

***Quasi-triangularization of matrix polynomials
over arbitrary fields***

Anguas, Luis M. and Dopico, Froilán M. and Hollister,
Richard and Mackey, D. Steven

2021

MIMS EPrint: **2021.18**

Manchester Institute for Mathematical Sciences
School of Mathematics

The University of Manchester

Reports available from: <http://eprints.maths.manchester.ac.uk/>

And by contacting: The MIMS Secretary
School of Mathematics
The University of Manchester
Manchester, M13 9PL, UK

ISSN 1749-9097

Quasi-Triangularization of Matrix Polynomials over Arbitrary Fields *

L.M. Anguas [†]

F. M. Dopico [‡]

R. Hollister [§]

D.S. Mackey [¶]

December 5, 2021

Abstract

In [19], Taslaman, Tisseur, and Zaballa show that any regular matrix polynomial $P(\lambda)$ over an algebraically closed field is spectrally equivalent to a triangular matrix polynomial of the same degree. When $P(\lambda)$ is real and regular, they also show that there is a real quasi-triangular matrix polynomial of the same degree that is spectrally equivalent to $P(\lambda)$, in which the diagonal blocks are of size at most 2×2 . This paper generalizes these results to regular matrix polynomials $P(\lambda)$ over *arbitrary fields* \mathbb{F} , showing that any such $P(\lambda)$ can be quasi-triangularized to a spectrally equivalent matrix polynomial over \mathbb{F} of the same degree, in which the largest diagonal block size is bounded by the highest degree appearing among all of the \mathbb{F} -irreducible factors in the Smith form for $P(\lambda)$.

1 Introduction

Triangularizations of matrix polynomials that preserve degree as well as the finite and infinite spectral structure via unimodular equivalence are essentially Schur-like forms for matrices whose entries are polynomials. These have been achieved over algebraically closed fields for regular quadratic matrix polynomials in [20], and for regular matrix polynomials of arbitrary degree in [19]. There are also some results in [19] on singular matrix polynomials, but these will not be addressed in this paper. Also in [19, 20], when the underlying field is \mathbb{R} , the authors show how to produce quasi-triangularizations with diagonal blocks of size at most 2×2 .

The goal of this paper is similar: given a regular matrix polynomial $P(\lambda)$, to show how to construct a regular quasi-triangular matrix polynomial with the same finite and infinite spectral structure as $P(\lambda)$, i.e., is spectrally equivalent to $P(\lambda)$, and has the same degree as $P(\lambda)$. However, this work is an extension of [19] in that the results presented here are for matrix polynomials over an *arbitrary field*. In order to achieve this generalization, though, the possibility of diagonal blocks of sizes even larger than 2×2 must be allowed. We show that a quasi-triangularization can always be constructed in which the sizes of the diagonal blocks do not exceed $k \times k$, where k is the highest degree among all of the irreducible factors of the invariant polynomials in the Smith form of the given polynomial matrix $P(\lambda)$. Note that the term *quasi-triangular* is used throughout this paper to refer to square matrices that are block upper (or lower) triangular with square blocks along the main diagonal, at least one of which has size 2×2 or larger. A matrix is *k-quasi-triangular* if the diagonal blocks are no larger than $k \times k$.

Here is a brief overview of the paper. After some preliminary discussion of concepts, notation, and terminology in Section 2, we begin in Section 3 by solving the quasi-triangular *realization problem* for finite spectral data over an arbitrary field \mathbb{F} . That is, we take as starting point a collection of finite spectral data rather than a matrix polynomial, and show how to construct a strictly regular *k*-quasi-triangular matrix polynomial over \mathbb{F} having exactly the given spectral data; here k is the largest degree among the irreducible divisors of the given data. The solution of this inverse problem is the main technical result of the paper; all other results depend on and follow from this. The central idea of the proof is to take the given spectral data, form the corresponding Smith form,

*This publication is part of the “Proyecto de I+D+i PID2019-106362GB-I00 financiado por MCIN/AEI/10.13039/501100011033”. It has been also funded by “Ministerio de Economía, Industria y Competitividad (MINECO)” of Spain through grants MTM-2015-65798-P and BES-2013-065688.

[†]Departamento de Matemática e Informática Aplicadas a las Ingenierías Civil y Naval, Escuela Técnica Superior de Ingenieros de Caminos, Canales y Puertos, Universidad Politécnica de Madrid, Profesor Aranguren, 3, 28040 Madrid, Spain (l.anguas@upm.es)

[‡]Departamento de Matemáticas, Universidad Carlos III de Madrid, Avda. Universidad 30, 28911 Leganés, Spain (dopico@math.uc3m.es)

[§]Department of Mathematics, University at Buffalo, Buffalo NY 14260 USA (rahollis@buffalo.edu)

[¶]Department of Mathematics, Western Michigan University, Kalamazoo MI 49008-5248 USA (steve.mackey@wmich.edu)

and then systematically “un-diagonalize” in a way that moves the matrix polynomial toward the desired degree, while maintaining quasi-triangularity. This is a proof technique used in [8], and then developed by [19]; some antecedents of this technique can also be found in [13]. We develop it further here, adapting it to the arbitrary field setting. Indeed, a number of nontrivial ingredients go into proving this quasi-triangular realization result by this method — majorization plays a role, as well as a new combinatorial lemma on the partitioning of integer multisets. In Section 4 we extend this realization result to include elementary divisors at ∞ in the given spectral data. In order to achieve this extension, we use the well-known tool of Möbius transformations [12], although once again we will need to do a significant amount of work to adapt them to work smoothly for matrix polynomials over arbitrary fields. At this point the signature result of the paper — the quasi-triangularization of any regular matrix polynomial over an arbitrary field in Theorem 4.9 — now follows easily. Finally, in Section 5 we investigate conditions for describing when exact triangularization is possible in the arbitrary field setting; we also display several families of examples, some that illustrate the sharpness of k as a general upper bound on the size of the diagonal blocks in quasi-triangularizations, and others that show that this upper bound k can sometimes be a huge overestimate of the diagonal block size that is actually attainable.

2 Preliminaries

In this paper, we deal with *polynomial matrices* (also referred to as *matrix polynomials*) over an arbitrary field \mathbb{F} , i.e., matrices whose entries are polynomials with coefficients from \mathbb{F} . In particular, we will be working only with *regular* matrix polynomials, that is, square matrix polynomials with determinant different from the zero polynomial. A polynomial matrix is *unimodular* if it is regular and has a (nonzero) constant determinant. Throughout the paper the set of natural numbers are denoted by \mathbb{N} , and includes zero; \mathbb{N}^+ is then the set of positive integers.

The Smith form is a canonical representation of matrix polynomials under *unimodular equivalence*, i.e., obtained by left and right multiplication by unimodular matrix polynomials. This form was first defined for integer matrices [18]. We will use the extension given in [6] for matrix polynomials:

Theorem 2.1 (Smith form).

Let $P(\lambda)$ be an $m \times n$ matrix polynomial over an arbitrary field \mathbb{F} . Then there exists $r \in \mathbb{N}$, and unimodular matrix polynomials $U(\lambda)$ and $V(\lambda)$ such that

$$U(\lambda)P(\lambda)V(\lambda) = \text{diag}(s_1(\lambda), \dots, s_{\min\{m,n\}}(\lambda)) =: S(\lambda)$$

where $s_i(\lambda) \in \mathbb{F}[\lambda]$, for $i = 1, \dots, \min\{m, n\}$, $s_1(\lambda), \dots, s_r(\lambda)$ are monic, $s_{r+1}(\lambda), \dots, s_{\min\{m,n\}}(\lambda)$ are identically-zero, and $s_i(\lambda)$ is a divisor of $s_{i+1}(\lambda)$ for $i = 1, \dots, r-1$. Moreover, the number r is equal to the rank of P , and the diagonal entries of the $m \times n$ matrix polynomial $S(\lambda)$ are uniquely determined by the multiplicative relations

$$s_1(\lambda)s_2(\lambda) \cdots s_j(\lambda) = \gcd\{\text{all } j \times j \text{ minors of } P(\lambda)\}, \text{ for } j = 1, \dots, r. \quad (2.1)$$

When $P(\lambda)$ is regular, then $r = m = n$, and $S(\lambda)$ is a nonsingular diagonal matrix.

The diagonal matrix $S(\lambda)$ featuring in this theorem is called the *Smith form* of $P(\lambda)$. The nonzero diagonal entries of $S(\lambda)$ are called the *invariant polynomials* of $P(\lambda)$, and their zeros are the *finite eigenvalues* of $P(\lambda)$. An invariant polynomial will be called *trivial* if it is identically equal to 1 and *nontrivial* otherwise.

A non-constant irreducible polynomial $\chi(\lambda) \in \mathbb{F}[\lambda]$ that divides some invariant polynomial of $P(\lambda)$ will be called an *irreducible divisor* of $P(\lambda)$. This new concept and terminology is adopted in this work because of the central role that will be played here by these objects. Given an invariant polynomial $s_i(\lambda)$ and an irreducible divisor $\chi(\lambda)$, or indeed any irreducible polynomial $\chi(\lambda)$, there is a unique natural number α_i (perhaps zero) such that

$$s_i(\lambda) = \chi(\lambda)^{\alpha_i} \widehat{s}_i(\lambda),$$

with $\widehat{s}_i(\lambda)$ not divisible by $\chi(\lambda)$. Any factor $\chi(\lambda)^{\alpha_i}$ with $\alpha_i > 0$ is traditionally called an *elementary divisor* [7] of $P(\lambda)$. The number α_i , whether it is zero or nonzero, is called the *partial multiplicity* of the irreducible $\chi(\lambda)$ with respect to the invariant polynomial $s_i(\lambda)$, while the sequence $(\alpha_1, \alpha_2, \dots, \alpha_r)$, with $0 \leq \alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_r$, is called the *partial multiplicity sequence* of $P(\lambda)$ at $\chi(\lambda)$, which we will denote by

$$\mathcal{PM}(P, \chi) := (\alpha_1, \alpha_2, \dots, \alpha_r). \quad (2.2)$$

Note that $\mathcal{PM}(P, \chi)$ may consist of all zeroes, and χ is an irreducible divisor of P exactly when some α_i is nonzero. The partial multiplicity sequence of a degree one irreducible divisor $\chi(\lambda) = \lambda - \lambda_0$ is also sometimes

referred to as the partial multiplicity sequence of the eigenvalue λ_0 . An abbreviated notation, $\mathcal{PM}(\chi)$, will also be used for the partial multiplicity sequence associated to $\chi(\lambda)$. This may be used when the underlying matrix polynomial is understood, but more commonly it will be employed when there is *no* matrix polynomial in the background at all, and the sequence being specified (and its association with χ) is part of a collection of given spectral data that is yet to be realized. To emphasize this role of being input data for a realization problem, we will sometimes instead write $\mathcal{PM}_{\text{given}}(\chi)$.

In using the notation (2.2), it is very useful to *abandon* a standard convention, i.e., the convention that only *monic* irreducibles should be allowed or considered. As far as partial multiplicities are concerned, there is no important difference between an irreducible $\chi(\lambda)$ and any nonzero scalar multiple $c\chi(\lambda)$. In particular, since for any invariant polynomial $s_i(\lambda)$ we have

$$s_i(\lambda) = \chi(\lambda)^{\alpha_i} \widehat{s}_i(\lambda) = (c\chi(\lambda))^{\alpha_i} \frac{\widehat{s}_i(\lambda)}{c^{\alpha_i}} = (c\chi(\lambda))^{\alpha_i} \widetilde{s}_i(\lambda)$$

where $c\chi$ and \widetilde{s}_i are coprime, we see that the definition of the partial multiplicity sequence $(\alpha_1, \dots, \alpha_r)$ is unaffected by allowing non-monic irreducibles. Thus we may consistently say that

$$\mathcal{PM}(P, c\chi) = \mathcal{PM}(P, \chi) \quad (2.3)$$

for any nonzero scalar $c \in \mathbb{F}$, and thereby simply ignore whether an irreducible is monic or not. In line with this, we will say that two \mathbb{F} -irreducible polynomials are *distinct irreducibles* if neither is a nonzero scalar multiple of the other.

Remark 2.2. In effect we are implicitly defining an equivalence relation on the set of \mathbb{F} -irreducible polynomials. That is, two irreducibles are equivalent (with respect to partial multiplicity sequences) if one is a nonzero scalar multiple of the other. Distinct irreducibles are then just representatives of different equivalence classes under this relation.

Remark 2.3. It is important to keep in mind that there are fields \mathbb{F} that support the presence of \mathbb{F} -irreducible polynomials of arbitrarily high degree. A simple example of this is the field \mathbb{Q} of rational numbers. Using the Eisenstein criterion [2], it is easy to see that the polynomial $x^n + p$ is \mathbb{Q} -irreducible for any prime number $p \geq 2$ and any $n \geq 1$.

Definition 2.4. The *finite spectral structure* of $P(\lambda)$ refers to any of the following equivalent (and easily interconvertible) collections of data:

- (a) the set of all distinct irreducible divisors of $P(\lambda)$, each equipped with their partial multiplicity sequence,
- (b) the multiset of all elementary divisors of $P(\lambda)$, together with the number $\text{rank } P$,
- (c) the multiset of all invariant polynomials of $P(\lambda)$, including the trivial ones.

(Recall that a multiset is like a set, but with repetitions allowed, i.e., a “set with multiplicities” [10].)

Remark 2.5. Note that the rank of a matrix polynomial is encoded in each partial multiplicity sequence by its length.

Some regular matrix polynomials have structure that is not completely captured by their finite spectral structure alone. To get the full story for these matrix polynomials, it is also necessary to include their “spectral structure at ∞ ”. For this, some additional terminology is needed. The *grade* of a matrix polynomial $P(\lambda)$ is a natural number g such that

$$P(\lambda) = P_0 + P_1\lambda + \dots + P_g\lambda^g,$$

with each $P_j \in \mathbb{F}^{n \times n}$. Note, however, that unlike in the definition of the degree d of $P(\lambda)$, there is no requirement here for the leading coefficient P_g to be nonzero. Thus we see that $g \geq d$ always holds, no matter what the choice of grade might be. And we emphasize that, in contrast with degree, grade is indeed a choice, although a very common choice is for g to be taken to be equal to d . The *grade g reversal* of $P(\lambda)$ is the matrix polynomial

$$(\text{rev}_g P)(\lambda) := \lambda^g P(1/\lambda),$$

and $P(\lambda)$ has an *eigenvalue at infinity* if $(\text{rev}_g P)(\lambda)$ has an eigenvalue at zero. Moreover, the partial multiplicity sequence of the eigenvalue at infinity for $P(\lambda)$ is, by definition, identical to the partial multiplicity sequence of

the eigenvalue zero for $(\text{rev}_g P)(\lambda)$. We will use $\mathcal{PM}(P, \infty)$ as temporary notation for the partial multiplicity sequence of P at ∞ , so that this definition can be expressed as

$$\mathcal{PM}(P, \infty) := \mathcal{PM}(\text{rev}_g P, \lambda).$$

However, later in Section 4 we will have reason to change this notation to something that is more consistent with the notation in (2.2), and also works more smoothly with Möbius transformations and their properties.

Definition 2.6. The *infinite spectral structure* of $P(\lambda)$ refers to the eigenvalue at infinity (if it exists), together with its partial multiplicity sequence.

The relationship of grade to degree (e.g., whether $g = d$ or $g > d$) is reflected in the infinite spectral structure, in particular in the first partial multiplicity at ∞ . Note that some of the following result appeared previously in [3, Lemma 2.17], but without proof.

Lemma 2.7. Suppose $P(\lambda)$ is any $m \times n$ matrix polynomial over a field \mathbb{F} , with rank r , degree d , and grade g . Let $(\alpha_1, \alpha_2, \dots, \alpha_r)$ with $0 \leq \alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_r$ be the partial multiplicity sequence of P at ∞ . Then

$$\alpha_1 = g - d, \quad (2.4)$$

and hence $g = d$ if and only if $\alpha_1 = 0$. (Equivalently, $g > d$ if and only if $\alpha_1 > 0$.) Furthermore, suppose $\tilde{P}(\lambda)$ is an $m \times n$ matrix polynomial that is entry-wise identical to $P(\lambda)$, but its grade has been chosen to be equal to its degree d . Then the partial multiplicity sequences at ∞ of $P(\lambda)$ and $\tilde{P}(\lambda)$ are related by a constant “shift” of $g - d$, that is,

$$\mathcal{PM}(P, \infty) = \mathcal{PM}(\tilde{P}, \infty) + (g - d) \cdot (1, 1, \dots, 1). \quad (2.5)$$

Proof. Begin by expressing $P(\lambda) = P_g \lambda^g + \dots + P_d \lambda^d + \dots + P_1 \lambda + P_0$, so that $\tilde{P}(\lambda) = P_d \lambda^d + \dots + P_1 \lambda + P_0$, with $P_d \neq 0$ and $P_j = 0$ for the $g - d$ leading coefficients with $j = d + 1, \dots, g$. Then the reversals of P and \tilde{P} are related by

$$\text{rev}_g P = \lambda^g P(1/\lambda) = \lambda^{g-d} \lambda^d P(1/\lambda) = \lambda^{g-d} (\text{rev}_d \tilde{P}). \quad (2.6)$$

Note that the constant term of $\text{rev}_d \tilde{P}$ is the *nonzero* matrix P_d .

Next observe that for any constant $\mu_0 \in \mathbb{F}$ and any matrix polynomial $Q(\lambda)$ over \mathbb{F} , the number of initial zero partial multiplicities in the sequence $\mathcal{PM}(Q, \chi)$ with $\chi(\lambda) = \lambda - \mu_0$ is equal to $\text{rank } Q(\mu_0)$. This can be easily seen from the Smith form

$$U(\lambda)Q(\lambda)V(\lambda) = \hat{S}(\lambda) \quad (2.7)$$

by evaluating (2.7) at μ_0 , then computing ranks to get $\text{rank}(Q(\mu_0)) = \text{rank}(\hat{S}(\mu_0))$. But any diagonal entry of $\hat{S}(\lambda)$ corresponding to a nonzero partial multiplicity of $\lambda - \mu_0$ will be zeroed out in $\hat{S}(\mu_0)$, so the remaining nonzero diagonal entries of $\hat{S}(\mu_0)$, there are $\text{rank } \hat{S}(\mu_0) = \text{rank } Q(\mu_0)$ many of them, will count the zero partial multiplicities of $\lambda - \mu_0$ for Q . Applying this to $Q(\lambda) = \text{rev}_d \tilde{P}$, we see that

$$\begin{aligned} (\text{no. of zero partial multiplicities for } \tilde{P} \text{ at } \infty) &= (\text{no. of zero partial multiplicities for } \text{rev}_d \tilde{P} \text{ at } (\lambda - 0)) \\ &= \text{rank}(\text{rev}_d \tilde{P}(0)) = \text{rank } P_d > 0, \end{aligned}$$

since P_d is nonzero. Thus $\mathcal{PM}(\tilde{P}, \infty) := \mathcal{PM}(\text{rev}_d \tilde{P}, (\lambda - 0)) = (0, \dots)$. Hence in the Smith form $\tilde{S}(\lambda)$ for $\text{rev}_d \tilde{P}$, i.e., $\tilde{S}(\lambda) = \text{diag}(\tilde{s}_1(\lambda), \tilde{s}_2(\lambda), \dots)$, we see that $\tilde{s}_1(\lambda)$ must be coprime to λ . Now from (2.6) it follows that the Smith form $S(\lambda)$ for $\text{rev}_g P$ is related to $\tilde{S}(\lambda)$ by

$$S(\lambda) = \lambda^{g-d} \tilde{S}(\lambda).$$

Hence we can immediately conclude that $\alpha_1 = g - d$, and that the “shift” relation in (2.5) holds. \square

The first major result in this paper will be concerned with a special class of regular matrix polynomials, as in the following definition.

Definition 2.8. Regular matrix polynomials that have *no* infinite spectral structure at all will be referred to as *strictly regular*.

For a regular $n \times n$ matrix polynomial $P(\lambda)$ of grade g , being strictly regular is equivalent to the leading coefficient P_g being *nonsingular* (and hence necessarily also that $g = d$), or equivalently to $\deg(\det P(\lambda))$ being equal to gn .

Definition 2.9. The *complete spectral data* of a matrix polynomial $P(\lambda)$ is the combination of the finite and infinite spectral structures. Two regular matrix polynomials with the same complete spectral data are said to be *spectrally equivalent*.

The notion of spectral equivalence was introduced in [3] as a way to compare matrix polynomials, both regular and singular, even if they have different sizes or degrees. Although the definition given in [3] is quite different than Definition 2.9, it was shown in [3] that for regular polynomials being spectrally equivalent is the same as having the same complete spectral data. Note that in [19], having the same complete spectral data is termed *strongly equivalent*.

It is important to keep firmly in mind the contrast between unimodular equivalence and spectral equivalence. Unimodular equivalence preserves all finite spectral structure, but carries with it the unfortunate possibility of altering any infinite spectral structure that might be present. We will see a concrete illustration of this phenomenon in Example 4.8, but a more detailed discussion of the possible effects of unimodular transformations on infinite spectral structure can be found in [3, Sect. 4.3]. In this paper the ultimate aim is to produce quasi-triangularizations that are spectrally equivalent to the given matrix polynomial, and not just unimodularly equivalent.

The following theorem, dubbed the Index Sum Theorem in [3], appeared for the first time in [21] for the complex field, and in [16] and [17] for the real field. This theorem describes the fundamental relationship between the grade, rank, and structural data (elementary divisors together with minimal indices) of *any* matrix polynomial, and was recently extended to arbitrary fields in [3]. Here we state the result only for regular matrix polynomials over an arbitrary field; the result in its full generality can be found in [3]. For more on the history of this theorem, as well as its connection to other fundamental results, see [1].

Theorem 2.10 (Index Sum Theorem for Regular Matrix Polynomials).

Let $P(\lambda)$ be a regular $n \times n$ matrix polynomial of degree d and grade g having the the complete spectral data:

- invariant polynomials $p_j(\lambda)$ of degrees δ_j , for $j = 1, \dots, n$,
- infinite partial multiplicities $\gamma_1, \dots, \gamma_n$,

where some of the degrees or partial multiplicities can be zero. Then the index sum σ satisfies the relation

$$\sigma := \sum_{j=1}^n \delta_j + \sum_{j=1}^n \gamma_j = gn. \quad (2.8)$$

If $P(\lambda)$ is strictly regular, so that all of the γ_i 's are zero and $g = d$, then (2.8) simplifies to just the relation

$$\sigma = \sum_{j=1}^n \delta_j = dn. \quad (2.9)$$

The index sum theorem plays a key role in the following theorem, which is a very special case of a much more general result in [4]. The corollary immediately following it will be used at a key moment in the proof of our first main result, Theorem 3.1.

Theorem 2.11 (Fundamental Realization Theorem for Strictly Regular Matrix Polynomials).

Let \mathbb{F} be an arbitrary field, and $n, d \in \mathbb{N}^+$ be given positive integers. Consider a collection of n monic polynomials $p_1(\lambda), \dots, p_n(\lambda)$ with coefficients in \mathbb{F} and respective degrees $\delta_1, \dots, \delta_n$, such that $p_j(\lambda)$ divides $p_{j+1}(\lambda)$ for $j = 1, \dots, n-1$. Then there exists a strictly regular $n \times n$ matrix polynomial $P(\lambda)$ over \mathbb{F} with degree d and invariant polynomials $p_1(\lambda), \dots, p_n(\lambda)$, if and only if (2.9) holds.

Corollary 2.12. Let \mathbb{F} be an arbitrary field, and $Q(\lambda)$ a regular $n \times n$ polynomial matrix over \mathbb{F} . Suppose that $\sigma := \deg(\det Q(\lambda))$ is divisible by n , in particular that $\sigma = dn$. Then there is a strictly regular $n \times n$ matrix polynomial $P(\lambda)$ over \mathbb{F} of degree d that is unimodularly equivalent to $Q(\lambda)$.

Proof. Let $q_1(\lambda), \dots, q_n(\lambda)$ be the invariant polynomials of $Q(\lambda)$, with degrees $\delta_1, \dots, \delta_n$, respectively. Then from the Smith form of $Q(\lambda)$ we know that $\sum_{j=1}^n \delta_j = \sigma = dn$, i.e., that (2.9) holds. Thus by Theorem 2.11 there exists a strictly regular $n \times n$ matrix polynomial $P(\lambda)$ over \mathbb{F} with degree d and the same invariant polynomials as $Q(\lambda)$. Now since $P(\lambda)$ and $Q(\lambda)$ have the same Smith form, they must be unimodularly equivalent. \square

An important concept, used in [4] and [19] for handling regular matrix polynomials that have nontrivial infinite spectral structure, is that of Möbius transformations of matrix polynomials. These transformations and their properties were studied in [12], and will be important for the results in Section 4.

Definition 2.13 (Möbius transformations of matrix polynomials).

Let $P(\lambda)$ be a matrix polynomial of grade g over the field \mathbb{F} , and suppose $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \mathbb{F}^{2 \times 2}$ is nonsingular. The matrix polynomial

$$\mathbf{M}_A(P)(\lambda) := (c\lambda + d)^g P\left(\frac{a\lambda + b}{c\lambda + d}\right) \quad (2.10)$$

is the *Möbius transformation* of $P(\lambda)$ with respect to A .

