Linear Time Interactive Certificates for the Minimal Polynomial and the Determinant of a Sparse Matrix

Jean-Guillaume Dumas
Univ. Grenoble Alpes
Laboratoire LJK, CNRS, umr 5224
51, av. des Mathématiques,
F38041 Grenoble, France
jean-guillaume.dumas@imag.fr

Erich Kaltofen
Department of Mathematics
North Carolina State University
Raleigh, NC 27695-8205, USA
kaltofen@math.ncsu.edu
www.kaltofen.us

Emmanuel Thomé
Caramba – Inria Nancy
Inria, CNRS (Loria), Univ. Lorraine
615, rue du jardin botanique
54600 Villers-lès-Nancy, France
emmanuel.thome@inria.fr

Gilles Villard
CNRS, U. Lyon, Laboratoire LIP
(CNRS, ENSL, Inria, UCBL)
46, allée d'Italie, F69364 Lyon
Cedex 07, France
gilles.villard@ens-lyon.fr

ABSTRACT

Computational problem certificates are additional data structures for each output, which can be used by a—possibly randomized—verification algorithm that proves the correctness of each output. In this paper, we give an algorithm that computes a certificate for the minimal polynomial of sparse or structured $n \times n$ matrices over an abstract field, of sufficiently large cardinality, whose Monte Carlo verification complexity requires a single matrix-vector multiplication and a linear number of extra field operations. We also propose a novel preconditioner that ensures irreducibility of the characteristic polynomial of the generically preconditioned matrix. This preconditioner takes linear time to be applied and uses only two random entries. We then combine these two techniques to give algorithms that compute certificates for the determinant, and thus for the characteristic polynomial, whose Monte Carlo verification complexity is therefore also linear.

1. INTRODUCTION

We consider a square sparse or structured matrix $A \in \mathbb{F}^{n\times n}$, where $\mathbb{F}$ is an exact field. By sparse or structured we mean that multiplying a vector by A requires fewer operations than a dense matrix-vector multiplication. The arithmetic cost to apply A is denoted by $\mu(A)$ which thus satisfies $\mu(A) \leq n(2n-1)$ (n^2 multiplications and $n(n-1)$ additions).

For such sparse matrices, Wiedemann's algorithm, together with some preconditioning [3], provides means to compute the rank, the minimal polynomial or the characteristic polynomial of sparse matrices, via the computation of the minimal polynomial of its associated sequence projected at random values [29].

The novelty of this paper is to provide an algorithm that computes a certificate for the minimal polynomial of sparse matrices. The Monte Carlo verification complexity of our certificate is linear in the input size and this certificate is composed of the minimal polynomial itself, of three other polynomials and of a vector.

The verification procedure used throughout this paper is that of *interactive certificates* with the taxonomy of [9]. Indeed, we consider a *Prover*, nicknamed *Peggy*, who will perform a computation and provide additional data structures. We also consider a *Verifier*, nicknamed *Victor*, who will check the validity of the result, faster than by just recomputing it. By *certificates* for a problem that is given by input/output specifications, we mean, as in [19, 20], *an input-dependent data structure and an algorithm that computes from that input and its certificate the specified output, and that has lower computational complexity than any known algorithm that does the same when only receiving the input. Correctness of the data structure is not assumed but validated by the algorithm.* By *interactive certificate*, we mean interactive proofs, similar to $\sum$-protocols (as in [6]) were 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 Victor to be fully convinced. Such proof systems 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 probability are < 1, as the protocols can always be run several times. Some of our certificates will also be *perfectly complete*, that is a true statement is never rejected by the Verifier.

All our certificates can be simulated non-interactively by Fiat-Shamir heuristic [11]: uniformly sampled random values produced by Victor are replaced by hashes of the input and of previous messages in the protocol. Complexities are preserved, as producing cryptographically strong pseudo-random bits by a cryptographic hash function (e.g., like the extendable output functions of the SHA-3 family defined in [2, 23]), is linear in the size of both its input and output. More precisely, we will need $O(n\log(n))$ cryptographically strong random bits for the minimal polynomial certificate of Figure 3, and only $O(\log(n))$ for the determinant, Figures 4 and 5.

ISSAC '16, July 19 - 22, 2016, Waterloo, ON, Canada

ISBN 978-1-4503-4380-0/16/07 . . . $15.00

DOI: http://dx.doi.org/10.1145/2930889.2930908

There may be two main ways to design such certificates. On the one hand, efficient protocols can be designed for delegating computational tasks. In recent years, generic protocols have been designed for circuits with polylogarithmic depth [15, 26]. The resulting protocols are interactive and their cost for the Verifier is usually only roughly proportional to the input size. They however can produce a non negligible overhead for the Prover and are restricted to certain classes of circuits. Variants with an amortized cost for the Verifier can also be designed, see for instance [24], quite often using relatively costly homomorphic routines. We here however want the Verifier to run faster than the Prover, so we discard amortized models where the Verifier is allowed to do a large amount of precomputations, that can be amortized only if, say, the same matrix is repeatedly used [4, 14].

On the other hand, dedicated certificates (data structures and algorithms that are publicly verifiable a posteriori, without interaction) have also been developed, e.g., for dense exact linear algebra [13, 20, 12]. There the certificate constitutes a proof of correctness of a result, not of a computation, and can thus also stand an independent, computation error-correcting verification. 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. In the current paper we give new problem-specific certificate with fast verification and negligible overhead for the Prover.

In exact linear algebra, the simplest problem with an optimal certificate is the linear system solution, LINSOLVE: for a matrix A and a vector b, checking that x is actually a solution is done by one multiplication of x by A. The cost of this check is similar to that of just enumerating all the non-zero coefficients of A. Thus certifying a linear system is reduced to multiplying a matrix by a vector: LINSOLVE≺MATVECMULT. More precisely, by A≺B, we mean that there exists certificates for A that use certificates for B whose verification times are essentially similar: $\text{Verif}(A) = \text{Verif}(B)^{1+o(1)}$. In [9], two reductions have been made: first, that the rank can be certified via certificates for linear systems; second, that the characteristic polynomial can be certified via certificates for the determinant: CHARPOLY≺DET and RANK≺LINSOLVE. The verification procedure for the rank is essentially optimal, it requires two matrix-vector products and $n^{1+o(1)}$ additional operations; while the verification of the characteristic polynomial after verification of a determinant is simply linear. No reduction, however, was given for the determinant. We bridge this gap in this paper. Indeed, we show here that the computation of the minimal polynomial can be checked in linear time by a single matrix-vector multiplication: MINPOLY≺MATVECMULT. Then we use Wiedemann's reduction of the determinant to the minimal polynomial, DET≺MINPOLY, [29, 17], and propose a more efficient preconditioning for the same reduction.

This paper comes with a companion paper [8] that solve similar problems but with different techniques. We nonetheless believe that they are of independent interest, as shown by the following comparison of their salient differences: that

- The paper [8] gives certificates for the Wiedemann sequence, while we here directly certify its minimal polynomial;
- Complexities for the Verifier time and the extra communications are linear here while they are increased by $(\log n)^{\Omega(1)}$ in [8];
- The verification in [8] requires a black box for the transposed matrix;
- The certificates here are interactive while in [8], one of the certificates is, up to our knowledge, the only known non-interactive protocol for the determinant with Prover complexity $n^{1.5+o(1)}$.

The paper is organized as follows. We first present in Section 2 a new multiplicative preconditioner that allows to check the determinant as a quotient of minors. In Section 3 we define Wiedemann's projected Krylov sequence and propose a Monte Carlo certificate for the minimal polynomial of this sequence in Section 4. We apply this with random projections in Section 5, which provides a certificate for the minimal polynomial of the matrix. In Section 6, we see that with a diagonal preconditioning, we obtain another certificate for the determinant. In Section 7, we then combine this idea with the preconditioner of Section 2 to obtain a more efficient certificate for the determinant. This can be combined with the characteristic polynomial reduction of [9], in order to provide also a linear time certificate for the characteristic polynomial of sparse or structured matrices.

2. A SIMPLE INTERACTIVE CERTIFICATE FOR DETERMINANT

We first present a new multiplicative preconditioner that enables the Prover to compute our simple certificate, which is based on the characteristic matrix of the companion matrix of the polynomial $z^n + \sigma$ (see (1) below).

