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. One reason for doing this is to motivate the abstract definitions to come, and another is to make available at least one example 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 chapter by discussing these as well as permutations in general.

Next, I'll look at the groups $\mathrm{GL}_n(k)$ and $\mathrm{SL}_n(k)$ for an arbitrary coefficient field k. 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 chapter I'll let $k = \mathfrak{k}$ 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.

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

Contents

Part I. The symmetric group

1. Permutations

Let $I = [1, n]$. A **permutation** of I is just an invertible map from I to itself. The permutations of I form a group $\mathfrak{S}_I$, also written as $\mathfrak{S}_n$.

There are several ways to express a permutation σ. One is just by writing the **permutation array** $[\sigma(\iota_i)]$. Another is as a list of **cycles**. For example, the cycle $(i_1 \,|\, i_2 \,|\, i_3)$ takes

$$\begin{aligned} i_1 &\longmapsto i_2 \\ i_2 &\longmapsto i_3 \\ i_3 &\longmapsto i_1 \,. \end{aligned}$$

If σ is a permutation of I then σ is said to **invert** (i, j) if $i < j$ but $\sigma(i) > \sigma(j)$. These 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. This requires $n(n-1)/2$ comparisons.

A **transposition** just interchanges two of the integers in I, and each **elementary transposition** $\langle j \rangle = (j \,|\, j+1)$ interchanges two neighbouring ones.

1.1. Lemma. *(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 $(j \,|\, j+1)$ inverts only the single pair $(j, j+1)$;
(d) there are just as many inversions of σ as of its inverse;
(e) there is a unique permutation which inverts every pair and hence has $n(n-1)/2$ inversions.

Proof. To prove (a), it suffices to prove (b). 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), 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 $\langle i \rangle$ inverts only the last pair.

For (d), σ inverts (i, j) if and only if σ^{-1} inverts $\big(\sigma(j), \sigma(i)\big)$.

For (e), take $\sigma(i) = n - i + 1$.

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 $\langle j \rangle \langle k \rangle$ is a cycle of order three, but otherwise $\langle j \rangle$ and $\langle k \rangle$ commute and the product has order two:

$$\begin{aligned} \langle i \rangle \langle i+1 \rangle &= (i \,|\, i+1)(i+1 \,|\, i+2) = (i \,|\, i+1 \,|\, i+2) \\ \langle j \rangle \langle k \rangle &= (j \,|\, j+1)(k \,|\, k+1) \\ &= \langle k \rangle \langle j \rangle \,. \end{aligned}$$

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. Something stronger is in fact true:

1.2. Proposition. *(a) The group $\mathfrak{S}_n$ is generated by the elementary transpositions;*
(b) the minimal length $\ell(\sigma)$ of an expression for the permutation σ as a product of elementary transpositions is equal to the number of its inversions.

Proof. Suppose σ to be a permutation other than the trivial one. Read along in the array $\big(\sigma(i)\big)$ until you find an inversion, $\sigma(i) > \sigma(i+1)$. Let $\tau = \sigma\,\langle i\rangle$. The array $\big(\tau(j)\big)$ is the same as that for σ, except that the positions of $\sigma(i)$ and $\sigma(i+1)$ have been swapped. So the number of inversions for τ is exactly one less than for σ itself—we have removed precisely the pair $(i, i+1)$ from the inversion list. 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. It does this when the permutation we are considering is the trivial one. So we get an equation

$$\sigma\langle i_1\rangle \dots \langle i_n\rangle = I, \quad \sigma = \langle i_n\rangle \dots \langle i_1\rangle\,.$$

This argument says that if $I(\sigma)$ is the set of inversions of σ and $(i, i+1)$ is one of them, then

$$I(\sigma\langle i\rangle) = I(\sigma) - \{(i, i+1)\}\,,$$

and σ may hence be written as a product of $|I(\sigma)|$ elementary transpositions. Hence $\ell(\sigma) \le |I(\sigma)|$. The converse is just as simple—if $(i, i+1)$ is not in $I(\sigma)$ then it is one for $\sigma\langle i\rangle$, so that

$$I(\sigma\langle i\rangle) = I(\sigma) \cup \{(i, i+1)\}\,.$$

This together with the previous equation imply that

$$|I(\sigma)| \le \ell(\sigma)\,.$$

This proof is, in effect, the well known **bubble sort** algorithm for sorting an array.

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

2. The geometry of permutations

If k is an arbitrary field, the group $\mathfrak{S}_n$ acts on the n-dimensional vector space k^n by permuting basis elements:

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

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

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

In other words, it also permutes the coordinates of a vector. We thus have an embedding of $\mathfrak{S}_n$ into $\mathrm{GL}_n(k)$.