Example 2.14. Note that the grade g reversal operation is an example of a Möbius transformation, specifically $\text{rev}_g P = \mathbf{M}_R(P)$ for $R = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$.

For conceptual clarity, it is important to keep in mind that for any fixed field \mathbb{F} and grade g , the formula in (2.10) actually defines a whole *family* of transformations, one for each matrix size $m \times n$. Once one has fixed the underlying field \mathbb{F} , grade g , input matrix size $m \times n$, and matrix A , then it is known [12] that the Möbius transformation \mathbf{M}_A defines a *bijection* on the set of all $m \times n$, grade g matrix polynomials over \mathbb{F} . However, even by staying within the confines of one of these bijections, the transformation $\mathbf{M}_A(P)$ does *not* in general preserve the degree of the input P ; degree may increase, decrease, or remain unchanged [12]. This lack of degree preservation happens even for the simplest size matrix polynomials, i.e., for scalar (1×1) polynomials. However, for scalar polynomials the following lemma describes one simple scenario, important for this paper, where degree preservation by Möbius transformations is guaranteed. The proof of this result will be postponed until Section 4.1, where a number of results about the properties of Möbius transformations and their effect on partial multiplicity sequences will be developed.

Lemma 2.15. Suppose \mathbb{F} is an arbitrary field, and $\chi(\lambda)$ is any \mathbb{F} -irreducible scalar polynomial with $\deg \chi \geq 2$. Let \mathbf{M}_A be the Möbius transformation associated with any $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in GL(2, \mathbb{F})$. Then with $\text{grade}(\chi)$ taken to be equal to $\deg(\chi)$, the transformation \mathbf{M}_A preserves both the degree and the \mathbb{F} -irreducibility of χ . That is, $\mathbf{M}_A(\chi)$ is \mathbb{F} -irreducible, and $\deg(\mathbf{M}_A(\chi)) = \deg(\chi)$.

3 Quasi-triangular realization of finite spectral data

Our first result is the construction of a quasi-triangular realization of a given list of finite spectral data, with a choice of degree that is compatible with the index sum theorem.

Theorem 3.1 (Quasi-Triangular Realization: Strictly Regular Case).

Suppose a list of m nonconstant monic polynomials $s_1(\lambda), \dots, s_m(\lambda)$ over an arbitrary field \mathbb{F} is given, satisfying the divisibility chain condition $s_1(\lambda) \mid s_2(\lambda) \mid \dots \mid s_m(\lambda)$. Let $\sigma := \sum_{i=1}^m \deg(s_i(\lambda))$, and define k to be the maximum degree among all of the \mathbb{F} -irreducible factors of the polynomials $s_i(\lambda)$ for $i = 1, \dots, m$. Then for any choice of nonzero $d, n \in \mathbb{N}$ such that $n \geq m$ and $dn = \sigma$, there exists an $n \times n$, degree d , strictly regular matrix polynomial $Q(\lambda)$ over \mathbb{F} that is k -quasi-triangular, and has exactly the given polynomials $s_1(\lambda), \dots, s_m(\lambda)$ as its nontrivial invariant polynomials, together with $n - m$ trivial invariant polynomials. In addition, $Q(\lambda)$ can always be chosen so that the degree of every entry in any off-diagonal block of $Q(\lambda)$ is strictly less than d .

Note that the two conditions in this theorem, i.e., that $n \geq m$ and $dn = \sigma$, are both necessary conditions for *any* strictly regular realization of the given data, whether that realization is quasi-triangular or not. The restriction $n \geq m$ simply says that the realization (and its Smith form) needs to be big enough to accommodate all of the nontrivial invariant polynomials. The condition $dn = \sigma$ is simply the Index Sum Theorem in the form (2.9) required of any strictly regular realization of data with the given σ .

Now before embarking on the extensive technical details of the proof of Theorem 3.1, which will occupy our attention for the rest of Section 3, it will be helpful to give a brief idea of the overall strategy of the argument. The first step in our process of quasi-triangular realization is to construct the Smith form corresponding to the given data and the choice of n , i.e.,

$$S(\lambda) = I_{(n-m) \times (n-m)} \oplus \text{diag}\{s_1(\lambda), s_2(\lambda), \dots, s_m(\lambda)\}.$$

We take this as our starting point, and begin changing $S(\lambda)$ by unimodular transformations, “un-diagonalizing” it and slowly turning it into the desired quasi-triangularization. The first phase of this un-diagonalization of the Smith form aims to shift irreducible factors around on the diagonal, in such a way as to try to make the degrees of all of the diagonal entries as close to the target degree d as possible. This is a proof idea pioneered in [8], and developed further in [19]. Although the diagonal form is sacrificed, in this first stage at least upper triangularity is maintained.

Now for some collections of spectral data, this phase may succeed in making the diagonal entries all have exactly the target degree d ; in this case a triangularization is obtained. As shown in [19], this can always be achieved when the underlying field is algebraically closed. However, for arbitrary fields we have to be satisfied with something less. The best that can be achieved in general is to rearrange the irreducible factors along the diagonal so that no two diagonal entries differ in degree by more than k (the maximum degree among all of the irreducible divisors of the given spectral data), and so that the vector of diagonal entries can be grouped into contiguous blocks D_1, \dots, D_ℓ for some $\ell \leq n$, each of size at most k , where the *average* degree of the entries in each D_j is exactly the target degree d .

In the second phase of construction, the upper triangular matrix attained so far is partitioned into blocks, with square (upper triangular) blocks down the diagonal that correspond to the contiguous blocks D_j of diagonal entries just mentioned. Each of these diagonal blocks is now “un-triangularized” so that it has exactly degree d . We now have a quasi-triangular realization in which all of the diagonal blocks have degree d , but the off-diagonal blocks have been left uncontrolled, and may still have degree larger than d .

The final step, then, is to visit each of the off-diagonal blocks in a “sweep by super-diagonals pattern”, using matrix polynomial division to force the degree of all of the off-diagonal blocks to be strictly less than d . This completes the construction of a k -quasi-triangular realization of degree d for the given spectral data.

Example 3.2. Throughout this paper, we will keep returning to a single illustrative example, continually developing it as we go. To start, consider the finite spectral data over the two-element field $\mathbb{F} = \mathbb{Z}_2$, consisting of the invariant polynomials

$$s_1(\lambda) = \phi\psi, \quad s_2(\lambda) = \chi\phi\psi, \quad s_3(\lambda) = \chi^2\phi\psi^2, \quad s_4(\lambda) = \chi^3\phi\psi^2, \quad \text{and} \quad s_5(\lambda) = \chi^3\phi^3\psi^4,$$

where $\chi(\lambda) = \lambda^4 + \lambda^3 + 1$, $\phi(\lambda) = \lambda^2 + \lambda + 1$, and $\psi(\lambda) = \lambda$ are \mathbb{F} -irreducible. The sum of the degrees of these invariant polynomials is $\sigma = 60$, so we may choose the target degree to be $d = 10$. As a consequence, the size of the realization must be 6×6 , and so the Smith form for our given data is

$$S(\lambda) = \text{diag}\{1, s_1(\lambda), s_2(\lambda), s_3(\lambda), s_4(\lambda), s_5(\lambda)\}.$$

The irreducible divisors together with their associated partial multiplicity sequences are

$$\begin{aligned} \mathcal{PM}_{\text{given}}(\chi) &= (0, 0, 1, 2, 3, 3) \\ \mathcal{PM}_{\text{given}}(\phi) &= (0, 1, 1, 1, 1, 3) \\ \mathcal{PM}_{\text{given}}(\psi) &= (0, 1, 1, 2, 2, 4), \end{aligned}$$

which provides an alternative way to present the given finite spectral data.

3.1 Coprime partitions, factor-counting vectors, and the unimodular transfer lemma

In this section we develop the tools needed to implement the first phase of the construction of a quasi-triangular realization from given finite spectral data. That is, we will see how it is possible to rearrange irreducible factors along the diagonal via unimodular equivalence, while at the same time maintaining upper triangularity. However, before we present these tools, some additional concepts and terminology will be required.

Let \mathcal{M} be the multiset of all of the \mathbb{F} -irreducible factors of a scalar polynomial $p(\lambda)$. A partition of \mathcal{M} into a disjoint union $\mathcal{M} = \mathcal{F} \sqcup \mathcal{G}$ of multisets \mathcal{F} and \mathcal{G} is called a *coprime partition* if every $f \in \mathcal{F}$ is coprime to every $g \in \mathcal{G}$. This is equivalent to saying that there is no irreducible factor of p that appears in both \mathcal{F} and \mathcal{G} . Given such a coprime partition, we can uniquely factor $p(\lambda)$ into

$$p(\lambda) = p_{\mathcal{F}}(\lambda) \cdot p_{\mathcal{G}}(\lambda),$$

where $p_{\mathcal{F}}(\lambda)$ denotes the product of all of the \mathbb{F} -irreducible factors in $p(\lambda)$ from \mathcal{F} , and $p_{\mathcal{G}}(\lambda)$ denotes the product of all of the \mathbb{F} -irreducible factors in $p(\lambda)$ from \mathcal{G} . We also denote by $|p(\lambda)|_{\mathcal{F}}$ the total *number* of \mathbb{F} -irreducible factors from \mathcal{F} in $p(\lambda)$, with a similar meaning for $|p(\lambda)|_{\mathcal{G}}$. Note that $|p_{\mathcal{F}}(\lambda)|_{\mathcal{F}} = |p(\lambda)|_{\mathcal{F}}$.

If $\mathbf{v}(\lambda)$ is a polynomial n -vector and $\mathcal{F} \sqcup \mathcal{G}$ is a coprime partition of the multiset of all of the \mathbb{F} -irreducible factors of all of the entries of $\mathbf{v}(\lambda)$, then the integer vector

$$|\mathbf{v}(\lambda)|_{\mathcal{F}} := (|v_1(\lambda)|_{\mathcal{F}}, \dots, |v_n(\lambda)|_{\mathcal{F}})$$

is the *factor-counting vector* of $\mathbf{v}(\lambda)$ with respect to \mathcal{F} . Also useful is the integer *degree vector*

$$\deg \mathbf{v} := (\deg v_1(\lambda), \dots, \deg v_n(\lambda)).$$

Finally, given an $n \times n$ matrix polynomial $P(\lambda)$, its main diagonal vector $\mathbf{p}(\lambda) := \text{diag } P(\lambda)$, and a coprime partition $\mathcal{F} \sqcup \mathcal{G}$ of the multiset \mathcal{M} of all of the \mathbb{F} -irreducible factors of the entries of $\mathbf{p}(\lambda)$, we define the *diagonal factor-counting vector* of $P(\lambda)$ with respect to \mathcal{F} to be the integer vector

$$\mathbf{d}_{\mathcal{F}}(P) := |\mathbf{p}(\lambda)|_{\mathcal{F}}. \quad (3.1)$$

With these concepts, terminology, and notation in hand, we can now describe and develop the tool for transferring irreducible factors along the diagonal of an upper triangular matrix polynomial.

Lemma 3.3 (2×2 Unimodular Transfer Lemma). *Let*

$$T(\lambda) = \begin{bmatrix} p(\lambda) & q(\lambda) \\ 0 & r(\lambda) \end{bmatrix}$$

be a regular 2×2 upper triangular matrix polynomial over an arbitrary field \mathbb{F} . Let $\mathcal{F} \sqcup \mathcal{G}$ be any coprime partition of the multiset of all of the \mathbb{F} -irreducible factors in the product $p(\lambda)r(\lambda)$. Let $m = |p|_{\mathcal{F}}$ and $n = |r|_{\mathcal{F}}$ so that $\mathbf{d}_{\mathcal{F}}(T) = (m, n)$. Then for any $\alpha, \beta \in \mathbb{N}$ such that $\alpha + \beta = m + n$ and $\min\{m, n\} \leq \alpha, \beta \leq \max\{m, n\}$, there exists a regular upper triangular matrix polynomial $\tilde{T}(\lambda)$ of the form

$$\tilde{T}(\lambda) = \begin{bmatrix} \tilde{p}_{\mathcal{F}}(\lambda) \cdot p_{\mathcal{G}}(\lambda) & \tilde{q}(\lambda) \\ 0 & \tilde{r}_{\mathcal{F}}(\lambda) \cdot r_{\mathcal{G}}(\lambda) \end{bmatrix} \quad (3.2)$$

with $\mathbf{d}_{\mathcal{F}}(\tilde{T}) = (\alpha, \beta)$, such that $T(\lambda)$ is unimodularly equivalent to $\tilde{T}(\lambda)$.

Proof. If $m = n$ there is nothing to do, so assume that $m \neq n$. Let $g(\lambda) := \gcd\{p(\lambda), q(\lambda), r(\lambda)\}$, and factor

$$p(\lambda) = g(\lambda) \cdot \hat{p}(\lambda) \quad \text{and} \quad r(\lambda) = g(\lambda) \cdot \hat{r}(\lambda). \quad (3.3)$$

Then the Smith form of $T(\lambda)$ is $\text{diag}\{g(\lambda), g(\lambda)\hat{p}(\lambda)\hat{r}(\lambda)\}$, so any \tilde{T} as in (3.2) that we construct that has this Smith form will be unimodularly equivalent to T . Here is how to construct many such \tilde{T} 's.

Begin by refining the factorizations of p and r in (3.3), in a way that is compatible with the given coprime partition $\mathcal{F} \sqcup \mathcal{G}$:

$$\begin{aligned} p(\lambda) &= g(\lambda) \cdot \hat{p}(\lambda) = (g_{\mathcal{F}}(\lambda)g_{\mathcal{G}}(\lambda)) \cdot (\hat{p}_{\mathcal{F}}(\lambda)\hat{p}_{\mathcal{G}}(\lambda)) \\ &= (g_{\mathcal{F}}(\lambda)\hat{p}_{\mathcal{F}}(\lambda)) \cdot (g_{\mathcal{G}}(\lambda)\hat{p}_{\mathcal{G}}(\lambda)) = p_{\mathcal{F}}(\lambda) \cdot p_{\mathcal{G}}(\lambda) \end{aligned} \quad (3.4)$$

and

$$\begin{aligned} r(\lambda) &= g(\lambda) \cdot \hat{r}(\lambda) = (g_{\mathcal{F}}(\lambda)g_{\mathcal{G}}(\lambda)) \cdot (\hat{r}_{\mathcal{F}}(\lambda)\hat{r}_{\mathcal{G}}(\lambda)) \\ &= (g_{\mathcal{F}}(\lambda)\hat{r}_{\mathcal{F}}(\lambda)) \cdot (g_{\mathcal{G}}(\lambda)\hat{r}_{\mathcal{G}}(\lambda)) = r_{\mathcal{F}}(\lambda) \cdot r_{\mathcal{G}}(\lambda). \end{aligned} \quad (3.5)$$

Since we wish to leave $p_{\mathcal{G}}(\lambda)$ and $r_{\mathcal{G}}(\lambda)$ completely undisturbed in going from T to \tilde{T} , and also to have $g(\lambda)$ present in both diagonal entries of \tilde{T} in order to preserve the Smith form, this means that the only room for maneuvering is with the factors in $\hat{p}_{\mathcal{F}}(\lambda)$ and $\hat{r}_{\mathcal{F}}(\lambda)$. So let $a(\lambda)$ and $b(\lambda)$ be *any* two polynomials (including possibly $a \equiv 1$ or $b \equiv 1$) over \mathbb{F} such that

$$a(\lambda) \cdot b(\lambda) = \hat{p}_{\mathcal{F}}(\lambda) \cdot \hat{r}_{\mathcal{F}}(\lambda), \quad (3.6)$$

and consider the polynomial matrix

$$\tilde{T}(\lambda) = \begin{bmatrix} g_{\mathcal{F}}(\lambda)a(\lambda) \cdot p_{\mathcal{G}}(\lambda) & g(\lambda) \\ 0 & g_{\mathcal{F}}(\lambda)b(\lambda) \cdot r_{\mathcal{G}}(\lambda) \end{bmatrix} = \begin{bmatrix} g(\lambda) \cdot a(\lambda)\hat{p}_{\mathcal{G}}(\lambda) & g(\lambda) \\ 0 & g(\lambda) \cdot b(\lambda)\hat{r}_{\mathcal{G}}(\lambda) \end{bmatrix}, \quad (3.7)$$

of the form in (3.2) with $\tilde{p}_{\mathcal{F}}(\lambda) = g_{\mathcal{F}}(\lambda)a(\lambda)$, $\tilde{q}(\lambda) = g(\lambda)$, and $\tilde{r}_{\mathcal{F}}(\lambda) = g_{\mathcal{F}}(\lambda)b(\lambda)$. Then the gcd of the entries of \tilde{T} is easily seen to be $g(\lambda)$, so the Smith form of $\tilde{T}(\lambda)$ is

$$\begin{aligned} \text{diag} \left\{ g(\lambda), g(\lambda)[a(\lambda)b(\lambda)] \cdot [\hat{p}_{\mathcal{G}}(\lambda)\hat{r}_{\mathcal{G}}(\lambda)] \right\} &= \text{diag} \left\{ g(\lambda), g(\lambda)[\hat{p}_{\mathcal{F}}(\lambda)\hat{r}_{\mathcal{F}}(\lambda)] \cdot [\hat{p}_{\mathcal{G}}(\lambda)\hat{r}_{\mathcal{G}}(\lambda)] \right\} \\ &= \text{diag} \{ g(\lambda), g(\lambda)\hat{p}(\lambda)\hat{r}(\lambda) \}, \end{aligned}$$

which is identical to the Smith form of $T(\lambda)$. Thus for any choice of $a(\lambda)$ and $b(\lambda)$ in (3.6), we have $\tilde{T} \sim T$. Letting $\delta = |g(\lambda)|_{\mathcal{F}}$, we see that the diagonal factor-counting vector $\mathbf{d}_{\mathcal{F}}(\tilde{T})$ for $\tilde{T}(\lambda)$ in (3.7) can be (α, β) for any $\alpha + \beta = m + n$ with $\delta \leq \alpha, \beta \leq m + n - \delta$. Since

$$\delta \leq \min(m, n) \leq \max(m, n) \leq m + n - \delta,$$

any (α, β) pair given in the statement of the lemma is always achievable for $\mathbf{d}_{\mathcal{F}}(\tilde{T})$. \square

This lemma for 2×2 triangular matrices can now be used on a triangular matrix of any size to “transfer” irreducible factors belonging to a family \mathcal{F} between any two *adjacent* diagonal entries, while maintaining triangularity, and without disturbing any of the factors that belong to the complementary family \mathcal{G} . This is done by embedding the 2×2 unimodular transformations provided by Lemma 3.3 into larger identity matrices.

Corollary 3.4. *Let $P(\lambda)$ be a regular $n \times n$ upper triangular polynomial matrix, and let $\mathcal{M} = \mathcal{F} \sqcup \mathcal{G}$ be a coprime partition of the multiset of all irreducible factors of the diagonal entries of $P(\lambda)$. Consider any 2×2 principal submatrix of $P(\lambda)$ with adjacent diagonal entries, i.e.,*

$$T(\lambda) = \begin{bmatrix} p_{ii}(\lambda) & p_{ij}(\lambda) \\ 0 & p_{jj}(\lambda) \end{bmatrix} \quad \text{with } j = i + 1,$$

and let $\mathbf{d}_{\mathcal{F}}(T) = (m, n)$. Then for any integers α, β such that $\alpha + \beta = m + n$ and $\min\{m, n\} \leq \alpha, \beta \leq \max\{m, n\}$, there exists a regular upper triangular matrix polynomial $\tilde{P}(\lambda)$ such that

- $\tilde{p}_{ii}(\lambda) = [\tilde{p}_{ii}(\lambda)]_{\mathcal{F}} \cdot p_{\mathcal{G}}(\lambda)$ with $|\tilde{p}_{ii}(\lambda)|_{\mathcal{F}} = \alpha$,
- $\tilde{p}_{jj}(\lambda) = [\tilde{p}_{jj}(\lambda)]_{\mathcal{F}} \cdot p_{\mathcal{G}}(\lambda)$ with $|\tilde{p}_{jj}(\lambda)|_{\mathcal{F}} = \beta$,
- $\tilde{p}_{kk}(\lambda) = p_{kk}(\lambda)$ for all $k \neq i, j$,
- $\tilde{P}(\lambda)$ is unimodularly equivalent to $P(\lambda)$.

Proof. Apply Lemma 3.3 to the submatrix T to get two 2×2 unimodular matrices $E(\lambda)$ and $F(\lambda)$ such that

$$E(\lambda)T(\lambda)F(\lambda) = \begin{bmatrix} [\tilde{p}_{ii}(\lambda)]_{\mathcal{F}} \cdot p_{\mathcal{G}}(\lambda) & \tilde{p}_{ij}(\lambda) \\ 0 & [\tilde{p}_{jj}(\lambda)]_{\mathcal{F}} \cdot p_{\mathcal{G}}(\lambda) \end{bmatrix}$$

with diagonal \mathcal{F} -factor-counting vector (α, β) . Then construct the $n \times n$ unimodular matrices

$$\hat{E}(\lambda) = \left[\begin{array}{c|c|c} I_{i-1} & & \\ \hline & E(\lambda) & \\ \hline & & I_{n-i-1} \end{array} \right] \quad \text{and} \quad \hat{F}(\lambda) = \left[\begin{array}{c|c|c} I_{i-1} & & \\ \hline & F(\lambda) & \\ \hline & & I_{n-i-1} \end{array} \right].$$

Transforming $P(\lambda)$ using these two unimodular matrices, we obtain the desired matrix $\tilde{P}(\lambda) = \hat{E}(\lambda)P(\lambda)\hat{F}(\lambda)$. \square

Example 3.5. Consider the Smith form $S(\lambda)$ from Example 3.2. The multiset \mathcal{M} of all of the irreducible factors in the entries of $S(\lambda)$ contains many copies of χ , ϕ , and ψ , but we can partition it into $\mathcal{M} = \mathcal{F}_1 \sqcup \mathcal{F}_2 \sqcup \mathcal{F}_3 \sqcup \mathcal{F}_4$, where \mathcal{F}_j contains all of the irreducible factors of degree j . Thus we see that \mathcal{F}_1 contains all of the copies of ψ , \mathcal{F}_2 contains all of the copies of ϕ , \mathcal{F}_3 is empty, and \mathcal{F}_4 contains all of the copies of χ . The diagonal factor-counting vectors of $S(\lambda)$ then are

$$\begin{aligned} \mathbf{d}_{\mathcal{F}_1}(S) &= (0, 1, 1, 2, 2, 4), & \mathbf{d}_{\mathcal{F}_3}(S) &= (0, 0, 0, 0, 0, 0), \\ \mathbf{d}_{\mathcal{F}_2}(S) &= (0, 1, 1, 1, 1, 3), & \mathbf{d}_{\mathcal{F}_4}(S) &= (0, 0, 1, 2, 3, 3). \end{aligned}$$

Note that we include $\mathbf{d}_{\mathcal{F}_3}(S)$ here not only for the sake of completeness, but also to show that having any of the partition multisets be empty is allowed. Later on, we will use Corollary 3.4 to rearrange the irreducible factors along the diagonal so that these diagonal factor-counting vectors will have the property of being *1-homogeneous*, a concept to be defined in the next section.

Note that these diagonal factor-counting vectors just happen to match up with the partial multiplicity sequences in this example. But this is not typical, and follows from there being at most one irreducible divisor in each piece of the given coprime partition of \mathcal{M} . For general coprime partitions, diagonal factor-counting vectors of Smith forms will each be a *sum* of partial multiplicity sequences.

3.2 Homogenization of natural vectors and un-diagonalizing the Smith form

With the ability to transfer irreducible factors along the diagonal (Corollary 3.4) now in our tool box, it is time to see how to employ that tool to rearrange the diagonal irreducible factors so as to make the new diagonal entries as close in degree to each other as possible. This is phase 1 of our quasi-triangular realization construction. It will be shown that the best that can be done in general is to make these diagonal entry degrees differ by at most k , where k is the highest degree among all of the irreducible factors along the diagonal. To facilitate the discussion of this process, we introduce the following two concepts. Note that vectors whose entries are natural numbers, in particular diagonal factor-counting vectors, appear frequently in this discussion; such vectors will be referred in brief as *natural vectors*. Keep in mind that in this paper the natural numbers \mathbb{N} *includes zero*.

Definition 3.6. A natural vector $\mathbf{v} = (v_1, \dots, v_r) \in \mathbb{N}^r$ is *k-homogeneous* if $|v_i - v_j| \leq k$ for any $1 \leq i, j \leq r$.

Definition 3.7. Let $\mathbf{v} = (v_1, \dots, v_r) \in \mathbb{N}^r$ with component sum $s = v_1 + v_2 + \dots + v_r$, and divide s by r to get $s = qr + t$, with $0 \leq t < r$. Then any permutation of the 1-homogeneous vector

$$(q, q, \dots, q, \underbrace{q+1, q+1, \dots, q+1}_{t \text{ copies}}) \in \mathbb{N}^r$$

is called a *homogenization* of \mathbf{v} .

Before addressing our primary objective, that is, the phase 1 rearrangement of diagonal irreducible factors, it will be useful to do a preliminary examination of the process of homogenizing natural vectors via two very simple operations that we will refer to as “interchange” and “compression”. We will see that these two operations on natural vectors are closely related to transformations of the diagonal of upper triangular polynomial matrices achievable by the Unimodular Transfer Corollary 3.4. The operations of interchange and compression also bring us into contact with the classical notion of majorization of vectors, which we recall next.