LEMMA 1.

$$\textit{Let } \Gamma(\sigma,\tau) = \begin{bmatrix} \tau & -1 & 0 & \dots & 0 \\ 0 & \tau & -1 & \ddots & \vdots \\ \vdots & 0 & \ddots & \ddots & 0 \\ 0 & & \ddots & \tau & -1 \\ \sigma & 0 & \dots & 0 & \tau \end{bmatrix} \in \mathbb{F}[\sigma,\tau]^{n\times n}, \qquad (1)$$

where τ and σ are variables. If $A \in \mathbb{F}^{n\times n}$ is non-singular, then $\det(\lambda I_n - A\,\Gamma(\sigma,\tau))$ is irreducible in $\mathbb{F}(\sigma,\tau)[\lambda]$.

PROOF. We observe that for $\Gamma(\sigma,\tau)$ in (1) we have $\det(\Gamma(\sigma,\tau)) = \tau^n + \sigma$, which is an irreducible polynomial in the bivariate polynomial domain $\mathbb{F}[\sigma,\tau]$. Next we consider the characteristic polynomial of $B(\sigma,\tau) = A\,\Gamma(\sigma,\tau)$, namely

$$\begin{aligned} c^{B(\sigma,\tau)}(\lambda) &= \det(\lambda I_n - B(\sigma,\tau)) \\ &= \lambda^n + c_{n-1}(\sigma,\tau)\lambda^{n-1} + \cdots \\ &\qquad + c_1(\sigma,\tau)\lambda \underbrace{\pm \overbrace{\det(A)}^{\in\mathbb{F} \text{ and } \neq 0} (\tau^n + \sigma)}_{c_0(\sigma,\tau)}. \qquad (2) \end{aligned}$$

We shall argue that the polynomial $c^{B(\sigma,\tau)}(\lambda)$ in (2) above is irreducible in $\mathbb{F}(\sigma,\tau)[\lambda]$. Because $B(\sigma,\tau)$ has linear forms in τ and σ as entries, $\deg_\tau(c_i) \leq n - i$ in (2). We now suppose that

$$g(\lambda,\sigma,\tau)\ h(\lambda,\sigma,\tau) = c^{B(\sigma,\tau)}(\lambda) \qquad (3)$$

is a non-trivial factorization in $\mathbb{F}[\lambda,\sigma,\tau]$. Then one of the degree-0 coefficients in λ of either g or h, which are $g(0,\sigma,\tau)$ or $h(0,\sigma,\tau)$, is a scalar multiple of $\tau^n + \sigma$ because $c_0(\sigma,\tau) = \pm\det(A)(\tau^n + \sigma) \neq 0$ (we assumed that A is non-singular) is irreducible in $\mathbb{F}[\sigma,\tau]$. Suppose $g(0,\sigma,\tau)$ is a scalar multiple of $\tau^n + \sigma$ and consequently $h(0,\sigma,\tau) \in \mathbb{F}$. Then $\deg_\tau(g) = n$ and therefore $\deg_\tau(h) = 0$, which means by (3) that h must divide all coefficients of the powers of τ in $c^{B(\sigma,\tau)}(\lambda)$ in (2). However, the term τ^n only occurs in $c_0(\sigma,\tau)$ and has coefficient $\pm\det(A)$ which is a non-zero field element, and therefore $h \in \mathbb{F}$, too. Note that $c^{B(\sigma,\tau)}(\lambda)$ has leading coefficient 1 in λ and therefore no non-trivial factor in $\mathbb{F}[\sigma,\tau]$

(one says it is primitive over $\mathbb{F}[\sigma, \tau]$). Then, by Gauss's Lemma, any factorization of $c^{B(\sigma,\tau)}(\lambda)$ in $\mathbb{F}(\sigma, \tau)[\lambda]$ can be rewritten as a factorization in $\mathbb{F}[\sigma, \tau][\lambda]$, and there is no no-trivial one, so the polynomial must be irreducible also in $\mathbb{F}(\sigma, \tau)[\lambda]$. □

The Prover must convince the Verifier that $\Delta = \det(A)$. Here is the 2 round interactive protocol. We allow a matrix A with $\det(A) = 0$ in which case the Prover may not be able to provide the same certificate for the determinant, but, with high probability, she cannot cheat the Verifier. The Prover can certify $\det(A) = 0$ by a vector $w \in \mathbb{F}^n$ with $w \neq 0^n$ and $Aw = 0^n$, which the Verifier can check. The matrix A is public.

Prover	*Communication*	*Verifier*
1. $B = A\,\Gamma(t, s)$ $t, s \in \mathbb{F}$ with $t^n + s \neq 0$,	$\xrightarrow{t,\, s}$	Checks $t^n + s \neq 0$,
2. $c^B(\lambda) = \det(\lambda I_n - B)$,	$\xrightarrow{c^B}$	
3. $C = [b_{i,j}]_{1\leq i,j\leq n-1}$, $c^C(\lambda) = \det(\lambda I_{n-1} - C)$, $t, s \in \mathbb{F}$ also with $\mathrm{GCD}(c^B, c^C) = 1$.	$\xrightarrow{c^C}$	Checks $\mathrm{GCD}(c^B(\lambda), c^C(\lambda)) = 1$
4.	$\xleftarrow{r_1}$	$r_1 \in S \subseteq \mathbb{F}$ random with $c^B(r_1) \neq 0$.
5. Computes w such that $(r_1 I_n - B)w = e_n = \begin{bmatrix}0\\ \vdots \\ 0\\ 1\end{bmatrix}$	$\xrightarrow{w}$	Checks $(r_1 I_n - B)w = e_n$, $w_n = c^C(r_1)/c^B(r_1)$.
6.		Returns $\det(A) = \dfrac{c^B(0)}{t^n + s}$.

Figure 1: A simple sparse determinant protocol

The interactive protocol is given in Figure 1. First we show it is complete, namely that if A is non-singular the Prover can choose $s, t \in \mathbb{F}$ such that $\mathrm{GCD}(c^B(\lambda), c^C(\lambda)) = 1$, provided $\mathbb{F}$ has sufficiently many elements. If A is singular, the Prover may not be able to do so, in which case she can communicate that $\det(A) = 0$ and a non-zero vector $w \in \mathbb{F}^n$, $w \neq 0^n$, with $Aw = 0^n$.

Let $c^{B(\sigma,\tau)}(\lambda) = \det(\lambda I_n - A\Gamma(\sigma, \tau))$, where $\Gamma(\sigma, \tau)$ is in (1), and let $C(\sigma, \tau) = [(B(\sigma, \tau))_{i,j}]_{1\leq i,j\leq n-1} \in \mathbb{F}[\sigma, \tau]^{(n-1)\times(n-1)}$. We have the non-zero Sylvester resultant

$$\rho(\sigma, \tau) = \mathrm{Res}_\lambda(c^{B(\sigma,\tau)}(\lambda), \det(\lambda I_{n-1} - C(\sigma, \tau))) \neq 0$$

because any non-trivial GCD in $\mathbb{F}(\sigma, \tau)[\lambda]$ would have to divide $c^{B(\sigma,\tau)}(\lambda)$, which is irreducible by Lemma 1. The Prover chooses t, s such that $\rho(t, s) \neq 0$, which for sufficiently large fields is possible by random selection. Note that $c^B(\lambda)$ for such choices of t, s may no longer be irreducible, but that the resultant of c^B and c^C is equal to $\rho(t, s) \neq 0$ (all polynomials have leading coefficient 1 in λ), hence $\mathrm{GCD}(c^B, c^C) = 1$.

The difficulty in certificates by interaction is the proof of soundness, that is, that the Verifier detects a dishonest Prover with high probability. Suppose that the Prover commits $H(\lambda) \neq c^B(\lambda)$ in place of $c^B(\lambda)$ and/or $h(\lambda) \neq c^C(\lambda)$ in place of $c^C(\lambda)$, with $\deg(H) = n$ and $\deg(h) = n - 1$ and both H and h with leading coefficient 1 in λ. The Prover might have chosen t, s such that $\mathrm{GCD}(c^B(\lambda), c^C(\lambda)) \neq 1$. Or she may have been unable to compute such t, s in the case when A is singular, and communicated the false H and h instead of presenting a linear column relation w as a certificate of singularity. In any case, because $h(\lambda)/H(\lambda)$ is a reduced fraction of polynomials with leading coefficient 1 and because $c^B(\lambda)$ and $c^C(\lambda)$ also have leading coefficient 1, we must have $h(\lambda)/H(\lambda) \neq c^C(\lambda)/c^B(\lambda)$, or equivalently $h(\lambda)c^B(\lambda) - H(\lambda)c^C(\lambda) \neq 0$.

