Introduction to admissible representations of p-adic groups

Bill Casselman
University of British Columbia
cass@math.ubc.ca

The structure of GL(n)

The structure of arbitrary reductive groups over a $\mathfrak{p}$-adic field $\mathfrak{k}$ is an intricate subject. But for GL_n, although things are still not trivial, results can be verified directly. and another is to make available at least one example One reason for doing this is to motivate the abstract definitions to come, for which the theory is not trivial and for which the abstract theory is not necessary.

An important role in representation theory is played by the linear transformations which just permute the basis elements of a vector space, and I therefore begin this essay by discussing these as well as permutations in general.

Next, I'll look at the groups $\mathrm{GL}_n(F)$ and $\mathrm{SL}_n(F)$ for an arbitrary coefficient field F. This is principally because at some point later on we shall want to make calculations in the finite group $\mathrm{GL}_n(\mathbb{F}_q)$ as well as the $\mathfrak{p}$-adic group $\mathrm{GL}_n(\mathfrak{k})$. But also because procedures to deal with $\mathfrak{p}$-adic matrices are very similar to those for matrices over arbitrary fields.

Then, in the third part of this essay I'll let $F = \mathfrak{k}$ and look at the extra structures that arise in $\mathrm{GL}_n(\mathfrak{k})$ and $\mathrm{SL}_n(\mathfrak{k})$. Much of this is encoded in the geometry of the buildings attached to them, which will be dealt with elsewhere. This part of the essay is much different from an earlier version, in which I grossly over-simplified things.

In any n-dimensional vector space over a field F, let ε_i be the i-th basis element in the standard basis ε of F^n. Its i-th coordinate is 1 and all others are 0.

Contents

Part I. The symmetric group

1. Permutations

A **permutation** of a finite set I is just an invertible map from I to itself. The permutations of I form a group $\mathfrak{S}_I$, which will be written as $\mathfrak{S}_n$ if $I = [1, n]$. In fact, I shall assume that $I = I_n$ unless specifically said otherwise.

There are several ways to express a permutation σ. One is just by writing the **ordered permuted array** $((\sigma(\iota_i)))$. Another is as a product of **cycles**. For example, the permutation whose array is $((3, 4, 5, 2, 1))$ is expressed as the cycle product $(1 : 3 : 5)(2 : 4)$.

INVERSIONS. If σ is a permutation of I then its **inversions** make up the set of pairs

$$I(\sigma) = \{(i, j) \,|\, i < j, \sigma(i) > \sigma(j)\}\,.$$

For reasons that will become clear, the number of inversions of σ is known as its **length**, and written as $\ell(\sigma)$. The inversions of a permutation σ can be read off directly from the permutation array of σ—one first lists all $\big(\sigma(1), \sigma(i)\big)$ with $1 < i$ and $\sigma(1) > \sigma(i)$, then all similar pairs $\big(\sigma(2), \sigma(i)\big)$, etc. For example, the number of inversions of $((3, 4, 5, 2, 1))$ is $2 + 2 + 2 + 1 = 7$.

Remark. There is a close relationship between nethods of counting inversions of a permutation, and methods of sorting an array of integers. The method of counting inversions described above requires $n(n-1)/2$ comparisons, and is essentially the well known **bubble sort**, which rearranges any array in increasing order. But recursive merge sorting (as explained in S5.2.4 of [Knuth:1973]) can compute the number of inversions in time proportional to $n \log n$. To get a rough idea of how this works, suppose you want to count inversions of the array $[2, 6, 3, 4, 5, 1]$. • You split it into two halves $[2, 6, 3]$ and $[4, 5, 1]$, and sort each of these, by recursion, obtaining inversion counts along the way. You thus arrive at a new pair, here $((2, 3, 6))$ and $((1, 4, 5))$.

• You then merge the two sorted halves into a single sorted list. This means looking in turn at their first elements, then removing the least of the two from its half and adding it to a new array. Here, you start off the new array with 1 because $1 < 2$, chosen from the second array. Since the choice is from the second array, 1 comes after 2, and there is at least one evident inversion. But at that point you know not only that $1 < 2$, but that 1 is less than every item in the first array, all of which are at least 2. So the inversion count increments by 3, the length of the first array.

∘ ———— ∘

A **transposition** $(i : j)$ in $\mathfrak{S}_n$ just interchanges i and j, and an **elementary transposition** $s_j = (j : j + 1)$ interchanges two neighbouring items. In general, multiplying a permutation on the right affects the locations in the permutation array, while multiplying on the left affects the items. Thus

$$((3, 4, 5, 1, 2))\,(2 : 3) = ((3, \boldsymbol{5}, \boldsymbol{4}, 1, 2))\,,$$

but

$$(2 : 3)\,((3, 4, 5, 1, 2)) = ((\boldsymbol{2}, 4, 5, 1, \boldsymbol{3}))\,.$$

1.1. Proposition. *(a) The only permutation with no inversions is the identity map;*
(b) if σ is any permutation other than the trivial one, there exists i such that $\sigma(i) > \sigma(i+1)$;
(c) an elementary transposition s_i inverts only the single pair $(i, i+1)$;
(d) an elementary transposition s_i permutes all pairs $j < k$ other than $i < i + 1$;
(e) if $\sigma(i) < \sigma(i+1)$, then

$$I(\sigma s_i) = s_i T(\sigma) \sqcap \{(i, i+1)\}\,;$$

(f) if $\sigma(i) > \sigma(i+1)$, then

$$I(\sigma s_i) = s_i(I(\sigma) - \{(i, i+1)\})\,.$$

(g)

$$I(\sigma) = \sigma I(\sigma^{-1})\,.$$

Proof. To prove (a), it suffices to prove (b), which is just a reformulation. Suppose $\sigma(i) < \sigma(i+1)$ for all $i < n$. If $k_i = \sigma(i+1) - \sigma(i)$ then $\sigma(n) - \sigma(1) = \sum_1^{n-1} k_i$, but since also $1 \le \sigma(1)$, $\sigma(n) \le n$ we must have all $k_i = 1$.

For (c), (d), (e), and (f) the pairs (j,k) with $j < k$ can be partitioned into those with (i) neither j nor k equal to i or $i+1$; (ii) $j < i$ or $j < i+1$; (iii) $i < k$ or $i+1 < k$; (iv) $i < i+1$. The transposition s_i inverts only the last pair.

For (g), σ inverts (i,j) if and only if σ^{-1} inverts $\big(\sigma(j), \sigma(i)\big)$. So $I(\sigma^{-1}) = \sigma I(\sigma)$. ∎

1.2. Corollary. *For any σ in $\mathfrak{S}_n$*

$$\ell(\sigma s_i) = \begin{cases} \ell(\sigma) + 1 & \text{if } \sigma(i) < \sigma(i+1) \\ \ell(\sigma) - 1 & \text{if } \sigma(i) > \sigma(i+1). \end{cases}$$

Consequently:

1.3. Corollary. *The map $\sigma \mapsto (-1)^{\ell(\sigma)}$ is a homomorphism from $\mathfrak{S}_n$ to $\{\pm 1\}$.*

Every simple cycle may be written as a product of transpositions:

$$(i_1 : i_2 : i_3 : \ldots : i_{m-1} : i_m) = (i_1 : i_2)(i_2 : i_3) \ldots (i_{m-1} : i_m)\,.$$

In particular, if j and k are neighbours—i.e. if $k = j \pm 1$—then $s_j s_k$ is a cycle of order three, but otherwise s_j and s_k commute and the product has order two:

$$\begin{aligned} s_i s_{i+1} &= (i : i+1)(i+1 : i+2) = (i : i+1 : i+2) \\ s_j s_k &= (j : j+1)(k : k+1) \\ &= s_k s_j\,. \end{aligned}$$

Since every cycle may be expressed as a product of transpositions and every permutation may expressed as a product of cycles, every permutation may also be expressed as a product of transpositions. Better:

1.4. Proposition. *Every permutation in $\mathfrak{S}_n$ may be expressed as a product of $\ell(\sigma)$ elementary transpositions.*

Proof. The proof will describe an algorithm to find such a product by induction on $\ell(\sigma)$.

There is nothing to prove if σ is the trivial permutation. So now suppose σ to be other than trivial. According to (b) of Proposition 1.1 it must have at least one inversion. Read along in the array $\big(\sigma(i)\big)$ until you find i with $\sigma(i) > \sigma(i+1)$. Let $\tau = \sigma\, s_i$. Then $(i, i+1)$ is not an inversion for τ, so by (d) of Proposition 1.1 the number of inversions for τ is exactly one less than for σ itself. We can keep on applying these swaps, at each point multiplying on the right by an elementary transposition. The number of inversions decreases at every step, and hence it must stop, which it does only when the permutation we are considering is trivial. So we get an equation

$$\sigma s_{i_1} \ldots\, s_{i_n} = I, \quad \sigma = s_{i_n} \ldots\, s_{i_1}\,.$$ ∎

1.5. Corollary. *The minimal length of an exspression of σ as a product of elementary transpositions is the same as $\ell(\sigma)$, the number of its inversions.*

Proof. The Proposition asserts that $\ell(\sigma) \le |I(\sigma)|$. But implies that $|I(\sigma)| \le \ell(\sigma)$. ∎

1.6. Corollary. *Two permutations σ and τ are equal if and only if $I(\sigma) = I(\tau)$.*

For example, the permutation with array $(n, n-1, \ldots, 1)$ is the unique permutation that inverts all pairs.

In the next section I'll exhibit a geometric explanation for the relationship between $\ell(\sigma)$ and $I(\sigma)$.

2. The geometry of permutations

The group $\mathfrak{S}_n$ acts on the n-dimensional vector space $\mathbb{R}^n$ by permuting basis elements:

$$\sigma\colon\ \varepsilon_i \longmapsto \varepsilon_{\sigma(i)}\ .$$

Its matrix w_σ is that with the corresponding permutation of the columns of I. Thus

$$\sigma\Big(\sum x_i\varepsilon_i\Big) = \sum x_i\varepsilon_{\sigma(i)} = \sum x_{\sigma^{-1}(i)}\varepsilon_i\ . \tag{2.1}$$

In other words, it permutes the coordinates of a vector. We thus have an embedding of $\mathfrak{S}_n$ into $\mathrm{GL}_n(\mathbb{R})$. There is also a dual representation on the space of linear functions on V—by definition

$$\langle\widehat{\sigma}(f), v\rangle = \langle f, \sigma^{-1}(v)\rangle\ .$$

Among these linear functions are the coordinate functions x_i, and then

$$\langle\widehat{\sigma}(x_i), v\rangle = \langle x_i, \sigma^{-1}(v)\rangle$$

so that $\widehat{\sigma}(x_i) = x_{\sigma(i)}$, in accordance with (2.1) .

A transposition $(i\colon j)$ swaps coordinates x_i and x_j, hence has determinant -1.

2.2. Proposition. *If σ can be represented as a product of m transpositions then $\det(w_\sigma) = (-1)^m$;*

This gives a second proof that the map $\sigma \mapsto (-1)^{\ell(\sigma)}$ is a homomorphism.

Proof. Because $\det(xy) = \det(x)\det(y)$. ∎

ROOTS. Some of the basic facts about combinatorics in $\mathfrak{S}_n$ and related groups are explained intuitively by applying the ridiculously simple fact that if f is a linear function on a real vector space V then the regions where $f < 0$ and $f > 0$ are separated by the hyperplane $f = 0$, which means that if $f(P) < 0$ and $f(Q) > 0$ any line from P to Q will possess exactly one point where f vanishes.

