On matrices commuting with their Frobenius

Fabian Gundlach* Béranger Seguin*

Abstract. The *Frobenius* of a matrix M with coefficients in $\overline{\mathbb{F}}_p$ is the matrix $\sigma(M)$ obtained by raising each coefficient to the p-th power. We consider the question of counting matrices with coefficients in $\mathbb{F}_q$ which commute with their Frobenius, asymptotically when q is a large power of p. We give answers for matrices of size 2, for diagonalizable matrices, and for matrices whose eigenspaces are defined over $\mathbb{F}_p$. Moreover, we explain what is needed to solve the case of general matrices. We also solve (for both diagonalizable and general matrices) the corresponding problem when one counts matrices M commuting with all the matrices $\sigma(M)$, $\sigma^2(M)$, ... in their Frobenius orbit.

MSC 2020: 14G17 · 15A27 · 14M15

Contents

1. Introduction

Throughout the paper, we fix a prime power p and an integer $n \geq 2$. For any field K, we denote by $\mathfrak{M}_n(K)$ the ring of $n \times n$-matrices with coefficients in K and by $\mathfrak{M}_n^{\mathrm{diag}}(K)$ the subset of matrices that are diagonalizable over the algebraic closure $\overline{K}$. We denote by σ the Frobenius automorphism of the $\mathbb{F}_p$-algebra $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ acting entrywise by $x \mapsto x^p$. The symbol q always denotes a power of p.

1.1. Main results

Consider the following four subsets of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$:

$$\mathfrak{X} = \left\{ M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p) \;\middle|\; M \text{ and } \sigma(M) \text{ commute} \right\}, \qquad \mathfrak{X}^{\mathrm{diag}} = \mathfrak{X} \cap \mathfrak{M}_n^{\mathrm{diag}}(\overline{\mathbb{F}}_p),$$

$$\mathfrak{X}_\infty = \left\{ M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p) \;\middle|\; M, \sigma(M), \sigma^2(M), \dots \text{ commute pairwise} \right\}, \qquad \mathfrak{X}_\infty^{\mathrm{diag}} = \mathfrak{X}_\infty \cap \mathfrak{M}_n^{\mathrm{diag}}(\overline{\mathbb{F}}_p).$$

In this paper, we estimate the asymptotic sizes of the intersections of these sets with $\mathfrak{M}_n(\mathbb{F}_q)$ as $q \to \infty$ (p and n are fixed, and q is a power of p). Letting $\mathbb{F}_q$ be any finite field containing $\mathbb{F}_p$, our main results are the following three theorems (the implied constants in the O-estimates are all independent of q):

*Universität Paderborn, Fakultät EIM, Institut für Mathematik, Warburger Str. 100, 33098 Paderborn, Germany. Emails: fabian.gundlach@uni-paderborn.de, math@beranger-seguin.fr.

Theorem 1.1 (cf. Theorem 3.5)**.** *We have* $|\mathfrak{X}_\infty^{\mathrm{diag}} \cap \mathfrak{M}_n(\mathbb{F}_q)| = p^{n^2-n} \cdot q^n + O_{p,n}(q^{n-1})$.

Theorem 1.2 (cf. Theorem 3.7, and Corollary 2.2 for the case $n = 2$)**.** *We have*

$$|\mathfrak{X}_\infty \cap \mathfrak{M}_n(\mathbb{F}_q)| = c_\infty(p,n) \cdot q^{\lfloor n^2/4\rfloor+1} + O_{p,n}\Big(q^{\lfloor n^2/4\rfloor}\Big),$$

where

$$c_\infty(p,2) = p^2+p+1, \qquad c_\infty(p,3) = p^6+p^5+3p^4+3p^3+3p^2+p+1,$$
$$c_\infty(p,n) = \binom{n}{n/2}_p \text{ if } n \geq 4 \text{ is even}, \qquad c_\infty(p,n) = 2\binom{n}{\lfloor n/2\rfloor}_p \text{ if } n \geq 5 \text{ is odd}.$$

(The Gaussian binomial coefficient $\binom{n}{k}_p$ *is the number of* k*-dimensional subspaces of* $\mathbb{F}_p^n$*.)*

Theorem 1.3 (cf. Theorem 4.17)**.** *We have*

$$|\mathfrak{X}^{\mathrm{diag}} \cap \mathfrak{M}_n(\mathbb{F}_q)| = c^{\mathrm{diag}}(p,n) \cdot q^{\lfloor n^2/3\rfloor+1} + O_{p,n}\Big(q^{\lfloor n^2/3\rfloor+1/2}\Big), \text{ where } c^{\mathrm{diag}}(p,n) = \begin{cases} p^2 & \text{if } n=2, \\ 2 & \text{if } n=4, \\ 1 & \text{if } n \notin \{2,4\}. \end{cases}$$

Lastly, we relate the exponent of q in the asymptotics of $|\mathfrak{X}\cap\mathfrak{M}_n(\mathbb{F}_q)|$ as $q \to \infty$ to the dimensions of intersections $\operatorname{Cent} M \cap \operatorname{Cl} M$, where $\operatorname{Cent} M$ and $\operatorname{Cl} M$ respectively denote the centralizer and conjugacy class of a matrix $M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$. More precisely, define for any $M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$ the integer

$$d(M) := (\text{number of distinct eigenvalues of } M) + \dim\big(\operatorname{Cent} M \cap \operatorname{Cl} M\big). \tag{1.1}$$

We prove a general statement (Proposition 5.8), which implies the following:

Theorem 1.4. *For any finite field* $\mathbb{F}_q \supseteq \mathbb{F}_p$*, we have* $|\mathfrak{X} \cap \mathfrak{M}_n(\mathbb{F}_q)| = |\mathfrak{X}^{\mathrm{diag}} \cap \mathfrak{M}_n(\mathbb{F}_q)| + O_{p,n}(q^{a_{p,n}})$*, where* $a_{p,n}$ *is the maximum value of* $d(M)$ *over non-diagonalizable matrices* $M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p) \setminus \mathfrak{M}_n^{\mathrm{diag}}(\overline{\mathbb{F}}_p)$.

Unfortunately, we are unable to compute $d(M)$ in general. This is related to the hard problem of classifying pairs of commuting matrices up to simultaneous conjugation. In Section 6, we deal with a special case where that problem is solved, in order to illustrate how the principle behind Proposition 5.8 may be applied. Specifically, we prove the following theorem about the set $\mathfrak{X}^{\mathrm{eig.}/\mathbb{F}_p}$ of matrices $M \in \mathfrak{X}$ whose eigenspaces are all defined over $\mathbb{F}_p$:

Theorem 1.5 (cf. Theorem 6.9)**.** *For any finite field* $\mathbb{F}_q \supseteq \mathbb{F}_p$*, we have*

$$\left|\mathfrak{X}^{\mathrm{eig.}/\mathbb{F}_p} \cap \mathfrak{M}_n(\mathbb{F}_q)\right| = c^{\mathrm{eig.}/\mathbb{F}_p}(p,n) \cdot q^{\lfloor n^2/4\rfloor+1} + O_{p,n}(q^{\lfloor n^2/4\rfloor}),$$

for specific constants $c^{\mathrm{eig.}/\mathbb{F}_p}(p,n)$*, given in Theorem 6.9.*

1.2. Outline and strategy

In Section 2, we quickly deal with the special case $n = 2$.

In Section 3, we prove Theorems 1.1 and 1.2 (Theorems 3.5 and 3.7) about $\mathfrak{X}_\infty^{\mathrm{diag}}$ and $\mathfrak{X}_\infty$. In both cases, we observe (see Lemma 3.1) that for any matrix $M \in \mathfrak{X}_\infty$, its Frobenius orbit $\big(\sigma^i(M)\big)_{i\geq 0}$ generates a commutative algebra of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ defined over $\mathbb{F}_p$, and consisting of simultaneously diagonalizable matrices when moreover $M \in \mathfrak{X}_\infty^{\mathrm{diag}}$. Hence, the statements boil down to studying such subalgebras. More specifically, we prove the two following results:

Theorem 1.6 (cf. Lemma 3.2(a) and Theorem 3.4)**.** *There are exactly* p^{n^2-n} *commutative* n*-dimensional subalgebras of* $\mathfrak{M}_n(\mathbb{F}_p)$ *formed of diagonalizable matrices, and none of higher dimension.*

Theorem 1.7 (cf. Theorem 3.6)**.** *Let $n \geq 3$ and let $c_\infty(p,n)$ be as in Theorem 1.2. There are exactly $c_\infty(p,n)$ commutative $(\lfloor n^2/4 \rfloor + 1)$-dimensional subalgebras of $\mathfrak{M}_n(\mathbb{F}_p)$, and none of higher dimension.*

(Theorem 1.7/Theorem 3.6 is a consequence of [Sch05].)

In Section 4, we prove Theorem 1.3 about $\mathfrak{X}^{\mathrm{diag}}$. Using the Lang–Weil bound, the claim reduces to the computation of geometric invariants of the constructible subset $\mathfrak{X}^{\mathrm{diag}} \subseteq \mathfrak{M}_n(\overline{\mathbb{F}}_p)$, namely its dimension and the number of its irreducible components of maximal dimension. To determine the top-dimensional irreducible components of $\mathfrak{X}^{\mathrm{diag}}$, we stratify this set according to how the eigenspaces of an element M intersect the eigenspaces of its Frobenius conjugate $\sigma(M)$, using quivers to encode this combinatorial information.

In Section 5, we show Proposition 5.8 (and thus Theorem 1.4). To relate the dimension of $\mathfrak{X}$ to the numbers $d(M)$ defined above, we stratify $\mathfrak{M}_n(\mathbb{F}_q)$ according to the shape of the Jordan normal form of matrices (i.e., the number of Jordan blocks of each size for each eigenvalue).

In Section 6, we prove Theorem 6.9, which counts matrices in $\mathfrak{X} \cap \mathfrak{M}_n(\mathbb{F}_q)$ whose eigenspaces are defined over $\mathbb{F}_p$. This special case lets us illustrate the principle described in Section 5, and is made accessible by the fact that classifying pairs of commuting matrices *whose eigenspaces coincide* is relatively easy (cf. Proposition 6.4/Lemma 6.5).

1.3. Motivation and related results

Our initial contact with this problem came from the role played by analogous counts in the distribution of wildly ramified extensions of the local function field $\mathbb{F}_q(\!(T)\!)$ (see [GS25, Propositions 4.6 and 4.9]). In [GS25, Lemmas 6.3, 6.4, 6.5], we have obtained estimates for the number of matrices commuting with their Frobenius (as well as with the Frobenius of their Frobenius, etc.) in a specific group of invertible matrices, namely the Heisenberg group $H_k(\mathbb{F}_q)$, and this has let us describe the distribution of $H_k(\mathbb{F}_p)$-extensions of function fields. We were led to generalize that question to more general matrices, and to study it for itself, after realizing that it was a deep and non-trivial problem.

A different point of view is that we are counting the $(\mathbb{F}_q, \sigma)$-points of the difference scheme defined by the difference equation $M\sigma(M) = \sigma(M)M$ (for $\mathfrak{X}$ and $\mathfrak{X}^{\mathrm{diag}}$). This makes our problem fit into the general framework of Hrushovski–Lang–Weil estimates as studied in [SV22, HHYZ2424]. Through that lens, our results may be seen as estimating invariants of these difference schemes, notably the "transformal dimension" which seems related to the exponent of q in our asymptotics. Alternatively, one can define the variety of pairs of commuting matrices (an irreducible subvariety of $\mathbb{A}^{2n^2}_{\mathbb{F}_p}$ which is well-studied, see e.g. [MT55, Ger61a, Gur92, GS00]) and describe the geometry (dimension, irreducible components, ...) of its intersection with the graph of σ (also a subvariety of $\mathbb{A}^{2n^2}_{\mathbb{F}_p}$). Our results may be seen as contributing to this description.

Another inspiration for studying this question comes from previous results about counting specific kinds of matrices over $\mathbb{F}_q$, cf. [FH58, Ger61b] (for nilpotent matrices), [BGS14] (for symmetric nilpotent matrices), [Sch08] (for the distribution of characteristic polynomials), [FF60] (for pairs of commuting matrices), [Hua23] (for mutually annihilating pairs of matrices), etc.

1.4. Terminology and conventions

A linear subspace $V \subseteq \overline{\mathbb{F}}_p^n$ is *defined over* $\mathbb{F}_{p^r}$ if it is σ^r-invariant, i.e., $\sigma^r(V) = V$. By Galois descent for vector spaces, this is equivalent to the vector space having a basis consisting of vectors in $\mathbb{F}_{p^r}^n$, i.e., to the existence of an isomorphism $V \simeq V' \otimes_{\mathbb{F}_{p^r}} \overline{\mathbb{F}}_p$ for the $\mathbb{F}_{p^r}$-vector space $V' = V \cap \mathbb{F}_{p^r}^n$.

Varieties. In this paper, the word *variety* always refers to a (classical) quasi-projective variety over $\overline{\mathbb{F}}_p$, i.e., a (Zariski) locally closed subset of $\mathbb{P}^r(\overline{\mathbb{F}}_p)$ for some $r \geq 1$. We do *not* assume that varieties are irreducible. We say that a variety $V \subseteq \mathbb{P}^r(\overline{\mathbb{F}}_p)$ is *defined over* $\mathbb{F}_p$ if it is σ-invariant, i.e., $\sigma(V) = V$ where $\sigma\colon \mathbb{P}^r(\overline{\mathbb{F}}_p) \to \mathbb{P}^r(\overline{\mathbb{F}}_p)$ is induced by $\sigma\colon x \mapsto x^p$. The *dimension* of a constructible subset of $\mathbb{P}^r(\overline{\mathbb{F}}_p)$ is the (Krull) dimension of its Zariski closure. A regular map $f\colon X \to Y$ between smooth varieties is *étale* if for every $x \in X$, the differential $D_x f\colon T_x X \to T_{f(x)} Y$ of f at x is an isomorphism of $\overline{\mathbb{F}}_p$-vector spaces.[1]

1.5. Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) — Project-ID 491392403 — TRR 358 (Project A4).

2. The case of 2×2 matrices

We first quickly deal with the case $n = 2$, as it is particularly simple to obtain an exact count in this case, and the behavior is different compared to larger values of n.

Proposition 2.1. *Assume that $n = 2$, and let $M \in \mathfrak{M}_2(\overline{\mathbb{F}}_p)$. The following are equivalent:*

(i) M is of the form $\lambda M' + \mu I_2$ with $\lambda, \mu \in \overline{\mathbb{F}}_p$ and $M' \in \mathfrak{M}_2(\mathbb{F}_p)$.

(ii) $M \in \mathfrak{X}_\infty$.

(iii) $M \in \mathfrak{X}$.

Proof. Clearly, (i) $\Rightarrow$ (ii) $\Rightarrow$ (iii). Assume (iii). If M is a scalar matrix, (i) is clear. Otherwise, the condition that $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ and $\sigma(M) = \left(\begin{smallmatrix} \sigma(a) & \sigma(b) \\ \sigma(c) & \sigma(d) \end{smallmatrix}\right)$ commute rewrites as the following system of equations:

$$\begin{cases} a\sigma(a) + b\sigma(c) &= a\sigma(a) + c\sigma(b) \\ a\sigma(b) + b\sigma(d) &= b\sigma(a) + d\sigma(b) \\ c\sigma(a) + d\sigma(c) &= a\sigma(c) + c\sigma(d) \\ c\sigma(b) + d\sigma(d) &= b\sigma(c) + d\sigma(d) \end{cases} \iff \begin{cases} b\sigma(c) &= c\sigma(b) \\ b\sigma(d-a) &= (d-a)\sigma(b) \\ c\sigma(d-a) &= (d-a)\sigma(c), \end{cases}$$

meaning that the point $[b : c : d-a] \in \mathbb{P}^2(\overline{\mathbb{F}}_p)$ is σ-invariant, so belongs to $\mathbb{P}^2(\mathbb{F}_p)$. Writing $(b, c, d-a) = \lambda(\beta, \gamma, \delta)$ with $\beta, \gamma, \delta \in \mathbb{F}_p$ and $\lambda \in \overline{\mathbb{F}}_p^\times$, we have (i) with $\mu = a$ and $M' = \left(\begin{smallmatrix} 0 & \beta \\ \gamma & \delta \end{smallmatrix}\right)$. $\square$

Corollary 2.2. *Assume that $n = 2$, and let $\mathbb{F}_q$ be a finite field containing $\mathbb{F}_p$. Then, $|\mathfrak{X} \cap \mathfrak{M}_n(\mathbb{F}_q)| = |\mathfrak{X}_\infty \cap \mathfrak{M}_n(\mathbb{F}_q)| = q + (p^2 + p + 1)(q-1)q$.*

Proof. Using Proposition 2.1, the size of $\mathfrak{X} \cap \mathfrak{M}_n(\mathbb{F}_q) = \mathfrak{X}_\infty \cap \mathfrak{M}_n(\mathbb{F}_q)$ is given by

$$\underbrace{q}_{\text{scalar matrices}} + \underbrace{(p^2+p+1)}_{\text{choices of } [b:c:d-a] \in \mathbb{P}^2(\mathbb{F}_p)} \cdot \underbrace{(q-1)}_{\substack{\text{choices of } (b,c,d-a) \in \mathbb{F}_q^3 \setminus \{(0,0,0)\} \\ \text{once } [b:c:d-a] \text{ is fixed}}} \cdot \underbrace{q}_{\text{choices of } a} \qquad \square$$

3. Matrices commuting with their whole Frobenius orbit

In this section, we determine the asymptotics of $|\mathfrak{X}_\infty^{\mathrm{diag}} \cap \mathfrak{M}_n(\mathbb{F}_q)|$ and $|\mathfrak{X}_\infty \cap \mathfrak{M}_n(\mathbb{F}_q)|$, i.e., we prove Theorems 3.5 and 3.7 (which are Theorems 1.1 and 1.2). For any field K, we call a subalgebra A of $\mathfrak{M}_n(K)$ *diagonalizable* if its elements are simultaneously diagonalizable over $\overline{K}$. In particular, a diagonalizable subalgebra is commutative. The sets $\mathfrak{X}_\infty^{\mathrm{diag}}$ and $\mathfrak{X}_\infty$ can be decomposed using the (finitely many) diagonalizable (resp. commutative) subalgebras of $\mathfrak{M}_n(\mathbb{F}_p)$:

[1] By Hilbert's Nullstellensatz, varieties form a category equivalent to that of reduced quasi-projective schemes over $\overline{\mathbb{F}}_p$. A variety is defined over $\mathbb{F}_p$ if and only if the corresponding reduced $\overline{\mathbb{F}}_p$-subscheme of $\mathbb{P}^r_{\overline{\mathbb{F}}_p}$ is obtained via extension of scalars of a geometrically reduced $\mathbb{F}_p$-subscheme of $\mathbb{P}^r_{\mathbb{F}_p}$. A regular map between smooth varieties is étale if and only if the corresponding morphism of reduced smooth quasi-projective schemes is étale.

Lemma 3.1. *We have*

$$\mathfrak{X}_\infty^{\mathrm{diag}} = \bigcup_{\substack{A \subseteq \mathfrak{M}_n(\mathbb{F}_p) \\ \text{diagonalizable} \\ \text{subalgebra}}} A \otimes_{\mathbb{F}_p} \overline{\mathbb{F}}_p \qquad \text{and} \qquad \mathfrak{X}_\infty = \bigcup_{\substack{A \subseteq \mathfrak{M}_n(\mathbb{F}_p) \\ \text{commutative} \\ \text{subalgebra}}} A \otimes_{\mathbb{F}_p} \overline{\mathbb{F}}_p.$$

Proof. The inclusions $\supseteq$ are clear: if $M \in A \otimes_{\mathbb{F}_p} \overline{\mathbb{F}}_p$ for a commutative subalgebra A of $\mathfrak{M}_n(\mathbb{F}_p)$, then the matrices $\sigma^i(M)$ for $i = 0, 1, \ldots$ all belong to the commutative algebra $A \otimes_{\mathbb{F}_p} \overline{\mathbb{F}}_p$, hence commute with each other. If moreover A is diagonalizable, then so is M.

For the inclusions $\subseteq$, consider any matrix $M \in \mathfrak{X}_\infty$. Since the matrices $M, \sigma(M), \ldots$ commute, they generate a commutative $\overline{\mathbb{F}}_p$-subalgebra R of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$. This subalgebra is σ-invariant, so by Galois descent we have $R = A \otimes_{\mathbb{F}_p} \overline{\mathbb{F}}_p$ for some commutative subalgebra A of $\mathfrak{M}_n(\mathbb{F}_p)$, proving the second equality. If moreover $M \in \mathfrak{X}_\infty^{\mathrm{diag}}$, then the commuting matrices $M, \sigma(M), \ldots$ are diagonalizable, hence they are simultaneously diagonalizable. Any common eigenbasis of these matrices is in fact a common eigenbasis of all matrices in $R = A \otimes_{\mathbb{F}_p} \overline{\mathbb{F}}_p$, so $A \subseteq \mathfrak{M}_n(\mathbb{F}_p)$ is a diagonalizable subalgebra. □

As a consequence of Lemma 3.1, describing the asymptotic sizes of $\mathfrak{X}_\infty^{\mathrm{diag}} \cap \mathfrak{M}_n(\mathbb{F}_q)$ (resp. of $\mathfrak{X}_\infty \cap \mathfrak{M}_n(\mathbb{F}_q)$) boils down to determining the dimension and the number of the maximal-dimensional diagonalizable (resp. commutative) subalgebras of $\mathfrak{M}_n(\mathbb{F}_p)$. This is done in Subsection 3.1 and Subsection 3.2, respectively.

