Smoothing the degree of regularity for polynomial systems

Samuel Jaques[2], Lars Ran[1], Simona Samardjiska[1], and Melvin Seitner[1]

[1] Radboud University, Nijmegen, Netherlands
[2] University of Waterloo, Waterloo, Canada
sejaques@uwaterloo.ca,{lran,simonas}@cs.ru.nl,melvin.seitner@ru.nl

Abstract. The complexity of many algebraic algorithms for solving non-linear polynomial systems of equations over finite fields, such as the XL (eXtended Linearization) algorithm or variants of the F4/F5 algorithms, is directly determined by the so called degree of regularity. In essence, we can form the Macaulay matrix at this degree, which can be thought of as the linearization of monomial multiples of the polynomials from the problem instance, and then solve the obtained linear system. The degree of regularity guarantees we can solve the system, however, it is a rather coarse parameter. This means that, sometimes, in order to provably show that we achieve overdeterminess, we end up with a heavily overdetermined Macaulay matrix which unnecessarily increases the computational cost.
To reduce this coarseness, and thus avoid the unnecessary computational cost, we propose a technique for "smoothing" the degree of regularity that can be seen as operating at a degree in-between two integer values. Instead of the full Macaulay matrix, we consider specific submatrices that we show are sufficient to solve the given system. Under a mild assumption that generalizes the notion of semi-regularity, which we experimentally verify for a range of parameters, we show that our approach smooths the complexity of XL for the MQ (Multivariate Quadratic) problem. Finally, we apply our results to a recent UOV attack and demonstrate improvement of a few bits.

Keywords: semi-regular sequences · degree of regularity · Macaulay matrices · XL algorithm · initial segment · algebraic system solving.

1 Introduction

Solving multivariate quadratic polynomial systems, known as the MQ problem, is a well-known NP-hard problem [32] underlying the security of numerous cryptographic schemes whether that be by design or as a result of algebraic modeling

This research has been partially supported by the Dutch government through the NWO grant OCNW.M.21.193 (ALPaQCa) and the NWO Quantum Technology grant NGF.1623.23.020 (PQStrong).

of the cryptosystem [35,21,37,17,9,6,36,27,26,1,31,2,42,44,43]. For such cryptographic settings, it is typical to assume the MQ problem has at most a single solution. In addition, many optimization and classification problems can be modeled as polynomial systems. Therefore, it is of particular importance to understand not only the asymptotic behavior of the state-of-the-art algorithms for solving it, but also their practical hardness for relevant parameters.

The best studied algorithms for the MQ problem are algebraic ones, dating back to the works of Buchberger [3] and Lazard [39]. Over the years, these received significant attention and refinements, resulting in several closely related algorithms – XL [19], F4/F5 [23,24], Crossbred [34] and their variants and extensions [50,49,20,15,41,7,25,8]. In essence, the operation of all of these can be described as performing linear algebra on Macaulay matrices. Our focus will be on improving the XL (eXtended Linearization) algorithm, as the analysis of our approach is the most natural in this case. However, it can also be applied to F4/F5 and Crossbred, but it requires a slightly different analysis. We briefly explain the inner working of the XL algorithm.

Given an instance $\mathcal{F} = (f_1, \ldots, f_m) \in \mathbb{F}_q[X_1, \ldots, X_m]$ of the MQ-problem over a finite field $\mathbb{F}_q$, we first form a Macaulay matrix at degree d whose columns are indexed by all monomials μ in the variables $X_1, \ldots, X_m$ up to degree d, and whose rows are indexed by $\nu \cdot f_i$ for all monomials ν in the variables up to degree $d-2$. The matrix is then populated with the corresponding coefficients of $\nu \cdot f_i$. Now, it is not hard to see that a solution to the polynomial system is necessarily described by a vector in the right kernel of the matrix. By checking all elements of this kernel, we can extract a set of all solutions. This extraction is hard if the dimension of the kernel, i.e., the right nullity of the matrix, is large. However, if d is sufficiently big so that the right nullity of the Macaulay matrix is close to 1, finding this right kernel efficiently solves the given polynomial system.

The difficulty in this seemingly simple approach is determining what minimal degree d is necessary to find a solution. A crude indication is that we need to have at least as many rows as columns. Yet, that is not enough as there will certainly be linear dependencies among the rows given by $f_i f_j - f_j f_i = 0$ once $d \geq 4$. For a certain type of homogeneous polynomial systems that have only trivial relations $f_i f_j - f_j f_i = 0$, namely semi-regular systems, we can quite accurately predict the required minimum d, which in this case is the degree of regularity d_{reg} of the system [4,5]. In this work, we focus on homogeneous systems and the XL algorithm. XL has the benefit of using sparse linear algebra algorithms such as Wiedemann's [47], as it operates directly on a sparse matrix at the degree at which the solution can be found. As this is d_{reg} for semi-regular systems, we choose to focus solely on this invariant in our analysis. For more general systems that are not necessarily semi-regular, other invariants related to the degree of regularity are discussed in more detail in [16].

A closer inspection of the values of the degree of regularity for various parameters reveal that its graph has a stepwise shape, i.e., we observe jumps of the degree between close values of the number of variables. This discreteness has a visible impact in practice – it may happen that the required degree produces

significantly overdetermined Macaulay matrix that unnecessarily increases the time and memory complexity.

In this paper, we ask the question of whether we can form a matrix that is in-between two consecutive-degree Macaulay matrices but which can be used in a similar manner as a Macaulay matrix to find a solution. In other words, we would like to smooth-out the discreteness of the degree of regularity. Note immediately that a simple partial enumeration can be considered in some sense to achieve the goal, and it has long been known that the approach can reduce the complexity especially for small fields. For instance, this is done in HybridF5 based on the F5 algorithm [7] and in FXL [19] based on the XL algorithm. Both first partially guess a few variables and then solve the smaller multivariate system, repeating until a solution is found.

1.1 Contribution

As a first contribution, we propose a smoothing technique that uses a submatrix of the degree d_{reg} Macaulay matrix that is sufficient for solving the system. This submatrix is given by a set of rows and their combined set of supporting columns. In contrast to the F5 criterion, which primarily serves to reduce the number of rows of the Macaulay matrix, our focus is on reducing the number of columns.

We develop a heuristic strategy for achieving optimality of the choice of submatrix of the Macaulay matrix. As indicated before, we assume instances of the MQ problem that have at most a single solution, as this is typically the case in cryptographic settings. In order to efficiently solve the system, we wish to find the least wide submatrix that has a nullity of at most 1. Thus, we would like to consider submatrices of minimal width among all submatrices of the same rank.

Since finding the rank of an arbitrary submatrix is not trivial, as a proxy for the rank we consider the height of this submatrix. We show, using a classical combinatorial result by Clements and Lindström [18], that choosing the rows indexed by initial segments with respect to the `grevlex` monomial ordering optimizes the number of columns of the submatrix with respect to the number of rows. This result leads us to consider Macaulay matrices of initial segments as the basis of our improvement on XL.

We then proceed to estimating the nullity of such matrices. This is necessary in order to be able to determine the complexity of our improvement, just as this is necessary for the entire Macaulay matrix for determining the complexity of XL. To achieve this, we first define the notion of monomial semi-regularity which naturally generalizes the standard notion of semi-regularity. Based on experimental evidence, we conjecture that generic sequences of polynomials satisfy this property.

Under this assumption, we prove an upper bound of the nullity of Macaulay matrices of initial segments. As a main tool in our analysis we use the relations between trivial syzygy matrices of initial segments to expose their structure, partition them and estimate their rank.

We perform an extensive set of experiments to verify the validity of our bound. Our experiments show that the bound becomes very quickly exact as the

overdeterminacy of the system increases. We further estimate the improvements over XL, and demonstrate that we achieve the smoothing effect. For overdetermined systems, we find parameters where we are able to more than half the width of the matrices used. For systems with twice as many restrictions as variables, we find parameters where we can reduce the width of the Macaulay matrix by a factor exceeding 16.

Furthermore, we compare, but also combine, with FXL to achieve further improvement in the complexity.

Finally, we adapt the approach to a recent attack on UOV [14] from [29]. This adaptation again allows us to consider in-between options of an integer parameter from the pre-existing framework. In this case, that parameter is not the degree of the Macaulay matrix, but the number k of variables that are fixed to 0. We show over 3 bits of improvement over [29] for some parameters.

2 Preliminaries

Denote $[n] = \{1, \dots, n\}$. For a set S, we denote by $\mathcal{P}S$ the set of all subsets of S, and by $\mathcal{P}_c S$ the set of all subsets of S of size c.

Throughout the paper, let $\mathbb{F}$ be a field and let $\mathcal{R} = \mathbb{F}[\mathbf{X}]$ be the ring of polynomials in the variables $\mathbf{X} = (X_1, \dots, X_n)$ over $\mathbb{F}$. We denote by π_i the projection map $\mathbb{F}[X_1, \dots, X_n] \xrightarrow{\pi_i} \mathbb{F}[X_1, \dots, X_i]$ with kernel $\langle X_{i+1}, \dots, X_n \rangle$. For vectors, $\mathbf{v} = (v_i)_{i \in I} \in \mathbb{F}^I$, we define its support $\mathrm{supp}(\mathbf{v})$ as the set of indices $i \in I$ where $v_i \neq 0$.

Let $\mathcal{M}(\mathbf{X})$ denote the set of monomials in $\mathcal{R}$ and denote by $\mathcal{M}_d(\mathbf{X})$ the set of monomials of degree d. If $\mathbf{X}$ is clear from context, we may write $\mathcal{M} := \mathcal{M}(\mathbf{X})$ and $\mathcal{M}_d := \mathcal{M}_d(\mathbf{X})$. We use a shorthand notation $\left(\!\binom{n}{d}\!\right) = \binom{n+d-1}{d}$ for the size of $\mathcal{M}_d$. When we write a monomial as product of variables, $\mu = X_{a_d} \cdots X_{a_1}$, we assume $n \geq a_d \geq \cdots \geq a_1 \geq 1$ unless otherwise indicated. We write $\mathcal{R}_d := \langle \mathcal{M}_d \rangle_{\mathbb{F}}$ for the $\mathbb{F}$-linear span of $\mathcal{M}_d$. Similarly, for a set of monomials $S \in \mathcal{M}$ we write $\mathcal{R}_S := \langle S \rangle_{\mathbb{F}}$.

We say that $f \in \mathcal{R}$ is homogeneous of degree d if $f \in \mathcal{R}_d$. We call an ideal homogeneous if it is generated by homogeneous polynomials. For $f \in \mathcal{R}$, we define its support $\mathrm{supp}(f) \subset \mathcal{M}$ as the smallest set $S \subset \mathcal{M}$ for which $f \in \mathcal{R}_S$. For a homogeneous ideal $I \subset \mathcal{R}$ we will define $I_d := I \cap \mathcal{R}_d$ and more generally $I_S := I \cap \mathcal{R}_S$. For a sequence of polynomials $\mathcal{F} = (f_1, \dots, f_m)$ in $\mathcal{R}$ we denote $\mathcal{F}_{(i)} = (f_1, \dots, f_i)$.

2.1 Regularity assumptions

Throughout this work, for a given sequence of polynomials $\mathcal{F} = (f_1, \dots, f_m)$ in $\mathcal{R}$, in order to make predictions, we would like to assume that there are no non-trivial relations among the f_i up to a certain degree. This is captured by the notion of semi-regularity. We will follow the definition given in [5].

Definition 1. *The degree of regularity of a homogeneous ideal I is given by*

$$d_{reg} := \min\{d \geq 0 \mid I_d = \mathcal{R}_d\} \cup \{\infty\}.$$

By extension, we call the degree of regularity of a homogeneous system $f_1, \ldots, f_m$ the degree of regularity of $\langle f_1, \ldots, f_m \rangle$. If the system is not clear from the context, we will write it explicitly as $d_{reg}(\mathcal{F})$.

Definition 2. *A sequence $f_1, \ldots, f_m$ of homogeneous polynomials in $\mathcal{R}$ is d-regular if*

$$gf_i \in \langle f_1, \ldots, f_{i-1} \rangle \text{ and } \deg(gf_i) < d \implies g \in \langle f_1, \ldots, f_{i-1} \rangle.$$

for all $1 \leq i \leq m$. We call a sequence of homogeneous polynomials $f_1, \ldots, f_m$ a semi-regular system if it is d_{reg}-regular.

In [28] it is conjectured that semi-regularity is a generic property among systems $\mathcal{F} = (f_1, \ldots, f_m)$ of degrees $d_1, \ldots, d_m$ in n variables. In more practical terms, this means that random systems are likely semi-regular.

2.2 Macaulay and relation matrices

Given a sequence of quadratic polynomials $\mathcal{F} = (f_1, \ldots, f_m)$ in $\mathcal{R}$ we can construct its Macaulay matrix [40] to investigate its behavior in higher degrees using linear algebra.

Definition 3. *Let $\mathcal{F} = (f_1, \ldots, f_m)$ be a homogeneous sequence of quadratic polynomials in $\mathcal{R}$. The degree d Macaulay matrix of $\mathcal{F}$ is the matrix $\mathrm{Mac}^d(\mathcal{F}) \in \mathbb{F}^{m\left(\!\binom{n}{d-2}\!\right) \times \left(\!\binom{n}{d}\!\right)}$ with*

- *rows labeled by (μ, i) with $i \in [m]$ and $\mu \in \mathcal{M}_{d-2}$,*
- *columns labeled by $\nu \in \mathcal{M}_d$,*
- *where the coefficient of $((\mu, i), \nu)$ given by the coefficient of ν in $\mu \cdot f_i$.*

The characteristic feature of the Macaulay matrix is that

$$\mathrm{rowspace}\,\mathrm{Mac}^d(\mathcal{F}) \cong \langle \mathcal{F} \rangle_d.$$

as $\mathbb{F}$ vector spaces. To get an even more fine-grained structure, we can also study its (higher order) relation matrices.

Definition 4. *Let $\mathcal{F} = (f_1, \ldots, f_m)$ be a homogeneous sequence of quadratic polynomials in $\mathcal{R}$. The degree d relation matrix of order $0 \leq c < m$ of $\mathcal{F}$ is the matrix $\mathrm{Rel}^d_c(\mathcal{F}) \in \mathbb{F}^{\binom{m}{c+1}\left(\!\binom{n}{d-2c-2}\!\right) \times \binom{m}{c}\left(\!\binom{n}{d-2c}\!\right)}$ with*

- *rows labeled by pairs (μ, S) with $\mu \in \mathcal{M}_{d-2-2c}$ and $S \in \mathcal{P}_{c+1}[m]$,*
- *columns labeled by pairs (ν, T) with $\nu \in \mathcal{M}_{d-2c}$ and $T \in \mathcal{P}_c[m]$,*
- *where the coefficient of $((\mu, \{i_1, \ldots, i_{c+1}\}), (\nu, \{i_1, \ldots, i_{c+1}\} \setminus \{i_k\}))$ is given by $(-1)^{k+1}$ times the coefficient of ν in $\mu \cdot f_k$, for any $k \in [c+1]$.*

It is easy to check that $\mathrm{Rel}_0^d(\mathcal{F})$ is exactly the Macaulay matrix $\mathrm{Mac}^d(\mathcal{F})$. For brevity we may write $\mathrm{Rel}_c^d := \mathrm{Rel}_c^d(\mathcal{F})$ if $\mathcal{F}$ is clear from the context. Furthermore, we write $\mathrm{Rel}^d := \mathrm{Rel}_1^d$.

