19

Linearly generated sequences and applications

In this chapter, we develop some of the theory of linearly generated sequences. As an application, we develop an efficient algorithm for solving sparse systems of linear equations, such as those that arise in the subexponential-time algorithms for discrete logarithms and factoring in Chapter 16. These topics illustrate the beautiful interplay between the arithmetic of polynomials, linear algebra, and the use of randomization in the design of algorithms.

19.1 Basic definitions and properties

Let F be a field, let V be an F-vector space, and consider an infinite sequence

$$S = (\alpha_0, \alpha_1, \alpha_2, \ldots),$$

where $\alpha_i \in V$ for $i = 0, 1, 2 \ldots$. We say that S is **linearly generated (over F)** if there exist scalars $a_0, \ldots, a_{k-1} \in F$ such that the following recurrence relation holds:

$$\alpha_{k+i} = \sum_{j=0}^{k-1} a_j \alpha_{j+i} \quad (\text{for } i = 0, 1, 2, \ldots).$$

In this case, all of the elements of the sequence S are determined by the initial segment $\alpha_0, \ldots, \alpha_{k-1}$, together with the coefficients $a_0, \ldots, a_{k-1}$ defining the recurrence relation.

The general problem we consider is this: how to determine the coefficients defining such a recurrence relation, given a sufficiently long initial segment of S. To study this problem, it turns out to be very useful to rephrase the problem slightly. Let $g \in F[\mathtt{X}]$ be a polynomial of degree, say, k, and write

$g = \sum_{j=0}^{k} g_j \mathtt{X}^j$. Next, define

$$g \star S := \sum_{j=0}^{k} g_j \alpha_j.$$

Then it is clear that S is linearly generated if and only if there exists a non-zero polynomial g such that

$$(\mathtt{X}^i g) \star S = 0 \quad (\text{for } i = 0, 1, 2, \ldots). \tag{19.1}$$

Indeed, if there is such a non-zero polynomial g, then we can take

$$a_0 := -(g_0/g_k),\ a_1 := -(g_1/g_k),\ \ldots,\ a_{k-1} := -(g_{k-1}/g_k)$$

as coefficients defining the recurrence relation for S. We call a polynomial g satisfying (19.1) a **generating polynomial** for S. The sequence S will in general have many generating polynomials. Note that the zero polynomial is technically considered a generating polynomial, but is not a very interesting one.

Let $G(S)$ be the set of all generating polynomials for S.

Theorem 19.1. *$G(S)$ is an ideal of $F[\mathtt{X}]$.*

Proof. First, note that for any two polynomials f, g, we have $(f + g) \star S = (f \star S) + (g \star S)$ — this is clear from the definitions. It is also clear that for any $c \in F$ and $f \in F[\mathtt{X}]$, we have $(cf) \star S = c \cdot (f \star S)$. From these two observations, it is immediately clear that $G(S)$ is closed under addition and scalar multiplication. It is also clear from the definition that $G(S)$ is closed under multiplication by $\mathtt{X}$; indeed, if $(\mathtt{X}^i f) \star S = 0$ for all $i \geq 0$, then certainly, $(\mathtt{X}^i(\mathtt{X}f)) \star S = (\mathtt{X}^{i+1} f) \star S = 0$ for all $i \geq 0$. But any non-empty subset of $F[\mathtt{X}]$ that is closed under addition, multiplication by elements of F, and multiplication by $\mathtt{X}$ is an ideal of $F[\mathtt{X}]$ (see Exercise 9.27). □

Since all ideals of $F[\mathtt{X}]$ are principal, it follows that $G(S)$ is the ideal of $F[\mathtt{X}]$ generated by some polynomial $\phi \in F[\mathtt{X}]$ — we can make this polynomial unique by choosing the monic associate (if it is non-zero), and we call this polynomial the **minimal polynomial of** S. Note that S is linearly generated if and only if $\phi \neq 0$.

We can now restate our main objective as follows: given a sufficiently long initial segment of a linearly generated sequence, determine its minimal polynomial.

***Example* 19.1.** Of course, one can always define a linearly generated sequence by simply choosing an initial sequence $\alpha_0, \alpha_1, \ldots, \alpha_{k-1}$, along with

the coefficients $g_0, \ldots, g_{k-1}$ of a generating polynomial $g := g_0 + g_1\mathtt{X} + \cdots + g_{k-1}\mathtt{X}^{k-1} + \mathtt{X}^k$. One can enumerate as many elements of the sequence as one wants by using storage for k elements of V, along with storage for the coefficients of g, as follows:

$(\beta_0, \ldots, \beta_{k-1}) \leftarrow (\alpha_0, \ldots, \alpha_{k-1})$
repeat
    output β_0
    $\beta' \leftarrow -\sum_{j=0}^{k-1} g_j \beta_j$
    $(\beta_0, \ldots, \beta_{k-1}) \leftarrow (\beta_1, \ldots, \beta_{k-1}, \beta')$
forever

Because of the structure of the above algorithm, linearly generated sequences are sometimes also called **shift register sequences**. Also observe that if F is a finite field, and V is finite dimensional, the value stored in the "register" $(\beta_0, \ldots, \beta_{k-1})$ must repeat at some point, from which it follows that the linearly generated sequence must be ultimately periodic (see definitions above Exercise 4.8). □

***Example* 19.2.** Linearly generated sequences can also arise in a natural way, as this example and the next illustrate. Let $E := F[\mathtt{X}]/(f)$, where $f \in F[\mathtt{X}]$ is a monic polynomial of degree $\ell > 0$, and let α be an element of E. Consider the sequence $S := (1, \alpha, \alpha^2, \cdots)$ of powers of α. For any polynomial $g = \sum_{j=0}^{k} g_j \mathtt{X}^j \in F[\mathtt{X}]$, we have

$$g \star S = \sum_{j=0}^{k} g_j \alpha^j = g(\alpha).$$

Now, if $g(\alpha) = 0$, then clearly $(\mathtt{X}^i g) \star S = \alpha^i g(\alpha) = 0$ for all $i \geq 0$. Conversely, if $(\mathtt{X}^i g) \star S = 0$ for all $i \geq 0$, then in particular, $g(\alpha) = 0$. Thus, g is a generating polynomial for S if and only if $g(\alpha) = 0$. It follows that the minimal polynomial ϕ of S is the same as the minimal polynomial of α over F, as defined in §17.5. Furthermore, $\phi \neq 0$, and the degree m of ϕ may be characterized as the smallest positive integer m such that $1, \alpha, \ldots, \alpha^m$ are linearly dependent; moreover, as E has dimension ℓ over F, we must have $m \leq \ell$. □

***Example* 19.3.** Let V be a vector space over F of dimension $\ell > 0$, and let $\tau : V \to V$ be an F-linear map. Let $\beta \in V$, and consider the sequence $S := (\alpha_0, \alpha_1, \ldots)$, where $\alpha_i = \tau^i(\beta)$; that is, $\alpha_0 = \beta$, $\alpha_1 = \tau(\beta)$, $\alpha_2 = \tau(\tau(\beta))$,

and so on. For any polynomial $g = \sum_{j=0}^{k} g_j \mathtt{X}^j \in F[\mathtt{X}]$, we have

$$g \star S = \sum_{j=0}^{k} g_j \tau^j(\beta),$$

and for any $i \geq 0$, we have

$$(\mathtt{X}^i g) \star S = \sum_{j=0}^{k} g_j \tau^{i+j}(\beta) = \tau^i \left(\sum_{j=0}^{k} g_j \tau^j(\beta) \right) = \tau^i(g \star S).$$

Thus, if $g \star S = 0$, then clearly $(\mathtt{X}^i g) \star S = \tau^i(g \star S) = \tau^i(0) = 0$ for all $i \geq 0$. Conversely, if $(\mathtt{X}^i g) \star S = 0$ for all $i \geq 0$, then in particular, $g \star S = 0$. Thus, g is a generating polynomial for S if and only if $g \star S = 0$. The minimal polynomial ϕ of S is non-zero and its degree m is at most ℓ; indeed, m may be characterized as the least non-negative integer such that $\beta, \tau(\beta), \ldots, \tau^m(\beta)$ are linearly dependent, and since V has dimension ℓ over F, we must have $m \leq \ell$.

