Irredundant Generating Sets of Finite Nilpotent Groups

Liang Ze Wong
May 2012
Advisor: Professor R. Keith Dennis

May 22, 2012

Abstract

It is a standard fact of linear algebra that all bases of a vector space have the same cardinality, namely the dimension of the vector space over its base field. If we treat a vector space as an additive abelian group, then this is equivalent to saying that an irredundant generating set for a vector space must have cardinality equal to the dimension of the vector space. The same is not true for groups in general. For example, even a relatively simple group like $\mathbb{Z}_6$ can be generated by either 1 or 2 elements ($\langle 1\rangle = \langle 5\rangle = \langle 2,3\rangle = \langle 4,3\rangle$). This paper seeks to count the number of irredundant generating sets for direct products of elementary abelian groups, which turns out to be easily generalizable to finite nilpotent groups. We first define a function that counts partitions of a disjoint union of sets such that each block of the partition . This function allows us to break down the problem such that we need only consider the direct summands of the group. Since these turn out to be finite vector spaces, we use linear algebraic methods to study their properties. By combining formulas from each vector space with the function that counts special partitions, we are able to count irredundant generating sets of the original group.

Acknowledgements

I am very grateful to professor R. Keith Dennis for introducing me to the question that this paper seeks to answer, for agreeing to supervise this thesis on such short notice, and for much mathematical advice and guidance that has proved invaluable whenever I had gotten stuck.

This thesis covers group theory, posets, q-analogues and linear algebra, and I am indebted to the professors here at Cornell who have introduced me to these subjects. It has been very fun seeing ideas from different fields interacting together. In particular, I would like to thank professors Louis Billera and Ed Swartz for teaching me to think combinatorially, and professor Dennis for introducing me to the links between combinatorics and group theory.

Contents

Chapter 1

Introduction

Definition 1.0.1 Given a group, G, and a finite subset $S \subseteq G$, S is an **irredundant generating k-set** of G if it satisfies the following properties:

1. $|S| = k$
2. $\langle S \rangle = G$
3. $\forall g \in S, \langle S \setminus g \rangle \lneq \langle S \rangle$

A subset which satisfies condition 2 is **generating**, while a subset that satisfies condition 3 is **irredundant**.

Let $r(G)$ denote the smallest k such that there exist an irredundant generating k-set, and $m(G)$ denote the largest such k. If G is a vector space of dimension n, an irredundant generating set is simply a basis. All bases of a vector space have the same cardinality, so $r(G) = m(G) = n$. For a general group, however, it is usually the case that $r(G) \lneq m(G)$. Tarski's irredundant basis theorem [9] states that if $r(G) \leq k \leq m(G)$, then there exist irredundant generating k-sets. We would like to answer the question, "How many?"

Definition 1.0.2 Let $\phi_k(G)$ denote the number of irredundant generating k-sets of G.

If G is a finite vector space of dimension d over a field $\mathbb{F}$, then $\phi_k(G) \neq 0$ iff $d = k$, and in particular, is equal to $\frac{1}{d!}|GL(\mathbb{F}, d)|$. However, determining $\phi_k(G)$ turns out to be very tricky for groups that are not vector spaces. One might think that a similarly easy result might hold for cyclic groups, but this is not the case. Consider any cyclic group of square-free order, $G = \mathbb{Z}_m, m = p_1 p_2 \dots p_n$, where $p_i, 1 \leq i \leq n$, are distinct primes. Then $\phi_1(G)$ counts the number of generators of G, which are simply the integers (mod m) that are relatively prime to m. So $\phi_1(G) = \phi(m)$, where the ϕ on the right is the Euler totient function.

At the other end of the spectrum, $\phi_n(G) = \phi(m)$ as well. This follows from the direct decomposition of G as $\mathbb{Z}_{p_1} \times \cdots \times \mathbb{Z}_{p_n}$. An element $g \in G$ can then be written $g = (g^1, \dots, g^n), g^i \in \mathbb{Z}_{p_i}$, where the superscripts are merely indices, not powers. It is not hard to show that any irredundant generating n-set of G can be written in the form $g_1, \dots, g_n$, where $g_i = (g_i^1, g_i^2, \dots, g_i^n)$ is such that $g_i^j \neq 0 \iff i = j$. The number of irredundant generating n-sets of G is then the number of ways to choose one non-zero element from each $\mathbb{Z}_{p_i}$, which turns out to be $\phi(p_1)\phi(p_2)\dots\phi(p_n) = \phi(m)$.

Things become more complicated for $1 < k < n$. This paper grew out of an attempt to find $\phi_k(G)$ for such groups. It turns out that the lattice of subgroups of $G \cong \mathbb{Z}_{p_1} \times \cdots \times \mathbb{Z}_{p_n}, p_i$ distinct primes, is isomorphic to B_n, the lattice of subsets of $[n]$, and the problem of finding $\phi_k(G)$ is closely related to finding the number of minimal k-covers of $[n]$.

Definition 1.0.3 Given $[n]$ and a finite subset $S \subseteq [n]$, S is a **minimal k-cover** of $[n]$ if it satisfies the following properties:

1. $|S| = k$
2. $\bigcup_{s \in S} s = [n]$
3. $\forall r \in S, \bigcup_{s \in S \setminus r} s \subsetneq G$

The definitions for minimal covers and irredundant generating sets are very similar. The notion of "minimal" corresponds to "irredundant", while "cover" corresponds to "generating". Let $\phi_k(n)$ denote the number of minimal k-covers of $[n]$. Formulas for $\phi_k(n)$ are presented in [3] and [5], although in different forms. We present here the formula from [3], since it is instructive for our purpose:

Theorem 1.0.4

$$\phi_k(n) = \sum_{i=k}^{n} \binom{n}{i} \left\{ {i \atop k} \right\} (2^k - k - 1)^{n-i} \tag{1.1}$$

where $\left\{ {n \atop k} \right\}$ *are the Stirling numbers of the second kind.*

A proof of the theorem, as well as some observations about its relation to $\phi_k(G)$ where G is cyclic of square-free order, will be given as a corollary at the end of the thesis. Here, it suffices to note that this theorem results from the following lemma:

Lemma 1.0.5

$$\phi_k(n) = \sum_{s \subseteq [n]} \left\{ {|s| \atop k} \right\} (2^k - k - 1)^{n-|s|} \tag{1.2}$$

Roughly speaking, the Stirling numbers account for the "irredundant" property, while $2^k - k - 1$ accounts for "generating". This thesis formalizes and generalizes these two quantities and presents an analogous formula for $\phi_k(G)$ when G is finite and is a direct product of elementary abelian groups. This result is then easily extended to finite nilpotent groups, because the quotient of a nilpotent group by its Frattini subgroup is a direct product of elementary abelian groups.

The outline of this thesis is as follows:

In Chapter 2, we present some basic definitions and properties of posets and lattices. We also introduce the Möbius function as well as the technique of Möbius inversion.

Chapter 3 focuses on partition lattices and formulas associated with them. We will introduce a generalization of partitions for disjoint unions of sets, which we call the product partition.

Chapter 4 is a study of finite vector spaces and their lattice of subspaces. We survey a few q-analogues of combinatorial identities on sets that can be applied to vector spaces over $\mathbb{F}_q$, where q is a prime. We also study the nullspaces of linear transformations over these vector spaces, and show that they can tell us "how dependent" a set of vectors is.

Chapter 5 starts by showing that a finite nilpotent group can be quotiented by its Frattini subgroup to get a direct product of elementary abelian groups. We then combine the results from Chapters 3 and 4 to obtain a formula for counting irredundant generating sets of this quotient. This is presented in the main theorem, Theorem 5.3.5.

Chapter 6 generalizes the proof of the main theorem to give a formula for counting irredundant generating sets in direct products of lattices. We will recover Theorem 1.0.4 as a corollary.

SAGE implementations and computational verification of some of the functions introduced in this paper can be found at

`http://sagenb.org/home/pub/4673/`

Chapter 2

Posets, Lattices and the Möbius Function

In this chapter, we summarize some standard combinatorial notions regarding posets, lattices, their direct products, the Möbius function on a poset and Möbius inversion. A comprehensive treatment of these topics can be found in standard combinatorics texts such as [8].

2.1 Basic Definitions

Definition 2.1.1 A **poset**, P, is a set with a binary relation (usually denoted $\leq$) that is

1. Reflexive: $x \leq x$
2. Anti-symmetric: $x \leq y$ and $y \leq x \implies x = y$
3. Transitive: $x \leq y$ and $y \leq z \implies x \leq z$

for all $x, y, z \in P$.

The least and greatest elements of P, if they exist, are denoted $\hat{0}_P$ and $\hat{1}_P$, respectively. In this thesis, all posets will be finite, and will contain $\hat{0}$ and $\hat{1}$ such that $\hat{0} \neq \hat{1}$.

