(DRAFT:) THE BARCODE THEOREM, JORDAN CANONICAL FORM (AND AN APPENDIX ON MODULES OVER PID'S)

JOEL FRIEDMAN

Contents

Disclaimer: The material may sketchy and/or contain errors, which I will elaborate upon and/or correct in class. For those not in CPSC 531F: use this material at your own risk...

1. The Barcode Theorem

The "barcode theorem" is a theorem in linear algebra that is an integral part of persistent homology, first discovered in [ELZ02, ELZ00]. Yet, the "barcode theorem" can be viewed as a general theorem in linear algebra, and specifically [CZCG04] a consequence of the structure of graded modules over a PID.

In the case of usual interest to us, this theorem results from the Jordan canonical form of a matrix, specifically that of a matrix, M, that is nilpotent, i.e., such that $M^k = 0$ for some sufficiently large integer k (equivalently 0 is the only eigenvalue of M).

Definition 1.1. Let $\mathbb{F}$ be a field. Let $n \geq 0$ be an integer. A string of $\mathbb{F}$-vector spaces of length $n+1$ refers to the data consisting of a sequence $V^0, \ldots, V^n$ of vector spaces, and linear maps $\mathcal{L}^i \colon V^i \to V^{i+1}$ for $i = 0, \ldots, n-1$. We often use the symbols $V^\cdot, \mathcal{L}^\cdot$ to refer collectively to $\{V^i\}_{0 \leq i \leq n}$ and $\{\mathcal{L}^i\}_{0 \leq i \leq n-1}$, and $(V^\cdot, \mathcal{L}^\cdot)$ to the string.

Date: Thursday 2nd January, 2025, at 12:25.

Research supported in part by an NSERC grant.

We may represent a string of vector spaces with the diagram:

$$V^0 \xrightarrow{\mathcal{L}^0} V^1 \xrightarrow{\mathcal{L}^1} \cdots \xrightarrow{\mathcal{L}^{n-1}} V^n.$$

Definition 1.2. In Definition 1.1, let i, j be integers with $0 \le i \le j \le n$. For $0 \le k \le n$, let $V_k = \mathbb{F}$ if $i \le k \le j$, and otherwise let $V_k = 0$. For $i \le k \le j-1$, let $\mathcal{L}^k$ be the identity map. We call this string the (i,j)-*bar*. As a diagram, the (i,j)-bar can be represented as:

$$0 \to \cdots \to 0 \to \mathbb{F} \to \cdots \to \mathbb{F} \to 0 \to \cdots \to 0,$$

where the first appearance of $\mathbb{F}$ is in V^i, and the last in V^j, and all morphisms $\mathbb{F} \to \mathbb{F}$ are the identity maps (there is only one morphism $0 \to \mathbb{F}$, and only one $\mathbb{F} \to 0$). By a *bar* we mean any (i,j)-bar.

Our main theorem states that any string of finite dimensional vector spaces is *isomorphic* to a *direct sum* of bars of the form in Definition 1.1. Hence we need to define the *direct sum* of two strings and a *morphism* (and *isomorphism*) from one string to another. The reader can likely guess the definitions. Before giving them, we remark that these definitions are well known in the literature. (These remarks can be skipped by the reader unfamiliar with or unenthusiastic about sheaf theory.)

Remark 1.3. Definition 1.1 is a *presheaf* on the category that is a directed path, endowed with the coarsest topology (*la topologie grossière*), see [sga72], Exposés I and II, or, for specifics, [Fri05], Theorem 2.1. This is the notion of morphism we will use, and we will also use the resulting notion of direct sum.

Remark 1.4. Definition 1.1 can also be understood as a sheaf (in the classic sense) over the topological space whose underlying set is $\{0, 1, \ldots, n\}$, and whose $n+2$ open subsets are those of the form $U_i = \{i, i+1, \ldots, n\}$ for some $0 \le i \le n+1$ (so $U_{n+1} = \emptyset$): this follows since each of the $n+1$ non-empty open subsets (U_i with $0 \le i \le n$) is *irreducible* in the sense of [Fri05], Theorem 2.1.

Definition 1.5. Let $\mathcal{S}_1 = (V_1^\cdot, \mathcal{L}_1^\cdot)$, $\mathcal{S}_2 = (V_2^\cdot, \mathcal{L}_2^\cdot)$ be strings of $\mathbb{F}$-vector spaces of the same length $n+1$. The *direct sum* of $\mathcal{S}_1$ and $\mathcal{S}_2$ is the string whose i-th vector space is $V_1^i \oplus V_2^i$, and whose i-th morphism is $\mathcal{L}_1^i \oplus \mathcal{L}_2^i$.

Hence the diagram representing the direct sum is:

$$V_1^0 \oplus V_2^0 \xrightarrow{\mathcal{L}_1^0 \oplus \mathcal{L}_2^0} \cdots \xrightarrow{\mathcal{L}_1^{n-1} \oplus \mathcal{L}_2^{n-1}} V_1^n \oplus V_2^n.$$

The direct sum of any set of strings is similarly defined.

Example 1.6. Let $n = 2$. The direct sum of the bar $(0,0)$, the bar $(2,2)$, 2 copies of the bar $(1,2)$, and the bar $(0,2)$ is visualized by the *barcode*

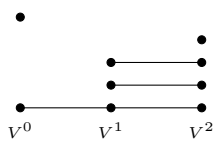

To formalize this, we label each bar with a unique letter from $A, \ldots, E$:

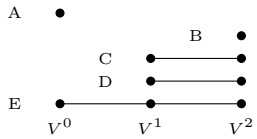

This describes V^0, V^1, V^2 as

$$V^0 = \mathbb{F}^{\{A,E\}} \to V^1 = \mathbb{F}^{\{C,D,E\}} \to V^2 = \mathbb{F}^{\{B,C,D,E\}},$$

where we understand the following convention: if S is a set, then $\mathbb{F}^S$ (as usual) refers to the $\mathbb{F}$-vector space of maps $S \to \mathbb{F}$; if S_1, S_2 are two sets (one usually thinks of $S_1, S_2 \subset T$ as subsets of an "ambient" set T), then one defines a "canonical map" $\mathcal{K}_{S_1 \to S_2} \colon \mathbb{F}^{S_1} \to \mathbb{F}^{S_2}$ taking $\mathbf{v} \in \mathbb{F}^{S_1}$ to the function that agrees on $\mathbf{v}$ on $S_1 \cap S_2$, and otherwise, i.e., on $S_2 \setminus S_1$, takes the value 0.