The previous example can be seen as a special case of this one, by taking V to be E, τ to be the α-multiplication map on E, and setting β to 1. □

The problem of computing the minimal polynomial of a linearly generated sequence can always be solved by means of Gaussian elimination. For example, the minimal polynomial of the sequence discussed in Example 19.2 can be computed using the algorithm described in §18.2. The minimal polynomial of the sequence discussed in Example 19.3 can be computed in a similar manner. Also, Exercise 19.3 below shows how one can reformulate another special case of the problem so that it is easily solved by Gaussian elimination. However, in the following sections, we will present algorithms for computing minimal polynomials for certain types of linearly generated sequences that are much more efficient than any algorithm based on Gaussian elimination.

Exercise 19.1. Show that the only sequence for which 1 is a generating polynomial is the "all zero" sequence.

Exercise 19.2. Let $S = (\alpha_0, \alpha_1, \ldots)$ be a sequence of elements of an F-vector space V. Further, suppose that S has non-zero minimal polynomial ϕ.

(a) Show that for any polynomials $g, h \in F[\mathtt{X}]$, if $g \equiv h \pmod{\phi}$, then $g \star S = h \star S$.

(b) Let $m := \deg(\phi)$. Show that if $g \in F[\mathtt{X}]$ and $(\mathtt{X}^i g) \star S = 0$ for $i = 0, \ldots, m-1$, then g is a generating polynomial for S.

Exercise 19.3. This exercise develops an alternative characterization linearly generated sequences. Let $S = (z_0, z_1, \ldots)$ be a sequence of elements of F. Further, suppose that S has minimal polynomial $\phi = \sum_{j=0}^{m} c_j \mathtt{X}^j$ with $m > 0$ and $c_m = 1$. Define the matrix

$$A := \begin{pmatrix} z_0 & z_1 & \cdots & z_{m-1} \\ z_1 & z_2 & \cdots & z_m \\ \vdots & \vdots & \ddots & \vdots \\ z_{m-1} & z_m & \cdots & z_{2m-2} \end{pmatrix} \in F^{m \times m}$$

and the vector

$$w := (z_m, \ldots, z_{2m-1}) \in F^{1 \times m}.$$

Show that

$$v = (-c_0, \ldots, -c_{m-1}) \in F^{1 \times m}$$

is the *unique* solution to the equation

$$vA = w.$$

Hint: show that the rows of A are linearly independent by making use of Exercise 19.2 and the fact that no polynomial of degree less than m is a generating polynomial for S.

Exercise 19.4. Suppose that you are given $a_0, \ldots, a_{k-1} \in F$ and $z_0, \ldots, z_{k-1} \in F$. Suppose that for all $i \geq 0$, we define

$$z_{k+i} := \sum_{j=0}^{k-1} a_j z_{j+i}.$$

Given $n \geq 0$, show how to compute z_n using $O(\mathrm{len}(n)k^2)$ operations in F.

Exercise 19.5. Let V be a vector space over F, and consider the set $V^{\times\infty}$ of all infinite sequences $(\alpha_0, \alpha_1, \ldots)$, where the α_i are in V. Let us define the scalar product of $g \in F[\mathtt{X}]$ and $S \in V^{\times\infty}$ as

$$g \cdot S = (g \star S, (\mathtt{X}g) \star S, (\mathtt{X}^2 g) \star S, \ldots) \in V^{\times\infty}.$$

Show that with this scalar product, $V^{\times\infty}$ is an $F[\mathtt{X}]$-module, and that a polynomial $g \in F[\mathtt{X}]$ is a generating polynomial for $S \in V^{\times\infty}$ if and only if $g \cdot S = 0$.

19.2 Computing minimal polynomials: a special case

We now tackle the problem of computing the minimal polynomial of a linearly generated sequence from a sufficiently long initial segment.

We shall first address a special case of this problem, namely, the case where the vector space V is just the field F. In this case, we have

$$S = (z_0, z_1, z_2, \ldots),$$

where $z_i \in F$ for $i = 0, 1, 2, \ldots$.

Suppose that we do not know the minimal polynomial ϕ of S, but we know an upper bound $M \geq 0$ on its degree. Then it turns out that the initial segment $z_0, z_1, \ldots z_{2M-1}$ completely determines ϕ, and moreover, we can very efficiently compute ϕ given the bound M and this initial segment. The following theorem provides the essential ingredient.

Theorem 19.2. *Let $S = (z_0, z_1, \ldots)$ be a sequence of elements of F, and define the reversed formal Laurent series*

$$z := \sum_{i=0}^{\infty} z_i \mathtt{X}^{-(i+1)} \in F((\mathtt{X}^{-1})),$$

whose coefficients are the elements of the sequence S. Then for any $g \in F[\mathtt{X}]$, we have $g \in G(S)$ if and only if $gz \in F[\mathtt{X}]$. In particular, S is linearly generated if and only if z is a rational function, in which case, its minimal polynomial is the denominator of z when expressed as a fraction in lowest terms.

Proof. Observe that for any polynomial $g \in F[\mathtt{X}]$ and any integer $i \geq 0$, the coefficient of $\mathtt{X}^{-(i+1)}$ in the product gz is equal to $\mathtt{X}^i g \star S$—just look at the formulas defining these expressions! It follows that g is a generating polynomial for S if and only if the coefficients of the negative powers of $\mathtt{X}$ in gz are all zero, which is the same as saying that $gz \in F[\mathtt{X}]$. Further, if $g \neq 0$ and $h := gz \in F[\mathtt{X}]$, then $\deg(h) < \deg(g)$—this follows simply from the fact that $\deg(z) < 0$ (together with the fact that $\deg(h) = \deg(g) + \deg(z)$). All the statements in the theorem follow immediately from these observations. □

By virtue of Theorem 19.2, we can compute the minimal polynomial ϕ of S using the algorithm in §18.5.2 for computing the numerator and denominator of a rational function from its reversed Laurent series expansion. More precisely, we can compute ϕ given the bound M on its degree, along with the first $2M$ elements $z_0, \ldots, z_{2M-1}$ of S, using $O(M^2)$ operations in F. Just for completeness, we write down this algorithm:

1. Run the extended Euclidean algorithm on inputs

$$a := \mathtt{X}^{2M} \quad \text{and} \quad b := z_0\mathtt{X}^{2M-1} + z_1\mathtt{X}^{2M-2} + \cdots + z_{2M-1},$$

   and let s', t' be as in Theorem 18.7, using $r^* := M$ and $t^* := M$.
2. Output $\phi := t'/\operatorname{lc}(t')$.

The characterization of linearly generated sequences provided by Theorem 19.2 is also very useful in other ways. For example, suppose the field F is finite. As we already saw in Example 19.1, any linearly generated sequence $S := (z_0, z_1, \ldots)$, where the z_i are in F, must be ultimately periodic. However, Theorem 19.2, together with the result of Exercise 18.13, tells us much more; for example, if the minimal polynomial ϕ of S is not divisible by $\mathtt{X}$, then S is purely periodic with period equal to the multiplicative order of $[\mathtt{X}]_\phi \in (F[\mathtt{X}]/(\phi))^*$.

19.3 Computing minimal polynomials: a more general case

Having dealt with the problem of finding the minimal polynomial of a sequence S of elements of F, we address the more general problem, where the elements of S lie in a vector space V over F. We shall only deal with a special case of this problem, but it is one which has useful applications:

- First, we shall assume that V has finite dimension $\ell > 0$ over F.
- Second, we shall assume that the sequence $S = (\alpha_0, \alpha_1, \ldots)$ has **full rank**, by which we mean the following: if the minimal polynomial ϕ of S over F has degree m, then the vectors $\alpha_0, \ldots, \alpha_{m-1}$ are linearly independent. The sequences considered in Examples 19.2 and 19.3 are of this type.
- Third, we shall assume that F is a finite field.

The Dual Space. To develop the theory behind the approach we are going to present, we need to discuss the **dual space** $\mathcal{D}_F(V)$ of V (over F), which consists of all F-linear maps from V into F. We may sometimes refer to elements of $\mathcal{D}_F(V)$ as **projections**. Now, as was discussed in §15.2, if we fix an ordered basis $\gamma_1, \ldots, \gamma_\ell$ for V, the elements of V are in one-to-one correspondence with the coordinate vectors $F^{1\times\ell}$, where the element $a_1\gamma_1 + \ldots + a_\ell\gamma_\ell \in V$ corresponds to the coordinate vector $(a_1, \ldots, a_\ell) \in F^{1\times\ell}$. The elements of $\mathcal{D}_F(V)$ are in one-to-one correspondence with $F^{\ell\times 1}$, where the map $\pi \in \mathcal{D}_F(V)$ corresponds to the column vector whose jth coordinate is $\pi(\gamma_j)$, for $j = 1, \ldots, \ell$. It is natural to call the column vector corresponding to π its **coordinate vector**. A projection $\pi \in \mathcal{D}_F(V)$ may

be evaluated at a point $\delta \in V$ by taking the product of the coordinate vector of δ with the coordinate vector of π.

