Yet Another Block Lanczos Algorithm: How To Simplify the Computation and Reduce Reliance on Preconditioners in the Small Field Case

Version With Proofs

Wayne Eberly[*]
Department of Computer Science
University of Calgary
2500 University Drive NW
Calgary, Alberta, Canada T2N 1N4
eberly@ucalgary.ca

ABSTRACT

A new block Lanczos algorithm for computations over small finite fields is presented and analysed. The algorithm can be used to solve a system of linear equations or sample uniformly from the null space whenever the number of nilpotent blocks with order at least two in the Jordan form of the given coefficient matrix is less than the block factor on the right. It can also be used to confirm that this matrix condition is *not* satisfied, in order to confirm that preconditioning of the given matrix is required.

Categories and Subject Descriptors

I.1.2 [**Symbolic and Algebraic Manipulation**]: Algorithms—*algebraic algorithms, analysis of algorithms*; F.2.1 [**Analysis of Algorithms and Problem Complexity**]: Numerical Algorithms and Problems—*computations in finite fields, computations in matrices*

General Terms

Algorithms, Reliability, Performance

Keywords

Randomized computations, computations over small fields, Lanczos algorithms

1. INTRODUCTION

Since the mid nineteen-eighties, Krylov-based algorithms have been used to solve systems of linear equations over finite fields or to sample from the null space of matrices over such fields, as needed to solve a variety of problems. A considerable amount of work has subsequently taken place to improve the efficiency and reliability of these methods; the LinBox home page, `www.linalg.org`, is a good source for additional references about this. These techniques have been effective when storage requirements prohibit the use of elimination-based methods and when other special-purpose techniques have not been available.

Various matrix properties have been assumed when proving the reliability of these methods. For computations over large fields these assumptions have not been problematic, because extremely simple and efficient matrix "preconditioners" can be used to establish the properties that are required; quite a few of these are presented in the report of Chen et al. [2]. Unfortunately the set of preconditioners available for computations over small fields is more limited – to my knowledge only a sparse preconditioner first described by Wiedemann has presently been analyzed (see [9], [2], [4]), and this is somewhat more costly than desirable.

At present, the set of problems that can be solved reliably over small fields by these techniques without preconditioning is quite limited: Villard [8], has demonstrated that if rectangular blocks are used (with the block size on the left exceeding the block size on the right by at least two) then a block Wiedemann algorithm can be used, reliably, to find a nonzero element of the null space of any singular matrix.

In this paper two additional matrix problems are considered, namely, the solution of a linear system $Ax = b$ (returning either a solution for the system or a certificate establishing that it is inconsistent) and the problem of the uniform and random selection of elements from the null space — a problem that must be solved for computations over F_2 when sieve-based algorithms for integer factorization are applied [1].

As their names suggest, Krylov-based algorithms perform computations over the *Krylov space* of a set of vectors. In particular, an algorithm using block size r on the right requires that an initial set of r vectors $v_1, v_2, \ldots, v_r$ is somehow provided, and the algorithm (either implicitly or explicitly) carries out a search in the Krylov space generated by these vectors, that is, the space spanned by the set of vectors $A^i v_j$ for $i \geq 0$ and $1 \leq j \leq r$. Virtually all of the Krylov-based algorithms that have been investigated to date require either that the vectors $v_1, v_2, \ldots, v_r$ are given as part of the input

[*] Research supported in part by Natural Sciences and Engineering Research Council of Canada grant OGP00089756.

ISSAC 2010 Munich, Germany

or that they are randomly generated.

Suppose now that k is a positive integer. We will say that a matrix A (with entries in a finite field F_q with q elements) is *k-derogatory* if the first k invariant factors of the matrix are divisible by x^2 or, equivalently, the Jordan normal form of A includes at least k nilpotent Jordan blocks with order at least two. We will call the matrix A *k-nonderogatory*, otherwise. If the above-mentioned vectors $v_1, v_2, \ldots, v_r$ are to be chosen randomly then it is easily demonstrated that it is necessary for the coefficient matrix A to be nonderogatory if a Krylov-based algorithm is to be used reliably to solve a linear system. This is also a necessary condition if one wishes to sample uniformly and randomly from the null space of a given matrix.

In this paper a new block Lanczos algorithm that uses rectangular blocks (such that the difference between the block size on the left and the block size on the right is at least logarithmic in the order of the coefficient matrix) is described. With high probability the Krylov space on the right is completely searched, during the computation, at low cost. The orthogonalization process (that is at the heart of any "Lanczos" algorithm) used here is somewhat simpler than those of the prior block Lanczos algorithms of Coppersmith [3] or Hovinen [6]. Unlike the simpler method of Montgomery [7] it does not require that the coefficient matrix be symmetric and its reliability can proved.

Henceforth F_q will denote the finite field with q elements and, for integers i and j, ${\mathsf{F}_q}^{i\times j}$ will denote the set of $i \times j$ matrices with entries in F_q.

2. ITERATION OVER A KRYLOV SPACE

Given a matrix $A \in {\mathsf{F}_q}^{n\times n}$, r vectors $v_1, v_2, \ldots, v_r \in {\mathsf{F}_q}^{n\times 1}$ and a positive integer δ, the algorithm described in this section will either traverse the Krylov space generated by $v_1, v_2, \ldots, v_r$ or will fail — the latter happening with probability less than $7q^{-\delta}$.

The algorithm begins with a Lanczos phase to generate the following pair of sets.

- The set S_1 consists of m_1 ordered pairs of vectors,

  $$S_1 = \{(\mu_1, \nu_1), (\mu_2, \nu_2), \ldots, (\mu_{m_1}, \nu_{m_1})\} \qquad (1)$$

  such that, for $1 \le i, j \le m_1$, $\mu_i, \nu_j \in {\mathsf{F}_q}^{n\times 1}$ and

  $$\mu_i^t A \nu_j = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \ne j. \end{cases} \qquad (2)$$

- The set S_2 is a set of m_2 vectors

  $$S_2 = \{\lambda_1, \lambda_2, \ldots, \lambda_{m_2}\} \qquad (3)$$

  such that

  $$A\lambda_i = 0 \qquad (4)$$

  for $1 \le i \le m_2$.

In particular, if successful, the algorithm returns sets S_1 and S_2, as described above such that

$$\nu_1, \nu_2, \ldots, \nu_{m_1}, \lambda_1, \lambda_2, \ldots, \lambda_{m_2}$$

span a large subspace of the Krylov space generated by $v_1, v_2, \ldots, v_r$.

The algorithm begins with a uniform and independent selection of a set of vectors $u_1, u_2, \ldots, u_\ell \in {\mathsf{F}_q}^{n\times 1}$ where

$$\ell \ge r + 2\left(\lceil \log_q n \rceil + \delta\right). \qquad (5)$$

```
1.   Initialize $S_1$, $S_2$, $s_L$, $s_R$, $L_0$ and $U$
     loop
2.     $s_R := s_R + 1$
3.     Initialize $R$
4.     Orthogonalize the vectors in $R$; if each of the
       resulting vectors is zero then break
       if ($|U| < \ell - \lceil \log_q n \rceil - \delta$) then
         if ($U \subseteq \{(s_L, 1), (s_L, 2), \ldots, (s_L, \ell)\}$) then
5.         $s_L := s_L + 1$
6.         Initialize the sequence of vectors $L_{s_L}$
7.         Orthogonalize $L_{s_L}$
8.         $U := U \cup \{(s_L, 1), (s_L, 2), \ldots, (s_L, \ell)\}$
         else
9.         Report failure and break
         end if
       end if
10.    Either update $S_1$, $S_2$ $U$, $R$ or proceed to
       elimination phase
     end loop
```

Figure 1: Lanczos Phase of the Main Algorithm

All subsequent steps of this algorithm are deterministic.

2.1 Details of the Algorithm

The major stages of the Lanczos phase are shown in Figure 1. Throughout its execution S_1 is a set of ordered pairs of vectors, and S_2 is a set of vectors (with sizes m_1 and m_2 respectively) satisfying the various conditions shown in Equations (1)–(4), above. The integers s_L and s_R represent the "stages" of the generation of vectors on the left and on the right that are currently in progress.

The algorithm maintains sequences $L_0, L_1, \ldots, L_{s_L}$ of vectors, each of length ℓ: At each point in the computation

$$L_i = \langle \sigma_{i,1}, \sigma_{i,2}, \ldots, \sigma_{i,\ell} \rangle \qquad (6)$$

where $\sigma_{i,j} \in {\mathsf{F}_q}^{n\times 1}$ for $0 \le i \le s_L$ and $1 \le j \le \ell$. Vectors in these sequences are either completely "processed" or "unprocessed." A set of ordered pairs of integers

$$U \subseteq \{(i, j) \mid 0 \le i \le s_L \text{ and } 1 \le j \le \ell\}$$

is used to keep track of the "unprocessed vectors:" For $0 \le i \le s_L$ and $1 \le j \le \ell$, the j^{th} vector $\sigma_{i,j}$ in the sequence L_i is "unprocessed" if and only if $(i, j) \in U$.

The algorithm also maintains a set R of vectors that is used to continue generation of the desired Krylov space, as described below.

2.1.1 Initialization

Sets S_1 and S_2 are initially empty, so that $m_1 = m_2 = 0$. Variables s_L and s_R are set to have values 0 and -1 respectively. L_0 is set to be the sequence

$$L_0 = \langle u_1, u_2, \ldots, u_\ell \rangle,$$

so that $\sigma_{0,i} = u_i$ for $1 \le i \le \ell$ and, since each of these vectors is initially unprocessed,

$$U = \{(0, 1), (0, 2), \ldots, (0, \ell)\}.$$

2.1.2 Initialization of R

Since s_R is incremented at the beginning of the first execution of the loop body, $s_R \ge 0$ whenever R is initialized at step 3. If $s_R = 0$ then R is set to include the input vectors:

$$R := \{v_1, v_2, \ldots v_r\} \qquad \text{if } s_R = 0.$$

Otherwise, $s_R \geq 1$ and R is set to include the product of A and each vector currently in R, that is,

$$R := \{A\nu \mid \nu \in R\} \qquad \text{if } s_R > 0.$$

2.1.3 Orthogonalization of Vectors in R

Each vector $\zeta \in R$ is updated at step 4,

$$\zeta := \zeta - \sum_{j=0}^{m_1} \left(\mu_j^t A\zeta\right) \nu_j \tag{7}$$

to ensure that $\mu_j^t A\zeta = 0$ for $1 \leq j \leq m_1$ and for all $\zeta \in R$. As is the case for other variants of the Lanczos algorithm we will reduce the cost of this computation by showing that $\mu_j^t A\zeta = 0$ for most values of j before this step is carried out.

2.1.4 Initialization of L_{s_L}

Since L_0 is initialized at step 1, and s_L is incremented at step 5 immediately before step 6, it is clear that $s_L \geq 1$ and the set L_{s_L-1} is defined each time step 6 is performed. The sequence L_{s_L} is initialized to be

$$L_{s_L} = \langle A^t\sigma_{s_L-1,1}, A^t\sigma_{s_L-1,2}, \ldots, A^t\sigma_{s_L-1,\ell}\rangle,$$

that is, for $1 \leq j \leq \ell$, the j^{th} entry $\sigma_{s_L,j}$ in the sequence L_{s_L} is initialized to be the product of A^t and the (current) j^{th} entry in the sequence L_{s_L-1}.

2.1.5 Orthogonalization of Vectors in L_{s_L}

For $1 \leq j \leq \ell$, the j^{th} vector $\sigma_{s_L,j}$ in the sequence L_{s_L} is updated at step 7,

$$\sigma_{s_L,j} := \sigma_{s_L,j} - \sum_{k=0}^{m_1} \left(\sigma_{s_L,j}^t A\nu_k\right) \mu_k \tag{8}$$

to ensure that $\sigma_{s_L,j}^t A\nu_k = 0$ for $1 \leq k \leq m_1$. Once again the cost of this computation can be reduced by showing that $\sigma_{s_L,j}^t A\nu_k = 0$ for most values of k before this step.

2.1.6 Updating S_1, S_2, U and R

Suppose that $R = \{\kappa_1, \kappa_2, \ldots, \kappa_{|R|}\}$ at this point.

Step 10 begins with the use of an elimination-based process to compute the dimension, s, of the vector space S that is spanned by the vectors

$$\{A\kappa_i \mid 1 \leq i \leq |R|\}$$

as well as a sequence of integers $i_1, i_2, \ldots, i_s$ such that

$$1 \leq i_1 < i_2 < \cdots < i_S \leq |R|$$

and the vectors $A\kappa_{i_1}, A\kappa_{i_2}, \ldots, A\kappa_{i_S}$ form a basis for S.

For each integer j such that $1 \leq j \leq |R|$ and such that $j \notin \{i_1, i_2, \ldots, i_S\}$, $A\kappa_j$ is written as a linear combination of the above vectors, that is, elements $c_{j,1}, c_{j,2}, \ldots, c_{j,s}$ of F_q are obtained such that

$$A\kappa_j = \sum_{h=1}^{s} c_{j,h} A\kappa_{i_h};$$

the vector

$$\widehat{\kappa_j} = \kappa_j - \sum_{h=1}^{s} c_{j,h}\kappa_{i_h},$$

which is in the null space of A, is added to S_2 at this point:

$$S_2 := S_2 \cup \{\widehat{\kappa_j} \mid 1 \leq j \leq |R| \text{ and } j \notin \{i_1, i_2, \ldots, i_S\}\}$$

Now let $B_R \in {\mathsf{F}_q}^{n\times s}$ be the matrix whose columns are the above vectors $\kappa_{i_1}, \kappa_{i_2}, \ldots, \kappa_{i_s}$. Note that the matrices B_R and AB_R each have full rank s.

Let $B_L \in {\mathsf{F}_q}^{n\times|U|}$ be a matrix whose columns are the vectors $\sigma_{i,j}$ (included in the sequences $L_0, L_1, \ldots, L_{s_L}$) such that $(i,j) \in U$, arranged so that $\sigma_{h,i}$ appears to the left of $\sigma_{j,k}$ whenever $h < j$.

Let t be the rank of $B_L^t AB_R$, which is clearly at most s. The algorithm continues by computing the lexicographically first linearly maximal independent subset of the rows of this matrix; let

$$I_R = \{(i_1, j_1), (i_2, j_2), \ldots, (i_t, j_t)\} \subseteq U$$

indicate the rows of the matrix that have been included in this subset of rows so that the k^{th} row selected is the row vector $\tau_k = \sigma_{i_k,j_k}^t AB_R$ for $1 \leq j \leq t$. Compute a maximal linearly independent

$$I_C = \{h_1, h_2, \ldots, h_t\} \subseteq \{1, 2, \ldots, s\}$$

as well; set $\widehat{B}_R \in {\mathsf{F}_q}^{n\times t}$ to be the submatrix of B_R with these columns and let $C \in {\mathsf{F}_q}^{t\times t}$ be the nonsingular submatrix of$B_L^t AB_R$ which rows from I_R and columns from I_C. S_1 is now updated by adding ordered pairs

$$(\mu_{m_1+1}, \nu_{m_1+1}), (\mu_{m_1+2}, \nu_{m_1+2}), \ldots, (\mu_{m_1+t}, \nu_{m_1+t}),$$

where m_1 is the size of the set before this update, $\mu_{m_1+k} = \sigma_{i_k,j_k}$, and where ν_{m_1+k} is the k^{th} column of the matrix $\widehat{B}_R \cdot C^{-1}$ for $1 \leq k \leq t$. It is easily seen, by the choice of the above vectors, that if the set S_1 satisfied conditions (1) and (2) before this update then it does so after, as well.

Suppose $s = t$, so that $\widehat{B}_R = B_R$. Set U is first updated by removing the positions of vectors that have now been fully "processed:" $U := U \setminus I$. The remaining "unprocessed" vectors are orthogonalized against the vectors that have been included in S_1, that is, each remaining vector $\sigma_{i,j}$ for $(i,j) \in U$ is updated as

$$\sigma_{i,j} := \sigma_{i,j} - \sum_{k=m_1+1}^{m_1+s} \left(\sigma_{i,j}^t A\nu_k\right) \mu_k$$

so that $\sigma_{i,j}^t A\nu_k = 0$ for each ordered pair $(\mu_k, \nu_k) \in S_1$.

Following the computation of these vectors, the algorithm updates the set R by setting R to be the set

$$R := \{\nu_{m_1+1}, \nu_{m_1+2}, \ldots, \nu_{m_1+s}\}.$$

On the other hand, if $t < s$ then each vector λ that is a column of B_R not included in $\widehat{B}_R$ is updated as

$$\lambda := \lambda - \sum_{k=m_1+1}^{m_1+t} \left(\mu_k^t A\lambda\right) \nu_k.$$

Set S_3 is now initialized to include the vectors λ that have been computed as above, so that $S_3 = \{\lambda_1, \lambda_2, \ldots, \lambda_{s-t}\}$ where $A\lambda_1, A\lambda_2, \ldots, A\lambda_{s-t}$ are linearly independent and so that $\mu_k^t A\lambda_h = 0$ for $1 \leq k \leq m_1$ and $1 \leq h \leq s - t$. The set R is now updated to

$$R := \{\nu_{m_1+1}, \nu_{m_1+2}, \ldots, \nu_{m_1+t}\} \cup S_3$$

and we proceed to the elimination phase of the algorithm.

2.2 Elimination Phase

The final phase of the algorithm consists of a loop in which we repeatedly update the sets R, S_2 and S_3 and the value of s_R as follows.

(i) For each vector $\lambda \in R$ we compute the value

$$\widehat{\lambda} = A\lambda - \sum_{k=1}^{m_1} \left(\mu_k^t A(A\lambda)\right) \nu_k \tag{9}$$

so that $\mu_k^t A\widehat{\lambda} = 0$ for $1 \leq k \leq m_1$. Set $\widehat{R}$ to be the set of vectors λ that have been produced.

(ii) Using an elimination-based process, $\widehat{R}$ is partitioned to form a pair of sets, R' and R'', so that the vectors in the set

$$\widehat{S} = \{A\varphi \mid \varphi \in S_3 \cup R'\}$$

are linearly independent, and so that, for each vector $\widehat{\lambda} \in R''$, $A\widehat{\lambda}$ is a linear combination of the vectors in $\widehat{S}$. Indeed, elements c_τ of F_q are computed such that

$$A\widehat{\lambda} = \sum_{\tau \in S_3 \cup R'} c_\tau A\tau,$$

and these are used to compute a value

$$\lambda^\star = \widehat{\lambda} - \sum_{\tau \in S_3 \cup R'} c_\tau \tau$$

in the null space of A.

All of the above values $\lambda^\star$ are now added to S_2. S_3 and R are replace by the sets $S_3 \cup R'$ and R', respectively, and s_R is incremented.

The process terminates when it is discovered that $R = \emptyset$.

2.3 Properties of Sets

Lemma 1. *The following properties are satisfied at the beginning of each execution of the loop body of the algorithm shown in Figure 1.*

(a) S_1 is a set of ordered pairs as shown at line (1) *satisfying the conditions given at line* (2).

(b) S_2 is a set of vectors in the null space of A.

(c) If S_1 is as shown at line (1) *then $\sigma_{i,j}^t A\nu_k = 0$ for every ordered pair of integers i and j such that $(i,j) \in U$, and for $1 \leq k \leq m_1$.*

Lemma 2. *At the beginning of each execution of the body of the loop in the algorithm in Figure 1, either $s_L = 0$ and*

$$U \subseteq \{(0,1), (0,2), \ldots, (0,\ell)\}$$

or $s_L \geq 1$ and

$$U \subseteq \{(s_L - 1, 1), (s_L - 1, 2), \ldots, (s_L, \ell), \\ (s_L, 1), (s_L, 2), \ldots, (s_L, \ell)\}.$$

Suppose that $S_3 = \emptyset$ before and at the beginning of the final execution of the loop body in the Lanczos phase.

Lemma 3. *The following properties are satisfied at the beginning of each execution of the body of the loop in the algorithm shown in Figure 1 (where S_1 is as shown at line* (1)*), and at the beginning of each execution of the loop body in the elimination phase:*

(a) $s_L \geq 0$, and the set of vectors

$$V_1 = \{\mu_i \mid 1 \leq i \leq m_1\} \cup \{\sigma_{i,j} \mid (i,j) \in U\}$$

spans the same vector space as the set of vectors

$$V_2 = \left\{ \left(A^t\right)^i u_j \mid 0 \leq i \leq s_L \text{ and } 1 \leq j \leq \ell \right\}.$$

(b) $s_R \geq -1$; if $s_R \geq 0$ then the set of vectors

$$W_1 = \{\nu_i \mid 1 \leq i \leq m_1\} \cup S_2 \cup S_3$$

spans the same vector space as the set of vectors

$$W_2 = \left\{ A^i v_j \mid 0 \leq i \leq s_R \text{ and } 1 \leq j \leq r \right\}.$$

Suppose that $S_3 = \emptyset$ before the final execution of the body of the loop in the Lanczos phase.

Lemma 4. *Suppose the algorithm does not report* `failure`*. Then the following properties are satisfied on termination.*

(a) The set $\{\nu_i \mid 1 \leq i \leq m_1\} \cup S_2 \cup S_3$ spans the Krylov space that is generated by $v_1, v_2, \ldots, v_r$.

(b) The set of vectors S_2 spans the intersection of the above Krylov space and the null space of A.

2.4 Reducing the Cost of Orthogonalizations

As in all Lanczos-based algorithms, the time and space used by this algorithm will be reduced by simplifying the orthogonalizations of vectors that are required.

Lemma 5. *Consider any execution of the algorithm shown in Figure 1.*

(a) Consider any vector ζ that is being orthogonalized at step 4 of the algorithm or during the elimination phase. If $j <= m_1 - 6\ell - r$ then $\mu_j^t A\zeta = 0$ before this orthogonalization is carried out.

(b) Consider any vector $\sigma_{s_L,j}$ that is being orthogonalized at step 7. If $k <= m_1 - 2\ell - 2r$ then $\sigma_{s_L,j}^t A\nu_k = 0$ before this orthogonalization is carried out.

Thus the orthogonalization steps at lines (7) and (9), and at (8), can respectively be replaced by the simpler (and cheaper) operations

$$\zeta := \zeta - \sum_{j=\min(1, m_1 - 6\ell - r + 1)}^{m_1} (\mu_j^t A\zeta)\nu_j \tag{10}$$

and

$$\sigma_{s_L,j} := \sigma_{s_L,j} - \sum_{k=\min(1, m_1 - 2\ell - 2r + 1)} (\sigma_{s_L,j}^t A\nu_k)\mu_k. \tag{11}$$

The vectors σ_{ij} must also be orthogonalized as part of step 10 as additional ordered pairs are added to S_1. The next lemma implies that $O(n\ell)$ operations over F_q are used as part of these steps to update this vector.

Lemma 6. *At most 3ℓ ordered pairs are added to S_1, after the creation of a vector $\sigma_{i,j}$ before either the algorithm ends, or the ordered pair (i,j) is removed from U (at which point, an updated version of $\sigma_{i,j}$ is used as the first entry of an ordered pair that is added to S_1).*

2.5 Bounding the Probability of Failure

Suppose now that the Krylov space of the given vectors $v_1, v_2, \dots, v_r$ has dimension d, and let e be the dimension of the intersection of this space and the null space of A.

Lemma 7. *Suppose that $n \geq 2$, $\delta \geq 2$, and the algorithm described above is run using a matrix $A \in \mathsf{F}_q{}^{n\times n}$ and vectors $v_1, v_2, \dots, v_r$.*

(a) The algorithm fails (by executing step 9 during the Lanczos phase) with probability at most $2q^{-\delta}$.

(b) If the algorithm does not report failure then the probability that $d-e-|S_1| > i$ on termination of the Lanczos phase is less than $2q^{-\delta} + 3q^{-i}$ for any integer $i \geq 1$.

Note that the value $d - e - |S_1|$ bounds both the size of the set S_3 that is computed during the elimination phase and the number of executions of the body of the loop in that phase of the algorithm.

