Error Correction in Fast Matrix Multiplication and Inverse

Daniel S. Roche

United States Naval Academy
Annapolis, Maryland, U.S.A.
roche@usna.edu

ABSTRACT

We present new algorithms to detect and correct errors in the product of two matrices, or the inverse of a matrix, over an arbitrary field. Our algorithms do not require any additional information or encoding other than the original inputs and the erroneous output. Their running time is softly linear in the number of nonzero entries in these matrices when the number of errors is sufficiently small, and they also incorporate fast matrix multiplication so that the cost scales well when the number of errors is large. These algorithms build on the recent result of Gasieniec et al. [18] on correcting matrix products, as well as existing work on verification algorithms, sparse low-rank linear algebra, and sparse polynomial interpolation.

KEYWORDS

error correction, matrix multiplication, linear algebra, sparse interpolation, algorithms

ACM Reference Format:

Daniel S. Roche. 2018. Error Correction in Fast Matrix Multiplication and Inverse. In *ISSAC '18: 2018 ACM International Symposium on Symbolic and Algebraic Computation*, July 16–19, 2018, New York, NY, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3208976.3209001>

1 INTRODUCTION

Efficiently and gracefully handling computational errors is critical in modern computing. Such errors can result from short-lived hardware failures, communication noise, buggy software, or even malicious third-party servers.

The first goal for fault-tolerant computing is *verification*. The groundbreaking work of Freivalds [17] presented a linear-time algorithm to verify the correctness of a single matrix product. Recently, efficient verification algorithms for a wide range of computational linear algebra problems have been developed [13, 15, 16, 27].

A higher level of fault tolerance is achieved through *error correction*, which is the goal of the present study. This is a strictly stronger model than verification, where we seek not only to identify when a result is incorrect, but also to compute the correct result if it is “close” to the given, incorrect one. Some recent error-correction problems considered in the computer algebra literature include Chinese remaindering [5, 19, 29], system solving [8, 28], and function recovery [12, 26].

A generic approach to correcting errors *in transmission* would be to use an error-correcting code, writing the result with some redundancy to support transmission over a noisy channel. But this

approach requires extra bandwidth, is limited to a fixed number of errors, and does not account for malicious alterations or mistakes in the computation itself.

Instead, we will use the information in the problem itself to correct errors. This is always possible by re-computing the result, and so the goal is always to correct a small number of errors in (much) less time than it would take to do the entire computation again. Our notion of error correction makes it possible to correct an unbounded number of mistakes that occur either in computation or transmission, either at random or by a malicious party.

Some other applications of error correction have nothing to do with alterations or mistakes. One example is sparse matrix multiplication, which can be solved with an error-correction algorithm by simply assuming the “erroneous” matrix is zero. Another example is computing a result over $\mathbb{Z}$ using Chinese remaindering modulo small primes, where the entries have varying bit lengths. The small entries are determined modulo the first few primes, and the remaining large entries are found more quickly by applying an error correction algorithm using (known) small entries.

1.1 Our results

Matrix multiplication with errors. The first problem we consider is the same as in Gasieniec et al. [18], correcting a matrix product. Given $A, B, C \in \mathbb{F}^{n \times n}$, the task is to compute the unique error matrix $E \in \mathbb{F}^{n \times n}$ such that $AB = C - E$.

Algorithm 5 solves this problem using

$$\tilde{O}\left(t + k \cdot \min\left(\left\lceil \frac{t}{r} \right\rceil, n/\min(r, \frac{k}{r})^{3-\omega}\right)\right)$$

field operations, where

- $t = \#A + \#B + \#C$ is the number of nonzero entries in the input matrices, at most $O(n^2)$;
- $k = \#E$ is the number of errors in the given product C ; and
- $r \leq n$ is the number of distinct rows (or columns) in which errors occur, whichever is larger.

An even more detailed complexity statement, incorporating rectangular matrices and a failure probability, can be found in Theorem 7.1. To understand the implications of this complexity, consider six cases, as summarized in Table 1.

The cost of our algorithm depends on the locations of the errors. If there are a constant number of errors per row or column, we can use sparse multiplication to achieve softly-linear complexity $\tilde{O}(t + n + k)$. This is similar to [18] except that we do not require the input matrices to be dense.

If $k > n$ errors are spread out among all rows and columns, the worst-case cost using dense multiplication is $\tilde{O}(k^{0.38}n^{1.63})$. As the number of errors grows to n^2 , this cost approaches that of normal matrix multiplication without error correction.

This paper is authored by an employee(s) of the United States Government and is in the public domain. Non-exclusive copying or redistribution is allowed, provided that the article citation is given and the authors and agency are clearly identified as its source.

ISSAC '18, July 16–19, 2018, New York, NY, USA

2018. ACM ISBN 978-1-4503-5550-6/18/07.

<https://doi.org/10.1145/3208976.3209001>

	Spread-out errors $r = \min(n, k)$		Compact errors $r = \sqrt{k}$
	$k \leq n$	$k > n$	
Sparse	$t + k + n$	$k \lceil \frac{t}{n} \rceil + n$	$\sqrt{k}t + k + n$
Dense	$t + kn$	$k^{\omega-2}n^{4-\omega}$ $k^{0.38}n^{1.63}$	$t + k^{(\omega-1)/2}n$ $t + k^{0.69}n$

Table 1: Soft-oh cost of Algorithm 5 in several cases, for correcting up to k errors (which appear in r rows) in the product of two $n \times n$ matrices with at most t nonzero entries. The two rows indicate whether sparse or dense rectangular multiplication is used as a subroutine.

In the other extreme, if the errors are isolated to a square (but not necessarily contiguous) submatrix, then the worst-case cost using dense multiplication is $\tilde{O}(t + k^{0.69}n)$. Again this scales to n^ω as $k \rightarrow n^2$, but the cost is better for $k < n^2$. This situation may also make sense in the context of distributed computing, where each node in a computing cluster might be assigned to compute one submatrix of the product.

Our algorithm is never (asymptotically) slower than that of [18] or the $O(n^\omega)$ cost of naïve recomputation, but it is faster in many cases: for example when the inputs are sparse, when there are many errors, and when the errors are compactly located.

Matrix inverse with errors. The second problem we consider is that of correcting computational errors in a matrix inverse. Formally, given $A, B \in F^{n \times n}$, the task is to compute the unique error matrix $E \in F^{n \times n}$ such that $A^{-1} = B + E$.

Algorithm 6 shows how to solve this problem using

$$\tilde{O}\left(t + nk / \min(r, \frac{k}{r})^{3-\omega} + r^\omega\right)$$

field operations, where the parameters t, k, n are the same as before, and r is the number of rows or columns which contain errors, whichever is *smaller*. This is the same complexity as our algorithm for matrix multiplication with errors, plus an additional r^ω term.

The additional r^ω term makes the situation here less fractured than with the multiplication algorithm. In the case of spread-out errors, the cost is $\tilde{O}(t + nk + k^\omega)$, which becomes simply $\tilde{O}(n^2 + k^\omega)$ when the input matrices are dense. The case of compact errors is better, and the complexity is just $\tilde{O}\left(t + nk^{(\omega-1)/2}\right)$, the same as that of our multiplication algorithm.