A transposition $(i\,|\,j)$ swaps coordinates x_i and x_j, hence has determinant -1. If a permutation can be represented as a product of ℓ transpositions then its sign is $(-1)^\ell$.

Now let $k = \mathbb{R}$ and $V = \mathbb{R}^n$. Assign to it the usual Euclidean inner product. Let $\alpha = \alpha_{i,j}$ be the linear function $x_i - x_j$ on V, and let $\alpha^\vee$ be the vector in v associated to it by Euclidean duality—it has i-th coordinate 1, j-coordinate -1, and

$$\langle \alpha, v\rangle = \alpha \bullet v$$

for all v in V. Let $\alpha_i = \alpha_{i,i+1}$ for $i = 0, \dots, n-1$. For reasons connected to the structure of the group GL_n, the $\alpha_{i,j}$ are called **roots**, and the α_i are called **elementary roots**. The ones with $i < j$ are called **positive roots**. For $i < j$ we have

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

The orthogonal projection of a vector v onto the line through $\alpha^\vee$ is

$$\left(\frac{\alpha^\vee \bullet v}{\alpha^\vee \bullet \alpha^\vee}\right)\alpha^\vee = \left(\frac{\alpha^\vee \bullet v}{2}\right)\alpha^\vee\,.$$

The transposition $s = (i \mid j)$ amounts to orthogonal reflection in the **root hyperplane** $x_i = x_j$ or $\alpha_{i,j} = 0$:

$$r_s v = v - (\alpha^\vee \bullet v)\alpha^\vee = v - \langle \alpha, v \rangle \alpha^\vee \,.$$

The coordinates of any vector may be permuted to a unique weakly decreasing array. In other words, the region $x_1 \geq x_2 \geq \ldots \geq x_n$ is a fundamental domain for $\mathfrak{S}_n$ in $\mathbb{R}^n$. 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 $\Theta \subseteq \Delta = \{0, \ldots n-1\}$—to Θ corresponds the face where $\alpha_i = 0$ for i in Θ, otherwise $\alpha_i > 0$. The open cone itself is $C_\emptyset$.

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 λ is positive if and only if $\lambda > 0$ on C.

Every root λ corresponds to a hyperplane $\lambda = 0$, 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 $\sigma^{-1}(\lambda) < 0$.

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

(a) the root $\sigma(\alpha_{i,j})$ is negative;
(b) the hyperplane $\alpha_{i,j} = 0$ 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

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

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

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

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

In other words, the number of inversions for 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 $s = \langle i \rangle$. 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 $\lambda = 0$ where $\lambda = w\alpha_s$. If wC is its terminal chamber the gallery corresponds to an expression for w as a product of elementary reflections.

2.2. 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. ∎

Later on, I shall need a certain factorization of permutations different from that by elementary transpositions. For $1 \le i \le n$ let W_i be the subgroup of $\mathfrak{S}_n$ made up of σ with (a) $\sigma(j) = j$ for $j < i$ and (b) $\sigma(j) < \sigma(k)$ for $i < j < k \le n$. In other words, the array $\big(\sigma(i)\big)$ is trivial up to i, non-decreasing beyond it. In particular, the group W_{n-1} is just the group of permutations of $\{n-1, n\}$.

2.3. Proposition. *The map $\sigma \mapsto \sigma(i)$ is a bijection of W_i with $[i, n]$.*

Proof. The condition on $\sigma(j)$ for $j < i$ means that the map from W_i into the group of permutations of $[i, n]$. Given that $\sigma(i) = k$, there is then a unique increasing sequence of the remaining numbers $\ell \neq k$ in $[i, n]$, so that there exists exactly one σ in W_i mapping $[i+1, n]$ into that sequence and taking i to k. ∎

2.4. Corollary. *Every σ in $\mathfrak{S}_n$ may be expressed as a unique product*

$$\sigma = \sigma_1 \dots \sigma_n$$

with σ_i in W_i.

Proof. Given σ, the Proposition tells us that exists a unique σ_1 in W_1 with $\sigma_1(1) = \sigma(1)$. Then $\sigma\sigma_1^{-1}$ fixes 1. There exists a unique σ_2 in W_2 with $\sigma_2(i) = \sigma\sigma_1^{-1}(i)$ for $i = 1, 2$. Etc. ∎

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:

$$\begin{aligned} (w_\sigma m)_{i,j} &= m_{\sigma^{-1}(i),j} \\ (m w_\sigma)_{i,j} &= m_{i,\sigma(j)} \end{aligned}$$

and hence

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

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:

3.1. 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 **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 $\lambda_{i,j} = \varepsilon_i - \varepsilon_j$.