Assign to $\mathbb{R}^n$ the usual Euclidean inner product. Let $\lambda = \lambda_{i,j}$ be the linear function $x_i - x_j$ on V. The functions $\lambda_{i,j}$ are called the **roots** of the group $\mathfrak{S}_n$. The ones with $i < j$ are called **positive roots**. The roots $\alpha_i = \lambda_{i,i+1}$ for $i = 0, \dots, n-1$ are called the **simple** roots. Every positive root is an integral linear combination of simple roots, with non-negative coefficients, since for $i < j$

$$\lambda_{i,j} = \alpha_i + \cdots + \alpha_{j-1}$$

The **dominant** root $\widetilde{\alpha} = \lambda_{1,n-1}$ is the sum of all simple roots. This terminology is motivated by a relationship with the structure of the Lie algebra of GL_n.

The group $\mathfrak{S}_n$ permutes the roots: σ takes $\lambda_{i,j}$ to $\lambda_{\sigma(i),\sigma(j)}$. The linear transformation $s_{\lambda_{i,j}}$ determined by the transposition $(i\colon j)$ amounts to Euclidean reflection in the hyperplane $x_i = x_j$. It fixes all points in this hyperplane and acts as scalar multiplication by -1 on the line perpendicular to it.

The coordinates of any vector may be permuted to a unique weakly decreasing array. In other words:

2.3. Proposition. *The region $\overline{C}$ in which $x_1 \geq x_2 \geq \ldots \geq x_n$ is a fundamental domain for the action of $\mathfrak{S}_n$ on $\mathbb{R}^n$.*

This is the same as the region where $\alpha_i \geq 0$ for $0 \leq i < n$, and its walls are open subsets of the root hyperplanes $\alpha_i = 0$. Its faces C_Θ are parametrized by subsets $\Theta \subseteq \Delta = \{0, \dots n-1\}$—to Θ corresponds the face where $\alpha_i = 0$ for i in Θ, otherwise $\alpha_i > 0$. The open cone itself is $C_\emptyset$.

Remark. There is a somewhat arbitrary choice involved here. Equally valid, and valuable in some circumstances, is that of weakly increasing arrays as the fundamental domain. I'll probably shift between these two choices without explicit comment.

∘ ———— ∘

Any vector in v is the permutation of a vector in a unique face of C. The vector space V is therefore partitioned into transforms of these faces labeled by subsets of Δ. The transforms of $C = C_\emptyset$ are the connected components of the complement of the root hyperplanes. They are called **chambers** of the partition. Every face of a chamber is the transform by $\mathfrak{S}_n$ of a unique face C_Θ of C, and in particular every wall is the transform of some unique C_{α_i}. A root γ is positive if and only if $\gamma > 0$ on C.

Every root γ corresponds to a hyperplane $\gamma = 0$, and conversely, every such **root hyperplane** is the zero set of two roots, exactly one of which is positive. A root λ separates C and $\sigma(C)$ if and only if $\lambda(C) > 0$ and $\lambda\big(\sigma(C)\big) < 0$, or equivalently if and only if $\lambda > 0$ and $\sigma^{-1}(\lambda) < 0$.

2.4. Proposition. *Suppose σ to be a permutation and $i < j$. The following are equivalent:*

(a) the root $\sigma(\lambda_{i,j})$ is negative;
(b) the hyperplane $\lambda_{i,j} = 0$ separates C from $\sigma(C)$;
(c) the pair (i, j) is an inversion for σ.

Proof. (a) and (b) are equivalent by definition. As for the remaining equivalence, the vector

$$\rho = (n, n-1, \dots, 2, 1)$$

lies in C, so $\gamma < 0$ if and only if $\langle \gamma, \rho \rangle < 0$. The i-th coordinate of ρ is $n - i + 1$. But then

$$\begin{aligned} \langle \sigma(\lambda_{i,j}), \rho \rangle &= \langle \lambda_{i,j}, \sigma^{-1}(\rho) \rangle \\ &= (n - \sigma(i) + 1) - (n - \sigma(j) + 1) \\ &= \sigma(j) - \sigma(i) \end{aligned}$$

so that $\sigma(\lambda_{i,j}) < 0$ if and only if $\sigma(i) > \sigma(j)$. ∎

In other words:

2.5. Theorem. *The number of inversions of σ is exactly the same as the number of root hyperplanes separating C from $\sigma(C)$.*

To understand more precisely how geometry explains that the length of σ is the cardinality of $|I(\sigma)|$ we must find a geometric interpretation for products of elementary transpositions.

This involves the notion of **galleries**. Two chambers are **neighbours** if they are distinct and share a wall. The neighbours of C are the sC for $s = s_i$. If w is an arbitrary permutation, then wC and wsC are neighbours, and all are of this kind. If s_1, s_2, etc. is a sequence of elementary transpositions, then C, s_1C, s_1s_2C, etc. is a chain of neighbouring chambers. A gallery is a chain of neighbouring chambers. Elements of $\mathfrak{S}_n$ transform one gallery into another. The wall separating wC and wsC is the w-transform of $\alpha_s = 0$, hence the root hyperplane the $\gamma = 0$ where $\gamma = w\alpha_s$. If wC is its terminal chamber the gallery corresponds to an expression for w as a product of elementary reflections.

2.6. Proposition. *Reduced expressions correspond to galleries that cross any given root hyperplane at most once.*

Proof. Suppose we are given a gallery that crosses a hyperplane twice—say the first crossing is from xC to xsC while the second is back again from $xsyC$ to $xsytC$. Applying x^{-1} to this segment, we see that we have a gallery from C to $sytC$. The two galleries syC and $sytC$ are neighbours, by assumption separated by the same hyperplane, must therefor be the root hyperplane $\alpha_s = 0$. Reflection by s must therefore interchange them, so syC and yt are the same chamber, hence $sy = yt$, $syt = y$, and the gallery with intermediate segment from xC to xyC has the same terminal as the original, but has length 2 less. So the original expression was not reduced.

On the other hand, any gallery from C to wC must cross every separating hyperplane, so the length is at least the number of such hyperplanes. ∎

3. Permutations in GL(n)

The matrix w_σ of the linear transformation associated to σ is the **permutation matrix** with i-th column equal to $\varepsilon_{\sigma(i)}$. Thus

$$(w_\sigma)_{i,j} = \begin{cases} 1 & \text{if } i = \sigma(j) \\ 0 & \text{otherwise.} \end{cases}$$

If $m = (m_{i,j})$ is an $n \times n$ matrix then multiplication on the left by the permutation matrix w permutes its rows, and multiplication on the right permutes its columns. Explicitly:

$$(w_\sigma m)_{i,j} = m_{\sigma^{-1}(i),j}$$
$$(m w_\tau)_{i,j} = m_{i,\tau(j)}$$

and hence

3.1. Lemma. *We have*

$$(w_\sigma m w_\tau^{-1})_{i,j} = m_{\sigma^{-1}(i),\tau^{-1}(j)} \, .$$

Conjugation by a permutation matrix permutes the diagonal entries of a diagonal matrix. Conversely, suppose that conjugation by x leaves stable the group of diagonal matrices. This means that $xax^{-1} = b$ is diagonal for all diagonal a. Thus $xa = bx$. In the 2×2 case, for example, we would have

$$\begin{bmatrix} x_{1,1} & x_{1,2} \\ x_{2,1} & x_{2,2} \end{bmatrix} \begin{bmatrix} a_1 & 0 \\ 0 & a_2 \end{bmatrix} = \begin{bmatrix} a_1 x_{1,1} & a_2 x_{1,2} \\ a_1 x_{2,1} & a_2 x_{2,2} \end{bmatrix} = \begin{bmatrix} b_1 & 0 \\ 0 & b_2 \end{bmatrix} \begin{bmatrix} x_{1,1} & x_{1,2} \\ x_{2,1} & x_{2,2} \end{bmatrix} = \begin{bmatrix} b_1 x_{1,1} & b_1 x_{1,2} \\ b_2 x_{2,1} & b_2 x_{2,2} \end{bmatrix} \, .$$

Then $x_{i,j} a_j = b_i x_{i,j}$ for all i, j. Since x is non-singular, in every row of x there exists at least one entry $x_{i,j(i)} \neq 0$. Thus for all i, j we have

$$\begin{aligned} a_j x_{i,j} &= b_i x_{i,j} \\ a_{j(i)} x_{i,j(i)} &= b_i x_{i,j(i)} \\ a_{j(i)} &= b_i \\ \left(\frac{a_j}{a_{j(i)}} \right) x_{i,j} &= x_{i,j} \end{aligned}$$

We are free to choose the a_j arbitrarily, so we see that $x_{i,j} = 0$ for $j \neq j(i)$. Therefore in each row only one entry is non-zero. In short, x is the product of permutation and diagonal matrices. All in all:

3.2. Proposition. *If A is the group of diagonal matrices in $G = \mathrm{GL}_n$ then the permutation matrices induce an isomorphism of $\mathfrak{S}_n$ with $N_G(A)/A$.*

The contents of this section elaborate in a special case what a later chapter will say about general root systems. The group A acts by conjugation on the Lie algebra of GL_n, which is the vector space of matrices M_n. It acts trivially on the diagonal matrices, and the complement decomposes into a direct sum of A-stable spaces $M_{i,j}$ ($i \neq j$) spanned by the single matrix $e_{i,j}$ with a single non-zero entry in row i, column j. The corresponding eigencharacter is a_i/a_j. Let $X^*(A)$ be the lattice of characters of A, the algebraic homomorphisms from A to the multiplicative group $\mathbb{G}_m$. It has as basis the characters

$$\varepsilon_i \colon \, a \longmapsto a_i \, .$$

Multiplication of characters is written additively—the character $a \mapsto a_i/a_j$ is $\gamma_{i,j} = \varepsilon_i - \varepsilon_j$.

Assign $V = X^*(A) \otimes \mathbb{R}$ the Euclidean norm in which the ε_i form an orthonormal basis. Let α_i for $1 \leq i < n$ be the root $\varepsilon_{i+1} - \varepsilon_i$. Every root can be written as a unique integral combination of the α_i. The dominant root is

$$\widetilde{\alpha} = \sum_{1 \leq i < n} \alpha_i = (1, \ldots, 1) \, .$$

The vector ρ lies in what is in these circumstances the **positive chamber** where all $\alpha_i \geq 0$ or equivalently

$$x_1 \leq x_2 \leq \ldots \leq x_n \ .$$

To each root $\gamma = \gamma_{p,q}$ corresponds an embedding, which I express as $\gamma^\vee$ in spite of the obvious conflict, of SL_2 into GL_n. The matrix

$$x = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

has as image the matrix y with

$$\begin{aligned} y_{i,j} &= \delta_{i,j} \quad \text{unless } \{i,j\} = \{p,q\} \\ y_{p,p} &= a, \quad y_{p,q} = b, \quad y_{q,p} = c, \quad y_{q,q} = d \ . \end{aligned}$$

The conflict of notation is that $\gamma^\vee$ now denotes an embedding of both the multiplicative group and of SL_2 into GL_n. The two uses are at least weakly consistent, since the embedding of the multiplicative group can be factored through SL_2. At any rate, there will be no serious problem since the two maps $\gamma^\vee$ can be distinguished by what sort of things they are applied to. (This is called *operator overloading* in programming.)

