COMPUTING EXTENSIONS OF FLAG-JUMP t-MODULES

DAWID E. KĘDZIERSKI, DOUGLAS KOELLMER, AND PIOTR KRASOŃ

Abstract. We study Anderson **t**-modules whose positive-degree coefficient matrices admit a flag-jump basis and extend the **t**-reduction algorithm to this setting. We prove that these **t**-modules are abelian, with motive rank equal to the sum of their jump degrees. If Φ has jump degrees $j_1, \dots, j_d$ and $\deg_\tau \Psi < \min_s j_s$, we compute the **t**-module structure on $\mathrm{Ext}^1_\tau(\Phi, \Psi)$, its subgroup $\mathrm{Ext}_{0,\tau}(\Phi, \Psi)$, and the associated exact sequence. The construction allows nonzero nilpotent constant terms and includes the strictly pure case. When Ψ is the Carlitz module, $N_\Phi = 0$, and all jump degrees are at least two, the construction yields a dual **t**-module. We give complexity bounds for the algorithms and implementations in SymPy and Mathematica.

1. Introduction

The arithmetic theory of Drinfeld modules and Anderson **t**-modules has acquired an increasingly algorithmic character in recent years. Explicit methods have been developed for computing endomorphism rings, isogenies, characteristic polynomials of Frobenius, and L-functions; see, for example, [2, 3, 5]. Algorithms for computing the motives and dual motives associated with Anderson **t**-modules have also been developed [15]. Alongside these theoretical developments, computer algebra implementations have made it possible to calculate invariants and examine examples that are difficult to handle by hand. Our work continues this computational approach by providing effective methods for homological constructions in the category of Anderson **t**-modules.

Extension groups play an important role in this category and in the construction of dual **t**-modules. Their description in terms of biderivations modulo inner biderivations provides a concrete starting point for computation. To determine the induced **t**-module structure explicitly, however, one needs finite-dimensional coordinates for extension classes and a procedure for computing the action in those coordinates. This requires a reduction algorithm together with proofs of termination and uniqueness of the resulting representatives.

Date: October 9, 2026.

2020 *Mathematics Subject Classification.* 11G09, 11R58, 18G50, 11Y16.

Key words and phrases. Anderson **t**-modules, extension groups, flag-jump bases, FJ presentations, algorithms and computational complexity.

Corresponding author: Piotr Krasoń.

Papanikolas and Ramachandran [16] used the biderivation description to establish a Weil–Barsotti formula for Drinfeld modules. Subsequently, explicit computations of extension structures for Drinfeld modules and certain special $\mathbf{t}$-modules were developed in [11]. The $\mathbf{t}$-reduction algorithm introduced in [6], with the correction in [7], extended these computations to strictly pure and certain triangular presentations. This method also enters the duality results of [12] and the homological constructions studied in [13].

The aim of this paper is to extend the range of presentations for which the $\mathbf{t}$-reduction algorithm can be carried out. In the strictly pure case, the highest coefficient matrix is invertible and supplies all the directions needed for reduction at a single degree. When this matrix is singular, the necessary directions may occur in different coefficient blocks. The problem is then to select suitable columns and arrange the reduction steps so that eliminating one term does not reintroduce terms already removed. We address this problem by introducing *flag-jump* (FJ) presentations, based on a compatibility condition between the coefficient blocks and their associated flag. The precise definition and its consequences are developed in Section 3.

Our main result shows that, if the source admits an FJ presentation and the degree of the target is smaller than every jump degree of the source, the $\mathbf{t}$-reduction algorithm determines a natural Anderson $\mathbf{t}$-module structure on the extension group. We provide algorithms for computing this structure, the subgroup represented by biderivations with zero constant term, and the terms and morphisms of the associated exact sequence.

The FJ condition includes strictly pure presentations and also applies to presentations with a singular highest coefficient matrix, including suitable block triangular examples. We prove that $\mathbf{t}$-modules admitting an FJ presentation are abelian and give an explicit basis of their associated motives (see Theorem 3.12). The jump degrees determine the motive rank of the source and, together with the dimension of the target, the dimension of the extension $\mathbf{t}$-module. We also describe the behavior of FJ presentations in exact sequences (see Proposition 3.8).

Under the additional hypotheses required for duality, our results also yield dual $\mathbf{t}$-modules, extending the setting treated in the strictly pure case. The FJ condition concerns the chosen presentation: an unsuccessful basis search does not rule out the existence of an isomorphic presentation satisfying this condition.

The computation proceeds in two stages. First, we search the positive coefficient blocks for a flag-jump basis and a compatible ordering of the column positions. This search uses rank constraints, linear independence tests, and precedence relations obtained from the selected columns. Once a compatible basis has been found, it determines the reduction steps used to compute the extension structure and the maps in its associated exact sequence. The basis search and full reduction are implemented in the accompanying software [10]. The exact-sequence construction in Section 6 is a linear postprocessing of the full action. Worked examples illustrate the

computation of Ext groups, their subgroups $\mathrm{Ext}_{0,\tau}$, and the inclusion maps from the input coefficient matrices.

The paper is organized as follows. Section 2 recalls the necessary background on $\mathbf{t}$-modules, extensions, and biderivations. Section 3 introduces flag-jump presentations, develops their main properties, and proves their abelianity. Section 4 presents the algorithms for finding a flag-jump basis and a compatible column ordering. Section 5 establishes the existence and uniqueness of normal representatives and constructs the induced $\mathbf{t}$-module structure on the extension group. Section 6 treats zero-constant extensions and gives algorithms for computing the associated exact sequence and its maps. Section 7 develops the construction of dual $\mathbf{t}$-modules and the corresponding Weil–Barsotti exact sequence. A running example illustrates these constructions throughout Sections 3–7. Section 8 analyzes the arithmetic complexity of the algorithms and discusses their implementation. In the appendix, we investigate exact sequences of $\mathbf{t}$-modules with middle term $\mathrm{Ext}^1_\tau(\Phi,\Psi)$ and quotient on which t acts as multiplication by θ. In particular, we establish conditions under which the kernel of the corresponding quotient map coincides with $\mathrm{Ext}_{0,\tau}(\Phi,\Psi)$.

2. Preliminaries

2.1. **Skew polynomials and presentations.** Let $q = p^s$, where p is prime and $s \geq 1$, and let L be a field of characteristic p. Fix an $\mathbb{F}_q$-algebra homomorphism

$$\iota : \mathbb{F}_q[t] \longrightarrow L, \qquad \theta := \iota(t).$$

For $a \in L$ and $r \geq 0$, put $a^{(r)} = a^{q^r}$. The same notation denotes entrywise Frobenius powers of matrices.

The skew polynomial ring $L\{\tau\}$ consists of polynomials with coefficients on the left and multiplication determined by $\tau a = a^q \tau$. For matrices of compatible sizes, this gives

$$(X\tau^r)(Y\tau^s) = XY^{(r)}\tau^{r+s}.$$

We write $\deg_\tau$ for the degree in τ, with $\deg_\tau 0 = -\infty$. For a matrix U over $L\{\tau\}$, its constant coefficient is denoted by U_0 or ∂U.

An Anderson t-module of dimension d over L is an $\mathbb{F}_q$-algebra homomorphism

$$\Phi : \mathbb{F}_q[t] \longrightarrow \mathrm{Mat}_d(L\{\tau\})$$

such that, writing $\Phi_t := \Phi(t)$,

$$\partial\Phi_t = \theta I_d + N_\Phi, \qquad N_\Phi^d = 0. \tag{2.1}$$

For a source Φ of dimension d and a target Ψ of dimension e, we write

$$\Phi_t = \theta I_d + N_\Phi + \sum_{r=1}^{n} A_r\tau^r, \quad \text{and} \quad \Psi_t = \theta I_e + N_\Psi + \sum_{r=1}^{m} B_r\tau^r.$$

We call n the degree of Φ_t and denote it as $\deg_\tau \Phi = n$.

A morphism $F : \Phi \longrightarrow \Psi$ is a matrix

$$F \in \mathrm{Mat}_{e\times d}(L\{\tau\})$$

satisfying

$$F\Phi_t = \Psi_t F.$$

Its differential is $\partial F = F_0$. For general background on Anderson t-modules, we refer to [1, 4, 8, 18].

A presentation is called *strictly pure* if its highest positive coefficient matrix is invertible.

2.2. **Biderivations and extensions.** Every $\mathbf{t}$-module Φ of dimension d induces an $\mathbb{F}_q[t]-$module structure on the space L^d. This is given by the following formula:

$$a * x := \Phi_a(x), \quad \text{for} \quad a \in \mathbb{F}_q[t], x \in L^d.$$

This $\mathbb{F}_q[t]-$module is called the Mordell-Weil group of Φ and is denoted as $\Phi(L^d)$. Analogously, every morphism of $\mathbf{t}$-modules induces a morphism of $\mathbb{F}_q[t]-$modules. In this way one obtains the inclusion functor from the category of $\mathbf{t}$-modules into the category of $\mathbb{F}_q[t]-$modules. This functor enables us to transfer the notion of exact sequences to the category of $\mathbf{t}$-modules in the following way. We say that the sequence of $\mathbf{t}$-modules

$$0 \longrightarrow \Phi \longrightarrow \Upsilon \longrightarrow \Psi \longrightarrow 0 \tag{2.2}$$

is exact if the corresponding sequence of $\mathbb{F}_q[t]-$modules

$$0 \longrightarrow \Phi(L^d) \longrightarrow \Upsilon(L^{e+d}) \longrightarrow \Psi(L^e) \longrightarrow 0$$

is exact. By $\mathrm{Ext}^1_\tau(\Psi, \Phi)$ we will denote the set of exact sequences of the form (2.2) modulo the isomorphic sequences, with group operation given by Baer addition.

A (Φ, Ψ)-biderivation is an $\mathbb{F}_q$-linear map $\delta : \mathbb{F}_q[t] \longrightarrow \mathrm{Mat}_{e\times d}(L\{\tau\})$ satisfying

$$\delta(ab) = \Psi(a)\delta(b) + \delta(a)\Phi(b) \quad \text{for} \quad a, b \in \mathbb{F}_q[t].$$

Following [6], we denote the space of biderivations by $\mathrm{Der}(\Phi, \Psi)$ and write $\delta_t := \delta(t)$. Evaluation at t gives an $\mathbb{F}_q$-linear isomorphism

$$\mathrm{Der}(\Phi, \Psi) \longrightarrow \mathrm{Mat}_{e\times d}(L\{\tau\}), \qquad \delta \longmapsto \delta_t.$$

Thus every matrix $V \in \mathrm{Mat}_{e\times d}(L\{\tau\})$ determines a unique biderivation with $\delta_t = V$.

The corresponding extension has presentation

$$X_t = \begin{bmatrix} \Phi_t & 0 \\ \delta_t & \Psi_t \end{bmatrix}.$$

Its constant coefficient differs from θI_{d+e} by a nilpotent matrix, so there is no additional restriction on $\partial\delta_t$.

For $U \in \mathrm{Mat}_{e\times d}(L\{\tau\})$, define the inner biderivation $\delta^{(U)}$ by its value at t: $\delta_t^{(U)} = U\Phi_t - \Psi_t U$. The image of the map

$$\delta^{(-)} : \mathrm{Mat}_{e\times d}(L\{\tau\}) \longrightarrow \mathrm{Der}(\Phi,\Psi), \qquad U \longmapsto \delta^{(U)},$$

is denoted by $\mathrm{Der}_{\mathrm{in}}(\Phi,\Psi)$. Under evaluation at t, this space is identified with the matrices $\delta_t^{(U)}$.

The biderivation description of extensions gives an isomorphism of $\mathbb{F}_q[t]$-modules

$$\mathrm{Ext}^1_\tau(\Phi,\Psi) \cong \frac{\mathrm{Der}(\Phi,\Psi)}{\mathrm{Der}_{\mathrm{in}}(\Phi,\Psi)} \cong \frac{\mathrm{Mat}_{e\times d}(L\{\tau\})}{\{\delta_t^{(U)} : U \in \mathrm{Mat}_{e\times d}(L\{\tau\})\}}.$$

In matrix notation, the action of t is

$$t[V] = [\Psi_t V] = [V\Phi_t],$$

because $V\Phi_t - \Psi_t V = \delta_t^{(V)} \in \mathrm{Der}_{\mathrm{in}}(\Phi,\Psi)$.

The map $U \mapsto \delta^{(U)}$ is additive and $\mathbb{F}_q$-linear, but need not be L-linear. Indeed, the positive-degree terms in $\Psi_t(\mu U)$ involve Frobenius powers of μ.

2.3. **Zero-constant biderivations.** Following [6], put

$$\mathrm{Der}_0(\Phi,\Psi) = \{\delta \in \mathrm{Der}(\Phi,\Psi) : \partial\delta_t = 0\}.$$

Its image in $\mathrm{Ext}^1_\tau(\Phi,\Psi)$ is

$$\mathrm{Ext}_{0,\tau}(\Phi,\Psi) = \frac{\mathrm{Der}_0(\Phi,\Psi)}{\mathrm{Der}_0(\Phi,\Psi) \cap \mathrm{Der}_{\mathrm{in}}(\Phi,\Psi)}. \tag{2.3}$$

A class belongs to this subgroup precisely when it admits a representative $V = \delta_t$ with $V_0 = 0$. It is stable under the action of t, since

$$(\Psi_t V)_0 = (\theta I_e + N_\Psi)V_0.$$

Then

$$\delta^{(U)} \in \mathrm{Der}_0(\Phi,\Psi) \quad \Longleftrightarrow \quad (\partial U)N_\Phi = N_\Psi(\partial U).$$

2.4. **Motives and abelianity.** Let Φ be a t-module of dimension d over L. Its associated t-motive is

$$M_\Phi = \mathrm{Hom}_{\mathbb{F}_q,L}(\mathbb{G}_a^d, \mathbb{G}_a) \cong L\{\tau\}^{1\times d},$$

where the homomorphisms are $\mathbb{F}_q$-linear morphisms of algebraic groups defined over L. The natural left $L\{\tau\}$-action, together with

$$t \cdot f = f \circ \Phi_t,$$

makes M_Φ a left $L[t]$-module. We call Φ *abelian* if M_Φ is finitely generated over $L[t]$. In this case M_Φ is free over $L[t]$, and its rank is called the rank of Φ.

3. Flag-jump presentations

3.1. **The running example.** Throughout this section, we illustrate the definitions with the pair

$$\Phi_t = \begin{bmatrix} \theta + \tau^2 & 0 \\ \tau^3 & \theta + \tau^2 + \tau^3 \end{bmatrix}, \qquad \Psi_t = \theta + \tau \tag{3.1}$$

over $L = \mathbb{F}_q(\theta)$, where θ is indeterminate. Here $\dim \Phi = 2$ and $\dim \Psi = 1$. The positive coefficient matrices of Φ_t are

$$A_1 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \qquad A_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \qquad A_3 = \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix}.$$

Thus the highest coefficient matrix has rank one, whereas

$$\Big[A_1 \mid A_2 \mid A_3\Big] = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix}$$

has rank two. Write $e_1 = (1,0)^T$ and $e_2 = (0,1)^T$.

3.2. **Column orders, flags, and jumps.** Let

$$A = \Big[A_1 \mid \cdots \mid A_n\Big], \qquad A_j = \Big[c_j^1, \ldots, c_j^d\Big], \qquad A_j \in \mathrm{Mat}_d(L). \tag{3.2}$$

For a permutation $\sigma \in S_d$, read the columns in the order

$$c_n^{\sigma(1)}, \ldots, c_n^{\sigma(d)}, c_{n-1}^{\sigma(1)}, \ldots, c_{n-1}^{\sigma(d)}, \ldots, c_1^{\sigma(1)}, \ldots, c_1^{\sigma(d)}. \tag{3.3}$$

The blocks are read in decreasing degree, with the same column order inside each block. Put $\rho = \sigma^{-1}$, so that $\rho(s)$ is the place occupied by position s in this order.

Running Example. In the example (3.1), the identity order gives

$$c_3^1,\ c_3^2,\ c_2^1,\ c_2^2,\ c_1^1,\ c_1^2.$$

For $\sigma = (1\,2)$, the order is

$$c_3^2,\ c_3^1,\ c_2^2,\ c_2^1,\ c_1^2,\ c_1^1.$$

In particular, $\rho(2) = 1$ and $\rho(1) = 2$.

Definition 3.1. Let $v_1, \ldots, v_{nd}$ be the columns in (3.3). The associated σ-flag is

$$W_0 \subseteq W_1 \subseteq W_2 \subseteq \cdots \subseteq W_{nd},$$

where $W_0 = \{0\}$ and $W_\ell = \mathrm{span}_L(v_1, \ldots, v_\ell)$ for $1 \le \ell \le nd$.

A column $v_\ell = c_j^s$ is a *σ-jump column* if $W_\ell \neq W_{\ell-1}$. We denote its position by (j, s) and write

$$J_\sigma(A) = \Big\{(j,s) \mid c_j^s \text{ is a } \sigma\text{-jump column}\Big\}.$$

Running Example. For $\sigma = (1\,2)$ in the example (3.1), the flag is

$$W_0 = \{0\}, \qquad W_1 = W_2 = W_3 = Le_2, \qquad W_4 = W_5 = W_6 = L^2.$$

The successive steps are

ℓ	column	v_ℓ	$\dim_L W_\ell$	jump
1	c_3^2	e_2	1	yes
2	c_3^1	e_2	1	no
3	c_2^2	e_2	1	no
4	c_2^1	e_1	2	yes
5	c_1^2	0	2	no
6	c_1^1	0	2	no

Hence $J_{(2,1)}(A) = \{(3,2),(2,1)\}$. In the identity order, the first and third columns give the jumps, so $J_{\mathrm{id}}(A) = \{(3,1),(2,1)\}$.

3.3. **Admissibility and the FJ type.**

Definition 3.2. The block matrix A is *σ-flag-jump admissible* if it has exactly d jump columns, with one in each original column position $s \in \{1,\dots,d\}$. Relabeling these columns by their original positions gives

$$J_\sigma(A) = \Big\{(j_1,1),\dots,(j_d,d)\Big\}.$$

The ordered basis

$$\mathbb{B}_\sigma(A) = \Big(c_{j_1}^1,\dots,c_{j_d}^d\Big)$$

is called the canonical *σ-flag-jump basis*. The permutation σ is a compatible order for this basis.