As before, the cost of this algorithm approaches the $O(n^\omega)$ cost of the naïve inverse computation as the number of errors k grows to n^2 .

Domains. The algorithms we present work over any field and require only that a field element of order larger than n can be found. Because our algorithms explicitly take $O(n)$ time to compute the powers of this element anyway, the usual difficulties with efficiently generating high-order elements in finite fields do not arise.

Over the integers $\mathbb{Z}$ or rationals $\mathbb{Q}$, the most efficient approach would be to work modulo a series of primes and use Chinese remaindering to recover the result. A similar approach could be used over polynomial rings.

Over the reals $\mathbb{R}$ or complex numbers $\mathbb{C}$, our algorithms work only under the unrealistic assumption of infinite precision. We have not considered the interesting question of how to recover from the two types of errors (noise and outliers) that may occur in finite precision.

1.2 Related work

Let ω be a constant between 2 and 3 such that two $n \times n$ matrices can be multiplied using $O(n^\omega)$ field operations. In practice, $\omega = 3$ for dimensions n up to a few hundred, after which Strassen-Winograd multiplication is used giving $\omega \leq 2.81$. The best asymptotic algorithm currently gives $\omega < 2.3728639$ [31].

Even though the practical value of ω is larger than the asymptotic one, matrix multiplication routines have been the subject of intense implementation development for decades, and highly-tuned software is readily available for a variety of architectures [14, 35]. Running times based on n^ω have practical as well as theoretical significance; the ω indicates where an algorithm is able to take advantage of fast low-level matrix multiplication routines.

Special routines for multiplying sparse matrices have also been developed. Writing t for the number of nonzero entries in the input, the standard row- or column-wise sparse matrix multiplication costs $O(tn)$ field operations. Yuster and Zwick [36] improved this to $\tilde{O}(n^2 + t^{0.697}n^{1.197})$. Later work also incorporates the number k of nonzeros in the product matrix. Lingas [32] gave an output-sensitive algorithm with running time $O(k^{0.186}n^2)$. Amossen and Pagh [1] have another with complexity $\tilde{O}(t^{0.667}k^{0.667} + t^{0.862}k^{0.408})$, an improvement when the outputs are not too dense. Pagh [33] developed a different approach with complexity $\tilde{O}(t + kn)$; that paper also includes a nice summary of the state of the art. Our multiplication with errors algorithm can also be used for sparse multiplication, and it provides a small improvement when the input and output are both not *too* sparse.

The main predecessor to our work is the recent algorithm of Gasieniec et al. [18], which can correct up to k errors in the product of two $n \times n$ matrices using $\tilde{O}(n^2 + kn)$ field operations. Their approach is based in part on the earlier work of Pagh [33] and makes clever use of hashing and fast Fourier transforms in order to achieve the stated complexity.

We not only improve on the complexity of that problem, but also consider related problem of correcting errors in a matrix inverse. Computing the inverse of a matrix (without errors) has the same asymptotic complexity $O(n^\omega)$ as fast matrix multiplication [9].

2 OVERVIEW

Multiplication with errors. A high-level summary of our multiplication algorithm is as follows:

- (1) Determine which rows in the product contain errors.
- (2) Write the unknown sparse matrix of errors as $E = C - AB$. Remove the rows with no errors from E, C , and A , so we have $E' = C' - A'B$.
- (3) Treating the rows of E' as sparse polynomials, use structured linear algebra and fast matrix multiplication to evaluate each row polynomial at a small number of points.
- (4) Use sparse polynomial interpolation to recover at least half of the rows from their evaluations.

(5) Update C and iterate $O(\log k)$ times to recover all k errors.

To determine the rows of E which are nonzero in Step 1, we apply a simple variant of Frievalds' verification algorithm: for a random vector $\mathbf{v}$, the nonzero rows of $E\mathbf{v}$ are probably the same as the nonzero rows of E itself.

The evaluation/interpolation approach that follows is reminiscent of that in Gasieniec et al. [18], with two important differences. First, using sparse polynomial techniques rather than hashing and FFTs means the complexity scales linearly with the number of nonzeros t in the input rather than the dense matrix size n^2 . Second, by treating the rows separately rather than recovering all errors at once, we are able to incorporate fast matrix multiplication so that our worst-case cost when all entries are erroneous is never more than $O(n^\omega)$, the same as that of naïve recomputation.

Section 4 provides details on the Monte Carlo algorithm to determine the rows and columns where errors occur (Step 1 above), Section 5 presents a variant of sparse polynomial interpolation that suits our needs for Step 4, and Section 6 explains how to perform the row polynomial evaluations in Step 3. A certain rectangular matrix multiplication in this step turns out to dominate the overall complexity. We connect all these pieces of the multiplication algorithm in Section 7.

Inverse with errors. Our algorithm for correcting errors in a matrix inverse is based on the multiplication algorithm, but with two complications. Writing A for the input matrix, B for the erroneous inverse, and E for the (sparse) matrix of errors in B , we have:

$$AE = I - AB, \quad EA = I - BA.$$

We want to compute $E\mathbf{v}$ for a random vector $\mathbf{v}$ in order to determine which rows of E are nonzero. This is no longer possible directly, but using the second formulation as above, we can compute $E\mathbf{A}\mathbf{v}$. Because A must be nonsingular, $\mathbf{A}\mathbf{v}$ has the same distribution as a random vector, so this approach still works.

The second complication is that removing zero rows of E does not change the dimension of the right-hand side, but only removes corresponding *columns* from A on the left-hand side. We make use of the recent result of Cheung et al. [11], which shows how to quickly find a maximal set of linearly independent rows in a sparse matrix. This allows us to find a small submatrix X of A such that $XE' = I' - A'B$, with E', I', A' being compressed matrices as before. Because the size of X depends on the number of errors, it can be inverted much more quickly than re-computing the entire inverse of A .

The dominating cost in most cases is the same rectangular matrix product as needed as in matrix multiplication with error correction. However, here we also have the extra cost of finding X and computing its inverse, which dominates the complexity when there are a moderate number of errors and their locations are spread out.

This algorithm — which depends on the subroutines of Sections 4 to 6 — is presented in detail in Section 8.

3 NOTATION AND PRELIMINARIES

The soft-oh notation $\tilde{O}(\dots)$ is the same as the usual big-oh notation but ignoring sub-logarithmic factors: $f \in \tilde{O}(g)$ if and only if $f \in O(g \log^{O(1)} g)$, for some runtime functions f and g .

We write $\mathbb{N}$ for the set of nonnegative integers and $\mathbb{F}_q$ for the finite field with q elements.

For a finite set $\mathcal{S}$, the number of elements in $\mathcal{S}$ is denoted $\#\mathcal{S}$, and $\mathbb{P}(\mathcal{S})$ is the powerset of $\mathcal{S}$, i.e., the set of subsets of $\mathcal{S}$.