Part II. The Bruhat decomposition

4. Gauss elimination

In this section, let

$$\begin{aligned} G &= \mathrm{GL}_n(F) \\ B &= \text{ the subgroup of upper triangular matrices} \\ N &= \text{ the upper triangular unipotent matrices} \\ \overline{N} &= \text{ the lower triangular unipotent matrices} \\ A &= \text{ the diagonal matrices} \\ W &= \mathfrak{S}_n \ . \end{aligned}$$

PERMUTED ECHELON FORM. A rectangular matrix is said to be in **permuted echelon form** if it has the following two properties:

- the last non-zero entry in each column is 1;
- all entries to the right of such an entry (which I call a **pivot**) vanish.

It may happen that there are no such entries in a column—i.e., that all its entries vanish.

For example, all such 2×1 matrices are of the form

$$\begin{bmatrix} \circ \\ \circ \end{bmatrix} \quad \begin{bmatrix} * \\ 1 \end{bmatrix} \quad \begin{bmatrix} 1 \\ \circ \end{bmatrix} \ .$$

The relevant 2×2 matrices are:

$$\begin{bmatrix} \circ & \circ \\ \circ & \circ \end{bmatrix} \quad \begin{bmatrix} * & \circ \\ 1 & \circ \end{bmatrix} \quad \begin{bmatrix} 1 & \circ \\ \circ & \circ \end{bmatrix} \quad \begin{bmatrix} \circ & * \\ \circ & 1 \end{bmatrix} \quad \begin{bmatrix} \circ & 1 \\ \circ & \circ \end{bmatrix} \quad \begin{bmatrix} * & 1 \\ 1 & \circ \end{bmatrix} \quad \begin{bmatrix} 1 & \circ \\ \circ & 1 \end{bmatrix} \ .$$

In all these, $*$ is arbitrary.

There is a simple recursive criterion for an $r \times c$ matrix to be in permuted echelon form. It follows immediately from the definition.

(a) Either the first column vanishes identically, or the last non-zero entry in the first column is equal to 1. In the first case, it is required that the matrix made up of the remaining columns, which is of smaller size than the original, be in permuted echelon form. In the second case, suppose this first non-zero entry to be in row i.
(b) All entries in row i and subsequent columns must vanish.
(c) The matrix extracted from it by removing the first column and row i must be in permuted echelon form.

This leads easily to an algorithm for executing the main result of this section:

4.1. Theorem. *If m is any matrix, then there exists a unique matrix in permuted echelon form that can be obtained from it by multiplying on the right by an invertible upper triangular matrix.*

Proof. I shall prove the existence here, but deal with uniqueness in a later section.

The proof given here, which amounts to an algorithm, is a straightforward variant of Gauss elimination.

Multiplying on the right by an invertible upper triangular matrix amounts to performing some combination of these two elementary column operations: (a) multiplying a column by a non-zero constant, or (b) adding to one column a multiple of an earlier one. Something similar is true for row operations effected by multiplying on the left. For example

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} 1 & x \\ \circ & 1 \end{bmatrix} = \begin{bmatrix} a & ax+b \\ c & cx+d \end{bmatrix}$$
$$\begin{bmatrix} 1 & x \\ \circ & 1 \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} a+cx & b+dx \\ c & d \end{bmatrix}.$$

I now proceed by induction on the number of columns, say k. If $k = 1$, just multiply the column by the inverse of the last non-zero entry. Otherwise, suppose m to have $k > 1$ columns. Locate the last non-zero entry in the first column, if there is one. (a) Multiply the column by its inverse to make it 1; (b) subtract suitable multiples of the first column from subsequent columns to make the last $k-1$ entries in the i-th row vanish. Then apply induction to the matrix made up of those column, with row i extracted. ∎

CLASSIFICATION. I now want to look more closely at the structure of matrices in permuted echelon form. It will help to have an example in view:

$$\begin{bmatrix} * & * & \circ & * & * \\ * & 1 & \circ & \circ & \circ \\ * & \circ & \circ & * & 1 \\ 1 & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & * & \circ \\ \circ & \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & \circ & \circ \end{bmatrix}.$$

• A matrix in permuted echelon form may have any number of vanishing columns—i.e. columns all of whose entries are 0, and they can be anywhere. They exert no restriction on the rest of the matrix, so in classifying matrices in permuted echelon form we need look only at matrices of size (say) $r \times c$ with no vanishing columns. Like this:

$$\begin{bmatrix} * & * & * & * \\ * & \mathbf{1} & \circ & \circ \\ * & \circ & * & \mathbf{1} \\ \mathbf{1} & \circ & \circ & \circ \\ \circ & \circ & * & \circ \\ \circ & \circ & \mathbf{1} & \circ \\ \circ & \circ & \circ & \circ \end{bmatrix}.$$

• In each column there is now a unique pivot, below which all entries vanish. These are emphasized in the matrix above. Let r_j be the row in which the pivot for coliumn j occurs. Associated to the matrix is therefore a permutation matrix σ:

$$\sigma_{i,j} = \begin{cases} 1 & \text{if } i = r_j \\ 0 & \text{otherwise.} \end{cases}$$

• Given a parmutation matrix σ of size $c \times c$, and a row assignment (r_j), it is easy to describe all possible $r \times c$ matrices of permuted echelon form to which it is associated. First of all, write down the embedded copy of σ, that follwos the rule just above. Then set every entry to the right of a pivot equal to 0, and likewise every entry below a pivot.

$$\begin{bmatrix} & & & \\ & \mathbf{1} & \circ & \circ \\ & \circ & & \mathbf{1} \\ \mathbf{1} & \circ & \circ & \circ \\ \circ & \circ & & \circ \\ \circ & \circ & \mathbf{1} & \circ \\ \circ & \circ & \circ & \circ \end{bmatrix}.$$

• This leaves a certain number of entries undetermined. These are arbitrary.

• To summarize: the basic form of a matrix in permuted echelon form is determined by its size, its rank (less than or equal to the number of columns), and the row assignments of its embedded pernutation matrix.

PARTIAL PERMUTATION MATRICES. To see what happens next, I'll look at an example. At the end of the process sketched above, we shall be looking at a matrix like this:

$$\begin{bmatrix} * & * & \circ & * & * \\ * & 1 & \circ & \circ & \circ \\ * & \circ & \circ & * & 1 \\ 1 & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & * & \circ \\ \circ & \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & \circ & \circ \end{bmatrix}.$$

I claim that we can now multiply this on the left by upper triangular unipotent matrices in order to obtain what I call a **partial permutation matrix**. This is a special kind of permuted echelon matrix, in which each column and each row has at most one non-zero entry, which is equal to 1.

This will involve applying elementary row operations, but we have to be careful about the order in which I do thihgs. To simplify things slightly, I may assume that none of the columns vanishes, because that will not affect the process significantly. So we start with

$$\begin{bmatrix} * & * & * & * \\ * & 1 & \circ & \circ \\ * & \circ & * & 1 \\ 1 & \circ & \circ & \circ \\ \circ & \circ & * & \circ \\ \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & \circ \end{bmatrix}.$$

In rough terms, we look at one pivot after another, starting at the bottom, and working up. When we are looking at a particular pivot, there may be non-zero entries in the pivot's column. By definition, these can only be above the pivot. We can make them vanish by applying a simple row operation to the matrix, multiplying bon the left by an upper triangular unipotent matrix. I claim that at the end we have a partial permutation matrix.

Why is this? If m is a matrix in parmuted echelon form, I'll call one of its rows **clean** if either (a) all its entries vanish, or (b) it is a row containing a pivot. I'll call a column clean if its only non-vanishing entry is a pivot. The proof of my claim reduces to the claim that when we are looking at a particular pivot all rows up to and including the row it lies in is clean, and all columns belonging to lower pivots are clean.

For example, starting with the matrix above, we get in turn

$$\begin{bmatrix} * & * & \circ & * \\ * & 1 & \circ & \circ \\ * & \circ & \circ & 1 \\ 1 & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ \\ \circ & \circ & \mathbf{1} & \circ \\ \circ & \circ & \circ & \circ \end{bmatrix} \longrightarrow \begin{bmatrix} \circ & * & \circ & * \\ \circ & 1 & \circ & \circ \\ \circ & \circ & \circ & 1 \\ \mathbf{1} & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ \\ \circ & \circ & \mathbf{1} & \circ \\ \circ & \circ & \circ & \circ \end{bmatrix} \longrightarrow \begin{bmatrix} \circ & * & \circ & \circ \\ \circ & 1 & \circ & \circ \\ \circ & \circ & \circ & \mathbf{1} \\ \mathbf{1} & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ \\ \circ & \circ & \mathbf{1} & \circ \\ \circ & \circ & \circ & \circ \end{bmatrix} \longrightarrow \begin{bmatrix} \circ & \circ & \circ & \circ \\ \circ & \mathbf{1} & \circ & \circ \\ \circ & \circ & \circ & \mathbf{1} \\ \mathbf{1} & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ \\ \circ & \circ & \mathbf{1} & \circ \\ \circ & \circ & \circ & \circ \end{bmatrix}.$$

I leave details of the proof of the following result as an exercise.

4.2. Theorem. *Given any matrix m, there exist an upper triangular matrix u, an upper triangular unipotent matrix v, and a partial permutation matrix w such that $m = vwu$.*

It will follow from results in the later discussion on flags that w is unique. In some cases, we shall also see uniqueness results for u and $\overline{v}$.

Recall that $G = \mathrm{GL}_n(F)$. The product $G \times G$ acts on left and right on the set of matrices of a given rank. The corollary asserts that this action is transitive.

We have exhibited a decomposition of the set of matrices into double cosets $B\backslash G/B$, in which B is the subgroup of upper triangular matrices. But if w_ℓ is the skew-symmetric permutation matrix swapping e_i with e_{n+1-i}, then $\overline{B} = w_\ell B w_\ell$ is the subgroup of lower triangular matrices, and we have

$$G = \bigsqcup_w BwBw_\ell = \bigsqcup_w Bw \cdot w_\ell Bw_\ell = \bigsqcup_w Bw\overline{B}\,.$$

In many situations this is particularly useful.

One final remark: the algorithm laid out above shows that an invertible matrix g is in the largest double coset $Bw_\ell B$ if and only if the square matrices extracted from the lower left of g are all invertible.

5. Flags

Fix a basis (e_i) of $E = F^n$. In this section, I'll prove the claim of uniqueness in Theorem 4.1.

A **flag** $\mathcal{V}$ in E is a weakly increasing sequence of subspaces

$$\mathcal{V}_1 \subseteq \mathcal{V}_2 \subseteq \cdots \subseteq \mathcal{V}_m\,.$$

I put no restriction on length or on repeats. The single space $\{0\}$, for example, is a flag. Even if repeated a number of times.

Any ordered set of vectors $v = (v_i)$ determines the flag $[\![v]\!]$, in which $[\![v]\!]_i$ is spanned by the v_j with $j \leq i$. In particular, the basis (e_i) determines the **standard flag** $\mathcal{E}$. The flag $\mathcal{T}(m)$ associated to any matrix m is that determined by the order of its columns, preceded for convenience by 0.

If $\mathcal{F}$ and $\mathcal{G}$ are two flags, of lengths f and g, the **coincidence matrix** of the pair is the $f \times g$ matrix h with

$$C_{i,j} = \dim \mathcal{F}_i \cap \mathcal{G}_j\,.$$

If $\mathcal{F}$ is a single flag, its coincidence matrix is that of $\mathcal{E}$ and $\mathcal{F}$.

The coincidence matrix of a flag $\mathcal{F}$ with respect to itself is the $n \times n$ matrix $(\min i, j)$. In any coincidence matrix C the entries in each row and column are weakly increasing, and that $C_{i,j} \leq \min i, j$ for all i, j.

