Mathematische Zeitschrift

Matrix Kloosterman sums modulo prime powers

M. Erdélyi[1,2] · Á. Tóth[3,4] · G. Zábrádi[1,3]

Abstract

We give optimal bounds for matrix Kloosterman sums modulo prime powers extending earlier work of the first two authors on the case of prime moduli. These exponential sums arise in the theory of the horocyclic flow on GL_n.

1 Introduction

The purpose of this paper is to give good upper bounds for the sums

$$K_n(A, B; p^k) = \sum_{X \in GL_n(\mathbb{Z}/p^k\mathbb{Z})} \psi((AX + X^{-1}B)/p^k), \qquad (1)$$

with given $A, B \in \mathbb{Z}^{n \times n}$ where for an $n \times n$ matrix X we let $\psi(X) = e^{2\pi i \operatorname{Tr} X}$. Note that $\psi(XY) = \psi(YX)$. A good upper bound may mean different things, it could be optimal, or somewhat crude but easily usable, and we will provide both.

These sums themselves are of independent interest. They arise naturally in certain equi-distribution problems and are natural analogues of the classical Kloosterman sums $\sum_{x \ (p)} e^{2\pi i(ax+bx^{-1})/p}$.

In our earlier paper [6] we dealt with the case $k = 1$. As usual upper bounds modulo a prime require the heavy machinery of some type of Weil cohomology. For higher prime powers the methods are usually of a very different sort, based on Taylor expansions, and occasionally referred to as the stationary phase [3]. One such example is provided by Salie's explicit evaluation of the one-dimensional case in [14]. We will prove that such a result holds generically even for $n \times n$ matrices, but generic here is much more restricted than being non-zero, or even invertible.

✉ G. Zábrádi
  gergely.zabradi@ttk.elte.hu

1 "Lendület" Analytic Number Theory and Representation Theory Research Group, HUN-REN Alfréd Rényi Institute of Mathematics, Budapest, Hungary

2 Budapest University of Technology and Economics, Budapest, Hungary

3 Institute of Mathematics, Eötvös Loránd University, Pázmány Péter sétány 1/C, H1117, Budapest, Hungary

4 "Lendület" Automorphic Research Group, HUN-REN Alfréd Rényi Institute of Mathematics, Budapest, Hungary

We summarize the main results. Clearly if p divides all matrix entries of A and B then one may clear appropriate powers and either arrive at a trivial sum, or one where one of A or B is different from the zero-matrix. From now on we will assume that this is the case and denote it as

$$\gcd(A, B, p) = 1.$$

First we have the following reductions to a counting argument.

Proposition 1.1 *Assume $\gcd(A, B, p) = 1$ and that $k > 1$.*

1. *If $k = 2l$ then*

$$K_n(A, B; p^k) = \sum_X \psi((AX + X^{-1}B)/p^k)$$

where the sum is over $X \in GL_n(\mathbb{Z}/p^k\mathbb{Z})$ satisfying $XAX \equiv B \mod p^l$.
2. *If $k = 2l + 1$ then*

$$K_n(A, B; p^k) = \frac{1}{p^{n^2}} \sum_X \psi((AX + X^{-1}B)/p^k)S_{A,B}(X; p)$$

where the sum is over $X \in GL_n(\mathbb{Z}/p^k\mathbb{Z})$, $XAX \equiv B \mod p^l$, and where

$$S_{A,B}(X; p) = \sum_{U \bmod p} \psi((SU + TU^2)/p),$$

with $S = (AX - X^{-1}B)/p^l$ and $T = AX$.

Remark 1.2 Note that if $p \neq 2$, then $S_{A,B}(X; p)$ is either 0 or equals $G(AX; p)$, for the generalized Gauss sum

$$G(T; p) = \sum_{U \bmod p} \psi(TU^2/p)$$

See below in Sect. 3.

The case when A is invertible modulo p is special and can be made more explicit, in complete analogue of Salie's evaluation of the classical one dimensional Kloosterman sum [14].

Corollary 1.3 1. $K_n(A, B; p^k) = 0$ *unless the invariant factors, (the Smith normal forms), of A and B agree up to mod p^l, where $l = [k/2]$. In particular if $k > 1$ and $\gcd(\det A, p) = 1$ (ie. A is invertible mod p) then $K_n(A, B; p^k) = 0$ unless $\gcd(\det B, p) = 1$ as well.*
2. *Assume that $\gcd(\det(AB), p) = 1$ and that AB is regular semisimple mod p, (i.e. all eigenvalues are different). Then*

   (a) *If $k = 2l$ then*

   $$K_n(A, B; p^k) = p^{kn^2/2} \sum_Y \psi(2Y/p^k)$$

   where the sum is over $Y \in GL_n(\mathbb{Z}/p\mathbb{Z})$, $Y^2 \equiv AB \mod p^k$.

(b) *If $k = 2l + 1$ then*

$$K_n(A, B; p^k) = \zeta \, p^{kn^2/2} \sum_Y \psi(2Y/p^k)$$

where the sum is over $Y \in GL_n(\mathbb{Z}/p\mathbb{Z})$, $Y^2 \equiv AB \mod p^k$, and where ζ is a p-th root of unity.

(c) *In particular we have $|K_n(A, B; p^k)| \leq 2^n \, p^{kn^2/2}$.*

Note that in the regular semisimple case we have square root cancellation. Also note that we do not assume that the eigenvalues of AB are defined over $\mathbb{F}_p$. Finally, this condition is generic, its complement is a Zariski closed set.

We now return to the non-generic cases. By Proposition 1.1 in order to bound the sums in 1 we need to bound

$$N^*(A, B; p^l) = \#\{X \in GL_n(\mathbb{Z}/p^l\mathbb{Z}) | AX \equiv X^{-1}B \mod p^l\}. \tag{2}$$

Note that

$$N^*(A, B; p^l) = N^*(XA, BX^{-1}; p^l)$$

for any $X \in GL_n(\mathbb{Z}/p^l\mathbb{Z})$ and so if $N^*(A, B; p^l)$ is not zero, then it equals $N^*(C, C; p^l)$ for $C = XA$, for any X for which $XA \equiv BX^{-1} \pmod{p^l}$.

For $l = 1$ it is possible (see Thm. 1.5 below) to describe $N^*(C, C; p)$ explicitly. This allows one to show that for a well defined exponent $e = e_C$ we have $m_n < N^*(C, C; p)/p^e < M_n$ for some absolute constants m_n, M_n that depend on n only. The exponent $e = e_C$ itself depends on the combinatorial type of the Jordan decomposition of C over an algebraic closure of $\mathbb{F}_p$.

As a first step towards this goal we have the following reduction.

Proposition 1.4 *For any $C \in \mathbb{F}_p^{n \times n}$ let $m_{C^2}(x) = \prod_{j=1}^r f_j(x)^{k_j} \in \mathbb{F}_p[x]$ be the minimal polynomial of C^2, where the $f_j(x) \in \mathbb{F}_p[x]$ are irreducible. Let $V_j = \ker f_j(C^2)^{k_j}$ and $C_j = C_{|V_j}$ be the restriction of C to V_j. Then we have*

$$N^*(C, C; p) = \prod_{j=1}^r N^*(C_j, C_j; p). \tag{3}$$

In case of a primary minimal polynomial the explicit counting formulas depend on the value of $\left(\frac{x}{f(x)}\right)$ where as usual [13] $\left(\frac{x}{f(x)}\right)$ is the quadratic residue symbol, defined for an irreducible polynomial f,

$$\left(\frac{g(x)}{f(x)}\right) = \begin{cases} 1 & \text{if } g \text{ is a non - zero square in } \mathbb{F}_p[x]/(f), \\ -1 & \text{if } g \text{ is a not a square in } \mathbb{F}_p[x]/(f), \\ 0 & \text{if } f | g \end{cases}$$

To state our main result we need to introduce further notation the details of which are presented in Sect. 3. For any partition $\lambda = [n_1, ..., n_k]$ we let N_λ be a nilpotent matrix with Jordan blocks of size $n_1, ..., n_k$. For $q = p^d$ let $\mathbb{F}_q$ be the field with q elements, and $Z_{GL_{|\lambda|}(\mathbb{F}_q)}(N_\lambda)$ be the centralizer of N_λ in the group $GL_{|\lambda|}(\mathbb{F}_q)$ where $|\lambda| := n_1 + \cdots + n_k$. (See Proposition 3.1).

If we have partitions μ, ν we let $\lambda = \mu + \nu$ be their join. Also the dual partition λ' of λ may be defined via the matrix N_λ as $\lambda' = [d_1, ..., d_k]$, $d_1 \geq d_2 \geq \cdots \geq d_k$ where

$$\dim \ker N_\lambda^j = d_1 + \cdots + d_j.$$

Theorem 1.5 *Assume that $p \neq 2$. Let C be a $m \times m$ matrix and assume that the minimal polynomial of C^2 is of the form $f(x)^k$, where $f(x) \in \mathbb{F}_p[x]$ is irreducible. We let $q = p^d$, $d = \deg f$, and λ be the partition of m/d with dual $\lambda' = [d_1, ..., d_k]$, $d_1 \geq d_2 \geq \cdots \geq d_k$ where*

$$\frac{1}{d} \dim_{\mathbb{F}_p} \ker f^j(C^2) = d_1 + \cdots + d_j.$$

1. *If $\left(\frac{x}{f(x)}\right) = 1$ we have $N^*(C, C, p) = N_+^*(\lambda, q)$, where $q = p^{\deg f}$ and*

$$N_+^*(\lambda, q) = \sum_{\lambda = \mu + \nu} \frac{\#Z_{GL_{|\lambda|}(\mathbb{F}_q)}(N_\lambda)}{\#Z_{GL_{|\mu|}(\mathbb{F}_q)}(N_\mu)\#Z_{GL_{|\nu|}(\mathbb{F}_q)}(N_\nu)} . \qquad (4)$$

2. *If $\left(\frac{x}{f(x)}\right) = -1$ then all d_j will be even, and $\lambda = \mu + \mu$ for some partition μ. Then we have $N^*(C, C, p) = N_-^*(\lambda, q)$ where*

$$N_-^*(\lambda, q) = \frac{\#Z_{GL_{|\lambda|}(\mathbb{F}_q)}(N_\lambda)}{\#Z_{GL_{|\mu|}(\mathbb{F}_{q^2})}(N_\mu)} . \qquad (5)$$

3. *If $f(x) = x$ we have*

$$N^*(C, C; p) = p^{\sum_{j=1}^k d_j^2} \prod_{j=1}^k \prod_{t_j=1}^{r_j} \left(1 - \frac{1}{p^{t_j}}\right) \qquad (6)$$

where we put $r_j := d_j - d_{j+1}$ ($j = 1, \ldots, k$ where $d_{k+1} := 0$), ie. r_j is the number of blocks of size $j \times j$ ($j = 1, \ldots, k$) in the Jordan normal form of C.

