

Faster algorithms for the characteristic polynomial

Clément Pernet

Arne Storjohann

David R. Cheriton School of Computer Science
University of Waterloo, Ontario,
Canada N2L 3G1
{cpernet,astorjoh}@uwaterloo.ca

ABSTRACT

A new randomized algorithm is presented for computing the characteristic polynomial of an $n \times n$ matrix over a field. Over a sufficiently large field the asymptotic expected complexity of the algorithm is $O(n^\theta)$ field operations, improving by a factor of $\log n$ on the worst case complexity of Keller-Gehrig's algorithm [8].

1. INTRODUCTION

Computing the characteristic polynomial of an $n \times n$ matrix A over a field F is a classical problem. Keller-Gehrig [8] gave three reductions of the problem to matrix multiplication. Let θ be an admissible exponent for the complexity of matrix multiplication: $O(n^\theta)$ operations from F are sufficient to multiply together two $n \times n$ matrices over F . In this paper all complexity bounds are in terms of field operations from F and we make the common assumption that $\theta > 2$.

Keller-Gehrig's third algorithm has cost $O(n^\theta)$ but only works for input matrices with restrictive genericity requirements. His first algorithm, a simplified version of the second, also only works for input matrices satisfying certain requirements. Of primary interest here is his second algorithm which works for all input matrices and has a worst case cost of $O(n^\theta \log n)$. The extra $\log n$ factor arises because the algorithm computes $A^2, A^4, A^8, \dots, A^{\lceil \log_2 n \rceil}$ using binary powering.

Computing the characteristic polynomial is closely related to other problem such as computing the minimal polynomial, testing two matrices for similarity, and computing the Frobenius canonical form. Known reductions to matrix multiplication for these problems, both deterministic [10, 11] and probabilistic [4, 5, 6], all have an extra $\log n$ factor in their worst case complexity bounds, arising because Keller-Gehrig's algorithm is used as a subroutine directly [5, 6, 10, 11] or because a logarithmic number of powers of A might be computed [4].

In this paper we combine ideas from [6, 8, 12] to get a

new randomized algorithm for computing the characteristic polynomial. If F has at least $2n^2$ elements our Las Vegas algorithm has expected cost $O(n^\theta)$, matching the lower bound for this problem. Unlike Keller-Gehrig's $O(n^\theta \log n)$ algorithm, we proceed in phases for $k = 1, 2, 3, \dots, n$ and thus converges arithmetically.

In Section 2 we introduce some notation and recall some facts about Krylov matrices. Section 3 gives a worked example of the new algorithm and offers an overview of Sections 4–6 which are devoted to presenting the algorithm and proving correctness. The new algorithm is not only of theoretical interest but also practical. In Section 7 we describe an implementation, present some timings, and compare with the previously most efficient implementations that we are aware of. Section 8 concludes.

2. NOTATION

We will frequently write matrices using a conformal block decomposition. A block is a submatrix comprised of a contiguous sequence of rows and columns. A block may be a single matrix entry or may have row or column dimension zero. The generic block label $*$ denotes that a block is possibly nonzero. Blocks that are necessarily zero are left unlabelled.

In this paper a companion matrix looks like

$$C_* = \begin{bmatrix} 0 & \cdots & 0 & * \\ 1 & \ddots & \vdots & \vdots \\ & \ddots & 0 & * \\ & & 1 & * \end{bmatrix} \in K^{k \times k}. \quad (1)$$

and the sizes of companion blocks in the Frobenius canonical form are monotonically nonincreasing. Companion blocks may have dimension zero. We use the label B_* to denote a block which has all entries zero except for possibly entries in the last column. The dimension of a block labelled B_* will be conformal with adjacent blocks.

For a square matrix $A \in K^{n \times n}$ and vector $v \in K^{n \times 1}$, let $K_A(v, d)$ denote the Krylov matrix

$$[v \mid Av \mid \cdots \mid A^{d-1}v] \in K^{n \times d}.$$

For $V \in K^{n \times j}$ we denote by $\text{Orb}_A(V)$ the subspace of K^n spanned by all the column vectors in $[V \mid AV \mid A^2V \mid \dots]$.