Example. If $\mathcal{F}$ is the flag associated to the reversed matrix

$$\begin{bmatrix} \circ & \circ & \circ & 1 \\ \circ & \circ & 1 & \circ \\ \circ & 1 & \circ & \circ \\ 1 & \circ & \circ & \circ \end{bmatrix}$$

which is that defined by the array $(0, e_4, e_3, e_2, e_1)$, its coincidence matrix is

$$\begin{bmatrix} \circ & \circ & \circ & \circ & 1 \\ \circ & \circ & \circ & 1 & 2 \\ \circ & \circ & 1 & 2 & 3 \\ \circ & 1 & 2 & 3 & 4 \end{bmatrix} . \tag{5.1}$$

$$\circ \text{————} \circ$$

The point of introducing flags is that two matrices of the same size determine the same flag if and only if one of them can be obtained from the other by multiplying it on the right by an invertible upper triangular matrix. Therefore Theorem 4.1 can be reformulated:

5.2. Proposition. *Every flag is represented by a unique matrix in permuted echelon form.*

Or, equivalently:

5.3. Proposition. *Suppose m to be any matrix. If c_1 and c_2 are matrices in permuted echelon form that are obtained from m through multiplication on the right by upper triangular unipotent matrices, they are equal.*

The proof will characterize certain properties of the matrix m by properties of the flag $\mathcal{T}(m)$. Let Γ be the coincidence matrix of $\mathcal{F}$, and then let

$$[\Delta\Gamma]_{i,j} = \Gamma_{i,j} - \Gamma_{i,j-1} \quad (j \geq 1) .$$

I call the matrix $\Delta\Gamma$ the **profile** of the flag $\mathcal{F}$.

For example, let

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

The basic fact is that a vector

$$e_i + \sum\nolimits_{j<i} c_j e_j$$

is in $\mathcal{C}_k$ if and only if $k \geq i$. Hence its coincidence matrix is

$$C = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 2 \\ 1 & 1 & 2 & 2 & 3 \\ 1 & 1 & 2 & 3 & 4 \\ 1 & 2 & 3 & 4 & 5 \end{bmatrix} ,$$

and

$$\Delta C = \begin{bmatrix} 0 & 0 & 0 & 0 & \boldsymbol{1} \\ \boldsymbol{1} & 0 & 0 & 0 & \boldsymbol{1} \\ \boldsymbol{1} & 0 & \boldsymbol{1} & 0 & \boldsymbol{1} \\ \boldsymbol{1} & 0 & \boldsymbol{1} & \boldsymbol{1} & \boldsymbol{1} \\ \boldsymbol{1} & \boldsymbol{1} & \boldsymbol{1} & \boldsymbol{1} & \boldsymbol{1} \end{bmatrix} .$$

Even better, a graph:

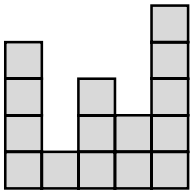

There is an obvious and simple relationship between the first and last matrices. This is a general fact. Suppose c to be any matrix in permuted column form. Recall that r_j is the row in which the pivot in column j appears (equal to 0 if there is no pivot). Define the matrix $\gamma = \gamma_c$ of the same size as c:

$$\gamma_{i,j} = \begin{cases} 1 & \text{if } i \geq r_j \\ 0 & \text{otherwise.} \end{cases}$$

The example above illustrates:

5.4. Lemma. *If m is a matrix in permuted column form, then γ_m is equal to $\Delta\Gamma_{\mathcal{T}(m)}$.*

I leave this as an exercise.

The point is that the shape of the permuted echelon form of a matrix depends only on the flag associated to it. As a consequence, any two matrices in permuted column form representing the same matrix have the same pivots (icluding 0 columns). To conclude the proof, it remains to show that if c_1 and c_2 are two matrices in permuted column form'with the same pivots, and $c_2 = c_1 u$ for some lower triangular matrix u, then $u = I$ and $c_2 = c_1$. This can be done by induction on the number of columns in the c_i. I leave it, too, as an exercise.

In any case, this concludes the proof of uniqueness in Theorem 4.1.

6. Unipotent groups

Suppose N to be a unipotent algebraic group N defined over a field F of characteristic 0. It possesses a filtration

$$N_n = \{1\} \subset N_{n-1} \subset \ldots \subset N_0 = N$$

by normal subgroups such that (a) there exists an isomorphism $\mathfrak{e}_i$ of N_i/N_{i+1} with F; (b) each quotient N_i/N_{i+1} is contained in the centre of N/N_{i+1}. An easy induction then shows every element can be expressed as a product of elements $\mathfrak{e}i(x_i)$

$$\mathfrak{e}_{n-1}(x_{n-1})\mathfrak{e}_{n-2}(x_{n-1}) \cdots \mathfrak{e}_0(x_0)\,.$$

expressed in decreasing order. I'll call it the **normal form** associated to the filtration and the splittings.

Example. Take N to be the group of unipotent upper triangular matrices in GL_n. There is a very simple filtration with an interpretation in terms of matrices. For a pair (i,j) with $1 \leq i < j \leq n$ let $r = r(i,j) = j(j-1)/2 - i$, let $N_{[r]}$ be the group made up of the matrices $\mathfrak{e}_{i,j}(x)$, and let

$$N_r = \prod\nolimits_{s \geq r} N_{[s]}\,.$$

Thus if $n = 3$

$$N_{[0]} = \left\{ \begin{bmatrix} 1 & * & \circ \\ \circ & 1 & \circ \\ \circ & \circ & 1 \end{bmatrix} \right\}, N_{[1]} = \left\{ \begin{bmatrix} 1 & \circ & \circ \\ \circ & 1 & * \\ \circ & \circ & 1 \end{bmatrix} \right\}, N_{[2]} = \left\{ \begin{bmatrix} 1 & \circ & * \\ \circ & 1 & \circ \\ \circ & \circ & 1 \end{bmatrix} \right\}, N_{[3]} = \left\{ \begin{bmatrix} 1 & \circ & \circ \\ \circ & 1 & \circ \\ \circ & \circ & 1 \end{bmatrix} \right\}.$$

Of course, given $0 \leq r < n(n-1)/2$ there exist unique (i,j) such that $r = j(j-1)/2 - i$, so that $N_{[r]}$ and N_r are well defined.

I leave it as an exerise to verify:

6.1. Lemma. *If n is the unipotent upper triangular matrix with $n_{i,j} = x_{i,j}$ then*

$$n = \prod\nolimits_{i,j} \mathfrak{e}_{i,j}(x_{i,j})\,,$$

arranged in order of decreasing $r(i,j)$.

For example:

$$\begin{bmatrix} 1 & x_{1,2} & x_{1,3} & x_{1,4} \\ \circ & 1 & x_{2,3} & x_{2,4} \\ \circ & \circ & 1 & x_{3,4} \\ \circ & \circ & \circ & 1 \end{bmatrix} = \begin{bmatrix} 1 & \circ & \circ & x_{1,4} \\ \circ & 1 & \circ & x_{2,4} \\ \circ & \circ & 1 & x_{3,4} \\ \circ & \circ & \circ & 1 \end{bmatrix} \begin{bmatrix} 1 & \circ & x_{1,3} & \circ \\ \circ & 1 & x_{2,3} & \circ \\ \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & 1 \end{bmatrix} \begin{bmatrix} 1 & x_{1,2} & \circ & \circ \\ \circ & 1 & \circ & \circ \\ \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & 1 \end{bmatrix},$$

where

$$\begin{bmatrix} 1 & \circ & \circ & x_{1,4} \\ \circ & 1 & \circ & x_{2,4} \\ \circ & \circ & 1 & x_{3,4} \\ \circ & \circ & \circ & 1 \end{bmatrix} = \begin{bmatrix} 1 & \circ & \circ & x_{1,4} \\ \circ & 1 & \circ & \circ \\ \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & 1 \end{bmatrix} \begin{bmatrix} 1 & \circ & \circ & \circ \\ \circ & 1 & \circ & x_{2,4} \\ \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & 1 \end{bmatrix} \begin{bmatrix} 1 & \circ & \circ & \circ \\ \circ & 1 & \circ & \circ \\ \circ & \circ & 1 & x_{3,4} \\ \circ & \circ & \circ & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & \circ & x_{1,3} & \circ \\ \circ & 1 & x_{2,3} & \circ \\ \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & 1 \end{bmatrix} = \begin{bmatrix} 1 & \circ & x_{1,3} & \circ \\ \circ & 1 & \circ & \circ \\ \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & 1 \end{bmatrix} \begin{bmatrix} 1 & \circ & \circ & \circ \\ \circ & 1 & x_{2,3} & \circ \\ \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & x_{1,2} & \circ & \circ \\ \circ & 1 & \circ & \circ \\ \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & 1 \end{bmatrix} = \begin{bmatrix} 1 & x_{1,2} & \circ & \circ \\ \circ & 1 & \circ & \circ \\ \circ & \circ & 1 & \circ \\ \circ & \circ & \circ & 1 \end{bmatrix}.$$

Revert to the case of a general unipotent group N. The following is well known, but a proof is hard to find in the literature.

6.2. Lemma. *Any element in N may be expressed as a product of elements in the $N_{[i]}$ in any order.*

Proof. The order is to be prescribed as a permutation σ of $[0,n)$—we want to write

$$\nu = \nu_0\nu_1\ldots\nu_{n-1}$$

where ν_i lies in $N_{[\sigma(i)]}$. Conversely, an element in $N_{[i]}$ lies in the place $\sigma^{-1}(i)$. For example, if $\sigma = (3,0,2,1)$ then we want an expression $x_3x_0x_2x_1$ with x_i in $N_{[i]}$.

The proof is an algorithm. We start with ν in normal form $\nu = \prod_i \nu_i$, in decreasing order. As the calculation proceeds, at the beginning of the m-th stage we shall have an expression

$$\nu = \kappa_m\lambda_m$$

where κ_m lies in N_m, expressed in normal form, and λ_m is a product of elements in the $N_{[i]}$ with $i \le m-1$. Thus $\kappa_m x_m^{-1}$ lies in N_{m+1} for some x in $N_{[m]}$ we can write

$$\nu = \kappa_m\lambda_m = \kappa_m x_m^{-1}\cdot x_m\lambda_m\,.$$

I have said that λ_m is a product of y_i in $N_{[i]}$, but in what order? In the order induced on $[0,m)$ by σ—i.e. so that the sequence $\sigma^{-1}(i)$ is decreasing.

The first step is trivial: we put ν in normal form κ_1x_0, setting $\lambda_1 = x_0$. Here κ_1 will be in normal form, so there exists x_1 in $N_{[1]}$ such that $\mu_2 = \kappa_1x_1^{-1}$ is in N_2. We write

$$\nu = \mu_2x_1\cdot x_0 = \mu_2\cdot x_1x_0\,.$$

But what we do now depends on σ. There are two cases. Either 1 precedes 0 in σ or it doesn't. If it does, we leave x_1x_0 as it is, and set $\kappa_2 = \mu_2$. If not, we write

$$x_1x_0 = x_1x_0x_1^{-1}x_0^{-1} \cdot x_0x_1 \, .$$

The commutator $x_1x_0x_1^{-1}x_0^{-1}$ lies in N_2, and we have

$$\nu = \mu_2 \cdot x_1x_0x_1^{-1}x_0^{-1} \cdot x_0x_1 = \kappa_2 \cdot x_0x_1 \, .$$

Repeat as needed.