Definition 2.1.2 Given $(P, \leq_P), (Q, \leq_Q)$ posets, a **poset homomorphism** from P to Q is a function $f : P \to Q$ that preserves order. i.e. for all $x, y \in P$,

$$x \leq_P y \implies f(x) \leq_Q f(y)$$

A **poset isomorphism** is a bijection $f : P \to Q$ such that f^{-1} is also order preserving. i.e. for all $x, y \in P$,

$$x \leq_P y \iff f(x) \leq_Q f(y)$$

Definition 2.1.3 Given $(P, \leq)$ a poset, the **dual poset** of P is the poset $(P, \leq_*)$ with order relation $\leq_*$ defined by

$$x \leq_* y \iff y \leq x$$

A poset is **self-dual** if it is isomorphic to its dual.

Definition 2.1.4 Given a poset P, a **subposet** of P is a subset $S \subseteq P$ with the order relation $\leq$ induced by that of P:

$$x \leq y \in S \iff x \leq y \in P$$

Definition 2.1.5 For $x, y \in P$, we may also define the **interval** $[x, y] := \{z \in P | x \leq z \leq y\}$, which is a subposet of P. In this interval, $\hat{0}_{[x,y]} = x$ and $\hat{1}_{[x,y]} = y$.

Definition 2.1.6 Given a poset, $(P, \leq)$, and elements $x < y \in P$, then y **covers** x if there does not exist $z \in P$ such that $x < z < y$. We denote this relation by $x \lessdot y$.

Definition 2.1.7 An element $x \in P$ containing $\hat{0}_P$ is an **atom** if it covers $\hat{0}_P$. Dually, an element $x \in P$ containing $\hat{1}_P$ is a **coatom** if it is covered by $\hat{1}_P$. We denote by $A(P)$ and $CoA(P)$ the set of atoms and coatoms of P, respectively.

Definition 2.1.8 A **chain** or **linear order** is a poset in which any two elements are comparable i.e. a totally ordered set. A subset $C \subseteq P$ is called a chain if C is a chain when regarded as a subposet of P. The length of a chain is $l(C) := |C| - 1$. A chain in P is **maximal** if it is not contained in any other chain in P.

The poset $[n] = \{1, 2, \ldots, n\}$ equipped with the standard ordering of integers is a chain of length $n - 1$. We may thus also say that C is a chain if C is isomorphic to $[n]$ where $n = |C| - 1$.

Definition 2.1.9 P is a **graded poset** if there exists a rank function, $\rho : P \to \mathbb{N}$ such that

- $x < y \implies \rho(x) < \rho(y)$
- $x \lessdot y \implies \rho(y) = \rho(x) + 1$

If P contains $\hat{0}_P$ and is graded, it is convenient to set

$$\rho(\hat{0}_P) = 0.$$

If P also contains $\hat{1}_P$, we may define the rank of P to be

$$\rho(P) := \rho(\hat{1}_P).$$

If P is such that all maximal chains have the same length, then it turns out that $\forall x \in P$, all maximal chains in the interval $[\hat{0}, x]$ have the same length. If we let $\rho(x) := l(C)$, where C is any maximal chain from $\hat{0}$ to x, then it is easy to see that ρ is a rank function on P. In fact, if P is finite and contains both $\hat{0}$ and $\hat{1}$, then P is graded iff all maximal chains of P have the same length [8].

Definition 2.1.10 A poset $(L, \leq)$ is a **lattice** if $\forall a, b \in L$

- $\exists a \vee b \in L$ such that $a \vee b \geq a, b$
- $\exists a \wedge b \in L$ such that $a \wedge b \leq a, b$

$a \vee b$ and $a \wedge b$ are called the **meet** (or least upper bound) and **join** (or greatest lower bound), respectively, of a and b.

If L is finite, then L contains $\hat{1}_L := \bigvee_{x \in L} x$ and $\hat{0}_L := \bigwedge_{x \in L} x$, respectively.

Definition 2.1.11 Let $(L_1, \wedge_1, \vee_1)$ and $(L_2, \wedge_2, \vee_2)$ be lattices, and let $f : L_1 \to L_2$ be a poset homomorphism. Then f is also a **lattice-homomorphism** if it preserves meets and joins. i.e.

$$f(x \wedge_1 y) = f(x) \wedge_2 f(y)$$
$$f(x \vee_1 y) = f(x) \vee_2 f(y)$$

One may check that a bijective lattice homomorphism is automatically a lattice isomorphism (note that this is not true if we replace "lattice" with "poset"). For L a lattice, we may also define the dual lattice, L^*, to be the lattice with $\wedge$ and $\vee$ reversed.

Definition 2.1.12 A **semi-modular lattice** is a finite graded lattice such that

$$\forall x, y \in L, \rho(x) + \rho(y) \geq \rho(x \vee y) + \rho(x \wedge y)$$

Definition 2.1.13 L is **atomistic** if every element is the join of a set of atoms. Dually, L is **coatomistic** if every element is the meet of a set of coatoms.

Definition 2.1.14 Γ is a **geometric lattice** if it is finite, atomistic and semi-modular.

The lattice of subsets, B_n, of a set $[n]$, the lattice of subspaces, $B_n(q)$ of $\mathbb{F}_q^n$, and the lattice of partitions, Π_n of $[n]$ all turn out to be geometric lattices [8].

2.2 Direct Products

Given posets, $P_1, P_2, \ldots, P_n$, we can also form their direct product in the following manner:

Definition 2.2.1 The **direct product**, $P = P_1, P_2, \ldots, P_n$ is the set $P_1 \times P_2 \times \cdots \times P_n$ equipped with the order relation:

$$(x_1, x_2, \ldots, x_n) \leq (y_1, y_2, \ldots, y_n) \iff x_i \leq y_i \text{ for all } 1 \leq i \leq n$$

If $P_1, \ldots P_n$ are posets with respective least and greatest elements $\hat{0}_i \neq \hat{1}_i$, then

$$\hat{0}_P := (\hat{0}_1, \ldots, \hat{0}_n)$$

$$\hat{1}_P := (\hat{1}_1, \ldots, \hat{1}_n)$$

are the least and greatest elements of $P = P_1 \times \cdots \times P_n$

Each P_i sits inside $P = P_1 \times \cdots \times P_n$ as the subposet $\{x_1\} \times \cdots \times P_i \times \cdots \times \{x_n\}$, where x_j is any element of P_j. If each P_i contains $\hat{0}_i \neq \hat{1}_i$, we may define the following inclusions:

$$\iota_i : P_i \to P, x \mapsto (\hat{0}_1, \ldots, x, \ldots, \hat{0}_n) \tag{2.1}$$

$$I_i : P_i \to P, x \mapsto (\hat{1}_1, \ldots, x, \ldots, \hat{1}_n) \tag{2.2}$$

where the x is in the i^{th} position.

It is easy to verify that for $(x_1, \ldots, x_n), (y_1, \ldots, y_n) \in P$, we have

$$(x_1, \ldots, x_n) \lessdot (y_1, \ldots, y_n) \iff \exists i \text{ such that } x_i \lessdot y_i \text{ and } x_j = y_j, \forall j \neq i$$

Under the above definition of covering in P, we have:

$$A(P) = \bigcup_{i=1}^{n} \iota_i(A(P_i))$$

$$CoA(P) = \bigcup_{i=1}^{n} I_i(CoA(P_i))$$

Further, if $P_1, \ldots P_n$ are graded, with respective rank functions ρ_i such that $\rho_i(\hat{0}_i) = 0$ for all i, then P is graded, with rank function:

$$\rho(x_1 \ldots, x_n) = \sum_{i=1}^{n} \rho_i(x_i)$$

2.3 The Möbius Function and Möbius Inversion

The Möbius function on a poset P with $\hat{0}$ and $\hat{1}$ is a map

$$\mu : P \times P \longrightarrow \mathbb{Z} \tag{2.3}$$

In fact, the range of the function can be any abelian group, but for our purposes, we set it to be $\mathbb{Z}$. The function is defined as the function that satisfies, for all $x, y \in P$:

$$\mu(x, x) = 1 \tag{2.4}$$

$$\sum_{x \leq z \leq y} \mu(x, z) = 0 \tag{2.5}$$

$$\sum_{x \leq z \leq y} \mu(z, y) = 0 \tag{2.6}$$

$$\tag{2.7}$$

This allows us to compute $\mu(x, y)$ recursively via either of the two formulas

$$\mu(x, y) = -\sum_{x \leq z < y} \mu(x, z), x < y \tag{2.8}$$

$$\mu(x, y) = -\sum_{x < z \leq y} \mu(z, y), x < y \tag{2.9}$$

The Möbius function is particularly useful when we let μ take values in $\mathbb{Z}$ because of the Möbius inversion formula, which can be found in most combinatorics texts, including [8].