FACT 1. *Let $A \in K^{n \times n}$ be arbitrary and $U \in K^{n \times n}$ be nonsingular. Then*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

Algorithm 2 CharPolyRec(A, n, k, x)

Require: A k -shifted form $A \in \mathbb{K}^{n \times n}$, an indeterminate x .

Ensure: return $\det xI - A$, or **fail**.

```
if  $n = k$  then
  Return  $\det(xI - A)$ 
else
   $(d_1, \dots, d_m) := \text{Extension}(A, k)$ 
  /* If the call to Extension fails then abort
  and return fail */
   $\bar{m} :=$  minimal index with  $d_{\bar{m}} < k + 1$ 
   $\bar{n} := d_1 + \dots + d_{\bar{m}}$ 
   $\begin{bmatrix} \bar{A} & B \\ C & D \end{bmatrix} := \text{Transform}(A, k, (d_1, \dots, d_m))$ 
  /* If  $C$  is not the zero matrix then abort and
  return fail */
  Return CharPolyRec( $\bar{A}, \bar{n}, k + 1, x$ )  $\times \det(xI - D)$ 
end if
```

THEOREM 2. Algorithm 2 (CharPolyRec) returns the characteristic polynomial of the input matrix or **fail**. The cost of the algorithm is $O(n^\theta)$.

PROOF. The complexity is deduced from the following arithmetic progression:

$$\sum_{k=1}^n k(n/k)^\theta = n^\theta \sum_{k=1}^n (1/k)^{\theta-1} = O(n^\theta)$$

since $\theta - 1 > 1$. \square

To ensure that the algorithm will only fail with a bounded probability, the input matrix A has to be preconditioned by a random similarity transformation. This gives the following algorithm.

Algorithm 3 CharPoly(A, n, x)

Require: A matrix $A \in \mathbb{K}^{n \times n}$, an indeterminate x .

Ensure: return $\det(xI - A)$, or **fail**.

```
/* Fail will be returned with probability at
most 1/2. We require  $\#\mathbb{K} \geq 2n^2$ . */
 $\Lambda :=$  a subset of  $\mathbb{K}$  with  $\#\Lambda \geq 2n^2$ 
Choose  $V \in \mathbb{K}^{n \times n}$  with entries uniformly and randomly
from  $\Lambda$ .
 $B := V^{-1}AV$  /* If  $V$  is singular then abort and
return fail */
Return CharPolyRec( $B, n, 1, x$ )
```

The probability analysis of Algorithm 3 (CharPoly) will be detailed in Section 6; the cost of the algorithm is obviously still $O(n^\theta)$ field operations.

6. PRECONDITIONING

Let $A \in \mathbb{K}^{n \times n}$ be an arbitrary matrix. In this subsection we prove that Algorithm 2 (CharPolyRec) will not fail when given as input the tuple $(B, n, 1, x)$, where $B = V^{-1}AV$ and V is filled with algebraically independent indeterminates. Upon specialization of the indeterminates with random field elements, as is done by Algorithm 3 (CharPoly), a bound of 1/2 on the probability of failure will follow due to the Schwartz-Zippel Lemma [9, 13].

The proof of the following theorem is similar to and inspired by [12, Proof of Proposition 6.1]. Note that for convenience we assume that the Frobenius form of A has n blocks,

some of which may be trivial (i.e., 0×0). In the statement of the theorem this means that some of the f_* and d_* may be zero.

THEOREM 3. Let $A \in \mathbb{K}^{n \times n}$ have Frobenius form with blocks of dimension $f_1 \geq \dots \geq f_n$, and let v_1, \dots, v_n be the columns of a matrix V filled with algebraically independent indeterminates. Suppose (d_1, \dots, d_n) is monotonically nonincreasing sequence of nonnegative integers. Then

$$K = [K_A(v_1, d_1) \mid \dots \mid K_A(v_n, d_n)]$$

has full column rank if and only if $\sum_{j=1}^i d_j \leq \sum_{j=1}^i f_j$ for all $1 \leq i \leq n$.