Remark 1.7. In the above example we have subsets $S_1, S_2, S_3 \subset T$ where $T = \{A, \ldots, E\}$, and a sequence

$$\mathbb{F}^{S_1} \to \mathbb{F}^{S_2} \to \mathbb{F}^{S_3}. \tag{1}$$

Note that since $E \in S_1, S_3$, E represents the $(0, 2)$ bar, and hence we also have $E \in S_2$. As a consequence, we have

$$\mathcal{K}_{S_1 \to S_3} = \mathcal{K}_{S_2 \to S_3} \mathcal{K}_{S_1 \to S_2}. \tag{2}$$

The above equation is extremely convenient, and so we will insist on it. By contrast, if we take $S_1 = S_3 = \{A\}$ and $S_2 = \emptyset$, then (1) is the string $\mathbb{F} \to 0 \to \mathbb{F}$ (we say "the" since there is only one way to define the arrows), and (2) doesn't hold. The fact that (2) doesn't hold is reflected in the fact that $\mathbb{F} \to 0 \to \mathbb{F}$ is not a single bar, but the direct sum of a $(0, 0)$ bar and a $(2, 2)$ bar.

In view of the above remark we make the following definition.

Definition 1.8. Let T be a set. We say that a sequence of subsets $S_0, \ldots, S_n \subset T$ is *bar-like* if:

(1) $S_0 \cup \ldots \cup S_n = T$; and
(2) for all $0 \leq i < j < k \leq n$ and all $t \in T$ we have $t \in S_i$ and $t \in S_k$ implies that $t \in S_j$.

Notice that in the definition above we have for any $0 \leq i < j < k \leq n$:

$$\mathcal{K}_{S_i \to S_k} = \mathcal{K}_{S_j \to S_k} \mathcal{K}_{S_i \to S_j},$$

and therefore for any $i < j$ we have

$$\mathcal{K}_{S_i \to S_j} = \mathcal{K}_{S_{j-1} \to S_j} \ldots \mathcal{K}_{S_i \to S_{i+1}}.$$

Proposition 1.9. *Let $\mathbb{F}$ be a field, and $t \geq 0$ be an integer. Let a string of $n+1$ $\mathbb{F}$-vector subspaces equal the direct sum of m bars. Then for a set, T, of cardinality m there are subsets $S_0, \ldots, S_n \subset T$ that are bar-like, and where*

(1) each element $t \in T$ corresponds to a (p_t, q_t)-bar where p_t, q_t are the unique integers satisfying $t \in S_i$ iff $p_t \leq i \leq q_t$;
(2) for each $i = 0, \ldots, n$ we have $V^i = \mathbb{F}^{S_i}$; and
(3) for each $i = 0, \ldots, n-1$ we have $\mathcal{L}^i = \mathcal{K}_{S_i \to S_{i+1}}$.

Conversely, for any finite set T and bar-like subsets $S_0, \dots, S_n \subset T$, there is a direct sum of $|T|$ bars that satisfies (1)–(3) above.

Definition 1.10. Let $\mathcal{S}_1 = (V_1^\cdot, \mathcal{L}_1^\cdot)$, $\mathcal{S}_2 = (V_2^\cdot, \mathcal{L}_2^\cdot)$ be strings of $\mathbb{F}$-vector spaces of the same length $n+1$. A *morphism* $\mathcal{S}_1 \to \mathcal{S}_2$ is a collection of maps $\mathcal{M}^i \colon V_1^i \to V_2^i$ that intertwine with the morphisms of $\mathcal{S}_1$ and $\mathcal{S}_2$ in the evident sense, i.e., for all $0 \le i \le n$, we have $\mathcal{M}^{i+1}\mathcal{L}_1^{i+1} = \mathcal{L}_2^i \mathcal{M}^i$ for all i.

Hence we can depict this morphism with a "commutative diagram":

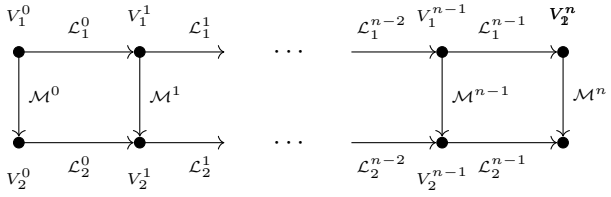

It is immediate that this morphism is an isomorphism (i.e., this morphism has an inverse morphism) iff each $\mathcal{M}^i$ is an isomorphism.

The main point of this article is the following theorem, and to give an algorithm in the general case.

Theorem 1.11. *Any string, $\mathcal{F}$, of length $n+1$ of finite dimensional vector spaces is isomorphic to a direct sum of bars. Moreover, for each $0 \le i \le j \le n$, the number of (i,j)-bars in this direct sum is independent of this direct sum.*

Concretely the above theorem can be viewed in a number of ways.

First, any string of $n+1$ $\mathbb{F}$-vector spaces $(V^\cdot, \mathcal{L}^\cdot)$ has a bar-like sequence of subsets $S_0, \dots, S_n$ of a set T and, for each $0 \le i \le n$, a bijection from S_i to a basis, X_i, of V^i, such that: for each $0 \le i \le n-1$, $\mathcal{L}^i$ is the unique linear map taking each $x \in X^i$ corresponding to an element of $s \in S_i$ to 0 if $s \in S_i \setminus S_{i+1}$, and otherwise to the $x' \in X^{i+1}$ corresponding to the element $s \in S_{i+1}$ (which therefore lies in $S_i \cap S_{i+1}$).

Second, for any string of $n+1$ $\mathbb{F}$-vector spaces $(V^\cdot, \mathcal{L}^\cdot)$ there is a direct sum of bars $(\tilde{V}^\cdot, \tilde{\mathcal{L}}^\cdot)$ and isomorphisms $\mathcal{M}^i \colon V^i \to \tilde{V}^i$, such that for each $0 \le i \le n-1$, $\mathcal{L}^i$ is given by $(\mathcal{M}^{i+1})^{-1}\tilde{\mathcal{L}}^i\mathcal{M}^i$.