3.1. Diagonalizable matrices

Lemma 3.2.

(a) Every diagonalizable subalgebra of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ has dimension at most n.

(b) There is a bijection between the set of n-dimensional diagonalizable subalgebras A of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ and the set of unordered n-tuples $\{E_1, \ldots, E_n\}$ of one-dimensional subspaces of $\overline{\mathbb{F}}_p^n$ such that $E_1 \oplus \ldots \oplus E_n = \overline{\mathbb{F}}_p^n$.

(c) An n-dimensional diagonalizable subalgebra A is defined over $\mathbb{F}_p$ if and only if the corresponding tuple $\{E_1, \ldots, E_n\}$ is σ-invariant, i.e., if there is a permutation $\pi \in \mathfrak{S}_n$ such that $\sigma(E_i) = E_{\pi(i)}$.

Proof. Let A be any diagonalizable subalgebra of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$, and pick a common eigenbasis $\mathcal{B} = (e_1, \ldots, e_n)$ of the matrices in A, so that every matrix in A is diagonal when expressed in $\mathcal{B}$. We immediately obtain (a), and we see that if A is n-dimensional, then it consists of all matrices which are diagonal with respect to $\mathcal{B}$. In this case, $\mathcal{B}$ is unique up to permutation and rescaling, as the spaces $\langle e_i \rangle$ are exactly the one-dimensional subspaces which are invariant under all matrices in A. Thus, $A \mapsto \{\langle e_1 \rangle, \ldots, \langle e_n \rangle\}$ defines a bijection as in (b). For (c), note that if $e_1, \ldots, e_n$ is a common eigenbasis of A, then $\sigma(e_1), \ldots, \sigma(e_n)$ is a common eigenbasis of $\sigma(A)$. Combined with this, (b) implies that A is fixed by σ if and only if σ permutes the eigenspaces $\langle e_1 \rangle, \ldots, \langle e_n \rangle$. □

Let $c_\infty^{\mathrm{diag}}(p, n)$ be the number of n-dimensional diagonalizable subalgebras of $\mathfrak{M}_n(\mathbb{F}_p)$. Distinguishing between the possible permutations π, and using the fact that $\mathfrak{S}_n$ acts freely on ordered tuples of pairwise distinct spaces, Lemma 3.2 immediately implies:

$$c_\infty^{\mathrm{diag}}(p, n) = \frac{1}{n!} \sum_{\pi \in \mathfrak{S}_n} |N(\pi)| \tag{3.1}$$

where $N(\pi)$ is the set of **ordered** tuples $(E_1, \ldots, E_n)$ of one-dimensional subspaces of $\overline{\mathbb{F}}_p^n$ such that $E_1 \oplus \ldots \oplus E_n = \overline{\mathbb{F}}_p^n$ and $\sigma(E_i) = E_{\pi(i)}$ for all $i = 1, \ldots, n$.

Lemma 3.3. *For any permutation $\pi \in \mathfrak{S}_n$, we have*

$$|N(\pi)| = \frac{|\mathrm{GL}_n(\mathbb{F}_p)|}{\prod_{C \text{ cycle in } \pi}(p^{|C|} - 1)}.$$

Proof. We show that $\mathrm{GL}_n(\mathbb{F}_p)$ acts transitively on $N(\pi)$, with stabilizers isomorphic to $\prod_C \mathbb{F}_{p^{|C|}}^\times$. The claim will then immediately follow using the orbit-stabilizer theorem.

Let $C_1, \dots, C_r$ be the cycles of π. For any $(E_1, \dots, E_n) \in N(\pi)$ and any cycle C_k of π, consider the subspace $F_k := \bigoplus_{i \in C_k} E_i$. Since π permutes the elements of the cycle C_k, this subspace F_k is by definition of $N(\pi)$ fixed by σ and hence defined over $\mathbb{F}_p$. Moreover, $\bigoplus_k F_k = \bigoplus_i E_i = \overline{\mathbb{F}}_p^n$.

The group $\mathrm{GL}_n(\mathbb{F}_p)$ acts transitively on the set of tuples $(F_1, \dots, F_r)$ of subspaces of $\mathbb{F}_p^n$ such that $F_1 \oplus \dots \oplus F_r = \mathbb{F}_p^n$ and $\dim F_k = |C_k|$ for all k, and the stabilizers for that action are isomorphic to $\prod_k \mathrm{GL}(F_k)$. It is hence sufficient to prove, for fixed subspaces $F_1', \dots, F_r'$ of $\mathbb{F}_p^n$, that the action of $\prod_k \mathrm{GL}(F_k' \otimes \overline{\mathbb{F}}_p)$ on the set of tuples $(E_1, \dots, E_n)$ of one-dimensional subspaces of $\overline{\mathbb{F}}_p^n$ such that $\sigma(E_i) = E_{\pi(i)}$ and $\bigoplus_{i \in C_k} E_i = F_k' \otimes \overline{\mathbb{F}}_p$ is transitive, with stabilizers isomorphic to $\prod_k \mathbb{F}_{p^{|C_k|}}^\times$. As that action is "block-diagonal", we can restrict our attention to a single cycle. We now assume that $\pi = (1, \dots, n)$.

We will show that we then have a ($\mathrm{GL}_n(\mathbb{F}_p)$-equivariant) bijection

$$f \colon \{\mathbb{F}_p\text{-basis } (a_1, \dots, a_n) \text{ of } \mathbb{F}_{p^n}\}/\mathbb{F}_{p^n}^\times \xrightarrow{\sim} N(\pi)$$

sending $[(a_1, \dots, a_n)]$ to the tuple $(E_1, \dots, E_n)$ where $E_1 = \langle (a_1, \dots, a_n) \rangle$ and $E_i = \sigma^{i-1}(E_1)$ for $i = 2, \dots, n$. Since the group $\mathrm{GL}_n(\mathbb{F}_p)$ acts simply transitively on the set of $\mathbb{F}_p$-bases of $\mathbb{F}_{p^n}$, it will then indeed act transitively on $N(\pi)$ with stabilizer isomorphic to $\mathbb{F}_{p^n}^\times$.

It remains to show that the map f is well-defined and bijective. For any $(E_1, \dots, E_n) \in N(\pi)$, we have $E_i = \sigma^{i-1}(E_1)$ for $i = 2, \dots, n$ and $\sigma^n(E_1) = E_1$, so E_1 must be generated by a vector with coordinates in $\mathbb{F}_{p^n}$. Moreover, if we define $E_1 = \langle (a_1, \dots, a_n) \rangle$ and $E_i = \sigma^{i-1}(E_1)$ for $i = 2, \dots, n$, then $E_1, \dots, E_n$ span $\overline{\mathbb{F}}_p^n$ if and only if the matrix $\left(\sigma^{i-1}(a_j)\right)_{i,j}$ is invertible, which is equivalent to $a_1, \dots, a_n$ forming an $\mathbb{F}_p$-basis of $\mathbb{F}_{p^n}$.[2] $\square$

Theorem 3.4 (cf. Theorem 1.6)**.** *We have $c_\infty^{\mathrm{diag}}(p, n) = p^{n^2 - n}$.*

Proof. For any partition of n with n_ℓ parts of size ℓ, there are exactly $n!/\prod_{\ell \geq 1} \ell^{n_\ell} n_\ell!$ permutations with n_ℓ cycles of length ℓ (the centralizer of any such permutation is isomorphic to $\prod_\ell (\mathbb{Z}/\ell\mathbb{Z})^{n_\ell} \rtimes \mathfrak{S}_{n_\ell}$). Hence, Equation (3.1) and Lemma 3.3 imply

$$\frac{c_\infty^{\mathrm{diag}}(p, n)}{|\mathrm{GL}_n(\mathbb{F}_p)|} = \sum_{\substack{\text{partition of } n \\ \text{with } n_\ell \text{ parts of size } \ell}} \frac{1}{\prod_{\ell \geq 1} \ell^{n_\ell}\, n_\ell!\, (p^\ell - 1)^{n_\ell}}.$$

As sizes of parts of partitions of n are characterized by the property $\sum_\ell \ell n_\ell = n$ (where $n_\ell \geq 0$ for all ℓ, and $n_\ell = 0$ for almost all ℓ), the right-hand side is the coefficient in front of X^n of the power series

$$\sum_{\substack{n_1, n_2, \ldots \geq 0 \\ \text{almost all } 0}} \prod_{\ell \geq 1} \frac{X^{\ell n_\ell}}{\ell^{n_\ell}\, n_\ell!\, (p^\ell - 1)^{n_\ell}} = \prod_{\ell \geq 1} \sum_{n \geq 0} \frac{X^{\ell n}}{\ell^n\, n!\, (p^\ell - 1)^n} = \prod_{\ell \geq 1} \exp\left(\frac{X^\ell}{\ell(p^\ell - 1)}\right) = \prod_{\ell \geq 1} \exp\left(\frac{p^{-\ell} X^\ell}{\ell(1 - p^{-\ell})}\right)$$

[2] If $\left(\sigma^{i-1}(a_j)\right)_{i,j}$ is singular, then there is a non-trivial linear combination $\sum_j \lambda_j \sigma^{i-1}(a_j) = 0$ with coefficients in $\mathbb{F}_{p^n}$ between its columns, which amounts to $\sum_j \sigma^i(\lambda_j) a_j = 0$ for all $i \in \{0, \dots, n-1\}$, so the vector $(a_1, \dots, a_n) \in (\mathbb{F}_{p^n})^n$ is orthogonal to the subspace $\mathrm{Span}_i\left(\sigma^i(\lambda_1, \dots, \lambda_n)\right) \subseteq (\mathbb{F}_{p^n})^n$; that subspace is σ-invariant, hence admits an $\mathbb{F}_p$-basis, in particular it contains a non-zero vector in $\mathbb{F}_p^n$, which implies that there is a non-trivial linear combination $\sum_j \mu_j a_j = 0$ with coefficients in $\mathbb{F}_p$. Conversely, if $a_1, \dots, a_n$ are linearly dependent over $\mathbb{F}_p$, then up to the action of $\mathrm{GL}_n(\mathbb{F}_p)$, we can assume that $a_n = 0$ and then $\left(\sigma^{i-1}(a_j)\right)_{i,j}$ is singular as its last column vanishes.

$$= \exp\left(\sum_{\substack{\ell\geq 1\\ k\geq 0}} \frac{p^{-\ell}X^\ell}{\ell} p^{-\ell k}\right) = \exp\left(-\sum_{k\geq 0}\ln(1-p^{-(1+k)}X)\right) = \prod_{k\geq 1}\frac{1}{1-p^{-k}X} = \prod_{k\geq 1}\sum_{i\geq 0} p^{-ki}X^i$$

$$= \sum_{n\geq 0}\left(\sum_{\substack{i_1,i_2,\ldots\geq 0\\ i_1+i_2+\ldots=n}} p^{-\sum_{k\geq 1} k i_k}\right)X^n = \sum_{n\geq 0}\sum_{s\geq n}\left|\left\{i_1,i_2,\ldots\geq 0 \,\middle|\, \begin{matrix} i_1+i_2+\ldots=n\\ \sum_{k\geq 1} k i_k = s\end{matrix}\right\}\right|\cdot p^{-s}X^n.$$

On the other hand:

$$\sum_{n\geq 0}\frac{p^{n^2-n}}{|\mathrm{GL}_n(\mathbb{F}_p)|}X^n = \sum_{n\geq 0}\frac{p^{\frac{n(n-1)}{2}}}{(p^n-1)\cdots(p-1)}X^n = \sum_{n\geq 0}\left(\prod_{k=1}^n\frac{p^{k-1}}{p^k-1}\right)X^n = \sum_{n\geq 0}\frac{1}{p^n}\left(\prod_{k=1}^n\frac{1}{1-p^{-k}}\right)X^n$$

$$= \sum_{n\geq 0}\frac{1}{p^n}\left(\prod_{k=1}^n\sum_{i\geq 0}p^{-ki}\right)X^n = \sum_{n\geq 0}\frac{1}{p^n}\left(\sum_{i_1,\ldots,i_n\geq 0}\prod_{k=1}^n p^{-k i_k}\right)X^n = \sum_{n\geq 0}\left(\sum_{i_1,\ldots,i_n\geq 0} p^{-(\sum_{k=1}^n k i_k + n)}\right)X^n$$

$$= \sum_{n\geq 0}\sum_{s\geq n}\left|\left\{i_1,\ldots,i_n\geq 0 \,\middle|\, \sum_{k=1}^n k i_k = s-n\right\}\right|\cdot p^{-s}X^n.$$

Therefore, the claim reduces to the following equality for all $s\geq n$:

$$\left|\left\{i_1,i_2,\ldots\geq 0 \,\middle|\, \begin{matrix} i_1+i_2+\ldots=n\\ \sum_{k\geq 1} k i_k = s\end{matrix}\right\}\right| = \left|\left\{i_1,\ldots,i_n\geq 0 \,\middle|\, \sum_{k=1}^n k i_k = s-n\right\}\right|.$$

We can interpret a list $(i_1,i_2,\ldots)$ such that $i_1+i_2+\ldots=n$ and $\sum_{k\geq 1} k i_k = s$ as a partition of s with exactly n (non-zero) parts (i_k is the number of parts of size k). Similarly, we can interpret a tuple $(i_1,\ldots,i_n)$ such that $\sum_{k=1}^n k i_k = s-n$ as a partition of $s-n$ whose parts all have size $\leq n$.

Consider a partition of s with exactly n parts. Removing 1 from each part turns this partition into a partition of $s-n$ with at most n parts. Then, taking conjugate partitions turns that partition into a partition of $s-n$ whose parts all have sizes $\leq n$. As both of these operations can be inverted, we have described a bijection between the two sets, proving the claim. □

Theorem 3.5 (cf. Theorem 1.1)**.** *For any finite field $\mathbb{F}_q\supseteq\mathbb{F}_p$, we have*

$$|\mathfrak{X}_\infty^{\mathrm{diag}}\cap\mathfrak{M}_n(\mathbb{F}_q)| = p^{n^2-n}\cdot q^n + O_{p,n}(q^{n-1}).$$

Proof. By Lemma 3.2(a) and Theorem 3.4, there are exactly $c_\infty^{\mathrm{diag}}(p,n) = p^{n^2-n}$ diagonalizable subalgebras of $\mathfrak{M}_n(\mathbb{F}_p)$ of dimension n and none of larger dimension. The claim thus follows from Lemma 3.1 by inclusion-exclusion. (For any n-dimensional subalgebra A of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ defined over $\mathbb{F}_p$, we have $|A\cap\mathfrak{M}_n(\mathbb{F}_q)| = q^n$, and for any two such subalgebras $A_1\neq A_2$, we have $|A_1\cap A_2\cap\mathfrak{M}_n(\mathbb{F}_q)|\leq q^{n-1}$.) □

3.2. General matrices

Let $n\geq 3$. We recall the definition of the Gaussian binomial coefficient

$$\binom{n}{k}_p := \frac{(p^n-1)\cdots(p^{n-k+1}-1)}{(p^k-1)\cdots(p-1)},$$

which is the number of k-dimensional subspaces of $\mathbb{F}_p^n$.

Theorem 3.6 (cf. Theorem 1.7)**.** *The maximal dimension of a commutative subalgebra of $\mathfrak{M}_n(\mathbb{F}_p)$ is $\lfloor n^2/4\rfloor+1$, and the number $c_\infty(p,n)$ of commutative subalgebras of that dimension is given by:*

$$c_\infty(p,3) = p^6+p^5+3p^4+3p^3+3p^2+p+1,$$

$$c_\infty(p,n) = \binom{n}{n/2}_p \text{ if } n\geq 4 \text{ is even,} \qquad c_\infty(p,n) = 2\binom{n}{\lfloor n/2\rfloor}_p \text{ if } n\geq 5 \text{ is odd.}$$

Proof. The commutative subalgebras of maximal dimension of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ were classified in [Sch05] (see also [Mir98]). In particular, they have dimension $\lfloor n^2/4 \rfloor + 1$.

We now explain how to parametrize them. For any subspace $V \subsetneq \overline{\mathbb{F}}_p^n$, let C_V be the linear subspace of matrices $A \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$ such that $\operatorname{im} A \subseteq V \subseteq \ker A$, and let C'_V be the algebra $C_V + \overline{\mathbb{F}}_p I_n$. The product of any two elements of C_V is zero, so C'_V is a commutative subalgebra of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$. Moreover, V can be recovered as the union of all images of nilpotent elements of C'_V, so the map $V \mapsto C'_V$ is injective. We have $\sigma(C'_V) = C'_{\sigma(V)}$, so the algebra C'_V is defined over $\mathbb{F}_p$ if and only if V is defined over $\mathbb{F}_p$. By [Sch05, Satz II and Satz III], when $n > 3$, the commutative subalgebras of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ of (maximal) dimension $\lfloor n^2/4 \rfloor + 1$ are exactly those of the form C'_V with $\dim V = \lfloor n/2 \rfloor$ or $\dim V = \lceil n/2 \rceil$. So, for $n > 3$, there are as many $(\lfloor n^2/4 \rfloor + 1)$-dimensional commutative subalgebras defined over $\mathbb{F}_p$ as there are choices for such a subspace V defined over $\mathbb{F}_p$, namely $\binom{n}{n/2}_p$ for even n and $\binom{n}{\lfloor n/2 \rfloor}_p + \binom{n}{\lceil n/2 \rceil}_p = 2\binom{n}{\lfloor n/2 \rfloor}_p$ for odd n. This proves the result for $n > 3$.

We now compute $c_\infty(p, 3)$. According to [Sch05, Satz II, Satz III and p. 76], there are five conjugacy classes (up to $\mathrm{GL}_3(\overline{\mathbb{F}}_p)$-conjugation) of three-dimensional commutative subalgebras of $\mathfrak{M}_3(\overline{\mathbb{F}}_p)$. In the following table, we list one representative A of each conjugacy class and the number $N(A)$ of subalgebras defined over $\mathbb{F}_p$ in the corresponding conjugacy class (the computations of $N(A)$ are detailed below the table):

	representative A	$N(A)$
(1)	$\left\{ \begin{pmatrix} \alpha & \beta & \gamma \\ & \alpha & \\ & & \alpha \end{pmatrix} \,\middle|\, \alpha, \beta, \gamma \in \overline{\mathbb{F}}_p \right\}$	$p^2 + p + 1$
(2)	$\left\{ \begin{pmatrix} \alpha & & \beta \\ & \alpha & \gamma \\ & & \alpha \end{pmatrix} \,\middle|\, \alpha, \beta, \gamma \in \overline{\mathbb{F}}_p \right\}$	$p^2 + p + 1$
(3)	$\left\{ \begin{pmatrix} \alpha & & \\ & \beta & \\ & & \gamma \end{pmatrix} \,\middle|\, \alpha, \beta, \gamma \in \overline{\mathbb{F}}_p \right\}$	p^6
(4)	$\left\{ \begin{pmatrix} \alpha & \beta & \\ & \alpha & \\ & & \gamma \end{pmatrix} \,\middle|\, \alpha, \beta, \gamma \in \overline{\mathbb{F}}_p \right\}$	$p^2(p^2 + p + 1)(p + 1)$
(5)	$\left\{ \begin{pmatrix} \alpha & \beta & \gamma \\ & \alpha & \beta \\ & & \alpha \end{pmatrix} \,\middle|\, \alpha, \beta, \gamma \in \overline{\mathbb{F}}_p \right\}$	$(p^2 + p + 1)(p + 1)(p - 1)$

Cases (1) and (2) correspond to the conjugacy classes $\{C'_V \mid V \subseteq \overline{\mathbb{F}}_p^3 \text{ one-dimensional}\}$ and $\{C'_V \mid V \subseteq \overline{\mathbb{F}}_p^3 \text{ two-dimensional}\}$, respectively, each of which contains $\binom{3}{1}_p = \binom{3}{2}_p = p^2 + p + 1$ subalgebras defined over $\mathbb{F}_p$ (see the arguments above for odd $n > 3$). Case (3) corresponds to the conjugacy class of diagonalizable subalgebras, which according to Theorem 3.4 contains p^6 subalgebras defined over $\mathbb{F}_p$. In cases (4) and (5), the $\mathrm{GL}_3(\overline{\mathbb{F}}_p)$-stabilizers S of A with respect to conjugation are respectively

$$\left\{ \begin{pmatrix} a & b & \\ & c & \\ & & d \end{pmatrix} \,\middle|\, a, c, d \in \overline{\mathbb{F}}_p^\times,\ b \in \overline{\mathbb{F}}_p \right\} \quad \text{and} \quad \left\{ \begin{pmatrix} a & b & c \\ & d & e \\ & & f \end{pmatrix} \,\middle|\, a, d, f \in \overline{\mathbb{F}}_p^\times,\ b, c, e \in \overline{\mathbb{F}}_p,\ \text{with } af = d^2 \right\}.$$

In both cases, we have $H^1(\mathrm{Gal}(\overline{\mathbb{F}}_p|\mathbb{F}_p), S) = \{1\}$,[3] so any algebra which is $\mathrm{GL}_3(\overline{\mathbb{F}}_p)$-conjugate to A and defined over $\mathbb{F}_p$ is actually $\mathrm{GL}_3(\mathbb{F}_p)$-conjugate to A.[4] The size of the $\mathrm{GL}_3(\mathbb{F}_p)$-conjugacy class is $|\mathrm{GL}_3(\mathbb{F}_p)|/|S \cap \mathrm{GL}_3(\mathbb{F}_p)|$, which is verified to be the number given in the table. Summing everything, we find that $c_\infty(p, 3) = p^6 + p^5 + 3p^4 + 3p^3 + 3p^2 + p + 1$. □

As in the proof of Theorem 3.5, we deduce from Lemma 3.1 and Theorem 3.6 the following theorem, which is Theorem 1.2 from the introduction (for $n \geq 3$):

[3] By [Ser79, Chap. X, §1, Exercise 2], the unit group of any algebra defined over $\mathbb{F}_p$ has trivial first Galois cohomology. This directly shows case (4), and case (5) follows by looking at the long exact sequence in cohomology arising from the short exact sequence $1 \to S \to T^\times \to \overline{\mathbb{F}}_p^\times \to 1$, where T is the algebra of upper triangular matrices with coefficients in $\overline{\mathbb{F}}_p$, and the homomorphism on the right is $\begin{pmatrix} a & b & c \\ & d & e \\ & & f \end{pmatrix} \mapsto afd^{-2}$.

