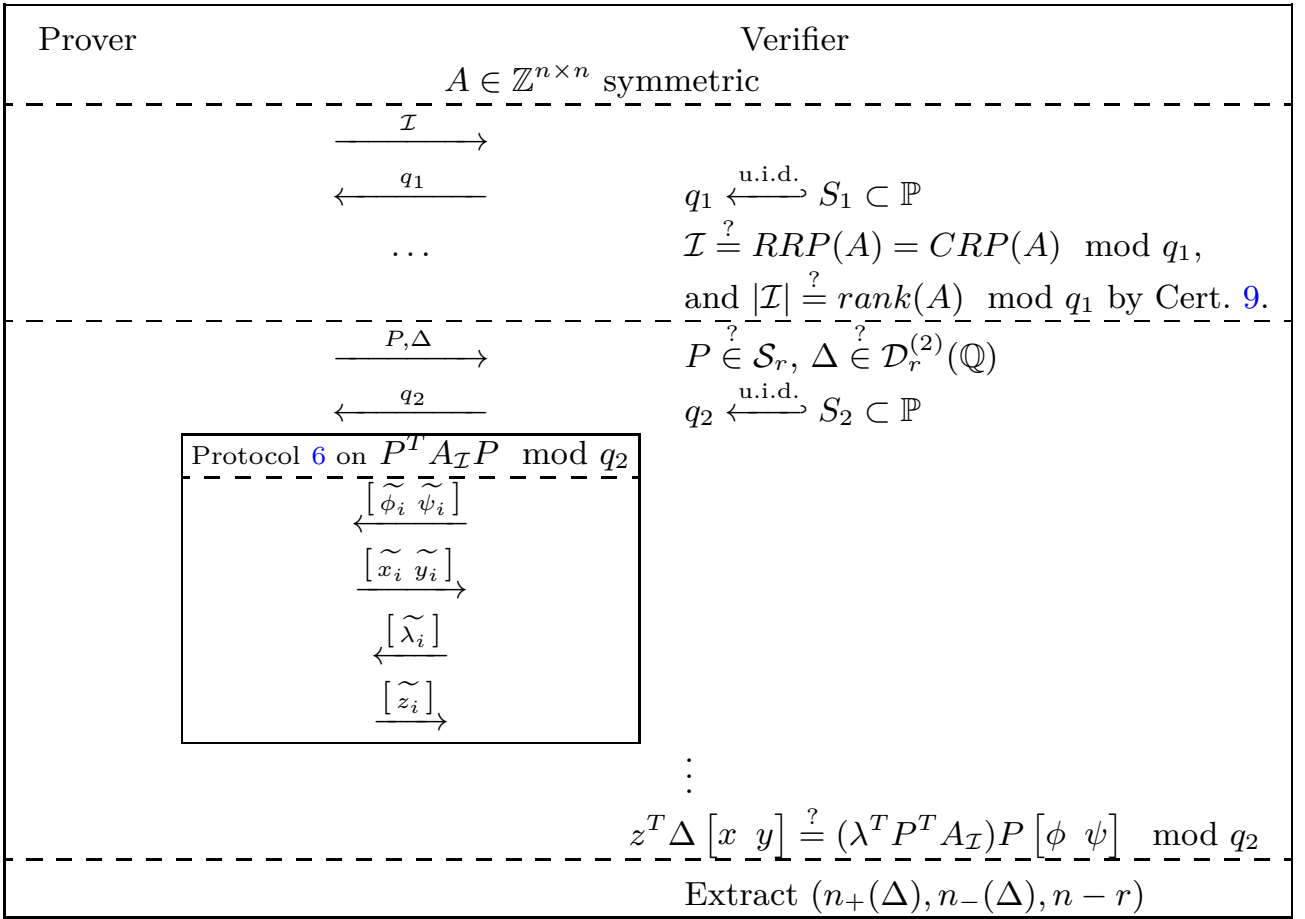


Figure 1: (Verifiable) signature computation on a single Intel Skylake core @3.4GHz.

For the verification, the block diagonal matrix Δ , and the permutation P are committed. The Verifier then randomly chooses a prime q and enters an interactive certification process for P and $\Delta \pmod q$ using Protocol 6, as shown on Protocol 12.