2.6 Summary

Theorem 1. *Let $A \in \mathsf{F}_q{}^{n\times n}$ and $v_1, v_2, \dots, v_r \in \mathsf{F}_q{}^{n\times 1}$. Let $\delta \geq 2$ be a positive integer. The algorithm described above can be used, with $\ell \geq r + 2(\lceil \log_q n\rceil + \delta)$, to compute a set of vectors that span the Krylov basis K generated by $v_1, v_2, \dots, v_r$, as well as bases for this space and for its intersection with the null space of A, failing with probability at most $2q^{-\delta}$.*

The number of applications of A or A^t to vectors during the execution of this algorithm is in $\Theta(n)$ in the worst case and, if $\delta \geq \log_q n$, then the expected number of additional operations in F_q used by the algorithm is in $\Theta(n^2\ell)$.

It also uses space to to store $\Theta(\ell)$ vectors, needed either for the orthogonalizations that are performed while the algorithm takes place or to support the applications of this algorithm that are described in the sequel.

On the other hand, if $\delta < \log_q n$ then a version of the algorithm that fails with probability at most $7q^{-\delta}$, and that uses the above numbers of operations in the worst case, can be obtained simply by terminating the elimination phase of the algorithm, and reporting `failure`, as soon as it is noted that $|S_3| > \delta$.

3. APPLICATIONS

In this section we will consider the application of the block Lanczos algorithm presented in the previous section. Suppose the Jordan normal form of $A \in \mathsf{F}_q{}^{n\times n}$ includes exactly m nilpotent Jordan blocks

$$\begin{bmatrix} 0 & 1 & & & & 0 \\ & 0 & 1 & & & \\ & & \ddots & & & \\ & & & 0 & 1 & \\ & & & & 0 & 1 \\ 0 & & & & & 0 \end{bmatrix}$$

with order at least two. Then there exists a nonsingular matrix $X \in \mathsf{F}_q{}^{n\times n}$ such that

$$A = X \begin{bmatrix} A_1 & & 0 \\ & A_2 & \\ 0 & & A_3 \end{bmatrix} X^{-1} \qquad (12)$$

where

- A_1 is a nonsingular matrix with order n_1 for some integer n_1 such that $0 \leq n_1 \leq n - 2m$.
- A_2 is a block diagonal matrix

$$A_2 = \begin{bmatrix} J_1 & & & 0 \\ & J_2 & & \\ & & \ddots & \\ 0 & & & J_m \end{bmatrix} \qquad (13)$$

  where J_h is a nilpotent Jordan block with order $n_{2,h} \geq 2$, so A_2 has order $n_2 = n_{2,1} + n_{2,2} + \cdots + n_{2,m} \geq 2m$.

- A_3 is a zero matrix with order $n_3 = n - n_1 - n_2 \geq 0$.

If e_i is the i^{th} standard unit vector, $\mathcal{V}_1$ is the vector space with basis $Xe_1, Xe_2, \dots, Xe_{n_1}$, the vector space $\mathcal{V}_2$ has basis $Xe_{n_1+1}, Xe_{n_1+2}, \dots, Xe_{n_1+n_2}$, and $\mathcal{V}_3$ has basis $Xe_{n_1+n_2+1}, Xe_{n_1+n_2+2}, \dots, Xe_n$, then, since the above matrix X is nonsingular, it is clear that $\mathsf{F}_q{}^{n\times 1}$ is the direct sum of the vector spaces $\mathcal{V}_1$, $\mathcal{V}_2$, and $\mathcal{V}_3$. Thus each vector $x \in \mathsf{F}_q{}^{n\times 1}$ can be written uniquely as the sum of vectors $\alpha \in \mathcal{V}_1$, $\beta \in \mathcal{V}_2$, and $\gamma \in \mathcal{V}_3$. It is also clear each of the vectors spaces $\mathcal{V}_1$, $\mathcal{V}_2$, and $\mathcal{V}_3$ is closed under multiplication by A. Furthermore A acts as an invertible operator in $\mathcal{V}_1$, a nilpotent operator in $\mathcal{V}_2$, and the zero operator in $\mathcal{V}_3$.

It is clear from Equations (12) and (13), above, that $\mathcal{V}_2$ is the Krylov space of a set of m vectors in $\mathcal{V}_2$; henceforth we will set $\omega_1, \omega_2, \dots, \omega_m$ to be a set of vectors in this space such that $\mathcal{V}_2$ is the Krylov space generated (using A) by $\omega_1, \omega_2, \dots, \omega_m$.

Lemma 8. *Let $v_i = \alpha_i + \beta_i + \gamma_i \in \mathsf{F}_q{}^{n\times n}$ such that $\alpha_i \in \mathcal{V}_1$, $\beta_i \in \mathcal{V}_2$, and $\gamma_i \in \mathcal{V}_3$, for $1 \leq i \leq r$.*

(a) The Krylov space K that is generated by $v_1, v_2, \dots, v_r$ includes both the Krylov space $K_1 \subseteq \mathcal{V}_1$ that is generated by $\alpha_1, \alpha_2, \dots, \alpha_r$ and the Krylov space $K_2 \subseteq \mathcal{V}_2$ that is generated by $A\beta_1, A\beta_2, \dots, A\beta_r$.

(b) Suppose that $\{\kappa_1, \kappa_2, \dots, \kappa_h\}$ is a basis for K_1 and that $\{\lambda_1, \lambda_2, \dots, \lambda_j\}$ is a basis for K_2, where K_1 and K_2 are as given in part (a), above. Then the set of vectors

$$\{\kappa_1, \kappa_2, \dots, \kappa_h\} \cup \{\lambda_1, \lambda_2, \dots, \lambda_j\} \cup \{\beta_1 + \gamma_1, \beta_2 + \gamma_2, \dots, \beta_r + \gamma_r\}$$

spans the Krylov space generated by $v_1, v_2, \dots, v_r$.

Lemma 9. *Suppose that $r \geq m + \Delta$ for a positive integer Δ, and that vectors $v_1, v_2, \dots, v_r$ are chosen uniformly and independently from $\mathsf{F}_q{}^{n\times 1}$. Then with probability at least $1 - 2q^{-\Delta}$ there exist elements $\lambda_1, \lambda_2, \dots \lambda_m$ of $\mathcal{V}_3$ such the Krylov space generated by $v_1, v_2, \dots, v_r$ includes the vectors*

$$\omega_1 + \lambda_1, \omega_2 + \lambda_2, \dots, \omega_m + \lambda_m.$$

In this case the Krylov space also includes $A\beta$ for all $\beta \in \mathcal{V}_2$.

3.1 Solving a Consistent Linear System

Suppose A is as in Equations (12) and (13) and consider the problem of finding a vector x such that $Ax = b$, for a given $b \in \mathsf{F}_q{}^{n\times 1}$, assuming that such a vector exists.

An algorithm for this computation will begin with an application of an augmentation of the algorithm using a set $v_1, v_2, \dots, v_r$ where $v_1 = b$ and for $v_2, v_3, \dots, v_r$ selected uniformly and independently from $\mathsf{F}_q{}^{n\times 1}$.

Note that if the given system is consistent then $b = \alpha + \beta$ such that $\alpha \in \mathcal{V}_1$, and $\beta \in \mathcal{V}_2$ is an element of the Krylov space generated by $A\omega_1, A\omega_2, \ldots, A\omega_m$, for $\omega_1, \omega_2, \ldots, \omega_m$ as described above. Furthermore, the the fact that A acts as an invertible operator in $\mathcal{V}_1$ can be used to establish that there Krylov space generated by α includes a value χ such that $A\chi = \alpha$. The following can now be established.

LEMMA 10. *Suppose that the system $Ax = b$ is consistent, the number m of nilpotent blocks in the Jordan normal form of A is at most $r - \Delta - 1$, for a positive integer $\Delta \geq 2$, and that vectors $v_1, v_2, \ldots, v_r$ are selected as described above. Then the Krylov space generated by $v_1, v_2, \ldots, v_r$ includes a solution for the system $Ax = b$ with probability at least $1 - 2q^{-\Delta}$.*

Suppose now that the Krylov space does contain a vector $\widehat{x}$ such that $A\widehat{x} = b$. Suppose that the algorithm from Section 2 is augmented so that it maintains an additional pair of vectors, x and res, that are initialized to have values 0 and b at step 1. Clearly

$$Ax + \mathit{res} = b; \tag{14}$$

the values x and res will be updated as the set S_1 is modified, in order to ensure that

$$\mu_i^t \mathit{res} = 0$$

for each ordered pair $(\mu_i, \nu_i) \in S_1$. In particular this be achieved by setting c_i to be $\mu_i^t b$ and then replacing x and res by the values $x + c_i\nu_i$ and $\mathit{res} - c_i A\nu_i$, respectively. It is easily checked Equation (14) is still satisfied and that $\mu_i^t A(Ax - b) = \mu_i^t \mathit{res} = 0$ after this update.

Let us compare the computed value x to the solution $\widehat{x}$ mentioned above, after all the ordered pairs in S_1 have been considered: It is clear that $\mu_i^t A(x - \widehat{x}) = \mu_i^t(Ax - b) = 0$ for $1 \leq i \leq m_1$. Both x and $\widehat{x}$ belong to the Krylov space that has been generated, so $x - \widehat{x}$ must be a linear combination of the vectors in $S_2 \cup S_3$. Since each vector in S_2 belongs to the null space of A, it follows that $Ax - b = Ax - A\widehat{x}$ is a linear combination of vectors $A\lambda$, for $\lambda \in S_3$.

Now recall from Section 2 that $S_3 = \{\lambda_1, \lambda_2, \ldots, \lambda_{m_3}\}$ (for $m_3 = |S_3|$) where the vectors $A\lambda_1, A\lambda_2, \ldots, A\lambda_{m_3}$ are linearly independent. Values $c_1, c_2, \ldots, c_{m_3} \in \mathsf{F}_q$ such that

$$(Ax - b) = c_1 A\lambda_1 + c_2 A\lambda_2 + \cdots + c_{m_3} A\lambda_{m_3}$$

can be found, by a simple elimination-based process, using $\Theta(nm_3^2)$ additional operations, and, after the value x is updated to be

$$x := x - c_1\lambda_1 - c_2\lambda_2 - \cdots - c_{m_3}\lambda_{m_3},$$

it is clear that $Ax = b$ as required.

As described above the above computation can be carried out using a modified version of the main algorithm, requiring storage of another two vectors (namely, x and res). Part (b) of Lemma 7 (which bounds the probability that m_3 exceeds a given size) can be used to establish that the expected number of additional operations over F_q that are required is in $\Theta(n^2)$.

3.2 Bounding the Number of Nilpotent Blocks

Next consider the problem of deciding whether A is m-derogatory for a given integer $m \geq 0$.

LEMMA 11. *Let $A \in \mathsf{F}_q{}^{n\times n}$. Then the Jordan normal form of A includes exactly m nilpotent blocks with order of at least two if and only if the intersection $\mathcal{Z}$ of the image of A and the null space of A is a vector space with dimension m.*

This suggests the following algorithm: Given an integer m and a positive integer $\Delta \geq 2$, generate $r = m + \Delta$ vectors $z_1, z_2, \ldots, z_r$ uniformly and independently, and then use an augmented version of the algorithm from Section 2 to compute a spanning set for the Krylov space generated by $v_1 = Az_1, v_2 = Az_2, \ldots, v_r = Az_r$.

This augmented version of the algorithm maintains a second pair of sets T_1 and T_2, such that $|T_1| = |S_1| = m_1$, $|T_2| = |S_2| = m_2$,

$$T_1 = \{\varphi_1, \varphi_2, \ldots, \varphi_{m_1}\}$$

such that $A\varphi_i = \nu_i$ (the second entry of the i^{th} ordered pair of vectors in S_1) for $1 \leq i \leq m_1$, and where

$$S_2 = \{A\zeta \mid \zeta \in T_2\}$$

as well. The algorithm also makes use of a set P of vectors such that $R = \{A\rho \mid \rho \in P\}$. These sets are easily maintained: T_1 and T_2 are initialized to be empty, along with S_1 and S_2, at the beginning of the algorithm. During the first execution of the loop body P is initialized to include $z_1, z_2, \ldots, z_r$ at the same time as R is initialized to include $v_1, v_2, \ldots, v_r$. No further (additional) multiplications of vectors by A or A^t are needed, since P can be replaced by R immediately before R is updated at step 3. Elements of R are updated at step 4 by adding linear combinations of $\nu_1, \nu_2, \ldots, \nu_{m_1}$; the corresponding elements of P can be updated by adding linear combinations of $\varphi_1, \varphi_2, \ldots, \varphi_{m_1}$ using the same multipliers. An inspection of the updating of R during the elimination confirms P can be updated at the same time as R, during this part of the algorithm, without additional applications of A or A^t as well. Finally, the sets P, T_1 and T_2 can be updated at the same time as R, S_1 and S_2 in the same way. The number of operations over F_q and the storage space needed by it are at most doubled.

On termination one should compute a maximal linearly independent set S_2' of the vectors in S_2, returning the corresponding subset of the vectors T_2' in T_2. Suppose that set of k vectors $\tau_1, \tau_2, \ldots, \tau_k$ is produced. It follows by their construction that $A\tau_1, A\tau_2, \ldots A\tau_k$ are linearly independent and that $A^2\tau_i = 0$ for $1 \leq i \leq k$.

LEMMA 12. *Suppose that the above algorithm is applied using a matrix $A \in \mathsf{F}_q{}^{n\times n}$.*

(a) If A has $\widehat{m} < r$ nilpotent blocks with order at least two in its Jordan normal form then $k \leq \widehat{m}$. The probability that $k < \widehat{m}$ (so that a basis for $\mathcal{Z}$ is not obtained) is at most $2q^{r-\widehat{m}}$.

(b) If A has $\widehat{m} \geq r$ nilpotent blocks with order at least two in its Jordan normal form then the probability that $k < m$ (so that it is not proved that A is m-derogatory) is at most $2q^{-\Delta}$.

Consequently, since $m \leq r - \Delta$, a basis for $\mathcal{Z}$ is obtained with probability at least $1 - 2q^{-\Delta}$ if A is m-nonderogatory. On the other hand, if A is m-derogatory then a certificate of this is obtained with probability at least $1 - 2q^{-\Delta}$, instead.

3.3 Sampling from the Null Space

Let $A \in \mathsf{F}_q{}^{n\times n}$. Suppose that A is a singular matrix that has exactly m nontrivial nilpotent blocks in its Jordan normal form and, indeed, that the value of m is known and that a basis for the intersection $\mathcal{Z}$ of the image of A and its null space has been obtained (perhaps, using the algorithm described in Section 3.2, above). The idealistic assumption that one can sample elements uniformly and independently from F_q will also be made.

Suppose that we are given an integer $d > 0$ and that we wish to generate a sequence

$$\zeta_1, \zeta_2, \dots, \zeta_d$$

such that $A\zeta_i = 0$ for $1 \le i \le d$ and such that every such sequence is generated with probability q^{-kd}, where k is the (generally unknown) dimension of the null space.

An algorithm to solve this problem will now be presented. The algorithm will either generate a sequence as described above, or it will detect and report `failure`. Since failure can be detected, and its probability can be bounded, a process that eventually produces a sequence of the above form, whose expected running time is small, can be obtained as a sequence of independent trials of the algorithm that will be described, ending with the first trial that succeeds.

The probability of failure can be bounded but, unfortunately, depends on the parameter d: The probability of failure is at most $2q^{-d}$. That noted, the probability of failure can be reduced to at most $2q^{-\Delta}$ for a given integer $\Delta > d$, while generating sequences with the desired probabilities, by applying the algorithm to produce a sequence of length Δ and then discarding the final $\Delta - d$ elements.

To begin, let $v_1, v_2, \dots, v_{m+d}$ be uniformly and independently selected from ${\mathsf{F}_q}^{n\times 1}$. Apply the algorithm in Section 2 with vectors $Av_1, Av_2, \dots, Av_{m+d}$ to check that the intersection of the Krylov space generated by these vectors and the null space of A has dimension m. Since this intersection is a subspace of $\mathcal{Z}$ its dimension is at most m, and part (a) of Lemma 12 implies that its dimension is equal to m, and this intersection is equal to $\mathcal{Z}$, with probability at least $1 - 2q^{-d}$; `failure` should be reported if this is not the case.

Provided that `failure` has not been reported, we should continue by running the algorithm in Section 2 again, with vectors $v_1, v_2, \dots, v_{m+d}$, and considering the set S_2 that has now been obtained. We should generate a basis for the space $\mathcal{X}$ spanned by these vectors by using them to extend the basis for $\mathcal{Z}$ that is already available. Since $\mathcal{X}$ has dimension at most $m + d$. and $\mathcal{Z} \subseteq \mathcal{X}$, this process results in at most d additional vectors. Since the null space of A is contained in the direct sum of $\mathcal{V}_2$ and $\mathcal{V}_3$, these vectors are

$$\beta_{m+1} + \gamma_1, \beta_{m+2} + \gamma_2, \dots, \beta_{m+e}, \gamma_e \tag{15}$$

where $0 \le e \le d$, $\beta_j \in \mathcal{V}_2$ for $m+1 \le j \le m+e$, and $\gamma_j \in \mathcal{V}_3$ for $1 \le j \le e$.

The process that has been described so far can be considered to be an "experiment" that implicitly defines a subspace of $\mathcal{V}_3$ with dimension at most d, namely, the space $\mathcal{W}$ with dimension e with basis $\gamma_1, \gamma_2, \dots, \gamma_e$ — these vectors are easily seen to be linearly independent, because the set of vectors including the given basis $\beta_1, \beta_2, \dots, \beta_m$ for $\mathcal{Z} \subseteq \mathcal{V}_2$ as well as the vectors at line (15), above are, by construction, linearly independent, and because β_{m+i} is a linear combination of $\beta_1, \beta_2, \dots, \beta_m$ for $1 \le i \le e$.

Of course, if a basis for $\mathcal{V}_3$ was available then we could also generate a subspace $\mathcal{W}$ of $\mathcal{V}_3$ with dimension at most d using a second "experiment," namely, by generating d linear combinations of the elements of this basis, uniformly and independently, and considering the subspace that is spanned by the vectors that have been generated.

A careful consideration of the elimination process that is needed to extend the basis in the first "experiment," as described above, and the observation that a given probability distribution can be described in multiple ways, establishes the following.

Lemma 13. *Let $\mathcal{W}$ be a subspace of $\mathcal{V}_3$ whose dimension at most d. The probability that $\mathcal{W}$ is generated, using the first of the experiments described above, is equal to the probability that $\mathcal{W}$ is generated using the second experiment, instead.*

Suppose that, as the next step in this computation, we wish to generate a sequence

$$\widehat{\zeta}_1 = \widehat{\beta}_1 + \widehat{\gamma}_1, \widehat{\zeta}_2 = \widehat{\beta}_2 + \widehat{\gamma}_2, \dots, \widehat{\zeta}_d = \widehat{\beta}_d + \widehat{\gamma}_d \tag{16}$$

where $\widehat{\beta}_1, \widehat{\beta}_2, \dots, \widehat{\beta}_d$ belong to $\mathcal{Z}$, and where $\widehat{\gamma}_1, \widehat{\gamma}_2, \dots, \widehat{\gamma}_d$ have been selected uniformly and independently from $\mathcal{V}_3$.

Consider, once again, the sequence of linearly independent vectors $\gamma_1, \gamma_2, \dots, \gamma_e \in \mathcal{V}_3$ that have been generated at the end of the previous step. Let $B_{\mathcal{W}} \in {\mathsf{F}_q}^{n\times e}$ be the matrix with these vectors as its columns and consider a matrix $C \in {\mathsf{F}_q}^{n\times d}$ whose column are to be selected from $\mathcal{V}_3$. The columns of any such matrix span exactly one vector space $\mathcal{W} \subseteq \mathcal{V}_3$ with dimension at most d. Indeed, the columns of a given matrix C span the same vector space $\mathcal{W}$ as the above vectors $\gamma_1, \gamma_2, \dots, \gamma_e$ if and only if there is a matrix $N \in {\mathsf{F}_q}^{e\times d}$ with maximal rank e such that $B_{\mathcal{W}}N = C$. Since we wish to generate such matrices C uniformly, and since $B_{\mathcal{W}}N_1 = B_{\mathcal{W}}N_2$ (for $N_1, N_2 \in {\mathsf{F}_q}^{e\times d}$) if and only if $N_1 = N_2$, a sequence of vectors as shown at line (16) can now be generated, following the determination of the space $\mathcal{W}$ as described above, by carrying out the following steps.

1. Generate a matrix N uniformly and randomly from the set of matrices in ${\mathsf{F}_q}^{e\times d}$ with rank e.

2. Set $\widehat{\zeta}_1, \widehat{\zeta}_2, \dots, \widehat{\zeta}_d$ to the columns of the product $B_{\mathcal{W}} \cdot N$.

The above matrix N is to be selected uniformly from the set of matrices in ${\mathsf{F}_q}^{e\times d}$ whose rows are linearly independent. Such a matrix can be generated by choosing the rows, one at a time; following the selection of the first i rows the $i+1^{\text{st}}$ can be selected by considering a sequence of uniformly and independently selected vectors from ${\mathsf{F}_q}^{1\times d}$. The first of these that is not a linear combination of the i rows that have already been selected should be used. Since the expected number of vectors that must be considered to generate each row is less than two, the expected number of elements that must be uniformly selected from F_q as part of this process is in $\Theta(de)$. An elimination-based process can be used to check the rank condition given above; the expected number of operations needed for this is in $O(d^3) \subseteq O(n^2 d)$. The cost of the matrix multiplication required for the second step is in $O(ned) \subseteq O(n^2 d)$ as well.

Recall our assumption that a basis for $\mathcal{Z}$ is available. The computation can be completed by using this basis to select a sequence of m vectors $\beta'_1, \beta'_2, \dots, \beta'_d$ uniformly and independently from $\mathcal{Z}$ — choosing these independently of the vectors $\widehat{\zeta}_1, \widehat{\zeta}_2, \dots, \widehat{\zeta}_d$ that have already been obtained. It is sufficient to set $\zeta_i = \widehat{\zeta}_i + \beta'_i$, for $1 \le i \le d$, to obtain a sequence of d vectors that are uniformly and independently selected from the null space of A, as desired. This final step requires the selection of $md \leq n^2$ values uniformly and independently from F_q, followed by $O(nmd) \subseteq O(n^2d)$ additional operations over this field.

3.4 Certifying Inconsistency

Consider next the problem of certifying that a given system $Ax = b$ is *not* consistent, that is, that b does not belong to the column space of A.

If the matrix A is m-nonderogatory then so its transpose and, as observed by Giesbrecht, Lobo and Saunders [5], the probability that $\mu^t b = 0$ is at most $1/q$ if the system of linear equations $Ax = b$ is inconsistent and μ is a uniformly and randomly chosen element of the null space of A^t.

An algorithm that certifies inconsistency of a system $Ax = b$, for an m-nonderogatory matrix A can now be obtained using the algorithms that have already been described above: One should first determine the number of nilpotent blocks with order at least two in the Jordan normal form of A^t, and compute a basis for the space $\mathcal{Z}$ corresponding to the matrix A^t, by an application of the algorithm described in Subsection 3.2. This should the followed by an application of the algorithm from Subsection 3.3 to produce a sequence of Δ vectors $\zeta_1, \zeta_2, \ldots, \zeta_\Delta$ that are uniformly and independently selected from the null space of A^t: If the previous algorithm did not fail (that is, it really did compute a basis for $\mathcal{Z}$), and the system $Ax = b$ is inconsistent, then a vector ζ such that $\zeta^t A = 0$ but $\zeta^t b \neq 0$, establishing that b is not in the column space of A, will have been produced with probability at least $1 - q^{-\Delta}$.

As described here the algorithm can "fail" for a variety of reasons. This can be addressed by combining it with an attempt to solve the given system using the algorithm in Subsection 3.1. The result is an algorithm — which should receive the matrix A, vector b, and a positive integer m — and which produces either a vector x such that $Ax = b$, a vector μ such that $\mu^t A = 0$ but $\mu^t b \neq 0$, a proof that the matrix A is m-derogatory, as described above, or where the only reason for `failure` is an unlucky choice of the random values that have been selected. Consequently repeated trials of this will eventually result in one of outputs (i), (ii), or (iii). In the event of (iii) a user should presumably try again with a larger value of m or apply a preconditioner in order to bring A into a more manageable form.