Proposition 2.3.1 (Möbius Inversion) *Let P be a finite poset with $\hat{0}, \hat{1}$. Let $f, g : P \to \mathbb{F}$, where $\mathbb{F}$ is a field. Then*

$$g(x) = \sum_{y \leq x} f(y), \text{ for all } x \in P \tag{2.10}$$

if and only if

$$f(x) = \sum_{y \leq x} \mu(y, x) g(y), \text{ for all } x \in P \tag{2.11}$$

Just as the Möbius function can be calculated recursively in two (equivalent) ways, the Möbius inversion formula has an alternative form:

Proposition 2.3.2 *Let P be a finite poset with $\hat{0}, \hat{1}$. Let $f, g : P \to \mathbb{F}$, where $\mathbb{F}$ is a field. Then*

$$g(x) = \sum_{y \geq x} f(y), \text{ for all } x \in P \tag{2.12}$$

if and only if

$$f(x) = \sum_{y \geq x} \mu(x, y) g(y), \text{ for all } x \in P \tag{2.13}$$

All we have changed is whether x or y is greater. Which formulation we use depends on what we are trying to compute. Often, f is the function we want. The formulas tell us that if we know either $\sum_{y \leq x} f(y)$ or $\sum_{y \geq x} f(y)$, then we have a formula for f.

It turns out that the Möbius function behaves nicely on direct products of posets.

Proposition 2.3.3 *Let $P_1, \ldots, P_n$ be posets with $\hat{0}_i, \hat{1}_i$, and Möbius functions μ_i. Let $P = P_1 \times \cdots \times P_n$. Then for $\mathbf{x} = (x_1, \ldots, x_n), \mathbf{y} = (y_1, \ldots, y_n), x_i, y_i \in P_i$,*

$$\mu(\mathbf{x}, \mathbf{y}) = \prod_{i=1}^{n} \mu(x_i, y_i) \tag{2.14}$$

A proof can be found in most standard combinatorics texts or texts on lattices, including [8] and [2].

2.4 Linear Orders

We end of this chapter by giving an example of a direct product of posets, and its Möbius function. This poset, however, is not a geometric lattice. First, we consider the linear order, $[n]$.

Proposition 2.4.1 *Let $L = [n]$, then for $x, y \in L$,*

$$\mu(x,y) = \begin{cases} 1 & \text{if } x = y \\ -1 & \text{if } y = x+1 \\ 0 & \text{otherwise} \end{cases} \tag{2.15}$$

We can form the direct product of n linear orders to obtain a set of n-tuples ordered by

$$(x_1, \ldots, x_n) \leq (y_1, \ldots, y_n) \iff x_i \leq y_i, \forall i$$

The Möbius function on this product is then:

$$\mu\big((x_1, \ldots, x_n), (y_1, \ldots, y_n)\big) = \begin{cases} (-1)^{\sum_{i=1}^{m}(y_i - x_i)} & \text{if } 0 \leq y_i - x_i \leq 1 \text{ for all } i \\ 0 & \text{otherwise} \end{cases}$$

Suppose we have functions f and g defined on $[m_1] \times \cdots \times [m_n]$ such that:

$$f(x_1, \ldots, x_n) = \sum_{i_1=1}^{x_1} \sum_{i_2=1}^{x_2} \cdots \sum_{i_n=1}^{x_n} g(i_1, \ldots, i_n)$$

Then this is equivalent to:

$$f(\mathbf{x}) = \sum_{\mathbf{1} \leq \mathbf{y} \leq \mathbf{x}} g(\mathbf{y})$$

where boldface indicates tuples, and $\mathbf{1} = (1, \ldots, 1)$. We can then apply Möbius inversion on $[m_1] \times \cdots \times [m_n]$ to obtain:

$$g(\mathbf{x}) = \sum_{\mathbf{1} \leq \mathbf{y} \leq \mathbf{x}} \mu(\mathbf{y}, \mathbf{x}) f(\mathbf{y})$$

Chapter 3

Partitions

In this chapter, we study the lattice of partitions of $[n]$, and then proceed to define a special kind of partition on a disjoint union of sets.

3.1 Partition Lattices

Definition 3.1.1 A **partition** of $[n] := \{1, 2, \ldots, n\}$ is a collection, $\pi := \{B_1, B_2, \ldots, B_k\}$ of non-empty subsets $\varnothing \neq B_i \subseteq [n]$ satisfying:

1. $\bigcup_{i=1}^{n} B_i = [n]$
2. $\forall i \neq j, B_i \cap B_j = \varnothing$

The subsets, B_i are called **blocks**, and the partition π is a k**-partition** if it contains k blocks. We write $|\pi| = k$.

Definition 3.1.2 The lattice of partitions of an n-set, denoted Π_n, is the set of partitions of $[n]$ with the ordering $\leq$ defined such that for partitions $\pi = \{B_1, B_2, \ldots, B_k\}$ and $\theta = \{C_1, \ldots, C_l\}$,

$$\pi \leq \theta \iff \forall i \in [k], \exists j \in [l] \text{ such that } B_i \subseteq C_j$$

This is called the ordering by refinement. Under this ordering, the least element $\hat{0}_{\Pi_n}$ is the partition with n blocks, all of which are singletons, and the greatest element $\hat{1}_{\Pi_n}$ is the partition consisting of the single block $[n]$. The rank function on this lattice is:

$$\rho(\pi) := n - |\pi| \tag{3.1}$$

Thus $\rho(\hat{0}_{\Pi_n}) = 0$ while $\rho(\hat{1}_{\Pi_n}) = n - 1$.

Theorem 3.1.3 *Let* $\left\{ {n \atop k} \right\}$ *denote the number of* k*-partitions of an* n*-set. Then*

$$\left\{ {n \atop k} \right\} = \frac{1}{k!} \sum_{j=0}^{k} (-1)^{k-j} \binom{k}{j} j^n.$$

The numbers $\left\{ {n \atop k} \right\}$ are the Stirling numbers of the second kind. They count the number of elements in Π_n of rank $n - k$. A proof of the above formula can be found in most texts for enumerative combinatorics, including [8].

Next we determine the Möbius function on Π_n for some fixed n. First we let $\mu_n := \mu(\hat{0}_{\Pi_n}, \hat{1}_{\Pi_n})$. Then from [8], we have that:

$$\mu_n = (-1)^{n-1}(n-1)! \tag{3.2}$$

For partitions $\theta \leq \pi$, if $\pi = \{B_1, \ldots, B_k\}$ is such that each B_i is partitioned into n_i blocks in θ, then it is clear that:

$$[\theta, \pi] \cong \Pi_{n_1} \times \cdots \times \Pi_{n_k} \tag{3.3}$$

So $\mu(\theta, \pi) = \mu_{n_1}\mu_{n_2}\ldots\mu_{n_k}$, and hence we see that determining μ_n for all $n \in \mathbb{N}$ is sufficient to determine μ on any partition lattice.

3.2 The Product Partition

In the foregoing discussion, we see that direct products of partition lattices occur naturally as sublattices of a single partition lattice. We introduce here a slightly different concept: the product partition of a disjoint union of sets.

We start with an n-tuple, $\mathbf{m} = (m_1, m_2, \ldots, m_n)$, and define sets $X^i = \{x_1^i, x_2^i, \ldots x_{m_i}^i\}, i \in [n]$ such that $|X^i| = m_i$. The m_i's are allowed to be zero, in which case $X^i = \varnothing$. We define also the set:

$$X^{\mathbf{m}} = \bigsqcup_{i=1}^{n} X^i$$

which is simply the (disjoint) union of the X^i's. Then $|X^{\mathbf{m}}| = |\mathbf{m}| := \sum_{i=1}^n m_i$. We want to consider k-partitions of $X^{\mathbf{m}}$ with the additional property that each block of the partition contains at most one element from each X^i.

First we note that any ordinary partition, π, of $X^{\mathbf{m}}$ (with no restrictions) induces a partition π_i on X^i via

$$a, b \in X^i \text{ are in the same block in } \pi \implies a, b \in X^i \text{ are in the same block of } \pi_i.$$

Alternatively, for each $i \in [n]$, if $B_1, \ldots, B_k$ are the blocks of π, define $B_j^i := B_j \cap X^i, j \in [k]$. Then π_i is the partition whose blocks are the non-empty B_j^i's.

Definition 3.2.1 Let Π_{m_i} denote the lattice of partitions of X^i for each $i \in [n]$, and $\Pi_{|\mathbf{m}|}$ be the lattice of all partitions of $X^{\mathbf{m}}$. Let $\Pi_{\mathbf{m}}$ denote the direct product of partition lattices:

$$\Pi_{\mathbf{m}} := \Pi_{m_1} \times \Pi_{m_2} \times \cdots \times \Pi_{m_n}$$

Elements of Π_{m_i} and $\Pi_{|\mathbf{m}|}$ are ordinary partitions. Elements $\boldsymbol{\pi} \in \Pi_{\mathbf{m}}, \boldsymbol{\pi} = (\pi_1, \ldots, \pi_n)$ are tuples of partitions, and will be distinguished from ordinary partitions by being typeset in boldface. We may also define the size of $\boldsymbol{\pi}$, $|\boldsymbol{\pi}| := \sum_{i=1}^n |\pi_i|$.