From [3, Theorem 5], we let $h = \log_2(\sqrt{n}^n \|A\|_\infty^n)$ be the logarithm of Hadamard's bound for the invariant factors of A . There cannot be more than h primes reducing the rank. Therefore it is possible to sample $c \cdot h$ distinct primes of magnitude bounded by $O(h \log(h))$ for any constant $c > 2$ and select q_1 from that set S_1 . Once the rank is certified, the Prover can compute the permutation and lift the diagonal. Finally the



Protocol 12: Certificate for the signature of a symmetric matrix

rational $PL\Delta L^T P^T$ factorization of the full rank matrix can be similarly verified modulo a prime q_2 . As for the determinant, no more than h primes can reduce the rank of Δ and q_2 can be selected from the same kind of set. We have proven:

Corollary 2. *For a symmetric matrix $A \in \mathbb{Z}^{n \times n}$, certificate 12 for its signature is sound and perfectly complete.*

The communication comprise that of the Certificate 6, the permutation matrix P , all of size n , as well as small primes bounded by h , and finally Δ . Just like that of the characteristic polynomial, the size of Δ can be quadratic and therefore the whole protocol is not linear. Thus a simpler quadratic certificate communicating the triangular matrix L modulo q_2 , and checking the decomposition $A - L\Delta L^T$ via Freivalds certificate might also work. But then the communication and Verifier time would always be quadratic. Instead, Protocol 12, just like the Protocol using the characteristic polynomial, is better if the size of the determinant is small, as then the size of Δ might be much less than that of L (for instance linear if the determinant is a constant). Protocol 12 is also interesting if $\mu(A)$ is less than quadratic.

6. Constant round certificates

When delegating computations, the network latency can make communication rounds expensive. It can therefore also be interesting not only to reduce the communication volume, but also the number of rounds. We therefore propose in this section a certificate with a constant number of rounds for triangular equivalence, still preserving Prover

efficiency as well as linear communication volume and Verifier cost. This applies then directly, as previously shown, to row or column rank profiles. However it fails to apply to the generic rank profile, at least in a straightforward manner, and we were unable to produce such a certificate in constant round for this task.

6.1. Representative Laurent polynomial of a matrix

Following a technique in [18], we first define the *representative Laurent polynomial*, $P_A(X)$ of an $m \times n$ matrix A as :

$$P_A(X) = [1 \ X \ X^2 \ \dots \ X^{m-1}] \cdot A \cdot \begin{bmatrix} 1 \\ X^{-1} \\ \vdots \\ X^{1-n} \end{bmatrix} = \sum_{i=1}^m \sum_{j=1}^n A_{i,j} X^{i-j}$$

Therefore, if a matrix is lower triangular, then its representative Laurent polynomial cannot have negative powers and it is therefore a polynomial of degree at most $m - 1$. The converse is not true, consider for instance an upper diagonal with two opposite coefficients : $A_{i,i+1} = 0$ for all i except $A_{1,2} = -A_{2,3}$. Generically, if one pre-multiplies A on the right by a random non-zero diagonal matrix, these cancellations will not occur as in general $d_1 A_{1,2} \neq -d_2 A_{2,3}$ unless $A_{1,2} = A_{2,3} = 0$.

6.2. Constant round triangular equivalence certificate

From this representation we can obtain a triangular equivalence certificate that requires only a constant number of rounds: the Prover commits that polynomial, then the Verifier will evaluate the polynomial at a random point and compare this to the actual projections. The counterpart is that the field size must be sufficiently large so that the polynomial identity testing does not fail. The full certificate is given in Protocol 13. It requires that the Prover solves a regular system (this is checked deterministically by reapplying the resulting vector), and a preconditioning by a diagonal matrix to prevent cancellations.