4. FURTHER WORK

This paper has been a kind of "demonstration of concept," in that it describes an algorithm that has not yet been implemented. It is possible that some of the results presented here could be improved and, indeed, this would certainly be desirable: While it is *sufficient* to store approximately 7ℓ pairs of vectors to be used for the orthogonalizations that are considered in Subsection 2.4 I do not know whether it is necessary and, of course, the number of vectors to be stored should be reduced if this is possible.

Indeed, it is not clear that the *algorithm* described in Section 2 is necessary: It is certainly plausible (but, to my knowledge, not yet verified) that an existing block Wiedemann algorithm could also be used to carry out the computations described in Section 3.

The small field preconditioner mentioned at the beginning of this paper is arguably more expensive than is desirable, but it also achieves a stronger matrix property than is used here, in that it ensures, with high probability, that the number of nontrivial invariant factors of the preconditioned matrix is small. Are there other, less expensive small field preconditioners, that are sufficient to ensure the weaker condition that the preconditioned matrix is nonderogatory?

A final theoretical question concerns a property of the null space of a matrix: Is it possible to discover the dimension of the null space (and, therefore, the rank of the given matrix) using only the fact that the matrix is nonderogatory, and without preconditioning? Note that the algorithm to sample from the null space in Section 3 does not require this value and, as far as I can tell, fails to provide any information that could be used to discover this in general. I suspect that the answer to this question is "no," but have no idea of how to prove this.

5. REFERENCES

[1] J. P. Buhler, J. H. W. Lenstra, and C. Pomerance. Factoring integers with the number field sieve. In *The Development of the Number Field Sieve*, volume 1554 of *Lecture Notes in Computer Science*, pages 50–94. Springer-Verlag, 1993.

[2] L. Chen, W. Eberly, E. Kaltofen, B. D. Saunders, W. J. Turner, and G. Villard. Efficient matrix preconditioners for black box linear algebra. *Linear Algebra and Its Applications*, 343–344:119–146, 2002.

[3] D. Coppersmith. Solving linear equations over GF(2): Block Lanczos algorithm. *Linear Algebra and Its Applications*, 192:33–60, 1993.

[4] W. Eberly. Reliable Krylov-based algorithms for matrix null space and rank. In *Proceedings, 2004 International Symposium on Symbolic and Algebraic Computation*, pages 127–134, Santander, Spain, 2004.

[5] M. Giesbrecht, A. Lobo, and B. D. Saunders. Certifying inconsistency of sparse linear systems. In *Proceedings, 1998 International Symposium on Symbolic and Algebraic Computation*, pages 113–119, University of Rostock, Germany, 1998.

[6] B. Hovinen and W. Eberly. A reliable block Lanczos algorithm over small finite fields. In *Proceedings, 2005 International Symposium on Symbolic and Algebraic Computation*, pages 177–184, Beijing, China, 2005.

[7] P. L. Montgomery. A block Lanczos algorithm for finding dependencies over GF(2). In *Advances in Cryptology—EUROCRYPT '95*, volume 921 of *Lecture Notes in Computer Science*, pages 106–120. Springer-Verlag, 1995.

[8] G. Villard. Further analysis of Coppersmith's block Wiedemann algorithm for the solution of sparse linear systems. In *Proceedings, 1997 International Symposium on Symbolic and Algebraic Computation*, pages 32–39, Maui, Hawaii, 1997.

[9] D. Wiedemann. Solving sparse linear equations over finite fields. *IEEE Transactions on Information Theory*, IT-32:54–62, 1986.

APPENDIX

The following sections will not included in the extended abstract in the conference proceedings but will be available in a full version of the paper that will be available at the author's web site and cited in conference version. They are included here for the convenience of referees.

A. PROOFS OF CLAIMS IN SECTION 2

A.1 Proof of Lemma 1

The claims in Lemma 1 follow by a reasonably straightforward induction on the number of executions of the loop body.

The claims are trivial for the first execution of the body of the loop, since S_1 and S_2 have each been set to be empty at this point.

Suppose now that $i \geq 1$ and that the part (a) of the claim is satisfied at the beginning of the i^{th} execution of the loop body. Note that during the execution of this body of the loop, zero or more additional ordered pairs are added, that is, S_1 is replaced by

$$S_1 \cup \{(\mu_{m_1+1}, \nu_{m_1+1}), (\mu_{m_1+2}, \nu_{m_1+2}), \ldots, (\mu_{m_1+s}, \nu_{m_1+s})\}$$

where the additional ordered pairs, shown, above, are computed during Step 10. Each of the vectors

$$\nu_{m_1+1}, \nu_{m_1+2}, \ldots, \nu_{m_1+s}$$

is computed in Step 10 as a linear combination of vectors in the set R, and it is assured at step 4 that these are orthogonal to $\mu_1, \mu_2, \ldots, \mu_{m_1}$. Consequently

$$\mu_i A^t \nu_j = 0 \quad \text{for } 1 \leq i \leq m_1 \text{ and } m_1 + 1 \leq j \leq m_1 + s.$$

Similarly, the vectors $\mu_{m_1+1}, \mu_{m_1+2}, \ldots, \mu_{m_1+s}$ are each selected in step 10 from the set of vectors $\sigma_{i,j}$ such that $(i,j) \in U$. The ordered pairs in U either belonged to U at the beginning of this execution of the loop body — in which case, it follows by the induction hypothesis (and part (d)) of the claim) that they are orthogonal to $\nu_1, \nu_2, \ldots, \nu_{m_1}$ — or they are introduced at step 8, in which case it is assured by the processing at the previous step that they are orthogonal to $\nu_1, \nu_2, \ldots, \nu_{m_1}$ as well. It follows that

$$\mu_i^t A \nu_j = 0 \quad \text{for } m_1 + 1 \leq i \leq m_1 + s \text{ and } 1 \leq j \leq m_1$$

as well.

Finally, the construction of $\nu_{m_1+1}, \nu_{m_1+2}, \ldots, \nu_{m_1+s}$ at step 10 ensures that, for $m_1 + 1 \leq i, j \leq m_1 + s$,

$$\mu_i^t A \nu_j = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

It follows the the above equations, and the inductive hypothesis, that the set of ordered pairs S_1 satisfies the condition at line (2) at the end of this execution of the loop body. The set also clearly satisfies this constraint at the beginning of the next execution as well; that is, part (a) of the claim is satisfied at the beginning of the next execution of the body of the loop.

Establishing part (b) is straightforward, since S_2 is only modified by adding a set of vectors that are clearly in the null space of A.

Finally, it should be noted that the set U is only modified at steps 8 and 10 of the algorithm. The orthogonalization process at step 7 (and the inductive hypothesis) ensure that $\sigma_{i,j}^t A \nu_k = 0$ for each pair of integers i and j such that $(i,j) \in U$, and for $1 \leq k \leq m_1$ following step 8. The orthogonalization process following the updating of U in step 10 ensures that $\sigma_{i,j}^t A \nu_k$ for each pair of integers i and j such that $(i,j) \in U$ and for $1 \leq k \leq m_1 + s$ at the end of step 10 as well. Since the set U is not changed between the end of one execution of the loop body and the beginning of the next, it follows that part (c) of the claim is also satisfied at the beginning of the next execution of the loop body, as required.

Note: It follows from the argument given above that the claims given as parts (a) and (b) of the lemma are also satisfied at the beginning of the execution of the elimination phase of the algorithm; this will be used to establish some of the results that follow.

A.2 Proof of Lemma 2

It is necessary and sufficient to show that either $s_L = 0$ and

$$U \subseteq \{(0,1), (0,2), \ldots, (0,\ell)\}$$

or $s_L > 0$ and

$$U \subseteq \{(s_L - 1, 1), (s_L - 1, 2), \ldots, (s_L, \ell), \\ (s_L, 1), (s_L, 2), \ldots, (s_L, \ell)\}.$$

at the beginning of the i^{th} execution of the body of the loop, for each integer i such that the loop body is executed i or more times. This is easily established by induction on i.

The above relationship clearly holds when $i = 1$, since S_L is initialized to be 0 and U is set to be

$$\{(0,1), (0,2), \ldots, (0,\ell)\}$$

just before the execution of the loop begins.

Suppose that $i \geq 2$, the loop body is executed at least i times, and the above relationship holds at the beginning of the i^{th} execution of the loop body. During this execution of the loop body, either the tests above step 5 succeed and steps 5–8 are executed, the first test succeeds and second fails, so that step 9 is executed, or the first test fails so that steps 5–9 are skipped.

In the first of these cases, the success of the tests before step 5 ensure that $s_L \geq 0$ and

$$U \subseteq \{(s_L, 1), (s_L, 2), \ldots, (s_L, \ell)\}$$

immediately before step 5. The execution of steps 5–8 (which begins with s_L being incremented) clearly ensures that $s_L \geq 1$ and

$$U \subseteq \{(s_L - 1, 1), (s_L - 1, 2), \ldots, (s_L, \ell), \\ (s_L, 1), (s_L, 2), \ldots, (s_L, \ell)\}.$$

after these steps are completed. No ordered pairs are added to U during step 10 (although some may be removed) so the above containment holds at the end of this execution of the loop body and the beginning of the next.

In the second of these cases, the execution of step 9 terminates the execution of the loop. Consequently there is no $i + 1^{\text{st}}$ execution of the loop body, and nothing to be proved.

In the final case (steps 5–9 are skipped), the value of s_L is not changed by this execution of the loop body and no ordered pairs are added to U (again, some may be removed at step 10). Consequently the above relationship holds at the end of this execution of the loop body, and at the beginning of the next execution of the loop body, as needed to complete the proof.

Note: A consideration of this argument can be used to establish a slightly more general result, which will be used in remaining proofs: The relationships concerning s_L and U also hold at the beginning of step 10 in the body of this loop and at the beginning of the elimination phase of this algorithm.

A.3 Proof of Lemmas 3 and 4

LEMMA 14. *Let i be a positive integer such that the body of the loop in the algorithm shown in Figure 1 is executed at least i times. Suppose that, at the beginning of the i^{th} execution of the body of the loop,*

- *$\mathcal{V}$ is the vector space spanned by the vectors $\sigma_{i,j}$ for $0 \le i \le s_L$ and $1 \le j \le \ell$, and*
- *$\mathcal{W}$ is the vector space spanned by the vectors $A^t\sigma_{s_L,j}$ such that $1 \le j \le \ell$ and $(s_L, j) \notin U$.*

Then

(a) $A^t\sigma_{i,j} \in \mathcal{V}$ for all integers i and j such that $0 \le i \le s_L - 2$ and $1 \le j \le \ell$.

(b) For each integer j such that $1 \le j \le \ell$ there exist vectors $\eta_j \in \mathcal{V}$ and $\theta_j \in \mathcal{W}$ such that $A^t\sigma_{s_L-1,j} = \eta_j + \theta_j$.

Note: In the following proof if it necessary to consider the values of several variables and sets as these are changed during an execution of the body of the loop. While appropriate notation will be used to keep track of these values for some variables and sets (so that, for example, the argument refers to sets U, $\widehat{U}$, and U', using each to represent the set assigned to the variable U at different points in the computation) this is not feasible for the vectors being processed: Each vector $\sigma_{i,j}$ will have several different values during the execution of the loop; the context of the claims made, below, should identify precisely which value is being referred to in each case.

PROOF. The result can be proved by induction on i. If $i = 1$ then $s_L = 0$ and claim is trivial. Indeed, it is trivial for each integer i such that $s_L = 0$ at the beginning of the i^{th} execution of the body of the loop.

Suppose now that i is any positive integer such that the above relationships hold at the beginning of the i^{th} execution of the loop body. These relationships also hold at the beginning of step 10 during this execution of the loop body, if step 10 is reached, provided that the above vector spaces $\mathcal{V}$ and $\mathcal{W}$ are adjusted to agree with the above definitions when s_L and U are modified.

To see that this is the case, note that either both tests preceding step 5 succeed, so that steps 5–8 are executed, the first test succeeds and the second fails, so that step 9 is executed, or the first test fails so that steps 5–9 are skipped.

Suppose that steps 5–8 are executed, so that s_L is incremented to have value $\widehat{s_L} = s_L + 1$. Let $\widehat{\mathcal{V}}$ be the vector space spanned by the vectors $\sigma_{i,j}$ for $0 \le i \le \widehat{s_L}$ and $1 \le j \le \ell$ after steps 5–8 have been carried out. Then, since $L_{\widehat{s_L}}$ is initialized to be the sequence

$$\langle A^t\sigma_{\widehat{s_L}-1,1}, A^t\sigma_{\widehat{s_L}-1,2}, \ldots, A^t\sigma_{\widehat{s_L}-1,\ell}\rangle$$

at step 5, and the orthogonalization at step 7 does not change the vector space spanned by the vectors $\sigma_{i,j}$ for $0 \le i \le \widehat{s_L} - 1 = s_L$ and $1 \le j \le \ell$, the vector space $\widehat{\mathcal{V}}$ contains the original vector spaces $\mathcal{V}$ and $\mathcal{W}$ as well as each vector $A^t\sigma_{\widehat{s_L}-1,j}$ for $1 \le j \le \ell$. As a result (again, since the vectors $\sigma_{i,j}$ are not changed by these steps, for $i < \widehat{s_L}$) one can see that $A^t\sigma_{i,j} \in \widehat{\mathcal{V}}$ for $0 \le i \le \widehat{s_L} - 1$ and $1 \le j \le \ell$ as needed to establish the analogue of claim (a), above.

Notice next that, following step 8, the set U is updated to be

$$\widehat{U} = U \cup \{(\widehat{s_L}, j) \mid 1 \le j \le \ell\}.$$

Let $\widehat{\mathcal{W}}$ be the vector space spanned by the vectors $A^t\sigma_{\widehat{s_L},j}$ such that $1 \le j \le \ell$ and $(\widehat{s_L}, j) \notin \widehat{U}$. Then, since $(\widehat{s_L}, j) \in \widehat{U}$ for all j, $\widehat{\mathcal{W}} = \{0\}$. However, as noted above, $A^t\sigma_{\widehat{s_L}-1,j} \in \widehat{\mathcal{V}}$, so one can choose $\eta_j = A^t\sigma_{\widehat{s_L}-1,j}$ and $\theta_j = 0$ for $1 \le j \le \ell$ in order to establish the analogue of claim (b).

Suppose next that step 9 is executed; then `failure` is reported and the computation ends at this point. Step 10 is not reached and, since there is no subsequent execution of the loop body, nothing must be proved.

The desired result is trivial in the remaining case (steps 5–9 are all skipped), since s_L, the vectors $\sigma_{i,j}$ for $0 \le i \le s_L$ and $1 \le j \le \ell$, and the vector spaces $\mathcal{V}$ and $\mathcal{W}$ are all unchanged. Set $\widehat{s_L} = s_L$, $\widehat{\mathcal{V}} = \mathcal{V}$, $\widehat{U} = U$, and $\widehat{\mathcal{W}} = \mathcal{W}$ in this case.

It remains to consider the effect of step 10. Let $\mathcal{V}'$ be the vector space spanned by the vectors $\sigma_{i,j}$ for $0 \le i \le \widehat{s_L}$ and $1 \le j \le \ell$ on completion of step 10; while some of these vectors are modified this is done as part of an orthogonalization process, so that the vector space spanned by these vectors is not changed: $\mathcal{V}' = \widehat{\mathcal{V}}$. Now if $0 \le i \le \widehat{s_L} - 2$ and $1 \le j \le \ell$ then, as observed at the end of Section A.2, $(i, j) \notin \widehat{U}$ at the beginning of the execution of step 10. Consequently if $i \le \widehat{s_L} - 2$ then neither the vector $\sigma_{i,j}$ nor $A^t\sigma_{i,j}$ is changed as step 10 is carried out. It follows that $A^t\sigma_{i,j} \in \mathcal{V}'$. Thus part (a) of the claim is satisfied once again at the beginning of the next execution of the body of the loop (if, indeed, this statement is reached), as needed.

It remains to argue that the analogue of claim (b) holds at the end of the execution of step 10. That is, we must show that if $1 \le j \le \ell$ then there exist vectors $\eta_j \in \mathcal{V}' = \widehat{\mathcal{V}}$ and $\theta_j \in \mathcal{W}'$ such that $A^t\sigma_{\widehat{s_L}-1,j} = \eta_j + \theta_j$, where $U' \subseteq \widehat{U}$ is the set that has replaced $\widehat{U}$ and where $\mathcal{W}'$ is the vector space spanned by the vectors $A^t\sigma_{\widehat{s_L},j}$ such that $\sigma_{\widehat{s_L},j} \notin U'$.

Recall that $\widehat{\mathcal{W}}$ and $\mathcal{W}'$ are each defined as the vector spaces spanned by sets of vectors. Since $\widehat{\mathcal{W}}$ is defined in terms of vectors that are not in $\widehat{U}$ none of these vectors is changed by the computations that follow. Each of these vectors is part of the spanning set used to define $\mathcal{W}'$ so that

$$\widehat{\mathcal{W}} \subseteq \mathcal{W}'.$$

Now, the orthogonalization process included in step 10 can be modelled as a sequence of updates

$$\sigma_{h,i} := \sigma_{h,i} + \alpha\sigma_{j,k}$$

where $\alpha \in \mathsf{F}_q$, $(h, i) \in U'$, and where $(j, k) \in \widehat{U} \setminus U'$. Let us consider the effect of each change on the decomposition of the vector $A^t\sigma_{\widehat{s_L}-1,m}$ for $1 \le m \le \ell$ whose existence is to be maintained.

Since

$$\begin{aligned}\widehat{U} \subseteq \{(\widehat{s_L}-1,1), (\widehat{s_L}-1,2), \ldots, (\widehat{s_L},\ell),\\ (\widehat{s_L},1), (\widehat{s_L},2), \ldots, (\widehat{s_L},\ell)\},\end{aligned}$$

$h \in \{\widehat{s_L}-1, \widehat{s_L}\}$ and $j \in \{\widehat{s_L}-1, \widehat{s_L}\}$. Each of these cases is considered separately below.

Case: $h = j = \widehat{s_L}-1$. In this case an update

$$\sigma_{\widehat{s_L}-1,i} := \sigma_{\widehat{s_L}-1,i} - \alpha\sigma_{\widehat{s_L}-1,k}$$

is performed, where $\alpha \in \mathsf{F}_q$, $(\widehat{s_L}-1,i) \in U'$ and $(\widehat{s_L}-1,k) \in \widehat{U}\backslash U'$, so that $i \neq k$. The spaces $\widehat{\mathcal{V}}$ and $\widehat{\mathcal{W}}$ are not changed by this update (that is, the sets of vectors used to define these span the same vector spaces as they did before the update took place), and the vectors $A^t\sigma_{\widehat{s_L}-1,t}$ are not modified for any integer t such that $1 \leq t \leq \ell$ and $t \neq i$. It is therefore sufficient in this case to confirm that $A^t\sigma_{\widehat{s_L}-1,i}$ can still be written as the sum of vectors in $\widehat{\mathcal{V}} = \mathcal{V}'$ and $\widehat{\mathcal{W}} \subseteq \mathcal{W}'$ after this update.

Both $A^t\sigma_{\widehat{s_L}-1,i}$ and $A^t\sigma_{\widehat{s_L}-1,k}$ can be written in this way *before* the update performed; suppose that

$$A^t\sigma_{\widehat{s_L}-1,i} = \eta_i + \theta_i \quad \text{and} \quad A^t\sigma_{\widehat{s_L}-1,j} = \eta_j + \theta_j$$

where $\eta_i, \eta_j \in \widehat{\mathcal{V}}$ and $\theta_i, \theta_j \in \widehat{\mathcal{W}}$ before the update. Then, following the update,

$$A^t\sigma_{\widehat{s_L}-1,i} = \eta_i' + \theta_i'$$

for $\eta_i' = \eta_i - \alpha\eta_j \in \widehat{\mathcal{V}}$ and $\theta_i' = \theta_i - \alpha\theta_j \in \widehat{\mathcal{W}}$, as required.

Case: $h = \widehat{s_L}$ and $j = \widehat{s_L}-1$. In this case an update

$$\sigma_{\widehat{s_L},i} := \sigma_{\widehat{s_L},i} - \alpha\sigma_{\widehat{s_L}-1,k}$$

is being performed, where $(\widehat{s_L},i) \in U' \subseteq \widehat{U}$. The vector $A^t\sigma_{\widehat{s_L},i}$ is not included in the spanning set that defines $\widehat{\mathcal{W}}$, so the set of vectors used to define $\widehat{\mathcal{W}}$ spans the same vector space as before. None of the vectors $A^t\sigma_{\widehat{s_L}-1,j}$ are changed, for any of the integers j such that $1 \leq j \leq \ell$, so the analogue of claim (b) continues to hold.

Case: $h = j = \widehat{s_L}$. In this case an update

$$\sigma_{\widehat{s_L},i} := \sigma_{\widehat{s_L},i} - \alpha\sigma_{\widehat{s_L},k}$$

is being performed, where $(\widehat{s_L},i) \in U'$ and $(\widehat{s_L},k) \in \widehat{U} \setminus U'$, so that $i \neq k$. Since $(\widehat{s_L},i) \in \widehat{U}$ and $(\widehat{s_L},k) \in \widehat{U}$, neither $A^t\sigma_{\widehat{s_L},i}$ nor $A^t\sigma_{\widehat{s_L},k}$ is included in the spanning set used to define $\widehat{\mathcal{W}}$, so this set of vectors spans the same vector space as it did before the update. As in the previous case, none of the vectors $A^t\sigma_{\widehat{s_L}-1,j}$ such that $1 \leq j \leq \ell$ is changed by this update, so the analogue of part (b) of the claim continues to hold, once again.

Case: $h = \widehat{s_L}-1$ and $j = \widehat{s_L}$. In this case an update

$$\sigma_{\widehat{s_L}-1,i} := \sigma_{\widehat{s_L}-1,i} - \alpha_{i,k}\sigma_{\widehat{s_L},k}$$

is being considered, where $\alpha_{i,k} \in \mathsf{F}_q$, $(\widehat{s_L}-1,i) \in U'$ and $(\widehat{s_L},k) \in \widehat{U} \setminus U'$. Unlike the previous cases we will consider the effects of all such updates at once: For every integer i such that $(\widehat{s_L}-1,i) \in U'$, the vector $\sigma_{\widehat{s_L}-1,i}$ is updated as follows:

$$\sigma_{\widehat{s_L}-1,i} := \sigma_{\widehat{s_L}-1,i} - \sum_{(\widehat{s_L},k)\in\widehat{U}\backslash U'} \alpha_{i,k}\sigma_{\widehat{s_L},k}.$$

Suppose that $A^t\sigma_{\widehat{s_L}-1,i} = \eta_i + \theta_i$ where $\eta_i \in \widehat{\mathcal{V}} = \mathcal{V}'$ and $\theta_i \in \widehat{\mathcal{W}}$ before these updates, for $1 \leq i \leq \ell$. Then, following the updates,

$$A^t\sigma_{\widehat{s_L}-1,i} = \eta_i + \theta_i'$$

where

$$\theta_i' = \theta_i - \sum_{(\widehat{s_L},k)\in\widehat{U}\backslash U'} \alpha_{i,k}A^t\sigma_{\widehat{s_L},k}.$$

Now it suffices to notice that $\theta_i' \in \mathcal{W}'$: $\theta_i \in \mathcal{W}'$ since $\theta_i \in \widehat{\mathcal{W}}$ and $\widehat{\mathcal{W}} \subseteq \mathcal{W}'$. Each of the vectors $A^t\sigma_{\widehat{s_L},k}$ such that $(\widehat{s_L},k) \in \widehat{U} \setminus U'$ belongs to $\mathcal{W}'$ as well, because each is included in the spanning set that defines $\mathcal{W}'$. Consequently $A^t\sigma_{\widehat{s_L}-1,i}$ is the sum of vectors in $\widehat{\mathcal{V}} = \mathcal{V}'$ and in $\mathcal{W}'$ as required.