Running Example. In the identity order, the example (3.1) has both jumps in position one. It is therefore not identity-admissible. For $\sigma = (1\,2)$, the jump positions are $(3,2)$ and $(2,1)$, so the matrix is σ-flag-jump admissible. Its canonical basis and basis matrix are

$$\mathbb{B}_\sigma(A) = \Big(c_2^1, c_3^2\Big) = (e_1, e_2), \qquad C = \Big[c_2^1, c_3^2\Big] = I_2. \tag{3.4}$$

Although c_3^2 is selected first, it is the second vector of the canonical basis because its original column position is two.

Definition 3.3. For a presentation $\Phi_t = \partial\Phi\tau^0 + \sum_{r=1}^{n} A_r\tau^r$, put

$$C^+(\Phi) = [A_1 \mid \cdots \mid A_n].$$

We say that Φ is of *type* $\mathrm{FJ}_{\sigma,N}$ if $C^+(\Phi)$ is σ-flag-jump admissible and

$$N = \min_{1\le s\le d} j_s. \tag{3.5}$$

The integers j_s are its *jump degrees*. We also put $J = \sum_{s=1}^{d} j_s$. A presentation is **FJ** if it is of this type for some σ and N.

Remark 3.4. The FJ condition depends on the chosen presentation and is not preserved under isomorphisms of t-modules. In this respect it is analogous to strict purity: both conditions concern the coefficient matrices in a chosen coordinate system.

For example, suppose that q is odd and consider

$$\Phi_t = \theta I_2 + \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \tau + \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \tau^2.$$

This is an FJ presentation: in the order $(1, 2)$, the jump columns occur in positions one and two, at degrees two and one, respectively. However, for

$$P = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \in \mathrm{GL}_2(\mathbb{F}_q),$$

the isomorphic presentation is

$$P^{-1}\Phi_t P = \theta I_2 + \frac{1}{2}\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \tau + \frac{1}{2}\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \tau^2.$$

For either order of the two positions, the first column read in the degree-two block gives the first jump. The column in the same position in the degree-one block gives the second jump, since it lies outside the span of the degree-two block. Thus both jumps occur in the same position, and no order is admissible.

We therefore use the term *FJ presentation*. The assertion that a t-module admits an FJ presentation is, of course, invariant under isomorphism.

Running Example. For the example (3.1),

$$j_1 = 2, \qquad j_2 = 3, \qquad N = 2, \qquad J = 5.$$

Thus Φ is of type $\mathrm{FJ}_{(1\,2),2}$. Moreover, $\deg_\tau \Psi_t = 1 < N$, so the pair satisfies the hypotheses of Theorem 5.3 as well as Theorem 5.6.

Example 3.5. Every strictly pure $\mathbf{t}$-module is FJ. Indeed, let Φ be strictly pure with leading coefficient matrix A_n. Since A_n is invertible, its d columns are linearly independent in any order. Consequently, all of them are selected before any lower block is examined, yielding $j_s = n$ for every s. Thus, every strictly pure presentation is of type $\mathrm{FJ}_{\sigma,n}$ for every permutation σ.

3.4. **The leading rows.** Recall that, for a matrix C, we denote its entry in position (a, b) by $C_{a\times b}$. With this convention, the symbol $*$ indicates that the corresponding index ranges over all possible values. Thus, $C_{a\times *}$ denotes the a-th row of C, while $C_{*\times b}$ denotes its b-th column.

For a selection of one column $b_s = c^s_{j_s}$ in each position, write

$$C = [b_1, \ldots, b_d], \qquad \lambda_s = \left(C^{-1}\right)_{s\times *},$$

assuming that C is invertible. For every $v \in L^d$, we have

$$v = \sum_{s=1}^{d} (\lambda_s \cdot v)\, b_s.$$

Thus $\lambda_s \cdot v$ is the coefficient of b_s in the expansion of v in the basis $(b_1, \ldots, b_d)$.

Proposition 3.6. *The selected columns form the canonical jump basis for $\sigma = \rho^{-1}$ if and only if, for every $s \in \{1, 2, \ldots, d\}$, the following equalities hold:*

$$\big(\lambda_s A_{j_s}\big)_s = 1, \quad \lambda_s A_r = 0 \quad \textit{for} \quad r > j_s, \quad \big(\lambda_s A_{j_s}\big)_l = 0 \quad \textit{for} \quad \rho(l) < \rho(s).$$

In this case, $\lambda_s \Phi_t$ has degree j_s, with leading coefficient one in position s and zero in every position preceding s in the σ-order.

Proof. Suppose first that the selected columns are the jump columns. Every column read before b_s lies in the span of the earlier jump columns, so its coefficient along b_s is zero. This applies to every column of A_r with $r > j_s$ and to the columns preceding s within A_{j_s}. Finally, $\lambda_s \cdot b_s = 1$.

Conversely, suppose that the equalities hold. It suffices to show that the selected columns form a basis compatible with the σ-flag. Expand a column c_r^l in the basis $b_1, \ldots, b_d$. A nonzero coefficient along b_s implies $j_s \geq r$; if $j_s = r$, it also implies $\rho(s) \leq \rho(l)$. Thus, every basis column appearing in the expansion precedes or coincides with c_r^l in the σ-order. If c_r^l is unselected, all basis columns appearing in its expansion strictly precede it. Each selected column is linearly independent of the preceding basis columns and therefore creates a jump. □

Under the assumptions of Proposition 3.6, the column $c_{j_s}^s$ is called the *pivot column* associated with $\lambda_s \Phi_t$, and the entry $(\lambda_s A_{j_s})_s = 1$ is called its *pivot*. If $B = (b_1, b_2, \ldots, b_d)$ is a jump basis, then the matrix

$$C = [c_{j_1}^1, \ldots, c_{j_d}^d]$$

will be called the *jump matrix*.

Running Example. In the example (3.1), $\sigma = (1\,2)$, $C = I_2$, so $\lambda_1 = (1, 0)$ and $\lambda_2 = (0, 1)$. The required identities are

$$\lambda_1 A_2 = (1, 0), \quad \lambda_1 A_3 = (0, 0), \quad \lambda_2 A_2 = (0, 1), \quad \lambda_2 A_3 = (1, 1).$$

The leading rows are therefore

$$\lambda_1 \Phi_t = (\theta + \tau^2, 0), \qquad \lambda_2 \Phi_t = (\tau^3, \theta + \tau^2 + \tau^3).$$

3.5. **Suffix ranks and the required number of columns.** The number of jumps contributed by each block can be computed before choosing a column order. Define

$$S_j = \operatorname{span}_L \Big\{ c_r^s : r \geq j,\ 1 \leq s \leq d \Big\}, \qquad s_j = \dim_L S_j, \tag{3.6}$$

with $S_{n+1} = \{0\}$ and $s_{n+1} = 0$. Thus s_j is the rank of the suffix matrix $[A_j \mid \cdots \mid A_n]$. Put

$$r_j = s_j - s_{j+1} \quad \text{for} \quad 1 \leq j \leq n. \tag{3.7}$$

For a full-row-rank block matrix, also put

$$k = \max \{ j : s_j = d \}. \tag{3.8}$$

Proposition 3.7. *Every canonical jump basis contains exactly r_j columns from A_j. Consequently, for an FJ presentation,*

$$N = k, \qquad J = \sum_{j=k}^{n} j r_j.$$

These quantities are independent of the compatible order.

Proof. After all blocks of degree at least j have been read, the selected columns form a basis of S_j. The block A_j therefore contributes $s_j - s_{j+1}$ new columns. Since $s_k = d > s_{k+1}$, block k contributes at least one jump, and no lower block contributes a jump. Hence the smallest jump degree is k. Summing the degrees of the selected columns gives the formula for J. □

Running Example. In the example (3.1),

$$S_3 = Le_2, \qquad S_2 = S_1 = L^2,$$

so the suffix ranks and required numbers are

j	s_j	r_j
1	2	0
2	2	1
3	1	1

Thus $k = 2$, $N = 2$, and $J = 2 \cdot 1 + 3 \cdot 1 = 5$. A jump basis must contain one column from A_3 and one from A_2. For example, (c_2^1, c_2^2) is an independent selection with one column in each original position, but it cannot be a jump basis because it uses no column from A_3.

Since $S_2 = L^2$, reading A_1 creates no further jumps. The basis search can therefore use the trimmed matrix

$$A_{\text{trim}} = [A_2 \mid A_3] = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 \end{bmatrix}.$$

The selected columns retain their original block indices two and three.

3.6. **Properties of FJ presentations.** We first describe the behaviour of FJ presentations in short exact sequences.

Proposition 3.8. *Let*

$$\delta: \quad 0 \longrightarrow \Psi \longrightarrow \Upsilon \longrightarrow \Phi \longrightarrow 0$$

be a short exact sequence defined by a biderivation δ.

(i) *Suppose that Φ and Ψ are of types $\mathrm{FJ}_{\sigma_1,N_1}$ and $\mathrm{FJ}_{\sigma_2,N_2}$, respectively. If $\deg_\tau \delta(t) < M := \min\{N_1, N_2\}$, then Υ is of type $\mathrm{FJ}_{\sigma,M}$, where σ is any shuffle of the two compatible orders.*

(ii) *Conversely, suppose that Υ is of type $\mathrm{FJ}_{\sigma,M}$ and $\deg_\tau \delta < M$. Then both Φ and Ψ admit FJ presentations with respect to the orders induced by σ.*

Proof. For every $j \geq M$, the coefficient of τ^j in Υ_t is block diagonal. Consequently, in this degree range, every successive column span is the direct sum of the corresponding spans for the two diagonal presentations. A column creates a jump precisely when its column in the corresponding diagonal block creates a jump.

Thus any shuffle preserving the two internal orders selects exactly the embedded jump columns of Φ_t and Ψ_t. These already form a basis of the whole coordinate space, so columns of degree less than M create no further jumps.

For the converse, all jumps occur in the block diagonal degree range. Restricting the column order to either diagonal block therefore selects its greedy jump columns. Since the whole presentation has exactly one jump in every position, the same holds in each diagonal block. □

Corollary 3.9. *Let*

$$0 = \Upsilon_0 \subset \Upsilon_1 \subset \cdots \subset \Upsilon_v = \Upsilon$$

be a filtration by t-submodules. Suppose that its factors $\Phi_a = \Upsilon_a/\Upsilon_{a-1}$ have FJ presentations with minimum jump degrees N_a, and put

$$M = \min_{1 \leq a \leq v} N_a.$$

If each extension $0 \longrightarrow \Upsilon_{a-1} \longrightarrow \Upsilon_a \longrightarrow \Phi_a \longrightarrow 0$ is represented by a biderivation δ_a satisfying $\deg_\tau \delta_a < M$, then the resulting presentation of Υ is FJ.

Proof. Apply Proposition 3.8 successively. □

Suppose that Υ admits a filtration whose factors are strictly pure t-modules $\Phi_1, \ldots, \Phi_v$. In coordinates adapted to this filtration, its presentation has the block lower triangular form

$$\Upsilon = \begin{bmatrix} \Phi_v & 0 & \cdots & 0 \\ \delta_{v-1\times v} & \Phi_{v-1} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \delta_{1\times v} & \delta_{1\times v-1} & \cdots & \Phi_1 \end{bmatrix}. \tag{3.9}$$

Here $\delta_{a\times b}$ denotes the off-diagonal block in the block row corresponding to Φ_a and the block column corresponding to Φ_b, for $a < b$.

Proposition 3.10. *Let Υ be a block-triangular $\mathbf{t}$-module of the form (3.9), with $\deg_\tau \Phi_a = n_a$. If $\deg_\tau \delta_{a\times b} < \min\{n_a, n_b\}$ for all $a < b$, then Υ_t is FJ with respect to every ordering of its column positions, and*

$$N = \min_a n_a, \qquad J = \sum_a d_a n_a.$$

Proof. Write $\Upsilon_t = \sum_{j=0}^{M} A_j \tau^j$ and read the coefficient matrices A_j in decreasing order of j. Before A_j is examined, the selected columns span precisely the coordinate blocks whose diagonal degrees are greater than j. At

degree j, each diagonal block with $n_a = j$ has an invertible coefficient matrix. The degree bounds make all off-diagonal blocks in the same block row or block column vanish at degree j and above. Its columns therefore give d_a independent new directions, in any order. Columns belonging to blocks with larger diagonal degrees lie in the span already obtained: their entries in coordinate blocks with diagonal degrees of at most j vanish. Columns belonging to blocks with smaller diagonal degrees vanish entirely at degree j.

Descending induction proves that the selection succeeds for every ordering of the column positions. Each block a contributes exactly d_a selected columns of degree n_a, so

$$N = \min_a n_a, \qquad J = \sum_a d_a n_a.$$

□

These statements connect FJ presentations with triangular t-modules. Recall that a triangular t-module admits a τ-composition series whose factors are Drinfeld modules; see [13, Section 4]. In coordinates adapted to this series, its matrix is triangular, with Drinfeld modules $\psi_1, \dots, \psi_v$ on the diagonal. Since a Drinfeld module is strictly pure and $\deg_\tau \psi_a(t) = \operatorname{rk} \psi_a$, the preceding proposition gives the following consequence.

Corollary 3.11. *Let $\Upsilon = (\underline{\psi}, \underline{\delta})$ be a triangular $\mathbf{t}$-module. If*

$$\deg_\tau \delta_{a\times b} < \min\{\operatorname{rk} \psi_a, \operatorname{rk} \psi_b\} \quad \text{for all} \quad a < b,$$

then Υ_t is FJ for every column order, with

$$j_a = \operatorname{rk} \psi_a, \qquad N = \min_a \operatorname{rk} \psi_a, \qquad J = \sum_a \operatorname{rk} \psi_a.$$

3.7. **Abelianity and motive rank.** We now show that the jump degrees determine the rank of the associated motive. We use the notation M_Φ introduced in the preliminaries.

Theorem 3.12. *Suppose that Φ_t is an FJ presentation with respect to the σ-order, with jump degrees $j_1, \dots, j_d$. Then M_Φ is free over $L[t]$, with basis*

$$\left\{ e_s^T \tau^r : 1 \le s \le d,\ 0 \le r < j_s \right\}, \tag{3.10}$$

where e_s is the s-th standard column vector. In particular, Φ is abelian and

$$\operatorname{rk} \Phi = J = \sum_{s=1}^{d} j_s.$$

Proof. Let C be the jump matrix, let λ_s be row s of C^{-1}, and put $N = \min_s j_s$. Define

$$R_\Phi = \left\{ (f_1, \dots, f_d) \in L\{\tau\}^{1\times d} : \deg_\tau f_s < j_s \text{ for every } s \right\}.$$

The rows in (3.10) form an L-basis of R_Φ.

Every row $g = \sum_k g_k \tau^k \in L\{\tau\}^{1\times d}$ has a unique expansion

$$g = \sum_{s,k} \mu_{s,k} \tau^k \lambda_s, \qquad (\mu_{1,k}, \dots, \mu_{d,k}) = g_k C^{(k)}.$$

Proposition 3.6 implies that $\mu_{s,k}\tau^k\lambda_s\Phi_t$, when nonzero, has leading degree $k+j_s$, coefficient $\mu_{s,k}$ in position s, and zero coefficients in all earlier positions of the σ-order.

For $g \neq 0$, put

$$M = \max_{\mu_{s,k}\neq 0} (k + j_s).$$

Among all pairs (s,k) with $\mu_{s,k} \neq 0$ and $k + j_s = M$, choose the one whose pivot position s occurs first in the fixed σ-order. Explicitly, if the positions are ordered as $\sigma(1), \dots, \sigma(d)$, choose $s = \sigma(r)$ with r as small as possible among these pairs. The chosen summand has coefficient $\mu_{s,k} \neq 0$ at degree M in position s. Every other summand of degree M has its pivot later in the σ-order, and hence has zero coefficient in position s; summands of smaller degree contribute nothing at degree M. Thus this leading term cannot be canceled, and $\deg_\tau(g\Phi_t) = M$.

Since $M = k + j_s \geq j_s$, this term also shows that $g\Phi_t \notin R_\Phi$. Moreover, $\deg_\tau g = \max_{\mu_{s,k}\neq 0} k$ and $j_s \geq N$. Consequently, for every $g \neq 0$,

$$g\Phi_t \neq 0, \qquad g\Phi_t \notin R_\Phi, \qquad \deg_\tau(g\Phi_t) \geq \deg_\tau g + N.$$

Using the same pivot columns, successively eliminate every coefficient in position s whose degree is at least j_s. Process degrees in decreasing order and, at each degree, positions in the σ-order. Proposition 3.6 ensures termination, giving a division

$$f = g\Phi_t + R, \qquad R \in R_\Phi.$$

Uniqueness follows from $g\Phi_t \notin R_\Phi$ for $g \neq 0$.

If $g \neq 0$, the leading pivot term of $g\Phi_t$ cannot be canceled by R, since its degree is at least the bound j_s at that position. Hence

$$\deg_\tau g \leq \deg_\tau f - N.$$

Since $N > 0$, repeated division therefore terminates and expresses f as

$$f = \sum_{\ell=0}^{v} R_\ell \Phi_t^\ell, \qquad R_\ell \in R_\Phi.$$

Since $R_\ell\Phi_t^\ell = t^\ell \cdot R_\ell$, the rows in (3.10) generate M_Φ over $L[t]$.

Finally, suppose that $\sum_{\ell=0}^{v} R_\ell \Phi_t^\ell = 0$. Uniqueness of division gives $R_0 = 0$. Right multiplication by Φ_t is injective, so removing one such factor and repeating gives $R_\ell = 0$ for every ℓ. The L-independence of the rows in (3.10) now proves their $L[t]$-independence. Their number is J. □

The proof uses only Frobenius powers of coefficients and therefore does not require L to be perfect. When L is perfect, we obtain the following corollary.

Corollary 3.13. *Suppose that Φ_t is an FJ presentation defined over a perfect field L. Then Φ is* **t***-finite.*

Proof. Since L is perfect, every abelian **t**-module over L is **t**-finite by [14, Theorem A]. □

4. Finding a jump basis

We now give Algorithms 1–6 for deciding whether $A = [A_1 \mid \cdots \mid A_n]$ admits an admissible column order. The search first selects one column from each original position, subject to independence and the required number of columns from each block. It then tests whether a single permutation, used within every block, makes these selected columns the jump columns. The running example will illustrate both tests and the need to undo an unsuccessful selection.

