



Research article**Ideals, duality, and automorphisms of multi-branch finite local rings****Sami H. Saif*, and Alhanouf Ali Alhomaiddhi***

Department of Mathematics, College of Science, King Saud University, P.O. Box 2455, Riyadh 11451, Saudi Arabia

* **Correspondence:** Email: ssaif1@ksu.edu.sa; aalhomaiddhi@ksu.edu.sa.

Abstract: Let $k = \mathbb{F}_q$, let $r \geq 1$, and let $\mathbf{n} = (n_1, \dots, n_r)$ with $n_i \geq 2$. We determine the ideal, annihilator, cyclic-module, and automorphism structures of the finite local k -algebra $R_{\mathbf{n}} = k[x_1, \dots, x_r] / \langle x_i^{n_i}, x_i x_j \ (i \neq j) \rangle$. Every proper ideal is represented uniquely by a minimal branch-cutoff vector and an admissible terminal subspace. Transition maps and coordinate saturation then give explicit containment, intersection, and sum formulas and reveal projective and Grassmannian fibers in the ideal lattice. A closed formula counts the ideals for every r and specializes to the chain and two-branch counts. From the same data, we compute minimal generating sets, a bivariate generator–rank enumerator, and a projective-torus parameterization of principal ideals. We also determine annihilators, double-annihilator closure, its exact defect, and the resulting distributive duality core. For cyclic quotients, this closure is the torsionless reflection, in the sense that the quotient by $I^{\perp\perp}$ is universal among maps into torsionless modules; it yields a complete classification of torsionless cyclic modules and shows that every nonzero proper cyclic quotient has bidual cokernel dimension $r^2 - 1$. Finally, we determine the full automorphism group together with its tangent, socle, and depth-two actions on ideals. The two-branch rings $k[u, v] / \langle u^\ell, v^\ell, uv \rangle$ are recovered as the first nontrivial specialization.

Keywords: finite local algebra; fiber product; ideal lattice; annihilator closure; cyclic module; automorphism group

Mathematics Subject Classification: 13E10, 13C13, 16W20, 06C05

1. Introduction

Finite-dimensional local algebras bring together three structures that are usually studied separately: the filtration, the socle, and the ideal lattice. In a finite chain ring, all three are controlled by one nilpotent direction: The radical is principal, the ideals are linearly ordered, and each ideal is a radical power. This rigidity underlies the detailed structure theory of finite chain rings developed in [1], together with the description of their unit groups in [2, 3] and of their automorphism groups in [4].

The picture changes when the embedding dimension exceeds one. Even if distinct radical directions annihilate one another, an ideal need not be the direct sum of its intersections with those directions. Terminal classes from several branches may be tied together by common generators. Such couplings introduce finite-geometric data into the ideal lattice, reduce minimal generator numbers, and alter both annihilator duality and automorphism actions. These effects are recorded neither by the individual branch chains nor by the dimensions alone.

Let $k = \mathbb{F}_q$, let $r \geq 1$, and let $\mathbf{n} = (n_1, \dots, n_r)$ with $n_i \geq 2$. We study

$$R_{\mathbf{n}} = k[x_1, \dots, x_r] / \langle x_i^{n_i}, x_i x_j \ (i \neq j) \rangle. \quad (1.1)$$

Equivalently, $R_{\mathbf{n}}$ is the iterated fiber product of the truncated polynomial rings $k[x_i] / \langle x_i^{n_i} \rangle$ over their common residue field. Its Jacobson radical has the branch decomposition $J = A_1 \oplus \dots \oplus A_r$, $A_i = x_i k[x_i] / \langle x_i^{n_i} \rangle$, and $A_i A_j = 0 \ (i \neq j)$. Each branch retains the one-dimensional filtration of a chain ring, while the common identity and residue field place the branches inside one local algebra. Thus, the family lies between finite chain rings and arbitrary finite local rings: It is rigid enough for a complete analysis but still exhibits higher embedding dimension, a multidimensional socle, failure of the Frobenius property, and geometric families of ideals.

Fiber products and related amalgamations of local rings have been studied primarily through transfer questions for homological and Gorenstein properties. Moore determined the structure of the cohomology algebra of a fiber product [5], and Nasseh and Sather-Wagstaff characterized the vanishing of Ext and Tor over fiber products [6]; see [7] for amalgamated duplications and [8] for connected sums. All of these are transfer results: An invariant of $S_1 \times_k \dots \times_k S_r$ is expressed through the corresponding invariants of the factors, and the ideal lattice is not an invariant of that kind. Each $\text{Idl}(S_i)$ is a chain, whereas $\text{Idl}(R_{\mathbf{n}})$ is modular but never distributive for $r \geq 2$ (Corollary 3.8), and by Corollary 3.10 the number of proper ideals exceeds the number $\prod_{i=1}^r n_i$ of branchwise ideals by a strictly positive polynomial in q , a surplus carried by ideals that no transfer principle over the factors can see. Those results describe when a property of the components is inherited by the amalgam, but they do not describe the ideals of the amalgam individually. The problem treated here is of the second kind. We determine the internal ideal structure of the specific family (1.1), including the full lattice operations, refined enumerators, annihilator closure, cyclic biduality, and automorphism actions. The point is not the branch decomposition itself, but the additional subspace data needed to reconstruct an ideal and the way that data changes under sums, annihilation, dualization, and symmetry.

The ideal theory of finite rings also provides broader motivation. For finite chain rings, Norton and Sălăgean developed residue–torsion methods for linear and cyclic codes [9]; Dinh and López-Permouth treated cyclic and negacyclic codes [10]; and Sălăgean studied the repeated-root case [11]. Each of these methods presupposes that the ideals of the coefficient ring form a chain, so none of them applies once $r \geq 2$. Wood’s module-duality framework explains why the Frobenius property is decisive for field-like duality over finite rings [12], and Corollary 2.3 below shows that this property fails for every genuinely multi-branch member of (1.1). Coefficient rings with several nilpotent directions have appeared in cyclic-code constructions, in most cases with nonvanishing mixed products, as in [13–16]; the split case relevant here, in which distinct radical directions annihilate one another, was treated for two branches of length two in [17]. That paper studies cyclic codes over the coefficient ring $R_{(2,2)} = k[u, v] / \langle u^2, v^2, uv \rangle$ and does not address the ideal-theoretic invariants considered here. Its ring is moreover the degenerate member of the family: There $J^2 = 0$, so every k -subspace of $J \cong k^2$ is an ideal, the cutoff vector carries

no information beyond the branch supports, and the branch–subspace calculus reduces to linear algebra, with $\kappa(I) \leq 1$, a duality core of three elements, constant bidual cokernel 3, and $\text{Aut}_k(R_{(2,2)}) \cong \text{GL}_2(k)$. Each of these degeneracies fails once $\max_i n_i \geq 3$ or $r \geq 3$, where the cutoff vector and the terminal subspace become independent, admissible subspaces of dimension greater than one occur, $\kappa(I)$ reaches $r - 1$, the duality core has $2 + \prod_i (n_i - 1)$ elements, the bidual cokernel is $r^2 - 1$, and the automorphism group acquires the monomial constraint of Theorem 6.3. The present paper therefore treats the intrinsic structure of the whole family $R_{\mathbf{n}}$, and the specialization $\ell = \iota = 2$ serves only as a consistency check. A second and opposite limitation appears in the constacyclic-code literature over rings written as $\mathbb{F}_{q^2} + u\mathbb{F}_{q^2}$: When u^2 is a unit, as in [18], where $u^2 = w^2$ for a generator w of $\mathbb{F}_{q^2}^\times$, the ring is not local. It carries a pair of orthogonal idempotents, splits as a product of two fields, and is therefore Frobenius, so its duality theory reduces to the field case componentwise. The family (1.1) lies at the opposite extreme: Every nonunit is nilpotent, no nontrivial idempotent exists, and the ideal lattice is not a product of lattices of smaller rings. The questions answered below, which ideals occur, what annihilation retains, and which cyclic modules embed into their biduals, are consequently inaccessible to idempotent decomposition. These works motivate the ring-theoretic foundations developed here, but no coding-theoretic claim is made without a separate study of polynomial quotients and code duality over $R_{\mathbf{n}}$.

Although multiplication in $R_{\mathbf{n}}$ is branchwise, the ideal lattice is not the Cartesian product of the branch chains. Once the branch intersections are fixed, terminal classes may be coupled by an additional subspace. For two branches this subspace is either zero or a non-coordinate projective point. For $r \geq 3$, higher-dimensional admissible subspaces occur, so the corresponding lattice fibers contain genuine Grassmannian strata. This transition is the source of the new multi-branch phenomena.

Our principal tool is a *branch–subspace calculus*. A proper ideal is represented canonically by a pair $(\mathbf{a}, W)$, where the cutoff vector $\mathbf{a}$ records the branch intersections and the admissible subspace W records the terminal coupling. Transition maps transport terminal data between cutoff layers, while coordinate saturation restores canonical form after addition. These operations turn ideal classification into an explicit model of the entire lattice and provide the input for the generator, annihilator, module-duality, and automorphism theories.

The contributions are organized in the order required by the later constructions.

- (i) Theorem 3.2 classifies all proper ideals by canonical cutoff–subspace data. Propositions 3.3 and 3.6 then determine the full order, intersection, and sum operations, while Proposition 3.7 identifies the projective and Grassmannian strata. Corollary 3.10 converts this into a closed count of $\text{Idl}(R_{\mathbf{n}})$ for every r , validated in Remark 3.11.
- (ii) Proposition 4.1 computes I/JI and a minimal generating set. This yields the compression invariant, the support-refined subspace count, the bivariate transform in Theorem 4.6, and the principal-ideal parameterization in Proposition 4.7.
- (iii) Theorem 5.2 determines annihilators from canonical data. Theorem 5.5 and Proposition 5.6 compute the closure, its defect, and its distributive fixed-point lattice. Proposition 5.14 identifies the same closure as the obstruction to cyclic bidual injectivity, and Theorems 5.16 and 5.19 classify torsionless cyclic quotients and establish the universal cokernel dimension $r^2 - 1$.
- (iv) Propositions 6.1 and 6.2 determine the linear image and the nonlinear kernel of the automorphism action. Theorem 6.3 assembles the full group, and Theorem 6.5 exhibits an explicit subgroup that acts trivially on ideals and determines the induced action on the depth-two sublattice.

These contributions are linked by four mechanisms: transition and saturation of terminal data; compression of branch heads by coupled generators; collapse of that geometry under annihilation and cyclic biduality; and the separation of automorphisms into tangent, socle, and depth-two shadows. The resulting theory therefore records what the branch decomposition alone cannot detect.

The dependencies among the main results are the following, each arrow meaning “is used in the proof of”:

$$\begin{aligned} \text{Thm. 3.2} &\longrightarrow \{\text{Props. 3.3–3.7, Cor. 3.10, Prop. 4.1, Thm. 5.2}\}, \\ \{\text{Props. 4.1, 4.4}\} &\longrightarrow \{\text{Prop. 4.3, Thm. 4.6, Prop. 4.7}\}, \\ \text{Thm. 5.2} &\longrightarrow \text{Thm. 5.5} \longrightarrow \{\text{Prop. 5.6, Prop. 5.14, Thm. 5.19}\}, \\ \text{Prop. 5.14} &\longrightarrow \{\text{Prop. 5.15, Thm. 5.16}\}, \\ \{\text{Props. 6.1, 6.2}\} &\longrightarrow \text{Thm. 6.3} \longrightarrow \{\text{Prop. 6.4, Thm. 6.5, Cor. 6.7}\}. \end{aligned}$$

For $r = 2$, the ring becomes $R_{(\ell, \iota)} \cong k[u, v]/\langle u^\ell, v^\iota, uv \rangle$, in the presentation customary for that case. Corollary 3.10 recovers the direct and mixed ideal families and gives $|\text{Idl}(R_{(\ell, \iota)})| = q(\ell - 1)(\iota - 1) + \ell + \iota$. This specialization is used as a consistency check, whereas all principal constructions and results are formulated for arbitrary r .

Section 2 fixes the notation and establishes the intrinsic branch profile together with the isomorphism classification. Four principal sections follow. Section 3 develops the branch–subspace calculus and the lattice operations. Section 4 treats generator sets and refined enumeration. Section 5 develops annihilator duality and its cyclic-module application. Section 6 determines the automorphism group and its induced actions. The display above records the dependencies among the main results; a reader interested only in the ideal lattice may stop after Corollary 3.10. All enumerative and structural formulas below were confirmed by exhaustive machine computation over the profiles recorded in the examples and remarks; the scripts are available on request.

2. Preliminaries and the branch profile

The first task is to recover the branch data intrinsically, before any ideal classification can depend on chosen generators. This section fixes the global notation, determines the local and graded structure of $R_{\mathbf{n}}$, and shows that the profile reconstructs the branch-length multiset. The resulting isomorphism classification supplies the cutoff ranges used in the ideal calculus and the equal-length blocks used later in the automorphism group.

Throughout, $k = \mathbb{F}_q$, where $q = p^m$, and $\mathbf{n} = (n_1, \dots, n_r)$ with $n_i \geq 2$. We use the same symbols x_i for the residue classes of the indeterminates in $R_{\mathbf{n}}$ and set $S_i = k[x_i]/\langle x_i^{n_i} \rangle$ and $A_i = x_i S_i$. For a finite commutative ring R , we write $\text{Idl}(R)$ for its lattice of ideals, ordered by inclusion, and $\text{Idl}(R)$ always includes the zero ideal and R itself.

The following proposition consolidates the presentation-dependent facts that are used throughout the paper.

Proposition 2.1. *Every element of $R_{\mathbf{n}}$ has a unique expression*

$$a + \sum_{i=1}^r \sum_{t=1}^{n_i-1} a_{i,t} x_i^t, \quad a, a_{i,t} \in k. \quad (2.1)$$

Moreover, $R_{\mathbf{n}} \cong S_1 \times_k \cdots \times_k S_r$, where the fiber product is taken with respect to the residue maps $S_i \rightarrow k$. In particular, $\dim_k R_{\mathbf{n}} = 1 + \sum_{i=1}^r (n_i - 1)$ and $|R_{\mathbf{n}}| = q^{1 + \sum_{i=1}^r (n_i - 1)}$. The ring is local with Jacobson radical $J = \langle x_1, \dots, x_r \rangle = A_1 \oplus \cdots \oplus A_r$, $R_{\mathbf{n}}/J \cong k$, and an element is a unit precisely when its constant term in (2.1) is nonzero. For $t \geq 1$,

$$J^t = \bigoplus_{i=1}^r \langle x_i^t, x_i^{t+1}, \dots, x_i^{n_i-1} \rangle, \quad (2.2)$$

with the i th summand interpreted as zero for $t \geq n_i$. Consequently, $\text{edim}(R_{\mathbf{n}}) = r$, $\text{Soc}(R_{\mathbf{n}}) = \text{Ann}(J) = \bigoplus_{i=1}^r kx_i^{n_i-1}$, and the nilpotency index of J is $\max_i n_i$. With $\deg x_i = 1$, the algebra is standard graded, the canonical surjection induces an isomorphism $\text{gr}_J(R_{\mathbf{n}}) \cong R_{\mathbf{n}}$, and the Hilbert series is $\text{Hilb}_{R_{\mathbf{n}}}(z) = 1 + \sum_{i=1}^r (z + z^2 + \cdots + z^{n_i-1})$.

Proof. The defining ideal is generated by monomials, so the images of the monomials 1 and x_i^t with $1 \leq t \leq n_i - 1$ form a k -basis; this gives (2.1) and the dimension count. The map $R_{\mathbf{n}} \rightarrow S_1 \times_k \cdots \times_k S_r$ sending x_i to the branch generator of S_i is a surjective k -algebra homomorphism between spaces of equal dimension, hence an isomorphism. An element with nonzero constant term a can be written as $a(1 - z)$ with $z \in J$ nilpotent, so it is invertible; conversely, every element of J is nilpotent. Thus, J is the set of nonunits, the ring is local, and $R_{\mathbf{n}}/J \cong k$. Since $A_i A_j = 0$ for $i \neq j$, powers of J are computed branchwise, which gives (2.2) and hence the socle, the embedding dimension, and the nilpotency index. Finally, the defining ideal is homogeneous for $\deg x_i = 1$, and the degree- t component is precisely a complement of J^{t+1} in J^t . Hence the J -adic associated graded algebra is canonically isomorphic to $R_{\mathbf{n}}$, and the displayed basis gives the Hilbert series. $\square$

The assertions in Proposition 2.1 are used below without further comment. To encode the intrinsic data compactly, we use the following profile.

Definition 2.2. The branch profile of $R_{\mathbf{n}}$ is the finite nonincreasing sequence $\beta_{\mathbf{n}}(t) = \dim_k J^t/J^{t+1} = |\{i : n_i > t\}|$, $t \geq 1$. For $d \geq 2$, write $\mu_{\mathbf{n}}(d) = |\{i : n_i = d\}|$ for the multiplicity of the branch length d .

The profile is the Hilbert function of the associated graded algebra in positive degrees. Its relevance here is that it can be inverted exactly:

$$\mu_{\mathbf{n}}(d) = \beta_{\mathbf{n}}(d-1) - \beta_{\mathbf{n}}(d), \quad d \geq 2, \quad (2.3)$$

where $\beta_{\mathbf{n}}(1) = r$ because every n_i is at least 2. Thus, the successive drops of the layers record where branches terminate. Equivalently, the Ferrers diagram with row lengths $n_i - 1$ has column lengths $\beta_{\mathbf{n}}(t)$; the branch-length multiset and the branch profile are conjugate partitions of the same integer $\dim_k J$.

The socle dimension gives the first structural consequence of the branch decomposition.

Corollary 2.3. The following conditions are equivalent:

$$\begin{aligned} R_{\mathbf{n}} \text{ is a chain ring} &\iff R_{\mathbf{n}} \text{ is Frobenius} \\ &\iff R_{\mathbf{n}} \text{ is zero-dimensional Gorenstein} &\iff r = 1. \end{aligned}$$

Hence every genuinely multi-branch case $r \geq 2$ is non-chain and non-Frobenius.

Proof. If $r = 1$, then $R_{\mathbf{n}}$ is a truncated polynomial chain ring and is therefore Frobenius and zero-dimensional Gorenstein. Conversely, Proposition 2.1 gives $\dim_k \text{Soc}(R_{\mathbf{n}}) = r$. For a finite commutative local algebra, the Frobenius and zero-dimensional Gorenstein conditions are both equivalent to a one-dimensional socle; see, for example, [1, Ch. XV] and [12, Sec. 3]. Hence either property forces $r = 1$. If $r \geq 2$, the branch ideals $\langle x_i \rangle$ and $\langle x_j \rangle$ are incomparable for $i \neq j$, so the ideal lattice is not a chain. This proves all equivalences. $\square$

To pass from k -algebra isomorphisms to abstract ring isomorphisms, we first identify the coefficient field intrinsically.

Lemma 2.4. *The coefficient field is characterized inside the abstract ring by*

$$k = \{y \in R_{\mathbf{n}} : y^q = y\}. \quad (2.4)$$

Consequently, every ring automorphism preserves k setwise and restricts to an element of $\text{Gal}(k/\mathbb{F}_p)$.

Proof. Write $y = a + z$ with $a \in k$ and $z \in J$. Since q is a power of the characteristic and $a^q = a$, $y^q - y = z^q - z = z(z^{q-1} - 1)$. If $y^q = y$, then $z = 0$, because z^{q-1} is nilpotent and hence $z^{q-1} - 1$ is a unit. Thus, $y \in k$, which proves (2.4). The set on the right-hand side is intrinsically defined, so it is preserved by every ring automorphism. $\square$

The profile inversion now yields the principal result of this section.

Theorem 2.5. *Let $\mathbf{n} = (n_1, \dots, n_r)$ and $\mathbf{n}' = (n'_1, \dots, n'_{r'})$, with all entries at least 2, and let both rings be defined over $k = \mathbb{F}_q$. The following are equivalent:*

- (a) $R_{\mathbf{n}} \cong R_{\mathbf{n}'}$ as rings;
- (b) $\beta_{\mathbf{n}}(t) = \beta_{\mathbf{n}'}(t)$ for every $t \geq 1$;
- (c) $r = r'$ and $\{n_1, \dots, n_r\} = \{n'_1, \dots, n'_{r'}\}$ as multisets.

Thus, the branch profile is a complete isomorphism invariant within the family.

Proof. A ring isomorphism preserves the Jacobson radical and all its powers. By Lemma 2.4, it maps k onto k and is therefore semilinear over k on every quotient; hence it preserves the dimensions $\dim_k J^t/J^{t+1}$, which proves (a) $\Rightarrow$ (b). Formula (2.3) shows that the profile determines the multiplicity of every branch length, and $\beta_{\mathbf{n}}(1) = r$ recovers the number of branches; hence (b) $\Rightarrow$ (c). Finally, equality of the branch-length multisets permits a reindexing for which $n_i = n'_i$ for every i ; the corresponding permutation of the variables induces an isomorphism, which proves (c) $\Rightarrow$ (a). $\square$

Remark 2.6. *The branch lengths $n_i - 1$ and the positive Loewy-layer dimensions $\beta_{\mathbf{n}}(t)$ are conjugate partitions of $\dim_k J$. Thus, the radical-layer dimensions reconstruct rather than merely constrain the presentation. In particular, $k[u, v]/\langle u^\ell, v^\ell, uv \rangle \cong k[u, v]/\langle u^{\ell'}, v^{\ell'}, uv \rangle$ if and only if $\{\ell, \iota\} = \{\ell', \iota'\}$. At the two boundary profiles, $r = 1$ gives a truncated chain ring, whereas $n_i = 2$ for every i gives $J^2 = 0$ and makes every k -subspace of J an ideal. Between these extremes, non-coordinate terminal subspaces supply the additional data developed in the next section.*

3. Ideal classification and the branch-subspace calculus

The intrinsic branch profile determines the ambient ring, but branch intersections alone do not determine an ideal. Terminal classes may couple several branches, and that additional finite-geometric datum must be retained under containment and lattice operations. This section introduces canonical cutoff-subspace coordinates, proves their uniqueness, and develops transition and saturation operations that compute containment, intersection, and sum. The resulting lattice model also isolates the projective and Grassmannian fibers that first appear beyond the two-branch case.