The relation matrix Rel^d encodes the trivial syzygies and we obtain

$$\operatorname{rowspace}\mathrm{Rel}^d \cong \langle f_i \otimes e_j - f_j \otimes e_i \mid 1 \le i < j \le m \rangle_{d-2}.$$

as subspace of $\mathcal{R}_{d-2}^m \cong \mathcal{R}_{d-2} \otimes \mathbb{F}^m$. Likewise, the higher order relation matrices encode the syzygies between syzygies and so forth. Note that this implies that $\mathrm{Rel}_c^d \circ \mathrm{Rel}_{c-1}^d = 0$ for any $c > 0$.[3]

We denote the space of syzygies of degree d and order c (resp. Koszul syzygies) by

$$\mathrm{Syz}_c^d := \operatorname{lker}\mathrm{Rel}_{c-1}^d \quad \text{and} \quad \mathrm{KSyz}_c^d := \operatorname{rowspace}\mathrm{Rel}_c^d\,.$$

where lker denotes the left kernel. We may refer to Koszul syzygies as trivial syzygies and to the first order syzygies simply as syzygies. We write $\mathrm{KSyz}^d := \mathrm{KSyz}_1^d$ and $\mathrm{Syz}^d := \mathrm{Syz}_1^d$. In general we have $\mathrm{KSyz}_c^d \subseteq \mathrm{Syz}_c^d$. If the sequence $(f_1, \ldots, f_m)$ is semi-regular and $d < d_{\mathrm{reg}}$, then it holds that $\mathrm{Syz}_c^d = \mathrm{KSyz}_c^d$ for any $c > 0$. Using inclusion-exclusion it follows that

$$\operatorname{null}\mathrm{Mac}^d = \sum_{c=0}^{m} (-1)^c \binom{m}{c} \left(\!\!\binom{n}{d-2c}\!\!\right). \tag{1}$$

which corresponds to the Hilbert series from [5, Proposition 6].

2.3 Monomial orders

The set $\mathcal{M}$ can be ordered in a way that is compatible with multiplication by using a monomial order. In this paper, we will work mainly in the `grevlex` order which is defined by $X_{a_d} \cdots X_{a_1} > X_{b_{d'}} \cdots X_{b_1}$ if and only if $d > d'$ or $d = d'$ and the first non-zero entry in $(a_d - b_d, \ldots, a_1 - b_1)$ is negative. Given a polynomial $f \in \mathcal{R}$ we define its leading monomial $\mathrm{LM}(f)$ to be the highest monomial in $\mathrm{supp}(f)$ with respect to `grevlex`.

In this paper, we will also need to specifically order the relation matrices Rel_c^d. For any d, c, we equip the set $\mathcal{M}_{d-2c} \times \mathcal{P}_c[m]$, which indexes the rows of Rel_{c-1}^d and columns of Rel_c^d, with the following order: $(\mu, S) < (\nu, T)$ if $\mu < \nu$ in the `grevlex` order or $\mu = \nu$ and the smallest element of $(S \cup T) \setminus (S \cap T)$ lies in S. We order the rows and columns in descending order so that we also get the matrix relation $\mathrm{Rel}_c^d \cdot \mathrm{Rel}_{c-1}^d = 0$.

[3] In fact, these maps constitute the Koszul complex of $\mathcal{R}/\langle\mathcal{F}\rangle$ by taking the differentials $d_c = (\mathrm{Rel}_c)^t$

$$\cdots \to \mathcal{R}^{\binom{m}{c+1}} \xrightarrow{d_c} \mathcal{R}^{\binom{m}{c}} \to \cdots \xrightarrow{d_1} \mathcal{R}^m \xrightarrow{d_0} \mathcal{R} \twoheadrightarrow \mathcal{R}/\langle\mathcal{F}\rangle.$$

3 Running XL on submatrices of Mac^d

To smooth-out the complexity of the XL algorithm, we consider submatrices $\mathbf{M}$ of Mac^d such that Mac^d can, up to reordering rows and columns, be partitioned as follows:

$$\mathrm{Mac}^d = \left[\begin{array}{c|c} * & * \\ \hline \mathbf{0} & \mathbf{M} \end{array}\right].$$

Let $C \subseteq \mathcal{M}_d$ denote the set indexing the columns of $\mathbf{M}$. Then, for any $\mathbf{v} \in \ker \mathrm{Mac}^d$, we have $(\mathbf{v}_i)_{i \in C} \in \ker \mathbf{M}$. If the nullity of $\mathbf{M}$ is sufficiently low, then we can solve the MQ-problem by finding the right kernel of $\mathbf{M}$, in the same manner the XL algorithm solves the problem by finding the right kernel of Mac^d. The complexity of this kernel finding depends heavily on the width $|C|$ of $\mathbf{M}$. Thus, it would be preferable to use submatrices $\mathbf{M}$ of Mac^d that have a low nullity and a small width.

For an $i \in [m]$ and $\mu \in \mathcal{M}_{d-2}$, the support of the row indexed by the pair (μ, i) is $\mathrm{supp}(\mu \cdot f_i) \subseteq \mu \cdot \mathcal{M}_2$. Thus, for a set of rows of Mac^d indexed by $S \times [m]$ for some subset $S \subseteq \mathcal{M}_{d-2}$, their combined support is at most $S \cdot \mathcal{M}_2$. To help describe this support, we introduce some terminology and notation.

Definition 5. *For an integer $c > 0$, let $\partial^c : \mathcal{PM} \to \mathcal{PM}$ be a mapping defined by $\partial^c(S) := S \cdot \mathcal{M}_c$, for every $S \subseteq \mathcal{M}$. We will refer to $\partial^c(S)$ as the cth shadow of S. Whenever $c = 1$ we may simply write ∂ instead of ∂^1.*

Note that one can describe ∂^c as the cth iteration of ∂. We can now rephrase our earlier observation: For any $S \subseteq M_{d-2}$, the support of the rows of Mac^d indexed by $S \times [m]$ is at most $\partial^2(S)$. Based on this observation, for $S \subseteq \mathcal{M}_{d-2}$, we denote by Mac^S the submatrix of Mac^d consisting of entries whose rows are labeled by $(\mu, i) \in S \times [m]$, and whose columns are labeled by $\nu \in \partial^2(S)$.

3.1 Initial segment induced submatrices of Mac^d

As indicated before, we would ideally use sets S that minimize $|\partial^2(S)|$ among all $T \subseteq \mathcal{M}_{d-2}$ that give matrices Mac^T of the same rank. Unfortunately, $\mathrm{rank}(\mathrm{Mac}^S)$ is rather hard to predict. Instead, we consider the height $|S| \cdot m$ of Mac^S as a proxy for its rank. This leaves us with the following question: What subsets $S \subseteq \mathcal{M}_{d-2}$ of any fixed size minimize $|\partial^2(S)|$? As we will show in this section, this minimum can be achieved by choosing $S \subseteq \mathcal{M}_{d-2}$ to be an initial segment as defined in the following definition:

Definition 6. *An initial segment of a totally ordered set X is a set of the form $\{x \in X \mid x \leq M\}$ for some largest element $M \in X$. We denote the initial segment of $\mathcal{M}_d$ with largest element μ as*

$$\mathcal{S}_{\leq \mu} = \{\nu \in \mathcal{M}_d \mid \nu \leq \mu\}.$$

We may also use the notation $\mathcal{S}_{<\mu} := \mathcal{S}_{\leq \mu} \setminus \{\mu\}$. We also regard the empty set as an initial segment, and may write $-\infty := \max(\varnothing)$ and $\mathcal{S}_{\leq -\infty} := \varnothing$.

To show that initial segments achieve this claimed minimum, we use a result by Clements and Lindström [18]. This result covers subsets $F_d \subseteq \mathbb{Z}^n_{\geq 0}$ whose elements sum to d, equipped with the lexicographical ordering. Through the bijection $\mathbb{Z}^n_{\geq 0} \to \mathcal{M}$ given by $(d_1, \ldots, d_n) \mapsto X_1^{d_1} \cdots X_n^{d_n}$, this corresponds to the reverse of the `grevlex`-ordering on $\mathcal{M}_d$. With that connection, we can rephrase [18, Corollary 1] as follows:

Lemma 1 ([18, Corollary 1]). *For $S \subseteq \mathcal{M}_d$, let $\mathrm{init}(S)$ denote the initial segment of $\mathcal{M}_d$ of size $|S|$. Then, for any $S \subseteq \mathcal{M}_d$, it holds that $\partial(\mathrm{init}(S)) \subseteq \mathrm{init}(\partial(S))$.*

Theorem 1. *Let $\mathcal{S}$ be an initial segment of $\mathcal{M}_d$. Then, for any $c > 0$, it holds that*

$$\left|\partial^c(\mathcal{S})\right| = \min_{\substack{T \subseteq \mathcal{M}_d \\ |T| = |\mathcal{S}|}} \left|\partial^c(T)\right|.$$

Proof. We prove this statement by induction over c. The base case $c = 1$ follows directly from Lemma 1. Now, let $T \subseteq \mathcal{M}_d$ be of size $|\mathcal{S}|$. Then, by Lemma 1, we have $\partial(\mathcal{S}) \subseteq \mathrm{init}(\partial(T))$, hence $\partial^c(\mathcal{S}) \subseteq \partial^{c-1}(\mathrm{init}(\partial(T)))$. Per the induction hypothesis, $|\partial^{c-1}(\mathrm{init}(\partial(T)))| \leq |\partial^{c-1}(\partial(T))| = |\partial^c(T)|$. We conclude that $\partial^c(\mathcal{S}) \leq |\partial^c(T)|$. □

This theorem shows that choosing $\mathcal{S}$ to be an initial segment optimizes the number of columns of $\mathrm{Mac}^{\mathcal{S}}$ with respect to the number of its rows. For this reason, we will limit our scope to submatrices of Mac^d of the form $\mathrm{Mac}^{\mathcal{S}}$, where $\mathcal{S} \subseteq \mathcal{M}$ is an initial segment. We will denote $\mathrm{Mac}^{\mathcal{S}_{\leq \mu}}$ as Mac^{μ}.

Note that $\mathcal{S}_{\leq X_1^{d-2}} = \mathcal{M}_{d-2}$ and $\mathrm{Mac}^{X_1^{d-2}} = \mathrm{Mac}^d$, so our approach encompasses the XL algorithm as a special case. More generally, consider monomials of the form $\mu_k = X_n^k X_1^{d-2-k}$. Then $\mathcal{S}_{\leq \mu_k} = X_n^k \cdot \mathcal{M}_{d-2-k}$. The entry of Mac^{μ_k} indexed by $((X_n^k \cdot \mu, i), X_n^k \cdot \nu)$, where $\mu \in \mathcal{M}_{d-2-k}, \nu \in \mathcal{M}_{d-k}$, is precisely the entry of Mac^{d-k} indexed by $((\mu, i), \nu)$. Hence, $\mathrm{Mac}^{\mu_k} = \mathrm{Mac}^{d-k}$. Thus, monomials $X_n X_1^{d-3} \leq \mu \leq X_1^{d-2}$, give intermediate options between the Macaulay matrices of degrees $d-1$ and d.

In addition to minimizing the size of $\partial^c(\mathcal{S})$, initial segments have another nice property regarding shadows that we will use later, namely, taking shadows preserves the property of being an initial segment.

Lemma 2. *For any $\mu \in \mathcal{M}_d$, it holds that $\partial^c(\mathcal{S}_{\leq \mu}) = \mathcal{S}_{\leq \mu \cdot X_1^c}$.*

Proof. Write $\mu = X_{a_1} \cdots X_{a_d}$. Let $\nu = X_{b_{c+d}} \cdots X_{b_1} \in \mathcal{S}_{\leq \mu \cdot X_1^c}$, meaning the leftmost non-zero entry of $(a_d - b_{c+d}, \ldots, a_1 - b_{c+1}, 1 - b_c, \ldots, 1 - b_1)$ is negative if it exists. Hence, the same applies to $(a_d - b_{c+d}, \ldots, a_1 - b_{c+1})$. Thus, $\nu' = X_{b_{c+d}} \cdots X_{b_{c+1}} \leq \mu$, and therefore $\nu = \nu' \cdot X_{b_c} \cdots X_{b_1} \in \partial^c(\mathcal{S}_{\leq \mu})$.

Conversely, let $\nu \cdot \xi \in \partial^c(\mathcal{S}_{\leq \mu})$, with $\nu = X_{b_c} \cdots X_{b_1} \in \mathcal{S}_{\leq \mu}$, and $\xi \in \mathcal{M}_c$. Write $\nu \cdot \xi = X_{b'_{c+d}} \cdots X_{b'_1}$. Since the indices $b'_{c+d} \geq \cdots \geq b'_1$ include $b_d \geq \cdots \geq b_1$, we have $b'_{c+i} \geq b_i$ for every $i \in [d]$. The leftmost non-zero entry of $(a_d - b_d, \ldots, a_1 - b_1)$ is negative if it exists, so the same holds for $(a_d - b'_{c+d}, \ldots, a_1 - b'_{c+1})$. That property extends to $(a_d - b'_{c+d}, \ldots, a_1 - b'_{c+1}, 1 - b'_c, \ldots, 1 - b'_1)$, since the entries $1 - b'_i$ for $i \in [c]$ cannot be positive. We conclude that $\nu \cdot \xi \leq \mu \cdot X_1^c$. □

4 Nullities of Macaulay matrices

In order to use the matrices Mac^μ in a generalization of the XL algorithm, we would like to be able to predict the nullity of Mac^μ, so that we can find a minimal $\mu \in \mathcal{M}_{d-2}$ for which we expect Mac^μ to have sufficiently low nullity.

As discussed in Section 3, for a monomial $\mu \in \mathcal{M}_{d-2}$, the Macaulay matrix Mac^d can be partitioned as follows:

$$\mathrm{Mac}^d = \left[\begin{array}{c|c} * & * \\ \hline \mathbf{0} & \mathrm{Mac}^\mu \end{array}\right].$$

From this partition, we see that any relation between the rows of Mac^μ corresponds uniquely to a relation between the $|\mathcal{S}_{\leq\mu}| \cdot m$ rows of Mac^d that contain Mac^μ. Thus, we can identify

$$\mathrm{Syz}^\mu := \operatorname{lker} \mathrm{Mac}^\mu \cong \mathrm{Syz}^d \cap \mathcal{R}^m_{\mathcal{S}_{\leq\mu}}.$$

Recall now that Syz^d has a subset KSyz^d of trivial syzygies. The overlap between these trivial syzygies and Syz^μ gives us a set KSyz^μ of trivial syzygies on Mac^μ. If we assume that all syzygies on Mac^μ are trivial, we can determine the nullity of Mac^μ by determining the number of trivial syzygies, since then

$$\operatorname{null} \mathrm{Mac}^\mu = \left|\partial^2(\mathcal{S}_{\leq\mu})\right| - m|\mathcal{S}_{\leq\mu}| + \dim(\mathrm{KSyz}^\mu). \tag{2}$$

The assumption that all syzygies on Mac^μ are trivial is elaborated upon in Section 4.1.