4.1. **Preprocessing by suffix ranks.** Algorithm 1 computes the suffix ranks from Subsection 3.5 and stops at the first full-rank suffix encountered while j decreases. This gives

$$k = \max\{j : s_j = d\} \quad \text{and} \quad \text{Required}[j] = r_j = s_j - s_{j+1}.$$

The entries below k remain zero, since the lower blocks contribute no jumps.

To retain the original block indices, we record the trimmed input as $A_{\text{trim}} = (k, [A_k \mid \cdots \mid A_n])$. All subsequent algorithms use these original block indices. The additional record $\text{Flag} = ((k, s_k), \ldots, (n, s_n))$ stores the suffix ranks. The successive column spans remain the flag defined in Section 3.

Algorithm 1 Preprocessing by suffix ranks

Input: $A = [A_1 \mid \cdots \mid A_n]$, with $A_j \in \text{Mat}_d(L)$.
Output: $(\text{Found}, A_{\text{trim}}, \text{Flag}, \text{Required})$.

procedure PREPROCESS(A)
  if $\operatorname{rank} A < d$ **then**
    return $(\text{False}, \varnothing, \varnothing, \varnothing)$
  end if
  $s_{n+1} \leftarrow 0$, Flag $\leftarrow ()$, Required $\leftarrow (0, \ldots, 0)$ of length n
  for $j = n, n-1, \ldots, 1$ **do**
    $s_j \leftarrow \operatorname{rank}[A_j \mid \cdots \mid A_n]$
    $\text{Required}[j] \leftarrow s_j - s_{j+1}$
    Append (j, s_j) to Flag
    if $s_j = d$ **then**
      $A_{\text{trim}} \leftarrow (j, [A_j \mid \cdots \mid A_n])$
      Reverse Flag to increasing block order
      return $(\text{True}, A_{\text{trim}}, \text{Flag}, \text{Required})$
    end if
  end for
end procedure

Running Example. For the example (3.1), Algorithm 1 computes $s_3 = 1$ and $s_2 = 2$, and then stops. Its output is

$$A_{\text{trim}} = \left(2, \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 \end{bmatrix}\right), \quad \text{Flag} \ = ((2,2),(3,1)), \quad \text{Required} = (0,1,1).$$

Every possible jump basis must therefore use one column from A_2 and one from A_3.

4.2. **Candidate lists.** For each original position i, Algorithm 2 lists its nonzero columns in decreasing block order. A candidate is recorded as (i, j, c_j^i), so its position and block index remain available throughout the search.

Algorithm 2 Candidate columns with original indices

Input: $A_{\text{trim}} = (k, [A_k \mid \cdots \mid A_n])$.
Output: Lists $L_1, \ldots, L_d$.

```
procedure Lists(A_trim)
    for i = 1, ..., d do
        L_i ← ()
        for j = n, n − 1, ..., k do
            if c_j^i ≠ 0 then
                Append (i, j, c_j^i) to L_i
            end if
        end for
    end for
    return (L_1, ..., L_d)
end procedure
```

Algorithm 3 orders the lists by increasing length, so the search considers positions with fewer choices first. Equal lengths retain their original order. This only determines the order in which selections are tried; the compatible column order σ is computed separately by Algorithm 4.

Algorithm 3 Ordering the candidate lists

Input: Lists $L_1, \ldots, L_d$.
Output: Ordered family OL.

```
procedure ShortFirst(L)
    Stably sort L_1, ..., L_d by increasing length
    return the resulting family OL
end procedure
```

Running Example. In the example (3.1), the lists are

$$L_1 = ((1,3,e_2),(1,2,e_1)), \quad L_2 \ = ((2,3,e_2),(2,2,e_2)).$$

Both have length two, so OL $= (L_1, L_2)$. The search therefore chooses a candidate for position one before choosing a candidate for position two, even though the eventual compatible order is $\sigma = (1\,2)$.

4.3. **Testing a compatible order.** In the algorithms, an independent selection is recorded as

$$\mathbb{B} = \big((1, j_1, b_1), \dots, (d, j_d, b_d)\big), \qquad b_s = c^s_{j_s}.$$

Its vectors form a basis $(b_1, \dots, b_d)$, with one column from each original position. Suppose that exactly r_j columns have been selected from each block A_j. Put

$$C = [b_1, \dots, b_d], \qquad B_j = \{s : j_s = j\}, \qquad D_j = C^{-1} A_j.$$

Thus $(D_j)_{l\times i}$ is the coefficient of b_l in the expansion of c^i_j. For each unselected column c^i_j, the condition

$$l \in B_j, \qquad i \notin B_j, \qquad (D_j)_{l\times i} \neq 0 \tag{4.1}$$

requires position l to precede position i within that block. Record this requirement as a directed edge $l \to i$ on the vertex set $\{1, \dots, d\}$. The same permutation must satisfy all these requirements, across all blocks.

Proposition 4.1. *The selected columns form the canonical jump basis for a permutation σ if and only if every edge $l \to i$ places l before i in σ. Consequently, a compatible order exists if and only if the directed graph is acyclic.*

Proof. The selected columns from blocks of index at least j are independent and their number is $\sum_{h=j}^{n} r_h = s_j$. They therefore form a basis of S_j. In particular, the expansion of any column of A_j uses only selected columns whose block indices are at least j. Those from higher blocks have already been read.

If the selection is the canonical jump basis, an unselected column lies in the span of the selected columns read earlier. Each of its nonzero coordinates in a selected column of the same block must therefore occur at an earlier position. This provides all the requirements in (4.1).

Conversely, suppose that σ satisfies every requirement. Each unselected column is then a combination of selected columns read earlier. It creates no jump. Each selected column is independent of the preceding selected columns and hence creates a jump. These are precisely the prescribed columns. The compatible orders are thus the topological orders of the graph, which exist exactly when it is acyclic, see [9]. □

Algorithm 4 constructs such an order by repeatedly removing a vertex with no incoming edge from the remaining vertices. If no such vertex exists, subsequent incoming edges produce a directed cycle. Choosing the smallest available vertex makes the returned order lexicographically smallest for the selected basis.

Algorithm 4 A compatible order for a selected basis

Input: A_{trim}; an independent selection $\mathbb{B}$ with the required block counts.
Output: (Found, σ).

procedure ORDER($A_{\mathrm{trim}}, \mathbb{B}$)
  Form $C = [b_1, \ldots, b_d]$ in original position order
  $C^{-1} \leftarrow \mathrm{Inverse}(C)$, $\mathcal{R} \leftarrow \varnothing$
  for $j = k, \ldots, n$ **do**
    $B_j \leftarrow \{s : j_s = j\}$, $D_j \leftarrow C^{-1} A_j$
    for $i \in \{1, \ldots, d\} \setminus B_j$ **do**
      for $l \in B_j$ **do**
        if $(D_j)_{l \times i} \neq 0$ **then**
          $\mathcal{R} \leftarrow \mathcal{R} \cup \{(l, i)\}$
        end if
      end for
    end for
  end for
  $S \leftarrow \{1, \ldots, d\}$, $\sigma \leftarrow ()$
  while $S \neq \varnothing$ **do**
    $T \leftarrow \{i \in S : \nexists l \in S \text{ with } (l, i) \in \mathcal{R}\}$
    if $T = \varnothing$ **then**
      return $(\mathrm{False}, \varnothing)$
    end if
    $i_0 \leftarrow \min T$
    Append i_0 to σ;
    $S \leftarrow S \setminus \{i_0\}$
  end while
  return (True, σ)
end procedure

Running Example. For the selection $(c_2^1, c_3^2) = (e_1, e_2)$ in the example (3.1), $C = I_2$ and

$$B_2 = \{1\}, \quad B_3 = \{2\}, \qquad D_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad D_3 = \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix}.$$

Block A_2 contributes no edge because $(D_2)_{1\times 2} = 0$. Block A_3 contributes $2 \to 1$, since $(D_3)_{2\times 1} = 1$. Thus position two must be read first: its selected column $c_3^2 = e_2$ must precede the equal, unselected column $c_3^1 = e_2$. Algorithm 4 removes vertex two and then vertex one, returning $\sigma = (1\,2)$.

4.4. **Selection and backtracking.** Algorithm 5 selects one candidate from each list by depth-first search. Its current path contains independent columns, and Used$[j]$ counts the columns already selected from block j. A candidate is accepted only if it preserves independence and does not exceed Required$[j]$. Independence can be tested by checking that adjoining its vector increases the rank of the path matrix by one.

At a path of length d, every required count is attained: the counts are bounded above by the requirements and both sum to d. The selection is then tested by Algorithm 4. If the test or a later selection fails, the last choice is undone and the next candidate is tried. The helper DFS below is local to SEARCH; it shares the inputs, the path, and the count array. At a completed path, it sorts a copy into original position order, keeping the path available for backtracking.

Algorithm 5 Depth-first selection with required block counts

Input: $A_{\mathrm{trim}}, \mathrm{OL}, \mathrm{Required}$.
Output: $(\mathrm{Found}, \mathbb{B}, \sigma)$.

procedure SEARCH($A_{\mathrm{trim}}, \mathrm{OL}, \mathrm{Required}$)
  $d \leftarrow \#\mathrm{OL}$, Path $\leftarrow ()$
  Used $\leftarrow (0, \ldots, 0)$ of length #Required
  return DFS(1)
end procedure
procedure DFS(ℓ)
  if $\ell > d$ **then**
    $\mathbb{B} \leftarrow$ a copy of Path, sorted by original position
    $(\mathrm{Found}, \sigma) \leftarrow$ ORDER($A_{\mathrm{trim}}, \mathbb{B}$)
    if Found **then**
      return $(\mathrm{True}, \mathbb{B}, \sigma)$
    end if
    return $(\mathrm{False}, \varnothing, \varnothing)$
  end if
  for $(i, j, c) \in \mathrm{OL}[\ell]$, in list order **do**
    if Used$[j] <$ Required$[j]$ and c is independent of the path **then**
      Append (i, j, c) to Path
      Used$[j] \leftarrow$ Used$[j] + 1$
      $R \leftarrow \mathrm{DFS}(\ell + 1)$
      if R is successful **then** ▷ verify by means of Algorithm 4.
        return R
      end if
      Used$[j] \leftarrow$ Used$[j] - 1$
      Remove the last entry of Path
    end if
  end for
  return $(\mathrm{False}, \varnothing, \varnothing)$
end procedure

Running Example. For the example (3.1), the first choice from L_1 is $(1, 3, e_2)$. The candidate $(2, 3, e_2)$ is then rejected because the required count for block three has already been reached. The other candidate, $(2, 2, e_2)$, is rejected because it is dependent on the selected vector e_2. The search therefore removes $(1, 3, e_2)$ and resets Used[3] to zero.

The next choice from L_1 is $(1, 2, e_1)$. Now $(2, 3, e_2)$ is accepted: the two vectors are independent and use one column from each required block. Algorithm 4 returns $\sigma = (1\,2)$, so the search succeeds with

$$\mathbb{B} = \big((1, 2, e_1), (2, 3, e_2)\big), \qquad \sigma = (1\,2). \tag{4.2}$$

Here the triples record the original indices; their vector components give the basis (e_1, e_2).

Algorithm 6 combines the preceding steps.

Algorithm 6 Finding a jump basis

Input: $A = [A_1 \mid \cdots \mid A_n]$, with $A_j \in \mathrm{Mat}_d(L)$.
Output: $(\mathrm{Found}, \mathbb{B}, \sigma)$.

```
1: procedure JumpBasis(A)
2:     (Found, A_trim, Flag, Required) ← Preprocess(A)
3:     if not Found then
4:         return (False, ∅, ∅)
5:     end if
6:     L ← Lists(A_trim)
7:     OL ← ShortFirst(L)
8:     return Search(A_trim, OL, Required)
9: end procedure
```

Theorem 4.2. *Algorithm 6 returns success if and only if A admits an admissible column order. On success, its selected columns are the canonical jump basis for the returned permutation and satisfy Proposition 3.6.*

Proof. If $\operatorname{rank} A < d$, there cannot be a jump basis with one column from each position. Otherwise, Proposition 3.7 shows that every such basis lies in the retained blocks and satisfies the required block counts. At a successful leaf, Algorithm 5 has selected one column from each position, the columns are independent, and all required counts are attained. Proposition 4.1 makes the returned order compatible.

Conversely, fix a canonical jump basis for an admissible order. Each selected column occurs in its candidate list. Every partial selection along these lists is independent and stays within the required block counts. Thus none are rejected by Algorithm 5. The finite search reaches this selection unless it has already found another valid one. Its graph has a compatible order, so Algorithm 4 succeeds. The row identities follow from Proposition 3.6. □

Running Example. Applying Algorithm 6 to $C^+(\Phi)$ in the example (3.1) returns the basis and permutation in (4.2). Reading the retained blocks in this order gives

$$(c_3^2, c_3^1, c_2^2, c_2^1) = (e_2, e_2, e_2, e_1),$$

whose successive ranks are $1, 1, 1, 2$. Its jump columns are exactly c_3^2 and c_2^1. Their original block indices give $j_1 = 2$, $j_2 = 3$, hence $N = 2$ and $J = 5$.

The next section uses this basis to reduce biderivations and compute the t-action for the same pair (Φ, Ψ).

5. Extensions and the reduction algorithm

We apply the t-reduction algorithm introduced by Głoch, Kędzierski and Krasoń in [6, Section 3]; see also [7]. We briefly recall its steps. The general correctness argument will not be repeated: our task is to construct a basis satisfying the requirements of Step 2.

Let Φ and Ψ have dimensions d and e, respectively. For $1 \leq i \leq e$ and $1 \leq s \leq d$, let $E_{i,s}$ denote the $e \times d$ matrix with 1 in row i, column s, and 0 elsewhere.

5.1. The steps of the t-reduction algorithm.

Step 1: Identify extension classes with matrices modulo inner biderivations, using

$$\mathrm{Ext}^1_\tau(\Phi, \Psi) \cong \frac{\mathrm{Der}(\Phi, \Psi)}{\mathrm{Der}_{\mathrm{in}}(\Phi, \Psi)}$$

and evaluation at t.

Step 2: Choose an L-basis $(U_{i\times s})$ of $\mathrm{Mat}_{e\times d}(L)$ satisfying the reduction requirements A–D:

A. the inner biderivations $\delta^{(c\tau^k U_{i\times j})}$ form a system of generators of the $\mathbb{F}_q-$linear space $\mathrm{Der}_{\mathrm{in}}(\Phi, \Psi)$;

B. the matrix $\delta_t^{(c\tau^k U_{i\times j})}$ has at the $i \times j$-entry a polynomial whose coefficient at the largest power of τ is equal to c.

C. for every biderivation $\delta \in \mathrm{Der}(\Phi, \Psi)$ there exists a sequence of inner biderivations $\delta^{(c_l\tau^{k_l} U_{i\times j})}$ where $c_l \in L$ and $k_l \in \mathbb{Z}_{\geq 0}$ for $l = 1, 2, \ldots, r$ such that the reduced biderivation:

$$\mathrm{red}(\delta) = \delta - \sum_{l=1}^{r} \delta^{(c_l\tau^{k_l} \cdot U_{i\times j})} \in \mathrm{Mat}_{e\times d}(L\{\tau\})_{<\mathcal{N}},$$

where $\mathcal{N} = \left[\deg_\tau \delta_t^{(c\tau^0 U_{i\times j})}\right]_{i,j}$. The reduced form $\mathrm{red}(\delta)$ will be called *normal*.

D. different reduced biderivations in $\mathrm{Mat}_{e\times d}(L\{\tau\})_{<\mathcal{N}}$ represent different elements in $\mathrm{Ext}^1_\tau(\Phi, \Psi)$.

Points A., B., C. and D. imply that there exists an isomorphism of $\mathbb{F}_q-$linear spaces:

$$\mathrm{Ext}^1_\tau(\Phi, \Psi) \cong \mathrm{Mat}_{e\times d}(L\{\tau\})_{<\mathcal{N}}.$$

Step 3: For the resulting bounds $n_{i\times s}$, choose the coordinate monomials

$$E_{i\times s}\tau^r, \qquad 0 \leq r < n_{i\times s}.$$

Step 4: Compute the action on each coordinate with a scalar parameter c:
$$t * [cE_{i\times s}\tau^r] = [\Psi_t(cE_{i\times s}\tau^r)].$$
Step 5: Reduce these matrices using the inner biderivations from Step 2.
Step 6: Express the reduced matrices in the coordinates from Step 3.
Step 7: Replace each additive coordinate polynomial $\sum_h a_h c^{q^h}$ by its linearization $\sum_h a_h \tau^h$.
Final step: Assemble these coordinate vectors as the columns of the matrix Π_t.

By [6, Proposition 3.1], the existence of a basis satisfying Step 2 ensures that this procedure gives an Anderson $\mathbf{t}$-module structure on the extension group.

For the remainder of this section, we fix the following notation and assumptions. Fix an FJ presentation Φ_t of dimension d, with compatible order $\sigma = \rho^{-1}$, jump matrix C, and jump degrees $j_1, \dots, j_d$. Suppose that
$$\dim \Psi = e \quad \text{and} \quad m = \deg_\tau \Psi_t < N := \min_s j_s.$$
Set
$$U_{i\times s} = E_{i\times s}C^{-1}, \qquad \mathcal{E} = \Big\{V \in \mathrm{Mat}_{e\times d}(L\{\tau\}) : \deg_\tau V_{i\times s} < j_s\Big\}.$$
Thus matrices $U_{i\times s}$ form an L-basis of $\mathrm{Mat}_{e\times d}(L)$.

5.2. **Verification of Step 2 for FJ presentations.** We start with the following lemma.

Lemma 5.1. *For $\mu \in L^\times$ and $k \geq 0$, the inner biderivation associated with $H = \mu\tau^k U_{i\times s}$ satisfies*
$$\delta_t^{(H)} = \mu E_{i\times s}\tau^{k+j_s} + \sum_{\rho(l)>\rho(s)} a_l E_{i\times l}\tau^{k+j_s} + R, \qquad \deg_\tau R < k + j_s,$$
for suitable $a_l \in L$. In particular, its degree is $k + j_s$.