2. The Main Idea in the Proof of the Barcode Theorem

Definition 2.1. Consider a string $\mathcal{F} = (V^\cdot, \mathcal{L}^\cdot)$ of vector spaces:

$$V^0 \xrightarrow{\mathcal{L}^0} V^1 \xrightarrow{\mathcal{L}^1} \cdots \xrightarrow{\mathcal{L}^{n-1}} V^n.$$

The *total space of* $\mathcal{F}$ is the pair $(\mathbf{V}, \mathcal{L})$ where

$$\mathbf{V} = V^0 \oplus \cdots \oplus V^n,$$

and $\mathcal{L} \colon \mathbf{V} \to \mathbf{V}$ is the map given by

$$\mathcal{L}(v_0, \dots, v_n) \to (0, \mathcal{L}^0 v_0, \dots, \mathcal{L}^{n-1} v_{n-1}).$$

Hence, in the definition above, $\mathcal{L}$ is nilpotent, i.e., $\mathcal{L}^n = 0$.

We next recall the algorithm to put $\mathcal{L}$ into its Jordan canonical form, which determines the barcode of $\mathcal{F} = (V^\cdot, \mathcal{L}^\cdot)$.

3. Jordan Canonical Form of a Nilpotent Linear Operator

In this section we review Jordan canonical form linear operator $\mathcal{L}\colon V \to V$, assuming that $\mathcal{L}$ is nilpotent, i.e., $\mathcal{L}^k = 0$ for some $k \in \mathbb{N}$. We will then apply this to the linear transformation $\mathcal{L}\colon \mathbf{V} \to \mathbf{V}$ in Definition 2.1.

[The general case of Jordan canonical form is not much harder, but we won't need it in this article.]

[Abstractly, the existence of Jordan canonical form can be viewed as a special case of the primary decomposition of a module over the PID $\mathbb{F}[x]$; see Appendix A.]

[This article shows the usefulness of Jordan canonical form in certain "applied settings," despite the fact that "almost all" matrices are diagonalizable.[1]]

So let $\mathcal{L}\colon V \to V$ be a linear transformation of an $\mathbb{F}$-vector space V. Further assume that $\mathcal{L}$ is nilpotent, i.e., $\mathcal{L}^k = 0$ for some $k \in \mathbb{N}$, and fix k to be the smallest such integer. It follows that all eigenvalues of $\mathcal{L}$ are 0; since $0 \in \mathbb{F}$ whether or not $\mathbb{F}$ is algebraically closed, we can put $\mathcal{L}$ into Jordan canonical form (whether or not $\mathbb{F}$ is algebraically closed). Let us review the algorithm.

A *Jordan chain*[2] *of length k generated by w of $\mathcal{L}$* is any sequence

$$w, \mathcal{L}w, \dots, \mathcal{L}^{k-1}w \tag{3}$$

such that $w \in V$ and $\mathcal{L}^{k-1}w \neq 0$. Then it is almost immediate that these elements are linearly independent: indeed, if $\alpha_0 w + \alpha_1 \mathcal{L}w + \cdots + \alpha_{k-1}\mathcal{L}^{k-1}w = 0$ for $\alpha_i \in \mathbb{F}$, and some $\alpha_i \neq 0$, then for the smallest i with $\alpha_i \neq 0$ we apply $\mathcal{L}^{k-1-i}$ to both sides of the equation and conclude that $\alpha_i \mathcal{L}^{k-1}w = 0$, which is impossible.

Next note that for a Jordan chain (3), if we restrict $\mathcal{L}$ to the subspace, V', of V spanned by $\mathcal{L}^{k-1}w, \mathcal{L}^{k-2}w, \dots \mathcal{L}w, w$ (in this order), then $\mathcal{L}$ is identified with the matrix $J \in \mathbb{F}^{n\times n}$ acting on column vectors (i.e., acting to the left of column vectors) where $J = J_k(0) \in \mathbb{F}^{k\times k}$ is the standard $k \times k$ Jordan block matrix for the eigenvalue λ, i.e.,

$$J_k(\lambda) \stackrel{\text{def}}{=} \begin{bmatrix} \lambda & 1 & & & \\ & \lambda & 1 & & \\ & & \ddots & \ddots & \\ & & & \lambda & 1 \\ & & & & \lambda \end{bmatrix} \tag{4}$$

(where a blank space implies a 0).

It follows that to write $\mathcal{L}$ in Jordan canonical form is to find $w_1, \dots, w_s \in V$ and $k_1, \dots, k_s \in \mathbb{N}$ such that (1) for each i, w_i generates a Jordan chain of length k_i,

---

[1] A matrix $M \in \mathbb{F}^{n\times n}$ (i.e., an $n \times n$ matrix with entries in $\mathbb{F}$), M, is necessarily diagonalizable (over $\overline{\mathbb{F}}$, the algebraic closure of $\mathbb{F}$) when its characteristic polynomial $p_M(x) = \det(Ix - M)$ has n distinct roots, i.e., its discriminant (i.e., the resultant of p_M and p'_M) is nonzero. [The 2×2 all zeros matrix is diagonalizable, and has characteristic polynomial x^2, so this condition is not necessary.] Hence there is a polynomial $Q = Q(M)$, of the entries of M, such that $Q(M) \neq 0$ implies that M is diagonalizable. Since $Q(M)$ is not identically zero (it is nonzero on a diagonal matrix with distinct diagonal elements in $\overline{\mathbb{F}}$), it follows that Q is not the zero polynomial. It follows that if $\mathbb{F} = \mathbb{R}, \mathbb{C}$, the set of non-diagonalizable matrices in $\mathbb{F}^{n\times n}$ is of measure 0. More generally, for any field $\mathbb{F}$, the set of non-diagonalizable matrices in $\overline{\mathbb{F}}^{n\times n}$ lies in a proper, Zariski closed subset, and therefore is "exceptional" in various senses (assuming $\mathbb{F}$ is infinite or sufficiently large) that we will not bother to specify.

[2] We would add "with respect to the eigenvalue 0" for a general $\mathcal{L}$, not assumed to be nilpotent, and the chain would be $w, (\mathcal{L} - \lambda)w, \dots, (\mathcal{L} - \lambda)^{k-1}w$ for an eigenvalue λ.

(2) $k_1 + \cdots + k_s = n$, and (3) the union of

$$\bigcup_{i=1}^{s} \{w_i, \mathcal{L}w_i, \ldots, \mathcal{L}^{k_i-1} w_i\}$$

is a basis for V.