The Verifier with probability

$$\geq 1 - \frac{\deg(c^B(hc^B - Hc^C))}{|S| - n} \geq 1 - \frac{3n - 2}{|S| - n}$$

chooses an r_1 such that and $c^B(r_1) \neq 0$ and $(hc^B - Hc^C)(r_1) \neq 0$, both of which $\Longrightarrow c^C(r_1)/c^B(r_1) \neq h(r_1)/H(r_1)$. The element r_1 satisfies $H(r_1) \neq 0$, hence the $-n$ in the denominator, roots of H are eliminated from selection set S. Note that $\deg(hc^B - Hc^C) \leq 2n - 2$ because the leading terms λ^{2n-1} of both products cancel.

Suppose now that r_1 is chosen with those properties. Then $\det(r_1 I_n - B) = c^B(r_1) \neq 0$, and by Cramer's rule

$$w_n = \frac{\det(r_1 I_{n-1} - C)}{\det(r_1 I_n - B)} = \frac{c^C(r_1)}{c^B(r_1)} \neq \frac{h(r_1)}{H(r_1)} \tag{4}$$

and the Verifier's last check fails. Therefore, if the last check also succeeds, with probability $\geq 1 - (3n - 2)/(|S| - n)$ we have $H = c^B$ and the Verifier has the correct determinant. Therefore the protocol is sound with high probability.

We do not fully analyze how fast the Prover could compute t, s, c^B, c^C, and w, in a modified protocol using additional preconditioners, as we will use the full Wiedemann technology for sparse and black box matrices in more efficient protocols below, which we have derived from Figure 1. The Verifier in Figure 1 checks a polynomial GCD, chooses a random field element, computes a matrix-times-vector product, and performs some arithmetic, which for a sparse A constitutes work more or less proportional to the input size. Again, below we will improve the Verifier complexity, for instance make the GCD verification $O(n)$. The GCD = 1 property of Step 3 will remain the fundamental ingredient for the soundness of our protocols.

3. THE WIEDEMANN SEQUENCE

Let $A \in \mathbb{F}^{n\times n}$ and $u, v \in \mathbb{F}^n$. The infinite sequence

$$(a_0, a_1, a_2, \ldots, a_i, \ldots) \text{ with } a_0{=}u^T v,\, a_i{=}u^T A^i v \text{ for } i{\geq}1, \tag{5}$$

is due to D. Wiedemann [29]. The sequence is linearly generated by the scalar minimal generating polynomial $f_v^{A,u}(\lambda) \in \mathbb{F}[\lambda]$, which is a factor of the minimal polynomial of the matrix A, the latter of which we denote by $f^A(\lambda)$. Both f^A and $f_v^{A,u}$ are defined as monic polynomials, that is, have leading coefficient equal to 1.

THEOREM 2. *Let $S \subseteq \mathbb{F}$ be of finite cardinality, which we denote by $|S| < \infty$, and let $u, v \in S^n$ be uniformly randomly sampled. Then the probability that $f_v^{A,u} = f^A$ is at least* $(1 - \deg(f^A)/|S|)^2 > 1 - 2n/|S|$ (cf. [16, Theorem 5]).

We now define the residue for the linear generator $f_v^{A,u}$.

DEFINITION 3. *Let $G_u^{A,v}(\lambda) = \sum_{i\geq0} a_i\lambda^{-1-i} \in \mathbb{F}[[\lambda^{-1}]]$ be the generating function of the Wiedemann sequence* (5). *Then we define the residue $\rho_u^{A,v}(\lambda) = f_u^{A,v}(\lambda)G_u^{A,v}(\lambda) \in \mathbb{F}[\lambda]$.*

LEMMA 4. *The residue $\rho_u^{A,v}$ in Definition 3 satisfies:*

$$\mathrm{GCD}(\rho_u^{A,v}(\lambda), f_u^{A,v}(\lambda)) = 1$$

PROOF. The field of quotients of the ring of power series in λ^{-1} is denoted by $\mathbb{F}((\lambda^{-1}))$, the ring of extended power series in λ^{-1}, whose elements can be represented as $\sum_{i\leq k} c_i\lambda^i$ for $c_i \in \mathbb{F}$ and $k \in \mathbb{Z}$. The residue $\rho_u^{A,v}(\lambda)$ is computed in $\mathbb{F}((\lambda^{-1}))$, but because

$f_u^{A,v}(\lambda)$ is a linear generator for $(a_i)_{i\geq 0}$, $\rho_u^{A,v}(\lambda)$ is a polynomial in $\mathbb{F}[\lambda]$ with $\deg(\rho_u^{A,v}) < \deg(f_u^{A,v})$. If the greatest common divisor $g(\lambda) = \mathrm{GCD}(\rho_u^{A,v}(\lambda), f_u^{A,v}(\lambda))$ is not trivial, then the equation $\rho_u^{A,v}/g = (f_u^{A,v}/g)G_u^{A,v}$ yields $f_u^{A,v}/g$ as a linear generator for $(a_i)_{i\geq 0}$ of lower degree than the degree of $f_u^{A,v}$, which violates the minimality of the linear generator $f_u^{A,v}$. □

4. A CERTIFICATE FOR THE LINEAR GENERATOR

We give the 2-rounds interactive protocol in Figure 2. The Prover must convince the Verifier that $f_u^{A,v}$ is indeed the Wiedemann generator for $(u^T A^i v)_{i\geq 0}$. The matrix A and vectors u, v are public.

Prover	*Communication*	*Verifier*
1. $H(\lambda) = f_u^{A,v}(\lambda)$, $h(\lambda) = \rho_u^{A,v}(\lambda)$.	$\xrightarrow{H,h}$	
2. $\phi, \psi {\in} \mathbb{F}[\lambda]$ with $\phi f_u^{A,v} {+} \psi \rho_u^{A,v} {=} 1$, $\deg(\phi) {\leq} \deg(\rho_u^{A,v}) {-} 1$, $\deg(\psi) {\leq} \deg(f_u^{A,v}) {-} 1$.	$\xrightarrow{\phi,\psi}$	
3.		Random $r_0 \in S \subseteq \mathbb{F}$. Checks $\mathrm{GCD}(H(\lambda), h(\lambda)){=}1$ by $\phi(r_0)H(r_0){+}\psi(r_0)h(r_0) \overset{?}{=\!=} 1.$
	$\xleftarrow{r_1}$	Random $r_1 \in S \subseteq \mathbb{F}$.
4. Computes w such that $(r_1 I_n - A)w{=}v$.	$\xrightarrow{w}$	Checks $(r_1 I_n - A)w \overset{?}{=\!=} v$ and $(u^T w)H(r_1) \overset{?}{=\!=} h(r_1)$. Returns $f_u^{A,v}(\lambda) = H(\lambda)$.

Figure 2: Certificate for $f_u^{A,v}$

In Step 4 Peggy may not be able to produce a vector w when $r_1 I_n - A$ is singular. However, she may instead convince Victor that the random choice of r_1 has led to a "failure". We investigate the case where the linear system is inconsistent more precisely. Let $f^{A,v}(\lambda)$ denote the minimal linear generator of the Krylov sequences of vectors $(A^i v)_{i\geq 0}$. We have that the minimal polynomial f^A of A is a multiple of $f^{A,v}$. Suppose now that $f^{A,v}(r_1) = 0$ and let $\lambda_2, \ldots, \lambda_m$ be the remaining roots of $f^{A,v}$ in the algebraic closure of $\mathbb{F}$. We obtain from $f^{A,v}(A)v = 0^n$ that $0^n = (\prod_{j=2}^m (A - \lambda_j I_n))(A - r_1 I_n)w = (\prod_{j=2}^m (A - \lambda_j I_n))(-v)$, in violation that $f^{A,v}(A)v$ constitutes the first linear dependence of the Krylov vectors $(A^i v)_{i\geq 0}$. Conversely, if $f^{A,v}(r_1) \neq 0$ then the system $(r_1 I - A)w = v$ is consistent with $w = (1/f^{A,v}(r_1))\, \vec{\rho}^{A,v}(r_1)$ (see (6) multiplied by $\lambda I_n - A$ in the soundness proof below). We just proved that the linear system $(r_1 I_n - A)$ is inconsistent with v if and only if $f^{A,v}(r_1) = 0$ (hence in particular $r_1 I_n - A$ is singular). Peggy could provide a non-zero vector in the right nullspace of $r_1 I_n - A$ as a proof that Victor has communicated a bad r_1. Since one can have $f_u^{A,v}(r_1) \neq 0$ when $f^{A,v}(r_1) = 0$ Victor cannot test the choice of r_1 before sending it even if the communicated $f_u^{A,v}$ is correct.