Consequently the analogue of part (b) is satisfied at the end of step 10 (provided that these computations are actually carried out, that is, that algorithm does not report `failure` and terminate instead). Consequently the claims also hold at the beginning of the $i+1^{\text{st}}$ execution of the loop body, if such an execution takes place. □

Note: Once again, it is easily checked from the above that relationship claimed to hold at the beginning of the execution of the loop body also holds later on. It also holds at the point when an elimination phase begins (if there is one), and this will be used in later proofs.

Lemma 15. *At the beginning of each execution of the body of the loop of the algorithm shown in Figure 1, $s_L \geq 0$ and (for S_1 as defined at line* (1) *and $m_1 = |S_1|$)*

$$\begin{aligned}\{\mu_i \mid 1 \leq i \leq m_1\} \cup \{\sigma_{i,j} \mid (i,j) \in U\} = \\ \{\sigma_{i,j} \mid 0 \leq i \leq s_L \text{ and } 1 \leq j \leq \ell\} \quad (17)\end{aligned}$$

This relationship also holds at the beginning of the elimination phase of the algorithm (if, indeed, there is one).

Proof. This is another straightforward proof by induction on the number of executions of the body of the loop.

At the beginning of the first execution of the loop body $S_1 = \emptyset$ so that $m_1 = 0$ and

$$\{\mu_i \mid 1 \leq i \leq m_1\} = \emptyset;$$

$s_L = 0$ and $U = \{(0,j) \mid 1 \leq j \leq \ell\}$, so that

$$\{\sigma_{i,j} \mid (i,j) \in U\} = \{\sigma_{i,j} \mid 0 \leq i \leq s_L \text{ and } 1 \leq j \leq \ell\}$$

so that equation (17) holds at this point.

Suppose that equation (17) is correct at the beginning of the i^{th} execution of the body of the loop for a nonnegative integer i. Either both tests before step 5 succeed during the i^{th} execution of the loop body, so that steps 5–8 are executed, the first test succeeds and the second fails, so that step 9 is executed, or the first test fails and steps 5–9 are skipped.

If steps 5–8 are executed then s_L is incremented to have value $\widehat{s_L} = s_L + 1$ during the execution of these steps. Neither S_1 nor any vector $\sigma_{i,j}$ such that $1 \leq i \leq s_L$ and $(i,j) \notin U$ is changed; and the ordered pairs

$$(\widehat{s_L},1), (\widehat{s_L},2), \ldots, (\widehat{s_L},\ell)$$

are added to U (which is otherwise unchanged). Consequently the equation at line (17) is satisfied once again, on completion of step 8, provided that the references to s_L, S_1 (and m_1), and U refer to the values that this integer variable and pair of sets have on the completion of this step.

If step 9 is executed then there is no $i+1^{\text{st}}$ execution of the body of the loop, because `failure` is reported and the computation ends, so nothing must be proved.

Finally, it is clear that the equation at line (17) is satisfied before the beginning of step 10 if steps 5–9 are skipped, since s_L, S_1 (and m_1), U, and the vectors $\sigma_{i,j}$ such that $0 \le i \le s_L$ and $1 \le j \le \ell$ are unchanged between the beginning of this execution of the loop body and the beginning of step 10 in this case.

It remains only to consider the execution of step 10. Note that s_L and the vectors $\sigma_{i,j}$ such that $0 \le i \le s_L$ and $1 \le j \le \ell$ are not changed by this step, so that the set

$$\{\sigma_{i,j} \mid 0 \le i \le s_L \text{ and } 1 \le j \le \ell\}$$

is not changed. On the other hand, S_1 is changed (if an $i+1^{\text{st}}$ execution of the loop body will take place): A set of vectors

$$(\mu_{m_1+1}, \nu_{m_1+1}), (\mu_{m_1+2}, \nu_{m_1+2}), \dots, (\mu_{m_1+2}, \nu_{m_1+s})$$

is added to this set, where

$$\mu_{m_1+1} = \sigma_{i_1,j_1}, \mu_{m_1+2} = \sigma_{i_2,j_2}, \dots, \mu_{m_1+s} = \sigma_{i_s,j_s}$$

for distinct ordered pairs

$$(i_1, j_1), (i_2, j_2), \dots, (i_s, j_s)$$

that belonged to U at the beginning of the execution of this step. These are removed from U (which is otherwise unchanged) if the loop body will be executed again; U is not changed at all if this is the final execution of this loop body. It follows that the execution of step 10 does not change the set

$$\{\mu_i \mid 1 \le i \le m_1\} \cup \{\sigma_{i,j} \mid (i,j) \in U\}$$

either — it merely moves a set of vectors from the second set forming the above union into the first. Thus the equation at line (17) is satisfied at the end of step 10 and at the beginning of the next execution of the loop body — if, indeed, it is executed again — or at the beginning of the elimination phase, as required to complete the proof. □

The next lemma concerns two unusual cases. The more usual situation is considered in the lemma that follows.

Lemma 16. *Suppose that body of the loop of the algorithm shown in Figure 1 is only executed once.*

(a) *Computation ends at step 4 if and only if $v_i = 0$ for $1 \le i \le r$.*

(b) *If computation ends at step 10 then, for $1 \le i \le r$, v_i is a linear combination of the vectors in the set*

$$\{\nu_i \mid 1 \le i \le |S_1|\} \cup S_2 \cup S_3$$

for the sets S_1, S_2, and S_3 as these are defined at the beginning of the execution phase of the algorithm.

Proof. This follows by inspection of the code and the description of each step found in Section 2. □

Lemma 17. *Let i be a positive integer such that $i \ge 2$ and the body of the loop is executed is executed at least i times. Let*

$$W_{1,i} = \{\nu_j \mid 0 \le j \le m_1\}$$

where $m_1 = |S_1|$, for S_1 as shown at line (1) and as defined at the beginning of the i^{th} execution of the loop. Let $S_{2,i}$ be the set S_2 as it is defined at the beginning of the i^{th} execution of the loop.

(a) *Suppose there is an $i+1^{st}$ execution of the loop as well, and let $W_{1,i+1}$ and $S_{2,i+1}$ be the corresponding sets defined at the beginning of the $i+1^{st}$ execution of the loop. Then $A\nu$ is a linear combination of the vectors in $W_{1,i+1} \cup S_{2,i+1}$ for every vector $\nu \in W_{1,i}$.*

(b) *If the i^{th} execution of the body of the loop ends at step 4 then $A\nu$ is a linear combination of the vectors in $W_{1,i} \cup S_{2,i}$ for all $\nu \in W_{1,i}$.*

(c) *Finally, if the i^{th} execution of the body of the loop ends at step 10 then $A\nu$ is a linear combination of the vectors in $\widehat{W}_1 \cup \widehat{S}_2 \cup \widehat{S}_3$ for every vector $\nu \in W_{1,i}$, where*

$$\widehat{W}_1 = \{\nu_j \mid 0 \le j \le |\widehat{S}_1|\}$$

for where $\widehat{S}_1$ (respectively, $\widehat{S}_2$ and $\widehat{S}_3$) is the set assigned to the variable S_1 (respectively, S_2 and S_3) at the beginning of the elimination phase of the algorithm.

Proof. Each part of the claim can be established by noticing, on inspection of the code, that if the body of the loop of the algorithm is executed at least k times, for $k \ge 2$, then there exist sets

$$X_2, X_3, \dots, X_k$$

of vectors such that

$$W_{1,j} = X_2 \cup X_3 \cup \cdots \cup X_j$$

for $2 \le j \le k$, that is, $W_{1,j} \subseteq W_{1,j+1}$ for $2 \le j \le k-1$, and that $S_{2,j} \subseteq S_{2,j+1}$ for $2 \le j \le k-1$ as well. Indeed (noting that both S_1 and S_2 are empty sets at the beginning of the first execution) we may set $W_{1,1} = S_{2,1} = \emptyset$, so the above inclusions are now satisfied when $j = 1$ as well.

(a) The first part of the claim is easily established by induction on i, using the above and by inspection of the code.

Suppose, for a basis, that $i = 2$ and that the body of the loop is executed at least three times. Suppose that ν is a vector in $W_{1,2}$; then one can see by inspection of the code that $A\nu$ is one of the vectors included in the initialization of set R when step 3 is reached, during the second execution of the body of the loop. Consequently this belongs to the vector space spanned by the vectors in $W_{1,2} \cup R$ at this point. The orthogonalization of vectors in R performed at step 4 does not change this vector space, so that $A\nu$ belongs to the vector space spanned by the vectors in $W_{1,2} \cup R$ at the end of this step as well. Since there is a third execution of the loop one can see by inspection of the code (including the updating of S_1 and S_2) that each vector in R, at the end of step 4, is a linear combination of the vectors in X_3 that are added to $W_{1,3}$ and the vectors added to $S_{2,2}$ to produce $S_{2,3}$. It follows that $A\nu$ is a linear combination of the vectors in $W_{1,3} \cup S_{2,3}$, as required.

Now suppose that $i \ge 2$ and the properties in part (a) of the claim hold for i. Since there is nothing further to prove if the body of the loop is not executed at least $i+2$ times, let us assume that an $i+2^{\text{nd}}$ execution takes place.

Let ν be a vector in $W_{1,i+1}$. Then either $\nu \in W_{1,i}$ or $\nu \in X_{i+1} = W_{1,i+1} \setminus W_{1,i}$.

In the first case it follows by the inductive hypothesis that $A\nu$ is a linear combination of vectors in $W_{1,i+1}$ and $S_{2,i+1}$. Since $W_{1,i+1} \subseteq W_{1,i+2}$ and $S_{2,i+1} \subseteq S_{2,i+2}$, it follows that $A\nu$ is in the linear span of the vectors in $W_{1,i+2} \cup S_{2,i+2}$ as required.

On the other hand, if $\nu \in X_{i+1}$ then $A\nu$ is one of the vectors included in the set R when step 3 is executed, as part of the $i+1^{\text{st}}$ execution of the loop body. The argument given in the basis can now be applied to conclude that $A\nu$ is in the linear span of the vectors in $W_{1,i+2} \cup S_{2,i+2}$ in this case as well.

(b) Suppose that $i \geq 2$ and the i^{th} execution of the body of the loop ends at step 4: Then all the vectors included in the set R at step 3 have been set to zero after the orthogonalization process in step 4. Let $\nu \in W_{1,i}$. Then, as observed in the proof of part (a), $A\nu$ is a linear combination of the vectors in $W_{1,i} \cup S_{2,i}$ if $\nu \in W_{1,i-1}$, and $A\nu$ is one of the vectors used to initialize R (in step 3) if $\nu \in W_{1,i} \setminus W_{1,i-1}$. Since the orthogonalization process updates every vector in R by adding a linear combination of the vectors in $W_{1,i}$, the fact that $A\nu$ has been replaced by zero implies that $A\nu$ is a linear combination of these vectors, as claimed.

(c) Finally, suppose that the i^{th} execution of the body of the loop ends at step 10 and is followed by the beginning of the elimination phase. Once again let $\nu \in W_{1,i}$; then either $\nu \in W_{1,i-1}$ or $\nu{\in}W_{1,i} \setminus W_{1,i-1}$. In the former case the claim follows because $A\nu \in W_{1,i} \cup S_{2,i}$, as argued above, and because $W_{1,i} \subseteq \widehat{W}_1$ and $S_{2,i} \subset \widehat{S}$. In the latter case $A\nu$ is one of values include in the set R at step 3, and an inspection of the orthogonalization process in step 4 and the update process at step 10 (prior to the beginning of the elimination phase) confirms that $A\nu$ is a linear combination of vectors in $\widehat{W}_1 \cup \widehat{S}_2 \cup \widehat{S}_3$ in this case as well. □

Lemma 18. *Let i be a positive integer such that the body of the loop in the elimination phase of the algorithm is executed at least i times. Let $W_1 = \{\nu_i \mid 1 \leq i \leq |S_1|\}$ for the set S_1 as it is defined at the end of the elimination phase, and let $S_{2,i}$ (respectively, $S_{3,i}$) be the sets of vectors assigned to the variable S_2 (respectively, S_3) at the beginning of the i^{th} execution of the loop body.*

(a) If there is also an $i+1^{st}$ execution of the body of this loop and $\nu \in W_1 \cup S_{3,i}$ then $A\nu$ is a linear combination of vectors in $W_1 \cup S_{2,i+1} \cup S_{3,i+1}$.

(b) If the algorithm ends immediately after the i^{th} execution of the body of the loop then, if S_2 and S_3 are as defined at end of this execution, and $\nu \in W_1 \cup S_3$ then $A\nu$ is a linear combination of the vectors in $W_1 \cup S_2 \cup S_3$.

Proof Sketch. Virtually the same argument as used to proved parts (a) and (b) of the previous lemma can be used to establish this one, as well: One makes use of the containments $S_{2,i-1} \subseteq S_{2,i}$ and $S_{3,i-1} \subseteq S_{3,i}$ and a consideration of the way that vectors in $S_{3,i} \setminus S_{3,i-1}$ are used to update the sets S_2 and S_3 during an i^{th} execution of the body of this loop. □

Proof of Lemma 3. Both claims can be proved by induction on the number of executions of the body of each in the loops in the algorithm described in Section 2.

(a) Consider first the body of the loop in the algorithm shown in Figure 1. The first claim is trivially satisfied at the beginning of the first execution of the body of this loop, since $S_1 = \emptyset$ (so that $m_1 = 0$), $s_L = 0$, $\sigma_{0,j} = u_j$ for $1 \leq j \leq \ell$, and $U = \{(0,j) \mid 1 \leq j \leq \ell\}$:

$$V_1 = V_2 = \{u_1, u_2, \ldots, u_\ell\}.$$

Suppose now that $i \geq 1$, the body of the loop is executed at least i times, and sets of vectors V_1 and V_2 span the same vector space at the beginning of the i^{th} execution of the loop.

During this execution of the loop, either both tests before step 5 succeed and steps 5–8 are executed, the first test succeeds and the second fails, so that step 9 is executed, or the first test fails and steps 5–9 are skipped.

Suppose first that both tests succeed and steps 5–8 are executed. In this case, s_L is incremented, and replaced by $\widehat{s_L} = s_L + 1$, so that the set of vectors V_2 mentioned in the claim is replaced by the set

$$\widehat{V_2} = \{(A^t)^i u_j \mid 0 \leq i \leq \widehat{s_L} \text{ and } 1 \leq j \leq \ell\}.$$

On the other hand, during this computation, each of the vectors $\sigma_{\widehat{s_L},j}$ is initialized to be equal to $A^t \sigma_{s_L,j}$ for $1 \leq j \leq \ell$, and the ordered pairs $(\widehat{s_L}, j)$ are added to U for $1 \leq j \leq \ell$. Notice that, by Lemma 15, $\sigma_{s_L,j} \in V_1$ for $1 \leq j \leq \ell$ so it is clear by the inductive hypothesis that $A^t\sigma_{s_L,j}$ is a linear combination of the vectors in $\widehat{V}_2$. Since the orthogonalization process at step 7 does not change the vector space spanned by the vectors being processed, it should be clear that steps 5–8 replace the set V_1 by another set of vectors, $\widehat{V}_1$ that spans a subspace of the vector space spanned by $\widehat{V}_2$.

Since $\widehat{V}_1$ includes all the vectors in V_1 it is clear, by the inductive hypothesis, that the vector $(A^t)^i v_j$ is in the linear span of the vectors in $\widehat{V_1}$, for $0 \leq i \leq s_L$ and $1 \leq j \leq \ell$. It is therefore necessary and sufficient to show that $(A^t)^{\widehat{s_L}} u_j = (A^t)^{s_L+1} u_j$ is in the linear span of these vectors, as well, in order to establish that the vectors in $\widehat{V_1}$ and in $\widehat{V_2}$ span the same vector space.

Consider the vector $(A^t)^{s_L} u_j$; since this is in V_2 it follows by the inductive hypothesis that this is in the linear span of the vectors in V_1 and the alternative characterization of this set given by Lemma 15 can be used to conclude that

$$(A^t)^{s_L} u_j = \omega_1 + \omega_2 + \omega_3$$

where

- ω_1 is a linear combination of the vectors $\sigma_{h,i}$ such that $0 \leq h \leq s_L - 2$ and $1 \leq i \leq \ell$,
- ω_2 is a linear combination of the vectors $\sigma_{s_L-1,i}$ such that $1 \leq i \leq \ell$, and
- ω_3 is a linear combination of the vectors $\sigma_{s_L,i}$ such that $1 \leq i \leq \ell$.

Consequently $(A^t)^{\widehat{s_L}} u_j = A^t\omega_1 + A^t\omega_2 + A^t\omega_3$. It follows by part (a) of Lemma 14 that $A^t\omega_1$ is an element of the vector space spanned by the vectors in V_1, so it belongs to the vector space spanned by the vectors in $\widehat{V_1}$ as well. Part (b) of Lemma 14, and the fact that the orthogonalization process at step 7 does not modify the vector space spanned, imply that $A^t\omega_2$ is in the vector space spanned by the vectors in $\widehat{V_1}$ as well. Since $A^t\sigma_{s_L,i}$ was used as the initial value for $\sigma_{\widehat{s_L},i}$, the above observations and characterization of ω_3 imply that $A^t\omega_3$ is in the vector space spanned by the vectors in $\widehat{V_1}$ too. Consequently $(A^t)^{\widehat{s_L}}\sigma_j$ is also in the vector space spanned by these vectors, so that the vectors in $\widehat{V_1}$ and in $\widehat{V_2}$ span the same vector space.

Suppose next that step 9 is executed; then there is no $i+1^{\text{st}}$ execution of the loop, or elimination phase, so there is nothing that must be proved.

Finally, if steps 5–9 are skipped entirely then the sets V_1 and V_2 are not changed, during the execution of the loop body, before step 10 is reached.

Thus the vector spaces (corresponding to) V_1 and V_2 immediately before the execution of step 10 span the same vector space, as required.

Now it suffices to note that, while step 10 includes an orthogonalization process that modifies some of the vectors in V_1 and V_2 this does not modify the vector space that either set of vectors spans. Since these sets of vectors are not changed between the end of step 10 and the beginning of the next execution of the loop body, it follows that the vector spaces V_1 and V_2 span the same vector space at the beginning of the $i+1^{\text{st}}$ execution of the loop body (if there is such an execution), or at the beginning of the elimination phase of the algorithm (if there is not). Since S_1 and U are not modified at all during the elimination phase, V_1 and V_2 are not changed either, so they span the same vectors space throughout the elimination phase.

(b) The proof of part (b) is similar to that of part (a): We proceed, again, by induction on the number of executions of the body of each of the loops in the algorithm. To begin, consider the body of the loop of the Lanczos phase of the algorithm as this is shown in Figure 1.

If the body of the loop is only executed once before the algorithm moves to the elimination phase then it is clear that the vectors in W_1 must span a subspace of that spanned by $W_2 = \{v_1, v_2, \ldots, v_r\}$, since each of the vectors in W_1 is a linear combination of the vectors in W_2. Part(b) of Lemma 16 now implies that W_1 and W_2 span the same vector space.

One can see by inspection of the code (echoing the argument used to establish the above lemma) that W_1 and W_2 span the same vector space, namely, the space spanned by $v_1, v_2, \ldots, v_r$, at the beginning of the second execution of the body of this loop whenever this loop is executed two or more times.

It is also clear, since $s_R = -1$ at the beginning of the first execution of the loop body and since s_R is incremented during every such execution, that $s_R = i-1$, and

$$W_2 = \{A^j v_k \mid 0 \le j \le i-1 \text{ and } 1 \le k \le r\}$$

at the beginning of the i^{th} execution of the loop body whenever this loop body is executed i or more times.

Finally, a containment relation is clear. Suppose that $i \ge 1$ and the body of the loop is executed i or more times; let $W_{1,i}$ (respectively, $W_{i,i+1}$) be the set W_1 as this is defined at the beginning of the i^{th} (respectively, $i+1^{\text{st}}$) execution of the body of the loop. Then $W_{1,i} \subseteq W_{1,i+1}$.

With all that noted, let i be an integer such that $i \ge 1$, the sets of vectors W_1 and W_2 that are defined at the beginning of the i^{th} execution of the body of the loop span the same vectors space, and suppose the body of the loop is executed at least $i+1$ times. Notice that the vector space spanned by $W_{1,i+1}$ is a subspace of the vector space spanned by the vectors in the set

$$W_{1,i} \cup \{A\gamma \mid \gamma \in W_{1,i}\}.$$

Since $W_{1,i}$ spans the same space as the set of vectors $A^j v_k$ such that $0 \le j \le i-1$ and $1 \le k \le \ell$ it follows from the above that $W_{1,i+1}$ must span a subspace of the space spanned by the vectors $A^j v_k$ for $0 \le j \le i$ and $1 \le k \le \ell$. It is necessary and sufficient to establish that $A^j v_k$ is in the space spanned by W_{i+1}, for $0 \le j \le i$ and $1 \le k \le \ell$, in order to establish that W_1 and W_2 span the same vector space at the beginning of the $i+1^{\text{st}}$ execution of the loop.

Suppose first that $0 \le j \le i-1$; then it follows by the inductive hypothesis that $A^j v_k$ is in the space spanned by $W_{1,i}$, and the containment of $W_{1,i}$ in $W_{1,i+1}$ establishes the desired result.

On the other hand, if $j = i$ and $i \ge 1$, as assumed above, then $A^j v_k = A(A^{j-1}v_k) = A\gamma$ for a vector γ that is in the vector space spanned by $W_{1,i}$. It follows by part (a) of Lemma 17 that $A\gamma$ is in the space spanned by $W_{1,i+1}$ for each vector $\gamma \in W_{1,i}$. All linear combinations of these vectors must be in the space spanned by $W_{1,i+1}$ so, in particular, $A^j v_k = A^i v_k$ must be in this space. It follows that W_1 and W_2 span the same vector space at the beginning of the $i+1^{\text{st}}$ execution of the body of the loop.

It follows by the same argument that W_1 and W_2 span the same vector space at the end of the final execution of this loop, so they also span the same vector space at the beginning of the elimination phase.

Precisely the same argument, using Lemma 18 in place of Lemma 17 (and recalling that S_2 is a subset of the null space of A) can be applied to establish that W_1 and W_2 span the same vector space at the beginning of every execution of the loop in the elimination phase of the algorithm, and on its termination. □

Proof of Lemma 4. Consider the set of vectors

$$\{\nu_i \mid \ \le i \le m_1\} \cup S_2 \cup S_3$$

mentioned in the claim; it follows by part (b) of Lemma 3 that each of the vectors in this set belongs to the Krylov space generated by $v_1, v_2, \ldots, v_r$.

(a) This follows by part (b) of Lemma 17 and part (b) of Lemma 18: Regardless of whether termination ends at step 4 of the body of the loop in the first phase or on completion of the loop in the second phase, the vector space spanned by

$$\{\nu_i \mid 1 \le i \le m_1\} \cup S_2 \cup S_3$$

is contained in the Krylov space, includes each of the vectors $v_1, v_2, \ldots, v_r$ and, by the above lemma, is closed under multiplication by A. The only such space is the Krylov space itself.

(b) Clearly, each of the vectors in S_2 is in the intersection of the Krylov space and the null space of A, so the vector space that is spanned by this set of vectors is a subspace of this intersection.