Proof. By Proposition 3.6, $U_{i\times s}A_r = 0$ for $r > j_s$, while $U_{i\times s}A_{j_s}$ has pivot entry 1 in position (i, s) and zero entries in earlier positions of the σ-order. The leading coefficient of $H\Phi_t$ is $\mu(U_{i\times s}A_{j_s})^{(k)}$, so Frobenius twisting preserves the pivot and these zeros. Since
$$\deg_\tau(\Psi_t H) \leq \deg_\tau \Psi_t + k < j_s + k,$$
the target contribution does not affect this coefficient. □

Let C be the jump matrix. Every matrix $H = \sum_k H_k\tau^k$ has a unique expansion
$$H = \sum_{i\times s,k} \mu_{i\times s,k}\tau^k U_{i\times s}, \quad \text{where} \quad [\mu_{i\times s,k}]_{i,s} = H_k C^{(k)}. \tag{5.1}$$

Lemma 5.2. *The map $H \mapsto \delta_t^{(H)}$ is injective, and $\mathcal{E} \cap \mathrm{Der}_{\mathrm{in}}(\Phi, \Psi) = \{0\}$.*

Proof. Let $H \neq 0$ and set $R = \max\{k+j_s : \mu_{i\times s,k} \neq 0\}$. Among the positions attaining this maximum, choose the first position s in σ, and choose a row i with $\mu_{i\times s,R-j_s} \neq 0$. By Lemma 5.1, the coefficient of $\delta_t^{(H)}$ in position (i,s) at degree R is $\mu_{i\times s,R-j_s}$. Terms with smaller leading degrees cannot contribute there. Terms with leading degree R whose positions follow s in σ have zero entry there, and the other rows do not affect row i at that degree. Thus this coefficient is nonzero. It follows that $\delta_t^{(H)} \neq 0$. Since $R \geq j_s$, it also follows that $\delta_t^{(H)} \notin \mathcal{E}$. □

Theorem 5.3. *The basis $(U_{i\times s})_{i,s}$ satisfies the requirements of Step 2 with bounds*

$$n_{i\times s} = j_s.$$

Consequently, the t-reduction algorithm endows $\mathrm{Ext}^1_\tau(\Phi, \Psi)$ with an Anderson t-module structure of dimension

$$eJ = e\sum_{s=1}^{d} j_s.$$

Its coordinate representatives are the matrices in $\mathcal{E}$, with coordinate basis

$$\boldsymbol{E} = \Big(E_{i\times s}\tau^r : 1 \leq i \leq e,\ 1 \leq s \leq d,\ 0 \leq r < j_s\Big). \tag{5.2}$$

Proof. We verify A–D of Step 2.

For A, every $H = \sum_k H_k\tau^k$ has a unique expansion of the form (5.1). Then additivity gives the corresponding expansion of its inner biderivation.

Conditions B and D follow directly from Lemmas 5.1 and 5.2, respectively.

For C, eliminate coefficients of degree $R \geq j_s$ by subtracting the inner biderivation determined by $\mu\tau^{R-j_s}U_{i\times s}$. Process degrees in decreasing order and, at each degree, positions in the σ-order. The lemma shows that each correction affects only later positions at the same degree and lower degrees. Thus reduction terminates in $\mathcal{E}$.

The conclusion now follows from [6, Proposition 3.1]. □

5.3. **Vanishing of morphisms for t-modules with FJ presentations.** The following theorem is an immediate consequence of Lemma 5.2.

Theorem 5.4. *Let Φ be a $\mathbf{t}$-module admitting an FJ presentation with jump degrees $j_1, \ldots, j_d$, and let Ψ be a $\mathbf{t}$-module satisfying $\deg_\tau \Psi_t < \min_s j_s$. Then*

$$\mathrm{Hom}_\tau(\Phi, \Psi) = 0.$$

Proof. Let $H : \Phi \to \Psi$ be a morphism of t-modules. Then $\delta^{(H)} = H\Phi_t - \Psi_t H = 0$, and hence $H = 0$ by Lemma 5.2. □

Running Example. For the example (3.1) we obtain

$$\mathrm{Hom}_\tau\left(\begin{bmatrix}\theta+\tau^2 & 0\\ \tau^3 & \theta+\tau^2+\tau^3\end{bmatrix}, \theta+\tau\right) = 0.$$

5.4. **Precomputation for the t-reduction algorithm.** The following pseudocode implements these steps using the FJ data computed in Section 4. We construct a unique representative of each extension class with entrywise degree bounds determined by the j_s. Its eJ coefficients provide coordinates in which we compute the induced action of t.

Algorithm 7 obtains the reduction data from the two presentations. The original block indices in the selected basis give the degrees j_s.

Algorithm 7 Preparing the reduction data

Input: $\mathbf{t}$-module presentations Φ_t, Ψ_t over L.
Output: Reduction data $\mathcal{D}$, or a failure with its reason.

procedure PREPARE(Φ_t, Ψ_t)
  Extract $A_1, \dots, A_n$ from Φ_t
  $(\text{Found}, \mathbb{B}, \sigma) \leftarrow$ JUMPBASIS($[A_1 \mid \cdots \mid A_n]$)
  if not Found **then**
    return failure: no admissible order for this presentation
  end if
  Form C and $j \leftarrow (j_s)$ from $\mathbb{B}$, in original position order
  Compute C^{-1}
  if $\deg_\tau \Psi_t \geq \min_s j_s$ **then**
    return failure: the strict degree gap does not hold
  end if
  Form $U_{i\times s} = E_{i\times s}C^{-1}$ and the ordered basis $\mathbf{E}$
  $U \leftarrow \left(U_{i\times s}\right)_{i,s}$
  return $\mathcal{D} = (\Phi_t, \Psi_t, C, C^{-1}, \sigma, j, U, \mathbf{E})$
end procedure

Running Example. Recall the example (3.1):

$$\Phi_t = \begin{bmatrix} \theta + \tau^2 & 0 \\ \tau^3 & \theta + \tau^2 + \tau^3 \end{bmatrix}, \qquad \Psi_t = \theta + \tau.$$

Here $C = I_2$, $\sigma = (1\,2)$, and $(j_1, j_2) = (2, 3)$. Thus $m = 1 < N = 2$, and

$$U_{1\times 1} = (1, 0), \qquad U_{1\times 2} = (0, 1).$$

The normal representatives are precisely

$$F = (a_0 + a_1\tau,\ b_0 + b_1\tau + b_2\tau^2).$$

For $U_{1\times 1}$ the leading coefficient is in position one at degree two. For $U_{1\times 2}$ it is in position two at degree three; a coefficient in position one may occur at the same degree, since position one follows position two in σ.

5.5. **Termination of reduction.** Let $\rho = \sigma^{-1}$. For an input matrix V, set $D = \max\{0, \deg_\tau V\}$, taking $D = 0$ when $V = 0$. The coefficients to be removed are indexed by

$$\mathcal{P}_D = \Big\{(R, s, i) : j_s \leq R \leq D, \quad 1 \leq s \leq d,\ 1 \leq i \leq e\Big\}.$$

We order this set by decreasing R, then increasing $\rho(s)$, and then increasing i.

Suppose the current coefficient indexed by (R, s, i) is μ. Subtracting $\delta_t^{(\mu\tau^{R-j_s}U_{i\times s})}$ sets this coefficient to zero. Lemma 5.1 shows that the subtraction creates no term of degree greater than R. At degree R, it changes only row i, and its other changes occur in positions later than s in σ. Every coefficient it can change in $\mathcal{P}_D$ therefore comes at or after the chosen one.

Algorithm 8 visits the coefficients in this order. When a coefficient is zero, its correction H is zero.

Algorithm 8 Reduction modulo inner biderivations

Input: $V \in \mathrm{Mat}_{e\times d}(L\{\tau\})$ and valid data $\mathcal{D}$.
Output: (F, W), with $F \in \mathcal{E}$ and $V - F = \delta_t^{(W)}$.

procedure REDUCE($V, \mathcal{D}$)
  $F \leftarrow V, \qquad W \leftarrow 0$
  $D \leftarrow \max\{0, \deg_\tau V\}$, with $D = 0$ if $V = 0$
  for all $(R, s, i) \in \mathcal{P}_D$, in the prescribed order **do**
    $\mu \leftarrow (\mathrm{coeff}_{\tau^R} F)_{i\times s}$
    $k \leftarrow R - j_s$,
    $H \leftarrow \mu\tau^k U_{i\times s}$
    $F \leftarrow F - (H\Phi_t - \Psi_t H)$
    $W \leftarrow W + H$
  end for
  return (F, W)
end procedure

Proposition 5.5. *Algorithm 8 terminates, using at most*

$$K(D) = e\sum_{s=1}^{d} \max\{0, D - j_s + 1\}$$

nonzero corrections. It returns $F \in \mathcal{E}$ and W satisfying $V - F = \delta_t^{(W)}$.

Proof. The input has no term of degree greater than D, and Lemma 5.1 shows that every subtraction preserves this bound. After an element of $\mathcal{P}_D$ has been processed, its coefficient is zero and all subsequent steps leave it zero. Each element is visited once, and there are exactly $K(D)$ elements. At the end, every coefficient of degree $R \geq j_s$ in position s is zero, so $F \in \mathcal{E}$.

The identity $V - F = \delta_t^{(W)}$ holds initially and is preserved by every update, since $H \mapsto \delta_t^{(H)}$ is additive. □

The total degree need not decrease at every step. What ensures termination is that the process cannot return to a coefficient already removed.

Running Example. The example (3.1) makes this distinction explicit. Take $V = (0, \mu\tau^3)$, with $\mu \neq 0$, and put $\Delta = \theta^q - \theta$. The ordered list

is

$$\mathcal{P}_3 = \Big((3,2,1),(3,1,1),(2,1,1)\Big).$$

The successive corrections are

$$H_1 = \mu U_{1,2}, \qquad H_2 = -\mu\tau U_{1,1}, \qquad H_3 = -\mu^q U_{1,1}.$$

Writing $F_0 = V$ and $F_a = F_{a-1} - \delta_t^{(H_a)}$ gives

$$F_1 = \Big(-\mu\tau^3,\ \mu^q\tau - \mu\tau^2\Big),$$

$$F_2 = \Big(-\mu^q\tau^2 + \mu\Delta\tau,\ \mu^q\tau - \mu\tau^2\Big),$$

$$F_3 = \Big((\mu\Delta - \mu^{q^2})\tau,\ \mu^q\tau - \mu\tau^2\Big).$$

The first step leaves the degree equal to three, but moves the reducible coefficient from position two to position one. The second step removes it at degree three and leaves a reducible coefficient at degree two in position one. The third removes that coefficient. The degree-two term in position two is allowed because $j_2 = 3$. Hence $F_3 \in \mathcal{E}$, and

$$V - F_3 = \delta_t^{(W)}, \qquad W = \mu U_{1,2} - \mu\tau U_{1,1} - \mu^q U_{1,1}.$$

5.6. **The action of t on extension classes.** The results presented in this subsection are specializations of Proposition 3.1 in [6]. We nevertheless include them here for the sake of clarity and to facilitate the exposition.

The action on extension classes is $t[V] = [\Psi_t V]$. In normal coordinates, it is therefore

$$F \longmapsto \mathrm{red}(\Psi_t F). \tag{5.3}$$

We compute its matrix by applying this formula to each normal coordinate separately.

We write the basis $\mathbf{E}$ defined in (5.2) as $\mathbf{E} = (g_1, \dots, g_M)$. For each g_β, introduce a formal scalar c and reduce $\Psi_t(cg_\beta)$. At each step, the correction uses addition, multiplication by fixed elements of L, and Frobenius powers of the current coefficient. These operations preserve additive polynomials in c. Since the elimination order is finite, we obtain

$$\mathrm{red}(\Psi_t(cg_\beta)) = \sum_{\alpha=1}^{M} f_{\alpha\beta}(c) g_\alpha, \qquad f_{\alpha\beta}(c) = \sum_{h\geq 0} a_{\alpha\beta,h} c^{q^h}, \tag{5.4}$$

with finite sums and $a_{\alpha\beta,h} \in L$. The corresponding entry of the action matrix is

$$(\Pi_t)_{\alpha\beta} = \sum_{h\geq 0} a_{\alpha\beta,h}\tau^h.$$

Thus the reduction of $\Psi_t(cg_\beta)$ determines column β of Π_t.

Algorithm 9 Computing the action on extension classes

Input: Valid reduction data $\mathcal{D}$, with ordered basis $\mathbf{E}$ of the form (5.2).
Output: $\Pi_t \in \mathrm{Mat}_M(L\{\tau\})$.

procedure ExtMatrix($\mathcal{D}$)
  $\Pi_t \leftarrow 0_{M\times M}$
  for $\beta = 1, \dots, M$ **do**
    $(F, W) \leftarrow \textsc{Reduce}(\Psi_t(cg_\beta), \mathcal{D})$, with formal c
    Express $F = \sum_{\alpha=1}^{M} f_{\alpha\beta}(c) g_\alpha$
    for $\alpha = 1, \dots, M$ **do**
      Write $f_{\alpha\beta}(c) = \sum_h a_{\alpha\beta,h} c^{q^h}$
      $(\Pi_t)_{\alpha\beta} \leftarrow \sum_h a_{\alpha\beta,h} \tau^h$
    end for
  end for
  return Π_t
end procedure

Theorem 5.6. *Let Φ be of type $\mathrm{FJ}_{\sigma,N}$ and let $\deg_\tau \Psi_t < N$. Normal representatives give $\mathrm{Ext}^1_\tau(\Phi, \Psi)$ a natural $\mathbf{t}$-module structure of dimension $M = eJ$. Algorithm 9 computes its action matrix Π_t. In the coordinate order fixed above,*

$$\partial \Pi_t = (\theta I_e + N_\Psi) \otimes I_J. \tag{5.5}$$

Proof. Theorem 5.3 identifies extension classes with the additive group of $\mathcal{E}$. The elimination formulas show that (5.3) is an additive polynomial map defined over L. It agrees with the action on extension classes, so Π_t gives an $\mathbb{F}_q[t]$-action on $\mathbb{G}_a^M$.

To compute its differential, take $g = E_{i,s}\tau^r$, with $r < j_s$. The part linear in c in $\Psi_t(cg)$ is

$$(\theta I_e + N_\Psi) c E_{i,s} \tau^r.$$

It is already normal: multiplication by this constant matrix changes only the row index, preserving s and r. Every other term contains c^{q^h} with $h \geq 1$. The elimination formulas preserve or increase this Frobenius index, so such terms do not contribute to the part linear in c. This proves (5.5). Since $N_\Psi \otimes I_J$ is nilpotent, the nilpotence condition (2.1) in the definition of an Anderson t-module holds. □

Running Example. For the example (3.1), the normal basis is

$$g_1 = \begin{bmatrix} 1 & 0 \end{bmatrix}, \quad g_2 = \begin{bmatrix} \tau & 0 \end{bmatrix}, \quad g_3 = \begin{bmatrix} 0 & 1 \end{bmatrix}, \quad g_4 = \begin{bmatrix} 0 & \tau \end{bmatrix}, \quad g_5 = \begin{bmatrix} 0 & \tau^2 \end{bmatrix}.$$

Three of the images are already normal:

$$\begin{aligned}\operatorname{red}\left(\Psi_t(cg_1)\right) &= \theta c g_1 + c^q g_2,\\ \operatorname{red}\left(\Psi_t(cg_3)\right) &= \theta c g_3 + c^q g_4,\\ \operatorname{red}\left(\Psi_t(cg_4)\right) &= \theta c g_4 + c^q g_5.\end{aligned}$$

For the second coordinate,

$$\Psi_t(cg_2) = \begin{bmatrix}\theta c\tau + c^q\tau^2 & 0\end{bmatrix}.$$

Subtracting the inner biderivation determined by $H = c^q U_{1,1}$ removes $c^q\tau^2$ and gives

$$\operatorname{red}\left(\Psi_t(cg_2)\right) = \left(\theta c + c^{q^2}\right) g_2.$$

For the fifth coordinate,

$$\Psi_t(cg_5) = \begin{bmatrix}0 & \theta c\tau^2 + c^q\tau^3\end{bmatrix}.$$

The degree-two term is normal. Applying the three-step reduction above to the degree-three term, with $\mu = c^q$, gives

$$\operatorname{red}\left(\Psi_t(cg_5)\right) = \left(\Delta c^q - c^{q^3}\right) g_2 + c^{q^2} g_4 + (\theta c - c^q) g_5,$$

where $\Delta = \theta^q - \theta$. Consequently, the action matrix is

$$\Pi_t = \begin{bmatrix}\theta & 0 & 0 & 0 & 0\\ \tau & \theta+\tau^2 & 0 & 0 & \Delta\tau - \tau^3\\ 0 & 0 & \theta & 0 & 0\\ 0 & 0 & \tau & \theta & \tau^2\\ 0 & 0 & 0 & \tau & \theta - \tau\end{bmatrix}, \qquad \Delta = \theta^q - \theta. \tag{5.6}$$

Its differential is θI_5, in agreement with (5.5), since $N_\Psi = 0$.

6. $\operatorname{Ext}_{0,\tau}(\Phi, \Psi)$ and the exact sequence

Throughout this section, we set the following notation and assumptions. Take an FJ presentation Φ_t of dimension d, with compatible order $\sigma = \rho^{-1}$, jump matrix C, and jump degrees $j_1, \ldots, j_d$. Suppose that

$$\dim \Psi = e \quad \text{and} \quad m = \deg_\tau \Psi_t < N := \min_s j_s.$$

Set $U_{i\times s} = E_{i\times s} C^{-1}$ and the normal coordinates $\boldsymbol{E}$ of the form (5.2), constructed in Section 5. Set

$$J = \sum_{s=1}^{d} j_s \qquad \text{and} \qquad M = eJ.$$

6.1. **The inclusion of $\mathrm{Ext}_{0,\tau}(\Phi,\Psi)$.** Recall that

$$\mathrm{Der}_0(\Phi,\Psi)=\Big\{\delta\in\mathrm{Der}(\Phi,\Psi):\delta(t)_0=0\Big\}$$

and

$$\mathrm{Ext}_{0,\tau}(\Phi,\Psi)=\frac{\mathrm{Der}_0(\Phi,\Psi)}{\mathrm{Der}_0(\Phi,\Psi)\cap\mathrm{Der}_{\mathrm{in}}(\Phi,\Psi)}.$$

The inclusion of biderivations induces an injective map

$$\iota:\mathrm{Ext}_{0,\tau}(\Phi,\Psi)\longrightarrow\mathrm{Ext}^1_\tau(\Phi,\Psi).$$

Indeed, a biderivation in Der_0 represents zero in the full extension module precisely when it belongs to $\mathrm{Der}_0\cap\mathrm{Der}_{\mathrm{in}}$. Its image consists of the extension classes admitting a representative V with $V_0=0$. Since

$$\big(\Psi_tV\big)_0=(\theta I_e+N_\Psi)V_0,$$

the action of t preserves this subgroup, and ι is $\mathbb{F}_q[t]$-linear.