Given the above, the protocol is *perfectly complete*: if the values $(f_u^{A,v}, \rho_u^{A,v})$ and the system solution $w = (1/f^{A,v}(r_1))\, \vec{\rho}^{A,v}(r_1)$ communicated by the Prover are correct, the Verifier always accepts $f_u^{A,v}$. In case that the system $(r_1 I_n - A)w = v$ is inconsistent, the Prover could communicate a Farkas certificate of inconsistency $\bar{w}$ with $\bar{w}^T(r_1 I_n - A) = 0$ and $\bar{w}^T v \neq 0$. When a Farkas certificate of inconsistency is sent, the Verifier then accepts the "correct" output that the choice of r_1 has led to "failure."

In Figure 2 we allow inconsistent systems to remain uncertified for two reasons: 1. A Farkas certificate requires additional work for the Prover, possibly needing a transposed matrix times vector procedure for a black box matrix A. 2. Monte Carlo algorithms are always fast and do not fail. Their output is correct with probability $\geq$ a given bound. Our certificates indeed have Monte Carlo randomized verification algorithms, that is, as we will prove below, the interactive protocol does not fail via system inconsistency in Step 4 and, when all checks succeed, the first pair of polynomials communicated is $f_u^{A,v}$ and $\rho_u^{A,v}$ with probability $\geq (1 - (2n-2)/|S|) \times (1 - (3n-1)/|S|)$.

Proof of soundness: Now suppose that the Prover commits $H(\lambda)$ in place of $f_u^{A,v}(\lambda)$ and/or $h(\lambda)$ in place of $\rho_u^{A,v}(\lambda)$, with $n \geq \deg(H) > \deg(h)$ and $h/H \neq \rho_u^{A,v}/f_u^{A,v}$. The Verifier with probability

$$\begin{aligned} &\geq 1 - (\deg(\phi H + \psi h))/|S| \\ &\geq 1 - (\max\{\deg(H) + \deg(\phi), \deg(h) + \deg(\psi)\})/|S| \\ &\geq 1 - (2n-2)/|S| \end{aligned}$$

exposes that $(\phi H + \psi h)(r_0) \neq 1$. Suppose now that the Verifier's check succeeds and that actually $\mathrm{GCD}(H, h) = 1$. The Verifier with probability

$$\begin{aligned} &\geq 1 - \left(\deg(f^A(H\rho_u^{A,v} - h f_u^{A,v}))\right)/|S| \\ &\geq 1 - (3n-1)/|S| \end{aligned}$$

chooses an r_1 such that $f^A(r_1) \neq 0$ (that is, $r_1 I_n - A$ is non-singular) and $(H\rho_u^{A,v} - h f_u^{A,v})(r_1) \neq 0$. The latter rewrites as $h(r_1) \neq H(r_1)\rho_u^{A,v}(r_1)/f_u^{A,v}(r_1)$. Note that $f^A(r_1) \neq 0 \Rightarrow f_u^{A,v}(r_1) \neq 0$. Suppose now that r_1 is chosen with those properties. As in the proof of Lemma 4 we can define

$$G^{A,v}(\lambda) = \sum_{i\geq 0} \lambda^{-(i+1)}(A^i v) \in \mathbb{F}^n[[\lambda^{-1}]],$$

$$\vec{\rho}^{A,v}(\lambda) = f^{A,v}(\lambda)\, G^{A,v}(\lambda) \in \mathbb{F}^n[\lambda].$$

We have $(\lambda I_n - A)\, G^{A,v}(\lambda) = v$ and

$$(\lambda I_n - A)^{-1} v = \frac{1}{f^{A,v}(\lambda)}\, \vec{\rho}^{A,v}(\lambda). \qquad (6)$$

Furthermore, $(u^T \vec{\rho}^{A,v})/f^{A,v} = \rho_u^{A,v}/f_u^{A,v}$, where the left-hand side is not necessarily a reduced fraction. Thus, $u^T w\, H(r_1) = u^T(r_1 I_n - A)^{-1} v\, H(r_1) = (\rho_u^{A,v}/f_u^{A,v})(r_1) \cdot H(r_1) \neq h(r_1)$, and the Verifier's last check fails; one can compare with (4), which has $u = v = e_n$ and $f_u^{B,v} = c^B$. Therefore, if the last check also succeeds, with probability $\geq (1 - (2n-2)/|S|)(1 - (3n-1)/|S|)$ we have $\mathrm{GCD}(H, h) = 1$ and $h/H = \rho_u^{A,v}/f_u^{A,v}$.

About the certificate complexity: Excluding the input matrix A and the output minimal polynomial $f_u^{A,v}$, the certificate in Figure 2 comprises $\rho_u^{A,v}$, monic of degree strictly less than n; then ϕ, ψ, of degree respectively less than $n-2$ and $n-1$; and finally w, a vector of $\mathbb{F}^n$. The extra communications for the certificate are thus less than $4n$ field elements. For the time complexity, evaluating ϕ, f, ψ, ρ, f and ρ requires less than $6(2n)$ field operations, Aw is one matrix-vector, of cost $\mu(A)$, and checking $(r_1 I_n - A)w$ (either for a solution, or as a non-zero vector in the right nullspace) requires an additional $2n$ operations, like $u^T w$. The multiplication of the vector by $f_u^{A,v}(r_1)$ requires finally n more multiplications in the field. We have thus proven Theorem 5 thereafter.

THEOREM 5. *If the size of the field is $\geq 3n$, the protocol in Figure 2 is sound and complete. The associated certificate requires less than $4n$ extra field elements and is verifiable in less than $\mu(A) + 17n$ field operations.*

Note that, with slightly larger fields, one can further reduce the complexity for the Verifier.

COROLLARY 6. *If the size of the field is $\geq 5n - 2$, there exists a sound and perfectly complete protocol for certifying $f_u^{A,v}$, whose associated certificate requires less than $4n$ extra field elements and is verifiable in less than $\mu(A) + 13n$ field operations.*

PROOF. For the complexity improvement, it suffices to chose $r_0 = r_1$. Then the last evaluations of $f_u^{A,v}(r_1)$ and $\rho_u^{A,v}(r_1)$ are already computed by the GCD check. Now r_0 must be such that $f^A(H\rho_u^{A,v} - hf_u^{A,v})(\phi H + \psi h - 1)\,(r_0) \neq 0$, but the latter is of degree $\deg(f^A) + \max\{\deg(H) + \deg(\phi), \deg(h) + \deg(\psi)\} + \max\{\deg(H) + \deg(\rho_u^{A,v}), \deg(h) + \deg(f_u^{A,v})\} \leq 3n - 1 + 2n - 2$. Furthermore, the analysis in the beginning of this section shows that the protocol of Figure 2 can be extended to be *perfectly* complete. □

REMARK 7 *(*CERTIFICATE OVER SMALL FIELDS*).* For the certificates of this paper to be sound, the underlying field cannot be too small: $\Omega(n)$ for the minimal polynomial certificates, $\Omega(n^2)$ for the following determinant or characteristic polynomial certificates. If the field is smaller, then a classical technique is to embed it in an extension of adequate size. Arithmetic operations in this extension field $L \supset \mathbb{F}$ of degree $O(\log n)$ will then cost $\log(n)^{1+o(1)}$ operations in the base field, but with $s \geq \log_2(1/\epsilon)$ matrix A times vectors $\in \mathbb{F}^n$ operations, that is over the base field, an incorrect w is exposed with probability $\geq 1 - 1/2^s \geq 1 - \epsilon$ for any constant $\epsilon > 0$.