The remainder of Section 4 is largely spent analyzing the nullity of Mac^μ under that assumption. In particularin Corollary 2, we express an upper bound for $\operatorname{null} \mathrm{Mac}^\mu$. This expression involves dimensions of standard trivial syzygy spaces $\mathrm{KSyz}^{d'}$ of systems $\pi_i\mathcal{F}$. In Corollaries 3 and 4, we discuss how to compute bounds for these dimensions. Combined, these three corollaries give a computable upper bound for $\operatorname{null} \mathrm{Mac}^\mu$.

4.1 Regularity assumptions

The degree of regularity and the notion of semi-regularity as given in definitions 1 and 2 are too coarse for our approach. Still, we would like to work with the same assumption that no non-trivial relations occur in Mac^μ unless they combinatorially have to. We will therefore introduce the refinements of these notions tailored to initial segments.

Given a system $\mathcal{F} = (f_1, \ldots, f_m)$ of homogeneous polynomials in $\mathcal{R}$, we associate a vector space to each monomial $\mu \in \mathcal{M}$ with degree $d = \deg \mu$ and $1 \leq \ell \leq m$ given by

$$J_{\leq\mu,\ell} := \langle \nu \cdot f_i \mid 1 \leq i \leq \ell,\ \nu \in \mathcal{M}_{d-\deg f_i},\ \mathrm{LM}(\nu \cdot f_i) \leq \mu$$

Definition 7. *Given a sequence of polynomials $f_1, \ldots, f_m$, we define the monomial of regularity μ_{reg} for the sequence as*

$$\mu_{reg} := \min\left\{\mu \in \mathcal{M} \mid J_{\leq \mu, m} = \mathcal{R}_{\mathcal{S}_{\leq \mu}}\right\} \cup \{\infty\}.$$

Definition 8. *A sequence $f_1, \ldots, f_m$ of homogeneous polynomials in $\mathcal{R}$ is called* μ-regular *if*

$$gf_i \in J_{\leq \mu, i-1} \text{ and } \mathrm{LM}(gf_i) < \mu \implies g \in \langle f_1, \ldots, f_{i-1}\rangle.$$

for all $1 \leq i \leq m$. Furthermore, we call the sequence monomial semi-regular *if it is μ_{reg}-regular.*

Remark 1. As a sanity check for the soundness of our definition of μ-regularity, note that $I_d = J_{\leq X_1^d, m}$ and hence μ-regular sequences are also $(\deg \mu)$-regular. However, if one were to find a sequence where $\deg \mu_{\mathrm{reg}} < d_{\mathrm{reg}}$, then the sequence could be monomial semi-regular but not semi-regular in general.

Conjecture 1. Generic sequences are monomial semi-regular.

Since monomial semi-regular systems form an open set, this only requires the set to be non-empty just as in the semi-regular case. We tested this for random homogeneous quadratic systems of m equations and n variables and found monomial semi-regular sequences for each tested pair

$$\begin{aligned}(n, m) \in \{(n, n) \mid n \in \{5, \ldots, 8\}\} \cup \{(n, n+1) \mid n \in \{5, \ldots, 10\}\} \\ \cup \{(n, 2n) \mid n \in \{5, \ldots, 15\}\}.\end{aligned}$$

Note that one could, in principle, define the monomial of regularity for any monomial order. However, we only conjecture the existence of monomial semi-regular sequences for the `grevlex` order.

Let us connect this to the Macaulay matrix Mac^{μ}. If we assume that $\mathcal{F}$ is a homogeneous quadratic system such that $\mathrm{LM}(f_i) = X_1^2$ for all i (which can be assumed with high probability for a random instance over large fields), then we have the following isomorphism:

$$J_{\leq \mu \cdot X_1^2, m} \cong \mathrm{rowspace}\, \mathrm{Mac}^{\mu}\,.$$

Furthermore, in this case, we know that $\mu_{\mathrm{reg}} = \mu \cdot X_1^2$ for some μ.

Remark 2. We note that, on the other hand, each instance of the MQ problem is linear equivalent to an instance $\mathcal{F} = (f_1, \ldots, f_m)$ where each f_i has a different leading monomial. Considering such an equivalent system might give a lower monomial of regularity, which is not necessarily divisible by X_1^2. While this might give lower complexity, it complicates the analysis, and our experiments on small parameters indicate only minimal benefit for the complexity.

4.2 Anti-shadows

Let $\mathcal{S}$ be an initial segment of $\mathcal{M}_{d-2}$ with largest element μ. As outlined in the introduction of Section 4, we want to find elements of $\mathrm{KSyz}^d = \operatorname{rowspace} \mathrm{Rel}^d$ whose support is contained in $\mathcal{S} \times [m]$, as those elements give trivial syzygies on Mac^μ. A simple approach to finding such trivial syzygies would be to forego considering the entire rowspace of Rel^d, and merely ask which rows of Rel^d have a support contained in $\mathcal{S} \times [m]$. To help describe such rows, we introduce some more notation and terminology.

Definition 9. *Given an integer $c > 0$, let $\partial^{-c} : \mathcal{PM} \to \mathcal{PM}$ be defined as follows: $\partial^{-c}(S) := \{\mu \in \mathcal{M} \mid \mu \cdot \mathcal{M}_c \subseteq S\}$. We refer to $\partial^{-c}(S)$ as the cth anti-shadow of S.*

Just as in the case of shadows (Definition 5), for any $c \geq 0$, the map ∂^{-c} can be described as the cth iteration of ∂^{-1}.

Lemma 3. *For $c > 0$, the support of any row of Rel_c^d indexed by an element of $\partial^{-2c}(\mathcal{S}) \times \mathcal{P}_{c+1}[m]$ is contained in $\partial^{2-2c}(\mathcal{S}) \times \mathcal{P}_c[m]$.*

Proof. Consider the row of Rel_c^d indexed by $(\mu, T) \in \partial^{-2c}(\mathcal{S}) \times \mathcal{P}_{c+1}[m]$. An entry of this row is non-zero if, and only if, its index is of the form $(\nu, T \setminus \{i\})$ for some $i \in T$ with $\nu \in \operatorname{supp}(\mu \cdot f_i)$. Since $\operatorname{supp}(f_i) \subseteq \mathcal{M}_2$, it holds that $\operatorname{supp}(\mu \cdot f_i) \subseteq \partial^2(\{\mu\})$. We have $\partial^{-2c}(\mathcal{S}) = \partial^{-2}(\partial^{2-2c}(\mathcal{S}))$, so $\mu \in \partial^{-2c}(\mathcal{S})$ implies $\partial^2(\{\mu\}) \subseteq \partial^{2-2c}(\mathcal{S})$. □

Let sRel_c^μ denote the submatrix of Rel_c^d consisting of the entries whose rows are indexed by $\partial^{-2c}(\mathcal{S}) \times \mathcal{P}_{c+1}[m]$, and whose columns are indexed by $\partial^{2-2c}(\mathcal{S}) \times \mathcal{P}_c[m]$. Note that, $\mathrm{sRel}_0^\mu = \mathrm{Mac}^\mu$. We denote $\mathrm{sRel}^\mu := \mathrm{sRel}_1^\mu$, whose columns are indexed by $\mathcal{S} \times [m]$. By Lemma 3 we can partition $\mathrm{sRel}_c^d(\mathcal{F})$ as follows:

$$\mathrm{Rel}_c^d = \left[\begin{array}{c|c} * & * \\ \hline \mathbf{0} & \mathrm{sRel}_c^\mu \end{array}\right].$$

Since $\mathrm{Rel}_{c+1}^d \cdot \mathrm{Rel}_c^d = 0$, we find that $\mathrm{sRel}_{c+1}^\mu \cdot \mathrm{sRel}_c^\mu = 0$. One might now hope that the rowspace of sRel^μ contains all trivial syzygies on Mac^μ, and that $\operatorname{rowspace} \mathrm{sRel}_{c+1}^\mu = \operatorname{lker} \mathrm{sRel}_c^\mu$ if $\mathcal{F}$ is monomial semi-regular and $\mu \cdot X_1^2 < \mu_{\text{reg}}$. Unfortunately, this is not the case in general. If we do assume this to be the case, we can use the inclusion-exclusion principle to derive an expression for $\operatorname{null} \mathrm{Mac}^\mu$ similar to Equation (1).

Theorem 2. *Let $\mathcal{F} \subseteq \mathcal{R}$ be an instance of the MQ-problem, and let $\mathcal{S}$ be an initial segment of $\mathcal{M}_{d-2}$ with largest element μ. Suppose that, for any $c \geq 0$, it holds that $\operatorname{lker} \mathrm{sRel}_c^\mu = \operatorname{rowspace} \mathrm{sRel}_{c+1}^\mu$. Then the nullity of Mac^μ can be given as*

$$\operatorname{null} \mathrm{Mac}^\mu = \sum_{c \geq 0} (-1)^c \binom{m}{c} \left|\partial^{2-2c}(\mathcal{S})\right|. \tag{3}$$

We will later give an expression for $|\partial^{2-2c}(\mathcal{S})|$ in Proposition 1, which allows us to rewrite this as follows:

$$\operatorname{null}\operatorname{Mac}^{\mu} = \sum_{c\geq 0}(-1)^c \binom{m}{c}\left(\left(\!\!\binom{n}{d-2c}\!\!\right) - \sum_{\substack{1\leq k\leq d\\ a_k>1}}\left(\!\!\binom{a_k-1}{k+2-2c}\!\!\right)\right). \tag{4}$$

We are now interested to see when the premise of Theorem 2, that for every $c \geq 0$ the left kernel of sRel_c^{μ} is given by $\mathrm{sRel}_{c+1}^{\mu}$, holds. In practice we find there are many choices of m, n, q and μ where Equation (4) gives correct predictions. Consider Figure 1, in which we compare Equation (4) to experimentally found nullities of Mac^{μ} for randomly chosen $\mathcal{F} \in \mathcal{R}_2^m$, for the parameters $n = 10, m = 11$ and $q = 31$. We observe three types of behavior on different intervals.

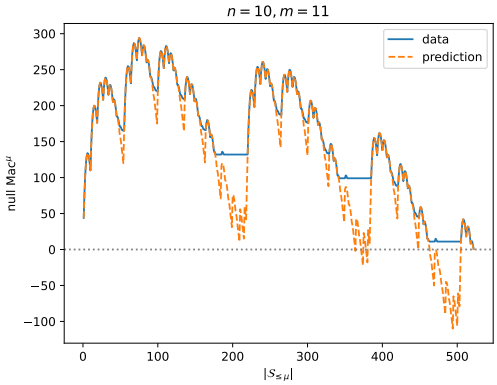

Fig. 1: Experimental Macaulay matrix nullities v.s. predictions from Theorem 2, Equation (4), for $n = 10$, $m = 11$, $q = 31$

(i) On many intervals, we find Equation (4) matches the data.
(ii) On some short intervals, we find discrepancies of sizes 1, 3, 6, ..., $\binom{m-1}{2}$ between Equation (4) and the experimentally found nullities. Each of these intervals end at some μ_0 divisible by X_1^2 but not X_1^3, with $\operatorname{null}\operatorname{Mac}^{\mu_0}$ being accurately given by Equation (4).
(iii) On each of the remaining intervals, we see the nullity largely sticks to some constant lower bound, with some spikes upwards. These intervals each end at some μ_0 divisible by X_1^3. The lower bound on this interval is $\operatorname{null}\operatorname{Mac}^{\mu_0}$, which is accurately given by Equation (4).

This behavior was consistently observed for every choice of (n, m) mentioned in Section 4.1. These experiments are discussed in more details in Section 5.

We give an explanation as to why the behavior described in point (ii) is expected. As mentioned in Remark 2, any instance of the MQ-problem is equivalent to an instance $\mathcal{F} = \{f_1, \ldots, f_m\} \subseteq \mathcal{R}_{\leq 2}$ where each polynomial f_i has a different leading monomial $\mathrm{LM}(f_i)$.

Let $\mu = X_{a_{d-2}} \cdots X_{a_1}$, and write $\mu_1 = X_{a_{d-2}} \cdots X_{a_3}$ and $\mu_2 = X_{a_2} X_{a_1}$, so that $\mu = \mu_1 \cdot \mu_2$. Suppose there are precisely $1 \leq r \leq m-2$ elements of $\mathcal{M}_2$ strictly larger then μ_2. Then there are at least $m-r$ elements $i \in [m]$ for which $\mathrm{LM}(f_i) < \mu_2$. Choose two such $i, j \in [m]$ with $i \neq j$, and consider the row of Rel^d indexed by $(\mu_1, \{i, j\})$. This row will have a support of

$$\big(\mathrm{supp}(\mu_1 \cdot f_j) \times \{i\}\big) \sqcup \big(\mathrm{supp}(\mu_1 \cdot f_i) \times \{j\}\big).$$

Since $\mu_1 \cdot \mathrm{LM}(f_i) < \mu_1 \cdot \mu_2 = \mu$, we have $\mathrm{supp}(\mu_1 \cdot f_i) \subseteq \mathcal{S}$. Likewise, it holds that $\mathrm{supp}(\mu_1 \cdot f_j) \subseteq \mathcal{S}$. Thus, the support of this row is a subset of $\mathcal{S} \times [m]$. However, as $\mu_1 \cdot X_1^2 > \mu_1 \cdot \mu_2 = \mu$, we have $\mu_1 \notin \partial^{-2}(\mathcal{S})$, so this row is not included in sRel^μ. Thus, Rel^d has at least $\binom{m-r}{2}$ rows whose support is a subset of $\mathcal{S} \times [m]$ which are not included in the submatrix sRel^μ. We note that these rows might still lie in $\mathrm{rowspace}\, \mathrm{sRel}^\mu$.

In Theorem 4, we give a concrete condition on m, n and μ that implies that $\mathrm{lker}\, \mathrm{sRel}_c^\mu = \mathrm{rowspace}\, \mathrm{sRel}_{c+1}^\mu$. This condition covers a significant portion of the parameters where we experimentally find that Equation (4) gives accurate predictions.

Sizes of shadows and anti-shadows In order to derive Equation (4) from Equation (3), we need to be able to compute the size of initial segments $\mathcal{S}_{\leq \mu}$, and their shadows and anti-shadows. We start with the size of $\mathcal{S}_{\leq \mu}$ itself.

Lemma 4. *Let $\mu = X_{a_d} \cdots X_{a_1}$. Then the size of the initial segment of $\mathcal{M}$ with μ as its largest element is*

$$\left|\mathcal{S}_{\leq \mu}\right| = \left(\!\!\binom{n}{d}\!\!\right) - \sum_{k=1}^{d} \left(\!\!\binom{a_k - 1}{k}\!\!\right).$$

Proof. Let us consider how many of monomials in $\mathcal{M}_d$ are strictly larger than μ. If $\nu = X_{b_d} \cdots X_{b_1} > \mu$, that means the first non-zero entry of $(b_d - a_d, \ldots, b_1 - a_1)$ is negative. Fixing the index k of this first non-zero entry $b_k - a_k$ implies $b_i = a_i$ for any $i > k$, and reduces choosing ν to choosing the degree k monomial $X_{b_k} \cdots X_{b_1}$ in the variables $(X_1, \cdots, X_{a_k - 1})$. Thus, $\mathcal{M}_d$ has $\sum_{k=1}^{d} \left(\!\!\binom{a_k - 1}{k}\!\!\right)$ elements ν strictly larger than μ. □

From Lemma 2, we know that ∂^c preserves the property of being an initial segment. We can show that the same holds for ∂^{-c}.