[4] If the algebra $U^{-1}AU$ is defined over $\mathbb{F}_p$ for some $U \in \mathrm{GL}_3(\overline{\mathbb{F}}_p)$, we obtain a 1-cocycle $\tau \mapsto U\tau(U)^{-1} \in S$. It must be a 1-coboundary $\tau \mapsto T\tau(T)^{-1}$ for some $T \in S$, so $U' := T^{-1}U$ lies in $\mathrm{GL}_3(\mathbb{F}_p)$, and then $U^{-1}AU = U'^{-1}AU'$.

Theorem 3.7. *Let $c_\infty(p,n)$ be as in Theorem 3.6. For any finite field $\mathbb{F}_q \supseteq \mathbb{F}_p$, we have*

$$|\mathfrak{X}_\infty \cap \mathfrak{M}_n(\mathbb{F}_q)| = c_\infty(p,n) \cdot q^{\lfloor n^2/4 \rfloor + 1} + O_{p,n}\Big(q^{\lfloor n^2/4 \rfloor}\Big).$$

4. Diagonalizable matrices commuting with their Frobenius

In this section, we determine the asymptotics of $|\mathfrak{X}^{\mathrm{diag}} \cap \mathfrak{M}_n(\mathbb{F}_q)|$, i.e., we prove Theorem 4.17 (which is Theorem 1.3). In Subsection 4.1, we associate to any such matrix a quiver $\mathcal{Q}$ encoding the dimensions of the intersections of the eigenspaces of M with those of $\sigma(M)$. This will let us write $\mathfrak{X}^{\mathrm{diag}}$ as a disjoint union of equidimensional constructible subsets $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}} \subseteq \mathfrak{M}_n(\overline{\mathbb{F}}_p) \simeq \overline{\mathbb{F}}_p^{n^2}$. In Subsection 4.2, we identify those quivers $\mathcal{Q}$ for which the dimension of $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}$ is maximal, and in Subsections 4.3 to 4.6, we compute the irreducible components of the corresponding sets $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}$, and we show that they are defined over $\mathbb{F}_p$. This allows us to prove Theorem 1.3 using the Lang–Weil bound in Subsection 4.7.

4.1. Diagonalizable matrices and their associated quivers

Balanced quivers. A *quiver* is a finite directed graph in which one also allows loops (from a vertex to itself) and multiple parallel edges. We say that a vertex of a quiver is *isolated* if there are no edges (including loops) having that vertex as either source or target. We say that a quiver is *balanced* if, for each vertex, equally many edges have that vertex as source and as target (i.e., in-degrees and out-degrees coincide). If $\mathcal{Q}$ is a quiver, we denote by $V(\mathcal{Q})$ the set of its vertices, and by $\mathcal{Q}(i,j)$ the set of edges $i \to j$ for any $i,j \in V(\mathcal{Q})$. Assuming that $\mathcal{Q}$ is balanced, we also define the degree $d_{\mathcal{Q}}(i) \coloneqq \sum_{j \in V(\mathcal{Q})} |\mathcal{Q}(i,j)| = \sum_{j \in V(\mathcal{Q})} |\mathcal{Q}(j,i)|$ of each vertex $i \in V(\mathcal{Q})$. We let Bal_n be the (finite) set of isomorphism classes of balanced quivers with no isolated vertices and n edges.

Quiver associated to a matrix. Let $M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$. For each eigenvalue λ of M, let E_λ be the eigenspace $\ker(M - \lambda I_n)$. Note that $\sigma(E_\lambda) = \ker(\sigma(M) - \sigma(\lambda) I_n)$ is the eigenspace of $\sigma(M)$ for the eigenvalue $\sigma(\lambda)$.

Definition 4.1. We associate to the matrix M a quiver $\mathcal{Q}_M$ defined as follows:

- its vertices are the eigenvalues λ of M;
- for any eigenvalues λ, μ, the number of edges $\lambda \to \mu$ is the dimension of $E_\lambda \cap \sigma(E_\mu)$.

Proposition 4.2. *Let $M \in \mathfrak{M}_n^{\mathrm{diag}}(\overline{\mathbb{F}}_p)$. Then, $M \in \mathfrak{X}^{\mathrm{diag}}$ if and only if the corresponding quiver $\mathcal{Q}_M$ has exactly n edges. In that case, $\mathcal{Q}_M \in \mathrm{Bal}_n$, and $\dim E_\lambda = d_{\mathcal{Q}_M}(\lambda)$ for all eigenvalues λ.*

Proof. Since $\bigoplus_\lambda E_\lambda = \overline{\mathbb{F}}_p^n$ and $\bigoplus_\lambda \sigma(E_\lambda) = \overline{\mathbb{F}}_p^n$, the spaces $E_\lambda \cap \sigma(E_\mu)$ are always linearly independent. The diagonalizable matrices M and $\sigma(M)$ commute if and only if they are simultaneously diagonalizable, i.e., if and only if

$$\bigoplus_{\lambda,\mu} \Big(E_\lambda \cap \sigma(E_\mu)\Big) = \overline{\mathbb{F}}_p^n,$$

meaning that the quiver $\mathcal{Q}_M$ has exactly n edges. In that case, for any eigenvalue λ of M, we have

$$\bigoplus_\mu \Big(E_\lambda \cap \sigma(E_\mu)\Big) = E_\lambda \simeq \sigma(E_\lambda) = \bigoplus_\mu \Big(E_\mu \cap \sigma(E_\lambda)\Big),$$

so the quiver is balanced and satisfies $d_{\mathcal{Q}_M}(\lambda) = \dim E_\lambda$ (in particular, it has no isolated vertices). □

The space of matrices having a given quiver. For any quiver $\mathcal{Q} \in \mathrm{Bal}_n$, we define the subset $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}} \subseteq \mathfrak{X}^{\mathrm{diag}}$ of matrices M such that $\mathcal{Q}_M \simeq \mathcal{Q}$.[5] Proposition 4.2 directly implies:

$$\mathfrak{X}^{\mathrm{diag}} = \bigsqcup_{\mathcal{Q} \in \mathrm{Bal}_n} \mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}. \tag{4.1}$$

We will show that each set $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}$ is constructible, so that, by the Lang–Weil estimates (cf. [LW54]), the leading term in the asymptotics of $|\mathfrak{X}^{\mathrm{diag}} \cap \mathfrak{M}_n(\mathbb{F}_q)|$ depends on the maximal dimension of $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}$ over quivers $\mathcal{Q} \in \mathrm{Bal}_n$, and on the number of irreducible components having that dimension that are defined over $\mathbb{F}_p$.

Fix a quiver $\mathcal{Q} \in \mathrm{Bal}_n$. In order to compute the geometric invariants of $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}$, we explain how to construct all the diagonalizable matrices M such that $\mathcal{Q}_M \simeq \mathcal{Q}$. For each vertex i of $\mathcal{Q}$, we must pick an eigenvalue λ_i and an eigenspace V_i, making sure that:

- the eigenvalues λ_i are distinct;
- the eigenspaces V_i are in direct sum, and together span the entire (n-dimensional) space;
- the dimension of $V_i \cap \sigma(V_j)$ equals the number of edges $i \to j$ in $\mathcal{Q}$.

For any finite-dimensional vector space V and any k, we denote by $\mathrm{Gr}_k(V)$ the Grassmannian variety parametrizing k-dimensional subspaces of V. This space has dimension $k(\dim V - k)$ if $0 \le k \le \dim V$ and is otherwise empty. (See for example [Har92, Lecture 6] for an introduction to Grassmannians.) We also write $\mathbb{P}(V) := \mathrm{Gr}_1(V)$ for the projective space parametrizing one-dimensional subspaces of V. We will repeatedly make use of the fact that for any k, l, m, the subset

$$\{(A, B) \in \mathrm{Gr}_k(V) \times \mathrm{Gr}_l(V) \mid \dim(A + B) = m\} \tag{4.2}$$

of $\mathrm{Gr}_k(V) \times \mathrm{Gr}_l(V)$ is locally closed, and that the maps defined on that set mapping (A, B) to $A + B \in \mathrm{Gr}_m(V)$ (resp. to $A \cap B \in \mathrm{Gr}_{k+l-m}(V)$) are regular. Moreover, for any $n, k \ge 0$, the following map is also regular:

$$\mathrm{Gr}_k(\overline{\mathbb{F}}_p^n) \to \mathrm{Gr}_k(\overline{\mathbb{F}}_p^n), \qquad A \mapsto \sigma(A).$$

Let $r = |V(\mathcal{Q})|$, say $V(\mathcal{Q}) = \{1, \dots, r\}$. We define the following two quasi-projective varieties:

- $\mathfrak{Y}_{\mathcal{Q}}$ is the variety of ordered tuples $(\lambda_1, \dots, \lambda_r)$ of distinct elements of $\overline{\mathbb{F}}_p$. It is a non-empty Zariski open subset of $\overline{\mathbb{F}}_p^r$, hence it is Zariski dense and its dimension is $r = |V(\mathcal{Q})|$.
- $\mathfrak{Z}_{\mathcal{Q}}$ is the (locally closed) subspace of $\mathrm{Gr}_{d_{\mathcal{Q}}(1)}(\overline{\mathbb{F}}_p^n) \times \cdots \times \mathrm{Gr}_{d_{\mathcal{Q}}(r)}(\overline{\mathbb{F}}_p^n)$ consisting of those tuples $(V_1, \dots, V_r)$ of subspaces of $\overline{\mathbb{F}}_p^n$ of dimensions $d_{\mathcal{Q}}(1), \dots, d_{\mathcal{Q}}(r)$ which together span $\overline{\mathbb{F}}_p^n$ and such that $\dim(V_i \cap \sigma(V_j)) = |\mathcal{Q}(i, j)|$ for all i, j.

Sending a pair $((\lambda_1, \dots, \lambda_r), (V_1, \dots, V_r)) \in \mathfrak{Y}_{\mathcal{Q}} \times \mathfrak{Z}_{\mathcal{Q}}$ to the diagonalizable matrix M with eigenvalues $\lambda_1, \dots, \lambda_r$ and corresponding eigenspaces $V_1, \dots, V_r$, we obtain a regular map

$$\mathfrak{Y}_{\mathcal{Q}} \times \mathfrak{Z}_{\mathcal{Q}} \to \mathfrak{M}_n(\overline{\mathbb{F}}_p) \tag{4.3}$$

whose image is exactly $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}$ by Proposition 4.2. In particular, $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}$ is a constructible subset of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ by Chevalley's theorem. The group $\mathrm{Aut}(\mathcal{Q})$ consisting of automorphisms of the quiver, i.e., of permutations of the vertices which preserve edge multiplicities, acts simply transitively on each fiber above a point of $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}$. Moreover, the Frobenius automorphism acts on the sets $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}$, $\mathfrak{Y}_{\mathcal{Q}}$, $\mathfrak{Z}_{\mathcal{Q}}$, and the map from Equation (4.3) is σ-equivariant.

To compute the dimension of $\mathfrak{Z}_{\mathcal{Q}}$, we use the following lemma:

[5] Be aware that this is *not* a quiver variety or a quiver Grassmannian in the usual sense.

Lemma 4.3. *Let $r \geq 1$. The map $\wp\colon \mathrm{GL}_n(\overline{\mathbb{F}}_p) \to \mathrm{GL}_n(\overline{\mathbb{F}}_p)$ given by $\wp(E) := E^{-1}\sigma^r(E)$ is étale and surjective. Moreover, $\mathrm{GL}_n(\mathbb{F}_{p^r})$ acts simply transitively on each fiber by left multiplication.*

Proof. More generally, for any $A \in \mathrm{GL}_n(\overline{\mathbb{F}}_p)$, consider the map $\wp_A\colon \mathrm{GL}_n(\overline{\mathbb{F}}_p) \to \mathrm{GL}_n(\overline{\mathbb{F}}_p)$ given by $\wp_A(E) := E^{-1}A\sigma^r(E)$. As $p = 0$ in $\overline{\mathbb{F}}_p$, the differential of σ^r is the zero map (at any point); by the product rule, the differential of $\wp_A$ at a matrix $E \in \mathrm{GL}_n(\overline{\mathbb{F}}_p)$ thus maps a tangent vector $\mathrm{d}E \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$ to $-E^{-1}\mathrm{d}EE^{-1}A\sigma^r(E) \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$. Hence, the differential of $\wp_A$ at every point E is a linear isomorphism, so $\wp_A$ is étale. Since domain and target have the same dimension and $\mathrm{GL}_n(\overline{\mathbb{F}}_p)$ is irreducible, this implies that $\wp_A$ is dominant for all A. The image of $\wp_A$ (which is dense, and constructible by Chevalley's theorem) then contains a non-empty open subset of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$, hence intersects the (dense) image of $\wp_{I_n} = \wp$. We have an equality $\wp_A(E_1) = \wp(E_2)$, implying that $A = \wp(E_2E_1^{-1})$. We have shown that the map $\wp$ is surjective.

Finally, we have $E^{-1}\sigma^r(E) = E'^{-1}\sigma^r(E')$ if and only if $E'E^{-1} \in \mathrm{GL}_n(\mathbb{F}_{p^r})$, so all non-empty fibers of $\wp$ are right $\mathrm{GL}_n(\mathbb{F}_{p^r})$-cosets. □

Lemma 4.4.

(a) *The space $\mathfrak{Z}_{\mathcal{Q}}$ is non-empty and has pure dimension $\sum_i d_{\mathcal{Q}}(i)^2 - \sum_{i,j} |\mathcal{Q}(i,j)|^2$, and the finite group $\mathrm{GL}_n(\mathbb{F}_p)$ acts transitively on the set of its irreducible components.*

(b) *Let $k \in V(\mathcal{Q})$ with $0 < |\mathcal{Q}(k,k)| < d_{\mathcal{Q}}(k)$. Consider the locally closed subset $\mathfrak{Z}_{\mathcal{Q},k} \subseteq \mathfrak{Z}_{\mathcal{Q}}$ consisting of those tuples $(V_1,\ldots,V_r) \in \mathfrak{Z}_{\mathcal{Q}}$ for which $V_k \cap \sigma(V_k)$ is defined over $\mathbb{F}_p$. This subset has strictly smaller dimension than $\mathfrak{Z}_{\mathcal{Q}}$.*

Proof.

(a) The formulas $U_{ij} := V_i \cap \sigma(V_j)$ and $V_i := \bigoplus_j U_{ij}$ define two inverse regular maps, showing that $\mathfrak{Z}_{\mathcal{Q}}$ is isomorphic to the subvariety $\widetilde{\mathfrak{Z}}_{\mathcal{Q}}$ of $\prod_{i,j} \mathrm{Gr}_{|\mathcal{Q}(i,j)|}(\overline{\mathbb{F}}_p^n)$ parametrizing tuples $(U_{ij})_{i,j\in[r]}$ of subspaces of $\overline{\mathbb{F}}_p^n$ satisfying the following three conditions: $\dim U_{ij} = |\mathcal{Q}(i,j)|$ for all $i,j \in V(\mathcal{Q})$, $\bigoplus_{i,j} U_{ij} = \overline{\mathbb{F}}_p^n$, and $\sigma(\bigoplus_j U_{ij}) = \bigoplus_j U_{ji}$ for all $i \in V(\mathcal{Q})$.

Define the $\overline{\mathbb{F}}_p$-vector spaces $C_{ij} := \overline{\mathbb{F}}_p^{|\mathcal{Q}(i,j)|}$ and $C := \bigoplus_{i,j} C_{ij}$. By definition, C is isomorphic to $\overline{\mathbb{F}}_p^n$. In order to parametrize tuples $(U_{ij})_{i,j} \in \widetilde{\mathfrak{Z}}_{\mathcal{Q}}$, we consider the surjective regular map

$$f\colon \mathrm{Isom}(C,\overline{\mathbb{F}}_p^n) \to \Big\{(U_{ij})_{i,j} \;\Big|\; \dim U_{ij} = |\mathcal{Q}(i,j)| \text{ and } \textstyle\bigoplus_{i,j} U_{ij} = \overline{\mathbb{F}}_p^n\Big\}, \quad E \mapsto \big(E(C_{ij})\big)_{i,j},$$

whose fibers are isomorphic to the variety

$$F := \prod_{i,j} \mathrm{GL}(C_{ij}), \quad \text{of dimension } \sum_{i,j} (\dim C_{ij})^2 = \sum_{i,j} |\mathcal{Q}(i,j)|^2.$$

For any $E \in \mathrm{Isom}(C,\overline{\mathbb{F}}_p^n)$, let $\sigma(E)$ be the $\overline{\mathbb{F}}_p$-linear isomorphism obtained as the composition $C \xrightarrow{\sigma^{-1}} C \xrightarrow{E} \overline{\mathbb{F}}_p^n \xrightarrow{\sigma} \overline{\mathbb{F}}_p^n$, where σ acts on C and on $\overline{\mathbb{F}}_p^n$ in the natural way. We have $\sigma(\bigoplus_j U_{ij}) = \bigoplus_j U_{ji}$ if and only if $\wp(E) := E^{-1}\sigma(E)$ sends $\bigoplus_j C_{ij}$ to $\bigoplus_j C_{ji}$, i.e., if and only if $\wp(E)$ lies in the irreducible variety

$$S := \prod_i \mathrm{Isom}\Big(\bigoplus_j C_{ij}, \bigoplus_j C_{ji}\Big), \quad \text{of dimension } \sum_i \Big(\sum_j \dim C_{ij}\Big)\Big(\sum_j \dim C_{ji}\Big) = \sum_i d_{\mathcal{Q}}(i)^2.$$

In other words, $\widetilde{\mathfrak{Z}}_{\mathcal{Q}} = f(\wp^{-1}(S))$. Together with Lemma 4.3, this implies that $\mathfrak{Z}_{\mathcal{Q}} \simeq \widetilde{\mathfrak{Z}}_{\mathcal{Q}} = f(\wp^{-1}(S))$ is non-empty and has pure dimension

$$\dim \mathfrak{Z}_{\mathcal{Q}} = \dim \wp^{-1}(S) - \dim F = \dim S - \dim F = \sum_i d_{\mathcal{Q}}(i)^2 - \sum_{i,j} |\mathcal{Q}(i,j)|^2$$

and that $\mathrm{GL}_n(\mathbb{F}_p)$ acts transitively on the set of its irreducible components.