We give the idea for $\mathbb{F}{=}\mathbb{Z}_2$ and $L{=}\mathbb{Z}_2[\theta]/(\chi(\theta))$ where χ is an irreducible polynomial in $\mathbb{Z}_2[\theta]$ with $\deg(\chi){=}k{+}1{=}O(\log n)$. The Prover returns $w{=}w^{[0]}{+}\theta w^{[1]}{+}\cdots{+}\theta^k w^{[k]}$ with $w^{[\kappa]}{\in}\mathbb{Z}_2^n$. For $v^{[0]}{+}\theta v^{[1]}{+}\cdots{+}\theta^k v^{[k]} = v{-}r_1 w$ with $v^{[\kappa]} \in \mathbb{Z}_2^n$ we must have $-Aw^{[\kappa]}{=}v^{[\kappa]}$ for all κ. Suppose that the Prover has sent a w that violates j of the equations, namely, $-Aw^{[\ell_i]}{=}y^{[\ell_i]}{\neq}v^{[\ell_i]}$ for $1 \leq i \leq j$. Then $|\{(c_1, \ldots, c_j){\in}\mathbb{Z}_2^j \mid \sum_{i=1}^j c_i y^{[\ell_i]}{\neq}\sum_{i=1}^j c_i v^{[\ell_i]}\}| \geq 2^{j-1}$. The inequality is immediate for $j{=}1, 2$. For $j \geq 3$ one gets for each $\sum_{i=1}^{j-2} c_i y^{[\ell_i]}{=}\sum_{i=1}^{j-2} c_i v^{[\ell_i]}$ the two unequal linear combinations for $c_{j-1}{=}1, c_j{=}0$ and $c_{j-1}{=}0, c_j{=}1$; one also gets for each $\sum_{i=1}^{j-2} c_i y^{[\ell_i]}{\neq}\sum_{i=1}^{j-2} c_i v^{[\ell_i]}$ the unequal linear combination for $c_{j-1}{=}c_j{=}0$; a second unequal combination is either for $c_{j-1} = 1$, $c_j = 0$ or $c_{j-1}{=}c_j{=}1$, because not both can be equal at the same time. Therefore $|\{(c_0, \ldots, c_k){\in}\mathbb{Z}_2^{k+1} \mid -A(\sum_{i=0}^k c_i w^{[i]}) \neq (\sum_{i=0}^k c_i v^{[i]})\ \}| \geq 2^k$ which shows that checking a random linear combination over $\mathbb{Z}_2$ exposes an incorrect w with probability $\geq 1/2$. The cost for the Verifier is thus $O(\mu(A)) + n^{1+o(1)}$ operations in $\mathbb{Z}_2$.

Complexity for the Prover: Theorem 5 places no requirement on how the Prover computes the required commitments. Coppersmith's block Wiedemann method [5], for instance, can be used.

THEOREM 8. *If the size of the field is $> 3n$ then the Prover can produce the certificate in Figure 2 in no more than $(1 + o(1))n$ multiplications of A times a vector in $\mathbb{F}^n$ and $n^{2+o(1)}$ additional operations in $\mathbb{F}$, using $n^{1+o(1)}$ auxiliary storage for elements in $\mathbb{F}$.*

PROOF. The Berlekamp-Massey algorithm can be applied on the sequence $(a_i)_{i \geq 0} = (u^T A^i v)_{i \geq 0}$ to recover $f_u^{A,v}$. However that requires (at most) $2n$ terms to be computed [29], and hence more multiplications of A times a vector than requested in the statement. Following [16, Sec. 7], the number of applications of A can be minimized using Coppersmith's block Wiedemann algorithm [5] with a special choice of parameters. We consider integers p and q (the blocking factors) such that $q = o(p)$, $p = n^{o(1)}$, and $1/q = o(1)$. For example, we may choose $p = (\log n)^2$ and $q = \log n$.

For the block Wiedemann algorithm we first take random $\mathcal{U}_0 \in \mathbb{F}^{n \times (p-1)}$ and $\mathcal{V}_0 \in \mathbb{F}^{n \times (q-1)}$ with entries sampled from $S \subseteq \mathbb{F}$, and construct $\mathcal{U} = [u, \mathcal{U}_0] \in \mathbb{F}^{n \times p}$ and $\mathcal{V} = [v, \mathcal{V}_0] \in \mathbb{F}^{n \times q}$. With $d_l = \lceil n/(p-1) \rceil$ and $d_r = \lceil n/(q-1) \rceil$, the rank of the block Hankel matrix $\mathrm{Hk}_{d_l, d_r}(A, \mathcal{U}_0, \mathcal{V}_0)$ in [18, (2.4)] is maximal with high probability, and equal to the dimension of the Krylov space $\mathcal{K}_{A, \mathcal{V}_0}$ (using $q \leq p$). Since the degree of the q-th (resp. the p-th) highest degree invariant factor of $\lambda - A$ has degree less than d_r (resp. d_l), less than d_r (resp. d_l) first Krylov iterates $(A^i v)_{i \geq 0}$ (resp. $(u^T A^i)_{i \geq 0}$) suffice for completing any basis of $\mathcal{K}_{A, \mathcal{V}_0}$ into a basis of $\mathcal{K}_{A, \mathcal{V}}$. (resp. $\mathcal{K}_{A, \mathcal{U}_0}$ and $\mathcal{K}_{A, \mathcal{U}}$). Therefore the rank of $\mathrm{Hk}_{d_l, d_r}(A, \mathcal{U}, \mathcal{V})$ must be maximal and equal to $\dim \mathcal{K}_{A, \mathcal{V}}$.

The determinant of a submatrix of $\mathrm{Hk}_{d_l, d_r}(A, \mathcal{U}, \mathcal{V})$ with maximal rank has degree at most $2n$ in the entries of $\mathcal{U}_0$ and $\mathcal{V}_0$. Using arguments analogous to those in [27, sec. 9.2] or [18, Sec. 3.2], by the DeMillo-Lipton/Schwartz/Zippel Lemma [7, 30, 25], we obtain that:

$$\text{Prob}\left(\text{rank}\, \mathrm{Hk}_{d_l, d_r}(A, \mathcal{U}, \mathcal{V}) = \dim \mathcal{K}_{A, \mathcal{V}}\right) \geq 1 - 2n/|S|. \quad (7)$$

If S contains sufficiently many elements, it follows [27, 18] that it is sufficient to compute the sequence of $p \times q$ matrices $\mathcal{A}_i = \mathcal{U}^T A^i \mathcal{V}$ up to $L = d_l + d_r$ terms, which can be achieved with $qL = n(1 + o(1))$ multiplications of A times a vector, and $n^{2+o(1)}$ additional operations in $\mathbb{F}$. Let

$$\mathcal{G}_{\mathcal{U}}^{A, \mathcal{V}}(\lambda) = \sum_{i \geq 0} \mathcal{A}_i \lambda^{-1-i} = \sum_{i \geq 0} \mathcal{U}^T A^i \mathcal{V} \lambda^{-1-i} \in \mathbb{F}^{p \times q}[[\lambda^{-1}]]$$

be the generating function for the block Wiedemann sequence. One can then compute a rational fraction description of $\mathcal{G}_{\mathcal{U}}^{A, \mathcal{V}}$, in the form of a pair of matrices $\mathcal{F}(\lambda) \in \mathbb{F}^{q \times q}[\lambda]$ and $\mathcal{R}(\lambda) \in \mathbb{F}^{p \times q}[\lambda]$ with degree $d_r = o(n)$, so that $\mathcal{F}$ is a minimal matrix generator for $(\mathcal{A}_i)_{i \geq 0}$, and

$$\mathcal{G}_{\mathcal{U}}^{A, \mathcal{V}}(\lambda) \mathcal{F}(\lambda) = \mathcal{R}(\lambda) \in \mathbb{F}^{p \times q}[\lambda].$$

The matrices $\mathcal{F}$ and $\mathcal{R}$ can be computed in time soft-linear in n [1], and since the dimensions are in $n^{o(1)}$, the matrix fraction $\mathcal{R}(\lambda) \times \mathcal{F}(\lambda)^{-1}$ is also obtained in soft-linear time. Any of the entries of this matrix fraction is a rational fraction description of a sequence corresponding to two chosen vectors among the blocking vectors. In particular, from the entry (1,1), Peggy obtains the rational fraction description $\rho_u^{A,v} / f_u^{A,v}$ of $(a_i)_{i \geq 0}$. The computation of ϕ and ψ follows via an extended Euclidean algorithm applied on $f_u^{A,v}$ and $\rho_u^{A,v}$.

A small variation of the above scheme solves the inhomogenous linear system in Step 4. Let $B = r_1 I_n - A$. With high probability, Victor chooses a r_1 such that $f^A(r_1) \neq 0$, hence B is non-singular. Together with (7) this gives the requirement $3n{+}1$ for the field size.