Proposition 3.2.2 *The map:*

$$\begin{aligned} \Psi : \Pi_{|\mathbf{m}|} &\longrightarrow \Pi_{\mathbf{m}} \\ \pi &\longmapsto (\pi_1, \pi_2, \ldots, \pi_n) \end{aligned}$$

where π_i is the partition induced by π on each X^i, is a lattice homomorphism.

Proof. For $\pi, \tau \in \Pi_{|\mathbf{m}|}$, let π_i, τ_i denote the partitions induced by π, τ, resp., on X^i. It suffices to show that $\forall i, \pi \leq \tau \implies \pi_i \leq \tau_i$. If $\pi \leq \tau$, each block of π is contained in some block of τ. Let $B^i \subseteq X^i$ be a block of π_i, then $B^i = B \cap X^i$ for some block B of π. $B \subseteq C$ for some block C of τ, so $B \cap X^i \subseteq C \cap X^i$. But $C \cap X^i$ is a block of τ_i, hence B^i is a subset of a block of τ_i, so $\pi_i \leq \tau_i$. $\square$

We are especially interested in partitions such that each block contains at most 1 element from each X^i. This means that for each X^i, all its elements are in different blocks, hence the partition induced on X^i is simply the least element of Π_{m_i}, which we denote by $\hat{0}_i$. Let $\hat{\mathbf{0}} := (\hat{0}_1, \ldots, \hat{0}_n) \in \Pi_{\mathbf{m}}$. We give such partitions of $X^{\mathbf{m}}$ a special name:

Definition 3.2.3 A partition $\pi \in \Pi_{\mathbf{m}}$ is a **product k-partition** if $|\pi| = k$ and $\Psi(\pi) = \hat{\mathbf{0}}$.

We are interested in the number of such partitions. Since we are only interested in k-partitions, let $\Pi^k_{\mathbf{m}}$ denote the k-partitions in $\Pi_{\mathbf{m}}$.

Definition 3.2.4

$$\left\{ {\mathbf{m} \atop k} \right\} := \left| \{ \pi \in \Pi^k_{\mathbf{m}} \big| \Psi(\pi) = \hat{\mathbf{0}} \} \right|$$

In order to get a formula for $\left\{ {\mathbf{m} \atop k} \right\}$, we define a few auxiliary functions on the lattice $\Pi_{\mathbf{m}}$. For $\boldsymbol{\pi} \in \Pi_{\mathbf{m}}, \boldsymbol{\pi} = (\pi_1, \ldots, \pi_n)$ let

$$\left\{ {\boldsymbol{\pi} \atop k} \right\}_{\geq} := \left| \left\{ \pi \in \Pi^k_{\mathbf{m}} \big| \Psi(\pi) \geq \boldsymbol{\pi} \right\} \right|$$

$$\left\{ {\boldsymbol{\pi} \atop k} \right\}_{=} := \left| \left\{ \pi \in \Pi^k_{\mathbf{m}} \big| \Psi(\pi) = \boldsymbol{\pi} \right\} \right|$$

Then it is clear that

$$\left\{ {\boldsymbol{\pi} \atop k} \right\}_{\geq} = \sum_{\boldsymbol{\tau} \geq \boldsymbol{\pi}} \left\{ {\boldsymbol{\tau} \atop k} \right\}_{=} \tag{3.4}$$

so by Möbius inversion,

$$\left\{ {\boldsymbol{\pi} \atop k} \right\}_{=} = \sum_{\boldsymbol{\tau} \geq \boldsymbol{\pi}} \mu(\boldsymbol{\pi}, \boldsymbol{\tau}) \left\{ {\boldsymbol{\tau} \atop k} \right\}_{\geq} \tag{3.5}$$

Proposition 3.2.5

$$\left\{ {\boldsymbol{\pi} \atop k} \right\}_{\geq} = \left\{ {|\boldsymbol{\pi}| \atop k} \right\}$$

Proof. The only requirement for $\pi \in \Pi^k_{\mathbf{m}}$ to be such that $\Psi(\pi) \geq \boldsymbol{\pi} = (\pi_1, \ldots, \pi_n)$ is that each block of each π_i in $\boldsymbol{\pi}$ is contained in a block of π. We may thus treat each block of each π_i as a single element. π is then a partition on the set of blocks, and since there are $|\boldsymbol{\pi}| = \sum_{i=1}^n |\pi_i|$ such blocks, the number of such partitions is $\left\{ {|\boldsymbol{\pi}| \atop k} \right\}$. □

Corollary 3.2.6

$$\left\{ {\mathbf{m} \atop k} \right\} = \sum_{\hat{\mathbf{0}} \leq \boldsymbol{\pi}} \mu(\hat{\mathbf{0}}, \boldsymbol{\pi}) \left\{ {|\boldsymbol{\pi}| \atop k} \right\} \tag{3.6}$$

Proof. This follows by noting that

$$\left\{ {\mathbf{m} \atop k} \right\} = \left\{ {\hat{\mathbf{0}} \atop k} \right\}_{=} \tag{3.7}$$

□

Remark 3.2.7 When $\mathbf{m} = \mathbf{1} = (1, 1, \ldots, 1)$, then each $\Pi_{m_i} = \Pi_1$, which consists of a single element $\hat{0}$. So the product of Π_{m_i}'s is also the single element, $\hat{\mathbf{0}} = (\hat{0}, \ldots, \hat{0})$. $\mu(\hat{\mathbf{0}}, \hat{\mathbf{0}}) = 1$, and $|(\hat{0}, \ldots, \hat{0})| = n$, giving

$$\left\{ {\mathbf{1} \atop k} \right\} = \left\{ {n \atop k} \right\} \tag{3.8}$$

Chapter 4

Subsets and Subspaces of $\mathbb{F}_q^n$

In this chapter, we study the lattice of subspaces, $B_n(q)$, of the finite vector space $\mathbb{F}_q^n$. We also attempt to enumerate subsets of $\mathbb{F}_q^n$ with certain properties. First, however, we introduce the lattice of subsets, B_n, of the set $[n]$, which will turn out to have close connections to $B_n(q)$.

Definition 4.0.8 The **lattice of subsets** of $[n]$, denoted B_n, is the set of subsets of $[n]$ with the following properties:

- $x \leq y \iff x \subseteq y$
- $x \vee y = x \cup y$
- $x \wedge y = x \cap y$
- $\rho(x) = |x|$

For $x, y \in B_n$ such that $x \leq y$, there is a bijective correspondence between subsets of y that contain x and subsets of $y - x$. It is not hard to see that this bijection is indeed a lattice isomorphism, so the interval $[x, y] \cong B_k$, where

$$k = |y - x| = |y| - |x| = \rho(y) - \rho(x). \tag{4.1}$$

Because $[x, y] \cong B_k$, we need only determine the Möbius function on B_n for all $n \in \mathbb{N}$. Let $\mu_n := \mu_{B_n}(\hat{0}, \hat{1})$, then in [8], we find that

Proposition 4.0.9

$$\mu_n = (-1)^n$$

The observation that $[x, y] \cong B_k$, where $k = \rho(y) - \rho(x)$ leads to the easy corollary:

Corollary 4.0.10 *Let $x, y \in B_n$ be such that $x \leq y$, then*

$$\mu_{B_n}(x, y) = (-1)^{\rho(y)-\rho(x)}$$

Remark 4.0.11 If f, g are functions from B_n to some field, $\mathbb{F}$, and are such that

$$f(x) = \sum_{y \subseteq x} g(y) \tag{4.2}$$

then Möbius inversion on B_n gives:

$$g(x) = \sum_{y \subseteq x} (-1)^{|y|-|x|} f(y) \tag{4.3}$$

which is the principle of inclusion-exclusion. Thus Möbius inversion can be thought of as a generalization of the principle of inclusion-exclusion.

We turn our attention now to $B_n(q)$.

Definition 4.0.12 Let q be a prime. The **lattice of subspaces** of the vector space $\mathbb{F}_q^n$, denoted $B_n(q)$, is the set of subspaces of $\mathbb{F}_q^n$ with the following properties:

- $x \leq y \iff x$ is a subspace of y
- $x \vee y = \text{Span}(x, y)$
- $x \wedge y = x \cap y$
- $\rho(x) = \dim_{\mathbb{F}_q}(x)$

Since subspaces and quotients of $\mathbb{F}_q^n$ are isomorphic to $\mathbb{F}_q^k$ for some $k \leq n$, any interval of $B_n(q)$ is isomorphic to $B_k(q)$ for some $k \leq n$. In particular, if $x \leq y$, then the interval $[x, y]$ is isomorphic to $B_k(q)$ where

$$k = \dim_{\mathbb{F}_q}(y/x) = \dim_{\mathbb{F}_q}(y) - \dim_{\mathbb{F}_q}(x) = \rho(y) - \rho(x). \tag{4.4}$$

As was the case for B_n, the fact that $[x, y] \cong B_k(q)$ for some k means we need only determine $\mu_{n,q} := \mu_{B_n(q)}(\hat{0}, \hat{1})$ for all $n \in \mathbb{N}$. We have from [8]:

Proposition 4.0.13

$$\mu_{n,q} = (-1)^n q^{\binom{n}{2}} \tag{4.5}$$

This leads to the easy corollary:

Corollary 4.0.14

$$\mu_{B_n(q)}(x, y) = (-1)^{\rho(y)-\rho(x)} q^{\binom{\rho(y)-\rho(x)}{2}} \tag{4.6}$$

4.1 q-Analogues

It turns out that many properties of $B_n(q)$ are analogous to properties of B_n. Indeed, if we ignore what it means, letting $q \to 1$ in many formulas for $B_n(q)$ give us formulas for B_n. Such formulas are called q-analogues. We present some examples that are relevant to our problem. Where possible, we will mention what a formula counts in the limit $q \to 1$. An interesting treatment of some applications of q-analogues can be found in [4].