(b) We reason as in (a). In terms of the notation above, the condition $\sigma(U_{kk}) = U_{kk}$ means that $\wp(E)$ must send C_{kk} to itself, so S must be replaced by the subset $S' := \{A \in S \mid A(C_{kk}) = C_{kk}\}$, and the claim reduces to showing that $\dim S' < \dim S$. We can describe S as the subset of the vector space $\mathrm{Hom}(C_{kk}, C_{kk}) \times \prod_{(i,j)\neq(k,k)} \mathrm{Hom}(C_{ij}, \bigoplus_{j'} C_{j'i}) \subseteq \mathrm{Hom}(C, C)$ formed of those endomorphisms which are invertible, so S' has the same dimension as that vector space, namely

$$\begin{aligned}\dim S' &= |\mathcal{Q}(k,k)|^2 + \sum_{(i,j)\neq(k,k)} |\mathcal{Q}(i,j)| \cdot d_{\mathcal{Q}}(i) \\ &= |\mathcal{Q}(k,k)|^2 - |\mathcal{Q}(k,k)| \cdot d_{\mathcal{Q}}(k) + \sum_{i,j} |\mathcal{Q}(i,j)| \cdot d_{\mathcal{Q}}(i) \\ &= -|\mathcal{Q}(k,k)| \cdot (d_{\mathcal{Q}}(k) - |\mathcal{Q}(k,k)|) + \sum_i d_{\mathcal{Q}}(i)^2 \\ &< \sum_i d_{\mathcal{Q}}(i)^2 = \dim S.\end{aligned}$$

□

Corollary 4.5. *The subset $\mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}} \subseteq \mathfrak{M}_n(\overline{\mathbb{F}}_p)$ is constructible, of pure dimension*

$$\dim \mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}} = |V(\mathcal{Q})| + \sum_{i \in V(\mathcal{Q})} d_{\mathcal{Q}}(i)^2 - \sum_{i,j \in V(\mathcal{Q})} |\mathcal{Q}(i,j)|^2.$$

Proof. Since every fiber of the surjection $\mathfrak{Y}_{\mathcal{Q}} \times \mathfrak{Z}_{\mathcal{Q}} \twoheadrightarrow \mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ is finite (of size $|\mathrm{Aut}(\mathcal{Q})|$), $\mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ is equidimensional and

$$\dim \mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}} = \dim \mathfrak{Y}_{\mathcal{Q}} + \dim \mathfrak{Z}_{\mathcal{Q}} \underset{\text{Lem. 4.4(a)}}{=} |V(\mathcal{Q})| + \sum_i d_{\mathcal{Q}}(i)^2 - \sum_{i,j} |\mathcal{Q}(i,j)|^2.$$

□

4.2. The octopus has maximal dimension

Corollary 4.5 and Equation (4.1) imply that the dimension of $\mathfrak{X}^{\mathrm{diag}}$ is the maximal dimension of $\mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ over quivers $\mathcal{Q} \in \mathrm{Bal}_n$, and give an explicit formula for the dimension of $\mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ in terms of the quiver $\mathcal{Q}$. We shall now compute this maximal dimension and describe the corresponding optimal quivers.

Proposition 4.6. *Let $n \geq 1$, and let $[n/3]$ be the (uniquely defined) integer closest to $n/3$. Then:*

$$\max_{\mathcal{Q} \in \mathrm{Bal}_n} \dim \mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}} = \left\lfloor \frac{n^2}{3} \right\rfloor + 1.$$

The maximum is reached by the following quiver with $[n/3] + 1$ vertices, which we call the octopus quiver *(with n edges) and denote by $\mathcal{O}_n$:*

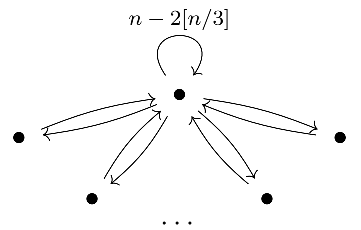

where the number on the top loop means that there are $n - 2[n/3]$ parallel loops from the central vertex to itself. Moreover, up to isomorphism:

- *When $n \notin \{2, 4\}$, there are no other quivers in Bal_n maximizing $\dim \mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$;*

- *When $n = 2$, there is a single additional optimal (non-connected) quiver, namely $\mathcal{O}_1 \sqcup \mathcal{O}_1$:*

- *When $n = 4$, there is a single additional optimal quiver, which we call the* dumbbell quiver*:*

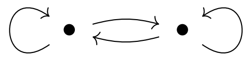

Proof. Corollary 4.5 gives a formula for $\dim \mathfrak{X}_{\mathcal{Q}}^{\text{diag}}$, reducing the proposition to a purely combinatorial statement. The proposed quivers do reach the proposed maximum, establishing the lower bound $\max_{\mathcal{Q}\in\text{Bal}_n} \dim \mathfrak{X}_{\mathcal{Q}}^{\text{diag}} \geq \lfloor n^2/3 \rfloor + 1$. We prove by induction on n that this is indeed the maximum, and that the quivers reaching that maximum are exactly the proposed ones. We leave aside the cases $n = 1$ and $n = 2$, which are easily checked. Let $n > 2$, and assume that for all $n' < n$ and for all $\mathcal{Q}' \in \text{Bal}_{n'}$ we have $\dim \mathfrak{X}_{\mathcal{Q}'}^{\text{diag}} \leq \lfloor n'^2/3 \rfloor + 1$. We consider a quiver $\mathcal{Q} \in \text{Bal}_n$ satisfying $\dim \mathfrak{X}_{\mathcal{Q}}^{\text{diag}} \geq \lfloor n^2/3 \rfloor + 1$.

We first show that $\mathcal{Q}$ is connected. For this, notice that $\dim \mathfrak{X}_{\mathcal{Q}}^{\text{diag}}$ is additive with respect to unions of vertex-disjoint quivers. By the induction hypothesis and since the function $\eta(n) := \lfloor n^2/3 \rfloor + 1$ is strictly superadditive on positive integers with the single exception of the equality $\eta(1)+\eta(1) = \eta(2)$, we cannot reach or beat $\lfloor n^2/3 \rfloor + 1$ if there are at least two connected components (recall that we have assumed $n > 2$).

Now, let ℓ be an integer, and consider a subquiver $C \subseteq \mathcal{Q}$ which is a union of any number of vertex-disjoint cycles whose lengths sum to ℓ (for example, C can be a single ℓ-cycle), thus consisting of ℓ vertices and ℓ edges. Removing from the quiver $\mathcal{Q}$ the edges of C and the vertices which have become isolated, we obtain a balanced quiver $\mathcal{Q} \setminus C$ with $n - \ell$ edges. We have, by Corollary 4.5:

$$\begin{aligned}
\dim \mathfrak{X}_{\mathcal{Q}}^{\text{diag}} - \dim \mathfrak{X}_{\mathcal{Q}\setminus C}^{\text{diag}} = {} & \underbrace{|\{i \in V(C) \mid d_{\mathcal{Q}}(i) = 1\}|}_{\text{vertices which have become isolated}} + \sum_{i \in V(C)} \left[d_{\mathcal{Q}}(i)^2 - (d_{\mathcal{Q}}(i) - 1)^2 \right] \\
& - \sum_{(i\to j)\in C} \left[|\mathcal{Q}(i,j)|^2 - (|\mathcal{Q}(i,j)| - 1)^2 \right] \\
= {} & |\{i \in V(C) \mid d_{\mathcal{Q}}(i) = 1\}| + 2 \sum_{i\in V(C)} d_{\mathcal{Q}}(i) - 2 \sum_{(i\to j)\in C} |\mathcal{Q}(i,j)|. \qquad (4.4)
\end{aligned}$$

By hypothesis, $\dim \mathfrak{X}_{\mathcal{Q}}^{\text{diag}} \geq \lfloor n^2/3 \rfloor + 1$. By the induction hypothesis, $\dim \mathfrak{X}_{\mathcal{Q}\setminus C}^{\text{diag}} \leq \lfloor (n-\ell)^2/3 \rfloor + 1$. Therefore:

$$\dim \mathfrak{X}_{\mathcal{Q}}^{\text{diag}} - \dim \mathfrak{X}_{\mathcal{Q}\setminus C}^{\text{diag}} \geq \left\lfloor \frac{n^2}{3} \right\rfloor - \left\lfloor \frac{(n-\ell)^2}{3} \right\rfloor \geq \frac{n^2 - 2}{3} - \frac{n^2 - 2n\ell + \ell^2}{3} = \frac{2n\ell - \ell^2 - 2}{3} \qquad (4.5)$$

We clearly have $|\{i \in V(C) \mid d_{\mathcal{Q}}(i) = 1\}| \leq \ell$. If $|\{i \in V(C) \mid d_{\mathcal{Q}}(i) = 1\}| = \ell$, then C is a union of connected components of $\mathcal{Q}$, hence $\mathcal{Q} = C$ as $\mathcal{Q}$ is connected (in particular, C is a single cycle in this case), so $\ell = r = n$, but then $\dim \mathfrak{X}_{\mathcal{Q}}^{\text{diag}} = n$ is less than $\lfloor n^2/3 \rfloor + 1$ since $n > 2$. Therefore, we actually have $|\{i \in V(C) \mid d_{\mathcal{Q}}(i) = 1\}| \leq \ell - 1$, and so:

$$|\{i \in V(C) \mid d_{\mathcal{Q}}(i) = 1\}| + 2 \sum_{i\in V(C)} d_{\mathcal{Q}}(i) - 2 \sum_{(i\to j)\in C} |\mathcal{Q}(i,j)| \leq (\ell - 1) + 2n - 2\ell = 2n - \ell - 1. \quad (4.6)$$

Combining Equations (4.4) to (4.6), we must then have:

$$2n - \ell - 1 \geq \frac{2n\ell - \ell^2 - 2}{3}.$$

Multiplying by 3 and rearranging, this becomes

$$\underbrace{(\ell - 2n)}_{<0}(\ell - 3) \geq 1,$$

which is only possible if $\ell \leq 2$. We have thus shown:

$$\text{There is no union of vertex-disjoint cycles of } \mathcal{Q} \text{ whose lengths sum to 3 or more.} \tag{C}$$

Since $\mathcal{Q}$ is balanced, it can be written as a union of (not necessarily disjoint) cycles. By (C), only 1-cycles (i.e., loops) and 2-cycles can occur. In particular, for any two vertices $i \neq j$, the number of edges $i \to j$ equals the number of edges $j \to i$. We use the notation $i \xleftrightarrow{\alpha} j$ as a shortcut for α edges $i \to j$ and α edges $j \to i$ (this still counts as 2α edges!). By (C), there can be at most two vertices with loops. We distinguish two cases:

Case 1: There are two vertices i, j with loops.

Then, (C) implies that any 2-cycle contains both i and j. As $\mathcal{Q}$ is connected, i and j are the only vertices and $\mathcal{Q}$ looks as follows:

$$\alpha \circlearrowright i \xleftrightarrow{\gamma} j \circlearrowleft \beta$$

where $\alpha + 2\gamma + \beta = n$ and $\alpha, \beta, \gamma \geq 1$ (this is only possible if $n \geq 4$). We have

$$\dim \mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}} = 2 + (\alpha+\gamma)^2 + (\beta+\gamma)^2 - \alpha^2 - 2\gamma^2 - \beta^2 = 2 + 2\gamma(\alpha+\beta) = 2 + 2\gamma(n-2\gamma) = -4\gamma^2 + 2n\gamma + 2.$$

This polynomial in γ reaches its real maximum at $\gamma = \frac{n}{4}$ with maximal value $\frac{n^2}{4} + 2$, which is strictly smaller than $\lfloor n^2/3 \rfloor + 1$ as soon as $n \geq 5$, contradicting the hypothesis. For $n = 4$, the only possibility is $\alpha = \beta = \gamma = 1$, corresponding to the dumbbell quiver.

Case 2: There is at most one vertex with loops.

If there is a vertex i with loops, then all 2-cycles must contain the vertex i according to (C). Otherwise, (C) still implies that any two 2-cycles must share a vertex, and since there cannot be a 3-cycle, all 2-cycles then share some common vertex i.[6] Either way, since $\mathcal{Q}$ is connected, we see that $\mathcal{Q}$ is of the following form:

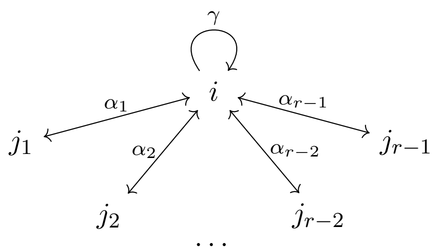

with $\gamma \geq 0$, $\alpha_1, \ldots, \alpha_{r-1} \geq 1$, and $\gamma + 2\sum_i \alpha_i = n$. We then have

$$\dim \mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}} = r + (n - \textstyle\sum_i \alpha_i)^2 + \sum_i \alpha_i^2 - (n - 2\sum_i \alpha_i)^2 - 2\sum_i \alpha_i^2 = r + 2n\sum_i \alpha_i - 3(\sum_i \alpha_i)^2 - \sum_i \alpha_i^2.$$

Let $S \coloneqq \sum_i \alpha_i$. We have $S \geq r - 1$ and $\sum_i \alpha_i^2 \geq S$, and thus:

$$\dim \mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}} \leq (S+1) + 2nS - 3S^2 - S = -3S^2 + 2nS + 1.$$

[6] Say $i \leftrightarrow j$ and $i \leftrightarrow k$ are two 2-cycles sharing a vertex i, with $j \neq k$. Any 2-cycle not containing the same vertex i would need to be $j \leftrightarrow k$. But then, $\mathcal{Q}$ would contain the 3-cycle $i \to j \to k \to i$, contradicting (C).

This upper bound is a quadratic polynomial in S whose real maximum is at $\frac{n}{3}$, thus reaching its integer maximum exactly when $S = [\frac{n}{3}]$, in which case this evaluates precisely to $\lfloor n^2/3 \rfloor + 1$. As, by hypothesis, we have $\dim \mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}} \geq \lfloor n^2/3 \rfloor + 1$, all inequalities above must be equalities. In particular, we have $S = r-1$ so $\alpha_1 = \ldots = \alpha_{r-1} = 1$, and $S = [n/3]$ so $r = S+1 = [n/3]+1$. The quiver $\mathcal{Q}$ is then precisely the octopus quiver. □

The following subsections are dedicated to describing the irreducible components of $\mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ for the quivers $\mathcal{Q}$ maximizing the dimension.

4.3. The special case $n = 2$

In the case $n = 2$, Proposition 4.6 shows that there are two isomorphism classes of quivers $\mathcal{Q} \in \mathrm{Bal}_2$ for which $\dim \mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ reaches the maximal value 2, namely:

$$\mathcal{O}_2 = \; 1 \rightleftarrows 2 \qquad \text{and} \qquad \mathcal{O}_1 \sqcup \mathcal{O}_1 = \; \circlearrowleft 1 \qquad 2 \circlearrowright$$

Proposition 4.7. *The two-dimensional set $\mathfrak{X}_{\mathcal{O}_2}^{\mathrm{diag}} \sqcup \mathfrak{X}_{\mathcal{O}_1 \sqcup \mathcal{O}_1}^{\mathrm{diag}}$ has exactly p^2 irreducible components, which are all fixed by σ.*

Proof. We let $X := \mathfrak{X}_{\mathcal{O}_2}^{\mathrm{diag}} \sqcup \mathfrak{X}_{\mathcal{O}_1 \sqcup \mathcal{O}_1}^{\mathrm{diag}}$. Note that $Y := \mathfrak{Y}_{\mathcal{O}_2} = \mathfrak{Y}_{\mathcal{O}_1 \sqcup \mathcal{O}_1}$ is the space of pairs of distinct elements of $\overline{\mathbb{F}}_p$. An element of $Z := \mathfrak{Z}_{\mathcal{O}_2} \sqcup \mathfrak{Z}_{\mathcal{O}_1 \sqcup \mathcal{O}_1}$ is a pair (V_1, V_2) of distinct one-dimensional subspaces of $\overline{\mathbb{F}}_p^2$ such that either $\sigma(V_1) = V_2$ and $\sigma(V_2) = V_1$ (for $\mathfrak{Z}_{\mathcal{O}_2}$), or $\sigma(V_1) = V_1$ and $\sigma(V_2) = V_2$ (for $\mathfrak{Z}_{\mathcal{O}_1 \sqcup \mathcal{O}_1}$); this can be summed up by saying that the unordered pair $\{V_1, V_2\}$ is σ-invariant. We have already counted such unordered pairs in Theorem 3.4 (cf. the bijection of Lemma 3.2), so we know that there are $p^{2^2-2} = p^2$ such pairs. (In this case, it is easier to distinguish between the two cases, giving $\frac{1}{2}(p^2 - p) + \frac{1}{2}(p^2 + p) = p^2$.) Thus, the set Z has size $2p^2$.

Both quivers have automorphism group $\mathrm{Aut}(\mathcal{Q})$ isomorphic to $\mathbb{Z}/2\mathbb{Z}$, corresponding to the permutation of the vertices 1 and 2. Thus, the maps of Equation (4.3) combine into a surjective regular σ-equivariant map $Y \times Z \twoheadrightarrow X$, whose fibers have size 2. Since Y is irreducible, the space $Y \times Z$ has $2p^2$ irreducible components, over which $\mathbb{Z}/2\mathbb{Z}$ acts freely (by swapping coordinates of both pairs), and moreover the $\mathbb{Z}/2\mathbb{Z}$-orbits are unions of σ-orbits (they form a single orbit for components coming from $\mathfrak{Z}_{\mathcal{O}_2}$, and two orbits for components coming from $\mathfrak{Z}_{\mathcal{O}_1 \sqcup \mathcal{O}_1}$). This implies that X has p^2 irreducible components, all of which are fixed by σ. □

4.4. A tool to prove irreducibility

We will repeatedly make use of the following lemma to prove the irreducibility of a variety:

Lemma 4.8. *Let $f \colon A \to B$ be a regular map between varieties. Assume that A is non-empty and has pure dimension d. Let $B_1, \ldots, B_s$ be locally closed subvarieties of B with $B = \bigsqcup_{i=1}^s B_i$ and for every $x \in B$, let F_x be a variety such that there is an injective regular map $\varphi_x \colon f^{-1}(x) \hookrightarrow F_x$. Assume that B_1 is irreducible, that F_x is irreducible for all $x \in B_1$, and that*

$$\begin{aligned} &\forall x \in B_1, && \dim F_x + \dim B_1 \leq d, \\ \forall i \in \{2, \ldots, s\}, \; &\forall x \in B_i, && \dim F_x + \dim B_i < d. \end{aligned}$$

Then, A is irreducible.

Proof. The assumptions imply that $\dim f^{-1}(B_i) < d$ for $i = 2, \ldots, s$. Hence, the (d-dimensional) irreducible components of A are in bijection with those of $A \setminus \bigcup_{i=2}^s f^{-1}(B_i)$. We can thus assume without loss of generality that $s = 1$, hence $B = B_1$.

Consider any irreducible component C of A. For generic $x \in \overline{f(C)}$, we have

$$d = \dim C \leq \dim(f^{-1}(x) \cap C) + \dim f(C) \leq \dim F_x + \dim B_1 \leq d,$$

so all inequalities have to be equalities: $\dim F_x + \dim B_1 = d$, the set $f(C)$ is dense in $B = B_1$ (recall that B_1 is irreducible), and the set $\varphi_x(f^{-1}(x) \cap C)$ (which is constructible by Chevalley's theorem) is dense in F_x (recall that F_x is irreducible), hence contains a non-empty open subset of F_x.

This implies that, for any two irreducible components C and C' of A and for generic $x \in B$, the set $\varphi_x(f^{-1}(x) \cap C \cap C') = \varphi_x(f^{-1}(x) \cap C) \cap \varphi_x(f^{-1}(x) \cap C')$ contains a non-empty open subset of F_x. We have shown that the fibers of the restricted map $f_{|C \cap C'} \colon C \cap C' \to B$ generically have dimension $\dim F_x$ (in particular, that restricted map is dominant), so $\dim(C \cap C') = \dim F_x + \dim B_1 = d$, which implies $C = C'$. □

4.5. The general case (the octopus variety)

Let $n \geq 3$, and let $\mathcal{Q}$ be the octopus quiver $\mathcal{O}_n$ (defined in Proposition 4.6). In this subsection, we show that $\mathfrak{Z}_{\mathcal{Q}}$ and $\mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ are irreducible (Proposition 4.11). For this purpose, we will need the following stratification of the Grassmannian:

Stratification of the Grassmannian. Let $0 \leq a \leq n$. We partition $\mathrm{Gr}_a(\overline{\mathbb{F}}_p^n)$ as follows

$$\mathrm{Gr}_a(\overline{\mathbb{F}}_p^n) = \bigsqcup_{0 \leq b \leq \min(a, n-a)} \mathfrak{T}_{a,b},$$

where $\mathfrak{T}_{a,b}$ is the subset of $\mathrm{Gr}_a(\overline{\mathbb{F}}_p^n)$ consisting of those a-dimensional subspaces $V \subseteq \overline{\mathbb{F}}_p^n$ such that $\dim(\sigma(V) + \sigma^{-1}(V)) = a + b$, or equivalently $\dim(\sigma(V) \cap \sigma^{-1}(V)) = a - b$.

Lemma 4.9. *For any $0 \leq b \leq \min(a, n-a)$, the strata $\mathfrak{T}_{a,b}$ satisfy the following properties:*

(a) $\mathfrak{T}_{a,b}$ is locally closed.

(b) $\mathfrak{T}_{a,b}$ is non-empty with pure dimension $b(n-b)$.