Theorem 11. *Let $A, B \in \mathbb{F}^{m \times n}$, $m \geq n$, and assume A is regular. Certificate 13 is sound, with probability larger than $1 - 2 \frac{n-1}{|S|}$ and perfectly complete. The Prover cost is dominated by one system solving, $O(mn^{\omega-1})$, the communication space is bounded by $3n + 1$ and the Verifier cost is bounded by $\mu(A) + \mu(B) + 7n$*

Proof. Let $x = D \begin{bmatrix} 1 \\ \lambda^{-1} \\ \vdots \\ \lambda^{1-n} \end{bmatrix}$. As A is regular, there is only one solution y to $Ay = Bx$, and

$y = Lx$. Therefore $[1 \ \lambda \ \dots \ \lambda^{n-1}] \cdot y = [1 \ \lambda \ \dots \ \lambda^{n-1}] \cdot LD \begin{bmatrix} 1 \\ \lambda^{-1} \\ \vdots \\ \lambda^{1-n} \end{bmatrix} = P_{LD}(\lambda)$ and the

protocol is correct. For the soundness:

Prover		Verifier
$A, B \in \mathbb{F}^{m \times n}$		
A is regular, $m \geq n$		
$\exists L$ lower triang. s.t. $AL = B$	$\xrightarrow{\exists L}$	
$g(X) = P_{LD}(X)$	\xleftarrow{D}	Choose $S \subset \mathbb{F}$
	$\xrightarrow{g(X)}$	$D \xleftarrow{\text{u.i.d.}} \mathcal{D}_n(S \setminus \{0\})$
	$\xleftarrow{\lambda}$	$g \stackrel{?}{\in} \mathbb{F}[X]_{\deg \leq n-1}$
		$\lambda \xleftarrow{\text{u.i.d.}} S$
y , s.t. $A \cdot y = B \cdot D \cdot \begin{bmatrix} 1 \\ \lambda^{-1} \\ \vdots \\ \lambda^{1-n} \end{bmatrix}$	\xrightarrow{y}	$A \cdot y \stackrel{?}{=} B \cdot D \cdot \begin{bmatrix} 1 \\ \lambda^{-1} \\ \vdots \\ \lambda^{1-n} \end{bmatrix}$
		$g(\lambda) \stackrel{?}{=} [1 \ \lambda \ \dots \ \lambda^{n-1}] \cdot y$

Protocol 13: Constant round linear communication certificate for triangular equivalence

- As A is regular, there is only one solution y to $Ay = Bx$, thus that check ensures that y is correct, unless not all columns in B are in the column space of A , which is handled as in the proof of Theorem 3.
- If L is not triangular then its upper part is not identically zero. Therefore by considering D as a diagonal matrix of indeterminates, at least one coefficient of negative degree of the representative rational fraction LD will be non identically zero. As those are of degree 1 in the indeterminates of D , for a random diagonal D , the representative rational fraction of LD will not be a polynomial with probability at least $1 - \frac{1}{|S|-1}$.
- If g is not a polynomial of degree at most $n - 1$, it is not the representative of a triangular matrix.
- If g is not the representative polynomial of LD then by the DeMillo-Lipton/Schwartz/Zippel lemma [2, 25, 23], its evaluation at λ will fail with probability $1 - 2\frac{n-1}{|S|}$ (since $X^{n-1}(g - P_{LD})(X)$ is a polynomial of degree at most $2(n - 1)$).

For the complexity, the Prover computes L , in $O(mn^{\omega-1})$. Then $P_{LD}(X)$ requires one

pass over the coefficients of L , and finally $y = LD \begin{bmatrix} 1 \\ \lambda^{-1} \\ \vdots \\ \lambda^{1-n} \end{bmatrix}$. The communication cost is

D , $g(X)$, y all of size n , and λ . The Verifier cost is, $\mu(A) + \mu(B)$ to apply A and B , as well as $2n - 3$ to compute $[1 \ \lambda \ \dots \ \lambda^{n-1}]$ and their inverses, $n - 2$ to multiply by D , $2(n - 1)$ to evaluate g , and $2(n - 1)$ to compute the dotproduct $[1 \ \lambda \ \dots \ \lambda^{n-1}] \cdot y$. \square

6.3. Constant round certificates for the row and column rank profiles

Now we can combine the lower rank Certificate 8, with the constant-round Certificate 13 for triangular equivalence, as a replacement of Certificate 4, within the column rank profile Certificate 9, in order to get the constant-round Certificate 14 for column rank profile. It remains Prover efficient, linear in communication volume and Verifier time.