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 $\lambda = \lambda_{p,q}$ corresponds an embedding, which I express as $\lambda^\vee$ 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 $\lambda^\vee$ 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 $\lambda^\vee$ 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 and Schubert cells

A rectangular matrix x is said to be in **permuted column echelon form** if it has the following two properties:

- every column has at least one non-zero entry, and the last non-zero entry in each column is 1;
- all entries to the right of the last non-zero entry in column i are zero.

Suppose x to be a matrix of dimensions $n \times c$. Let σ be the map from $[1, c]$ to $[1, n]$ such that the last non-zero entry in column i is in row $\sigma(i)$. In these circumstances, the map σ is called the row assignment of x. The second property above is that $x_{i,j} = 0$ for $j > \sigma(i)$. It is easy to see that

- if $c = 1$, the matrix x will be in permuted column echelon form if and only if its last non-zero entry is 1;
- if $c > 1$, it will be in permuted column echelon form if and only if (i) the matrix of its first $c - 1$ columns, which I call its **prefix matrix**, is in permuted column echelon form; and (ii) the last non-zero entry in column c is 1, and $a_{\sigma(i),c} = 0$ for $1 \le i < c$, where σ is the row assignment of the prefix matrix.

Hence that if an $n \times c$ matrix is in permuted column echelon form then its row assignment is injective and its columns are linearly independent.

The matrices of permuted echelon form may be classified roughly by the row assignment, the embedding σ of $[1, c]$ into $[1, n]$. Given σ, the conditions on a matrix with row assignment σ are:

- $a_{\sigma(j),j} = 1$;
- $a_{i,j} = 0$ for $i > \sigma(j)$;
- $a_{\sigma(j),k} = 0$ for $k > j$.

Among invertible 3×3 matrices, for example, we have the following types (where an asterisk is an arbitrary entry):

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 & 0 \\ 0 & * & 1 \\ 0 & 1 & 0 \end{bmatrix}, \quad \begin{bmatrix} * & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$
$$\begin{bmatrix} * & * & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad \begin{bmatrix} * & 1 & 0 \\ * & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} * & * & 1 \\ * & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}. \tag{4.1}$$

Here is the basic result of this section:

4.2. Theorem. *If M is an $n \times c$ matrix of column rank c, then there exists a unique matrix in permuted column form, which can be obtained from M by multiplying on the right by an invertible upper triangular matrix.*

Of course the upper triangular matrix is also unique.

Proof. The proof, which amounts to an algorithm, is a straightforward variation of the familiar 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 (or, of course, subtracting from) one column a multiple of a previous one. For example

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} 1 & x \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} a & ax + b \\ c & cx + d \end{bmatrix}.$$

The proof of both existence and uniqueness is by induction on the number of columns c. If $c = 1$ then we multiply the first column by the inverse of the last non-zero entry in that column. This is certainly the only way to turn M into permuted column echelon form.

Suppose M to have $c > 1$ columns, and that we have modified the first $c - 1$ columns into permuted echelon form. We first subtract from the c-th column suitable multiples of the first $c - 1$ columns to get condition (b) satisfied for the entries in the last column. Then we multiply the last column by the inverse of the last non-zero entry. This gives existence.

For uniqueness, this is trivial if $c = 1$. Otherwise, let V_i be the vector subspace spanned by the first i columns of M. The subspaces V_i remain invariant under right multiplication by an upper triangular matrix. Assume inductively that the first $c - 1$ columns of the final matrix are uniquely determined. The column vector x we calculate for the c-th column is then determined up to scalar multiplication by the conditions (a) that $x_i = 0$ if $i > \sigma(j)$ (b) x is in V_c. The normalization condition then determines it uniquely.

Suppose M to be an invertible $n \times n$ matrix in permuted column echelon form, and that the last non-zero entry in column j is in row $\sigma(j)$. The map $j \mapsto \sigma(j)$ is a permutation σ of $[1, n]$, and the matrix w_σ will also be in permuted column echelon form. Since multiplication on the left by a unipotent upper triangular matrix amounts to adding to various rows linear combinations of subsequent rows, it is simple to see that we can multiply M by an upper triangular matrix to obtain w_σ. Therefore:

4.3. Corollary. *If M is an invertible $n \times n$ matrix, it may be factored uniquely as sp, with s in permuted column echelon form and p upper triangular.*

I shall now reformulate this slightly by characterizing diferently the matrices in permuted column echelon form.