Let $C_n = \prod_{i=1}^r \{1, \dots, n_i\}$ be the cutoff poset, ordered coordinatewise. For $\mathbf{a} = (a_1, \dots, a_r) \in C_n$, define the direct cutoff ideal $D_{\mathbf{a}} = \bigoplus_{i=1}^r \langle x_i^{a_i} \rangle$, where $\langle x_i^{n_i} \rangle = 0$. Its active set is $S(\mathbf{a}) = \{i : a_i \geq 2\}$. The terminal space at $\mathbf{a}$ is $E_{\mathbf{a}} = \bigoplus_{i \in S(\mathbf{a})} k e_i^{\mathbf{a}}, e_i^{\mathbf{a}} = x_i^{a_i-1} + D_{\mathbf{a}} \in R_n/D_{\mathbf{a}}$. Thus, $E_{\mathbf{a}}$ is the subspace annihilated by J in the branch quotient $J/D_{\mathbf{a}}$.

Definition 3.1. A subspace $W \leq E_{\mathbf{a}}$ is called $\mathbf{a}$ -admissible if

$$W \cap k e_i^{\mathbf{a}} = 0 \quad \text{for every } i \in S(\mathbf{a}). \quad (3.1)$$

Let $\widetilde{W}$ be the unique lift of W to $\bigoplus_{i \in S(\mathbf{a})} k x_i^{a_i-1}$. The ideal datum $(\mathbf{a}, W)$ determines the k -subspace

$$I(\mathbf{a}, W) = D_{\mathbf{a}} \oplus \widetilde{W}. \quad (3.2)$$

The exclusion of the coordinate lines in (3.1) is a minimality condition. If $e_i^{\mathbf{a}} \in W$, then $x_i^{a_i-1} \in I(\mathbf{a}, W)$ and the i th cutoff must be lowered. Thus, admissibility is what makes the cutoff vector intrinsic and unique.

Theorem 3.2. The assignment $(\mathbf{a}, W) \mapsto I(\mathbf{a}, W)$ is a bijection from the set of pairs consisting of $\mathbf{a} \in C_n$ and an $\mathbf{a}$ -admissible subspace $W \leq E_{\mathbf{a}}$ onto the proper ideals of R_n . For the ideal $I = I(\mathbf{a}, W)$, one has

$$I \cap A_i = \langle x_i^{a_i} \rangle \quad (1 \leq i \leq r), \quad (3.3)$$

and

$$\dim_k I = \sum_{i=1}^r (n_i - a_i) + \dim_k W. \quad (3.4)$$

Proof. Let I be a proper ideal. Since R_n is local, $I \subseteq J$. For every i , the branch intersection $I \cap A_i$ is stable under multiplication by $k[x_i]$ and hence is uniquely of the form $\langle x_i^{a_i} \rangle$ for some $1 \leq a_i \leq n_i$. Therefore, $D_{\mathbf{a}} \subseteq I$ for a unique cutoff vector $\mathbf{a}$.

Write an element of I as $y = \sum_i f_i(x_i)$ with $f_i(x_i) \in A_i$. Multiplication by x_i gives $x_i f_i(x_i) \in I \cap A_i = \langle x_i^{a_i} \rangle$. Modulo $\langle x_i^{a_i} \rangle$, the kernel of multiplication by x_i on $A_i/\langle x_i^{a_i} \rangle$ is zero if $a_i = 1$ and is spanned by $e_i^{\mathbf{a}}$ if $a_i \geq 2$. Hence, $I/D_{\mathbf{a}}$ is a subspace W of $E_{\mathbf{a}}$. It cannot contain a coordinate line, since otherwise $x_i^{a_i-1} \in I \cap A_i$, contradicting the minimality of a_i . Thus, W is admissible and $I = I(\mathbf{a}, W)$.

Conversely, let $W \leq E_{\mathbf{a}}$ be admissible. For $w = \sum_i c_i x_i^{a_i-1} \in \widetilde{W}$, the defining relations give $x_j w = c_j x_j^{a_j} \in D_{\mathbf{a}}$ ($1 \leq j \leq r$). Therefore, $J\widetilde{W} \subseteq D_{\mathbf{a}}$, and since $R_n = k \oplus J$, the k -subspace $I(\mathbf{a}, W)$ is an ideal. Admissibility gives $I(\mathbf{a}, W) \cap A_i = \langle x_i^{a_i} \rangle$, which proves uniqueness of $\mathbf{a}$; the quotient by $D_{\mathbf{a}}$ then recovers W . Finally, (3.4) follows from the direct sum in (3.2).

Finally, the bijection is onto the *proper* ideals only: Every $I(\mathbf{a}, W)$ lies in J , the largest one being $I((1, \dots, 1), 0) = J$, whereas R_n is the unique ideal not contained in J , so exactly one element of $\text{Idl}(R_n)$ is omitted and every count below has the form $|\text{Idl}(R_n)| = 1 + \#\{(\mathbf{a}, W)\}$. $\square$

The representation theorem becomes a lattice model only after the terminal spaces at different cutoffs are connected. Suppose that $\mathbf{a} \geq \mathbf{b}$ coordinatewise. Define the cutoff transition map $\tau_{\mathbf{a},\mathbf{b}} : E_{\mathbf{a}} \longrightarrow E_{\mathbf{b}}$,

$$\tau_{\mathbf{a},\mathbf{b}}(e_i^{\mathbf{a}}) = \begin{cases} e_i^{\mathbf{b}}, & a_i = b_i, \\ 0, & a_i > b_i. \end{cases}$$

Note that $a_i = b_i$ and $i \in S(\mathbf{a})$ force $i \in S(\mathbf{b})$, so the map is well defined. These maps satisfy $\tau_{\mathbf{a},\mathbf{a}} = \text{id}_{E_{\mathbf{a}}}$, $\tau_{\mathbf{a},\mathbf{c}} = \tau_{\mathbf{b},\mathbf{c}}\tau_{\mathbf{a},\mathbf{b}}$ ($\mathbf{a} \geq \mathbf{b} \geq \mathbf{c}$), and they form the transition system that records which terminal coordinates survive when a direct cutoff ideal is enlarged.

Proposition 3.3. *Let $I(\mathbf{a}, W)$ and $I(\mathbf{b}, W')$ be proper ideals. Then*

$$I(\mathbf{a}, W) \subseteq I(\mathbf{b}, W') \iff \mathbf{a} \geq \mathbf{b} \text{ and } \tau_{\mathbf{a},\mathbf{b}}(W) \subseteq W'. \quad (3.5)$$

Consequently, the proper ideal lattice is isomorphic to the poset of admissible branch–subspace data equipped with the order on the right-hand side of (3.5).

Proof. If $I(\mathbf{a}, W) \subseteq I(\mathbf{b}, W')$, then (3.3) implies $\langle x_i^{a_i} \rangle \subseteq \langle x_i^{b_i} \rangle$ for every i ; hence $\mathbf{a} \geq \mathbf{b}$. Reduction modulo $D_{\mathbf{b}}$ sends the terminal class $e_i^{\mathbf{a}}$ to $e_i^{\mathbf{b}}$ when $a_i = b_i$ and to zero when $a_i > b_i$. Thus, the induced map on terminal data is exactly $\tau_{\mathbf{a},\mathbf{b}}$, and inclusion of the ideals forces $\tau_{\mathbf{a},\mathbf{b}}(W) \subseteq W'$.

Conversely, if these two conditions hold, then $D_{\mathbf{a}} \subseteq D_{\mathbf{b}}$ and every lifted element of W has, modulo $D_{\mathbf{b}}$, a terminal class in W' . Therefore, $I(\mathbf{a}, W) \subseteq I(\mathbf{b}, W')$. $\square$

The containment criterion describes the ideal lattice as an admissible-subspace system over the reversed cutoff poset. The next construction is needed because sums of admissible data need not remain admissible.

Definition 3.4. *Let $U \leq E_{\mathbf{a}}$ be an arbitrary subspace, not necessarily admissible. Its coordinate set is $\Omega_{\mathbf{a}}(U) = \{i \in S(\mathbf{a}) : \ker e_i^{\mathbf{a}} \subseteq U\}$. Define a new cutoff vector $\mathbf{a}^U$ by*

$$a_i^U = \begin{cases} a_i - 1, & i \in \Omega_{\mathbf{a}}(U), \\ a_i, & i \notin \Omega_{\mathbf{a}}(U), \end{cases}$$

so that $\mathbf{a} \geq \mathbf{a}^U$, and put $U^\circ = \tau_{\mathbf{a},\mathbf{a}^U}(U)$. The pair $(\mathbf{a}^U, U^\circ)$ is called the coordinate saturation of the raw datum $(\mathbf{a}, U)$.

The transition map $\tau_{\mathbf{a},\mathbf{a}^U}$ deletes exactly the coordinates indexed by $\Omega_{\mathbf{a}}(U)$ and preserves all remaining terminal coordinates, so no separate notation is required for saturation. Coordinate saturation has the following canonical normalization property.

Proposition 3.5. *For every $\mathbf{a} \in C_{\mathbf{n}}$ and every subspace $U \leq E_{\mathbf{a}}$, the saturated subspace U° is $\mathbf{a}^U$ -admissible and*

$$D_{\mathbf{a}} + \widetilde{U} = I(\mathbf{a}^U, U^\circ). \quad (3.6)$$

Thus, coordinate saturation is the unique normalization that converts arbitrary terminal data into canonical ideal data.

Proof. Let $\Omega = \Omega_{\mathbf{a}}(U)$. Since $ke_i^{\mathbf{a}} \subseteq U$ for every $i \in \Omega$, $D_{\mathbf{a}U} = D_{\mathbf{a}} + \sum_{i \in \Omega} kx_i^{a_i-1} \subseteq D_{\mathbf{a}} + \widetilde{U}$. After these coordinate directions are absorbed into the direct cutoff ideal, the remaining terminal information is precisely U° , which proves (3.6).

It remains to verify admissibility. Every coordinate indexed by Ω has zero image under $\tau_{\mathbf{a}, \mathbf{a}U}$, so U° has zero component along $e_i^{\mathbf{a}U}$ for every $i \in \Omega$ that survives in $S(\mathbf{a}^U)$. Suppose that, for some $j \notin \Omega$, the line $ke_j^{\mathbf{a}U}$ were contained in U° . Then some element of U would differ from $e_j^{\mathbf{a}}$ by a vector in the span of the coordinate lines indexed by Ω . That span is already contained in U , so $e_j^{\mathbf{a}} \in U$, contradicting $j \notin \Omega$. Hence, U° is admissible, and uniqueness follows from Theorem 3.2. $\square$

We can now express the lattice operations directly in canonical coordinates. For $\mathbf{a}, \mathbf{b} \in C_n$, let $\mathbf{a} \vee \mathbf{b}$ and $\mathbf{a} \wedge \mathbf{b}$ denote the coordinatewise maximum and minimum, $(\mathbf{a} \vee \mathbf{b})_i = \max\{a_i, b_i\}$, $(\mathbf{a} \wedge \mathbf{b})_i = \min\{a_i, b_i\}$.

Proposition 3.6. *Let $I = I(\mathbf{a}, W)$ and $K = I(\mathbf{b}, W')$ be proper ideals. Then the meet is*

$$I \cap K = I(\mathbf{a} \vee \mathbf{b}, Y), \quad Y = \tau_{\mathbf{a} \vee \mathbf{b}, \mathbf{a}}^{-1}(W) \cap \tau_{\mathbf{a} \vee \mathbf{b}, \mathbf{b}}^{-1}(W'), \quad (3.7)$$

and Y is automatically $(\mathbf{a} \vee \mathbf{b})$ -admissible.

For the join, set

$$U = \tau_{\mathbf{a}, \mathbf{a} \wedge \mathbf{b}}(W) + \tau_{\mathbf{b}, \mathbf{a} \wedge \mathbf{b}}(W') \leq E_{\mathbf{a} \wedge \mathbf{b}}. \quad (3.8)$$

If $((\mathbf{a} \wedge \mathbf{b})^U, U^\circ)$ is the coordinate saturation of $(\mathbf{a} \wedge \mathbf{b}, U)$, then

$$I + K = I((\mathbf{a} \wedge \mathbf{b})^U, U^\circ). \quad (3.9)$$

Proof. The branch intersection of $I \cap K$ with A_i is $\langle x_i^{a_i} \rangle \cap \langle x_i^{b_i} \rangle = \langle x_i^{\max\{a_i, b_i\}} \rangle$, so the cutoff vector of the meet is $\mathbf{a} \vee \mathbf{b}$. A terminal class in $E_{\mathbf{a} \vee \mathbf{b}}$ belongs to the quotient of $I \cap K$ precisely when its images under the two transition maps belong to W and W' , respectively. This gives (3.7). The resulting subspace is admissible because the branch intersections of $I \cap K$ are exactly the displayed cutoff ideals.

For the sum, one has $D_{\mathbf{a}} + D_{\mathbf{b}} = D_{\mathbf{a} \wedge \mathbf{b}}$. Modulo this direct ideal, the terminal classes contributed by I and K are exactly the two images in (3.8), because the lifted components dropped by the transition maps already lie in $D_{\mathbf{a} \wedge \mathbf{b}}$. Hence, $I + K = D_{\mathbf{a} \wedge \mathbf{b}} + \widetilde{U}$. Applying Proposition 3.5 yields (3.9). $\square$

The saturation step in the join formula is essential. New coordinate lines can appear when two mixed subspaces are added, and each such line lowers the corresponding branch cutoff. Thus, the lattice join is not obtained by taking only the coordinatewise minimum and the vector-space sum. This is the first operation in which the branch and geometric data interact nontrivially.

Proposition 3.7. *For $\mathbf{a} \in C_n$, define $a_i^- = \max\{1, a_i - 1\}$. Then the map $U \mapsto D_{\mathbf{a}} + \widetilde{U}$ is a lattice isomorphism from the full subspace lattice of $E_{\mathbf{a}}$ onto the interval $[D_{\mathbf{a}}, D_{\mathbf{a}^-}] \subseteq \text{Idl}(R_n)$. Under this isomorphism, the ideals whose canonical cutoff vector is exactly $\mathbf{a}$ correspond precisely to the admissible subspaces of $E_{\mathbf{a}}$.*

Proof. The quotient $D_{\mathbf{a}^-}/D_{\mathbf{a}}$ is naturally identified with $E_{\mathbf{a}}$. Ideals between $D_{\mathbf{a}}$ and $D_{\mathbf{a}^-}$ are exactly the k -subspaces of this quotient, because the quotient is annihilated by J . This proves the interval isomorphism. The final assertion is exactly the minimality condition encoded by admissibility. $\square$

The interval description now yields the global lattice properties.

Corollary 3.8. *The ideal lattice of $R_{\mathbf{n}}$ is ranked by $\text{rk}(I) = \dim_k I$, with $R_{\mathbf{n}}$ of rank $\dim_k R_{\mathbf{n}}$. It is modular for every r , and it is distributive if and only if $r = 1$. The rank-one elements of the interval $[D_{\mathbf{a}}, D_{\mathbf{a}^-}]$ form the projective space $\text{PG}(|S(\mathbf{a})| - 1, q)$; if $|S(\mathbf{a})| \geq 3$, then higher-dimensional Grassmannian strata occur in the same interval.*

Proof. If $I \subseteq K$ is a cover, then K/I is a simple module over the local ring $R_{\mathbf{n}}$ and is therefore isomorphic to the residue field k . Hence every cover raises the k -dimension by one, so $\dim_k I$ is a rank function. The ideal lattice is a submodule lattice of the regular module and is therefore modular. If $r = 1$, it is the chain of radical powers. If $r \geq 2$, choose $\mathbf{a}$ with $a_i = 2$ for two indices. By Proposition 3.7, the corresponding interval contains the subspace lattice of k^2 , which contains the nondistributive diamond M_3 . The statements about projective and Grassmannian strata follow from the same interval isomorphism. $\square$

The next example exhibits both a higher-dimensional coupling and the cutoff change forced by saturation.

Example 3.9. *Let $k = \mathbb{F}_3$, let $\mathbf{n} = (4, 3, 3)$, and take $\mathbf{a} = (3, 2, 2)$. Then $D_{\mathbf{a}} = \langle x_1^3, x_2^2, x_3^2 \rangle$ and $E_{\mathbf{a}} = ke_1 \oplus ke_2 \oplus ke_3$, where $e_1 = x_1^2 + D_{\mathbf{a}}$, $e_2 = x_2 + D_{\mathbf{a}}$, and $e_3 = x_3 + D_{\mathbf{a}}$. The plane $W = \text{span}_{\mathbb{F}_3}\{e_1 + e_3, e_2 + e_3\}$ contains no coordinate line and is therefore admissible. It defines the ideal $I(\mathbf{a}, W) = \langle x_1^3, x_2^2, x_3^2, x_1^2 + x_3, x_2 + x_3 \rangle$, which has dimension 5 over $\mathbb{F}_3$ by (3.4). The two-dimensional parameter W is a coupling phenomenon with no analogue for two branches.*

The join operation can force a cutoff change. Let $W_1 = \text{span}\{e_1 + e_2\}$ and $W_2 = \text{span}\{e_1 - e_2\}$. Both are admissible, but $W_1 + W_2 = \text{span}\{e_1, e_2\}$ contains two coordinate lines. Coordinate saturation therefore gives $I(\mathbf{a}, W_1) + I(\mathbf{a}, W_2) = D_{(2,1,2)} = \langle x_1^2, x_2, x_3^2 \rangle$. This illustrates why the saturation operation is indispensable in the lattice calculus.

The rank-two specialization collapses the Grassmannian fibers to projective points, and the classification also yields a closed ideal count for every r . For $s \geq 0$ let $\mathcal{A}_s(q)$ be the number of subspaces of k^s containing no coordinate line, that is, the number of admissible subspaces of an s -dimensional terminal space; refining it by dimension and exact support through Proposition 4.4 and (5.11) gives $\mathcal{A}_s(q) = \sum_{\theta=0}^s \binom{s}{\theta} F_{\theta}(q)$, so that $\mathcal{A}_0 = \mathcal{A}_1 = 1$, $\mathcal{A}_2 = q$, $\mathcal{A}_3 = q(2q - 1)$, and $\mathcal{A}_4 = q(q^3 + 3q^2 - 4q + 1)$.

Corollary 3.10. (i) *For every $\mathbf{n} = (n_1, \dots, n_r)$ with $n_i \geq 2$,*

$$|\text{Idl}(R_{\mathbf{n}})| = 1 + \sum_{S \subseteq [r]} \mathcal{A}_{|S|}(q) \prod_{i \in S} (n_i - 1). \quad (3.10)$$

Among these, the $\prod_{i=1}^r n_i$ direct ideals $D_{\mathbf{a}}$ are exactly those with zero terminal subspace, so the coupled ideals contribute the surplus $\sum_{|S| \geq 2} (\mathcal{A}_{|S|}(q) - 1) \prod_{i \in S} (n_i - 1)$, strictly positive as soon as $r \geq 2$.

(ii) *Let $r = 2$ and $\mathbf{n} = (\ell, \iota)$, so that $R_{(\ell, \iota)} \cong k[u, v]/\langle u^{\ell}, v^{\iota}, uv \rangle$ under $u \mapsto x_1$ and $v \mapsto x_2$. Then every admissible terminal subspace is either zero or a non-coordinate line, so the proper ideals are exactly the direct ideals $\langle x_1^a, x_2^b \rangle$, $1 \leq a \leq \ell$, $1 \leq b \leq \iota$, and the mixed ideals $\langle x_1^a, x_2^b, x_1^{a-1} + \lambda x_2^{b-1} \rangle$, $2 \leq a \leq \ell$, $2 \leq b \leq \iota$, $\lambda \in k^{\times}$, and (3.10) becomes*

$$|\text{Idl}(R_{(\ell, \iota)})| = q(\ell - 1)(\iota - 1) + \ell + \iota. \quad (3.11)$$

Proof. By Theorem 3.2 the proper ideals are partitioned by their cutoff vector, and the number of admissible subspaces of $E_{\mathbf{a}}$ depends only on $\dim_k E_{\mathbf{a}} = |S(\mathbf{a})|$, so it equals $\mathcal{A}_{|S(\mathbf{a})|}(q)$. For a fixed $S \subseteq [r]$ the cutoff vectors with $S(\mathbf{a}) = S$ are those with $a_i \in \{2, \dots, n_i\}$ for $i \in S$ and $a_i = 1$ otherwise, so there are $\prod_{i \in S} (n_i - 1)$ of them; summing over S and adding $R_{\mathbf{n}}$ gives (3.10). The direct ideals are those with $W = 0$, one for each of the $\prod_i n_i$ cutoff vectors, and $\prod_{i=1}^r n_i = \sum_{S \subseteq [r]} \prod_{i \in S} (n_i - 1)$ gives the surplus, which is positive because $\mathcal{A}_s(q) = 1$ for $s \leq 1$ while $\mathcal{A}_s(q) \geq q$ for $s \geq 2$.