Prover		Verifier
	$A \in \mathbb{F}^{m \times n}$	
$\mathcal{J} = (c_1, \dots, c_r)$ CRP of A	$\xrightarrow{\mathcal{J}}$	Choose $S \subset \mathbb{F}$
β s.t. $A\beta = \nu$	<div style="border: 1px solid black; padding: 5px; width: fit-content; margin: auto;"> <div style="text-align: center;">Protocol 8</div> <div style="text-align: center;"> $\xleftarrow{\nu}$ $\xrightarrow{\beta}$ </div> </div>	$\alpha = \mathcal{E}_{m, \mathcal{J}}(\xleftarrow{\text{u.i.d.}} S^r)$ $\nu = A\alpha$ $\beta \stackrel{?}{=} \alpha$
$V = \text{Diag}(v_1, \dots, v_n)W$ (see (10))	\xleftarrow{v}	$v \xleftarrow{\text{u.i.d.}} S^n$
Γ upper tri. s.t. $A_{*, \mathcal{J}}\Gamma = AV$		$W = [\mathbf{1}_{i < c_{j+1}}] \in \mathbb{F}^{n \times r}$
	<div style="border: 1px solid black; padding: 5px; width: fit-content; margin: auto;"> <div style="text-align: center;">Protocol 13</div> <div style="text-align: center;">on $A_{*, \mathcal{J}}$ and $B = AV$</div> <div style="text-align: center;"> \xleftarrow{D} $\xrightarrow{g(X)}$ $\xleftarrow{\lambda}$ \xrightarrow{y} </div> </div>	
		$z = \text{Diag}(v_i)WD \begin{bmatrix} 1 \\ \lambda^{-1} \\ \vdots \\ \lambda^{1-r} \end{bmatrix}$ $z_{c_j} = z_{c_j} - y_j, j = 1..r$ $Az \stackrel{?}{=} 0$ $g(\lambda) \stackrel{?}{=} \begin{bmatrix} 1 & \lambda & \dots & \lambda^{r-1} \end{bmatrix} y$

Protocol 14: Constant-round certificate for the column rank profile

Corollary 3. For an $m \times n$ matrix of rank r , Certificate 14 is sound and perfectly complete. It requires 3 rounds, a volume of communication of $m + n + 5r + 1$ and less than $2\mu(A) + n + 9r$ operations for the Verifier.

7. Conclusion

A summary of our contributions is given in Table 3, to be compared with the state of the art in Table 2.

Algorithm	Rounds	Prover		Communication	Probabilistic		$ S $
		Determ.	Time		Verifier	Time	
[15] over [1]	No	No	$\tilde{O}(r^\omega + \mu(A))$	$\tilde{O}(r^2 + m + n)$	$\tilde{O}(r^2 + \mu(A))$	≥ 2	
RANK [3]	2	No	$O(n(\mu(A) + n))$ or $O(mnr^{\omega-2})$	$O(m+n)$	$2\mu(A) + \tilde{O}(m+n)$	$\Omega(\min\{m, n\} \log(mn))$	
[9]	2	Yes	$O(mnr^{\omega-2})$	$O(n+r)$	$O(\mu(A) + n)$	≥ 2	
CRP/ [15] over [24]	No	No	$\tilde{O}(r^\omega + m + n + \mu(A))$	$\tilde{O}(r^2 + m + n)$	$\tilde{O}(r^2 + m + n + \mu(A))$	$\Omega(\min\{m, n\} \log(mn))$	
RRP [15] over [14]	No	Yes	$O(mnr^{\omega-2})$	$\tilde{O}(mn)$	$\tilde{O}(mn)$	≥ 2	
RPM [15] over [8]	No	No	$\tilde{O}(r^\omega + m + n + \mu(A))$	$\tilde{O}(r^2 + m + n)$	$\tilde{O}(r^2 + m + n + \mu(A))$	$\Omega(\min\{m, n\} \log(mn))$	
[15] over [4]	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	2	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 ($\omega =$ matrix multiplication exponent > 2)

Algorithm	Rounds	Prover		Communication	Probabilistic		$ S $
		Determ.	Time		Verifier	Time	
CRP/RRP	§ 2.2	No	Yes	$O(mnr^{\omega-2})$	$O(r(m+n))$	$O(r(m+n)) + \mu(A)$	≥ 2
	§ 4.2	$O(n)$	Yes	$O(mnr^{\omega-2})$	$O(m+n)$	$2\mu(A) + O(m+n)$	≥ 2
	§ 6.3	3	Yes	$O(mnr^{\omega-2})$	$O(m+n)$	$2\mu(A) + O(m+n)$	$\geq 2n - 1$
RPM	§ 2.3	No	Yes	$O(mnr^{\omega-2})$	$O(r(m+n))$	$O(r(m+n)) + \mu(A)$	≥ 2
	§ 4.3	$O(n)$	Yes	$O(mnr^{\omega-2})$	$O(m+n)$	$4\mu(A) + O(m+n)$	$\Omega(n)$
DET	§ 4.1 & PLUQ	$O(n)$	Yes	$O(n^\omega)$	$O(n)$	$\mu(A) + O(n)$	$\Omega(n)$