Definition 3.1. By a *Jordan basis* for a nilpotent linear operator $\mathcal{L}\colon V \to V$ we mean any pair of sequences $w_1, \ldots, w_s \in V$ and $k_1, \ldots, k_s \in \mathbb{N}$ satisfying (1)–(3) in the previous paragraph.

Example 3.2. Let

$$L = \begin{bmatrix} 0 & 1 & \\ 0 & 0 & \\ & & 0 \end{bmatrix}$$

(where a blank space implies a 0). Then L is a block diagonal matrix with a $J_2(0)$ block and a $J_1(0)$ block. If $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ are the standard basis vectors, then $\mathbf{e}_1, \mathbf{e}_3$ are eigenvectors (with L acting to the left of column vectors). Also, $\mathbf{e}_2, \mathcal{L}\mathbf{e}_2$ and $\mathbf{e}_3$ are two Jordan chains, where $\mathcal{L}$ is the operator on $\mathbb{F}^3$ expressed as column vectors via the basis $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$.

We say that a Jordan chain in (3) is *maximal* if $\mathcal{L}^k w = 0$ and if there is no w' such that $\mathcal{L}w' = w$. It is easy to see — and helpful for intuition — to note that any Jordan basis for $\mathcal{L}$ must consist of maximal Jordan chains.

It is now easy to give an algorithm for finding a Jordan basis for a nilpotent operator $\mathcal{L}\colon V \to V$. The point is that you want to find the longest Jordan chains of the Jordan basis first.

So let $k \in \mathbb{N}$ be the largest integer with $\mathcal{L}^{k-1}$ nonzero. Let $u_1, \ldots, u_r$ be a basis for the image of $\mathcal{L}^{k-1}$; then, by definition, there exist $w_1, \ldots, w_r$ such that $\mathcal{L}^{k-1} w_i = u_i$ for all i. We easily see that each w_i generates a (maximal) chain of length k, and the vectors $B_k = \{\mathcal{L}^j w_i\}$ with i, j ranging over $1 \le i \le s$ and $0 \le j \le k-1$ are linearly independent, similarly to the above argument. Since $\mathcal{L}^k = 0$, we see that all chains are of length at most k, and that any chain of length k is generated by a w that is a linear combination of $w_1, \ldots, w_r$ above (since $\mathcal{L}^{k-1} w$ must be a linear combination of $u_1, \ldots, u_r$ above).

We next find the maximal chains of length $k-1$ whose elements are linearly independent from B_k above: consider the image of $\mathcal{L}^{k-2}$, which clearly contains $\mathcal{L}^{k-2} w_i$ and $\mathcal{L}^{k-1} w_i$ for all $1 \le i \le r$; since these $2r$ vectors are linearly independent, we can choose $u'_1, \ldots, u'_{r'-1}$ to complete these vectors to a basis for the image of $\mathcal{L}^{k-2}$; we then choose w'_i such that $\mathcal{L}^{k-2} w'_i = u'_i$ for all i. We easily show that

$$B_k = \{\mathcal{L}^j w_i \mid 1 \le i \le r,\ 0 \le j \le k-1\}$$

and

$$B_{k-1} = \{\mathcal{L}^j w'_i \mid 1 \le i \le r',\ 0 \le j \le k-2\}$$

are disjoint subsets whose union is linearly independent.

Next we repeat this step to find vectors B_{k-2}, independent of $B_k \cup B_{k-1}$ and coming from chains of length $k-2$. We similarly find vectors $B_{k-3}, B_{k-4}, \ldots B_1$. Since B_1 is a basis of vectors in the image of $\mathcal{L}^0 = \mathrm{id}_V$, i.e., all of V, that completes $B_2 \cup \ldots \cup B_k$, we have that $B_1 \cup \ldots \cup B_k$ is a basis for all of V. The union over all i of the chains of length i arising from the B_i is therefore a Jordan basis.

Example 3.3. In Example 3.2, the image of $\mathcal{L}$ on $\mathbb{F}^3$ is the span of $\mathbf{e}_1$; since $\mathcal{L}\mathbf{e}_2 = \mathbf{e}_1$, this gives us the chain $\mathbf{e}_2, \mathcal{L}\mathbf{e}_2 = \mathbf{e}_1$, so $B_2 = \{\mathbf{e}_2, \mathbf{e}_1\}$. Moreover, B_2 is determined up to scalar multiple. Then B_1 consists of a single element, which may be any vector $\gamma_1\mathbf{e}_1 + \gamma_3\mathbf{e}_3$ with $\gamma_1, \gamma_3 \in \mathbb{F}$ with $\gamma_3 \neq 0$. Notice that if we started by looking for Jordan chains of length 1, i.e., eigenvectors, there is 2-dimensional possible space. If we take $\mathbf{e}_1 + \mathbf{e}_3$ and $\mathbf{e}_3$ as such a basis, there is no way to extend either of these "backwards" to make one of them a chain of length 2. This is why we start by finding the longest Jordan chains and then find successively shorter ones.

Remark 3.4. Despite the problem arising by starting with shorter chains and extending them backwards, identified in Example 3.3, one can still roughly do this, using one trick. Namely, in Section 3.1 of [HJ85] one first finds a basis with respect to which $\mathcal{L}$ is an upper triangular matrix, using the Schur decomposition; hence the diagonal is all 0's. From there one does an inductive argument, reducing the $n \times n$ case (assuming the matrix is upper triangular with 0's on the diagonal) to the $(n-1)\times(n-1)$ case (see Subsection 3.1.5 there). So provided that $\mathcal{L}$ is already written in upper triangular form, one can start with short chains and progressively look for longer (or new) ones.

Remark 3.5. Note that the total space $(\mathbf{V}, \mathcal{L})$ of a string of vector spaces is already an upper triangular matrix: indeed, choose arbitrary bases $B^0, \ldots, B^n$ for the respective vector spaces $V^0, \ldots, V^n$; then $\mathcal{L}$ with respect to $B^n, \ldots, B^0$ is a block matrix whose only nonzero blocks are those just above (or to the right of) the main diagonal. Hence in Remark 3.4 we can skip the Schur decomposition step.

Remark 3.6. There are likely very many algorithms to find a barcode decomposition of a string of vector spaces, and I currently (December 2024) don't know what is known here for a general string and/or strings arising in homology. However, given the previous two remarks, I'm guessing there are a lot of options, depending on the precise features of the string of vector spaces.