7. In GL(n)

In previous sections, we have seen that M_n is the union of double cosets BwN, in which w is a partial permutation matrix, and is unique. In this section, I'll assume that m is invertible, in which case w is in fact a permutation matrix. This will allow us to say more about the right-hand factor N.

I'll begin by demonstrating the problem. If $w \neq I$, as we shall see, the group $N^w = w^{-1}Nw \cap N$ is not trivial, and if n is in it then $wn = wnw^{-1} \cdot w$, so the expression bwn in BwN is not unique. In other words, the element n here is *a priori* only an element of the quotient $N_w\backslash N$. However, we can find a section of this quotient. Let $N_w = w^{-1}\overline{N}w \cap N$.

7.1. Lemma. *The product map*

$$N^w \times N_w \longrightarrow N$$

is bijective.

This implies immediately:

7.2. Theorem. *Every m in GL_n possesses a unique factorization $m = bwn$ with n in N_w.*

Proof. It will be convenient to use the language of root systems. Let $\mathfrak{a}$ be the Lie algebra of the group A, the vector space of diagonal matrices, and $\mathfrak{n}$ be the Lie algebra of upper triangular unipotent matrices. Let ε_i be the linear function on on $\mathfrak{a}$ taking a to its i-th entry $a_{i,i}$.

The group A acts by conjugation on the Lie algebra $\mathfrak{gl}_n$, the Lie algebra $\mathfrak{a}$ by the adjoint action. The Lie algebra $\mathfrak{gl}_n$ is the direct sum of eigenspaces. One of these is $\mathfrak{a}$, and the rest have dimension 1. which is then a direct sum of eigenspaces. Its Lie algebra $\mathfrak{a}$ acts by the adjoint action, through the Lie bracket. For each pair (i, j) with $i \neq j$ let $e_{i,j}$ be the elementary matrix with entries vanishing except at i, j, where it is 1. Then

$$[a, e_{i,j}] = (a_i - a_j)\, e_{i,j}$$

for all a in $\mathfrak{a}$. In other words, $e_{i,j}$ is an eigenvector with eigencharacter $\lambda_{i,j} = \varepsilon_i - \varepsilon_j$. This character of $\mathfrak{a}$ is called a **root**. Those with $i < j$ are positive, the others negative. The Lie algebra $\mathfrak{n}$ is a sum of root spaces for positive roots, its opposite $\overline{n}$ for negative.

Each root corresponds also to an embedding of F into $\mathfrak{gl}_n$:

$$\mathfrak{e}_\lambda(x) = I + xe_\lambda \, .$$

Let N_λ be its image.

Of course something similar holds for $\overline{N}$ and negative roots. The group W acts by conjugation on $\mathfrak{gl}_n$, and permutes the roots: $\langle w\lambda, a\rangle = \langle \lambda, w^{-1}aw\rangle$, and

$$w\mathfrak{e}_\lambda w^{-1} = \mathfrak{e}_{w\lambda} \, .$$

Acording to Lemma 6.2, then,

$$N_w = \prod_{\substack{\lambda > 0 \\ w^{-1}\lambda > 0}} N_\lambda \, .$$

I call the expression $x = bw_\sigma n$ the **Bruhat normal form** of x.

8. Parabolic subgroups

The group GL_n acts on flags in the obvious way. If a flag is represented by a matrix M, the matrix g takes M to gM. Since GL_n acts transitively on $F^n - \{0\}$, an induction argument shows that it acts transitively on princial flags in F^n. The stabilizer of a flag is a **parabolic** subgroup. The stabilizer of the standard principal flag is the Borel subgroup of lower triangular matrices.

8.1. Lemma. *A parabolic subgroup is its own normalizer.*

Proof. If P is the stabilizer of the flag $\mathcal{F}$, then g in GL_n, $gPg^{-1} = P$ if and only if $g\mathcal{F} = \mathcal{F}$. ∎

If n_i is the array of differences between the dimension of F_i and that of F_{i-1}, then $n = n_1 + n_2 + \cdots + n_r$—i.e. the n_i form a partition of n. For example, if the partition is $n = 1 + (n-1)$ we are looking at lines in projective space. If the partition is $n = 1 + \cdots + 1$ then the flag is principal.

For example, if we write $6 = 1 + 2 + 3$ we get the parabolic subgroup of matrices

$$\left[\begin{array}{c|cc|ccc} * & * & * & * & * & * \\ \hline 0 & * & * & * & * & * \\ 0 & * & * & * & * & * \\ \hline 0 & 0 & 0 & * & * & * \\ 0 & 0 & 0 & * & * & * \\ 0 & 0 & 0 & * & * & * \end{array}\right]$$

From the partition of n we obtain a direct sum decomposition

$$F^n = \oplus F^{n_i}$$

and an associated embedding of the product $\prod \mathrm{GL}_{n_i}(F)$ into $\mathrm{GL}_n(F)$, *via* disjoint $n_i \times n_i$ blocks along the diagonal. Another way to express the condition on matrices m in P_Θ is to say that the non-zero entries in m lie either above the diagonal, or in one of these blocks.

Mapping p in P to the sequence of its block diagonal matrices defines a homomorphism from P_Θ to $M_\Theta = \prod \mathrm{GL}_{n_i}(k)$. The kernel consists of the subgroup N_Θ of unipotent upper triangular matrices with only $n_i \times n_i$ identity matrices along the diagonal. The group P_Θ is the semi-direct product of M_Θ and N_Θ.

8.2. Proposition. *Every parabolic subgroup is conjugate to exactly one standard one.*

8.3. Proposition. *If $F = \mathfrak{k}$ then every quotient $P\backslash G$ is compact.*

Proof. If the partition is $n = 1 + (n-1)$ then the quotient $P\backslash G$ is isomorphic to projective space $\mathbb{P}(\mathfrak{k})$. We shall see in a moment that this is compact. It follows by induction that $P\backslash G$ is compact if P is the group of upper triangular matrices, and from this in turn for an arbitrary P.

Why is $\mathbb{P}(\mathfrak{k})$ compact? If (x_i) is a non-zero vector then it is projectively equivalent to (x_i/μ) where $\mu = x_m$ is the coordinate with maximum $\mathfrak{p}$-adic norm. But the set of all points (x_i) with $x_m = 1$ and $|x_i| \leq 1$ is compact. ∎

9. The Bruhat order

The double coset $Pw_\ell P$ is in almost any sense the largest of the double cosets. This can be made precise. Let $\overline{B} = A\overline{N}$ be the parabolic subgroup of lower triangular matrices, **opposite** to B. Then $\overline{B} = w_\ell B w_\ell$, and $B^{opp}B = w_\ell B w_\ell B$, the left translate of $Bw_\ell B$ by w_ℓ. For any $n \times n$ matrix let $X_{(r)}$ be the $r \times r$ matrix made up from its first r rows and r columns.

9.1. Proposition. *Suppose that x is an $n \times n$ matrix. It can be factored as $x = \nu a u$ with ν in $\overline{N}$ and u in N if and only if every one of the n matrices $X_{(r)}$ has non-zero determinant.*

Proof. This follows from a simple variation of the algorithm described above, applying induction. At step k of Gauss reduction the first diagonal entry must be invertible. But by induction we have at this step the factorization of the $k \times k$ sub-matrix, and the product of all the diagonal entries up to the k-th is its determinant. ∎

The matrices with an νau factorization are hence the complement in GL_n of a finite union of zero sets of polynomials. If the field k has a topology—for example, if it the $\mathfrak{p}$-adic field $\mathfrak{k}$—they form an open set.

Suppose that y is a permutation with the reduced expression $y = s_1 \dots x_n$. The element x is said to be in the closure $\overline{x}$ of y, or $x \le y$, if x is a product of a subsequence of the s_i.

9.2. Proposition. *The closure in G of the double coset $C(y)$ is the disjoint union of the cosets $C(x)$ with $x \le y$.*

Proof. By induction on the length of y. For $y = 1$ the assertion is trivial, and for $y = s$ it follows from what we have calculated for GL_2, since $P_s = \overline{C(s)}$ is the parabolic subgroup which is a product of B and a copy of GL_2 embedded along the diagonal.

It suffices now, using induction and , to show that if $\ell(ws) = \ell(w) + 1$ then $\overline{C(ws)} = \overline{C(w)}\,\overline{C(s)}$, or equivalently to show that $\overline{C(w)}\,\overline{C(s)}$ is closed. Both G/P and G/B are compact. If C is closed in G and $CB = C$, then the image of C in G/P is closed, and so is its inverse image in G, which is CP.

One consequence of this is that the closure of a permutation x does not depend on any particular reduced expression for it. This can be shown also by purely combinatorial arguments.

9.3. Proposition. *If σ and π are two permutations, $\pi \le \sigma$ if σ is obtained from π by a sequence of transpositions (i, j) with $i < j$ and i occurring to the left of j in the array $\pi(i)$.*

Also, following Deodhar, for an array $(\sigma(i))$, let $\langle \sigma_i \rangle$ be the array written in increasing order. I say one array (a_i) is less than or equal to (b_i) if $a_i \le b_i$ for all i. Then $\pi \le \sigma$ if and only if the initial sequence of π is less than or equal to that of σ.

Part III. p-adic fields

10. Lattices

I now take up the special features of $\mathrm{GL}_n(\mathfrak{k})$ where

$$\begin{aligned} \mathfrak{k} &= \text{a } \mathfrak{p}\text{-adiv field} \\ \mathfrak{o} &= \text{its ring of integers} \\ \varpi &= \text{a generator of } \mathfrak{p} \\ \mathbb{F}_q &= \mathfrak{o}/\mathfrak{p} \\ |x| &= q^{-m} \text{ if } x/\varpi^m \text{ is in } \mathfrak{o}^\times. \end{aligned}$$

Just about everything I'll say applies also to the quotient field of a Dedekind domain R, but with the ring $\mathfrak{o}$ of integers replaced by the localization of R at a prime ideal $\mathfrak{p}$. The simplest example would be $\mathbb{Z}_{(p)}$, the ring of all rational numbers a/b with b relatively prime to p. This is useful for playing around with programs.

A vector (x_i) in $\mathfrak{o}^n$ is called **primitive** if one of the x_i is a unit.

10.1. Lemma. *A vector $x = (x_1, \dots, x_n)$ in $\mathfrak{o}^n$ is an element of an $\mathfrak{o}$-basis of $\mathfrak{o}^n$ if and only if it is primitive.*

Proof. If x_i is a unit, then the set of vectors obtained by substituting x for ε_i in the standard basis is again a basis.

A **lattice** in $\mathfrak{k}^n$ is any finitely generated $\mathfrak{o}$-module in $\mathfrak{k}^n$ that spans $\mathfrak{k}^n$. The **standard lattice** is $\mathcal{L}_0 = \mathfrak{o}^n$.

10.2. Proposition. (Principal divisors) *If L is a lattice in $\mathfrak{k}^n$ then there exists a basis $e_1, e_2, \dots, e_n$ of $\mathfrak{o}^n$ and integers*

$$m_1 \le m_2 \le \cdots \le m_n$$

such that $\varpi^{m_1} e_1, \varpi^{m_2} e_2, \dots, \varpi^{m_n} e_r$ form a basis of L. The integers $m_1, m_2, \dots, m_n$ are uniquely determined.

Proof. The proof will be constructive.

The Proposition will follow by induction from a more general fact. Suppose ℓ_1, ℓ_2, ... , ℓ_r to be any finite set of elements of V. Choose a coordinate system, and let L be the matrix with columns ℓ_i. Then there exist matrices k_1, k_2 in $\mathrm{GL}_n(\mathfrak{o})$ and $\mathrm{GL}_k(\mathfrak{o})$ such that $D = k_1 L k_2$ is semi-diagonal with entries $D_{i,i} = \varpi^{m_i}$ such that $m_i \leq m_{i+1}$, and that D is unique.