One may also impose a vector space structure on $\mathcal{D}_F(V)$, in a very natural way: for $\pi, \pi' \in \mathcal{D}_F(V)$, the map $\pi + \pi'$ sends $\delta \in V$ to $\pi(\delta) + \pi'(\delta)$, and for $c \in F$, the map $c\pi$ sends $\delta \in V$ to $c\pi(\delta)$. By the observations in the previous paragraph, $\mathcal{D}_F(V)$ is an F-vector space of dimension ℓ; indeed, the sum and scalar multiplication operations on $\mathcal{D}_F(V)$ correspond to analogous operations on coordinate vectors.

One last fact we need about the dual space is the following:

Theorem 19.3. *Let V be an F-vector space of finite dimension $\ell > 0$. For any linearly independent vectors $\delta_1, \ldots, \delta_m \in V$, and any $a_1, \ldots, a_m \in F$, there exists $\pi \in \mathcal{D}_F(V)$ such that $\pi(\delta_i) = a_i$ for $i = 1, \ldots, m$.*

Proof. Fix any ordered basis for V, and let M be the $m \times \ell$ matrix whose ith row is the coordinate vector of δ_i with respect to this ordered basis. Let v be the $m \times 1$ column vector whose ith coordinate is a_i. As the δ_i are linearly independent, the rows of M must also be linearly independent. Therefore, the F-linear map that sends $w \in F^{\ell \times 1}$ to $Mw \in F^{m \times 1}$ is surjective. It follows that any solution w to the equation $v = Mw$ is the coordinate vector of a map $\pi \in \mathcal{D}_F(V)$ that satisfies the requirements of the theorem. □

That completes our digression on the dual space. We now return to the problem of computing the minimal polynomial ϕ of the linearly generated sequence $S = (\alpha_0, \alpha_1, \ldots)$. Assume we have a bound M on the degree of ϕ. As we are assuming S has full rank, we may assume that $M \leq \ell$. For any $\pi \in \mathcal{D}_F(V)$, we may consider the projected sequence $S_\pi = (\pi(\alpha_0), \pi(\alpha_1), \ldots)$. Observe that ϕ is a generating polynomial for S_π; indeed, for any polynomial $g \in F[\mathtt{X}]$, we have $g \star S_\pi = \pi(g \star S)$, and hence, for all $i \geq 0$, we have $(\mathtt{X}^i\phi) \star S_\pi = \pi((\mathtt{X}^i\phi) \star S) = \pi(0) = 0$. Let $\phi_\pi \in F[\mathtt{X}]$ denote the minimal polynomial of S_π. Since ϕ_π divides any generating polynomial of S_π, and since ϕ is a generating polynomial for S_π, it follows that ϕ_π is a divisor of ϕ.

This suggests the following algorithm for efficiently computing the minimal polynomial of S:

Algorithm MP:

$g \leftarrow 1 \in F[\mathtt{X}]$
repeat
    choose $\pi \in \mathcal{D}_F(V)$ at random
    compute the first $2M$ terms of the projected sequence S_π
    use the algorithm in §19.2 to compute the minimal polynomial
        ϕ_π of S_π
    $g \leftarrow \mathrm{lcm}(g, \phi_\pi)$
until $g \star S = 0$
output g

A few remarks on the above procedure are in order:

- in every iteration of the main loop, g is the least common multiple of a number of divisors of ϕ, and hence is itself a divisor of ϕ;
- under our assumption that S has full rank, and since g is a monic divisor of ϕ, if $g \star S = 0$, we may safely conclude that $g = \phi$;
- under our assumption that F is finite, choosing a random element π of $\mathcal{D}_F(V)$ amounts to simply choosing at random the entries of the coordinate vector of π, relative to some ordered basis for V;
- we also assume that elements of V are represented as coordinate vectors, so that applying a projection $\pi \in \mathcal{D}_F(V)$ to a vector in V takes $O(\ell)$ operations in F;
- similarly, adding two elements of V, or multiplying an element of V times a scalar, takes $O(\ell)$ operations in F.

Based on the above observations, it follows that when the algorithm halts, its output is correct, and that the cost of each loop iteration is $O(M\ell)$ operations in F. The remaining question to be answered is this: what is the expected number of iterations of the main loop? The answer to this question is $O(1)$, which leads to a total expected cost of Algorithm MP of $O(M\ell)$ operations in F.

The key to establishing that the expected number of iterations of the main loop is constant is provided by the following theorem.

Theorem 19.4. *Let $S = (\alpha_0, \alpha_1, \ldots)$ be a linearly generated sequence over the field F, where the α_i are elements of a vector space V of finite dimension $\ell > 0$. Let ϕ be the minimal polynomial of S over F, let $m := \deg(\phi)$, and assume that S has full rank (i.e., $\alpha_0, \ldots, \alpha_{m-1}$ are linearly independent).*

Under the above assumptions, there exists a surjective F-linear map $\sigma : \mathcal{D}_F(V) \to F[\mathtt{X}]_{<m}$ such that for all $\pi \in \mathcal{D}_F(V)$, the minimal polynomial ϕ_π

of the projected sequence $S_\pi := (\pi(\alpha_0), \pi(\alpha_1), \ldots)$ *satisfies*

$$\phi_\pi = \frac{\phi}{\gcd(\sigma(\pi), \phi)}.$$

Recall that $F[\mathtt{X}]_{<m}$ denotes the m-dimensional vector space of polynomials in $F[\mathtt{X}]$ of degree less than m.

Proof. While the statement of this theorem looks a bit complicated, its proof is quite straightforward, given our characterization of linearly generated sequences in Theorem 19.2 in terms of rational functions. We build the linear map σ as the composition of two linear maps, σ_0 and σ_1.

Let us define the map

$$\begin{aligned} \sigma_0 : \quad \mathcal{D}_F(V) &\to F(\!(\mathtt{X}^{-1})\!) \\ \pi &\mapsto \sum_{i=0}^{\infty} \pi(\alpha_i)\mathtt{X}^{-(i+1)}. \end{aligned}$$

We also define the map σ_1 to be the ϕ-multiplication map on $F(\!(\mathtt{X}^{-1})\!)$, that is, the map that sends $z \in F(\!(\mathtt{X}^{-1})\!)$ to $\phi \cdot z \in F(\!(\mathtt{X}^{-1})\!)$. The map σ is just the composition $\sigma = \sigma_1 \circ \sigma_0$. It is clear that both σ_0 and σ_1 are F-linear maps, and hence, so is σ.

First, observe that for $\pi \in \mathcal{D}_F(V)$, the series $z := \sigma_0(\pi)$ is the series associated with the projected sequence S_π, as in Theorem 19.2. Let ϕ_π be the minimal polynomial of S_π. Since ϕ is a generating polynomial for S, it is also a generating polynomial for S_π. Therefore, Theorem 19.2 tells us that

$$h := \sigma(\pi) = \phi \cdot z \in F[\mathtt{X}]_{<m},$$

and that ϕ_π is the denominator of z when expressed as a fraction in lowest terms. Now, we have $z = h/\phi$, and it follows that $\phi_\pi = \phi/\gcd(h, \phi)$ is this denominator.

Second, the hypothesis that $\alpha_0, \ldots, \alpha_{m-1}$ are linearly independent, together with Theorem 19.3, implies that $\dim_F(\mathrm{img}(\sigma_0)) \geq m$. Also, observe that σ_1 is an injective map (indeed, it is surjective as well). Therefore, $\dim_F(\mathrm{img}(\sigma)) \geq m$. In the previous paragraph, we observed that $\mathrm{img}(\sigma) \subseteq F[\mathtt{X}]_{<m}$, and since $\dim_F(F[\mathtt{X}]_{<m}) = m$, we may conclude that $\mathrm{img}(\sigma) = F[\mathtt{X}]_{<m}$. That proves the theorem. □

Given the above theorem, we can analyze the expected number of iterations of the main loop of Algorithm MP.