Definition 4.1.1 The **q-integer** of n is

$$(n)_q := \frac{q^n - 1}{q - 1} = 1 + q + \cdots + q^{n-1}$$

Proposition 4.1.2 *The number of elements of rank* 1 *in* $B_n(q)$ *is* $(n)_q$.

Proof. The rank 1 elements of $B_n(q)$ are the 1-dimensional subspaces of $\mathbb{F}_q^n$. There are $q^n - 1$ non-zero elements in $\mathbb{F}_q^n$, and each 1-dimensional subspace is generated by $q - 1$ of these non-zero elements, so there are $\frac{q^n-1}{q-1} = (n)_q$ such subspaces. □

Remark 4.1.3 When $q \to 1$, the elements of rank 1 in B_n are simply the singletons, of which there are $n = (n)_1$.

Definition 4.1.4 The **q-factorial** of n is

$$(n)_q! := \begin{cases} 1 & \text{if } n = 0 \\ (n)_q \times (n-1)_q \times \cdots \times (1)_q & \text{otherwise} \end{cases}$$

Unfortunately, the q-factorial does not count any nice quantity of $\mathbb{F}_q^n$ or $B_n(q)$. The normal factorial counts the number of permutations of $[n]$. This may be interpreted as the number of automorphisms of $[n]$ i.e. the number of bijections from $[n]$ to itself. We might thus expect $(n)_q!$ to count the number of automorphisms of $\mathbb{F}_q^n$. Automorphisms of $\mathbb{F}_q^n$ are invertible linear transformations from $\mathbb{F}_q^n$ to itself. These are precisely the elements of $GL(n,q)$, whose cardinality is a multiple of $(n)_q!$.

Proposition 4.1.5

$$|GL(n,q)| = (n)_q! \cdot \phi(q)^n \cdot q^{\binom{n}{2}}$$

Proof. The size of $GL(n,q)$ is the number of $n \times n$ matrices over $\mathbb{F}_q$ with linearly independent columns. The first column can be any non-zero vector in $\mathbb{F}_q^n$, of which there are $q^n - 1$. The second column can be any vector in $\mathbb{F}_q^n$ that is not in the span of the vector in the first column. The first column spans a 1-dimensional subspace with q elements, so there are $q^n - q$ choices for the second column. Subsequently, the number of vectors that can be placed in the $(i+1)^{\text{th}}$ column is the number of vectors of $\mathbb{F}_q^n$ that are not in the span of the first i columns. These i columns span an i-dimensional subspace with q^i elements, so the number of choices for the $(i+1)^{\text{th}}$ column is $q^n - q^i = q^i(q^{n-i} - 1)$, and so on. The number of matrices with independent columns is thus:

$$(q^n - 1) \times (q^n - q) \times \cdots \times (q^n - q^{n-1}) = (q^n - 1) \times (q^{n-1} - 1) \cdot q \times \cdots \times (q-1) \cdot q^{n-1} \tag{4.7}$$

Recalling that $(k)_q = \frac{q^k - 1}{q-1}$, the above becomes:

$$(n)_q \cdot \phi(q) \times (n-1)_q \cdot \phi(q) \cdot q \times \cdots \times (1)_q \cdot \phi(q) \cdot q^{n-1} = (n)_q! \cdot \phi(q)^n \cdot q^{1+2+\cdots+(n-1)} \tag{4.8}$$

Noting that $1 + 2 + \cdots + (n-1) = \binom{n}{2}$ yields the result. □

Corollary 4.1.6 *The number of ordered bases of $\mathbb{F}_q^n$ is $(n)_q! \cdot \phi(q)^n \cdot q^{\binom{n}{2}}$.*

Proof. This follows from noting that the columns of an invertible $n \times n$ matrix over $\mathbb{F}_q$ form an ordered basis of $\mathbb{F}_q^n$, and vice versa. □

A basis of $\mathbb{F}_q^n$ is a maximally independent subset of $\mathbb{F}_q^n$. In order to see what this might be analogous to in the case $q \to 1$, we need to define a notion of "independece" for $[n]$. We define a more general notion of "independence" on $B_n(q)$:

Definition 4.1.7 A k-tuple $(x_1, \ldots, x_k), x_i \in B_n(q)$ is **independent** if

$$\forall i, \left(\bigvee_{j \neq i} x_j \right) \wedge x_i = \hat{0} \tag{4.9}$$

Otherwise, it is called **dependent**.

Remark 4.1.8 For general k, the definition above allows the x_i's to be subspaces of any dimension, but when $k = n$, each x_i must be a 1-dimensional subspace.

Definition 4.1.9 A k-tuple $(\mathbf{x}_1, \ldots, \mathbf{x}_k), \mathbf{x}_i \in \mathbb{F}_q^n$ is independent if $\big(\langle \mathbf{x}_1 \rangle, \ldots, \langle \mathbf{x}_k \rangle\big)$ is independent. Otherwise it is dependent.

It is easy to see that for elements of $\mathbb{F}_q^n$, the notions of linear independence and independence as defined above coincide. We have seen that the proof of Proposition 4.1.5 gives:

Proposition 4.1.10 *The number of independent n-tuples with entries from $\mathbb{F}_q^n$ is*

$$(n)_q! \cdot \phi(q)^n \cdot q^{\binom{n}{2}}.$$

When $q = 1, k = n$, an independent n-tuple of $\mathbb{F}_1^n$ is simply an n-tuple whose entries (from $[n]$) are all distinct. The proposition says there are $(n)_1! \cdot \phi(1)^n \cdot 1^{\binom{n}{2}} = n!$ such n-tuples, which is what we expect.

A slight modification of the proof of Proposition 4.1.5 yields:

Proposition 4.1.11 *The number of independent n-tuples with entries from $B_n(q)$ is*

$$(n)_q! \cdot q^{\binom{n}{2}}.$$

Proof. As we remarked above, all entries must be 1-dimensional subspaces. The proof is similar to that of Proposition 4.1.5, with columns corresponding to entries in the n-tuple. But for each 1-dimensional subspace, its $\phi(q)$ non-zero elements generate the same subspace, so the number of choices for each entry should be divided by $\phi(q)$. □

Remark 4.1.12 When $q \to 1$, a tuple of elements from $[n]$ or of rank 1 elements from B_n is independent exactly when all entries of the tuple are distinct. The number of independent n-tuples with entries from either $[n]$ or B_n is thus $n!$.

We have seen two analogues for the quantities counted by $n!$, both of which take the form $(n)_q!$ with some extra factors. These extra factors cancel out in the q-analogue of the binomial coefficient:

Definition 4.1.13 The **q-binomial coefficient** is

$$\binom{n}{k}_q = \frac{(n)_q!}{(k)_q!(n-k)_q!}$$

Proposition 4.1.14 *The number of elements of rank k in $B_n(q)$ is $\binom{n}{k}_q$.*

Proof. Elements of rank k in $B_n(q)$ are simply the k-dimensional subspaces of $\mathbb{F}_q^n$. Each k-dimensional subspace of $\mathbb{F}_q^n$ can be specified by choosing a basis for it. Using a similar argument as the proof of Proposition 4.1.5, the number of ways to choose k linearly independent vectors in $\mathbb{F}_q^n$ is:

$$\begin{aligned}(q^n - 1) \times (q^n - q) \times \cdots \times (q^n - q^{k-1}) &= \frac{(q^n - 1) \times (q^n - q) \times \cdots \times (q^n - q^{n-1})}{(q^n - q^k) \times (q^n - q^{k+1}) \times \cdots \times (q^n - q^{n-1})} \\ &= \frac{(n)_q! \cdot \phi(q)^n \cdot q^{1+2+\cdots+(n-1)}}{(n-k)_q! \cdot \phi(q)^{n-k} \cdot q^{k+(k+1)+\cdots+(n-1)}} \qquad (4.10)\end{aligned}$$

But each k-dimensional subspace is isomorphic to $\mathbb{F}_q^k$, and will be specified by any k linearly independent vectors of $\mathbb{F}_q^k$. The number of k linearly independent vectors that give the same subspace is $(k)_q! \cdot \phi(q)^k \cdot q^{1+2+\cdots+(k-1)}$, which is given by setting $k = n$ in Proposition 4.4.4. Dividing by this number gives:

$$\begin{aligned}\frac{1}{(k)_q! \cdot \phi(q)^k \cdot q^{1+2+\cdots+(k-1)}} \cdot \frac{(n)_q! \cdot \phi(q)^n \cdot q^{1+2+\cdots+(n-1)}}{(n-k)_q! \cdot \phi(q)^{n-k} \cdot q^{k+(k+1)+\cdots+(n-1)}} &= \frac{(n)_q!}{(k)_q!(n-k)_q!} \\ &= \binom{n}{k}_q\end{aligned}$$

□

If we want to consider the number of ways to choose k independent elements from $\mathbb{F}_q^n$, we can count the number of ways to choose a k-dimensional subspace of $\mathbb{F}_q^n$, and multiply this by the number of ways to choose an unordered basis from that subspace.