We are now ready to state our main bounds. For a refined statement we need the stable rank of a matrix A defined as $\mathrm{rk}_\infty A = \lim_{m\to\infty} \mathrm{rk}\, A^m$.

Theorem 1.6 *Assume that $p \neq 2$ and $l \geq 1$.*

1. *Let $r = \mathrm{rk}(C \bmod p)$, $r_\infty = \mathrm{rk}_\infty(C \bmod p)$,*

$$N^*(C, C; p^l) \leq 2^{r_\infty} p^{e(l,n,r,r_\infty)}$$

   where

$$e(l, n, r, r_\infty) = (n - r)^2 + (r - r_\infty)^2 + r_\infty^2/2 + (l - 1)\left((n - r)(n - r_\infty) + r_\infty^2/2\right).$$

2. *Assume that $r = \mathrm{rk}(A \bmod p) = \mathrm{rk}(B \bmod p) > 0$. We have*

$$N^*(A, B; p^l) \leq 2^r \begin{cases} p^{ln(n-r)} & \text{if } r \leq n/2 \\ p^{ln^2/2} & \text{if } n/2 < r \leq n \end{cases}. \qquad (7)$$

3. *In particular if $n > 1$ and $\gcd(A, B, p) = 1$ then $N^*(A, B; p^l) \leq 2^n p^{l(n^2-n)}$. If $\gcd(\det A, \det B, p) = 1$ then $N^*(A, B; p^l) \leq 2^n p^{ln^2/2}$.*

Remark 1.7 The case $p = 2$ is special, in view of (14). Assume $AX = X^{-1}B = C \bmod p$. If $C \neq I_n$, in particular, if $AB \neq I_n \bmod 2$, the bound $2^{l(n^2-n)}$ still holds.

We also need the following general bound for the sum $S_{A,B}(X; p)$.

Proposition 1.8 *Assume that A, B are such that there exists X so that $AX \equiv X^{-1}B \bmod p^l$. Let*

$$S_{A,B}(X; p) = \sum_{U \bmod p} \psi((SU + TU^2)/p)$$

with $S = (AX - X^{-1}B)/p^l$ and $T = AX \bmod p$. Assume that $p \neq 2$. We have that

$$|S_{A,B}(X; p)| \leq p^{(n-r)(n-r_\infty)+r_\infty^2/2}$$

where $r = \mathrm{rk}\, T$, $r_\infty = \mathrm{rk}_\infty T$. In particular, we always have $|S_{A,B}(X; p)| \leq p^{n^2-n}$ and under the additional assumption that A is invertible, we have $|S_{A,B}(X; p)| \leq p^{n^2/2}$.

In view of Proposition 1.1 as a corollary of the above we have the following

Theorem 1.9 *Assume that $n > 1$ and the matrices A, B are not both $0 \bmod p$. We then have the following bounds.*

1. *If $k = 1$, by* [6]

$$|K_n(A, B; p)| \leq \begin{cases} 2p^{n^2-n+1} & \text{for all } A, B \\ 4p^{3n^2/4} & \text{if } \gcd(\det A, \det B, p) = 1 \\ 4p^{n^2/2} & \text{if } \gcd(\det B, p) = 1, \text{ and } AB^{-1} \text{ regular semisimple.} \end{cases}$$

2. *If $k > 1$ we have*

$$|K_n(A, B; p^k)| \leq 2^n \begin{cases} p^{kn^2-\lceil \frac{k}{2}\rceil n} & \text{for all } A, B \\ p^{kn^2-\lceil \frac{k}{2}\rceil n^2} & \text{if } \gcd(\det A, \det B, p) = 1. \end{cases}$$

The paper is organized as follows. First in section 2 we prove Proposition 1.1, this then gives the optimal bounds for the generic situation. In the next section we list some facts concerning partitions, the Sylvester equation and multivariable Gauss sums. These will be used in the following sections. First in Sect. 4 we give upper bounds for the number of solutions of various quadratic equations in matrices modulo a prime. In the last section Sect. 5 we then prove the estimates in the last three statements above.

While this work was in progress El-Baz, Lee and Strömbergsson [5] independently arrived to quantitavely similar bounds in their work on the equidistribution of rational points on horocycles. While there are some overlaps the main results are different in nature.

2 Reduction to counting

In this section we prove Proposition 1.1 and its corollaries. As in the statement we need to deal with the case of even and odd exponents separately.

2.1 The case p^{2l}

Let $k = 2l$. For any unit $U \in GL_n(\mathbb{Z}/p^k\mathbb{Z})$ $K_n(A, B; p^k) = \sum_{X \in GL_n(\mathbb{Z}/p^k\mathbb{Z})} \psi((AXU + U^{-1}X^{-1}B)/p^k)$.

Let $H < GL_n(\mathbb{Z}/p^k\mathbb{Z})$ be the subgroup of matrices U such that $U \equiv I$ (p^l). Explicitly we have

$$H = \{I + U_1 p^l \mid U_1 \ (p^l)\}.$$

Since $(I + U_1 p^l)^{-1} = I - U_1 p^l$, we have that

$$K_n(A, B; p^k) = \frac{1}{p^{ln^2}} \sum_{X \in GL_n(\mathbb{Z}/p^k\mathbb{Z})} \psi((AX + BX^{-1})/p^k) \sum_{U_1 \bmod p^l} \psi((AXU + BUX^{-1})/p^l)).$$

As noted above, $\psi(BUX^{-1}) = \psi(X^{-1}BU)$, and so the inner sum

$$\sum_{U_1 \bmod p^l} \psi((AX - X^{-1}B)U/p^l))$$

vanishes, unless $AX \equiv X^{-1}B \bmod p^l$. This proves the first claim in Proposition 1.1.

2.2 The case p^{2l+1}

Let $k = 2l + 1$. We again use the subgroup H defined above which in this case consists of matrices $U = I + U_1 p^l + U_2 p^{2l}$ where U_1 (resp. U_2) runs on $(\mathbb{Z}/p^l\mathbb{Z})^{n \times n}$ (resp. on $(\mathbb{Z}/p\mathbb{Z})^{n \times n}$), with inverse

$$U^{-1} = I - U_1 p^l + (U_1^2 - U_2)p^{2l}.$$

Therefore

$$K_n(A, B; p^k) = \frac{1}{p^{n^2(l+1)}} \sum_{X \in GL_n(\mathbb{Z}/p^k\mathbb{Z})} \sum_U \psi((AXU + X^{-1}BU^{-1})/p^k)$$

where $U = I + U_1 p^l + U_2 p^{2l}$ is such that U_1 will run mod p^l and U_2 will run mod p.

Now fix $X \in GL_n(\mathbb{Z}/p^k\mathbb{Z})$, and consider

$$\sum_U \psi((AX(I + U_1 p^l + U_2 p^{2l}) + X^{-1}B(I - U_1 p^l + (U_1^2 - U_2)p^{2l}))/p^k)$$

$$= \psi((AX + BX^{-1})/p^k)S_1(X)S_2(X)$$

where

$$S_1(X) = \sum_{U_1 \ (p^l)} \psi((AX - X^{-1}B)U_1 + X^{-1}BU_1^2 p^l)/p^{l+1})$$

and

$$S_2(X) = \sum_{U_2 \ (p)} \psi((AX - X^{-1}B)U_2/p).$$

Note that $S_2 = 0$ unless $AX \equiv X^{-1}B \bmod p$ in which case $S_2 = p^{n^2}$. S_1 is a Gauss sum in matrices, albeit a very special one. By the condition from S_2 we have that $AX - X^{-1}B = pM$ for some integral matrix M, then we have

$$S_1(X) = \sum_{U_1 \ (p^l)} \psi((MU_1 + TU_1^2)/p^l)$$

where $T \equiv X^{-1}Bp^{l-1} \equiv AXp^{l-1} \bmod p^l$. This gives the claim when $l = 1$. For $l > 1$ note that in view of $pT \equiv 0 \bmod p^l$ we have for any V that

$$\sum_{U_1 \ (p^l)} \psi((MU_1 + TU_1^2)/p^l) = \sum_{U_1 \ (p^l)} \psi((M(U_1 + pV) + T(U_1 + pV)^2)/p^l)$$

$$= \psi(MV/p^{l-1}) \sum_{U_1 \ (p^l)} \psi((MU + TU^2)/p^l) \,.$$

and so that $S_1(X) = \psi(MV/p^{l-1})S_1(X)$. A suitable choice of V shows that $S_1(X) = 0$ unless $M \equiv 0 \mod p^{l-1}$. In the original matrices A, B this is equivalent to $AX \equiv X^{-1}B \mod p^l$ in which case $S_1 = p^{(l-1)n^2} S_{A,B}(X; p)$. This gives the second claim of Proposition 1.1.