The number of nonzero entries in a matrix $M \in \mathbb{F}^{m \times n}$ is written as $\#M$. Note that $\#M \leq mn$, and if M has rank k then $\#M \geq k$.

For a univariate polynomial $f \in \mathbb{F}[x]$, $\text{supp}(f) \subseteq \mathbb{N}$ denotes the exponents of nonzero terms of f . Note that $\#\text{supp}(f) \leq \deg f + 1$.

We assume that the number of field operations to multiply two (dense) polynomials in $\mathbb{F}[x]$ with degrees less than n is $\tilde{O}(n)$. This is justified by the generic algorithm from Cantor and Kaltofen [10] which gives $O(n \log n \log \log n)$, or more recent results of Harvey et al. [20] which improve this to $O(n \log n 8^{\log^* n})$ for finite fields.

As discussed earlier, we take ω with $2 \leq \omega \leq 3$ to be any feasible exponent of matrix multiplication. By applying fast matrix multiplication in blocks, it is also possible to improve the speed of rectangular matrix multiplication in a straightforward way. In particular:

FACT 3.1. *The product of any $m \times \ell$ matrix times a $\ell \times n$ matrix can be found using $O(m\ell n / \min(m, \ell, n)^{3-\omega})$ ring operations.*

Note that it might be possible to improve this for certain values of m, ℓ, n using [30], but that complexity is much harder to state and we have not investigated such an approach.

We restate two elementary facts on sparse matrix multiplication.

FACT 3.2. *For any $M \in \mathbb{F}^{m \times n}$ and $\mathbf{v} \in \mathbb{F}^n$, the matrix-vector product $M\mathbf{v}$ can be computed using $O(\#M)$ field operations.*

COROLLARY 3.3. *For any $A \in \mathbb{F}^{m \times \ell}$ and $B \in \mathbb{F}^{\ell \times n}$, their product AB can be computed using $O(\#A \cdot n)$ field operations.*

4 IDENTIFYING NONZERO ROWS AND COLUMNS

Our algorithms for error correction in matrix product and inverse computation both begin by determining which rows and columns contain errors. This is accomplished in turn by applying a random row or column vector to the unknown error matrix E .

For now, we treat the error matrix E as a *black box*, i.e., an oracle matrix which we can multiply by a vector on the right or left-hand side. The details of the black box construction differ for the matrix multiple and inverse problem, so we delay those until later.

The idea is to apply a row or column vector of random values to the unknown matrix E , and then examine which entries in the resulting vector are zero. This is exactly the same idea as the classic Monte Carlo verification algorithm of Frievalds [17], which we extend to which rows or columns in a matrix are nonzero. This black-box procedure is detailed in Algorithm 1 FindNonzeroRows.

The correctness of FindNonzeroRows depends on the parameter ϵ . The running time depends on this ϵ , the size of $\mathbb{F}$, and the cost of performing matrix-vector products with the black-box matrix M .

LEMMA 4.1. *FindNonzeroRows always returns a list of nonzero row indices in M . With probability at least $1 - \epsilon$, it returns the index of every nonzero row in M .*

PROOF. Consider a single row $\mathbf{u}^T$ of M . If $\mathbf{u}$ is zero, then the corresponding row of $M\mathbf{V}$ will always be zero; hence every index returned must be a nonzero row.

Algorithm 1: FindNonzeroRows($V \mapsto MV, \epsilon$)

Input: Black box for right-multiplication by an unknown matrix $M \in \mathbb{F}^{m \times n}$ and error bound $\epsilon \in \mathbb{R}, 0 < \epsilon < 1$

Output: Set $\mathcal{J} \subseteq \{0, \dots, m-1\}$ that with probability at least $1 - \epsilon$ consists of the indices of nonzero rows in M

- 1 $\ell \leftarrow \lceil \log_{\#F} \frac{m}{\epsilon} \rceil$
- 2 $V \leftarrow$ matrix in $\mathbb{F}^{n \times \ell}$ with uniform random entries from $\mathbb{F}$
- 3 $B \leftarrow MV$ using the black box
- 4 **return** Indices of rows of B which are not all zeros

Next assume $\mathbf{u}^T$ has at least one nonzero entry, say at index i , and consider one row $\mathbf{v} \in \mathbb{F}^n$ of the random matrix V . Whatever the other entries in $\mathbf{u}$ and $\mathbf{v}$ are, there is exactly one value for the i 'th entry of $\mathbf{v}$ such that $\mathbf{u}^T \mathbf{v} = 0$. So the probability that $\mathbf{u}^T \mathbf{v} = 0$ is $\frac{1}{\#F}$, and by independence the probability that $\mathbf{u}^T V = \mathbf{0}^{1 \times \ell}$ is $1/(\#F)^\ell$.

Because M has at most m nonzero rows, the union bound tells us that the probability *any one* of the corresponding rows of MV is zero is at most $m/(\#F)^\ell$. From the definition of ℓ , this ensures a failure probability of at most ϵ . $\square$

Because we mostly have in mind the case that $\mathbb{F}$ is a finite field, we assume in FindNonzeroRows and in the preceding proof that it is possible to sample uniformly from the entire field $\mathbb{F}$. However, the same results would hold if sampling from any fixed subset of $\mathbb{F}$ (perhaps increasing the size of ℓ if the subset is small).

5 BATCHED LOW-DEGREE SPARSE INTERPOLATION

Our algorithms use techniques from sparse polynomial interpolation to find the locations and values of erroneous entries in a matrix product or inverse.

Consider an unknown univariate polynomial $f \in \mathbb{F}[x]$ with degree less than n and at most s nonzero terms. The nonzero terms of f can be uniquely recovered from evaluations $f(\theta^0), \dots, f(\theta^{2s-1})$ at $2s$ consecutive powers [3, 24].

Doing this efficiently in our context requires a few variations to the typical algorithmic approach. While the state of the art seen in recent papers such as [2, 22, 23, 25] achieves softly-linear complexity in terms of the sparsity t , there is either a higher dependency on the degree, a restriction to certain fields, or both.

Here, we show how to take advantage of two unique aspects of the problem in our context. First, we will interpolate a *batch* of sparse polynomials at once, on the same evaluation points, which allows for some useful precomputation. Secondly, the maximum degree of any polynomial we interpolate is n , the column dimension of the unknown matrix, and we are allowed linear-time in n .

The closest situation in the literature is that of van der Hoeven and Lecerf [21], who show how to recover a sparse polynomial if the exponents are known. In our case, we can tolerate a much larger set containing the possible exponents by effectively amortizing the size of this exponent set over the group of batched sparse interpolations.

Recall the general outline of Prony's method for sparse interpolation of $f \in \mathbb{F}[x]$ from evaluations $(f(\theta^i))_{0 \leq i < 2s}$:

- (1) Find the minimum polynomial $\Gamma \in \mathbb{F}[z]$ of the sequence of evaluations.
- (2) Find the roots $v_1, \dots, v_s$ of Γ .
- (3) Compute the discrete logarithm base θ of each v_i to determine the exponents of f .
- (4) Solve a transposed Vandermonde system built from v_i and the evaluations to recover the coefficients of f .