4. Proof of the Barcode Theorem

Let notation be as in Definition 2.1. Since $\mathcal{L}$ is nilpotent, we will use the algorithm in the previous section. So let $k \in \mathbb{N}$ be the smallest natural number with $\mathcal{L}^k = 0$, and let $\mathbf{u}_1, \ldots, \mathbf{u}_r$ be a basis for the image of $\mathcal{L}^{k-1}$; let $\mathbf{w}_1, \ldots, \mathbf{w}_r$ be such that $\mathbf{u}_i = \mathcal{L}^{k-1}\mathbf{w}_i$.

Say that a nonzero element $\mathbf{u} \in \mathbf{V}$ is *purely of degree* d if $\mathbf{u} = (u^0, \ldots, u^n)$ and $u^i = 0$ for $i \neq d$. Clearly any nonzero element of $\mathbf{V}$ can be uniquely written as a sum of elements purely of degrees $d_1 < d_2 < \ldots < d_t$ with $0 \leq d_1 < \ldots < d_t \leq n$; we call each such summand a *pure component of* $\mathbf{u}$. Clearly if $\mathbf{u}_i$ is in the image of $\mathcal{L}^{k-1}$, then so is each pure component of $\mathbf{u}_i$. It follows from the "basis exchange theorem" that we can replace the basis $\mathbf{u}_1, \ldots, \mathbf{u}_r$ of the image of $\mathcal{L}^{k-1}$ with one where each $\mathbf{u}_j$ is purely of some degree. Then if $\mathcal{L}^{k-1}\mathbf{w}_j = \mathbf{u}_j$ and $\mathbf{u}_j$ is purely of degree d, then the same holds with $\mathbf{w}_j$ replaced by its pure component of degree $d-(k-1)$.

This gives us a set B_k as in the previous section, which is the union of chains of length k each of which is generated by a w_i that is purely ofsome degree. We easily see that for any d, the dimension of the image of V^d in V^{d+k-1} is precisely the number of w_i that are purely of degree d; hence the number of $\mathbf{w}_1, \ldots, \mathbf{w}_r$ that

are purely of some degree d depends only on d and not the particular choice of $\mathbf{w}_1, \ldots, \mathbf{w}_r$.

In this way we similarly generate $B_{k-1}, B_{k-2}, \ldots, B_1$, which give a Jordan basis for $\mathcal{L}$. Moreover the number of Jordan chains of a given length k' in $B_{k'}$ beginning in an element purely in any given degree is independent of the choice of elements in $B_{k'}$. But the decomposition of $\mathbf{V}$ into Jordan chains generated by elements of $\mathbf{V}$, each of which is purely of some degree, is clearly the same thing as a barcode decomposition.

Appendix A. Modules over PID's

[I've really written this section mostly to jog my memory, since I learned most of these theorems in a class in the early 1980's, and haven't really used them much since...]

It is well known that any finitely generated abelian group is a sum of groups, each of which is of the form $\mathbb{Z}/m\mathbb{Z}$ for integers $m \geq 0$ (so either $m = 0$, which gives $\mathbb{Z}$, or $m \neq 0$, in which case $\mathbb{Z}/m\mathbb{Z}$ is a finite, cyclic group, and the case $m = 1$ can be omitted). Furthermore, when $m \geq 2$, $\mathbb{Z}/m\mathbb{Z}$ can be written as a sum of its *primary* parts, i.e., as a direct sum over $i \in [r]$ of $\mathbb{Z}/p_i^{n_i}$, where $m = p_1^{n_1} \cdots p_r^{n_r}$ is the prime factorization of m. It turns out that the usual proof of the previous paragraph generalizes to modules over any PID ($\mathbb{Z}$ is a PID), and an abelian group is the same thing as a $\mathbb{Z}$-module). This structure theorem gives a proof — although not necessarily the best algorithm — for determining the Jordan canonical form of a square matrix.

Here we will outline some of the main these ideas needed. I will use part of Lang's textbook, *Algebra*, [Lan02], specifically Section III.7, "Modules over Principal Rings."

I assume that you know what is meant by a commutative ring, R (we assume $1 \neq 0$), and an R-module (when R is a field, an R-module is just an R vector space). An *ideal* in R is a subset of R that is also an R-module. A ring, R, is called an *integrity ring* or *ring of integrity*[3] if it has no zero divisors; e.g., $R = \mathbb{Z}/6\mathbb{Z}$ is not of integrity, since $2 \cdot 3 = 0$ but $2, 3 \neq 0$; similarly, for $R = \mathbb{F}[x, y]/(xy)$, $xy = 0$ but $x, y \neq 0$. A PID[4] is a commutative ring, R, of integrity such that every ideal of R is *principal*, i.e., of the form $(a) \stackrel{\text{def}}{=} aR$ for some $a \in R$.

The simplest (interesting) examples of PID's are $\mathbb{Z}$ and $\mathbb{F}[x]$ for a field, $\mathbb{F}$. In algebraic geometry it might be useful to notice that any localization of a PID is again a PID. It is a famous result that if R is the ring of integers over a finite extension of the rationals, then R is a PID iff R is a unique factorization domain, and there are only finitely many such R.[5]

[3]This term is truer to the German *Integritätsring* or the French *anneau intègre*; Lang [Lan02] uses the term *entire rings*; a common English term for these rings is *integral domains*, which seems like a needlessly confusing mistranslation of the likely original German (see https://math.stackexchange.com/questions/45945/where-does-the-term-integral-domain-come-from); see also pages 91-92 II.2 of [Lan02]. In particular, being an "integral domain" has nothing to do with R being integrally closed over its field of fractions...

[4]Principal integral domain, where "integral domain" has the problematic name as explained in the previous footnote.

[5]Also, it is not hard to prove Fermat's last theorem regarding $x^n + y^n = z^n$ for the finitely many n such that $\mathbb{Z}[\zeta_n]$ is a PID (or unique factorization domain), where ζ_n is a primitive n-th root of unity).

[For intuition on matters below, it may be useful to note that the set of (multiplicative) units in $\mathbb{Z}$ is finite, but is only finitely generated in a general number field. Also, the set of units in $\mathbb{F}[x]$ is $\mathbb{F}^\times = \mathbb{F} \setminus \{0\}$.]