Suppose now that λ is an element of the intersection of the Krylov space and the null space of A. Then it follows by part (a) that

$$\lambda = \sum_{h=1}^{m_1} c_h \nu_h + \kappa + \omega \tag{18}$$

where $c_1, c_2, \dots, c_{m_1} \in \mathsf{F}_q$ and where κ is a linear combination of the vectors in S_2, and where ω is a linear combination of the vectors in S_3. Now it follows by the orthogonality relations given at line (2), the fact that each vector in S_2 is in the null space of A, and the fact that $\mu_i^t A \eta = 0$ for $1 \le i \le |S_1|$ and $\eta \in S_3$, that, for $1 \le i \le m_1$,

$$\begin{aligned}
\mu_i^t A \lambda &= \mu_i^t a \left(\sum_{h=1}^{m_1} c_h \nu_h \right) + \mu_i^t A \kappa + \mu_i^t A \omega \\
&= \sum_{h=1}^{m_1} c_h \mu_i^t A \nu_h + 0 \\
&\qquad\qquad (\text{since } A\kappa = 0 \text{ and } \mu_i^t A \omega = 0) \\
&= c_i \mu_i^t A \nu_i \qquad (\text{since } \mu_i^t A \nu_h = 0 \text{ if } h \ne i) \\
&= c_i.
\end{aligned}$$

On the other hand, $\mu_i^t A \lambda = 0$, since λ is in the null space of A. Thus $c_i = 0$ for $1 \le i \le m_1$, so that $\lambda = \kappa + \omega$ is a linear combination of the vectors in S_2 and S_3.

Suppose that $S_3 = \{\zeta_1, \zeta_2, \dots, \zeta_{m_3}\}$ so that

$$\omega = d_1 \zeta_1 + d_2 \zeta_2 + \dots, + d_{m_3} \zeta_{m_3}$$

for $d_1, d_2, \dots, d_{m_3} \in \mathsf{F}_q$. Since κ and λ are both in the null space of A, $\omega = \lambda - \kappa$ is in the null space as well, so

$$d_1 A \zeta_1 + d_2 A \zeta_2 + \dots + d_{m_1} A \zeta_{m_1} = A\omega = 0.$$

However, it follows by the construction of S_3 that the vectors $A\zeta_1, A\zeta_2, \dots, A\zeta_{m_1}$ are linearly independent, so that $d_1 = d_2 = \dots = d_{m_1} = 0$ and $\omega = 0$ as well. Thus $\lambda = \kappa$ so that λ is a linear combination of the vectors in S_2, as claimed. □

A.4 Proof of Lemmas 5 and 6

Lemma 19. *Consider the value of s_L and the set S_1 at the beginning of any execution of the loop body in the algorithm shown in Figure 1, and at the beginning of the elimination phase.*

If i is an integer such that $0 \le i \le s_L - 4$ and $1 \le j \le \ell$ then

(a) $\sigma_{i,j} = \mu_k$ for some ordered pair $(\mu_k, \nu_k) \in S_1$.

(b) $A^t \sigma_{ij}$ is a linear combination of the first entries in the ordered pairs that have been included in S_1:

$$A^t \sigma_{i,j} = \sum_{h=1}^{m_1} c_h \mu_h$$

where $c_h \in \mathsf{F}_q$ for $1 \le h \le m_1$.

Furthermore, if $c_h \ne 0$, for $1 \le h \le m_1$, then

$$\mu_h = \sigma_{u,v}$$

for integers u and v such that $1 \le v \le \ell$ and $0 \le u \le i + 3 \le s_L - 1$.

Proof. Recall that, for $0 \le i \le s_L$ and $1 \le j \le \ell$, either $\sigma_{i,j} \in \{\mu_1, \mu_2, \dots, \mu_{m_1}\}$ or $(i,j) \in U$. It follows by Lemma 2, and the remarks at the end of Section A.2, that $\sigma_{i,j} \in \{\mu_1, \mu_2, \dots, \mu_{m_1}\}$ whenever $i \le s_L - 2$, implying part (a). Indeed, part (a) is also satisfied at the beginning of a previous execution of the loop body, namely the first execution of the loop body such that $s_L = i + 2$. We will consider the processing of vectors between this execution of the loop body and the current one in order to complete the proof.

Let $\widehat{m_1}$ be the size of the set S_1 at the beginning of this previous execution of the loop body; then, at this point, S_1 consists of the ordered pairs (μ_h, ν_h) for $1 \le h \le \widehat{m_1}$. It follows again by Lemma 2 that $(i,j) \notin U$ at this point so that, in fact, $\sigma_{i,j} = \mu_k$ where $k \le \widehat{m_1}$. Furthermore, if $\widehat{U}$ is the set U as it is defined at the beginning of this execution of the loop body then part (a) of Lemma 14 can be used to establish that

$$A^t \sigma_{i,j} = \sum_{h=1}^{\widehat{m_1}} \widehat{c_h} \mu_h + \sum_{(h,k) \in \widehat{U}} d_{h,k} \sigma_{h,k}$$

where $\widehat{c_h} \in \mathsf{F}_q$ for $1 \le h \le \widehat{m_1}$ and $d_{h,k} \in \mathsf{F}_q$ for $(h,k) \in \widehat{U}$ as well. Since s_L currently has value $i+2$, μ_h must be equal to $\sigma_{u,v}$ where $u \le i+2$ at this point, for $1 \le h \le \widehat{m_1}$.

Consider the effect of subsequent executions of the loop body, up to the beginning of the execution when the value of s_L will be incremented to $i+3$. The orthogonalizations that are performed update vectors by adding linear combinations of the vectors μ_h for ordered pairs (μ_h, ν_h) in S_1. While additional ordered pairs can be added to S_1, they must all be of the form (μ_h, ν_h) where $\mu_h = \sigma_{u,v}$, for $u \le i + 2$, since these are the only vectors that are available before s_L is incremented. Now, since s_L is incremented during this execution of the body of the loop, steps 5–8 of the algorithm are executed, so that the tests prior to step 5 succeeded, and it is must be the case that the set U is a subset of

$$\{(i+2,1), (i+2,2), \dots, (i+2,\ell)\}$$

at this point. Consequently, if m_1' is the size of the set S_1 at this point then S_1 now consists of the first m_1' ordered pairs (μ_h, ν_h) for $1 \le h \le m_1'$ that will also be in S_1 at the beginning of subsequent executions of the loop body, and $A^t \sigma_{i,j}$ can now be expressed as

$$A^t \sigma_{i,j} = \sum_{h=1}^{m_1'} c_h' \mu_h + \sum_{k=1}^{\ell} d_k' \sigma_{i+2,k}$$

where $c_h' \in \mathsf{F}_q$ for $1 \le h \le m_1'$, $d_k' \in \mathsf{F}_q$ for $1 \le k \le \ell$ and, furthermore, each d_k' is only nonzero if $(i+2,k)$ is included in the set U at the beginning of this execution of the loop body.

Next consider the effect of subsequent executions of the loop body, up to the beginning of the execution when the value of s_L will be incremented to $i+4$. The only ordered pairs added to S_1 must be of the form (μ_h, ν_h) where $\mu_h = \sigma_{u,v}$ for $u \le i+3$ since no other vectors $\sigma_{u,v}$ are yet available. Orthogonalizations update vectors by adding linear combinations of the first entries of ordered pairs that have already been added to S_1. Finally, since step 5 of the loop body will be executed once again, all ordered pairs $(i+2,k)$ for $1 \le k \le \ell$ must have been removed from the set U at this point. Consequently, if m_1^* is the size of S_1 at this point then each of the ordered pairs (μ_h, ν_h) for $1 \le h \le m_1^*$ in S_1 at this point will also be included in S_1 during subsequent executions of the loop body, and

$$A^t \sigma_{i,j} = \sum_{h=1}^{m_1^*} c_h^* \mu_h$$

where $c_h^* \in \mathsf{F}_q$ for $1 \le h \le m_1^*$.

Now, since the claim in the statement of this lemma concerns a a vector $\sigma_{i,j}$ where i is less than or equal to the current value of s_L minus four, the claim concerns an execution of the loop body that follows the one described above. Consequently $m_1 \ge m_1^*$ at this point and the equation given in part (b) can be obtained by setting c_h to be c_h^* for $1 \le h \le m_1^*$ and setting c_h to be zero for $m_1^* + 1 \le h \le m_1$. □

Lemma 20. *Consider the set*

$$S_1 = \{(\mu_1 \nu_1), (\mu_2, \nu_2), \dots, (\mu_{m_1}, \nu_{m_1})\}$$

at any point during the execution of the algorithm in Figure 1 and at the beginning of the elimination phase of the algorithm. Suppose that $\mu_k = \sigma_{i_k, j_k}$ for $1 \le k \le m_1$.

Then, if k_1 and k_2 are integers such that $1 \le k_1 < k_2 \le m_1$, then $i_{k_2} \ge i_{k_1} - 1$.

Proof. Notice that if $k_1 < k_2$ then the ordered pair (μ_{k_2}, ν_{k_2}) was added to S_1 either during the same execution of the loop body as (μ_{k_1}, ν_{k_1}), or during an execution that followed the inclusion of (μ_{k_1}, ν_{k_1}). For all integers i_h and j_h the ordered pair (μ_h, ν_h) such that $\mu_h = \sigma_{i_h, j_h}$ is added to S_1 during the same execution of the loop body as the one in which the ordered pair (i_h, j_h) is removed from the set U.

The claim now follows from Lemma 2 and inspection of the code in Figure 1: The set U cannot include any ordered pairs (u, v) such that $u \le i_{k_1} - 2$ at any execution of the loop body that either includes the introduction of (μ_{k_1}, ν_{k_1}) to S_1 or follows it. □

Corollary 1. *Once again, consider the above set S_1 at the beginning of any execution of the loop body of the algorithm in Figure 1 and at the beginning of the elimination phase of the algorithm.*

If i is an integer such that $0 \le m_1 - 6\ell$ then $A^t \mu_i$ is a linear combination of the vectors $\mu_1, \mu_2, \dots, \mu_{i+6\ell-1}$.

Proof. We claim first that if $i \le m_1 - 6\ell$ then $\mu_i = \sigma_{j_i, k_i}$ for integers j_i and k_i such that $j_i \le s_L - 4$ and $1 \le k_i \le \ell$. Suppose otherwise: then $j_i \ge s_L - 3$ and it follows by Lemma 20, above, and an inspection of the code, that

$$\mu_{i+c} = \sigma_{j_{i+c}, k_{i+c}} \qquad \text{for } 1 \le c \le 6\ell$$

where $(j_{i+1}, k_{i+1}), (j_{i+2}, k_{i+2}), \dots, (j_{i+6\ell}, k_{i+6\ell})$ are *distinct* ordered pairs of the form (u, v), where $s_L - 4 \le u \le s_L$ and $1 \le v \le \ell$. Since only 5ℓ such ordered pairs exist this is clearly impossible.

It now follows by Lemma 19 (again, setting $\mu_i = \sigma_{j_i, k_i}$) that

$$A^t \sigma_{i,j} = \sum_{h=1}^{m_1} c_h \mu_h,$$

where $c_h \in \mathsf{F}_q$ for $1 \le h \le m_1$ and where c_h is only nonzero if $\mu_h = \sigma_{u,v}$ for integers u and v such that $1 \le u \le j_i + 3$ and $1 \le v \le \ell$.

Consider the largest integer w such that $1 \le w \le m_1$ and $\mu_w = \sigma_{i_w, j_w}$ where $i_w \le j_i + 3$. One can see, again, by Lemma 20 that each of the vectors

$$\mu_i, \mu_{i+1}, \dots, \mu_w$$

must be equal to one of the ordered pairs $\sigma_{u,v}$ where $j_i - 1 \le u \le j_i + 4$ and $1 \le v \le \ell$ and, again, no such ordered pair corresponds to more than one of the above vectors. Since only 6ℓ such ordered pairs exist it follows that $w - i \le 6\ell - 1$; the claim now follows by Lemma 19. □

Lemma 21. *Consider the sets S_1 and S_2 at the beginning of any execution of the body of the loop in the algorithm in Figure 1, and at the beginning of the elimination phase. If i is an integer such that $1 \le i \le m_1 - r$ then*

$$A\nu_i = \phi_i + \tau_i$$

where ϕ_i is a linear combination of the vectors $\nu_1, \nu_2, \dots, \nu_j$, for $j = \min(m_1, i + 2r - 1)$, and where τ_i is a linear combination of the vectors in S_2.

Proof. Note that at most r ordered pairs are added to the set S_1 during each execution of the loop body so, since $i \le m_1 - r$, the loop body has been executed at least once more since the ordered pair (μ_i, ν_i) was introduced. Indeed, the size of S_1 was at most $i + r - 1 \le m_1 - 1$ at the end of the execution of the loop body that included this ordered pair in S_1. A subsequent execution of the loop body must therefore have run to completion (without terminating the computation but, possibly, leading to the elimination phase) since at least one more ordered pair was subsequently added to S_1.

The result now follows by Lemma 17, above, and the observation that at most r more ordered pairs have been added to S_1 during this next execution of the loop body — so that the size of S_1 is at most $i + 2r - 1$ at this point. □

Proof of Lemma 5. The claims are trivial for the first execution of the body of the loop of the algorithm shown in Figure 1, since $S_1 = \emptyset$ and $m_1 = 0$ at this point.

(a) Consider first the orthogonalizations at step 4 of the body of the loop shown in Figure 1. Suppose that ζ is one of the vectors orthogonalized at step 4, and consider any vector μ_j such that $1 \le j \le m_1 - 6\ell - r$. It follows by Corollary 1, above, that

$$A^t \mu_j = \sum_{h=1}^{j+6\ell-1} c_h \mu_h$$

for $c_1, c_2, \dots, c_{j+6\ell-1} \in \mathsf{F}_q$.

On the other hand, $\zeta = A\nu_k$ where $m_1 - r + 1 \le k \le r_1$, since ν_k is the second entry of one of the ordered pairs added to S_1 during the previous execution of the loop body.

It follows that

$$\begin{aligned} \mu_j^t A \zeta &= \mu_j^t A^2 \nu_k \\ &= (A^t \mu_j)^t A \nu_k \\ &= \sum_{h=1}^{j+6\ell-1} c_h \mu_h^t A \nu_k \\ &= 0, \end{aligned}$$

by the condition at line (2), since $h \le j + 6\ell - 1 \le m_1 - r - 1 < k$.

Similarly, if a value ζ is being orthogonalized as shown at line (9) during the elimination phase of the algorithm, then either $\zeta = A\nu_i$ for an integer i such that $m_i - r + 1 \le i \le m_1$, or $\zeta = A\tau$ for $\tau \in S_3$. The same argument as above can be used to establish that $\mu_j^t A\zeta = 0$ for $1 \le j \le m_1 - 6\ell - r$ in either case.

(b) Suppose next that $\sigma_{s_L,j}$ is one of the vectors that is being orthogonalized during an execution of step 7 in the algorithm shown in Figure 1; then $\sigma_{s_L,j} = A^t\sigma_{s_{L-1},j}$ before this step begins.

Suppose that $(s_L - 1, j) \in U$ at this point. Then an examination of the algorithm (including the details of step 10) confirms that $\sigma^t_{s_L-1,j}A\nu_h = 0$ for $1 \le h \le m_1$ at this point in the computation.

Suppose instead that $(s_L - 1, j) \notin U$; then $\sigma_{s_L-1,j} = \mu_h$, for some integer h between 1 and m_1, that is, $\sigma_{s_L-1,j}$ is the first entry of an ordered pair that has been included in S_1.

Consider now the *smallest* integer h such that $\mu_h = (s_L - 1, k)$ for any integer k. It follows by Lemma 20, above, that each one of the vectors

$$\mu_h, \mu_{h+1}, \dots, \mu_{m_1}$$

must be equal to one of the vectors

$$\sigma_{s_L-2,1}, \sigma_{s_L-2,2}, \dots, \sigma_{s_L-2,\ell}, \\ \sigma_{s_L-1,1}, \sigma_{s_L-1,2}, \dots, \sigma_{s_L-1,\ell}.$$

The orthogonality conditions given at line (2) clearly imply that $\mu_h, \mu_{h+1}, \dots, \mu_{m_1}$ are distinct, so it must be the case that $m_1 - h + 1 \le 2\ell$, that is, $h \ge m_1 - 2\ell + 1$.

It follows that if $(s_L - 1, j) \notin U$ then $\sigma_{s_L-1,j} = \mu_h$ for some integer h such that $m_1 - 2\ell + 1 \le h \le m_1$. In this case it is clear (once again, by the orthogonality conditions at line (2)) that

$$\sigma^t_{s_L-1,j}A\nu_k = 0$$

for every integer k such that $1 \le k \le m_1 - 2\ell$.

Now consider an integer k such that $k \le m_1 - 2\ell - 2r$. It follows by Lemma 21 that there exist elements $c_1, c_2, \dots, c_{k+2r-1}$ of F_q and a vector τ in the linear span of the vectors in S_2 such that

$$A\nu_k = \left(\sum_{h=1}^{k+2r-1} c_h\nu_h\right) + \tau.$$

Since each vector in S_2 is in the null space of A, $A\tau = 0$. It now follows that

$$\begin{aligned}\sigma^t_{s_L,j}A\nu_k &= (A^t\sigma_{s_L-1,j})^tA\nu_k \\ &= \sigma^t_{s_L-1,j}A(A\nu_k) \\ &= \left(\sum_{h=1}^{k+2r-1} c_h\sigma^t_{s_L-1,j}A\nu_h\right) + \sigma^t_{s_L-1,j}A\tau \\ &= 0.\end{aligned}$$

This follows, in particular, by the fact that $h < m_1 + 2\ell$ if $h \le k + 2r - 1$, so that $\sigma^t_{s_L-1,j}A\nu_h$ as noted above, and by the fact that τ is in the null space of A. □

Proof of Lemma 6. This is a consequence of Lemma 20 and the fact that if $i = 0$ then $U = \emptyset$ at the beginning of the execution of the body of the loop in which $\sigma_{i,j}$ is initialized, while

$$U \subseteq \{(i-1,1), (i-1,2), \dots, (i-1,\ell)\}$$

at the beginning of this execution of the loop body if $i > 0$.

Suppose first that the ordered pair (i, j) is eventually removed from the set U; the current version of the vector $\sigma_{i,j}$ is used as the first entry of an ordered pair (μ_h, ν_h) that is added to S_1 during the same execution of the body of the loop. If S_1 had size m_1 at the beginning of the execution of the loop body in which $\sigma_{i,j}$ was initialized then it is clear that $h > m_1$ and one can see by the above lemma that each of the values

$$\mu_{m_1+1}, \mu_{m_1+2}, \dots, \mu_h$$

must be equal to $\sigma_{u,v}$ where $1 \le v \le \ell$ and $v \in \{i-1, i, i+1\}$. Furthermore each of these 3ℓ vectors $\sigma_{u,v}$ can be used as the first entry of at most one of these ordered pairs. Consequently $h - m_1 \le 3\ell$ as claimed.

Suppose, on the other hand, that 3ℓ or more additional ordered pairs are eventually added to S_1 without using $\sigma_{i,j}$ as the first entry in an ordered pair; then it follows by the above counting argument that a vector $\sigma_{u,v}$ must be used as the first entry of an ordered pair added to S_1 for $u \ge i + 2$. However this is impossible, since (i, j) will not have been removed from the set U before this: The second test before step 5 will fail, and the Lanczos phase of the algorithm will terminate, before vectors $\sigma_{u,v}$ such that $u \ge i + 2$ are initialized at all. □

A.5 Proof of Lemma 7

The proof of Lemma 7 is rather long. It can be split into the following stages, each of which will be handled in one of the subsections that follows.

- To begin it will be shown that (because of the use of a different block size on the left and the right) it can be assumed, without loss of generality, that the matrix A has at most r invariant factors that are different from 1 or x. This assumption is defended in Subsection A.5.1, below.
- It can also be shown that, at each point during the execution of the body of the loop shown in Figure 1, either failure or a premature movement to the elimination phase can only take place if a matrix $V_L^tAV_R$ is rank-deficient, where the columns of V_L are an initial set of the vectors

  $$u_1, u_2, \dots, u_\ell, A^tu_1, A^tu_2, \dots, A^tu_\ell, \\ (A^t)^2u_1, (A^t)^2u_2, \dots, (A^t)^2u_\ell, \dots$$

  and where the columns of V_R are an initial set of the vectors

  $$v_1, v_2, \dots, v_r, Av_1, Av_2, \dots, Av_r, \\ A^2v_1, A^2v_2, \dots, A^2v_r, \dots$$

  This will be shown in Subsection A.5.2.
- Finally, the above can be applied with previously established bounds [10] to bound the probability of failure of the algorithm as well as the probability of premature termination of the Lanczos phase, in order to establish the lemma. This is carried out in Subsection A.5.3.

A.5.1 Assumption Concerning Invariant Factors

For the rest of this proof we will assume that A has at most r invariant factors that are different from 1 or x; the goal of this section is to defend that assumption.

Suppose now that the Krylov space of the given vectors $v_1, v_2, \dots, v_r$ has dimension d; let $x_1, x_2, \dots, x_d \in \mathsf{F}_q{}^{n\times 1}$ be a basis for this Krylov space. Recall that the Krylov space is (by definition) closed under multiplication by A, so that Ax_i is a linear combination of $x_1, x_2, \dots, x_d$ for each integer i such that $1 \le i \le d$.

Let $y_1, y_2, \dots, y_{n-d}$ be a completion of this as a basis for $\mathsf{F}_q{}^{n\times 1}$, that is, suppose that the vectors

$$x_1, x_2, \dots, x_d, y_1, y_2, \dots, y_{n-d} \in \mathsf{F}_q{}^{n\times 1}$$

are linearly independent. We may now define a matrix $\widehat{A} \in \mathsf{F}_q{}^{n\times n}$ as an operator on these vectors. Let

$$\widehat{A}x_i = Ax_i \qquad \text{for } 1 \le i \le d \tag{19}$$

and let

$$\widehat{A}y_j = 0 \qquad \text{for } 1 \le j \le n-d. \tag{20}$$

Linearity can be used to define $\widehat{A}z$ for any other vector $z \in \mathsf{F}_q{}^{n\times 1}$.

A consideration of a "rational Jordan form" of a matrix (block diagonal, with blocks that are companion matrices of powers of irreducible polynomials) establishes that if a matrix has more than r invariant factors that are different from 1 or x then its image is not contained in a Krylov space of r or fewer vectors. This can be used to argue the following.

Fact 1. *The above matrix $\widehat{A}$ has at most r invariant factors that are different from 1 or x.*

Lemma 22. *Let $\alpha \in \mathsf{F}_q{}^{n\times 1}$ be any element of the Krylov space generated by the vectors $v_1, v_2, \dots, v_r$ using the matrix A.*

(a) $f(A)\alpha = f(\widehat{A})\alpha$ *for any polynomial* $f \in \mathsf{F}_q[x]$.

(b) $(f(A^t)\beta)^t\alpha = (f(\widehat{A}^t)\beta)^t\alpha$ *for every polynomial* $f \in \mathsf{F}_q[x]$ *and for every vector* $\beta \in \mathsf{F}_q{}^{n\times 1}$.

(c) $(f(A^t)\beta)^tA\alpha = (f(\widehat{A}^t)\widehat{A}\alpha$ *for every polynomial* $f \in \mathsf{F}_q[x]$ *and every vector* $\beta \in \mathsf{F}_q{}^{n\times 1}$, *as well.*

Proof. Let α be any element of the Krylov space generated by the vectors $v_1, v_2, \dots, v_r$. Then α is a linear combination of the vectors $x_1, x_2, \dots, x_d$ (comprising a basis for this Krylov space) that are used to define $\widehat{A}$ at Equations (19) and (20), so it is clear that $A\alpha \in \widehat{A}\alpha$.

(a) Since the Krylov space generated by $v_1, v_2, \dots, v_r$ is closed under multiplication by A it is easily proved by induction on e that $A^i\alpha = \widehat{A}^i\alpha$ for every element α of this Krylov space and for every integer $i \ge 0$ as well.

Now let $f \in \mathsf{F}_q[x]$; then

$$f = c_ex^e + c_{e-1}x^{e-1} + \cdots + c_1x + c_0$$

for an integer $e \ge 0$ and for $c_e, c_{e-1}, \dots, c_1, c_0 \in \mathsf{F}_q$, so that

$$\begin{aligned} f(A)\alpha &= \sum_{h=0}^{e} c_hA^h\alpha \\ &= \sum_{h=0}^{e} c_h\widehat{A}^h\alpha \\ &= f(\widehat{A})\alpha, \end{aligned}$$

as claimed.