First of all, we may as well assume that the degree m of ϕ is greater than 0, as otherwise, we are sure to get ϕ in the very first iteration. Let $\pi_1, \ldots, \pi_s$

be the random projections chosen in the first s iterations of Algorithm MP. By Theorem 19.4, the polynomials $\sigma(\pi_1), \ldots, \sigma(\pi_s)$ are uniformly and independently distributed over $F[\mathtt{X}]_{<m}$, and we have $g = \phi$ at the end of loop iteration s if and only if $\gcd(\phi, \sigma(\pi_1), \ldots, \sigma(\pi_s)) = 1$.

Let us define $\Lambda_F^\phi(s)$ to be the probability that $\gcd(\phi, f_1, \ldots, f_s) = 1$, where $f_1, \ldots, f_s$ are randomly chosen from $F[\mathtt{X}]_{<m}$. Thus, the probability that we have $g = \phi$ at the end of loop iteration s is equal to $\Lambda_F^\phi(s)$. While one can analyze the quantity $\Lambda_F^\phi(s)$, it turns out to be easier, and sufficient for our purposes, to analyze a different quantity. Let us define $\Lambda_F^m(s)$ to be the probability that $\gcd(f_1, \ldots, f_s) = 1$, where $f_1, \ldots, f_s$ are randomly chosen from $F[\mathtt{X}]_{<m}$. Clearly, $\Lambda_F^\phi(s) \geq \Lambda_F^m(s)$.

Theorem 19.5. *If F is a finite field of cardinality q, and m and s are positive integers, then we have*

$$\Lambda_F^m(s) = 1 - 1/q^{s-1} + (q-1)/q^{sm}.$$

Proof. For any positive integer n, let U_n be the set of all tuples of polynomials $(f_1, \ldots, f_s) \in F[\mathtt{X}]_{<n}^{\times s}$ with $\gcd(f_1, \ldots, f_s) = 1$, and let $u_n = |U_n|$. First, let h be any monic polynomial with $k := \deg(h) < n$. The set $U_{n,h}$ of all s-tuples of polynomials of degree less than n whose gcd is h is in one-to-one correspondence with U_{n-k}, via the map that sends $(f_1, \ldots, f_s) \in U_{n,h}$ to $(f_1/h, \ldots, f_s/h) \in U_{n-k}$. As there are q^k possible choices for h of degree k, we see that the set $V_{n,k}$, consisting of tuples $(f_1, \ldots, f_s) \in F[\mathtt{X}]_{<n}^{\times s}$ with $\deg(\gcd(f_1, \ldots, f_s)) = k$, has cardinality $q^k u_{n-k}$. Every non-zero element of $F[\mathtt{X}]_{<n}^{\times s}$ appears in exactly one of the sets $V_{n,k}$, for $k = 0, \ldots, n-1$. Taking into account the zero polynomial, it follows that

$$q^{sn} = 1 + \sum_{k=0}^{n-1} q^k u_{n-k}, \tag{19.2}$$

which holds for all $n \geq 1$. Replacing n by $n-1$ in (19.2), we obtain

$$q^{s(n-1)} = 1 + \sum_{k=0}^{n-2} q^k u_{n-1-k}, \tag{19.3}$$

which holds for all $n \geq 2$, and indeed, holds for $n = 1$ as well. Subtracting q times (19.3) from (19.2), we deduce that for $n \geq 1$,

$$q^{sn} - q^{sn-s+1} = 1 + u_n - q,$$

and rearranging terms:

$$u_n = q^{sn} - q^{sn-s+1} + q - 1.$$

Therefore,

$$\Lambda_F^m(s) = u_m/q^{sm} = 1 - 1/q^{s-1} + (q-1)/q^{sm}. \quad \square$$

From the above theorem, it follows that for $s \geq 1$, the probability P_s that Algorithm MP runs for more than s loop iterations is at most $1/q^{s-1}$. If T is the total number of loop iterations, then

$$\mathsf{E}[T] = \sum_{i \geq 1} \mathsf{P}[T \geq i] = 1 + \sum_{s \geq 1} P_s \leq 1 + \sum_{s \geq 1} 1/q^{s-1} = 1 + \frac{q}{q-1} = O(1).$$

Let us summarize all of the above analysis with the following:

Theorem 19.6. *Let S be a sequence of elements of an F-vector space V of finite dimension $\ell > 0$ over F, where F is a finite field. Assume that S is linearly generated over F with minimal polynomial $\phi \in F[\mathtt{X}]$ of degree m, and that S has full rank (i.e., the first m elements of S are linearly independent). Then given an upper bound M on m, along with the first $2M$ elements of S, Algorithm MP correctly computes ϕ using an expected number of $O(M\ell)$ operations in F.*

We close this section with the following observation. Suppose the sequence S is of the form $(\beta, \tau(\beta), \tau^2(\beta), \ldots)$, where $\beta \in V$ and $\tau : V \to V$ is an F-linear map. Suppose that with respect to some ordered basis for V, elements of V are represented as elements of $F^{1\times\ell}$, and elements of $\mathcal{D}_F(V)$ are represented as elements of $F^{\ell\times 1}$. The linear map τ also has a corresponding representation as a matrix $A \in F^{\ell\times\ell}$, so that evaluating τ at a point α in V corresponds to multiplying the coordinate vector of α on the right by A. Now, suppose $\beta \in V$ has coordinate vector $b \in F^{1\times\ell}$ and that $\pi \in \mathcal{D}_F(V)$ has coordinate vector $c^\top \in F^{\ell\times 1}$. Then if $\tilde{S}$ is the sequence of coordinate vectors of the elements of S, we have

$$\tilde{S} = (bA^i)_{i=0}^{\infty} \quad \text{and} \quad S_\pi = (bA^i c^\top)_{i=0}^{\infty}.$$

This more concrete, matrix-oriented point of view is sometimes useful; in particular, it makes quite transparent the symmetry of the roles played by β and π in forming the projected sequence.

Exercise 19.6. If $|F| = q$ and $\phi \in F[\mathtt{X}]$ is monic and factors into monic irreducible polynomials in $F[\mathtt{X}]$ as $\phi = p_1^{e_1} \cdots p_r^{e_r}$, show that

$$\Lambda_F^\phi(1) = \prod_{i=1}^{r} (1 - q^{-\deg(p_i)}) \geq 1 - \sum_{i=1}^{r} q^{-\deg(p_i)}.$$

From this, conclude that the probability that Algorithm MP terminates

after just one loop iteration is $1 - O(m/q)$, where $m = \deg(\phi)$. Thus, if q is very large relative to m, it is highly likely that Algorithm MP terminates after just one iteration of the main loop.

19.4 Solving sparse linear systems

Let V be a vector space of finite dimension $\ell > 0$ over a finite field F, and let $\tau : V \to V$ be an F-linear map. The goal of this section is to develop time- and space-efficient algorithms for solving equations of the form

$$\tau(\gamma) = \delta; \tag{19.4}$$

that is, given τ and $\delta \in V$, find $\gamma \in V$ satisfying (19.4). The algorithms we develop will have the following properties: they will be probabilistic, and will use an expected number of $O(\ell^2)$ operations in F, an expected number of $O(\ell)$ evaluations of τ, and space for $O(\ell)$ elements of F. By an "evaluation of τ," we mean the computation of $\tau(\alpha)$ for some $\alpha \in V$.

We shall assume that elements of V are represented as coordinate vectors with respect to some fixed ordered basis for V. Now, if the matrix representing τ with respect to the given ordered basis is sparse, having, say, $\ell^{1+o(1)}$ non-zero entries, then the space required to represent τ is $\ell^{1+o(1)}$ elements of F, and the time required to evaluate τ is $\ell^{1+o(1)}$ operations in F. Under these assumptions, our algorithms to solve (19.4) use an expected number of $\ell^{2+o(1)}$ operations in F, and space for $\ell^{1+o(1)}$ elements of F. This is to be compared with standard Gaussian elimination: even if the original matrix is sparse, during the execution of the algorithm, most of the entries in the matrix may eventually be "filled in" with non-zero field elements, leading to a running time of $\Omega(\ell^3)$ operations in F, and a space requirement of $\Omega(\ell^2)$ elements of F. Thus, the algorithms presented here will be much more efficient than Gaussian elimination when the matrix representing τ is sparse.