For (ii), a terminal space has dimension at most two. An admissible subspace cannot be the whole space, which contains both coordinate lines, and cannot be a coordinate line; its only nonzero possibility is therefore a line spanned by $e_1 + \lambda e_2$ with $\lambda \in k^\times$. Theorem 3.2 gives the displayed families, and uniqueness of the canonical data makes them disjoint. There are $\ell\iota$ proper direct ideals, including 0, and $(q-1)(\ell-1)(\iota-1)$ mixed ones, so the total is $1 + \ell\iota + (q-1)(\ell-1)(\iota-1) = q(\ell-1)(\iota-1) + \ell + \iota$, which is what (3.10) returns with $\mathcal{A}_2(q) = q$. $\square$

Part (ii) locates the corollary relative to the literature: For $\ell = \iota = 2$ it returns $q + 4$, reproducing the lattice of [17] without adding to it, and it shows that the count is not multiplicative over the branches, the excess $(q-1)(\ell-1)(\iota-1)$ over the product $\ell\iota$ of the branch counts being the contribution of the projective points, that is, the two-branch shadow of the Grassmannian strata of Corollary 3.8.

Remark 3.11. Formula (3.10) is a convenient test of Theorem 3.2, because it can be checked against three independent descriptions of the lattice. For $r = 1$ it returns $n_1 + 1$, the chain count, and for $r = 2$ it returns (3.11), which part (ii) obtains by a different argument. In the square-zero profile $J^2 = 0$, so every k -subspace of $J \cong k^r$ is an ideal and $|\text{Idl}(R_{\mathbf{n}})| = 1 + G_r(q)$, where $G_r(q) = \sum_{d=0}^r \binom{r}{d}_q$ is the Galois number; on the other hand, (3.10) returns $1 + \sum_{s=0}^r \binom{r}{s} \mathcal{A}_s(q)$. The two agree, that is, $\sum_{s=0}^r \binom{r}{s} \mathcal{A}_s(q) = G_r(q)$, which is precisely the statement that every subspace of k^r has a unique coordinate saturation in the sense of Proposition 3.5. Independently of these three checks, we enumerated all ideals directly, as the J -stable subspaces of J , for the profiles of Table 1; in every case the canonical datum recovered from an ideal is unique and admissible and reproduces the ideal, and the total agrees with (3.10).

Table 1. Exhaustive verification of Theorem 3.2 and Corollary 3.10; “enumerated” counts all J -stable subspaces of J together with $R_{\mathbf{n}}$.

$\mathbf{n}$	(3, 3)	(4, 3)	(2, 2)	(2, 2, 2)	(3, 2, 2)	(3, 3, 2)
q	2, 3	2, 3	5	2, 3	2	2
enumerated	14, 18	19, 25	9	17, 29	28	47
(3.10)	14, 18	19, 25	9	17, 29	28	47

Example 3.12. Let $k = \mathbb{F}_2$ and $\mathbf{n} = (4, 3, 2, 2)$, so that $r = 4$ and $\dim_k R_{\mathbf{n}} = 8$; this is a boundary case in three respects at once, since the field is as small as possible, two branches have the minimal length $n_i = 2$, and one branch is cut off at $a_i = n_i$. Take $\mathbf{a} = (4, 2, 2, 2)$, so that $D_{\mathbf{a}} = \langle x_2^2 \rangle$ and $E_{\mathbf{a}} = ke_1 \oplus ke_2 \oplus ke_3 \oplus ke_4$ with $e_1 = x_1^3 + D_{\mathbf{a}}$ and $e_i = x_i + D_{\mathbf{a}}$ for $i = 2, 3, 4$; only the second branch has $a_i < n_i$, so $B(\mathbf{a}) = \{2\}$ in the notation of Section 4.

For $W = \text{span}_{\mathbb{F}_2}\{e_2 + e_3, e_1 + e_4\}$ the nonzero vectors are $e_2 + e_3$, $e_1 + e_4$, and $e_1 + e_2 + e_3 + e_4$, so W contains no coordinate line and $T(W) = \{1, 2, 3, 4\}$. Hence $I = I(\mathbf{a}, W) = \langle x_2 + x_3, x_1^3 + x_4 \rangle$ has $\dim_k I = 1 + 2 = 3$ by (3.4), the direct part arising through $x_2(x_2 + x_3) = x_2^2$, and (4.4) gives $\mu(I) = 2$

and $\kappa(I) = 2$: Four terminal coordinates are carried by two generators. Over $\mathbb{F}_2$ the projective torus of Proposition 4.7 is trivial by (4.13), so this coupling is a consequence of $r \geq 3$ and not of a large ground field.

Saturation becomes visible in a join. The admissible lines $W_1 = \text{span}_{\mathbb{F}_2}\{e_3 + e_4\}$ and $W_2 = \text{span}_{\mathbb{F}_2}\{e_2 + e_3 + e_4\}$ give $I(\mathbf{a}, W_1) = \langle x_3 + x_4, x_2^2 \rangle$ and $I(\mathbf{a}, W_2) = \langle x_2 + x_3 + x_4 \rangle$ by (4.5), and their terminal sum U contains $(e_3 + e_4) + (e_2 + e_3 + e_4) = e_2$. Hence $\Omega_{\mathbf{a}}(U) = \{2\}$, the second cutoff drops to 1, and Proposition 3.5 predicts $I(\mathbf{a}, W_1) + I(\mathbf{a}, W_2) = I((4, 1, 2, 2), \text{span}_{\mathbb{F}_2}\{e_3 + e_4\}) = \langle x_2, x_3 + x_4 \rangle$ of dimension 3, as direct computation confirms: The two mixed generators sum to x_2 , and with x_2 the whole second branch enters. Without saturation the datum $(\mathbf{a}, U)$ would describe the same subspace but would not be canonical, and its cutoff vector would no longer record $I \cap A_2 = \langle x_2 \rangle$. The annihilators of these ideals are computed in Example 5.23.

4. Minimal generating sets and enumeration

The preceding section classifies ideals and their lattice operations. We now study how terminal coupling affects minimal generators and rank. This section computes I/JI , derives canonical minimal generating sets and a compression invariant, and then obtains a bivariate generator–rank transform together with a projective-torus parameterization of principal ideals.

For $\mathbf{a} \in C_{\mathbf{n}}$, retain $S(\mathbf{a}) = \{i : a_i \geq 2\}$ and define $B(\mathbf{a}) = \{i : a_i < n_i\}$. For an $\mathbf{a}$ -admissible subspace $W \leq E_{\mathbf{a}}$, let

$$T(W) = \text{supp}(W) = \{i \in S(\mathbf{a}) : \pi_i(W) \neq 0\}, \quad (4.1)$$

where $\pi_i : E_{\mathbf{a}} \rightarrow ke_i^{\mathbf{a}}$ is the coordinate projection. Thus, $T(W)$ records the branches involved in the terminal coupling, whereas $B(\mathbf{a})$ records the nonzero direct branches. We write $\theta = |T(W)|$ and $d = \dim_k W$.

Proposition 4.1. *Let $I = I(\mathbf{a}, W)$ be a proper ideal of $R_{\mathbf{n}}$. Then*

$$JI = \bigoplus_{i \in B(\mathbf{a}) \cap T(W)} \langle x_i^{a_i} \rangle \oplus \bigoplus_{i \in B(\mathbf{a}) \setminus T(W)} \langle x_i^{a_i+1} \rangle, \quad (4.2)$$

where $\langle x_i^{n_i} \rangle = 0$. Consequently, there is a natural k -vector-space decomposition

$$I/JI \cong W \oplus \bigoplus_{i \in B(\mathbf{a}) \setminus T(W)} k(x_i^{a_i} + JI). \quad (4.3)$$

In particular, the minimal number of generators of I is

$$\mu(I) = \dim_k(I/JI) = \dim_k W + |B(\mathbf{a}) \setminus T(W)|. \quad (4.4)$$

More precisely, if $w_1, \dots, w_d$ is a basis of W and $\widetilde{w}_1, \dots, \widetilde{w}_d$ are the canonical lifts, then

$$\mathcal{G}(\mathbf{a}, W) = \{\widetilde{w}_1, \dots, \widetilde{w}_d\} \cup \{x_i^{a_i} : i \in B(\mathbf{a}) \setminus T(W)\} \quad (4.5)$$

is a minimal generating set of I .

Proof. From $I = D_{\mathbf{a}} \oplus \widetilde{W}$, one has $JI = JD_{\mathbf{a}} + J\widetilde{W}$. The direct part is computed branchwise. Since products between distinct branches vanish,

$$JD_{\mathbf{a}} = \bigoplus_{i \in B(\mathbf{a})} \langle x_i^{a_i+1} \rangle. \quad (4.6)$$

Now let $w = \sum_{j \in S(\mathbf{a})} c_j e_j^{\mathbf{a}} \in W$, with canonical lift $\widetilde{w} = \sum_j c_j x_j^{a_j-1}$. For each i the defining relations give

$$x_i \widetilde{w} = c_i x_i^{a_i}. \quad (4.7)$$

If $i \notin T(W)$, then every $w \in W$ has i th coordinate zero, so the i th branch contributes nothing to $J\widetilde{W}$. If $i \in T(W)$, there exists $w \in W$ with $c_i \neq 0$. When $a_i < n_i$, (4.7) shows that $x_i^{a_i} \in JI$, and hence the entire branch ideal $\langle x_i^{a_i} \rangle$ is contained in JI . When $a_i = n_i$, the product in (4.7) is zero and there is no direct branch part to contribute. Combining these observations with (4.6) proves (4.2).

The right-hand side of (4.2) lies in $D_{\mathbf{a}}$. Therefore, $\widetilde{W} \cap JI = 0$, and the image of $\widetilde{W}$ in I/JI is naturally identified with W . For $i \in B(\mathbf{a}) \cap T(W)$, the entire branch ideal $\langle x_i^{a_i} \rangle$ lies in JI and contributes nothing to the quotient. For $i \in B(\mathbf{a}) \setminus T(W)$, the corresponding quotient is $\langle x_i^{a_i} \rangle / \langle x_i^{a_i+1} \rangle \cong k$, generated by the class of $x_i^{a_i}$. Distinct branch quotients are independent. This proves (4.3) and hence (4.4), using the standard local identity $\mu(I) = \dim_k(I/JI)$.

It remains to verify the explicit generating set. The lifted basis elements generate $\widetilde{W}$. If $i \in B(\mathbf{a}) \cap T(W)$, choose a basis combination $w \in W$ whose i th coordinate is nonzero. Multiplying its lift by x_i yields a nonzero scalar multiple of $x_i^{a_i}$, so the ideal generated by the lifted basis contains $\langle x_i^{a_i} \rangle$. For $i \in B(\mathbf{a}) \setminus T(W)$, the generator $x_i^{a_i}$ is explicitly included. Thus, (4.5) generates every summand of $D_{\mathbf{a}}$ together with all of $\widetilde{W}$, and hence generates I . Its image in I/JI is exactly the basis displayed in (4.3), so no proper subset can generate I . Therefore, the set is minimal. $\square$

The formula separates two different sources of generators. A basis of W provides the coupled generators, while the branches not met by $T(W)$ retain independent direct heads. This leads to an intrinsic measure of the generator saving produced by branch coupling.

Definition 4.2. For $I = I(\mathbf{a}, W)$, define its head support and coupling compression by

$$H(I) = B(\mathbf{a}) \cup T(W), \quad \kappa(I) = |T(W)| - \dim_k W. \quad (4.8)$$

The compression identity isolates the exact saving produced by the coupling subspace.

Proposition 4.3. For every proper ideal $I = I(\mathbf{a}, W)$,

$$\mu(I) = |H(I)| - \kappa(I). \quad (4.9)$$

If $W \neq 0$, then

$$1 \leq \kappa(I) \leq |T(W)| - 1. \quad (4.10)$$

The upper bound is attained precisely when $\dim_k W = 1$.

Proof. Since $|H(I)| = |B(\mathbf{a}) \setminus T(W)| + |T(W)|$, formula (4.9) follows from (4.4). Suppose $W \neq 0$ and put $\theta = |T(W)|$ and $d = \dim_k W$. By definition, W is contained in the coordinate space $k^{T(W)}$ and has full support there, so $d \leq \theta$. If $d = \theta$, then $W = k^{T(W)}$, which contains every coordinate line, contradicting admissibility. Hence, $1 \leq d \leq \theta - 1$, and therefore $1 \leq \theta - d \leq \theta - 1$. This is (4.10). Finally, $\theta - d = \theta - 1$ if and only if $d = 1$. $\square$

Thus, nonzero terminal coupling always reduces the number of generators below the number of branch-head coordinates that occur in the ideal. A line with support on θ branches produces the maximal saving $\theta - 1$, while higher-dimensional couplings interpolate between this extremal compression and the uncoupled situation.

We next count the possible coupling spaces while retaining both their dimension and their exact coordinate support. This support refinement is essential for generator enumeration, because the term $|B(\mathbf{a}) \setminus T(W)|$ cannot be recovered from $\dim_k W$ alone.

For integers $\theta, d \geq 0$, let $C_{\theta,d}(q)$ denote the number of d -dimensional subspaces $W \leq k^\theta$ satisfying both

$$W \cap ke_i = 0 \quad \text{and} \quad \pi_i(W) \neq 0 \quad (1 \leq i \leq \theta). \quad (4.11)$$

Thus, W is admissible and has full coordinate support. We use the convention that Gaussian coefficients vanish outside their natural range.

Proposition 4.4. *For all $\theta, d \geq 0$,*

$$C_{\theta,d}(q) = \sum_{\substack{\alpha, \gamma \geq 0 \\ \alpha + \gamma \leq \theta}} (-1)^{\alpha + \gamma} \binom{\theta}{\alpha} \binom{\theta - \alpha}{\gamma} \left[\begin{matrix} \theta - \alpha - \gamma \\ d - \alpha \end{matrix} \right]_q. \quad (4.12)$$

In particular, $C_{0,0}(q) = 1, C_{\theta,0}(q) = C_{\theta,\theta}(q) = 0$ ($\theta \geq 1$), and

$$C_{\theta,1}(q) = (q - 1)^{\theta - 1} \quad (\theta \geq 2). \quad (4.13)$$

Proof. Let $V = k^\theta$ with coordinate lines ke_i and coordinate hyperplanes $\ker \pi_i$. Condition (4.11) excludes the two families of bad events $\mathcal{E}_i = \{W : ke_i \subseteq W\}$ and $\mathcal{F}_i = \{W : \pi_i(W) = 0\}$. Apply inclusion–exclusion simultaneously to these 2θ events. Choose a set $A \subseteq [\theta]$ of line-containment conditions and a set $G \subseteq [\theta]$ of hyperplane-containment conditions. If $A \cap G \neq \emptyset$, then no subspace can satisfy all selected conditions, since $ke_i \not\subseteq \ker \pi_i$. Assume $A \cap G = \emptyset$, and put $\alpha = |A|$ and $\gamma = |G|$. The selected conditions are equivalent to

$$U_A = \bigoplus_{i \in A} ke_i \subseteq W \subseteq \mathcal{H}_G = \bigcap_{j \in G} \ker \pi_j. \quad (4.14)$$

Here $\dim U_A = \alpha$ and $\dim \mathcal{H}_G = \theta - \gamma$, and $U_A \subseteq \mathcal{H}_G$ because $A \cap G = \emptyset$. Passing to the quotient $\mathcal{H}_G/U_A$, the subspaces in (4.14) correspond bijectively to $(d - \alpha)$ -dimensional subspaces of a vector space of dimension $\theta - \alpha - \gamma$. Their number is therefore $\left[\begin{matrix} \theta - \alpha - \gamma \\ d - \alpha \end{matrix} \right]_q$. The number of ordered disjoint choices of (A, G) with these cardinalities is $\binom{\theta}{\alpha} \binom{\theta - \alpha}{\gamma}$. Inclusion–exclusion therefore gives (4.12).

For $\theta = 0$, the zero subspace is the unique possibility. If $\theta \geq 1$, the zero subspace does not have full support, while the whole space contains all coordinate lines, which proves the two vanishing statements. Finally, a line has full support precisely when it is spanned by a vector all of whose coordinates are nonzero. There are $(q - 1)^\theta$ such vectors and each line contains $q - 1$ of them, which proves (4.13). $\square$

Since (4.12) is an alternating double sum, we record two hand computations that let the reader reproduce it. For $\theta = 2$ and $d = 1$ the projective line $\text{PG}(1, q)$ has $q + 1$ points, two of which are the coordinate lines, so $C_{2,1}(q) = (q + 1) - 2 = q - 1$, in agreement with (4.13). For $\theta = 3$ and $d = 2$ a plane is the kernel of a functional $\varphi = (\alpha_1, \alpha_2, \alpha_3)$ determined up to a scalar; one has $\ker \varphi \cap ke_i = 0$ if and only if $\alpha_i \neq 0$, while $\pi_i(\ker \varphi) = 0$ would force $\varphi \in ke_i^*$, which is already excluded. Hence (4.11)

holds exactly when all three coordinates of φ are nonzero, and $C_{3,2}(q) = (q-1)^3/(q-1) = (q-1)^2$. Expanding (4.12) confirms this: Its six nonvanishing terms, indexed by (α, γ) , are $\begin{bmatrix} 3 \\ 2 \end{bmatrix}_q$, $-3\begin{bmatrix} 2 \\ 2 \end{bmatrix}_q$, $-3\begin{bmatrix} 2 \\ 1 \end{bmatrix}_q$, $6\begin{bmatrix} 1 \\ 1 \end{bmatrix}_q$, $3\begin{bmatrix} 1 \\ 0 \end{bmatrix}_q$, and $-3\begin{bmatrix} 0 \\ 0 \end{bmatrix}_q$, with sum $q^2 - 2q + 1$; the terms with $\alpha = 1, \gamma = 2$ and $\alpha = 3, \gamma = 0$ vanish because $\begin{bmatrix} 0 \\ 1 \end{bmatrix}_q = \begin{bmatrix} 0 \\ -1 \end{bmatrix}_q = 0$, and overlooking them is the likeliest source of an arithmetic slip. The same duality argument gives $C_{\theta, \theta-1}(q) = (q-1)^{\theta-1}$ for every $\theta \geq 2$, so the values are symmetric at both ends but not in between, $C_{4,2}(q) = (q-1)^2(q^2 + 3q - 1)$ being no power of $q-1$. Table 2 lists all values for $\theta \leq 4$, together with the totals $F_\theta(q)$ of (5.11) and the unrestricted counts $\mathcal{A}_\theta(q)$; every entry was confirmed by direct enumeration of the subspaces of k^θ for $q = 2, 3, 5$.

Table 2. Support-refined admissible subspace counts; entries outside the displayed range vanish.

θ	$C_{\theta,1}$	$C_{\theta,2}$	$C_{\theta,3}$	$F_\theta(q)$	$\mathcal{A}_\theta(q)$
0	—	—	—	1	1
1	0	—	—	0	1
2	$q-1$	0	—	$q-1$	q
3	$(q-1)^2$	$(q-1)^2$	0	$2(q-1)^2$	$q(2q-1)$
4	$(q-1)^3$	$(q-1)^2(q^2 + 3q - 1)$	$(q-1)^3$	$(q-1)^2(q^2 + 5q - 3)$	$q(q^3 + 3q^2 - 4q + 1)$

We now combine the branch–subspace classification with Proposition 4.4. Put $N_{\mathbf{n}} = \dim_k R_{\mathbf{n}} = 1 + \sum_{i=1}^r (n_i - 1)$. For $\mathbf{a} \in C_{\mathbf{n}}$, define the four cutoff statistics $\Delta(\mathbf{a}) = \sum_{i=1}^r (n_i - a_i)$, $s_{\mathbf{a}} = |S(\mathbf{a})|$, $b_{\mathbf{a}} = |B(\mathbf{a}) \cap S(\mathbf{a})|$ and $c_{\mathbf{a}} = |B(\mathbf{a}) \setminus S(\mathbf{a})|$. The last statistic counts exactly the indices for which $a_i = 1$.

Definition 4.5. Define $\mathcal{G}_{\mathbf{n}}(u, z) = \sum_{I \in \text{Idl}(R_{\mathbf{n}})} u^{\mu(I)} z^{\dim_k I}$, where $\mu(0) = 0$ and $\mu(R_{\mathbf{n}}) = 1$. Here u and z are formal indeterminates marking, respectively, the minimal number of generators and the k -dimension of an ideal.

Combining the canonical ideal classification with the support-refined count produces the main enumerator of this section.

Theorem 4.6. For $\mathbf{a} \in C_{\mathbf{n}}$, set

$$\mathcal{G}_{\mathbf{a}}(u, z) = z^{\Delta(\mathbf{a})} \sum_{\theta=0}^{s_{\mathbf{a}}} \left[\left(\sum_{d=0}^{\theta} C_{\theta,d}(q) (uz)^d \right) \left(\sum_{h=0}^{\theta} \binom{b_{\mathbf{a}}}{h} \binom{s_{\mathbf{a}} - b_{\mathbf{a}}}{\theta - h} u^{b_{\mathbf{a}} + c_{\mathbf{a}} - h} \right) \right], \quad (4.15)$$

with the convention that ordinary binomial coefficients vanish outside their natural range. Both inner sums are indexed by the same θ : The first records the geometry of a coupling of exact support size θ , and the second how that support sits among the $s_{\mathbf{a}}$ terminal coordinates. Then

$$\mathcal{G}_{\mathbf{n}}(u, z) = \sum_{\mathbf{a} \in C_{\mathbf{n}}} \mathcal{G}_{\mathbf{a}}(u, z) + uz^{N_{\mathbf{n}}}. \quad (4.16)$$

Consequently, the coefficient of $u^g z^e$ is the number of ideals of k -dimension e requiring exactly g generators. The specialization $\mathcal{G}_{\mathbf{n}}(1, z)$ is the rank-generating polynomial of the ideal lattice, $\mathcal{G}_{\mathbf{n}}(u, 1)$ records the complete generator spectrum, and $\mathcal{G}_{\mathbf{n}}(1, 1) = |\text{Idl}(R_{\mathbf{n}})|$.