Proposition 4.1.15 *The number of ways to choose k independent elements from $\mathbb{F}_q^n$ is*

$$\binom{n}{k}_q \frac{(k)_q!}{k!} \cdot \phi(q)^k \cdot q^{\binom{k}{2}}.$$

Proof. $\binom{n}{k}_q$ is the number of ways to choose a k-dimensional subspace and $(k)_q! \cdot \phi(q)^k \cdot q^{\binom{k}{2}}$ is the number of ordered bases for that subspace. Dividing by $k!$ gives us the number of unordered bases. □

Remark 4.1.16 When $q \to 1$, the number of ways to choose k independent elements from $[n]$ is the number of ways to choose k distinct elements from $[n]$ i.e. $\binom{n}{k}_1 = \binom{n}{k}$, which agrees with our intuition.

4.2 Irredundancy and Essentiality

We have defined the notion of independence in a lattice, which is a generalization of linear independence in a vector space. Here, we define a slightly different notion:

Definition 4.2.1 A k-tuple $(x_1, \ldots, x_k), x_i \in B_n(q)$ is **irredundant** if

$$\forall i, x_i \not\leq \left(\bigvee_{j \neq i} x_j \right).$$

Otherwise, it is called **redundant**.

Definition 4.2.2 An k-tuple $(\mathbf{x}_1, \ldots, \mathbf{x}_k), \mathbf{x}_i \in \mathbb{F}_q^n$ is irredundant if

$$\forall i, \mathbf{x}_i \notin \langle \mathbf{x}_j, j \neq i \rangle .$$

Otherwise, it is redundant.

Alternatively, $(\mathbf{x}_1, \ldots, \mathbf{x}_k), \mathbf{x}_i \in \mathbb{F}_q^n$ is irredundant if $(\langle \mathbf{x}_1 \rangle, \ldots, \langle \mathbf{x}_k \rangle)$ is irredundant. It is easy to see that an independent tuple is always irredundant. The converse is true for tuples of atoms:

Proposition 4.2.3 *Let $S = (x_1, \ldots, x_k), x_i \in A\big(B_n(q)\big)$ be a k-tuple of atoms. Then*

$$S \text{ is independent} \iff S \text{ is irredundant.}$$

If the x_i's are not atoms, the converse no longer holds, as a tuple of irredundant subspaces with dimensions greater than 1 might intersect in a common subspace of lower rank. Tuples of elements, however, correspond to tuples of atoms, so irredundance and independence for tuples of elements are equivalent. From here on, unless stated otherwise, k-tuples will be understood as having entries in $\mathbb{F}_q^n$, not $B_n(q)$. We will also refer to the linear span of the entries of a k-tuple as simply the span of the k-tuple.

In a k-tuple of elements, it might be the case that for some i but not all, we have $\mathbf{x}_i \notin \langle \mathbf{x}_j, j \neq i \rangle$. If the tuple is such that $\mathbf{x}_i \notin \langle \mathbf{x}_j, j \neq i \rangle$ for some i, the tuple is redundant and dependent. But not all redundant tuples are equally redundant, and neither are the entries of a redundant tuple. We want to single out the entries for which the irredundancy condition holds.

Definition 4.2.4 Given a tuple $(\mathbf{x}_1, \ldots, \mathbf{x}_k), \mathbf{x}_i \in \mathbb{F}_q^n$, an entry $\mathbf{x}_i$ is **essential** if

$$\mathbf{x}_i \notin \langle \mathbf{x}_j, j \neq i \rangle$$

An index $i \in [k]$ for which $\mathbf{x}_i$ is essential is called an **essential index**.

Note that the essentiality of an entry is defined relative to the tuple it belongs to. With this definition, a tuple is irredundant iff all its elements are essential. We may now classify tuples by their essential entries:

Definition 4.2.5 Given a k-tuple $(\mathbf{x}_1, \ldots, \mathbf{x}_k), \mathbf{x}_i \in \mathbb{F}_q^n$, its **essential index set** is the set of essential indices:

$$\epsilon(\mathbf{x}_1, \ldots, \mathbf{x}_k) := \{i | \mathbf{x}_i \text{ is essential}\} \subseteq [k] \tag{4.11}$$

From the previous section, if $k \leq n$, the number of k-tuples whose essential index set is the whole of $[k]$ is

$$\binom{n}{k}_q \frac{(k)_q!}{k!} \cdot \phi(q)^k \cdot q^{\binom{k}{2}}.$$

We want to generalize this by counting the number of k-tuples whose essential index set is some subset $\mathcal{I} \subseteq [k]$. Further, we are only interested in k-tuples that span a subspace of dimension n, so we must have $k \geq n$.

Definition 4.2.6 For $\mathcal{I} \subseteq [k]$, let $F_q^n(k, \mathcal{I})$ denote the number of k-tuples of elements from $\mathbb{F}_q^n$ that span $\mathbb{F}_q^n$, and whose essential index sets are exactly $\mathcal{I}$.

Clearly $F_q^n(k, \mathcal{I})$ depends only on the cardinality of $\mathcal{I}$ and not on $\mathcal{I}$ itself. If $|\mathcal{I}| = i$, then $F_q^n(k, \mathcal{I}) = F_q^n(k, [i])$. We modify our function such that it depends only on i:

Definition 4.2.7 For $i \leq k$, let $F_q^n(k, i)$ denote the number of k-tuples of elements from $\mathbb{F}_q^n$ that span $\mathbb{F}_q^n$, and whose essential index sets are exactly $[i]$.

We shall develop the formulas needed to calculate $F_q^n(k, i)$ in Section 4.4.

4.3 Dual Spaces

Before we proceed, we recall a few definitions and results concerning the dual space, V^* of a vector space V. All vector spaces we consider will be finite dimensional.

Definition 4.3.1 Given a vector space, V, over a field $\mathbb{F}$, a **linear functional** of V is a linear map

$$f : V \to \mathbb{F}$$

Linear functionals are often represented as row vectors, while the vectors they act upon (i.e. the vectors of V) are represented as column vectors. The action of a linear functional f on a vector $\mathbf{v}$ is given by the matrix product $f\mathbf{v}$, where f is a row vector and $\mathbf{v}$ is a column vector. Writing linear functionals as row vectors makes it obvious that the set of linear functionals on V is in fact a vector space.

Definition 4.3.2 Given a vector space, V, the **dual space** of V, denoted V^* is the space of all linear functionals on V.

We distinguish the linear functionals

$$f^i : V \to \mathbb{F}, f^i(\mathbf{e}_j) = \delta_{ij}, \tag{4.12}$$

where $\mathbf{e}_j$ are the standard basis (column) vectors of V with a 1 in the j^{th} place and 0 everywhere else. We may then write the f^i's as (row) vectors with a 1 in the i^{th} place and 0 everywhere else. These f^i's form a basis of V^*.

We are especially interested in subspaces of linear functionals that vanish on subspaces of V.

Definition 4.3.3 Let U be a subspace of V, then the **annihilator** of U is the subspace of V^*

$$U^o := \{f \in V^* | f(\mathbf{v}) = 0 \text{ for all } \mathbf{v} \in U\}$$

This notion applies dually to subspaces of V^*:

Definition 4.3.4 Let U^* be a subspace of V^*, then the annihilator of U^* is the subspace of V

$$U^{*o} := \{\mathbf{v} \in V | f(\mathbf{v}) = 0 \text{ for all } f \in U^*\}$$

If V is finite dimensional, then for $U \subseteq V$ and $U^* \subseteq V^*$,

$$U^{oo} \cong U \tag{4.13}$$

$$U^{*oo} \cong U^* \tag{4.14}$$

and we may identify U^{oo} with U and U^{*oo} with U^*.