We hasten to point out that the algorithms presented here may be more efficient than Gaussian elimination in other cases, as well. All that matters is that τ can be evaluated using $o(\ell^2)$ operations in F and/or represented using space for $o(\ell^2)$ elements of F—in either case, we obtain a time and/or space improvement over Gaussian elimination. Indeed, there are applications where the matrix of the linear map τ may not be sparse, but nevertheless has special structure that allows it to be represented and evaluated in subquadratic time and/or space.

We shall only present algorithms that work in two special, but important, cases:

- the first case is where τ is invertible,
- the second case is where τ is not invertible, $\delta = 0$, and a non-zero solution γ to (19.4) is required (i.e., we are looking for a non-zero element of $\ker(\tau)$).

In both cases, the key will be to use Algorithm MP in §19.3 to find the minimal polynomial ϕ of the linearly generated sequence

$$S := (\alpha_0, \alpha_1, \ldots), \quad (\alpha_i = \tau^i(\beta),\ i = 0, 1, \ldots), \tag{19.5}$$

where β is a suitably chosen element of V. From the discussion in Example 19.3, this sequence has full rank, and so we may use Algorithm MP. We may use $M := \ell$ as an upper bound on the degree of ϕ (assuming we know nothing more about τ and β that would allow us to use a smaller upper bound). In using Algorithm MP in this application, note that we do not want to store $\alpha_0, \ldots, \alpha_{2\ell-1}$—if we did, we would not satisfy our stated space bound. Instead of storing the α_i in a "warehouse," we use a "just in time" strategy for computing them, as follows:

- In the body of the main loop of Algorithm MP, where we calculate the values $a_i := \pi(\alpha_i)$, for $i = 0 \ldots 2\ell - 1$, we perform the computation as follows:

  $\alpha \leftarrow \beta$
  for $i \leftarrow 0$ to $2\ell - 1$ do
  $\quad a_i \leftarrow \pi(\alpha),\ \ \alpha \leftarrow \tau(\alpha)$

- In the test at the bottom of the main loop of Algorithm MP, if $g = \sum_{j=0}^{k} g_j \mathtt{X}^j$, we compute $\nu := g \star S \in V$ as follows:

  $\nu \leftarrow 0,\ \ \alpha \leftarrow \beta$
  for $j \leftarrow 0$ to k do
  $\quad \nu \leftarrow \nu + g_j \cdot \alpha,\ \ \alpha \leftarrow \tau(\alpha)$

  Alternatively, one could use a Horner-like algorithm:

  $\nu \leftarrow 0$
  for $j \leftarrow k$ down to 0 do
  $\quad \nu \leftarrow \tau(\nu) + g_j \cdot \beta$

With this implementation, Algorithm MP uses an expected number of $O(\ell^2)$ operations in F, an expected number of $O(\ell)$ evaluations of τ, and space for $O(\ell)$ elements of F. Of course, the "warehouse" strategy is faster than the "just in time" strategy by a constant factor, but it uses about ℓ times as much space; thus, for large ℓ, using the "just in time" strategy is a very good time/space trade-off.

The invertible case. Now consider the case where τ is invertible, and we want to solve (19.4) for a given $\delta \in V$. We may as well assume that $\delta \neq 0$, since otherwise, $\gamma = 0$ is the unique solution to (19.4). We proceed as follows. First, using Algorithm MP as discussed above, compute the minimal polynomial ϕ of the sequence S defined in (19.5), using $\beta := \delta$. Let $\phi = \sum_{j=0}^{m} c_j \mathtt{X}^j$, where $c_m = 1$ and $m > 0$. Then we have

$$c_0\delta + c_1\tau(\delta) + \cdots + c_m\tau^m(\delta) = 0. \tag{19.6}$$

We claim that $c_0 \neq 0$. To prove the claim, suppose that $c_0 = 0$. Then applying τ^{-1} to (19.6), we would obtain

$$c_1\delta + \cdots + c_m\tau^{m-1}(\delta) = 0,$$

which would imply that $\phi/\mathtt{X}$ is a generating polynomial for S, contradicting the minimality of ϕ. That proves the claim.

Since $c_0 \neq 0$, we can apply τ^{-1} to (19.6), and solve for $\gamma = \tau^{-1}(\delta)$ as follows:

$$\gamma = -c_0^{-1}(c_1\delta + \cdots + c_m\tau^{m-1}(\delta)).$$

To actually compute γ, we use the same "just in time" strategy as was used in the implementation of the computation of $g \star S$ in Algorithm MP, which costs $O(\ell^2)$ operations in F, $O(\ell)$ evaluations of τ, and space for $O(\ell)$ elements of F.

The non-invertible case. Now consider the case where τ is not invertible, and we want to find non-zero vector $\gamma \in V$ such that $\tau(\gamma) = 0$. The idea is this. Suppose we choose an arbitrary, non-zero element β of V, and use Algorithm MP to compute the minimal polynomial ϕ of the sequence S defined in (19.5), using this value of β. Let $\phi = \sum_{j=0}^{m} c_j \mathtt{X}^j$, where $m > 0$ and $c_m = 1$. Then we have

$$c_0\beta + c_1\tau(\beta) + \cdots + c_m\tau^m(\beta) = 0. \tag{19.7}$$

Let

$$\gamma := c_1\beta + \cdots c_m\tau^{m-1}(\beta).$$

We must have $\gamma \neq 0$, since $\gamma = 0$ would imply that $\lfloor \phi/\mathtt{X} \rfloor$ is a non-zero generating polynomial for S, contradicting the minimality of ϕ. If it happens that $c_0 = 0$, then equation (19.7) implies that $\tau(\gamma) = 0$, and we are done. As before, to actually compute γ, we use the same "just in time" strategy as was used in the implementation of the computation of $g \star S$ in Algorithm MP, which costs $O(\ell^2)$ operations in F, $O(\ell)$ evaluations of τ, and space for $O(\ell)$ elements of F.

The above approach fails if $c_0 \neq 0$. However, in this "bad" case, equation (19.7) implies that $\beta = -c_0^{-1}\tau(\gamma)$; that is, $\beta \in \mathrm{img}(\tau)$. One way to avoid such a "bad" β is to randomize: as τ is not surjective, the image of τ is a subspace of V of dimension strictly less than ℓ, and therefore, a *randomly* chosen β lies in the image of τ with probability at most $1/|F|$. So a simple technique is to choose repeatedly β at random until we get a "good" β. The overall complexity of the resulting algorithm will be as required: $O(\ell^2)$ expected operations in F, $O(\ell)$ expected evaluations of τ, and space for $O(\ell)$ elements of F.

As a special case of this situation, consider the problem that arose in Chapter 16 in connection with algorithms for computing discrete logarithms and factoring. We had to solve the following problem: given an $\ell \times (\ell - 1)$ matrix M with entries in a finite field F, containing $\ell^{1+o(1)}$ non-zero entries, find a non-zero vector $v \in F^{1\times\ell}$ such that $vM = 0$. To solve this problem, we can augment the matrix M, adding an extra column of zeros, to get an $\ell \times \ell$ matrix M'. Now, let $V = F^{1\times\ell}$ and let τ be the F-linear map on V that sends $\gamma \in V$ to $\gamma M'$. A non-zero solution γ to the equation $\tau(\gamma) = 0$ will provide us with the solution to our original problem; thus, we can apply the above technique directly, solving this problem using $\ell^{2+o(1)}$ expected operations in F, and space for $\ell^{1+o(1)}$ elements of F. As a side remark, in this particular application, we can choose a "good" β in the above algorithm without randomization: just choose $\beta := (0, \ldots, 0, 1)$, which is clearly not in the image of τ.

19.5 Computing minimal polynomials in $F[\mathtt{X}]/(f)$ (II)

Let us return to the problem discussed in §18.2: F is a field, $f \in F[\mathtt{X}]$ is a monic polynomial of degree $\ell > 0$, and $E := F[\mathtt{X}]/(f) = F[\eta]$, where $\eta := [\mathtt{X}]_f$; we are given an element $\alpha \in E$, and want to compute the minimal polynomial $\phi \in F[\mathtt{X}]$ of α over F. As discussed in Example 19.2, this problem is equivalent to the problem of computing the minimal polynomial of the sequence

$$S := (\alpha_0, \alpha_1, \ldots) \quad (\alpha_i := \alpha^i,\ i = 0, 1, \ldots),$$

and the sequence has full rank; therefore, we can use Algorithm MP in §19.3 directly to solve this problem, assuming F is a finite field.