(c) If $b > 0$, then $\mathfrak{T}_{a,b}$ is irreducible.

(If $n \not\equiv 2 \pmod 3$, we will only need the "trivial" special case $b = \min(a, n-a)$ of point (c).)

Proof. Let X be the set of pairs (V, V') of a-dimensional subspaces of $\overline{\mathbb{F}}_p^n$ with $\dim(V + V') = a + b$ (equivalently, $\dim(V \cap V') = a - b$). The subset $X \subseteq \mathrm{Gr}_a(\overline{\mathbb{F}}_p^n) \times \mathrm{Gr}_a(\overline{\mathbb{F}}_p^n)$ is locally closed, cf. Equation (4.2).

(a) Since $\dim(\sigma(V) + \sigma^{-1}(V)) = \dim(\sigma^2(V) + V)$, the set $\mathfrak{T}_{a,b}$ is locally closed as the pullback of X under the regular map $V \mapsto (V, \sigma^2(V))$.

(b) We use a similar strategy as in the proof of Lemma 4.4. First note that $\mathfrak{T}_{a,b}$ is isomorphic to the variety $\widetilde{\mathfrak{T}}_{a,b}$ of those pairs $(V, V') \in X$ satisfying $\sigma^2(V) = V'$.

Let $C_1 \coloneqq \overline{\mathbb{F}}_p^b$, $C_2 \coloneqq \overline{\mathbb{F}}_p^{a-b}$, $C_3 \coloneqq \overline{\mathbb{F}}_p^b$, $C_4 \coloneqq \overline{\mathbb{F}}_p^{n-a-b}$, and $C \coloneqq C_1 \oplus C_2 \oplus C_3 \oplus C_4$. We parametrize pairs $(V, V') \in X$ via the regular map

$$f \colon \mathrm{Isom}(C, \overline{\mathbb{F}}_p^n) \to X, \qquad E \mapsto (E(C_1 \oplus C_2), E(C_2 \oplus C_3)).$$

This map is surjective and its fibers are isomorphic to the variety

$$\begin{aligned} F &\coloneqq \{E \in \mathrm{GL}(C) \mid E(C_1 \oplus C_2) = C_1 \oplus C_2 \text{ and } E(C_2 \oplus C_3) = C_2 \oplus C_3\} \\ &= \{E \in \mathrm{GL}(C) \mid E(C_1) \subseteq C_1 \oplus C_2 \text{ and } E(C_2) = C_2 \text{ and } E(C_3) \subseteq C_2 \oplus C_3\} \end{aligned}$$

of dimension

$$\begin{aligned}\dim F &= \dim C_1 \cdot \dim(C_1 \oplus C_2) + (\dim C_2)^2 + \dim C_3 \cdot \dim(C_2 \oplus C_3) + \dim C_4 \cdot \dim C \\ &= ba + (a-b)^2 + ba + (n-a-b)n \\ &= a^2 + (n-a)n - b(n-b).\end{aligned}$$

(That a generic linear endomorphism $E\colon C \to C$ such that $E(C_1) \subseteq C_1 \oplus C_2$, $E(C_2) \subseteq C_2$ and $E(C_3) \subseteq C_2 \oplus C_3$ satisfies $E(C_2) = C_2$ and is invertible follows from the fact that F is Zariski open in the vector space of such endomorphisms, and is non-empty as it contains the identity.)

Let $E \in \mathrm{Isom}(C, \overline{\mathbb{F}}_p^n)$ and let $(V, V') = f(E)$. We have $\sigma^2(V) = V'$ if and only if the automorphism $\wp(E) := E^{-1}\sigma^2(E) \in \mathrm{GL}(C)$ (with $\sigma^2(E)$ defined analogously to $\sigma(E)$ in the proof of Lemma 4.4) lies in the irreducible variety

$$S := \{A \in \mathrm{GL}(C) \mid A(C_1 \oplus C_2) = C_2 \oplus C_3\}$$

of dimension

$$\dim S = \dim(C_1 \oplus C_2) \cdot \dim(C_2 \oplus C_3) + \dim(C_3 \oplus C_4) \cdot \dim C = a^2 + (n-a)n.$$

(As above, generic invertibility comes from the fact that S is non-empty, as it contains the invertible map $C_1 \oplus C_2 \oplus C_3 \oplus C_4 \to C_1 \oplus C_2 \oplus C_3 \oplus C_4, (x, y, z, w) \mapsto (z, y, x, w)$.)

By Lemma 4.3, $\wp^{-1}(S)$ is non-empty of pure dimension $\dim S = a^2 + (n-a)n$. In particular, $\mathfrak{T}_{a,b} \simeq \widetilde{\mathfrak{T}}_{a,b} = f(\wp^{-1}(S))$ is non-empty of pure dimension

$$\dim \wp^{-1}(S) - \dim F = a^2 + (n-a)n - a^2 - (n-a)n + b(n-b) = b(n-b).$$

(c) We use downward induction on a. (The case $a = n$ is vacuous.) The case $b = \min(a, n-a)$ is clear since by (b), $\mathfrak{T}_{a,b}$ is then a subvariety of dimension $a(n-a)$ of the irreducible variety $\mathrm{Gr}_a(\overline{\mathbb{F}}_p^n)$ of dimension $a(n-a)$, hence it is dense, hence itself irreducible. We can therefore assume that $0 < b < \min(a, n-a)$.

We are going to apply Lemma 4.8 to the regular map

$$f\colon \mathfrak{T}_{a,b} \to \bigsqcup_{0 \le c \le \min(a+b,\, n-a-b)} \mathfrak{T}_{a+b,c} = \mathrm{Gr}_{a+b}(\overline{\mathbb{F}}_p^n)$$

sending V to $W := \sigma^2(V) + V$. For any $W \in \mathfrak{T}_{a+b,c}$, the fiber $f^{-1}(W)$ is contained in the set of a-dimensional subspaces V of the $(a+b-c)$-dimensional vector space $W \cap \sigma^{-2}(W)$. In particular, the fiber is empty unless $a \le a+b-c$, so $c \le b$.

Let $0 \le c \le \min(b,\, n-a-b)$ and $W \in \mathfrak{T}_{a+b,c}$. The fiber $f^{-1}(W)$ embeds into the irreducible variety $\mathrm{Gr}_a(W \cap \sigma^{-2}(W)) \simeq \mathrm{Gr}_a(\overline{\mathbb{F}}_p^{a+b-c})$ and we have

$$\begin{aligned}&\dim \mathfrak{T}_{a,b} - \dim \mathrm{Gr}_a(\overline{\mathbb{F}}_p^{a+b-c}) - \dim \mathfrak{T}_{a+b,c} \\ &\overset{\text{(b)}}{=} b(n-b) - a(b-c) - c(n-c) \\ &= (b-c)(n-a-b-c).\end{aligned}$$

The right-hand side is positive for all c except $c = \min(b, n-a-b)$, for which it is zero. For this value of c, the assumption $0 < b < \min(a, n-a)$ implies that $a+b > a$ and $c > 0$, so $\mathfrak{T}_{a+b,c}$ is irreducible by the induction hypothesis.

The claim follows by applying Lemma 4.8 to the regular map f. $\square$

Remark 4.10. For $b = 0$, the variety $\mathfrak{T}_{a,0}$ consists of those a-dimensional subspaces $V \subseteq \overline{\mathbb{F}}_p^n$ such that $\sigma(V) = \sigma^{-1}(V)$, or, equivalently, of the finitely many a-dimensional subspaces of $\overline{\mathbb{F}}_p^n$ defined over $\mathbb{F}_{p^2}$. In particular, $\mathfrak{T}_{a,0}$ is not irreducible unless $a \in \{0, n\}$.

Proposition 4.11. *Let $n \geq 3$ and $k = [n/3]$. Let $\mathcal{Q} = \mathcal{O}_n$ be the octopus quiver with $k+1$ vertices and n edges. Then, the sets $\mathfrak{Z}_{\mathcal{Q}}$ and $\mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ are irreducible.*

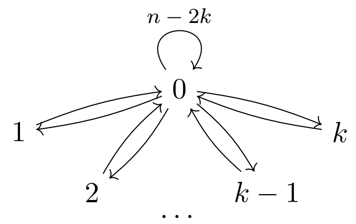

Proof. By Lemma 4.4(a), we have

$$\dim \mathfrak{Z}_{\mathcal{Q}} = \sum_i d_{\mathcal{Q}}(i)^2 - \sum_{i,j} |\mathcal{Q}(i,j)|^2 = (n-k)^2 + k - (n-2k)^2 - 2k = k(2n - 3k - 1).$$

Points in $\mathfrak{Z}_{\mathcal{Q}}$ correspond to tuples $(V_0, V_1, \ldots, V_k)$ of subspaces of $\overline{\mathbb{F}}_p^n$ of respective dimensions $n - k, 1, \ldots, 1$ together spanning $\overline{\mathbb{F}}_p^n$ such that $\dim(V_0 \cap \sigma(V_0)) = n - 2k$ and $V_1, \ldots, V_k \subseteq \sigma(V_0) \cap \sigma^{-1}(V_0)$ and $V_i \neq \sigma(V_j)$ for all $i, j \in \{1, \ldots, k\}$. We are going to apply Lemma 4.8 to the regular map

$$f\colon \mathfrak{Z}_{\mathcal{Q}} \to \bigsqcup_{0 \leq b \leq \min(n-k,k)} \mathfrak{T}_{n-k,b} = \mathrm{Gr}_{n-k}(\overline{\mathbb{F}}_p^n)$$

sending $(V_0, V_1, \ldots, V_k)$ to V_0.

Let $0 \leq b \leq \min(n-k, k)$ and consider an arbitrary $V_0 \in \mathfrak{T}_{n-k,b}$. The fiber $f^{-1}(V_0)$ consists of tuples $(V_1, \ldots, V_k)$ of linearly independent one-dimensional subspaces of the $(n-k-b)$-dimensional vector space $\sigma(V_0) \cap \sigma^{-1}(V_0)$. In particular, the fiber is empty unless $k \leq n - k - b$, i.e., $b \leq n - 2k$. We now assume that $b \leq n - 2k$. The fiber $f^{-1}(V_0)$ embeds into the irreducible variety $\Big(\mathbb{P}\big(\sigma(V_0) \cap \sigma^{-1}(V_0)\big)\Big)^k \simeq \Big(\mathbb{P}(\overline{\mathbb{F}}_p^{n-k-b})\Big)^k$, and by Lemma 4.9(b) we have

$$\begin{aligned} &\dim \mathfrak{Z}_{\mathcal{Q}} - \dim\Big(\mathbb{P}(\overline{\mathbb{F}}_p^{n-k-b})\Big)^k - \dim \mathfrak{T}_{n-k,b} \\ &= k(2n - 3k - 1) - k(n - k - b - 1) - b(n - b) \\ &= (n - 2k - b)(k - b). \end{aligned}$$

The right-hand side is positive for all b except $b = \min(n - 2k, k)$, for which it is zero. For this value of b, the assumption $n \geq 3$ together with the definition $k = [n/3]$ imply that $b > 0$, so $\mathfrak{T}_{n-k,b}$ is irreducible by Lemma 4.9(c). By Lemma 4.8, the variety $\mathfrak{Z}_{\mathcal{Q}}$ is irreducible. Since $\mathfrak{Y}_{\mathcal{Q}}$ and $\mathfrak{Z}_{\mathcal{Q}}$ are irreducible, so is their product and therefore so is the image $\mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$. $\square$

4.6. The special case $n = 4$ (the dumbbell variety)

When $n = 4$, Proposition 4.6 shows that there are two isomorphism classes of quivers $\mathcal{Q} \in \mathrm{Bal}_4$ such that $\dim \mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ reaches the maximal value 6, namely the octopus quiver $\mathcal{O}_4$ (for which $\mathfrak{X}_{\mathcal{O}_4}^{\mathrm{diag}}$ is irreducible by Proposition 4.11), and the dumbbell quiver $\mathcal{Q}$:

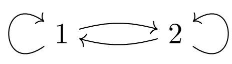

The goal of this subsection is to prove:

Proposition 4.12. *When $\mathcal{Q}$ is the dumbbell quiver, the sets $\mathfrak{Z}_{\mathcal{Q}}$ and $\mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ are irreducible.*

By Lemma 4.4(a), the set $\mathfrak{Z}_{\mathcal{Q}}$ has pure dimension 4. The points of $\mathfrak{Z}_{\mathcal{Q}}$ correspond to pairs $V = (V_1, V_2)$ of two-dimensional subspaces of $\overline{\mathbb{F}}_p^4$ such that $V_1 \oplus V_2 = \overline{\mathbb{F}}_p^4$ and $\dim(V_i \cap \sigma(V_j)) = 1$ for each $i, j \in \{1, 2\}$. For any $V = (V_1, V_2) \in \mathfrak{Z}_{\mathcal{Q}}$, define the one-dimensional vector spaces

$$L_1(V) \coloneqq V_1 \cap \sigma(V_1) \qquad \text{and} \qquad L_2(V) \coloneqq V_2 \cap \sigma(V_2),$$

the three-dimensional vector space

$$W(V) \coloneqq V_1 + \sigma(V_2),$$

and the vector spaces

$$U(V) \coloneqq W(V) \cap \sigma(W(V)),$$

and

$$M(V) \coloneqq U(V) \cap \sigma(U(V)) = W(V) \cap \sigma(W(V)) \cap \sigma^2(W(V)).$$

Since $W(V)$ has codimension 1 in $\overline{\mathbb{F}}_p^4$, we have $\dim U(V) \geq 2$ and $\dim M(V) \geq 1$. The space $W(V)$ is not defined over $\mathbb{F}_p$ as otherwise we would have $V_1 + V_2 \subseteq W(V) \subsetneq \overline{\mathbb{F}}_p^4$. This implies that $U(V) \subsetneq W(V)$ is two-dimensional.

Note that

$$L_1(V) \subseteq U(V) \qquad \text{and} \qquad L_2(V) \subseteq \sigma^{-1}(U(V)). \tag{4.7}$$

Strategy. Our strategy of proof for Proposition 4.12 is as follows: we show that for a generic element V of any irreducible component of $\mathfrak{Z}_{\mathcal{Q}}$, none of the subspaces $L_1(V)$, $L_2(V)$, $U(V)$, $M(V)$ are defined over $\mathbb{F}_p$. Disregarding those "exceptional" V for which any of these subspaces are defined over $\mathbb{F}_p$, we show that $M(V)$ is one-dimensional, and that the fibers of the map $V \mapsto M(V)$ embed into one-dimensional subvarieties of $\mathbb{P}^1(\overline{\mathbb{F}}_p) \times \mathbb{P}^1(\overline{\mathbb{F}}_p)$. Using Newton polygons, we show that these one-dimensional varieties are generically irreducible. Finally, we conclude using Lemma 4.8.

Lemma 4.13. *Consider the regular map*

$$\lambda\colon \mathfrak{Z}_{\mathcal{Q}} \to \mathbb{P}(\overline{\mathbb{F}}_p^4) \times \mathbb{P}(\overline{\mathbb{F}}_p^4), \qquad V \mapsto (L_1(V), L_2(V)).$$

Let F_λ be the closed subset of $\mathbb{P}(\overline{\mathbb{F}}_p^4) \times \mathbb{P}(\overline{\mathbb{F}}_p^4)$ corresponding to pairs (L_1, L_2) such that at least one of L_1 or L_2 is defined over $\mathbb{F}_p$, and let $\mathfrak{Z}'_{\mathcal{Q}} \coloneqq \mathfrak{Z}_{\mathcal{Q}} \setminus \lambda^{-1}(F_\lambda)$. Then:

(a) The closed subset $\lambda^{-1}(F_\lambda)$ of $\mathfrak{Z}_{\mathcal{Q}}$ is at most three-dimensional.

(b) For any $V \in \mathfrak{Z}'_{\mathcal{Q}}$, we have:

- *(i) $V_i = L_i(V) \oplus \sigma^{-1}(L_i(V))$ for each $i \in \{1, 2\}$.*
- *(ii) $U(V) + \sigma^{-1}(U(V)) + \sigma^{-2}(U(V)) = \overline{\mathbb{F}}_p^4$.*
- *(iii) The vector space $M(V)$ is one-dimensional.*

Proof.

(a) As both cases are symmetric, we can focus on the preimage of the space of pairs where L_1 is defined over $\mathbb{F}_p$. By Lemma 4.4(b), this preimage has dimension strictly less than $\dim \mathfrak{Z}_{\mathcal{Q}} = 4$.

(b) (i) By definition, $V_i \supseteq L_i(V) + \sigma^{-1}(L_i(V))$. By hypothesis, $L_i(V)$ is a one-dimensional space not defined over $\mathbb{F}_p$, so it has trivial intersection with $\sigma^{-1}(L_i(V))$, so the right-hand side is a direct sum and has dimension $2 = \dim V_i$, so the inclusion is an equality.

(ii) Combining (i) with Equation (4.7), we obtain $V_1 + V_2 \subseteq U(V) + \sigma^{-1}(U(V)) + \sigma^{-2}(U(V))$. The left-hand side is $\overline{\mathbb{F}}_p^4$ since $V \in \mathfrak{Z}_{\mathcal{Q}}$.

(iii) From (i), we see that the two-dimensional vector space $U(V)$ is not defined over $\mathbb{F}_p$. Thus, the vector space $M(V) = U(V) \cap \sigma(U(V))$ is at most one-dimensional, but it is also at least one-dimensional since it equals $W(V) \cap \sigma(W(V)) \cap \sigma^2(W(V))$ and $\dim W(V) = 3$. $\square$

Consider the regular map

$$\upsilon \colon \mathfrak{Z}'_Q \to \mathrm{Gr}_2(\overline{\mathbb{F}}_p^4), \qquad V \mapsto U(V).$$

If $U = \langle v, u \rangle$ is any two-dimensional subspace of $\overline{\mathbb{F}}_p^4$, then Equation (4.7) shows that there is a regular map

$$\varphi_{v,u} \colon \upsilon^{-1}(U) \to \mathbb{P}^1(\overline{\mathbb{F}}_p) \times \mathbb{P}^1(\overline{\mathbb{F}}_p), \qquad V \mapsto ([r_1 : s_1], [r_2 : s_2])$$

uniquely characterized by

$$L_1(V) = \langle r_1 v + s_1 u \rangle \qquad \text{and} \qquad L_2(V) = \langle r_2 \sigma^{-1}(v) + s_2 \sigma^{-1}(u) \rangle. \tag{4.8}$$

This map $\varphi_{v,u}$ is injective as $V_i = L_i(V) \oplus \sigma^{-1}(L_i(V))$ by Lemma 4.13(b)(i).

Let S be the (dense open) subset of $\overline{\mathbb{F}}_p^4$ consisting of those $m \in \overline{\mathbb{F}}_p^4$ for which the vectors $\sigma^i(m)$ for $i = 0, \ldots, 3$ are linearly independent, and let $g \colon S \to \overline{\mathbb{F}}_p^4$ be the map sending m to the unique tuple $(c_0, \ldots, c_3) \in \overline{\mathbb{F}}_p^4$ satisfying $\sigma^4(m) = \sum_{i=0}^{3} c_i \sigma^i(m)$. The map g is regular by Cramer's rule. Finally, for any $\underline{c} = (c_0, \ldots, c_3) \in \overline{\mathbb{F}}_p^4$, define the following (one-dimensional) closed subset $D_{\underline{c}} \subseteq \mathbb{P}^1(\overline{\mathbb{F}}_p) \times \mathbb{P}^1(\overline{\mathbb{F}}_p)$:

$$D_{\underline{c}} := \left\{ ([r_1 : s_1], [r_2 : s_2]) \;\middle|\; -c_0 r_1^{p+1} r_2^{p+1} + c_1 r_1^{p+1} r_2^p s_2 - c_2 r_1^{p+1} s_2^{p+1} + c_3 r_1^p s_1 s_2^{p+1} + s_1^{p+1} s_2^{p+1} = 0 \right\} \tag{4.9}$$

Lemma 4.14. *Consider the regular map (see Lemma 4.13(b)(iii))*

$$\mu \colon \mathfrak{Z}'_Q \to \mathbb{P}(\overline{\mathbb{F}}_p^4), \qquad V \mapsto M(V).$$

Let F_μ be the closed (finite) subset of $\mathbb{P}(\overline{\mathbb{F}}_p^4)$ corresponding to subspaces M which are defined over $\mathbb{F}_p$, and let $\mathfrak{Z}''_Q := \mathfrak{Z}'_Q \setminus \mu^{-1}(F_\mu)$.

(a) If $M \in F_\mu$, then the closed subset $\mu^{-1}(M)$ of $\mathfrak{Z}'_Q$ is at most three-dimensional.

(b) If $M = \langle m \rangle \in \mathbb{P}(\overline{\mathbb{F}}_p^4) \setminus F_\mu$, then the closed subset $\mu^{-1}(M)$ of $\mathfrak{Z}'_Q$ is at most one-dimensional. More specifically, if $\mu^{-1}(M)$ is non-empty, then m lies in S and there is an injective regular map $\mu^{-1}(M) \hookrightarrow D_{g(m)}$ (where $D_{g(m)}$ is as in Equation (4.9)).