Lemma 5. *Let $\mathcal{S} \subseteq \mathcal{M}_d$ be an initial segment, and let $c > 0$. Then $\partial^{-c}(\mathcal{S})$ is an initial segment of $\mathcal{M}_{d-c}$.*

Proof. If $\partial^{-c}(\mathcal{S})$ is empty, it is the initial segment $\mathcal{S}_{\leq -\infty}$. Assume that $\partial^{-c}(\mathcal{S})$ is non-empty, and let $\nu = \max(\partial^{-c}(\mathcal{S}))$. Then $\partial^{-c}(\mathcal{S}) \subseteq \mathcal{S}_{\leq \nu}$. By Lemma 2, we have $\partial^c(\mathcal{S}_{\leq \nu}) = \mathcal{S}_{\leq \nu \cdot X_1^n}$. Since $\nu \in \partial^{-c}(\mathcal{S})$, it holds that $\nu \cdot X_1^c \in \mathcal{S}$, hence $\partial^c(\mathcal{S}_{\leq \nu}) = \mathcal{S}_{\leq \nu \cdot X_1^n} \subseteq \mathcal{S}$. Therefore, we have $\mathcal{S}_{\leq \nu} \subseteq \partial^{-c}(\mathcal{S})$. We conclude that $\mathcal{S}_{\leq \nu} = \partial^{-c}(\mathcal{S})$. □

Since an initial segment $\mathcal{S}$ is uniquely determined by $\max(\mathcal{S})$, it makes sense to view ∂^c as a map $\mathcal{M}_d \sqcup \{-\infty\} \to \mathcal{M}_{d+c} \sqcup \{-\infty\}$. For $\mu \in \mathcal{M}_d \sqcup \{-\infty\}$ and $c \in \mathbb{Z}$, we define $\partial^c(\mu) := \max(\partial^c(\mathcal{S}_{\leq \mu}))$, so that $\partial^c(\mathcal{S}_{\leq \mu}) = \mathcal{S}_{\leq \partial^c(\mu)}$.

Lemma 6. *Let $\mu = X_{a_d} \cdots X_{a_1}$ and $c \in \mathbb{Z}$. Then $\partial^c(\mu)$ can be described as follows:*

(i) If $c \geq 0$, then $\partial^c(\mu) = \mu \cdot X_1^c$
(ii) If $-d \leq c < 0$, and $a_{-c} = 1$, then $\partial^c(\mu) = X_{a_d} \cdots X_{a_{1-c}}$. If $c = -d$, we define $X_{a_d} \cdots X_{a_{1-c}}$ to be 1.
(iii) If $-d \leq c < 0$, and $a_{-c} > 1$, then $\partial^c(\mu)$ is the monomial in $\mathcal{M}_{d-c}$ directly preceding $X_{a_d} \cdots X_{a_{1-c}}$, or $-\infty$ if no such preceding monomial exists.
(iv) If $c < -d$, then $\partial^c(\mu) = -\infty$.

Proof. Statement (i) follows directly from Lemma 2. If $c < -d$, then $\partial^c(\mathcal{S}_{\leq \mu})$ is clearly empty, so statement (iv) holds as well.

Now, suppose $-d \leq c < 0$. For any $\nu \in \mathcal{M}_{d-c}$, it holds that $\partial^{-c}(\{\nu\}) \subseteq \mathcal{S}_{\leq \mu}$ if, and only if, $\max(\partial^{-c}(\{\nu\})) = \nu \cdot X_1^{-c} \leq \mu$. Thus, we can describe $\partial^c(\mu)$ as the largest monomial to satisfy $\partial^c(\mu) \cdot X_1^{-c} \leq \mu$, or $-\infty$ if no such monomial exists. If μ is divisible by X_1^{-c}, then this is simply μ / X_1^{-c}, which proves statement (ii).

Suppose finally that μ is not divisible by X_1^{-c}, meaning $a_{-c} > 1$. Write $\mu' = X_{a_d} \cdots X_{a_{1-c}}$. Then $\mu' \cdot X_1^{-c} > \mu$, so $\mu' \notin \partial^c(\mathcal{S}_{\leq \mu})$. For any $\nu < \mu'$ of degree $d - c$, it holds that $\nu \cdot X_1^{-c} < \mu$, hence $\nu \in \partial^c(\mathcal{S}_{\leq \mu})$. We conclude that $\partial^c(\mathcal{S}_{\leq \mu}) = \mathcal{S}_{< \mu'}$, which proves statement (iii). □

Combined, Lemmas 4 and 6 allow us to easily express the sizes of shadows and anti-shadows of initial segments.

Proposition 1. *Let $\mu \in \mathcal{M}_d$. Write $\mu = X_{a_d} \cdots X_{a_{\ell+1}} \cdot X_1^{\ell}$, where $a_d \geq \cdots \geq a_{\ell+1} > 1$. Then, for any $c \in \mathbb{Z}$, the initial segment $\partial^c(\mathcal{S}_{\leq \mu})$ has size*

$$\left|\partial^c(\mathcal{S}_{\leq \mu})\right| = \left(\!\!\binom{n}{d+c}\!\!\right) - \sum_{k=\ell+1}^{d} \left(\!\!\binom{a_k - 1}{k+c}\!\!\right).$$

Proof. If $c < -d$, then $\partial^c(\mathcal{S}_{\leq \mu})$ is empty, and the claim holds. If $c \geq -\ell$, then, by statements (i) and (ii) of Lemma 6, we have $\partial^c(\mu) = X_{a_d} \cdots X_{a_{\ell+1}} \cdot X_1^{\ell+c}$. Lemma 4 now gives us the claimed expression.

Now, suppose $-d \leq c < -\ell$. Lemma 6(iii), then tells us that $\partial^c(\mathcal{S}) = \mathcal{S}_{< \mu'}$ for $\mu' = X_{a_d} \cdots X_{a_{1-c}}$. Thus, by Lemma 4, we have

$$\left|\partial^c(\mathcal{S}_{\leq \mu})\right| = \left|\mathcal{S}_{\leq \mu'}\right| - 1 = \left(\!\!\binom{n}{d+c}\!\!\right) - 1 - \sum_{k=1-c}^{d} \left(\!\!\binom{a_k - 1}{k+c}\!\!\right).$$

Since $\sum_{k=\ell+1}^{-c} \left(\!\!\binom{a_k-1}{k+c}\!\!\right) = \left(\!\!\binom{a_{-c}-1}{0}\!\!\right) = 1$, we obtain the result. □

4.3 Partitioning $\mathrm{Rel}^d(\mathcal{F})$

Following Equation (2), if we can upper bound the dimension of KSyz^μ, then we can upper bound the nullity of Mac^μ for monomial semi-regular sequences as long as $\mu \cdot X_1^2 < \mu_{\mathrm{reg}}$. In Section 4.2, we described a submatrix sRel^μ of Rel^d that gives a subspace of KSyz^μ. However, this often does not describe all of KSyz^μ. To better study the structure of KSyz^μ, we give a partition of the cth order trivial syzygy matrix Rel_c^d which includes sRel_c^μ.

For the remainder of Section 4.3 we fix a quadratic system $\mathcal{F}$ and an initial segment $\mathcal{S}$ with top element $\mu = X_{a_{d-2}} \cdots X_{a_1}$. For $0 \leq \ell \leq d-2$, we define $\mu_\ell := X_{a_{d-2}} \cdots X_{a_{\ell+1}} X_1^\ell$ so that $\mu = \mu_0 \leq \mu_1 \leq \cdots \leq \mu_{d-2} = X_1^{d-2}$.

Lemma 7. *For $1 \leq \ell \leq d-2$, we can partition $\mathrm{sRel}_c^{\mu_\ell}$ as follows:*

$$\mathrm{sRel}_c^{\mu_\ell} = \left[\begin{array}{c|c} \mathrm{Rel}_c^{\ell+2}\left(\pi_{a_\ell - 1}\mathcal{F}\right) & * \\ \hline 0 & \mathrm{sRel}_c^{\mu_{\ell-1}} \end{array}\right].$$

Proof. If $a_\ell = 1$, then $\mu_{\ell-1} = \mu_\ell$ and the lemma trivially holds. We assume, from here on, that $a_\ell > 1$. By Lemma 3, we know any row of $\mathrm{sRel}_c^{\mu_\ell}$ that overlaps with $\mathrm{sRel}_c^{\mu_{\ell-1}}$ has no non-zero entries outside of that overlap.

For the remainder of this proof, we write $\mathcal{S}_\ell := \mathcal{S}_{\leq \mu_\ell}$. Furthermore, we write $\nu := X_{a_d} \cdots X_{a_{\ell+1}}$, so that $\mu_\ell = \nu \cdot X_1^\ell$, and $\mu_{\ell-1} = \nu \cdot X_{a_\ell} X_1^{\ell-1}$. The rows of $\mathrm{sRel}_c^{\mu_\ell}$ that do not overlap with $\mathrm{sRel}_c^{\mu_{\ell-1}}$ are those indexed by elements of $\left(\partial^{-2c}(\mathcal{S}_\ell) \setminus \partial^{-2c}(\mathcal{S}_{\ell-1})\right) \times \mathcal{P}_c[m]$. We claim that

$$\partial^{-2c}(\mathcal{S}_\ell) \setminus \partial^{-2c}(\mathcal{S}_{\ell-1}) = \nu \cdot \mathcal{M}_{\ell-2c}(X_1, \ldots, X_{a_\ell - 1}), \tag{5}$$

We prove this for a few distinct cases. If $d < 2c$, then clearly both the LHS and RHS of Equation (5) are empty. Now, suppose $d \geq 2c > \ell$. Then the RHS of Equation (5) is again clearly empty. Furthermore, by Lemma 6(iii), we have $\partial^{-2c}(\mathcal{S}_{\ell-1}) = \partial^{-2c}(\mathcal{S}_\ell) = \mathcal{S}_{<X_{a_d} \cdots X_{a_{2c+1}}}$. Thus, the LHS of Equation (5) is empty here as well.

Next, suppose $2c = \ell$. Then $\partial^{-2c}(\mathcal{S}_\ell) = \mathcal{S}_{\leq \nu}$ by Lemma 6(ii), and $\partial^{-2c}(\mathcal{S}_{\ell-1}) = \mathcal{S}_{<\nu}$ by Lemma 6(iii). Thus, the LHS of Equation (5) is $\{\nu\}$, which is the same as the RHS.

Finally, suppose $2c < \ell$. Lemma 6(ii) then tells us $\partial^{-2c}(\mu_{\ell-1}) = \nu \cdot X_{a_\ell} X_1^{\ell-1-2c}$, and $\partial^{-2c}(\mu_\ell) = \nu \cdot X_1^{\ell-2c}$. Thus, we find that

$$\begin{aligned} \partial^{-2c}(\mathcal{S}_\ell) \setminus \partial^{-2c}(\mathcal{S}_{\ell-1}) &= \{\xi \in \mathcal{M}_{d-2c} \mid \partial^{-2c}(\mu_{\ell-1}) < \xi \leq \partial^{-2c}(\mu_\ell)\} \\ &= \nu \cdot \{\xi \in \mathcal{M}_{\ell-2c} \mid X_{a_\ell} X_1^{\ell-2c-1} < \xi \leq X_1^{\ell-2c}\} \\ &= \nu \cdot \mathcal{M}_{\ell-2c}(X_1, \ldots, X_{a_\ell - 1}). \end{aligned}$$

Thus, as claimed, the rows of $\mathrm{sRel}_c^{\mu_\ell}$ that do not overlap with $\mathrm{sRel}_c^{\mu_{\ell-1}}$ are those indexed by $(\nu \cdot \mathcal{M}_{\ell-2c}(X_1, \ldots, X_{a_\ell - 1})) \times \mathcal{P}_{c-1}[m]$. By a similar argument, the columns of $\mathrm{sRel}_c^{\mu_\ell}$ that do not overlap with $\mathrm{sRel}_c^{\mu_{\ell-1}}$ are those indexed by $(\nu \cdot \mathcal{M}_{\ell+2-2c}(X_1, \ldots, X_{a_\ell - 1})) \times \mathcal{P}_{c-1}[m]$.

The entry of $\mathrm{sRel}_c^{\mu_\ell}$ indexed by $((\nu\cdot\zeta, S), (\nu\cdot\xi, T))$ is precisely the same as the entry of $\mathrm{Rel}_c^{\ell+2}$ indexed by $((\zeta, S), (\xi, T))$. The submatrix of $\mathrm{Rel}_c^{\ell+2}$, consisting of those entries indexed by $((\zeta, S), (\xi, T))$ where $\zeta, \xi \in \mathcal{M}(X_1, \ldots, X_{a_\ell - 1})$ is precisely $\mathrm{Rel}_c^{\ell+2}(\pi_{a_\ell - 1}\mathcal{F})$. Thus, the submatrix of $\mathrm{sRel}_c^{\mu_\ell}$ consisting of entries from those rows and columns that do not overlap with $\mathrm{sRel}_c^{\mu_{\ell-1}}$ is equal to $\mathrm{Rel}_c^{\ell+2}(\pi_{a_\ell - 1}\mathcal{F})$. □

Combined, this gives us a partition of $\mathrm{sRel}_c^{\mu_{d-2}} = \mathrm{Rel}_c^d$ that includes $\mathrm{sRel}_c^{\mu_0} = \mathrm{sRel}_c^{\mu}$. To save space, we denote $\mathbf{R}_{c,a}^d = \mathrm{Rel}_c^d(\pi_{a-1}\mathcal{F})$.

Corollary 1. *The trivial syzygy matrix allows for the following partition:*

$$\mathrm{Rel}_c^d = \left[\begin{array}{c|c|c|c|c} \mathbf{R}_{c,a_{d-2}}^d & * & * & * & * \\ \hline \mathbf{0} & \mathbf{R}_{c,a_{d-3}}^{d-1} & * & * & * \\ \hline \mathbf{0} & \mathbf{0} & \ddots & * & * \\ \hline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{R}_{c,a_1}^3 & * \\ \hline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathrm{sRel}_c^{\mu} \end{array}\right].$$

With this decomposition of the relation matrix we are now able to express the dimension of KSyz^{μ} in terms of the dimensions of standard trivial syzygy spaces.