We have that $\mathcal{F}' = \mathcal{F}(r_1 - \lambda)$ is a minimal matrix generator for the sequence of $p \times q$ matrices $(\mathcal{U}^T B^i \mathcal{V})_{i \geq 0}$ (computing $\mathcal{F}'$ is a soft-linear operation). Since B is non-singular we have $\det \mathcal{F}'(0) \neq 0$. With column operations on $\mathcal{F}'$, we can arrange so that one column becomes $[f_0, 0, \ldots, 0]^T$ with $f_0 \neq 0$. $\mathcal{V}$ times the latter being $f_0.v$, following the inhomogeneous case in [5, Sec. 8], this suffices to solve the linear system $Bw = v$, using $d_r = o(n)$ more matrix-vector multiplications and $n^{2+o(1)}$ operations in $\mathbb{F}$. □

5. A CERTIFICATE FOR THE MINIMAL POLYNOMIAL

With random vectors u and v, we address the certification of the minimal polynomial f^A. Indeed, using Wiedemann's study [29, Sec. VI] or the alternative approach in [17, 16], we know that for random u and v, $f_u^{A,v}$ is equal to f^A with high probability (see Theorem 2). This is shown in Figure 3.

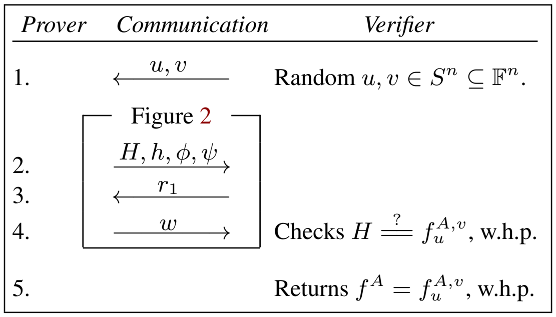

Prover	Communication	Verifier
1.	$\xleftarrow{u,v}$	Random $u, v \in S^n \subseteq \mathbb{F}^n$.
	Figure 2	
2.	$\xrightarrow{H,h,\phi,\psi}$	
3.	$\xleftarrow{r_1}$	
4.	$\xrightarrow{w}$	Checks $H \overset{?}{=\!=} f_u^{A,v}$, w.h.p.
5.		Returns $f^A = f_u^{A,v}$, w.h.p.

Figure 3: Certificate for f^A with random projections

PROPOSITION 9. *The protocol in Figure 3 is sound and complete.*

PROOF. First, the protocol is complete. Indeed, the result is correct means that $H = f^A$. Thus, if $f^A = f_u^{A,v}$, then completeness is guaranteed by the completeness in Theorem 5. Otherwise, H is a proper multiple of $f_u^{A,v}$ and $H = f^A \neq f_u^{A,v}$. If u, v are randomly chosen by Victor, this happens only with low probability, thanks to Theorem 2.

Second, for the soundness, if the result is incorrect, then $H \neq f^A$. If also $H \neq f_u^{A,v}$, then similarly, H will make the certificate of Figure 2 fail with high probability, by the soundness in Theorem 5. Otherwise, $H = f_u^{A,v} \neq f^A$, but if u, v are randomly chosen by Victor, this happens only with low probability, thanks to Theorem 2. □

COROLLARY 10. *If the size of the field is $\geq 5n-2$, there exists a sound and perfectly complete protocol for certifying f^A, whose associated certificate requires less than $8n$ extra field elements and is verifiable in less than $2\mu(A) + 26n$ field operations.*

PROOF. We add some work for the Prover, and double the certificate. First Peggy has to detect that the projections given by Victor reveal only a proper factor of the minimal polynomial. Second she needs to prove to Victor that he was wrong. For this, Peggy:

- Computes the minimal polynomial f^A of the matrix and the minimal polynomial $f_u^{A,v}$.
- If $f^A \neq f_u^{A,v}$, Peggy searches for projections $(\hat{u}, \hat{v})$, such that $\deg(f_{\hat{u}}^{A,\hat{v}}) > \deg(f_u^{A,v})$.
- Peggy then starts two certificates of Figure 2, one for (u, v), one for $(\hat{u}, \hat{v})$.

In case of success of the latter two certificates, Victor is convinced that both $f_u^{A,v}$ and $f_{\hat{u}}^{A,\hat{v}}$ are correct. But as the latter polynomial has a degree strictly larger than the former, he is also convinced that his projections can not reveal f^A. Complexities are given by applying Corollary 6 twice. □

REMARK 11 *(*CERTIFICATE FOR THE RANK*).* If A is non-singular, the certificate of [9, Figure 2] can be used to certify non-singularity: for any random vector b proposed by Victor, Peggy can solve the system $Aw = b$ and return w as a certificate. Now, if A is singular, a similar idea as for the minimal polynomial can be used to certify the rank: precondition the matrix A into a modified matrix B whose minimal polynomial is $f(x)x$ and characteristic polynomial is $f(x)x^k$, where $f(0) \neq 0$. As a consequence $\text{rank}(A) = n - k$. For instance, if A is symmetric, such a PRECONDCYCNIL preconditioner can be a non-singular diagonal matrix D if the field is sufficiently large [3, Theorem 4.3]. Otherwise, $A^T D_2 A$ is symmetric, with D_2 another diagonal matrix. Then the minimal polynomial certificate can be applied to $B = DA^T D_2 AD$ [10]. Comparing with the certificate for the rank in [9, Corollary 3], this new certificate saves a logarithmic factor in the verification time, but requires the field to be larger (from $\Omega(n)$ to $\Omega(n^2)$).

6. CERTIFICATE FOR THE DETERMINANT WITH DIAGONAL PRECONDITIONING

First of all, if A is singular, Peggy may not be able to produce the desired certificate. In which case she can communicate that $\det(A) = 0$ and produce a non-zero vector in the kernel: $w \in \mathbb{F}^n$, $w \neq 0^n$, with $Aw = 0^n$.

We thus suppose in the following that A is non-singular.

The idea of [29, Theorem 2] is to precondition the initial invertible matrix A into a modified matrix B whose characteristic polynomial is square-free, and whose determinant is an easily computable modification of that of A. For instance, such a PRECONDCYC preconditioner can be a non-singular diagonal matrix D if the field is sufficiently large [3, Theorem 4.2]:

$$\text{Prob}\left(\deg\left(f^{DA}\right) = n\right) \geq 1 - \frac{n(n-1)}{2|S|} \tag{8}$$

To certify the determinant, it is thus sufficient for Peggy to chose a non-singular diagonal matrix D and two vectors u, v such that $\deg(f_u^{DA,v}) = n$ and then to use the minimal polynomial certificate for $f_u^{DA,v}$, as shown on Figure 4.

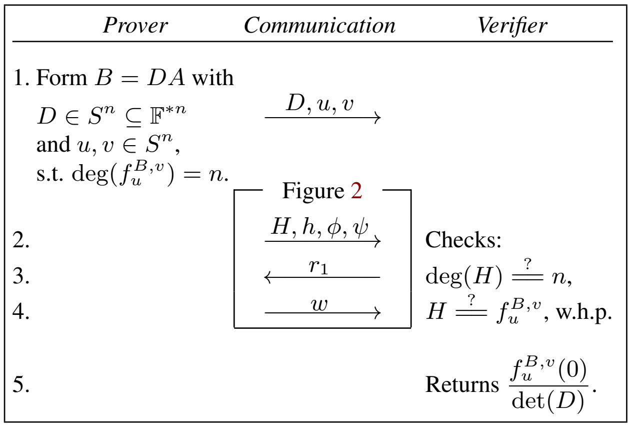

Prover	Communication	Verifier
1. Form $B = DA$ with $D \in S^n \subseteq \mathbb{F}^{*n}$ and $u, v \in S^n$, s.t. $\deg(f_u^{B,v}) = n$.	$\xrightarrow{D,u,v}$	
	Figure 2	
2.	$\xrightarrow{H,h,\phi,\psi}$	Checks:
3.	$\xleftarrow{r_1}$	$\deg(H) \overset{?}{=\!=} n$,
4.	$\xrightarrow{w}$	$H \overset{?}{=\!=} f_u^{B,v}$, w.h.p.
5.		Returns $\frac{f_u^{B,v}(0)}{\det(D)}$.

Figure 4: Determinant certificate with Diagonal preconditioning for a non-singular matrix

REMARK 12. Note that in the minimal polynomial sub-routine of Figure 4, Peggy can actually choose D, u, v, since the check on the degree of $f_u^{B,v}$ prevents bad choices for D, u, v. Victor could also select them himself, the overall volume of communications would be unchanged, but he would have to perform more work, namely selecting $3n + 2$ random elements instead of just 1.