The form of a matrix in permuted echelon form associated to a permutation matrix w_σ suggests multiplying it on the right with the inverse $w_{\sigma^{-1}}$. For example

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & * & 1 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & * \\ 0 & 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} * & * & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 1 & * & * \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

As this suggests, *in these circumstances $xw_{\sigma^{-1}}$ will always be an upper triangular unipotent matrix.* We can in fact specify just what kind of a unipotent matrix it will be.

Let N be the group of upper triangular unipotent matrices, $\overline{N}$ that of lower triangular unipotent matrices, and for any permutation matrix $w = w_\sigma$ let

$$N_w = N \cap w\overline{N}w^{-1}$$

In other words, the upper triangular unipotent matrix n lies in N_w if and only if the matrix $w^{-1}nw$ lies in $\overline{N}$. Since $(w^{-1}nw)_{i,j} = n_{\sigma(i),\sigma(j)}$, this means in all

$$n_{i,i} = 1 \tag{a}$$
$$n_{i,j} = 0 \quad \text{if } i > j \tag{b}$$
$$(w_\sigma^{-1} n w_\sigma)_{i,j} = n_{\sigma(i),\sigma(j)} = 0 \quad \text{if } i < j \tag{c}$$

and hence

$$n_{i,j} \neq 0 \text{ implies that } i < j \text{ and } \sigma^{-1}(i) > \sigma^{-1}(j)\,.$$

In other words, the non-zero $n_{i,j}$ are among those where (i, j) is an inversion of σ^{-1}. The dimension of N_w is therefore equal to $\ell(\sigma^{-1}) = \ell(\sigma)$. For example, N_1 is trivial, whereas N_{w_ℓ} is all of N.

4.4. Proposition. *An invertible $n \times n$ matrix x is in permuted column echelon form with σ as its associated permutation if and only if xw_σ^{-1} lies in N_w.*

Proof. Suppose for the moment that $x = (x_{i,j})$ is any matrix, σ a permutation. Then $y = xw_\sigma^{-1}$ is obtained by permuting columns of X, and more precisely

$$y_{i,j} = x_{i,\sigma^{-1}(j)}\,.$$

The conditions that x is in permuted column form are

$$x_{\sigma(j),j} = 1 \tag{a'}$$
$$x_{i,j} = 0 \text{ if } i > \sigma(j) \tag{b'}$$
$$x_{i,j} = 0 \text{ if } i = \sigma(k) \text{ and } j > k\,, \tag{c'}$$

which are exactly the equivalents of (a)–(c) for y.

4.5. Proposition. (Bruhat decomposition) *Let P be the group of upper triangular matrices. Every element of GL_n may be factored uniquely as $nw_\sigma p$ with σ in $\mathfrak{S}_n$, n in N_w, and p in P.*

Proof. Uniqueness follows from the previous result, Theorem 4.2, but it can also be shown directly. It comes down to showing that if u and v are upper triangular matrices and w_1 and w_2 permutation matrices such that $w_1^{-1}uw_2 = v$, then $w_1 = w_2$.

Let $w_i = w_{\sigma_i}$. Then

$$v_{i,j} = u_{\sigma_1(i),\sigma_2(j)}$$

and u and v are both upper triangular matrices with non-zero diagonal entries. Therefore

$$\sigma_1(i) \geq \sigma_2(i), \quad \sigma_2^{-1}\sigma_1(i) \geq i$$

for all i which implies that $\sigma_2^{-1}\sigma_1$ is trivial.

We shall see other reasons for uniqueness later on.

I call the expression $x = nw_\sigma p$ the **Bruhat normal form** of x.

There is another way to understand the structure of the double coset PwP. First of all since P factors as $AN = NA$ and permutation matrices conjugate diagonal matrices among themselves, $PwP = NwP$, and the group N acts transitively on PwP/P. The isotropy subgroup of wP is $N \cap wNw^{-1}$. The point is that group N_w is complementary to $N \cap wNw^{-1}$, and maps canonically onto the quotient $N/N \cap wNw^{-1}$. This is the content of:

4.6. Proposition. *If x and y are permutation matrices then the product map is a bijection of $N_x \times xN_yx^{-1}$ with N_{xy} if $\ell(xy) = \ell(x) + \ell(y)$.*

Proof. This is an easy induction on the length of y.

4.7. Corollary. *For any permutation matrix w*

$$N_w wP = PwP\,.$$

5. Flags

There is a geometrical interpretation of the calculation in the last section. A **flag** $\mathcal{F}$ in k^n is an increasing sequence of subspaces

$$F_0 = \{0\} \subset F_1 \subset F_2 \subset \cdots \subset F_c$$

It is called a **principal flag** if each dimension increment is one. As mentioned in the proof, any $n \times c$ matrix of rank c determines a flag

$$V_0 = \{0\} \subset V_1 \subset \ldots \subset V_c$$

where V_i is the space in k^n spanned by the first i columns. This is a principal flag of rank c. Multiplying the matrix on the right by any invertible upper triangular matrix preserves the flag. The Proposition describes a bijection between the set of principal flags of rank c and the $n \times c$ matrices in permuted column echelon form.