Proof. Fix a cutoff vector $\mathbf{a}$. By Theorem 3.2, the ideals with this cutoff are in bijection with the $\mathbf{a}$ -admissible subspaces $W \leq E_{\mathbf{a}}$. Suppose that $T(W)$ has cardinality θ and meets $B(\mathbf{a}) \cap S(\mathbf{a})$ in exactly h indices. The support is chosen by selecting h coordinates among the $b_{\mathbf{a}}$ active direct heads and $\theta - h$ coordinates among the remaining $s_{\mathbf{a}} - b_{\mathbf{a}}$ terminal coordinates. Hence the number of possible support sets is $\binom{b_{\mathbf{a}}}{h} \binom{s_{\mathbf{a}} - b_{\mathbf{a}}}{\theta - h}$.

After a support set T is fixed, an admissible subspace with exact support T and dimension d is, after relabeling coordinates, exactly a subspace counted by $C_{\theta,d}(q)$. Its ideal has dimension $\dim_k I(\mathbf{a}, W) = \Delta(\mathbf{a}) + d$ by (3.4). Moreover, $|B(\mathbf{a}) \setminus T(W)| = (b_{\mathbf{a}} - h) + c_{\mathbf{a}}$, because $c_{\mathbf{a}}$ direct heads lie outside the terminal space and $b_{\mathbf{a}} - h$ active direct heads are not met by the chosen support. The generator formula (4.4) therefore gives $\mu(I(\mathbf{a}, W)) = d + b_{\mathbf{a}} + c_{\mathbf{a}} - h$. Thus, the contribution of all ideals having the fixed statistics (θ, d, h) is $C_{\theta,d}(q) \binom{b_{\mathbf{a}}}{h} \binom{s_{\mathbf{a}} - b_{\mathbf{a}}}{\theta - h} u^{d+b_{\mathbf{a}}+c_{\mathbf{a}}-h} z^{\Delta(\mathbf{a})+d}$. Summing over h, d , and θ yields (4.15). Summing over all cutoff vectors counts every proper ideal exactly once. The only omitted ideal is the whole ring, whose dimension is $N_{\mathbf{n}}$ and whose minimal number of generators is 1; adding $uz^{N_{\mathbf{n}}}$ gives (4.16). The coefficient and specialization statements follow directly from the definition. $\square$

The transform (4.16) refines both the total ideal count and the rank polynomial by distinguishing ideals of the same dimension with different generator complexity. It also separates the cutoff combinatorics from the geometry of exact-support admissible subspaces through the coefficients $C_{\theta,d}(q)$.

To describe the principal stratum, let $\emptyset \neq T \subseteq [r]$ and define the projective torus $\mathbb{T}_T(k) = (k^{\times})^T / k^{\times}$, where common scalar multiplication is factored out, so that $|\mathbb{T}_T(k)| = (q - 1)^{|T|-1}$. The coefficient of u then admits the following geometric description.

Proposition 4.7. *For every nonempty $T \subseteq [r]$, choose a valuation vector $\mathbf{v} = (v_i)_{i \in T}$, $1 \leq v_i \leq n_i - 1$, and a point $[\mathbf{c}] = [(c_i)_{i \in T}] \in \mathbb{T}_T(k)$. Put*

$$P(T, \mathbf{v}, [\mathbf{c}]) = \left\langle \sum_{i \in T} c_i x_i^{v_i} \right\rangle. \quad (4.17)$$

Then the assignment $(T, \mathbf{v}, [\mathbf{c}]) \mapsto P(T, \mathbf{v}, [\mathbf{c}])$ is a bijection onto the nonzero proper principal ideals of $R_{\mathbf{n}}$. Moreover,

$$\dim_k P(T, \mathbf{v}, [\mathbf{c}]) = \sum_{i \in T} (n_i - v_i) - (|T| - 1). \quad (4.18)$$

Proof. Let $g \in J$ be nonzero and write it uniquely as $g = \sum_{i=1}^r f_i(x_i)$, where $f_i(x_i) \in A_i$. Let T be the set of indices for which $f_i \neq 0$. For each $i \in T$, write $f_i(x_i) = c_i x_i^{v_i} + x_i^{v_i+1} h_i(x_i)$, $c_i \in k^{\times}$, $1 \leq v_i \leq n_i - 1$. If $T = \{i\}$, then $f_i = x_i^{v_i} \zeta_i(x_i)$ for a unit ζ_i of the truncated branch ring, and therefore $\langle g \rangle = \langle x_i^{v_i} \rangle$. This is $P(T, \mathbf{v}, [\mathbf{c}])$, and the projective torus has one point.

Assume $|T| \geq 2$. Define a cutoff vector by

$$a_i = \begin{cases} v_i + 1, & i \in T, \\ n_i, & i \notin T, \end{cases} \quad (4.19)$$

and let $W = k(\sum_{i \in T} c_i e_i^{\mathbf{a}}) \leq E_{\mathbf{a}}$. Every coefficient c_i is nonzero, so $T(W) = T$. Since $|T| \geq 2$, the line W is not a coordinate line and is therefore admissible. Multiplication of g by x_i yields $x_i g = x_i f_i(x_i) = c_i x_i^{v_i+1} + x_i^{v_i+2} h_i(x_i) = x_i^{v_i+1} \zeta_i(x_i)$, where $\zeta_i(x_i)$ is a unit. Hence, $\langle x_i^{v_i+1} \rangle \subseteq \langle g \rangle$ for every

$i \in T$, so $D_{\mathbf{a}} \subseteq \langle g \rangle$. Modulo $D_{\mathbf{a}}$, the class of g is $\sum_{i \in T} c_i x_i^{v_i}$, which spans the lift of W . Therefore, $\langle g \rangle = I(\mathbf{a}, W) = P(T, \mathbf{v}, [\mathbf{c}])$. This also shows that higher terms in the branch components of a generator do not affect its principal ideal.

Conversely, let $I = I(\mathbf{a}, W)$ be a nonzero proper principal ideal. By Proposition 4.1, $\mu(I) = 1$. If $W = 0$, then (4.4) gives $|B(\mathbf{a})| = 1$, so $I = \langle x_i^{a_i} \rangle$ for a unique branch i ; this is obtained from (4.17) with $T = \{i\}$ and $v_i = a_i$. If $W \neq 0$, then (4.4) forces $\dim_k W = 1$ and $B(\mathbf{a}) \subseteq T(W)$. Since coordinates outside $T(W)$ cannot support a nonzero direct branch, one has $a_i = n_i$ for $i \notin T(W)$. Writing the unique terminal line as $W = k \left(\sum_{i \in T(W)} c_i e_i^{\mathbf{a}} \right)$, $c_i \neq 0$, and setting $v_i = a_i - 1$ gives exactly $I = P(T(W), \mathbf{v}, [\mathbf{c}])$. Thus, the map is surjective.

For injectivity, equality of two principal ideals implies equality of their canonical branch–subspace data by Theorem 3.2. In the singleton case, this recovers the unique branch and valuation. In the multi-branch case, it recovers the support T , each cutoff $v_i + 1$, and the terminal line, hence the projective class $[\mathbf{c}]$. Therefore, the parameterization is injective. Finally, if $|T| = 1$, the ideal dimension is $n_i - v_i$. If $|T| \geq 2$, then (3.4) and (4.19) give $\dim_k P = \sum_{i \in T} (n_i - v_i - 1) + 1 = \sum_{i \in T} (n_i - v_i) - (|T| - 1)$, which proves (4.18) in all cases. $\square$

Summing the toric parameters by ideal dimension gives the principal-ideal rank polynomial.

Corollary 4.8. Let $\mathcal{P}_{\mathbf{n}}^{\circ}(z) = \sum_{\substack{0 \neq I \subseteq R_{\mathbf{n}} \\ I \text{ principal}}} z^{\dim_k I}$. Then

$$\mathcal{P}_{\mathbf{n}}^{\circ}(z) = \sum_{\emptyset \neq T \subseteq [r]} (q-1)^{|T|-1} z^{1-|T|} \prod_{i \in T} (z + z^2 + \cdots + z^{n_i-1}). \quad (4.20)$$

Including 0 and $R_{\mathbf{n}}$, the principal-ideal rank polynomial is

$$1 + \mathcal{P}_{\mathbf{n}}^{\circ}(z) + z^{N_{\mathbf{n}}}. \quad (4.21)$$

In particular, the total number of principal ideals is

$$2 + \sum_{\emptyset \neq T \subseteq [r]} (q-1)^{|T|-1} \prod_{i \in T} (n_i - 1) = 2 + \frac{\prod_{i=1}^r (1 + (q-1)(n_i - 1)) - 1}{q-1}. \quad (4.22)$$

Proof. Fix T . By Proposition 4.7, every valuation vector contributes $(q-1)^{|T|-1}$ principal ideals. Put $h_i = n_i - v_i$. As v_i runs from 1 to $n_i - 1$, the integer h_i runs from 1 to $n_i - 1$, and (4.18) becomes $\dim_k P = \sum_{i \in T} h_i - (|T| - 1)$. Therefore, the contribution of the fixed support T to the rank polynomial is $(q-1)^{|T|-1} z^{1-|T|} \prod_{i \in T} (z + z^2 + \cdots + z^{n_i-1})$, which proves (4.20). The zero ideal and the whole ring contribute the two additional terms in (4.21). Setting $z = 1$ gives the subset sum in (4.22). Expanding $\prod_{i=1}^r (1 + (q-1)(n_i - 1))$ and dividing its nonconstant part by $q-1$ gives the product form. $\square$

The following example shows that the compression can have dimension greater than one and cannot be reduced to pairwise couplings.

Example 4.9. Let $k = \mathbb{F}_3$, let $\mathbf{n} = (4, 4, 3, 2)$, and take $\mathbf{a} = (3, 3, 2, 2)$. Let $W = \text{span}_{\mathbb{F}_3} \{e_1 + e_3 + e_4, e_2 + e_3 - e_4\} \leq E_{\mathbf{a}}$. The subspace W has dimension 2 and full support $T(W) = \{1, 2, 3, 4\}$, and contains no coordinate line; hence it is admissible. Here $B(\mathbf{a}) = \{1, 2, 3\}$ and $B(\mathbf{a}) \setminus T(W) = \emptyset$, so (4.4) and (4.8)

give $\mu(I(\mathbf{a}, W)) = 2$ and $\kappa(I(\mathbf{a}, W)) = 4 - 2 = 2$. Although four terminal branch coordinates occur, only two generators are required. Explicitly, $I(\mathbf{a}, W) = \langle x_1^2 + x_3 + x_4, x_2^2 + x_3 - x_4 \rangle$, because multiplication by the branch variables produces x_1^3 , x_2^3 , and x_3^2 , while the fourth branch is already terminal. This example exhibits generator compression that is impossible in a product-of-chains description. It also has no two-branch analogue: A four-coordinate, two-dimensional coupling saves two generators, whereas by Corollary 3.10 a nonzero two-branch coupling is necessarily a line and saves exactly one.

Example 4.10. For $\mathbf{n} = (3, 3)$ and $q = 3$, evaluating (4.15) at each of the nine cutoff vectors and summing gives

$$\mathcal{G}_{(3,3)}(u, z) = 1 + 4uz + 6uz^2 + 2uz^3 + uz^5 + u^2z^2 + 2u^2z^3 + u^2z^4,$$

while for $\mathbf{n} = (3, 2, 2)$ and $q = 2$,

$$\mathcal{G}_{(3,2,2)}(u, z) = 1 + 7uz + 4uz^2 + uz^5 + 7u^2z^2 + 6u^2z^3 + u^3z^3 + u^3z^4.$$

The specializations behave as predicted. Here $\mathcal{G}(1, 1)$ equals 18 and 28, matching (3.11) and Table 1; in the first case $\mathcal{G}(u, 1) = 1 + 13u + 4u^2$ gives 13 nonzero principal ideals, in agreement with the count $2 + ((1 + 2 \cdot 2)^2 - 1)/2 = 14$ of (4.22) once the zero ideal is included; and the coefficient of u^3 occurs only in the second, as it must, since by (4.4) three generators require $r \geq 3$. The refinement separates ideals that the rank polynomial cannot: Of the four ideals of dimension 3 in the first case, two are principal, namely $\langle x_1 + \lambda x_2 \rangle$ with $\lambda \in \mathbb{F}_3^\times$, and two are not, namely $\langle x_1, x_2^2 \rangle$ and $\langle x_1^2, x_2 \rangle$. Both polynomials, and those of the profiles (3, 2), (3, 3), (4, 3), (2, 2, 2), and (3, 2, 2) over $\mathbb{F}_2$, agree coefficientwise with a direct computation in which all ideals were listed and $\mu(I) = \dim_k I - \dim_k JI$ evaluated for each.

The generator theory therefore retains both the dimension and the exact coordinate position of the terminal coupling. The next section shows that annihilation behaves differently: It discards the internal position of that subspace and retains only its support.

5. Annihilator duality and cyclic modules

The generator theory retains the full terminal subspace, whereas annihilation remembers only its support. This section determines the resulting annihilator closure, its defect, fibers, and fixed-point lattice, and then applies the same structure to cyclic quotients and biduality.

For $I \in \text{Idl}(R_{\mathbf{n}})$, write $I^\perp = \text{Ann}(I) = \{y \in R_{\mathbf{n}} : yI = 0\}$. This notation refers only to multiplication in the ring; no Frobenius or character-module duality is assumed. For a canonical ideal $I(\mathbf{a}, W)$, recall that $T(W)$ denotes the coordinate support of W . We begin with the standard annihilator identities needed for the explicit closure calculation.

Lemma 5.1. For all ideals $I, K \in \text{Idl}(R_{\mathbf{n}})$,

$$I \subseteq K \implies K^\perp \subseteq I^\perp, \quad (5.1)$$

$$I \subseteq I^{\perp\perp}, \quad (5.2)$$

$$I^{\perp\perp\perp} = I^\perp, \quad (5.3)$$

$$(I + K)^\perp = I^\perp \cap K^\perp. \quad (5.4)$$

Consequently, $\text{cl}_{\text{Ann}} : \text{Idl}(R_{\mathbf{n}}) \longrightarrow \text{Idl}(R_{\mathbf{n}})$, $\text{cl}_{\text{Ann}}(I) = I^{\perp\perp}$, is an order-preserving, extensive, and idempotent closure operator.

Proof. The implication in (5.1) follows directly from the definition of an annihilator. If $x \in I$ and $y \in I^\perp$, then $yx = 0$; commutativity therefore gives $x \in I^{\perp\perp}$, which proves (5.2). Applying the order-reversing map $(-)^\perp$ to (5.2) yields $I^{\perp\perp\perp} \subseteq I^\perp$, while (5.2) applied to $I^\perp$ gives the reverse inclusion. This proves (5.3). Finally, an element annihilates $I + K$ if and only if it annihilates both I and K , which is (5.4). The asserted closure properties follow from (5.1)–(5.3). $\square$

The properties above are standard. The substantive point for R_n is that the closure can be computed explicitly from canonical ideal data.

Theorem 5.2. *Let $0 \neq I = I(\mathbf{a}, W) \subsetneq R_n$, $\mathbf{a} = (a_1, \dots, a_r)$, be in canonical form. Define*

$$\varepsilon_i(W) = \begin{cases} 1, & i \in T(W), \\ 0, & i \notin T(W), \end{cases}$$

and put

$$b_i = \max\{1, n_i - a_i + \varepsilon_i(W)\} \quad (1 \leq i \leq r). \quad (5.5)$$

Then

$$I^\perp = D_{\mathbf{b}} = \bigoplus_{i=1}^r \langle x_i^{b_i} \rangle. \quad (5.6)$$

In particular, the annihilator of every nonzero proper ideal is a direct branch ideal, and it depends on W only through its coordinate support $T(W)$.

Proof. Let $y \in I^\perp$ and write $y = \alpha + z$ with $\alpha \in k$ and $z \in J$. If $\alpha \neq 0$, then y is a unit; since $I \neq 0$, a unit cannot annihilate I . Hence $\alpha = 0$ and every element of $I^\perp$ lies in J . Using the branch decomposition, write $y = \sum_{i=1}^r f_i(x_i)$, $f_i(x_i) \in A_i$. For a nonzero f_i , let $v_i(f_i)$ be the least exponent of x_i occurring in f_i , and put $v_i(0) = \infty$.

We first impose the condition $yD_{\mathbf{a}} = 0$. If $a_i < n_i$, then the i th component of $D_{\mathbf{a}}$ is generated by $x_i^{a_i}$, and $f_i(x_i)x_i^{a_i} = 0$ if and only if $v_i(f_i) + a_i \geq n_i$, that is, $f_i \in \langle x_i^{n_i - a_i} \rangle$. If $a_i = n_i$, the i th component of $D_{\mathbf{a}}$ is zero and imposes no condition beyond $f_i \in A_i = \langle x_i \rangle$. Thus, annihilation of the direct part is equivalent to

$$f_i \in \langle x_i^{\max\{1, n_i - a_i\}} \rangle \quad (1 \leq i \leq r). \quad (5.7)$$

We next impose $y\widetilde{W} = 0$. Let $w = \sum_{j \in S(\mathbf{a})} c_j e_j^{\mathbf{a}} \in W$ and $\widetilde{w} = \sum_{j \in S(\mathbf{a})} c_j x_j^{a_j - 1}$ be its canonical lift. Since products between distinct branches vanish,

$$y\widetilde{w} = \sum_{j \in S(\mathbf{a})} c_j f_j(x_j) x_j^{a_j - 1}. \quad (5.8)$$

The summands in (5.8) lie in the direct sum $A_1 \oplus \cdots \oplus A_r$, so the product is zero if and only if every branch component is zero.

Suppose $i \in T(W)$. By definition of support, there exists $w \in W$ whose i th coordinate c_i is nonzero. Applying (5.8) to this w shows that $f_i(x_i)x_i^{a_i - 1} = 0$, which is equivalent to $v_i(f_i) + a_i - 1 \geq n_i$, that is,

$$f_i \in \langle x_i^{n_i - a_i + 1} \rangle. \quad (5.9)$$

Conversely, if (5.9) holds for every $i \in T(W)$, then each branch component in (5.8) vanishes for every $w \in W$. Indices outside $T(W)$ contribute nothing, because all corresponding coordinates of W are zero. Therefore, (5.9) is exactly the additional condition required to annihilate $\widetilde{W}$.

Combining (5.7) and (5.9), the allowable i th branch components are precisely $f_i(x_i) \in \langle x_i^{\max\{1, n_i - a_i + \varepsilon_i(W)\}} \rangle = \langle x_i^{b_i} \rangle$. The branch components are independent, so their direct sum is exactly $D_{\mathbf{b}}$. This proves (5.6). $\square$

The collapse formula determines the image of the annihilator map.

Corollary 5.3. *The annihilators of the nonzero proper ideals are precisely the direct ideals*

$$\{D_{\mathbf{b}} : 1 \leq b_i \leq n_i - 1 \text{ for every } i\}. \quad (5.10)$$

Equivalently, they are exactly the direct branch ideals lying between the socle and the maximal ideal: $\text{Soc}(R_{\mathbf{n}}) \subseteq D_{\mathbf{b}} \subseteq J$.

Proof. For a nonzero proper ideal, the cutoffs in (5.5) satisfy $1 \leq b_i \leq n_i - 1$, since $\varepsilon_i(W) = 1$ forces $a_i \geq 2$. Hence every annihilator belongs to the displayed family. Conversely, fix $\mathbf{b}$ with $1 \leq b_i \leq n_i - 1$ and set $a_i = n_i - b_i$. Then $1 \leq a_i \leq n_i - 1$, and Theorem 5.2 with $W = 0$ gives $\text{Ann}(D_{\mathbf{a}}) = D_{\mathbf{b}}$. The interval description follows from $\text{Soc}(R_{\mathbf{n}}) = D_{(n_1-1, \dots, n_r-1)}$ and $J = D_{(1, \dots, 1)}$. $\square$

Corollary 5.3 identifies the image of annihilation but not its fibers. The following proposition determines exactly how different cutoff vectors and terminal geometries collapse to the same direct ideal.

For $\mathbf{b}$ satisfying (5.10), put $Q(\mathbf{b}) = \{i : b_i = 1\}$. For $\theta \geq 0$, define

$$F_{\theta}(q) = \begin{cases} 1, & \theta = 0, \\ 0, & \theta = 1, \\ \sum_{d=1}^{\theta-1} C_{\theta,d}(q), & \theta \geq 2, \end{cases} \quad (5.11)$$

where $C_{\theta,d}(q)$ is the support-refined Grassmannian coefficient from Proposition 4.4.