If we use the "just in time" strategy in the implementation of Algorithm MP, as was used in §19.4, we get an algorithm that computes the minimal polynomial of α using $O(\ell^3)$ expected operations in F, but space for just $O(\ell^2)$ elements of F. Thus, in terms of space, this approach is far superior

to the algorithm in §18.2, based on Gaussian elimination. In terms of time complexity, the algorithm based on linearly generated sequences is a bit slower than the one based on Gaussian elimination (but only by a constant factor). However, if we use any subquadratic-time algorithm for polynomial arithmetic (see §18.6 and §18.7), we immediately get an algorithm that runs in subcubic time, while still using linear space. In the exercises below, you are asked to develop an algorithm that computes the minimal polynomial of α using just $O(\ell^{2.5})$ operations in F, at the expense of requiring space for $O(\ell^{1.5})$ elements of F—this algorithm does not rely on fast polynomial arithmetic, and can be made even faster if such arithmetic is used.

Exercise 19.7. Let $f \in F[\mathtt{X}]$ be a monic polynomial of degree $\ell > 0$ over a field F, and let $E := F[\mathtt{X}]/(f)$. Also, let $\eta := [\mathtt{X}]_f \in E$. For computational purposes, we assume that elements of E and $\mathcal{D}_F(E)$ are represented as coordinate vectors with respect to the usual "polynomial" basis $1, \eta, \ldots, \eta^{\ell-1}$. For $\beta \in E$, let M_β denote the β-multiplication map on E that sends $\alpha \in E$ to $\alpha\beta \in E$, which is an F-linear map from E into E.

(a) Show how to compute—given as input the polynomial f defining E, along with a projection $\pi \in \mathcal{D}_F(E)$ and an element $\beta \in E$—the projection $\pi \circ M_\beta \in \mathcal{D}_F(E)$, using $O(\ell^2)$ operations in F.

(b) Show how to compute—given as input the polynomial f defining E, along with a projection $\pi \in \mathcal{D}_F(E)$, an element $\alpha \in E$, and a parameter $k > 0$—all of the k values

$$\pi(1), \pi(\alpha), \ldots, \pi(\alpha^{k-1})$$

using just $O(k\ell + k^{1/2}\ell^2)$ operations in F, and space for $O(k^{1/2}\ell)$ elements of F. Hint: use the same hint as in Exercise 18.4.

Exercise 19.8. Let $f \in F[\mathtt{X}]$ be a monic polynomial over a finite field F of degree $\ell > 0$, and let $E := F[\mathtt{X}]/(f)$. Show how to use the result of the previous exercise, as well as Exercise 18.4, to get an algorithm that computes the minimal polynomial of $\alpha \in E$ over F using $O(\ell^{2.5})$ expected operations in F, and space for $O(\ell^{1.5})$ operations in F.

Exercise 19.9. Let $f \in F[\mathtt{X}]$ be a monic polynomial of degree $\ell > 0$ over a field F (not necessarily finite), and let $E := F[\mathtt{X}]/(f)$. Further, suppose that f is irreducible, so that E is itself a field. Show how to compute the minimal polynomial of $\alpha \in E$ over F *deterministically*, satisfying the following complexity bounds:

(a) $O(\ell^3)$ operations in F and space for $O(\ell)$ elements of F;

(b) $O(\ell^{2.5})$ operations in F and space for $O(\ell^{1.5})$ elements of F.

19.6 The algebra of linear transformations (∗)

Throughout this chapter, one could hear the whispers of the algebra of linear transformations. We develop some of the aspects of this theory here, leaving a number of details as exercises. It will not play a role in any material that follows, but it serves to provide the reader with a "bigger picture."

Let F be a field and V be a non-trivial F-vector space. We denote by $\mathcal{L}_F(V)$ the set of all F-linear maps from V into V. Elements of $\mathcal{L}_F(V)$ are called **linear transformations**. We can make $\mathcal{L}_F(V)$ into an F-vector space by defining addition and scalar multiplication as follows: for $\tau, \tau' \in \mathcal{L}_F(V)$, define $\tau + \tau'$ to be the map that sends $\alpha \in V$ to $\tau(\alpha) + \tau'(\alpha)$; for $c \in F$ and $\tau \in \mathcal{L}_F(V)$, define $c\tau$ to be the map that sends $\alpha \in V$ to $c\tau(\alpha)$.

Exercise 19.10. (a) Verify that with addition and scalar multiplication defined as above, $\mathcal{L}_F(V)$ is an F-vector space.

(b) Suppose that V has finite dimension $\ell > 0$. By identifying elements of $\mathcal{L}_F(V)$ with $\ell \times \ell$ matrices over F, show that $\mathcal{L}_F(V)$ has dimension ℓ^2.

As usual, for $\tau, \tau' \in \mathcal{L}_F(V)$, the composed map, $\tau \circ \tau'$ that sends $\alpha \in V$ to $\tau(\tau'(\alpha))$ is also an element of $\mathcal{L}_F(V)$ (verify). As always, function composition is associative (i.e., for $\tau, \tau', \tau'' \in \mathcal{L}_F(V)$, we have $\tau \circ (\tau' \circ \tau'') = (\tau \circ \tau') \circ \tau''$); however, function composition is not in general commutative (i.e., we may have $\tau \circ \tau' \neq \tau' \circ \tau$ for some $\tau, \tau' \in \mathcal{L}_F(V)$). For any $\tau \in \mathcal{L}_F(V)$ and an integer $i \geq 0$, the map τ^i (i.e., the i-fold composition of τ) is also an element of $\mathcal{L}_F(V)$. Note that for any $\tau \in \mathcal{L}_F(V)$, the map τ^0 is by definition just the identity map on V.

For any $\tau \in \mathcal{L}_F(V)$, and for any polynomial $f \in F[\mathtt{X}]$, with $f = \sum_i a_i \mathtt{X}_i$, we denote by $f(\tau)$ the linear transformation

$$f(\tau) := \sum_i a_i \tau^i.$$

Exercise 19.11. Verify the following properties of $\mathcal{L}_F(V)$. For all $\tau, \tau', \tau'' \in \mathcal{L}_F(V)$, for all $c \in F$, and all $f, g \in F[\mathtt{X}]$:

(a) $\tau \circ (\tau' + \tau'') = \tau \circ \tau' + \tau \circ \tau''$;

(b) $(\tau' + \tau'') \circ \tau = \tau' \circ \tau + \tau'' \circ \tau$;

(c) $c(\tau \circ \tau') = (c\tau) \circ \tau' = \tau \circ (c\tau')$;

(d) $f(\tau) \circ g(\tau) = (fg)(\tau) = g(\tau) \circ f(\tau)$;

(e) $f(\tau) + g(\tau) = (f + g)(\tau)$.

Under the addition operation of the vector space $\mathcal{L}_F(V)$, and defining multiplication on $\mathcal{L}_F(V)$ using the "$\circ$" operation, we get an algebraic structure that satisfies all the properties of Definition 9.1, with the exception of property (v) of that definition (commutativity). Thus, we can view $\mathcal{L}_F(V)$ as a *non-commutative* ring with unity (the identity map acts as the multiplicative identity).

For a fixed $\tau \in \mathcal{L}_F(V)$, we may consider the subset of $\mathcal{L}_F(V)$,

$$F[\tau] := \{f(\tau) : f \in F[\mathtt{X}]\},$$

which does in fact satisfy all the properties of Definition 9.1. Moreover, we can view F as a subring of $F[\tau]$ by identifying $c \in F$ with $c\tau^0 \in F[\tau]$. With this convention, for $f \in F[\mathtt{X}]$, the expression $f(\tau)$ has its usual meaning as the value of f evaluated at the point τ in the extension ring $F[\tau]$ of F. Let ϕ_τ is the minimal polynomial of τ over F, so that $F[\tau]$ is isomorphic as an F-algebra to $F[\mathtt{X}]/(\phi_\tau)$. We can also characterize ϕ_τ as follows (verify):

> if there exists a non-zero polynomial $f \in F[\mathtt{X}]$ such that $f(\tau) = 0$, then ϕ_τ is the monic polynomial of least degree with this property; otherwise, $\phi_\tau = 0$.

Another way to characterize ϕ is as follows (verify):

> ϕ_τ is the minimal polynomial of the sequence $(1, \tau, \tau^2, \ldots)$.

Note that ϕ_τ is never 1 — this follows from the assumption that V is non-trivial.