Proof. The proofs of (a) and (b) are very similar, the main difference being that for fixed M, in (b), there is only one possible vector space $U(V)$, whereas in (a), there is a two-dimensional set of possible vector spaces $U(V)$.

(a) Since M is defined over $\mathbb{F}_p$, we pick a σ-invariant generator $m \in (M \cap \mathbb{F}_p^4) \setminus \{0\}$ of M. For any $V \in \mu^{-1}(M)$, the two-dimensional vector space $U(V)$ contains M by definition. As $\{U \in \mathrm{Gr}_2(\overline{\mathbb{F}}_p^4) \mid M \subseteq U\} \simeq \mathbb{P}(\overline{\mathbb{F}}_p^4/M)$ is two-dimensional, it suffices to show that the image of the injective map $\varphi_{m,u} \colon \upsilon^{-1}(U) \to \mathbb{P}^1(\overline{\mathbb{F}}_p) \times \mathbb{P}^1(\overline{\mathbb{F}}_p)$ is at most one-dimensional for any $U = \langle m, u \rangle$ containing M. Let $U = \langle m, u \rangle$ be a two-dimensional subspace of $\overline{\mathbb{F}}_p^4$ containing M, and assume that $\upsilon^{-1}(U)$ is non-empty.

By Lemma 4.13(b)(ii), this implies $\overline{\mathbb{F}}_p^4 = U + \sigma^{-1}(U) + \sigma^{-2}(U) = \sigma^{-2}(\langle m, u, \sigma(u), \sigma^2(u) \rangle)$, so the vectors $m, u, \sigma(u), \sigma^2(u)$ form a basis of $\overline{\mathbb{F}}_p^4$. Write $\sigma^3(u) = \sum_{i=0}^{2} c_i \sigma^i(u) + c_3 m$ with $c_0, \ldots, c_3 \in \overline{\mathbb{F}}_p$.

For any $V \in \upsilon^{-1}(U)$, letting $\varphi_{m,u}(V) = ([r_1 : s_1], [r_2 : s_2])$, since $\sigma(V_1) \cap V_2 \neq 0$, we must have $\sigma^3(V_1) \cap \sigma^2(V_2) \neq 0$, where according to Lemma 4.13(b)(i) and Equation (4.8):

$$\begin{aligned}\sigma^3(V_1) &= \sigma^3(L_1(V)) + \sigma^2(L_1(V)) = \langle r_1^{p^3} m + s_1^{p^3}\sigma^3(u), \quad r_1^{p^2} m + s_1^{p^2}\sigma^2(u)\rangle \\ &= \langle (r_1^{p^3} + c_3 s_1^{p^3})m + c_0 s_1^{p^3} u + c_1 s_1^{p^3}\sigma(u) + c_2 s_1^{p^3}\sigma^2(u), \quad r_1^{p^2} m + s_1^{p^2}\sigma^2(u)\rangle, \\ \sigma^2(V_2) &= \sigma^2(L_2(V)) + \sigma(L_2(V)) = \langle r_2^{p^2} m + s_2^{p^2}\sigma(u), \quad r_2^p m + s_2^p u\rangle.\end{aligned}$$

Writing everything in terms of the basis $(m, u, \sigma(u), \sigma^2(u))$, this means that the matrix

$$\begin{pmatrix} r_1^{p^3} + c_3 s_1^{p^3} & c_0 s_1^{p^3} & c_1 s_1^{p^3} & c_2 s_1^{p^3} \\ r_1^{p^2} & & & s_1^{p^2} \\ r_2^{p^2} & & s_2^{p^2} & \\ r_2^p & s_2^p & & \end{pmatrix}$$

must be singular, so its determinant must vanish. This determinant is a non-zero polynomial in r_1, s_1, r_2, s_2 (it always involves the summand $r_1^{p^3} s_1^{p^2} s_2^{p^2+p}$), which shows that the image of $\varphi_{m,u}$ in $\mathbb{P}^1(\overline{\mathbb{F}}_p) \times \mathbb{P}^1(\overline{\mathbb{F}}_p)$ is indeed at most one-dimensional.

(b) Let $M = \langle m \rangle \in \mathbb{P}(\overline{\mathbb{F}}_p^4) \setminus F_\mu$. If $\mu^{-1}(M)$ is empty, the claims are clear, so we assume that $\mu^{-1}(M)$ is non-empty. For any $V \in \mu^{-1}(M)$, we have $U(V) \supseteq M(V) + \sigma^{-1}(M(V))$ by definition; since the one-dimensional space $M(V) = M$ is not defined over $\mathbb{F}_p$ and since $U(V)$ is two-dimensional, we in fact have

$$U(V) = M(V) \oplus \sigma^{-1}(M(V)) = \langle m, \sigma^{-1}(m)\rangle,$$

so $\mu^{-1}(M) = \upsilon^{-1}(U)$ where $U \coloneqq \langle m, \sigma^{-1}(m)\rangle$.

By Lemma 4.13(b)(ii), and because $\upsilon^{-1}(U)$ is non-empty, we have $\overline{\mathbb{F}}_p^4 = U + \sigma^{-1}(U) + \sigma^{-2}(U) = \sigma^{-3}(\langle m, \sigma(m), \sigma^2(m), \sigma^3(m)\rangle)$, so the vectors $m, \sigma(m), \sigma^2(m), \sigma^3(m)$ form a basis of $\overline{\mathbb{F}}_p^4$, i.e., m lies in S. Let $(c_0, \ldots, c_3) \coloneqq g(m) \in \overline{\mathbb{F}}_p^4$, so that by definition $\sigma^4(m) = \sum_{i=0}^3 c_i \sigma^i(m)$.

For any $V \in \mu^{-1}(M)$, letting $\varphi_{m,\sigma^{-1}(m)}(V) = ([r_1 : s_1], [r_2 : s_2])$, since $\sigma(V_1) \cap V_2 \neq 0$, we must have $\sigma^4(V_1) \cap \sigma^3(V_2) \neq 0$, where according to Lemma 4.13(b)(i) and Equation (4.8):

$$\begin{aligned}\sigma^4(V_1) &= \sigma^4(L_1(V)) + \sigma^3(L_1(V)) = \langle r_1^{p^4}\sigma^4(m) + s_1^{p^4}\sigma^3(m), \quad r_1^{p^3}\sigma^3(m) + s_1^{p^3}\sigma^2(m)\rangle \\ &= \langle c_0 r_1^{p^4} m + c_1 r_1^{p^4}\sigma(m) + c_2 r_1^{p^4}\sigma^2(m) + (c_3 r_1^{p^4} + s_1^{p^4})\sigma^3(m), \quad r_1^{p^3}\sigma^3(m) + s_1^{p^3}\sigma^2(m)\rangle, \\ \sigma^3(V_2) &= \sigma^3(L_2(V)) + \sigma^2(L_2(V)) = \langle r_2^{p^3}\sigma^2(m) + s_2^{p^3}\sigma(m), \quad r_2^{p^2}\sigma(m) + s_2^{p^2} m\rangle.\end{aligned}$$

Writing everything in terms of the basis $(m, \sigma(m), \sigma^2(m), \sigma^3(m))$, this means that the matrix

$$\begin{pmatrix} c_0 r_1^{p^4} & c_1 r_1^{p^4} & c_2 r_1^{p^4} & c_3 r_1^{p^4} + s_1^{p^4} \\ & & s_1^{p^3} & r_1^{p^3} \\ & s_2^{p^3} & r_2^{p^3} & \\ s_2^{p^2} & r_2^{p^2} & & \end{pmatrix}$$

must be singular, so its determinant

$$-c_0 r_1^{p^4+p^3} r_2^{p^3+p^2} + c_1 r_1^{p^4+p^3} r_2^{p^3} s_2^{p^2} - c_2 r_1^{p^4+p^3} s_2^{p^3+p^2} + c_3 r_1^{p^4} s_1^{p^3} s_2^{p^3+p^2} + s_1^{p^4+p^3} s_2^{p^3+p^2}$$

must vanish. Letting $\tau \colon \mathbb{P}^1(\overline{\mathbb{F}}_p) \times \mathbb{P}^1(\overline{\mathbb{F}}_p) \to \mathbb{P}^1(\overline{\mathbb{F}}_p) \times \mathbb{P}^1(\overline{\mathbb{F}}_p)$ be the bijective regular map $([r_1 : s_1], [r_2 : s_2]) \mapsto ([r_1^{p^3} : s_1^{p^3}], [r_2^{p^2} : s_2^{p^2}])$, we have shown that the image of the injective regular map $\tau \circ \varphi_{m,\sigma^{-1}(m)} \colon \mu^{-1}(M) \to \mathbb{P}^1(\overline{\mathbb{F}}_p^4) \times \mathbb{P}^1(\overline{\mathbb{F}}_p^4)$ is contained in $D_{(c_0,\ldots,c_3)} = D_{g(m)}$. $\square$

Lemma 4.15. *There is a non-empty open subset $O' \subseteq \overline{\mathbb{F}}_p^4$ such that, for all $\underline{c} \in O'$, the closed subset $D_{\underline{c}} \subseteq \mathbb{P}^1(\overline{\mathbb{F}}_p) \times \mathbb{P}^1(\overline{\mathbb{F}}_p)$ is irreducible.*

Proof. Let f be the following bihomogeneous polynomial in the variables r_1, s_1, r_2, s_2, with coefficients in $\mathbb{F}_p(c_0, \dots, c_3)$:

$$f = -c_0 r_1^{p+1} r_2^{p+1} + c_1 r_1^{p+1} r_2^p s_2 - c_2 r_1^{p+1} s_2^{p+1} + c_3 r_1^p s_1 s_2^{p+1} + s_1^{p+1} s_2^{p+1}$$

Let $L = \overline{\mathbb{F}_p(c_0, \dots, c_3)}$. By [Stacks, Lemma 0559], it suffices to prove that the subscheme of $\mathbb{P}^1_L \times \mathbb{P}^1_L$ defined by $f = 0$ is irreducible, i.e., that f is irreducible as a bihomogeneous polynomial over L. We will show this by specializing to $c_0 = 0$. Assume by contradiction that there are non-constant bihomogeneous polynomials $g, h \in L[r_1, s_1, r_2, s_2]$ such that $f = gh$. Let v be an extension of the c_0-adic valuation on $\mathbb{F}_p(c_0, \dots, c_3)$ to L and let $\mathfrak{p} \subset \mathcal{O} \subset L$ be the corresponding maximal ideal and valuation ring. We have $\mathcal{O}/\mathfrak{p} = \overline{\mathbb{F}_p(c_1, \dots, c_3)}$. Since the coefficients of f lie in $\mathcal{O}$, we can by Gauss' lemma assume without loss of generality that the coefficients of g and h also lie in $\mathcal{O}$.[7]

The *Newton polygon* $\mathrm{NP}(a) \subset \mathbb{R}^2$ of a bihomogeneous polynomial $a = \sum_{i,j} k_{ij} r_1^i s_1^{d-i} r_2^j s_2^{e-j}$ with coefficients in an integral domain is the convex hull of the points $(i, j) \in \mathbb{Z}^2_{\geq 0}$ with $k_{ij} \neq 0$. For any two such polynomials a, b, the Newton polygon $\mathrm{NP}(ab)$ is the Minkowski sum of $\mathrm{NP}(a)$ and $\mathrm{NP}(b)$.

Over $\mathcal{O}/\mathfrak{p} = \overline{\mathbb{F}_p(c_1, \dots, c_3)}$, we have

$$(f \bmod \mathfrak{p}) = c_1 r_1^{p+1} r_2^p s_2 - c_2 r_1^{p+1} s_2^{p+1} + c_3 r_1^p s_1 s_2^{p+1} + s_1^{p+1} s_2^{p+1}.$$

The Newton polygon of f is the (solid) triangle with corners $(0,0)$, $(p+1, 0)$, $(p+1, p+1)$ and the Newton polygon of $(f \bmod \mathfrak{p})$ is the (dashed) triangle with corners $(0,0)$, $(p+1,0)$, $(p+1,p)$.

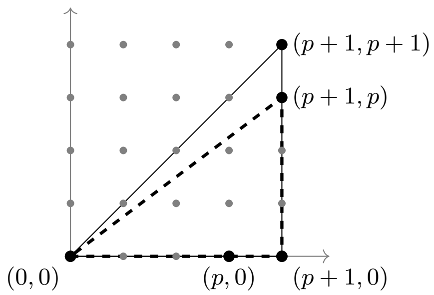

The line segment $[(0,0), (p+1, p)]$ contains no integer lattice points other than its endpoints. Since $\mathrm{NP}(f \bmod \mathfrak{p}) = \mathrm{NP}(g \bmod \mathfrak{p}) + \mathrm{NP}(h \bmod \mathfrak{p})$ and the corners of the Newton polygons $\mathrm{NP}(g \bmod \mathfrak{p})$ and $\mathrm{NP}(h \bmod \mathfrak{p})$ are non-negative integer lattice points, it follows that the Newton polygon of one of the factors (say $\mathrm{NP}(g \bmod \mathfrak{p})$) contains a translate of that line segment. Moreover, as all other edges of $\mathrm{NP}(f \bmod \mathfrak{p})$ are either horizontal or vertical, so are the other edges of $\mathrm{NP}(g \bmod \mathfrak{p})$. The only possibility is that $\mathrm{NP}(g \bmod \mathfrak{p}) = \mathrm{NP}(f \bmod \mathfrak{p})$, and then $\mathrm{NP}(g) \supseteq \mathrm{NP}(g \bmod \mathfrak{p}) = \mathrm{NP}(f \bmod \mathfrak{p})$.

We have $\mathrm{NP}(f) = \mathrm{NP}(g) + \mathrm{NP}(h)$, but the triangle $\mathrm{NP}(f)$ does not contain any proper translate of $\mathrm{NP}(f \bmod \mathfrak{p}) \subseteq \mathrm{NP}(g)$, so $\mathrm{NP}(h) = \{(0,0)\}$, i.e., h is a monomial of the form $k s_1^d s_2^e$. Clearly, such a monomial can only divide f if $d = e = 0$, so h must be constant. ☐

Corollary 4.16. *There is a dense open subset O of $\mathbb{P}(\overline{\mathbb{F}}_p^4)$ such that for all $M \in O$, there is an injective regular map from the fiber $\mu^{-1}(M)$ to a one-dimensional irreducible variety.*

[7] Let a and b be the smallest valuations of coefficients of g and h, respectively. Considering the lexicographically minimal monomials whose coefficients have these valuations and expanding the product, one can see that some coefficient of gh has valuation $a + b$. Since all coefficients of $f = gh$ lie in $\mathcal{O}$, this means that $a + b \geq 0$. Dividing g by an element of valuation a and multiplying h by the same element, we can ensure that the coefficients of g and h lie in $\mathcal{O}$.

Proof. All fibers of the map g are finite since they are cut out by the non-trivial polynomial equations $m_i^{p^4} = \sum_{i=0}^{3} c_i m_i^{p^i}$ in the coordinates $m_1, \ldots, m_4$ of m. Since $\dim S = 4 = \dim \overline{\mathbb{F}}_p^4$ and $\overline{\mathbb{F}}_p^4$ is irreducible, this implies that g is dominant. We have seen in Lemma 4.14(b) that for any $m \in S$ (in particular, $\langle m \rangle$ is not defined over $\mathbb{F}_p$), there is an injective regular map $\mu^{-1}(\langle m \rangle) \to D_{g(m)}$. (This is obviously true if $\mu^{-1}(\langle m \rangle)$ is empty.) Now, let O' be as in Lemma 4.15, so that $D_{g(m)}$ is irreducible when $g(m) \in O'$. The claim follows, taking O to be any dense open subset of the image of $g^{-1}(O') \subseteq S \subseteq \overline{\mathbb{F}}_p^4$ under the regular map $\overline{\mathbb{F}}_p^4 \to \mathbb{P}(\overline{\mathbb{F}}_p^4)$, $m \mapsto \langle m \rangle$. (The preimage $g^{-1}(O')$ is non-empty and open since O' is non-empty and open and g is dominant. Hence, its (constructible) image in $\mathbb{P}(\overline{\mathbb{F}}_p^4)$ is dense, so it contains a dense open subset.) $\square$

Proof of Proposition 4.12. The set $\mathfrak{Z}_{\mathcal{Q}}$ has pure dimension 4 by Lemma 4.4(a). Thus, Lemma 4.13(a) and Lemma 4.14(a) imply that the inclusions $\mathfrak{Z}''_{\mathcal{Q}} \subseteq \mathfrak{Z}'_{\mathcal{Q}} \subseteq \mathfrak{Z}_{\mathcal{Q}}$ are dense, so it suffices to prove that $\mathfrak{Z}''_{\mathcal{Q}}$ is irreducible. For this, fix O as in Corollary 4.16 (which is three-dimensional and whose complement is at most two-dimensional) and apply Lemma 4.8 to the map $\mu\colon \mathfrak{Z}''_{\mathcal{Q}} \to \mathbb{P}(\overline{\mathbb{F}}_p^4) \setminus F_\mu$. (The fiber $\mu^{-1}(M)$ embeds in a one-dimensional variety by Lemma 4.14(b), and that variety can be taken to be irreducible when $x \in O$ by Corollary 4.16.) $\square$

4.7. Conclusion

Theorem 4.17 (cf. Theorem 1.3)**.** *For any finite field $\mathbb{F}_q \supseteq \mathbb{F}_p$, we have*

$$|\mathfrak{X}^{\mathrm{diag}} \cap \mathfrak{M}_n(\mathbb{F}_q)| = c^{\mathrm{diag}}(p,n) \cdot q^{\lfloor n^2/3 \rfloor + 1} + O_{p,n}\Big(q^{\lfloor n^2/3 \rfloor + 1/2}\Big), \text{ where } c^{\mathrm{diag}}(p,n) = \begin{cases} p^2 & \text{if } n = 2, \\ 2 & \text{if } n = 4, \\ 1 & \text{if } n \notin \{2,4\}. \end{cases}$$

Proof. We have seen above that $\mathfrak{X}^{\mathrm{diag}}$ is a disjoint union of the finitely many constructible σ-invariant subsets $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}$. For all quivers $\mathcal{Q}$ with $\dim \mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}} \leq \lfloor n^2/3 \rfloor$, we have $|\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}} \cap \mathfrak{M}_n(\mathbb{F}_q)| = O_{p,n}(q^{\lfloor n^2/3 \rfloor})$ by the Schwarz–Zippel bound [LW54, Lemma 1]. Proposition 4.6 classifies the remaining quivers and shows that they all satisfy $\dim \mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}} = \lfloor n^2/3 \rfloor + 1$. In Propositions 4.7, 4.11 and 4.12, we have computed the number of irreducible components of $\mathfrak{X}^{\mathrm{diag}}_{\mathcal{Q}}$ in these cases, shown that they are all fixed by σ, and that the total number of irreducible components of dimension $\lfloor n^2/3 \rfloor + 1$ is precisely $c^{\mathrm{diag}}(p,n)$. The claim then follows from the Lang–Weil bound [LW54, Theorem 1]. $\square$

5. Towards general matrices commuting with their Frobenius

In this section, we relate the size of $\mathfrak{X} \cap \mathfrak{M}_n(\mathbb{F}_q)$ to the numbers $d(M)$ defined in Equation (1.1)., i.e., we prove Proposition 5.8 (which implies Theorem 1.4). To this end, we associate to any matrix in $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ a *Jordan shape*, encoding the sizes of all Jordan blocks associated to the eigenvalues.

Jordan shapes. A *Jordan shape* of size n is a pair $\mathcal{S} = (V, e)$ consisting of a finite set V and a map $e\colon V \times \mathbb{N} \to \mathbb{Z}_{\geq 0}$ such that $e(i,1) \geq 1$ and $e(i,1) \geq e(i,2) \geq \cdots$ for all $i \in V$ and such that $\sum_{i \in V} \sum_{k \geq 1} e(i,k) = n$. An *isomorphism* between Jordan shapes $\mathcal{S} = (V,e)$ and $\mathcal{S}' = (V', e')$ is a bijection $\pi\colon V \to V'$ such that $e(i,k) = e'(\pi(i), k)$ for all $i \in V$ and $k \geq 1$. We let JS_n be the (finite) set of isomorphism classes of Jordan shapes of size n.

Definition 5.1. To any matrix $M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$, we associate a Jordan shape $\mathcal{S}_M = (V_M, e_M)$ of size n as follows: the set V_M consists of the eigenvalues of M; for each eigenvalue λ and each $k \geq 1$, we let

$$e_M(\lambda, k) := \dim\Big(\ker(M - \lambda I_n)^k / \ker(M - \lambda I_n)^{k-1}\Big),$$

be the number of Jordan blocks of size at least k for this eigenvalue.