Table 3: This paper's contributions ($\omega =$ matrix multiplication exponent > 2)

We have provided certificates that can save overall computational time for the Provers and an order of magnitude in terms of communication volume or number of rounds. Table 1 compares linear and quadratic communications, as well as sub-cubic (PLUQ, CHARPOLY) or quadratic matrix operations. These results show 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.

One key idea in our contribution is to certify the existence of a triangular matrix in an equivalence relation, by having an n round protocol where data dependency matches the triangular shape of the unknown matrix factor. This approach was successfully adapted to the certificate of generic rank profileness, where now two triangular unknown triangular factors are considered, in the LU decomposition.

Mulmuley’s Laurent’s polynomial representation of a matrix successfully replaces the former technique to certify triangular equivalence, and consequently row or column rank profiles, reducing the number of rounds from linear to constant. However, we were unable to adapt this technique for the certificate for generic rank profileness, and consequently for certifying a rank profile matrix.

The use of symmetric Gaussian elimination allowed us to achieve a more practical certificate for the signature of symmetric integer matrices. Even though it is based on LDLT certificates with linear communication modulo a prime, the diagonal of rational eigenvalues remains quadratic in size, and full precision was required to recover their sign. Designing a linear communication, Prover efficient protocol to certify the signature is the other major open problem which should be investigated in the continuation of this work.

Note added October 8, 2020: Added in Table 2 and 3 that $\omega > 2$.