Definition 3.8 (Majorization [14]).

For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, let \mathbf{x}' and \mathbf{y}' denote the permutations of those vectors in which the entries have been arranged in decreasing order. We say that \mathbf{x} *majorizes* \mathbf{y} , or \mathbf{y} *is majorized by* \mathbf{x} , and write $\mathbf{x} \succeq \mathbf{y}$, if

$$\sum_{i=1}^{\ell} x'_i \geq \sum_{i=1}^{\ell} y'_i \quad \text{for } \ell = 1, 2, \dots, n, \quad (3.8)$$

with equality when $\ell = n$. Clearly this definition also applies without change when restricted to integer vectors $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^n$, or even to natural vectors $\mathbf{x}, \mathbf{y} \in \mathbb{N}^n$.

Remark 3.9. It is useful to keep in mind the basic intuition about majorization, i.e., that if \mathbf{x} majorizes \mathbf{y} , then the entries of \mathbf{x} are more “spread out” than those of \mathbf{y} . This can be seen, at least in part, by examining the two extreme components of the ordered vectors \mathbf{x}' and \mathbf{y}' . From the $\ell = 1$ inequality in (3.8) we have $x'_1 \geq y'_1$, and by combining the $\ell = n$ and $\ell = n - 1$ parts of (3.8) we see that $y'_n \geq x'_n$. Thus we have $x'_1 \geq y'_1 \geq \dots \geq y'_n \geq x'_n$, displaying some of the greater dispersion of the entries of \mathbf{x}' .

It is a classical result of the study of majorization (attributed in [14] to a 1903 article of Muirhead [15]) that for natural vectors \mathbf{x} and \mathbf{y} , \mathbf{x} majorizes \mathbf{y} ($\mathbf{x} \succeq \mathbf{y}$) if and only if \mathbf{x} can be transformed into \mathbf{y} by a finite sequence of operations known variously as *transfers*, *Dalton transfers*, or more recently (and colorfully) as *Robin Hood transfers* – think “rob from the rich and give to the poor”. Such an operation takes any two components of a natural vector \mathbf{v} , say v_i and v_j , and replaces them by natural numbers α and β that are closer together in size; more specifically, where $\alpha + \beta = v_i + v_j$ and $|\alpha - \beta| \leq |v_i - v_j|$. (Note that in his proof of this result, Muirhead uses only transfers in which v_i and v_j change by exactly 1, i.e., $\alpha = v_i \pm 1$ and $\beta = v_j \mp 1$.)

For our purposes, we need to have a Muirhead-like result for converting a natural vector into a homogenization of itself. However, since we will ultimately need to implement these conversions by unimodular transformations that preserve upper triangularity (see Corollary 3.4), it is essential that we limit our operations on natural vectors to ones that act only on *adjacent entries* of a vector, which is more restrictive than the Robin Hood transfers used by Muirhead. Thus we introduce the following two operations acting on natural vectors:

- **Interchange** of *adjacent* components:

$$(v_1, \dots, v_i, v_{i+1}, \dots, v_r) \rightsquigarrow (v_1, \dots, v_{i+1}, v_i, \dots, v_r) \in \mathbb{N}^r$$

- **Compression** of *adjacent* components:

$$(v_1, \dots, v_{i-1}, v_i, v_{i+1}, v_{i+2}, \dots, v_r) \rightsquigarrow (v_1, \dots, v_{i-1}, \alpha, \beta, v_{i+2}, \dots, v_r) \in \mathbb{N}^r,$$

where $\alpha + \beta = v_i + v_{i+1}$ and $|\alpha - \beta| < |v_i - v_{i+1}|$.

Equivalently, we can instead require that $\alpha + \beta = v_i + v_{i+1}$ and $\min\{v_i, v_{i+1}\} < \alpha, \beta < \max\{v_i, v_{i+1}\}$.

Note that we use the new terminology “interchange” and “compression” here for these more limited operations, to try to avoid any confusion with the more flexible Robin Hood transfers. We now state and prove the modified version of the classic Muirhead result, specialized to the transformation of a natural vector into a homogenization of itself. As an immediate consequence of Lemma 3.10 and the Muirhead theorem, we see that any natural vector \mathbf{v} majorizes any homogenization of itself.

Lemma 3.10 (Homogenization Lemma).

Consider any $\mathbf{v} = (v_1, \dots, v_r) \in \mathbb{N}^r$, with component sum $s = v_1 + \dots + v_r$, and average component value $a = s/r$. Let $q \in \mathbb{N}$ be an integer such that the average a is contained in the closed interval $[q, q+1]$. Then by a finite sequence of interchanges and at most $r-1$ compressions, \mathbf{v} can be transformed into a homogenization of itself, comprised of only q and $q+1$ components, ordered arbitrarily. Writing $s = qr + t$ with $0 \leq t \leq r$, then in the homogenization of \mathbf{v} there are exactly t copies of $q+1$ and $r-t$ copies of q .

Proof. The proof will proceed by a pair of inductions, each on the length r of the vector \mathbf{v} . The first induction considers only vectors \mathbf{v} whose average component value a is an integer, either q or $q+1$. In this case we will see that the homogenization of \mathbf{v} is the vector (a, a, \dots, a) . The second induction, which builds on the result of the first, considers the vectors \mathbf{v} where a is *not* an integer, i.e., $q < a < q+1$. For this second case the homogenization will have both q and $q+1$ entries.

Part 1 (a is an integer)

The case $r = 1$ is trivial, and the case $r = 2$ is also very easy – just do a single ($r-1 = 1$) compression to produce the homogenization (a, a) . So now assume that every k -vector $\mathbf{v} \in \mathbb{N}^k$ with $k < r$ and integer average a can be finitely transformed into a homogenization (a, a, \dots, a) using at most $k-1$ compressions. Let $\mathbf{v} \in \mathbb{N}^r$ be any r -vector with average value $a = q$. If \mathbf{v} already has any component equal to q , then finitely many interchanges move this component to the end, producing $\mathbf{v} \rightsquigarrow (\tilde{\mathbf{v}}, q)$. If not, then one can find at least one component larger than q and one component smaller than q . By interchanges make these components adjacent, and then a compression on these two components can produce at least one q component. Finitely many more interchanges then moves this newly-produced q to the end, again producing $\mathbf{v} \rightsquigarrow (\tilde{\mathbf{v}}, q)$. In either case, the $(r-1)$ -vector $\tilde{\mathbf{v}}$ has average component value $\tilde{a} = (s - q)/(r-1) = (ar - a)/(r-1) = a = q$. By the induction hypothesis the vector $\tilde{\mathbf{v}}$ can now be transformed by finitely many interchanges and at most $r-2$ compressions into the vector (a, a, \dots, a) , and thus we have the desired conclusion for the r -vector \mathbf{v} . The same argument works, mutatis mutandis, for vectors $\mathbf{v} \in \mathbb{N}^r$ with average value $a = q+1$.

Part 2 (a is *not* an integer, i.e., $q < a < q+1$ and $0 < t < r$)

Here the base case for the induction cannot be $r = 1$; it must be $r = 2$, with $a = q + \frac{1}{2}$. In this $r = 2$ case, then, a single compression can produce the vector $(q, q+1)$, as desired. Again as in Part 1, now assume that every k -vector $\mathbf{v} \in \mathbb{N}^k$ with $k < r$ and average $q < a < q+1$ can be finitely transformed into a homogenized version containing only q and $q+1$ components, using at most $k-1$ compressions. Let $\mathbf{v} \in \mathbb{N}^r$ be any r -vector with non-integer average value $q < a < q+1$. If \mathbf{v} already has any component equal to q or $q+1$, then finitely many interchanges moves this component to the end, producing $\mathbf{v} \rightsquigarrow (\tilde{\mathbf{v}}, q)$ or $(\tilde{\mathbf{v}}, q+1)$. If no component of \mathbf{v} is q or $q+1$, then there must be at least one component larger than $q+1$ and one component smaller than q . Make these components adjacent by interchanges, and compress these two components to produce at least one q or $q+1$ entry (possibly both, or perhaps even two q 's or two $q+1$'s). Move this q or $q+1$ entry to the end by interchanges, once again transforming \mathbf{v} into the form $(\tilde{\mathbf{v}}, q)$ or $(\tilde{\mathbf{v}}, q+1)$. The $(r-1)$ -vector $\tilde{\mathbf{v}}$ now has a component average

$$\tilde{a} = \begin{cases} \frac{s-q}{r-1} = \frac{qr+t-q}{r-1} = q + \frac{t}{r-1} \leq q+1 & \text{if } \mathbf{v} \rightsquigarrow (\tilde{\mathbf{v}}, q), \\ \frac{s-(q+1)}{r-1} = \frac{qr+t-q-1}{r-1} = q + \frac{t-1}{r-1} \geq q & \text{if } \mathbf{v} \rightsquigarrow (\tilde{\mathbf{v}}, q+1). \end{cases}$$

If \tilde{a} is q or $q+1$, then we are done by Part 1. Otherwise $q < \tilde{a} < q+1$, and we are done by the induction hypothesis of Part 2.

This completes the inductive argument that a homogenization of \mathbf{v} can always be achieved. The expression of the component sum s in the form $s = qr + t$ with $0 \leq t \leq r$ then uniquely determines the number of q and

$q + 1$ entries in the homogenization. A final sequence of interchanges can put the entries of the homogenization into any desired order, thus completing the proof of the lemma. \square

With the Homogenization Lemma in hand, we return to the first phase of our quasi-triangular realization process, the “un-diagonalizing” of a Smith form. The next result addresses the combinatorial essence of this problem, showing how it is possible to take a multiset of irreducible polynomials and distribute them among the entries of a vector in a way that minimizes the degree differences, and at the same time produces a viable configuration for the diagonal vector of an upper triangular un-diagonalized Smith form. The corollary immediately following shows that the target diagonal produced by Lemma 3.11 is in fact reachable from the Smith form via the type of triangularity-preserving unimodular transformations developed in Corollary 3.4.

Lemma 3.11. *Let \mathbb{F} be an arbitrary field, and consider a finite multiset \mathcal{M} of \mathbb{F} -irreducible polynomials of degree less than or equal to k . Also let n be any positive integer. Let $\mathcal{M} = \mathcal{F}_1 \sqcup \cdots \sqcup \mathcal{F}_k$ be the coprime partition of \mathcal{M} in which \mathcal{F}_j contains all of the \mathbb{F} -irreducible factors in \mathcal{M} of degree j . Then there exists a polynomial n -vector $\mathbf{p}(\lambda)$ with all nonzero entries such that*

- *the multiset of all of the \mathbb{F} -irreducible factors of all the entries of $\mathbf{p}(\lambda)$ is exactly \mathcal{M} ,*
- *the degree vector $\deg \mathbf{p} := (\deg p_1(\lambda), \dots, \deg p_n(\lambda))$ is k -homogeneous,*
- *and each of the factor-counting vectors $|\mathbf{p}(\lambda)|_{\mathcal{F}_j}$ for $j = 1, \dots, k$ is 1-homogeneous.*

Proof. For notational convenience, we first append enough copies of the constant polynomial 1 to the multiset \mathcal{M} so that the total number m of elements in \mathcal{M} is a multiple of n , i.e., $|\mathcal{M}| =: m = (q + 1)n$ with $q \geq 0$. The presence of these copies of 1 will have no impact on the vector \mathbf{p} that is constructed, but the exposition will be simplified by including them.

Next order the elements of \mathcal{M} into a list \mathcal{L} so that the sequence of degrees is decreasing. That is, let $\mathcal{L} = [a_1(\lambda), a_2(\lambda), \dots, a_m(\lambda)]$, where $\deg a_\ell \geq \deg a_{\ell+1}$ for every $\ell = 1, 2, \dots, m-1$. Now partition the list \mathcal{L} into $q + 1$ contiguous sublists of n polynomials each, and stack them up as in the following diagram. (Keep in mind that $qn + n = m$.)

$$\begin{array}{rcl}
 \mathcal{L}_0 & = & [a_1, \quad a_2, \quad \dots, a_n] \\
 \mathcal{L}_1 & = & [a_{n+1}, \quad a_{n+2}, \quad \dots, a_{2n}] \\
 & \vdots & \\
 \mathcal{L}_j & = & [a_{jn+1}, \quad a_{jn+2}, \quad \dots, a_{j(n+n)}] \\
 & \vdots & \\
 \mathcal{L}_q & = & [a_{qn+1}, \quad a_{qn+2}, \quad \dots, a_{qn+n}] \\
 & & \textcolor{red}{v_1(\lambda)} \quad \textcolor{red}{v_2(\lambda)} \quad \dots \quad \textcolor{red}{v_n(\lambda)}
 \end{array}$$

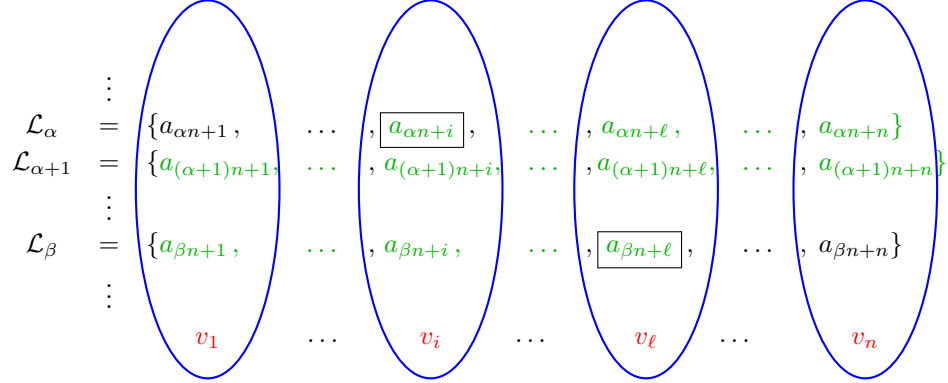
The n component polynomials of the desired vector $\mathbf{p}(\lambda)$ are formed by taking the column-wise products indicated by the blue ovals. For example, $p_1(\lambda) := a_1(\lambda)a_{n+1}(\lambda) \cdots a_{qn+1}(\lambda)$, $p_2(\lambda) := a_2(\lambda)a_{n+2}(\lambda) \cdots a_{qn+2}(\lambda)$, etc. It is clear by construction that the multiset of all of the \mathbb{F} -irreducible factors of all the entries of $\mathbf{p}(\lambda)$ is exactly \mathcal{M} . All that remains is to see why this $\mathbf{p}(\lambda)$ has the other two desired properties.

To see why the degree vector $\deg \mathbf{p}$ is k -homogeneous, first observe that the component degrees are decreasing, i.e., $\deg p_1 \geq \deg p_2 \geq \cdots \geq \deg p_n$. This follows immediately from the elements of the list \mathcal{L} being in decreasing order. Consequently the largest degree difference in \mathbf{p} will be between the first and last components $p_1(\lambda)$ and $p_n(\lambda)$. But we have

$$\begin{aligned}
 \deg p_1 - \deg p_n &= (\deg a_1 + \deg a_{n+1} + \cdots + \deg a_{qn+1}) - (\deg a_n + \deg a_{2n} + \cdots + \deg a_{qn+n}) \\
 &= \deg a_1 - [(\deg a_n - \deg a_{n+1}) + (\deg a_{2n} - \deg a_{2n+1}) + \cdots + (\deg a_{qn} - \deg a_{qn+1}) + \deg a_m] \\
 &\leq k,
 \end{aligned}$$

since $\deg a_1(\lambda) \leq k$, and the expression inside the brackets is a sum of non-negative numbers (due to the list \mathcal{L} being in decreasing-degree order). Hence the degree vector $\deg \mathbf{p}(\lambda)$ is k -homogeneous.

Finally, consider the factor-counting vectors $|\mathbf{p}(\lambda)|_{\mathcal{F}_j}$. To see why these vectors are all 1-homogeneous, observe that because the irreducible polynomials are listed in \mathcal{L} in order of decreasing degree, the degree j irreducibles in \mathcal{L} form a contiguous segment of \mathcal{L} , and hence are distributed among consecutive sublists \mathcal{L}_δ in a manner analogous to the green entries in the following diagram. On the other hand, if $\alpha < \beta$, then



The first instance of a degree j irreducible is the boxed entry $a_{\alpha n+i}$ with $1 \leq i \leq n$ in some sublist \mathcal{L}_α , and the last degree j irreducible is the boxed entry $a_{\beta n+l}$, where $\alpha \leq \beta$.

If $\alpha = \beta$, then $i \leq \ell$, and all of the degree j irreducibles in \mathcal{L} are in the sublist \mathcal{L}_α , so the factor-counting vector $|\mathbf{p}(\lambda)|_{\mathcal{F}_j}$ has only 0 and 1 entries; this certainly constitutes a 1-homogeneous vector. On the other hand, if $\alpha < \beta$ then we may have either $\ell < i - 1$, $\ell = i - 1$, or $\ell > i - 1$. Consider each of these possibilities in turn:

- $\ell < i - 1$: In this case the combined contribution of the degree j entries in \mathcal{L}_α and \mathcal{L}_β to $|\mathbf{p}(\lambda)|_{\mathcal{F}_j}$ is the n -vector

$$(1, 1, \dots, \underbrace{1}_\ell, 0, \dots, 0, \underbrace{1}_i, \dots, 1),$$

while the contribution of the sublists \mathcal{L}_δ with $\alpha < \delta < \beta$ is a constant vector with n entries all equal to $\beta - \alpha - 1$. The sum of these two vectors is $|\mathbf{p}(\lambda)|_{\mathcal{F}_j}$, and is clearly 1-homogeneous.

- $\ell = i - 1$: Now the contribution to $|\mathbf{p}(\lambda)|_{\mathcal{F}_j}$ from \mathcal{L}_α and \mathcal{L}_β combined is just the constant n -vector $(1, 1, \dots, 1)$, which together with the contribution from the sublists between \mathcal{L}_α and \mathcal{L}_β gives a constant vector for $|\mathbf{p}(\lambda)|_{\mathcal{F}_j}$, with entries all equal to $\beta - \alpha$. This is certainly 1-homogeneous, indeed even 0-homogeneous.
- $\ell > i - 1$: In this final case the combined contribution of the degree j entries in \mathcal{L}_α and \mathcal{L}_β to $|\mathbf{p}(\lambda)|_{\mathcal{F}_j}$ is the n -vector

$$(1, \dots, 1, \underbrace{2}_i, 2, \dots, 2, \underbrace{2}_\ell, 1, \dots, 1).$$

Together with the constant vector contribution from the sublists between \mathcal{L}_α and \mathcal{L}_β , we again see that $|\mathbf{p}(\lambda)|_{\mathcal{F}_j}$ is 1-homogeneous. □

Remark 3.12. Note that k -homogeneity for $\mathbf{p}(\lambda)$ in Lemma 3.11 is the best possible general result here, as illustrated by the following simple example. Suppose \mathcal{M} contains *only* irreducible polynomials of degree k , say ℓ of them, and $n > \ell$. Then clearly the n -vector $\mathbf{p}(\lambda)$ that minimizes the degree differences has $|\mathbf{p}(\lambda)|_{\mathcal{F}_k} = (k, k, \dots, k, 0, \dots, 0)$, and k -homogeneity cannot be improved upon in this situation.

We return to our running illustration of the results of this paper, as begun earlier in Examples 3.2 and 3.5. The next example demonstrates the application of Lemma 3.11 to this data.

Example 3.13. Recall the 6×6 Smith form from Example 3.2, i.e.,

$$S(\lambda) = \text{diag}\{1, \phi\psi, \chi\phi\psi, \chi^2\phi\psi^2, \chi^3\phi\psi^2, \chi^3\phi^3\psi^4\},$$

and consider the coprime partition of the list of irreducible divisors as in Example 3.5, with

$$\mathcal{F}_1 = \{\underbrace{\psi, \psi, \psi, \psi, \psi, \psi, \psi, \psi, \psi, \psi}_{10}\}, \quad \mathcal{F}_2 = \{\underbrace{\phi, \phi, \phi, \phi, \phi, \phi, \phi}_{7}\}, \quad \mathcal{F}_4 = \{\underbrace{\chi, \chi, \chi, \chi, \chi, \chi, \chi, \chi, \chi}_{9}\},$$

and an \mathcal{F}_3 that is empty. Following the proof of Lemma 3.11 gives us five sublists, each of length 6:

$$\begin{aligned}\mathcal{L}_0 &= \{\chi, \chi, \chi, \chi, \chi, \chi\} \\ \mathcal{L}_1 &= \{\chi, \chi, \chi, \phi, \phi, \phi\} \\ \mathcal{L}_2 &= \{\phi, \phi, \phi, \phi, \psi, \psi\} \\ \mathcal{L}_3 &= \{\psi, \psi, \psi, \psi, \psi, \psi\} \\ \mathcal{L}_4 &= \{\psi, \psi, 1, 1, 1, 1\}.\end{aligned}$$

Then taking the products going down the columns, we get $\mathbf{p}(\lambda) = (p_1(\lambda), \dots, p_6(\lambda))$ with

$$p_1(\lambda) = \chi^2 \phi \psi^2, \quad p_2(\lambda) = \chi^2 \phi \psi^2, \quad p_3(\lambda) = \chi^2 \phi \psi, \quad p_4(\lambda) = \chi \phi^2 \psi, \quad p_5(\lambda) = \chi \phi \psi^2, \quad p_6(\lambda) = \chi \phi \psi^2.$$

The vector $\mathbf{p}(\lambda)$ has factor-counting vectors

$$\begin{aligned}|\mathbf{p}(\lambda)|_{\mathcal{F}_1} &= (2, 2, 1, 1, 2, 2) \\ |\mathbf{p}(\lambda)|_{\mathcal{F}_2} &= (1, 1, 1, 2, 1, 1) \\ |\mathbf{p}(\lambda)|_{\mathcal{F}_3} &= (0, 0, 0, 0, 0, 0) \\ |\mathbf{p}(\lambda)|_{\mathcal{F}_4} &= (2, 2, 2, 1, 1, 1),\end{aligned}$$

which are all 1-homogeneous, and a 4-homogeneous degree vector $\deg \mathbf{p} = (12, 12, 11, 9, 8, 8)$, just as guaranteed by Lemma 3.11.

We now have the tools needed to reach the next milestone in our construction of a quasi-triangular realization of given finite spectral data. The following corollary is the main result to carry forward into the next stages of this construction.

Corollary 3.14 (Un-diagonalizing the Smith form).

Let \mathbb{F} be an arbitrary field, and consider any regular $n \times n$ diagonal polynomial matrix $S(\lambda)$ over \mathbb{F} that is in Smith form. Let \mathcal{M} be the multiset of all of the \mathbb{F} -irreducible factors of all of the invariant polynomials in $S(\lambda)$, and let k be the maximum degree among all elements of \mathcal{M} . Consider also the coprime partition $\mathcal{M} = \mathcal{F}_1 \sqcup \dots \sqcup \mathcal{F}_k$, in which each \mathcal{F}_j contains all of the \mathbb{F} -irreducible factors in \mathcal{M} of degree j . Then there is an upper triangular polynomial matrix $T(\lambda)$ that is unimodularly equivalent to $S(\lambda)$, with diagonal degree vector $\deg(\text{diag } T(\lambda))$ that is k -homogeneous, and such that each diagonal factor-counting vector $\mathbf{d}_{\mathcal{F}_j}(T)$ is 1-homogeneous.

Proof. Use the given multiset \mathcal{M} of \mathbb{F} -irreducible polynomials as input to Lemma 3.11. The output vector $\mathbf{p}(\lambda)$ from that Lemma is now the target diagonal for the desired upper triangular $T(\lambda)$. Since each factor-counting vector $|\mathbf{p}(\lambda)|_{\mathcal{F}_j}$ is 1-homogeneous, the Homogenization Lemma 3.10 guarantees that the transition from the vector $\text{diag } S(\lambda)$ to the vector $\mathbf{p}(\lambda)$ can be achieved using a finite number of compressions and interchanges. But Corollary 3.4 gives us the means to implement all of these compressions and interchanges as unimodular transformations applied to $S(\lambda)$. Doing this then converts $S(\lambda)$ into the desired upper triangular $T(\lambda)$. \square

3.3 A combinatorial lemma

In this short section we focus on establishing a new combinatorial property of “tightly packed” *integer* multisets, that is, multisets that contain more, perhaps even many more elements than the width of the interval into which they are packed. This property will enable us to permute the diagonal entries of the $T(\lambda)$ from Corollary 3.14 in preparation for the final phase of our quasi-triangular realization construction. Note that for ease of expression, in this section we use the word “list” as a synonym for multiset; however, nothing about any ordering of these lists is of any relevance for the development here.

Lemma 3.15 (Homogeneous Partitioning Property).

Let $\mathcal{I} = [j, j+k]$ be a closed interval with integer endpoints and length $k \geq 1$, and consider a list $\mathcal{L} = \{n_1, n_2, \dots, n_m\}$ of m integers, all in \mathcal{I} . If the average value of all of the entries in \mathcal{L} is an integer $\mu \in \mathcal{I}$, then \mathcal{L} can be partitioned into sublists $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_\ell$ such that the number of entries in each \mathcal{S}_i does not exceed k , and the average value of each sublist is exactly μ .

To help prove this result, we need another lemma characterizing the solution set of a certain diophantine equation in two variables.