There is a direct way to see, at least, where σ comes from. Suppose $\mathcal{F}_*$ to be the standard flag in k^n and suppose $\mathcal{F}$ to be any other flag. For V_i one of the components of $\mathcal{F}$, consider the map taking j in $[0, n]$ to $\dim V_i \cap V_{*,j}$. This is a non-decreasing function of i whose graph I call the **profile** of V_i (relative to $\mathcal{F}_*$). In going from the profile of V_i to that of V_{i+1} a new layer is added to that of V_i, starting at some point $n(i)$. *The sequence $n(i)$ is the permutation σ.* The assembly of profiles matches exactly the algorithm producing a matrix in column echelon form.

The following pictures show the profiles for $\sigma = [5, 2, 3, 1, 4]$, showing the new row added at each step.

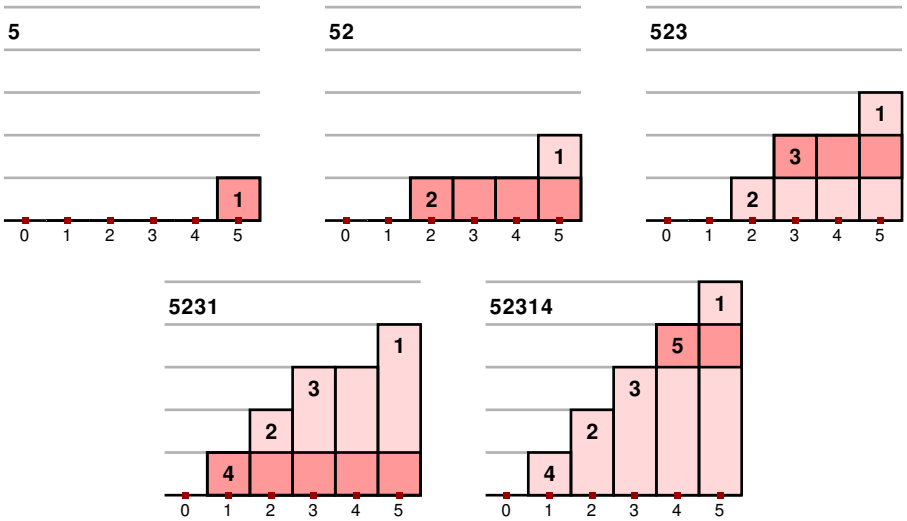

6. Calculating with the Bruhat normal form

Previous sections tell us that every element of $G = \mathrm{GL}_n$ can be expressed uniquely as $g = nwp$ with w in W and n in N_w. This representation of elements of G can be used as a convenient normal form, the **Bruhat normal form**. In particular, all group operations can be carried out strictly in terms of it. This is especially important since the same technique works for all reductive groups, when a good matrix representation might not be available.

One consequence of the Bruhat decomposition is that G is generated by P and W. Multiplication in the group comes down to a few questions. Suppose we are given $g = nwp$. **(a)** *For p_* in P, what is p_*g?* That's too easy. The important point is that the product will lie again in N_wwP. **(b)** *Given w in W, what is wg?* In particular, can we find the Bruhat normal form of this product? We know how to factor any w in W into a product of elementary transpositions, so this reduces to a simpler question. **(c)** *Given the elementary transposition s, what is the Bruhat normal form of sg?* This is more interesting.

Let $s = s_\alpha$, where $\alpha = \alpha_{i,i+1}$. Suppose $g = nwp$ with n in N_w. There are two cases.

Suppose first that $sw > w$. $ws > w$, which means that $w\alpha > 0$. We can write $p = an_\alpha n^*$ with n_α in N_α and n_* in $\prod_{\beta\neq\alpha} N_\beta$. But then

$$sg = sn_wwan_\alpha n_* = n_wwan_\alpha ss^{-1}n_*s = n_wwan_\alpha sn_{**}$$

where $n_{**} = s^{-1}n_*s$ lies again in N. But then we continue

$$wan_\alpha s = w \cdot an_\alpha a^{-1} \cdot as = wn_{*,\alpha}s \cdot s^{-1}as = wn_{*,\alpha}a_* = wn_{*,\alpha}w^{-1} \cdot wsa_* = n_{*,w\alpha}wsa$$

but under the assumption $ws > w$ we know that $N_{ws} = N_wN_{w\alpha}$, so $n_{*,w\alpha}$ lies in N_w.