THEOREM 13. *If the size of the field is $\geq \max\{\frac{1}{2}n^2 - \frac{1}{2}n, 5n - 2\}$, the protocol for the determinant of a non-singular matrix in Figure 4 is sound and complete. The associated certificate requires less than $8n$ extra field elements and is verifiable in less than*

$\mu(A) + 15n$ field operations. If the size of the field is $\geq n^2 + n + 5$, with high probability the Prover can produce it with no more than 2 minimal polynomial computations, and 1 system solving.

PROOF. First, if the Prover is honest, H is correctly checked to be $f_u^{B,v}$, as the minimal polynomial is complete. Then, by definition, $f_u^{B,v}$ is a factor of the characteristic polynomial of B. But if its degree is n, then it is the characteristic polynomial. Therefore its unit coefficient is the determinant of B and the certificate of Figure 4 is complete.

Second, for the soundness. If $H \neq f_u^{B,v}$, then the minimal polynomial certificate will most probably fail, by the soundness of the minimal polynomial certificate, as $|S| \geq (5n - 2)$, from Corollary 6. Otherwise, the degree check enforces that $f_u^{B,v}$ is the characteristic polynomial.

Now, for the complexity, with respect to the minimal polynomial certificate and Theorem 5, this certificate requires an extra diagonal matrix D. As $u, v, f_u^{B,v}$ are not input/output anymore, the extra communications grow from $4n$ field elements to $8n$. The verification procedure is similar except that verifying a linear system solution with B requires $n + \mu(A)$ operations and that $det(D)$ has to be computed, hence the supplementary $2n$ field operations. Finally, for the Prover, in order to find suitable vectors and diagonal matrices, Peggy can select them randomly in S^n and try $f_u^{DA,v}$ until $\deg(f_u^{DA,v}) = n$. It will succeed with the joint probabilities of Theorem 2 and Equation (8). For $n \geq 2$, as soon as $|S| \geq n^2 + n + 5$, the lower bound of the probability of success $(1 - n(n-1)/(2|S|))(1 - 2n/|S|)$ is higher than $1/2$ and the expected number of trials for the Prover is less than 2. □

Note that the certificate (as well as that of next section) can be made perfectly complete, by using the perfectly complete certificate for the minimal polynomial of Corollary 10.

7. DETERMINANT WITH GAMMA PRECONDITIONING

In order to compute the determinant via a minimal polynomial, the diagonal preconditioning ensures that it is equal to the characteristic polynomial, as the latter is square-free with factors of degree 1. With the preconditioning of Section 2, we can differently ensure that the characteristic polynomial is irreducible. It has the same effect, that it equals the minimal polynomial, but it also enforces that it has no smaller degree factors. Therefore, for any non-zero u and v, as a non-singular matrix is non-zero, the only possibility for $f_u^{A,v}$ is to be of degree n. Thus, we can give in Figure 5 an improved certificate. It chooses pre-determined vectors u and v to be as simple as possible: the canonical vector $e_1 = [1, 0, \ldots, 0]^T$.

THEOREM 14. *If the size of the field is $\geq \max\{n^2 - n, 5n - 2\}$, the protocol for the determinant of a non-singular matrix in Figure 5 is sound and complete. The associated certificate requires less than $5n$ extra field elements and is verifiable in less than $\mu(A) + 13n + o(n)$ field operations. If the size of the field is $\geq 2n^2 - 2n$, with high probability the Prover can produce it with no more than 2 minimal polynomial computations, and 1 system solving.*

PROOF. Completeness is given by the same argument as for Theorem 13. It is similar for the soundness, provided that $H_n(s,t) \neq 0$ implies that $\deg(f_{e_1}^{B,e_1}) = n$.

Therefore, let $C^{B(\sigma,\tau)}(\lambda) = \det(\lambda I_n - A\,\Gamma(\sigma,\tau))$, where $\Gamma(\sigma,\tau)$ is in (1). Then, for $e_1 = [1, 0, \ldots, 0]^T$ the first canonical vector, let $G_1^{B(\sigma,\tau)}(\lambda) = e_1^T(\lambda I_n - A)^{-1}e_1$, and let $\rho_1^{(\sigma,\tau)}(\lambda) = C^{B(\sigma,\tau)}(\lambda)G_1^{B(\sigma,\tau)}(\lambda)$.

	Prover	*Communication*	*Verifier*
1.	Form $B = A\,\Gamma(s,t)$ with $s, t \in \mathbb{F}$ s.t. $\Gamma(s,t)$ non-singular ($\Leftrightarrow t^n + s \neq 0$) and $H_n(s,t) \neq 0$. (see (1) and (9))	$\xrightarrow{s,t}$	Checks $t^n + s \neq 0$.
		Figure 2	
2.		$\xrightarrow{H,h,\phi,\psi}$	Checks:
3.		$\xleftarrow{r_1}$	$\deg(H) \overset{?}{=\!=} n$,
4.		$\xrightarrow{w}$	$H \overset{?}{=\!=} f_{e_1}^{B,e_1}$, w.h.p.
5.			Returns $\dfrac{f_{e_1}^{B,e_1}(0)}{t^n + s}$.

Figure 5: Determinant certificate with Gamma preconditioning for a non-singular matrix

As $C^{B(\sigma,\tau)}$ is monic and $e_1^T e_1 = 1$, then $G_1^{B(\sigma,\tau)}$ and $\rho_1^{(\sigma,\tau)}$ are not identically zero. Further, $C^{B(\sigma,\tau)}$ is irreducible by Lemma 1 but the minimal polynomial for the projected sequence must divide C^B. Therefore the sequence has C^B for minimal polynomial. Now, this sequence is denoted by $(s_i(\sigma,\tau)) = (e_1^T B(\sigma,\tau)^i e_1)$, and let $H_n(\sigma,\tau) \in \mathbb{F}(\sigma,\tau)^{n\times n}$:

$$H_n(\sigma,\tau) = \begin{bmatrix} s_0(\sigma,\tau) & s_1(\sigma,\tau) & \ldots & s_{n-1}(\sigma,\tau) \\ s_1(\sigma,\tau) & s_2(\sigma,\tau) & \ldots & s_n(\sigma,\tau) \\ \vdots & & \ddots & \vdots \\ s_{n-1}(\sigma,\tau) & s_n(\sigma,\tau) & \ldots & s_{2n-2}(\sigma,\tau) \end{bmatrix}. \qquad (9)$$

Since the minimal polynomial of the sequence is of degree n, $H_n(\sigma,\tau)$ is non-singular [18, Eq. (2.6)]. Thus, any s, t with $\det(H_n(s,t)) \neq 0$ will yield $f_{e_1}^{B(s,t),e_1} = C^{B(s,t)}$. Now for the complexities:

- The volume of communication is reduced, from $8n$ to $5n$.
- With respect to Theorem 13, the cost for Victor is slightly improved: an application by D replaced by an application by $\Gamma(s,t)$, a dot-product with u replaced by one with e_1, and n operations for $\det(D)$ are replaced by less than $2\lceil\log_2(n)\rceil + 1$ to compute $t^n + s$. The overall verification cost thus decreases from $\mu(A) + 15n + o(n)$ to $\mu(A) + 13n + o(n)$;
- For Peggy, applying the diagonal scaling costs n per iteration, while applying $\Gamma(s,t)$ costs $2n$ operations per iteration; but applying a random u^T costs about $2n$ operations per iteration, while applying e_1^T is just selecting the first coefficient, so her cost is slightly improved. Then, too choose such s, t, Peggy can try uniformly sampled elements in S, and see whether $\deg(f_{e_1}^{B(s,t),e_1}) = n$. Since $\deg(\Gamma(\sigma,\tau)) = 1$, we have that $\deg(s_i(\sigma,\tau)) \leq i$ and $\deg(\det(H_n(\sigma,\tau))) \leq n(n-1)$. Hence, by the DeMillo-Lipton/Schwartz/Zippel Lemma

$$\text{Prob}\left(\det(H_n(s,t)) \neq 0\right) \geq 1 - \frac{n(n-1)}{|S|}. \qquad (10)$$

As soon as $|S| \geq 2n(n-1)$, the probability of success for Peggy is thus larger than $1/2$ and the expected number of trials is less than 2. □

We gather the differences between the protocols of Figure 4 and 5 in Table 1. The two certificate differ mostly only in the preconditioning. But this allows to gain a lot of randomization: the number of random field elements per try to sample for Peggy is reduced

Certificates for the determinant of sparse matrices		
Preconditioner	§ 6 : D	§ 7 : $\Gamma(t,s)$
Verifier	$\mu(A) + 15n + o(n)$	$\mu(A) + 13n + o(n)$
Communications	$8n$	$5n$
Random elements	$3n + 2$	3
Prover	$\textsc{MinPoly}(n) + \textsc{LinSys}(n)$	
Field size	$\geq \frac{1}{2}(n^2 - n)$	$\geq n^2 - n$

Table 1: Summary of the complexity bounds of the certificates presented in this paper for the determinant of sparse matrices (n is the dimension, $\mu(A)$ bounds the cost of one matrix-vector product, $\textsc{MinPoly}(n)$ (resp. $\textsc{LinSys}(n)$) is the cost of computing the minimal polynomial of a sequence (resp. of solving a linear system).

from $3n$ to only 2. As the size of the set is $\Omega(n^2)$, this reduces the number of random bits from $O(n\log(n))$ to $O(\log(n))$.