Lemma 3.16. *Let a and b be positive integers, with $d := \gcd\{a, b\}$. Then the set of integer solutions (x, y) of the equation $ax = by$ consists of all the integer multiples of the pair (\tilde{b}, \tilde{a}) , where $\tilde{a} = a/d$ and $\tilde{b} = b/d$.*

Proof. We can immediately reduce to the equation $\tilde{a}x = \tilde{b}y$, which has exactly the same solution set as $ax = by$. Since \tilde{a} and \tilde{b} are relatively prime, for any solution pair (x, y) it must be true that $\tilde{a}|y$ and $\tilde{b}|x$. Letting $x = \tilde{b}k$ and $y = \tilde{a}\ell$, then $\tilde{a}\tilde{b}k = \tilde{b}\tilde{a}\ell$ implies that $k = \ell$. Thus the solutions of $\tilde{a}x = \tilde{b}y$ (and hence also of $ax = by$) are exactly the integer multiples of (\tilde{b}, \tilde{a}) . \square

We now prove the Homogeneous Partitioning Property.

Proof. (of Lemma 3.15) One strategy to achieve this homogeneous partitioning is to first translate all the given data (i.e., the interval \mathcal{I} and the list \mathcal{L}) by any fixed constant $c \in \mathbb{Z}$, solve the translated problem, and then “un-translate” the solution back to the original location. Thus it suffices to solve the problem for the case when $\mu = 0$ and the interval \mathcal{I} is $[e_\ell, e_r]$, with endpoints $e_\ell \leq 0$ and $e_r \geq 0$ such that $e_r - e_\ell = k$. The goal in this more specialized scenario, then, is to partition \mathcal{L} into sublists of no more than k entries each, such that the *sum* (equals the average) of the entries of each sublist \mathcal{S}_i is zero.

The proof for this special scenario proceeds by an induction on m , the number of elements in the list \mathcal{L} . (Note that this induction can be easily converted into an algorithm for computing the desired partition.) The base case(s) for this induction are all m such that $1 \leq m \leq k$, for which the result is trivially true. So now suppose that the homogeneous partitioning property holds for all number lists satisfying the hypotheses of Lemma 3.15 with $m \leq h$, and consider a list \mathcal{L} of $m = h + 1$ integers with values in $\mathcal{I} = [e_\ell, e_r]$ and average value $\mu = 0$. There are now three cases to consider:

- (a) *Some $n_i \in \mathcal{L}$ is equal to zero.*

In this case we can split off the singleton sublist $\mathcal{S}_1 = \{n_i\}$ from \mathcal{L} , leaving a smaller list $\hat{\mathcal{L}}$ with $m = h$ entries, and average value $\hat{\mu} = 0$. Applying the inductive hypothesis to $\hat{\mathcal{L}}$ completes the homogeneous partitioning of \mathcal{L} .

- (b) *No element of \mathcal{L} is zero, and all elements of \mathcal{L} lie at the endpoints e_ℓ and e_r of \mathcal{I} .*

Suppose there are α copies of e_ℓ and β copies of e_r , so that $\alpha + \beta = h + 1 > k$, and $\alpha e_\ell + \beta e_r = 0$, or equivalently $(-e_\ell)\alpha = e_r\beta$. Let $d = \gcd\{-e_\ell, e_r\}$. Applying Lemma 3.16 to the equation $ax = by$ with $a = -e_\ell$ and $b = e_r$, we see that the solution $(x, y) = (\alpha, \beta)$ to $ax = by$ is an integer multiple of $(\frac{e_r}{d}, \frac{-e_\ell}{d})$. Thus we can completely partition \mathcal{L} into sublists \mathcal{S}_i , each consisting of $\frac{e_r}{d}$ copies of e_ℓ and $\frac{-e_\ell}{d}$ copies of e_r . Since $\frac{e_r}{d} + \frac{-e_\ell}{d} = \frac{k}{d}$, each of these sublists has $\frac{k}{d} \leq k$ elements, and sum zero, as desired.

- (c) *No element of \mathcal{L} is zero, but there is some n_{i_1} from \mathcal{L} in the open interval $\tilde{\mathcal{I}} = (e_\ell, e_r)$.*

Begin building a sublist \mathcal{S} with the given (nonzero) element n_{i_1} in $\tilde{\mathcal{I}}$. Pick from among the remaining elements of \mathcal{L} to update $\mathcal{S} = \{n_{i_1}, n_{i_2}, \dots, n_{i_j}, \dots\}$, and keep track of the “partial sums” $\sigma_j := \sum_{\ell=1}^j n_{i_\ell}$ as you go to see if a zero sum has been achieved.

At each stage, the element n_{i_j} to be appended to \mathcal{S} is chosen to be any one of the remaining elements of \mathcal{L} that have a sign *opposite* to that of σ_{j-1} , in order to try to drive the partial sum value to zero. Observe that there must always exist such an “opposite-sign” element remaining in \mathcal{L} , since otherwise the sum (hence also the average) of all of the elements in \mathcal{L} would not be zero. Another consequence of this opposite-sign strategy is that the σ_j values can never be equal to either endpoint e_ℓ or e_r of the interval \mathcal{I} ; each σ_j must be one of the $k - 1$ integers in the interior of \mathcal{I} , i.e. in $\tilde{\mathcal{I}}$. To see why this is so, first observe that $\sigma_1 = n_{i_1}$ is in $\tilde{\mathcal{I}}$ by construction. For the passage from σ_j to σ_{j+1} with $j \geq 1$, there are three scenarios:

$$(i) \ e_\ell < \sigma_j < 0, \quad (ii) \ \sigma_j = 0, \quad \text{or} \quad (iii) \ 0 < \sigma_j < e_r.$$

In case (i), σ_{j+1} can be at most e_r larger than σ_j , so $\sigma_{j+1} \in \tilde{\mathcal{I}}$. In case (ii), the construction will cease, and there will be no σ_{j+1} . And in case (iii), σ_{j+1} can be at most $|e_\ell|$ smaller than σ_j , so once again $\sigma_{j+1} \in \tilde{\mathcal{I}}$.

Now carry on the building up of the sublist \mathcal{S} using the “opposite-sign” strategy, until either (a) a partial sum $\sigma_j = 0$ with $j \leq k - 1$ is attained, or (b) the sublist contains $k - 1$ elements with *every* partial sum $\sigma_1, \dots, \sigma_{k-1}$ being nonzero. If (a) occurs, then split off the sublist $\mathcal{S} = \{n_{i_1}, n_{i_2}, \dots, n_{i_j}\}$ from \mathcal{L} , and the remaining sublist $\hat{\mathcal{L}}$ can be homogeneously partitioned by the inductive hypothesis. On the other hand, if (b) occurs, then the $k - 1$ nonzero partial sums must have some repetitions, since there are only $k - 2$

nonzero integers in the open interval $\widetilde{\mathcal{I}}$. So suppose that $\sigma_{i_p} = \sigma_{i_j}$ for some $p < j \leq k - 1$, and let $\widehat{\mathcal{S}} := \{n_{i_{p+1}}, \dots, n_{i_j}\}$, with at most $k - 2$ elements. Observe that the sum of the elements in $\widehat{\mathcal{S}}$ is

$$\sum_{\ell=p+1}^j n_{i_\ell} = \left(\sum_{\ell=1}^j n_{i_\ell} - \sum_{\ell=1}^p n_{i_\ell} \right) = \sigma_{i_j} - \sigma_{i_p} = 0,$$

so splitting off the sublist $\widehat{\mathcal{S}}$ from \mathcal{L} starts the homogeneous partitioning, leaving a remaining sublist $\widehat{\mathcal{L}}$ that can be homogeneously partitioned by the inductive hypothesis.

The result for the general interval $\mathcal{I} = [j, j + k]$ now follows by translation. \square

Remark 3.17. A closer examination of the proof of the Homogeneous Partitioning property indicates that the presence of sublists of “full length” k in a homogeneous partitioning may be somewhat rare. In most of the scenarios for splitting off a sublist \mathcal{S} from the main list \mathcal{L} , the length of the split-off sublist is strictly less than k . In fact, the *only* scenario that can force a sublist to have length k is very special; with all data translated so that $\mu = 0$, *all* of the elements of \mathcal{L} must be at the endpoints e_ℓ and e_r , *and* these endpoints must be relatively prime.

Another feature of the homogeneous partitioning problem that is hinted at in the proof is that often these partitions are *not unique*. Indeed, in the partitioning procedure described in the proof, in particular for case (c), there may be many arbitrary choices that can be made, all of which lead to an admissible partitioning. In addition, there may be other homogeneous partitionings that cannot be generated from the procedure in the proof at all, no matter what choices are made.

Remark 3.18. It is worth noting that Lemma 3.15 can be extended to apply to the situation in which the average $\mu \in \mathcal{I}$ is *not* necessarily an integer. For this general case the statement looks like the following:

General Homogeneous Partitioning Property

Let $\mathcal{I} = [j, j + k]$ be a closed interval with *integer* endpoints and length $k \geq 1$, and consider a list $\mathcal{L} = \{n_1, n_2, \dots, n_m\}$ of m integers, all in \mathcal{I} . Let $\mu \in \mathcal{I}$ denote the average value of all of the entries in \mathcal{L} . Then \mathcal{L} can be partitioned into sublists $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_\ell$ such that the number of entries in each \mathcal{S}_i does not exceed k , and the average value of each sublist is in the closed interval $[\lfloor \mu \rfloor, \lceil \mu \rceil]$.

Note that this more general partitioning result reduces to exactly Lemma 3.15 when the average μ is an integer. However, since we will not need the general property for this paper, we omit the proof.

Example 3.19. Recall that in Example 3.13 we took the diagonal vector of a Smith form $S(\lambda)$, and rearranged the \mathbb{F} -irreducible factors via Lemma 3.11 to obtain a vector $\mathbf{p}(\lambda)$ with average degree 10 and 4-homogeneous degree vector $\deg \mathbf{p} = (12, 12, 11, 9, 8, 8)$. Lemma 3.15 now guarantees that there is a partitioning of $\deg \mathbf{p}$ (and a corresponding partitioning of the entries of \mathbf{p} itself) into sublists of size at most 4, where *each* sublist also has average degree 10. Applying the procedure described in the proof of the Lemma produces the partition $(12, 8 \mid 12, 8 \mid 11, 9)$ for $\deg \mathbf{p}$, and the corresponding rearrangement and partition

$$(p_1(\lambda), p_6(\lambda) \mid p_2(\lambda), p_5(\lambda) \mid p_3(\lambda), p_4(\lambda)) = (\chi^2 \phi \psi^2, \chi \phi \psi^2 \mid \chi^2 \phi \psi^2, \chi \phi \psi^2 \mid \chi^2 \phi \psi, \chi \phi^2 \psi) \quad (3.9)$$

of the entries of $\mathbf{p}(\lambda)$. It is interesting to note that there are several pathways through the partitioning procedure for this example, but all of them lead to the partitioning in (3.9). However, there are two other homogeneous partitionings of $\deg \mathbf{p}$ with sublist size at most $k = 4$, neither of which can be generated by the procedure of the Lemma. They are $(12, 8, 12, 8 \mid 11, 9)$ and $(12, 8 \mid 12, 8, 11, 9)$.

3.4 Un-triangularizing $T(\lambda)$

One final tool is needed in order to complete the construction of our k -quasi-triangular realization. After our realization is initially brought into quasi-triangular form, it may happen that the off-diagonal blocks have increased in degree beyond the target for the final matrix polynomial, since no control of these blocks has even been attempted in the early stages of the construction. Thus we need some method to bring these degrees back within the target range. Furthermore, it is important to do this by unimodular transformations, so that the desired finite spectral structure will not be spoiled. Lemma 3.22, a generalization of Lemma 2.4 from [20], shows how to achieve this goal.

Two preliminaries are needed for Lemma 3.22. First it will be helpful to recall the notion of the “diagonals” of a matrix, and more generally the block-diagonals of a block-partitioned matrix.

Definition 3.20 (Diagonals of a matrix).

The ℓ^{th} -diagonal of a matrix A is the set of entries a_{ij} of A such that $j - i = \ell$. (Note that the 0^{th} -diagonal is conventionally known as the main diagonal of A .) If A is block-partitioned with blocks A_{ij} , then the ℓ^{th} -block-diagonal of A is the collection of blocks A_{ij} with $j - i = \ell$.

A second important background fact concerns the division of matrix polynomials. We recall now this fundamental result for the convenience of the reader.

Lemma 3.21 (Division of matrix polynomials).

Suppose $A(\lambda)$ is any $m \times n$ matrix polynomial over an arbitrary field \mathbb{F} . Furthermore, let $B(\lambda)$ and $C(\lambda)$ be any two strictly regular matrix polynomials over \mathbb{F} , with size $m \times m$ and $n \times n$, respectively. Then we have:

- (a) Left division by B : There exist unique $m \times n$ matrix polynomials $Q_\ell(\lambda)$ and $R_\ell(\lambda)$ such that

$$A(\lambda) = B(\lambda)Q_\ell(\lambda) + R_\ell(\lambda),$$

and either $R_\ell(\lambda) = 0$, or $R_\ell(\lambda)$ is nonzero with $\deg R_\ell < \deg B$. The matrices $Q_\ell(\lambda)$ and $R_\ell(\lambda)$ are called the left quotient and left remainder of $A(\lambda)$, respectively, upon (left) division by $B(\lambda)$.

- (b) Right division by C : There exist unique $m \times n$ matrix polynomials $Q_r(\lambda)$ and $R_r(\lambda)$ such that

$$A(\lambda) = Q_r(\lambda)C(\lambda) + R_r(\lambda),$$

and either $R_r(\lambda) = 0$, or $R_r(\lambda)$ is nonzero with $\deg R_r < \deg C$. The matrices $Q_r(\lambda)$ and $R_r(\lambda)$ are called the right quotient and right remainder of $A(\lambda)$, respectively, upon (right) division by $C(\lambda)$.

Proof. See [7, Ch. 4], [9, Sect. 6.3], or [11, Sect. 7.2]. □

Lemma 3.22 (Degree reduction of off-diagonal blocks).

Let $T(\lambda) \in \mathbb{F}[\lambda]^{n \times n}$ be a block upper triangular matrix polynomial, partitioned into blocks such that all of the s diagonal blocks $T_{ii}(\lambda) \in \mathbb{F}[\lambda]^{n_i \times n_i}$ are strictly regular. Then $T(\lambda)$ is unimodularly equivalent to a block upper triangular matrix polynomial $\tilde{T}(\lambda)$ with exactly the same diagonal blocks as $T(\lambda)$, and with off-diagonal blocks satisfying

$$\deg \tilde{T}_{ij}(\lambda) < \min\{\deg T_{ii}(\lambda), \deg T_{jj}(\lambda)\} \quad (3.10)$$

for $1 \leq i < j \leq s$.

Proof. Let us begin by focusing on a single fixed but arbitrary off-diagonal block $T_{ij}(\lambda)$ with $i < j$, and showing how to reduce its degree by a unimodular transformation so as to satisfy the condition (3.10). It may be that (3.10) is already satisfied; in this case do nothing. Otherwise, suppose first that $\deg T_{ii} \leq \deg T_{jj}$. Then by Lemma 3.21 we can divide $T_{ij}(\lambda)$ by $T_{ii}(\lambda)$ on the left to obtain

$$T_{ij}(\lambda) = T_{ii}(\lambda)Q_{ij}(\lambda) + R_{ij}(\lambda) \text{ with } R_{ij}(\lambda) = 0 \text{ or } \deg R_{ij}(\lambda) < \deg T_{ii}(\lambda),$$

and hence $T_{ij}(\lambda) - T_{ii}(\lambda)Q_{ij}(\lambda) = R_{ij}(\lambda)$. Now define the unimodular matrix

$$V_{ij}(\lambda) := \begin{bmatrix} I_a & & & \\ & I_{n_i} & -Q_{ij}(\lambda) & \\ & & I_b & \\ & & & I_{n_j} \\ & & & & I_c \end{bmatrix}, \text{ where } a = \sum_{m=1}^{i-1} n_m, b = \sum_{m=i+1}^{j-1} n_m, c = \sum_{m=j+1}^s n_m.$$

Multiplying $T(\lambda)$ on the right by $V_{ij}(\lambda)$ has the effect of an elementary block-column operation, which replaces the block $T_{ij}(\lambda)$ by $\tilde{T}_{ij}(\lambda) = R_{ij}(\lambda)$, thus satisfying condition (3.10) at this one location. In addition we see that the diagonal blocks of $T(\lambda)$ remain unchanged by this block-column operation. Indeed, the *only* other blocks of $T(\lambda)$ that may even possibly be affected are the ones directly above the (i, j) -block, i.e., blocks $T_{\ell j}(\lambda)$ with $\ell \leq i$.

On the other hand, if $\deg T_{ii} > \deg T_{jj}$ then something analogous can be done. In this case use Lemma 3.21 to divide $T_{ij}(\lambda)$ by $T_{jj}(\lambda)$ on the right to obtain

$$T_{ij}(\lambda) = \hat{Q}_{ij}(\lambda)T_{jj}(\lambda) + \hat{R}_{ij}(\lambda) \text{ with } \hat{R}_{ij}(\lambda) = 0 \text{ or } \deg \hat{R}_{ij}(\lambda) < \deg T_{jj}(\lambda),$$

and hence $T_{ij}(\lambda) - \widehat{Q}_{ij}(\lambda)T_{jj}(\lambda) = \widehat{R}_{ij}(\lambda)$. Define the unimodular matrix

$$U_{ij}(\lambda) := \begin{bmatrix} I_a & & & \\ & I_{n_i} & & \\ & & I_b & -\widehat{Q}_{ij}(\lambda) \\ & & & I_{n_j} \\ & & & & I_c \end{bmatrix}, \quad \text{where} \quad a = \sum_{m=1}^{i-1} n_m, \quad b = \sum_{m=i+1}^{j-1} n_m, \quad c = \sum_{m=j+1}^s n_m,$$

and multiply $T(\lambda)$ on the left by $U_{ij}(\lambda)$. This has the effect of an elementary block-row operation, which replaces the block $T_{ij}(\lambda)$ by $\widehat{T}_{ij}(\lambda) = \widehat{R}_{ij}(\lambda)$, once again satisfying condition (3.10) at this one location. In the product $U_{ij}(\lambda)T(\lambda)$, the only other blocks of $T(\lambda)$ that can possibly be affected are those to the right of the (i, j) -block, i.e., blocks $T_{i\ell}(\lambda)$ with $\ell \geq j$. Altogether, then, the only possible collateral damage that can be inflicted on $T(\lambda)$ by this degree reduction of the (i, j) -block is to the blocks of $T(\lambda)$ in the L-shaped region marked in (3.11).

$$T(\lambda) = \begin{bmatrix} \diamond & & & & & \\ & \ddots & & & & \\ & & T_{ii} & & & \\ & & & T_{ij} & * & * & * & * & * & * & * & * \\ & & & & \ddots & & & & & & & \\ & & & & & T_{jj} & & & & & & \\ & & & & & & \ddots & & & & & \\ & & & & & & & \ddots & & & & \\ & & & & & & & & \ddots & & & \\ & & & & & & & & & \diamond & & \end{bmatrix} \quad (3.11)$$

It is important to keep this L-shape firmly in mind as we see how to order the block degree reductions so that no individual block's degree reduction spoils a block that has already had its degree reduced.

The key observation here is that whenever a block $T_{ij}(\lambda)$ has its degree reduced by the procedure described above, then all of the other blocks in $T(\lambda)$ that are affected by that reduction lie on “higher” diagonals of $T(\lambda)$ than the diagonal of $T_{ij}(\lambda)$. More precisely, if $T_{ij}(\lambda)$ lies on the ℓ^{th} -block-diagonal of $T(\lambda)$, then all other blocks affected by that degree reduction lie on an m^{th} -block-diagonal with $m > \ell$. Thus a single *sweep-by-diagonals* through the off-diagonal blocks will have the desired effect of achieving condition (3.10) on *all* off-diagonal blocks simultaneously. To do this kind of sweep, first target each of the blocks in the 1^{th} -block-diagonal for degree reduction, in any order. After visiting each block in the 1^{th} -block-diagonal, go next to the 2^{th} -block-diagonal and do a degree reduction on each of these blocks, again in any order. Because of the key observation, none of the degree reductions done in the 1^{th} -block-diagonal will be spoiled by the degree reductions done in the 2^{th} -block-diagonal, and none of the 2^{th} -block-diagonal degree reductions will spoil each other. Continue in this manner, moving up one block-diagonal at a time, until each off-diagonal block has been visited exactly once. At this point, the desired matrix polynomial $\tilde{T}(\lambda)$ will have been attained. \square

We finally have all the tools we will need to “un-triangularize” $T(\lambda)$ into a degree- d , k -quasi-triangular realization of the original list of finite spectral data. For convenience, we recall the statement of Theorem 3.1 here.

Theorem (Quasi-Triangular Realization: Strictly Regular Case).

Suppose a list of m nonconstant monic polynomials $s_1(\lambda), \dots, s_m(\lambda)$ over an arbitrary field \mathbb{F} is given, satisfying the divisibility chain condition $s_1(\lambda) \mid s_2(\lambda) \mid \dots \mid s_m(\lambda)$. Let $\sigma := \sum_{i=1}^m \deg(s_i(\lambda))$, and define k to be the maximum degree among all of the \mathbb{F} -irreducible factors of the polynomials $s_i(\lambda)$ for $i = 1, \dots, m$. Then for any choice of nonzero $d, n \in \mathbb{N}$ such that $n \geq m$ and $dn = \sigma$, there exists an $n \times n$, degree d , strictly regular matrix polynomial $Q(\lambda)$ over \mathbb{F} that is k -quasi-triangular, and has exactly the given polynomials $s_1(\lambda), \dots, s_m(\lambda)$ as its nontrivial invariant polynomials, together with $n - m$ trivial invariant polynomials. In addition, $Q(\lambda)$ can always be chosen so that the degree of every entry in any off-diagonal block of $Q(\lambda)$ is strictly less than d .

Proof. From the given spectral data, begin by constructing the $n \times n$ Smith form

$$S(\lambda) = \text{diag} \left\{ \underbrace{1, \dots, 1}_{n-m}, s_1(\lambda), \dots, s_m(\lambda) \right\}.$$

Now we can use the tools developed in Corollary 3.14, Lemma 3.15, Corollary 3.4, Corollary 2.12, and Lemma 3.22 to build the desired k -quasi-triangular realization of the spectral data contained in $S(\lambda)$, in the following five steps:

- Use $S(\lambda)$ as input to Corollary 3.14 to generate an upper triangular $T(\lambda)$ that is unimodularly equivalent to $S(\lambda)$, and has a diagonal degree vector $\deg(\text{diag } T(\lambda))$ that is k -homogeneous.
- Use the natural vector $\deg(\text{diag } T(\lambda))$, with average value $\mu = d$, as input to Lemma 3.15, and find a homogeneous partitioning of the degrees of the diagonal entries of $T(\lambda)$ into ℓ sublists. The corresponding partitioning of the diagonal entries themselves, with the elements of the entry sublists arranged into ℓ contiguous groups, provides a target for the rearrangement of the entries on the diagonal of $T(\lambda)$.
- Implement this rearrangement of the diagonal entries of $T(\lambda)$ via the triangularity-preserving unimodular transformations of $T(\lambda)$ provided by Corollary 3.4, each of which has the effect of simply performing an interchange of adjacent diagonal entries. (Of course some of the off-diagonal entries are being changed in the process, but we do not try to keep any control of them at this stage of the construction.) Denote the resulting upper triangular matrix by $\hat{T}(\lambda)$, and partition this $\hat{T}(\lambda)$ into blocks

$$\hat{T}(\lambda) = \begin{bmatrix} \hat{T}_{11}(\lambda) & \hat{T}_{12}(\lambda) & \cdots & \hat{T}_{1\ell}(\lambda) \\ & \hat{T}_{22}(\lambda) & \cdots & \hat{T}_{2\ell}(\lambda) \\ & & \ddots & \vdots \\ & & & \hat{T}_{\ell\ell}(\lambda) \end{bmatrix}, \quad (3.12)$$

so that each (upper triangular and square) diagonal block $\hat{T}_{jj}(\lambda)$ has diagonal entries that correspond to the sublist \mathcal{S}_j of the homogeneous partitioning of $\deg(\text{diag } T(\lambda))$. For each j , if we denote the size of the block $\hat{T}_{jj}(\lambda)$ by $n_j \times n_j$, then from the Homogeneous Partitioning property we know that $n_j \leq k$, and the average degree of the diagonal entries of $\hat{T}_{jj}(\lambda)$ is d . So the sum of the degrees of the diagonal entries of $\hat{T}_{jj}(\lambda)$ is dn_j .