2.3 The regular semisimple case

The proof of Corollary 1.3 Note that if $X \in GL_n(\mathbb{Z}/p^k\mathbb{Z})$, $AX \equiv X^{-1}B \mod p^l$, and A is invertible then so is B. Moreover if $Y = AX$ then $Y^2 \equiv AB \mod p^l$. Assume now that $Y^2 \equiv AB \mod p^l$ and let $X = A^{-1}Y$. Then $AX \equiv Y \mod p^l$ and $X^{-1}B \equiv Y^{-1}AB \equiv Y \mod p^l$. If $AB = Y^2$ is regular semisimple then all the eigenvalues of Y are different and no two of them sum to 0. This is exactly the condition (see Sect. 3.3) to modify Y by adding a suitable $p^l Z$ in such a way that $Y^2 \equiv AB$ hold mod p^k as well. In this case

$$K_n(A, B; p^k) = K_n(Y, Y; p^k)$$

and the claim is an easy corollary of the calculations done in previous two subsections and the regularity of Y^2.

3 Technical background

3.1 Partitions

A partition of an integer n is an ordered set $\lambda = [n_1, n_2, \ldots, n_r]$, $n_1 \geq n_2 \geq \cdots \geq n_r > 0$, of integers satisfying $\sum_i n_i = n$. We will write $n = |\lambda|$. If λ and μ are two partitions, $\lambda + \mu$ is the partition obtained by taking the parts of λ and μ together (and ordering them). We denote by $[n]$ the partition with one part n. In general if a number j appears r_j times in λ, the sequence $[..., j, ..., j, ...]$ will be replaced by $[..., j^{r_j}, ...]$, so for example the partition with n parts all equal to 1 is written as $[1^n]$.

Given a partition λ its associated Young (Ferrer) diagram has r rows with $n_1, n_2, ...n_r$ boxes in each row. For example for $\lambda = [4, 3, 1]$ we have the diagram
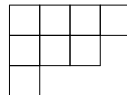
The transpose of the diagram of λ is also a Young diagram of a partition λ' called the conjugate or dual partition to λ which may be described as follows. Let $r_i = r_i(\lambda)$ be the number of parts of λ which are equal to $i \geq 1$ and $d_i = \sum_{j \geq i} r_j$. Then

$$\lambda' = [d_1, d_2, \ldots] \tag{8}$$

which has the diagram
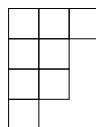
in our example $\lambda = [4, 3, 1]$.

3.2 Centralizers in $GL_n(\mathbb{F}_q)$

At first assume $p \neq 2$ and let N be the nilpotent transformation of an $\mathbb{F}_q$ vector space V of dimension n. Then V becomes an $\mathbb{F}_q[T]$-module where T acts via N, $Tv = Nv$. Such modules are isomorphic to the module $V_\lambda = \oplus_j \mathbb{F}_q[T]/(T^{n_j})$, for some partition $\lambda = [n_1, \ldots, n_k]$, $n_1 + \cdots + n_k = n$ which is unique by the structure theorem of finitely generated modules over principal ideal domains. To show the partition λ associated to N we will use the notation $N = N_\lambda$.

Note that the the dual partition λ' arises from considering $d_i = \dim(\ker(N^i)) - \dim(\ker(N^{i-1}))$. To see this assume we switch to the matrix point of view and assume that $N \in M_n(\mathbb{F}_q)$ is a nilpotent matrix over $\mathbb{F}_q$ with r_j blocks of size $j \times j$ ($j = 1, \ldots, n_1$) in the Jordan normal form of N. Let $d_i = \sum_{j \geq i} r_j$. as above. Then it is easy to see that d_1 is the number blocks, which also equals $\dim \ker N$. The claim then follows from an easy inductive argument. One can alternatively define

$$d_i = r_i + r_{i+1} + \cdots + r_k = \dim(\ker(N) \cap \mathrm{Im}(N^{i-1})).$$

Finally we will need the order of the centralizer of unipotent elements in $GL_n(\mathbb{F}_q)$. Note that the centralizer of the unipotent $I + N$ is the same as the centralizer of the nilpotent transformation N.

Proposition 3.1 *Let $N = N_\lambda = [n_1, ..., n_k]$ with dual partition $\lambda' = [d_1, ..., d_{n_1}]$. Then the centralizer of N has cardinality*

$$\#Z_{GL_n(\mathbb{F}_q)}(N) = \left(\prod_{j=1}^{n_1} (q^{r_j} - 1)(q^{r_j} - q) \ldots (q^{r_j} - q^{r_j - 1}) \right) \cdot q^{\sum_{j=1}^{n_1} (d_j^2 - r_j^2)}$$

$$= q^{\sum_{j=1}^{k} d_j^2} \prod_{j=1}^{k} \prod_{t_j=1}^{r_j} \left(1 - \frac{1}{q^{t_j}} \right).$$

Proof This is Corollary IV.I.8 in [16]. □

Remark 3.2 Note that if we define $\phi_r(T) = \prod_{j=1}^{r}(1 - T^j)$ and let $\phi_\lambda(T) = \prod \phi_{r_i(\lambda)}(T)$ then the statement can be rewritten as

$$\#Z_{GL_n(\mathbb{F}_q)}(N) = q^{\sum_{j=1}^{k} d_j^2} \phi_\lambda(1/q).$$

3.3 Sylvester's equation

Assume that A is $m \times m$, B is $n \times n$ and X and C are $m \times n$ matrices. The matrix equation

$$AX - XB = C,$$

called Sylvester's equation [17], has a rich literature over the real, or complex fields in view of the important role it plays in various applications. (See. e.g. [4].) There are two important questions here, existence of solutions, and a description of all solutions $\{X \mid AX - XB = C\}$.

For our task of estimating $K_n(A, B; p^k)$ we will concentrate on estimating the number of solutions. If the field of coefficients is $\mathbb{F}_q$ for some p-power q, then the number of solutions is clearly either 0 or q^d, where

$$d = d_{A,B} = \dim\{X \mid AX - XB = 0\} \leq mn. \tag{9}$$

While the bound by mn is trivial, in the case when $A = \lambda I_n$, $B = \lambda I_m$ for the same scalar λ, one has $d_{A,B} = mn$.

Note that we may interpret the equation via linear transformations. To do so let $W = \mathbb{F}_q^m$, $V = \mathbb{F}_q^n$ viewed as column vectors. Both W and V become $\mathbb{F}_q[T]$-modules via mapping T to A and B respectively.

If $AX = XB$ then $g(A)X = Xg(B)$ for any polynomial $g \in \mathbb{F}_q[T]$ and so X gives rise to a module homomorphism from W to V which we denote $\hom_{\mathbb{F}_q[T]}(W, V)$.

For an irreducible polynomial $f \in \mathbb{F}_q[T]$ let

$$V_{f^e} = \{v \in V \mid f^e(A)v = 0\},$$

and similarly for W. The f-primary component of V is $\cup_{e=1}^{\infty} V_{f^e}$, which we denote by V_{f^∞}.

Clearly if $X \in \hom_{\mathbb{F}_q[T]}(W, V)$ then

$$X(W_{f^e}) \subset V_{f^e} \tag{10}$$

which implies

$$\hom_{\mathbb{F}_q[T]}(W, V) = \bigoplus_f \hom_{\mathbb{F}_q[T]}(W_{f^\infty}, V_{f^\infty})$$

the sum over $f \in \mathbb{F}_q[T]$ irreducible. Since the problem is linear, we may go to a finite field extension if needed and then assume that the eigenvalues of A, B are in $\mathbb{F}_q$. For $f = T - \lambda$ and $e \geq 1$ we let

$$d_e(A - \lambda) = \dim V_{(T-\lambda)^e} = \dim \ker(A - \lambda I)^e, \; d_\infty(A - \lambda) = d_n(A - \lambda). \tag{11}$$

As an immediate corollary of the trivial bound (9 we get the following

$$d_{A,B} \leq \sum_{\lambda \in \mathbb{F}_q} d_\infty(A - \lambda)d_\infty(B - \lambda). \tag{12}$$

Note that that for A semisimple the inequality above becomes an equality, showing that the bound is sharp.

The application for us involves the special case when $B = -A$. If A is invertible the above bound is sufficient, but the nilpotent case needs a more refined version given in the following lemma.

Lemma 3.3 *Assume that $A \in M_m(\mathbb{F}_q)$, $B \in M_n(\mathbb{F}_q)$ are nilpotent. Let*

$$k = \dim \ker(A), \quad l = \dim \ker(B),$$

Then

$$\dim\{X \in M_{m,n}(\mathbb{F}_q) \mid AX = XB\} \leq \frac{kn + ml}{2}.$$

Proof As above let $W = \mathbb{F}_q^m$, $V = \mathbb{F}_q^n$ viewed as $\mathbb{F}_q[T]$ modules via mapping T to A and B. Then

$$W \simeq \bigoplus_{i=1}^k \mathbb{F}_q[T]/(T^{m_i}), \quad V \simeq \bigoplus_{j=1}^l \mathbb{F}_q[T]/(T^{n_j}).$$

for some partitions $(m_1, ..., m_k)$, $(n_1, ..., n_l)$ of m and n respectively (with k, l as defined above).