Suppose $r = 1$,so that L is a column vector. We can express $L = \varpi^m \lambda$ with λ a primitive vector in $\mathfrak{o}^n$.

Let $m_1 < \infty$ be the greatest integer such that ϖ^{m_1} divides all the coordinates of the ℓ_i. Then there exists at least one ℓ_i with coordinate $\varpi^{m_1} u$, u a unit in $\mathfrak{o}$. By a suitable swap of columns and rows, we may move this to upper left in the matrix. By suitable integral column and row operations, we may arrange it so that all other entries in the first row and column vanish. These row and column operations amount to multiplication on left and right by integral invertible matrices. Apply the induction to the lower right $(n-1) \times (r-1)$ submatrix.

As for uniqueness, ϖ^{m_1} is the greatest common divisor of all the matrix entries, $\varpi^{m_1+m_2}$ is the greatest common divisor of the 2×2 minor determinants, etc. Changing the basis of L amounts to multiplying this matrix on the right by a matrix in $\mathrm{GL}_n(\mathfrak{o})$ and doesn't change these characterizations.

10.3. Corollary. *All lattices are free $\mathfrak{o}$-modules.*

10.4. Corollary. *The group $\mathrm{GL}_n(\mathfrak{k})$ acts transitively on the set of lattices.*

10.5. Corollary. *If L and M are any two lattices there exists a basis (e_i) of L and an array (m_i) of integers with $m_i \leq m_{i+1}$ such that $(\varpi^{m_i} e_i)$ is a basis of M.*

The stabilizer of $\mathfrak{o}^n$ is $\mathrm{GL}_n(\mathfrak{o})$.

Let A be the group of diagonal matrices, $\mathfrak{A}$ the subgroup of the

$$\varpi^m = \mathrm{diag}(\varpi^{m_i})$$

whose diagonal entries are powers of ϖ, $\mathfrak{A}^{--}$ the subset of ϖ^m with $m_i \leq m_{i+1}$.

10.6. Corollary. *(Cartan decomposition) Every matrix g in $\mathrm{GL}_n(\mathfrak{k})$ can be expressed as*

$$g = \gamma_1 \, a \, \gamma_2$$

where γ_1 and γ_2 are in $\mathrm{GL}_n(\mathfrak{o})$ and a is in $\mathfrak{A}^{--}$. The element a is unique.

10.7. Proposition. *The compact open subgroup $\mathrm{GL}_n(\mathfrak{o})$ acts transitively on the space of principal flags in $\mathfrak{k}^n$.*

Proof. If we are given a principal flag (V_i) we can find in the line $V_1 \cap \mathfrak{o}^n$ a primitive vector, which by Lemma 10.1 is part of a basis, so we can find γ with $\gamma V_1 = \mathfrak{k}\varepsilon_1$. From now on we may assume that $V_1 = \mathfrak{k}\varepsilon_1$, and proceed by induction. Look next at V/V_1, which is given the γ-transform of the original flag modulo V_1. Any element of $\mathrm{GL}_n(\mathfrak{o})$ which transforms this flag into the standard one can be lifted to an element of $\mathrm{GL}_n(\mathfrak{o})$ leaving ε_1 fixed.

10.8. Corollary. *If $K = \mathrm{GL}_n(\mathfrak{o})$ then $G = KP = PK$ for any parabolic subgroup P.*

Since G/P is compact, so is K. But this follows from the more elementary fact that K is the projective limit. of the finite groups $\mathrm{GL}_m(\mathfrak{o}/\mathfrak{p}^m)$.

11. Volumes

Let $K = \mathrm{GL}_n(\mathfrak{o})$ and choose a Haar measure on G such that the volume of KI is equal to 1. Suppose a to be a diagonal matrix. *What is the volume of the open set KaK in GL_n?* Equivalently, what is $|KaK/K|$? The nature of the formula depends very definitely on a. For example, suppose $n = 2$. If t is the identity matrix then $KtK = K$ and its volume is that of K, but if

$$t = \begin{bmatrix} 1 & 0 \\ 0 & \varpi^m \end{bmatrix}$$

with $m > 0$ then the volume of KtK is equal to

$$(1 + q^{-1})q^m$$

times the volume of K. This is because the map $g \mapsto gL$ induces a bijection of KtK/K with lines in $\mathbb{P}^1(\mathfrak{o}/\mathfrak{p}^m)$. This difference in qualitative behaviour remains valid for all n, as we shall now see.

First comes a simpler question. The group $\mathrm{GL}_n(\mathfrak{o})$ maps canonically onto the finite group $\mathrm{GL}_n(\mathfrak{o}/\mathfrak{p}^m)$. How large is this group?

(1) Suppose $m = 1$. Then we are looking at $G = \mathrm{GL}_n(\mathbb{F}_q)$. Let τ_n be the number of elements in G. For $n = 1$ this is $\mathbb{F}_q^\times$. Thus $\tau_1 = (q - 1)$, for example. For $n > 1$, the group G acts transitively on $\mathbb{F}^n - \{0\}$, and the isotropy subgroup of $(1, 0, \ldots, 0)$ contains exactly those matrices m with $m_{1,1} = 1$ and $m_{i,1} = 0$ for $i > 1$. The elements $m_{1,i}$ are arbitrary, and the lower $(n - 1) \times (n - 1)$ matrix must be non-singular. The order of the isotropy group is therefore order $q^{n-1}\tau_{n-1}$, and we have for τ_n the recursive formula

$$\tau_n = \begin{cases} q - 1 & \text{if } n = 1 \\ (q^n - 1)q^{n-1}\tau_{n-1} & \text{otherwise} \end{cases}$$

from which we calculate by induction

$$\tau_n = (q^n - 1) \ldots (q - 1)q^{n(n-1)/2}$$

It is a polynomial in q of order q^{n^2}.

(2) The group $\mathrm{GL}_n(\mathfrak{o}/\mathfrak{p}^m)$ fits into an exact sequence

$$1 \to I + \mathfrak{p}M_{n-1}(\mathfrak{o}/\mathfrak{p}^m) \to \mathrm{GL}_n(\mathfrak{o}/\mathfrak{p}^m) \to \mathrm{GL}_n(\mathbb{F}_q) \to 1$$

The order of $\mathrm{GL}_n(\mathfrak{o}/\mathfrak{p}^m)$ is therefore the product of the $q^{n^2(m-1)}$ and the τ_n.

Now define a sort of normalized size

$$\gamma(\mathrm{GL}_n) = \frac{\tau_n}{q^{n^2}} = \left(1 - q^{-n}\right)\left(1 - q^{-(n-1)}\right) \ldots \left(1 - q^{-1}\right) .$$

11.1. Proposition. *Suppose*

$$m_1 \leq m_2 \leq \ldots \leq m_n$$

and let

$$a = \varpi^m .$$

Suppose that the integers n_i are the lengths of constant runs in the sequence (m_i), so that $n_1 + \cdots + n_k = n$ and

$$m_1 = \cdots = m_{n_1} < m_{n_1+1} = \cdots = m_{n_1+n_2} < m_{n_1+n_2+1} = \cdots$$

Let

$$M = M_a = \mathrm{GL}_{n_1} \times \mathrm{GL}_{n_2} \times \ldots \times \mathrm{GL}_{n_k}$$

and define

$$\gamma(M) = \prod \gamma(\mathrm{GL}_{n_i}) \ .$$

We embed M as diagonal blocks in GL_n. Then

$$|KaK/K| = \frac{\gamma(\mathrm{GL}_n)}{\gamma(M)}\ \delta_\emptyset(a) \ .$$

Here $\delta_\emptyset$ is the character

$$\prod_{j<i} |a_j/a_i|$$

of the group of diagonal matrices (a_i).

Proof. The map $k \mapsto kaK/K$ induces a bijection of $K/K\cap aKa^{-1}$ with KaK/K. The group $K\cap aKa^{-1}$ consists of all integral invertible matrices g with $g_{i,j} \equiv 0 \mod \varpi^{m_i-m_j}$. There is no condition when $m_i \le m_j$, so this is in effect a restriction only when $i > j$. Fix $m \ge m_n$, so that $m \ge m_i$ for all i. Then the index of $K \cap aKa^{-1}$ in K is the same as that of the index of $K_m \cap aK_ma^{-1}$ in K_m where $K_m = \mathrm{GL}_n(\mathfrak{o}/\mathfrak{p}^m)$.

Let P be the parabolic subgroup corresponding to the partition of n. The cardinality of K_m is $\tau_n\, q^{n^2(m-1)}$. What is that of $K_m\cap aK_ma^{-1}$? The image of $K_m\cap aK_ma^{-1}$ modulo $\mathfrak{p}$ is the parabolic subgroup $P(\mathbb{F}_q)$of $\mathrm{GL}_n(\mathbb{F})$, which has cardinality $\prod_{1\le i\le k}\tau_{n_i}\prod_{1\le i<j\le k} q^{n_in_j}$. The kernel of the projection is the subgroup of matrices in $K_m \cap aK_ma^{-1}$ congruent to I modulo $\mathfrak{p}$. It has cardinality

$$\prod_{1\le i\le j\le k} q^{n_in_j(m-1)} \prod_{1\le j<i\le k} q^{n_in_j(m-(\mu_i-\mu_j))} = \prod_{1\le i\le j\le k} q^{-n_in_j} \prod_{1\le i,j\le k} q^{n_in_jm} \prod_{1\le j<i\le n} q^{-(m_i-m_j)}$$

where μ_i is the common value of m_j in the block of n_i. The cardinality of $K_m \cap aK_ma^{-1}$ is therefore

$$\prod_{1\le i\le k}\tau_{n_i} \prod_{1\le i<j\le k} q^{n_in_j} \prod_{1\le i\le j\le k} q^{-n_in_j} \prod_{1\le i,j\le k} q^{n_in_jm} \prod_{1\le j<i\le n} q^{-(m_i-m_j)}$$
$$= q^{n^2m} \prod_{1\le i\le k} \tau_{n_i} q^{-n_i^2} \prod_{1\le j<i\le n} q^{-(m_i-m_j)}$$

and the cardinality of the quotient is therefore

$$\frac{\tau_n q^{n^2(m-1)}}{q^{n^2m}\prod_{1\le i\le k}\tau_{n_i}q^{-n_i^2}\prod_{1\le j<i\le n}q^{-(m_i-m_j)}} = \frac{\gamma(\mathrm{GL}_n)}{\gamma(M)} \prod_{1\le j<i\le n} q^{m_i-m_j} \ .$$

12. The affine permutation group

As preliminary material for the next section, I'll introduce now the analogue of the symmetric group $\mathfrak{S}_n$ relevant to the structure of the $\mathfrak{p}$-adic group $G = \mathrm{GL}_n(\mathfrak{k})$.

THE GROUPS IN PLAY. The $\mathfrak{p}$-adic analogue of the group A of diagonal matrices is the compact group $A(\mathfrak{o})$. I recall that There are two analogues of the Weyl group. One is

$$\mathfrak{W} = N_G(A(\mathfrak{o}))/A(\mathfrak{o}) = N_G(A)/A(\mathfrak{o})\,.$$

It contains the quotient $\mathfrak{A} = A/A(\mathfrak{o})$, which is isomorphic to $\mathbb{Z}^n$. It also contains $\mathfrak{S}_n$, and in fact it is the semi-direct product $\mathfrak{A} \rtimes \mathfrak{S}_n$.