For $H=\sum_k H_k\tau^k$, the constant coefficient of the inner biderivation is

$$\Big(\delta_t^{(H)}\Big)_0=H_0N_\Phi-N_\Psi H_0.$$

Thus $\delta^{(H)}\in\mathrm{Der}_0(\Phi,\Psi)$ precisely when $H_0N_\Phi=N_\Psi H_0$.

6.2. **The exact sequence.** Recall that Section 5 identifies the full extension group with the normal-coordinate space $\mathcal{E}$ and determines the matrix Π_t representing the $\mathbf{t}$-action.

To determine which extension classes belong to $\mathrm{Ext}_{0,\tau}$, we must determine which constant coefficients can be removed by an inner biderivation. Put

$$W=\mathrm{Mat}_{e\times d}(L),\qquad \mathcal{T}:W\longrightarrow W,\qquad \mathcal{T}(H_0)=H_0N_\Phi-N_\Psi H_0.$$

For $\delta_t^{(H)}=H\Phi_t-\Psi_tH$, one has

$$\Big(\delta_t^{(H)}\Big)_0=\mathcal{T}(H_0).$$

Thus equivalent representatives have constant coefficients that differ by an element of $\mathrm{im}\,\mathcal{T}$. The constant-coordinate information that remains is

$$Q=W/\mathrm{im}\,\mathcal{T},\qquad h=\dim_L Q=ed-\mathrm{rank}\,\mathcal{T}.$$

Thus h is the number of independent constant-coordinate obstructions to membership in $\mathrm{Ext}_{0,\tau}$. It is computed by the rank of the linear map $\mathcal{T}$.

Left multiplication by N_Ψ preserves $\mathrm{im}\,\mathcal{T}$, because

$$N_\Psi\mathcal{T}(H_0)=\mathcal{T}(N_\Psi H_0).$$

It therefore induces a nilpotent endomorphism $\overline{N}_\Psi$ of Q:

$$\overline{N}_\Psi(V+\mathrm{im}\,\mathcal{T})=N_\Psi V+\mathrm{im}\,\mathcal{T}.$$

Theorem 6.1. *Under the assumptions above, there is an exact sequence of* $\mathbf{t}$*-modules*

$$0 \longrightarrow \operatorname{Ext}_{0,\tau}(\Phi,\Psi) \xrightarrow{\iota} \operatorname{Ext}^1_{\tau}(\Phi,\Psi) \xrightarrow{p} (Q, \theta\,\mathrm{id}_Q + \overline{N}_\Psi) \longrightarrow 0,$$

where $p([V]) = V_0 + \operatorname{im}\mathcal{T}$*. The dimensions of the three terms are* $M-h$*,* M*, and* h*, respectively.*

Proof. The formula for $\left(\delta_t^{(H)}\right)_0$ shows that p is well defined. It is surjective because every matrix in W is a constant representative of a biderivation.

A class $[V]$ lies in $\ker p$ precisely when $V_0 = \mathcal{T}(H_0)$ for some constant matrix H_0. Subtracting $\delta_t^{(H_0)}$ then gives a zero-constant representative of the same class. Conversely, every zero-constant representative maps to zero. Hence $\ker p = \operatorname{Ext}_{0,\tau}(\Phi,\Psi)$. Moreover, $(\Psi_t V)_0 = (\theta I_e + N_\Psi)V_0$, so p respects the asserted $\mathbf{t}$-action on Q.

In the normal coordinates of Section 5, p is an L-linear surjection of additive algebraic groups. Its kernel is the linear space

$$\mathcal{E}_0 = \{F \in \mathcal{E} : F_0 \in \operatorname{im}\mathcal{T}\},$$

of dimension $M-h$. This space is preserved under the additive polynomial $\mathbf{t}$-action on $\mathcal{E}$, since $(\Psi_t F)_0 \in \operatorname{im}\mathcal{T}$ whenever $F_0 \in \operatorname{im}\mathcal{T}$. The differential of the restricted action is the restriction of $\partial\Pi_t$ to this invariant space. By Theorem 5.6, $\partial\Pi_t - \theta I_M$ is nilpotent, so its restriction is nilpotent as well. Thus $\mathcal{E}_0$ is an Anderson $\mathbf{t}$-module. The nilpotence of $\overline{N}_\Psi$ also makes Q a $\mathbf{t}$-module with the displayed action. Both maps in the exact sequence are morphisms of $\mathbf{t}$-modules. □

6.3. **Coordinates for the inclusion and the restricted action.** Theorem 6.1 characterizes $\operatorname{Ext}_{0,\tau}$ inside the full normal-coordinate space. We now turn this characterization into explicit coordinates.

Choosing a basis. Let $\mathcal{E}_+$ be the space of normal representatives with zero constant coefficient. Its coordinate monomials are

$$\boldsymbol{E}_+ = \Big\{E_{i,s}\tau^k, \qquad 1 \le i \le e, \quad 1 \le s \le d, \quad 1 \le k < j_s\Big\}. \tag{6.1}$$

There are $\ell = M - ed$ such monomials. Fix an order for them. All constant matrices are normal, so $\mathcal{E} = W \oplus \mathcal{E}_+$. Choose a basis $V_1, \dots, V_r$ of $\operatorname{im}\mathcal{T}$, where $r = ed - h$.

Proposition 6.2. *The space representing* $\operatorname{Ext}_{0,\tau}$ *in full normal coordinates is*

$$\mathcal{E}_0 = \operatorname{im}\mathcal{T} \oplus \mathcal{E}_+.$$

A coordinate basis is

$$\mathbf{E}_0 = (V_1, \dots, V_r) \cup \boldsymbol{E}_+, \tag{6.2}$$

with the positive-degree monomials in the fixed order above. Its cardinality is $r + \ell = M - h$.

Proof. By Theorem 6.1, the only condition on a full normal representative is $F_0 \in \operatorname{im}\mathcal{T}$. Its positive-degree coordinates are unrestricted. The direct sum and basis follow. □

We now specify the basis $V_1, \dots, V_r$ of $\operatorname{im} T$ used in Proposition 6.2. Recall that

$$W = \operatorname{Mat}_{e\times d}(L), \qquad T(H_0) = H_0 N_\Phi - N_\Psi H_0, \qquad r = ed - h.$$

For $H_0 = \left[h_{i\times s}\right]_{i,s} \in W$, let $\operatorname{vec}(H_0) \in L^{ed}$ be obtained by stacking its columns:

$$\operatorname{vec}(H_0) = \left[h_{1,1}, \dots, h_{e,1}, h_{1,2}, \dots, h_{e,d}\right]^T.$$

Thus $\operatorname{vec}(E_{i\times s})$ is the standard coordinate vector with index $i + e(s-1)$. In these coordinates, T is represented by

$$D = N_\Phi^T \otimes I_e - I_d \otimes N_\Psi, \qquad \operatorname{vec}\left(T(H_0)\right) = D \operatorname{vec}(H_0).$$

Use Gaussian elimination to find indices $k_1, \dots, k_r$ of original columns of D that form a basis of its column space, and set

$$B = \left[D_{*\times k_1} \mid \cdots \mid D_{*\times k_r}\right] \in \operatorname{Mat}_{ed\times r}(L).$$

For each a, write $k_a = i_a + e(s_a - 1)$ and define

$$H_a = E_{i_a \times s_a}, \qquad V_a = T(H_a) = H_a N_\Phi - N_\Psi H_a.$$

Then $B_{*\times a} = \operatorname{vec}(V_a)$, and $V_1, \dots, V_r$ form the chosen basis of $\operatorname{im} T$. The matrices H_a are therefore selected standard basis matrices of W; the matrices V_a are their images under T.

Each V_a is normal, and $[V_a] \in \operatorname{Ext}_{0,\tau}(\Phi, \Psi)$. Indeed, regarding H_a as a constant matrix over $L\{\tau\}$, put

$$\delta_t^{(H_a)} = H_a \Phi_t - \Psi_t H_a.$$

Its constant coefficient is $T(H_a) = V_a$. Hence

$$[V_a] = \left[V_a - \delta_t^{(H_a)}\right] \qquad \text{and} \qquad \left(V_a - \delta_t^{(H_a)}\right)_0 = 0.$$

Thus the class admits a representative with zero constant coefficient, even though its normal representative may have a nonzero constant coefficient.

Coordinates on the constant quotient. Choose a matrix $Z \in \operatorname{Mat}_{ed\times h}(L)$ whose columns complete those of B to a basis of L^{ed}. Then

$$A = \left[B \ Z\right] \in \operatorname{GL}_{ed}(L).$$

The matrices obtained by reshaping the columns of Z into $e \times d$ matrices represent a basis of $Q = W/\operatorname{im} T$. Write the inverse in row blocks as

$$A^{-1} = \begin{bmatrix} U \\ P_0 \end{bmatrix}, \qquad U \in \operatorname{Mat}_{r\times ed}(L), \qquad P_0 \in \operatorname{Mat}_{h\times ed}(L).$$

Since $A^{-1}A = I_{ed}$, block multiplication gives

$$I_{ed} = \begin{bmatrix} I_r & 0 \\ 0 & I_h \end{bmatrix} = \begin{bmatrix} U \\ P_0 \end{bmatrix} \left[B \ Z\right] = \begin{bmatrix} UB & UZ \\ P_0 B & P_0 Z \end{bmatrix}.$$

In particular, $\ker P_0 = \operatorname{im} B$.

Matrices of the inclusion and quotient maps. Let $\ell = M - ed$. Order the full normal coordinates with the constant matrix units first, in column-vectorization order, followed by the fixed basis of $\mathcal{E}_+$. Express Π_t in this same order. A full normal representative then has coordinate vector $\binom{c}{x}$, where $c \in L^{ed}$ records its constant coefficient and $x \in L^\ell$ records its positive-degree coefficients.

By Proposition 6.2, a class belongs to $\mathrm{Ext}_{0,\tau}(\Phi, \Psi)$ precisely when $c \in \operatorname{im} B$. Writing $c = Bu$, its coordinates in the chosen basis of $\mathcal{E}_0$ are $\binom{u}{x}$. Therefore the inclusion into the full normal-coordinate space has matrix

$$K = \begin{bmatrix} B & 0 \\ 0 & I_\ell \end{bmatrix}, \qquad K\begin{pmatrix} u \\ x \end{pmatrix} = \begin{pmatrix} Bu \\ x \end{pmatrix}.$$

The quotient map sends $\binom{c}{x}$ to $P_0 c$, so its matrix is

$$P = \begin{bmatrix} P_0 & 0 \end{bmatrix}.$$

Finally, set

$$R = \begin{bmatrix} U & 0 \\ 0 & I_\ell \end{bmatrix}.$$

For vectors in the image of K, this matrix recovers the $\mathcal{E}_0$ coordinates:

$$R\begin{pmatrix} Bu \\ x \end{pmatrix} = \begin{pmatrix} u \\ x \end{pmatrix}.$$

The matrices K, P, and R have dimensions

$$K \in \mathrm{Mat}_{M\times(M-h)}(L), \qquad P \in \mathrm{Mat}_{h\times M}(L), \qquad R \in \mathrm{Mat}_{(M-h)\times M}(L),$$

and satisfy

$$RK = I_{M-h}, \qquad PK = 0, \qquad \ker P = \operatorname{im} K.$$

Here R is a linear coordinate map; it need not be a morphism of $\mathbf{t}$-modules or provide a splitting of the exact sequence.

The inclusion and quotient maps. Reorder the full normal basis $\boldsymbol{E}$ with the constant matrix units first in column-vectorization order, followed by the fixed basis $\boldsymbol{E}_+$ of $\mathcal{E}_+$. Express Π_t in the same order. Let K, P, and R be as defined above.

Proposition 6.3. *In the bases above, the inclusion and quotient maps have matrices K and P, respectively. Their $\mathbf{t}$-action matrices are*

$$\Pi_{0,t} = R\Pi_t K, \qquad Q_t = \theta I_h + P_0(I_d \otimes N_\Psi)Z.$$

The identities

$$\Pi_t K = K\Pi_{0,t}, \qquad P\Pi_t = Q_t P$$

hold.

Proof. Left multiplication by N_Ψ has constant-coordinate matrix $I_d \otimes N_\Psi$. Applying this matrix to the quotient representatives Z and projecting by P_0 gives the formula for Q_t. Theorem 6.1 gives $P\Pi_t = Q_t P$.

Using $P\Pi_t = Q_t P$ and $PK = 0$, we obtain

$$P\Pi_t K = Q_t P K = 0.$$

Thus $\Pi_t K$ takes values in $\ker P = \operatorname{im} K$. Since $RK = I$, the restriction of KR to $\operatorname{im} K$ is the identity. Consequently,

$$K(R\Pi_t K) = (KR)\Pi_t K = \Pi_t K.$$

The restricted action, expressed in the basis determined by the columns of K, therefore has matrix $\Pi_{0,t} = R\Pi_t K$. In particular, $\Pi_t K = K\Pi_{0,t}$. □

Computing the exact sequence. The following four procedures compute the coordinate form of the exact sequence. The input consists of valid FJ reduction data $\mathcal{D}$ for (Φ, Ψ), satisfying the degree gap, and the full ordered normal basis $\mathbf{E} = (g_1, \dots, g_M)$ of the form (5.2). Algorithm 9 computes the full action matrix Π_t in this basis.

This procedure first finds a basis of the removable constant coefficients. It then completes this basis to become a basis of all constant matrices. The added columns represent a basis of Q.

Algorithm 10 Constant-coordinate data

Input: Nilpotent matrices $N_\Phi \in \mathrm{Mat}_d(L)$, $N_\Psi \in \mathrm{Mat}_e(L)$.

Output: (h, B, Z, U, P_0), where $\begin{bmatrix} B & Z \end{bmatrix}^{-1} = \begin{bmatrix} U \\ P_0 \end{bmatrix}$.

procedure ConstantData(N_Φ, N_Ψ)
  $n_c \leftarrow ed$
  $D \leftarrow N_\Phi^T \otimes I_e - I_d \otimes N_\Psi$
  $I \leftarrow$ pivot-column indices of D, using exact Gaussian elimination
  $B \leftarrow D_{*\times I}, \qquad r \leftarrow \#I, \qquad h \leftarrow n_c - r$
  $J_B \leftarrow$ pivot-column indices of B^T
  $J_Z \leftarrow \{1, \dots, n_c\} \setminus J_B$, in increasing order
  $Z \leftarrow (I_{n_c})_{*\times J_Z}$,
  $A \leftarrow \begin{bmatrix} B & Z \end{bmatrix}, \qquad G \leftarrow A^{-1}$
  $U \leftarrow G_{1:r\times *}, \qquad P_0 \leftarrow G_{r+1:n_c\times *}$
  return (h, B, Z, U, P_0)
end procedure

The basis of $\mathrm{Ext}_{0,\tau}$ consists of the matrices represented by the columns of B, followed by all positive-degree full normal monomials. The quotient basis is represented by the columns of Z. We first form the coordinate maps with constants first, then translate them back to the original full coordinate order.

Algorithm 11 Bases, inclusion, and projection

Input: Full ordered basis $\mathbf{E}$ of the form (5.2), dimensions d, e, and (h, B, Z, U, P_0) from Algorithm 10.

Output: Bases $\mathbf{E}_0, \mathbf{E}_Q$ and matrices K, R, P in the original full coordinates.

procedure SequenceCoordinates($\mathbf{E}, d, e, h, B, Z, U, P_0$)
  $M \leftarrow \#\mathbf{E}$, $n_c \leftarrow ed$, $\ell \leftarrow M - n_c$, $r \leftarrow n_c - h$
  $I_c \leftarrow$ indices of $E_{i,s}\tau^0$ in $\mathbf{E}$, ordered by s, then i
  $I_+ \leftarrow$ indices of all positive-degree g_β, in their original order
  $\rho \leftarrow (I_c, I_+)$ ▷ concatenation of the two index lists
  $S \leftarrow (I_M)_{*\times\rho}$ ▷ columns of I_M in the order ρ
  $\mathbf{E}_0 \leftarrow \left(\operatorname{vec}^{-1}\left(B_{*\times a}\right)\right)_{a=1}^{r}$ followed by $(g_\beta)_{\beta\in I_+}$
  $\mathbf{E}_Q \leftarrow \left(\operatorname{vec}^{-1}\left(Z_{*\times b}\right) + \operatorname{im}\mathcal{T}\right)_{b=1}^{h}$
  $\widehat{K} \leftarrow \operatorname{diag}(B, I_\ell)$, $\widehat{R} \leftarrow \operatorname{diag}(U, I_\ell)$, $\widehat{P} \leftarrow [P_0\ 0_{h\times\ell}]$
  $K \leftarrow S\widehat{K}$, $R \leftarrow \widehat{R}S^T$, $P \leftarrow \widehat{P}S^T$
  return $(\mathbf{E}_0, \mathbf{E}_Q, K, R, P)$
end procedure

Here, vec^{-1} denotes the inverse of the column-wise vectorization map.

The matrix S converts the reordered full coordinates to the original full coordinates. The products involving S are implemented by row or column permutations. The index lists I_c and I_+ are obtained in a single scan of the stored basis. Because it is a permutation matrix, $S^{-1} = S^T$. The symbol $\operatorname{diag}(B, I_\ell)$ denotes a rectangular block matrix when B is rectangular.

The following procedure computes the action on $\operatorname{Ext}_{0,\tau}$ by applying the full action to its embedded coordinates and recovering its basis coordinates with R. The quotient action uses only N_Ψ.

Algorithm 12 The two induced actions

Input: $\Pi_t = \sum_{a=0}^{s_\tau} C_a\tau^a$ in the original full basis; K, R from Algorithm 11; d; h, Z, P_0 from Algorithm 10; and N_Ψ.