Theorem 3. *Let $\mathcal{S} \subseteq \mathcal{M}_{d-2}$ be an initial segment with largest element $\mu = X_{a_{d-2}} \cdots X_{a_1}$, with $\mu \cdot X_1^2 < \mu_{reg}$. Then*

$$\dim \mathrm{KSyz}^{\mu} \leq \dim \mathrm{KSyz}^d - \sum_{1 \leq k \leq d-2} \dim \mathrm{KSyz}^{k+2}(\pi_{a_k - 1}\mathcal{F}).$$

Proof. For simplicity, we write $R_\ell := \mathcal{R}_{\mathcal{S}_{\leq \mu_\ell}}$ for the remainder of this proof. Recall that $\mathrm{KSyz}^{\mu_\ell} = \mathrm{KSyz}^d \cap R_\ell^m$ and note that $\operatorname{rowspace} \mathrm{sRel}^{\mu_\ell} \hookrightarrow \mathrm{KSyz}^{\mu_\ell}$. Then, by Lemma 7, we find

$$\begin{aligned} \mathrm{KSyz}^{\ell+2}(\pi_{a_\ell - 1}\mathcal{F}) &\cong (\operatorname{rowspace} \mathrm{sRel}^{\mu_\ell}) \,/\, \left(R_{\ell-1}^m \cap \operatorname{rowspace} \mathrm{sRel}^{\mu_\ell}\right) \\ &\hookrightarrow \mathrm{KSyz}^{\mu_\ell} \,/\, \left(R_{\ell-1}^m \cap \mathrm{KSyz}^{\mu_\ell}\right) = \mathrm{KSyz}^{\mu_\ell} \,/\, \mathrm{KSyz}^{\mu_{\ell-1}}. \end{aligned}$$

This leads to the following inequality:

$$\dim \mathrm{KSyz}^{\mu_\ell} - \dim \mathrm{KSyz}^{\mu_{\ell-1}} \geq \dim \mathrm{KSyz}^{\ell+2}(\pi_{a_\ell - 1}\mathcal{F})$$

The result now follows by considering the telescoping sum and noting that $\mathrm{KSyz}^{\mu_{d-2}} = \mathrm{KSyz}^d$. □

Corollary 2. *Let $\mathcal{S} \subseteq \mathcal{M}_{d-2}$ be an initial segment with largest element $\mu = X_{a_{d-2}} \cdots X_{a_1}$, with $\mu \cdot X_1^2 < \mu_{reg}$. Then*

$$\operatorname{null} \mathrm{Mac}^{\mu} \leq |\partial^2(\mathcal{S})| - m|\mathcal{S}| + \dim \mathrm{KSyz}^d - \sum_{1 \leq k \leq d-2} \dim \mathrm{KSyz}^{k+2}(\pi_{a_k - 1}\mathcal{F}).$$

4.4 Bounding the rank of KSyz

In order to make full use of Corollary 2, we would like strong bounds on the dimensions of the vector spaces $\mathrm{KSyz}^{k+2}(\pi_{a_k-1}\mathcal{F})$. Unfortunately, after projecting, $k+2$ might be much higher than the degree of regularity of $\pi_{a_k-1}\mathcal{F}$ and our usual formulas might break. However, when we assume that each $\pi_{a_k-1}\mathcal{F}$ is again semi-regular, we can still provide some bounds.

Lemma 8. *Let $\mathcal{F} = (f_1, \ldots, f_m)$ be a homogeneous quadratic sequence of polynomials in $\mathcal{R}$. Then for any $c > 0$ and d we have:*

1. $\dim \mathrm{KSyz}_c^d \geq \dim \mathrm{KSyz}_{c-1}^{d-2} \mathcal{F}_{(m-1)} + \dim \mathrm{KSyz}_c^d \mathcal{F}_{(m-1)}$
2. $\dim \mathrm{Syz}_{c+1}^d \leq \dim \mathrm{Syz}_c^{d-2} \mathcal{F}_{(m-1)} + \dim \mathrm{Syz}_{c+1}^d \mathcal{F}_{(m-1)}$

Proof. We use the following decomposition of Rel_c,

$$\mathrm{Rel}_c : \mathcal{R}^{\binom{m-1}{c-1}} \oplus \mathcal{R}^{\binom{m-1}{c}} \to \mathcal{R}^{\binom{m-1}{c}} \oplus \mathcal{R}^{\binom{m-1}{c+1}}$$

given in degree d by

$$\mathrm{Rel}_c^d = \left[\begin{array}{c|c} \mathrm{Rel}_{c-1}^{d-2} \mathcal{F}_{(m-1)} & * \\ \hline \mathbf{0} & \mathrm{Rel}_c^d \mathcal{F}_{(m-1)} \end{array}\right].$$

Here, the rowspace (resp. column space) is decomposed by whether $m \in S$ or $m \notin S$ where (S, μ) is the label of the row (resp. column). Since Syz_{c+1}^d and KSyz_c^d are, respectively, the left kernel and rowspace of Rel_c^d, the result follows. □

In general, this gives us the following practical bound:

Corollary 3. *Let $\mathcal{F}$ be a sequence of homogeneous quadratic polynomials in $\mathcal{R}$, then*

$$\dim \mathrm{KSyz}^d \mathcal{F} \geq \operatorname{rank} \mathrm{Mac}^{d-2} \mathcal{F}_{(m-1)} + \dim \mathrm{KSyz}^d \mathcal{F}_{(m-1)}.$$

In low enough degrees, we can be more precise.

Lemma 9. *Let $\mathcal{F}$ be a sequence of homogeneous quadratic polynomials in $\mathcal{R}$ such that $\mathcal{F}_{(m-1)}$ is d'-regular. Then, for $c > 0$ and $d < d' + 2(c-1)$, we have*

$$\mathrm{Syz}_c^d = \mathrm{KSyz}_c^d.$$

Proof. Note that d-regularity implies that $\mathrm{Syz}_1^d \mathcal{F}_i = \mathrm{KSyz}_1^d \mathcal{F}_i$ for $d < d'$ and any $1 \leq i < m$. Also note that $\mathrm{Syz}_c^d \mathcal{F}_{(0)} = \mathrm{KSyz}_c^d \mathcal{F}_{(0)} = 0$ for any c, d. We proceed by induction on c and m. By Lemma 8 and the induction hypothesis, we obtain

$$\dim \mathrm{Syz}_c^d - \dim \mathrm{KSyz}_c^d \leq \dim \mathrm{Syz}_c^d \mathcal{F}_{(m-1)} - \dim \mathrm{KSyz}_c^d \mathcal{F}_{(m-1)}.$$

Hence $\dim \mathrm{Syz}_c^d \leq \dim \mathrm{KSyz}_c^d$ and the result then follows from $\mathrm{KSyz}_c^d \subseteq \mathrm{Syz}_c^d$. □

Corollary 4. *Let $\mathcal{F}$ be a homogeneous quadratic sequence in $\mathcal{R}$ such that $\mathcal{F}_{(m-1)}$ is semi-regular. If $d < d_{reg}(\mathcal{F}_{(m-1)}) + 2$ then*

$$\dim \mathrm{KSyz}^d(\mathcal{F}) = \sum_{c\geq 2} (-1)^c \binom{m}{c} \left(\!\!\binom{n}{d-2c}\!\!\right) .$$

Proof. First, note that $\dim \mathrm{KSyz}^d_{c-1} = \binom{m}{c}\left(\!\!\binom{n}{d-2c}\!\!\right) - \dim \mathrm{Syz}^d_c$ by rank-nullity. By Lemma 9 we find $\dim \mathrm{Syz}^d_c = \dim \mathrm{KSyz}^d_c$ for any $c \geq 2$. Hence, we get the required alternating sum. □

Together, Corollaries 3 and 4, give us exactly the tools that we need to use Corollary 2 on sequences for which each $\pi_i\mathcal{F}$ is semi-regular again. In Figure 2, we plotted the upper bound described by these corollaries against the experimental nullities for $n = 10$, $m = 11$ and $n = 17$, $m = 34$ in $\mathbb{F}_{251}$. These experiments are discussed in more detail in Section 5.

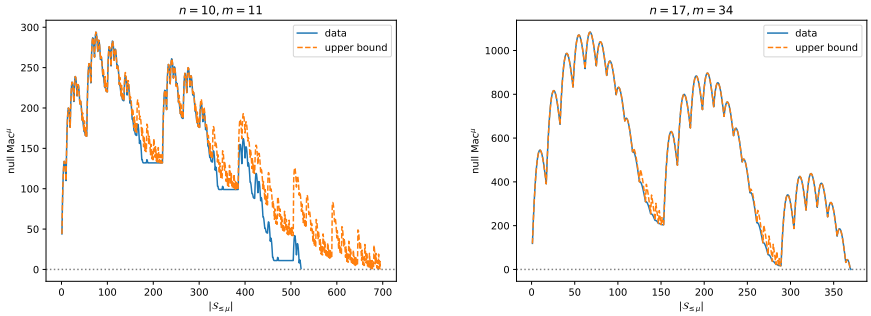

Fig. 2: Comparison experimental nullity with upper bound on the nullity.

If it additionally holds that $k < d_{\mathrm{reg}}(\pi_{a_k-1}\mathcal{F}_{(m-1)})$ for every $1 \leq k \leq d-2$ where $a_k > 1$, one can use Corollary 4 to compute the bound from Corollary 2 exactly. The resulting expression precisely matches Equation (4) from Theorem 2. In fact, under these conditions, we can show that the premise of Theorem 2 is met, which allows us to predict the nullity exactly.

Theorem 4. *Let $\mathcal{F} = f_1, \ldots, f_m$ be a monomial semi-regular quadratic sequence in $\mathcal{R}$ such that $\pi_i\mathcal{F}$ is semi-regular for each $1 \leq i \leq n$. Let $\mu = X_{a_{d-2}} \cdots X_{a_1}$. Suppose $\mu \cdot X_1^2 < \mu_{reg}$, and that*

$$k < d_{reg}(\pi_{a_k-1}\mathcal{F}_{(m-1)}) \quad \text{for all } 1 \leq k \leq d-2 \text{ such that } a_k > 1. \tag{6}$$

Then, for any $c \geq 1$ it holds that

$$\mathrm{lker}\left(\mathrm{sRel}^{\mu}_{c-1}\right) = \mathrm{rowspace}\left(\mathrm{sRel}^{\mu}_{c}\right).$$

Proof. As in Section 4.3, we denote $\mu_\ell := X_{a_{d-2}} \cdots X_{a_{\ell+1}} X_1^\ell$ for $0 \leq \ell \leq d-2$, so that $\mu = \mu_0 \leq \cdots \leq \mu_{d-2} = X_1^{d-2}$. Note that, if Equation (6) holds for μ, then it holds for each μ_ℓ.

Per the assumption that $\mu \cdot X_1^2 < \mu_{\text{reg}}$, the left kernel of sRel_{c-1}^{μ} is given by the elements of rowspace Rel_c^d whose support is contained in $\partial^{2-2c}(\mathcal{S}_{\leq \mu}) \times \mathcal{P}_c[m]$. We will proof, by backwards induction over ℓ, that any element of rowspace Rel_c^d whose support is contained in $\partial^{2-2c}(\mathcal{S}_{\leq \mu_\ell}) \times \mathcal{P}_c[m]$ lies in the rowspace of $\text{sRel}_c^{\mu_\ell}$. If $\ell = d-2$, then $\text{sRel}_c^{\mu_\ell} = \text{Rel}_c^d$, so the base case holds.

Suppose that the claim holds for some fixed $0 < \ell \leq d-2$, and that $a_\ell > 1$ so that $\mu_{\ell-1} \neq \mu_\ell$. Let $\mathbf{v} \in$ rowspace Rel_c^d, and suppose $\text{supp}(\mathbf{v}) \subseteq \partial^{2-2c}(\mathcal{S}_{\leq \mu_{\ell-1}}) \times \mathcal{P}_c[m]$. This implies $\text{supp}(\mathbf{v}) \subseteq \partial^{2-2c}(\mathcal{S}_{\leq \mu_\ell}) \times \mathcal{P}_c[m]$, so the induction hypothesis tells us that $\mathbf{v} \in$ rowspace $\text{sRel}_c^{\mu_\ell}$.

Choose $\mathbf{w}$ so that $\mathbf{w}^\top \text{Rel}_c^{\mu_\ell} = \mathbf{v}$, and write $\mathbf{w}'$ for the entries of $\mathbf{w}$ not indexed by $\partial^{-2c}(\mathcal{S}_{\leq \mu_{\ell-1}}) \times \mathcal{P}_{c+1}[m]$. By Lemma 7, the fact that $\mathbf{v}$ is only supported in the columns containing $\text{sRel}_c^{\mu_{\ell-1}}$ implies that $\mathbf{w}' \in \text{lker}\, \text{Rel}_c^{\ell+2}(\pi_{a_\ell - 1}(\mathcal{F}))$.

Together with Equation (6), Lemma 9 implies that $\text{lker}\, \text{Rel}_c^{\ell+2}(\pi_{a_\ell-1}(\mathcal{F}))$ is the rowspace of $\text{Rel}_{c+1}^{\ell+2}(\pi_{a_\ell-1}(\mathcal{F}))$, which is a submatrix of $\text{sRel}_{c+1}^{\mu_\ell}$ by Lemma 7. Thus, any element of $\text{lker}\, \text{Rel}_c^{\ell+2}(\pi_{a_\ell-1}(\mathcal{F}))$ can be extended into an element of rowspace $\text{sRel}_{c+1}^{\mu_\ell} \subseteq \text{lker}\, \text{sRel}_c^{\mu_\ell}$.

In particular, this applies to $\mathbf{w}'$. Thus, there exists a $\mathbf{w}'' \in \text{lker}\, \text{sRel}_c^{\mu_\ell}$ that agrees with $\mathbf{w}$ on any entries not indexed by $\partial^{-2c}(\mathcal{S}_{\leq \mu_{\ell-1}}) \times \mathcal{P}_c[m]$. We conclude that $(\mathbf{w} - \mathbf{w}'')\, \text{sRel}_c^{\mu_\ell} = \mathbf{w}\, \text{sRel}_c^{\mu_\ell} = \mathbf{v}$ lies in the rowspace of $\text{Rel}_c^{\mu_{\ell-1}}$. □

5 Comparison and experiments

Now that we have the tools in place to bound the monomial of regularity, we study how this new approach compares to classic XL and hybridization variants like FXL. Furthermore, we compare the bound on the monomial of regularity with the actual monomial of regularity in experiments.

Since we will be solving quadratic systems, there is one caveat that we must address. For systems that have a solution, the degree (and monomial) of regularity must be ∞ since I_d can never equal $\mathcal{R}_d$. This does mean that random systems with a planted solution will not be monomial semi-regular. For semi-regularity this is addressed in [38] where the authors introduce the generalized degree of regularity. We expect a similar generalization can be made for the monomial of regularity. However, we will be pragmatic and assume that random systems are μ-regular where μ is the lowest such that null $\text{Mac}^\mu = 1$. We experimentally verified this for the same (n, m) pairs as in Section 4.1.

To test our predictions we ran some experiments using `sagemath` [45]. The code for experiments and predictions can be found at:

https://github.com/LarsMath/monomial-of-regularity

In these experiments, we construct random quadratic systems with a planted solution. Then, we construct the Mac^μ sequentially, computing its nullity, until

we reach 1. At these points we verified that we could extract the planted solution. In Table 1, we compare the amount of columns $\left|\partial^2(\mathcal{S}_{\leq\mu})\right|$ for the actual monomial of regularity, our prediction for the monomial of regularity and the columns of XL in degree d_{reg}. For the prediction we used Corollaries 2, 3, and 4. It seems that the bigger the difference $m-n$, the more precise the predictions are, with $m=2n$ being correct in all tested cases. In this regime, we only included the results with the biggest improvements over XL.

n	m	$\tilde{d}_{\text{reg}}$	experimental	prediction	XL
5	5	5	100	119	126
6	6	6	321	427	462
7	7	7	1105	1561	1716
8	8	8	3822	5785	6435
8	9	5	701	770	792
9	10	6	1938	2057	3003
10	11	6	4255	4914	5005
11	12	7	11798	12922	19448
9	11	5	1034	1126	1287
10	12	6	2967	3002	5005
11	13	6	5815	6025	8008
12	14	6	10373	12285	12376
11	16	5	2137	2137	3003
12	18	5	3042	3045	4368
13	19	5	5270	5326	6188
14	21	5	7565	7705	8568
9	18	4	284	284	495
10	20	4	501	501	715
16	32	5	6722	6722	15504
17	34	5	11441	11441	20349

Table 1: Number of columns of Mac^{μ} for lowest μ with $\operatorname{null}\text{Mac}^{\mu}=1$ for different values of m, n over $\mathbb{F}_{251}$.