Note that any element X of $\hom_{\mathbb{F}_q[T]}(\mathbb{F}_q[T]/(T^b), \mathbb{F}_q[T]/(T^a))$ is determined by the value X on 1 mod T^b and so by (10)

$$\dim \hom_{\mathbb{F}_q[T]}(\mathbb{F}_q[T]/(T^b), \mathbb{F}_q[T]/(T^a)) = \min(a, b) \tag{13}$$

This gives

$$\dim\{X \in M_{m,n}(\mathbb{F}_q) \mid AX = XB\} = \sum_{i,j=1}^{k,l} \min(m_i, n_j) \le \sum_{i,j=1}^{k,l} \frac{m_i + n_j}{2} \le \frac{kn + ml}{2}.$$

$\square$

Corollary 3.4 *The trivial bound (12) can be strengthened to*

$$\dim\{X \bmod p \mid AX + XA = 0\} \le \sum_{\lambda \in \mathbb{F}_q} d_1(A - \lambda) d_\infty(A + \lambda). \tag{14}$$

3.4 Generalities on multivariable Gauss sums.

We start with a general setup on $V = \mathbb{F}_p^m$. Let $F(x) = Q(x) + L(x)$ with a quadratic form Q, and a linear form L on V and define

$$G(F; p) = \sum_{x \in V} e(F(x)/p)$$

where $e(z) = e^{2\pi i z}$. If $L = 0$ and $p \ne 2$, the sum $G(Q; p)$ is easy to evaluate after diagonalizing Q; it is a product of trivial factors and Gauss sums. The case when $p = 2$ is slightly more involved [8], but still explicit.

For our use in what follows, some of the details are relevant, and so we sketch these. Assume that $p \ne 2$ when we have that Q comes from a bilinear form

$$B(x, y) = Q(x + y) - Q(x) - Q(y)$$

so that $Q(x) = \frac{1}{2}B(x, x)$. B gives rise to a Riesz map $R : V \to V^*$, $R : y \mapsto (R_y : x \mapsto B(x, y))$, which may not be surjective if B is degenerate. Still we have the following dichotomy.

Proposition 3.5 *Assume that $p \ne 2$. Let $F(x) = Q(x) + L(x)$, where Q is a quadratic and L is a linear form on V.*

1. *If L has a Riesz-representative, i.e. $L = R_y$ for some y, then $G(F; p) = e(-Q(y)/p)G(Q; p)$.*
2. *If L does not have a Riesz-representative then $G(F; p) = 0$.*

Proof The first statement is trivial. To see the second, note that L does not have a Riesz representative if and only if ker L does not contain

$$V^\perp = \{v \in V \mid B(v, y) = 0 \text{ for all } y \in V\}.$$

Therefore there is y, such that $L(y) \ne 0$ but $B(x, y) = 0$ for all x, and so that $Q(x + y) = Q(x) + Q(y)$, and in particular $Q(y) = 0$. But then

$$G(F; p) = \sum_x e(F(x + y)/p) = \sum_x e((Q(x + y) + L(x + y))/p) = e(L(y)/p)G(F; p)$$

showing $G(F; p) = 0$ since $e(L(y)/p) \ne 1$. $\square$

The folllowing is an easy corollary of the evaluation of the standard Gauss sum [1].

Corollary 3.6 *Assume that $F = Q + L$ and that L has a Riesz representative, $L = R_y$. Write $Q = Q_0 \perp Q_1$ where Q_0 is totally isotropic, and Q_1 non-degenarate. Let $\dim Q_1 = r$. Then*

$$G(F) = (\det Q_1/p) \, e(-Q(y)/p) \, p^{n-r/2},$$

where $(\cdot/p)$ is Legendre's symbol.

4 The equation $AX \equiv X^{-1}B$ to prime modulus

4.1 Preliminary observations

We will concentrate on the case when the equation $AX = X^{-1}B$ is solvable. As noted above we may then simply assume that $C = AX_0 = X_0^{-1}B$ for some fixed solution X_0 and consider the equation $CX = X^{-1}C$. For this equation we start with the proof of Proposition 1.4.

Proof of Proposition 1.4 For any solution X of the equation $CX = X^{-1}C$ we also have $XC = CX^{-1}$ by multiplying from the left by X^{-1} and from the right by X. So we compute $C^2X = C(CX) = CX^{-1}C = XC^2$, ie. X and C^2 commute. Let us write $m_{C^2}(x) \in \mathbb{F}_p[x]$ for the minimal polynomial of C^2 and write it as $m_{C^2}(x) = \prod_{j=1}^r f_j(x)^{k_j}$ where $f_j(x) \in \mathbb{F}[x]$ is irreducible.

Then we may decompose $\mathbb{F}_p^n = V_1 \oplus \cdots \oplus V_r$ as the direct sum of generalized eigenspaces where $V_j := \ker(f_j(C^2)^{k_j})$. Since X and C^2 commute, V_j is an X-invariant subspace in $\mathbb{F}_p^n$ for $1 \le j \le r$: indeed, for any $v \in V_j$ we have $f_j(C^2)^{k_j}v = 0$ whence $f_j(C^2)^{k_j}(Xv) = Xf_j(C^2)^{k_j}v = 0$ showing $Xv \in \ker(f_j(C^2)^{k_j}) = V_j$. By a similar argument we also deduce that V_j is also C-invariant (as C also commutes with C^2). Therefore restricting the identity $CX = X^{-1}C$ to the subspace V_j we deduce $C_jX_j = X_j^{-1}C_j$ where C_j (resp. X_j) is the restriction of C (resp. of X) to V_j. On the other hand, whenever we have matrices $X_j \in GL(V_j)$ with $C_jX_j = X_j^{-1}C_j$ ($1 \le j \le r$) then we may form the block matrix X from the matrices X_j to obtain a solution of the equation $CX = X^{-1}C$. So the number of solutions of the equation $CX = X^{-1}C$ is the product of the number of solutions on each V_j.

This proves the proposition.

Now assume C is invertible. Since X is also invertible $XCX = C$ if and only if $Y = CX$ satisfies $Y^2 = C^2$ and we will count the number of solutions to this equation under the assumption that C^2 has minimal polynomial $f^k(x)$, with $f(x) \ne x$ irreducible. For this recall that any linear transformation T has a unique multiplicative Jordan–Chevalley decomposition as

$$T = T_s T_u = T_u T_s \tag{15}$$

where T_s is semisimple, T_u is unipotent. If $Y = Y_s Y_u$ and $Y^2 = C^2$ then

$$Y_s^2 = C_s^2 \text{ and } Y_u^2 = C_u^2.$$

In case $p \ne 2$, we can immediately infer that $Y_u = C_u$ from the following

Lemma 4.1 *Assume that $Z_1, Z_2 \in GL_n(\mathbb{F}_p)$ are unipotent such that $Z_1^2 = Z_2^2$. If $p \ne 2$ then $Z_1 = Z_2$.*

Proof For simplicity we use the simple property that for any unipotent element Z in $GL_n(\mathbb{F}_p)$ we have $Z^{p^r} = I$ for some r. Since p is odd

$$Z_1 = (Z_1^2)^{(p^r+1)/2} = (Z_2^2)^{(p^r+1)/2} = Z_2.$$

$\square$

Therefore estimating $N^*(A, B; p)$ is reduced to estimating

$$n(C, C; p) = \{Y \in M_n(\mathbb{F}_p) \mid Y^2 = C^2, YC_u = C_uY\}. \tag{16}$$

Our last observation is now the following

Lemma 4.2 *Put $V := \mathbb{F}_p^n$ and assume that the minimal polynomial of $C^2 \colon V \to V$ is $m_{C^2}(x) = f^k(x)$ where $f \in \mathbb{F}_p[x]$ is irreducible and let $q = p^{\deg f}$. Then V has the structure of an $\mathbb{F}_q$-vectorspace such that all $C, X, Y, Y_s, Y_u \colon V \to V$ as above are $\mathbb{F}_q$-linear for any invertible solution X of the equation $XCX = C$. Further, Y_s^2 is an $\mathbb{F}_q$-scalar multiple of the identity.*

Proof We assume that $f(x) \neq x$ otherwise the claim is trivial. Note that $\mathbb{F}_p[x]/(f)$ is isomorphic to the field $\mathbb{F}_q$ with $q = p^{\deg f}$ elements. We choose such an isomorphism and let α denote the image of x in $\mathbb{F}_q$, so that $\mathbb{F}_q = \mathbb{F}_p[\alpha]$. The ring $\mathbb{F}_p[x]/(f^k)$ is then isomorphic to $\mathbb{F}_q[t]/((t-\alpha)^k)$ and since $C_s^2 \in \mathbb{F}_p[C^2]$ is a semisimple element, it may be identified with $\alpha \in \mathbb{F}_q^*$.

Therefore the action of $\mathbb{F}_p[C_s^2]$ on V gives an $\mathbb{F}_q$-linear structure and since X, Y, Y_s, Y_u all commute with C_s^2 they may be viewed as an $\mathbb{F}_q$-linear transformation. Finally, we have $Y_s^2 = C_s^2 = \alpha I$. $\square$

4.2 The proof of Theorem 1.5 in the invertible cases

Let $V := \mathbb{F}_p^m$ and $C \colon V \to V$ be an invertible $\mathbb{F}_p$-linear map such that we have $m_{C^2}(x) = f^k(x)$ for the minimal polynomial of C^2 with some irreducible polynomial $f(x) \in \mathbb{F}_p[x]$. By Lemma 4.2 we even have an $\mathbb{F}_q$-linear structure on V (with $q := p^{\deg f}$) such that both C and any solution X to the equation $CX = X^{-1}C$ are $\mathbb{F}_q$-linear. Further, C has the Jordan–Chevalley decomposition $C = C_u C_s = C_s C_u$ with C_u unipotent and C_s semisimple. At first assume $p \neq 2$ and let N be the nilpotent transformation $C_u - I$. Then V becomes an $\mathbb{F}_q[T]$-module where T acts via N, $Tv = Nv$. Such modules are isomorphic to the module

$$V_\lambda = \oplus_j \mathbb{F}_q[T]/(T^{n_j}), \tag{17}$$

for some partition $\lambda = [n_1, .., n_k]$, $n_1 + \cdots + n_k = n$ which is unique by the structure theorem of finitely generated modules over PIDs. To show the partition λ associated to N we will use the notation $N = N_\lambda$.