PROOF. The “only if” direction follows because for any block X of i vectors, even a generic block $X = [v_1 \mid \dots \mid v_i]$, the dimension of $\text{Orb}_A(X)$ is at most $\sum_{j=1}^i f_j$.

To prove the other direction we specialize the indeterminates in the vectors v_i . In particular, it will be sufficient to construct a full column rank matrix

$$K = [K_1 \mid \dots \mid K_n]$$

over \mathbb{K} such that each K_i is in Krylov form and has dimension d_i , $1 \leq i \leq n$. Consider a change of basis matrix $U \in \mathbb{K}^{n \times n}$ such that $U^{-1}AU$ is in Frobenius form. Then

$$U = [K_A(u_1, f_1) \mid \dots \mid K_A(u_n, f_n)]$$

is nonsingular. Let

$$\bar{K} = [\bar{K}_1 \mid \dots \mid \bar{K}_n]$$

be the submatrix of U such that each \bar{K}_i has the form

$$\bar{K}_i = [K_A(u_i, \min(f_i, d_i)) \mid E_i],$$

where E_i has dimension $d_i - \min(f_i, d_i)$, and the columns of E_1, E_2, \dots, E_n are filled with unused columns of U , using the columns in order from left to right. Then \bar{K} has full column rank and each \bar{K}_i has the correct dimension. Our goal now is to demonstrate the existence of an invertible matrix T such that $K = \bar{K}T$ has the desired form. We will construct $T = I + \sum_{i=1}^n (T_i - I)$ where each T_i is unit upper triangular. For all i with $d_i \leq f_i$ no transformation of \bar{K}_i is required: set $T_i = I$. If $f_i < d_i$ then

$$\bar{K}_i = [K_A(u_i, f_i) \mid K_A(A^{s_1}u_{j_1}, t_1) \mid \dots \mid K_A(A^{s_k}u_{j_k}, t_k)]$$

where, by construction of the E_i , we have $j_1 < j_2 < \dots < j_k$, $t_l = f_{j_l} - s_l$ for $1 \leq l \leq k - 1$, and $t_k \leq f_k$. Using the property $\sum_{j=1}^i d_j \leq \sum_{j=1}^i f_j$ we have $j_k < i$. Since (d_1, \dots, d_n) is monotonically nondecreasing and $K_A(v_l, d_l)$ is a submatrix of \bar{K}_i for $1 \leq l \leq k$, it follows that

$$s_l \geq d_i \text{ for } 1 \leq l \leq k. \quad (4)$$

We can write \bar{K}_i as the sum of the following $k + 1$ matrices:

$$\bar{K}_i = [K_A(u_i, f_i) \mid 0, \dots, 0] \quad (5)$$

$$+ \sum_{l=1}^{k-1} [0, \dots, 0 \mid K_A(A^{s_l}u_{j_l}, f_{j_l} - s_l) \mid 0, \dots, 0] \quad (6)$$

$$+ [0, \dots, 0 \mid K_A(A^{s_k}u_{j_k}, t_k)] \quad (7)$$

To bring the matrix in (5) to Krylov form we may add suitable linear combinations of the first f_i columns to the last $d_i - f_i$ columns to obtain

$$[K_A(u_i, f_i) \mid K_A(A^{f_i}u_i, d_i - f_i)].$$

This is possible since the i 'th invariant subspace has dimension f_i . Denote by $T_i^{(1)}$ the unit upper triangular matrix which effects this transformation on \bar{K} .

Now consider the matrix in (7). The Krylov space needs to be extended on the left to fill in the zero columns as follows:

$$\left[K_A(A^s u_{j_k}, s_k - s) \mid K_A(A^{s_k} u_{j_k}, t_k) \right].$$

From (4) we may conclude that $s \geq 0$. Since $K_A(A^s u_{j_k}, s_k - s)$ is a submatrix of $\left[\bar{K}_1 \mid \cdots \mid \bar{K}_{i-1} \right]$, we need only copy former to latter columns. Denote by $T_i^{(2)}$ the unit upper triangular matrix which effects the copying on these columns. Similarly, there exists a unit upper triangular matrix $T_i^{(3)}$ which extends the Krylov sequence of the matrix in (6) to the left and right. Let $T_i = T_i^{(1)} + T_i^{(2)} + T_i^{(3)}$. \square

In the following corollary the matrix A and V are as in Theorem 3, that is, $A \in \mathbb{K}^{n \times n}$ has Frobenius form with blocks of dimension $f_1 \geq f_2 \geq \cdots \geq f_n$ and V is an $n \times n$ matrix filled with indeterminates. The corollary follows as a result of Fact 1.2.