- By Corollary 2.12 we know that each *diagonal* block $\hat{T}_{jj}(\lambda)$ is unimodularly equivalent to a strictly regular polynomial matrix $P_{jj}(\lambda)$ of degree d , which is (probably) no longer upper triangular. Let these equivalences be denoted by

$$P_{jj}(\lambda) := \hat{U}_{jj}(\lambda) \hat{T}_{jj}(\lambda) \hat{V}_{jj}(\lambda),$$

where each $\hat{U}_{jj}(\lambda)$ and $\hat{V}_{jj}(\lambda)$ is $n_j \times n_j$ and unimodular. Now define the $n \times n$ block-diagonal unimodular matrices

$$\hat{U}(\lambda) := \text{diag}\{\hat{U}_{11}(\lambda), \hat{U}_{22}(\lambda), \dots, \hat{U}_{\ell\ell}(\lambda)\} \quad \text{and} \quad \hat{V}(\lambda) := \text{diag}\{\hat{V}_{11}(\lambda), \hat{V}_{22}(\lambda), \dots, \hat{V}_{\ell\ell}(\lambda)\},$$

and apply them to the full $\hat{T}(\lambda)$, to get

$$P(\lambda) := \hat{U}(\lambda) \hat{T}(\lambda) \hat{V}(\lambda) = \begin{bmatrix} P_{11}(\lambda) & P_{12}(\lambda) & \cdots & P_{1\ell}(\lambda) \\ & P_{22}(\lambda) & \cdots & P_{2\ell}(\lambda) \\ & & \ddots & \vdots \\ & & & P_{\ell\ell}(\lambda) \end{bmatrix}. \quad (3.13)$$

This matrix $P(\lambda)$ is now k -quasi-triangular, with strictly regular diagonal blocks each of degree d . But $P(\lambda)$ as a whole may not yet be of degree d , because the off-diagonal blocks have not been kept under any control at all.

- The final step brings the degrees of the off-diagonal blocks back under control, while at the same time not disturbing the diagonal blocks in the process. Lemma 3.22 applied to $P(\lambda)$ achieves this, reducing the degree of each off-diagonal block to be strictly less than d , and leaving the diagonal blocks unchanged, to obtain the final desired realization

$$Q(\lambda) = \begin{bmatrix} Q_{11}(\lambda) & Q_{12}(\lambda) & \cdots & Q_{1\ell}(\lambda) \\ & Q_{22}(\lambda) & \cdots & Q_{2\ell}(\lambda) \\ & & \ddots & \vdots \\ & & & Q_{\ell\ell}(\lambda) \end{bmatrix}. \quad (3.14)$$

with $Q_{jj}(\lambda) = P_{jj}(\lambda)$ for $j = 1, \dots, \ell$.

This polynomial matrix $Q(\lambda)$ is k -quasi-triangular, has degree d , and is unimodularly equivalent to the original $S(\lambda)$, and hence has exactly the given finite spectral data. To see that $Q(\lambda)$ is strictly regular, we regard $\text{grade}(Q)$ as being equal to the degree d ; by Lemma 2.7, any other choice will force $Q(\lambda)$ to have nontrivial infinite spectral structure. Since by construction the sum of the degrees of all of the invariant polynomials of $Q(\lambda)$ is dn , the Index Sum Theorem 2.10 immediately shows that the sum of the partial multiplicities of $Q(\lambda)$ at infinity must be zero, and hence that $Q(\lambda)$ is strictly regular. \square

Example 3.23. We now bring the extended illustrative example (started in Example 3.2 and continuing through Examples 3.5, 3.13, and 3.19) to a culmination, using the proof of Theorem 3.1 to complete the construction of a strictly regular, k -quasi-triangular realization (over the field $\mathbb{F} = \mathbb{Z}_2$) of the given finite spectral data from back in Example 3.2. Recall that $k = 4$, the target size is 6×6 with degree 10, and the three irreducible divisors in the original spectral data are $\chi(\lambda) = \lambda^4 + \lambda^3 + 1$, $\phi(\lambda) = \lambda^2 + \lambda + 1$, and $\psi(\lambda) = \lambda$.

In Example 3.19, we found a homogeneous partitioning $(12, 8 \mid 12, 8 \mid 11, 9)$ of diagonal degrees, and corresponding permutation of diagonal entries to give us the target diagonal

$$(\chi^2\phi\psi^2, \chi\phi\psi^2, \chi^2\phi\psi^2, \chi\phi\psi^2, \chi^2\phi\psi, \chi\phi^2\psi) \quad (3.15)$$

for the upper triangular matrix (3.12) in our construction. Now the theory we have developed guarantees that the $S(\lambda)$ in Example 3.2 can be unimodularly transformed into an upper triangular $\hat{T}(\lambda)$ such that $\text{diag } \hat{T}(\lambda)$ is exactly the vector in (3.15). And furthermore, that this transformation can be implemented as a finite sequence of embedded 2×2 unimodular transformations acting only on adjacent diagonal entries. It would be very tedious to display all of these transformations, and the resulting upper triangular $\hat{T}(\lambda)$ is very likely to have a densely populated upper triangular part. So instead, for ease of exposition we exhibit an alternative $\hat{T}(\lambda)$ with the desired diagonal vector that is not only sparse, but is also easily checked to be unimodularly equivalent to $S(\lambda)$.

$$\hat{T}(\lambda) = \left[\begin{array}{cc|cc|cc} \chi^2\phi\psi^2 & \chi\phi\psi & 0 & 0 & 0 & 0 \\ 0 & \chi\phi\psi^2 & 0 & 0 & 0 & 1 \\ \hline & & \chi^2\phi\psi^2 & 0 & 0 & 0 \\ & & 0 & \chi\phi\psi^2 & \phi\psi & 0 \\ \hline & & & & \chi^2\phi\psi & 0 \\ & & & & 0 & \chi\phi^2\psi \end{array} \right]$$

Observe that this $\hat{T}(\lambda)$ has been partitioned into blocks as in (3.12), in a manner that conforms to the partitioning of the diagonal vector (3.9) arising from the homogeneous partitioning of the diagonal degrees.

Each of the diagonal blocks of $\hat{T}(\lambda)$ can now be transformed into 2×2 blocks having degree 10, via simple unimodular transformations. Define the unimodular matrices

$$U_1 := \begin{bmatrix} 1 & \phi+1 \\ 0 & 1 \end{bmatrix}, \quad U_2 := U_1, \quad U_3 := \begin{bmatrix} 1 & \psi \\ 0 & 1 \end{bmatrix}, \quad V_1 := \begin{bmatrix} 1 & 0 \\ \psi^2 & 1 \end{bmatrix}, \quad V_2 := V_1, \quad \text{and} \quad V_3 := U_3^T.$$

Since the underlying field here is $\mathbb{F} = \mathbb{Z}_2$, we then have

$$\begin{aligned} U_1 \hat{T}_{11} V_1 &= \chi\phi\psi \cdot \begin{bmatrix} 1 & \phi+1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \chi\psi & 1 \\ 0 & \psi \end{bmatrix} \begin{bmatrix} 1 & 0 \\ \psi^2 & 1 \end{bmatrix} = \chi\phi\psi \cdot \begin{bmatrix} \phi+1 & \phi\psi+\psi+1 \\ \psi^3 & \psi \end{bmatrix}, \\ U_2 \hat{T}_{22} V_2 &= \chi\phi\psi^2 \cdot \begin{bmatrix} 1 & \phi+1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \chi & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ \psi^2 & 1 \end{bmatrix} = \chi\phi\psi^2 \cdot \begin{bmatrix} 1 & \phi+1 \\ \psi^2 & 1 \end{bmatrix}, \\ U_3 \hat{T}_{33} V_3 &= \chi\phi\psi \cdot \begin{bmatrix} 1 & \psi \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \chi & 0 \\ 0 & \phi \end{bmatrix} \begin{bmatrix} 1 & 0 \\ \psi & 1 \end{bmatrix} = \chi\phi\psi \cdot \begin{bmatrix} \psi^2+1 & \phi\psi \\ \phi\psi & \phi \end{bmatrix}, \end{aligned}$$

which are each readily seen to have degree 10.

Applying these transformations collectively to all of $\hat{T}(\lambda)$ via $U := \text{diag}\{U_1, U_2, U_3\}$ and $V := \text{diag}\{V_1, V_2, V_3\}$ produces the 2-quasi-triangular realization

$$P(\lambda) := U \hat{T}(\lambda) V = \left[\begin{array}{cc|cc|cc} \chi\phi\psi(\phi+1) & \chi\phi\psi(\phi\psi+\psi+1) & 0 & 0 & \psi(\phi+1) & \phi+1 \\ \chi\phi\psi^4 & \chi\phi\psi^2 & 0 & 0 & \psi & 1 \\ \hline & & \chi\phi\psi^2 & \chi\phi\psi^2(\phi+1) & \phi\psi(\phi+1) & 0 \\ & & \chi\phi\psi^4 & \chi\phi\psi^2 & \phi\psi & 0 \\ \hline & & & & \chi\phi\psi(\psi^2+1) & \chi\phi^2\psi^2 \\ & & & & \chi\phi^2\psi^2 & \chi\phi^2\psi \end{array} \right].$$

with diagonal blocks that are all of degree 10. Observe that all of the off-diagonal blocks of $P(\lambda)$ have degree strictly less than 10. Hence we can skip the final step (i.e., using Lemma 3.22 to reduce the degrees of the off-diagonal blocks) of the general procedure, and declare that our final degree 10, strictly regular, 2-quasi-triangular realization $Q(\lambda)$ is identical to $P(\lambda)$ in this example.

Finally, note that there are at least two ways to see that our final realization $Q(\lambda)$ is indeed strictly regular. One is by an index sum argument – taking $\deg Q$ to be equal to $\deg Q = 10$, we see that there is no room in the index sum constraint (2.8) for any infinite partial multiplicities to be nonzero, hence $Q(\lambda)$ must be strictly regular. The second way is simply to examine the leading coefficient of $Q(\lambda)$ as a matrix polynomial, i.e., the matrix coefficient of λ^{10} . That matrix is easily seen to be just $\text{diag}(R, R, R)$, where R is the 2×2 matrix $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$; clearly this coefficient is nonsingular. Hence $Q(\lambda)$ is strictly regular (see comments just after Definition 2.8).

4 Including infinite spectral data in a realization

In this section we extend the range of our quasi-triangular realization construction to handle spectral data arising from a general regular matrix polynomial — i.e., data that may now include nontrivial infinite spectral structure. The main tool that allows us to achieve this extension are Möbius transformations, employing a now well-used technique; see [4], [12], and [19]. In a nutshell, the strategy of this technique is to use an appropriately chosen Möbius transformation to translate a realization problem involving a mixture of finite and infinite spectral data into one that has *only* finite spectral data, solve this strictly regular realization problem, and then translate the solution back to the original spectral data using the inverse Möbius transformation.

However, before we are able to implement this strategy, it will be necessary to reexamine one of the fundamental properties of Möbius transformations from [12], and recast it in a way that enables us to work with these transformations more smoothly in the context of arbitrary fields \mathbb{F} , in particular in the presence of higher degree \mathbb{F} -irreducible divisors in spectral data. This will be the task of Section 4.1. With this reimagined property of Möbius transformations in hand, in Section 4.2 we then solve the quasi-triangular realization problem with nontrivial infinite spectral structure. This gives us the final result needed to now easily prove in Section 4.3 the featured result of the paper, the quasi-triangularization of arbitrary regular matrix polynomials.

Before embarking on this final stage, though, it is worth pointing out several subtle issues that arise in the course of carrying out this strategy. The first concerns an extra hypothesis that is only needed if the field \mathbb{F} is finite, and even then only to exclude very special types of spectral data set. When \mathbb{F} is finite, then it is possible that *every* element of \mathbb{F} , as well as ∞ , may appear in a given spectral data set as an eigenvalue. If that happens, then any possible Möbius transformation that one might attempt to use will simply map the set of eigenvalues bijectively to itself, and thus will not be able to transform the given spectral data into one that has only finite spectral data. To exclude this one problematic scenario, it has been necessary to include an additional hypothesis, i.e., that there must be some element of \mathbb{F} that is *not* an eigenvalue in the given spectral data. Of course for any infinite field \mathbb{F} , this hypothesis always holds, and so has no impact on the range of spectral data sets to which the argument applies. But for finite fields, having this one additional condition satisfied is sufficient to make the rest of the argument work smoothly. (When this condition is violated, it is not known whether quasi-triangular realizations exist or not.)

A second issue concerns the reduction of off-diagonal degrees to be less than the target degree, as was done in the strictly regular case. When there is nontrivial spectral structure at ∞ , such a reduction may not be possible. The difficulty is that this reduction is achieved via unimodular transformations, which may alter the spectral structure at ∞ . Indeed, Example 4.8 in Section 4.2 gives a concrete illustration of how not only the infinite spectral structure, but even the overall degree of the realization itself, can be spoiled by doing the kind of degree reduction of off-diagonal blocks described in Lemma 3.22.

4.1 Möbius transformation of spectral structure over an arbitrary field

We begin this section by reviewing the effect of Möbius transformations on the spectral structure of a matrix polynomial, in the manner developed in [12]. In that paper the emphasis was on scalars in the ambient field \mathbb{F} (plus ∞) as potential eigenvalues, and so the development was best adapted to algebraically closed fields. However, when working with matrix polynomials over arbitrary fields, with irreducible divisors of higher degree, the formulation in [12] can be rather inconvenient, even a bit clumsy to use. Thus our aim in this section is to reformulate the relationship between partial multiplicity sequences and Möbius transformations in such a way that it works smoothly and efficiently over all fields, and is well-adapted to our extended notion of partial multiplicity sequences for irreducible divisors of any degree, as defined in (2.2). We begin, though, with a brief review of the relevant concepts, notation, and results from [12], in particular Theorem 5.3 from that paper.

If $P(\lambda)$ is a matrix polynomial over \mathbb{F} , and $\mu_0 \in \mathbb{F}_\infty$ is any scalar in the “extended” field $\mathbb{F}_\infty := \mathbb{F} \cup \{\infty\}$, then in [12] the partial multiplicity sequence associated with μ_0 was denoted by $\mathcal{J}(P, \mu_0)$. Note that the letter \mathcal{J} was chosen there because the partial multiplicity sequences for P for *all* scalars $\mu_0 \in \mathbb{F}_\infty$ was collectively termed the “Jordan characteristic” of P . In the current paper, the partial multiplicity sequence $\mathcal{J}(P, \mu_0)$ is associated with the degree one irreducible polynomial $\lambda - \mu_0$, rather than with the scalar μ_0 . Thus we have the following equivalence between notations:

$$\mathcal{J}(P, \mu_0) \equiv \mathcal{PM}(P, \lambda - \mu_0). \quad (4.1)$$

Continuing the review of [12] — to any nonsingular $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ over \mathbb{F} there is an associated *Möbius function* $\mathbf{m}_A(\lambda) : \mathbb{F}_\infty \rightarrow \mathbb{F}_\infty$, which is a bijection of scalars given by the formula

$$\mathbf{m}_A(\lambda) := \frac{a\lambda + b}{c\lambda + d},$$

and a Möbius transformation on matrix polynomials of grade g given by

$$\mathbf{M}_A(P(\lambda)) := (c\lambda + d)^g P(\mathbf{m}_A(\lambda)).$$

Note that this is exactly the same as in Definition 2.13. Now a key result from [12] establishes a simple relationship between the Jordan characteristics of P and $\mathbf{M}_A(P)$.

Theorem 4.1 (Theorem 5.3 from [12]).

Suppose $P(\lambda)$ is an $m \times n$ matrix polynomial of grade g over an arbitrary field \mathbb{F} , and $A \in GL(2, \mathbb{F})$ is a nonsingular matrix with associated Möbius function $\mathbf{m}_A(\lambda)$ and Möbius transformation \mathbf{M}_A . Then for any $\mu_0 \in \mathbb{F}_\infty$,

$$\mathcal{J}(\mathbf{M}_A(P), \mu_0) = \mathcal{J}(P, \mathbf{m}_A(\mu_0)), \quad (4.2)$$

or equivalently,

$$\mathcal{J}(\mathbf{M}_A(P), \mathbf{m}_{A^{-1}}(\mu_0)) = \mathcal{J}(P, \mu_0). \quad (4.3)$$

The formulas in Theorem 4.1 have one glaring drawback when trying to capture the effect of a Möbius transformation on the complete spectral structure of a matrix polynomial P over a field \mathbb{F} that is *not* algebraically closed. The problem is that the irreducible divisors of P may all have degree 2 or larger, and thus have *no* eigenvalues in \mathbb{F}_∞ at all. In this scenario the formulas in Theorem 4.1 tell us nothing at all about the spectral effects of the Möbius transformation. At least not directly. The only way to recover any information about the spectral effects of \mathbf{M}_A would be to pass to the algebraic closure $\overline{\mathbb{F}}$, which though feasible, may become an unwelcome annoyance. It would be very useful to instead have some analog of Theorem 4.1 that applies directly to partial multiplicity sequences associated with irreducible divisors of any degree, i.e., to the $\mathcal{PM}(P, \chi)$ defined earlier in (2.2). There is indeed such an analog, and the goal of the remainder of this section is to establish this extension of Theorem 4.1. We claim that the following relationship holds:

$$\mathcal{PM}(\mathbf{M}_A(P), \mathbf{M}_A(\chi)) = \mathcal{PM}(P, \chi). \quad (4.4)$$

To properly interpret this formula, however, two conventions must be observed:

- On the left-hand side of (4.4), $\mathbf{M}_A(P)$ is taken with respect to the specified grade for P , but $\mathbf{M}_A(\chi)$ should always be taken with grade χ equal to $\deg \chi$, with just the one exception described next.
- If $\chi(\lambda) = \beta$ is any *nonzero* constant, then χ is to be regarded as a *grade one* polynomial, i.e., as $\alpha\lambda + \beta$ with $\alpha = 0$.

Note that the naive intuition underlying this second convention is that the “root” of $0\lambda + \beta$ is $\lambda = -\beta/0 = \infty$, and so $0\lambda + \beta$ can play the role of a “polynomial stand-in” for an eigenvalue at ∞ . Thus we will now use the notation $\mathcal{PM}(P, 0\lambda + \beta)$ to replace the earlier temporary notation $\mathcal{PM}(P, \infty)$ for the partial multiplicity sequence associated with an eigenvalue at ∞ . We will see in Theorem 4.6 and its proof that this now makes the formula (4.4) internally consistent and universally applicable to all partial multiplicity sequences, both for finite and for infinite spectral structure. Some additional motivation and justification for this (perhaps unexpected?) second convention is given in Remark 4.2.

Remark 4.2. Let \mathcal{G}_1 denote the set of all *nonzero* grade one scalar (i.e., 1×1) polynomials $L(\lambda) = \alpha\lambda + \beta$ over \mathbb{F} . It is well known [12] that any Möbius transformation \mathbf{M}_A defines a bijection $\mathcal{G}_1 \rightarrow \mathcal{G}_1$, and the associated Möbius

function \mathbf{m}_A defines a bijection $\mathbb{F}_\infty \rightarrow \mathbb{F}_\infty$ on the extended field of scalars \mathbb{F}_∞ . The purpose of this remark is to briefly explore the parallelism between the bijection \mathbf{M}_A and the bijection $\mathbf{m}_{A^{-1}}$ (not \mathbf{m}_A).

What underlies this parallelism is the simple calculation of $\mathbf{M}_A(\lambda - \mu_0)$, which also plays a central role in the proof of Theorem 4.6. Using Definition 2.13, it is not hard to see that $\mathbf{M}_A(\lambda - \mu_0)$ is (usually) a nonzero scalar multiple of $\lambda - \mathbf{m}_{A^{-1}}(\mu_0)$, with the one possible exception of being just a nonzero scalar. Now recall that in the context of discussing spectral data, we are associating the degree one \mathbb{F} -irreducible $(\lambda - \mu_0)$ with the eigenvalue μ_0 . So we see that the mapping of degree one irreducibles by \mathbf{M}_A is mirrored by the mapping of eigenvalues in \mathbb{F} by $\mathbf{m}_{A^{-1}}$, at least for most μ_0 .

To make this parallelism more precise, let us define a map $\rho : \mathcal{G}_1 \rightarrow \mathbb{F}_\infty$ that sends $\alpha\lambda + \beta \in \mathcal{G}_1$ with $\alpha \neq 0$ to its root $-\frac{\beta}{\alpha} \in \mathbb{F}$, i.e., define $\rho(\alpha\lambda + \beta) := -\frac{\beta}{\alpha}$. This specifies ρ on most of the desired domain \mathcal{G}_1 , and enables us to express the parallelism in the diagram

$$\begin{array}{ccc} \mathcal{G}_1 & \xrightarrow{\mathbf{M}_A} & \mathcal{G}_1 \\ \rho \downarrow & & \downarrow \rho \\ \mathbb{F}_\infty & \xrightarrow{\mathbf{m}_{A^{-1}}} & \mathbb{F}_\infty \end{array} \quad . \quad (4.5)$$

For all elements of \mathcal{G}_1 with $\alpha \neq 0$ (with at most one exception), we see that the diagram (4.5) commutes. But what about elements of \mathcal{G}_1 with $\alpha = 0$ (and hence $\beta \neq 0$), and the one possible exceptional case when $\mathbf{M}_A(\lambda - \mu_0)$ is just a nonzero scalar? Is there a “natural” way to complete the definition of the map ρ to the whole domain \mathcal{G}_1 , and to do so in such a way that (4.5) commutes? There is indeed a unique way to achieve this, and that is to define $\rho(0\lambda + \beta) := \infty \in \mathbb{F}_\infty$, for any $\beta \neq 0$. Thus we see from another direction why it is natural to associate the grade 1 polynomial $0\lambda + \beta$ with an eigenvalue at ∞ . Note that the commuting of the diagram (4.5) holds for any $A \in GL(2, \mathbb{F})$ and its associated Möbius transformation and function.

The simplest (and perhaps most revealing) example of (4.5) uses \mathbf{M}_R with $R = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ i.e., $\mathbf{M}_R = \text{rev}_1$. In this case with $\alpha \neq 0$ and $\beta \neq 0$ we have $\mathbf{M}_R(\alpha\lambda + \beta) = \beta\lambda + \alpha$ with reciprocal roots $-\frac{\beta}{\alpha}$ and $-\frac{\alpha}{\beta}$, respectively, and also $\mathbf{m}_{R^{-1}}(\mu_0) = \mathbf{m}_R(\mu_0) = \frac{1}{\mu_0}$ is the reciprocal map. So clearly everything in the diagram commutes when α and β are *both* nonzero. But if $\alpha = 0$, then with our convention we have $\mathbf{M}_R(0\lambda + \beta) = \beta\lambda$, with roots ∞ and 0 , which are commonly regarded in the literature as being “reciprocal”. Thus we see directly when $A = R$ that (4.5) commutes for all elements of \mathcal{G}_1 .

It is also useful to note that (4.5) still commutes if we pass to the equivalence classes discussed earlier in Remark 2.2. That is, let us declare that two \mathbb{F} -irreducible polynomials are equivalent if one is a nonzero scalar multiple of the other, i.e., $\chi \sim c\chi$ for any nonzero scalar $c \in \mathbb{F}$. Letting $\tilde{\mathcal{G}}_1$ denote the resulting set of equivalence classes of grade one polynomials over \mathbb{F} , then both of the maps \mathbf{M}_A and ρ respect these equivalence classes, and so induce well-defined quotient maps $\tilde{\mathbf{M}}_A$ and $\tilde{\rho}$, which are both bijections. This, then, gives us the following commutative diagram in which all of the mappings are bijective.

$$\begin{array}{ccc} \tilde{\mathcal{G}}_1 & \xrightarrow{\tilde{\mathbf{M}}_A} & \tilde{\mathcal{G}}_1 \\ \tilde{\rho} \downarrow & & \downarrow \tilde{\rho} \\ \mathbb{F}_\infty & \xrightarrow{\mathbf{m}_{A^{-1}}} & \mathbb{F}_\infty \end{array} \quad (4.6)$$

Finally, note that with our conventions, the scenario $\mathbf{M}_A(\chi) = \beta$ *cannot* arise in (4.4) if the \mathbb{F} -irreducible χ has $\deg(\chi) \geq 2$; it can only occur if $\chi \in \mathcal{G}_1$. This is not obvious, but follows from a property to be proved in Lemma 4.4, namely that for any \mathbb{F} -irreducible χ with $\deg(\chi) \geq 2$, and $\mathbf{M}_A(\chi)$ taken with respect to degree, then $\deg \mathbf{M}_A(\chi) = \deg(\chi)$.

Before proceeding to prove (4.4) in Theorem 4.6, we first establish some preliminary lemmas. The first of these lemmas appeared in [12] as Corollary 3.24a, but we recall it here together with its simple proof for the convenience of the reader.

Lemma 4.3 (Basic product property of Möbius transformations).

Let $p, q \in \mathbb{F}[\lambda]$ be nonzero scalar polynomials, and let $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in GL(2, \mathbb{F})$. Letting the grades of p , q , and pq all be chosen to be equal to their degrees, then we have $\mathbf{M}_A(pq) = \mathbf{M}_A(p)\mathbf{M}_A(q)$.

Proof. Let $d_1 = \deg p$ and $d_2 = \deg q$. Then from Definition 2.13 we have

$$\begin{aligned} \mathbf{M}_A(pq) &= (c\lambda + d)^{d_1+d_2} p(\mathbf{m}_A(\lambda)) q(\mathbf{m}_A(\lambda)) \\ &= (c\lambda + d)^{d_1} p(\mathbf{m}_A(\lambda)) \cdot (c\lambda + d)^{d_2} q(\mathbf{m}_A(\lambda)) = \mathbf{M}_A(p) \mathbf{M}_A(q). \end{aligned}$$

□

The next result was mentioned in Section 2 under the name Lemma 2.15. We recall it here and provide the proof postponed from earlier.