Steps 1 and 4 can be performed over any field using fast multipoint evaluation and interpolation in $\tilde{O}(s)$ field operations [24]. However, the other steps depend on the field and in particular Step 3 can be problematic over large finite fields.

To avoid this issue, we first pre-compute all possible roots θ^{e_i} for any minpoly Γ_i in any of the batch of r sparse interpolations that will be performed. Although the set of possible roots is larger than the degree of any single minpoly Γ_i , in our case the set is always bounded by the dimension of the matrix, so performing this precomputation once is worthwhile.

5.1 Batched root finding

While fast root-finding procedures over many fields have already been developed, we present a simple but efficient root-finding procedure for our specific case that has the advantage of running in softly-linear time over any coefficient field. The general idea is to first compute a coprime basis (also called a gcd-free basis) for the minpolys Γ_i , then perform multi-point evaluation with the precomputed list of possible roots.

The details of procedure FindRoots are in Algorithm 2 below. In the algorithm, we use a *product tree* constructed from a list of t polynomials. This is a binary tree of height $O(\log t)$, with the original polynomials at the leaves, and where every internal node is the product of its two children. In particular, the root of the tree is the product of all t polynomials. The total size of a product tree, and the cost of computing it, is softly-linear in the total size of the input polynomials; see Borodin and Munro [6] for more details.

LEMMA 5.1. *Given a set $Q \subseteq \mathbb{F}$ with of size $\#Q = c$ and a list of r polynomials $\Gamma_1, \dots, \Gamma_r \in \mathbb{F}[z]$ each with degree at most s , Algorithm 2 FindRoots correctly determines all roots of all Γ_i 's which are in Q using $\tilde{O}(rs + c)$ field operations*

PROOF. Algorithms 18.1 and 21.2 of Bernstein [4] are deterministic and compute (respectively) a coprime basis and then a factorization of the Γ_i 's according to the basis elements.

The polynomials $M_{i,j}$ form a standard product tree from the polynomials in the coprime basis. For any element of the product tree, its roots must be a subset of the roots of its parent node. So the multi-point evaluations on Line 9 determine, for every polynomial in the tree, which elements of Q are roots of that polynomial.

The cost of computing the coprime basis using Algorithm 18.1 in [4] is $\tilde{O}(rs)$ with at least a $\log^5(rs)$ factor. This gives the first term in the complexity (and explains our use of soft-oh notation throughout). This cost dominates the cost of constructing the product tree and factoring the Γ_i 's on Line 11.

Using the standard product tree algorithm [6], the cost of multi-point evaluation of a degree- n polynomial at each of n points is $\tilde{O}(n)$ field operations. In our algorithm, this happens at each level down the product tree. At each of $\lceil \log_2 t \rceil \in O(\log(rs))$ levels there are

Algorithm 2: FindRoots($Q, [\Gamma_1, \dots, \Gamma_r]$)

Input: Set of possible roots $Q \subseteq F$ and r polynomials $\Gamma_1, \dots, \Gamma_r \in F[z]$
Output: A list of sets $\{\mathcal{R}_1, \dots, \mathcal{R}_r\} \in \mathbb{P}(Q)^r$ such that $g_i(\alpha_j) = 0$ for every $1 \leq i \leq r$ and $\alpha_j \in \mathcal{R}_i$

```

1  $\{g_0, \dots, g_{t-1}\} \leftarrow$  coprime basis of  $\{\Gamma_i\}_{1 \leq i \leq r}$  using
   Algorithm 18.1 of Bernstein [4]
/* Compute product tree from bottom-up */
2 for  $j \leftarrow 0, 1, \dots, t-1$  do  $M_{0,j} \leftarrow g_j$ 
3 for  $i \leftarrow 1, 2, \dots, \lceil \log_2 t \rceil$  do
4   for  $j \leftarrow 0, 1, \dots, \lceil t/2^i \rceil - 1$  do
5      $M_{i,j} \leftarrow M_{i-1,2j} \cdot M_{i-1,2j+1}$ 
/* Find roots of basis elements */
6  $S_{\lceil \log_2 t \rceil + 1, 0} \leftarrow Q$ 
7 for  $i \leftarrow \lceil \log_2 t \rceil, \lceil \log_2 t \rceil - 1, \dots, 0$  do
8   for  $j \leftarrow 0, 1, \dots, \lceil t/2^i \rceil - 1$  do
9     Evaluate  $M_{i,j}$  at each point in  $S_{i+1, \lfloor j/2 \rfloor}$  using fast
       multipoint evaluation
10     $S_{i,j} \leftarrow$  roots of  $M_{i,j}$  found on previous step
/* Determine roots of original polynomials */
11 Compute exponents  $(e_{i,j})_{1 \leq i \leq r, 0 \leq j < t}$  such that
     $\Gamma_i = g_0^{e_{i,0}} \dots g_{t-1}^{e_{i,t-1}}$  for all  $i$ , using algorithm 21.2 of
    Bernstein [4]
12 for  $i \leftarrow 1, 2, \dots, r$  do  $\mathcal{R}_i \leftarrow \bigcup_{0 \leq j < t, e_{i,j} \geq 1} S_{0,j}$ 

```

at most c points (since the polynomials at each level are relatively prime). The total degree at each level is at most $\deg \prod_i \Gamma_i \leq rs$. If $rs > c$, then this cost is already bounded by that of determining the coprime basis. Otherwise, the multi-point evaluations occur in $O(c/(rs))$ batches of rs points each, giving the second term in the complexity statement. $\square$

5.2 Batched sparse interpolation algorithm

We are now ready to present the full algorithm for performing simultaneous sparse interpolation on r polynomials based on $2s$ evaluations each.

Note that, in the case of very small fields with no order- n elements, this approach can still be used by working over an extension of order at most $O(\log n)$. In our application n is actually the dimension of the erroneous matrix, so this does not change the overall complexity by more than a logarithmic factor.

THEOREM 5.2. *Let F be a field with an element $\theta \in F$ whose multiplicative order is at least n , and suppose $f_1, \dots, f_r \in F[x]$ are unknown univariate polynomials. Given bounds $n, s \in \mathbb{N}$, a set $S \subseteq \mathbb{N}$ of possible exponents, and evaluations $f_i(\theta^j)$ for all $1 \leq i \leq r$ and $0 \leq j < 2s$, Algorithm 3 MultiSparseInterp requires $\tilde{O}(rs + \#S \log n)$ field operations in the worst case.*

The algorithm correctly recovers every polynomial f_i such that $\# \text{supp} f_i \leq s$, $\text{supp} f_i \subseteq S$, and $\deg f_i < n$.