Output: $\Pi_{0,t}$ and Q_t.

procedure SequenceActions($\Pi_t, K, R, d, h, Z, P_0, N_\Psi, \theta$)
  $\Pi_{0,t} \leftarrow 0$, $K_{\mathrm{tw}} \leftarrow K$
  for $a = 0, \ldots, s_\tau$ **do**
    $C_{0,a} \leftarrow RC_aK_{\mathrm{tw}}$, $\Pi_{0,t} \leftarrow \Pi_{0,t} + C_{0,a}\tau^a$
    $K_{\mathrm{tw}} \leftarrow K_{\mathrm{tw}}^{(1)}$, unless $a = s_\tau$
  end for
  $\overline{N} \leftarrow P_0(I_d \otimes N_\Psi)Z$, $Q_t \leftarrow \theta I_h + \overline{N}$
  return $(\Pi_{0,t}, Q_t)$
end procedure

The final procedure combines the preceding constructions and checks the coordinate identities certifying exactness and compatibility with the **t**-actions.

Algorithm 13 Assembly and verification of the exact sequence

Input: Valid reduction data $\mathcal{D}$ for (Φ, Ψ), including the full normal basis $\mathbf{E}$ and the field operations.

Output: Bases and action matrices of the three terms, the inclusion K, the projection P, and h.

procedure ExactSequence($\mathcal{D}$)
  Read $\Phi, \Psi, \theta, d, e, \mathbf{E}$ from $\mathcal{D}$
  $N_\Phi \leftarrow \partial\Phi_t - \theta I_d$
  $N_\Psi \leftarrow \partial\Psi_t - \theta I_e$
  $(h, B, Z, U, P_0) \leftarrow$ ConstantData(N_Φ, N_Ψ)
  $(\mathbf{E}_0, \mathbf{E}_Q, K, R, P) \leftarrow$ SequenceCoordinates($\mathbf{E}, d, e, h, B, Z, U, P_0$)
  $\Pi_t \leftarrow$ ExtMatrix($\mathcal{D}$) ▷ from Algorithm 9
  $(\Pi_{0,t}, Q_t) \leftarrow$ SequenceActions($\Pi_t, K, R, d, h, Z, P_0, N_\Psi, \theta$)
  $M \leftarrow \#\mathbf{E}, \qquad M_0 \leftarrow M - h$
  assert $RK = I_{M_0}$ and $PK = 0$
  assert $\operatorname{rank} K = M_0$ and $\operatorname{rank} P = h$
  Write $\Pi_t = \sum_a C_a \tau^a$
  $\Pi_{0,t} = \sum_a C_{0,a} \tau^a$
  $s_\tau \leftarrow \max\{0, \deg_\tau \Pi_t\}$; missing coefficients are zero
  for $a = 0, \dots, s_\tau$ **do**
    assert $C_a K^{(a)} = K C_{0,a}$
  end for
  assert $PC_0 = Q_t P$ and $PC_a = 0$ for every $a \geq 1$
  assert $(C_{0,0} - \theta I_{M_0})^e = 0$
  assert $(Q_t - \theta I_h)^e = 0$
  return $(h, \mathbf{E}_0, \mathbf{E}, \mathbf{E}_Q, \Pi_{0,t}, \Pi_t, Q_t, K, P)$
end procedure

6.4. **The structure of the quotient.** Proposition 6.3 gives the quotient matrix in coordinates chosen by Gaussian elimination. Its isomorphism type has a simpler description: it depends only on the Jordan block lengths of N_Φ and N_Ψ. Nilpotent Jordan form exists over L without extending the field. Write the block lengths as

$$a_1, \dots, a_u \quad \text{for } N_\Phi, \qquad b_1, \dots, b_v \quad \text{for } N_\Psi.$$

Theorem 6.4. *Set $c_{\alpha\beta} = \min(a_\alpha, b_\beta)$. Then the quotient in Theorem 6.1 is*

$$(Q, \theta \operatorname{id}_Q + \overline{N}_\Psi) \simeq \bigoplus_{\alpha=1}^{u} \bigoplus_{\beta=1}^{v} \left(\mathbb{G}_a^{c_{\alpha\beta}}, \theta I_{c_{\alpha\beta}} + J_{c_{\alpha\beta}} \right),$$

where J_c is a nilpotent Jordan block of length c. In particular,

$$h = \sum_{\alpha=1}^{u} \sum_{\beta=1}^{v} \min(a_\alpha, b_\beta).$$

Proof. Put both nilpotents in Jordan form and decompose W into the rectangular matrix spaces belonging to source and target block pairs. Both $\mathcal{T}$ and left multiplication by N_Ψ preserve this decomposition. We may treat one pair of block lengths a, b.

On $\mathrm{Mat}_{b\times a}(L)$ let x denote right multiplication by J_a and y denote left multiplication by J_b. They commute. Starting with $E_{b,1}$ gives

$$y^j x^i E_{b,1} = E_{b-j,1+i}, \qquad 0 \le i < a, \quad 0 \le j < b.$$

These are all the matrix units. Hence this rectangular space, with its two operators, is the cyclic module $L[x,y]/(x^a, y^b)$.

The map $\mathcal{T}$ is multiplication by $x - y$. Passing to its cokernel imposes the relation $x = y$, so on this pair

$$\operatorname{coker} \mathcal{T} \simeq L[x,y]/(x^a, y^b, x-y) \simeq L[z]/(z^{\min(a,b)}).$$

Since left multiplication by N_Ψ corresponds to multiplication by y, its induced action on the quotient is multiplication by z. In the ordered basis $z^{c-1}, z^{c-2}, \ldots, 1$, with $c = \min(a,b)$, its matrix is J_c. Summing over all block pairs proves the assertion. □

Thus the quotient can have nontrivial nilpotent action despite having no positive powers of τ. Each block of length c is an iterated extension of c copies of $(\mathbb{G}_a, \theta)$.

Corollary 6.5. *For $d, e > 0$, the action on Q is scalar if and only if $N_\Phi = 0$ or $N_\Psi = 0$. If both vanish, then $h = ed$ and $\dim \mathrm{Ext}_{0,\tau}(\Phi, \Psi) = eJ - ed$.*

Proof. The quotient nilpotent vanishes precisely when every $\min(a_\alpha, b_\beta)$ is one. If one of the original nilpotents is zero, all its block lengths are one. If both are nonzero, each has a block of length at least two, and that pair contributes a quotient block of length at least two. The dimension formulas follow from Theorem 6.1. □

For example, if $N_\Phi = N_\Psi = J_2$, then $h = 2$ and the quotient is $(\mathbb{G}_a^2, \theta I_2 + J_2)$. The kernel of $\mathcal{T}$ is spanned by I_2 and J_2. Testing individual matrix units detects only $E_{12} = J_2$ and misses $I_2 = E_{11} + E_{22}$. This explains why the quotient dimension must be computed by linear algebra rather than by counting individual directions of a chosen basis.

Running Example. In the example (3.1) we have $N_\Phi = N_\Psi = 0$, $C = I_2$, and $(j_1, j_2) = (2, 3)$. Consequently $\mathcal{T} = 0$ on $W = L^{1\times 2}$, so

$$Q = W, \qquad h = 2, \qquad M = 5, \qquad M_0 = 3, \qquad Q_t = \theta I_2.$$

Theorem 6.1 therefore applies directly. In Algorithm 10, B is empty, $Z = I_2$, and $P_0 = I_2$. There is no constant basis vector in the kernel.

Every inner biderivation has zero constant coefficient. Thus the full normal representative

$$F = (a_0 + a_1\tau,\ b_0 + b_1\tau + b_2\tau^2)$$

belongs to $\mathrm{Ext}_{0,\tau}$ exactly when $a_0 = b_0 = 0$. Its coordinates are therefore given by

$$\mathbf{E}_0 = (f_1, f_2, f_3) = \big((\tau, 0), (0, \tau), (0, \tau^2)\big).$$

The action can be checked by hand. Set $\Delta = \theta^q - \theta$. For f_1, set $H_1 = (c^q, 0)$. Then

$$\delta_t^{(H_1)} = \Big(c^q\tau^2 - c^{q^2}\tau, 0\Big).$$

Subtracting it from $\Psi_t(cf_1)$ gives $(\theta c + c^{q^2})f_1$. The image $\Psi_t(cf_2) = \theta c f_2 + c^q f_3$ is already normal. For f_3, put $\mu = c^q$ and $H = (-\mu\tau - \mu^q, \mu)$. Direct skew polynomial multiplication gives

$$\delta_t^{(H)} = ((-\Delta\mu + \mu^{q^2})\tau, \mu\tau^2 + \mu\tau^3 - \mu^q\tau).$$

Subtracting this from $\Psi_t(cf_3) = (0, \theta c\tau^2 + c^q\tau^3)$ yields a normal matrix. We obtain

$$\begin{aligned}
\mathrm{red}(\Psi_t(cf_1)) &= (\theta c + c^{q^2})f_1,\\
\mathrm{red}(\Psi_t(cf_2)) &= \theta c f_2 + c^q f_3,\\
\mathrm{red}(\Psi_t(cf_3)) &= (\Delta c^q - c^{q^3})f_1 + c^{q^2} f_2 + (\theta c - c^q)f_3.
\end{aligned}$$

Hence

$$\Pi_{0,t} = \begin{bmatrix} \theta + \tau^2 & 0 & \Delta\tau - \tau^3 \\ 0 & \theta & \tau^2 \\ 0 & \tau & \theta - \tau \end{bmatrix}.$$

Its differential is θI_3, so it satisfies the nilpotence condition (2.1) in the definition of an Anderson $\mathbf{t}$-module.

In the full coordinate order $(a_0, a_1, b_0, b_1, b_2)$, Algorithm 11 gives

$$K = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \qquad P = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}, \qquad R = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Thus $\iota(u, v, w) = (0, u, 0, v, w)$ and $p(a_0, a_1, b_0, b_1, b_2) = (a_0, b_0)$. The projection is surjective. If $Px = 0$, then $x_1 = x_3 = 0$ and $x = K(x_2, x_4, x_5)^T$. Hence $\ker P = \mathrm{im}\, K$, and R recovers the unique kernel coordinates.

Selecting columns $2, 4, 5$ of the full matrix (5.6), and selecting its rows $1, 3$, gives

$$\Pi_t K = K\Pi_{0,t}, \qquad P\Pi_t = \theta P.$$

These are direct identities of the displayed matrices; no software calculation is needed. The exact sequence of Theorem 6.1 is

$$0 \longrightarrow (\mathbb{G}_a^3, \Pi_{0,t}) \xrightarrow{K} (\mathbb{G}_a^5, \Pi_t) \xrightarrow{P} (\mathbb{G}_a^2, \theta I_2) \longrightarrow 0.$$

7. Dual $\mathbf{t}$-modules

Let $\mathcal{C}$ denote the Carlitz module,

$$\mathcal{C}_t = \theta + \tau.$$

Duality for Drinfeld modules was introduced by Taguchi [17]; its description in terms of extension groups was developed in [16]. The corresponding construction for strictly pure $\mathbf{t}$-modules is studied in [12]. We use this extension interpretation for the flag-jump presentations considered here.

Definition 7.1. For a $\mathbf{t}$-module Φ with $N_\Phi = 0$, define its dual by

$$\Phi^\vee := \mathrm{Ext}_{0,\tau}(\Phi, \mathcal{C}). \tag{7.1}$$

This definition gives an $\mathbb{F}_q[t]$-module. When Φ is of type $\mathrm{FJ}_{\sigma,N}$ with $N \geq 2$, the inequality

$$\deg_\tau \mathcal{C}_t = 1 < N$$

allows us to apply the reduction results of the preceding sections. The resulting normal coordinates give the dual a $\mathbf{t}$-module structure, whose action matrix is computed by Algorithm 12 with $\Psi = \mathcal{C}$, after the full action has been computed.

Theorem 7.2 (Weil–Barsotti formula)**.** *Let Φ be of type $\mathrm{FJ}_{\sigma,N}$, with $N \geq 2$ and $N_\Phi = 0$. Write $d = \dim \Phi$ and $J = \sum_{s=1}^{d} j_s$.*

(i) *There is an exact sequence of $\mathbf{t}$-modules*

$$0 \longrightarrow \Phi^\vee \xrightarrow{\iota_\Phi} \mathrm{Ext}^1_\tau(\Phi, \mathcal{C}) \xrightarrow{p_\Phi} (\mathbb{G}_a^d, \theta I_d) \longrightarrow 0. \tag{7.2}$$

The dimensions of the three terms are $J - d$, J, and d, respectively, and $\partial \Phi_t^\vee = \theta I_{J-d}$.

(ii) *Let Ψ be of type $\mathrm{FJ}_{\sigma',N'}$, with $N' \geq 2$ and $N_\Psi = 0$, and let $f : \Phi \to \Psi$ be a morphism. Put $e = \dim \Psi$. Pullback of extensions induces a morphism of $\mathbf{t}$-modules*

$$f^\vee : \Psi^\vee \longrightarrow \Phi^\vee$$

and a commutative diagram with exact rows:

$$\begin{array}{ccccccc} 0 & \longrightarrow & \Psi^\vee & \xrightarrow{\iota_\Psi} & \mathrm{Ext}^1_\tau(\Psi, \mathcal{C}) & \xrightarrow{p_\Psi} & (\mathbb{G}_a^e, \theta I_e) \longrightarrow 0 \\ & & \downarrow{\scriptstyle f^\vee} & & \downarrow{\scriptstyle f^*} & & \downarrow{\scriptstyle (\partial f)^T} \\ 0 & \longrightarrow & \Phi^\vee & \xrightarrow{\iota_\Phi} & \mathrm{Ext}^1_\tau(\Phi, \mathcal{C}) & \xrightarrow{p_\Phi} & (\mathbb{G}_a^d, \theta I_d) \longrightarrow 0. \end{array} \tag{7.3}$$

Here $\partial f = f_0$ is the constant coefficient matrix of f. The right vertical map is written in column coordinates; in row coordinates it is right multiplication by ∂f.

Proof. The degree gap with $\mathcal{C}$ gives a $\mathbf{t}$-module structure on $\mathrm{Ext}^1_\tau(\Phi, \mathcal{C})$ by Theorem 5.6. Since $N_\Phi = N_\mathcal{C} = 0$, the constant-coordinate map $\mathcal{T}$ is zero on $\mathrm{Mat}_{1\times d}(L)$. Its quotient is this entire d-dimensional space, with scalar action θI_d. Thus Theorem 6.1 gives (7.2).

Every inner biderivation has constant coefficient zero. Consequently, the normal representatives of $\Phi^\vee$ have ordered basis

$$\mathbf{E}_0^\Phi = \big(E_{1,s}\tau^r : 1 \le s \le d,\ 1 \le r < j_s\big).$$

Its cardinality is $\sum_s (j_s - 1) = J - d$. The full normal space has dimension J, and taking the constant coefficient gives the quotient of dimension d. The differential of the full action is θI_J; its restriction to the dual coordinates is therefore θI_{J-d}. This proves (i).

For (ii), represent f by its matrix in $\mathrm{Mat}_{e\times d}(L\{\tau\})$, so that $f\Phi_t = \Psi_t f$. If an extension of Ψ by $\mathcal{C}$ is represented by the row V, define

$$f^*([V]) = [Vf]. \tag{7.4}$$

This is well defined on extension classes. Indeed, for an inner representative $H\Psi_t - \mathcal{C}_t H$,

$$(H\Psi_t - \mathcal{C}_t H)f = (Hf)\Phi_t - \mathcal{C}_t(Hf).$$

Hence pullback sends inner biderivations to other inner biderivations. Moreover,

$$(Vf)_0 = V_0 \partial f,$$

so it preserves Der_0 and induces $f^\vee$. Moreover $(\mathcal{C}_t V)f = \mathcal{C}_t(Vf)$, which proves compatibility with the $\mathbf{t}$-action.

These maps are additive polynomial maps in normal coordinates. For each normal basis element g for $\mathrm{Ext}^1_\tau(\Psi,\mathcal{C})$, reduce cgf for the pair $(\Phi,\mathcal{C})$, with formal c. The resulting coordinate polynomials determine the corresponding column of f^*. Restriction to the dual coordinates gives $f^\vee$.

The left square in (7.3) commutes by the definition of $f^\vee$. For the right square, taking constant coefficients gives

$$p_\Phi(f^*([V])) = V_0 \partial f.$$

This is the asserted quotient map, proving (ii). □

Remark 7.3. For composable morphisms $f : \Phi \to \Psi$ and $g : \Psi \to \Omega$ as in the theorem, pullback satisfies $(g \circ f)^* = f^* \circ g^*$ and $(\mathrm{id}_\Phi)^* = \mathrm{id}$. Since $f^\vee$ is induced by the restriction $f^* : \ker p_\Psi \to \ker p_\Phi$, where p_Ψ and p_Φ are the quotient maps in (7.3), these identities give

$$(g \circ f)^\vee = f^\vee \circ g^\vee, \qquad (\mathrm{id}_\Phi)^\vee = \mathrm{id}_{\Phi^\vee}.$$

7.1. **The dual of the running example.** For the example (3.1), $\Psi = \mathcal{C}$. Thus the module computed in Section 6 is $\Phi^\vee$ for

$$\Phi_t = \begin{bmatrix} \theta + \tau^2 & 0 \\ \tau^3 & \theta + \tau^2 + \tau^3 \end{bmatrix}.$$

The jump degrees $(j_1, j_2) = (2, 3)$ contribute one and two dual coordinates, respectively. Hence

$$\dim \Phi^\vee = (2-1) + (3-1) = 3.$$

With representatives $(u\tau, v\tau + w\tau^2)$ its action matrix is

$$\Phi_t^\vee = \begin{bmatrix} \theta + \tau^2 & 0 & \Delta\tau - \tau^3 \\ 0 & \theta & \tau^2 \\ 0 & \tau & \theta - \tau \end{bmatrix}, \quad \text{where} \quad \Delta = \theta^q - \theta. \tag{7.5}$$

The Weil–Barsotti sequence is

$$0 \longrightarrow \Phi^\vee \xrightarrow{\iota_\Phi} \operatorname{Ext}^1_\tau(\Phi, \mathcal{C}) \xrightarrow{p_\Phi} (\mathbb{G}_a^2, \theta I_2) \longrightarrow 0.$$

Its normal-coordinate form is

$$0 \longrightarrow (\mathbb{G}_a^3, \Phi_t^\vee) \xrightarrow{K} (\mathbb{G}_a^5, \Pi_t) \xrightarrow{P} (\mathbb{G}_a^2, \theta I_2) \longrightarrow 0,$$

where K and P are the inclusion and projection matrices from Section 6.