We now sketch the proof of the main structure theorems for Abelian groups, which are the same things as $\mathbb{Z}$-modules. We prove this in the general context of R-modules for any PID, R. We begin by following Lang [Lan02] (Section III.7, Modules over Principal Rings).

Lemma A.1. *Let R be a PID, and M be a free R-module with basis $e_1, \ldots, e_n$ (i.e., each element of R can be written uniquely as $r_1e_1 + \ldots + r_ne_n$; you can also think of e_i as the i-th standard basis vector of R^n). Let $F \subset M$ be any R-submodule. Then there are $f_1, \ldots, f_n \in F$ such that F is a free R-module generated by the nonzero f_i. (We allow f_i to be 0 for convenience of sketching the proof.)*

Proof. For each $i = 1, \ldots, n$, consider the $a \in R$ such that

$$\text{there exist } b_1, \ldots, b_{i-1} \in R \text{ such that } b_1e_1 + \ldots + b_{i-1}e_{i-1} + ae_i \in F.$$

The set of such a is easily seen to be an ideal of R, and hence equal to (a_i) for some $a_i \in R$. If $a_i \neq 0$, choose any b_{ji} such that

$$f_i = b_{1i}e_1 + \ldots + b_{i-1,i}e_i + a_ie_i \in F;$$

if $a_i = 0$ we set $f_i = 0$. We easily verify — using the triangularity of matrix relating the f_i to the e_i where $a_i \neq 0$ — that these are the f_i we seek. □

From here we will diverge from the proof in Lang's book. (The proof below is essentially the proof I first learned, following a textbook of Jacobson on algebra...)

If $L \in R^{m\times n}$, i.e., L is an $m \times n$ matrix with entries in R, an *elementary row operation on L* is any way of

(1) choosing $\alpha, \beta, \gamma, \delta \in R$ such that
$$\begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix}$$
is invertible (which, using some facts from below, is easily seen to be equivalent to $\alpha\delta - \beta\gamma$ being a (multiplicative) unit it R);
(2) choosing $i, j \in [m]$ with $i \neq j$;
(3) then replacing the i-th and j-th rows of L with, respectively, α times the i-th row plus β times the j-th row, and γ times the i-th row plus δ times the j-th row.

We make the similar definition for *elementary column operations*.

[In order to accommodate the situation where $m = 1$, i.e., L has only one row, we could allow ourselves to multiply any row by any unit of R. (However, we won't really be interested in row operations when $m = 1$.)]

Hence, each elementary row operation on a matrix in the usual sense, when $R = \mathbb{F}$ is a field, is an elementary row operation in the above sense; moreover, when $R = \mathbb{F}$ is a field, every elementary row operation in the above sense can be achieved by at most four of the elementary row operations in the usual sense (left to the reader).

Lemma A.2 (Smith normal form)**.** *Let R be a PID, and let $L \in R^{n\times n}$, i.e., L is an $n \times n$ matrix with entries in R. Then one can perform elementary row and column operations on L to obtain a diagonal matrix L' whose diagonal elements*

are $b_1, \dots, b_n$, with $(b_n) \subset (b_{n-1}) \subset \cdots \subset (b_1)$ (we also write this as $b_1|b_2|\dots|b_n$, writing $a|b$ to mean $(b) \subset (a)$).

We remark that Smith normal form is more typically defined for general matrices rather than only square matrices, i.e., $L \in R^{m\times n}$ with $m \neq n$; we personally prefer to think of the square case; it always suffices to consider only the square case, at the expense of adding 0's to a non-square matrix to make it square (see also its application below).

Below we sketch a proof of the above lemma. Before doing so, we make a number of remarks showing that PID's work similarly to $\mathbb{Z}$ (itself a PID) in many ways, such as: (I don't think we really need (7,8) below, but they are a lot of fun...)

(1) If $(c_1) \subset (c_2) \subset \cdots$ is an increasing infinite sequence of ideals of R (all ideals are principal), then $c_i = c_{i+1} = \dots$ for some i (proof: the union of these ideals is an ideal, therefore equal to (c) with $c \in (c_i)$ for some i, and the result follows). More generally, a ring R is called *Noetherian* if any increasing sequence of ideals in R eventually stabilizes;[6] see Lang's textbook [Lan02], Chapter X (Noetherian Rings and Modules), Section 1 (Basic Criteria).)
(2) For any $a, b \in R$, $(a)(b) = (ab)$ (proof: very easy).
(3) If for $a, b \in R$ are nonzero and $(a) = (b)$, then a, b differ by multiplicative units in R (proof: by definition $b = a\alpha$ and $a = b\beta$ for some $\alpha, \beta \in R$, and hence $a = a\alpha\beta$, so $a(1 - \alpha\beta) = 0$, so, by the integrity of R, $1 - \alpha\beta = 0$ so α, β have multiplicative inverses — each other — in R).
(4) For any $a, b \in R$ we write $a|b$ if $(b) \subset (a)$; in this case we have $a = bc$ for some $c \in R$ (by definition, since $b \in (b) \subset (a)$, so $b \in (a)$ so, by definition $b = ac$ for some $c \in R$).
(5) We say that $a, b \in R$ are *relatively prime* if $(a, b) = R$, or, equivalently, $a\alpha + b\beta = 1$ for some $\alpha, \beta \in R$. In this case $(ab) \subset (a) \cap (b)$ holds with equality (proof: any element of $(a) \cap (b)$ is of the form $aA = bB$, and so $aA\beta = bB\beta = (1 - a\alpha)B$, and hence $B = aA\beta + a\alpha B$ and so $a|B$ so $B = a\gamma \in (a)$ and so $bB \in (ab)$).
(6) For any $a, b \in R$, one can define their GCD (greatest common divisor) as any $c \in R$ such that $(c) = (a, b)$ $(c) = (a, b)$ (i.e., c generates the ideal generated by a and b), and c is uniquely defined up to a unit of R. We then have $(a) \subset (c)$ so $a = cd_1$ and similarly $b = cd_2$ for $d_1, d_2 \in R$, and hence $ab = cL$, where $L = cd_1d_2$ (so $c|L$).
(7) The L in the previous item can rightfully be called the LCM (least common multiple) of a, b since $(a) \cap (b) = (L)$ (proof: since $(c) = (a, b) = c(d_1, d_2)$, we have $c \in c(d_1, d_2)$ so $1 \in (d_1, d_2)$ so d_1, d_2 are relatively prime. Hence $(d_1) \cap (d_2) = (d_1d_2)$, and hence $(a) \cap (b) = c((d_1) \cap (d_2)) = (cd_1d_2) = (L)$).
(8) Hence $(a)(b) = (a, b)((a) \cap (b))$, where the right-hand-side is the product of ideals corresonding to the GCD and LCM of (a) and (b).
(9) Say that $(a, b) = (c)$. Then $c = a\alpha + b\beta$ for some $\alpha, \beta \in R$; since $(a) \subset (a, b) = (c)$, $c = a\gamma$ for some $\gamma \in R$, and similarly $c = b\delta$. It follows that