PROOF. From Lemma 5.1, the cost of FindRoots is $\tilde{O}(rs + \#S)$. This dominates the cost of the minpoly and transposed Vandermonde computations, which are both $\tilde{O}(rs)$ [24]. Using binary

Algorithm 3: MultiSparseInterp(r, n, s, S, θ, Y)

Input: Bounds $r, n, s \in \mathbb{N}$, set $S \subseteq \{0, \dots, n-1\}$ of possible exponents, high-order element $\theta \in F$, and evaluations $Y_{i,j} \in F$ such that $f_i(\theta^j) = Y_{i,j}$ for all $1 \leq i \leq r$ and $0 \leq j < 2s$
Output: Nonzero coefficients and corresponding exponents of $f_i \in F[x]$ or 0 indicating failure, for each $1 \leq i \leq r$

```

1 for  $i \leftarrow 1, 2, \dots, r$  do
2    $\Gamma_i \leftarrow$  MinPoly( $Y_{i,0}, \dots, Y_{i,2s-1}$ ) using fast
   Berlekamp-Massey algorithm
3  $\mathcal{R}_1, \dots, \mathcal{R}_r \leftarrow$  FindRoots( $\{\theta^d \mid d \in S\}, \Gamma_1, \dots, \Gamma_r$ )
4 for  $i \leftarrow 1, 2, \dots, r$  do
5   if  $\#\mathcal{R}_i \neq \deg \Gamma_i$  then  $f_i \leftarrow 0$ 
6   else
7      $t_i \leftarrow \#\mathcal{R}_i$ 
8      $\{e_{i,1}, \dots, e_{i,t_i}\} \leftarrow \{d \in S \mid \theta^d \in \mathcal{R}_i\}$ 
9      $[a_{i,1}, \dots, a_{i,t_i}] \leftarrow$  solution of transposed
       Vandermonde system from roots  $[\theta^{e_{i,1}}, \dots, \theta^{e_{i,t_i}}]$ 
       and right-hand vector  $[Y_{i,0}, \dots, Y_{i,t_i-1}]$ 
10     $f_i \leftarrow a_{i,1}x^{e_{i,1}} + \dots + a_{i,t_i}x^{e_{i,t_i}}$ 

```

powering, computing the set $\{\theta^d \mid d \in S\}$ on Line 3 requires $O(\#S \log n)$ field operations.

If f_i has at most $t_i \leq s$ nonzero terms, then the minpoly Γ_i computed on Line 2 is a degree- t_i polynomial with exactly t_i distinct roots. If $\text{supp}(f_i) \subseteq S$, then by Lemma 5.1, FindRoots correctly finds all these roots. Because these are the only terms in f_i , solving the transposed Vandermonde system correctly determines all corresponding coefficients; see [24] or [21] for details. $\square$

6 PERFORMING EVALUATIONS

Consider the matrix multiplication with errors problem, recovering a sparse error matrix $E \in F^{r \times n}$ according to the equation $E = C - AB$. In this section, we show how to treat each row of E as a sparse polynomial and evaluate each of these at consecutive powers of a high-order element θ , as needed for the sparse interpolation algorithm from Section 5. The same techniques will also be used in the matrix inverse with errors algorithm.

Consider a column vector of powers of an indeterminate, $\mathbf{x} = [1, x, x^2, \dots, x^{n-1}]^T$. The matrix-vector product $E\mathbf{x}$ then consists of m polynomials, each degree less than n , and with a 1-1 correspondence between nonzero terms of the polynomials and nonzero entries in the corresponding row of E .

Evaluating every polynomial in $E\mathbf{x}$ at a point $\theta \in F$ is the same as multiplying E by a vector $\mathbf{v} = [1, \theta, \theta^2, \dots, \theta^{n-1}]^T$. Evaluating each polynomial in $E\mathbf{x}$ at the first $2s$ powers $1, \theta, \theta^2, \dots, \theta^{2s-1}$ allows for the unique recovery of all t -sparse rows, according to Theorem 5.2. This means multiplying E times an evaluation matrix $V \in F^{n \times 2s}$ such that $V_{i,j} = \theta^{ij}$.

Consider a single row vector $\mathbf{u} \in F^{1 \times n}$ with c nonzero entries. Multiplying $\mathbf{u}V$ means evaluating a c -sparse polynomial at the first $2s$ powers of θ ; we can view it as removing the $n - c$ zero entries from $\mathbf{u}$ to give $\mathbf{u}' \in F^{1 \times c}$, and removing the same $n - c$ rows from

Algorithm 4: DiffEval(A, B, C, θ, s)

Input: Matrices $A \in \mathbb{F}^{r \times \ell}$, $B \in \mathbb{F}^{\ell \times n}$, $C \in \mathbb{F}^{r \times n}$, high-order element $\theta \in \mathbb{F}$, and integer $s \in \mathbb{N}$

Output: Matrix $Y = (C - AB)V \in \mathbb{F}^{r \times 2s}$, where V is the $r \times 2s$ evaluation matrix given by $(\theta^{ij})_{0 \leq i < n, 0 \leq j < 2s}$.

- 1 Pre-compute θ^j for all indices j of nonzero columns in B or C
- 2 $Y_C \leftarrow CV$ using the pre-computed θ^j 's and fast transposed Vandermonde applications row by row
- 3 $Y_B \leftarrow BV$ similarly
- 4 $Y_{AB} \leftarrow AY_B$ using dense or sparse multiplication, whichever is faster
- 5 **return** $Y_C - Y_{AB}$

V to get $V' \in \mathbb{F}^{c \times 2s}$, and then computing the product $u'V'$. This evaluation matrix with removed rows V' is actually a *transposed Vandermonde matrix* built from the entries θ^i for each index i of a nonzero entry in u .

An explicit algorithm for this transposed Vandermonde matrix application is given in Section 6.2 of Bostan et al. [7], and they show that the complexity of multiplying $u'V'$ is $\tilde{O}(c + s)$ field operations, and essentially amounts to a power series inversion and product tree traversal. (See also Section 5.1 of van der Hoeven and Lecerf [21] for a nice description of how this approach is used for fast evaluation of a sparse polynomial at consecutive powers.) We repeat this for each row in a given matrix M to compute AV .

Now we are ready to show how to compute $EV = (C - AB)V$ efficiently, using the approach just discussed to compute CV and BV before multiplying A times BV and performing a subtraction. As we will show, the rectangular multiplication of A times BV is the only step which is not softly-linear time, and this dominates the complexity.

LEMMA 6.1. *Algorithm 4 DiffEval always returns the correct matrix of evaluations $(C - AB)V$ and uses*

$$\tilde{O}\left(\#B + \#C + n + s \cdot \min\left(\#A + r, \frac{r\ell}{\min(r, s, \ell)^{3-\omega}}\right)\right)$$

field operations.

PROOF. Correctness comes from the discussion above and the correctness of the algorithm in section 6.2 of Bostan et al. [7] which is used on Lines 2 and 3.

The cost of pre-computing all powers θ^j using binary powering is $O(n \log n)$ field operations, since the number of nonzero columns in B or C is at most n . Using these precomputed values, each transposed Vandermonde apply costs $\tilde{O}(k + s)$, where k is the number of nonzero entries in the current row. Summing over all rows of C and B gives $\tilde{O}(\#B + \#C + rs + \ell s)$ for these steps.