Proposition 5.4. *Fix $\mathbf{b}$ with $1 \leq b_i \leq n_i - 1$. A nonzero proper canonical ideal $I(\mathbf{a}, W)$ satisfies $\text{Ann}(I(\mathbf{a}, W)) = D_{\mathbf{b}}$ if and only if, with $T = T(W)$, its cutoff coordinates satisfy*

$$a_i = \begin{cases} n_i - b_i + 1, & i \in T, \\ n_i - b_i, & i \notin T \text{ and } b_i \geq 2, \\ n_i - 1 \text{ or } n_i, & i \notin T \text{ and } b_i = 1, \end{cases} \quad (5.12)$$

and W is an $\mathbf{a}$ -admissible subspace with exact support T . Consequently, the size of the fiber over $D_{\mathbf{b}}$ is

$$|\{0 \neq I \subsetneq R_{\mathbf{n}} : \text{Ann}(I) = D_{\mathbf{b}}\}| = \sum_{T \subseteq [r]} 2^{|\mathcal{Q}(\mathbf{b}) \setminus T|} F_{|T|}(q) - \mathbf{1}_{\mathbf{b}=(1, \dots, 1)}. \quad (5.13)$$

Proof. By Theorem 5.2, the equality $\text{Ann}(I(\mathbf{a}, W)) = D_{\mathbf{b}}$ is equivalent, coordinate by coordinate, to

$$b_i = \max\{1, n_i - a_i + \varepsilon_i(W)\}. \quad (5.14)$$

Let $T = T(W)$. If $i \in T$, then $\varepsilon_i(W) = 1$ and $a_i \leq n_i$, so $n_i - a_i + 1 \geq 1$; hence the maximum in (5.14) is unnecessary and $a_i = n_i - b_i + 1$. If $i \notin T$ and $b_i \geq 2$, then $\varepsilon_i(W) = 0$ and the value of the maximum is strictly greater than 1, so $b_i = n_i - a_i$ and $a_i = n_i - b_i$. Finally, suppose $i \notin T$ and $b_i = 1$. Then $\max\{1, n_i - a_i\} = 1$, which is equivalent to $n_i - a_i \leq 1$. Since $a_i \leq n_i$, this gives the two possibilities $a_i = n_i - 1$ and $a_i = n_i$. This proves (5.12); reversing the argument proves its sufficiency.

Fix a support set T . For each index in $Q(\mathbf{b}) \setminus T$, the last line of (5.12) gives two independent cutoff choices, and all other cutoffs are forced. Hence, there are $2^{|Q(\mathbf{b}) \setminus T|}$ cutoff vectors associated with T . For each such vector, the number of admissible subspaces with exact support T depends only on $|T|$ and equals $F_{|T|}(q)$ by (5.11). Summing over T gives the first term in (5.13).

There is one exceptional datum. If $\mathbf{b} = (1, \dots, 1)$, $T = \emptyset$, and $a_i = n_i$ for every i , then $W = 0$ and $I(\mathbf{a}, 0) = 0$. Its annihilator is the whole ring rather than $D_{(1, \dots, 1)} = J$, because Theorem 5.2 was stated for nonzero ideals. This unique zero datum must be removed, which produces the final indicator term. No other datum in the sum represents the zero ideal. $\square$

We now determine the closure $I \mapsto I^{\perp\perp}$ explicitly. This reveals how much of an ideal is lost under annihilation and then restored by a second annihilation.

Theorem 5.5. Let $0 \neq I = I(\mathbf{a}, W) \subsetneq R_{\mathbf{n}}$, and define $\mathbf{c} = (c_1, \dots, c_r)$ by

$$c_i = \begin{cases} a_i - 1, & i \in T(W), \\ a_i, & i \notin T(W) \text{ and } a_i < n_i, \\ n_i - 1, & i \notin T(W) \text{ and } a_i = n_i. \end{cases} \quad (5.15)$$

Then

$$I^{\perp\perp} = D_{\mathbf{c}}. \quad (5.16)$$

In particular, double annihilation removes all terminal-subspace data. Define the annihilator defect by $\delta_{\text{Ann}}(I) = \dim_k I^{\perp\perp} - \dim_k I$ and the absent-branch set by $Z(I) = \{i : a_i = n_i, i \notin T(W)\}$. Then

$$\delta_{\text{Ann}}(I) = |T(W)| - \dim_k W + |Z(I)| = \kappa(I) + |Z(I)|, \quad (5.17)$$

where $\kappa(I)$ is the coupling-compression invariant from (4.8).

Proof. By Theorem 5.2, $I^{\perp} = D_{\mathbf{b}}$, where $b_i = \max\{1, n_i - a_i + \varepsilon_i(W)\}$. The subspace parameter of $D_{\mathbf{b}}$ is zero. Applying the same theorem a second time gives $c_i = \max\{1, n_i - b_i\}$. We evaluate this expression in the three possible cases.

If $i \in T(W)$, then $b_i = n_i - a_i + 1$, including the boundary case $a_i = n_i$; hence $c_i = n_i - b_i = a_i - 1$. If $i \notin T(W)$ and $a_i < n_i$, then $b_i = n_i - a_i$, and therefore $c_i = a_i$. If $i \notin T(W)$ and $a_i = n_i$, then $b_i = 1$ and $c_i = n_i - 1$. This proves (5.15) and (5.16).

Using (3.4), $\dim_k I = \sum_{i=1}^r (n_i - a_i) + \dim_k W$ and $\dim_k I^{\perp\perp} = \sum_{i=1}^r (n_i - c_i)$. For each $i \in T(W)$, replacing a_i by $c_i = a_i - 1$ increases the direct branch dimension by 1. For each $i \in Z(I)$, replacing the absent cutoff $a_i = n_i$ by $c_i = n_i - 1$ also contributes one new dimension. All remaining coordinates are unchanged. Since the terminal subspace W disappears, we obtain $\dim_k I^{\perp\perp} - \dim_k I = |T(W)| + |Z(I)| - \dim_k W$, which is the first equality in (5.17). The second follows from $\kappa(I) = |T(W)| - \dim_k W$. $\square$

Formula (5.17), together with Proposition 4.3, links two a priori different phenomena. The term $\kappa(I)$ measures the number of generators saved by terminal coupling, whereas $\delta_{\text{Ann}}(I)$ measures the dimension restored by double annihilation. Thus, coupling compression reappears exactly as the geometric part of the double-annihilator defect, while the additional term $|Z(I)|$ records branches absent from the ideal altogether.

Proposition 5.6. *Let $\mathfrak{D}_{\mathbf{n}} = \{I \in \text{Idl}(R_{\mathbf{n}}) : I^{\perp\perp} = I\}$. Then*

$$\mathfrak{D}_{\mathbf{n}} = \{0, R_{\mathbf{n}}\} \cup \{D_{\mathbf{a}} : 1 \leq a_i \leq n_i - 1 \text{ for every } i\}. \quad (5.18)$$

Moreover:

- (i) $\mathfrak{D}_{\mathbf{n}}$ is a bounded distributive sublattice of $\text{Idl}(R_{\mathbf{n}})$;
- (ii) annihilation restricts to an order-reversing involution of $\mathfrak{D}_{\mathbf{n}}$;
- (iii) for $1 \leq a_i \leq n_i - 1$,

$$D_{\mathbf{a}}^{\perp} = D_{\mathbf{n}-\mathbf{a}}, \quad \mathbf{n} - \mathbf{a} = (n_1 - a_1, \dots, n_r - a_r); \quad (5.19)$$

- (iv) for all $I, K \in \mathfrak{D}_{\mathbf{n}}$,

$$(I + K)^{\perp} = I^{\perp} \cap K^{\perp}, \quad (I \cap K)^{\perp} = I^{\perp} + K^{\perp}. \quad (5.20)$$

Thus, $\mathfrak{D}_{\mathbf{n}}$ is the largest part of the ideal lattice on which multiplication annihilation acts as a perfect lattice duality.

Proof. The ideals 0 and $R_{\mathbf{n}}$ are fixed by double annihilation, because $0^{\perp} = R_{\mathbf{n}}$ and $R_{\mathbf{n}}^{\perp} = 0$. Let $0 \neq I(\mathbf{a}, W) \subseteq R_{\mathbf{n}}$. By Theorem 5.5, its double annihilator is a direct ideal; therefore, the equality $I^{\perp\perp} = I$ forces $W = 0$. With $W = 0$, formula (5.15) becomes $c_i = a_i$, if $a_i < n_i$, and $n_i - 1$, when $a_i = n_i$. Hence, $D_{\mathbf{a}}$ is fixed if and only if $a_i < n_i$ for every i , which proves (5.18). Equivalently, the same conclusion follows from (5.17): Admissibility implies $|T(W)| - \dim_k W > 0$ whenever $W \neq 0$, so zero defect occurs exactly when $W = 0$ and $Z(I) = \emptyset$.

For two direct ideals in the core, $D_{\mathbf{a}} \cap D_{\mathbf{b}} = D_{\mathbf{a} \vee \mathbf{b}}$ and $D_{\mathbf{a}} + D_{\mathbf{b}} = D_{\mathbf{a} \wedge \mathbf{b}}$, in the notation of Proposition 3.6. Since all coordinates of $\mathbf{a}$ and $\mathbf{b}$ lie between 1 and $n_i - 1$, the same is true of their coordinatewise maxima and minima. Thus, the proper nonzero part of $\mathfrak{D}_{\mathbf{n}}$ is a product of finite chains, with the coordinate orders reversed. Adjoining the global bottom 0 and top $R_{\mathbf{n}}$ preserves distributivity, which proves (i).

For a core ideal $D_{\mathbf{a}}$, Theorem 5.2 gives $D_{\mathbf{a}}^{\perp} = D_{(n_1 - a_1, \dots, n_r - a_r)}$, because $1 \leq a_i \leq n_i - 1$. The resulting cutoff vector again lies in the same range, and applying the operation twice returns $\mathbf{a}$. Together with the exchange of 0 and $R_{\mathbf{n}}$, this proves (ii) and (iii). The first identity in (5.20) is (5.4). The second follows by applying the order-reversing involution to the first identity, or directly from the coordinate formulas for meet, join, and (5.19). This proves (iv). $\square$

The coordinate description gives the size of the closed sublattice.

Corollary 5.7. *The number of double-annihilator-closed ideals is $|\mathfrak{D}_{\mathbf{n}}| = 2 + \prod_{i=1}^r (n_i - 1)$.*

Proof. By Proposition 5.6, the duality core consists of 0 , the whole ring, and the direct ideals $D_{\mathbf{a}}$ with $1 \leq a_i \leq n_i - 1$ for every i . The cutoff coordinates may be chosen independently, which gives $\prod_i (n_i - 1)$ nonzero proper core ideals. $\square$

The fixed points of annihilation itself are more restrictive than the fixed points of double annihilation; the following corollary determines them.

Corollary 5.8. *A self-annihilating ideal $I = I^\perp$ exists if and only if every branch length n_i is even. When it exists, it is unique and equals $\mathfrak{h}_{\mathbf{n}} = D_{\mathbf{n}/2} = \langle x_1^{n_1/2}, \dots, x_r^{n_r/2} \rangle$. Its cardinality satisfies*

$$|\mathfrak{h}_{\mathbf{n}}|^2 = q^{r-1} |R_{\mathbf{n}}|. \quad (5.21)$$

Consequently, the familiar Frobenius relation $|I| |I^\perp| = |R_{\mathbf{n}}|$ holds for this fixed point if and only if $r = 1$.

Proof. Suppose $I = I^\perp$. The cases $I = 0$ and $I = R_{\mathbf{n}}$ are impossible. For a nonzero proper ideal, Theorem 5.2 shows that $I^\perp$ is direct; hence $I = D_{\mathbf{a}}$ for some $\mathbf{a}$. No coordinate a_i can equal n_i , because then the i th coordinate of the annihilator cutoff would be 1, whereas the i th component of $D_{\mathbf{a}}$ would be zero. Thus, $1 \leq a_i \leq n_i - 1$, and (5.19) gives $D_{\mathbf{a}} = D_{\mathbf{n}-\mathbf{a}}$. Uniqueness of direct cutoff data implies $2a_i = n_i$ for every i . Therefore, every n_i must be even and $a_i = n_i/2$. Conversely, if every n_i is even, then (5.19) directly gives $D_{\mathbf{n}/2}^\perp = D_{\mathbf{n}/2}$.

Finally, $\dim_k \mathfrak{h}_{\mathbf{n}} = \sum_{i=1}^r \left(n_i - \frac{n_i}{2}\right) = \frac{1}{2} \sum_{i=1}^r n_i$ and $\dim_k R_{\mathbf{n}} = \sum_{i=1}^r n_i - r + 1$. Hence, $2 \dim_k \mathfrak{h}_{\mathbf{n}} - \dim_k R_{\mathbf{n}} = r - 1$, and exponentiating with base q gives (5.21). $\square$

Concretely, for $\mathbf{n} = (4, 4)$ one has $\dim_k R_{\mathbf{n}} = 7$ and $\mathfrak{h}_{(4,4)} = D_{(2,2)} = \langle x_1^2, x_2^2 \rangle = \text{span}_k \{x_1^2, x_1^3, x_2^2, x_2^3\}$ of dimension 4: An element $y = f_1(x_1) + f_2(x_2)$ annihilates it precisely when $f_i(x_i)x_i^2 = 0$, that is, when $f_i \in \langle x_i^2 \rangle$, so $\mathfrak{h}^\perp = \mathfrak{h}$ directly. Here $|\mathfrak{h}|^2 = q^8 = q \cdot q^7$, whereas the Frobenius relation would require q^7 , the extra factor being the excess socle dimension $r - 1 = 1$. For $\mathbf{n} = (2, 4)$ one gets $\mathfrak{h}_{(2,4)} = \langle x_1, x_2^2 \rangle$ with $|\mathfrak{h}|^2 = q^6 = q \cdot q^5$; this ideal is not a power of J , since $J = \langle x_1, x_2 \rangle$ and $J^2 = \langle x_2^2 \rangle$ give $J^2 \subsetneq \mathfrak{h}_{(2,4)} \subsetneq J$, so that, unlike in the chain case, the self-annihilating ideal is a radical power only when all the n_i are equal. For $\mathbf{n} = (2, 2, 4, 4)$ one has $\mathfrak{h} = \langle x_1, x_2, x_3^2, x_4^2 \rangle$ and $|\mathfrak{h}|^2 = q^{12} = q^3 \cdot q^9$, again as predicted. The parity hypothesis is not decorative: For $\mathbf{n} = (3, 4)$ the equation $2a_1 = 3$ has no solution, and annihilation interchanges the closest candidates $D_{(1,2)}$ and $D_{(2,2)}$ instead of fixing either.

The condition $I \subseteq I^\perp$ is equivalent to $I^2 = 0$ and is the multiplicative analogue of self-orthogonality. The canonical containment criterion turns this condition into explicit branch inequalities.

Corollary 5.9. *Let $I = I(\mathbf{a}, W)$ be a proper ideal. Then the following are equivalent:*

- (i) $I^2 = 0$;
- (ii) $I \subseteq I^\perp$;
- (iii) the cutoff coordinates satisfy

$$2a_i \geq n_i \quad (i \notin T(W)), \quad 2a_i \geq n_i + 2 \quad (i \in T(W)). \quad (5.22)$$

Proof. Commutativity gives the equivalence of (i) and (ii): The condition $I^2 = 0$ means that every element of I annihilates every element of I , which is exactly $I \subseteq \text{Ann}(I)$.

Assume first that $I \neq 0$; the zero ideal satisfies all three conditions. By Theorem 5.2, $I^\perp = D_{\mathbf{b}} = I(\mathbf{b}, 0)$ with $b_i = \max\{1, n_i - a_i + \varepsilon_i(W)\}$. The containment criterion of Proposition 3.3 gives

$$I(\mathbf{a}, W) \subseteq I(\mathbf{b}, 0) \iff \mathbf{a} \geq \mathbf{b} \text{ and } \tau_{\mathbf{a}, \mathbf{b}}(W) = 0. \quad (5.23)$$

Let $i \notin T(W)$. The transition condition imposes no restriction in the i th coordinate, and $a_i \geq b_i$ is equivalent to $a_i \geq \max\{1, n_i - a_i\}$. Since $a_i \geq 1$, this reduces to $2a_i \geq n_i$.

Now let $i \in T(W)$, so that $b_i = n_i - a_i + 1$. Because the projection of W onto the i th coordinate is nonzero, the transition map can annihilate that coordinate only if $a_i > b_i$; equality would preserve the i th terminal coordinate and make $\tau_{\mathbf{a},\mathbf{b}}(W)$ nonzero. Therefore, $a_i > b_i = n_i - a_i + 1$, which, for integers, is equivalent to $2a_i \geq n_i + 2$. Conversely, the inequalities in (5.22) imply $a_i \geq b_i$ for every i and the strict inequality $a_i > b_i$ on every coordinate in $T(W)$. The transition map therefore annihilates every coordinate in the support of W , so $\tau_{\mathbf{a},\mathbf{b}}(W) = 0$. By (5.23), $I \subseteq I^\perp$. This proves the equivalence. $\square$

The three-branch ideal used earlier now displays the loss of internal terminal geometry under annihilation.

Example 5.10. Let $k = \mathbb{F}_3$ and $\mathbf{n} = (4, 3, 3)$, and consider the ideal from Example 3.9: $\mathbf{a} = (3, 2, 2)$ and $W = \text{span}_{\mathbb{F}_3}\{e_1 + e_3, e_2 + e_3\}$. Here $T(W) = \{1, 2, 3\}$ and $\dim_k W = 2$. Theorem 5.2 gives

$$I(\mathbf{a}, W)^\perp = D_{(2,2,2)} = \langle x_1^2, x_2^2, x_3^2 \rangle,$$

while Theorem 5.5 gives $I(\mathbf{a}, W)^{\perp\perp} = D_{(2,1,1)} = \langle x_1^2, x_2, x_3 \rangle$ and $\delta_{\text{Ann}}(I) = |T(W)| - \dim_k W = 3 - 2 = 1$. Thus, a two-dimensional coupling plane and all of its projective position data collapse to a direct ideal after one annihilation, and the single lost dimension is exactly the coupling-compression invariant. Unlike the two-branch case, where a nonzero coupling is only a projective point, this example shows that annihilation also erases the internal geometry of a higher-dimensional plane.

Remark 5.11. For $r \geq 2$, the ring $R_{\mathbf{n}}$ is non-Frobenius by Corollary 2.3, and annihilation is not an involution on the full ideal lattice. The failure is measured exactly by (5.17). Nevertheless, annihilation becomes an involutive lattice anti-automorphism on the canonical distributive sublattice $\mathfrak{D}_{\mathbf{n}}$. The duality core therefore isolates the largest branchwise part of the ideal lattice on which a Frobenius-like annihilator symmetry survives despite the global socle dimension r .

The annihilator theory has a direct module-theoretic interpretation for cyclic quotients. For an $R_{\mathbf{n}}$ -module M , write $M^\vee = \text{Hom}_{R_{\mathbf{n}}}(M, R_{\mathbf{n}})$ and let $\eta_M : M \rightarrow M^{\vee\vee}$, $\eta_M(m)(f) = f(m)$, be the canonical evaluation map. We call M *torsionless* if η_M is injective and *reflexive* if it is an isomorphism, and we put $M_I = R_{\mathbf{n}}/I$.

For M_I , double annihilation determines the torsionless closure: The quotient $I^{\perp\perp}/I$ is the kernel of η_{M_I} , of dimension $\delta_{\text{Ann}}(I)$, whereas every nonzero proper cyclic quotient has bidual cokernel dimension $r^2 - 1$. Thus, injectivity reflects the ideal-specific branch-subspace data, whereas surjectivity depends only on the embedding dimension r . We begin by identifying $M_I^\vee$ with an ideal of the ambient ring.

Proposition 5.12. For every ideal $I \in \text{Idl}(R_{\mathbf{n}})$, the map

$$\Phi_I : I^\perp \longrightarrow M_I^\vee, \quad \Phi_I(y)(\varrho + I) = \varrho y, \quad (5.24)$$

is an isomorphism of $R_{\mathbf{n}}$ -modules; hence $M_I^\vee \cong I^\perp$. If $0 \neq I = I(\mathbf{a}, W) \subsetneq R_{\mathbf{n}}$, then $M_I^\vee \cong D_{\mathbf{b}}$, $b_i = \max\{1, n_i - a_i + \varepsilon_i(W)\}$. Thus, the dual of a cyclic quotient depends on the terminal subspace W only through its coordinate support $T(W)$.

Proof. Let $y \in I^\perp$. If $\varrho + I = \varsigma + I$, then $\varrho - \varsigma \in I$, and therefore $(\varrho - \varsigma)y = 0$. Hence, the formula in (5.24) is well defined, and it is plainly $R_{\mathbf{n}}$ -linear.

If $\Phi_I(y) = 0$, then $y = \Phi_I(y)(1 + I) = 0$, so Φ_I is injective. Conversely, let $f \in M_I^\vee$ and put $y = f(1 + I)$. For every $a \in I$, $ay = af(1 + I) = f(a + I) = f(0) = 0$, so $y \in I^\perp$. Moreover, for every $\varrho \in R_n$, $f(\varrho + I) = \varrho f(1 + I) = \varrho y = \Phi_I(y)(\varrho + I)$, so $f = \Phi_I(y)$. This proves surjectivity and hence the first assertion. The canonical description for a nonzero proper ideal follows immediately from Theorem 5.2. $\square$

Proposition 5.12, together with the triple-annihilator identity (5.3), converts the annihilator fibers of Proposition 5.4 into fibers of module dualization and places every cyclic dual in the double-annihilator-closed part of the ideal lattice.