It possesses a well defined homomorphism onto $\mathbb{Z}$, taking

$$m \times \sigma \longmapsto \sum m_i\,.$$

The second group $\widetilde{\mathfrak{S}}_n$ is its kernel. It generated by $\mathfrak{S}_n$ and the subgroup $\mathfrak{A}_0$ of m in $\mathbb{Z}^n$ with $\sum m_i = 0$, and it is the semi-direct product of $\mathfrak{A}_0$ and $\mathfrak{S}_n$.

This second group is the affine Weyl group associated to the root system of GL_n. The affine roots in this case are the affine functions $\lambda + k$, in which λ is a root for GL_n—one of the functions $x_i - x_j$. The affine Weyl group is that generated by the orthogonal reflections in the lines $\lambda + k = 0$.

The following figure demonstrates that at least such reflections are in the group $\widetilde{\mathfrak{S}}_n$, since reflection in the hyperplane $\lambda = 1$ is the same as the reflection s_λ in the hyperplane $\lambda = 0$ followed by translation by $\lambda^\vee$.

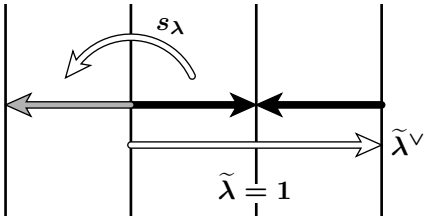

Let $\mathcal{C}$ be the region where $x_i \le x_{i+1}$, $x_n - x_1 \le 1$ in V_0.

Draw somewhere the lines of intersection of $x_i = 0$ with V_0.

12.1. Lemma. *Given any point x in V, there exists a product w of reflections s_i such that $w(x)$ lies in $\mathcal{C}$. The affine Weyl group is generated by the reflections s_i together with the reflection s_0 in the hyperplane $\widetilde{\alpha} = 1$.*

I recall that

$$\widetilde{\alpha} = \alpha_1 + \cdots + \alpha_{n-1} = \varepsilon_1 - \varepsilon_n$$

is the dominant root.

Proof. The region $\mathcal{C}$ contains the point $\rho = (1/n, 2/n, \ldots, n/n)$. If w lies in $\widetilde{\mathfrak{S}}_n$, then

Affine reflection:

$$v \longmapsto v - (\langle \lambda, v\rangle - k)\lambda^\vee$$

In the figure below, the partition of the plane into translations of a fundamental domain for $n = 3$ is shown.

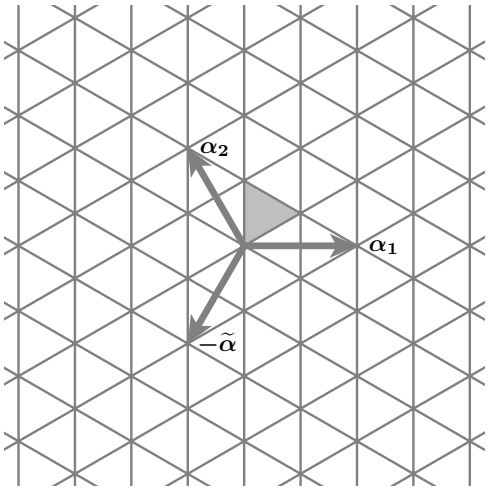

The analogue of roots in for this situation case are the A root is **positive** if it is non-negative on the chosen fundamental domain.

Every element of $\mathfrak{W}$ can be factored as a product of elements of $\mathcal{S}_n$ and $\mathfrak{A}$. But there is a more interest factorization possible. For arbitrary n, let ω be the $n \times n$ matrix with ϖ at lower left, and 1 along the superdiagonal. If $n = 3$, for example

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

Its determinant is $(-1)^{n-1}\varpi$.

12.2. Proposition. *Every element x of $\mathfrak{W}$ can be factored as $x = \omega^n w$, with w in $\widetilde{\mathfrak{S}}$.*

GEOMETRY. The group $\mathfrak{W}$ acts naturally on $\mathbb{R}^n$. The group $\mathfrak{A}$ acts by translations, and $\mathfrak{S}$ by permutations. The group $\widetilde{\mathfrak{S}}$ acts on the subspace with $\sum x_i = 0$. These groups act discretely as affine transformations. It is natural to ask, what is a fundamental domain?

A fundamental domain for $\mathbb{Z}^n$ is the unit cube $0 \le x_i \le 1$. The group $\mathfrak{S}_n$ acts on this, so *a fundamental domain for the action of $\mathfrak{S}_n$ on this cube is also a fundamental domain for $\mathfrak{W}$ acting on $\mathbb{R}^n$*. I'll choose this to be the intersection $\mathcal{C}$ of the unit cube with the fundamental domain for $\mathfrak{S}_n$ in $\mathbb{R}^n$, where $x_i \le x_{i+1}$ for $1 \le i < n$. The figure on the left below illustrates the case of GL_2. The group $\mathfrak{S}_2$ contains the single reflection in the line $y = x$. My choice of fundamental domain is shaded.

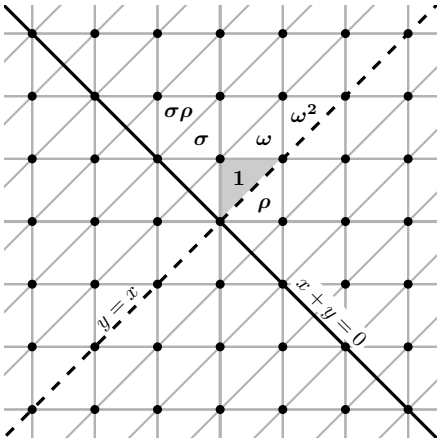

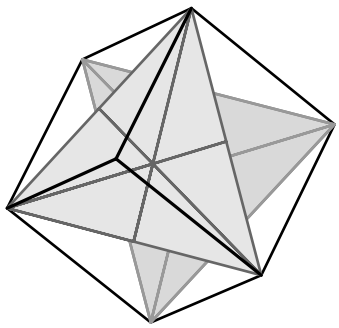

The transforms of $\mathcal{C}$ are called **alcoves**. Each alcove is in fact the transform of $\mathcal{C}$ by a unique element of $\mathfrak{W}$, and this is how some of the alcoves in the figure are labeled.

In the figure, the transformation ρ is in $\mathfrak{S}_n$, and amounts to reflection in the line $y = x$. The transformation σ is reflection in $y = x+1$, and the composite $\sigma\rho$ is the same as translation by $(-1, 1)$. The transformation ω is a kind of Euclidean motion known as a twisted translation, and ω^2 is indeed translation by $(1, 1)$.

I repeat, the fundamental domain $\mathcal{C}$ I have chosen is where $0 \le x_i \le 1$ for all i, and $x_i \le x_{i+1}$ for $1 \le i < n$. The figure on the right above attempts to show something of what happens for GL_3. In general:

12.3. Proposition. *The fundamental domain $\mathcal{C}$ is the convex hull of the points*

$$\begin{aligned}\delta_0 &= (0, \dots, 0, 0, 0)\\ \delta_1 &= (0, \dots, 0, 0, 1)\\ \delta_2 &= (0, \dots, 0, 1, 1)\\ &\dots\\ \delta_n &= (1, \dots, 1, 1, 1)\,.\end{aligned}$$

That is to say, the last i coordinates of δ_i are 1 and the rest 0. In particular, the domain is a simplex of dimension n.

Proof. Because if $x = (x_i)$ lies in $\mathcal{C}$ then

$$\varepsilon_i = \delta_{n-i+1} - \delta_{n-i}$$

and hence

$$\begin{aligned}x &= x_1\varepsilon_1 + \dots + x_n\varepsilon_n\\ &= x_1(\delta_n - \delta_{n-1}) + x_2(\delta_{n-1} - \delta_{n-2}) + \dots + x_n(\delta_1 - \delta_0)\\ &= x_1\delta_n + (x_2 - x_1)\delta_{n-1} + \dots + (x_n - x_{n-1})\delta_1\,.\end{aligned}$$

The length of $\omega^n\sigma$ with σ in $\mathfrak{S}_n$ is by definition $\ell(w)$. This is adequately justified by:

12.4. Proposition. *The length of σ is the same as the number of root hyperplanes between the fundamental domain $\mathcal{C}$ and $\sigma(\mathcal{C})$.*

13. Iwahori subgroups

This section is concerned with a generalization of the Cartan decomposition Corollary 10.6. Let $K = \mathrm{GL}_n(\mathfrak{o})$, $W = \mathfrak{S}n$, B the Borel subgroup of upper triangular matrices. Recall that ϖ is a generator of $\mathfrak{p}$.

IWAHORI FACTORIZATIONS. Let I be the inverse image in $\mathrm{GL}_n(\mathfrak{o})$ of the group of upper triangular matrices in $\mathrm{GL}_n(\mathbb{F}_q)$, those for which all entries $g_{i,j}$ with $i > j$ lie in $\mathfrak{p}$. An **Iwahori subgroup** of $\mathrm{GL}_n(\mathfrak{k})$ is any conjugate of I.

If P is a parabolic subgroup of G, $\overline{P}$ opposite to P, and K a compact open subgroup of G, define

$$\begin{aligned}M &= P \cap \overline{P}\\ K_{\overline{N}} &= K \cap \overline{N}\\ K_M &= K \cap M\\ K_N &= K \cap N\end{aligned}$$

The group K is said to possess an **Iwahori factorization** with respect to (P, P^{opp}) if the product map

$$K_{\overline{N}} \times K_M \times K_N \to K$$

is a bijection.

I recall that the standard parabolic subgroups of $\mathrm{GL}_n(\mathfrak{k})$ are those stabilizing the standard flags

$$0 \subset \mathfrak{k}^{n_1} \subset \mathfrak{k}^{n_1+n_2} \subset \ldots \subset \mathfrak{k}^n$$

They are those parabolic subgroups containing the Borel group of upper triangular matrices, and their opposites are their transposes.

13.1. Proposition. *Every element g of the Iwahori group I has a unique factorization with respect to any standard parabolic subgroup.*

Proof. It suffices to show this for the Borel subgroup. In this case it follows from Proposition 9.1, since the determinant of an element of an Iwahori subgroup is a unit in $\mathfrak{o}$.

13.2. Proposition. *For any pair $(P, \overline{P})$ there exists a sequence of compact open subgroups K forming a basis of neighbourhoods of the identity, each possessing an Iwahori factorization with respect to P.*

Proof. For $\mathrm{GL}_n(\mathfrak{k})$ we choose the sequence $\mathrm{GL}_n(\mathfrak{p}^m)$. The Iwahori factorization follows from .

THE IWAHORI DECOMPOSITION. The group $A(\mathfrak{o})$ plays a role in the $\mathfrak{p}$-adic group analogous to that of A in the algebraic group. The normalizer $N_G\big(A(\mathfrak{o})\big)$ of $A(\mathfrak{o})$ in G is the same as the normalizer of A, the semi-direct product of A and the Weyl group W. The quotient $\mathfrak{A} = A/A(\mathfrak{o})$, may be identified with the group of diagonal matrices whose entries are powers of ϖ. It is isomorphic to $\mathbb{Z}^n$. The quotient

$$\mathfrak{W} = N_G\big(A(\mathfrak{o})\big)/A(\mathfrak{o})$$

contains $\mathfrak{A}$ as a normal subgroup, and is in fact the semidirect product $\mathfrak{A} \rtimes W$. It is the $\mathfrak{p}$-adic analogue of W.