Lemma 4.4. *Suppose \mathbb{F} is an arbitrary field, and $\chi(\lambda)$ is any \mathbb{F} -irreducible scalar polynomial with $\deg \chi \geq 2$. Let \mathbf{M}_A be the Möbius transformation associated with any $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in GL(2, \mathbb{F})$. Then with $\text{grade}(\chi)$ taken to be equal to $\deg(\chi)$, the transformation \mathbf{M}_A preserves both the degree and the \mathbb{F} -irreducibility of χ . That is, $\mathbf{M}_A(\chi)$ is \mathbb{F} -irreducible, and $\deg(\mathbf{M}_A(\chi)) = \deg(\chi)$.*

Proof. That \mathbf{M}_A preserves \mathbb{F} -irreducibility was shown in [12, Corollary 3.24]. To see why degree is preserved, we look at $\chi(\lambda)$ over the algebraic closure $\overline{\mathbb{F}}$ of \mathbb{F} . Over $\overline{\mathbb{F}}$ we may factor χ completely into linear factors, i.e., $\chi(\lambda) = k(\lambda - r_1)(\lambda - r_2) \cdots (\lambda - r_d)$, where $d = \deg \chi \geq 2$ and $k \in \mathbb{F}$. Now observe that for χ to be \mathbb{F} -irreducible, all of these roots r_i must be in $\overline{\mathbb{F}} \setminus \mathbb{F}$, otherwise χ would be \mathbb{F} -reducible. Computing $\mathbf{M}_A(\chi)$ with grade equal to degree d , and using the multiplicative property of Möbius transformations in Lemma 4.3, we obtain

$$\mathbf{M}_A(\chi) = k \mathbf{M}_A(\lambda - r_1) \cdot \mathbf{M}_A(\lambda - r_2) \cdots \mathbf{M}_A(\lambda - r_d),$$

where each of these transformations $\mathbf{M}_A(\lambda - r_j)$ is computed with respect to grade 1. Using the definition, we easily see that

$$\mathbf{M}_A(\lambda - r_j) = (a - cr_j)\lambda + (b - dr_j). \quad (4.7)$$

Now consider two cases: $c = 0$ and $c \neq 0$. If $c = 0$, then since A is nonsingular we must have $a \neq 0$, so each of the factors $\mathbf{M}_A(\lambda - r_j)$ will have degree one, and $\mathbf{M}_A(\chi)$ will have degree $d = \deg \chi$. On the other hand, if $c \neq 0$, then once again each of these $\mathbf{M}_A(\lambda - r_j)$ will have degree one, and hence $\deg(\mathbf{M}_A(\chi)) = d = \deg \chi$, unless one of the roots r_j satisfies $a - cr_j = 0$, i.e., $r_j = \frac{a}{c}$. This, however, is impossible, since $\frac{a}{c} \in \mathbb{F}$, but $r_j \in \overline{\mathbb{F}} \setminus \mathbb{F}$. □

Lemma 4.5. *Suppose \mathbb{F} is an arbitrary field, and $\widehat{\mathbb{F}} \supseteq \mathbb{F}$ is any field extension of \mathbb{F} . Then polynomials $p, q \in \mathbb{F}[\lambda]$ are coprime in $\mathbb{F}[\lambda]$ if and only if they are coprime in $\widehat{\mathbb{F}}[\lambda]$. More specifically, if $p, q \in \mathbb{F}[\lambda]$ are distinct \mathbb{F} -irreducible polynomials, then as polynomials over the algebraic closure $\overline{\mathbb{F}}$, p and q have no roots in common.*

Proof. Suppose p and q are coprime in $\mathbb{F}[\lambda]$. Then p and q satisfy the Bezout identity $pr + qt \equiv 1$, for some polynomials $r, t \in \mathbb{F}[\lambda]$. This identity still holds for p, q, r, t viewed as polynomials over $\widehat{\mathbb{F}}$, so p and q are coprime in $\widehat{\mathbb{F}}[\lambda]$. Conversely, if p and q are coprime in $\widehat{\mathbb{F}}[\lambda]$, then they have no nontrivial common factor in $\widehat{\mathbb{F}}[\lambda]$, and hence no nontrivial common factor in the smaller ring $\mathbb{F}[\lambda]$, and thus are coprime in $\mathbb{F}[\lambda]$. The statement about \mathbb{F} -irreducible polynomials p and q now follows immediately, since *distinct* \mathbb{F} -irreducibles are necessarily coprime in $\mathbb{F}[\lambda]$. □

We are now in a position to prove the main result of this section, i.e., the formula that describes the effect of Möbius transformations on spectral structure. It is important to emphasize that this formula holds in complete generality for *all* matrix polynomials (regular or singular, of any size, over arbitrary fields), for *all* irreducible divisors (of any degree), and for *all* Möbius transformations \mathbf{M}_A , whatever the underlying (nonsingular) matrix A . Note also that the proof makes extensive use of the previously-known formulas in Theorem 4.1.

Theorem 4.6 (Effect of Möbius on spectral structure over arbitrary fields).

Let $P(\lambda)$ be any grade g matrix polynomial over \mathbb{F} , where \mathbb{F} is an arbitrary field. Also let $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in GL(2, \mathbb{F})$ be any nonsingular matrix over \mathbb{F} , with associated Möbius transformation \mathbf{M}_A . Then for any \mathbb{F} -irreducible scalar polynomial $\chi(\lambda)$, including the grade one polynomial $(0\lambda + \beta)$ with $\beta \neq 0$, we have

$$\mathcal{PM}(\mathbf{M}_A(P), \mathbf{M}_A(\chi)) = \mathcal{PM}(P, \chi). \quad (4.8)$$

(Here $\mathbf{M}_A(P)$ is taken with respect to grade g , while each $\mathbf{M}_A(\chi)$ is taken with grade equal to $\deg \chi$, with the sole exception of the grade one $\chi(\lambda) = 0\lambda + \beta$, with $\beta \neq 0$.)

Proof. We consider three cases, depending on the nature of the \mathbb{F} -irreducible $\chi(\lambda)$ involved. In all cases we have $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in GL(2, \mathbb{F})$, and $A^{-1} = \frac{1}{\det A} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$. In the following calculations, keep in mind that we are free (because of (2.3)) to alter \mathbb{F} -irreducibles by (nonzero) scalar multiples, whenever convenient. We will also make repeated use of the formula (4.3) from Theorem 4.1, as well as the notation conversion formula (4.1), without ever explicitly mentioning that they are being used.

(a) $[\chi = 0\lambda + \beta, \text{ with } \beta \neq 0 \text{ and } \text{grade}(\chi) = 1]$ Given the above conventions, we have

$$\begin{aligned} \mathcal{PM}(P, \chi) &= \mathcal{J}(P, \infty) = \mathcal{J}(\mathbf{M}_A(P), \mathbf{m}_{A^{-1}}(\infty)) \\ &= \mathcal{J}(\mathbf{M}_A(P), \frac{d}{-c}) = \mathcal{PM}(\mathbf{M}_A(P), c\lambda + d) = \mathcal{PM}(\mathbf{M}_A(P), \mathbf{M}_A(\chi)), \end{aligned}$$

as desired.

(b) $[\chi = \lambda - \mu_0, \text{ with } \mu_0 \in \mathbb{F} \text{ and } \text{grade}(\chi) = \text{deg}(\chi) = 1]$

Recall from (4.7) that $\mathbf{M}_A(\lambda - \mu_0) = (a - c\mu_0)\lambda + (b - d\mu_0)$. Splitting this into two cases, depending on whether $a - c\mu_0$ is zero or not, we have

$$\mathbf{M}_A(\chi) = \begin{cases} (a - c\mu_0)[\lambda - \mathbf{m}_{A^{-1}}(\mu_0)] & \text{if } a - c\mu_0 \neq 0, \\ (b - d\mu_0) & \text{if } a - c\mu_0 = 0. \end{cases}$$

So if $a - c\mu_0 \neq 0$, then

$$\begin{aligned} \mathcal{PM}(P, \chi) &= \mathcal{J}(P, \mu_0) = \mathcal{J}(\mathbf{M}_A(P), \mathbf{m}_{A^{-1}}(\mu_0)) \\ &= \mathcal{PM}(\mathbf{M}_A(P), \lambda - \mathbf{m}_{A^{-1}}(\mu_0)) = \mathcal{PM}(\mathbf{M}_A(P), \mathbf{M}_A(\chi)). \end{aligned}$$

On the other hand, if $a - c\mu_0 = 0$ (i.e., $\mu_0 = \frac{a}{c} = \mathbf{m}_A(\infty)$), then $b - d\mu_0 \neq 0$ because of the nonsingularity of A , and we have

$$\begin{aligned} \mathcal{PM}(P, \chi) &= \mathcal{J}(P, \mu_0) = \mathcal{J}(\mathbf{M}_A(P), \mathbf{m}_{A^{-1}}(\mu_0)) \\ &= \mathcal{J}(\mathbf{M}_A(P), \infty) = \mathcal{PM}(\mathbf{M}_A(P), 0\lambda + (b - d\mu_0)) = \mathcal{PM}(\mathbf{M}_A(P), \mathbf{M}_A(\chi)). \end{aligned}$$

(c) $[\text{grade}(\chi) = \text{deg}(\chi) \geq 2]$

In this argument we will go back and forth between \mathbb{F} and its algebraic closure $\overline{\mathbb{F}}$, exploiting the fact that P , \mathbf{M}_A , and χ can be viewed either as objects over \mathbb{F} or as objects over $\overline{\mathbb{F}}$.

Viewing $\chi(\lambda)$ as a polynomial in $\overline{\mathbb{F}}[\lambda]$, we may factor it completely into linear factors

$$\chi(\lambda) = \beta(\lambda - r_1)^{\delta_1}(\lambda - r_2)^{\delta_2} \cdots (\lambda - r_k)^{\delta_k}, \quad 0 \neq \beta \in \mathbb{F}, \quad (4.9)$$

where the roots r_i are *distinct*, and all in $\overline{\mathbb{F}} \setminus \mathbb{F}$. (See proof of Lemma 4.4.) Denoting the j^{th} partial multiplicity of χ in P by $\alpha_j := [\mathcal{PM}(P, \chi)]_j$, so that the j^{th} invariant polynomial s_j of P can be expressed as

$$s_j(\lambda) = \chi^{\alpha_j} \tilde{s}_j(\lambda) \quad (4.10)$$

with χ coprime to \tilde{s}_j , we can now easily see that over $\overline{\mathbb{F}}$ we have $[\mathcal{PM}(P, \lambda - r_i)]_j = \delta_i \alpha_j$, for each $1 \leq i \leq k$. This follows from (4.9) and (4.10), together with the fact that *no additional copies* of $(\lambda - r_i)$ can arise from \tilde{s}_j , by Lemma 4.5. Applying the result of part (b) of this proof (over the field $\overline{\mathbb{F}}$), we now have that

$$[\mathcal{PM}(\mathbf{M}_A(P), \mathbf{M}_A(\lambda - r_i))]_j = [\mathcal{PM}(P, \lambda - r_i)]_j = \delta_i \alpha_j, \quad (4.11)$$

for all $1 \leq i \leq k$ and $1 \leq j \leq \text{rank } P$. Note also from part(b) that $\mathbf{M}_A(\lambda - r_i) = (a - cr_i)[\lambda - \mathbf{m}_{A^{-1}}(r_i)]$ is always degree one, since $a - cr_i \neq 0$ ($a, c \in \mathbb{F}$, but $r_i \in \overline{\mathbb{F}} \setminus \mathbb{F}$). (Although it is not strictly needed for this argument, it is also helpful to keep in mind that the transformation \mathbf{M}_A is *bijective* on the set of grade one polynomials over $\overline{\mathbb{F}}$, so that $\mathbf{M}_A(\lambda - r_i)$ is distinct from any other $\mathbf{M}_A(\lambda - \mu_0)$.)

Gathering together all of the partial multiplicities in (4.11), we now see that the j^{th} invariant polynomial of $\mathbf{M}_A(P)$ can be written, over $\overline{\mathbb{F}}$, as

$$\begin{aligned} m_j(\lambda) &= [\mathbf{M}_A(\lambda - r_1)]^{\delta_1 \alpha_j} [\mathbf{M}_A(\lambda - r_2)]^{\delta_2 \alpha_j} \cdots [\mathbf{M}_A(\lambda - r_k)]^{\delta_k \alpha_j} \cdot \hat{m}_j(\lambda) \\ &= \left[[\mathbf{M}_A(\lambda - r_1)]^{\delta_1} [\mathbf{M}_A(\lambda - r_2)]^{\delta_2} \cdots [\mathbf{M}_A(\lambda - r_k)]^{\delta_k} \right]^{\alpha_j} \cdot \hat{m}_j(\lambda), \end{aligned}$$

where $\hat{m}_j(\lambda)$ is coprime in $\overline{\mathbb{F}}[\lambda]$ to each of the degree one factors $\mathbf{M}_A(\lambda - r_i)$. But by the product property of Möbius transformations in Lemma 4.3 we also have

$$\begin{aligned} &[\mathbf{M}_A(\lambda - r_1)]^{\delta_1} [\mathbf{M}_A(\lambda - r_2)]^{\delta_2} \cdots [\mathbf{M}_A(\lambda - r_k)]^{\delta_k} \\ &= \mathbf{M}_A[(\lambda - r_1)^{\delta_1}] \mathbf{M}_A[(\lambda - r_2)^{\delta_2}] \cdots \mathbf{M}_A[(\lambda - r_k)^{\delta_k}] \\ &= \mathbf{M}_A[(\lambda - r_1)^{\delta_1} (\lambda - r_2)^{\delta_2} \cdots (\lambda - r_k)^{\delta_k}] = \mathbf{M}_A[\chi(\lambda)] / \beta. \end{aligned}$$

So the invariant polynomial $m_j(\lambda)$ for $\mathbf{M}_A(P)$, which is the same over \mathbb{F} or over $\overline{\mathbb{F}}$ (since Smith forms are invariant under field extensions), can be expressed as

$$m_j(\lambda) = \left[\mathbf{M}_A[\chi(\lambda)] \right]^{\alpha_j} \tilde{m}_j(\lambda),$$

where $\mathbf{M}_A[\chi(\lambda)]$ and $\tilde{m}_j(\lambda)$ are coprime in $\overline{\mathbb{F}}[\lambda]$. Since $\mathbf{M}_A[\chi(\lambda)]$, and hence also $\tilde{m}_j(\lambda)$, is a polynomial over \mathbb{F} , by Lemma 4.5 they are also coprime in $\mathbb{F}[\lambda]$. Also we know from Lemma 4.4 and the discussion surrounding it that $\mathbf{M}_A[\chi(\lambda)]$ is an \mathbb{F} -irreducible polynomial of the same degree as $\chi(\lambda)$. Thus it is well-defined to speak of the partial multiplicities of $\mathbf{M}_A(P)$ at $\mathbf{M}_A(\chi)$, and we see that over \mathbb{F} we have

$$[\mathcal{PM}(\mathbf{M}_A(P), \mathbf{M}_A(\chi))]_j = \alpha_j = [\mathcal{PM}(P, \chi)]_j$$

for each $1 \leq j \leq \text{rank } P$, i.e., for each individual partial multiplicity. Hence we also have the equality

$$\mathcal{PM}(\mathbf{M}_A(P), \mathbf{M}_A(\chi)) = \mathcal{PM}(P, \chi)$$

for the whole partial multiplicity sequence, as desired. □

4.2 Quasi-triangular realization with infinite spectral structure

With the new tools for working with Möbius transformations developed in Section 4.1, we are now able to show how to construct a quasi-triangular realization for spectral data that may now include nontrivial structure at ∞ . An example to illustrate this construction follows immediately after the theorem.

Theorem 4.7 (Quasi-Triangular Realization with Eigenvalue at ∞).

Let \mathbb{F} be an arbitrary field, and suppose a list of m invariant polynomials $s_1(\lambda), \dots, s_m(\lambda)$ over \mathbb{F} forming a divisibility chain $s_1(\lambda) | s_2(\lambda) | \dots | s_m(\lambda)$ is given, together with a nonempty list of ℓ nonzero partial multiplicities at infinity $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_\ell$. Let

$$\sigma := \sum_{i=1}^m \deg(s_i(\lambda)) + \sum_{j=1}^{\ell} \alpha_j$$

be the index sum for this data, and define k to be the maximum degree among all of the \mathbb{F} -irreducible divisors of $s_m(\lambda)$. Suppose also that the field \mathbb{F} contains some scalar $\omega \in \mathbb{F}$ such that $s_m(\omega) \neq 0$. Then for any choice of nonzero $g, n \in \mathbb{N}$ such that $n \geq \max\{m, \ell\}$ and $gn = \sigma$, there exists an $n \times n$, grade g matrix polynomial $Q(\lambda)$ over \mathbb{F} that is k -quasi-triangular, has exactly the given invariant polynomials $s_1(\lambda), \dots, s_m(\lambda)$ with all other $n - m$ invariant polynomials equal to 1, and partial multiplicity sequence at infinity $(0, \dots, 0, \alpha_1, \dots, \alpha_\ell)$. In addition, $\deg Q(\lambda) = g$ if and only if $n > \ell$.

Proof. Begin by taking the given spectral data and converting it into the “first form” described in Definition 2.4, i.e., into a list of \mathbb{F} -irreducible divisors χ_j (for $j = 1, t$), each equipped with a partial multiplicity sequence $\mathcal{PM}_{\text{given}}(\chi_j)$ of length n . Recall that the partial multiplicities at ∞ are recorded as $\mathcal{PM}_{\text{given}}(\beta)$ for some (any) nonzero $\beta \in \mathbb{F}$. Of course, doing this conversion may require adjoining some additional initial trivial invariant polynomials, and also perhaps some additional initial zero partial multiplicities at ∞ , as needed to fill out the length n .

Next we design some Möbius transformations \mathbf{M}_A that will interchange the spectral roles of ∞ and the “special” scalar $\omega \in \mathbb{F}$ that has been assumed to exist. (The “special” property of ω is that $\mathcal{PM}(\lambda - \omega) = (0, 0, \dots, 0)$.) We claim that the \mathbf{M}_A defined by the nonsingular matrix

$$A = \begin{cases} \begin{bmatrix} \omega & 0 \\ 1 & -\omega \end{bmatrix} & \text{if } \omega \neq 0, \\ \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} & \text{if } \omega = 0, \end{cases}$$

achieves this goal. As a mapping on grade one scalar polynomials, this \mathbf{M}_A is a bijection, and it is straightforward to check that

$$\mathbf{M}_A(\lambda - \omega) = \begin{cases} \omega^2 & \text{if } \omega \neq 0 \\ 1 & \text{if } \omega = 0 \end{cases} \quad \text{and} \quad \mathbf{M}_A(\beta) = \begin{cases} \beta(\lambda - \omega) & \text{if } \omega \neq 0 \\ \beta\lambda = \beta(\lambda - \omega) & \text{if } \omega = 0 \end{cases},$$

for any nonzero $\beta \in \mathbb{F}$. Thus we see that for any ω , the transformation \mathbf{M}_A effects an interchange of the roles of ω and ∞ as eigenvalues, while all other finite scalars $\mu \in \mathbb{F}$ (represented by degree one polynomials $\lambda - \mu$) are simply permuted in some fashion, the details of which are irrelevant here.

Using this \mathbf{M}_A , we now define a new collection of spectral data that has no eigenvalue at ∞ . For each irreducible divisor χ_j in the original data, we replace it by $\mathbf{M}_A(\chi_j)$, but assign to it the same partial multiplicity sequence that χ_j has in the given spectral data. (Here each $\mathbf{M}_A(\chi_j)$ is taken with grade equal to degree.) In other words, we declare that

$$\mathcal{PM}(\mathbf{M}_A(\chi_j)) := \mathcal{PM}_{\text{given}}(\chi_j).$$

By Lemma 4.4 we know that each $\mathbf{M}_A(\chi_j)$ is \mathbb{F} -irreducible with $\deg \mathbf{M}_A(\chi_j) = \deg \chi_j$.

Similarly, for the eigenvalue at ∞ in the original given data, we replace it by $\mathbf{M}_A(0\lambda + 1) = \lambda - \omega$, and assign to it the partial multiplicity sequence that ∞ has in the given spectral data. (Here $0\lambda + 1$ is viewed as a grade one polynomial.) In other words, we declare that

$$\mathcal{PM}(\lambda - \omega) = \mathcal{PM}(\mathbf{M}_A(0\lambda + 1)) := \mathcal{PM}_{\text{given}}(0\lambda + 1).$$

Thus we have specified a new collection \mathcal{C} of purely finite spectral data, that is, a list of \mathbb{F} -irreducible divisors $\{\mathbf{M}_A(\chi_1), \mathbf{M}_A(\chi_2), \dots, \mathbf{M}_A(\chi_t), \mathbf{M}_A(0\lambda + 1) = (\lambda - \omega)\}$, together with assigned partial multiplicity sequences for each. Since the irreducible divisors in \mathcal{C} have the same degrees as their partners from the original spectral data (the same grade in the case of partners $0\lambda + 1$ and $\lambda - \omega$), as well as the same partial multiplicities, then the index sum σ is the same for \mathcal{C} as it was for the original data. Thus we may use the same values of n , g , and k for \mathcal{C} as was used for the original data.

Now by Theorem 3.1 there exists an $n \times n$, degree g matrix polynomial $P(\lambda)$ over \mathbb{F} that is k -quasi-triangular, and has exactly the spectral data in \mathcal{C} . In other words,

$$\mathcal{PM}(P(\lambda), \mathbf{M}_A(\chi_j)) = \mathcal{PM}_{\text{given}}(\chi_j), \quad (4.12)$$

$$\text{and} \quad \mathcal{PM}(P(\lambda), \mathbf{M}_A(0\lambda + 1)) = \mathcal{PM}_{\text{given}}(0\lambda + 1). \quad (4.13)$$

Applying Theorem 4.6 using the Möbius transformation $\mathbf{M}_{A^{-1}}$ (with respect to grade g) to (4.12) and (4.13), and defining $Q(\lambda) := \mathbf{M}_{A^{-1}}(P)$, we have

$$\begin{aligned} \mathcal{PM}(Q(\lambda), \chi_j) &= \mathcal{PM}(\mathbf{M}_{A^{-1}}(P), \chi_j) = \mathcal{PM}(\mathbf{M}_{A^{-1}}(P), \mathbf{M}_{A^{-1}}(\mathbf{M}_A(\chi_j))) \\ &= \mathcal{PM}(P(\lambda), \mathbf{M}_A(\chi_j)) = \mathcal{PM}_{\text{given}}(\chi_j), \end{aligned}$$

and

$$\begin{aligned} \mathcal{PM}(Q(\lambda), (0\lambda + 1)) &= \mathcal{PM}(\mathbf{M}_{A^{-1}}(P), (0\lambda + 1)) = \mathcal{PM}(\mathbf{M}_{A^{-1}}(P), \mathbf{M}_{A^{-1}}(\mathbf{M}_A(0\lambda + 1))) \\ &= \mathcal{PM}(P(\lambda), \mathbf{M}_A(0\lambda + 1)) = \mathcal{PM}_{\text{given}}(0\lambda + 1), \end{aligned}$$

Thus the matrix polynomial $Q(\lambda)$ is the desired quasi-triangular realization of the given spectral data. The relation between the degree and the grade of $Q(\lambda)$ follows immediately from Lemma 2.7. \square

Example 4.8. Consider the following irreducible divisors along with their given partial multiplicity sequences:

Irreducible divisor	$\mathcal{PM}_{\text{given}}$
$\eta(\lambda) := \lambda^4 + \lambda + 1$	$(0, 0, 1, 2, 3, 3)$
$\phi(\lambda) = \lambda^2 + \lambda + 1$	$(0, 1, 1, 1, 1, 3)$
∞ , i.e., “ $0\lambda + 1$ ”	$(0, 1, 1, 2, 2, 4)$

with index sum 60. (Recall that the grade one matrix polynomial $0\lambda + 1$ is our stand-in for ∞ , so the third partial multiplicity sequence records the desired infinite spectral structure.) Thus we may legitimately choose to seek a quasi-triangular realization of this data with grade $g = 10$ and size $n = 6$.