Corollary 7.4. *Let* $c = \max\limits_{\alpha,\beta} \min(a_\alpha, b_\beta)$. *Then* $\operatorname{Ext}^1_\tau(\Phi, \Psi)$ *admits a filtration of* $\mathbf{t}$*-modules*

$$\operatorname{Ext}_{0,\tau}(\Phi, \Psi) = F_0 \subseteq F_1 \subseteq \cdots \subseteq F_c = \operatorname{Ext}^1_\tau(\Phi, \Psi),$$

whose successive quotients satisfy

$$F_r/F_{r-1} \simeq \big(\mathbb{G}_a^{m_r}, \theta I_{m_r}\big), \qquad 1 \le r \le c,$$

where

$$m_r = \#\{\alpha : a_\alpha \ge r\} \cdot \#\{\beta : b_\beta \ge r\}.$$

Proof. Let

$$\pi : \operatorname{Ext}^1_\tau(\Phi, \Psi) \longrightarrow Q$$

be the projection from Theorem 6.1. Set

$$Q_r = \ker\big(\overline{N}_\Psi^r\big), \qquad F_r = \pi^{-1}(Q_r), \qquad 0 \le r \le c.$$

Since the t-action on Q is given by $\theta\,\mathrm{id}_Q + \overline{N}_\Psi$, each Q_r is a $\mathbf{t}$-submodule. Thus the $\mathbf{t}$-modules F_r form the asserted filtration.

By exactness, π induces isomorphisms

$$F_r/F_{r-1} \simeq Q_r/Q_{r-1}.$$

Since $\overline{N}_\Psi(Q_r) \subseteq Q_{r-1}$, the induced $\mathbf{t}$-action on these quotients is scalar.

By Theorem 6.4, $\overline{N}_\Psi$ has one Jordan block of length $\min(a_\alpha, b_\beta)$ for each pair (α, β). Consequently,

$$\begin{aligned} \dim_L(Q_r/Q_{r-1}) &= \#\{(\alpha,\beta) : \min(a_\alpha, b_\beta) \ge r\} \\ &= \#\{\alpha : a_\alpha \ge r\} \cdot \#\{\beta : b_\beta \ge r\} \\ &= m_r. \end{aligned}$$

Therefore,

$$F_r/F_{r-1} \simeq (\mathbb{G}_a^{m_r}, \theta I_{m_r}),$$

as claimed. □

8. Arithmetic complexity

8.1. **The arithmetic model.** We count additions, multiplications, inversions, and exact zero tests in L at unit cost, and use classical Gaussian elimination and matrix multiplication. Coefficient access and the comparisons used to organize the search are also included. Write f_q for the cost of one application of $x \mapsto x^q$ in L. Repeated squaring gives the arithmetic bound $\mathsf{f}_q = O(\log q)$; a representation with a Frobenius primitive may give a smaller cost.

Retain the notation

$$d = \dim \Phi, \qquad e = \dim \Psi, \qquad n = \deg_\tau \Phi_t, \qquad m = \deg_\tau \Psi_t,$$

and assume $m < \min_s j_s$ for Algorithms 7–13. Put

$$j_{\max} = \max_s j_s, \qquad J = \sum_s j_s, \qquad M = eJ.$$

Matrix units and normal basis elements are stored by their indices. Skew polynomials are stored by their coefficient matrices, and a formal additive coefficient

$$f(c) = \sum_{a=0}^{H} f_a c^{q^a}$$

is stored as $(f_0, \dots, f_H)$. Thus its storage and arithmetic depend on $H + 1$.

8.2. **The jump-basis search.** Suppose preprocessing retains $b = n - k + 1$ blocks. As written, Algorithm 1 first computes the rank of a $d \times nd$ matrix and then recomputes the ranks of suffixes containing $1, \dots, b$ blocks. Its cost is therefore

$$O((n + b^2)d^3).$$

If the initial rank is deficient, it stops after $O(nd^3)$ operations. Maintaining an echelon basis while adjoining blocks would reduce the preprocessing bound to $O(nd^3)$; the displayed bound applies to the successive rank computations in the pseudocode.

For a selected basis, Algorithm 4 computes C^{-1} and the b products $C^{-1}A_j$. This costs $O(bd^3)$. The graph has at most $d(d-1)$ edges, and the successive scans for a vertex with no incoming edge cost $O(d^3)$ additional elementary tests. Algorithm 2 examines bd vectors of length d, giving $O(bd^2)$ zero tests. Algorithm 3 uses $O(d \log d)$ comparisons of list lengths.

Let $a_1 \le \cdots \le a_d$ be the ordered list lengths, let ν be the number of candidates examined by the depth-first search, and let ℓ be the number of completed selections tested by Algorithm 4. Then

$$\nu \le \sum_{r=1}^{d} \prod_{i=1}^{r} a_i, \qquad \ell \le \min\left\{ \prod_{i=1}^{d} a_i,\ \frac{d!}{\prod_{j=k}^{n} \mathrm{Required}[j]!} \right\}. \tag{8.1}$$

The second bound uses the required block counts: each original position has at most one candidate in a given block. Recomputing ranks for independence

tests gives $O(\nu d^3)$, while the completed selections cost $O(\ell b d^3)$. Define

$$\mathsf{S} = (n + b^2)d^3 + \nu d^3 + \ell b d^3. \tag{8.2}$$

The bounds for Algorithms 1–7 are summarized below.

Algorithm and procedure	Arithmetic upper bound
1. Preprocess	$O((n + b^2)d^3)$
4. Order	$O(bd^3)$
2. Lists	$O(bd^2)$
3. ShortFirst	$O(d \log d)$ comparisons
5. Search	$O(\nu d^3 + \ell b d^3)$
6. JumpBasis	$O(\mathsf{S})$
7. Prepare	$O(\mathsf{S} + d^3 + ed + M)$

Since $a_i \le b$, (8.1) gives a bound exponential in d. For fixed d, it is polynomial in b. Sorting the lists and enforcing the block counts can reduce the visited tree substantially. Incremental echelon tests can replace νd^3 by νd^2; the bounds above use the direct rank tests described in Section 4.

8.3. **Full reduction and Frobenius indices.** Let D be an upper bound on the degree of an input to Algorithm 8. The number of visited positions is

$$\mathsf{K}(D) = e \sum_{s=1}^{d} \max\{0, D - j_s + 1\} \le ed(D + 1).$$

The visited positions include those whose coefficient is zero; their prescribed order is generated by nested loops.

Let λ_s be row s of the jump-basis matrix C^{-1}. For a correction $H = \mu\tau^k U_{i,s}$, the products are

$$H\Phi_t = \sum_{r=0}^{j_s} \mu E_{i,s}(C^{-1}A_r)^{(k)}\tau^{k+r},$$

$$\Psi_t H = \sum_{v=0}^{m} \mu^{q^v} B_v E_{i,s}(C^{-1})^{(k+v)}\tau^{k+v}.$$

Here $A_0 = \theta I_d + N_\Phi$ and $B_0 = \theta I_e + N_\Psi$. The first product changes one row, and each coefficient of the second is an outer product. Precompute $C^{-1}A_r$ and their required twists, together with the twists of C^{-1}. Since $m < j_s$, for an input of degree at most D, no index greater than D is needed. The shared cache costs

$$\mathsf{A}(D) = (n + 1)d^3 + (D + 1)(n + 2)d^2\mathsf{f}_q.$$

After caching, a correction with scalar coefficient costs

$$O(\kappa), \qquad \kappa = d(n + 1) + ed(m + 1) + m\mathsf{f}_q.$$

This includes updating the accumulated inner matrix.

For a formal additive coefficient, let s bound the number of stored Frobenius indices in every intermediate coefficient. Scalar products, additions and coefficient twists then cost at most s times their scalar counterparts. Define

$$\mathsf{R}(D, s) = s\big(ed(D+1) + \mathsf{K}(D)\kappa\big).$$

One full reduction costs $O(\mathsf{A}(D) + \mathsf{R}(D, s))$; the cache is calculated once for all inputs sharing the same bound D. For coefficients in L, take $s = 1$. Each correction increases the largest formal index by at most m. Starting with an index of at most H therefore gives

$$H_{\mathrm{out}} \le H + m\mathsf{K}(D).$$

Only the Frobenius indices are stored, not all ordinary powers of the formal parameter.

8.4. **The full action and the exact sequence.** For Algorithm 9, put

$$D_* = m + j_{\max} - 1, \qquad s_* = 1 + m\big(1 + \mathsf{K}(D_*)\big).$$

Each input $\Psi_t(cg_\beta)$ has a degree of at most D_* and an initial formal index of at most m. Constructing all M columns and storing all entries of the resulting coefficient matrices, including zero entries, gives

$$T_9 = \mathsf{A}(D_*) + M\mathsf{R}(D_*, s_*) + M^2 s_*. \tag{8.3}$$

In particular, Π_t has at most s_* coefficient matrices.

For Algorithms 10–13 write

$$n_c = ed, \qquad r = \operatorname{rank} \mathcal{T}, \qquad h = n_c - r, \qquad M_0 = M - h,$$
$$s_\Pi = 1 + \max\{0, \deg_\tau \Pi_t\} \le s_*.$$

The bounds below use dense coefficient matrices and classical matrix arithmetic. They start with prepared reduction data.

Algorithm 10 performs echelon computations on D and B^T, and inverts the $n_c \times n_c$ matrix $[B\ Z]$. Its cost is $O(n_c^3)$. The nonpivot row indices of B provide the complement, so this construction requires only two echelon computations and one inversion. Algorithm 11 scans the stored basis and forms the coordinate maps. Permutations are implemented by index reordering. Writing the dense matrices K, R, P costs $O(M^2)$; no dense multiplication by the permutation matrix is required.

For each full coefficient C_a, the product $RC_aK^{(a)}$ costs

$$O(M_0M^2 + M_0^2M) = O(M^2M_0).$$

Successive entrywise twists of K cost $O(s_\Pi M M_0 \mathsf{f}_q)$ in total. Dense computation of $P_0(I_d \otimes N_\Psi)Z$ costs $O(n_c^2 h + n_c h^2)$. Thus Algorithm 12 costs

$$T_{12} = s_\Pi M^2 M_0 + s_\Pi M M_0 \mathsf{f}_q + n_c^2 h + n_c h^2.$$

This bound includes writing the $s_\Pi M_0^2$ restricted-action coefficients and the h^2 quotient entries. In particular, it is bounded by $O(s_\Pi M^3 + s_\Pi M^2 \mathsf{f}_q + n_c^3)$. The inclusion and projection are constant matrices and need no further reduction of representatives.

The assertions in Algorithm 13 have a cost that must also be counted. Rank and constant-product checks cost $O(M^3)$. Checking all coefficient intertwining identities costs $O(s_\Pi M^3)$, together with the Frobenius chain if it is not reused. Repeated squaring for the two nilpotence checks costs $O((M_0^3 + h^3)\log(e+1))$. A uniform verification bound is therefore

$$V_{13} = s_\Pi M^3 + s_\Pi M M_0 \mathsf{f}_q + (M_0^3 + h^3)\log(e+1). \tag{8.4}$$

If the full action is not supplied, Algorithm 13 also pays the cost T_9. The resulting bounds are

Procedure	Bound inside $O(\cdot)$
8. Reduce	$\mathsf{A}(D) + \mathsf{R}(D,1)$
9. ExtMatrix	T_9
10. ConstantData	n_c^3
11. SequenceCoordinates	M^2
12. SequenceActions	T_{12}
13. ExactSequence	$T_9 + n_c^3 + M^2 + T_{12} + V_{13}$

When Π_t is already available, omit T_9 from the last row. Sparse storage can reduce the output and multiplication costs, but the displayed bounds make no sparsity assumption. The dense postprocessing uses $O(s_\Pi M^2 + n_c^2)$ field elements, including its full input and output matrices. Successive twists of K are generated one at a time, rather than storing an additional matrix for every Frobenius index.

Running Example. In the example (3.1) $d = 2$, $e = 1$, $m = 1$, $(j_1, j_2) = (2, 3)$, and $\mathcal{T} = 0$. Hence $n_c = 2$, $r = 0$, $h = 2$, $M = 5$, and $M_0 = 3$. Also $D_* = 3$, $\mathsf{K}(3) = 3$, and $s_* = 5$. The displayed full action actually has degree three, so $s_\Pi = 4$. The inclusion K and its left inverse R are coordinate selections with entries in $\mathbb{F}_q$. Every twist fixes them. Consequently the restricted action is obtained by extracting rows and columns $2, 4, 5$ from Π_t, rather than performing dense products or running another reduction. The quotient is θI_2.

After the jump-basis search, the work in Algorithms 7–13 is polynomial in the dimensions, degree bounds, and f_q. These are arithmetic, not bit-complexity, bounds. Over $\mathbb{F}_q(\theta)$, for example, θ^{q^a} has degree q^a in θ; coefficient growth must be accounted for separately.

The accompanying software *T-Ext Explorer* [10] provides the basis search and full reduction used by Algorithms 1–9. Algorithms 10–13 compute the coordinate maps of the exact sequence, the induced actions, and their verification.

Appendix A. Partial reduction of zero-constant biderivations

In this appendix, we study exact sequences of $\mathbf{t}$-modules with middle term $\mathrm{Ext}_\tau^1(\Phi, \Psi)$ and quotient on which t acts as multiplication by θ. We

identify conditions under which the kernel of the quotient map coincides with $\mathrm{Ext}_{0,\tau}(\Phi,\Psi)$.

The following notation and assumptions apply throughout this appendix.

Fix an FJ presentation Φ_t of dimension d, with compatible order $\sigma=\rho^{-1}$, jump matrix C, and jump degrees $j_1,\dots,j_d$. Suppose that

$$\dim \Psi = e \quad \text{and} \quad m = \deg_\tau \Psi_t < N := \min_s j_s.$$

Set

$$J=\sum_{s=1}^{d} j_s \qquad \text{and} \qquad M = eJ.$$

Put

$$W=\mathrm{Mat}_{e\times d}(L), \qquad \mathcal{T}(H_0)=H_0N_\Phi - N_\Psi H_0, \qquad \widetilde{K}=\ker \mathcal{T}.$$

Choose an L-linear subspace $\widetilde{B}\subseteq\widetilde{K}$ such that $N_\Psi\widetilde{B}\subseteq\widetilde{B}$. The action on the quotient $\widetilde{K}/\widetilde{B}$ is given by $\theta\,\mathrm{id}_{\widetilde{K}/\widetilde{B}}+\nu$, where $\nu(H_0+\widetilde{B})=N_\Psi H_0+\widetilde{B}$. Partial reduction means allowing all positive-degree corrections, but allowing constant corrections only from $\widetilde{B}$. More precisely, define

$$\mathcal{P}_{\widetilde{B}}=\Big\{H\in\mathrm{Mat}_{e\times d}(L\{\tau\}) : H_0\in\widetilde{B}\Big\}, \qquad \mathcal{I}_{\widetilde{B}}=\Big\{\delta^{(H)} : H\in\mathcal{P}_{\widetilde{B}}\Big\},$$

and set

$$\mathcal{E}_{\widetilde{B}}=\mathrm{Der}_0(\Phi,\Psi)/\mathcal{I}_{\widetilde{B}}.$$

Here $\delta_t^{(H)}=H\Phi_t-\Psi_tH$. The inclusion $\mathcal{I}_{\widetilde{B}}\subseteq\mathrm{Der}_0(\Phi,\Psi)$ follows from $\big(\delta_t^{(H)}\big)_0=\mathcal{T}(H_0)$. Taking $\widetilde{B}=\widetilde{K}$ gives $\mathcal{E}_{\widetilde{K}}=\mathrm{Ext}_{0,\tau}(\Phi,\Psi)$. Notice that the map

$$\pi_{\widetilde{B}}:\mathcal{E}_{\widetilde{B}}\to\mathrm{Ext}_{0,\tau}(\Phi,\Psi), \qquad \pi_{\widetilde{B}}\big([\delta]_{\widetilde{B}}\big)=[\delta]_{\widetilde{K}}$$

is surjective. By the definitions of $\mathcal{P}_{\widetilde{K}}$, $\mathcal{I}_{\widetilde{K}}$ and $\mathcal{E}_{\widetilde{K}}$ we see that the map

$$j_{\widetilde{B}}:\big(\widetilde{K}/\widetilde{B},\theta\,\mathrm{id}_{\widetilde{K}/\widetilde{B}}+\nu\big)\to\mathcal{E}_{\widetilde{B}}, \qquad j_B(H_0+\widetilde{B})=\Big[\delta^{(H_0)}\Big]_{\widetilde{B}}$$

is injective.

Theorem A.1. *Let $\widetilde{h}=\dim_L\widetilde{K}$, and $\widetilde{b}=\dim_L\widetilde{B}$. Normal representatives endow $\mathcal{E}_{\widetilde{B}}$ with an Anderson $\mathbf{t}$-module structure of dimension $M-\widetilde{b}$. There is an exact sequence of $\mathbf{t}$-modules*

$$0\longrightarrow\Big(\widetilde{K}/\widetilde{B},\theta\,\mathrm{id}_{\widetilde{K}/\widetilde{B}}+\nu\Big)\xrightarrow{j_{\widetilde{B}}}\mathcal{E}_{\widetilde{B}}\xrightarrow{\pi_{\widetilde{B}}}\mathrm{Ext}_{0,\tau}(\Phi,\Psi)\longrightarrow 0.$$

The kernel has dimension $\widetilde{h}-\widetilde{b}$, and the last term has dimension $M-\widetilde{h}$.

Proof. First, $N_\Psi\widetilde{K}\subseteq\widetilde{K}$, since $\mathcal{T}(N_\Psi H)=N_\Psi\mathcal{T}(H)$. Since also $N_\Psi\widetilde{B}\subseteq\widetilde{B}$, ν is well defined. Moreover, $\nu^e=0$, because $N_\Psi^e=0$. Hence $\Big(\widetilde{K}/\widetilde{B},\theta\,\mathrm{id}+\nu\Big)$ is an Anderson $\mathbf{t}$-module: its action is $\mathbb{F}_q$-linear and its differential minus $\theta\,\mathrm{id}$ is nilpotent.