(b) Part (b) follows from part (a), since

$$(f(A^t)\beta)^t\alpha = \beta^t(f(A)\alpha) = \beta^t(f(\widehat{A})\alpha) = (f(\widehat{A}^t)\beta)^t\alpha$$

as claimed.

(c) Finally, the claim in part (c) is easily established using part (b) and the polynomial $g = x \cdot f$. □

Notice that it follows by part (a) of the above lemma that the Krylov space generated by $v_1, v_2, \dots, v_r$, using matrix A, is the same as the vector space generated by these vectors using matrix $\widehat{A}$ instead. As noted before that, $\widehat{A}$ has at most r invariant factors that are different from 1 or x. The above assumption can therefore be justified by establishing that the algorithm in Figure 1 behaves in essentially the same way using $\widehat{A}$ as the given coefficient matrix as it does when A is the given matrix instead. The remaining lemmas in this section establish that this is the case.

Now let $u_1, u_2, \dots, u_\ell \in \mathsf{F}_q{}^{n\times 1}$ and consider the following executions of the algorithm shown in Figure 1 on inputs $v_1, v_2, \dots, v_r$, when $u_1, u_2, \dots, u_\ell$ have been chosen as the used to initialize the sequence L_0 at step 1:

(i) The algorithm is executed using the matrix A.

(ii) The algorithm is executed using the matrix $\widehat{A}$.

Let s_L, s_R, S_1, S_2, μ_i, ν_i, L_i, $\sigma_{i,j}$ R and U denote the values that are maintained during execution (i) of the algorithm, as these are described in Section 2, and let $\widehat{s}_L$, $\widehat{s}_R$ $\widehat{S}_1$, $\widehat{S}_2$, $\widehat{\mu}_i$, $\widehat{\nu}_i$, $\widehat{L}_i$, $\widehat{\sigma}_{i,j}$, $\widehat{R}$ and $\widehat{L}$ denote the corresponding values that are maintained during execution (ii), instead.

Lemma 23. *Let $i \ge 1$. Then the body of the loop in the algorithm shown in Figure 1 is executed i or more times during execution* (i) *of the algorithm, as described above, if and only if it is executed i or more times during execution* (ii) *times, as well. The values maintained by the algorithms are related as follows at the beginning of the i^{th} execution of the loop body, whenever the loop body is executed i or more times.*

(a) $\widehat{s}_L = s_L$ *and* $\widehat{s}_R = s_R$.

(b) $|\widehat{S}_1| = |S_1|$, $\widehat{\nu}_i = \nu_i$ *for* $1 \le i \le |S_1|$, *and there exist polynomials* $f_{i,j} \in \mathsf{F}_q[x]$ *for* $1 \le i \le |S_1|$ *and* $1 \le j \le \ell$ *such that*

$$\mu_i = f_{i,1}(A^t)u_1 + f_{i,2}(A^t)u_2 + \cdots + f_{i,\ell}(A^t)u_\ell$$

and

$$\widehat{\mu}_i = f_{i,1}(\widehat{A}^t)u_1 + f_{i,2}(\widehat{A}^t)u_2 + \cdots + f_{i,\ell}(\widehat{A}^t)u_\ell$$

as well.

(c) $\widehat{S}_2 = S_2$.

(d) *For $1 \leq i \leq s_L$ and $1 \leq j \leq \ell$ there exist polynomials $f_{i,j,k}$ for $1 \leq k \leq \ell$ such that*

$$\sigma_{i,j} = f_{i,j,1}(A^t)u_1 + f_{i,j,2}(A^t)u_2 + \cdots + f_{i,j,\ell}(A^t)u_\ell$$

and

$$\widehat{\sigma}_{i,j} = f_{i,j,1}(\widehat{A}^t)u_1 + f_{i,j,2}(\widehat{A}^t)u_2 + \cdots + f_{i,j,\ell}(\widehat{A}^t)u_\ell$$

as well.

(e) $\widehat{R} = R$.

(f) *$\widehat{U} = U$ and, for every ordered pair (i, j) such that $1 \leq i \leq s_L$ and $1 \leq j \leq \ell$ such that $(i, j) \notin U$, and for every integer k such that $1 \leq k \leq |S_1|$, $\sigma_{i,j} = \mu_k$ if and only if $\widehat{\sigma}_{i,j} = \widehat{\mu}_k$.*

PROOF. This follows by a straightforward induction on i, using the results of Lemma 22 and inspection of the algorithm shown in Figure 1. □

LEMMA 24. *The following properties are satisfied by executions* (i) *and* (ii) *of the algorithm shown in Figure 1, where these are as described above.*

(a) *Each execution uses the same number of executions of the body of the loop.*

(b) *If i is an integer such that the body of the loop is executed at least i times (for either execution of the algorithm) then the first execution fails by executing step 9 during the i^{th} execution of the loop body if and only if the second does so, as well.*

(c) *If i is an integer such that the body of the loop is executed at least i times (for either execution of the algorithm) then the algorithm terminates at step 4 of the i^{th} execution of the body of the loop, during execution* (i)*, if and only if at does so during execution* (ii) *as well.*

(d) *If i is an integer such that the body of the loop is executed at least i times (for either execution of the algorithm) then the algorithm moves to the elimination phase immediately after the execution of step 10, in the i^{th} execution of the loop body as part of execution* (i) *of the algorithm, if and only if it does as part of execution* (ii) *as well.*

PROOF. This can be established by a continuation of the argument used to prove Lemma 23, above: Notice that, by the results of Lemmas 22 and 23, precisely the same conditions are checked in order to decide whether the algorithm should be terminated during the final execution of the loop body at (or immediately before) steps 4, 9, and 10, so the same decisions are made in each case. □

A continuation of this argument establishes the following as well.

LEMMA 25. *Suppose that executions* (i) *and* (ii) *of the algorithm, described above, proceed to an elimination phase. Let*

$$S_1 = \{(\mu_1, \nu_1), (\mu_2, \nu_2), \ldots, (\mu_{m_1}, \nu_{m_1})\}$$

be the set of ordered pairs of vectors computed (as "S_1") by execution (i) *of the algorithm, at beginning of the elimination phase, and let*

$$\widehat{S_1} = \{(\widehat{\mu}_1, \widehat{\nu}_1), (\widehat{\mu}_2, \widehat{\nu}_2), \ldots, (\widehat{\mu}_{\widehat{m_1}}, \widehat{\nu}_{\widehat{m_1}})\}$$

be the corresponding set of ordered pairs of vectors computed (as "S_1") by execution (ii) *of the algorithm, at the beginning of the elimination phase.*

Then $m_1 = \widehat{m_1}$ and, indeed, $\nu_i = \widehat{\nu}_i$ for $1 \leq i \leq m_1$.

A.5.2 Characterization of Failure or Premature Termination

For $j \geq 0$, let $V_{L,j} \in \mathsf{F}_q{}^{n \times (j+1)\ell}$ be the matrix with columns

$$u_1, u_2, \ldots, u_\ell, A^t u_1, A^t u_2, \ldots, A^t u_\ell, \ldots \\ (A^t)^j u_1, (A^t)^j u_2, \ldots, (A^t)^j u_\ell$$

Let $\widehat{r}$ be the rank of A and set $\widehat{V}_L \in \mathsf{F}_q{}^{n \times \widehat{r}}$ be the matrix whose columns are the initial $\widehat{r}$ columns of the above matrix $V_{L,\lceil r/\ell \rceil}$.

Let $V_{R,j} \in \mathsf{F}_q{}^{n \times (j+1)r}$ be the matrix with columns

$$v_1, v_2, \ldots, v_r, Av_1, Av_2, \ldots, Av_r, \ldots \\ A^j v_1, A^j v_2, \ldots, A^j v_r.$$

LEMMA 26. *Let $i \geq 1$ such that the body of the loop in the algorithm shown in Figure 1 is executed at least i times. Suppose, furthermore, that*

$$|U| < \ell - \lceil \log_q n \rceil - \delta$$

at the beginning of the i^{th} execution of the loop body, so that either steps 5–8 or step 9 are executed. Then $i \geq 2$, $\ell s_L < \widehat{r}$, and either

$$U \subseteq \{(s_L, 1), (s_L, 2), \ldots, (s_L, \ell)\}$$

(so that step 9 is not executed) or the matrix

$$V^t_{L,s_L-1} A V_{R,s_R-1} \in \mathsf{F}_q{}^{\ell s_L \times r s_R}$$

has rank less than ℓs_L.

PROOF. To begin notice that, by construction, $|U| = \ell$ at the beginning of the first execution of the loop, so that $i \geq 2$ if $|U| < \ell - \lceil \log_q n \rceil - \delta$. Furthermore it is clear, by inspection of the code that, whenever the tests following step 4 are executed,

$$|S_1| = (s_L + 1)\ell - |U|,$$

so that

$$\ell s_L = |S_1| - \ell + |U| < |S_1| - \lceil \log_q n \rceil - \delta.$$

Since the orthogonality relations at line (2) imply that the vectors $A\nu_1, A\nu_2, \ldots, A\nu_{m_1}$ are linearly independent, $|S_1| = m_1$ is less than or equal to the rank $\widehat{r}$ of A and it follows by the above that $\ell s_L < \widehat{r}$ as well.

Now consider the sets S_1 and U as they were defined at an earlier execution of the body of the loop, namely, the execution when the value of the variable "s_L" was first set to be equal to one less than its value at the beginning of the i^{th} execution of the body of the loop. It follows by part (a) of Lemma 3, and by Lemma 15, that the columns of the above matrix V_{L,s_L-1} span the same vector space as is spanned by the vectors

$$\{\sigma_{j,k} \mid 0 \leq j \leq s_L - 1 \text{ and } 1 \leq k \leq \ell\},$$

where the vectors in this set are as defined at the end of this earlier execution of the body of the loop (and the beginning of the one that follows it).

Consider once again the (later) execution of the loop body that is mentioned in the statement of the lemma and, in particular, the set S_1 as it is defined when the tests after step 4 are carried out. It follows by part (b) of Lemma 3, the fact that $Av = 0$ for all $v \in S_2$, and the fact that $S_3 = \emptyset$ at this point, that (for $m_1 = |S_1|$) the vectors

$$A\nu_1, A\nu_2, \ldots, A\nu_{m_1}$$

span the same vector space as the columns of the matrix AV_{R,s_R-1}.

Suppose now that the matrix $V_{L,s_L-1}^t AV_{R,s_R-1}$ has full rank, ℓs_L. Then it follows by the above observations and inspection of the code — notably including the greedy way in which values are chosen to be removed from U — that each of the ordered pairs (j,k) for $0 \le j \le s_L - 1$ and $1 \le k \le \ell$ must have been removed from U and used to define the first entry of an ordered pair (μ_h, ν_h) that has been added to S_1 before the beginning of the i^{th} execution of the body of the loop. Consequently

$$U \subseteq \{(s_L, 1), (s_L, 2), \ldots, (s_L, \ell)\}$$

at this point, as required to establish the above claim. □

Lemma 27. *Let $i \ge 1$, suppose that the loop in the algorithm shown in Figure 1 is executed at least i times, and that step 10 is reached during the i^{th} execution of the loop body. Consider the values s_L and s_R, and the sets S_1, S_2 and S_3 as they are defined on completion of this step of the algorithm.*

(a) Suppose that $(s_L+1)\ell$ is less than or equal to the rank $\widehat{r}$ of A at this point. Then either the matrix $V_{L,s_L}^t AV_{R,s_R}$ has rank less than that of AV_{R,s_R} at this point, or the body of the loop will be executed at least once more.

(b) Suppose, instead, that $(s_L + 1)\ell$ is greater than $\widehat{r}$. If the matrices $\widehat{V}_L AV_{R,s_R}$ and AV_{R,s_R} have the same rank then the body of the loop will be executed at least once more. Otherwise the elimination phase will begin after the completion of this step; the size of the set S_1 will be greater than or equal to the rank of $\widehat{V}_L A^t V_{R,s_R}$ when the elimination phase begins.

Proof. The argument needed to establish this lemma resembles the one used to prove the previous one.

(a) Suppose first that $(s_L + 1)\ell$ is less than or equal to the rank of A. Once again, it follows by Lemma 3 that the columns in the matrix V_{L,s_L} span the same vector space as the vectors

$$\{\mu_j \mid 1 \le j \le m_1\} \cup U, \tag{21}$$

while the columns of the matrix AV_R span the same vector space as the vectors

$$\{A\nu_j \mid 1 \le j \le m_1\} \cup \{A\tau \mid \tau \in R\}, \tag{22}$$

for S_1, m_1 and R as they are defined at the beginning of this execution of step 10. Consequently $V_{L,s_L}^t AV_{R,s_R}$ has the same rank as the matrix $W_L^t W_R$, where W_L has the columns shown at line (21) and W_R has the columns shown, and it suffices to consider $W_L^t W_R$ instead.

Note next that $W_L^t W_R$ is block diagonal: Its top left block is the identity matrix with order m_1, while its bottom right block is a matrix whose columns include those of the matrix $B_L^t AB_R$ that is considered as part of step 10 of the algorithm, and that has the same rank as this matrix.

The claim now follows immediately by a consideration of the details of step 10, since the Lanczos phase of the algorithm is only terminated at this point if $B_L^t AB_R$ is rank deficient.

(b) The claim for the case that $(s_L + 1)\ell > \widehat{r}$ follows by a repetition and continuation of this argument, using the fact that the columns of the matrix $\widehat{V}_L$ are a subset of those of the matrix V_{L,s_L} that is defined at this point, and by an inspection of the details of step 10. □

A.5.3 Completion of the Proof

The next lemma is easily proved and allows the results from [10] to be easily applied.

Lemma 28. *Let $C \in \mathsf{F}_q{}^{n\times d}$ be a matrix with full rank $d \le n$ and let $w_1, w_2, \ldots, w_{n-d}$ be a basis for the set of vectors*

$$N = \{w \in \mathsf{F}_q{}^{n\times 1} \mid w^t C = 0\}$$

Let $y_1, y_2, \ldots, y_k \in \mathsf{F}_q{}^{n\times 1}$. Then the vectors

$$y_1^t C, y_2^t C, \ldots, y_k^t C$$

are linearly independent in $\mathsf{F}_q{}^{d\times 1}$ if and only if the vectors

$$w_1, w_2, \ldots, w_{n-d}, y_1, y_2, \ldots, y_k$$

are linearly independent in $\mathsf{F}_q{}^{n\times 1}$.

Proof. Suppose first that the vectors $y_1^t C, y_2^t C, \ldots, y_k^t C$ are linearly dependent in $\mathsf{F}_q{}^{d\times 1}$, so that there exist elements $\alpha_1, \alpha_2, \ldots, \alpha_k$ of F_q, not all zero, such that

$$\alpha_1 y_1^t C + \alpha_2 y_2^t C + \cdots + \alpha_k y_k^t C = 0.$$

It follows that $\alpha_1 y_1 + \alpha_2 y_2 + \cdots + \alpha_k y_k \in N$, so that this vector is a linear combination of $w_1, w_2, \ldots, w_{n-d}$ and, since at least one of $\alpha_1, \alpha_2, \ldots, \alpha_k$ is nonzero it clearly follows that the set of vectors

$$y_1, y_2, \ldots, y_k, w_1, w_2, \ldots, w_{n-d}$$

is linearly dependent.

Suppose, conversely, that the set of vectors

$$y_1, y_2, \ldots, y_k, w_1, w_2, \ldots, w_{n-d}$$

is linearly dependent, so there exist elements $\alpha_1, \alpha_2, \ldots, \alpha_k$, $\beta_1, \beta_2, \ldots, \beta_{n-d}$ in F_q, not all zero, such that

$$\alpha_1 y_1 + \alpha_2 y_2 + \cdots + \alpha_k y_k + \beta_1 w_1 + \beta_2 w_2 + \cdots + \beta_{n-d} w_{n-d} = 0.$$

Furthermore, at least one of $\alpha_1, \alpha_2, \cdots + \alpha_k$ must be nonzero, since it is given that the set of vectors $w_1, w_2, \ldots, w_{n-d}$ is linearly independent. Now, since $w_i^t C = 0$ for $1 \le i \le n - d$, it follows that

$$\alpha_1 y_1^t C + \alpha_2 y_2^t C + \cdots + \alpha_k y_k^t C = 0$$

as well, establishing that that the set of vectors

$$y_1^t C, y_2^t C, \ldots, y_k^t C$$

is also linearly dependent, as required. □

The following is a result of a part of Lemma 3.11 of [10] using the notation in this paper.

Lemma 29 (Eberly and Hovinen [10]). *Suppose the matrix $A \in \mathsf{F}_q{}^{n\times n}$ has at most r invariant factors that are different from 1 or x and has rank $\widehat{r}$. Let $x_1, x_2, \ldots, x_h$ be a set of linearly independent vectors in $\mathsf{F}_q{}^{n\times 1}$.*

Suppose that $u_1, u_2, \dots u_\ell$ are chosen uniformly and independently from ${\mathsf{F}_q}^{n\times 1}$ and, for an integer k such that $1 \le k \le \widehat{r}$, let S be the set of vectors consisting of $x_1, x_2, \dots, x_h$ along with the first k of the vectors

$$\begin{aligned} &u_1, u_2, \dots, u_\ell, A^t u_1, A^t u_2, \dots, A^t u_\ell,\\ &\qquad\qquad (A^t)^2 u_1, (A^t)^2 u_2, \dots, (A^t)^2 u_\ell, \dots \end{aligned}$$

(a) If $h + k < n$ then the above set S of vectors is linearly dependent with probability at most

$$\frac{1 + 2q^{1+r-\ell}}{(q-1)q^{n-(h+k)}} < 2q^{k+k-n}.$$

(b) If $h + k > n$ and $1 \le j \le n$ then the above set S of vectors spans a subspace of ${\mathsf{F}_q}^{n\times 1}$, whose dimension is at most $n - j$, with probability at most

$$\begin{aligned} &\frac{1}{(q-1)q^{(k+k-n)+j-1}} + \frac{2q^{r-\ell+1}}{(q-1)q^{j-1}}\\ &\qquad\qquad \le \frac{1}{q^{(h+k-n)+j-1}} + \frac{2q^{1-2\delta}}{n^2}. \end{aligned}$$

A final lemma concerning the size of the set U will also be helpful

Lemma 30.

$$|U| \ge \ell - r - \lceil \log_q n \rceil - \delta \ge \lceil \log_q n \rceil + \delta$$

at the beginning of the each execution of the body of the loop of the algorithm shown in Figure 1 and at the beginning of the elimination phase, and

$$|U| \ge \ell - \lceil \log_q n \rceil - \delta$$

immediately before step 10 in each execution of the body of the loop in the algorithm in Figure 1.

Proof. This follows by a straightforward induction on the number i of executions of the loop body, of the algorithm shown in Figure 1, that have already taken place. The claim is satisfied when $i = 0$ since $|U| = \ell$ at the beginning of the first execution of the loop body and immediately before the first execution of step 10.

Suppose $i \ge 1$, that $|U| \ge \ell - r - \log_q n\rceil - \delta$ after the first i executions of the loop body, and that an $i + 1^{\text{st}}$ execution of the loop body also takes place.

If $|U| < \ell - \lceil \log_q n \rceil - \delta$ then steps 5–8 are executed as part of the $i + 1^{\text{st}}$ execution of the loop body, adding another ℓ elements to U, so that

$$\begin{aligned} |U| &\ge (\ell - r - \lceil \log_q n \rceil - \delta) + \ell\\ &\ge \ell - \rceil \log_q n \rceil - \delta \end{aligned}$$

immediately before the execution of step 10. On the other hand, if the size of U is initially greater than $\ell - \lceil \log_q n \rceil - \delta$ at the beginning of the $i + 1^{\text{st}}$ execution of the loop body, so that steps 5–9 are skipped, then it is also at least at large immediately before step 10 as well.

Now it suffices to note that at most r elements are removed from U when step 10 is carried out so that

$$|U| \ge \ell - r - \lceil \log_q n \rceil - \delta$$

at the beginning of the $i + 2^{\text{nd}}$ execution of the loop body once again, if there is one, or at the beginning of the elimination phase, otherwise.

The fact that $\ell - r - \lceil \log_q n \rceil - \delta \ge \lceil \log_q n \rceil + \delta$ follows by the choice of ℓ. □

Proof of Lemma 7. Consider now the execution of the algorithm shown in Figure 1 when it is given a matrix $A \in {\mathsf{F}_q}^{n\times n}$ and vectors $v_1, v_2, \dots, v_r$. As explained in Subsection A.5.1 this computation will be indistinguishable from a computation in which A is replaced by the matrix $\widehat{A}$ that is considered there. An examination of $\widehat{A}$ confirms that this matrix has at most r invariant factors that are different from 1 or x and, furthermore, that this matrix has rank $d - e$ where d is the dimension of the Krylov space generated by $v_1, v_2, \dots, v_r$ and e is the dimension of the intersection of this space and the null space of A. We will therefore assume, without loss of generality, that the given matrix A has these properties.

(a) Let i be an integer such that $i \ge 1$ and the body of the loop of the algorithm in Figure 1 is executed at least i times. It follows by Lemma 26 that if the algorithm fails during the i^{th} execution of the loop body, by executing step 9, then the matrix $V^t_{L,s_L-1} A V_{R,s_R-1}$ must have rank less than ℓs_L, where s_L and s_R are as defined at the point when the test before step 9 is carried out, and with $V_{L,s_L-1} \in {\mathsf{F}_q}^{n\times s_L}$ and $V_{R,s_R-1} \in {\mathsf{F}_q}^{n\times s_R}$ as at the beginning of Subsection A.5.2. Part (a) of Lemma 26 also establishes that ℓs_L is less than the rank of A at this point.

Now a consideration of part (b) of Lemma 3, the fact that S_2 is a subset of the null space of A, and the fact that $S_3 = \emptyset$ at this point, confirm that the matrix $A V_{R,s_R}$ has rank $|S_1|$. indeed, its columns span the same vector space as the vectors $A\nu_i$ for $1 \le i \le |S_1|$ and it follows by the orthogonality relations at line (2) that the latter set of vectors is linearly independent.

Let $w_1, w_2, \dots, w_{n-|S_1|}$ be a basis for the set

$$w \in {\mathsf{F}_q}^{n\times 1} \mid w^t A V_{R,s_R} = 0\}.$$

It follows by Lemma 28 that the set S of vectors that include $w_1, w_2, \dots, w_{n-|S_1|}$ and the columns of V_{L,s_L-1} must span a vector space with dimension less than $n - |S_1| + \ell s_L$, that is, this set of vectors must be linearly dependent.

One can also see by part (a) of Lemma 3 that $|S_1| + |U| = \ell(s_L + 1)$ at this point in the computation.

It follows that the above set S of vectors has size

$$\begin{aligned} n - |S_1| + \ell s_L &= n - (\ell(s_L + 1) - |U|) + \ell s_L\\ &= n - \ell + |U|\\ &< n - \lceil \log_q n \rceil - \delta, \end{aligned}$$

since $|U| < \ell - \lceil \log_q n \rceil - \delta$ in this point, and it follows by an application of part (a) of Lemma 29 with $k = \ell s_L$ and $h = n - |S_1|$ that, if $u_1, u_2, \dots, u_\ell$ are chosen uniformly and independently, then that algorithm fails at this point with probability at most $2q^{-\lceil \log_q n \rceil - \delta} \le (2q^{-\delta})/n$.

Clearly the body of the loop is executed at most n times, so the probability that the algorithm fails, by executing step 9, at any point at all, can now be bounded by $2q^{-\delta}$ as claimed.

(b) Now let i be a positive integer such that $0 \le i \le d - e$ where, as noted above, A is assumed to have rank $d - e$.

We wish to bound the probability that $|S_1| < d-e-i$ on termination of the Lanczos phase of the algorithm, when the algorithm does not report failure.

Notice that, since failure is not reported, such an execution of the Lanczos phase of the algorithm must end either at step 4 or at step 10. Termination at step 4 can be ruled out because this only happens when each vector $A\nu_i$ can be expressed as a linear combination of vectors in

$$\{\nu_j \mid 1 \leq j \leq |S_1|\} \cup S_2,$$

and it can be argued in this case that $|S_1| = d-e$.