By Lemmas 4.1 and 4.2 it is enough to count elements in the set

$$\mathcal{R}_\alpha(\lambda) = \{Y \in GL_n(\mathbb{F}_q) \mid Y^2 = \alpha I, YN_\lambda = N_\lambda Y\}, \tag{18}$$

ie. we have $N^*(A, B; p) = n(C, C; p) = \#\mathcal{R}_\alpha(\lambda)$ where $\alpha \in \mathbb{F}_q^*$ denotes the unique eigenvalue of C^2 as an $\mathbb{F}_q$-linear.

We start with the case when $\left(\frac{x}{f(x)}\right) = 1$ when we have $\alpha = \beta^2$ for some $\beta \in \mathbb{F}_q^*$. We have the following

Lemma 4.3 *Let $\mathcal{S}(\lambda)$ denote set of pairs (U_+, U_-) of $\mathbb{F}_q[T]$-submodules $U_+, U_- \leq V$ such that $V = U_+ \oplus U_-$. The maps*

$$Y \mapsto (U_{Y,+}, U_{Y,-})$$
$$Y_U(u_+ + u_-) = \beta(u_+ - u_-) \leftarrow V = U_+ \oplus U_-$$

are inverse bijections between $\mathcal{R}_\alpha(\lambda)$ and $\mathcal{S}(\lambda)$. Here $u_+ \in U_+$, $u_- \in U_-$ and $U_{Y,+}$ (resp. $U_{Y,-}$) denotes the β-eigenspace (resp. $-\beta$-eigenspace) of Y.

Proof Since $p \neq 2$, $Y^2 = \alpha I$ implies Y is semisimple whence V is the direct sum of the two eigenspaces of Y. Moreover, these eigenspaces are N-invariant, ie. they are $\mathbb{F}_q[T]$-submodules. Conversely, given such a decomposition $V = U_+ \oplus U_-$, we have $Y_U^2 = \alpha I$ and Y_U commutes with N. □

By the theorem of elementary divisors, for any decomposition $V = U_+ \oplus U_-$ of the $\mathbb{F}_q[T]$-module V, λ is the sum of the multisets μ and ν where μ (resp. ν) is the partition of $\dim U_+$ (resp. of $\dim U_-$) corresponding to the restriction of N to U_+ (resp. to U_-). Therefore we may write $\mathcal{S}(\lambda)$ as the union of

$$\mathcal{S}(\lambda, \mu, \nu) = \left\{ (U_+, U_-) \mid U_+, U_- \leq_{\mathbb{F}_q[T]} V, U_+ \oplus U_- = V, U_+ \cong \bigoplus_{j \in \mu} \mathbb{F}_q[T]/(T^j), U_- \cong \bigoplus_{j \in \nu} \mathbb{F}_q[T]/(T^j) \right\}$$

where μ runs over the multisets included in λ and $\nu = \lambda - \mu$ is the difference.

Lemma 4.4 *For any decomposition $\lambda = \mu + \nu$ the centralizer $Z_{GL_{|\lambda|}(\mathbb{F}_q)}(N)$ of N acts transitively on the set $\mathcal{S}(\lambda, \mu, \nu)$.*

Proof Assume we are given two decompositions $U_+ \oplus U_- = V = U'_+ \oplus U'_-$ in $\mathcal{S}(\lambda, \mu, \nu)$. Then we have the isomorphisms $U_+ \cong \bigoplus_{j \in \mu} \mathbb{F}_q[T]/(T^j) \cong U'_+$ and $U_- \cong \bigoplus_{j \in \nu} \mathbb{F}_q[T]/(T^j) \cong U'_-$ of $\mathbb{F}_q[T]$-modules. Taking the direct sum of these two isomorphisms we obtain an automorphism $g \colon V = U_+ \oplus U_- \to U'_+ \oplus U'_- = V$. Being $\mathbb{F}_q[T]$-linear it means that g lies in $Z_{GL_{|\lambda|}(\mathbb{F}_q)}(N)$ when viewed as an $\mathbb{F}_q$-linear transformation. □

Proposition 4.5 *If $\alpha \in (\mathbb{F}_q^*)^2$ then*

$$n(C, C; p) = \#\mathcal{R}_\alpha(\lambda) = \sum_{\lambda = \mu + \nu} \frac{\#Z_{GL_{|\lambda|}(\mathbb{F}_q)}(N_\lambda)}{\#Z_{GL_{|\mu|}(\mathbb{F}_q)}(N_\mu) \#Z_{GL_{|\nu|}(\mathbb{F}_q)}(N_\nu)}.$$

Proof By Lemma 4.3 we obtain $\#\mathcal{R}_\alpha(\lambda) = \#\mathcal{S}(\lambda) = \sum_{\lambda = \mu + \nu} \#\mathcal{S}(\lambda, \mu, \nu)$. The statement follows from Lemma 4.4 noting that the stabilizer of a given decomposition $V = U_+ \oplus U_-$ in $\mathcal{S}(\lambda, \mu, \nu)$ equals

$$Z_{GL_{|\lambda|}(\mathbb{F}_q)}(N_\lambda) \cap (GL_{|\mu|}(\mathbb{F}_q) \times GL_{|\nu|}(\mathbb{F}_q)) = Z_{GL_{|\mu|}(\mathbb{F}_q)}(N_\mu) \times Z_{GL_{|\nu|}(\mathbb{F}_q)}(N_\nu).$$

□

Now assume $\left(\frac{x}{f(x)} \right) = -1$, so we have $\alpha = \beta^2$ for some $\beta \in \mathbb{F}_{q^2}^* \backslash \mathbb{F}_q^*$. Put σ for the nontrivial element in $\mathrm{Gal}(\mathbb{F}_{q^2}/\mathbb{F}_q)$. Then $\varphi = \sigma \otimes I \colon \mathbb{F}_{q^2} \otimes_{\mathbb{F}_q} V$ is a σ-semilinear map (ie. $\mathbb{F}_q[T]$-linear with $\varphi(\beta v) = \sigma(\beta)\varphi(v) = -\beta\varphi(v)$). Further, for a $\mathbb{F}_{q^2}[T]$-module homomorphism $f \colon \mathbb{F}_{q^2} \otimes_{\mathbb{F}_q} V \to \mathbb{F}_{q^2} \otimes_{\mathbb{F}_q} V$ there exists an $\mathbb{F}_q[T]$-module homomorphism $\widetilde{f} \colon V \to V$ with $f = 1 \otimes \widetilde{f}$ if and only if f commutes with φ, ie. $f \circ \varphi = \varphi \circ f$.

Lemma 4.6 *Assume $\alpha \notin (\mathbb{F}_q^*)^2$. Then the centralizer $Z_{GL_{|\lambda|}(\mathbb{F}_q)}(N_\lambda)$ acts transitively on the set $\mathcal{R}_\alpha(\lambda)$ (by conjugation).*

Proof Let $Y, Y' \in \mathcal{R}_\alpha(\lambda)$, ie. $Y, Y' \colon V \to V$ are both $\mathbb{F}_q[T]$-linear isomorphisms with $Y^2 = \alpha I = Y'^2$. Put $U \le \mathbb{F}_{q^2} \otimes_{\mathbb{F}_q} V$ (resp. $U' \le \mathbb{F}_{q^2} \otimes_{\mathbb{F}_q} V$) for the β-eigenspace of $1 \otimes Y$ (resp. of $1 \otimes Y'$). Then $\varphi(U)$ (resp. $\varphi(U')$) is the $\sigma(\beta) = -\beta$-eigenspace of Y (resp. of Y'). In particular, we have $U \oplus \varphi(U) = \mathbb{F}_{q^2} \otimes_{\mathbb{F}_q} V = U' \oplus \varphi(U')$. Now if $U \cong \bigoplus_j \mathbb{F}_{q^2}[T]/(T^{m_j})$ for some partition $\mu = [m_1, \ldots, m_s]$ then we have the isomorphism

$$\varphi(U) \cong \bigoplus_j \mathbb{F}_{q^2}[T]/(\sigma(T)^{m_j}) \cong \bigoplus_j \mathbb{F}_{q^2}[T]/(T^{m_j}) \cong U$$

of $\mathbb{F}_{q^2}[T]$-modules whence $\lambda = \mu + \mu$. Similarly, $U' \cong \varphi(U')$. By the structure theorem for finitely generated modules over the PID $\mathbb{F}_{q^2}[T]$, we must have $U \cong U'$ as $\mathbb{F}_{q^2}[T]$-modules, as well. Taking such an isomorphism $S \colon U \to U'$ we also define $S(\varphi(u)) := \varphi(S(u))$ on $\varphi(U)$ giving rise to an $\mathbb{F}_{q^2}[T]$-linear automorphism $S \colon \mathbb{F}_{q^2} \otimes_{\mathbb{F}_q} V = U \oplus \varphi(U) \to U' \oplus \varphi(U') = \mathbb{F}_{q^2} \otimes_{\mathbb{F}_q} V$ that satisfies $S(1 \otimes Y)S^{-1} = 1 \otimes Y'$. Moreover, S descends to a map $\widetilde{S} \colon V \to V$ (such that $S = 1 \otimes \widetilde{S}$) since it commutes with φ. Finally, $\widetilde{S}$ satisfies $\widetilde{S}Y\widetilde{S}^{-1} = Y'$ and lies in the centralizer of N_λ as it is $\mathbb{F}_q[T]$-linear. $\qquad\square$

Proposition 4.7 *Assume $\alpha \notin (\mathbb{F}_q^*)^2$. Then we have $\lambda = \mu + \mu$ for some partition μ and*

$$n(C, C; p) = \#\mathcal{R}_\alpha(\lambda) = \frac{\#Z_{GL_{|\lambda|}(\mathbb{F}_q)}(N_\lambda)}{\#Z_{GL_{|\mu|}(\mathbb{F}_{q^2})}(N_\mu)}.$$

Proof By Lemma 4.6 $\mathcal{R}_\alpha(\lambda)$ is the conjugacy class of C_s in $Z_{GL_{|\lambda|}(\mathbb{F}_q)}(N_\lambda)$. A moments thought shows that we may define an $\mathbb{F}_{q^2}$-linear structure on V where C_s acts via multiplication by β where the $\mathbb{F}_{q^2}$-linear maps are exactly those which are $\mathbb{F}_q$-linear and commute with C_β. In particular, the centralizer of C_β in $Z_{GL_{|\lambda|}(\mathbb{F}_q)}(N_\lambda)$ equals $Z_{GL_{|\mu|}(\mathbb{F}_{q^2})}(N_\mu)$. $\qquad\square$

This leads to formula (5).