The most significant step is the rectangular multiplication of $A \in \mathbb{F}^{r \times \ell}$ times $Y_B \in \mathbb{F}^{\ell \times s}$. Using sparse multiplication, the cost is $O(\#A \cdot s)$, by Corollary 3.3. Using dense multiplication, the cost is $O(r\ell s / \min(r, \ell, s)^{\omega-3})$ according to Fact 3.1.

Because all the relevant parameters $\#A, s, r, \ell$ are known before the multiplication, we assume that the underlying software makes the best choice among these two algorithms.

Note that rs and ℓs are definitely dominated by the cost of the dense multiplication. In the sparse multiplication case, the situation is more subtle when $\#A < \ell$. But in this case, we just suppose that rows of BV which are not used in the sparse multiplication $A(BV)$ are never computed, so the complexity is as stated. $\square$

7 MULTIPLICATION WITH ERROR CORRECTION ALGORITHM

We now have all the necessary components to present the complete multiplication algorithm, Algorithm 5 MultiplyEC. The general idea is to repeatedly use FindNonzeroRows to determine the locations on nonzeros in E , then DiffEval and MultiSparseInterp to recover half of the nonzero rows at each step.

Setting the target sparsity of each recovered row to $s = \lceil 2k/r \rceil$, where r is the number of nonzero rows, ensures that at least half of the rows are recovered at each step. To get around the fact that the number of errors k is initially unknown, we start with an initial guess of $k = 1$, and double this guess whenever fewer than half of the remaining rows are recovered.

The procedure is probabilistic of the Monte Carlo type *only* because of the Monte Carlo subroutine FindNonzeroRows to determine which rows are erroneous. The other parts of the algorithm — computing evaluations and recovering nonzero entries — are deterministic based on the bounds they are given.

THEOREM 7.1. *Algorithm 5 MultiplyEC finds all errors in C with probability at least $1 - \epsilon$ and always uses*

$$\tilde{O}\left(m + \left\lceil \log_{\#F} \frac{1}{\epsilon} \right\rceil (n + t) + k \cdot \min\left(\left\lceil \frac{t}{r} \right\rceil, \ell / \min(\ell, r, \frac{k}{r})^{3-\omega}\right)\right)$$

field operations, where k is the actual number of errors in the given product C . Otherwise, it uses $\tilde{O}(m\ell n / \min(m, \ell, n)^{3-\omega})$ field operations and may return an incorrect result.

PROOF. The only probabilistic parts of the algorithm are the calls to FindNonzeroRows, which can only fail by incorrectly returning too few rows. If this never happens, then each iteration through the while loop either (a) discovers that the value of k was too small and increases it, or (b) recovers half of the rows or columns of E .

So the total number of iterations of the while loop if the calls to FindNonzeroRows are never incorrect is at most $\lceil \log_2 k \rceil + \lceil \log_2 m \rceil + \lceil \log_2 n \rceil$. By the union bound, from the way that ϵ' is computed and the fact that there are two calls to FindNonzeroRows on every iteration, the probability that FindNonzeroRows is never incorrect is at least $1 - \epsilon$.

In this case, the rest of the correctness and the running time follow directly from Lemmas 4.1 and 6.1 and Theorem 5.2.

If a call to FindNonzeroRows returns an incorrect result, two bad things can happen. First, if $\mathcal{J}$ is too small during some iteration, all of the evaluations may be incorrect, leading to k being increased. In extremely unlucky circumstances, this may happen $O(\log(rn))$ times until the product is naïvely recomputed on Line 19, leading to the worst-case running time.

The second thing that can occur when a call to FindNonzeroRows fails is that it may incorrectly report all errors have been found,

Algorithm 5: MultiplyEC($A, B, C, \theta, \epsilon$)

Input: Matrices $A \in \mathbb{F}^{m \times \ell}$, $B \in \mathbb{F}^{\ell \times n}$, $C \in \mathbb{F}^{m \times n}$, high-order element $\theta \in \mathbb{F}$, and error bound $0 < \epsilon < 1$
Output: Matrix $E \in \mathbb{F}^{m \times n}$ such that, with probability at least $1 - \epsilon$, $AB = C - E$

```

1  $k \leftarrow 1$ 
2  $E \leftarrow \mathbf{0}^{m \times n}$ 
3  $\epsilon' \leftarrow \epsilon / (4 \lceil \log_2(mn) \rceil + 1)$ 
4  $\mathcal{J} \leftarrow \text{FindNonzeroRows}(V \mapsto CV - A(BV), \epsilon')$ 
5 while  $\#\mathcal{J} \geq 1$  do
6   if  $\#\text{FindNonzeroRows}(V \mapsto$ 
7      $(V^T(C - E) - (V^T A)B)^T, \epsilon') > \#\mathcal{J}$  then
8     Transpose  $C$  and  $E$ , swap and transpose  $A$  and  $B$ ,
9     replace  $\mathcal{J}$ 
10     $r \leftarrow \#\mathcal{J}$ 
11     $A' \leftarrow$  submatrix of  $A$  from rows in  $\mathcal{J}$ 
12     $C' \leftarrow$  submatrix of  $(C - E)$  from rows in  $\mathcal{J}$ 
13     $s \leftarrow \lceil 2(k - \#E)/r \rceil$ 
14     $Y \leftarrow \text{DiffEval}(A', B, C', \theta, s)$ 
15     $f_1, \dots, f_r \leftarrow \text{MultiSparseInterp}(r, n, s, \theta, Y)$ 
16    for  $i \leftarrow 1, 2, \dots, r$  do
17      Set  $(\mathcal{J}_i, e)$ th entry of  $E$  to  $c$  for each term  $cx^e$  of  $f_i$ 
18       $\mathcal{J} \leftarrow \text{FindNonzeroRows}(V \mapsto (C - E)V - A(BV), \epsilon')$ 
19      if  $\#\mathcal{J} > r/2$  then
20         $k \leftarrow 2k$ 
21        if  $k \geq 2n\#J$  then return  $C - AB$ 
22      foreach  $i \in \mathcal{J}$  do
23        Clear entries from row  $i$  of  $E$  added on this iteration
24 return  $E$ 

```

leading to the small $1 - \epsilon$ chance that the algorithm returns an incorrect result. $\square$

8 INVERSE ERROR CORRECTION ALGORITHM

As discussed in Section 2, our algorithm for correcting errors in a matrix inverse follows mostly the same outline as that for correcting a matrix product, with two important changes. These are the basis for the next two lemmas.

LEMMA 8.1. *For any matrices $E, A \in \mathbb{F}^{n \times n}$ where A is invertible, a call to $\text{FindNonzeroRows}(V \mapsto EAV, \epsilon)$ correctly returns the nonzero rows of E with probability at least $1 - \epsilon$.*

PROOF. Recall from the proof of Lemma 4.1 that the correctness of FindNonzeroRows depends on applying a matrix $V \in \mathbb{F}^{n \times \ell}$ of random entries to the unknown matrix E .