Let $\widehat{S}_1$ be the set of ordered pairs included in the set "S_1" immediately before an execution of step 10, and let $\widehat{U}$ be the set of values included in "U" at this point as well. As explained in Subsection 2.1.6, and elaborated upon in the proof of Lemma 27, the Lanczos phase of the algorithm is terminated during this execution of step 10, on an attempt to add a set of s additional ordered pairs to this set, where AV_{R,s_R} has rank $|\widehat{S}_1| + s$ at this point. In particular, the Lanczos phase is terminated on the discovery that only t such ordered pairs can be added for $t < s$.

One of the following three cases must hold.

(i) The Lanczos phase ends during an execution of the body of the loop shown in Figure 1, in the manner described above, when $\ell(s_L+1) < d-e$.

(ii) The Lanczos phase ends during an execution of the body of the loop, as described above, when $\ell(s_L + 1) \geq d-e$ and when $|\widehat{S}_1| + s \leq d-e-i$, where s is the value mentioned above.

(iii) The Lanczos phase ends during an execution of the body of the loop, as described above, on the first execution of the body of the loop such that $\ell(s_L + 1) \geq d-e$ and when $|\widehat{S}_1| + s > d-e-i$, where s is the value mentioned above.

Note that $|S_1| > d+e-i$ if the Lanczos phase proceeds past any of the points mentioned above.

Consider an execution of the body of the loop in which $\ell(s_L+1) \leq d-e$ on completion of step 10, so that case (i) holds. It follows by part (a) of Lemma 27 that the matrix $V_{L,s_L}^t AV_{R,s_R}$ has rank less than that of AV_{R,s_R} after step 10 is executed. At this point, the matrix AV_{R,s_R} has rank $|\widehat{S}_1|+s$. Let $w_1, w_2, \dots, w_{n-|\widehat{S}_1|-s}$ be a basis for the space

$$N = \{w \in \mathsf{F}_q{}^{n\times 1} \mid w^t AV_{R,s_R} = 0\}; \qquad (23)$$

it now follows by Lemma 28 that the Lanczos phase can only end at this point if the set S of vectors that include the above vectors, along with the columns of V_{L,s_L+1}, span a vector space with dimension less than n.

Now notice that the size of this set of vectors is

$$\begin{aligned}
\ell(s_L+1) &+ n - |\widehat{S}_1| - s \\
&= \widehat{S}_1| + |\widehat{U}| + n - |\widehat{S}_1| - s \\
&= n + |\widehat{U}| - s \\
&\geq n + |\widehat{U}| - r && \text{(since } s \leq r) \\
&\geq n + \ell - \lceil \log_q n \rceil - \delta - r && \text{(by Lemma 30)} \\
&\geq n + \lceil \log_q n \rceil + \delta && \text{(by the choice of } \ell).
\end{aligned}$$

Part (b) of Lemma 29 can now be applied, with $h = n - |\widehat{S}_1| - s$, $k = \ell(s_L+1)$, and $j = 1$, to conclude that the probability that the elimination phase begins after this execution of the loop body is at most

$$\frac{q^{-\delta}}{n} + \frac{q^{1-2\delta}}{n^2} \leq \frac{2q^{-\delta}}{n}.$$

Next consider an execution of the body of the loop in which case (ii) is applicable. In follows by part (b) of Lemma 27 that the matrix $\widehat{V}_L^t AV_{R,s_R}$ has rank less than that of AV_{R,s_R}, for $\widehat{V}_L \in \mathsf{F}_q{}^{n\times(d-e)}$ as described at the beginning of Subsection A.5.2.

Once again let $w_1, w_2, \dots, w_{n-|\widehat{S}_1|-s}$ be a basis for the set N shown at line (23) above. Lemma 28 can be applied, once again, to conclude that the set S of vectors that includes these along with the columns of $\widehat{V}_L$ must span a vector space with dimension less than n.

Since case (ii) is applicable $|\widehat{S}_1| + s \leq d-e-i$. Suppose, in particular, that $|\widehat{S}_1| + s = d-e-i-\Delta$ for a nonnegative integer Δ. Then the above set of vectors S has size at most

$$\begin{aligned}
(d-e) &+ n - |\widehat{S}_1| + s \\
&= (d-e) + n - (d-e-i-\Delta) \\
&= n + i + \Delta.
\end{aligned}$$

Part (b) of Lemma 29 can now be applied, with $h = n - |\widehat{S}_1| - s$, $k = d-e$, and $j = 1$, conclude that Lanczos phase ends after this execution of step 10 with probability at most

$$\frac{1}{(q-1)q^{i+\Delta}} + \frac{2q^{1-2\delta}}{n^2} \leq q^{-i-\Delta} + \frac{2q^{-\delta}}{n}.$$

Finally, consider case (iii). Let

$$\Delta = (|\widehat{S}_1| + s) - (d-e-i)$$

so that Δ is a positive integer and, if S is the set of vectors described in the consideration of the previous case then S has size

$$(d-e) + (n - |\widehat{S}_1| - s) = n + i - \Delta.$$

It follows by part (b) of Lemma 27 that the rank of $\widehat{V}^t AV_{R,s_R}$ must be less than that of AV_{R,s_R} and, furthermore (since the size of $|S_1|$ is supposed to be less than $d-e-i$) that S must span a vector space with dimension less than

$$\begin{aligned}
(n - |\widehat{S}_1| - s) &+ (d-e-i) \\
&= (n - (d-e-i+\Delta)) + (d-e-i) \\
&= n - \Delta.
\end{aligned}$$

Part (b) of Lemma 29 can now be applied once again, with $h = n - |\widehat{S}_1| - s$, $k = d-e$, and $j = \Delta+1$, to conclude that the Lanczos phase ends at this point, with $|S_1| < d-e-i$, with probability at most

$$\frac{1}{q^{(i-\Delta)+\Delta+1-1}} + \frac{2q^{1-2\delta}}{n^2} \leq q^{-i} + \frac{2q^{-\delta}}{n}.$$

Having considered these cases it remains only to notice that step 10 is executed at most n times. Case (ii) cannot arise more than once for any choice of the nonnegative integer Δ that is mentioned in its analysis, as given above, while case (iii) is only possible once. The above bounds can now be added combined to conclude that $|S_1|$ has size less than $d-e-i$, on termination of the Lanczos phase of the computation, with probability at most

$$n \cdot \left(\frac{2q^{-\delta}}{n}\right) + \sum_{\Delta \geq 0} q^{-i-\Delta} + q^{-i} \leq 2q^{-\delta} + 3q^{-i},$$

as claimed. □

A.6 Proof of Theorem 1

Consider an execution of the algorithm described in Section 2, with vectors $u_1, u_2, \ldots, u_\ell$ chosen uniformly and independently from ${\mathsf{F}_q}^{n\times 1}$ for $\ell \geq r + 2(\lceil \log_q n \rceil + \delta)$.

The bound on failure included in Theorem 1 follows by part (a) of Lemma 7.

The applications of the matrix A^t to vectors is easily limited to the initialization of vectors in the sequences L_{s_L} at step 6 of the algorithm shown in Figure 1, for $s_L > 0$, and to the orthogonalizations of these vectors at step 7: Notice that the orthogonalization of $\sigma_{s_L,j}$ shown at line (11) can be carried out by applying A^t to $\sigma_{s_L,j}$ and using the step

$$\sigma_{s_L,j} := \sigma_{s_L,j} - \sum_{k=\min(1,m_1-2\ell-2r+1)} ((A^t\sigma_{s_L,j})^t \nu_k)\mu_k.$$

As a result A^t is applied (at most) twice for each vector that is eventually included as the first entry of an ordered pair included in the set S_1 or that remains in the set U. The number of applications of A^t used here is at most

$$2(|S_1| + |U|) < 2n + 4\ell,$$

where the sets S_1 and U are as defined at the end of the Lanczos phase of the algorithm.

The remaining initializations, updates, and orthogonalizations of vectors can be carried out by applying A at most three times for each vector in one of the sets R that is to be processed, during the Lanczos phase, provided that step 10 is implemented sensibly: A first application is needed to initialize each vector in the set R at step 3; a second application to every such vector is sufficient for the orthogonalization step (as shown by line (10) at step 4; and A is applied to each vector $\kappa \in R$ once again at the beginning of step 10. An examination of the details of step 10 confirms that each each vector $A\kappa$ can — and should be — maintained for each $\kappa \in R$ throughout this step, and that this can be used to update the elements of the set U that are described there without further applications of either A or A^t.

Since each of the vectors R is included either as the second entry of an ordered pair in S_1, or added to S_2 or S_3, one can see that the number of applications of A to vectors during the Lanczos phase is at most $3(|S_1| + |S_2| + |S_3|)$ where S_1, S_2 and S_3 are as defined at the end of the Lanczos phase of the computation. An examination of the details of the elimination phase confirms that the number of applications of A to vectors during the entire computation is at most $3(|S_1| + |S_2| + |S_3|)$, for S_1, S_2 and S_3 as they are defined on termination, as well.

It is easy to see by inspection of the code that the size of S_2 never exceeds r — a vector is only added to S_2 when the size of the set R is decreased. On the other hand, $|S_1|+|S_3|$ is less than or equal to the dimension of the Krylov space generated by $v_1, v_2, \ldots, v_r$ and this is clearly at most n. Thus $|S_1| + |S_2| + |S_3| \leq n + r$, so at most $3n + 3r$ applications of A to vectors are required.

A similar analysis can be used to establish that $O(n\ell(n + \ell)) = O(n^2\ell)$ additional arithmetic operations over F_q are used during the Lanczos phase, because $O(n\ell)$ operations are needed to initialize and process each vector before either this vector is used in an ordered pair in S_1 or added to one of the sets S_2 or S_3, or the Lanczos phase ends.

In order to bound the expected number of additional operations over F_q used by the elimination phase, notice that if m_3 is the size of S_3 on termination then $O(n\ell m_3)$ operations are used for orthogonalization steps and $O(nm_3^2)$ steps are used for the applications of Gaussian elimination that are required. Consequently the number of additional operations needed is in $O(nm_3(\ell + m_3))$. Now if $m_3 \leq \lceil \log_q n \rceil$ then this bound is clearly in $O(n^2\ell)$ as needed. It is never more than cubic in n since m_3 cannot be greater than n. One can now see, by an application of part (b) of of Lemma 7, that if $\delta \geq \lceil \log_q n \rceil$ then the probability that $|S_3|$ exceeds $\lceil \log_q \rceil$ is at most $5/n$. It follows from the above that the expected number of additional operations used by the elimination phase is at most linear in

$$1 \cdot \ell n^2 + \tfrac{5}{n} \cdot n^3 \in O(n^2\ell),$$

as claimed.

The remark that follows the statement of the theorem follows by a straightforward extension of the above analysis, applying the probability bounds given in Lemma 7.

B. PROOFS OF CLAIMS IN SECTION 3

B.1 Proofs of Lemmas 8 and 9

Proof of Lemma 8. Consider the representation of A as shown at lines (12) and (13). Since the matrixA_1 is nonsingular its minimal polynomial f is not divisible by x. On the other hand, A_2 is nilpotent so that its minimal polynomial g is a power of x: $g = x^h$ for $h \geq 2$. Since A_3 is the zero matrix it is clear that $g(A_3) = 0$ as well.

Consequently the above polynomials f and g are relatively prime, so there exist polynomials $h_1, h_2 \in \mathsf{F}_q[x]$ such that

$$h_1 \equiv \begin{cases} 1 \bmod f \\ 0 \bmod g \end{cases} \qquad \text{and} \qquad h_2 \equiv \begin{cases} 0 \bmod f \\ 1 \bmod g. \end{cases}$$

Now consider vectors $v_i = \alpha_i + \beta_i + \gamma_i$ for $1 \leq i \leq r$, where $\alpha_i \in \mathcal{V}_1$, $\beta_i \in \mathcal{V}_2$, and $\gamma_i \in \mathcal{V}_2$. Let K be the Krylov space generated by $v_1, v_2, \ldots, v_r$ and let K_1 be the vector space generated by $\alpha_1, \alpha_2, \ldots, \alpha_r$, as in the statement of the lemma. It will be useful to consider a "Krylov space" that is not mentioned in the statement of the lemma, as well: let $\widehat{K}_2$ be the Krylov space generated by the vectors $\beta_1 + \gamma_1, \beta_2 + \gamma_2, \ldots, \beta_r + \gamma_r$.

Notice that, by the descriptions of $\mathcal{V}_1$, $\mathcal{V}_2$ and $\mathcal{V}_3$ preceding the claim, $h_1(v_i) = \alpha_i$ for $1 \leq i \leq r$, so that $\alpha_1, \alpha_2, \ldots, \alpha_r \in K$. Since K is closed under multiplication by A if follows that $K_1 \subseteq K$.

Similarly, $h_2(v_i) = \beta_i + \gamma_i$ for $1 \leq i \leq r$, and $\widehat{K}_2 \subseteq K$ as well.

On the other hand, $v_i = \alpha_i + (\beta_i + \gamma_i)$ so that v_i can be written as a sum of elements of K_1 and $\widehat{K}_2$. Indeed, it is clear that every element of the Krylov space K can be expressed in this way.

Finally, notice that $K_1 \cap \widehat{K}_2 = \{0\}$. Thus K is the direct sum of K_1 and $\widehat{K}_2$.

(a) Notice that, since $A\gamma = 0$ for all $\gamma \in \mathcal{V}_3$, $A(\beta_i + \gamma_i) = A\beta_i$ for $1 \leq i \leq r$. Consequently $K_2 \subseteq \widehat{K}_2$ and the claim follows from the containments that have been established above.

(b) The second claim follows from the fact that K is the direct sum of K_1, and $\widehat{K}_2$, and that the set of vectors
$$\{\lambda_1, \lambda_2, \ldots, \lambda_j\} \cup \{\beta_1 + \gamma_1, \beta_2 + \gamma_2, \ldots, \beta_r + \gamma_r\}$$
mentioned in the statement of the lemma is a basis for $\widehat{K}_2$. □

Notice that the following — which will be used again — has also been proved.

Lemma 31. *If $v_i = \alpha_1 + \beta_1 + \gamma_1$ where $\alpha_i \in \mathcal{V}_1$, $\beta_i \in \mathcal{V}_2$ and $\gamma_i \in \mathcal{V}_3$ for $1 \leq i \leq r$, then the Krylov space $\widehat{K}_2$ that is generated by the vectors $\beta_1+\gamma_1, \beta_2+\gamma_2, \ldots, \beta_r+\gamma_r$ is a subset of the Krylov space K that is generated by $v_1, v_2, \ldots, v_r$.*

A consideration of the decomposition of the matrix A_2 shown at line (13) confirms that there exist values
$$\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_m \in \mathcal{V}_2$$
such that
$$A^{n_{2,i}-1}\widehat{\omega}_i \neq 0 = A^{n_{2,i}}\widehat{\omega}_i \qquad \text{for } 1 \leq i \leq m$$
and the Krylov space generated by $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_m$ is equal to $\mathcal{V}_2$. Indeed, $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_m$ can be chosen as generators of the Jordan blocks in the decomposition shown at line (13). It is clear that the set of vectors $A^j\widehat{\omega}_i$ such that $0 \leq j \leq n_{2,i}-1$ and $1 \leq i \leq m$ forms a basis for $\mathcal{V}_2$. The next lemma follows directly from this.

Lemma 32. *Every element $\beta \in \mathcal{V}_2$ can be written uniquely as*
$$\varphi_1(A)\widehat{\omega}_1 + \varphi_2(A)\widehat{\omega}_2 + \cdots + \varphi_m(A)\widehat{\omega}_m$$
for polynomials $\varphi_1, \varphi_2, \ldots, \varphi_m \in \mathsf{F}_q[x]$ such that the degree of φ_i is less than $n_{2,i}$ for $1 \leq i \leq m$.

Proof of Lemma 9. A slightly stronger result will be established: It will be shown that the probability that
$$\omega_1 + \lambda_1, \omega_2 + \lambda_2 + \ldots, \omega_m + \lambda_m$$
are not all found in the Krylov space, for some choice of $\lambda_1, \lambda_2, \ldots, \lambda_m \in \mathcal{V}_3$, is at most
$$q^{-\Delta-1} + q^{-\Delta-2} + \cdots + q^{-\Delta-m}.$$

Recall that each vector v_i can be written (uniquely) as $\alpha_i+\beta_i+\gamma_i$ where $\alpha_i \in \mathcal{V}_1$, $\beta_i \in \mathcal{V}_2$, and $\gamma_i \in \mathcal{V}_3$ for $1 \leq i \leq r$.

With that in mind notice that, by Lemma 31, it is sufficient to show that $\omega_1+\lambda_1, \omega_2+\lambda_2, \ldots, \omega_m+\lambda_m$ are each contained in $\widehat{K}_2$ for some choice of $\lambda_1, \lambda_2, \ldots, \lambda_m \in \mathcal{V}_3$, where $\widehat{K}_2$ is as described in the above lemma.

Indeed, since $\lambda_1, \lambda_2, \ldots, \lambda_m$ can be chosen freely from $\mathcal{V}_3$, it is sufficient to consider the probability that the vectors $\omega_1, \omega_2, \ldots, \omega_m$ are each included in the Krylov space generated by $\beta_1, \beta_2, \ldots, \beta_r$ — for if ω_i is in the Krylov space generated by $\beta_1, \beta_2, \ldots, \beta_r$, for $1 \leq i \leq m$, then
$$\omega_i = \psi_{i,1}(A)\beta_1 + \psi_{i,2}(A)\beta_2 + \cdots + \psi_{i,r}(A)\beta_r$$
for some choice of polynomials $\psi_{i,1}, \psi_{i,2}, \ldots, \psi_{i,r} \in \mathsf{F}_q[x]$, implying that $\omega_i + \lambda_i$ is in the Krylov space generated by $\beta_1 + \gamma_1, \beta_2 + \gamma_2, \ldots, \beta_r + \gamma_r$, for the values
$$\lambda_i = \psi_{i,1}(A)\gamma_1 + \psi_{i,2}(A)\gamma_2 + \cdots + \psi_{i,r}(A)\gamma_r \in \mathcal{V}_3,$$
for $1 \leq i \leq r$, as well.

Furthermore we may assume without loss of generality that $\omega_i = \widehat{\omega}_i$ for $1 \leq i \leq m$, for the values $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_m$ that are discussed above: If $\omega_1, \omega_2, \ldots, \omega_m$ is any set of vectors in $\mathcal{V}_2$ generating $\mathcal{V}_2$ as a Krylov space, then the Krylov space generated by $\beta_1, \beta_2, \ldots, \beta_m$ includes $\omega_1, \omega_2, \ldots, \omega_m$ if and only if this space includes $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_m$.

Notice now that if $m = 0$ then there is nothing to prove.

If $m = 1$ then it follows by Lemma 32 that (since the vectors $\beta_1, \beta_2, \ldots, \beta_r$ are chosen uniformly and independently from $\mathcal{V}_2$)
$$\beta_1 = \varphi_1(A)\widehat{\omega}_1, \beta_2 = \varphi_2(A)\widehat{\omega}_1, \ldots, \beta_r = \varphi_r(A)\widehat{\omega}_1$$
where $\varphi_1, \varphi_2, \ldots, \varphi_r$ are chosen uniformly and independently from the set of polynomials with degree less than $n_{2,1}$ in $\mathsf{F}_q[x]$. Now
$$\varphi_1(0) = \varphi_2(0) = \cdots = \varphi_r(0) = 0$$
(and, consequently, $\widehat{\omega}_1$ is not in the Krylov space generated by $\beta_1, \beta_2, \ldots, \beta_r$) with probability $q^{-r} \leq q^{-\Delta-1}$. On the other hand, if $\varphi_i(0) \neq 0$ for some integer i such that $1 \leq i \leq r$ then the polynomials φ_i and $x^{n_{2,1}}$ are relatively prime and there exists a polynomial $\psi \in \mathsf{F}_q[x]$ such that
$$\psi \times \varphi_i \equiv 1 \bmod x^{n_{2,1}}.$$
In this case $\psi(A)\beta_i = \psi(A)\varphi_i(A)\widehat{\omega}_1 = \widehat{\omega}_1$ so that $\widehat{\omega}_1$ is in the Krylov space generated by $\beta_1, \beta_2, \ldots, \beta_r$ as required.

Suppose next that $m \geq 2$; one can see, by Lemma 32, that
$$\beta_i = \beta_i' + \varphi_{m,i}(A)\widehat{\omega}_m, \tag{24}$$
for $1 \leq i \leq r$, where the values $\beta_1', \beta_2', \ldots, \beta_r'$ are chosen uniformly and independently from the Krylov space generated by $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_{m-1}$, and where $\varphi_{m,1}, \varphi_{m,2}, \ldots, \varphi_{m,r}$ are chosen uniformly and independently from the set of polynomials with degree less than $n_{2,m}$ in $\mathsf{F}_q[x]$. Now
$$\varphi_{m,1}(0) = \varphi_{m,2}(0) = \cdots = \varphi_{m,r}(0) = 0$$
(and $\widehat{\omega}_m$ is not in the Krylov space generated by the vectors $\beta_1, \beta_2, \ldots, \beta_r$), with probability $q^{-r} \leq q^{-m-\Delta}$.

Suppose, instead, that $\varphi_{m,i}(0) \neq 0$ for some integer i. Reordering $v_1, v_2, \ldots, v_r$ (and $\beta_1, \beta_2, \ldots, \beta_r$) we may assume without loss of generality that $i = r$ and that $\varphi_{m,r}(0) \neq 0$.

It follows again that $\varphi_{m,r}$ and $x^{2,m}$ are relatively prime, so that there exists a polynomial $\psi_m \in \mathsf{F}_q[x]$ with degree less than $n_{2,m}$ such that $\psi_m \times \varphi_{m,r} \equiv 1 \bmod x^{n_{2,m}}$, and the Krylov space generated by $\beta_1, \beta_2, \ldots, \beta_r$ certainly includes
$$\psi_m(A)\beta_r = \psi_m(A)(\varphi_{m,r}(A)\widehat{\omega}_r + \beta_r') = \widehat{\omega}_m + \beta_r'', \tag{25}$$
where $\beta_r'' = \psi_m(A)\beta_r'$ is an element of the Krylov space generated by $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_{m-1}$ depending only on β_r, so that it is clearly chosen independently of $\beta_1, \beta_2, \ldots, \beta_{r-1}$.

Notice next that, by equations (24) and (25), the Krylov space generated by $\beta_1, \beta_2, \ldots, \beta_r$ also includes the values

$$\begin{aligned}\widehat{\beta}_i &= \beta_i - \varphi_{m,i}(A)\psi_m(A)\beta_r \\ &= (\beta_i' + \varphi_{m,i}(A)\widehat{\omega}_m) - \varphi_{m,i}(A)(\widehat{\omega}_m + \beta_r'') \\ &= \beta_i' - \varphi_{m,i}(A)\beta_r''\end{aligned}$$

for $1 \leq i \leq r-1$. $\varphi_{m,i}(A)\beta_r''$ is an element of the Krylov space generated by $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_{m-1}$, since β_r'' is.

Since $\beta_1', \beta_2', \ldots, \beta_{r-1}'$ are chosen uniformly and independently from the Krylov space generated by $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_{m-1}$ and, furthermore, since these are chosen independently of $\varphi_{m,i}(A)\beta_r''$, it follows that the above values $\widehat{\beta}_1, \widehat{\beta}_2, \ldots, \widehat{\beta}_{r-1}$ are chosen uniformly and independently from the Krylov space generated by $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_{m-1}$ as well.

Proceeding inductively on m, we may now conclude that the probability that $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_{m-1}$ are *not* in the Krylov space generated by $\widehat{\beta}_1, \widehat{\beta}_2, \ldots, \widehat{\beta}_{r-1}$ is at most

$$q^{-\Delta-1} + q^{-\Delta-2} + \cdots + q^{-\Delta-m+1}.$$

Now, if $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_{m-1}$ *are* all members of this Krylov space then they are certainly in the Krylov space generated by $\beta_1, \beta_2, \ldots, \beta_r$ as well. It has already been noted that this Krylov space includes $\widehat{\omega}_m + \beta_r''$, for a member β_r'' of the Krylov space generated by $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_{m-1}$, so it is clear that $\widehat{\omega}_m$ is also contained in the Krylov space generated by $\beta_1, \beta_2, \ldots, \beta_r$.