Corollary 4.8 *Assume that C has minimal polynomial $f(x)^k$ where $f(x) \ne x$ is irreducible and $p \ne 2$. Then we have $n(C, C, p) < \frac{q^2+1}{q^2-1} q^{\lfloor \frac{n^2}{2 \deg^2 f} \rfloor} = \frac{q^2+1}{q^2-1} p^{\deg f \lfloor \frac{n^2}{2 \deg^2 f} \rfloor} \le \frac{p^2+1}{p^2-1} p^{\frac{n^2}{2}}$ and if $k = 1$ then there exists a constant $0 < c(q) < 1$ (with $\lim_{q \to \infty} c(q) = 1$) such that*

$$c(q)q^{\lfloor \frac{n^2}{2 \deg^2 f} \rfloor} < n(C, C; p) < \frac{q^2+1}{q^2-1} q^{\lfloor \frac{n^2}{2 \deg^2 f} \rfloor}.$$

In particular

$$n(C, C; p) \le 2p^{n^2/2 \deg f}.$$

Proof Since $n(C, C, p)$ is the number of square roots of C_s^2 commuting with C_u, the case $C_u = I$ gives an upper bound for the number of solutions in general. So we may assume $N = 0$. Put $n_1 := \dim_{\mathbb{F}_q} V = \frac{n}{\deg f}$. Then in the split case we compute

$$n(C, C; p) = \sum_{j=0}^{n_1} \frac{\#GL_{n_1}(\mathbb{F}_q)}{\#GL_j(\mathbb{F}_q)\#GL_{n_1-j}(\mathbb{F}_q)}$$

$$= \sum_{j=0}^{n_1} \frac{(q^{n_1} - 1) \ldots (q^{n_1} - q^{n_1-1})}{(q^j - 1) \ldots (q^j - q^{j-1})(q^{n_1-j} - 1) \ldots (q^{n_1-j} - q^{n_1-j-1})}$$

$$= \sum_{j=0}^{n_1} q^{n_1^2 - j^2 - (n_1-j)^2} \frac{(1 - 1/q^{n_1}) \dots (1 - 1/q)}{(1 - 1/q^j) \dots (1 - 1/q)(1 - 1/q^{n_1-j}) \dots (1 - 1/q)}$$

$$< \sum_{j=0}^{n_1} q^{n_1^2 - j^2 - (n_1-j)^2} < q^{\lfloor \frac{n_1^2}{2} \rfloor} (1 + \sum_{j=1}^{\infty} \frac{2}{q^{2j}}) = \frac{q^2 + 1}{q^2 - 1} q^{\lfloor \frac{n_1^2}{2} \rfloor}.$$

On the other hand, we have

$$n(C, C; p) = \sum_{j=0}^{n_1} q^{n_1^2 - j^2 - (n_1-j)^2} \frac{(1 - 1/q^{n_1}) \dots (1 - 1/q)}{(1 - 1/q^j) \dots (1 - 1/q)(1 - 1/q^{n_1-j}) \dots (1 - 1/q)}$$

$$> q^{n_1^2 - \lfloor \frac{n_1}{2} \rfloor^2 - \lceil \frac{n_1}{2} \rceil^2} (1 - 1/q^{n_1}) \dots (1 - 1/q) > c(q) q^{\lceil \frac{n_1^2}{2} \rceil}$$

with constant $c(q) = \prod_{j=1}^{\infty} (1 - 1/q^j)$ that clearly satisfies $\lim_{q \to \infty} c(q) = 1$. $\qquad \square$

Finally, assume $p = 2$. Since the 2-Frobenius is bijective on finite fields of characteristic 2, C_s^2 has a unique square root $Y_s = C_s$. So we need to count the square roots of the unipotent matrix C_u^2 or equivalently the square roots of the nilpotent matrix $C_u^2 + I = (C_u + I)^2$.

Lemma 4.9 *Assume that q is a power of* 2. *For any integer $n > 0$ we have the identification*

$$\mathbb{F}_q[T]/(T^n) \cong \mathbb{F}_q[T^2]/((T^2)^{\lfloor \frac{n}{2} \rfloor}) \oplus \mathbb{F}_q[T^2]/((T^2)^{\lceil \frac{n}{2} \rceil})$$

as $\mathbb{F}_q[T^2]$-modules.

Proof This amounts to the fact that the square of a nilpotent Jordan block of size n splits into two blocks of size $\lfloor \frac{n}{2} \rfloor$ and $\lceil \frac{n}{2} \rceil$. $\qquad \square$

Proposition 4.10 *Assume q is a power of* 2. *Then the number of solutions of the matrix equation $Y^2 = C^2$ equals*

$$\sum_{\mu} \frac{\#Z_{GL_n(\mathbb{F}_q)}(N_\lambda^2)}{\#Z_{GL_n(\mathbb{F}_q)}(N_\mu)}$$

where $\mu = [m_1, \dots, m_k]$ runs on the set of partitions such that

$$\left[\lfloor \frac{m_1}{2} \rfloor, \lceil \frac{m_1}{2} \rceil, \dots, \lfloor \frac{m_k}{2} \rfloor, \lceil \frac{m_k}{2} \rceil \right] = \left[\lfloor \frac{n_1}{2} \rfloor, \lceil \frac{n_1}{2} \rceil, \dots, \lfloor \frac{n_k}{2} \rfloor, \lceil \frac{n_k}{2} \rceil \right].$$

Proof By Lemma 4.9 N_μ^2 is similar to N_λ^2 if and only if

$$\left[\lfloor \frac{m_1}{2} \rfloor, \lceil \frac{m_1}{2} \rceil, \dots, \lfloor \frac{m_k}{2} \rfloor, \lceil \frac{m_k}{2} \rceil \right] = \left[\lfloor \frac{n_1}{2} \rfloor, \lceil \frac{n_1}{2} \rceil, \dots, \lfloor \frac{n_k}{2} \rfloor, \lceil \frac{n_k}{2} \rceil \right].$$

So for each such μ we are reduced to determine the cardinality of the fiber at N_λ^2 of the map

$$\{\text{conjugacy class of } N_\mu\} \to \{\text{conjugacy class of } N_\lambda^2\}$$
$$M \mapsto M^2.$$

However, all the fibers of the above map have the same cardinality by conjugation, so the number of solutions to the above equation is

$$\frac{\#\{\text{conjugacy class of } N_\mu\}}{\#\{\text{conjugacy class of } N_\lambda^2\}} = \frac{\#Z_{GL_n(\mathbb{F}_q)}(N_\lambda^2)}{\#Z_{GL_n(\mathbb{F}_q)}(N_\mu)}.$$

$\qquad \square$

4.3 The proof of Theorem 1.5 in the nilpotent case

Lemma 4.11 *Assume $XCX = C$ for some $C \in M_n(\mathbb{F}_p)$ and $X \in \mathrm{GL}_n(\mathbb{F}_p)$. Then for any integer $j \geq 1$ the subspaces $\ker(C^j)$ and $\mathrm{Im}(C^j)$ are X-invariant.*

Proof Since X is invertible we may write $CX = X^{-1}C$ and $XC = CX^{-1}$, so by induction on j we deduce $C^j X = X^{(-1)^j} C^j$ and $XC^j = C^j X^{(-1)^j}$. Therefore if $v \in \ker(C^j)$ then we have $C^j Xv = X^{(-1)^j} C^j v = 0$, ie. $Xv \in \ker(C^j)$. On the other hand, we compute $XC^j w = C^j X^{(-1)^j} w \in \mathrm{Im}(C^j)$ for any $w \in \mathbb{F}_p^n$. $\square$

Proposition 4.12 *Let $C \in M_n(\mathbb{F}_p)$ be a nilpotent matrix such that there are r_j blocks of size $j \times j$ ($j = 1, \ldots, k$) in the Jordan normal form of C. Then the number of solutions of the equation $XCX = C$ in $X \in \mathrm{GL}_n(\mathbb{F}_p)$ equals*

$$\left(\prod_{j=1}^{k} (p^{r_j} - 1)(p^{r_j} - p) \ldots (p^{r_j} - p^{r_j - 1}) \right) \cdot p^{\sum_{j=1}^{k} (d_j^2 - r_j^2)} = p^{\sum_{j=1}^{k} d_j^2} \prod_{j=1}^{k} \prod_{t_j=1}^{r_j} \left(1 - \frac{1}{p^{t_j}} \right)$$

where we put $d_i = r_i + r_{i+1} + \cdots + r_k = \dim(\ker(C) \cap \mathrm{Im}(C^{i-1})) = \dim(\ker(C^i)) - \dim(\ker(C^{i-1}))$.