It is easy to see that if V happens to be finite dimensional, with $\ell := \dim_F(V)$, then by Exercise 19.10, $\mathcal{L}_F(V)$ has dimension ℓ^2. Therefore, there must be a linear dependence among $1, \tau, \ldots, \tau^{\ell^2}$, which implies that the minimal polynomial of τ is non-zero with degree at most ℓ^2. We shall show below that in this case, the minimal polynomial of τ actually has degree at most ℓ.

For a fixed $\tau \in \mathcal{L}_F(V)$, we can define a "scalar multiplication" operation $\odot$, that maps $f \in F[\mathtt{X}]$ and $\alpha \in V$ to

$$f \odot \alpha := f(\tau)(\alpha) \in V;$$

that is, if $f = \sum_i a_i \mathtt{X}^i$, then

$$f \odot \alpha = \sum_i a_i \tau^i(\alpha).$$

EXERCISE 19.12. Show that the scalar multiplication $\odot$, together with the usual addition operation on V, makes V into an $F[\mathtt{X}]$-module; that is, show that for all $f, g \in F[\mathtt{X}]$ and $\alpha, \beta \in V$, we have

$$f \odot (g \odot \alpha) = (fg) \odot \alpha, \ (f + g) \odot \alpha = f \odot \alpha + g \odot \alpha,$$
$$f \odot (\alpha + \beta) = f \odot \alpha + f \odot \beta, \ 1 \odot \alpha = \alpha.$$

Note that each choice of τ gives rise to a different $F[\mathtt{X}]$-module structure, but all of these structures are extensions of the usual vector space structure, in the sense that for all $c \in F$ and $\alpha \in V$, we have $c \odot \alpha = c\alpha$.

Now, for fixed $\tau \in \mathcal{L}_F(V)$ and $\alpha \in V$, consider the $F[\mathtt{X}]$-linear map $\rho_{\tau,\alpha} : F[\mathtt{X}] \to V$ that sends $f \in F[\mathtt{X}]$ to $f \odot \alpha = f(\tau)(\alpha)$. The kernel of this map must be a submodule, and hence an ideal, of $F[\mathtt{X}]$; since every ideal of $F[\mathtt{X}]$ is principal, it follows that $\ker(\rho_{\tau,\alpha})$ is the ideal of $F[\mathtt{X}]$ generated by some polynomial $\phi_{\tau,\alpha}$, which we can make unique by insisting that it is monic or zero. We call $\phi_{\tau,\alpha}$ the **minimal polynomial of α under τ**.We can also characterize $\phi_{\tau,\alpha}$ as follows (verify):

> if there exists a non-zero polynomial $f \in F[\mathtt{X}]$ such that $f(\tau)(\alpha) = 0$, then $\phi_{\tau,\alpha}$ the monic polynomial of least degree with this property; otherwise, $\phi_{\tau,\alpha} = 0$.

Another way to characterize $\phi_{\tau,\alpha}$ is as follows (verify):

> $\phi_{\tau,\alpha}$ is the minimal polynomial of the sequence
>
> $$(\alpha, \tau(\alpha), \tau^2(\alpha), \ldots).$$

Note that since $\phi_\tau(\tau)$ is the zero map, we have

$$\phi_\tau \odot \alpha = \phi_\tau(\tau)(\alpha) = 0,$$

and hence $\phi_\tau \in \ker(\rho_{\tau,\alpha})$, which means that $\phi_{\tau,\alpha} \mid \phi_\tau$.

Now consider the image of $\rho_{\tau,\alpha}$, which we shall denote by $\langle\alpha\rangle_\tau$. As an $F[\mathtt{X}]$-module, $\langle\alpha\rangle_\tau$ is isomorphic to $F[\mathtt{X}]/(\phi_{\tau,\alpha})$. In particular, if $\phi_{\tau,\alpha}$ is non-zero and has degree m, then $\langle\alpha\rangle_\tau$ is a vector space of dimension m over F; indeed, the vectors $\alpha, \tau(\alpha), \ldots, \tau^{m-1}(\alpha)$ form a basis for $\langle\alpha\rangle_\tau$ over F; moreover, m is the smallest non-negative integer such that $\alpha, \tau(\alpha), \ldots, \tau^m(\alpha)$ are linearly dependent.

Observe that for any $\beta \in \langle\alpha\rangle_\tau$, we have $\phi_{\tau,\alpha} \odot \beta = 0$; indeed, if $\beta = f \odot \alpha$, then

$$\phi_{\tau,\alpha} \odot (f \odot \alpha) = (\phi_{\tau,\alpha} f) \odot \alpha = f \odot (\phi_{\tau,\alpha} \odot \alpha) = f \odot 0 = 0.$$

In the following three exercises, τ is an element of $\mathcal{L}_F(V)$, and $\odot$ is the associated scalar multiplication that makes V into an $F[\mathtt{X}]$-module.

Exercise 19.13. Let $\alpha \in V$ have minimal polynomial $f \in F[\mathtt{X}]$ under τ, and let $\beta \in V$ have minimal polynomial $g \in F[\mathtt{X}]$ under τ. Show that if $\gcd(f, g) = 1$, then

(a) $\langle \alpha \rangle_\tau \cap \langle \beta \rangle_\tau = \{0\}$, and

(b) $\alpha + \beta$ has minimal polynomial $f \cdot g$ under τ.

Exercise 19.14. Let $\alpha \in V$. Let $q \in F[\mathtt{X}]$ be a monic irreducible polynomial such that $q^e \odot \alpha = 0$ but $q^{e-1} \odot \alpha \neq 0$ for some integer $e \geq 1$. Show that q^e is the minimal polynomial of α under τ.

Exercise 19.15. Let $\alpha \in V$, and suppose that α has minimal polynomial $f \in F[\mathtt{X}]$ under τ, with $f \neq 0$. Let $g \in F[\mathtt{X}]$. Show that $g \odot \alpha$ has minimal polynomial $f/\gcd(f, g)$ under τ.

We are now ready to state the main result of this section, whose statement and proof are analogous to that of Theorem 8.40:

Theorem 19.7. *Let $\tau \in \mathcal{L}_F(V)$, and suppose that τ has non-zero minimal polynomial ϕ. Then there exists $\beta \in V$ such that the minimal polynomial of β under τ is ϕ.*

Proof. Let $\odot$ be the scalar multiplication associated with τ. Let $\phi = p_1^{e_1} \cdots p_r^{e_r}$ be the factorization of ϕ into monic irreducible polynomials in $F[\mathtt{X}]$.

First, we claim that for each $i = 1, \ldots, r$, there exists $\alpha_i \in V$ such that $\phi/p_i \odot \alpha_i \neq 0$. Suppose the claim were false: then for some i, we would have $\phi/p_i \odot \alpha = 0$ for all $\alpha \in V$; however, this means that $(\phi/p_i)(\tau) = 0$, contradicting the minimality property in the definition of the minimal polynomial ϕ. That proves the claim.

Let $\alpha_1, \ldots, \alpha_r$ be as in the above claim. Then by Exercise 19.14, each $\phi/p_i^{e_i} \odot \alpha_i$ has minimal polynomial $p_i^{e_i}$ under τ. Finally, by part (b) of Exercise 19.13, the vector

$$\beta := \phi/p_1^{e_1} \odot \alpha_1 + \cdots + \phi/p_r^{e_r} \odot \alpha_r$$

has minimal polynomial ϕ under τ. □

Theorem 19.7 says that if τ has minimal polynomial ϕ of degree $m \geq 0$, then there exists $\beta \in V$ such that

$$\beta, \tau(\beta), \ldots, \tau^{m-1}(\beta)$$

are linearly independent. From this, it immediately follows that:

Theorem 19.8. *If V has finite dimension $\ell > 0$, then for any $\tau \in \mathcal{L}_F(V)$, the minimal polynomial of τ is non-zero of degree at most ℓ.*

We close this section a simple observation. Let V be an arbitrary, non-trivial $F[\mathtt{X}]$-module with scalar multiplication $\odot$. Restricting the scalar multiplication from $F[\mathtt{X}]$ to F, we can naturally view V as an F-vector space. Let $\tau : V \to V$ be the map that sends $\alpha \in V$ to $\mathtt{X} \odot \alpha$. It is easy to see that $\tau \in \mathcal{L}_F(V)$, and that for all polynomials $f \in F[\mathtt{X}]$, and all $\alpha \in V$, we have $f \odot \alpha = f(\tau)(\alpha)$. Thus, instead of starting with a vector space and defining an $F[\mathtt{X}]$-module structure in terms of a given linear map, we can go the other direction, starting from an $F[\mathtt{X}]$-module and obtaining a corresponding linear map. Furthermore, using the language introduced in Examples 14.14 and 14.15, we see that the $F[\mathtt{X}]$-exponent of V is the ideal of $F[\mathtt{X}]$ generated by the minimal polynomial of τ, and the $F[\mathtt{X}]$-order of any element $\alpha \in V$ is the ideal of $F[\mathtt{X}]$ generated by the minimal polynomial of α under τ. Theorem 19.7 says that there exists an element in V whose $F[\mathtt{X}]$-order is equal to the $F[\mathtt{X}]$-exponent of V, assuming the latter is non-zero.