Under the same assumption, the same argument run in reverse shows that every element of $N_{ws}wP$ can be written as a product of elements in N_wwP and $N_\alpha sP$.

Now let's look at the case $x = ws < w$. Here we can write $w = xs > x$. Thus as we have just proven every element of N_wwP as a product g_xg_s of elements in N_xxP and $N_\alpha sP$. But then

$$ws = g_xg_ss, \quad g_ss = n_\alpha sps$$

and

$$n_\alpha sps = n_\alpha san_{*,\alpha}ss^{-1}n_*s$$

where $s^{-1}n_*s$ lies in N. We are therefore reduced to the case $G = \mathrm{GL}_2$, which I leave as an exercise. Applying the previous case to x and s, we see in the end an explicit way to represent gs as an element of either $N_x x P$ or $N_w w P$.

Define $C(w)$ to be the double coset PwP. I have shown, among other things:

6.1. Proposition. *For any s in S and w in W*

$$C(w)C(s) = \begin{cases} C(ws) & \ell(ws) > \ell(w) \\ C(w) \cup C(ws) & \textit{otherwise} \end{cases}$$

We shall see later similar results for any reductive group. It will have a Bruhat decomposition, and this will give rise to a very explicit way to represent elements of the group. There will again be explicit formulas describing the group structure purely in these terms. What is unusual about GL_n is that its representation by matrices allows in particular a very convenient way to describe the unipotent elements. In general, an explicit and somewhat complicated description in terms of generators and relations is required to replace matrix multiplication. This will be in a way a substitute for the matrix representation.

At some point later on we will need an explicit calculation for GL_2.

6.2. Proposition. *For $x \neq 0$*

$$\begin{bmatrix} 1 & 0 \\ x & 1 \end{bmatrix} = \begin{bmatrix} 1 & x^{-1} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} x & 0 \\ 0 & -x^{-1} \end{bmatrix} \begin{bmatrix} 1 & x^{-1} \\ 0 & 1 \end{bmatrix}$$

7. 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.

7.1. Proposition. *Suppose that x is an $n \times n$ matrix. It can be factored as $x = \nu a n$ 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 $\nu a u$ 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 $x \leq y$, if x is a product of a subsequence of the s_i.

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

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 Proposition 6.1, 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.

7.3. Proposition. *If σ and π are two permutations, $\pi \leq \sigma$ 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 $a_i \leq b_i$ for all i. Then $\pi \leq \sigma$ if and only if the initial sequence of π is less than or equal to that of σ.

8. Parabolic subgroups

If n_i is the of difference 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. The stabilizer of a flag is called a **parabolic subgroup**. If the partition is $n = 1 + \cdots + 1$ then the flag is called **principal**, and the stabilizer is a **Borel subgroup**.

8.1. Proposition. *The group GL_n acts transitively on flags associated to the same partition.*

Proof. It acts transitively on bases of k^n, and every flag is one associated to some compatible basis.

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

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

and an associated embedding of the product $\prod \mathrm{GL}_{n_i}(k)$ into $\mathrm{GL}_n(k)$, *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. *Any parabolic subgroup P is its own normalizer in GL_n.*

8.4. Proposition. *If $k = \mathfrak{k}$ 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.

Part III. p-adic fields

9. Lattices

I now take up the special features of $\mathrm{GL}_n(\mathfrak{k})$ where $\mathfrak{k}$ is a $\mathfrak{p}$-adic field. Just about everything I'll say in this case 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.

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

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

9.1. 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.

9.2. 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 $\mathfrak{o}^n$, an integer $r \geq 0$, and integers $m_1, m_2, \dots, m_r$ with*

$$m_1 \leq m_2 \leq \cdots \leq m_r$$

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

Recall that ϖ is a generator of $\mathfrak{p}$.

Proof. The proof proceeds by induction on n, and will be constructive.

Suppose $\ell_1, \ell_2, \dots, \ell_k$ are given generators of L. It can happen that all ℓ_i are zero, in which case we are through. Otherwise 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 whose coordinate is of the form $\varpi^{m_1} u$ with u a unit in $\mathfrak{o}$, which we may as well assume to be ℓ_1. According to Lemma 9.1, there exists a basis of $\mathfrak{o}^n$ whose first element is $e_1 = \varpi^{-m_1}\ell_1$. Choose coordinates on $\mathfrak{k}^n$ corresponding to this basis. By subtraction of a multiple of ℓ_1, we may assume that $\ell_2, \dots, \ell_k$ have first coordinate 0. In effect, L is now the direct sum of the rank one $\mathfrak{o}$-module generated by ℓ_1 and the module L_* generated by the rest of the ℓ_i. The latter may be considered a lattice in $\mathfrak{k}^{n-1}$.