We describe the algebraic realization of $\mathcal{E}_{\widetilde{B}}$. Let $\mathcal{P}_+ = \{H : H_0 = 0\}$. Reduction using corrections from $\mathcal{P}_+$ gives unique normal representatives in

$$\mathcal{V} = \operatorname{span}_L \{E_{i,s}\tau^r : 1 \le i \le e,\ 1 \le s \le d,\ 1 \le r \le j_s\}.$$

The pivot argument used for full reduction applies with each threshold j_s replaced by $j_s + 1$. Thus $\operatorname{Der}_0 / \{\delta^{(H)} : H \in \mathcal{P}_+\}$ is realized by $\mathcal{V} \simeq \mathbb{G}_a^M$. Its $\mathbf{t}$-action is obtained by reducing $\Psi_t V$. Since $(\theta I_e + N_\Psi)V$ is already in normal form and all remaining terms involve positive Frobenius powers of the coordinates of V, the differential of the induced $\mathbf{t}$-action acts on each coefficient by $\theta I_e + N_\Psi$. Consequently, its nilpotent part has an exponent of at most e. Its difference from θI_M is therefore nilpotent.

For completeness, the quotient by the constant corrections from $\widetilde{B}$ also has polynomial coordinates. Write red_+ for this reduction, and let

$$\Gamma(H_0) = \operatorname{red}_+ \left(\delta_t^{(H_0)} - \mathcal{T}(H_0)\right), \qquad H_0 \in W.$$

Let $\beta : \mathcal{V} \to W$ extract the coefficients in positions (i, s) at degrees j_s. The leading pivots of $\delta_t^{(U_{i\times s})}$ have coefficient 1. In decreasing degree and then in the compatible column order, all other retained boundary coefficients occur later. Consequently $f = \beta\Gamma$ is a triangular additive polynomial automorphism of W. If $\mathcal{V}_< = \ker\beta$, the map

$$V \longmapsto (f^{-1}\beta(V),\ V - \Gamma(f^{-1}\beta(V)))$$

is an additive polynomial isomorphism $\mathcal{V} \simeq W \oplus \mathcal{V}_<$. For $H_0 \in \widetilde{B}$, one has $\mathcal{T}(H_0) = 0$, so the image of its inner correction corresponds to $(H_0, 0)$. The quotient is therefore isomorphic as an additive algebraic group to

$$\left(W/\widetilde{B}\right) \oplus \mathcal{V}_< \simeq \mathbb{G}_a^{M-\widetilde{b}}.$$

It is stable under t, by

$$\Psi_t \delta_t^{(H)} = \delta_t^{(\Psi_t H)}, \qquad (\Psi_t H)_0 = (\theta I_e + N_\Psi) H_0.$$

Its differential is the induced operator on the quotient of $\operatorname{Lie}(\mathcal{V})$. The nilpotence condition in the definition of an Anderson $\mathbf{t}$-module therefore passes to $\mathcal{E}_{\widetilde{B}}$.

Every element of its kernel is represented by $\delta^{(H)}$ with $H_0 \in \widetilde{K}$. Since $H - H_0 \in \mathcal{P}_+$, its class equals $\left[\delta^{(H_0)}\right]_{\widetilde{B}}$. Conversely, every such class maps to zero. The injectivity of $H \mapsto \delta_t^{(H)}$ under the FJ degree-gap assumptions shows that

$$\left[\delta^{(H_0)}\right]_{\widetilde{B}} = 0 \quad \Longleftrightarrow \quad H_0 \in \widetilde{B}.$$

This proves exactness and the dimension formulas.

Finally, the same identity for $\Psi_t \delta_t^{(H)}$ gives

$$t\, j_{\widetilde{B}}\left(H_0 + \widetilde{B}\right) = j_{\widetilde{B}}\left((\theta I_e + N_\Psi) H_0 + \widetilde{B}\right).$$

Thus $j_{\widetilde{B}}$ is a morphism of $\mathbf{t}$-modules. The quotient map $\pi_{\widetilde{B}}$ also respects t.

□

The choice of individual basic directions. With $U_{i\times s} = E_{i\times s}C^{-1}$, put

$$B_{\mathrm{bas}} = \mathrm{span}_L \left\{ U_{i\times s} \mid U_{i\times s}N_\Phi = N_\Psi U_{i\times s} \right\}.$$

In the strictly lower triangular presentation convention, $N_\Psi B_{\mathrm{bas}} = 0$. Indeed, $U_{i\times s}N_\Phi$ is supported in row i, whereas $N_\Psi U_{i\times s}$ is supported in rows strictly following i; their equality forces both matrices to vanish. The theorem therefore applies to this choice. One may write $\widetilde{\mathcal{E}}_0 = \mathcal{E}_{B_{\mathrm{bas}}}$. Its kernel in the displayed sequence is

$$\left(\widetilde{K}/B_{\mathrm{bas}},\ \theta\,\mathrm{id} + \nu\right),$$

which measures the relations arising from linear combinations of basic directions that are not already included in B_{bas}. Even though $N_\Psi B_{\mathrm{bas}} = 0$, the induced operator ν on $\widetilde{K}/B_{\mathrm{bas}}$ need not vanish.

The constructions in [11, Theorem 9.2] and [6, Theorem 4.1] can be interpreted in terms of the partial reduction described above. More precisely, reduction along individual basic directions yields the intermediate $\mathbf{t}$-module $\widetilde{\mathcal{E}}_0 = \mathcal{E}_{B_{\mathrm{bas}}}$, which was identified there with $\mathrm{Ext}_{0,\tau}(\Phi, \Psi)$. This identification need not hold in general, since linear combinations of basic directions may give rise to additional constant inner relations. The precise relation is given by the exact sequence

$$0 \longrightarrow \left(\widetilde{K}/B_{\mathrm{bas}}, \theta\,\mathrm{id} + \nu\right) \longrightarrow \widetilde{\mathcal{E}}_0 \longrightarrow \mathrm{Ext}_{0,\tau}(\Phi, \Psi) \longrightarrow 0.$$

Thus, the earlier identification holds precisely when $B_{\mathrm{bas}} = \widetilde{K}$. In particular, this is the case when both N_Φ and N_Ψ vanish.

A.1. **Partial basis and nontrivial example.** The following example distinguishes a count of individual coordinate intertwiners from the dimension of the full intertwining space. It also shows why the action on the constant quotient need not be scalar. Let $L = \mathbb{F}_q(\theta)$ and put

$$J_2 = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \qquad \Phi_t = \begin{bmatrix} \theta + \tau^2 & 1 \\ 0 & \theta + \tau^2 \end{bmatrix}, \qquad \Psi_t = \begin{bmatrix} \theta & 1 + \tau \\ 0 & \theta \end{bmatrix}.$$

Thus $N_\Phi = N_\Psi = J_2$. The leading coefficient of Φ_t is I_2, so we may take $C = I_2$, $\sigma = (1, 2)$ and $(j_1, j_2) = (2, 2)$. The degree gap holds: $\deg_\tau \Psi_t = 1 < 2$. The full normal representatives have degree less than two in every entry. Consequently, $M = 8$, with ordered basis

$$\mathcal{E} = (E_{1\times1}, E_{1\times1}\tau, E_{1\times2}, E_{1\times2}\tau, E_{2\times1}, E_{2\times1}\tau, E_{2\times2}, E_{2\times2}\tau). \tag{A.1}$$

The extra τ in Ψ_t acts on a scalar coefficient c by $c \mapsto c^q$. Terms of degree two are reduced using inner biderivations. In the order (A.1), the full action

is

$$\Pi_t = \begin{bmatrix} \theta I_4 & A(\tau) \\ 0 & \theta I_4 \end{bmatrix}, \qquad A(\tau) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ \tau & 1 & 0 & 0 \\ 0 & -\tau & 1 & 0 \\ 0 & 0 & \tau & 1 \end{bmatrix}.$$

For example, the degree-two term in $\Psi_t(cE_{2\times 1}\tau)$ is $c^q E_{1\times 1}\tau^2$, whose normal reduction is $-c^q E_{1\times 2}$. This accounts for the entry $-\tau$ in $A(\tau)$.

For $H_0 = \begin{bmatrix} x & y \\ z & w \end{bmatrix}$, the constant map is

$$\mathcal{T}(H_0) = H_0 J_2 - J_2 H_0 = \begin{bmatrix} -z & x-w \\ 0 & z \end{bmatrix}.$$

With column vectorization $\mathrm{vec}(H_0) = (x, z, y, w)^T$, its matrix is

$$D = \begin{bmatrix} 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & 0 \end{bmatrix}.$$

Hence

$$\operatorname{rank} D = 2, \qquad \ker \mathcal{T} = LI_2 \oplus LE_{1\times 2}, \qquad \widetilde{h} = 4 - \operatorname{rank} D = 2.$$

Among the four individual matrix units, only $E_{1\times 2}$ lies in this kernel. The second kernel direction, $I_2 = E_{1\times 1} + E_{2\times 2}$, is a linear combination of two units whose images are nonzero. Thus the rule $\varepsilon_{i\times s} = 0$ when $E_{i\times s}J_2 = J_2 E_{i\times s}$, and 1 otherwise, gives

$$\varepsilon = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}.$$

The number of zero entries is one, whereas the full kernel dimension is two. A dimension calculation using the individual count would predict $8 - 1 = 7$; the correct dimension of $\mathrm{Ext}_{0,\tau}(\Phi, \Psi)$ is $8 - \widetilde{h} = 6$.

This discrepancy is visible directly in the candidate partial basis. The shifted degree ranges $1 \le r < j_s + \varepsilon_{i,s}$ give

$$(E_{1\times 1}\tau, E_{1\times 1}\tau^2, E_{1\times 2}\tau, E_{2\times 1}\tau, E_{2\times 1}\tau^2, E_{2\times 2}\tau, E_{2\times 2}\tau^2).$$

These seven representatives do not give seven independent additive coordinates. For a constant matrix H_0,

$$\delta_t^{(H_0)} = H_0\Phi_t - \Psi_t H_0 = \mathcal{T}(H_0) + H_0\tau^2 - J_2 H_0^{(1)}\tau.$$

In particular, for every $c \in L$,

$$\delta_t^{(cI_2)} = cI_2\tau^2 - c^q E_{1\times 2}\tau.$$

Consequently,

$$\left[cE_{1\times 1}\tau^2\right] + \left[cE_{2\times 2}\tau^2\right] - \left[c^q E_{1\times 2}\tau\right] = 0. \tag{A.2}$$

In the displayed candidate order, the additive parameter vector $(0, c, -c^q, 0, 0, 0, c)$ therefore represents zero. The missing intertwiner I_2 produces this one-parameter redundancy, missed by counting individual coordinate directions. The family parametrizes the subgroup but is not a coordinate basis; the partial basis applicability test rejects it.

The image of the constant map is

$$\operatorname{im}\mathcal{T} = LE_{12} \oplus L(E_{22} - E_{11}).$$

The rank-based construction consequently gives the ordered basis

$$\mathcal{E}_0 = (E_{1\times2}, E_{2\times2} - E_{1\times1}, E_{1\times1}\tau, E_{1\times2}\tau, E_{2\times1}\tau, E_{2\times2}\tau). \tag{A.3}$$

The first two normal representatives have nonzero constant terms, but these constants are removable by inner biderivations. Indeed,

$$[cE_{1\times2}] = \left[-cE_{1\times1}\tau^2\right], \qquad [c(E_{2\times2}-E_{1\times1})] = \left[c^qE_{1\times1}\tau - cE_{2\times1}\tau^2\right] \quad (c \in L).$$

Both classes therefore admit representatives with zero constant term, as required by the definition of $\mathrm{Ext}_{0,\tau}$. In the order (A.3), the restricted action is

$$\Pi_{0,t} = \begin{bmatrix} \theta & 1 & 0 & 0 & -\tau & 0 \\ 0 & \theta & 0 & 0 & 0 & 0 \\ 0 & 0 & \theta & 0 & 1 & 0 \\ 0 & \tau & 0 & \theta & 0 & 1 \\ 0 & 0 & 0 & 0 & \theta & 0 \\ 0 & 0 & 0 & 0 & 0 & \theta \end{bmatrix}.$$

The inclusion, expressed in the full normal coordinates, is the constant map

$$K(u_1, \ldots, u_6)^T = (-u_2, u_3, u_1, u_4, 0, u_5, u_2, u_6)^T. \tag{A.4}$$

Choose $[E_{21}]$ and $[E_{22}]$ as the ordered basis of $\mathrm{Mat}_2(L)/\operatorname{im}\mathcal{T}$. For $X = \begin{bmatrix} x & y \\ z & w \end{bmatrix}$, its quotient coordinates are $(z, x+w)^T$. In the full coordinate order (A.1), the projection is therefore

$$P(x_1, \ldots, x_8)^T = (x_5, x_1 + x_7)^T. \tag{A.5}$$

The induced action of J_2 sends $[E_{2\times1}]$ to $[E_{1\times1}] = [E_{2\times2}]$ and $[E_{2\times2}]$ to $[E_{1\times2}] = 0$. Thus the quotient action is

$$Q_t = \begin{bmatrix} \theta & 0 \\ 1 & \theta \end{bmatrix}, \qquad Q_t \neq \theta I_2. \tag{A.6}$$

The maps above satisfy

$$\Pi_t K = K\Pi_{0,t}, \qquad P\Pi_t = Q_t P, \qquad PK = 0.$$

Moreover, K is injective, P is surjective and $\ker P = \operatorname{im} K$. These facts follow immediately from (A.4)–(A.5). These identities are understood in the ring $L\{\tau\}$. The entries of K and P are $0, 1, -1$, hence Frobenius-fixed,

so the identities can be checked coefficientwise despite the τ-terms in the action matrices. We obtain the exact sequence of $\mathbf{t}$-modules

$$0 \longrightarrow (\mathbb{G}_a^6, \Pi_{0,t}) \xrightarrow{K} (\mathbb{G}_a^8, \Pi_t) \xrightarrow{P} (\mathbb{G}_a^2, Q_t) \longrightarrow 0. \tag{A.7}$$

In terms of extension classes, its first two terms are $\mathrm{Ext}_{0,\tau}(\Phi, \Psi)$ and $\mathrm{Ext}^1_\tau(\Phi, \Psi)$, respectively. The differentials of Π_t, $\Pi_{0,t}$ and Q_t are θI plus square-zero nilpotent matrices. The computations and the dependence (A.2) remain valid in every characteristic, including characteristic two.

This example shows that the full rank calculation and the induced quotient action are essential: the number of individual coordinate intertwiners need not equal h, and the constant quotient need not have scalar $\mathbf{t}$-action.

References

[1] G. W. Anderson, *t-motives*, Duke Math. J. **53** (1986), no. 2, 457–502.

[2] C. Armana, E. Berardini, X. Caruso, A. Leudière, J. Nardi, and F. Pazuki, *A computational approach to Drinfeld modules*, preprint, 2026, arXiv:2601.02162.

[3] D. Ayotte, X. Caruso, A. Leudière, and J. Musleh, *Drinfeld modules in SageMath*, ACM Commun. Comput. Algebra **57** (2023), no. 2, 65–71, arXiv:2305.00422.

[4] W. D. Brownawell and M. A. Papanikolas, *A rapid introduction to Drinfeld modules, t-modules, and t-motives*, in *t-Motives: Hodge Structures, Transcendence and Other Motivic Aspects*, EMS Publishing House, 2020, 3–30.

[5] S. Garai and M. Papikian, *Computing endomorphism rings and Frobenius matrices of Drinfeld modules*, J. Number Theory **237** (2022), 145–164.

[6] F. Głoch, D. E. Kędzierski, and P. Krasoń, *Algorithms for determination of t-module structures on some extension groups*, Int. J. Number Theory **21** (2025), no. 8, 1889–1922, doi:10.1142/S1793042125500927.

[7] F. Głoch, D. E. Kędzierski, and P. Krasoń, *Corrigendum: Algorithms for determination of t-module structures on some extension groups*, Int. J. Number Theory **22** (2026), no. 9, 2175–2178, doi:10.1142/S1793042126920028.

[8] D. Goss, *Basic Structures of Function Field Arithmetic*, Springer, Berlin, 1996.

[9] A. B. Kahn, *Topological Sorting of Large Networks*, Communications of the ACM, Volume 5 (1962), 558–562.

[10] D. E. Kędzierski, D. Koellmer, and P. Krasoń, *T-Ext Explorer: Computing extensions of flag-jump t-modules*, software, version `1.0.0`, Zenodo, 2026, doi:10.5281/zenodo.24081823.

[11] D. E. Kędzierski and P. Krasoń, *On Ext^1 for Drinfeld modules*, J. Number Theory **256** (2024), 97–135.

[12] D. E. Kędzierski and P. Krasoń, *Weil–Barsotti formula for t-modules*, J. Number Theory **278** (2026), 1–25.

[13] D. E. Kędzierski and P. Krasoń, *Homological methods in the generalization of Drinfeld modules*, preprint, arXiv:2512.07607, version 3, 2026.

[14] A. Maurischat, *Abelian equals A-finite for Anderson A-modules*, Ann. Inst. Fourier, online first (2026), doi:10.5802/aif.3754.

[15] A. Maurischat, *Determining t-motives and dual t-motives in Anderson's theory*, preprint, arXiv:2505.12779, version 2, 2026.

[16] M. A. Papanikolas and N. Ramachandran, *A Weil–Barsotti formula for Drinfeld modules*, J. Number Theory **98** (2003), 407–431.

[17] Y. Taguchi, *A duality for finite t-modules*, J. Math. Sci. Univ. Tokyo **2** (1995), 563–588.

[18] D. S. Thakur, *Function Field Arithmetic*, World Scientific, Singapore, 2004.

Institute of Mathematics, University of Szczecin, ul. Wielkopolska 15, 70-451 Szczecin, Poland
Email address: dawidedmundkedzierski@gmail.com

Independent Researcher
Email address: dougkoellmer@hotmail.com

Department of Mathematics, Ariel University, Ariel 40700, Israel
Email address: piotrkras26@gmail.com