- [1] Ho Yee Cheung, Tsz Chiu Kwok, and Lap Chi Lau. Fast Matrix Rank Algorithms and Applications. *Journal of the ACM*, 60(5):31:1–31:25, October 2013. ISSN 0004-5411. doi: [10.1145/2528404](https://doi.org/10.1145/2528404).
- [2] Richard A. DeMillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Inf. Process. Letters*, 7(4):193–195, June 1978. doi: [10.1016/0020-0190\(78\)90067-4](https://doi.org/10.1016/0020-0190(78)90067-4).
- [3] Jean-Guillaume Dumas and Erich Kaltofen. Essentially optimal interactive certificates in linear algebra. In Katsusuke Nabeshima, editor, *ISSAC’2014*, pages 146–153. ACM Press, New York, July 2014. doi: [10.1145/2608628.2608644](https://doi.org/10.1145/2608628.2608644).
- [4] Jean-Guillaume Dumas, Clément Pernet, and Ziad Sultan. Simultaneous computation of the row and column rank profiles. In Manuel Kauers, editor, *ISSAC’2013*, pages 181–188. ACM Press, New York, June 2013. doi: [10.1145/2465506.2465517](https://doi.org/10.1145/2465506.2465517).
- [5] Jean-Guillaume Dumas, Erich Kaltofen, Emmanuel Thomé, and Gilles Villard. Linear time interactive certificates for the minimal polynomial and the determinant of a sparse matrix. In Xiao-Shan Gao, editor, *ISSAC’2016*, pages 199–206. ACM Press, New York, July 2016. ISBN 978-1-4503-4380-0. doi: [10.1145/2930889.2930908](https://doi.org/10.1145/2930889.2930908).
- [6] Jean-Guillaume Dumas, Erich Kaltofen, Gilles Villard, and Lihoong Zhi. Polynomial time interactive proofs for linear algebra with exponential matrix dimensions and scalars given by polynomial time circuits. In Safey El Din [22], pages 125–132. doi: [10.1145/3087604.3087640](https://doi.org/10.1145/3087604.3087640).
- [7] Jean-Guillaume Dumas, David Lucas, and Clément Pernet. Certificates for triangular equivalence and rank profiles. In Safey El Din [22], pages 133–140. doi: [10.1145/3087604.3087609](https://doi.org/10.1145/3087604.3087609).
- [8] Jean-Guillaume Dumas, Clément Pernet, and Ziad Sultan. Fast computation of the rank profile matrix and the generalized Bruhat decomposition. *Journal of Symbolic Computation*, 83:187–210, November–December 2017. doi: [10.1016/j.jsc.2016.11.011](https://doi.org/10.1016/j.jsc.2016.11.011).
- [9] Wayne 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] 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 *LNCS*, pages 186–194. Springer-Verlag, 1987, 11–15 August 1986. URL <http://www.cs.rit.edu/~jjk8346/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](https://doi.org/10.1007/3-540-09526-8_5).
- [12] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In Cynthia Dwork, editor, *STOC'2008*, pages 113–122. ACM Press, May 2008. ISBN 978-1-60558-047-0. doi: [10.1145/1374376.1374396](https://doi.org/10.1145/1374376.1374396).
- [13] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: Interactive proofs for muggles. *J. ACM*, 62(4):27:1–27:64, 2015. doi: [10.1145/2699436](https://doi.org/10.1145/2699436).
- [14] C.-P. Jeannerod, C. Pernet, and A. Storjohann. Rank-profile revealing gaussian elimination and the CUP matrix decomposition. *Journal of Symbolic Computation*, 56:46–68, 2013. doi: [10.1016/j.jsc.2013.04.004](https://doi.org/10.1016/j.jsc.2013.04.004).
- [15] Erich L. Kaltofen, Michael Nehring, and B. David Saunders. Quadratic-time certificates in linear algebra. In Anton Leykin, editor, *ISSAC'2011*, pages 171–176. ACM Press, New York, June 2011. ISBN 978-1-4503-0675-1. doi: [10.1145/1993886.1993915](https://doi.org/10.1145/1993886.1993915).
- [16] The LINBOX group. FFLAS-FFPACK 2.3.1, November 2017. URL <http://linbox-team.github.io/fflas-ffpack>.
- [17] The LINBOX group. LINBOX 1.5.1, November 2017. URL <http://linalg.org>.
- [18] K Mulmuley. A Fast Parallel Algorithm to Compute the Rank of a Matrix over an Arbitrary Field. In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing, STOC '86*, pages 338–339, New York, NY, USA, 1986. ACM. ISBN 978-0-89791-193-1. doi: [10.1145/12130.12164](https://doi.org/10.1145/12130.12164).
- [19] Edward W. Ng, editor. *EUROSAM '79, International Symposium on Symbolic and Algebraic Computation, Marseille, France, June 1979, Proceedings*, volume 72 of *LNCS*, 1979. Springer. ISBN 3-540-09519-5. doi: [10.1007/3-540-09519-5](https://doi.org/10.1007/3-540-09519-5).
- [20] Clément Pernet and Arne Storjohann. Faster algorithms for the characteristic polynomial. In Christopher 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).
- [21] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. In Daniel Wichs and Yishay Mansour, editors, *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 49–62. ACM, 2016. ISBN 978-1-4503-4132-5. doi: [10.1145/2897518.2897652](https://doi.org/10.1145/2897518.2897652).
- [22] Mohab Safey El Din, editor. *ISSAC'2017, Proceedings of the 2017 ACM International Symposium on Symbolic and Algebraic Computation, Kaiserslautern, Deutschland, July 2017*. ACM Press, New York.
- [23] Jacob T. Schwartz. Probabilistic algorithms for verification of polynomial identities. In Ng [19], pages 200–215. ISBN 3-540-09519-5. doi: [10.1007/3-540-09519-5_72](https://doi.org/10.1007/3-540-09519-5_72).
- [24] Arne Storjohann and Shiyun Yang. A Relaxed Algorithm for Online Matrix Inversion. In Kazuhiro Yokoyama, editor, *ISSAC'2015*, pages 339–346. ACM Press, New York, July 2015. ISBN 978-1-4503-3435-8. doi: [10.1145/2755996.2756672](https://doi.org/10.1145/2755996.2756672).
- [25] Richard Zippel. Probabilistic algorithms for sparse polynomials. In Ng [19], pages 216–226. ISBN 3-540-09519-5. doi: [10.1007/3-540-09519-5_73](https://doi.org/10.1007/3-540-09519-5_73).