Adding together the bounds on probabilities that are mentioned above, we see that the probability that $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_m$ are not all in the Krylov space generated by $\beta_1, \beta_2, \ldots, \beta_r$ is at most

$$q^{-\Delta-1} + q^{-\Delta-2} + \cdots + q^{-\Delta-m+1} + q^{-\Delta-m}$$

as claimed.

Clearly $q^{-\Delta-1} + q^{-\Delta-2} + \cdots + q^{-\Delta-m} \leq 2q^{-\Delta}$, so the probability bound included in the statement of the lemma is correct.

Finally notice that, if $\widehat{\omega}_i$ is in the Krylov space generated by $\beta_1, \beta_2, \ldots, \beta_r$ for $1 \leq i \leq m$, then

$$\widehat{\omega}_i = \xi_1(A)\beta_1 + \xi_2(A)\beta_2 + \cdots + \xi_r(A)\beta_r$$

for polynomials $\xi_1, \xi_2, \ldots, \xi_r \in \mathsf{F}_q[x]$. It clearly follows (since $\xi_i(A)A = A\xi_i(A)$ for $1 \leq i \leq r$) that

$$\begin{aligned}&A\xi_1(A)(\beta_1 + \gamma_1) + A\xi_2(A)(\beta_2 + \gamma_2) + \\ &\qquad\qquad \cdots + A\xi_r(A)(\beta_r + \gamma_r) \\ &= A\xi_1(A)\beta_1 + A\xi_2(A)\beta_2 + \cdots + A\xi_r(A)\beta_r \\ &= A\widehat{\omega}_i.\end{aligned}$$

Thus the Krylov space includes $A\widehat{\omega}_1, A\widehat{\omega}_2, \ldots, A\widehat{\omega}_m$ and, indeed, $A\beta$ for all $\beta \in \mathcal{V}_2$, as claimed. □

B.2 Proof of Lemma 10

Let $b = \alpha + \beta$ for $\alpha \in \mathcal{V}_1$ and $\beta \in \mathcal{V}_2$ as noted prior to the claim.

The claim is trivial if $m = 0$, since $\mathcal{V}_2 = \{0\}$ in this case: $b = \alpha$ and A_1 is nonsingular. Consequently if $f_1 \in \mathsf{F}_q[x]$ is the minimal polynomial of A_1 then $f_1(0) \neq 0$, $\gcd(f_1, x) = 1$, and there exists a polynomial $g \in \mathsf{F}_q[x]$ such that $xg \equiv 1 \bmod f_1$. In this case, since $v_1 = Ab$, $g(A)v_1 = g(A)Ab = b$, establishing that b is in the desired Krylov space.

Suppose instead that $m > 0$, so that A_2 is nonzero and nilpotent. Let f_1 be the minimal polynomial of A_1 as above, and let f_2 be the minimal polynomial of A_2, so that $f_2 = x^h$ for an integer $h \geq 2$. Again, $\gcd(f_1, f_2) = 1$ in $\mathsf{F}_q[x]$ and there exists a polynomial $g \in \mathsf{F}_q[x]$ such that $gf_2 = x^h g \equiv 1 \bmod f_1$. In this case it suffices to note that

$$g(A)A^{h-1}v_1 = g(A)A^{h-1}(A\alpha + A\beta) = \alpha,$$

establishing that α is guaranteed to be in the Krylov space generated by $v_1, v_2, \ldots, v_r$.

It remains only to note that if the system $Ax = b$ is consistent, for $b = \alpha + \beta$ and $\beta \in \mathcal{V}_2$, then β must belong to the Krylov space generated by $A\omega_1, A\omega_2, \ldots, A\omega_m$, for the set of vectors $\omega_1, \omega_2, \ldots, \omega_m$, generating $\mathcal{V}_2$ as a Krylov space, that are mentioned before the claim. Consequently $\mathcal{V}_2$ includes a vector χ such that $A\chi = \beta$. Furthermore, if the Krylov space generated by $v_1, v_2, \ldots, v_r$ includes vectors

$$\omega_1 + \lambda_1, \omega_2 + \lambda_2, \ldots, \omega_m + \lambda_m$$

for any choice of values $\lambda_1, \lambda_2, \ldots, \lambda_m \in \mathcal{V}_3$, then the Krylov space also includes the value $\chi' = \chi + \lambda$ for some $\lambda \in \mathcal{V}_3$. Consequently $A\chi' = A\chi + A\lambda = \beta + 0 = \beta$, and the Krylov space generated by $v_1, v_2, \ldots, v_r$ includes a vector x such that $Ax = b$.

The result now follows by Lemma 9, which bounds the probability that the vector space generated by $v_1, v_2, \ldots, v_r$ does not include vectors of the above form.

B.3 Proofs of Lemmas 11 and 12

Proof of Lemma 11. Suppose that the Jordan normal form of A includes exactly m nilpotent blocks with order at least two; then A is as shown at line (12) where A_1 is nonsingular, A_3 is a zero matrix. and A_2 is as shown at line (13).

Let the vectors spaces $\mathcal{V}_1$, $\mathcal{V}_2$, and $\mathcal{V}_3$ be as described at the beginning of Section 3, so that A acts as an invertible operator (with coefficient matrix A_1) on $\mathcal{V}_1$, as a nilpotent operator (with coefficient matrix A_2) on $\mathcal{V}_2$, and where $A\mathcal{V}_3 = \{0\}$. Now ${\mathsf{F}_q}^{n \times 1} = \mathcal{V}_1 \oplus \mathcal{V}_2 \oplus \mathcal{V}_3$ and it suffices to note that the image of A is a subspace of $\mathcal{V}_1 \oplus \mathcal{V}_2$, while the null space of A is a subspace of $\mathcal{V}_2 \oplus \mathcal{V}_3$. Consequently the intersection of the image and the null space of A is a subspace of $\mathcal{V}_2$.

Consider now the vectors $\widehat{\omega}_1, \widehat{\omega}_2, \ldots, \widehat{\omega}_m$ that generate $\mathcal{V}_2$ as a Krylov space that are described following Lemma 31. It is clear that, since $n_{2,i} \geq 2$ for $1 \leq i \leq m$, the set

$$A^{n_{2,1}-1}\widehat{\omega}_1, A^{n_{2,2}-1}\widehat{\omega}_2, \ldots, A^{n_{2,m}-1}\widehat{\omega}_m \in \mathcal{V}_2$$

is a linearly independent set of vectors in the image of A and, indeed, these form a basis for the intersection of $\mathcal{V}_2$ and the null space of A. Consequently these form a basis for the intersection of the image of A and the null space of A as well, as needed to prove the claim. □

Proof of Lemma 12. Suppose that $\tau_1, \tau_2, \ldots, \tau_k$ are as described prior to the statement of the lemma, so that the set of vectors $A\tau_1, A\tau_2, \ldots, A\tau_k$ is linearly independent and so that $A^2\tau_1 = A^2\tau_2 = \cdots = A^2\tau_k = 0$.

(a) Since the set of vectors $A\tau_1, A\tau_2, \ldots, A\tau_k$ is linearly independent, and these vectors belong to the intersection of the image of A and its null space, this intersection clearly has dimension at least k. The claimed inequality $k \leq \widehat{m}$ now follows by Lemma 11.

Notice now that $k < \widehat{m}$ only if the Krylov space generated by $z_1, z_2, \dots, z_r$ does not include a set of vectors

$$\omega_1 + \lambda_1, \omega_2 + \lambda_2, \dots, \omega_{\widehat{m}} + \lambda_{\widehat{m}}$$

where $\lambda_1, \lambda_2, \dots, \lambda_{\widehat{m}} \in \mathcal{V}_3$ and where $\omega_1, \omega_2, \dots, \omega_{\widehat{m}}$ are vectors that generate $\mathcal{V}_2$ as a Krylov space. The probability bound that is included in part (a) is therefore a consequence of Lemma 9.

(b) The probability bound included in part (b) can be established by an application of the technique that was used to prove Lemma 9: Notice that if A has $\widehat{m} \geq r = m + \Delta$ nilpotent blocks with order at least two in its Jordan normal form, then there exist vectors $\widehat{\omega}_1, \widehat{\omega}_2, \dots, \widehat{\omega}_{\widehat{m}} \in \mathcal{V}_2$, and integers $n_{2,1}, n_{2,2}, \dots, n_{2,\widehat{m}}$ that are each greater than two, that are as described (with $\widehat{m}$ replacing m) in the text that follows Lemma 31. Consequently each element of $\mathcal{V}_2$ can be uniquely expressed as described in Lemma 32, with $\widehat{m}$ replacing m once again.

Now let $z_i = \alpha_i + \beta_i + \gamma_i$ where $\alpha_i \in \mathcal{V}_1$, $\beta_i \in \mathcal{V}_2$, and $\gamma_i \in \mathcal{V}_3$ for $1 \leq i \leq r$, and notice that, since $v_1, v_2, \dots, v_r$ are chosen uniformly and independently from $\mathsf{F}_q{}^{n \times 1}$,

$$\beta_i = \varphi_{i,1}(A)\widehat{\omega}_1 + \varphi_{i,2}(A)\widehat{\omega}_2 + \cdots + \varphi_{i,\widehat{m}}(A)\widehat{\omega}_{\widehat{m}},$$

where each polynomial $\varphi_{i,j}$ is chosen uniformly from the set of polynomials with degree less than $n_{2,j}$, for $1 \leq i \leq r$ and $1 \leq j \leq \widehat{m}$ and, furthermore, each polynomial $\varphi_{i,j}$ is chosen independently from the set of polynomials $\varphi_{s,t}$ such that $1 \leq s \leq r$, $1 \leq t \leq \widehat{m}$, and $i \neq s$ or $j \neq t$.

The probability that $\varphi_{r,j}(0) = 0$ for every integer j such that $1 \leq j \leq \widehat{m}$ is $q^{-\widehat{m}}$. Suppose, now, that this is not the case, and pick an integer j such that $1 \leq j \leq \widehat{m}$ and $\varphi_{r,j}(0) \neq 0$. Then $\gcd(\varphi_{r,j}, x^{n_{2,j}}) = 1$, and one can argue as in the proof of Lemma 9 that $\widehat{\omega}_j$ is in the Krylov space generated by β_r and, furthermore, that the Krylov space generated by $\beta_1, \beta_2, \dots, \beta_r$ is the same as the Krylov space generated by $\widehat{\beta}_1, \widehat{\beta}_2, \dots, \widehat{\beta}_{r-1}$ and β_r, where $\widehat{\beta}_1, \widehat{\beta}_2, \dots, \widehat{\beta}_{r-1}$ are chosen uniformly and independently from the Krylov space generated by the vectors $\widehat{\omega}_h$ such that $1 \leq h \leq \widehat{m}$ and $h \neq j$.

Proceeding inductively, as in the proof of Lemma 9, one can establish that the intersection of the Krylov space generated by $A\beta_{r-m+1}, A\beta_{r-m+2}, \dots, A\beta_r$ and the null space of A has dimension less than r with probability at most

$$q^{m-1-\widehat{m}} + q^{m-\widehat{m}} + \cdots + q^{-\widehat{m}} \leq q^{m-\widehat{m}}.$$

This implies that intersection of the Krylov space generated by $v_1 = Az_1, v_2 = Az_2, \dots, v_r = Az_r$ and the null space of A has dimension less than r with at most this probability as well. It now suffices to notice that $q^{m-\widehat{m}} \leq q^{m-r} = q^{-\Delta} \leq 2q^{-\Delta}$. □

B.4 Proof of Lemma 13

This can be established by a modification of the argument used to prove Lemma 9.

Suppose, once again, that $v_i = \alpha_i + \beta_i + \gamma_i$ where $\alpha_i \in \mathcal{V}_1$, $\beta_i \in \mathcal{V}_2$, and $\gamma_i \in \mathcal{V}_3$ for $1 \leq i \leq m + d$. As indicated in part (b) of Lemma 8, $\beta_1 + \gamma_1, \beta_2 + \gamma_2, \dots, \beta_{m+d} + \gamma_{m+d}$ all belong to the Krylov space K that is generated by $v_1, v_2, \dots, v_{m+d}$ and, furthermore, the intersection of K and the null space of A is contained in the Krylov space generated by these vectors. We may therefore proceed using $\beta_1 + \gamma_1, \beta_2 + \gamma_2, \dots, \beta_{m+d} + \gamma_{m+d}$ instead of $v_1, v_2, \dots, v_{m+d}$.

If $m = 0$ then there is nothing to prove: The two "experiments" described in the lemma are identical.

Suppose instead that $m = 1$. Since `failure` was not reported before this step, the Krylov space generated by $Av_1, Av_2, \dots, Av_{d+1}$ includes a nonzero vector in the intersection of the image of A and the null space, so that the Krylov space generated by

$$A(\beta_1 + \gamma_1), A(\beta_2 + \gamma_2), \dots, A(\beta_{d+1} + \gamma_{d+1})$$

— that is, by $A\beta_1, A\beta_2, \dots, A\beta_{d+1}$ — includes such a vector as well. Consequently, at least one of $A\beta_1, A\beta_2, \dots, A\beta_{d+1}$ must be nonzero.

Since these vectors are all in $\mathcal{V}_2$, there exist nonnegative integers $k_1, k_2, \dots, k_{d+1}$ such that $k_i = 0$ if $A\beta_i = 0$ and such that $k_i > 0$ and $A^{k_i - 1}(A\beta_i) \neq 0 = A^{k_i}(A\beta_i)$ otherwise, for $1 \leq i \leq d + 1$. Reordering $v_1, v_2, \dots, v_{d+1}$ (and $\beta_1 + \gamma_1, \beta_2 + \gamma_2, \dots, \beta_{d+1} + \gamma_{d+1}$) as needed, we may assume without loss of generality that $k_{d+1} \geq k_i$ for $1 \leq i \leq d$, so that $k_{d+1} > 0$.

Since $m = 1$ (so that the matrix A_2 consists of a single nilpotent Jordan block) and $k_{d+1} \geq k_i$ for $1 \leq i \leq d$, there exist polynomials $\varphi_1, \varphi_2, \dots, \varphi_d \in \mathsf{F}_q[x]$, which depend only on $\beta_1, \beta_2, \dots, \beta_{d+1}$ (so that, in particular, they are independent of $\gamma_1, \gamma_2, \dots, \gamma_{d+1}$), such that $\varphi_i(A)\beta_{d+1} = \beta_i$ for $1 \leq i \leq d$.

Now, the Krylov space that is generated by $\beta_1 + \gamma_1, \beta_2 + \gamma_2, \dots, \beta_{d+1} + \gamma_{d+1}$ is clearly the same as the Krylov space generated by $\beta_{d+1} + \gamma_{d+1}$ and by the vectors

$$(\beta_i + \gamma_i) - \varphi_i(A)(\beta_{d+1} + \gamma_{d+1}),$$

for $1 \leq i \leq d$. Note also that

$$(\beta_i + \gamma_i) - \varphi_i(A)(\beta_{d+1} + \gamma_{d+1}) = \gamma_i - \varphi_i(A)\gamma_{d+1} \in \mathcal{V}_3$$

for $1 \leq i \leq d$; let $\gamma_i' = \gamma_i - \varphi_i(A)\gamma_{d+1}$ for $1 \leq i \leq d$.

Recall that the vectors $\gamma_1, \gamma_2, \dots, \gamma_d$ are chosen uniformly and independently from $\mathcal{V}_3$ and, furthermore, that they are chosen independently of $\varphi_1, \varphi_2, \dots, \varphi_d \in \mathsf{F}_q[x]$ or $\gamma_{d+1} \in \mathcal{V}_3$. It follows that the vectors $\gamma_1', \gamma_2', \dots, \gamma_d'$ are chosen uniformly and independently from $\mathcal{V}_3$ as well — that is, every sequence of vectors of this form is selected with probability $|\mathcal{V}_3|^{-d}$.

Now it suffices to note that, by the choice of k_{d+1}, the vector $\widehat{\beta} = A^{k_{d+1}}(\beta_{d+1} + \gamma_{d+1}) = A^{k_{d+1}-1}(A\beta_{d+1})$ is a nonzero element of $\mathcal{V}_2$ such that $A\widehat{\beta} = 0$, so that $\widehat{\beta}$ is in the intersection of the null space of A and the Krylov space generated by $Av_1, Av_2, \dots, Av_{d+1}$. Indeed (again, since $m = 1$), the single vector $\widehat{\beta}$ forms a basis for this space, while the vectors $\widehat{\beta}, \gamma_1', \gamma_2', \dots, \gamma_d'$ span the intersection of K and the null space of A. This establishes the claim when $m = 1$.

If $m \geq 2$ then we begin as before by defining the nonnegative integers $k_1, k_2, \dots, k_{m+d}$ by setting k_i to be 0 if $A\beta_i = 0$ and by choosing $k_i > 0$ such that $A^{k_i - 1}(A\beta_i) \neq 0 = A^{k_i}(A\beta_i)$ otherwise. Reordering $v_1, v_2, \dots, v_{m+d}$ as needed, we may assume without loss of generality that $k_{m+d} \geq k_i$ for $1 \leq i \leq m + d - 1$. Once again, since the Krylov space generated by $A\beta_1, A\beta_2, \dots, A\beta_{m+d}$ includes a nonzero element of the null space of A, these vectors are not all zero and $k_{m+d} \geq 1$.

At this point a consideration of the structure of the matrix A_2 can be used to conclude that there exists a set of vectors $\omega_1, \omega_2, \ldots, \omega_m \in \mathcal{V}_2$ and nonnegative integers $\ell_1, \ell_2, \ldots, \ell_m$ such that the vectors $\omega_1, \omega_2, \ldots, \omega_m$ generate $\mathcal{V}_2$ as a Krylov space, and such that

$$A^{\ell_1}\omega_1, A^{\ell_2}\omega_2, \ldots, A^{\ell_m}\omega_m$$

is a basis for the intersection of $\mathcal{V}_2$ and the null space of A. Furthermore, these values can be chosen in such a way that

$$A^{k_{m+d}}\beta_{m+d} = A^{\ell_m}\omega_m + \sum_{i=1}^{m-1} A^{\ell_i}\omega_i$$

for values $\delta_1, \delta_2, \ldots, \delta_{m-1} \in \mathsf{F}_q$, and such that β_i can be written as

$$\beta_i = \beta_i' + \varphi_i(A)\beta_{m+d}$$

where β_i' is contained in the Krylov space generated by $\omega_1, \omega_2, \ldots, \omega_{m-1}$ and where $\varphi_i \in \mathsf{F}_q[x]$ for $1 \le i \le m+d-1$ — for it would not be the case that $k_{m+d} \ge k_i$ for $1 \le i \le m+d-1$, otherwise. The vectors $\beta_1', \beta_2', \ldots, \beta_{m+d-1}'$ and the polynomials $\varphi_1, \varphi_2, \ldots, \varphi_{m+d-1}$ depend only on the vectors $\beta_1, \beta_2, \ldots, \beta_{m+d}$, so they are clearly chosen independently of $\gamma_1, \gamma_2, \ldots, \gamma_{m+d}$.

Note that the Krylov space generated by $\beta_1 + \gamma_1, \beta_2 + \gamma_2, \ldots, \beta_{m+d} + \gamma_{m+d}$ is the same as the vector space generated by $\beta_{m+d} + \gamma_{m+d}$ and the vectors

$$\begin{aligned}
&\beta_i + \gamma_i - \varphi_i(A)(\beta_{m+d} + \gamma_{m+d}) \\
&\quad = \beta_i' + \varphi_i(A)\beta_{m+d} + \gamma_i - \varphi_i(A)\beta_{m+d} - \varphi_i(A)\gamma_{m+d} \\
&\quad = \beta_i' + \gamma_i',
\end{aligned}$$

for $\gamma_i' = \gamma_i - \varphi_i(A)\gamma_{m+d} \in \mathcal{V}_3$.

Since $\gamma_1, \gamma_2, \ldots, \gamma_{m+d-1}$ are chosen uniformly and independently from $\mathcal{V}_3$ and, furthermore, they are chosen independently of the polynomials $\varphi_1, \varphi_2, \ldots, \varphi_{m+d-1}$ and the vector γ_{m+d}, it is clear that $\gamma_1', \gamma_2', \ldots, \gamma_{m+d-1}'$ are chosen uniformly and independently from $\mathcal{V}_3$ as well (that is, every such sequence is selected with probability $|\mathcal{V}_3|^{1-m-d}$).

Clearly $A(\beta_i' + \gamma_i') = A\beta_i'$ for $1 \le i \le m+d-1$ and the Krylov space generated by these vectors is contained in the Krylov space generated by $\omega_1, \omega_2, \ldots, \omega_{m-1}$ — for β_i' is included in this space for $1 \le i \le m+d-1$. On the other hand, one can see by the choice of the vectors β_i' that the Krylov space generated by $A\beta_1', A\beta_2', \ldots, A\beta_{m+d-1}', A\beta_{m_d}$ is the same as the Krylov space generated by the vectors $A\beta_1, A\beta_2, \ldots, A\beta_{m+d}$, and we know that the intersection of this space and the null space of A has dimension m. It must therefore be the case that the intersection of the Krylov space generated by $A\beta_1', A\beta_2', \ldots, A\beta_{m+d-1}'$ and the null space of A must be a vector space with dimension $m-1$ — in particular, it must be the intersection of the Krylov space generated by $A\omega_1, A\omega_2, \ldots, A\omega_{m-1}$ and the null space of A.

As noted above, $\gamma_1', \gamma_2', \ldots, \gamma_{m+d-1}'$ are selected uniformly and independently from $\mathcal{V}_3$. Proceeding inductively (on m), we may now conclude that the intersection of the Krylov space generated by the vectors $A\beta_1', A\beta_2', \ldots, A\beta_{m+d-1}'$ and the null space of A has a basis $\widehat{\beta}_1, \widehat{\beta}_2, \ldots, \widehat{\beta}_{m-1}$, while the intersection of the Krylov space generated by $\beta_1' + \gamma_1', \beta_2' + \gamma_2', \ldots, \beta_{m+d-1}' + \gamma_{m+d-1}'$ and the null space of A is spanned by a sequence of vectors $\widehat{\beta}_1, \widehat{\beta}_2, \ldots, \widehat{\beta}_{m-1}, \gamma_1'', \gamma_2'', \ldots, \gamma_d''$, for $\gamma_1'', \gamma_2'', \ldots, \gamma_d'' \in \mathcal{V}_3$, such that every possible choice of the vectors $\gamma_1'', \gamma_2'', \ldots, \gamma_d''$ is obtained with probability $|\mathcal{V}_3|^{-d}$.

It remains only to notice that the set of vectors

$$A^{k_{m+d}}\beta_{m+d}, \widehat{\beta}_1, \widehat{\beta}_2, \ldots, \widehat{\beta}_{m-1}$$

forms a basis for the intersection of the Krylov space generated by $A\beta_1, A\beta_2, \ldots, A\beta_{m+d}$ and the null space of A, while the vectors

$$A^{k_{m+d}}\beta_{m+d}, \widehat{\beta}_1, \widehat{\beta}_2, \ldots, \widehat{\beta}_{m-1}, \gamma_1'', \gamma_2'', \ldots, \gamma_d''$$

span the intersection of the Krylov space generated by $\beta_1 + \gamma_1, \beta_2 + \gamma_2, \ldots, \beta_{m+d} + \gamma_{m+d}$ and the null space of A, to complete the proof.

C. REFERENCES

[10] W. Eberly and B. Hovinen. Bounding the nullities of random block Hankel matrices: An alternative approach. Technical Report 2005-779-10, Department of Computer Science, University of Calgary, 2005. Available online at `www.cpsc.ucalgary.ca/~eberly/Research/publications.php`.