The first step in constructing such a realization is to translate the data using an appropriate Möbius transformation. Since 0 is *not* an eigenvalue in the given spectral data, we may take $\omega = 0$ as our “special” value in the field $\mathbb{F} = \mathbb{Z}_2$, which will be the recipient of the spectral data at ∞ . The Möbius transformation(s) \mathbf{M}_R given by the matrix $R = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$, i.e., reversal, achieves this goal. Note that this reversal is taken with respect to grade equalling degree when applied to the irreducible divisors (with the exception of the grade *one* $0\lambda + 1$), and rev_g with $g = 10$ when applied to the matrix polynomial as a whole. The transformed spectral data is

Transformed irreducible divisor	$\mathcal{PM}_{\text{given}}$
$\mathbf{M}_R(\eta) = \lambda^4 + \lambda^3 + 1 = \chi(\lambda)$	$(0, 0, 1, 2, 3, 3)$
$\mathbf{M}_R(\phi) = \phi(\lambda)$	$(0, 1, 1, 1, 1, 3)$
$\mathbf{M}_R(0\lambda + 1) = \lambda = \psi(\lambda)$	$(0, 1, 1, 2, 2, 4)$

Observe that this is the same data as was used in Examples 3.2-3.23, so an appropriate realization for this transformed data is the 2-quasi-triangular (grade 10, strictly regular, and degree 10) realization from Example 3.23:

$$Q(\lambda) = \left[\begin{array}{cc|cc|cc} \chi\phi\psi(\phi+1) & \chi\phi(\chi+\psi+1) & 0 & 0 & \psi(\phi+1) & \phi+1 \\ \chi\phi\psi^4 & \chi\phi\psi^2 & 0 & 0 & \psi & 1 \\ \hline & & \chi\phi\psi^2 & \chi\phi\psi^2(\phi+1) & \phi\psi(\phi+1) & 0 \\ & & \chi\phi\psi^4 & \chi\phi\psi^2 & \phi\psi & 0 \\ \hline & & & & \chi\phi\psi(\psi^2+1) & \chi\phi^2\psi^2 \\ & & & & \chi\phi^2\psi^2 & \chi\phi^2\psi \end{array} \right].$$

To complete the construction of the realization for the original spectral data, we need to apply the inverse for the Möbius transformation \mathbf{M}_R used at the beginning, i.e., $\mathbf{M}_{R^{-1}} = \mathbf{M}_R$. The result is the desired realization

$$\tilde{Q}(\lambda) = \mathbf{M}_R(Q) = \text{rev}_{10}(Q) = \left[\begin{array}{cc|cc|cc} \eta\phi(\phi+1) & \eta\phi(\lambda^3 + \lambda + 1) & 0 & 0 & \lambda^8 + \lambda^7 & \lambda^9 + \lambda^8 \\ \eta\phi & \eta\phi\lambda^2 & 0 & 0 & \lambda^9 & \lambda^{10} \\ \hline 0 & 0 & \eta\phi\lambda^2 & \eta\phi(\lambda+1) & \phi\lambda^5(\lambda+1) & 0 \\ 0 & 0 & \eta\phi & \eta\phi\lambda^2 & \phi\lambda^7 & 0 \\ \hline 0 & 0 & 0 & 0 & \eta\phi(\lambda^3 + \lambda) & \eta\phi^2 \\ 0 & 0 & 0 & 0 & \eta\phi^2 & \eta\phi^2\lambda \end{array} \right].$$

where $\eta := \mathbf{M}_R(\chi) = \text{rev}_4(\chi)$.

Observe that $\text{grade}(\tilde{Q}) = 10$ by construction, and since α_1 (the first partial multiplicity at ∞) is zero, by Lemma 2.7 we know that $\deg \tilde{Q}$ must be the same as $\text{grade} \tilde{Q}$. Indeed $\deg \tilde{Q}$ is 10, but only barely. There is *only one* entry, the (2, 6) entry, that has degree 10. Consequently the leading coefficient of $\tilde{Q}(\lambda)$ has rank one, or equivalently rank deficiency 5, which is consistent with the eigenvalue at ∞ having geometric multiplicity 5, as specified in the given spectral data.

Observe also that both the (1, 1) and (3, 3) blocks have degree 9. So it would be possible to use Lemma 3.22 and either of these diagonal blocks to reduce the degree of the off-diagonal (1, 3)-block to be strictly less than 9 via unimodular equivalence. This would preserve the finite spectral structure, but would have the unwanted side effect of spoiling the infinite spectral structure. Since degree and grade would no longer be equal, by Lemma 2.7 we would have $\alpha_1 > 0$, contrary to the given spectral data at ∞ .

4.3 Quasi-triangularization of regular matrix polynomials

In the following signature result of this paper, the goal is to quasi-triangularize a given polynomial matrix, rather than to construct a quasi-triangular realization of some given data. That is, we will start with an arbitrary regular polynomial matrix $P(\lambda)$, and show that there must always be a quasi-triangular matrix $Q(\lambda)$ with the same degree, grade, and complete spectral data as $P(\lambda)$.

Theorem 4.9 (Quasi-Triangularization).

Suppose $P(\lambda)$ is a regular $n \times n$ matrix polynomial of grade g and degree d , over an arbitrary field \mathbb{F} . Define k to be the maximum degree among all of the \mathbb{F} -irreducible divisors of $P(\lambda)$. If $P(\lambda)$ is not strictly regular, further suppose that there is some constant $\omega \in \mathbb{F}$ such that $s_n(\omega) \neq 0$, where $s_n(\lambda)$ is the n^{th} invariant polynomial of $P(\lambda)$. (I.e., there is some $\omega \in \mathbb{F}$ that is not in the spectrum of P .) Then there exists a regular k -quasi-triangular matrix polynomial $Q(\lambda)$ over \mathbb{F} that has exactly the same size, grade, degree, and complete spectral data as $P(\lambda)$. When $P(\lambda)$ is strictly regular, then the k -quasi-triangularization $Q(\lambda)$ is strictly regular, and may be chosen to have the additional property that all off-diagonal blocks have degree strictly less than d .

Proof. From $P(\lambda)$, extract the complete spectral data, size, degree, and grade. If $P(\lambda)$ is strictly regular, then use this data together with Theorem 3.1 to construct the desired $Q(\lambda)$ with $\text{grade } Q = \deg Q = \deg P = \text{grade } P$.

If $P(\lambda)$ is not strictly regular, then define a new matrix polynomial $\tilde{P}(\lambda)$ that is entrywise identical to $P(\lambda)$, but with $\text{grade } \tilde{P}$ chosen to be equal to $d = \deg P$. (If $g = d$ to begin with, then \tilde{P} is identical to P in every way.) Note that the degrees of P and \tilde{P} as well as their finite spectral structures are the same, even though their grades

may be different. Now if $\mathcal{PM}(P, 0\lambda + 1) = (\alpha_1, \dots, \alpha_n)$ is the partial multiplicity sequence of P at ∞ , then by Lemma 2.7 the partial multiplicity sequence of \tilde{P} at ∞ is shifted by $\alpha_1 = g - d$ from that of P , i.e.,

$$\mathcal{PM}(\tilde{P}, 0\lambda + 1) = \mathcal{PM}(P, 0\lambda + 1) - \alpha_1 \cdot (1, 1, \dots, 1) = (0, \alpha_2 - \alpha_1, \dots, \alpha_n - \alpha_1).$$

Now we can use Theorem 4.7 to construct a k -quasi-triangular realization $\tilde{Q}(\lambda)$ for the complete spectral data, size, degree, and grade of this $\tilde{P}(\lambda)$, for which we have $\deg \tilde{Q} = \text{grade } \tilde{Q} = \text{grade } \tilde{P} = \deg \tilde{P} = \deg P = d$. Finally, we define a k -quasi-triangular matrix polynomial $Q(\lambda)$ that is entrywise identical to $\tilde{Q}(\lambda)$, and so has $\deg Q = \deg \tilde{Q} = d$, but now has grade chosen to be g . The infinite spectral data of Q is shifted by $\alpha_1 = g - d$ from that of \tilde{Q} , so we have

$$\begin{aligned} \mathcal{PM}(Q, 0\lambda + 1) &= \mathcal{PM}(\tilde{Q}, 0\lambda + 1) + \alpha_1 \cdot (1, 1, \dots, 1) \\ &= \mathcal{PM}(\tilde{P}, 0\lambda + 1) + \alpha_1 \cdot (1, 1, \dots, 1) = (\alpha_1, \alpha_2, \dots, \alpha_n) = \mathcal{PM}(P, 0\lambda + 1). \end{aligned}$$

Thus we see that $Q(\lambda)$ has exactly the same size, degree, grade, and complete spectral data as $P(\lambda)$, and so is the desired quasi-triangularization of $P(\lambda)$. \square

Remark 4.10. It is worth emphasizing that the relationship between $P(\lambda)$ and the quasi-triangularization $Q(\lambda)$ is stronger than just unimodular equivalence. Theorem 4.9 guarantees the existence of a *spectrally equivalent* k -quasi-triangularization for any regular matrix polynomial over an arbitrary field.

5 More on diagonal block sizes

In this final section we explore the range of possibilities for diagonal block sizes in degree-preserving quasi-triangularizations. We have shown that every regular matrix polynomial $P(\lambda)$ over an arbitrary field admits a spectrally equivalent degree-preserving k -quasi-triangularization, where k is the highest degree among the irreducible divisors of $P(\lambda)$. But is that really the best possible general result? Could it be that there is a smaller bound on diagonal block sizes of quasi-triangularizations that holds for all regular matrix polynomials? Section 5.1 addresses this issue, exhibiting a family of examples that shows that the k in Theorem 4.9 is indeed the best possible general bound for diagonal block sizes.

By contrast, in Section 5.2 we probe the opposite end of the size range of diagonal block sizes, trying to determine when it is possible to achieve diagonal blocks that are all 1×1 , i.e., when it is possible to just plain *triangularize* in a spectrally equivalent and degree-preserving way. Although we have not even come close to completely settling this question, we are at least able to identify some scenarios where a necessary and sufficient condition for triangularizability can be found, and some other more general scenarios where a condition sufficient to guarantee the existence of a triangularization can be given.

5.1 Sharpness of the upper bound k

The matrix polynomials described in Example 5.1 show that Theorem 4.9 provides the best possible general bound on diagonal block sizes. This is done via an infinite family of examples where it can be proved that every possible quasi-triangularization has all of its diagonal blocks of size $k \times k$ or larger, where k is the largest degree among all of the irreducible divisors.

Example 5.1. Consider any strictly regular $n \times n$ matrix polynomial $P(\lambda)$ with degree $d \geq 2$, such that $P(\lambda)$ has exactly one irreducible divisor $\chi(\lambda)$, and $k = \deg(\chi) \geq 2$ is coprime to d . Since $mk = dn$ for some $m \in \mathbb{N}$ by the index sum theorem, we must then also have $k|n$.

Suppose \mathbb{F} is any field that supports such an \mathbb{F} -irreducible polynomial $\chi(\lambda)$ with $k = \deg(\chi) \geq 2$; e.g., $\mathbb{F} = \mathbb{Q}$ has such a $\chi(\lambda)$ for any $k \geq 2$ at all. Then there are infinitely many choices of d and n that satisfy the conditions mentioned above, i.e., that k is coprime to d and $k|n$. For any of these choices of $k, d, n, \chi(\lambda)$, and the field \mathbb{F} , the Fundamental Realization Theorem 2.11 guarantees the existence of a matrix polynomial $P(\lambda)$ over \mathbb{F} as described above, i.e., one that is strictly regular, $n \times n$, degree d , and with just one irreducible divisor χ . Thus there are infinitely many matrix polynomials encompassed by the discussion in this Example.

Now suppose that $Q(\lambda)$ is *any* degree- d quasi-triangularization of $P(\lambda)$; that is, Q is block-upper-triangular, has degree d , and is unimodularly equivalent to P . Suppose Q has ℓ diagonal blocks $Q_{11}(\lambda), Q_{22}(\lambda), \dots, Q_{\ell\ell}(\lambda)$

with sizes $n_i \times n_i$, respectively, so that $n = \sum_{i=1}^{\ell} n_i$. Now each block $Q_{ii}(\lambda)$ has degree at most d , so that $\deg(\det Q_{ii}) \leq dn_i$. Thus we have

$$dn = \deg(\det P) = \deg(\det Q) = \deg\left(\prod_{i=1}^{\ell} \det Q_{ii}\right) = \sum_{i=1}^{\ell} \deg(\det Q_{ii}) \leq \sum_{i=1}^{\ell} dn_i = dn.$$

But this means that each inequality $\deg(\det Q_{ii}) \leq dn_i$ must actually be an equality, so $\deg(\det Q_{ii}) = dn_i \geq 2$, and hence $Q_{ii}(\lambda)$ is not a constant block. Since $\det Q = \alpha \chi^m$ for some nonzero scalar $\alpha \in \mathbb{F}$, that means that χ must divide each $\det Q_{ii}(\lambda)$. Thus k divides each dn_i , and hence also divides each n_i since k and d are coprime. But this means that $k \leq n_i$ for each $1 \leq i \leq \ell$, and so $Q(\lambda)$ is at best k -quasi-triangular.

Finally, note that the discussion in Remark 3.17 is relevant to this example. In that remark, it was pointed out that there is only one scenario in which the homogeneous partitioning procedure forces there to be sublists with exactly k elements, which then later lead to $k \times k$ diagonal blocks in the quasi-triangularization. It is not hard to show that this example falls exactly under this scenario, so our procedure will necessarily produce a quasi-triangularization with *all* of its diagonal blocks being $k \times k$. This, of course, is not equivalent to *proving* that no quasi-triangularization with any diagonal block of size smaller than $k \times k$ can exist, as we have done above. But it certainly is completely consistent with that result.

When d and k are coprime, then Example 5.1 has shown that the “best” quasi-triangularization that can be attained may sometimes be forced to have all of its diagonal blocks with size $k \times k$. However, there are many matrix polynomials that have quasi-triangularizations with much smaller diagonal blocks than the general upper bound of $k \times k$. Indeed we have seen this already in Example 3.23, where we had $k = 4$, but were able to construct a 2-quasi-triangularization. The next example gives a whole family of matrix polynomials that show that the gap between this general upper bound k and the actual smallest realizable diagonal block size for quasi-triangularizations can be arbitrarily large. The discussion in Section 5.2 provides further examples of this phenomenon.

Example 5.2. Consider an arbitrary target degree d and irreducible polynomial $\chi(\lambda)$ with $k = \deg(\chi) = 2d$. Note that there are many fields \mathbb{F} that support the presence of such high degree \mathbb{F} -irreducible polynomials, e.g., $\mathbb{F} = \mathbb{Q}$. For any such \mathbb{F} -irreducible $\chi(\lambda)$ of degree k , there is a unique way to express it in the form $\chi(\lambda) = \lambda^d p(\lambda) + q(\lambda)$, where $\deg p = d$ and $\deg q = d - 1$. Then letting

$$X(\lambda) = \begin{bmatrix} p(\lambda) & -1 \\ q(\lambda) & \lambda^d \end{bmatrix},$$

we see that $\deg X = d$, $\det X = \chi(\lambda)$, and the Smith form of $X(\lambda)$ is just $\text{diag}(1, \chi(\lambda))$. This implies that the Smith form of $X^6(\lambda)$ must be of the form $\text{diag}(\chi^\ell, \chi^m)$, where $\ell + m = 6$.

Now let $P(\lambda)$ be any strictly regular 12×12 matrix polynomial over \mathbb{F} with degree d , that has the Smith form $\text{diag}(I_{10}, \chi^\ell, \chi^m)$. By the Fundamental Realization Theorem 2.11 such matrix polynomials must exist. Consider next the matrix polynomial

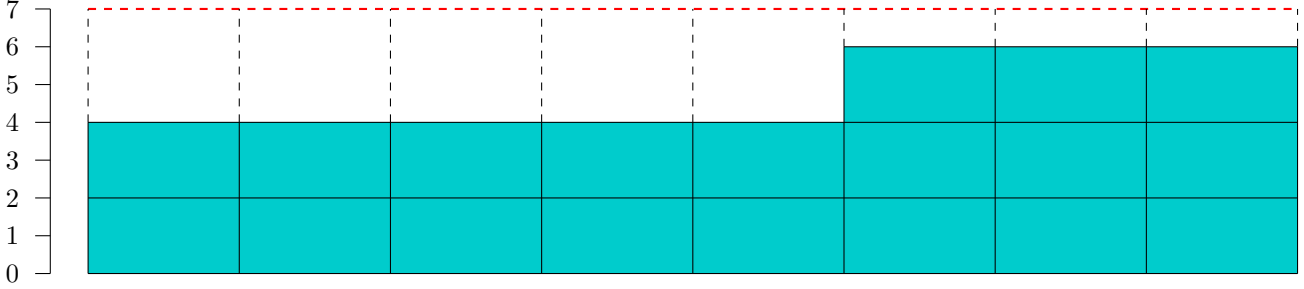
$$Q(\lambda) = \begin{bmatrix} X(\lambda) & I_2 & & & \\ & X(\lambda) & I_2 & & \\ & & X(\lambda) & I_2 & \\ & & & X(\lambda) & I_2 \\ & & & & X(\lambda) \end{bmatrix} \in \mathbb{F}[\lambda]^{12 \times 12}.$$

This $Q(\lambda)$ has degree d , and it is not hard to show that it is unimodularly equivalent to $\text{diag}(I_{10}, X^6(\lambda))$, and hence also to $P(\lambda)$. Thus $Q(\lambda)$ is a degree-preserving 2-quasi-triangularization of $P(\lambda)$. This $Q(\lambda)$ has very much smaller diagonal blocks than what is guaranteed by the general result in Theorem 4.9, with a gap (of $k - 2$) in the size of diagonal blocks between the general upper bound and those actually occurring in $Q(\lambda)$, a gap that can be arbitrarily large.

5.2 When triangularizing is possible

A natural question to ask is “When is it possible to triangularize?”, or in other words, when can we guarantee the existence of a degree-preserving quasi-triangularization in which all diagonal blocks are 1×1 ? To answer this, let’s begin with an example where all of the irreducible divisors have degree 2 or less, as in a real matrix polynomial.

Example 5.3. Suppose that we are trying to build a *triangularization* of a strictly regular matrix polynomial with degree $d = 7$ and size $n = 8$, and there are a total of $m_1 = 18$ degree-1 irreducible factors and $m_2 = 19$ degree-2 irreducible factors in the Smith form, for a total degree sum of 56. One strategy to build an appropriate target diagonal is to first spread the degree-2 factors out as much as possible, and then try to fill in the resulting gaps with the degree-1 factors.



Let \mathbf{f}_1 and \mathbf{f}_2 denote the degree-1 and degree-2 factor-counting vectors of the Smith form, respectively. If \mathbf{v}_2 is the homogenization of \mathbf{f}_2 , then we know that it is possible to spread out the 19 degree-2 factors along the diagonal to realize this \mathbf{v}_2 via unimodular transformations, using Lemma 3.10 and Corollary 3.4. This homogenization is visualized in the diagram above, where each column displays the contents of a diagonal entry location, each box stands for an irreducible factor, and the height of each box displays the degree of that factor, in this case a height/degree of 2. To achieve a triangularization, we need to have each diagonal entry have degree 7, so our goal is to populate each column in the diagram with boxes up to exactly height 7. The amount of remaining space between the top of the current stack of cyan blocks and the red dashed line in each column will be called a degree gap, and the vector containing all of the degree gaps will be called the gap vector \mathbf{g} . In this example the gap vector is $\mathbf{g} = (3, 3, 3, 3, 3, 1, 1, 1)$. So what remains is to try to distribute the 18 degree-1 irreducible factors (height-1 boxes) so as to exactly fill these gaps. In other words, we need to try to convert \mathbf{f}_1 into the gap vector \mathbf{g} . Now our only means to move these degree-1 factors around is to use the tools from Corollary 3.4, which correspond to compressions and interchanges of the entries of the factor-counting vector. But such actions can only convert \mathbf{f}_1 into a vector that it majorizes. (Recall the classical Muirhead theorem discussed right after Remark 3.9.) Thus this strategy will succeed in producing an appropriate diagonal for an achievable triangularization whenever $\mathbf{f}_1 \succeq \mathbf{g}$. Consequently we see that this majorization condition is a sufficient condition to guarantee the triangularizability of a strictly regular matrix polynomial whenever all irreducible divisors are of degree at most two.

In fact, though, the condition $\mathbf{f}_1 \succeq \mathbf{g}$ is also a necessary condition for triangularizability, and so gives a characterization in this scenario, as will be seen as an immediate consequence of the following development. We begin with some background lemmas. The first of these lemmas uses an alternative definition for majorization of vectors, that is equivalent to the one given earlier in Definition 3.8.

Definition 5.4 (Majorization [14]).

For vectors \mathbf{x}, \mathbf{y} in \mathbb{R}^n (or in \mathbb{Z}^n), let \mathbf{x}'' and \mathbf{y}'' denote the permutations of those vectors in which the entries have been arranged in *increasing* order. We say that \mathbf{x} *majorizes* \mathbf{y} , or \mathbf{y} *is majorized by* \mathbf{x} , and write $\mathbf{x} \succeq \mathbf{y}$, if

$$\sum_{i=1}^{\ell} x''_i \leq \sum_{i=1}^{\ell} y''_i \quad \text{for } \ell = 1, 2, \dots, n, \quad (5.1)$$

with equality when $\ell = n$.

Lemma 5.5. Suppose $T(\lambda)$ is a regular triangular matrix polynomial over an arbitrary field \mathbb{F} , and has the Smith form $S(\lambda)$. Let $\mathcal{F} \sqcup \mathcal{G}$ be any coprime partition of the multiset of all of the \mathbb{F} -irreducible factors of the invariant polynomials in $S(\lambda)$, equivalently of all of the \mathbb{F} -irreducible factors of the diagonal entries of $T(\lambda)$. Then the majorization relations

$$\mathbf{d}_{\mathcal{F}}(S) \succeq \mathbf{d}_{\mathcal{F}}(T) \quad (5.2)$$

hold for the diagonal factor-counting vectors of $S(\lambda)$ and $T(\lambda)$, with respect to every \mathcal{F} from every such coprime partition.

Proof. For convenience, let us introduce some notation to ease the discussion. Let

$$\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \dots, \sigma_n) := \mathbf{d}_{\mathcal{F}}(S)$$

be an abbreviation for the diagonal factor-counting vector $\mathbf{d}_{\mathcal{F}}(S)$. Note that $\sigma_1 \leq \sigma_2 \leq \dots \leq \sigma_n$ because of the divisibility chain property of invariant polynomials, but the entries of $\mathbf{d}_{\mathcal{F}}(T)$ may not be in any such order. So let $\tau = (\tau_1, \tau_2, \dots, \tau_n)$ be the permutation of $\mathbf{d}_{\mathcal{F}}(T)$ such that $\tau_1 \leq \tau_2 \leq \dots \leq \tau_n$. Now from (2.1) we know that $s_{11}(\lambda) | t_{jj}(\lambda)$ for $j = 1, \dots, n$, so $|s_{11}(\lambda)|_{\mathcal{F}} \leq |t_{jj}(\lambda)|_{\mathcal{F}}$ for $j = 1, \dots, n$, and thus $\sigma_1 \leq \tau_1 = \min_j |t_{jj}(\lambda)|_{\mathcal{F}}$.

Using (2.1) again, we have that $[s_{11}(\lambda)s_{22}(\lambda)] [t_{ii}(\lambda)t_{jj}(\lambda)]$ for $i, j = 1, \dots, n$ with $i \neq j$. Hence

$$|s_{11}(\lambda)s_{22}(\lambda)|_{\mathcal{F}} \leq |t_{ii}(\lambda)t_{jj}(\lambda)|_{\mathcal{F}}$$

for $i, j = 1, \dots, n$ with $i \neq j$, and consequently $\sigma_1 + \sigma_2 \leq \tau_1 + \tau_2$. In a similar manner, we see from (2.1) that for each $1 \leq \ell \leq n-1$ we have

$$|s_{11}(\lambda)s_{22}(\lambda) \cdots s_{\ell\ell}(\lambda)|_{\mathcal{F}} \leq |t_{i_1 i_1}(\lambda)t_{i_2 i_2}(\lambda) \cdots t_{i_\ell i_\ell}(\lambda)|_{\mathcal{F}}$$

for all ℓ -tuples $(i_1, i_2, \dots, i_\ell)$ with distinct entries and $i_1, i_2, \dots, i_\ell = 1, \dots, n$. Thus we have

$$\sigma_1 + \sigma_2 + \cdots + \sigma_\ell \leq \tau_1 + \tau_2 + \cdots + \tau_\ell$$

for each $1 \leq \ell \leq n-1$. Finally, since $\det S = c \cdot \det T$ for some nonzero scalar $c \in \mathbb{F}$, we have

$$|s_{11}(\lambda)s_{22}(\lambda) \cdots s_{nn}(\lambda)|_{\mathcal{F}} = |t_{11}(\lambda)t_{22}(\lambda) \cdots t_{nn}(\lambda)|_{\mathcal{F}},$$

and hence

$$\sigma_1 + \sigma_2 + \cdots + \sigma_n = \tau_1 + \tau_2 + \cdots + \tau_n.$$

By Definition 5.4, then, we have $\sigma \succeq \tau$, and hence that (5.2) holds. \square

The next result shows that if the irreducible divisors of a matrix polynomial have *only two different degrees*, and one of those is degree 1, then having a triangularization of any kind implies that there must also exist a triangularization in which the highest degree irreducible factors are indeed "spread out as much as possible" as in Example 5.3, i.e., where their factor-counting vector is 1-homogeneous.

Lemma 5.6. *Suppose a strictly regular matrix polynomial $P(\lambda)$ of degree d has a triangularization $Q_0(\lambda)$ of degree d . Further suppose that the multiset \mathcal{M} of all of the \mathbb{F} -irreducible factors in the Smith form for $P(\lambda)$ contains only two distinct degrees, 1 and k , for some $k \geq 2$. Let $\mathcal{M} = \mathcal{F}_1 \sqcup \mathcal{F}_k$ be the coprime partition in which \mathcal{F}_j contains all of the irreducible factors in \mathcal{M} of degree j for $j = 1, k$. Then $P(\lambda)$ has a triangularization $T(\lambda)$ of degree d in which the diagonal factor-counting vector $\mathbf{d}_{\mathcal{F}_k}(T)$ is 1-homogeneous.*

Proof. If $\mathbf{d}_{\mathcal{F}_k}(Q_0)$ is already 1-homogeneous, take $T(\lambda) = Q_0(\lambda)$ and then of course we are done. So suppose that $\kappa = (\kappa_1, \kappa_2, \dots, \kappa_n) := \mathbf{d}_{\mathcal{F}_k}(Q_0)$ is not 1-homogeneous. The argument will consist of a procedure showing how to convert $Q_0(\lambda)$ by a finite sequence of triangularizations for $P(\lambda)$ into a degree- d triangularization $T(\lambda)$ that has the desired 1-homogeneity property.