If $n = 1$, we are through, while if $n > 1$ we can apply induction and assume the Proposition to be true for L_*.

As for uniqueness, choose the ℓ_i to make up a basis of L. If we write out the matrix whose columns are the vectors ℓ_i then 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 $\mathrm{GL}_r(\mathfrak{o})$ and doesn't change these characterizations.

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

The group $\mathrm{GL}_n(\mathfrak{k})$ acts on the set of all full lattices, considered as subsets of $V = \mathfrak{k}^n$. The stabilizer of $\mathfrak{o}^n$ is $\mathrm{GL}_n(\mathfrak{o})$. According to Corollary 9.3, the group $\mathrm{GL}_n(\mathfrak{k})$ acts transitively on full lattices.

Let A be the group of diagonal matrices, A^{++} those whose entries are ϖ^{m_i} with $m_i \leq m_{i+1}$.

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

$$g = \gamma_1 \, t \, \gamma_2$$

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

Proof. The group $\mathrm{GL}_n(\mathfrak{k})$ acts on the right on bases of $\mathfrak{k}^n$, and the group $\mathrm{GL}_n(\mathfrak{o})$ acts on the right on bases of any full lattice. Given g in $\mathrm{GL}_n(\mathfrak{k})$, apply Proposition 9.2 to the lattice $g\,\mathfrak{o}^n$. We get a basis (ℓ_i) of $\mathfrak{o}^n$ such that $\ell\, t = (\varpi^{m_i}\ell_i)$ is a basis of $g\,\mathfrak{o}^n$, where $t = \mathrm{diag}(\varpi^{m_i})$. There exists some γ_1 in $\mathrm{GL}_n(\mathfrak{o})$ such that $\varepsilon\gamma_1 = \ell$, and there exists another γ_2 in $\mathrm{GL}_n(\mathfrak{o})$ such that $\varepsilon g = \ell t \gamma_2$. But then

$$\varepsilon\, g = \varepsilon\, \gamma_1\, t\, \gamma_2$$

where ε is the standard basis. ∎

9.5. 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 9.1 is part of a basis, so we can find γ with $\gamma V_1 = \mathfrak{k}\varepsilon_1$. From now on we may assume $V_1 = \mathfrak{k}\varepsilon_1$. Look next on 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. ∎

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

This proves again that $P\backslash G$ is compact for parabolic subgroups P.

10. Volumes

Let $K = \mathrm{GL}_n(\mathfrak{o})$. We shall next compute the volume of the open set KtK in GL_n explicitly under the assumption that the volume of K is 1. Equivalently, we shall compute $|KtK/K|$. The nature of the formula depends very definitely on t. 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

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

10.1. Proposition. *Suppose*

$$m_1 \le m_2 \le \dots \le m_n$$

and let

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

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

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

Let

$$M = M_t = \mathrm{GL}_{n_1} \times \mathrm{GL}_{n_2} \times \dots \times \mathrm{GL}_{n_k}$$

and define

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

We embed M as diagonal blocks in GL_n. Then

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

Here $\delta_\emptyset$ is the character

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

of the group of diagonal matrices (a_i).

Proof. The map $k \mapsto ktK/K$ induces a bijection of $K/K \cap tKt^{-1}$ with KtK/K. The group $K \cap tKt^{-1}$ consists of all integral invertible matrices g with $g_{i,j} \equiv 0 \bmod \varpi^{m_i - m_j}$. 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 $K \cap tKt^{-1}$ in K is the same as that of the index of $K_m \cap tK_mt^{-1}$ 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 $K_m \cap tK_mt^{-1}$? The image of $K_m \cap tK_mt^{-1}$ 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 $K_m \cap tK_mt^{-1}$ 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 $K_m \cap tK_mt^{-1}$ 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

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

11. Iwahori subgroups

Let I be the inverse image in $\mathrm{GL}_n(\mathfrak{o})$ of the group of upper triangular matrices in $\mathrm{GL}_n(\mathbb{F}_q)$. An **Iwahori subgroup** of $\mathrm{GL}_n(\mathfrak{k})$ is any conjugate of I.

11.1. Proposition. *We have the factorization $G = IWP$.*

Proof. Use the Bruhat factorization for $\mathrm{GL}_n(\mathbb{F}_q)$. This tells us that $K = IW(K \cap P)$. Apply Corollary 9.6.

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.

11.2. 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 7.1, since the determinant of an element of an Iwahori subgroup is a unit in $\mathfrak{o}$.