If $u^1, \ldots, u^k$ is a basis of $U^* \subseteq V^*$, then any $f \in U^*$ can be expressed as a linear combination of these basis elements, and it is easy to see that for a given $\mathbf{v} \in V$,

$$f(\mathbf{v}) = 0 \text{ for all } f \in U^* \iff u^i(\mathbf{v}) = 0 \text{ for all } i \in [k]. \tag{4.15}$$

This gives us:

$$\begin{aligned} U^{*o} &= \{\mathbf{v} \in V | u^i(\mathbf{v}) = 0 \text{ for all } i \in [k]\} \\ &= \bigcap_{i=1}^{k} \{\mathbf{v} \in V | u^i(\mathbf{v}) = 0\} \\ &= \bigcap_{i=1}^{k} \ker u^i \end{aligned} \tag{4.16}$$

We may let the u^i's be rows of a matrix, T, and treat T as a linear transformation from V to itself. Since the rank of T is equal to the dimension of the row-span of T, we have

$$\text{rank } T = \dim U^* \tag{4.17}$$

Further,

$$\begin{aligned} \ker T &= \{\mathbf{v} \in V | T(\mathbf{v}) = \mathbf{0}\} \\ &= \{\mathbf{v} \in V | u^i(\mathbf{v}) = 0 \text{ for all } i \in [k]\} \\ &= U^{*o} \end{aligned} \tag{4.18}$$

Then by the rank-nullity theorem, rank $T + \dim \ker T = \dim V$, so

$$\dim U^* + \dim U^{*o} = \dim V \tag{4.19}$$

Replacing U^* by U^o for some subspace $U \subseteq V$, and recalling the identification $U^{oo} = U$, we get the dual version of this statement:

$$\dim U^o + \dim U = \dim V \tag{4.20}$$

Let $V = \mathbb{F}_q^n$, and let $V^* = \mathbb{F}_q^{n*}$ be its dual space. Both are n-dimensional vector spaces over $\mathbb{F}_q$, hence are isomorphic. If we define $B_n^*(q)$ to be the lattice of subspaces of V^*, then $B_n(q)$ is isomorphic to $B_n^*(q)$, via

$$^* : U \mapsto U^* \tag{4.21}$$

where U^* is the subspace of $\mathbb{F}_q^{n*}$ whose elements, as row-vectors, are transposes of the elements of U, as column vectors.

However, we also have an isomorphism from $B_n(q)$ to the dual of $B_n^*(q)$, given by the annihilator:

$$^o : U \mapsto U^o \tag{4.22}$$

This is an isomorphism from $B_n(q)$ to the dual of $B_n^*(q)$ in that we have

Proposition 4.3.5 *For all $x, y \in B_n(q)$,*

$$(x \vee y)^o = x^o \wedge y^o \tag{4.23}$$
$$(x \wedge y)^o = x^o \vee y^o \tag{4.24}$$

Proof. Let $\{\mathbf{x}_1, \ldots, \mathbf{x}_k\}, \{\mathbf{y}_1, \ldots, \mathbf{y}_l\}$ be bases of x and y, respectively. Then

$$\begin{aligned} x^o &= \{f \in \mathbb{F}_q^{n*} | f(\mathbf{x}_i) = \mathbf{0} \text{ for all } i \in [k]\} \\ &= \bigcap_{i=1}^{k} \{f \in \mathbb{F}_q^{n*} | f(\mathbf{x}_i) = \mathbf{0}\} \\ &= \bigcap_{i=1}^{k} \langle \mathbf{x}_i \rangle^o \end{aligned} \tag{4.25}$$

Likewise, $y^o = \bigcap_{i=1}^{l} \langle \mathbf{y}_i \rangle^o$. Since $\{\mathbf{x}_1, \ldots, \mathbf{x}_k, \mathbf{y}_1, \ldots, \mathbf{y}_l\}$ is a basis of $x \vee y$, we have

$$\begin{aligned} (x \vee y)^o &= \left(\bigcap_{i=1}^{k} \langle \mathbf{x}_i \rangle^o \right) \cap \left(\bigcap_{i=1}^{k} \langle \mathbf{x}_i \rangle^o \right) \\ &= x^o \wedge y^o \end{aligned} \tag{4.26}$$

which proves the first statement. For the second statement, we use the isomorphism from $B_n(q)$ to $B_n^*(q)$ and apply the first statement to $x^*, y^* \in B_n^*(q)$:

$$(x^* \vee y^*)^o = x^{*o} \wedge y^{*o} \tag{4.27}$$

Since x^o, y^o are elements of $B_n^*(q)$, we get:

$$(x^o \vee y^o)^o = x^{oo} \wedge y^{oo} \tag{4.28}$$

Taking the annihilator of both sides, and recalling the identification $x^{oo} = x, \forall x \in B_n(q)$, we get

$$x^o \vee y^o = (x \wedge y)^o.$$

□

Recall that for $x \in B_n(q), q \neq 1$, we have $\rho(x) = \dim x$. This allows us to rewrite (4.20) as

$$\rho(x^o) + \rho(x) = n. \tag{4.29}$$

Remark 4.3.6 An analogue of the "annihilator" for the case where $q \to 1$ is the set-complement of subsets of $[n]$. If we define

$$^o : U \mapsto [n] \setminus U \tag{4.30}$$

where $U \subseteq [n]$, then $^o : B_n \to B_n^*$ is a lattice isomorphism, where B_n^* is the lattice dual of B_n. One may check that all the results regarding the annihilator apply to B_n.

4.4 Nullspaces and Essentiality

We now return to the problem of deriving a formula for $F_q^n(k, i)$. To achieve this, we require some linear algebraic properties of $\mathbb{F}_q^n$. Since the set $[n]$ is not a vector space, the results described in this section will not hold when we let $q \to 1$. We will assume in this section that q is prime.

If we let $\mathbf{x} \in \mathbb{F}_q^n$ be written as column vectors,

$$\mathbf{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix},$$

we may treat a k-tuple, $T = (\mathbf{x}_1, \ldots, \mathbf{x}_k)$, as a matrix whose columns are the entries of the tuple:

$$\begin{pmatrix} x_{11} & \cdots & x_{k1} \\ \vdots & \ddots & \vdots \\ x_{1n} & \cdots & x_{kn} \end{pmatrix}.$$

Such a matrix can be treated as a linear transformation, $T : \mathbb{F}_q^k \to \mathbb{F}_q^n$, and we may speak of its linear algebraic properties. One particular feature of interest is its nullspace, $\ker T$. It turns out that many properties of $(\mathbf{x}_1, \ldots, \mathbf{x}_k)$ as a k-tuple may be written as properties of its nullspace! In the rest of the section, we use T to denote a k-tuple as well as the matrix and linear transformation associated to it.

Proposition 4.4.1 *The k-tuple, T, spans $\mathbb{F}_q^n \iff \ker T$ has dimension $k - n$ (over $\mathbb{F}_q$).*

Proof. The tuple T spans $\mathbb{F}_q^n$ iff the rank of T is n. By the rank-nullity theorem, this happens iff $\ker T$ has dimension $k - n$. □

Recall the linear functionals

$$f^i : \mathbb{F}_q^k \to \mathbb{F}_q, f^i(\mathbf{e}_j) = \delta_{ij},$$

Since each f^i has rank 1, $\ker f^i$ is a $(k-1)$-dimensional subspace of $\mathbb{F}_q^k$.

Proposition 4.4.2 *For $i \in [k]$, T a k-tuple (with entries in $\mathbb{F}_q^n$),*

$$i \in \epsilon(T) \iff \ker T \subseteq \ker f^i.$$

Proof. The i^{th} column, $\mathbf{x}_i$, of $T = (\mathbf{x}_1, \ldots, \mathbf{x}_k), \mathbf{x}_j \in \mathbb{F}_q^n$ is essential iff it is not in the span of the other columns of T. This is equivalent to saying

$$\forall a_1, \ldots, a_k \in \mathbb{F}_q, \sum_{j=1}^{k} a_j \mathbf{x}_j = \mathbf{0} \implies a_i = 0 \tag{4.31}$$

which may also be written:

$$T \begin{pmatrix} a_1 \\ \vdots \\ a_k \end{pmatrix} = \mathbf{0} \implies a_i = 0 \tag{4.32}$$

This gives us

$$i \in \epsilon(T) \iff \forall \mathbf{a} \in \ker T, a_i = 0. \tag{4.33}$$

But the set of vectors $\mathbf{a} \in \mathbb{F}_q^k$ such that $a_i = 0$ is precisely $\ker f^i$, yielding the result. □

As a corollary, we get a characterization of $\epsilon(T)$:

Corollary 4.4.3 *For $i \in [k]$, T a k-tuple, $[i] = \epsilon(T)$ iff*

$$\ker T \subseteq \bigwedge_{j \in [i]} \ker f^j$$

and $\forall l \notin [i]$,

$$\ker T \not\subseteq \bigwedge_{j \in [i] \cup l} \ker f^j$$

Proof. This follows easily from Proposition 4.4.2, and noting that the negation of the condition in the Proposition gives:

$$i \notin \epsilon(T) \iff \ker T \not\subseteq \ker f^i. \tag{4.34}$$

□

We have completely characterized the tuples, T, that are counted by $F_q^n(k,i)$ in terms of their nullspaces, $\ker T$. For a fixed subspace, $N \subseteq \mathbb{F}_q^k$ the number of matrices, T, such that $\ker T = N$ is determined only by the dimension of N.