$$c = c(\gamma\alpha + \delta\beta),$$

[6] Much of algebraic geometry can be written more concisely when working with Noetherian rings, and Hartshorne's celebrated textbook exploits this. For example, any ideal in a Noetherian ring is finitely generated. Of course, if you really need to work with spaces locally modeled by the spectra of rings like $R = \mathbb{F}[x_1, x_2, \dots]$, then you might think otherwise...

and so by the integrity of R, $\gamma\alpha + \delta\beta = 1$. Hence

$$\det \begin{bmatrix} \alpha & \beta \\ -\delta & \gamma \end{bmatrix} = 1,$$

and

$$\begin{bmatrix} \alpha & \beta \\ -\delta & \gamma \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} c \\ -\delta a + \gamma b \end{bmatrix} = \begin{bmatrix} c \\ -\delta\gamma c + \gamma\delta c \end{bmatrix} = \begin{bmatrix} c \\ 0 \end{bmatrix}$$

It follows that the column vector $[a\ b]$ differs from $[c\ 0]$ by an elementary column operation.

(10) By induction on the preceding remark, if $L \in R^{1\times n}$, i.e., L is a matrix with a single row and n columns, then there are elementary column operations on L taking it to a matrix $[c\ 0\ 0\ \dots\ 0]$, where c is the GCD of the entries of L.

Proof of Lemma A.2. First we describe a procedure to find a matrix equivalent to L whose top left entry is c, and the other entries in the first row and column of L are all 0's: to do so, according to (2) above, we may perform column operations on L to yield a matrix L' whose first row is of the form $[c'\ 0\ 0\ \dots\ 0]$. If all entries in the first column of L' are divisible by c', we can clear out all entries under the top c' in the first row, which finishes the procedure (with $c = c'$). If not, then setting c'' to be the GCD of the entries of the first row of L', we have $(c') \subset (c'')$ with strict containment; we then perform the analogous row operations on L' to get a matrix L'' whose first column is the transpose of $[c''\ 0\ \dots\ 0]$. If c'' divides all the entries of the first column of L'' we are done; otherwise we go back to column operations on L'' to get its first row to be $[c'''\ 0\ 0\ \dots\ 0]$ where $(c'') \subset (c''')$ with strict containment. Continuing in this way, we get strict containments $(c') \subset (c'') \subset (c''') \subset \cdots$ which must terminate by (3) above.

Having finished with the first step, we inductively perform elementary row and column operations on rows and columns numbers 2 through n, so that the we get an equivalent matrix all of whose first two rows and columns are 0 except on the diagonal. By induction we can continue until the original matrix L is equivalent to a diagonal matrix L_{diag}, whose diagonal elements are $c_1, c_2, \dots, c_n$.

Now we wish to bring L_{diag} to an equivalent diagonal matrix whose diagonal entries are $b_1, \dots, b_n$ with $b_i | b_{i+1}$ for all i.

First, we show that the top left 2×2 block of L_{diag} can be brought to a diagonal matrix whose first diagonal entry is the GCD of c_1, c_2; the second diagonal entry will also (or therefore) by divisible by this GCD. (We leave this to the reader.) We then proceed similarly with the 2×2 minor built from rows and columns 1 and 3, and then columns 1 and 4 and so on; this gives an equivalent diagonal matrix whose diagonal elements are $b_1, c_2', c_3', \dots, c_n'$ where b_1 is the GCD of $c_1, \dots, c_n$ and b_1 divides $c_2', \dots, c_n'$. Proceeding similarly with the lower right $(n-1) \times (n-1)$ submatrix, and so on, we eventually get an equivalent diagonal matrix whose entries are $b_1, \dots, b_n$ with $b_i | b_{i+1}$ for all i. □

(I'm not claiming that the algorithm given in this proof is particularly efficient...)

Lemma A.3. *If R is a PID, and M is an R-module generated by n generators, then there are $b_1, \dots, b_n$ such that*

$$R = (b_1) \oplus (b_2) \oplus \cdots \oplus (b_n)$$

and $(b_n) \subset (b_{n-1}) \subset \cdots \subset (b_1)$ *(one also typically writes this is* $b_1|b_2|\dots|b_n$*, where* $a|b$ *is its usual meaning in a PID, namely* $(b) \subset (a)$*, which coincides with its usual meaning in* $\mathbb{Z}$ *or* $\mathbb{F}[x]$*).*

Proof. Let $e_1, \dots, e_n$ generate M. Let $S \subset R^n$ be the subset of $\mathbf{s} = (s_1, \dots, s_n) \in R^n$ such that

$$s_1 e_1 + \cdots + s_n e_n \in R$$

(hence S describes "all relations" among $e_1, \dots, e_n$ in R). We easily see that S is an R-sub-module of R^n; by Lemma A.1, S is free and generated by $f_1, \dots, f_m \in S$ where $m \le n$. Let L be the matrix $n \times n$ matrix with entries in R, whose first m columns are $f_1, \dots, f_m$, and whose remaining ones are all 0's (if $m < n$ one alternatively one can take $L \in R^{n \times m}$ and use Smith normal form in the more general (not necessarily square) case). Then M is isomorphic to $R^n/\text{image}(L)$. If L' is obtained from any sequence of elementary row and column operations (R is a ring, so we allow multiplying a row/column only by a multiplicative unit in R), then we still have M is isomorphic to $R^n/\text{image}(L')$, since elementary column operations don't change the image of a matrix, and elementary row operations reflect a change of basis in R^n. After sufficiently many elementary column and row operations we may bring L into Smith normal form, i.e., we may write L as a diagonal matrix with elements $b_1, b_2, \dots, b_n$ with $b_i | b_{i+1}$ for all $i \in [n-1]$. This gives the desired isomorphism. □

Next we give the version of the above lemma in terms of *primary* ideals of R.