COROLLARY 1. *Let $B := V^{-1}AV$ and k satisfy $2 \leq k \leq n$. The lexicographically maximal sequence (d_1, \dots, d_n) of nonnegative integers such that:*

- $d_i \leq k$ for all $1 \leq i \leq n$, and
- $K = \left[K_B(e_1, d_1) \mid \cdots \mid K_B(e_n, d_n) \right]$ has full column rank,

will satisfy $d_1 + \cdots + d_n = n$ and can be written as

$$(d_1, \dots, d_n) = (k, \dots, k, d_{\bar{m}}, f_{\bar{m}+1}, f_{\bar{m}+2}, \dots, f_n)$$

with $k > d_{\bar{m}} \geq f_{\bar{m}+1}$. Moreover,

$$K^{-1}BK = \left[\begin{array}{c|c} \bar{A} & B \\ \hline & D \end{array} \right]$$

is in shifted Hessenberg form, where \bar{A} is in $(k+1)$ -shifted form of dimension $\bar{n} = d_1 + \cdots + d_{\bar{m}}$, and D is in Hessenberg form.

Each entry of $VK = \left[K_A(v_1, d_1) \mid \cdots \mid K_A(v_n, d_n) \right]$ is a linear combination of indeterminates of V . It follows that the determinant of VK is a nonzero polynomial in the indeterminates of V with total degree at most n .

Let $K_1 = I_n$ and K_i be the matrix of Corollary 1 for $k = i$, $2 \leq i \leq n$. Given as input $(B, n, 1, x)$, Algorithm 2 (**CharPolyRec**) will perform a change of basis at each step and compute the structured Krylov extension $K_{i-1}^{-1}K_i$ for $i = 2, 3, \dots, n$. Let Δ be the product of the determinant of V and each matrix VK_i . Then Δ is a nonzero polynomial of total degree bounded by n^2 . The next result now follows from the Schwartz-Zippel lemma.

THEOREM 4. *Algorithm 2 (**CharPoly**) will return fail with probability at most $1/2$.*

We remark that the randomized Frobenius form algorithms in [5, 6] rely on the fact that the diagonal blocks in the Hessenberg form $K_n^{-1}BK_n$ will be those of the Frobenius form of A , and thus require that the determinant of the single matrix VK_n not vanish upon specialization of V with random field elements.

7. IMPLEMENTATION

In this section we discuss an implementation of the new characteristic polynomial algorithm that is modified to perform the preconditioning step more efficiently in practice. Actually, the algorithm is adaptive and involves a parameter that is highly architecture-dependant and must be set experimentally. We present experiments comparing the practical performance of our implementation with several other softwares.

The implementation we describe here makes use of the FFLAS-FFPACK library¹. This C++ library provides the efficient basic routines such as matrix multiplications and LQUP decomposition that make use of the level 3 BLAS numerical routines [1, 2].

7.1 Efficient preconditioning

Although it does not affect the asymptotic complexity, the preconditioning phase $V^{-1}AV$ of Algorithm 3 (**CharPoly**) is expensive in practice. This preconditioning phase can also be achieved by modifying Algorithm 2 (**CharPolyRec**) to compute the first Krylov extension using random vectors from Λ instead of identity vectors.