Corollary 5.13. *Let $0 \neq I, K \subsetneq R_n$. Then $M_I^\vee \cong M_K^\vee$ if and only if $I^\perp = K^\perp$. Consequently, the duals of all nonzero proper cyclic quotients fall into exactly $\prod_{i=1}^r (n_i - 1)$ isomorphism classes, indexed by the direct ideals $D_{\mathbf{b}}$ with $1 \leq b_i \leq n_i - 1$.*

Proof. If $I^\perp = K^\perp$, the conclusion follows from Proposition 5.12. Conversely, write $I^\perp = D_{\mathbf{b}}$ and $K^\perp = D_{\mathbf{c}}$. By Corollary 5.3, all coordinates of $\mathbf{b}$ and $\mathbf{c}$ lie between 1 and $n_i - 1$. If $D_{\mathbf{b}} \cong D_{\mathbf{c}}$ as modules, then their module annihilators are equal, and (5.19) gives $\text{Ann}(D_{\mathbf{b}}) = D_{\mathbf{n}-\mathbf{b}}$ and $\text{Ann}(D_{\mathbf{c}}) = D_{\mathbf{n}-\mathbf{c}}$. Uniqueness of direct cutoff data therefore implies $\mathbf{b} = \mathbf{c}$. The count follows from the classification of the image of the annihilator map. $\square$

The next proposition converts double annihilation into the exact kernel and image of bidual evaluation.

Proposition 5.14. *For every ideal $I \in \text{Idl}(R_n)$, the evaluation map fits into the exact sequence*

$$0 \longrightarrow I^{\perp\perp}/I \longrightarrow M_I \xrightarrow{\eta_{M_I}} M_I^{\vee\vee}. \quad (5.25)$$

Moreover,

$$\ker(\eta_{M_I}) = I^{\perp\perp}/I \quad (5.26)$$

and

$$\text{im}(\eta_{M_I}) \cong R_n/I^{\perp\perp}. \quad (5.27)$$

Consequently,

$$M_I \text{ is torsionless} \iff I = I^{\perp\perp}. \quad (5.28)$$

For $0 \neq I \subsetneq R_n$,

$$\dim_k \ker(\eta_{M_I}) = \delta_{\text{Ann}}(I) = \kappa(I) + |Z(I)|. \quad (5.29)$$

Proof. By Proposition 5.12, we may identify $M_I^\vee$ with $I^\perp$. Under this identification, $\eta_{M_I}(\varrho + I) : I^\perp \longrightarrow R_n, y \longmapsto \varrho y$. Therefore,

$$\begin{aligned} \varrho + I \in \ker(\eta_{M_I}) &\iff \varrho y = 0 \text{ for every } y \in I^\perp \\ &\iff \varrho \in \text{Ann}(I^\perp) = I^{\perp\perp}. \end{aligned}$$

Since $I \subseteq I^{\perp\perp}$, the kernel is exactly $I^{\perp\perp}/I$, which proves (5.25) and (5.26). The first isomorphism theorem then gives $\text{im}(\eta_{M_I}) \cong (R_n/I)/(I^{\perp\perp}/I) \cong R_n/I^{\perp\perp}$, which proves (5.27).

The evaluation map is injective precisely when its kernel is zero, equivalently when $I = I^{\perp\perp}$; this is (5.28). Finally, for a nonzero proper ideal, $\dim_k \ker(\eta_{M_I}) = \dim_k I^{\perp\perp} - \dim_k I = \delta_{\text{Ann}}(I)$, and the formula of Theorem 5.5 yields the final equality in (5.29). $\square$

Thus, the annihilator defect is the exact dimension of the obstruction to torsionlessness. Its decomposition $\delta_{\text{Ann}}(I) = \kappa(I) + |Z(I)|$ separates the contribution of terminal coupling from that of completely absent branches.

Proposition 5.15. *For every ideal $I \in \text{Idl}(R_n)$, define $\mathfrak{t}(M_I) = R_n/I^{\perp\perp}$. Then $\mathfrak{t}(M_I)$ is torsionless, and the canonical epimorphism $\pi_I : M_I \longrightarrow \mathfrak{t}(M_I), \varrho + I \longmapsto \varrho + I^{\perp\perp}$, has kernel $I^{\perp\perp}/I$. Moreover, $\mathfrak{t}$ is a reflection in the strict sense: For every torsionless R_n -module N , not assumed cyclic or finitely generated, and every homomorphism $f : M_I \rightarrow N$, there is a unique $g : \mathfrak{t}(M_I) \rightarrow N$ with $f = g \circ \pi_I$, that is,*

$$\pi_I^* : \text{Hom}_{R_n}(\mathfrak{t}(M_I), N) \longrightarrow \text{Hom}_{R_n}(M_I, N) \quad \text{is bijective.} \quad (5.30)$$

Thus, π_I is universal among homomorphisms from M_I into torsionless R_n -modules. We state the conclusion as this universal property rather than as an adjunction, since a torsionless R_n -module need not be cyclic. In particular, if $I \subseteq K$ and R_n/K is torsionless, then $I^{\perp\perp} \subseteq K$ and $M_I \rightarrow R_n/K$ factors uniquely through $\mathfrak{t}(M_I)$. Furthermore, $\text{im}(\eta_{M_I}) \cong \mathfrak{t}(M_I)$.

Proof. The ideal $I^{\perp\perp}$ is fixed by double annihilation, since Lemma 5.1 gives $(I^{\perp\perp})^{\perp\perp} = I^{\perp\perp\perp\perp} = I^{\perp\perp}$. Hence Proposition 5.14 shows that $R_n/I^{\perp\perp}$ is torsionless. The kernel of π_I consists exactly of the classes represented by $I^{\perp\perp}$, and is therefore $I^{\perp\perp}/I$.

A submodule $N' \leq N$ of a torsionless module is again torsionless, since N is torsionless exactly when every $0 \neq n \in N$ admits $h \in N^\vee$ with $h(n) \neq 0$, and such an h restricts to N' . Now let N be torsionless, let $f : M_I \rightarrow N$ be linear, and put $n = f(1 + I)$. For $a \in I$ one has $an = f(a + I) = 0$, so $I \subseteq \text{Ann}(n)$; the cyclic submodule $R_n n \cong R_n / \text{Ann}(n)$ is torsionless, whence $\text{Ann}(n) = \text{Ann}(n)^{\perp\perp}$ by (5.28), and applying (5.1) twice to $I \subseteq \text{Ann}(n)$ gives $I^{\perp\perp} \subseteq \text{Ann}(n)$. Therefore, f kills $\ker \pi_I = I^{\perp\perp}/I$ and factors as $f = g \circ \pi_I$, uniquely because π_I is surjective, which is (5.30); the last assertion is the case $N = R_n/K$ with f the quotient map. The final assertion is (5.27). $\square$

Proposition 5.15 identifies double annihilation as a torsionless reflection, in the strict categorical sense of (5.30); its essential image on cyclic modules is the duality core $\mathfrak{D}_n$, by Theorem 5.16 below. This universal property is the module-theoretic payoff of the annihilator closure and prepares the classification below.

Theorem 5.16. *For an ideal $I \in \text{Idl}(R_n)$, the following conditions are equivalent:*

- (i) $M_I = R_n/I$ is torsionless;
- (ii) $I \in \mathfrak{D}_n$;
- (iii) M_I embeds into R_n ;
- (iv) M_I is isomorphic to a principal ideal of R_n .

For a nonzero proper ideal, these conditions hold precisely when $I = D_a$, $1 \leq a_i \leq n_i - 1$ ($1 \leq i \leq r$). In that case, for every $\mathbf{c} = (c_1, \dots, c_r) \in (k^\times)^r$, put $y_{\mathbf{a}, \mathbf{c}} = \sum_{i=1}^r c_i x_i^{n_i - a_i}$. Then

$$\text{Ann}(y_{\mathbf{a}, \mathbf{c}}) = D_a, \quad (5.31)$$

and

$$R_n/D_a \longrightarrow R_n, \quad \varrho + D_a \longmapsto \varrho y_{\mathbf{a}, \mathbf{c}}, \quad (5.32)$$

is an injective homomorphism with image $\langle y_{\mathbf{a}, \mathbf{c}} \rangle$. For fixed $\mathbf{a}$, the distinct principal ideals arising in this way are parameterized by the projective torus $(k^\times)^r/k^\times$, and hence number $(q-1)^{r-1}$.

Proof. The equivalence of (i) and (ii) follows from Proposition 5.14 and the definition of $\mathfrak{D}_{\mathbf{n}}$. By Proposition 5.6, the nonzero proper ideals in the duality core are exactly the direct ideals $D_{\mathbf{a}}$ with $1 \leq a_i \leq n_i - 1$.

Assume that $I = D_{\mathbf{a}}$ has this form. Every exponent $n_i - a_i$ lies between 1 and $n_i - 1$. Let $y = y_{\mathbf{a}, \mathbf{c}}$. If $z = \sum_i f_i(x_i) \in J$, then, because different branches annihilate one another, $zy = \sum_{i=1}^r c_i f_i(x_i) x_i^{n_i - a_i}$. The direct-sum decomposition of J shows that $zy = 0$ if and only if $f_i(x_i) x_i^{n_i - a_i} = 0$ for every i , which is equivalent to $f_i(x_i) \in \langle x_i^{a_i} \rangle$ for every i . A nonzero constant term cannot annihilate the nonzero element y . Hence, $\text{Ann}(y) = D_{\mathbf{a}}$, which proves (5.31). The map in (5.32) is well defined, its kernel is $\text{Ann}(y)/D_{\mathbf{a}} = 0$, and its image is the principal ideal $\langle y \rangle$. Thus, (ii) implies both (iii) and (iv). The boundary cases $I = 0$ and $I = R_{\mathbf{n}}$ correspond, respectively, to the identity embedding of $R_{\mathbf{n}}$ and to the zero module.

Condition (iii) implies (iv), because the image of a cyclic module under an $R_{\mathbf{n}}$ -linear map is a cyclic ideal. Conversely, suppose that $M_I \cong \langle y \rangle$ for some $y \in R_{\mathbf{n}}$. The annihilator of M_I is I , whereas the annihilator of $\langle y \rangle$ is $\text{Ann}(y)$; therefore, $I = \text{Ann}(y)$. By Lemma 5.1, $I^{\perp\perp} = \text{Ann}(\text{Ann}(\text{Ann}(y))) = \text{Ann}(y) = I$. Thus, (iv) implies (ii), which completes the equivalence.

Finally, the valuation vector of $y_{\mathbf{a}, \mathbf{c}}$ is $(n_i - a_i)_{i=1}^r$ and its branch support is all of $[r]$. The toric parameterization in Proposition 4.7 shows that two such elements generate the same principal ideal precisely when their coefficient vectors differ by a common nonzero scalar. Hence the distinct realizations are parameterized by $(k^\times)^r / k^\times$ and their number is $(q - 1)^{r-1}$. $\square$

The classification immediately yields a dimension-refined enumerator.

Corollary 5.17. *Distinct ideals define nonisomorphic cyclic quotients. Consequently, the number of isomorphism classes of torsionless cyclic $R_{\mathbf{n}}$ -modules, including 0 and $R_{\mathbf{n}}$, is*

$$2 + \prod_{i=1}^r (n_i - 1). \quad (5.33)$$

Their dimension enumerator is

$$\mathcal{T}_{\mathbf{n}}(z) = 1 + z^{N_{\mathbf{n}}} + z^{1-r} \prod_{i=1}^r (z + z^2 + \cdots + z^{n_i-1}), \quad (5.34)$$

where the coefficient of z^e is the number of torsionless cyclic modules of k -dimension e .

Proof. For every ideal I one has $\text{Ann}_{R_{\mathbf{n}}}(R_{\mathbf{n}}/I) = I$; therefore, isomorphic cyclic quotients have equal defining ideals. By Theorem 5.16, the nonzero proper torsionless cyclic quotients are $R_{\mathbf{n}}/D_{\mathbf{a}}$ with $1 \leq a_i \leq n_i - 1$. Their number is $\prod_i (n_i - 1)$, and adding the zero and regular modules gives (5.33). Moreover, $\dim_k(R_{\mathbf{n}}/D_{\mathbf{a}}) = N_{\mathbf{n}} - \sum_{i=1}^r (n_i - a_i) = 1 - r + \sum_{i=1}^r a_i$. Summing over all admissible cutoff vectors gives the product in (5.34); the terms 1 and $z^{N_{\mathbf{n}}}$ account for the zero and regular modules. $\square$

The next result determines the entire bidual, not only the kernel of the evaluation map. It exhibits a second, uniform obstruction to Frobenius-like duality.

Proposition 5.18. *Let $0 \neq I \subsetneq R_{\mathbf{n}}$ and write $I^{\perp} = D_{\mathbf{b}}$, $1 \leq b_i \leq n_i - 1$. For each i , put $B_i = \langle x_i^{b_i} \rangle$ and define $\mathbf{c}^{(i)} = (c_1^{(i)}, \dots, c_r^{(i)})$ by*

$$c_j^{(i)} = \begin{cases} b_i, & j = i, \\ n_j - 1, & j \neq i. \end{cases} \quad (5.35)$$

Then

$$\mathrm{Hom}_{R_{\mathbf{n}}}(B_i, R_{\mathbf{n}}) \cong D_{\mathbf{c}^{(i)}}, \quad (5.36)$$

and

$$M_I^{\vee\vee} \cong \bigoplus_{i=1}^r D_{\mathbf{c}^{(i)}}. \quad (5.37)$$

Under this isomorphism, the evaluation map is

$$\eta_{M_I}(\varrho + I) = (\varrho x_1^{b_1}, \dots, \varrho x_r^{b_r}). \quad (5.38)$$

Proof. Since every $b_i < n_i$, one has the internal direct-sum decomposition $D_{\mathbf{b}} = B_1 \oplus \dots \oplus B_r$. Therefore, $M_I^{\vee\vee} \cong \mathrm{Hom}_{R_{\mathbf{n}}}(D_{\mathbf{b}}, R_{\mathbf{n}}) \cong \bigoplus_{i=1}^r \mathrm{Hom}_{R_{\mathbf{n}}}(B_i, R_{\mathbf{n}})$, where finite direct products and direct sums coincide.

Fix i . The cyclic module $B_i = R_{\mathbf{n}} x_i^{b_i}$ has defining relations given by $\mathrm{Ann}(x_i^{b_i})$. An $R_{\mathbf{n}}$ -linear map $f : B_i \rightarrow R_{\mathbf{n}}$ is therefore determined by $y = f(x_i^{b_i})$, and such a value is permissible exactly when $\mathrm{Ann}(x_i^{b_i})y = 0$. Thus, evaluation at the generator gives an isomorphism $\mathrm{Hom}_{R_{\mathbf{n}}}(B_i, R_{\mathbf{n}}) \cong \mathrm{Ann}(\mathrm{Ann}(x_i^{b_i})) = B_i^{\perp\perp}$. The ideal B_i has canonical cutoff b_i on the i th branch and absent cutoff n_j on every branch $j \neq i$. Applying Theorem 5.5 gives $B_i^{\perp\perp} = D_{\mathbf{c}^{(i)}}$ with $\mathbf{c}^{(i)}$ as in (5.35). This proves (5.36) and (5.37).

Finally, under the identification $M_I^{\vee} = D_{\mathbf{b}}$, the evaluation functional associated with $\varrho + I$ is multiplication by ϱ on $D_{\mathbf{b}}$. Its restriction to B_i sends the generator $x_i^{b_i}$ to $\varrho x_i^{b_i}$. Hence its coordinate in $\mathrm{Hom}_{R_{\mathbf{n}}}(B_i, R_{\mathbf{n}}) \cong D_{\mathbf{c}^{(i)}}$ is $\varrho x_i^{b_i}$, which proves (5.38). $\square$

The branchwise decomposition now isolates the ambient contribution to the bidual and yields the second main cyclic-module result.

Theorem 5.19. *Let $0 \neq I \subsetneq R_{\mathbf{n}}$ and set $M = M_I$. Then*

$$\dim_k M^{\vee} = \dim_k M + r - 1 - \delta_{\mathrm{Ann}}(I), \quad (5.39)$$

$$\dim_k M^{\vee\vee} = \dim_k M + r^2 - 1 - \delta_{\mathrm{Ann}}(I), \quad (5.40)$$

$$\dim_k \ker(\eta_M) = \delta_{\mathrm{Ann}}(I), \quad (5.41)$$

$$\dim_k \mathrm{coker}(\eta_M) = r^2 - 1. \quad (5.42)$$

The first three quantities retain the ideal-dependent annihilator defect, whereas the cokernel dimension is independent of I . Thus, bidual evaluation separates into an ideal-sensitive injectivity defect and an ambient surjectivity defect determined solely by the embedding dimension r .

Proof. Write $I = I(\mathbf{a}, W)$, put $d = \dim_k W$, and let $Z(I)$ be the absent-branch set from Theorem 5.5. By Proposition 5.12, $M^{\vee} \cong D_{\mathbf{b}}$, where $\mathbf{b}$ is given by (5.5). We compute $\dim_k D_{\mathbf{b}} = \sum_i (n_i - b_i)$ coordinatewise.

If $i \in T(W)$, then $b_i = n_i - a_i + 1$, so $n_i - b_i = a_i - 1$. If $i \notin T(W)$ and $a_i < n_i$, then $b_i = n_i - a_i$, so $n_i - b_i = a_i$. If $i \in Z(I)$, then $a_i = n_i$, $b_i = 1$, and $n_i - b_i = a_i - 1$. Therefore,

$$\dim_k M^{\vee} = \sum_{i=1}^r a_i - |T(W)| - |Z(I)|. \quad (5.43)$$

On the other hand, using (3.4), $\dim_k M = N_{\mathbf{n}} - \dim_k I = 1 - r + \sum_{i=1}^r a_i - d$. Subtracting this expression from (5.43) gives $\dim_k M^\vee - \dim_k M = r - 1 - (|T(W)| - d + |Z(I)|)$, and the term in parentheses is $\delta_{\text{Ann}}(I)$ by (5.17), which proves (5.39).

By Proposition 5.18, $\dim_k M^{\vee\vee} = \sum_{i=1}^r \dim_k D_{\mathbf{c}^{(i)}}$, where $\dim_k D_{\mathbf{c}^{(i)}} = (n_i - b_i) + \sum_{j \neq i} (n_j - (n_j - 1)) = n_i - b_i + r - 1$. Hence, $\dim_k M^{\vee\vee} = \sum_{i=1}^r (n_i - b_i) + r(r - 1) = \dim_k M^\vee + r(r - 1)$, and substituting (5.39) yields (5.40).

Equation (5.41) is (5.29). Finally, $\dim_k \text{im}(\eta_M) = \dim_k M - \dim_k \ker(\eta_M) = \dim_k M - \delta_{\text{Ann}}(I)$, and therefore

$$\begin{aligned} \dim_k \text{coker}(\eta_M) &= \dim_k M^{\vee\vee} - \dim_k \text{im}(\eta_M) \\ &= (\dim_k M + r^2 - 1 - \delta_{\text{Ann}}(I)) - (\dim_k M - \delta_{\text{Ann}}(I)) = r^2 - 1, \end{aligned}$$

which proves (5.42). $\square$

The universal cokernel formula gives the reflexivity dichotomy.

Corollary 5.20. *For every nonzero proper cyclic quotient M_I , $\dim_k \text{coker}(\eta_{M_I}) = r^2 - 1$. Consequently:*

- (i) *if $r \geq 2$, no nonzero proper cyclic quotient is reflexive;*
- (ii) *if $r = 1$, every cyclic quotient is reflexive.*

Hence, $r^2 - 1$ is the exact bidual-surjectivity obstruction separating the chain case from the genuinely multi-branch case.

Proof. The universal cokernel formula is (5.42). If $r \geq 2$, then $r^2 - 1 > 0$, so the evaluation map of every nonzero proper cyclic quotient fails to be surjective. If $r = 1$, the ring is a finite chain ring and every ideal belongs to the duality core; Proposition 5.14 therefore gives injectivity, while (5.42) gives surjectivity. The zero and regular modules are reflexive in every case. $\square$

The next example separates the ideal-sensitive kernel from the ambient cokernel in a concrete three-branch quotient.

Example 5.21. *Let $k = \mathbb{F}_3$ and $\mathbf{n} = (4, 3, 3)$. Consider first the direct ideal $D_{\mathbf{a}} = D_{(2,1,2)}$. Since every coordinate of $\mathbf{a}$ lies between 1 and $n_i - 1$, the cyclic module $M = R_{\mathbf{n}}/D_{(2,1,2)}$ is torsionless. Taking $y = x_1^2 + x_2^2 + x_3$ gives $\text{Ann}(y) = D_{(2,1,2)}$, $M \cong \langle y \rangle \subseteq R_{\mathbf{n}}$ and $\dim_k M = 1 - 3 + (2 + 1 + 2) = 3$. Here $\delta_{\text{Ann}}(D_{\mathbf{a}}) = 0$, and Theorem 5.19 gives*

$$\begin{aligned} \dim_k M^\vee &= 5, & \dim_k M^{\vee\vee} &= 11, \\ \dim_k \ker(\eta_M) &= 0, & \dim_k \text{coker}(\eta_M) &= 8. \end{aligned}$$

Thus, M embeds into its bidual but is not reflexive.

Now, let $I = I(\mathbf{a}', W)$ be the coupled ideal from Example 5.10, where $\mathbf{a}' = (3, 2, 2)$ and $W = \text{span}_{\mathbb{F}_3}\{e_1 + e_3, e_2 + e_3\}$. There $\delta_{\text{Ann}}(I) = 1$, and since $\dim_k I = 5$, one has $\dim_k M_I = 3$. Therefore,

$$\begin{aligned} \dim_k M_I^\vee &= 4, & \dim_k M_I^{\vee\vee} &= 10, \\ \dim_k \ker(\eta_{M_I}) &= 1, & \dim_k \text{coker}(\eta_{M_I}) &= 8. \end{aligned}$$

The coupled terminal plane creates the one-dimensional obstruction to torsionlessness, whereas the eight-dimensional obstruction to bidual surjectivity is imposed uniformly by the three-branch ambient ring. In the chain specialization $r = 1$, this universal cokernel vanishes, so the same calculation isolates the precise point at which reflexivity fails.