Proof First of all, we have $C^k = 0$ and $\dim_{\mathbb{F}_p}(\ker(C) \cap \mathrm{Im}(C^i)) = \sum_{j=i+1}^{k} r_j$. By lemma 4.11 the flag $0 \leq \ker(C) \cap \mathrm{Im}(C^{k-1}) \leq \cdots \leq \ker(C) \cap \mathrm{Im}(C^i) \leq \cdots \leq \ker(C) \cap \mathrm{Im}(C) \leq \ker(C)$ in $\ker(C)$ must be X-invariant for any solution X of the equation $XCX = C$. The set of such maps $X_1 := X_{|\ker(C)} \colon \ker(C) \to \ker(C)$ is the parabolic subgroup $P_{(r_1, \ldots, r_k)}$ of $\mathrm{GL}_{r_1 + \cdots + r_k}(\mathbb{F}_p)$ of type $(r_1, \ldots, r_k)$ which has cardinality

$$\#P_{(r_1, \ldots, r_k)} = \left(\prod_{j=1}^{k} (p^{r_j} - 1)(p^{r_j} - p) \ldots (p^{r_j} - p^{r_j - 1}) \right) \cdot p^{\sum_{1 \leq i < j \leq k} r_i r_j}.$$

Lemma 4.13 *For any $2 \leq t \leq k$ and $X_1 \in P_{(r_1, \ldots, r_k)}$ the number of extensions of X_1 to a one-to-one linear map $X_t \colon \ker(C^t) \to \ker(C^t)$ satisfying*

(i) $XCX = C$ *and*

(ii)t $\ker(C^i) \cap \mathrm{Im}(C^j)$ *is X_t-invariant for all $1 \leq i \leq t$ and $0 \leq j \leq k$*

equals

$$\prod_{2 \leq j \leq t} p^{r_j(r_1 + \cdots + r_k) + r_{j+1}(r_2 + \cdots + r_k) + \cdots + r_k(r_{k-j+1} + \cdots + r_k)}.$$

Proof We proceed by induction on t. Assume we have a map $X_{t-1} \colon \ker(C^{t-1}) \to \ker(C^{t-1})$ satisfying (i) and $(ii)_{t-1}$ and pick an element $v \in \ker(C^t) \cap \mathrm{Im}(C^j)$ not lying in $\ker(C^{t-1}) + \mathrm{Im}(C^{j+1})$. We need to choose $X_t v \in \ker(C^t) \cap \mathrm{Im}(C^j)$ so that $CX_t v = X_{t-1}^{-1} Cv$ is satisfied since $Cv \in \ker(C^{t-1}) \cap \mathrm{Im}(C^{j+1})$ on which subspace the map X_t^{-1} is already defined by X_{t-1}^{-1} (as X_{t-1} is one-to-one). Moreover, $X_{t-1}^{-1} Cv$ lies in $\ker(C^{t-1}) \cap \mathrm{Im}(C^{j+1})$ by assumption $(ii)_{t-1}$. In particular, there exists a vector $w \in \mathrm{Im}(C^j)$ such that $Cw = X_{t-1}^{-1} Cv$ and $w \in \ker(C^t)$ (as we have $C^t w = C^{t-1} X_{t-1}^{-1} Cv = 0$). Further, w is unique upto $\ker(C) \cap \mathrm{Im}(C^j)$, so the possible values of $X_t v$ is exactly $w + (\ker(C) \cap \mathrm{Im}(C^j))$ which has cardinality $\#(\ker(C) \cap \mathrm{Im}(C^j)) = p^{r_{j+1} + \cdots + r_k}$. Finally, we let v run on the lift of a basis of the quotient

space $(\ker(C^t) \cap \mathrm{Im}(C^j))/(\ker(C^{t-1}) + \mathrm{Im}(C^{j+1}))$ for any $j = k - t, k - t - 1, \ldots, 1, 0$ (noting $\mathrm{Im}(C^{k-t+1}) \subseteq \ker(C^{t-1})$) we deduce that the number of extensions of X_{t-1} to a map $X_t \colon \ker(C^t) \to \ker(C^t)$ satisfying (i) and $(ii)_t$ is

$$\prod_{j=0}^{k-t} \#(\ker(C) \cap \mathrm{Im}(C^j))^{\dim_{\mathbb{F}_p}(\ker(C^t) \cap \mathrm{Im}(C^j))/(\ker(C^{t-1}) + \mathrm{Im}(C^{j+1}))}$$

$$= p^{r_t(r_1 + \cdots + r_k) + r_{t+1}(r_2 + \cdots + r_k) + \cdots + r_k(r_{k-t+1} + \cdots + r_k)}$$

as we have $\dim_{\mathbb{F}_p}(\ker(C^t) \cap \mathrm{Im}(C^j))/(\ker(C^{t-1}) + \mathrm{Im}(C^{j+1})) = r_{j+t}$. $\qquad\square$

The statement follows from the above lemma by taking $t = k$: the number of solutions of $XCX = C$ in invertible X equals

$$\#P_{(r_1, \ldots, r_k)} \prod_{2 \le j \le k} p^{r_j(r_1 + \cdots + r_k) + r_{j+1}(r_2 + \cdots + r_k) + \cdots + r_k(r_{k-j+1} + \cdots + r_k)}$$

$$= \left(\prod_{j=1}^{k} (p^{r_j} - 1)(p^{r_j} - p) \ldots (p^{r_j} - p^{r_j - 1}) \right) \cdot p^{\sum_{j=2}^{k}(j-1)r_j^2 + \sum_{1 \le i < j \le k} 2 i r_i r_j}$$

$$= \left(\prod_{j=1}^{k} (p^{r_j} - 1)(p^{r_j} - p) \ldots (p^{r_j} - p^{r_j - 1}) \right) \cdot p^{\sum_{j=1}^{k}(d_j^2 - r_j^2)}$$

as claimed. $\qquad\square$

Corollary 4.14 *Assume C is nilpotent. Then we have* $n(C, C; p) \le p^{\sum_{j=1}^{k} d_j^2} \le p^{\mathrm{rk}(C)^2 + (n - \mathrm{rk}(C))^2}$.

Proof Using Proposition 4.12 we compute

$$n(C, C; p) = \left(\prod_{j=1}^{k} (p^{r_j} - 1)(p^{r_j} - p) \ldots (p^{r_j} - p^{r_j - 1}) \right) \cdot p^{\sum_{j=1}^{k}(d_j^2 - r_j^2)}$$

$$\le \left(\prod_{j=1}^{k} (p^{r_j})^{r_j} \right) \cdot p^{\sum_{j=1}^{k}(d_j^2 - r_j^2)} = p^{\sum_{j=1}^{k} d_j^2} \le p^{d_1^2 + (n - d_1)^2} = p^{\mathrm{rk}(C)^2 + (n - \mathrm{rk}(C))^2}$$

by noting $d_1 = \dim \ker(C) = n - \mathrm{rk}(C)$ and $\sum_{j=1}^{k} d_j = n$. $\qquad\square$

Remark 4.15 For fixed n and $p \to \infty$ the above upper estimate $p^{\sum_{j=1}^{k} d_j^2}$ is in fact the order of magnitude of $n(C, C, p)$:

$$\frac{n(C, C; p)}{p^{\sum_{j=1}^{k} d_j^2}} = \prod_{j=1}^{k} \prod_{t_j=1}^{r_j} \left(1 - \frac{1}{p^{t_j}} \right) > \left(1 - \frac{1}{p} \right)^n .$$

5 The proofs of the bounds

5.1 The equation $AX \equiv X^{-1}B$ modulo prime powers

Assume that $A, B \in M_n(\mathbb{Z})$. We are interested in estimating the size of the affine variety $V_{A,B}(p^l)$ where

$$V_{A,B}(p^l) = \{X \in GL_n(\mathbb{Z}/p^l\mathbb{Z}) \mid AX = X^{-1}B\}. \tag{19}$$

We will collect elements of $V_{A,B}(p^{l+1})$ according to their image in $V_{A,B}(p^l)$. The final push down to $l = 1$ will play a special role and we let

$$V_{A,B}^{(C)}(p^l) = \{X \in V_{A,B}(p^l) \mid AX \equiv X^{-1}B \equiv C \bmod p\}. \tag{20}$$

Let now $X_0 \in V_{A,B}^{(C)}(p^l)$ be given. Then all $X \in V_{A,B}(p^{l+1})$ such that $X \equiv X_0 \bmod p^l$ may be written as $X = X_0(I + p^l Y)$, for some $Y \bmod p$. The goal is to bound Y for which (19) also holds mod p^{l+1}. This leads to

$$AX_0 Y + Y X_0^{-1} B \equiv (X_0^{-1}B - AX_0)/p^l \bmod p.$$

Since $X_0 \in V_{A,B}^{(C)}(p^l)$ we have that Y is a solution to the Sylvester equation

$$CY + YC \equiv (X_0^{-1}B - AX_0)/p^l \bmod p. \tag{21}$$

Note that the equation above might have no solution, or exactly as many solution as

$$CY = -YC \bmod p$$

for which we may apply Proposition 3.3 and its corollary. This gives

Lemma 5.1 *Let* $\mathrm{rk}\, C = r$, $\mathrm{rk}_\infty C = r_\infty$. *If* $l \geq 1$ *and* $p \neq 2$ *then*

$$\#V_{A,B}^{(C)}(p^{l+1}) \leq p^{(n-r)(n-r_\infty)+r_\infty^2/2}\#V_{A,B}^{(C)}(p^l).$$

Proof This is merely a restatement of Lemma 3.3 and (12). First note that

$$\dim \ker C = n - r \text{ and } \dim \ker C^n = n - r_\infty.$$

To simplify the contribution of the non-zero eigenvalues in (12) use that

$$2d_\infty(A - \lambda)d_\infty(A + \lambda) \leq \frac{\big(d_\infty(A - \lambda) + d_\infty(A + \lambda)\big)^2}{2}$$

and that for $a_1, ..., a_k$ positive integers, $a_1^2 + \cdots + a_k^2 \leq (a_1 + \cdots + a_k)^2$. $\square$

This estimate is wasteful since the solution set could be empty. However, this will suffice for us.