For higher n, m we cannot compute the actual monomial of regularity, but we can still predict it. In Figure 3 we plot the amount of columns for this approach and for classic XL for parameter sets with either $m=n+1$, $m=2n$ or $m=5n$. One can clearly see the jumps in classic XL that are caused by the discreteness of d_{reg}. In that sense, the monomial of regularity really smooths this complexity.

In Table 2, we compare the amount of columns for this approach and for classic XL. We consider parameter sets with either $m=n$, $m=n+1$, $m=2n$ or $m=5n$. Within those restrictions, we focus on those choices of n where the degree of regularity increases, as those are the parameters where our approach gives the greatest improvement.

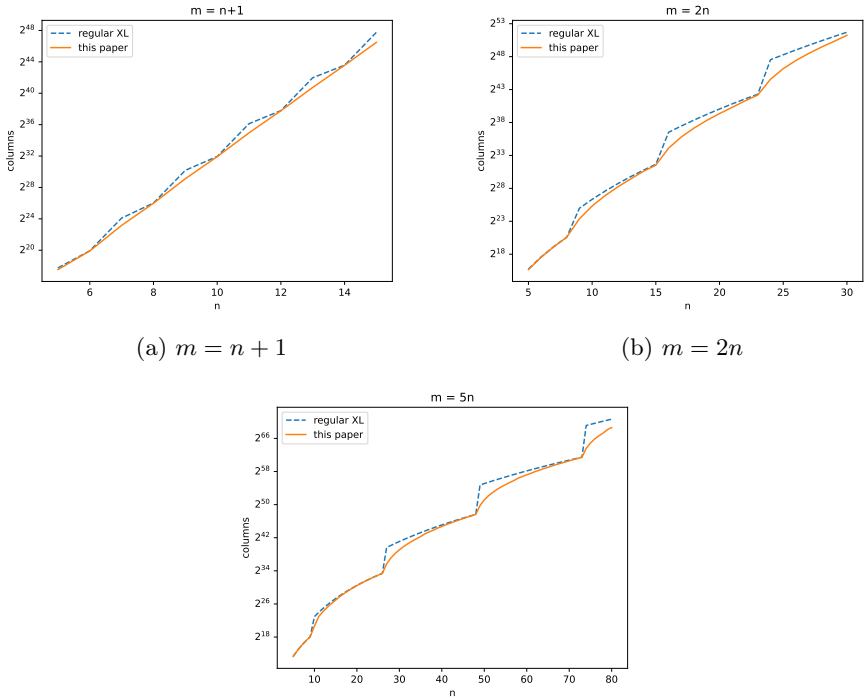

(a) $m = n + 1$ (b) $m = 2n$

(c) $m = 5n$

Fig. 3: Predicted amount of columns compared to classic XL for $m = n + 1$, $m = 2n$ and $m = 5n$.

5.1 Comparing to hybridization techniques

When the XL algorithm was introduced in [19], the authors provided a hybrid variant FXL wherein one fixes variables before solving. This can also be considered to "smooth" complexity. However, the complexity of FXL is dependent on the field size, whereas our optimization is agnostic to it. Furthermore, if we allow for a similar hybridization technique, our approach is strictly more general.[4]

Assuming both FXL and our approach are implemented using Wiedemann's algorithm for sparse linear algebra [47], we compare these algorithms using the following complexity estimates (in field operations) for the field $\mathbb{F}_q$. The results of this comparison can be found in Table 3. For FXL we estimate a cost of

$$\min\left\{3q^k \left(\!\!\binom{n-k}{2}\!\!\right)\left(\!\!\binom{n-k}{d}\!\!\right)^2 : 0 \leq k \leq n,\ 2 \leq d, \operatorname{null} \operatorname{Mac}^d = 1\right\}, \tag{7}$$

[4] In [20], the authors introduce another hybridization technique, namely post-XL kernel enumeration. In our complexity analysis we always found $\operatorname{null} \operatorname{Mac}^\mu = 1$ to be optimal, so we leave this out of further analysis.

n	m	q	d	$\log_2(\lvert\mathcal{M}_d\rvert)$	$\log_2(\lvert\partial^2\mathcal{S}_{\leq\mu}\rvert)$	n	m	q	d	$\log(\lvert\mathcal{M}_d\rvert)$	$\log(\lvert\partial^2\mathcal{S}_{\leq\mu}\rvert)$
9	9	31	9	36.2	35.9	24	48	31	6	47.5	44.6
10	10	31	10	40.4	40.0	33	66	31	7	58.5	55.2
11	11	31	11	44.5	44.1	41	82	31	8	68.3	64.7
12	12	31	12	48.6	48.2	50	100	31	9	78.5	74.7
13	13	31	13	52.7	52.3	60	120	31	10	89.0	85.2
14	14	31	14	56.8	56.4	69	138	31	11	98.9	94.8
15	15	31	15	60.9	60.5	79	158	31	12	109.1	105.0
16	16	31	16	65.0	64.6	88	176	31	13	118.8	114.5
17	17	31	17	69.1	68.7	98	196	31	14	128.9	124.5
18	18	31	18	73.2	72.7	108	216	31	15	138.9	134.5
13	14	31	8	42.0	40.7	10	50	31	3	22.9	20.7
15	16	31	9	47.8	46.5	27	135	31	4	39.6	35.7
17	18	31	10	53.5	52.2	49	245	31	5	54.7	49.8
19	20	31	11	59.2	57.9	74	370	31	6	69.1	63.6
21	22	31	12	64.9	63.5	100	500	31	7	82.9	77.0
23	24	31	13	70.6	69.2	127	635	31	8	96.4	90.2
25	26	31	14	76.3	74.8	155	775	31	9	109.8	103.4
27	28	31	15	81.9	80.5	183	915	31	10	123.1	116.4
29	30	31	16	87.6	86.1	212	1060	31	11	136.3	129.5
31	32	31	17	93.2	91.7	242	1210	31	12	149.6	142.7

Table 2: Predicted amount of columns compared to classic XL

and for our approach a cost of

$$\min\left\{3q^k\left(\binom{n-k}{2}\right)\left|\partial^2(\mathcal{S}_{\leq\mu})\right|^2 : k\leq n,\ \mu\in\pi_{n-k}\mathcal{M}, \operatorname{null}\operatorname{Mac}^\mu = 1\right\}. \quad (8)$$

6 Applications to recent attacks on UOV

We now turn to applying our technique to a concrete example, namely to a recently introduced attack on the UOV signature scheme. Note that similar analysis can be made to apply the smoothing to other algebraic attacks.

The Unbalanced Oil and Vinegar scheme (UOV) [35,14] is a signature scheme whose construction is based on the hardness of the MQ problem. The public key in this scheme is an MQ map $\mathcal{F}$ which is constructed such that its kernel includes a secret o-dimensional subspace called the oil-space. Three variants of UOV stand out, MAYO [11], QR-UOV [30], and SNOVA [46]. These all share the UOV construction, but also introduce additional structure.

One method to attack schemes from the UOV family is to straightforwardly use XL, or similar algorithms, to solve the underlying MQ problem. In [43], a novel approach was introduced that attacks instances of UOV over fields of characteristic 2 by using Macaulay matrices in the exterior algebra instead. Subsequently, in [33] this approach was extended to work for odd characteristic in

n	m	q	FXL			this paper		
			cost	k	d	cost	k	$\lvert\mathcal{S}_{\leq\mu}\rvert$
10	10	31	35.1	1	6	34.0	1	274
11	11	31	36.9	1	6	36.8	1	675
12	12	31	41.1	1	7	39.9	1	1676
13	13	31	42.7	1	7	42.7	2	609
14	14	31	45.0	2	6	44.9	2	1325
15	15	31	48.5	1	8	47.8	2	3145
16	16	31	50.7	2	7	50.4	2	7026
17	17	31	54.3	1	9	53.3	2	17823
18	18	31	56.4	2	8	55.8	2	39471
19	19	31	58.3	3	7	58.3	3	15358
20	20	31	62.1	2	9	61.1	2	219794
21	21	31	64.0	3	8	63.5	3	77922
22	22	31	67.7	2	10	66.4	3	201442

n	m	q	FXL			this paper		
			cost	k	d	cost	k	$\lvert\mathcal{S}_{\leq\mu}\rvert$
10	10	251	38.2	1	6	37.1	1	274
11	11	251	39.9	1	6	39.9	1	675
12	12	251	44.1	1	7	42.9	1	1676
13	13	251	45.8	1	7	45.7	1	4222
14	14	251	49.9	1	8	48.7	1	10410
15	15	251	51.5	1	8	51.5	1	26594
16	16	251	55.7	1	9	54.4	1	65354
17	17	251	57.3	1	9	57.3	1	168542
18	18	251	61.5	1	10	60.2	1	413699
19	19	251	63.0	1	10	63.0	1	1074062
20	20	251	67.2	1	11	65.8	1	2636337
21	21	251	68.7	1	11	68.7	1	6878509
22	22	251	72.9	1	12	71.5	1	16893514

Table 3: Comparison with FXL over $\mathbb{F}_q$ using formulas Equation (7) and Equation (8). All the costs are given in field operations $\log_2$.

the p-reduced algebra. Later, [29] built upon this by considering Macaulay matrices over the p^ℓ-reduced algebra which they called the p^ℓ-truncated polynomial rings. The perspective of 2-truncated polynomial rings is the most natural for our optimization.

6.1 2-truncated rings

Let us first briefly recall the framework from [29]. We only discuss the approach with 2-truncated rings, instead of more general p^ℓ-truncated rings.

Let $\mathcal{M}_d^{(2)}(\mathbf{X}) \subseteq \mathcal{M}_d(\mathbf{X})$ denote the set of degree d monomials whose exponent for each variable is at most 1. We may write $\mathcal{M}_d^{(2)} := \mathcal{M}_d^{(2)}(\mathbf{X})$ if $\mathbf{X}$ is clear from context. Furthermore, we write $\mathcal{M}^{(2)} := \bigcup_{d\geq 0}^{n} \mathcal{M}_d^{(2)} \subseteq \mathcal{M}$, which gives a basis for the $\mathbb{F}_q$-module $\mathcal{R}^{(2)} := \mathbb{F}_q[X_1, \ldots, X_n]/\langle X_1^2, \cdots, X_n^2\rangle$ which is isomorphic to the exterior algebra.

We define a map $v_d : \mathbb{F}_q^n \to \mathbb{F}_q^{\left(\!\binom{n}{d}\!\right)}$ by $v_d(x_1, \ldots, x_n)_i = x_{a_d} \cdots x_{a_1}$, where $X_{a_1} \cdots X_{a_d}$ is the ith smallest element of $\mathcal{M}_d$ in the `grevlex`-ordering, so that

$$\operatorname{Mac}^d v_d(\mathbf{x}) = \big(\big(\mu f_i\big)(\mathbf{x})\big)_{\mu\in\mathcal{M}_{d-2}, f_i\in\mathcal{F}}.$$

Let $\mathcal{F} = (f_1, \ldots, f_m)$ be an MQ map over a field of characteristic 2, with an o-dimensional oil-space O. Then, for any $\mathbf{x} \in O$, it holds that $v_d(\mathbf{x}) \in \ker(\operatorname{Mac}^d)$.

If $d \leq o$, then $\langle v_d(\mathbf{x}) \mid \mathbf{x} \in O\rangle_{\mathbb{F}_q} \subseteq \ker(\operatorname{Mac}^d)$ will contain an $\binom{o}{d}$-dimensional subspace of vectors whose support is restricted to those entries corresponding to $\mathcal{M}_d^{(2)}$. Supposing that we have found this subspace, we can use $\binom{o}{d} - 1$ linear restrictions to isolate an element which can be used to recover one or more

vectors from O. Finding a single vector from the oil-space significantly reduces the complexity of finding the entire oil-space.

To help describe this subspace, we adapt our previous definitions of Macaulay and relation matrices to 2-truncated rings.

Definition 10. *The 2-truncated degree d Macaulay matrix of $\mathcal{F}$ is the matrix* $\mathrm{Mac}^{(2),d}(\mathcal{F}) \in \mathbb{F}^{m\binom{n}{d-2}\times\binom{n}{d}}$ *with*

- *rows labeled by pairs (μ, i) with $\mu \in \mathcal{M}^{(2)}_{d-2}$ and $i \in [m]$,*
- *columns labeled by ν with $\nu \in \mathcal{M}^{(2)}_{d}$,*
- *where the coefficient of $((\mu, i), \nu)$ is given by the coefficient of ν in $\mu \cdot f_i$.*

The aforementioned $\binom{o}{d}$-dimensional subspace of $\langle v_d(\mathbf{x}) \mid \mathbf{x} \in O\rangle_{\mathbb{F}_q} \subseteq \ker(\mathrm{Mac}^d)$, whose support is restricted to those entries corresponding to $\mathcal{M}^{(2)}_d$, is found in the right kernel of $\mathrm{Mac}^{(2),d}$.

The left kernel of this matrix includes the rowspace of the matrix given by the following definition:

Definition 11. *The 2-truncated degree d relation matrix of $\mathcal{F}$ is the matrix* $\mathrm{Rel}^{(2),d}(\mathcal{F}) \in \mathbb{F}^{\binom{\binom{m}{2}}{}\binom{n}{d-4}\times m\binom{n}{d-2}}$ *with*

- *rows labeled by pairs (μ, S) with $\mu \in \mathcal{M}^{(2)}_{d-4}$ and S a multiset of 2 elements from $[m]$,*
- *columns labeled by pairs (ν, i) with $\nu \in \mathcal{M}^{(2)}_{d-2}$ and $i \in [m]$,*
- *where the coefficient of $((\mu, S), (\nu, i))$ is given by the coefficient of ν in $\mu \cdot f_j$ if $S = \{i, j\}$, and 0 otherwise.*

If $\mathcal{F}$ is clear from context, we may write $\mathrm{Mac}^{(2),d} := \mathrm{Mac}^{(2),d}(\mathcal{F})$ and $\mathrm{Rel}^{(2),d} := \mathrm{Rel}^{(2),d}(\mathcal{F})$. We write $\mathrm{KSyz}^{(2),d} := \operatorname{rowspace}\mathrm{Rel}^{(2),d} \subseteq \operatorname{lker}\mathrm{Mac}^{(2),d}$ for the space of trivial syzygies.

The definition of homogeneity translates easily to $\mathcal{R}^{(2)}$. The notion of semi-regularity and the degree of regularity d_{reg} for such systems are introduced in [5]. Note that these definitions are for polynomial rings over $\mathbb{F}_2$ but can be easily generalized for polynomial rings over $\mathbb{F}_{2^\ell}$ by extension of scalars. To avoid confusion on the base field, we will rename the property.

Definition 12. *A sequence $f_1, \ldots, f_m$ of homogeneous polynomials in $\mathcal{R}^{(2)}$ is Boolean d-regular over if*

$$gf_i \in \langle f_1, \ldots, f_{i-1}\rangle \text{ and } \deg(gf_i) < d \implies g \in \langle f_1, \ldots, f_i\rangle.$$

for all $1 \le i \le m$. We call a sequence of homogeneous polynomials $f_1, \ldots, f_m$ a Boolean semi-regular system if it is Boolean d_{reg}-regular.