So depending on one's mood, one can place emphasis either on the linear map τ, or just talk about $F[\mathtt{X}]$-modules without mentioning any linear maps.

Exercise 19.16. Let $\tau \in \mathcal{L}_F(V)$ have non-zero minimal polynomial ϕ of degree m, and let $\phi = p_1^{e_1} \cdots p_r^{e_r}$ be the factorization of ϕ into monic irreducible polynomials in $F[\mathtt{X}]$. Let $\odot$ be the scalar multiplication associated with τ. Show that $\beta \in V$ has minimal polynomial ϕ under τ if and only if $\phi/p_i \odot \beta \neq 0$ for $i = 1, \ldots, r$.

Exercise 19.17. Let $\tau \in \mathcal{L}_F(V)$ have non-zero minimal polynomial ϕ. Show that τ is an invertible map if and only if $\mathtt{X} \nmid \phi$.

Exercise 19.18. Let F be a finite field, and let V have finite dimension $\ell > 0$ over F. Let $\tau \in \mathcal{L}_F(V)$ have minimal polynomial ϕ, with $\deg(\phi) = m$ (and of course, by Theorem 19.8, we have $m \leq \ell$). Suppose that $\alpha_1, \ldots, \alpha_s$ are randomly chosen elements of V. Let g_j be the minimal polynomial of α_j under τ, for $j = 1, \ldots, s$. Let Q be the probability that $\operatorname{lcm}(g_1, \ldots, g_s) = \phi$. The goal of this exercise is to show that $Q \geq \Lambda_F^{\phi}(s)$, where $\Lambda_F^{\phi}(s)$ is as defined in §19.3.

(a) Using Theorem 19.7 and Exercise 19.15, show that if $m = \ell$, then $Q = \Lambda_F^{\phi}(s)$.

(b) Without the assumption that $m = \ell$, things are a bit more challenging. Adopting the matrix-oriented point of view discussed at the end of §19.3, and transposing everything, show that

– there exists $\pi \in \mathcal{D}_F(V)$ such that the sequence $(\pi \circ \tau^i)_{i=0}^{\infty}$ has minimal polynomial ϕ, and
– if, for $j = 1, \ldots, s$, we define h_j to be the minimal polynomial of the sequence $(\pi(\tau^i(\alpha_j)))_{i=0}^{\infty}$, then the probability that $\operatorname{lcm}(h_1, \ldots, h_s) = \phi$ is equal to $\Lambda_F^{\phi}(s)$.

(c) Show that $h_j \mid g_j$, for $j = 1, \ldots, s$, and conclude that $Q \geq \Lambda_F^{\phi}(s)$.

Exercise 19.19. Let $f, g \in F[\mathtt{X}]$ with $f \neq 0$, and let $h := f/\gcd(f, g)$. Show that $g \cdot F[\mathtt{X}]/(f)$ and $F[\mathtt{X}]/(h)$ are isomorphic as $F[\mathtt{X}]$-modules.

Exercise 19.20. In this exercise, you are to derive the **fundamental theorem of finite dimensional $F[\mathtt{X}]$-modules**, which is completely analogous to the fundamental theorem of finite abelian groups. Both of these results are really special cases of a more general decomposition theorem for modules over a principal ideal domain. Let V be an $F[\mathtt{X}]$-module. Assume that as an F-vector space, V has finite dimension $\ell > 0$, and that the $F[\mathtt{X}]$-exponent of V is generated by the monic polynomial $\phi \in F[\mathtt{X}]$ (note that $1 \leq \deg(\phi) \leq \ell$). Show that there exist monic, non-constant polynomials $\phi_1, \ldots, \phi_t \in F[\mathtt{X}]$ such that

- $\phi_i \mid \phi_{i+1}$ for $i = 1, \ldots, t-1$, and
- V is isomorphic, as an $F[\mathtt{X}]$-module, to the direct product of $F[\mathtt{X}]$-modules

$$V' := F[\mathtt{X}]/(\phi_1) \times \cdots \times F[\mathtt{X}]/(\phi_t).$$

Moreover, show that the polynomials $\phi_1, \ldots, \phi_t$ satisfying these conditions are uniquely determined, and that $\phi_t = \phi$. Hint: one can just mimic the proof of Theorem 8.44, where the exponent of a group corresponds to the $F[\mathtt{X}]$-exponent of an $F[\mathtt{X}]$-module, and the order of a group element corresponds to the $F[\mathtt{X}]$-order of an element of an $F[\mathtt{X}]$-module — everything translates rather directly, with just a few minor, technical differences, and the previous exercise is useful in proving the uniqueness part of the theorem.

Exercise 19.21. Let us adopt the same assumptions and notation as in Exercise 19.20, and let $\tau \in \mathcal{L}_F(V)$ be the map that sends $\alpha \in V$ to $\mathtt{X} \odot \alpha$. Further, let $\sigma : V \to V'$ be the isomorphism of that exercise, and let $\tau' \in \mathcal{L}_F(V')$ be the $\mathtt{X}$-multiplication map on V'.

(a) Show that $\sigma \circ \tau = \tau' \circ \sigma$.

(b) From part (a), derive the following: there exists an ordered basis for V over F, with respect to which the matrix representing τ is the

"block diagonal" matrix

$$T = \begin{pmatrix} C_1 & & & \\ & C_2 & & \\ & & \ddots & \\ & & & C_t \end{pmatrix},$$

where each C_i is the companion matrix of ϕ_i (see Example 15.1).

Exercise 19.22. Let us adopt the same assumptions and notation as in Exercise 19.20.

(a) Using the result of that exercise, show that V is isomorphic, as an $F[\mathtt{X}]$-module, to a direct product of $F[\mathtt{X}]$-modules

$$F[\mathtt{X}]/(p_1^{e_1}) \times \cdots \times F[\mathtt{X}]/(p_r^{e_r}),$$

where the p_i are monic irreducible polynomials (not necessarily distinct) and the e_i are positive integers, and this direct product is unique up to the order of the factors.

(b) Using part (a), show that there exists an ordered basis for V over F, with respect to which the matrix representing τ is the "block diagonal" matrix

$$T' = \begin{pmatrix} C'_1 & & & \\ & C'_2 & & \\ & & \ddots & \\ & & & C'_r \end{pmatrix},$$

where each C'_i is the companion matrix of $p_i^{e_i}$.

Exercise 19.23. Let us adopt the same assumptions and notation as in Exercise 19.20.

(a) Suppose $\alpha \in V$ corresponds to $([f_1]_{\phi_1}, \ldots, [f_t]_{\phi_t}) \in V'$ under the isomorphism of that exercise. Show that the $F[\mathtt{X}]$-order of α is generated by the polynomial

$$\operatorname{lcm}(\phi_1/\gcd(f_1, \phi_1), \ldots, \phi_t/\gcd(f_t, \phi_t)).$$

(b) Using part (a), give a short and simple proof of the result of Exercise 19.18.

19.7 Notes

Berlekamp [15] and Massey [60] discuss an algorithm for finding the minimal polynomial of a linearly generated sequence that is closely related to the one presented in §19.2, and which has a similar complexity. This connection between Euclid's algorithm and finding minimal polynomials of linearly generated sequences has been observed by many authors, including Mills [64], Welch and Scholtz [102], and Dornstetter [35].

The algorithm presented in §19.3, is due to Wiedemann [103], as are the algorithms for solving sparse linear systems in §19.4, as well as the statement and proof outline of the result in Exercise 19.18.

Our proof of Theorem 19.5 is based on an exposition by Morrison [65].

Using fast matrix and polynomial arithmetic, Shoup [91] shows how to implement the algorithms in §19.5 so as to use just $O(\ell^{(\omega+1)/2})$ operations in F, where ω is the exponent for matrix multiplication (see §15.6), and so $(\omega + 1)/2 < 1.7$.