Remark 5.22. Formula (5.42) is independent of the ideal and decomposes as $r^2 - 1 = (r - 1) + r(r - 1)$. The first term is the excess of the socle dimension over the Frobenius value 1; the second arises from the r branch summands of (5.37), each of which acquires under $\text{Hom}(-, R_{\mathbf{n}})$ the $r - 1$ terminal socle directions of the remaining branches. The cokernel therefore depends only on r , and not on q , $\mathbf{n}$, or I .

The kernel, by contrast, is ideal-dependent and obeys the sharp bound

$$0 \leq \delta_{\text{Ann}}(I) \leq r - 1 \quad (0 \neq I \subsetneq R_{\mathbf{n}}). \quad (5.44)$$

Indeed, by (5.17), if $W = 0$, then $\kappa(I) = 0$, while $Z(I) = \{i : a_i = n_i\}$ is a proper subset of $[r]$ because $I \neq 0$; and if $W \neq 0$, then $\kappa(I) \leq |T(W)| - 1$ by (4.10) and $Z(I) \subseteq [r] \setminus T(W)$. The bound is attained in Example 5.23. Bidual evaluation therefore separates an ideal-sensitive kernel defect, ranging over $\{0, 1, \dots, r - 1\}$, from the constant ambient cokernel $r^2 - 1$.

These formulas were also verified computationally for the profiles (4), (3, 2), (3, 3), (4, 3), (2, 2, 2), and (3, 2, 2) over $\mathbb{F}_2$, and (3, 3) over $\mathbb{F}_3$, by computing the relevant Hom-spaces directly for every nonzero proper ideal. All the identities (5.39)–(5.42) were confirmed, including the chain case $r = 1$, where the cokernel vanishes.

Example 5.23. We continue Example 3.12, with $k = \mathbb{F}_2$, $\mathbf{n} = (4, 3, 2, 2)$, and $\mathbf{a} = (4, 2, 2, 2)$, in order to test (5.17) where both of its terms are active. For $I = I(\mathbf{a}, W)$ with $W = \text{span}_{\mathbb{F}_2}\{e_2 + e_3, e_1 + e_4\}$ one has $T(W) = \{1, 2, 3, 4\}$, $\dim_k I = 3$, and $Z(I) = \emptyset$, so (5.5) and (5.15) give

$$I^\perp = D_{(1,2,1,1)} = \langle x_1, x_2^2, x_3, x_4 \rangle, \quad I^{\perp\perp} = D_{(3,1,1,1)} = \langle x_1^3, x_2, x_3, x_4 \rangle,$$

of dimensions 6 and 5, whence $\delta_{\text{Ann}}(I) = 5 - 3 = 2 = \kappa(I) + |Z(I)| = 2 + 0$. For the smaller coupling $I' = I(\mathbf{a}, \text{span}_{\mathbb{F}_2}\{e_2 + e_3\})$ one has $\dim_k I' = 2$, support $\{2, 3\}$, $\kappa(I') = 1$, and $Z(I') = \{1, 4\}$, since $a_1 = n_1$ and $a_4 = n_4$ while neither index lies in the support. Theorem 5.2 returns the same annihilator $D_{(1,2,1,1)}$, an instance of the fibers of Proposition 5.4, so $I'^{\perp\perp} = D_{(3,1,1,1)}$ again and $\delta_{\text{Ann}}(I') = 5 - 2 = 3 = 1 + 2$: Both terms are needed, and by (5.44) the value $3 = r - 1$ is the largest possible. The two cyclic quotients consequently have the same dual and the same bidual, of dimensions 6 and 18, and the same cokernel $r^2 - 1 = 15$, but kernels 2 and 3; the ideal-sensitive part of biduality is exactly δ_{Ann} . All these values agree with a direct computation of $\text{Hom}(M, R_{\mathbf{n}})$.

The preceding results reveal two distinct sources of non-Frobenius behavior. The closure $I \mapsto I^{\perp\perp}$ controls torsionlessness, with defect determined by the canonical branch–subspace data, whereas every nonzero proper cyclic quotient has the universal bidual cokernel of dimension $r^2 - 1$. Thus, the kernel obstruction is ideal-dependent, while the failure of reflexivity is governed solely by the embedding dimension r .

6. Automorphisms

The preceding classifications are expressed in canonical branch coordinates. This section determines how automorphisms act on these data by computing the tangent image and the nonlinear kernel,

assembling the full automorphism group, and describing its socle and depth-two actions. This separates the freely mixing square-zero directions from the monomially rigid long branches and identifies a subgroup acting trivially on the depth-two ideal lattice.

For $d \geq 2$, put $\mathcal{I}_d = \{i : n_i = d\}$ and $r_d = |\mathcal{I}_d|$. Throughout this section, let $\mathcal{Q} = \mathcal{I}_2$, $\mathcal{L} = [r] \setminus \mathcal{Q}$, and $r_+ = |\mathcal{L}| = r - r_2$. With $\bar{x}_i = x_i + J^2$, define $V_2 = \text{span}_k\{\bar{x}_i : i \in \mathcal{Q}\}$ and $V_+ = \text{span}_k\{\bar{x}_i : i \in \mathcal{L}\}$, so that $J/J^2 = V_2 \oplus V_+$. For $d \geq 3$, let $\text{Mon}_{r_d}(k) = (k^\times)^{r_d} \rtimes \mathfrak{S}_{r_d}$ denote the group of monomial matrices on the branches of length d . Empty factors are understood to be trivial.

Every k -algebra automorphism preserves J and therefore induces an invertible linear transformation of J/J^2 . We write $\rho_1 : \text{Aut}_k(R_{\mathbf{n}}) \rightarrow \text{GL}_k(J/J^2)$ for this representation and call it the *tangent shadow*.

The first automorphism result determines the full linear image on the first layer.

Proposition 6.1. *Relative to the decomposition $J/J^2 = V_2 \oplus V_+$, the image of ρ_1 is*

$$L_{\mathbf{n}} = \left\{ \begin{pmatrix} A & B \\ 0 & \Pi \end{pmatrix} : \begin{array}{l} A \in \text{GL}(V_2), \\ B \in \text{Hom}_k(V_+, V_2), \\ \Pi \in \prod_{d \geq 3} \text{Mon}_{r_d}(k) \end{array} \right\}. \quad (6.1)$$

Equivalently, square-zero branches admit arbitrary linear changes of basis, any long first-layer direction may be sheared by a square-zero vector, and the long branches may only be rescaled and permuted among branches of the same nilpotency length. In particular,

$$|L_{\mathbf{n}}| = q^{r_2 r_+} |\text{GL}_{r_2}(q)| \prod_{d \geq 3} (q-1)^{r_d} r_d!. \quad (6.2)$$

Proof. Let $\varphi \in \text{Aut}_k(R_{\mathbf{n}})$ and write $\varphi(x_i) \equiv \sum_{j=1}^r a_{ji} x_j \pmod{J^2}$, so that (a_{ji}) is the matrix of $\rho_1(\varphi)$. Since φ is an automorphism, this matrix is invertible.

We first compare source and target nilpotency lengths. Fix a source index i and put $d = n_i$. The relation $x_i^d = 0$ gives $\varphi(x_i)^d = 0$. Suppose that $a_{ji} \neq 0$ for some target branch with $n_j > d$. The component of $\varphi(x_i)$ in the j th branch has the form $a_{ji} x_j + \text{terms of } x_j\text{-degree at least } 2$, and its d th power has lowest nonzero term $a_{ji}^d x_j^d$. Because $d < n_j$, the class x_j^d is nonzero, and no term from another branch can cancel it, since distinct branches have zero product and are linearly independent. This contradicts $\varphi(x_i)^d = 0$. Hence,

$$a_{ji} \neq 0 \implies n_j \leq n_i. \quad (6.3)$$

In particular, if $i \in \mathcal{Q}$, then the i th column is supported entirely on the rows indexed by $\mathcal{Q}$. Relative to $V_2 \oplus V_+$, the matrix is therefore block upper triangular, $\begin{pmatrix} A & B \\ 0 & C \end{pmatrix}$, and its invertibility implies that $A \in \text{GL}(V_2)$ and $C \in \text{GL}(V_+)$.

We next use the relations $x_i x_h = 0$ for $i \neq h$. Fix a long target branch $j \in \mathcal{L}$, so $n_j \geq 3$. Taking the j th branch component of $\varphi(x_i) \varphi(x_h) = 0$ and comparing the coefficient of x_j^2 gives $a_{ji} a_{jh} = 0$, ($i \neq h$), since the degree-two term in that branch is exactly $a_{ji} a_{jh} x_j^2$ and $x_j^2 \neq 0$. Thus, each row of the long block C contains at most one nonzero entry. Since C is invertible, every row and every column contains exactly one nonzero entry, so C is monomial.

Let σ be the permutation of the long indices determined by this monomial matrix, so that the source branch $i \in \mathcal{L}$ has its unique long first-layer component in the target branch $\sigma(i)$. By (6.3), $n_{\sigma(i)} \leq n_i$ for

every $i \in \mathcal{L}$. Since σ is a permutation of $\mathcal{L}$, summing these inequalities gives equality of the two sums; hence equality holds termwise, and $n_{\sigma(i)} = n_i$ for every $i \in \mathcal{L}$. Therefore, C is block monomial with respect to the length classes $\mathcal{I}_d$ with $d \geq 3$, that is, $C \in \prod_{d \geq 3} \text{Mon}_{r_d}(k)$. No relation restricts the block B : Vectors in V_2 square to zero and annihilate all of J , so they may be added freely to the first-layer images of long branches.

It remains to prove that every matrix in (6.1) is realized. Choose A , B , and a length-preserving monomial transformation Π . For each long index i , write the corresponding monomial image as $\lambda_i x_{\sigma(i)}$, where $\lambda_i \in k^\times$ and $n_{\sigma(i)} = n_i$. Define a substitution by

$$x_i \mapsto \begin{cases} \sum_{j \in Q} A_{ji} x_j, & i \in Q, \\ \lambda_i x_{\sigma(i)} + \sum_{j \in Q} B_{ji} x_j, & i \in \mathcal{L}. \end{cases}$$

The image of a square-zero generator lies in the square-zero space V_2 and therefore squares to zero. The image of a long generator has one long component of the same nilpotency length plus a square-zero component, so its n_i th power is zero. Images of distinct generators multiply to zero, because their long components lie in distinct branches and every square-zero component annihilates J . The substitution therefore defines a k -algebra endomorphism. Its tangent matrix is invertible, so its images generate J/J^2 ; Nakayama's lemma implies that they generate J , and the endomorphism is consequently surjective. Since R_n is finite, it is an automorphism. This proves (6.1).

Finally, the block B contributes $q^{r_2 r_+}$ choices, the block A contributes $|\text{GL}_{r_2}(q)|$, and the length- d monomial block contributes $(q-1)^{r_d} r_d!$. Multiplying these independent factors gives (6.2). $\square$

For $d \geq 2$, let $\mathcal{N}_d(k)$ be the truncated substitution group

$$\mathcal{N}_d(k) = \{y \mapsto y + \gamma_2 y^2 + \cdots + \gamma_{d-1} y^{d-1} : \gamma_i \in k\},$$

with composition taken modulo y^d ; thus, $|\mathcal{N}_d(k)| = q^{d-2}$ and $\mathcal{N}_2(k) = 1$. Define the additive cross-socle space $\mathcal{X}_n = \{\mathbf{c} = (c_{ji})_{j \in \mathcal{L}, i \neq j} : c_{ji} \in k\}$, with $\dim_k \mathcal{X}_n = (r-1)r_+$. An element $\mathbf{c} \in \mathcal{X}_n$ determines the substitution $\gamma_{\mathbf{c}}(x_i) = x_i + \sum_{\substack{j \in \mathcal{L} \\ j \neq i}} c_{ji} x_j^{n_j-1}$.

After determining the linear image, we identify precisely the automorphisms that are invisible on the tangent space.

Proposition 6.2. *Let $\mathcal{U}_n = \ker \rho_1$. Every element of $\mathcal{U}_n$ has a unique normal form*

$$x_i \mapsto x_i + f_i(x_i) + \sum_{\substack{j \in \mathcal{L} \\ j \neq i}} c_{ji} x_j^{n_j-1}, \quad f_i(x_i) \in x_i^2 k[x_i] / \langle x_i^{n_i} \rangle. \quad (6.4)$$

Moreover, the branch substitutions commute with the cross-socle translations, and

$$\mathcal{U}_n \cong \left(\prod_{i=1}^r \mathcal{N}_{n_i}(k) \right) \times (\mathcal{X}_n, +). \quad (6.5)$$

Consequently,

$$|\mathcal{U}_n| = q^{\sum_{i=1}^r (n_i-2) + (r-1)r_+}. \quad (6.6)$$

Proof. Let $\varphi \in \mathcal{U}_{\mathbf{n}}$. Since φ induces the identity on J/J^2 , each generator has the form $\varphi(x_i) = x_i + h_i$ with $h_i \in J^2$. Using the direct branch decomposition of J^2 , write uniquely $h_i = f_i(x_i) + \sum_{j \neq i} h_{ji}(x_j)$, where $f_i(x_i) \in x_i^2 k[x_i]/\langle x_i^{n_i} \rangle$ and $h_{ji}(x_j) \in x_j^2 k[x_j]/\langle x_j^{n_j} \rangle$.

Fix distinct indices i and j . The relation $\varphi(x_i)\varphi(x_j) = 0$ has j th branch component equal to $h_{ji}(x_j)(x_j + f_j(x_j))$. The element $x_j + f_j(x_j)$ has x_j -adic valuation one and can be written as $x_j \zeta_j(x_j)$ with $\zeta_j(0) = 1$. Hence, $\zeta_j(x_j)$ is a unit in $k[x_j]/\langle x_j^{n_j} \rangle$, and the annihilator of $x_j + f_j(x_j)$ inside the j th branch is the one-dimensional terminal space $kx_j^{n_j-1}$. Therefore, $h_{ji}(x_j) = c_{ji}x_j^{n_j-1}$ for some $c_{ji} \in k$. If $n_j = 2$, then J^2 has no nonzero j th branch component, so necessarily $c_{ji} = 0$. This proves that every element of the kernel has the form (6.4), and uniqueness follows from the direct branch decomposition.

Conversely, consider a substitution of the form (6.4) and put $y_i = x_i + f_i(x_i) + \sum_{j \neq i} c_{ji}x_j^{n_j-1}$. The same-branch term $x_i + f_i(x_i)$ is divisible by x_i , and its n_i th power is zero. Every cross term is a socle element and therefore has square zero, and products involving distinct branches vanish. It follows that $y_i^{n_i} = 0$. If $i \neq h$, then every nonzero same-branch contribution to $y_i y_h$ is a product of $x_j + f_j(x_j)$ with a multiple of $x_j^{n_j-1}$, or a product of two terminal socle elements; all such products are zero. Thus, $y_i y_h = 0$, and the substitution respects all defining relations. The induced map on J/J^2 is the identity, so by Nakayama's lemma the resulting endomorphism is surjective and therefore bijective. Hence every parameter choice in (6.4) gives an element of $\mathcal{U}_{\mathbf{n}}$.

It remains to establish the group decomposition. Branch substitutions on distinct branches commute. A branch substitution $x_j \mapsto x_j + f_j(x_j)$ fixes the terminal element $x_j^{n_j-1}$, because $(x_j + f_j(x_j))^{n_j-1} = x_j^{n_j-1}$; every other term has degree at least n_j . Conversely, inserting a sum of cross-socle terms into a polynomial f_i of degree at least two produces no additional term, since all mixed products vanish and each terminal socle term has square zero. Therefore, every branch substitution commutes with every cross-socle translation.

If $\mathbf{c}, \mathbf{c}' \in \mathcal{X}_{\mathbf{n}}$, then $\gamma_{\mathbf{c}} \circ \gamma_{\mathbf{c}'} = \gamma_{\mathbf{c}+\mathbf{c}'}$, because each terminal element $x_j^{n_j-1}$ with $j \in \mathcal{L}$ satisfies $n_j - 1 \geq 2$ and is therefore fixed by every cross-socle translation. Thus, the cross-socle subgroup is naturally the additive group of $\mathcal{X}_{\mathbf{n}}$, and the normal form is a direct product as in (6.5).

Finally, the branch substitution factors contribute $q^{\sum_i(n_i-2)}$, and $\dim_k \mathcal{X}_{\mathbf{n}} = (r-1)r_+$, since each long target branch $j \in \mathcal{L}$ carries $r-1$ independent cross coefficients. This proves (6.6). $\square$

The tangent image and the nonlinear kernel now assemble into the full automorphism group.

Theorem 6.3. *The tangent representation gives a split exact sequence $1 \longrightarrow \mathcal{U}_{\mathbf{n}} \longrightarrow \text{Aut}_k(R_{\mathbf{n}}) \xrightarrow{\rho_1} L_{\mathbf{n}} \longrightarrow 1$. Hence,*

$$\text{Aut}_k(R_{\mathbf{n}}) \cong \mathcal{U}_{\mathbf{n}} \rtimes L_{\mathbf{n}}. \quad (6.7)$$

If $k = \mathbb{F}_{p^m}$, then

$$\text{Aut}(R_{\mathbf{n}}) \cong (\mathcal{U}_{\mathbf{n}} \rtimes L_{\mathbf{n}}) \rtimes \text{Gal}(k/\mathbb{F}_p). \quad (6.8)$$

In particular,

$$|\text{Aut}_k(R_{\mathbf{n}})| = q^{\sum_i(n_i-2)+(r-1)r_+ + r_2 r_+} |\text{GL}_{r_2}(q)| \prod_{d \geq 3} (q-1)^{r_d} r_d!, \quad (6.9)$$

$$|\text{Aut}(R_{\mathbf{n}})| = m |\text{Aut}_k(R_{\mathbf{n}})|. \quad (6.10)$$

Proof. By definition, the kernel of ρ_1 is $\mathcal{U}_n$, and Proposition 6.1 identifies its image with L_n . The realization constructed in the proof of that proposition assigns to every matrix $P \in L_n$ a linear substitution $s(P)$ preserving the defining relations $x_i^{n_i} = 0$ and $x_i x_j = 0$ ($i \neq j$), as verified in the second step below.

We verify that this assignment is a group homomorphism and a right inverse to ρ_1 .

First, L_n is a group: By (6.1) its elements are the block matrices $\begin{pmatrix} A & B \\ 0 & \Pi \end{pmatrix}$ relative to $J/J^2 = V_2 \oplus V_+$, and

$$\begin{pmatrix} A & B \\ 0 & \Pi \end{pmatrix} \begin{pmatrix} A' & B' \\ 0 & \Pi' \end{pmatrix} = \begin{pmatrix} AA' & AB' + B\Pi' \\ 0 & \Pi\Pi' \end{pmatrix}, \quad (6.11)$$

where $\Pi\Pi'$ is again monomial and length-preserving and AA' again invertible, while inverses have the same shape because $\begin{pmatrix} A & B \\ 0 & \Pi \end{pmatrix}^{-1} = \begin{pmatrix} A^{-1} & -A^{-1}B\Pi^{-1} \\ 0 & \Pi^{-1} \end{pmatrix}$.

Second, write $P = (P_{ji}) \in L_n$ in the coordinates of Proposition 6.1, so that P_{ji} is the coefficient of x_j in the image of x_i , and let $s(P)(x_i) = \sum_{j=1}^r P_{ji}x_j$ be the linear substitution realizing P . The block shape (6.1) says precisely that $P_{ji} \neq 0$ forces $n_j \leq n_i$, and that a row indexed by a branch with $n_j \geq 3$ has at most one nonzero entry. The first condition gives $s(P)(x_i)^{n_i} = \sum_j P_{ji}^{n_i} x_j^{n_i} = 0$, since a nonzero P_{ji} forces $x_j^{n_i} = 0$; the second gives $s(P)(x_i) s(P)(x_{i'}) = \sum_j P_{ji} P_{ji'} x_j^2 = 0$ for $i \neq i'$, since $x_j^2 = 0$ when $n_j = 2$ while at most one of the two coefficients is nonzero when $n_j \geq 3$. Hence $s(P)$ respects both families of defining relations, and it is bijective because its induced map on J/J^2 is the invertible matrix P .

Third, for $P, Q \in L_n$ one has $(s(P) \circ s(Q))(x_i) = \sum_j Q_{ji} \sum_l P_{lj} x_l = \sum_l (PQ)_{li} x_l = s(PQ)(x_i)$, and two k -algebra maps agreeing on $x_1, \dots, x_r$ are equal, so s is a group homomorphism; being linear, $s(P)$ induces P on J/J^2 , that is, $\rho_1 \circ s = \text{id}_{L_n}$. Hence, s splits the exact sequence in the statement, and (6.7) follows, with $\mathcal{U}_n = \ker \rho_1$ normal.