Let κ_i and κ_j be the minimum and maximum entries in κ , so $\kappa_j - \kappa_i \geq 2$. Now by a finite sequence of interchanges we can arrange that the corresponding diagonal entries of $Q_0(\lambda)$ are adjacent, say in the (i, i) and $(i+1, i+1)$ locations, and these interchanges can be implemented by unimodular transformations using Corollary 3.4. This gives us a new degree d triangularization $\tilde{Q}_0(\lambda)$ for $P(\lambda)$, with a new $\tilde{\kappa} = \mathbf{d}_{\mathcal{F}_k}(\tilde{Q}_0)$ with $\tilde{\kappa}_i = \kappa_i$ and $\tilde{\kappa}_{i+1} = \kappa_j$. Now we do a compression of the degree k factors in these two adjacent entries, again via a unimodular transformation from Corollary 3.4, decreasing the maximum $\tilde{\kappa}_{i+1}$ by one and increasing the minimum $\tilde{\kappa}_i$ by one. This gives us a triangular matrix polynomial $\hat{Q}_0(\lambda)$ that is no longer degree d , although it is still unimodularly equivalent to $P(\lambda)$; all diagonal entries have degree d , except for the i^{th} and $(i+1)^{\text{th}}$, which now have degrees $d+k$ and $d-k$, respectively. We can now restore degree d on these adjacent diagonal entries, by doing a compression of the degree 1 factors, again using a unimodular transformation from Corollary 3.4. The i^{th} diagonal entry of $\hat{Q}_0(\lambda)$ has $d - k\kappa_i$ degree-1 factors, while the $(i+1)^{\text{th}}$ diagonal entry has $d - k\kappa_j$ degree-1 factors. Since $(d - k\kappa_i) > (d - k\kappa_j)$, with a degree difference of at least $2k$, we can do a compression of the degree-1 factors where the i^{th} diagonal entry loses k degree-1 factors, and the $(i+1)^{\text{th}}$ diagonal entry gains k degree-1 factors. This gives a triangular matrix polynomial $\check{Q}_0(\lambda)$ that is unimodularly equivalent to $P(\lambda)$, and has all diagonal entries with degree d again, but the off-diagonal entries may now have degree larger than d . This is remedied by using Lemma 3.22, with all diagonal blocks taken to be of size 1×1 . This finally gives us a degree- d triangularization $Q_1(\lambda)$ of $P(\lambda)$, with $\mathbf{d}_{\mathcal{F}_k}(Q_1)$ one step closer to being 1-homogeneous than $\mathbf{d}_{\mathcal{F}_k}(Q_0)$ was.

Of course if $\mathbf{d}_{\mathcal{F}_k}(Q_1)$ is now 1-homogeneous, then we take $T(\lambda) = Q_1(\lambda)$ and we are done. If not, we repeat the above procedure on $Q_1(\lambda)$ to produce a new degree- d triangularization $Q_2(\lambda)$ with diagonal factor-counting vector $\mathbf{d}_{\mathcal{F}_k}(Q_2)$ that is even closer to being 1-homogeneous. Continuing this, we generate a sequence Q_0, Q_1, Q_2, \dots of degree- d triangularizations for $P(\lambda)$, which in finitely many steps must eventually produce a degree- d triangularization $T(\lambda)$ for which $\mathbf{d}_{\mathcal{F}_k}(T)$ is 1-homogeneous. \square

Proposition 5.7. *Let $P(\lambda)$ be a strictly regular $n \times n$ polynomial matrix of degree d over a field \mathbb{F} . Let $S(\lambda)$ be the Smith form of $P(\lambda)$, and assume that all irreducible divisors of $P(\lambda)$ are degree 1 or degree k , where $k \geq 2$. Let $\mathcal{F}_1 \sqcup \mathcal{F}_k$ be the same coprime partition of the multiset of all \mathbb{F} -irreducible factors in $S(\lambda)$ as in Lemma 5.6. Organize the degree- k factors into a vector of n polynomials $\mathbf{q}(\lambda) := (q_1(\lambda), q_2(\lambda), \dots, q_n(\lambda))$ in the same way as in the proof of Lemma 3.11 (and as in the diagram for Example 5.3), i.e., so that $|\mathbf{q}(\lambda)|_{\mathcal{F}_k}$ is 1-homogeneous. Define the degree gaps $g_i := d - \deg(q_i)$, and the corresponding gap vector $\mathbf{g} := (g_1, g_2, \dots, g_n)$. (Note that some of the g_i may be negative.) Then $P(\lambda)$ has a triangularization of degree d if and only if the majorization condition $\mathbf{d}_{\mathcal{F}_1}(S) \succeq \mathbf{g}$ holds.*

Proof. (\Rightarrow) Suppose $P(\lambda)$ has a triangularization of degree d . Then by Lemma 5.6, $P(\lambda)$ has a triangularization $T(\lambda)$ of degree d in which the diagonal factor-counting vector $\mathbf{d}_{\mathcal{F}_k}(T)$ is 1-homogeneous. Applying Lemma 5.5 to the coprime partition $\mathcal{F}_1 \sqcup \mathcal{F}_k$, we then have that $\mathbf{d}_{\mathcal{F}_1}(S) \succeq \mathbf{d}_{\mathcal{F}_1}(T)$. But in this triangularization $T(\lambda)$, it is easy to see that the diagonal factor-counting vector $\mathbf{d}_{\mathcal{F}_1}(T)$ is exactly the same as the gap vector \mathbf{g} . Thus $\mathbf{d}_{\mathcal{F}_1}(S) \succeq \mathbf{g}$, as desired.

(\Leftarrow) Now conversely, suppose that the majorization condition $\mathbf{d}_{\mathcal{F}_1}(S) \succeq \mathbf{g}$ holds. Starting from $S(\lambda)$, we know from Corollary 3.4 and Lemma 3.10 that we can spread out the irreducible factors in \mathcal{F}_k along the diagonal via unimodular transformations so as to form a triangular matrix polynomial $\tilde{T}(\lambda)$ such that $\mathbf{d}_{\mathcal{F}_k}(\tilde{T}) = |\mathbf{q}(\lambda)|_{\mathcal{F}_k}$ is 1-homogeneous, and $\mathbf{d}_{\mathcal{F}_1}(\tilde{T}) = \mathbf{d}_{\mathcal{F}_1}(S)$. Since $\mathbf{d}_{\mathcal{F}_1}(S) \succeq \mathbf{g}$, there exists a finite sequence of interchanges and compressions of adjacent diagonal entries that will turn $\mathbf{d}_{\mathcal{F}_1}(S)$ into \mathbf{g} . Implementing this sequence via unimodular transformations from Corollary 3.4 will produce a triangular matrix polynomial $T(\lambda)$ such that $\mathbf{d}_{\mathcal{F}_k}(T) = \mathbf{d}_{\mathcal{F}_k}(\tilde{T}) = |\mathbf{q}(\lambda)|_{\mathcal{F}_k}$ is 1-homogeneous, $\mathbf{d}_{\mathcal{F}_1}(T) = \mathbf{g}$, $\deg T = d$, and $T(\lambda)$ is unimodularly equivalent to $S(\lambda)$, and hence also to $P(\lambda)$. In other words, $T(\lambda)$ is the desired degree- d triangularization of $P(\lambda)$. \square

Remark 5.8. Necessary and sufficient conditions for the triangularizability of strictly regular *real* matrix polynomials were given in [19, Theorem 4.9]. Note that Proposition 5.7 recovers this result from [19] as the special case when $k = 2$.

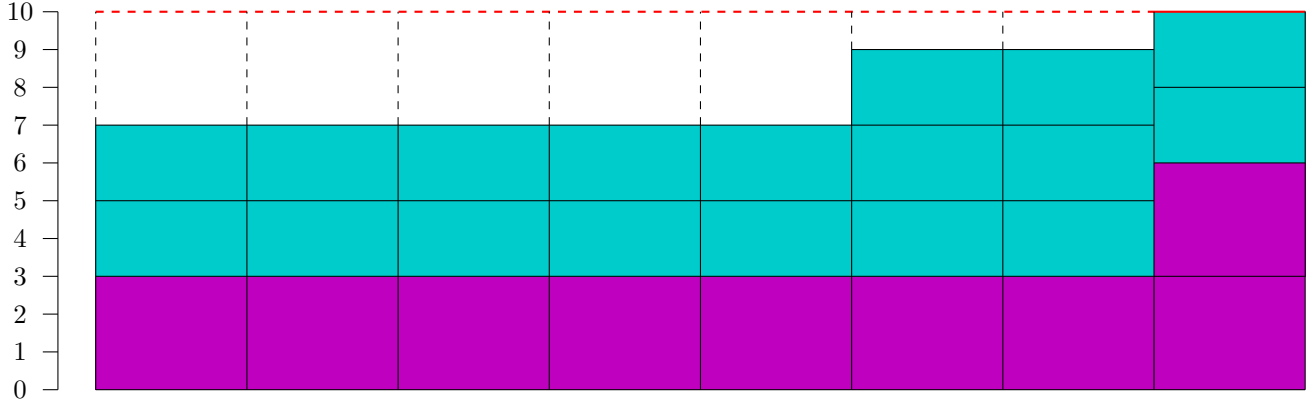
The result of Proposition 5.7 can be extended to a slightly more general scenario, still involving irreducible divisors with only two degrees, but no longer tied to requiring one of those degrees to be 1. This more general scenario is essentially just a “scaled” version of the one discussed in the Proposition.

Corollary 5.9. *Let $P(\lambda)$ be a strictly regular $n \times n$ polynomial matrix of degree d over a field \mathbb{F} . Let $S(\lambda)$ be the Smith form of $P(\lambda)$, and assume that all irreducible divisors of $P(\lambda)$ are degree ℓ or degree k , where $k > \ell \geq 1$. Also assume that $\ell|k$ and $\ell|d$. Let $\mathcal{F}_\ell \sqcup \mathcal{F}_k$ be the coprime partition of the multiset of all \mathbb{F} -irreducible factors in $S(\lambda)$ where \mathcal{F}_ℓ contains all of the degree- ℓ factors, and \mathcal{F}_k contains all of the degree- k factors. Organize the degree- k factors into a vector of n polynomials $\mathbf{q}(\lambda) := (q_1(\lambda), q_2(\lambda), \dots, q_n(\lambda))$ in the same way as in the proof of Lemma 3.11, i.e., so that $|\mathbf{q}(\lambda)|_{\mathcal{F}_k}$ is 1-homogeneous. Define the degree gaps $g_i := d - \deg(q_i)$, and the corresponding gap vector $\mathbf{g} := (g_1, g_2, \dots, g_n)$. (Note that some of the g_i may be negative.) Then $P(\lambda)$ has a triangularization of degree d if and only if the majorization condition $\ell \cdot \mathbf{d}_{\mathcal{F}_\ell}(S) \succeq \mathbf{g}$ holds.*

Proof. The scenario described in this corollary can be viewed as a scaled version of the one handled in Proposition 5.7. From the divisibility assumptions $\ell|k$ and $\ell|d$, let $k = \kappa\ell$ and $d = \delta\ell$, with $1 < \kappa < \delta$. Then by viewing ℓ as the basic “unit” of degree, the scenario of this corollary is just like that of Proposition 5.7 with 1, k and d replaced by 1, κ , and δ . \square

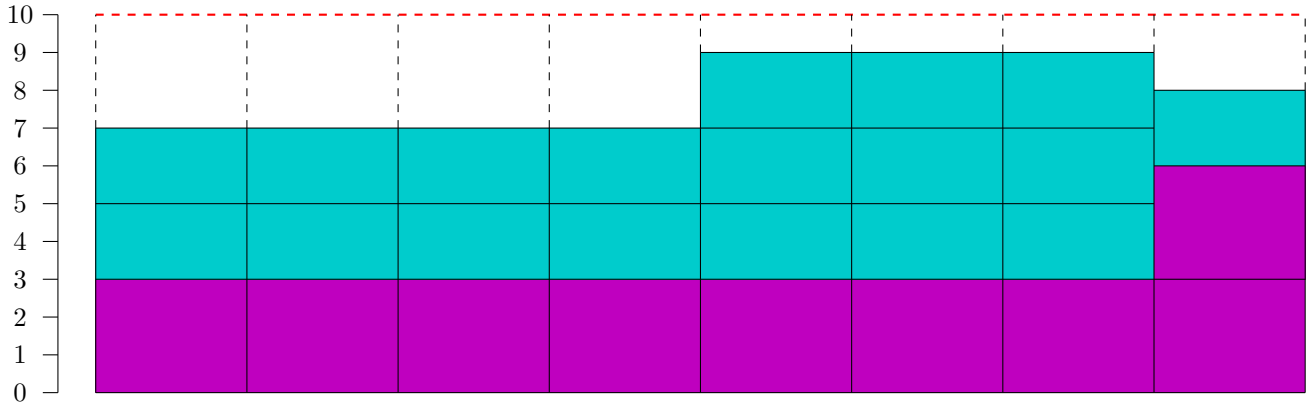
The results in Proposition 5.7 and its Corollary 5.9 show that the triangularization question is still somewhat tractable when there are no more than two different degrees among all of the irreducible divisors that are present. However, when irreducible divisors have three or more degrees, the picture gets much more involved, with some significant combinatorial complexity now possible. The strategy guiding Example 5.3 is still viable, though, and sometimes is able to provide sufficient conditions for guaranteeing that a triangularization is possible, although these conditions may no longer be necessary. To see why this is the case, we consider a few more examples, this time with irreducible divisors of degrees 1, 2, and 3.

Example 5.10. For this example we aim for degree $d = 10$ and size $n = 8$, but this time with $m_3 = 9$ degree-3 factors, $m_2 = 18$ degree-2 factors, and $m_1 = 17$ degree-1 factors, for a total degree sum of 80. The sufficient condition is the same as before, i.e., the gap vector must be majorized by the degree-1 factor-counting vector, but the gap vector is defined slightly differently to how it was done in Example 5.3. To determine the gap vector, begin by spreading out the degree-3 and degree-2 factors in a way similar to Lemma 3.11, as pictured in the following diagram. Note that we will definitely be able to do this by unimodular transformations, since both the degree-3 and the degree-2 factor-counting vectors in the diagram are 1-homogeneous, and we have already seen that factor-counting vectors can always be homogenized.



The gap vector for this configuration is $\mathbf{g} = (3, 3, 3, 3, 3, 1, 1, 0)$, so in order for a triangularization (with this particular degree-3 and degree-2 configuration) to be guaranteed to exist, this vector must be majorized by the degree-1 factor-counting vector for the diagonal of the Smith form. For instance, if the degree-1 factor-counting vector in the Smith form is $\mathbf{f}_1 = (0, 0, 0, 0, 2, 5, 5, 5)$, then there is a triangularization, since $\mathbf{f}_1 \succeq \mathbf{g}$. On the other hand if the Smith form's degree-1 factor-counting vector is $\tilde{\mathbf{f}}_1 = (1, 1, 1, 2, 3, 3, 3, 3)$, then a triangularization may still exist, but it cannot be guaranteed to exist by this pathway since $\tilde{\mathbf{f}}_1 \not\succeq \mathbf{g}$.

However, if we modify the layout of the degree-2 factors just a little bit, then we can see that the degree-1 factor-counting vector $\tilde{\mathbf{f}}_1$ will admit a triangularization. Let us shift one (cyan) height-2 block from the eighth column to the fifth column, as in the diagram.



But is this new configuration actually reachable by unimodular transformations? Now we have a new condition; the degree-2 factor-counting vector in the Smith form must majorize $(2, 2, 2, 2, 3, 3, 3, 1)$, the degree-2 factor-counting vector in this new configuration. Assuming that this new condition is satisfied, we still need the degree-1 factor-counting vector to majorize the new gap vector $\tilde{\mathbf{g}} = (3, 3, 3, 3, 1, 1, 1, 2)$. For the particular degree-1 factor-counting vector $\tilde{\mathbf{f}}_1$ that failed before, though, everything is now fine, since $\tilde{\mathbf{f}}_1$ and $\tilde{\mathbf{g}}$ are just permutations of each other.

So in order to guarantee the existence of a triangularization using this new configuration of degree-3 and degree-2 factors, in general we will need *two* majorization conditions to be satisfied. One can now easily imagine the combinatorial nightmare that will almost certainly accompany any effort to devise general conditions that

are necessary for triangularization in the arbitrary field setting. This is why we have contented ourselves with only a brief discussion of simple sufficient conditions for triangularizability when there are at least three degrees of irreducible divisor present. We leave the investigation of necessary conditions for triangularizability for further research.

This brings us now to our final result, which gives a generalized sufficient condition for guaranteeing the existence of a triangularization over an arbitrary field.

Proposition 5.11. *Let $P(\lambda)$ be a strictly regular $n \times n$ polynomial matrix of degree d over a field \mathbb{F} . Let $S(\lambda)$ be the Smith form of $P(\lambda)$, and assume that all irreducible divisors are degree k or less. Organize the degree-2 through degree- k factors into n polynomials $q_1(\lambda), q_2(\lambda), \dots, q_n(\lambda)$ in the same way as in the proof of Lemma 3.11, and define the degree gaps $g_i := d - \deg(q_i)$. (Note that some of the g_i may be negative.) If the degree-1 factor-counting vector for the diagonal of $S(\lambda)$ majorizes the gap vector $\mathbf{g} := (g_1, g_2, \dots, g_n)$, then $P(\lambda)$ has a triangularization of degree d .*

Proof. It is possible to employ the techniques pictured in Examples 5.3 and 5.10 (i.e., following the pattern of the proof of Lemma 3.11 in distributing all irreducible factors of degree two and higher, and then filling in the rest of the available spaces with all of the remaining degree-1 factors) to design a target diagonal in which all of the entries have degree d . Note that the degree- ℓ factor-counting vectors for $2 \leq \ell \leq k$ in this target diagonal are all 1-homogeneous, and thus are definitely all realizable by spreading out the irreducible divisors in the Smith form using Corollary 3.4 and Lemma 3.10. The majorization hypothesis about the degree-1 factor-counting vector for the diagonal of $S(\lambda)$ then suffices to imply that the degree-1 factor-counting vector for the target diagonal, i.e., the gap vector \mathbf{g} , can also be realized using Corollary 3.4. Once all of these factor-counting vectors for the target diagonal are realized, we will have attained the desired degree- d triangularization for $P(\lambda)$. \square

In the statement of Proposition 5.11 it was noted that under the given conditions, it is possible for the gap vector \mathbf{g} to have negative entries. Whenever this occurs, then it is impossible for any conceivable degree-1 factor-counting vector for the Smith form to majorize \mathbf{g} , since all entries of a factor-counting vector are non-negative. In this scenario, then, Proposition 5.11 tells us nothing about the existence or non-existence of a triangularization. Other arrangements of the higher degree irreducible factors along the diagonal may still lead to a triangularization, as illustrated in Example 5.10.

Remark 5.12. Note that the condition in Proposition 5.11 for ensuring triangularizability can be adapted to regular matrix polynomials $P(\lambda)$ having nontrivial infinite spectral structure. First apply a Möbius transformation to transform $P(\lambda)$ into a matrix polynomial $Q(\lambda)$ with only finite spectral structure, i.e., into a strictly regular matrix polynomial, as was done in the proof of Theorem 4.7. Since any Möbius transformation preserves the degree of any irreducible divisor of degree two or higher by Lemma 4.4, all of the degree- ℓ factor-counting vectors of $Q(\lambda)$ will be exactly the same as those of $P(\lambda)$, except for $\ell = 1$. The partial multiplicities at ∞ for $P(\lambda)$ will turn into partial multiplicities for $Q(\lambda)$ at some degree-1 irreducible $\lambda - \omega$, hence the degree-1 factor-counting vector for the Smith form of $Q(\lambda)$ will be equal to the sum of the degree-1 factor-counting vector for the Smith form of $P(\lambda)$ together with the partial multiplicity sequence for $P(\lambda)$ at ∞ . In other words, the infinite partial multiplicities effectively get included with all of the degree-1 irreducible factors. We can now apply the majorization condition in Proposition 5.11 (or Proposition 5.7) to determine if a triangularization for $Q(\lambda)$ is guaranteed. If it is, then the inverse Möbius transformation applied to this triangularization for $Q(\lambda)$ provides a spectrally equivalent triangularization for $P(\lambda)$. Note that this generalization to all regular matrix polynomials appears in [19, 20] for real matrix polynomials.

6 Conclusion

This work has shown that any regular matrix polynomial $P(\lambda)$ over an arbitrary field \mathbb{F} is spectrally equivalent to a k -quasi-triangular matrix polynomial over \mathbb{F} of the same size and degree, where k is the largest degree among all of the irreducible divisors of $P(\lambda)$. This extends and generalizes the earlier work in [19], which found triangularizations and 2-quasi-triangularizations for regular matrix polynomials over algebraically closed fields and the real field \mathbb{R} . We have also shown that for any field \mathbb{F} , this k is the best possible general bound on the diagonal block sizes of quasi-triangularizations that holds for all regular matrix polynomials over \mathbb{F} .

Several new tools and results were developed in order to achieve this extension to arbitrary fields. Among these are:

- a technique to allow the flexible but controlled movement of individual irreducible factors up and down the diagonal of a triangular polynomial matrix via unimodular transformations,
- a new homogeneous partitioning property of “tightly packed” integer multisets,
- a reformulation of the interaction of Möbius transformations with spectral data, in a way that makes it easier to work with higher degree irreducible divisors.

A number of issues remain to be settled, especially ones related to the size of diagonal blocks in quasi-triangularizations. Although we know that these block sizes need never be any larger than k , and that sometimes they are all forced to be of size exactly k , very often quasi-triangularizations can be found with diagonal block sizes much smaller than the upper bound k . For given spectral data, can one predict how small the diagonal blocks can be made in a quasi-triangularization, and indeed when these blocks can all be made 1×1 , i.e., when can we actually triangularize? Some limited results were given along these lines, but much about this question still remains open.

References

- [1] L.M. ANGUS, F.M. DOPICO, R. HOLLISTER, AND D.S. MACKEY. Van Dooren’s index sum theorem and rational matrices with prescribed structural data. *SIAM J. Matrix Anal. Appl.*, 40(2) (2019), p.720–738.
- [2] M. ARTIN. *Algebra*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1991.
- [3] F. DE TERÁN, F.M. DOPICO, AND D.S. MACKEY. Spectral equivalence of matrix polynomials and the index sum theorem. *Linear Algebra Appl.*, 459 (2014), p.264–333.
- [4] F. DE TERÁN, F.M. DOPICO, AND P. VAN DOOREN. Matrix polynomials with completely prescribed eigenstructure. *SIAM J. Matrix Anal. Appl.*, 36 (2015), p.302–328.
- [5] F.M. DOPICO, D.S. MACKEY, AND P. VAN DOOREN. Product realizations of structural data for matrix polynomials. In preparation.
- [6] G. FROBENIUS. Theorie der linearen Formen mit ganzen Coefficienten. *J. Reine Angew. Math. (Crelle)* 86 (1878) p.146–208.
- [7] F.R. GANTMACHER. *The Theory of Matrices*. Chelsea Pub. Co., New York, 1959.
- [8] I. GOHBERG, P. LANCASTER, AND L. RODMAN. *Matrix Polynomials*. Academic Press, New York-London, 1982.
- [9] T. KAILATH. *Linear Systems*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1980.
- [10] D.E. KNUTH. *The Art of Computer Programming. Vol. 2: Seminumerical Algorithms*. Addison-Wesley, Reading, Massachusetts, 1981, 2nd Ed.
- [11] P. LANCASTER AND M. TISMENETSKY. *The Theory of Matrices*. Academic Press, London, 1985.
- [12] D.S. MACKEY, N. MACKEY, C. MEHL, AND V. MEHRMANN. Möbius transformations of matrix polynomials. *Linear Algebra Appl.*, 470 (2015), p.120–184.
- [13] E. MARQUES DE SÁ. On the diagonals of integral matrices. *Czech. Math. J.*, 30(2) (1980), p.207–212.
- [14] A.W. MARSHALL, I. OLKIN, AND B.C. ARNOLD. *Inequalities: Theory of Majorization and Its Applications*. Springer, 2011.
- [15] R.F. MUIRHEAD. Some methods applicable to identities and inequalities of symmetric algebraic functions of n letters. *Proc. Edinburgh Math. Soc.*, 21 (1903), p.144–157.
- [16] W.H.L. NEVEN AND C. PRAAGMAN. Column reduction of polynomial matrices. *Linear Algebra Appl.*, 188/189 (1993), p.569–589.
- [17] C. PRAAGMAN. Invariants of polynomial matrices. In I. Landau (Ed.), *Proceedings of the First European Control Conference*. Grenoble, 1991, INRIA, 1991, p.1274–1277.

- [18] H.J.S. SMITH. On systems of linear indeterminate equations and congruences. *Phil. Trans. Roy. Soc. London*, 151 (1861), p.293–326.
- [19] L. TASLAMAN, F. TISSEUR, AND I. ZABALLA. Triangularizing matrix polynomials. *Linear Algebra and Its Applications*, 439 (2013), p.1679–1699.
- [20] F. TISSEUR AND I. ZABALLA. Triangularizing quadratic matrix polynomials. *SIAM J. Matrix Anal. Appl.*, 34 (2013), p.312–337.
- [21] G. VERGHESE, P. VAN DOOREN, AND T. KAILATH. Properties of the system matrix of a generalized state-space system. *Internat. J. Control*, 30 (1979), p.235–243.