It is conjectured that random systems are Boolean semi-regular for most m, n. For such systems, [5] shows that the syzygies are exactly given by the relation matrix for $d < d_{\mathrm{reg}}$ and that the following alternating sum coming from the Hilbert series holds:

$$\operatorname{null}\mathrm{Mac}^{(2),d} = \sum_{c=0}^{m}(-1)^c \left(\!\!\binom{m}{c}\!\!\right)\binom{n}{d-2c}.$$

The attack against UOV Note that, again, a solution to the system implies that it cannot be Boolean semi-regular. Still, we expect UOV systems to behave semi-regularly up to the degree when a solution can be efficiently found. Let $d(m, n, o)$ denote lowest $d > 0$ such that the RHS of the above expression becomes less then $\binom{o}{d}$.

One can choose to place $k \leq o$ linear restrictions, for example by fixing the first k variables to be 0. This gives us the MQ map $\pi_{n-k}\mathcal{F}$ in $n-k$ variables, with an oil space of dimension at least $o-k$. If $\pi_{n-k}\mathcal{F}$ satisfies the same semi-regularity assumption we made for $\mathcal{F}$, and $d_k := d(m, n-k, o-k) \leq o-k$, then we can apply the approach discussed above to $\pi_{n-k}\mathcal{F}$. In [29], for all cases where 2-truncated Macaulay matrices are used, i.e., $p^\ell = 2$, the optimal choice for this parameter k was found to be such that d_k is precisely $o-k$.

If $n < 3o$, then one can construct an MQ map with n variables, $3m-2$ linear restrictions and an oilspace of size $3o-n$. Elements from the oilspace of this derived MQ-map can be used to find an element from the original oilspace. This approach is known as the intersection attack [22], while applying algebraic methods directly to the MQ-map is known as the reconciliation attack [10]. The approach discussed in this section can be applied to both attacks.

6.2 Smoothing the number of fixed variables

Our goal in this section is to improve the 2-truncated attack by using submatrices of $\mathrm{Mac}^{(2),d}$, much like we did for XL for the generic MQ-problem. We start by adapting some earlier definitions to this setting.

For $\mu \in \mathcal{M}_d^{(2)}$, let $\mathcal{S}_{\leq\mu}^{(2)}$ denote the initial segment of $\mathcal{M}_d^{(2)}$ with largest element μ. From here on, for $c \geq 0$, we denote by $\partial^c : \mathcal{PM}^{(2)} \to \mathcal{PM}^{(2)}$ the mapping defined by $\partial^c(S) := (S \cdot \mathcal{M}_c) \cap \mathcal{M}^{(2)}$. Furthermore, if $\mu \in \mathcal{M}^{(2)}$, we define $\partial^c(\mu) := \max \partial^c(\{\mu\})$ so that $\partial^c(\mathcal{S}_{\leq\mu}^{(2)}) = \mathcal{S}_{\leq\partial^c(\mu)}^{(2)}$.

For $\mu \in \mathcal{M}_{d-2}^{(2)}$, we denote by $\mathrm{Mac}^{(2),\mu}$ the submatrix of $\mathrm{Mac}^{(2),d}$ consisting of those entries indexed by $((\mu, i), \nu)$ such that $\mu \in \mathcal{S}_{\leq\mu}^{(2)}$ and $\nu \in \partial^2(\mathcal{S}_{\leq\mu}^{(2)})$. The 2-truncated Macaulay matrix can then be partitioned as follows:

$$\mathrm{Mac}^{(2),d} = \left[\begin{array}{c|c} * & * \\ \hline \mathbf{0} & \mathrm{Mac}^{(2),\mu} \end{array}\right].$$

From this partition, we see that $\mathrm{lker}(\mathrm{Mac}^{(2),\mu})$ can be identified with a subspace of $\mathrm{lker}\,\mathrm{Mac}^{(2),d}$. We define the space of trivial syzygies on $\mathrm{Mac}^{(2),\mu}$ as $\mathrm{KSyz}(\mathrm{Mac}^{(2),\mu}) := \mathrm{lker}(\mathrm{Mac}^{(2),\mu}) \cap \mathrm{KSyz}^{(2),d}$.

Recall that, in earlier sections, the matrix Mac^d had the Macaulay matrices of all smaller degrees as submatrices. Specifically, for $k \leq d-2$, the submatrix corresponding to the initial segment $X_n^k \cdot \mathcal{M}_{d-k-2} \subseteq \mathcal{M}_{d-2}$ is precisely Mac^{d-k}. Thus, when considering submatrices Mac^μ of Macaulay matrices Mac^d, the precise value of d was mostly unimportant.

Here, that choice of degree is more significant. First of all, the degree may not be larger then the dimension of the oil-space. Additionally, for $k > 0$, it does not generally hold that $\mathrm{Mac}^{(2),d-k}$ is a submatrix of $\mathrm{Mac}^{(2),d}$. Instead, consider the monomial $\mu_k := \min(\mathcal{M}_k^{(2)}) \cdot \max(\mathcal{M}_{d-2-k}^{(2)}) = X_n \cdots X_{n-k+1} \cdot X_{d-2-k} \cdots X_1$, which gives the initial segment $\mathcal{S}_{\leq \mu_k}^{(2)} = X_n \cdots X_{n-k+1} \cdot \mathcal{M}_{d-2-k}^{(2)}(X_1, \ldots, X_{n-k})$. The corresponding submatrix of $\mathrm{Mac}^{(2),d}$ is

$$\mathrm{Mac}^{(2),\mu_k}(\mathcal{F}) = \mathrm{Mac}^{(2),d-k}(\pi_{n-k}\mathcal{F}).$$

We see that choosing μ to be sufficiently small effectively removes variables from $\mathrm{Mac}^{(2),\mu}$, replicating the role that the parameter k plays in the framework from [29]. Note, in particular, that for any k such that $d_k = o - k$, the Macaulay matrix $\mathrm{Mac}^{(2),d_k}(\pi_{n-k}\mathcal{F})$ is a submatrix of $\mathrm{Mac}^{(2),o}(\mathcal{F})$. Such k were consistently found to be optimal for the 2-truncated attack in [29]. This suggests it is optimal to use submatrices of $\mathrm{Mac}^{(2),o}$ specifically, and not of any lower degree matrix.

We described our original approach as allowing for intermediate options of the degree. The approach here can instead be said to allow for intermediate options of the number k of fixed variables, as choosing $\mu \in \mathcal{M}_o^{(2)}$ such that $\mu_{k+1} < \mu \leq \mu_k$ gives a submatrix of $\mathrm{Mac}^{(2),o-k}(\pi_{n-k}\mathcal{F})$ that itself has $\mathrm{Mac}^{(2),o-k-1}(\pi_{n-k-1}\mathcal{F})$ as a submatrix.

In the upcoming analysis, it is generally more convenient to assume that $\deg(\mu) = d_{\mathrm{reg}} - 2$. Thus, in the actual analysis, we use $\pi_{n-k}\mathcal{F}$ in place of $\mathcal{F}$.

6.3 Nullities of truncated Macaulay matrices

We define the Boolean monomial of regularity in a similar fashion as before. Given a system $\mathcal{F} = (f_1, \ldots, f_m)$ of homogeneous polynomials in $\mathcal{R}^{(2)}$, we associate vector spaces to each monomial $\mu \in \mathcal{M}_d^{(2)}$ $1 \leq \ell \leq m$ given by

$$J_{\leq \mu,\ell} := \langle \nu \cdot f_i \mid 1 \leq i \leq \ell,\ \nu \in \mathcal{M}_{d-\deg f_i},\ \mathrm{LM}(\nu \cdot f_i) \leq \mu \rangle_{\mathbb{F}_q}.$$

$$J_{\leq \mu,\ell}^{(2)} := \langle \nu \cdot f_i \mid 1 \leq i \leq \ell,\ \nu \in \mathcal{M}_{d-\deg f_i}^{(2)},\ \mathrm{LM}(\nu \cdot f_i) \leq \mu \rangle_{\mathbb{F}_q}.$$

Definition 13. *Given a sequence of polynomials $f_1, \ldots, f_m \in \mathcal{R}^{(2)}$ with degree of regularity d_{reg}, we define the Boolean monomial of regularity μ_{reg} for the sequence as*

$$\mu_{reg} := \min \left\{ \mu \in \mathcal{M}_{d_{reg}}^{(2)} \;\middle|\; J_{\leq \mu,m}^{(2)} = \mathcal{R}_{\mathcal{S}_{\leq \mu}}^{(2)} \right\} \cup \{\infty\}.$$

The next definition is the only one that really changes as it should incorporate the fact that $f_i^2 = 0$.

Definition 14. *A sequence $f_1, \ldots, f_m$ of homogeneous polynomials in $\mathcal{R}$ is called Boolean* μ-regular *if*

$$g f_i \in J_{\leq \mu,i-1}^{(2)} \text{ and } \mathrm{LM}(g f_i) < \mu \implies g \in \langle f_1, \ldots, f_i \rangle.$$

for all $1 \leq i \leq m$. *Furthermore, we call the sequence* monomial semi-regular *if it is* μ_{reg}*-regular.*

Much like in section 4.1, we conjecture that generic sequences in $\mathcal{R}^{(2)}$ are monomial semi-regular, which only requires the set to be non-empty. We tested random homogeneous quadratic systems for the same parameters (m, n) mentioned in section 4.1, and again found monomial semi-regular sequences for each tested pair.

Remark 3. In section 4.1 we showed that $J_{\leq \mu \cdot X_1^2, m} = \operatorname{rowspace} \operatorname{Mac}^{\mu}$ for systems with $\operatorname{LM}(f_i) = X_1^2$ for all i. Here we cannot do the same, as the equality $\mu \cdot \operatorname{LM}(f) = \operatorname{LM}(\mu \cdot f)$ does not hold in the 2-truncated algebra. However, it still holds that $\operatorname{rowspace} \operatorname{Mac}^{(2),\mu} \subseteq J^{(2)}_{\leq X_1 X_2 \mu, m}$ with a small gap. For practical reasons we will therefore work with $\operatorname{rowspace} \operatorname{Mac}^{(2),\mu}$.

Next, we aim to give upper bounds for the nullity of $\operatorname{Mac}^{(2),\mu}$, similar to sections 4.3 and 4.4. We start with the following result which can be proven analogously to Corollary 2.

Corollary 5. *Let* $\mathcal{S} \subseteq \mathcal{M}^{(2)}_{d-2}$ *be an initial segment with largest element* $\mu = X_{a_{d-2}} \cdots X_{a_1}$, *with* $\partial^2(\mu) < \mu_{reg}$. *Then*

$$\operatorname{null} \operatorname{Mac}^{(2),\mu} \leq |\partial^2 \mathcal{S}| - m|\mathcal{S}| + \dim \operatorname{KSyz}^{(2),d} - \sum_{1 \leq k \leq d-2} \dim \operatorname{KSyz}^{(2),k+2}(\pi_{a_k - 1} \mathcal{F}).$$

Next, we give a result analogous to Corollary 3.

Corollary 6. *Let* $\mathcal{F}$ *be a sequence of homogeneous quadratic polynomials in* $\mathcal{R}$, *then*

$$\dim \operatorname{KSyz}^{(2),d} \mathcal{F} \geq \operatorname{rank} \operatorname{Mac}^{(2),d-2} \mathcal{F} + \dim \operatorname{KSyz}^{(2),d} \mathcal{F}_{(m-1)}.$$

Proof. We use the following decomposition of $\operatorname{Rel}^{(2),d}$:

$$\operatorname{Rel}^{(2),d} : \mathcal{R}^{(2)} \oplus (\mathcal{R}^{(2)})^{m-1} \to (\mathcal{R}^{(2)})^{m} \oplus (\mathcal{R}^{(2)})^{\left(\!\binom{m-1}{2}\!\right)},$$

given in degree d by

$$\operatorname{Rel}^{(2),d} = \left[\begin{array}{c|c} \operatorname{Mac}^{(2),d-2} \mathcal{F} & * \\ \hline \mathbf{0} & \operatorname{Rel}^{(2),d} \mathcal{F}_{(m-1)} \end{array}\right].$$

Here, the rowspace, is decomposed by whether or not $m \in S$, where (μ, S) is the label of the row, while the columspace is decomposed by whether or not $i = m$, where (ν, i) is the label of the column. Since $\operatorname{KSyz}^{(2),d}$ is the rowspace of $\operatorname{Rel}^{(2),d}$, the result follows. □

If we assume Boolean semi-regularity for $\mathcal{F}_{(i)}$ for all $i \in [m]$, then Corollaries 5 and 6 give us an upper bound of $\operatorname{null} \operatorname{Mac}^{(2),\mu}$ if $\partial^2(\mu) < \mu_{\text{reg}}$. For any $\mu \in \mathcal{M}^{(2)}_o$, if that upper bound is negative, we can conclude that $\mu_{\text{reg}} \leq \partial^2(\mu)$.

6.4 Complexity Estimates

In this section, we will compare the complexity of our approach to [29]. Let k be the highest integer such that $d_k = o - k$. This matches the optimal number of variables fixed in [43] for all parameter sets, and in [29] for those parameter sets where the 2-truncated approach was used. We write $d = d_k$.

In order to estimate the complexity, we determine a $\mu = X_{a_{d-2}} \cdots X_{a_1} \in \mathcal{M}^{(2)}_{d-2}(X_1, \ldots, X_{n-k})$ such that the upper bound for $\operatorname{null} \operatorname{Mac}^{(2),\mu}(\pi_{n-k}\mathcal{F})$ computed using Corollaries 5 and 6 is negative. We choose $a_{d-2}, \cdots, a_1$ in sequence. After determining $a_d, \ldots, a_{\ell+1}$, we choose $a_\ell < a_{\ell+1}$ as the largest index such that, for $\mu_\ell = X_{a_{d-2}} \cdots X_{a_\ell} \cdot X_{d-2-\ell} \cdots X_1$, the computed upper bound for $\operatorname{null} \operatorname{Mac}^{(2),\mu_\ell}(\pi_{n-k}\mathcal{F})$ is negative.

The estimated number of $\mathbb{F}_q$ operations needed to find a kernel element of $\operatorname{Mac}^{(2),\mu}(\pi_{n-k}\mathcal{F})$ using Wiedemann's algorithm [48] is

$$3\binom{n}{2}\left|\partial^2(\mathcal{S}^{(2)}_{\leq\mu})\right|^2.$$

In line with [13,11], we multiply this by $2(\log_2 q)^2 + \log_2 q$ to estimate the number of required gates.

We computed the estimates for both the reconciliation attack and the intersection attack for the parameter sets of UOV from [13], as well as the reconciliation attack for $\texttt{MAYO}_2$, the only parameter set from [12] for which [29] uses 2-truncated rings. In table 4, we compare these estimates.

For all these parameter sets, we found that our choice of $\mu = X_{a_{d-2}} \cdots X_{a_1}$ satisfies $a_{d-2} > n - k$. This means that the submatrix $\operatorname{Mac}^{(2),\mu}(\pi_{n-k}\mathcal{F})$ of $\operatorname{Mac}^{(2),o-k}(\pi_{n-k}\mathcal{F})$ itself contains $\operatorname{Mac}^{(2),o-k-1}(\pi_{n-k-1}\mathcal{F})$ as a submatrix. The improvement compared to [43] and [29] largely depends on how close μ is to $X_{d-2} \cdots X_1$. We see this reflected in the table. For some parameters, we find several bits of improvement, such as 3.6 bits for the intersection attack for the `UOV-Is` parameters. For others, we find only marginal improvement.