Our heuristic for this preconditioning step is to compute a block Krylov matrix $M = [U|AU|\dots|A^{c-1}U]$ where U is formed by $\lceil n/c \rceil$ random vectors, for some parameter c . If this matrix is non singular, then the matrix $M^{-1}AM$ will be in c -shifted form (up to row and column permutations) and Algorithm 2 (**CharPolyRec**) can be called with shift parameter $k = c$ instead of $k = 1$. If $r = \text{rank}(M) < n$ then the linearly independent columns of M can be completed into a non singular matrix \bar{M} by adding $n - r$ columns at the end, and we obtain the block upper triangular matrix

$$\bar{M}^{-1}A\bar{M} = \left[\begin{array}{c|c} H_c & * \\ \hline & R \end{array} \right]$$

where the $r \times r$ matrix H_c is in c -shifted form (up to row and column permutations). The characteristic polynomial of this matrix is computed by two recursive calls on the diagonal blocks H_c and R . Algorithm 4 (**CharPoly**) gives the algorithm with this modified preconditioning step.

Further explanations on the completion of M into \bar{M} using the LQUP decomposition can be found in [3]. Note that again, only c columns of the matrix H_c have to be computed, which makes the computation of B much cheaper.

As c gets larger, the slices of the block Krylov matrix K become smaller. In the extreme case $c = n$, the algorithm computes the usual Krylov matrix of only one vector. In this case, the algorithm is equivalent to the algorithm **LU-Krylov** presented in [3, algorithm 2.2]. Assuming $\theta = 3$ the leading constant of algorithm **LU-Krylov** is competitive ($2.66n^3$) but the algorithm does not fully exploit matrix multiplication. At the opposite, the case $c = 2$ corresponds to Algorithm 3 (**CharPoly**): it reduces the problem fully to matrix multiplication. The preconditioning parameter c makes it possible to balance the computation between these two algorithms.

Figure 1 displays the computation time of the algorithm for different values of c . Three matrices of order 5000 are used: they differ in the number of blocks in their Frobenius form. For $c < 55$, the timings are decreasing when

¹This library is available online at <http://www-ljk.imag.fr/membres/Jean-Guillaume.Dumas/FFLAS> or within the LinBox library <http://www.linalg.org>

Algorithm 4 CharPoly(A, n, x)

Require: A matrix $A \in \mathbb{K}^{n \times n}$, an indeterminate x , a preconditioning parameter c .

Ensure: $\det(xI - A)$, or fail.

/ Fail will be returned with probability at*

*most 1/2 if $\#\mathbb{K} > 2n^2$ */*

$\Lambda :=$ a subset of \mathbb{K} with $\#\Lambda \geq 2n^2$

$m := \lceil n/c \rceil$

Choose $V \in \mathbb{K}^{n \times m}$ with entries uniformly and randomly from Λ .

Compute the $n \times (c \lceil n/c \rceil)$ matrix

$$M = [V|AV|\dots|A^{c-1}V]$$

Compute (L, Q, U, P) , the LQUP decomposition of M^T .

Let $r = \text{rank}(M^T)$

$$\overline{M} := \left[MQ \left[\begin{array}{c|c} I_r & 0 \end{array} \right] \middle| P^T \left[\begin{array}{c} 0 \\ I_{n-r} \end{array} \right] \right]$$

$$B := \overline{M}^{-1} A \overline{M} = \begin{bmatrix} H_c & * \\ & R \end{bmatrix}$$

Return CharPolyRec(H_c, n, c, x) \times CharPolyRec($R, n, 0, x$)




Figure 1: Finding the optimal preconditioning parameter c for matrices of order 5000, Itanium2-64 1.3Ghz, 192Gb

c increases, which shows the advantage of using the block Krylov preconditioning for a large enough value for c . Then the timings increase again for larger c . In these cases, the dominant operation is the computation of the block Krylov matrix M by many matrix multiplications of uneven dimensions. The matrix multiplication routine used will be more efficient for computing one $n \times n$ by $n \times n$ product rather than $c n \times n$ by $n \times n/c$ products, due to both the level 3 BLAS behaviour and the use of sub-cubic matrix multiplication. The optimal value $c = 55$ gives here the best timings. This value is not only depending the matrix dimension, but also on the architecture and the BLAS that are used, since it is linked with the ratio between the efficiency of the matrix vector product and the matrix matrix multiplication.