The coefficient field is intrinsic by Lemma 2.4, so every ring automorphism restricts to an element of $\text{Gal}(k/\mathbb{F}_p)$, and the kernel of this restriction map is $\text{Aut}_k(R_n)$. Conversely, every field automorphism extends coefficientwise to R_n while fixing the branch generators: In the canonical form (2.1) the twist $\widehat{\sigma}(a + \sum_{i,t} a_{i,t} x_i^t) = \sigma(a) + \sum_{i,t} \sigma(a_{i,t}) x_i^t$ is a ring automorphism, because the defining relations are monomial with coefficients in the prime field, and $\widehat{\sigma\tau} = \widehat{\sigma} \widehat{\tau}$, so $\sigma \mapsto \widehat{\sigma}$ is a section of the restriction map. Hence the restriction sequence also splits, which gives (6.8). Combining Proposition 6.2 with (6.2) gives (6.9); multiplying by $|\text{Gal}(k/\mathbb{F}_p)| = m$ gives (6.10). $\square$

Table 3 records an independent check of the order formula: For each profile we enumerated all tuples $(y_1, \dots, y_r) \in J^r$ with $y_i^{n_i} = 0$ and $y_i y_j = 0$ ($i \neq j$) whose linear part is invertible, and compared the count with (6.9). The last two columns are used in Remark 6.6.

Table 3. Brute-force verification of (6.9) and the kernel K_n of the action of $\text{Aut}_k(R_n)$ on the ideal lattice compared with the subgroup $\mathcal{N}_n$ of Theorem 6.5.

$\mathbf{n}$	q	enumerated $ \text{Aut}_k $	(6.9)	$ K_n $	$ \mathcal{N}_n $
(2, 2)	2	6	6	1	1
(2, 2)	3	48	48	2	1
(2, 2, 2)	3	11232	11232	2	1
(3, 2)	2	8	8	2	2
(2, 4)	2	16	16	4	4
(3, 3)	2	32	32	4	4
(3, 3)	3	648	648	9	9
(4, 3)	2	32	32	8	8

Put $\Sigma_2 = \text{span}_k\{x_i : i \in \mathcal{Q}\}$, $\Sigma_d = \text{span}_k\{x_i^{d-1} : i \in \mathcal{I}_d\}$ ($d \geq 3$), and $\Sigma_+ = \bigoplus_{d \geq 3} \Sigma_d$, so that $\text{Soc}(R_{\mathbf{n}}) = \Sigma_2 \oplus \Sigma_+$. For $d \geq 3$, let $\Lambda_d = \{\lambda^{d-1} : \lambda \in k^\times\} \leq k^\times$ be the subgroup of $(d-1)$ st powers, and let $\text{Mon}_{r_d}(\Lambda_d)$ denote the monomial matrices whose nonzero entries belong to Λ_d .

The socle representation records a different quotient of the same automorphism group. It governs the orbits of socle ideals and is applied in Corollary 6.7 to settle those orbits completely in the square-zero profile.

Proposition 6.4. *The image of the natural representation $\rho_{\text{soc}} : \text{Aut}_k(R_{\mathbf{n}}) \longrightarrow \text{GL}_k(\text{Soc}(R_{\mathbf{n}}))$ consists exactly of the block matrices*

$$\Gamma_{\mathbf{n}} = \left\{ \begin{pmatrix} A' & 0 \\ \Xi & \Pi' \end{pmatrix} : \begin{array}{l} A' \in \text{GL}(\Sigma_2), \\ \Xi \in \text{Hom}_k(\Sigma_2, \Sigma_+), \\ \Pi' \in \prod_{d \geq 3} \text{Mon}_{r_d}(\Lambda_d) \end{array} \right\}. \quad (6.12)$$

Consequently, $|\Gamma_{\mathbf{n}}| = q^{r_2 r_+} |\text{GL}_{r_2}(q)| \prod_{d \geq 3} \left(\frac{q-1}{\gcd(q-1, d-1)} \right)^{r_d} r_d!$. Since every k -subspace of the socle is an ideal, the automorphism-orbit problem for socle ideals is exactly the orbit problem for $\Gamma_{\mathbf{n}}$ on the subspace lattice of $\Sigma_2 \oplus \Sigma_+$.

Proof. Take $\varphi \in \text{Aut}_k(R_{\mathbf{n}})$ and use the decomposition from Theorem 6.3. If $i \in \mathcal{Q}$, then $x_i^{n_i-1} = x_i$, and its image has an arbitrary invertible square-zero component coming from the block A of the tangent shadow. The cross-socle part of $\mathcal{U}_{\mathbf{n}}$ may add an arbitrary element of Σ_+ to each such image. Therefore, the images of the square-zero socle basis vectors have the form prescribed by the blocks A' and Ξ .

Now let $i \in \mathcal{I}_d$ with $d \geq 3$. The tangent shadow sends x_i to $\lambda_i x_{\sigma(i)} + v_i$ + terms of degree at least two, where $v_i \in V_2$, the permutation σ preserves $\mathcal{I}_d$, and $\lambda_i \in k^\times$. The nonlinear kernel may further add terminal socle terms from other long branches. Raising this expression to the power $d-1$ eliminates every square-zero and cross-socle summand, and higher same-branch terms also disappear because their contribution has degree at least d . Thus, $\varphi(x_i^{d-1}) = \varphi(x_i)^{d-1} = \lambda_i^{d-1} x_{\sigma(i)}^{d-1}$. Hence the action on Σ_d is monomial with nonzero coefficients in Λ_d . No long socle vector acquires a square-zero component, which explains the zero upper-right block in (6.12).

Conversely, choose a matrix in (6.12). The block A' is realized by a linear change of the square-zero generators. Each nonzero entry of the monomial block on Σ_d belongs to Λ_d , so it has the form λ^{d-1} for some $\lambda \in k^\times$ and is realized by rescaling and permuting the corresponding length- d branch generators. Finally, the block Ξ is realized by adding arbitrary long socle elements to the square-zero generators through cross-socle translations. Therefore, every displayed block matrix occurs.

The order formula follows because $|\Lambda_d| = (q-1)/\gcd(q-1, d-1)$. Finally, $J \text{Soc}(R_{\mathbf{n}}) = 0$, so every k -subspace of the socle is an ideal. The action of $\text{Aut}_k(R_{\mathbf{n}})$ on those ideals factors through ρ_{soc} and has image $\Gamma_{\mathbf{n}}$, which proves the orbit statement. $\square$

Let $\mathcal{N}_{\mathbf{n}} = \prod_{i=1}^r \mathcal{N}_{n_i}(k) \leq \mathcal{U}_{\mathbf{n}}$ be the subgroup of branchwise substitutions, with all cross-socle coefficients zero, and let $\text{Idl}_{\geq 2}(R_{\mathbf{n}}) = \{I \in \text{Idl}(R_{\mathbf{n}}) : I \subseteq J^2\}$ be the depth-two ideal sublattice.

The final theorem exhibits nonlinear symmetries that disappear on ideals and gives the effective action below the first radical layer.

Theorem 6.5. *The following statements hold.*

- (i) *Every element of $\mathcal{N}_{\mathbf{n}}$ fixes every ideal of $R_{\mathbf{n}}$.*

- (ii) Every element of the full nonlinear kernel $\mathcal{U}_{\mathbf{n}}$ fixes every ideal in $\text{Idl}_{\geq 2}(R_{\mathbf{n}})$.
- (iii) Hence the action of $\text{Aut}_k(R_{\mathbf{n}})$ on the depth-two ideal lattice factors through the length-preserving monomial group $\prod_{d \geq 3} \text{Mon}_{r_d}(k)$.

More explicitly, let $I(\mathbf{a}, W) \subseteq J^2$ and let a tangent-shadow element send each long branch i to $\lambda_i x_{\sigma(i)}$ modulo the square-zero space and J^2 . Define $\mathbf{a}'$ by $a'_{\sigma(i)} = a_i (i \in \mathcal{L}), a'_j = 2 (j \in \mathcal{Q})$, and define $\Psi_{\mathbf{a}} : E_{\mathbf{a}} \rightarrow E_{\mathbf{a}'}, \Psi_{\mathbf{a}}(e_i^{\mathbf{a}}) = \lambda_i^{a_i-1} e_{\sigma(i)}^{\mathbf{a}'}$ on the coordinates $i \in \mathcal{L}$ with $a_i \geq 3$, which are the only ones that can occur in W . Then

$$\varphi(I(\mathbf{a}, W)) = I(\mathbf{a}', \Psi_{\mathbf{a}}(W)). \quad (6.13)$$

Proof. Let $\psi \in \mathcal{N}_{\mathbf{n}}$. On the i th branch, $\psi(x_i) = x_i + f_i(x_i) = x_i \zeta_i(x_i), \zeta_i(0) = 1$. Therefore, $\zeta_i(x_i)$ is a unit and $\psi(\langle x_i^{a_i} \rangle) = \langle x_i^{a_i} \rangle$ for every cutoff a_i . Moreover, if $a_i \geq 2$, then $\psi(x_i^{a_i-1}) \equiv x_i^{a_i-1} \pmod{\langle x_i^{a_i} \rangle}$. Consequently, ψ fixes the direct cutoff ideal $D_{\mathbf{a}}$ and induces the identity on the terminal space $E_{\mathbf{a}}$. By the uniqueness of the canonical branch–subspace representation in Theorem 3.2, $\psi(I(\mathbf{a}, W)) = I(\mathbf{a}, W)$ for every canonical pair. This proves (i).

Now let $\gamma_{\mathbf{c}}$ be a cross-socle translation and suppose that $I(\mathbf{a}, W) \subseteq J^2$. Then every nonzero direct branch generator $x_i^{a_i}$ occurring in $D_{\mathbf{a}}$ has exponent at least two. In addition, if the coordinate $e_i^{\mathbf{a}}$ occurs in W , its lift $x_i^{a_i-1}$ also has exponent at least two; otherwise the ideal would contain a nonzero first-layer class and could not lie in J^2 . For every exponent $e \geq 2$, $\gamma_{\mathbf{c}}(x_i)^e = x_i^e$, because the added cross-socle terms square to zero and annihilate the i th branch. Thus, $\gamma_{\mathbf{c}}$ fixes every generator in a canonical generating set for $I(\mathbf{a}, W)$. Combining this with part (i) and the direct-product decomposition of $\mathcal{U}_{\mathbf{n}}$ proves (ii).

It follows that the action on $\text{Idl}_{\geq 2}(R_{\mathbf{n}})$ factors through the tangent shadow. The blocks A and B in (6.1) are supported on square-zero first-layer coordinates, and they disappear after taking powers of exponent at least two. Hence only the length-preserving monomial block acts on a depth-two ideal, which proves the factorization in (iii).

Finally, the image of $x_i^{a_i}$ is a unit multiple of $x_{\sigma(i)}^{a_i}$, so the direct cutoff vector is permuted as stated. On the terminal quotient, $\varphi(x_i^{a_i-1}) \equiv \lambda_i^{a_i-1} x_{\sigma(i)}^{a_i-1} \pmod{D_{\mathbf{a}'}}$. Therefore, the induced map on the terminal subspace is $\Psi_{\mathbf{a}}$, and the uniqueness of canonical data gives (6.13). $\square$

Remark 6.6. This section determines the full automorphism group of $R_{\mathbf{n}}$: its order, its tangent image, its nonlinear kernel, the two semidirect decompositions, and the induced socle action. For the action on ideals, Theorem 6.5 identifies a subgroup $\mathcal{N}_{\mathbf{n}}$ of order $q^{\sum_i (n_i-2)}$ that fixes every ideal and determines the action on the depth-two sublattice, and in the square-zero profile the action is described completely by Corollary 6.7.

The full orbit decomposition of $\text{Idl}(R_{\mathbf{n}})$ and the exact kernel $K_{\mathbf{n}}$ of the action remain open in general. One always has $\mathcal{N}_{\mathbf{n}} \subseteq K_{\mathbf{n}}$, but equality may fail: If $n_i = 2$ for every i , then $\mathcal{N}_{\mathbf{n}} = 1$, whereas $K_{\mathbf{n}} = k^{\times}$, since the scalars are exactly the elements of $\text{GL}(J)$ fixing every subspace of J . Thus, Theorem 6.5 identifies a trivially acting subgroup, not necessarily the whole kernel. In every non-square-zero profile of Table 3, however, we found $K_{\mathbf{n}} = \mathcal{N}_{\mathbf{n}}$, which suggests that equality holds whenever some $n_i \geq 3$; determining $K_{\mathbf{n}}$ in general, and the corresponding orbit structure, is left open.

The square-zero profile is the unique case in which the tangent action is the full general linear group, and it is also the case in which the socle shadow determines the orbits outright.

Corollary 6.7. *If $n_i = 2$ for every i , then $R_{\mathbf{n}} = k \oplus J$, $J^2 = 0$, and $\text{Aut}_k(R_{\mathbf{n}}) \cong \text{GL}_r(k)$. Every proper ideal is a k -subspace of J , and two proper ideals are $\text{Aut}_k(R_{\mathbf{n}})$ -equivalent if and only if they have the same dimension. Thus, the full ideal lattice has exactly $r + 2$ automorphism orbits: one for each proper-ideal dimension $0, 1, \dots, r$, and one for the whole ring.*

Proof. Here $\mathcal{Q} = [r]$, $\mathcal{L} = \emptyset$, and every product of two radical elements is zero, so $J = \text{Soc}(R_{\mathbf{n}}) = \Sigma_2$ and $\Sigma_+ = 0$. Proposition 6.1 gives $\text{Aut}_k(R_{\mathbf{n}}) = \text{GL}(J)$. In (6.12) the blocks Ξ and Π' are empty, so $\Gamma_{\mathbf{n}} = \text{GL}(\Sigma_2) = \text{GL}(J)$, and the final assertion of Proposition 6.4 identifies the automorphism-orbit problem for the proper ideals with the $\text{GL}(J)$ -orbit problem on the subspace lattice of J . A linear automorphism preserves dimension and the general linear group is transitive on subspaces of a fixed dimension, so two proper ideals lie in the same orbit exactly when their dimensions agree. The unit ideal is fixed separately, which gives the stated count. $\square$

The smallest profile exhibiting all three symmetry layers is described next.

Example 6.8. *Let $\mathbf{n} = (2, 3, 3)$, so that $r_2 = 1$, $r_+ = 2$, and $r_3 = 2$. The tangent shadow is $L_{\mathbf{n}} \cong k^2 \rtimes (k^\times \times \text{Mon}_2(k))$, with $|L_{\mathbf{n}}| = 2q^2(q-1)^3$. The nonlinear kernel has order $|\mathcal{U}_{\mathbf{n}}| = q^{(0+1+1)+(3-1) \cdot 2} = q^6$, and therefore $|\text{Aut}_k(R_{\mathbf{n}})| = 2q^8(q-1)^3$. By Proposition 6.4, the two length-three terminal directions are scaled on the socle only by squares in $k^\times$, whereas the square-zero socle direction may acquire an arbitrary component in their span. On the depth-two ideal lattice, by contrast, the entire nonlinear kernel and the square-zero shear disappear, and only the monomial symmetry of the two length-three branches remains. Thus, the tangent, socle, and depth-two actions retain genuinely different portions of the same automorphism group, and the profile $(2, 3, 3)$ is the smallest one in which square-zero mixing, equal-length long-branch permutation, and depth-two rigidity occur simultaneously.*

7. Conclusions

We developed a unified structural theory for the multi-branch local algebras $R_{\mathbf{n}}$. The profile reconstructs the branch-length multiset, while the branch–subspace calculus records the additional data that the branch decomposition does not detect. Transition maps and coordinate saturation convert containment, intersection, and sum into explicit operations on canonical ideal data; they also expose the projective and Grassmannian fibers responsible for modularity and for the failure of distributivity. Corollary 3.10 turns the classification into a closed count of the ideal lattice for every r .

The same coordinates control quantitative and duality invariants. We obtained minimal generating sets, a compression formula, refined generator–rank enumerators, and a projective-torus parameterization of principal ideals. Annihilation then exhibits the complementary phenomenon: It erases the position and dimension of a terminal coupling while retaining its coordinate support. The resulting double-annihilator closure has an explicit defect and a distributive fixed-point lattice.

For cyclic quotients, this closure has a module-theoretic meaning: It is the torsionless reflection, universal among maps into torsionless modules, and its defect is exactly the kernel dimension of bidual evaluation. The torsionless cyclic modules are classified by the duality core and realized as full-support principal ideals. The universal cokernel dimension $r^2 - 1$ separates the chain case from every genuinely multi-branch case and provides a concrete categorical consequence of the ideal classification.

Finally, the automorphism group separates into a tangent image, an explicit nonlinear kernel, and a Galois action. Its tangent, socle, and depth-two shadows retain different portions of the symmetry, with

square-zero mixing at the first layer and monomial rigidity on long branches. The two-branch family is recovered as the rank-two specialization, whereas higher-dimensional terminal subspaces account for the new behavior when $r \geq 3$. Natural continuations include the determination of the full kernel of the action on the ideal lattice (Remark 6.6), the orbit problem on selected ideal strata, and the extension of the bidual analysis from cyclic quotients to higher-rank modules.

Author contributions

Sami H. Saif: Conceptualization, Methodology, Formal analysis, Investigation, Writing–review and editing, Writing–original draft; Alhanouf Ali Alhomaiddhi: Data curation, Validation, Visualization, Writing–review and editing. Both authors have read and approved the final version of the manuscript for publication.

Use of Generative-AI tools declaration

The authors declare that they have not used Artificial Intelligence (AI) tools in the creation of this article.

Acknowledgments

The authors would like to thank the Ongoing Research Funding program, (ORF-2026-871), King Saud University, Riyadh, Saudi Arabia.

Conflict of interest

The authors declare that there are no conflicts of interest.

References

1. B. R. McDonald, *Finite Rings with Identity*, New York: Marcel Dekker, 1974.
2. X. D. Hou, K. H. Leung, S. L. Ma, On the groups of units of finite commutative chain rings, *Finite Fields Appl.*, **9** (2003), 20–38. [https://doi.org/10.1016/S1071-5797\(02\)00003-5](https://doi.org/10.1016/S1071-5797(02)00003-5)
3. S. Alabiad, Y. Alkhamees, Recapturing the structure of group of units of any finite commutative chain ring, *Symmetry*, **13** (2021), 307. <https://doi.org/10.3390/sym13020307>
4. S. Alabiad, Y. Alkhamees, On automorphism groups of finite chain rings, *Symmetry*, **13** (2021), 681. <https://doi.org/10.3390/sym13040681>
5. W. F. Moore, Cohomology over fiber products of local rings, *J. Algebra*, **321** (2009), 758–773. <https://doi.org/10.1016/j.jalgebra.2008.10.015>
6. S. Nasseh, S. Sather-Wagstaff, Vanishing of Ext and Tor over fiber products, *Proc. Amer. Math. Soc.*, **145** (2017), 4661–4674. <https://doi.org/10.1090/proc/13633>
7. M. D’Anna, M. Fontana, An amalgamated duplication of a ring along an ideal: The basic properties, *J. Algebra Appl.*, **6** (2007), 443–459. <https://doi.org/10.1142/S0219498807002326>

8. H. Ananthnarayan, L. L. Avramov, W. F. Moore, Connected sums of Gorenstein local rings, *J. Reine Angew. Math.*, **667** (2012), 149–176. <https://doi.org/10.1515/CRELLE.2011.132>
9. G. H. Norton, A. Sălăgean, On the structure of linear and cyclic codes over a finite chain ring, *Appl. Algebra Engrg. Commun. Comput.*, **10** (2000), 489–506. <https://doi.org/10.1007/PL00012382>
10. H. Q. Dinh, S. R. López-Permouth, Cyclic and negacyclic codes over finite chain rings, *IEEE Trans. Inform. Theory*, **50** (2004), 1728–1744. <https://doi.org/10.1109/TIT.2004.831789>
11. A. Sălăgean, Repeated-root cyclic and negacyclic codes over a finite chain ring, *Discrete Appl. Math.*, **154** (2006), 413–419. <https://doi.org/10.1016/j.dam.2005.03.016>
12. J. A. Wood, Duality for modules over finite rings and applications to coding theory, *Amer. J. Math.*, **121** (1999), 555–575. <https://doi.org/10.1353/ajm.1999.0024>
13. B. Yildiz, S. Karadeniz, Cyclic codes over $\mathbb{F}_2 + u\mathbb{F}_2 + v\mathbb{F}_2 + uv\mathbb{F}_2$, *Des. Codes Cryptogr.*, **58** (2011), 221–234. <https://doi.org/10.1007/s10623-010-9399-3>
14. S. T. Dougherty, S. Karadeniz, B. Yildiz, Cyclic codes over R_k , *Des. Codes Cryptogr.*, **63** (2012), 113–126. <https://doi.org/10.1007/s10623-011-9539-4>
15. R. Sobhani, M. Molakarimi, Some results on cyclic codes over the ring $R_{2,m}$, *Turk. J. Math.*, **37** (2013), 1061–1074. <https://doi.org/10.3906/mat-1211-20>
16. P. K. Kewat, B. Ghosh, S. Pattanayak, Cyclic codes over the ring $\mathbb{Z}_p[u, v]/\langle u^2, v^2, uv - vu \rangle$, *Finite Fields Appl.*, **34** (2015), 161–175. <https://doi.org/10.1016/j.ffa.2015.01.005>
17. S. H. Saif, A. A. Alhomaiddhi, Cyclic codes over a split local ring of type (2, 2): Structure, Gray images, and distance analysis, *Mathematics*, **14** (2026), 2019. <https://doi.org/10.3390/math14112019>
18. T. Bag, H. Q. Dinh, K. Abdukhalikov, A. K. Upadhyay, W. Yamaka, Constacyclic codes over $\mathbb{F}_{q^2}[u]/\langle u^2 - w^2 \rangle$ and their application in quantum code construction, *J. Appl. Math. Comput.*, **68** (2022), 3821–3834. <https://doi.org/10.1007/s12190-021-01692-w>



AIMS Press

©2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)