I shall call an $n \times n$ matrix X with $r \leq n$, of rank r, **Iwahori-reduced** if it has this property: *every column and every row has exactly one non-zero entry, and that of the form ϖ^n.* Elements of the group $\widetilde{W}$ may thus be identified with Iwahori-reduced $n \times n$ matrices. The group $\widetilde{W}$ acts as affine transformations on $\mathbb{Z}^n$, and is called the **affine permutation group**.

For $\mathrm{GL}_2(k)$, for example, the Iwahori-reduced matrices are these:

$$\begin{bmatrix} \varpi^k & 0 \\ 0 & \varpi^\ell \end{bmatrix}, \quad \begin{bmatrix} 0 & \varpi^k \\ \varpi^\ell & 0 \end{bmatrix}.$$

I begin with some results useful in a moment:

13.3. Proposition. *We have the factorizations $K = IWP(\mathfrak{o})$ and $G = IWP$.*

Proof. The Bruhat factorization for $\mathrm{GL}_n(\mathbb{F}_q)$ tells us that $K = IWP(\mathfrak{o})$. Apply Corollary 10.8.

Here is a $\mathfrak{p}$-adic version of the Bruhat decomposition:

13.4. Proposition. *Every element g of $\mathrm{GL}_n(\mathfrak{k})$ factors uniquely as $g = \iota_1 \omega \iota_2$ where each ι_i is in I and ω is in $\mathfrak{W}$.*

The proof will be constructive. To replace the elementary row and column operations of the Bruhat factorization, we require here certain **elementary Iwahori operations** on columns:

- add to a column d a multiple xc of a previous column c by some x in $\mathfrak{o}$;
- add to a column c a multiple xd of a subsequent column by x in $\mathfrak{p}$;
- multiply a column by a unit in $\mathfrak{o}$;

and also on rows:

- add to a row c a multiple xd of a subsequent row d with x in $\mathfrak{o}$;
- add to a row d a multiple xc of a previous row with x in $\mathfrak{p}$;
- multiply a row by a unit in $\mathfrak{o}$;

Each of these column (row) operations amounts to right (resp. left) multiplication by what I'll call an Iwahori matrix .

Here are some examples:

$$\begin{bmatrix} 1 & x \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} u + xv \\ v \end{bmatrix}$$
$$\begin{bmatrix} 1 & 0 \\ \varpi x & 1 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} u \\ \varpi xu + v \end{bmatrix}$$
$$\begin{bmatrix} u & v \end{bmatrix} \begin{bmatrix} 1 & x \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} u & xu + v \end{bmatrix}$$
$$\begin{bmatrix} u & v \end{bmatrix} \begin{bmatrix} 1 & 0 \\ \varpi x & 1 \end{bmatrix} = \begin{bmatrix} u + \varpi xv & v \end{bmatrix} .$$

The proof will now proceed by explaining exactly what elementary Iwahori operations will carry out the reduction we want. This will take several steps.

Step 1. Thje starting point is the Cartan factorization of an element of $\mathrm{GL}_n(\mathfrak{k})$, say $g = k_1 a k_2$, with a in A^{++}, which we do know how how to compute explicitly.

Step 2. We also know how to apply Proposition 13.3 to k_1 and k_2, to get

$$g = \iota_1 w_1 b_1 a b_2 w_2 \iota_2 .$$

We may ignore the outer factors, and must now embed $w_1 b_1 a b_2 w_2$ into some $IawI$.

Step 3. We can write the inner product in the form

$$w_1 a \cdot a^{-1} b_1 a b_2 w_2 .$$

Since a is in A^{++}, $a_0^{-1} p_1 a$ lies in $B(\mathfrak{o})$. We can express this in the form

$$w_1 a b w_2 .$$

I switch to get the form

$$a w_1 b w_2 .$$

Step 4. In order to apply induction, it will be best to formulate a slight generalization of the problem we are facing. Suppose a in A, x and y in W, ι in I. *How can we embed $ax\iota y$ in some $I\omega I$ with ω in $\mathfrak{W}$?*

Step 5. We can factor $\iota = n\overline{b}$ with $\overline{b}$ in $\overline{B}$. Then $\overline{b}y = y \cdot y^{-1}\overline{b}y$ and $y^{-1}\overline{b}y$ is in I. So we are back to the case $axny$ with n in $N(\mathfrak{o})$.

Step 6. We go by induction on $m = \ell(y)$. If $y = 1$, there is no problem, since awn lies in $\mathfrak{W}I$.

Otherwise, y can be expressed as sz with $\ell(z) = m - 1$ and s equal to some s_i. We now rewrite in the form

$$aws \cdot s^{-1} n s \cdot z .$$

Denoting ws by w, we are next going to embed $aw \cdot s^{-1} n s$ in some $I\omega I$. I switch this to the form $wa \cdot s^{-1} n s$, and we can apply the induction hypothesis.

Step 7. We can factor n as uv, with u in the same copy of GL_2 as s, and $s^{-1} v s$ in I. So we are reduced to the case wan with n in that copy of GL_2.

Step 8. The element a can be factored as $a_1 a_2$, in which $a_1 \in A$ lies in the copy of GL_2 in which s lies, and a_1 commutes with this copy. We can also write $w = w_1 w_2$, with w_2 in that copy of GL_2 and

$\ell(w) \geq \ell(w_1)$. We are now considering $w_1 a_1 \cdot w_2 w_2 \cdot \overline{n}$, in which the last three terms lie in the copy of GL_2, a_1 commutes with that copy, and $w_1 \alpha > 0$. Suppose we can factor $w_2 a_2 \overline{n}$ as $\iota_1 \omega \iota_2$ in GL_2. Then

$$w_1 a_1 \iota_1 \omega \iota_2 = w_1 \iota_1 a_1 \omega \iota_2 = w_1 \iota_1 w_{12}^{-1} \cdot w_1 a_1 \omega \iota_2 \,,$$

and we are essentially finished (ready to apply induction, as I mentioned).

This concludes the first half of the claim.

I'll deal with uniqueness in the next section.

14. The affine Bruhat decomposition

Remark. What is special about the group GL_n is that there is a very efficient way to find the Cartan decomposition. For arbitrary reductive groups there is a very general way to carry out Iwahori factorizations that involves finding reduced factorizations of elements in the analogue of $\mathfrak{W}$. Here we are only required to factorize elements of W, which is in general far less work. For all split groups, the problem is reduced to a computation in GL_2, as it is here.

Very generally, express $IsI \cdot ItI$ for s, t simple reflections in $\widetilde{\mathfrak{S}}$.

As we shall see, the Iwahori factorization is a basic tool in the representation theory of reductive $\mathfrak{p}$-adic groups.

15. The building

LATTICE FLAGS. A **principal lattice flag** will be a sequence of lattices

$$L_0 \subset L_1 \subset \ldots \subset L_n$$

which reduces to aprincipal flag in $L_n/\varpi L_n$. That is to say (a) each L_i/L_{i-1} is isomorphic to $\mathbb{F} = \mathfrak{o}/\mathfrak{p}$ and (b) $L_0 = \varpi L_n$. Hence:

15.1. Lemma. *The group $\mathrm{GL}_n(\mathfrak{k})$ acts transitively on the set of principal lattice flags. The stabilizers of principal lattice flags are Iwahori subgroups.*

Reformulation of Proposition 13.4. Given two lattice flags ...

For $0 \leq i \leq n$ set

$$\mu_i = (0, \ldots, 0, 1, \ldots 1)$$

so that the j-th cooordinate of mu_i is 0 for $j < i$, and 1 for $j \geq i$. **The convex hull is a fundamental domain for $\mathfrak{S}_n$ on the unit cube, since we can permute to where $x_i \leq x_{i+1}$.**

Given a basis $e = (e_i)$, the lattice flag $\mathcal{LF}e = (L_\bullet$ determined by it is that for which L_i is the span of

$$\varpi^{\mu_i} e = \{e_1, \ldots, e_i, \varpi e_{i+1}, \ldots, \varpi e_n\} \,.$$

The basis e is said to be adapted to $L_\bullet$. Such a basis is unique up to multiplication on the right by a matrix in the Iwahori subgroup—upper triangular modulo $\mathfrak{p}$.

Now to prove uniqueness in Proposition 13.4. We follow roughly an argument similar that for uniqueness in the Bruhat decomposition can be used. This version uses the volume of $L_i \cap \mathcal{L}_j$ to construct the profile, where $\mathcal{L}_j$ is the standard lattice flag.

Alcoves for GL_n correspond to lattice flags (not obvious). Projection from GL_n to PGL_n. *What is my goal?* **I want to give a second way to explain $G = I\mathfrak{W}I$. I need a good description of $I/I \cap wIw^{-1}$. Somewhere, point out the complex is that of lattice flags stable under diagonal matrices. Associate some GL_2 to each root, even affine.**

Later: length and sizes of quotients of Iwahori groups.

NORMS. Suppose V to be a vector space of dimension n over $\mathfrak{k}$. A **norm** on V is a function $\|v\|$ from V to $\mathbb{R}_{\geq 0}$ with the following properties:

(a) for all c in $\mathfrak{k}$, v in V,

$$\|cv\| = |c|\,\|v\|\,;$$

(b) for all u, v

$$\|u+v\| \le \sup \|u\|, \|v\|\,;$$

(c) $\|v\| = 0$ if and only if $v = 0$.

Supose L to be a lattice. Define

$$\|v\|_L = \sup_{v/c \in L} |c|\,.$$

15.2. Lemma. *The function $\|v\|_L$ is a norm on V.*

For example, fix the standard basis e_i of $\mathfrak{k}^n$. Then for the lattice $\varpi^m \mathfrak{o}^n$ with basis $(\varpi^{m_i} e_i)$ we have

$$\|v\| = \sup_i |v_i/\varpi^{m_i}| = \sup |v_i| q^{m_i}\,.$$

It is often convenient to use instead the additive norm $\mathrm{ord}(x) = \log_q |x|$, so that this becomes

$$\mathrm{ord}_m(v) = \inf_i (\mathrm{ord}(v_i) + m_i)\,.$$

This suggests the following generalization. For any r in $\mathbb{R}^n$ define

$$\mathrm{ord}_r(v) = \inf_i (\mathrm{ord}(v_i) + r_i)\,,$$

and then set $\|v\|_r = q^{-\mathrm{ord}_r(v)}$. If the r_i are all integers then this is the norm associated to the lattice $\varpi^r \mathfrak{o}^n$.

15.3. Proposition. *For any r in $\mathbb{R}^n$ the function $\|v\|_r$ is a norm. Every norm is the one associated to some basis and some r in $\mathbb{R}^n$.*

Neiher basis nor r is unique. We shall understand this much better shortly.

15.4. Proposition. *For r in the convex hull of the μ_i the norm $\|v\|_{e,r}$ is independent of the adapted basis e.*

This gives a triangulation of $\mathbb{R}^n$, since *every point lies inside the transform of this simplex by a unique element of $\mathfrak{W}$.*

Unit cube, $\mathfrak{S}_n$.

Part IV. References

1. Armand Borel and George D. Mostow (editors), **Algebraic groups and discontinuous subgroups**, in *Proceedings of Symposia in Pure Mathematics* **IX**, American Mathematical Society, 1966.

2. Nagayoshi Iwahori, 'Generalized Tits systems', pp. 71–83 in [Borel-Mostow:1966].

3. Donald E. Knuth, **Sorting and searching**, volume III of **The art of computer programming**, Addison-Wesley, 1973.

4. Ian G. Macdonald, **Spherical functions on a group of p-adic type**, University of Madras, 1971.