We say that an ideal (p) of R (i.e., $p \in R$) is *prime* if (p) is a maximal ideal[7] of R strictly contained in R; an ideal of R is *primary* if it is of the form $(p)^n$ for some prime p and some $n \in \mathbb{N}$.

We now wish to show that for any non-unit $c \in R$, we have

(1) $(c) = (p_1)^{n_1} \dots (p_r)^{n_r}$ for some primes $p_1, \dots, p_r$ where $(p_i) \ne (p_j)$ for $i \ne j$, and $n_1, \dots, n_r \in \mathbb{N}$;

(2) if so, then

$$R/(c) \simeq \bigoplus_{i=1}^{r} R/(p_i^{n_i}), \tag{5}$$

which we call a *primary decomposition of* $R/(c)$ (which is unique up to permuting the $p_i^{n_i}$).

For (1), let (p_1) be any maximal ideal containing (c); there must be a largest n such that $(c) \subset (p_1)^n$, for otherwise we violate the Noetherian property of R. Since $p_1^{n_1} | c$ we have $c = p_1^{n_1} c'$, where either c' is a unit (in which case we are done), or c' is a non-unit in R with $p_1 \not| c'$. If so, choose a prime (p_2) with $(c') \subset (p_2)$, and so we have $c' = p_2^{n_2} c''$ with $p_1 \not| c''$ and $p_2 \not| c''$. We continue similarly; this process has to end, since R is Noetherian. This establishes (1).

To do so, first note that for any $a, b \in R$ there are natural maps $R/(ab)$ to both $R/(a)$ and $R/(b)$, and hence a natural map

$$\phi \colon R/(ab) \to R/(a) \oplus R/(b). \tag{6}$$

We claim that if a, b are relatively prime, then ϕ is an isomorphism: indeed, the kernel of ϕ is zero, since $(a) \cap (b) = (ab)$. To show that ϕ is onto, it suffices

[7] Of course, for general rings, R, an ideal I is prime if $ab \in I$ implies either $a \in I$ or $b \in I$; in the context of a PID, any prime ideal is also a maximal ideal.

to show that both $(1,0)$ and $(0,1)$ are in the image of ϕ; since a, b are relatively prime, $\alpha a + \beta b = 1$ for some $\alpha, \beta \in R$, and hence $1 - \alpha a$ in $R/(ab)$ maps to $(1,0) \in R(a) \oplus R(b)$; hence $(1,0)$ is in the image of ϕ, and similarly so is $(0,1)$.

Since ϕ in (6) is an isomorphism when a, b are relatively prime, we apply this inductively to conclude (5).

In this way we get the following corollary to Lemma A.3.

Corollary A.4. *Let R be a PID, and let M be a finitely generated module over R. Then M is isomorphic to the direct sum of some finite number of copies of R and modules $R/(p_i^{n_i})$ where p_i is a prime of R, $n_i \in \mathbb{N}$ and i ranges over a (possibly empty) set I.*

Remark A.5. The proof in Lang's textbook [Lan02] goes "the other way," first directly proving Corollary A.4 on its own (first splitting M as a direct sum of a purely torsion part and then a torsion free part), and then deducing Lemma A.3 as a consequence.

Remark A.6. If $\mathcal{L}\colon V \to V$ is a linear operator on an $\mathbb{F}$-vector space, then $\mathcal{L}$ gives rise to an action of $\mathbb{F}[x]$ on V where $p(x)$ acts on V as the linear operator $p(L)$. If $\mathbb{F}$ is algebraically closed, then every prime ideal in $\mathbb{F}[x]$ is necessarily of the form $(x - \lambda)$ for some $\lambda \in \mathbb{F}$. Each primary idea of $\mathbb{F}[x]$ is therefore of the form $\mathbb{F}[x]/(x-\lambda)^k$ for some $k \in \mathbb{N}$, which corresponds to a $k \times k$ Jordan block with eigenvalue λ, i.e., to $J_k(\lambda)$ in the notation (4).

Remark A.7. Consider Remark A.6 when $\mathbb{F}$ is not algebraically closed. Then every prime ideal is of the form $\mathbb{F}[x]/(p(x))$ where $p = p(x)$ is a monic, irreducible polynomial over $\mathbb{F}$. So, for example, when $\mathbb{F} = \mathbb{R}$, we get ideals where $p(x)$ is a degree two polynomial with imaginary, complex conjugate roots, as well as $p(x) = x - \alpha$ with $\alpha \in \mathbb{R}$.

References

[CZCG04] Gunnar Carlsson, Afra Zomorodian, Anne Collins, and Leonidas Guibas, *Persistence barcodes for shapes*, Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing (New York, NY, USA), SGP '04, Association for Computing Machinery, 2004, p. 124–135.

[ELZ00] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian, *Topological persistence and simplification*, 41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000), IEEE Comput. Soc. Press, Los Alamitos, CA, 2000, pp. 454–463. MR 1931842

[ELZ02] ———, *Topological persistence and simplification*, vol. 28, 2002, Discrete and computational geometry and graph drawing (Columbia, SC, 2001), pp. 511–533. MR 1949898

[Fri05] Joel Friedman, *Cohomology of Grothendieck topologies and lower bounds in Boolean complexity*, preprint, 70 pages. Available at http://arxiv.org/abs/cs/0512008.

[HJ85] Roger A. Horn and Charles R. Johnson, *Matrix analysis*, Cambridge University Press, Cambridge, 1985. MR 832183

[Lan02] Serge Lang, *Algebra*, third ed., Graduate Texts in Mathematics, vol. 211, Springer-Verlag, New York, 2002. MR 1878556 (2003e:00003)

[sga72] *Théorie des topos et cohomologie étale des schémas. Tome 1: Théorie des topos*, Springer-Verlag, Berlin, 1972, Séminaire de Géométrie Algébrique du Bois-Marie 1963–1964 (SGA 4), Dirigé par M. Artin, A. Grothendieck, et J. L. Verdier. Avec la collaboration de N. Bourbaki, P. Deligne et B. Saint-Donat, Lecture Notes in Mathematics, Vol. 269. MR 50 #7130

Department of Computer Science, University of British Columbia, Vancouver, BC V6T 1Z4, CANADA.
Email address: jf@cs.ubc.ca
URL: http://www.cs.ubc.ca/~jf