The proof of Theorem 1.6 When $l = 1$ the bound for $N^*(C, C; p)$ follows from Theorem 1.5 together with exact formulae in Proposition 3.1 which gave Corollaries 4.8 and 4.14.

To see this, note that we may decompose C (over the ground field $\mathbb{F}_p$) as a block matrix with one block invertible of size r_∞ and one block nilpotent of size $n - r_\infty$.

For the invertible part we have the upper bound $2^{r_\infty} p^{r_\infty^2/2}$ using Corollary 4.8 for each irreducible factor $\neq X$ of the minimal polynomial of C^2, noting that there at most r_∞ such factors, and applying $a_1^2 + \cdots + a_k^2 \leq (a_1 + \cdots + a_k)^2$ for positive integers $a_1, ..., a_k$.

The nilpotent block has rank $r - r_\infty$ and so by Corollary 4.14 we have the upper bound $p^{(n-r)^2 + (r-r_\infty)^2}$ from which

$$N^*(C, C; p) \leq 2^{r_\infty} p^{(n-r)^2 + (r-r_\infty)^2 + r_\infty^2/2}. \tag{22}$$

This proves the first statement in case $l = 1$. Whenever $l > 1$ we use Lemma 5.1 inductively to get that

$$N^*(C, C; p^l) \leq 2^{r_\infty} p^{e(l,n,r,r_\infty)}$$

where

$$e(l, n, r, r_\infty) = (n - r)^2 + (r - r_\infty)^2 + r_\infty^2/2 + (l - 1)\left((n-r)(n-r_\infty) + r_\infty^2/2\right).$$

In order to prove the second statement, note that unless $N^*(A, B; p^l) = 0$, we find a common value $C := AX_0 = X_0^{-1}B$ such that $N^*(A, B; p^l) = N^*(C, C; p^l)$ and put $r_\infty := r_\infty(C)$. Moreover, from Lemma 4.1 the value of r_∞ is the same for any of the C-s that arise. So we are bound to estimate $e(l, n, r, r_\infty)$.

To simplify the exponent assume first that $n/2 \leq r \leq n$. A calculation shows that the maximum of the function

$$(1 - x)^2 + (x - y)^2 + y^2/2 + (l - 1)\left((1 - x)(1 - y) + y^2/2\right)$$

on the domain

$$D = \{(x, y) \mid 1/2 \leq x \leq 1, \ 0 \leq y \leq x\}. \tag{23}$$

is $l/2$, proving the claim in this case.

For $r < n/2$ we use that $(n - r)(n - r_\infty) + r_\infty^2/2 \leq n(n - r)$, and so

$$e(l, n, r, r_\infty) \leq ln(n - r)$$

in view of $(n - r)^2 + r^2 - n(n - r) = r(2r - n) \leq 0$.

This establishes both bounds in (7).

Finally to prove the universal bound $N^*(A, B; p^l) \leq 2^n p^{l(n^2 - n)}$ note that it holds trivially for $l \geq 2$, since $n^2 - n \geq n^2/2$.

For $N^*(A, B; p)$ start with the bound in (22). Note that if $r = n$ then also $r_\infty = n$ and so it is enough to prove that, for $0 \leq r_\infty \leq r$, $1 \leq r \leq n - 1$, we have

$$(n - r)^2 + (r - r_\infty)^2 + r_\infty^2/2 \leq n^2 - n.$$

However since $0 \leq r_\infty \leq r$ we have

$$(r - r_\infty)^2 + r_\infty^2/2 \leq r^2$$

and for $1 \leq r \leq n - 1$

$$(n - r)^2 + r^2 \leq n^2 - 2n + 2.$$

Finally $n^2 - 2n + 2 \leq n^2 - n$ holds since $n \geq 2$. $\qquad\square$

5.2 Gauss sums of matrices

There are various ways exponential sums with quadratic functions of the entries of an $n \times n$ matrix arise. For example in the theory of Siegel modular forms $Q(X) = \text{Tr } X^t A X$, and the associated Gauss sums play an important role see e.g [18]. These have a very different flavor

than ours, as the tensor properties allow one to diagonalize A, which immediately yields a diagonalization of the quadratic form $Q(x_{11}, x_{12}, \ldots, x_{nn})$. This approach is not directly applicable to our situation since we have $Q(X) = \text{Tr}\, T X^2$ for some matrix T. While this case appeared in the literature, see e.g. [7] our treatment is based directly on Proposition 3.5 and its corollary 3.6.

The proof of Proposition 1.8 We have to estimate the sum

$$S_{A,B}(X; p) = \sum_{U \bmod p} \psi((SU + TU^2)/p),$$

where X is such that $AX \equiv X^{-1}B \bmod p^l$, and where $S = (AX - X^{-1}B)/p^l$ and $T = AX \bmod p$. This is clearly a general Gauss sum. To apply Corollary 3.6 let $B(U, Y) = Q(U + Y) - Q(U) - Q(Y) = \text{Tr}((TY + YT)U)$ be the associated bilinear form.

We have that either $S_{A,B}(X; p) = 0$ or there exists Y such that $\text{Tr}\, SU = B(U, Y)$ for some Y in which case

$$S_{A,B}(X; p) = \left(\frac{\det(Q_1)}{p}\right) e^{-2\pi i Q(Y)/p} p^{n^2 - R}$$

where R is the rank of the quadratic form $Q(X) = \text{Tr}(TX^2)$.

Whether there exists Y such that $\text{Tr}\, SU = B(U, Y)$ for all U is again determined by the solubility of a Sylvester equation

$$TY + YT = S.$$

Moreover it implies that the rank of Q is $R = n^2 - K$, where $K = \#\{Y : TY + YT \equiv 0 \bmod p\}$. This is estimated as in Lemma 5.1 using Lemma 3.3 which gives the claim.

5.3 Bounding $K_n(A, B; p^k)$

The proof of Theorem 1.9 The case of $k = 1$ was handled in [6].

When $k = 2l$ we have $|K_n(A, B; p^k)| \leq p^{ln^2} N^*(A, B; p^l)$ by Proposition 1.1 noting that any mod p^l solution of the matrix equation $AX \equiv X^{-1}B \pmod{p^l}$ has exactly p^{ln^2} lifts to $GL_n(\mathbb{Z}/p^k\mathbb{Z})$.

Similarly, in case $k = 2l + 1$ ($l \geq 1$) we deduce $|K_n(A, B; p^k)| \leq p^{ln^2} N^*(A, B; p^l)$ $\max_X |S_{A,B}(X; p)|$ from Proposition 1.1. The estimate for the Gauss sum $S_{A,B}(X; p)$ is given in Proposition 1.8 while the estimate for $N^*(A, B; p^l)$ is in Theorem 1.6 both in the general case and under the stronger assumption $\gcd(\det A, \det B, p) = 1$.

References

1. Berndt, B.C., Williams, K.S., Evans, R.J.: Gauss and Jacobi sums. Wiley (1998)
2. Cochrane, T., Zheng, Z.: A survey on pure and mixed exponential sums modulo prime powers. In: Surveys in number theory. CRC Press (2009)
3. Dabrowski, R., Fisher, B.: A stationary phase formula for exponential sums over $\mathbb{Z}/p^m\mathbb{Z}$ and applications to $GL(3)$-Kloosterman sums. Acta Arith. **80**(1), 1–48 (1997)
4. Duan, G.-R.: Generalized Sylvester equations: unified parametric solutions. CRC Press (2019)
5. El Baz, D., Min, L., Andreas, S.: Effective equidistribution of rational points on expanding horospheres. arXiv:2212.07408
6. Erdélyi, M., Árpád, T.: "Matrix Kloosterman sums", to appear in Algebra and Number Theory. arXiv:2109.00762
7. Kuroda, M.: Quadratic Gauss sums on matrices. Linear Algebra Appl. **384**, 187–198 (2004)
8. Milnor, J.W., Dale, H.: Symmetric bilinear forms, vol. 73. Springer, Berlin (1973)
9. Lee, M., Marklof, J.: Effective equidistribution of rational points on expanding horospheres. Int. Math. Res. Not. **2018**(21), 6581–6610 (2018)
10. Morrison, K.E.: Integer sequences and matrices over finite fields. J. Integer Seq. **9**(2), 3 (2006)
11. Penrose, R.: A generalized inverse for matrices. In: Mathematical proceedings of the Cambridge philosophical society **51**(3), Cambridge University Press (1955)
12. Potter, J.E.: Matrix quadratic solutions. SIAM J. Appl. Math. **14**(3), 496–501 (1966)
13. Rosen, M.: Number theory in function fields, **210**. Springer (2002)
14. Salié, H.: Über die Kloostermanschen Summen $S(u, v; q)$. Mathematische Zeitschrift **34**(1), 91–109 (1932)
15. Stanley, R.P.: Enumerative Combinatorics Volume 1 second edition. Cambridge Studies in Advanced Mathematics (2011)
16. Springer, T.A., Robert, S.: "Conjugacy classes" in: Seminar on Algebraic Groups and Related Finite Groups (The Institute for Advanced Study, Princeton, N.J., 1968/69), pages 167–266. In: Lecture Notes in Mathematics, Vol. 131. Springer, Berlin (1970)
17. Sylvester, J.J.: Sur l'équation en matrices px= xq. CR Acad. Sci. Paris **99**(2), 67–71 (1884)
18. Walling, L.H.: Explicitly realizing average Siegel theta series as linear combinations of Eisenstein series. Ramanujan J **47**(3), 475–499 (2018)