Two matrices M and M' are conjugate if and only if they have equal Jordan shapes, i.e., $V_M = V_{M'}$ and $e_M = e_{M'}$. Two matrices having *isomorphic* Jordan blocks, by contrast, may not have the same eigenvalues (for instance, M and $\sigma(M)$ always have isomorphic Jordan shapes via $\pi\colon \lambda \mapsto \sigma(\lambda)$).

The space of matrices with a given Jordan shape commuting with their Frobenius. For any Jordan shape $\mathcal{S} \in \mathrm{JS}_n$, we define the subset $\mathfrak{X}_{\mathcal{S}} \subseteq \mathfrak{X}$ of matrices $M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$ such that $\mathcal{S}_M \simeq \mathcal{S}$ and such that M commutes with $\sigma(M)$. Clearly,

$$\mathfrak{X} = \bigsqcup_{\mathcal{S} \in \mathrm{JS}_n} \mathfrak{X}_{\mathcal{S}}.$$

Remark 5.2. The sets $\mathfrak{X}_{\mathcal{S}}$ for $\mathcal{S} \in \mathrm{JS}_n$ defined here are related to the constructible sets $\mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ for $\mathcal{Q} \in \mathrm{Bal}_n$ defined in Subsection 4.1 as follows: if the shape $\mathcal{S} = (V, e)$ corresponds to diagonalizable matrices (meaning that $e(i,2) = 0$ for all $i \in V$), then $\mathfrak{X}_{\mathcal{S}}$ is the union of the sets $\mathfrak{X}_{\mathcal{Q}}^{\mathrm{diag}}$ over all quivers $\mathcal{Q} \in \mathrm{Bal}_n$ whose vertex set $V(\mathcal{Q})$ is V and whose degrees satisfy $d_{\mathcal{Q}}(i) = e(i,1)$ for all $i \in V$.

For any matrix $M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$, denote by $\mathrm{Cent}\, M$ its centralizer and by $\mathrm{Cl}\, M$ its conjugacy class. Note that $\mathrm{Cent}\, M$ is a subalgebra of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ and that $\mathrm{Cl}\, M$ is a constructible subset of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$.

Now, fix a shape $\mathcal{S} = (V, e)$, say with $V = \{1, \dots, r\}$. Let $\mathfrak{Y}_{\mathcal{S}} \subseteq \overline{\mathbb{F}}_p^r$ be the (open) subset formed of tuples $\lambda = (\lambda_1, \dots, \lambda_r)$ of distinct elements of $\overline{\mathbb{F}}_p$. For any $\lambda \in \mathfrak{Y}_{\mathcal{S}}$, we define a matrix $A_{\mathcal{S},\lambda}$ of shape $\mathcal{S}$ as follows: $A_{\mathcal{S},\lambda}$ is the matrix in Jordan normal form having $e(i,k) - e(i,k+1)$ Jordan blocks of size k associated to each eigenvalue λ_i, where we put the Jordan blocks for eigenvalue λ_i before those for eigenvalue λ_j if $i < j$, and we order blocks with the same eigenvalue by their size.

Lemma 5.3. *For any $\lambda, \lambda' \in \mathfrak{Y}_{\mathcal{S}}$, we have $\mathrm{Cent}\, A_{\mathcal{S},\lambda} = \mathrm{Cent}\, A_{\mathcal{S},\lambda'}$. We denote the corresponding subalgebra of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ by $\mathrm{Cent}\, \mathcal{S}$.*

Proof. For any $i \in \{1, \dots, n\}$, the generalized eigenspace of $A_{\mathcal{S},\lambda}$ with eigenvalue λ_i is also the generalized eigenspace of $A_{\mathcal{S},\lambda'}$ with eigenvalue λ'_i. Denote this common generalized eigenspace by G_i. We have $A_{\mathcal{S},\lambda'} v = A_{\mathcal{S},\lambda} v + (\lambda'_i - \lambda_i) v$ for all $v \in G_i$. The claim follows since any matrix commuting with $A_{\mathcal{S},\lambda}$ or $A_{\mathcal{S},\lambda'}$ preserves the generalized eigenspaces. □

Remark 5.4. The centralizer $\mathrm{Cent}\, \mathcal{S}$ admits an explicit description (some coefficients have to vanish, and some others must be equal), see [Gan53, Chap. VIII, §2]. Its dimension is $\sum_{i=1}^r \sum_{k \geq 1} e(i,k)^2$.

Corollary 5.5. *The set of matrices $U \in \mathrm{GL}_n(\overline{\mathbb{F}}_p)$ such that $A_{\mathcal{S},\lambda}$ commutes with $U A_{\mathcal{S},\widetilde{\lambda}} U^{-1}$ does not depend on the choice of $\lambda, \widetilde{\lambda} \in \mathfrak{Y}_{\mathcal{S}}$. We denote this closed subset of $\mathrm{GL}_n(\overline{\mathbb{F}}_p)$ by $\mathfrak{D}_{\mathcal{S}}$.*

Proof. This follows from Lemma 5.3 due to the following equivalences:

$$\begin{array}{ccc} A_{\mathcal{S},\lambda} \text{ commutes with } U A_{\mathcal{S},\widetilde{\lambda}} U^{-1} & \Longleftrightarrow & U A_{\mathcal{S},\widetilde{\lambda}} U^{-1} \in \mathrm{Cent}\, \mathcal{S} \text{ (independent of } \lambda) \\ \Updownarrow & & \\ U^{-1} A_{\mathcal{S},\lambda} U \text{ commutes with } A_{\mathcal{S},\widetilde{\lambda}} & \Longleftrightarrow & U^{-1} A_{\mathcal{S},\lambda} U \in \mathrm{Cent}\, \mathcal{S} \text{ (independent of } \widetilde{\lambda}) \end{array}$$

□

Proposition 5.6. *For any Jordan shape $\mathcal{S} = (V, e)$, the set $\mathfrak{X}_{\mathcal{S}}$ is a constructible subset of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ of dimension $|V| + \dim \mathfrak{D}_{\mathcal{S}} - \dim \mathrm{Cent}\, \mathcal{S}$.*

Proof. As before, we may assume that $V = \{1, \dots, r\}$. Let M be a matrix such that we have an isomorphism $\pi\colon \mathcal{S} \to \mathcal{S}_M$. Then, taking $\lambda := (\pi(1), \dots, \pi(r))$, we see that M must be conjugate to $A_{\mathcal{S},\lambda}$. Write $M = U A_{\mathcal{S},\lambda} U^{-1}$. Then, M commutes with $\sigma(M) = \sigma(U) A_{\mathcal{S},\sigma(\lambda)} \sigma(U)^{-1}$ if and only if $A_{\mathcal{S},\lambda}$ commutes with $(U^{-1}\sigma(U)) A_{\mathcal{S},\sigma(\lambda)} (U^{-1}\sigma(U))^{-1}$, i.e., if and only if $\wp(U) := U^{-1}\sigma(U)$ lies in $\mathfrak{D}_{\mathcal{S}}$. We have shown that the regular map

$$\mathfrak{Y}_{\mathcal{S}} \times \wp^{-1}(\mathfrak{D}_{\mathcal{S}}) \to \mathfrak{M}_n(\overline{\mathbb{F}}_p), \qquad (\lambda, U) \mapsto U A_{\mathcal{S},\lambda} U^{-1}$$

has image $\mathfrak{X}_{\mathcal{S}}$. In particular, $\mathfrak{X}_{\mathcal{S}}$ is constructible by Chevalley's theorem. Each fiber is the union of $|\mathrm{Aut}(\mathcal{S})|$ sets of the form $\{(\lambda, US) \mid S \in (\mathrm{Cent}\,\mathcal{S})^\times\}$ where $(\lambda, U) \in \mathfrak{Y}_{\mathcal{S}} \times \wp^{-1}(\mathfrak{D}_{\mathcal{S}})$, hence has dimension $\dim(\mathrm{Cent}\,\mathcal{S})^\times = \dim \mathrm{Cent}\,\mathcal{S}$. By Lemma 4.3, we have $\dim \wp^{-1}(\mathfrak{D}_{\mathcal{S}}) = \dim \mathfrak{D}_{\mathcal{S}}$. Thus,

$$\dim \mathfrak{X}_{\mathcal{S}} = \dim \mathfrak{Y}_{\mathcal{S}} + \dim \wp^{-1}(\mathfrak{D}_{\mathcal{S}}) - \dim \mathrm{Cent}\,\mathcal{S} = |V| + \dim \mathfrak{D}_{\mathcal{S}} - \dim \mathrm{Cent}\,\mathcal{S}. \qquad \square$$

Lemma 5.7. *For any matrix $M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$ with Jordan shape $\mathcal{S}_M \simeq \mathcal{S}$, the subset* $\mathrm{Cent}\,M \cap \mathrm{Cl}\,M$ *of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ has pure dimension* $\dim \mathfrak{D}_{\mathcal{S}} - \dim \mathrm{Cent}\,\mathcal{S}$.

Proof. Replacing M by a conjugate, we can assume without loss of generality that $M = A_{\mathcal{S},\lambda}$ for some $\lambda \in \mathfrak{Y}_{\mathcal{S}}$. Then, the regular map

$$\mathfrak{D}_{\mathcal{S}} \to \mathfrak{M}_n(\overline{\mathbb{F}}_p), \qquad U \mapsto U A_{\mathcal{S},\lambda} U^{-1}$$

has image $\mathrm{Cent}\,M \cap \mathrm{Cl}\,M$, and each fiber is a left coset of $(\mathrm{Cent}\,\mathcal{S})^\times$. $\square$

For any matrix $M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$, let

$$d(M) := (\text{number of distinct eigenvalues of } M) + \dim\big(\mathrm{Cent}\,M \cap \mathrm{Cl}\,M\big).$$

Proposition 5.8. *Let $\mathcal{S}$ be any Jordan shape of size n and let $M \in \mathfrak{M}_n(\overline{\mathbb{F}}_p)$ be any matrix with $\mathcal{S}_M \simeq \mathcal{S}$. Then, there is an integer $c \geq 1$ and a finite field $\mathbb{F}_{q_0} \supseteq \mathbb{F}_p$ such that:*

(a) $|\mathfrak{X}_{\mathcal{S}} \cap \mathfrak{M}_n(\mathbb{F}_q)| \leq c \cdot q^{d(M)} + O_{p,n}(q^{d(M)-1/2})$ *for all finite fields $\mathbb{F}_q \supseteq \mathbb{F}_p$.*

(b) $|\mathfrak{X}_{\mathcal{S}} \cap \mathfrak{M}_n(\mathbb{F}_q)| = c \cdot q^{d(M)} + O_{p,n}(q^{d(M)-1/2})$ *for all finite fields $\mathbb{F}_q \supseteq \mathbb{F}_{q_0}$.*

Proof. By Proposition 5.6 and Lemma 5.7, the constructible set $\mathfrak{X}_{\mathcal{S}}$ has dimension $d(M)$. The claims follow from the Lang–Weil bound [LW54, Theorem 1], where c is the number of $d(M)$-dimensional irreducible components of $\mathfrak{X}_{\mathcal{S}}$, and $\mathbb{F}_{q_0}$ is any finite field over which these irreducible components are all defined. $\square$

Theorem 1.4 follows from Proposition 5.8 by summing over all shapes corresponding to non-diagonalizable matrices.

Remark 5.9. We do not know whether for any $n \geq 3$, there is a non-diagonalizable matrix M for which $d(M)$ is larger than or equal to the exponent $\lfloor n^2/3 \rfloor + 1$ we obtained for diagonalizable matrices in Theorem 1.3. The largest value which we have been able to obtain for nilpotent matrices is $d(M) = \lfloor n(n-1)/3 \rfloor + 1$, for the nilpotent matrix M with one Jordan block of size $[n/3] + 1$ and $n - [n/3] - 1$ Jordan blocks of size 1.

Remark 5.10. Some computations of $\dim\big(\mathrm{Cent}\,\mathcal{S} \cap \mathrm{Cl}\,A_{\mathcal{S},\lambda}\big)$ exist in the literature, centered mostly around the nilpotent case (i.e., $r = 1$, $\lambda = (0)$). In particular, in that case, an upper bound is given by the dimension of the space of nilpotent matrices in $\mathrm{Cent}\,\mathcal{S}$, that is $\sum_{k \geq 1} e(0,k)^2 - e(0,1)$, and equality holds if and only if $\mathcal{S}$ is *self-large*, meaning that $e(0,k) - e(0,k+2) \leq 1$ for all k, i.e., any two distinct Jordan blocks have sizes differing by at least 2. (In that case, a generic nilpotent matrix in $\mathrm{Cent}\,\mathcal{S}$ automatically has shape $\mathcal{S}$.) We refer to [Pan08] for details concerning this case.

6. Matrices with eigenspaces defined over $\mathbb{F}_p$ and commuting with their Frobenius

In this section, in order to illustrate the principle described in Section 5, we deal with a special case: the set $\mathfrak{X}^{\mathrm{eig.}/\mathbb{F}_p}$ of matrices $M \in \mathfrak{X}$ whose eigenspaces $\ker(M - \lambda I_n)$ are all defined over $\mathbb{F}_p$. Specifically, we determine the asymptotics of $|\mathfrak{X}^{\mathrm{eig.}/\mathbb{F}_p} \cap \mathfrak{M}_n(\mathbb{F}_q)|$, i.e., we prove Theorem 6.9 (which is Theorem 1.5).

This case is made accessible by the following observation:

Lemma 6.1. *Let A and B be two commuting matrices in $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$.*

(i) If $\ker A \subseteq \ker B$, then $\ker A^k \subseteq \ker B^k$ for all $k \geq 1$.

(ii) If $\ker(A - \lambda I_n) = \ker(B - \lambda I_n)$ for all $\lambda \in \overline{\mathbb{F}}_p$, then $\ker(A - \lambda I_n)^k = \ker(B - \lambda I_n)^k$ for all $\lambda \in \overline{\mathbb{F}}_p$ and $k \geq 1$. In particular, the matrices A and B are conjugate.

Proof. We prove (i) by induction on k: the case $k = 1$ is clear. Let $k \geq 2$ and assume that $\ker A^{k-1} \subseteq \ker B^{k-1}$. Let $x \in \ker A^k$. Then, $A(x) \in \ker A^{k-1} \subseteq \ker B^{k-1}$, so $AB^{k-1}(x) = B^{k-1}A(x) = 0$, so $B^{k-1}(x) \in \ker A \subseteq \ker B$, so $B^k(x) = 0$.

For (ii), we reason for a fixed λ. Subtracting λI_n from A and B, we may assume that $\lambda = 0$. The inclusion $\ker A^k \subseteq \ker B^k$ and the reverse inclusion then both follow from (i). □

Corollary 6.2. *If $M \in \mathfrak{X}^{\mathrm{eig.}/\mathbb{F}_p}$, then the generalized eigenspaces $\ker(M - \lambda_i I_n)^k$ of M are all defined over $\mathbb{F}_p$.*

Proof. The space $\ker(M - \lambda_i I_n)^k$ is defined over $\mathbb{F}_p$ if and only if $\ker(M - \lambda_i I_n)^k = \ker(\sigma(M) - \sigma(\lambda_i)I_n)^k$. Since $M \in \mathfrak{X}^{\mathrm{eig.}/\mathbb{F}_p}$, the matrices $M - \lambda_i I_n$ and $\sigma(M) - \sigma(\lambda_i)I_n$ commute and have equal kernels (this is the case $k = 1$). Both inclusions between $\ker(M - \lambda_i I_n)^k$ and $\ker(\sigma(M) - \sigma(\lambda_i)I_n)^k$ then follow from Lemma 6.1(i). □

For each Jordan shape $\mathcal{S} = (\{1, \dots, r\}, e)$, let $\mathfrak{X}_{\mathcal{S}}^{\mathrm{eig.}/\mathbb{F}_p}$ be the subset of $\mathfrak{X}^{\mathrm{eig.}/\mathbb{F}_p}$ consisting of those matrices whose Jordan shape is isomorphic to $\mathcal{S}$. Note that $\mathfrak{X}^{\mathrm{eig.}/\mathbb{F}_p} = \bigsqcup_{\mathcal{S} \in \mathrm{JS}_n} \mathfrak{X}_{\mathcal{S}}^{\mathrm{eig.}/\mathbb{F}_p}$.

Proposition 6.3. *Let $\mathcal{S} = (\{1, \dots, r\}, e)$ be a Jordan shape, and let $\lambda = (\lambda_1, \dots, \lambda_r) \in \mathfrak{Y}_{\mathcal{S}}$. The set $\mathfrak{X}_{\mathcal{S}}^{\mathrm{eig.}/\mathbb{F}_p}$ is a non-empty constructible subset of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$ of pure dimension $r + \dim \mathfrak{E}_{\mathcal{S},\lambda}$, where $\mathfrak{E}_{\mathcal{S},\lambda}$ is the following locally closed subset of $\mathfrak{M}_n(\overline{\mathbb{F}}_p)$:*

$$\mathfrak{E}_{\mathcal{S},\lambda} \coloneqq \left\{ B \in \operatorname{Cent} \mathcal{S} \;\middle|\; \ker(B - \lambda_i I_n) = \ker(A_{\mathcal{S},\lambda} - \lambda_i I_n) \text{ for all } 1 \leq i \leq r \right\}.$$

Proof. The eigenspace $E_i \coloneqq \ker(A_{\mathcal{S},\lambda} - \lambda_i I_n)$ is by definition defined over $\mathbb{F}_p$. If $M = U A_{\mathcal{S},\lambda} U^{-1}$, then the eigenspace $\ker(M - \lambda_i I_n) = U(E_i)$ is defined over $\mathbb{F}_p$ if and only if $(U^{-1}\sigma(U))(E_i) = E_i$. Letting $\mathfrak{D}'_{\mathcal{S}}$ be the set of matrices $U \in \mathfrak{D}_{\mathcal{S}}$ such that $U(E_i) = E_i$ for all $i \in \{1, \dots, r\}$, the same proof as Proposition 5.6 shows that $\mathfrak{X}_{\mathcal{S}}^{\mathrm{eig.}/\mathbb{F}_p}$ has dimension $r + \dim \mathfrak{D}'_{\mathcal{S}} - \dim \operatorname{Cent} \mathcal{S}$. Thus, it suffices to prove that $\mathfrak{E}_{\mathcal{S},\lambda}$ has pure dimension $\dim \mathfrak{D}'_{\mathcal{S}} - \dim \operatorname{Cent} \mathcal{S}$. Note that $\mathfrak{E}_{\mathcal{S},\lambda} \subseteq \operatorname{Cl} A_{\mathcal{S},\lambda}$ by Lemma 6.1(ii). The computation is then analogous to the proof of Lemma 5.7. □

We now compute the dimension of $\mathfrak{E}_{\mathcal{S},\lambda}$:

Proposition 6.4. *Consider a shape $\mathcal{S} = (\{1, \dots, r\}, e)$ and a tuple $\lambda = (\lambda_1, \dots, \lambda_r) \in \mathfrak{Y}_{\mathcal{S}}$. Then:*

(i) We have an isomorphism of varieties $\mathfrak{E}_{\mathcal{S},\lambda} \simeq \prod_i \mathfrak{E}_{\mathcal{S}_i,\lambda_i}$, where $\mathcal{S}_i \coloneqq \big(\{i\}, (i,k) \mapsto e(i,k)\big)$ is the subshape for the eigenvalue λ_i.

(ii) $\dim \mathfrak{E}_{\mathcal{S},\lambda} = \sum_{i=1}^{r} \sum_{k \geq 1} e(i,k) \cdot e(i,k+1)$

Proof.

(i) Let $B \in \mathfrak{E}_{\mathcal{S},\lambda}$. Since B commutes with $A_{\mathcal{S},\lambda}$, it preserves the generalized eigenspace G_{λ_i} for each eigenvalue λ_i, inducing maps $B_i \colon G_{\lambda_i} \to G_{\lambda_i}$ which are easily checked to belong to $\mathfrak{E}_{\mathcal{S}_i,\lambda_i}$. We have $\bigoplus_{i=1}^{r} G_{\lambda_i} = \overline{\mathbb{F}}_p^n$, so B can be reconstructed from the restricted maps $B_i \colon G_{\lambda_i} \to G_{\lambda_i}$. We have described two inverse regular maps.

(ii) By (i), we reduce to the case $r = 1$. Without loss of generality (subtracting λI_n from everything), we have $\lambda = 0$. Then, the claim amounts to Lemma 6.5 below with $A = A_{\mathcal{S},0}$. □

Lemma 6.5. *Let A be a nilpotent endomorphism of an n-dimensional vector space V. Let $e_A(k) :=$ $\dim\ker A^k - \dim\ker A^{k-1}$ and let $\mathfrak{C}_A := \{B \in \mathrm{Cent}(A) \mid \ker B = \ker A\}$. Then:*

$$\dim \mathfrak{C}_A = \sum_{k\geq 1} e_A(k)\cdot e_A(k+1).$$

Proof. We actually show that the linear subspace $\mathfrak{C}'_A := \{B \in \mathrm{Cent}\,A \mid \ker B \supseteq \ker A\}$ has the announced dimension. Since $\mathfrak{C}_A$ is an open subset of $\mathfrak{C}'_A$ (it is defined by the non-vanishing of certain determinants) and is non-empty (it contains A), it is Zariski dense and the result shall follow.