REMARK 15 *(*CERTIFICATE FOR THE CHARACTERISTIC POLYNOMIAL*).* The certificate for the determinant can be combined with the characteristic polynomial reduction of [9, Figure 1]. As this reduction is linear for Victor, this then provides now also a *linear* time verification procedure for the characteristic polynomial:

1. Peggy sends c^A as the characteristic polynomial;
2. Victor now sends back a random point $r \in S \subseteq \mathbb{F}$;
3. Peggy and Victor enter a determinant certificate for $rI - A$;
4. Once convinced, Victor checks that $\det(rI - A) \overset{?}{=\!=} c^A(r)$.

For a random $r \in S$, in the determinant sub-certificate, $rI - A$ will be non-singular if $c^A(r) \neq 0$, hence with probability $\geq 1 - n/|S|$. Then $\det(rI - A)$ is certified using the certificate of Figure 5.

Best known algorithms for computing the characteristic polynomial, using either quadratic space and fast matrix multiplication or linear space [28], have cost $O(n^{2+\alpha})$ for some $\alpha > 0$. The characteristic polynomial with $\mu(A) = O(n)$ is thus an example of a problem whose worst-case complexity bound is super-quadratic in the verification cost $O(n)$.

8. REFERENCES

[1] B. Beckerman and G. Labahn. A uniform approach for the fast computation of matrix-type Padé approximants. *SIAM Journal on Matrix Analysis and Applications*, 15(3):804–823, July 1994.

[2] G. Bertoni, J. Daemen, M. Peeters, and G. Assche. Sponge-based pseudo-random number generators. In S. Mangard and F.-X. Standaert, editors, *CHES 2010*, pages 33–47. Springer, 2010.

[3] L. Chen, W. Eberly, E. L. Kaltofen, B. D. Saunders, W. J. Turner, and G. Villard. Efficient matrix preconditioners for black box linear algebra. *Linear Algebra Appl.*, 343-344:119–146, 2002. www.math.ncsu.edu/ kaltofen/bibliography/02/CEKSTV02.pdf.

[4] K.-M. Chung, Y. T. Kalai, and S. P. Vadhan. Improved delegation of computation using fully homomorphic encryption. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 483–501. Springer, 2010.

[5] D. Coppersmith. Solving homogeneous linear equations over GF(2) via block Wiedemann algorithm. *Mathematics of Computation*, 62(205):333–350, Jan. 1994.

[6] R. J. F. Cramer. *Modular Design Of Secure Yet Practical Cryptographic Protocols*. PhD thesis, U. van Amsterdam, Jan. 1997.

[7] R. A. DeMillo and R. J. Lipton. A probabilistic remark on algebraic program testing. *Inf. Process. Letters*, 7(4):193–195, June 1978.

[8] J.-G. Dumas, E. Kaltofen, and E. Thomé. Certificates for the verification of Wiedemann's Krylov sequence. Technical report, arXiv:, Jan. 2016. arxiv.org/abs/1507.01083.

[9] J.-G. Dumas and E. L. Kaltofen. Essentially optimal interactive certificates in linear algebra. In K. Nabeshima, editor, *ISSAC'2014*, pages 146–153. ACM Press, New York, July 2014. www.math.ncsu.edu/ kaltofen/bibliography/14/DuKa14.pdf.

[10] W. Eberly and E. Kaltofen. On randomized Lanczos algorithms. In Küchlin [21], pages 176–183. www.math.ncsu.edu/ kaltofen/bibliography/97/EbKa97.pdf.

[11] 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 *Lecture Notes in Computer Science*, pages 186–194. Springer, 11–15 Aug. 1986.

[12] D. Fiore and R. Gennaro. Publicly verifiable delegation of large polynomials and matrix computations, with applications. In *ACM CCS '12*, pages 501–512, New York, NY, USA, 2012. ACM.

[13] R. Freivalds. Fast probabilistic algorithms. In J. Bečvář, editor, *Mathematical Foundations of Computer Science 1979*, volume 74 of *Lecture Notes in Computer Science*, pages 57–69, Olomouc, Czechoslovakia, Sept. 1979. Springer.

[14] C. Gentry, J. Groth, Y. Ishai, C. Peikert, A. Sahai, and A. Smith. Using fully homomorphic hybrid encryption to minimize non-interative zero-knowledge proofs. *Journal of Cryptology*, pages 1–24, 2014.

[15] 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.

[16] E. Kaltofen. Analysis of Coppersmith's block Wiedemann algorithm for the parallel solution of sparse linear systems. *Mathematics of Computation*, 64(210):777–806, Apr. 1995. www.math.ncsu.edu/ kaltofen/bibliography/95/Ka95_mathcomp.pdf.

[17] E. Kaltofen and V. Pan. Processor efficient parallel solution of linear systems over an abstract field. In *ACM SPAA '91*, pages 180–191, New York, N.Y., 1991. ACM Press. www.math.ncsu.edu/ kaltofen/bibliography/91/KaPa91.pdf.

[18] E. Kaltofen and G. Villard. On the complexity of computing determinants. *Computational Complexity*, 13(3):91–130, 2005. www.math.ncsu.edu/ kaltofen/bibliography/04/KaVi04_2697263.pdf.

[19] E. L. Kaltofen, B. Li, Z. Yang, and L. Zhi. Exact certification in global polynomial optimization via sums-of-squares of rational functions with rational coefficients. *Journal of Symbolic Computation*, 47(1):1–15, Jan. 2012. www.math.ncsu.edu/ kaltofen/bibliography/09/KLYZ09.pdf.

[20] 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. www.math.ncsu.edu/ kaltofen/bibliography/11/KNS11.pdf.

[21] W. W. Küchlin, editor. *ISSAC'97*. ACM Press, New York, July 1997.

[22] E. W. Ng, editor. *EUROSAM '79, International Symposium on Symbolic and Algebraic Computation, Marseille, France, June 1979, Proceedings*, volume 72 of *Lecture Notes in Computer Science*. Springer, 1979.

[23] NIST. *FIPS Publication 202: SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions*, Aug. 2015.

[24] B. Parno, J. Howell, C. Gentry, and M. Raykova. Pinocchio: Nearly practical verifiable computation. In *IEEE SP '13*, pages 238–252, Washington, DC, USA, 2013. IEEE Computer Society.

[25] J. T. Schwartz. Probabilistic algorithms for verification of polynomial identities. In Ng [22], pages 200–215.

[26] J. Thaler. Time-optimal interactive proofs for circuit evaluation. In R. Canetti and J. A. Garay, editors, *CRYPTO '13*, volume 8043 of *Lecture Notes in Computer Science*, pages 71–89. Springer, 2013.

[27] G. Villard. Further analysis of Coppersmith's block Wiedemann algorithm for the solution of sparse linear systems. In Küchlin [21], pages 32–39. perso.ens-lyon.fr/gilles.villard/BIBLIOGRAPHIE/PDF/issac97.pdf.

[28] G. Villard. Computing the Frobenius normal form of a sparse matrix. In *CASC'00*, pages 395–407. Springer, 2000. perso.ens-lyon.fr/gilles.villard/BIBLIOGRAPHIE/PDF/casc.pdf.

[29] D. H. Wiedemann. Solving sparse linear equations over finite fields. *IEEE Transactions on Information Theory*, 32(1):54–62, Jan. 1986.

[30] R. Zippel. Probabilistic algorithms for sparse polynomials. In Ng [22], pages 216–226.