					Reconciliation attack			Intersection attack		
	n	m	o	q	k	[43]	This paper	k	[29]	This paper
`UOV-Ip`	112	44	44	256	26	141.4	139.4	7	128.4	126.2
`UOV-Is`	160	64	64	16	42	176.5	176.4	16	158.5	154.9
`UOV-III`	184	72	72	256	47	207.5	206.4	14	181.5	181.3
`UOV-V`	244	96	96	256	64	258.4	257.9	22	223.3	223.1
$\texttt{MAYO}_2$	118	108	10	16	4	113.1	109.7	-	-	-

Table 4: Complexity comparison between [43,29] and our approach for parameters from [13,12].

7 Future work

Our paper leaves a number of open questions. Tightening the upper bounds for $\operatorname{null} \mathrm{Mac}^{\mu}$ would give a better estimate of μ_{reg}.

We have mainly focused on generalizing XL to allow for parameters in-between integral degrees. It may be possible to similarly generalize other variants of XL such as F4 [23], F5 [24] and Crossbred [34]. In F4/F5, one could design a pair-selection strategy that only considers S-polynomials of a certain support. This then complements the F5 criterion which is mainly focused on removing rows instead of columns. For Crossbred one could apply our optimization both in the kernel finding in the first part and in solving the smaller Macaulay matrix in the second part. Techniques used in this paper might also be applied to systems of higher degrees.

As covered in Section 3.1, the choice to limit our scope to initial segments is based on viewing the height of Mac^S as a proxy for its rank. If one could characterize sets $S \subseteq \mathcal{M}_{d-2}$ that optimize the width of Mac^S with respect to its rank or nullity, under some semi-regularity assumption, this may give further improvement.

Finally, the definition of $J_{\leq\mu}$, used to define the monomial of regularity, depends on the choice of generators $\mathcal{F}$. Instead, one could define $J_{\leq\mu}(I_2)$ or even $J_{\leq\mu}(I_{d'})$, which is an often strict superset of $J_{\leq\mu}(\mathcal{F})$ still contained in $\mathcal{R}_{\mathcal{S}_{\leq\mu}}$. This gives something akin to degree falls. Such methods would require more extensive analysis.

References

1. Aragon, N., Gaborit, P., Hauteville, A., Ruatta, O., Zémor, G.: Low rank parity check codes: New decoding algorithms and applications to cryptography. IEEE Trans. Inf. Theory **65**(12), 7697–7717 (2019). https://doi.org/10.1109/TIT.2019.2933535, https://doi.org/10.1109/TIT.2019.2933535
2. Arora, S., Ge, R.: New algorithms for learning in presence of errors. In: ICALP (1). pp. 403–415 (2011)
3. B., B.: Ein algorithmus zum auffinden der basiselemente des restklassenringes nach einem nulldimensionalen polynomideal. Ph. D. Thesis, Math. Inst., University of Innsbruck (1965), https://cir.nii.ac.jp/crid/1570291224270729344
4. Bardet, M., Faugère, J., Salvy, B.: On the complexity of the F5 Gröbner basis algorithm. Journal of Symbolic Computation **70**, 49–70 (2015), https://hal.inria.fr/hal-01064519/document
5. Bardet, M., Faugère, J.C., Salvy, B., Yang, B.Y.: Asymptotic behaviour of the degree of regularity of semi-regular polynomial systems. https://api.semanticscholar.org/CorpusID:10699714
6. Benadjila, R., Feneuil, T., Rivain, M.: MQ on my Mind: Post-Quantum Signatures from the Non-Structured Multivariate Quadratic Problem . In: 2024 IEEE 9th European Symposium on Security and Privacy (EuroS&P). pp. 468–485. IEEE Computer Society, Los Alamitos, CA, USA (Jul 2024). https://doi.org/10.1109/EuroSP60621.2024.00032, https://doi.ieeecomputersociety.org/10.1109/EuroSP60621.2024.00032

7. Bettale, L., Faugère, J.C., Perret, L.: Hybrid approach for solving multivariate systems over finite fields. Journal of Mathematical Cryptology **volume 3**(issue 3), 177–197 (2009), http://www-polsys.lip6.fr/~jcf/Papers/JMC2.pdf
8. Bettale, L., Faugère, J., Perret, L.: Solving polynomial systems over finite fields: improved analysis of the hybrid approach. In: van der Hoeven, J., van Hoeij, M. (eds.) Proceedings of the 37th International Symposium on Symbolic and Algebraic Computation – ISSAC '12. pp. 67–74. ACM (2012), https://hal.inria.fr/hal-00776070/document
9. Beullens, W.: Sigma Protocols for MQ, PKP and SIS, and Fishy Signature Schemes. In: EUROCRYPT (3). Lecture Notes in Computer Science, vol. 12107, pp. 183–211. Springer (2020)
10. Beullens, W.: Improved cryptanalysis of uov and rainbow. In: Canteaut, A., Standaert, F.X. (eds.) Advances in Cryptology – EUROCRYPT 2021. pp. 348–373. Springer International Publishing, Cham (2021)
11. Beullens, W.: Mayo: Practical post-quantum signatures from oil-and-vinegar maps. In: AlTawy, R., Hülsing, A. (eds.) Selected Areas in Cryptography. pp. 355–376. Springer International Publishing, Cham (2022)
12. Beullens, W., Campos, F., Celi, S., Hess, B., Kannwischer, M.J.: MAYO. Tech. rep., National Institute of Standards and Technology (2024), available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures
13. Beullens, W., Chen, M., Ding, J., Gong, B., Kannwischer, M.J., Patarin, J., Peng, B., Schmidt, D., Shih, C., Tao, C., Yang, B.: UOV — Unbalanced Oil and Vinegar. Tech. rep., National Institute of Standards and Technology (2024), available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures
14. Beullens, W., Chen, M.S., Hung, S.H., Kannwischer, M.J., Peng, B.Y., Shih, C.J., Yang, B.Y.: Oil and Vinegar: Modern Parameters and Implementations. Cryptology ePrint Archive, Report 2023/059 (Jun 2023), https://ia.cr/2023/059, accepted for publication at TCHES'23
15. Buchmann, J.A., Ding, J., Mohamed, M.S.E., Mohamed, W.S.A.E.: Mutantxl: Solving multivariate polynomial equations for cryptanalysis. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2009)
16. Caminata, A., Gorla, E.: Solving degree, last fall degree, and related invariants. Journal of Symbolic Computation **114**, 322–335 (2023). https://doi.org/https://doi.org/10.1016/j.jsc.2022.05.001, https://www.sciencedirect.com/science/article/pii/S0747717122000487
17. Chen, M.S., Hülsing, A., Rijneveld, J., Samardjiska, S., Schwabe, P.: From 5-pass $\mathcal{MQ}$-based identification to $\mathcal{MQ}$-based signatures. In: Cheon, J.H., Takagi, T. (eds.) Advances in Cryptology – ASIACRYPT 2016. vol. 10032, pp. 135–165 (2016), http://eprint.iacr.org/2016/708
18. Clements, G., Lindström, B.: A generalization of a combinatorial theorem of macaulay. Journal of Combinatorial Theory **7**(3), 230–238 (1969). https://doi.org/https://doi.org/10.1016/S0021-9800(69)80016-5, https://www.sciencedirect.com/science/article/pii/S0021980069800165
19. Courtois, N., Klimov, A., Patarin, J., Shamir, A.: Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In: Preneel, B. (ed.) Advances in Cryptology — EUROCRYPT 2000. pp. 392–407. Springer Berlin Heidelberg, Berlin, Heidelberg (2000)
20. Courtois, N.T., Patarin, J.: About the xl algorithm over gf(2). In: Joye, M. (ed.) Topics in Cryptology — CT-RSA 2003. pp. 141–157. Springer Berlin Heidelberg, Berlin, Heidelberg (2003)

21. Ding, J., Schmidt, D.: Rainbow, a New Multivariable Polynomial Signature Scheme. In: Ioannidis, J., Keromytis, A.D., Yung, M. (eds.) ACNS. Lecture Notes in Computer Science, vol. 3531, pp. 164–175 (2005)
22. Ding, J., Yang, B.Y., Chen, C.H.O., Chen, M.S., Cheng, C.M.: New differential-algebraic attacks and reparametrization of rainbow. In: Bellovin, S.M., Gennaro, R., Keromytis, A., Yung, M. (eds.) Applied Cryptography and Network Security. pp. 242–257. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
23. Faugère, J.C.: A new efficient algorithm for computing Gröbner bases (F4). Journal of Pure and Applied Algebra **139**, 61–88 (1999), http://www-polsys.lip6.fr/~jcf/Papers/F99a.pdf
24. Faugère, J.C.: A new efficient algorithm for computing Gröbner bases without reduction to zero (F5). In: Proceedings of the 2002 International Symposium on Symbolic and Algebraic Computation ISSAC. pp. 75–83. ACM Press (2002)
25. Faugère, J.C., Din, M.S.E., Spaenlehauer, P.J.: Gröbner bases of bihomogeneous ideals generated by polynomials of bidegree (1, 1): Algorithms and complexity. J. Symb. Comput. **46**(4), 406–437 (2011)
26. Faugère, J., Gauthier-Umaña, V., Otmani, A., Perret, L., Tillich, J.: A Distinguisher for High-Rate McEliece Cryptosystems. IEEE Trans. Inf. Theory **59**(10), 6830–6844 (2013). https://doi.org/10.1109/TIT.2013.2272036, https://doi.org/10.1109/TIT.2013.2272036
27. Faugère, J.C., dit Vehel, F.L., Perret, L.: Cryptanalysis of MinRank. In: Wagner, D. (ed.) CRYPTO. Lecture Notes in Computer Science, vol. 5157, pp. 280–296. Springer (2008)
28. Fröberg, R.: An inequality for hilbert series of graded algebras. MATHEMATICA SCANDINAVICA **56**, 117–144 (Dec 1985). https://doi.org/10.7146/math.scand.a-12092, https://www.mscand.dk/article/view/12092
29. Furue, H., Ikematsu, Y.: Key recovery attacks on UOV using p^l-truncated polynomial rings. Cryptology ePrint Archive, Paper 2026/298 (2026), https://eprint.iacr.org/2026/298
30. Furue, H., Ikematsu, Y., Kiyomura, Y., Takagi, T.: A new variant of unbalanced oil and vinegar using quotient ring: Qr-uov. In: Tibouchi, M., Wang, H. (eds.) Advances in Cryptology – ASIACRYPT 2021. pp. 187–217. Springer International Publishing, Cham (2021)
31. Gaborit, P., Ruatta, O., Schrek, J.: On the complexity of the rank syndrome decoding problem. IEEE Trans. Inf. Theory **62**(2), 1006–1019 (2016). https://doi.org/10.1109/TIT.2015.2511786, https://doi.org/10.1109/TIT.2015.2511786
32. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman (1979)
33. Jin, Y., Pan, Y., He, X., Gong, B., Ding, J.: Security analysis on UOV families with odd characteristics: Using symmetric algebra. Cryptology ePrint Archive, Paper 2025/1137 (2025), https://eprint.iacr.org/2025/1137
34. Joux, A., Vitse, V.: A Crossbred Algorithm for Solving Boolean Polynomial Systems. In: Kaczorowski, J., Pieprzyk, J., Pomykała, J. (eds.) Number-Theoretic Methods in Cryptology. pp. 3–21. Springer International Publishing, Cham (2018)
35. Kipnis, A., Patarin, J., Goubin, L.: Unbalanced Oil and Vinegar Signature Schemes. In: Advances in Cryptology – EUROCRYPT '99. Lecture Notes in Computer Science, vol. 1592, pp. 206–222. Springer (1999)
36. Kipnis, A., Shamir, A.: Cryptanalysis of the HFE Public Key Cryptosystem by Relinearization. In: CRYPTO '99. LNCS, vol. 1666, pp. 19–30. Springer (1999)

37. KoichiSakumoto, Shirai, T., Hiwatari, H.: Public-key identification schemes based on multivariate quadratic polynomials. In: Rogaway, P. (ed.) Advances in Cryptology – CRYPTO 2011. vol. 6841, pp. 706–723 (2011), https://www.iacr.org/archive/crypto2011/68410703/68410703.pdf
38. Kudo, M., Yokoyama, K.: The solving degrees for computing gröbner bases of affine semi-regular polynomial sequences. Cryptology ePrint Archive, Paper 2024/528 (2024), https://eprint.iacr.org/2024/528
39. Lazard, D.: Gröbner bases, gaussian elimination and resolution of systems of algebraic equations. In: European Conference on Computer Algebra. pp. 146–156. Springer (1983)
40. Macaulay, F.: The Algebraic Theory of Modular Systems. Cambridge tracts in mathematics and mathematical physics, no. 19, Cambridge University Press (1916)
41. Mohamed, M.S.E., Mohamed, W.S.A.E., Ding, J., Buchmann, J.: MXL2: Solving Polynomial Equations over GF(2) Using an Improved Mutant Strategy. In: Buchmann, J., Ding, J. (eds.) PQCrypto. Lecture Notes in Computer Science, vol. 5299, pp. 203–215. Springer (2008)
42. Perlner, R., Smith-Tone, D.: Rainbow Band Separation is Better than we Thought. Cryptology ePrint Archive, Report 2020/702 (2020), https://eprint.iacr.org/2020/702
43. Ran, L.: Wedges, Oil, and Vinegar. In: Daemen, J., Thomé, E. (eds.) Advances in Cryptology – EUROCRYPT 2026. pp. 363–388. Springer Nature Switzerland, Cham (2026)
44. Ran, L., Samardjiska, S.: Rare structures in tensor graphs. In: Chung, K.M., Sasaki, Y. (eds.) Advances in Cryptology – ASIACRYPT 2024. pp. 66–96. Springer Nature Singapore, Singapore (2025)
45. The Sage Developers: SageMath, the Sage Mathematics Software System (Version 10.6) (2025), `https://www.sagemath.org`
46. Wang, L.C., Tseng, P.E., Kuan, Y.L., Chou, C.Y.: A simple noncommutative UOV scheme. Cryptology ePrint Archive, Paper 2022/1742 (2022), https://eprint.iacr.org/2022/1742
47. Wiedemann, D.: Solving sparse linear equations over finite fields. IEEE Transactions on Information Theory **32**(1), 54–62 (1986). https://doi.org/10.1109/TIT.1986.1057137
48. Wiedemann, D.: Solving sparse linear equations over finite fields. IEEE transactions on information theory **32**(1), 54–62 (1986)
49. Yang, B.Y., Chen, J.M.: All in the XL Family: Theory and Practice. In: Park, C., Chee, S. (eds.) ICISC '04. Lecture Notes in Computer Science, vol. 3506, pp. 67–86. Springer (2004)
50. Yang, B., Chen, J.: Theoretical analysis of XL over small fields. In: Information Security and Privacy. pp. 277–288 (2004). https://doi.org/10.1007/978-3-540-27800-9_24, http://dx.doi.org/10.1007/978-3-540-27800-9_24, http://www.iis.sinica.edu.tw/papers/byyang/2386-F.pdf