Note that the algorithm gets faster as the dimension of the largest block decreases.

7.2 Timing comparisons

We now compare the running time of our implementation of Algorithm 4 CharPoly with that of other state of the art implementations of characteristic polynomial algorithms. The routine LU-Krylov, available in the FFLAS-FFPACK and LinBox, libraries was shown to be the most efficient implementation in most cases [3].

For all the following experiments, we used the finite field $\mathbb{Z}/(547909)$. On one hand, it is large enough to ensure a high probability of success; none of the computations returned fail. On the other hand, the field size is small enough so that the FFLAS-FFPACK routines can make efficient use of the level 3 BLAS subroutines, using delayed modular reductions with the 53 bits of the double mantissa.

n	LU-Krylov	New algorithm
200	0.024	0.032
300	0.06s	0.088s
500	0.248s	0.316s
750	1.084s	1.288s
1000	2.42s	2.296s
5000	267.6s	153.9s
10 000	1827s	991s
20 000	14 652s	7097s
30 000	48 887s	24 928s

Table 1: Computation time for 1 Frobenius block matrices, Itanium2-64 1.3Ghz, 192Gb

Table 1 presents the timings for the computation of the characteristic polynomial of matrices having only one block on their Frobenius form. The preconditioning parameter c has been set to 100 for these experiments. The new algorithm improves the computation time of LU-Krylov for matrices of order not less than 1000. For matrices of order 30 000, the improvement factor is about 47.6%, due to the fact that the new algorithm fully reduces to matrix multiplication and can better exploit the level 3 BLAS efficiency. Figure 2 presents these timings in a log scale graph. The

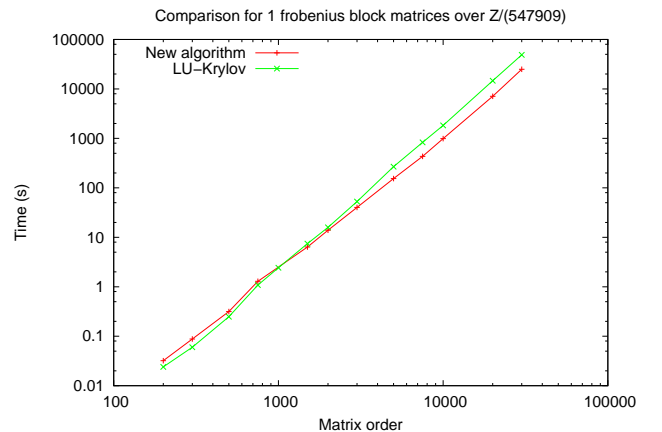


Figure 2: Timing comparison between the new algorithm and LU-Krylov, logarithmic scales, Itanium2-64 1.3Ghz, 192Gb

slopes of the two lines, which corresponds to the exponent of their complexity, are both close to 3. However, the slope of

the new algorithm is slightly lower, indicating the effective use of sub-cubic matrix multiplication for this computation.

n	magma-2.11	LU-Krylov	New algorithm
100	0.010s	0.005s	0.006s
300	0.830s	0.294s	0.105s
500	3.810s	1.316s	0.387s
800	15.64s	4.663s	1.387s
1000	29.96s	10.21s	2.755s
1500	102.1s	33.36s	7.696s
2000	238.0s	79.13s	17.91s
3000	802.0s	258.4s	61.09s
5000	3793s	1177s	273.4s
7500	MT	4209s	991.4s
10 000	MT	8847s	2080s

Table 2: Computation time for 1 Frobenius block matrices, Athlon 2200, 1.8Ghz, 2Gb

MT: Memory thrashing

Lastly table 2 gives a comparison with magma-2.11². Here again, our new implementation improves the computation time of this software, with a gain factor of about 13.8 for $n = 5000$. Moreover, its better memory management makes it possible to compute with larger matrices. On this machine, the efficiency ratio between matrix-vector and matrix multiplication is much lower than on the Itanium2. Therefore the new algorithm gets already faster for dimensions over 300.