We reason by induction on the dimension n of V. Since A is nilpotent, $\mathrm{im}\,A$ has strictly smaller dimension than V, and $\overline{A} := A|_{\mathrm{im}\,A}$ is a nilpotent endomorphism of $\mathrm{im}\,A$. Moreover,

$$\begin{aligned} e_{\overline{A}}(k) &= \dim(\ker A^k \cap \mathrm{im}\,A) - \dim(\ker A^{k-1}\cap \mathrm{im}\,A) = \dim A(\ker A^{k+1}) - \dim A(\ker A^k) \\ &= (\dim\ker A^{k+1} - \dim\ker A) - (\dim\ker A^k - \dim\ker A) = e_A(k+1), \end{aligned}$$

so $\dim \mathfrak{C}'_{\overline{A}} = \sum_{k\geq 2} e_A(k)\cdot e_A(k+1)$ by the induction hypothesis. It therefore suffices to show that the linear map $f\colon \mathfrak{C}'_A \to \mathfrak{C}'_{\overline{A}}$ sending B to its restriction $B|_{\mathrm{im}\,A}$ is surjective and that its kernel has dimension $e_A(1)\cdot e_A(2)$.

Consider an endomorphism $\overline{B}\colon \mathrm{im}\,A \to \mathrm{im}\,A$ in $\mathfrak{C}'_{\overline{A}}$. The fiber $f^{-1}(\overline{B})$ consists of those endomorphisms $B\colon V\to V$ whose restriction to $\mathrm{im}\,A$ is $\overline{B}$, which vanish on $\ker A$, and such that the following diagram commutes:

$$\begin{tikzcd} \mathrm{im}\,A \arrow[d, "\overline{B}"'] & V \arrow[l, two heads, "A"'] \arrow[d, dashed, "B"] \\ \mathrm{im}\,A & V \arrow[l, two heads, "A"] \end{tikzcd}$$

We pick a complement C of $\mathrm{im}\,A + \ker A$ in V. Since $\overline{B} \in \mathfrak{C}'_{\overline{A}}$, restriction to C defines a bijection between $f^{-1}(\overline{B})$ and the set of linear maps $B'\colon C \to V$ such that the following diagram commutes:

$$\begin{tikzcd} \mathrm{im}\,A \arrow[d, "\overline{B}"'] & C \arrow[l, "A"'] \arrow[d, dashed, "B'"] \\ \mathrm{im}\,A & V \arrow[l, two heads, "A"] \end{tikzcd}$$

In particular, the fibers are non-empty (the map $\overline{B}\circ A$ factors through the surjection $A\colon V \twoheadrightarrow \mathrm{im}\,A$), so f is surjective. Taking $\overline{B} = 0$, we see that the kernel of f is isomorphic to the vector space of linear maps $B'\colon C \to \ker A$, of dimension $\dim\ker A\cdot \dim C$. The claim follows since $\dim\ker A = e_A(1)$ and

$$\begin{aligned} \dim C &= \dim V - \dim(\mathrm{im}\,A + \ker A) = \dim\mathrm{im}\,A + \dim\ker A - \dim(\mathrm{im}\,A+\ker A) \\ &= \dim(\mathrm{im}\,A\cap\ker A) = \dim A(\ker A^2) = \dim\ker A^2 - \dim\ker A = e_A(2). \end{aligned} \qquad \square$$

Proposition 6.6. *The maximal value of $\dim \mathfrak{X}_{\mathcal{S}}^{\mathrm{eig.}/\mathbb{F}_p} = r + \sum_{i=1}^r \sum_{k\geq 1} e(i,k)\cdot e(i,k+1)$ over shapes $\mathcal{S}$ of size n is $\lfloor n^2/4\rfloor + 1$, and it is reached exactly for the following shapes (up to isomorphism), where we represent a shape $\mathcal{S} = (\{1,\ldots,r\}, e)$ by the tuple $((e(1,1), e(1,2),\ldots), \ldots, (e(r,1),\ldots))$, omitting the trailing zeros:*

n	optimal shapes
2	$((1,1))$, $((1),(1))$
3	$((2,1))$, $((1,1,1))$, $((1,1),(1))$, $((1),(1),(1))$
$2m,\ m\geq 2$	$((m,m))$
$2m+1,\ m\geq 2$	$((m+1,m))$, $((m,m,1))$

Proof. First, we consider only shapes with $r = 1$. Let $\mathcal{S} = (\{1\}, e)$, and let s be such that $e(1,s) \neq 0$ and $e(1,s+1) = 0$. We have

$$\dim \mathfrak{X}_{\mathcal{S}}^{\mathrm{eig.}/\mathbb{F}_p} = 1 + \sum_{k=1}^{s-1} e(1,k)e(1,k+1) \leq 1 + \sum_{k=1}^{s-1} e(1,1)e(1,k+1) = 1 + e(1,1)\cdot(n - e(1,1)),$$

with equality if and only if $e(1,1) = e(1,2) = \ldots = e(1,s-1)$. Since $e(1,1)$ is an integer, $1 + e(1,1)\cdot(n-e(1,1))$ has maximal value $1 + \lfloor n/2 \rfloor \cdot \lceil n/2 \rceil = 1 + \lfloor n^2/4 \rfloor$, reached exactly when $e(1,1) \in \{\lfloor n/2 \rfloor, \lceil n/2 \rceil\}$. If n is even, only the shape $\left((\frac{n}{2}, \frac{n}{2})\right)$ gives equality. If n is odd, distinguishing between the two possible values of $e(1,1)$ gives the two equality cases with $r = 1$.

Now, consider the case of a general shape $\mathcal{S} = (\{1,\ldots,r\}, e)$. By the case $r = 1$, we have

$$r + \sum_{i=1}^{r} \sum_{k\geq 1} e(i,k)\cdot e(i,k+1) = \sum_{i=1}^{r}\left(1 + \sum_{k\geq 1} e(i,k)\cdot e(i,k+1)\right) \leq \sum_{i=1}^{r}\left(1 + \left\lfloor \frac{\left(\sum_{k\geq 1} e(i,k)\right)^2}{4} \right\rfloor\right).$$

However, the function $\eta(n) := \lfloor n^2/4 \rfloor + 1$ is strictly superadditive except for the equalities $\eta(1) + \eta(1) = \eta(2)$ and $\eta(1) + \eta(2) = \eta(1) + \eta(1) + \eta(1) = \eta(3)$. Therefore, we must have $r = 1$ if $n > 3$, and the cases $n \in \{2,3\}$ are quickly dealt with. $\square$

It remains only to obtain estimates for $|\mathfrak{X}_{\mathcal{S}}^{\mathrm{eig.}/\mathbb{F}_p} \cap \mathfrak{M}_n(\mathbb{F}_q)|$ when $\mathcal{S}$ is one of the optimal shapes of Proposition 6.6. For this, we are going to need the following two lemmas:

Lemma 6.7. *Let $a \geq 1$ and let $\vec{v}, \vec{w} \in \mathbb{F}_q^a$ be non-zero vectors. The number of matrices $N \in \mathrm{GL}_a(\mathbb{F}_q)$ satisfying $N\vec{w} = \sigma(N)\vec{v}$ is $q^{a(a-1)} + O_{p,a}(q^{a(a-1)-1})$ if $\vec{v}$ and $\sigma(\vec{w})$ are linearly independent, and $O_{p,a}(q^{a(a-1)})$ otherwise.*

Proof. Assume first that $\vec{v}$ and $\sigma(\vec{w})$ are linearly independent. Replacing $(\vec{v}, \vec{w})$ by $(\sigma(U)\vec{v}, U\vec{w})$ for an appropriate $U \in \mathrm{GL}_a(\mathbb{F}_q)$, we can assume without loss of generality that $\vec{v} = \vec{e}_1$ and $\vec{w} = \vec{e}_2$ are the first two standard basis vectors. Then, $N\vec{w} = \sigma(N)\vec{v}$ means that the second column of N is deduced from the first column by applying σ. The number of invertible matrices satisfying this condition is as claimed.

Now, assume that $\sigma(\vec{w}) = \lambda\vec{v}$ for some $\lambda \in \mathbb{F}_q^{\times}$. Replacing $(\vec{v}, \vec{w})$ by $(\sigma(U)\vec{v}, U\vec{w})$ for an appropriate matrix $U \in \mathrm{GL}_a(\mathbb{F}_q)$, we can assume that $\vec{v} = \vec{e}_1$ and $\vec{w} = \sigma^{-1}(\lambda)\vec{e}_1$. The condition $N\vec{w} = \sigma(N)\vec{v}$ then leaves at most $p^a = O_{p,n}(1)$ options for the first column of N. $\square$

Lemma 6.8. *Let $m \geq 1$. For any filtration of linear subspaces $0 = V_0 \subseteq \cdots \subseteq V_s = \overline{\mathbb{F}}_p^m$, where each V_k is defined over $\mathbb{F}_p$, the number of (nilpotent) matrices $M \in \mathfrak{M}_m(\mathbb{F}_q)$ commuting with $\sigma(M)$ and such that $\ker M^k = V_k$ for all $k \in \{1,\ldots,s\}$ only depends on q and on the numbers $e(k) := \dim(V_k/V_{k-1})$. We denote this count by $w_q(e(1),\ldots,e(s))$ (we omit trailing zeros in the notation, i.e., this means $e(k) = 0$ for $k \geq s+1$). Moreover:*

(a) For any $m \geq 1$, we have $w_q(m) = 1$.

(b) For any $a \geq b \geq 1$ with $a + b = m$, we have $w_q(a,b) = q^{ab} + O_{p,a,b}(q^{ab-1})$.

(c) For any $a \geq 1$ with $2a+1 = m$, we have $w_q(a,a,1) = q^{a(a+1)} + O_{p,a}(q^{a(a+1)-1})$.

Proof. Conjugating by an element of $\mathrm{GL}_m(\mathbb{F}_p)$, we can assume without loss of generality that each V_k is generated by the first $\dim V_k = e(1) + \cdots + e(k)$ standard basis vectors of $\mathbb{F}_p^m$. In particular, this proves the well-definedness of $w_q(e(1),\ldots,e(s))$.

(a) That $e(1) = m$ implies that $V_1 = \overline{\mathbb{F}}_p^m$, and only the zero matrix satisfies $\ker M = V_1 = \overline{\mathbb{F}}_p^m$.

(b) The condition $\ker M^k = V_k$ for all $k \in \{1,2\}$ means that M is of the form $M = \left(\begin{smallmatrix} 0 & N \\ 0 & 0 \end{smallmatrix}\right)$ for some $a \times b$ matrix N of rank b. If M is of this form, then so is $\sigma(M)$ and they automatically commute. The number of such matrices N with coefficients in $\mathbb{F}_q$ is $q^{ab} + O_{p,a,b}(q^{ab-1})$.

(c) The condition $\ker M^k = V_k$ for all $k \in \{1,2,3\}$ means that M is of the form $M = \left(\begin{smallmatrix} 0 & N & \vec{u} \\ 0 & 0 & \vec{v} \\ 0 & 0 & 0 \end{smallmatrix}\right)$ for some invertible $a \times a$ matrix $N \in \mathrm{GL}_a(\mathbb{F}_q)$, some column vector $\vec{u} \in \mathbb{F}_q^a$, and some nonzero column vector $\vec{v} \in \mathbb{F}_q^a$. If M is of this form, it commutes with $\sigma(M)$ if and only if $N\sigma(\vec{v}) = \sigma(N)\vec{v}$. Taking $\vec{w} := \sigma(\vec{v})$, the claim then follows from Lemma 6.7 by summing over all possible pairs of vectors $\vec{u} \in \mathbb{F}_q^a$ and $\vec{v} \in \mathbb{F}_q^a \setminus \{0\}$, as $\vec{v}$ and $\sigma(\vec{w}) = \sigma^2(\vec{v})$ are linearly independent if and only if $\langle \vec{v} \rangle$ is not defined over $\mathbb{F}_{p^2}$, which is the generic case. □

Theorem 6.9. *For any finite field $\mathbb{F}_q \supseteq \mathbb{F}_p$, we have*

$$|\mathfrak{X}^{\mathrm{eig.}/\mathbb{F}_p} \cap \mathfrak{M}_n(\mathbb{F}_q)| = c^{\mathrm{eig.}/\mathbb{F}_p}(p,n) \cdot q^{\lfloor n^2/4 \rfloor + 1} + O_{p,n}(q^{\lfloor n^2/4 \rfloor}),$$

where

$$c^{\mathrm{eig.}/\mathbb{F}_p}(p,2) = \tfrac{1}{2}(p+2)(p+1), \qquad c^{\mathrm{eig.}/\mathbb{F}_p}(p,3) = \tfrac{1}{6}(p^2+p+1)(p^4+7p^3+6p^2+6p+12),$$

$$c^{\mathrm{eig.}/\mathbb{F}_p}(p,n) = \binom{n}{n/2}_p \quad \text{if } n \geq 4 \text{ is even,}$$

$$c^{\mathrm{eig.}/\mathbb{F}_p}(p,n) = \binom{n}{\lfloor n/2 \rfloor}_p + \binom{n}{\lfloor n/2 \rfloor}_p \cdot \binom{\lceil n/2 \rceil}{1}_p \quad \text{if } n \geq 5 \text{ is odd.}$$

Proof. For any Jordan shape $\mathcal{S}$ which is not listed in Proposition 6.6, we have $\dim \mathfrak{X}_{\mathcal{S}}^{\mathrm{eig.}/\mathbb{F}_p} \leq \lfloor n^2/4 \rfloor$ and therefore $|\mathfrak{X}_{\mathcal{S}}^{\mathrm{eig.}/\mathbb{F}_p} \cap \mathfrak{M}_n(\mathbb{F}_q)| = O_{p,n}(q^{\lfloor n^2/4 \rfloor})$ by the Schwarz–Zippel bound [LW54, Lemma 1].

Now, let $\mathcal{S} = (V, e)$ be one of the Jordan shapes listed in Proposition 6.6. To construct a matrix $M \in \mathfrak{X}_{\mathcal{S}}^{\mathrm{eig.}/\mathbb{F}_p}$, we choose its $|V|$ (distinct) eigenvalues λ_i and the corresponding generalized eigenspaces G_i of dimension $d(i) := \sum_{k \geq 1} e(i,k)$ for all i (which must be defined over $\mathbb{F}_p$ by Corollary 6.2), modulo the automorphisms of $\mathcal{S}$. There are $q^{|V|} + O_{p,n}(q^{|V|-1})$ choices for the eigenvalues and $|\mathrm{GL}_n(\mathbb{F}_p)| / \prod_{i \in V} |\mathrm{GL}_{d(i)}(\mathbb{F}_p)|$ choices for the generalized eigenspaces (as one shows using the orbit-stabilizer theorem). For each i, we then need to choose the filtration of subspaces $V_{i,k} := \ker(M - \lambda_i I_n)^k$ (each defined over $\mathbb{F}_p$), satisfying $0 = V_{i,0} \subseteq \cdots \subseteq V_{i,s_i} = G_i$, with $\dim(V_{i,k}/V_{i,k-1}) = e(i,k)$. The group $\mathrm{GL}_{d(i)}(\mathbb{F}_p)$ acts transitively on such flags. Describing the stabilizer of a given flag (by induction on s_i) and using the orbit-stabilizer theorem, one shows that the number of such flags for each i is

$$\frac{|\mathrm{GL}_{d(i)}(\mathbb{F}_p)|}{\prod_{k \geq 1} |\mathrm{GL}_{e(i,k)}(\mathbb{F}_p)| \cdot \prod_{k > l \geq 1} p^{e(i,k) \cdot e(i,l)}}.$$

Finally, we need to choose for each i the restriction of $M - \lambda_i I_n$ to the generalized eigenspace G_i. We estimated the number $w_q(e(i,1), e(i,2), \ldots, e(i,s_i))$ of choices for this restriction in Lemma 6.8. For any Jordan shape $\mathcal{S} = (V, e)$, we then obtain

$$|\mathfrak{X}_{\mathcal{S}}^{\mathrm{eig.}/\mathbb{F}_p} \cap \mathfrak{M}_n(\mathbb{F}_q)| = \frac{|\mathrm{GL}_n(\mathbb{F}_p)| \cdot (q^{|V|} + O_{p,n}(q^{|V|-1}))}{|\mathrm{Aut}(\mathcal{S})|} \cdot \prod_{i \in V} \frac{w_q(e(i,1), e(i,2), \ldots, e(i,s_i))}{\prod_{k \geq 1} |\mathrm{GL}_{e(i,k)}(\mathbb{F}_p)| \cdot \prod_{k > l \geq 1} p^{e(i,k) \cdot e(i,l)}}.$$

The claim follows by summing over all the shapes listed in Proposition 6.6 and using the formulas given in Lemma 6.8. □

References

[BGS14] Andries E. Brouwer, Rod Gow, and John Sheekey. Counting symmetric nilpotent matrices. *The Electronic Journal of Combinatorics*. 21(2). 2014. DOI: 10.37236/3810.

[FF60] Walter Feit and Nathan J. Fine. Pairs of commuting matrices over a finite field. *Duke Mathematical Journal*. 27, pp. 91–94. 1960. DOI: 10.1215/S0012-7094-60-02709-5.

[FH58] Nathan J. Fine and Israel N. Herstein. The probability that a matrix be nilpotent. *Illinois Journal of Mathematics*. 2, pp. 499–504. 1958.

[Gan53] Felix R. Gantmakher. *The Theory of Matrices, Vol. 1*. 1953.

[Ger61a] Murray Gerstenhaber. On dominance and varieties of commuting matrices. *Annals of Mathematics. Second Series*. 73, pp. 324–348. 1961. DOI: 10.2307/1970336.

[Ger61b] Murray Gerstenhaber. On the number of nilpotent matrices with coefficients in a finite field. *Illinois Journal of Mathematics*. 5, pp. 330–333. 1961.

[GS00] Robert M. Guralnick and B. A. Sethuraman. Commuting pairs and triples of matrices and related varieties. *Linear Algebra and its Applications*. 310(1-3), pp. 139–148. 2000. DOI: 10.1016/S0024-3795(00)00065-3.

[GS25] Fabian Gundlach and Béranger Seguin. Counting two-step nilpotent wildly ramified extensions of function fields. 2025. arXiv: 2502.18207.

[Gur92] Robert M. Guralnick. A note on commuting pairs of matrices. *Linear and Multilinear Algebra*. 31(1-4), pp. 71–75. 1992. DOI: 10.1080/03081089208818123.

[Har92] Joe Harris. *Algebraic Geometry: A First Course*. volume 133 of *Graduate Texts in Mathematics*. Springer-Verlag, New York. 1992. DOI: 10.1007/978-1-4757-2189-8. ISBN: 0-387-97716-3.

[HHYZ2424] Martin Hils, Ehud Hrushovski, Jinhe Ye, and Tingxiang Zou. Lang-Weil Type Estimates in Finite Difference Fields. 2024. arXiv: 2406.00880.

[Hua23] Yifeng Huang. Mutually annihilating matrices, and a Cohen-Lenstra series for the nodal singularity. *Journal of Algebra*. 619, pp. 26–50. 2023. DOI: 10.1016/j.jalgebra.2022.11.021.

[LW54] Serge Lang and André Weil. Number of points of varieties in finite fields. *American Journal of Mathematics*. 76, pp. 819–827. 1954. DOI: 10.2307/2372655.

[Mir98] Maryam Mirzakhani. A simple proof of a theorem of Schur. *American Mathematical Monthly*. 105(3), pp. 260–262. 1998. DOI: 10.2307/2589084.

[MT55] Theodore S. Motzkin and Olga Taussky. Pairs of matrices with property L. II. *Transactions of the American Mathematical Society*. 80, pp. 387–401. 1955. DOI: 10.2307/1992996.

[Pan08] Dmitri I. Panyushev. Two results on centralisers of nilpotent elements. *Journal of Pure and Applied Algebra*. 212(4), pp. 774–779. 2008. DOI: 10.1016/j.jpaa.2007.07.003.

[Sch05] Issai Schur. Zur Theorie der vertauschbaren Matrizen. *Journal für die reine und angewandte Mathematik*. 130, pp. 66–76. 1905. URL: https://eudml.org/doc/149219.

[Sch08] Eric Schmutz. Splitting fields for characteristic polynomials of matrices with entries in a finite field. *Finite Fields and their Applications*. 14(1), pp. 250–257. 2008. DOI: 10.1016/j.ffa.2007.10.001.

[Ser79] Jean-Pierre Serre. *Local Fields*. volume 67 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin. 1979. ISBN: 0-387-90424-7.

[Stacks] The Stacks Project. https://stacks.math.columbia.edu.

[SV22] Kadattur Vasudevan Shuddhodan and Yakov Varshavsky. The Hrushovski-Lang-Weil estimates. *Algebraic Geometry*. 9(6), pp. 651–687. 2022. DOI: 10.14231/AG-2022-020.