In the current formulation, instead of applying the random matrix V directly to E , instead the product AV is applied. But because A is nonsingular, there is a 1-1 correspondence between the set of all matrices $V \in \mathbb{F}^{n \times \ell}$ and matrices the set $\{AV \mid V \in \mathbb{F}^{n \times \ell}\}$. That is, applying a random matrix to EA is the same as applying a random matrix to E itself, and therefore the same results hold. $\square$

Algorithm 6: InverseEC(A, B, θ, ϵ)

Input: Matrices $A, B \in \mathbb{F}^{n \times n}$, high-order element $\theta \in \mathbb{F}$, and error bound $0 < \epsilon < 1$
Output: Matrix $E \in \mathbb{F}^{n \times n}$ such that, with probability at least $1 - \epsilon$, $A^{-1} = B + E$

```

1  $k \leftarrow 1$ 
2  $E \leftarrow \mathbf{0}^{n \times n}$ 
3  $\epsilon' \leftarrow \epsilon / (8 \lceil \log_2 n \rceil + 1)$ 
4  $\mathcal{J} \leftarrow \text{FindNonzeroRows}(V \mapsto V - B(AV), \epsilon')$ 
5 while  $\#\mathcal{J} \geq 1$  do
6   if  $\#\text{FindNonzeroRows}(V \mapsto V - A((B + E)V), \epsilon') > \#\mathcal{J}$ 
7     then
8     Transpose  $A$ ,  $B$ , and  $E$ , and replace  $\mathcal{J}$ 
9      $r \leftarrow \#\mathcal{J}$ 
10     $A' \leftarrow$  submatrix of  $A$  from columns in  $\mathcal{J}$ 
11     $\mathcal{J}' \leftarrow$  set of  $r$  linearly independent rows in  $A'$ 
12    according to Lemma 8.2
13     $I', A'' \leftarrow$  submatrices of  $I, A$  from the rows chosen in
14     $\mathcal{J}'$ 
15     $X^{-1} \leftarrow$  inverse of submatrix of  $A'$  according to
16    Lemma 8.2
17     $s \leftarrow \lceil 2(k - \#E)/r \rceil$ 
18     $Y \leftarrow \text{DiffEval}(A'', B, I', \theta, s)$ 
19     $Y' \leftarrow X^{-1}Y$ 
20     $f_1, \dots, f_r \leftarrow \text{MultiSparseInterp}(r, n, s, \theta, Y')$ 
21    for  $i \leftarrow 1, 2, \dots, r$  do
22      Set  $(\mathcal{J}_i, e)$ th entry of  $E$  to  $c$  for each term  $cx^e$  of  $f_i$ 
23       $\mathcal{J} \leftarrow \text{FindNonzeroRows}(V \mapsto V - (B + E)(AV), \epsilon')$ 
24      if  $\#\mathcal{J} > r/2$  then
25         $k \leftarrow 2k$ 
26        if  $k \geq 2n\#J$  then return  $A^{-1} - B$ 
27      foreach  $i \in \mathcal{J}$  do
28        Clear entries from row  $i$  of  $E$  added on this iteration
29 return  $E$ 

```

LEMMA 8.2. *Given any rank- r matrix $A \in \mathbb{F}^{n \times r}$, it is possible to compute a matrix $X \in \mathbb{F}^{r \times r}$ formed from a subset of the rows of A , and its inverse X^{-1} , using $\tilde{O}(\#A + r^\omega)$ field operations.*

PROOF. This is a direct consequence of [11, Theorem 2.11], along with the classic algorithm of [9] for fast matrix inversion. Alternatively, one could use the approach of [34] to compute the lexicographically minimal set of linearly independent rows in A , as well as a representation of the inverse, in the same running time. $\square$

The resulting algorithm for matrix inversion with errors is presented in Algorithm 6 InverseEC.

THEOREM 8.3. *Algorithm 6 InverseEC finds all errors in B with probability at least $1 - \epsilon$, and always uses*

$$\tilde{O}\left(\left\lceil \log_{\#F} \frac{1}{\epsilon} \right\rceil t + nk / \min(r, \frac{k}{r})^{3-\omega} + r^\omega\right)$$

field operations, where k is the actual number of errors in the given product C . Otherwise, it uses $\tilde{O}(n^\omega)$ field operations and may return an incorrect result.

PROOF. In this algorithm, we make use of two different formulas for E :

$$EA = I - BA \quad (1)$$

$$AE = I - AB \quad (2)$$

The first formula (1) is used to determine the nonzero rows of E on Lines 4 and 19, and the second formula (2) is used to evaluate the rows of XE on Line 14.

The main difference in this algorithm compared to `MultiplyEC` is the need to find r nonzero rows of A' to constitute the matrix X and compute its inverse X^{-1} . According to Lemma 8.2 these both cost $\tilde{O}(r^\omega)$, which gives the additional term in the complexity statement. Note that this also eliminates the utility of sparse multiplication in all cases, simplifying the complexity statement somewhat compared to that of `MultiplyEC`.

The rest of the proof is identical to that of Theorem 7.1. $\square$

ACKNOWLEDGMENTS

This work was performed while the author was graciously hosted by the Laboratoire Jean Kuntzmann at the Université Grenoble Alpes. Thanks to many colleagues for fruitful discussions on these topics, especially Clément Pernet and Jean-Guillaume Dumas.

This work was supported in part by the National Science Foundation under grants 1319994 (https://www.nsf.gov/awardsearch/showAward?AWD_ID=1319994) and 1618269 (https://www.nsf.gov/awardsearch/showAward?AWD_ID=1618269).