11.3. 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 Proposition 4.5.

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. But the quotient

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

is now the semidirect product of $A/A(\mathfrak{o})$ and W. This plays the role of W. The group $A/A(\mathfrak{o})$ is isomorphic to $\mathbb{Z}^n$, a free group of rank n over $\mathbb{Z}$. It may be identified with the group of diagonal matrices whose entries are powers of ϖ. Call an $n \times r$ matrix X with $r \le n$, 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:

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

We want now to prove a $\mathfrak{p}$-adic version of the Bruhat decomposition—that every element g of $\mathrm{GL}_n(\mathfrak{k})$ factors uniquely as $g = b_1 \widetilde{w}\, b_2$ where b_i is an element of I and $\widetilde{w}$ is in $\widetilde{W}$.

To go with this claim is an algorithm involving **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:

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

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

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

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

Theorem. *Any $n \times r$ matrix of column rank r, with $r \leq n$, can be reduced by elementary Iwahori row and column operations to a unique Iwahori-reduced matrix.*

Proof. By induction on r.

For $r = 1$, suppose the vector is of the form

$$\begin{bmatrix} \varpi^n x \\ \varpi^n u \\ \varpi^{n+1} y \end{bmatrix}$$

with x, y integral, u a unit, the entry $u\varpi^n$ in the i-th row. First multiply the i-row by $1/u$ to get the middle entry equal to ϖ^n. Then left multiplication by the matrix that is unipotent except in the i-column and whose i-th column is

$$\begin{bmatrix} -x \\ 1 \\ -\varpi y \end{bmatrix}$$

will reduce the vector to

$$\begin{bmatrix} 0 \\ \varpi^n \\ 0 \end{bmatrix} .$$

Suppose now that the first r columns have been reduced, and look at column $r + 1$. Let R be the set of rows with a non-zero entry in the first r columns. Choose i in the complement of R as in the first step, so that $|a_{j,1}| \leq |a_{i,1}|$ for $j \leq i$ and not in R and $|a_{j,1}| < |a_{i,1}|$ for $j > i$ and not in R. Apply Iwahori row operations to get zeroes in all but row i and the rows in R. Let's look at just one of the rows in R as well as row i. We have one of these configurations: of non-zero entries in these rows:

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

where u is a unit.

That is to say, we are now reduced to essentially 2×2 problem. Let's suppose we are in the first case ...

$$\begin{bmatrix} \varpi^k & u\varpi^n \\ 0 & \varpi^\ell \end{bmatrix}.$$

If $k \leq n$ we can subtract $u\varpi^{n-k}$ times the first column from the second to get

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

If $\ell \leq n$ we can subtract $u\varpi^{n-\ell}$ times the second row from the first to get the same matrix.

So now we may assume $k > n$ and $\ell > n$. Subtract from the second row $\varpi^{\ell-n}/u$ times the first. This gives

$$\begin{bmatrix} \varpi^k & u\varpi^n \\ -\varpi^{k+\ell-n}/u & 0 \end{bmatrix}.$$

Divide the second column by u, multiply the second row by $-u$:

$$\begin{bmatrix} \varpi^k & \varpi^n \\ \varpi^{k+\ell-n} & 0 \end{bmatrix}.$$

Subtract ϖ^{k-n} times the second column from the first to get

$$\begin{bmatrix} 0 & \varpi^n \\ \varpi^{k+\ell-n} & 0 \end{bmatrix},$$

which is Iwahori reduced.

... and then in the second:

$$\begin{bmatrix} 0 & \varpi^\ell \\ \varpi^k & u\varpi^n \end{bmatrix}.$$

where u is a unit.

If $n > \ell$ or $n \geq k$ this can be reduced to

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

So now we assume $n \leq \ell$, $n < k$. Subtract ϖ^{k-n}/u times the second column from the first:

$$\begin{bmatrix} -\varpi^{k+\ell-n}/u & \varpi^\ell \\ 0 & u\varpi^n \end{bmatrix}.$$

Multiply and divide by u:

$$\begin{bmatrix} \varpi^{k+\ell-n} & \varpi^\ell \\ 0 & \varpi^n \end{bmatrix}.$$

Subtract $\varpi^{\ell-n}$ times the second row from the first:

$$\begin{bmatrix} \varpi^{k+\ell-n} & 0 \\ 0 & \varpi^n \end{bmatrix}.$$

This concludes the first part of the proof.

As for uniqueness, an argument similar that for uniqueness in the Bruhat decomposition can be used. This version uses the volume of $L_i \cap L_{*,j}$ to construct the profile, where $L_{*,j}$ is the standard lattice flag.