8. CONCLUSIONS

We remark that the characteristic polynomial algorithm we have presented can easily be modified to compute the entire Frobenius form by checking some divisibility conditions of the polynomials induced by the blocks in the computed Hessenberg form. The additional cost is bounded by $O(n^\theta)$ since $\theta > 2$. Thus, we obtain a Las Vegas algorithm for computing the Frobenius form of a matrix over field that has expected cost $O(n^\theta)$.

To ensure a probability of success at least 1/2, we require that the ground field have at least $2n^2$ elements. If the field is too small we can work over an extension but a better solution (currently) would be the apply an alternative algorithm such as LU-Krylov discussed in the previous section for computing the characteristic polynomial, or the Frobenius form algorithm of Eberly [4].

For comparison, Eberly's Las Vegas Frobenius form algorithm has expected cost $O(n^\theta \log n)$, no restrictions on the field size, and it computes a similarity transform matrix as well as the form itself. Our algorithm has expected cost $O(n^\theta)$, requires the ground field to have size at least $2n^2$, and does not recover a similarity transform matrix in the same time.

On the one hand, recovery of a similarity transform matrix is undoubtedly useful for various applications [5]. On the other hand, for problems such as computing the minimal polynomial or testing two matrices for similarity the Frobenius form itself will suffice.

The main open problem we identify is to eliminate the condition on the field size while maintaining the cost bound $O(n^\theta)$: ideally the algorithm could be derandomized entirely. The currently fastest deterministic algorithm has cost $O(n^\theta (\log n)(\log \log n))$ [10, 11].

9. REFERENCES

- [1] J.-G. Dumas, T. Gautier, and C. Pernet. Finite field linear algebra subroutines. In *Proc. Int'l. Symp. on Symbolic and Algebraic Computation: ISSAC '02*, pages 63–74. ACM Press, New York, 2002.
- [2] J.-G. Dumas, P. Giorgi, and C. Pernet. Finite field linear algebra package. In J. Gutierrez, editor, *Proc. Int'l. Symp. on Symbolic and Algebraic Computation: ISSAC '04*, pages 119–126. ACM Press, New York, 2004.
- [3] J.-G. Dumas, C. Pernet, and Z. Wan. Efficient computation of the characteristic polynomial. In M. Kauers, editor, *Proc. Int'l. Symp. on Symbolic and Algebraic Computation: ISSAC '05*, pages 140–147. ACM Press, New York, 2005.
- [4] W. Eberly. Asymptotically efficient algorithms for the Frobenius form. Technical report, Department of Computer Science, University of Calgary, 2000.
- [5] M. Giesbrecht. *Nearly Optimal Algorithms for Canonical Matrix Forms*. PhD thesis, University of Toronto, 1993.
- [6] M. Giesbrecht. Nearly optimal algorithms for canonical matrix forms. *SIAM Journal of Computing*, 24:948–969, 1995.
- [7] O. Ibarra, S. Moran, and R. Hui. A generalization of the fast LUP matrix decomposition algorithm and applications. *Journal of Algorithms*, 3:45–56, 1982.
- [8] W. Keller-Gehrig. Fast algorithms for the characteristic polynomial. *Theoretical Computer Science*, 36:309–317, 1985.
- [9] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27:701–717, 1980.
- [10] A. Storjohann. *Algorithms for Matrix Canonical Forms*. PhD thesis, Swiss Federal Institute of Technology, ETH-Zurich, 2000.
- [11] A. Storjohann. Deterministic computation of the Frobenius form (Extended Abstract). In *Proc. 42nd Annual Symp. Foundations of Comp. Sci.*, pages 368–377, Los Alamitos, California, 2001. IEEE Computer Society Press.
- [12] G. Villard. A study of Coppersmith's block Wiedemann algorithm using matrix polynomials. Technical Report RR 975-I-M, IMAG Grenoble France, April 1997.
- [13] R. Zippel. Probabilistic algorithms for sparse polynomials. In *Proc. EUROSAM 79*, pages 216–226, Marseille, 1979.

²We are grateful to the Medicis computing center hosted by the CNRS STIX lab : medicis.polytechnique.fr/medicis for the possibility of running magma on their machines