REFERENCES

- [1] R. R. Amossen and R. Pagh. Faster join-projects and sparse matrix multiplications. In *ICDT '09*, pages 121–126, New York, NY, USA, 2009. ACM. doi:10.1145/1514894.1514909.
- [2] A. Arnold, M. Giesbrecht, and D. S. Roche. Faster sparse multivariate polynomial interpolation of straight-line programs. *J. Symbolic Computation*, 2015. doi:10.1016/j.jsc.2015.11.005.
- [3] M. Ben-Or and P. Tiwari. A deterministic algorithm for sparse multivariate polynomial interpolation. In *STOC '88*, pages 301–309, New York, NY, USA, 1988. ACM. doi:10.1145/62212.62241.
- [4] D. J. Bernstein. Factoring into coprimes in essentially linear time. *J. Algorithms*, 54(1):1–30, 2005. doi:10.1016/j.jalgor.2004.04.009.
- [5] J. Böhm, W. Decker, C. Fieker, and G. Pfister. The use of bad primes in rational reconstruction. *Math. Comp.*, 84(296):3013–3027, 2015. doi:10.1090/mcom/2951.
- [6] A. Borodin and I. Munro. *The computational complexity of algebraic and numeric problems*. Number 1 in Elsevier Computer Science Library; Theory of Computation Series. American Elsevier Pub. Co., New York, 1975.
- [7] A. Bostan, G. Lecerc, and E. Schost. Tellegen's principle into practice. In *ISSAC '03*, pages 37–44. ACM, 2003. doi:10.1145/860854.860870.
- [8] B. Boyer and E. L. Kaltofen. Numerical linear system solving with parametric entries by error correction. In *SNC '14*, pages 33–38. ACM, 2014. doi:10.1145/2631948.2631956.
- [9] J. R. Bunch and J. E. Hopcroft. Triangular factorization and inversion by fast matrix multiplication. *Mathematics of Computation*, 28(125):231–236, 1974. URL <http://www.jstor.org/stable/2005828>.
- [10] D. G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28:693–701, 1991. doi:10.1007/BF01178683.
- [11] H. Y. Cheung, T. C. Kwok, and L. C. Lau. Fast matrix rank algorithms and applications. *J. ACM*, 60(5):31:1–31:25, Oct. 2013. doi:10.1145/2528404.
- [12] M. T. Comer, E. L. Kaltofen, and C. Pernet. Sparse polynomial interpolation and Berlekamp/Massey algorithms that correct outlier errors in input values. In *ISSAC '12*, pages 138–145, New York, NY, USA, 2012. ACM. doi:10.1145/2442829.2442852.
- [13] J.-G. Dumas and E. Kaltofen. Essentially optimal interactive certificates in linear algebra. In *ISSAC '14*, pages 146–153. ACM, 2014. doi:10.1145/2608628.2608644.
- [14] J.-G. Dumas, P. Giorgi, and C. Pernet. Dense linear algebra over word-size prime fields: the FFLAS and FFPACK packages. *ACM Trans. Math. Softw.*, 35:19:1–19:42, October 2008. doi:10.1145/1391989.1391992.
- [15] J.-G. Dumas, E. Kaltofen, E. Thomé, and G. Villard. Linear time interactive certificates for the minimal polynomial and the determinant of a sparse matrix. In *ISSAC '16*, pages 199–206. ACM, 2016. doi:10.1145/2930889.2930908.
- [16] J.-G. Dumas, D. Lucas, and C. Pernet. Certificates for triangular equivalence and rank profiles. In *ISSAC '17*, pages 133–140. ACM, 2017. doi:10.1145/3087604.3087609.
- [17] R. Freivalds. Fast probabilistic algorithms. In J. Bečvář, editor, *Mathematical Foundations of Computer Science 1979*, pages 57–69. Springer Berlin Heidelberg, 1979.
- [18] L. Gasieniec, C. Levcopoulos, A. Lingas, R. Pagh, and T. Tokuyama. Efficiently correcting matrix products. *Algorithmica*, 79(2):428–443, Oct. 2017. doi:10.1007/s00453-016-0202-3.
- [19] O. Goldreich, D. Ron, and M. Sudan. Chinese remaindering with errors. In *STOC '99*, pages 225–234. ACM, 1999. doi:10.1145/301250.301309.
- [20] D. Harvey, J. van der Hoeven, and G. Lecerc. Faster polynomial multiplication over finite fields. *J. ACM*, 63(6):52:1–52:23, Jan. 2017. doi:10.1145/3005344.
- [21] J. van der Hoeven and G. Lecerc. On the bit-complexity of sparse polynomial and series multiplication. *J. Symbolic Computation*, 50:227–0254, 2013. doi:10.1016/j.jsc.2012.06.004.
- [22] Q. Huang and X. Gao. Faster deterministic sparse interpolation algorithms for straight-line program multivariate polynomials. *CoRR*, abs/1709.08979, 2017. URL <http://arxiv.org/abs/1709.08979>.
- [23] S. M. M. Javadi and M. Monagan. Parallel sparse polynomial interpolation over finite fields. In *PASCO '10*, pages 160–168, New York, NY, USA, 2010. ACM. doi:10.1145/1837210.1837233.
- [24] E. Kaltofen and L. Yagati. Improved sparse multivariate polynomial interpolation algorithms. In P. Gianni, editor, *Symbolic and Algebraic Computation*, volume 358 of *Lecture Notes in Computer Science*, pages 467–474. Springer Berlin / Heidelberg, 1989. doi:10.1007/3-540-51084-2_44.
- [25] E. L. Kaltofen. Fifteen years after DSC and WLSS2: What parallel computations I do today [invited lecture at PASCO 2010]. In *PASCO '10*, pages 10–17, New York, NY, USA, 2010. ACM. doi:10.1145/1837210.1837213.
- [26] E. L. Kaltofen and Z. Yang. Sparse multivariate function recovery from values with noise and outlier errors. In *ISSAC '13*, pages 219–226. ACM, 2013. doi:10.1145/2465506.2465524.
- [27] E. L. Kaltofen, M. Nehring, and B. D. Saunders. Quadratic-time certificates in linear algebra. In *ISSAC '11*, pages 171–176. ACM, 2011. doi:10.1145/1993886.1993915.
- [28] E. L. Kaltofen, C. Pernet, A. Storjohann, and C. Waddell. Early termination in parametric linear system solving and rational function vector recovery with error correction. In *ISSAC '17*, pages 237–244. ACM, 2017. doi:10.1145/3087604.3087645.
- [29] M. Khonji, C. Pernet, J.-L. Roch, T. Roche, and T. Stalinski. Output-sensitive decoding for redundant residue systems. In *ISSAC '10*, pages 265–272. ACM, 2010. doi:10.1145/1837934.1837985.
- [30] F. Le Gall. Faster algorithms for rectangular matrix multiplication. In *FOCS '12*, pages 514–523, Oct. 2012. doi:10.1109/FOCS.2012.80.
- [31] F. Le Gall. Powers of tensors and fast matrix multiplication. In *ISSAC '14*, pages 296–303, New York, NY, USA, 2014. ACM. doi:10.1145/2608628.2608664.
- [32] A. Lingas. A fast output-sensitive algorithm for boolean matrix multiplication. In A. Fiat and P. Sanders, editors, *Algorithms - ESA 2009*, pages 408–419. Springer Berlin Heidelberg, 2009. doi:10.1007/978-3-642-04128-0_37.
- [33] R. Pagh. Compressed matrix multiplication. *ACM Trans. Comput. Theory*, 5(3):9:1–9:17, Aug. 2013. doi:10.1145/2493252.2493254.
- [34] A. Storjohann and S. Yang. A relaxed algorithm for online matrix inversion. In *ISSAC '15*, pages 339–346, New York, NY, USA, 2015. ACM. doi:10.1145/2755996.2756672.
- [35] Q. Wang, X. Zhang, Y. Zhang, and Q. Yi. Augem: Automatically generate high performance dense linear algebra kernels on x86 cpus. In *Proc. Int'l Conference on High Performance Computing, Networking, Storage and Analysis*, SC '13, pages 25:1–25:12. ACM, 2013. doi:10.1145/2503210.2503219.
- [36] R. Yuster and U. Zwick. Fast sparse matrix multiplication. *ACM Trans. Algorithms*, 1(1):2–13, July 2005. doi:10.1145/1077464.1077466.