Proposition 4.4.4 *For $N \subseteq \mathbb{F}_q^k$, $\dim N = k - n$, the number of $n \times k$ matrices T such that $\ker T = N$ is*

$$(n)_q! \cdot \phi(q)^n \cdot q^{\binom{n}{2}}$$

Proof. By the rank-nullity theorem, $\dim \ker T = k - n \implies \operatorname{rank} T = n$. The rank of T is the dimension of its column span, which is equal to the dimension of its row span. Since T has n rows, these n rows must form a basis of the n dimensional row span of T. Proposition 4.1.5 shows that the number of such ordered bases (of $\mathbb{F}_q^n$) is $(n)_q! \cdot \phi(q)^n \cdot q^{\binom{n}{2}}$. □

All that remains is to count the number of nullspaces, $N \subseteq \mathbb{F}_q^k$ such that $\ker T = N \implies \epsilon(T) = [i]$. From Corollary 4.4.3, such an N must be contained in $\bigwedge_{j\in[i]} \ker f^j$, but not in any $\bigwedge_{j\in[i]\cup l} \ker f^j$ for $l \notin [i]$. The subspaces $\ker f^j$ have the following property:

Lemma 4.4.5 *For all $\mathcal{I} \subseteq [k], |\mathcal{I}| = i$,*

$$\dim\left(\bigwedge_{j\in\mathcal{I}} \ker f^j\right) = k - i$$

We will prove this lemma by proving a more general statement about collections of subspaces that have this property:

Definition 4.4.6 A collection of s maximal subspaces, $V_1, \ldots, V_s \subset \mathbb{F}_q^k, s \leq k$, is in **general position** if $\forall \mathcal{I} \subseteq [s], |\mathcal{I}| = i$,

$$\dim\left(\bigwedge_{j\in\mathcal{I}} V_j\right) = k - i.$$

Proposition 4.4.7 *A collection, $V_1, \ldots, V_s$, of maximal subspaces of $\mathbb{F}_q^k$ are in general position iff $\forall \mathcal{I} \subseteq [s], |\mathcal{I}| = i$,*

$$\dim\left(\bigvee_{j\in\mathcal{I}} V_j^o\right) = i.$$

Proof. By (4.20), $\dim\left(\bigwedge_{j\in\mathcal{I}} V_j\right) + \dim\left(\bigwedge_{j\in\mathcal{I}} V_j\right)^o = k$. Applying Proposition 4.3.5, we get:

$$\left(\bigwedge_{j\in\mathcal{I}} V_j\right)^o = \bigvee_{j\in\mathcal{I}} V_j^o, \tag{4.35}$$

which yields the proposition. □

The lemma is proved by observing that $(\ker f^j)^o = \langle f^j \rangle$ and that the f^j's are linearly independent, so

$$\dim\left(\bigvee_{j\in\mathcal{I}} (\ker f^j)^o\right) = \dim\left\langle f^j | j \in \mathcal{I}\right\rangle = i. \tag{4.36}$$

Let $d = k-i$, $V = \bigwedge_{j\in[i]} \ker f^j$, then V is isomorphic to $\mathbb{F}_q^d$ and for any $l \notin [i], V_l := \bigwedge_{j\in[i]\cup l} \ker f^j$ is isomorphic to a $(k-i-1)$-dimensional (maximal) subspace of V. In fact, $V_{i+1}, \ldots, V_k$ are in general position. We want to count the number of subspaces, $N \subseteq V$, such that $\dim N = k-n \leq d$ and N is not contained in any V_l.

Definition 4.4.8 Fix $n \geq k$, let $V = \mathbb{F}_q^n$ and let $V_1, \ldots, V_n$ be a collection of maximal subspaces in general position. Let $\begin{bmatrix[n]\\k\end{bmatrix[]}_q$ denote the number of k-dimensional subspaces of $\mathbb{F}_q^n$ that are not contained in any V_i.

Remark 4.4.9 Note that this definition does not depend on the maximal subspaces chosen, since any collection of maximal subspaces can be transformed into any other by a suitable linear transformation.

We introduce a few auxiliary functions in order to apply Möbius inversion on B_n.

Definition 4.4.10 For a subspace $U \subset V$, let

$$\sigma(U) := \left\{i \in [n] \middle| U \subseteq V_i\right\}$$

For $\mathcal{I} \subseteq [n]$, let

$$f_{\supseteq}^k(\mathcal{I}) := \left|\{U \subseteq V | \dim U = k, \sigma(U) \supseteq \mathcal{I}\}\right|$$
$$f_{=}^k(\mathcal{I}) := \left|\{U \subseteq V | \dim U = k, \sigma(U) = \mathcal{I}\}\right|$$

Then $f_{\supseteq}^k(\mathcal{I}) = \sum_{\mathcal{J}\supseteq\mathcal{I}} f_{=}^k(\mathcal{J})$, so by Möbius inversion on B_n, we have

$$\begin{aligned} f_{=}^k(\mathcal{I}) &= \sum_{\mathcal{J}\supseteq\mathcal{I}} \mu_n(\mathcal{I},\mathcal{J}) f_{\supseteq}^k(\mathcal{J}) \\ &= \sum_{\mathcal{J}\supseteq\mathcal{I}} (-1)^{|\mathcal{J}|-|\mathcal{I}|} f_{\supseteq}^k(\mathcal{J}) \end{aligned} \tag{4.37}$$

Since $\begin{bmatrix[n]\\k\end{bmatrix[]}_q$ counts the number of k-dimensional subspaces, U, such that $\sigma(U) = \varnothing$, we get:

$$\begin{aligned} \begin{bmatrix[n]\\k\end{bmatrix[]}_q &= f_{=}^k(\varnothing) \\ &= \sum_{\mathcal{J}} (-1)^{|\mathcal{J}|} f_{\supseteq}^k(\mathcal{J}) \end{aligned} \tag{4.38}$$

All that remains is to find $f_{\supseteq}^k(\mathcal{J})$:

Lemma 4.4.11

$$f_{\supseteq}^k(\mathcal{I}) = \binom{n-|\mathcal{I}|}{k}_q$$

Proof. The intersection $\bigcup_{i\in\mathcal{I}} V_i$ is a subspace of dimension $n-|\mathcal{I}|$, and $f_{\supseteq}^k(\mathcal{I})$ counts the number of ways to choose a k-dimensional subspace of $\bigcup_{i\in\mathcal{I}} V_i \cong \mathbb{F}_q^{n-|\mathcal{I}|}$ with no other restrictions. □

Corollary 4.4.12

$$\begin{bmatrix}\begin{bmatrix} n \\ k \end{bmatrix}\end{bmatrix}_q = \sum_{j=0}^{n} (-1)^j \binom{n}{j} \binom{n-j}{k}_q$$

Proof. This lemma together with (4.38) gives the formula:

$$\begin{bmatrix}\begin{bmatrix} n \\ k \end{bmatrix}\end{bmatrix}_q = \sum_{\mathcal{J}} (-1)^{|\mathcal{J}|} \binom{n-|\mathcal{J}|}{k}_q$$

But since each summand depends only on $|\mathcal{J}|$, we may sum over $|\mathcal{J}| = j$. There are $\binom{n}{j}$ subsets of cardinality j, so multiplying by this factor gives the corollary. □

Putting all of this together, we get (finally):

Theorem 4.4.13

$$F_q^n(k,i) = (n)_q! \cdot \phi(q)^n \cdot q^{\binom{n}{2}} \cdot \begin{bmatrix}\begin{bmatrix} k-i \\ k-n \end{bmatrix}\end{bmatrix}_q$$

Proof. We want the number of matrices, T, such that $\epsilon(T) = [i]$. By Corollary 4.4.3, this is equivalent to counting the number of matrices whose nullspaces satisfy the conditions in the Corollary. Corollary 4.4.12 gives the number of such nullspaces, and for a fixed nullspace, Proposition 4.4.4 gives the number of matrices with this nullspace. □

Remark 4.4.14 A key part of the derivation of $F_q^n(k,i)$ involved transferring the problem from $\mathbb{F}_q^n$ to the dual space of $\mathbb{F}_q^k$ and the lattice $B_k(q)$. This suggest that an analogous formula for the case where $q \to 1$ might be found using the lattice B_k instead of B_n, and using the set-complement in place of the annihilator.

Chapter 5

Finite Nilpotent Groups

Finally, we return to our original problem: finding the number of irredundant generating k-sets of finite nilpotent groups. Nilpotent groups are "almost abelian", and this property allows us to reduce the problem to that of finding irredundant generating sets of groups that are direct products of elementary abelian groups.

5.1 The Lattice of Subgroups

A lot of information about a group can be found through its lattice of subgroups. In this section, we develop the necessary results concerning the lattice of subgroups of G that we will need to find $\phi_k(G)$.

Definition 5.1.1 The **lattice of subgroups**, $L(G)$, of a group G is the set of subgroups of G equipped with $\leq, \wedge, \vee$ such that for $x, y \in L(G)$,

- $x \leq y \iff x$ is a subgroup of y
- $x \vee y = \langle x, y \rangle$
- $x \wedge y = x \cap y$

Definition 5.1.2 The **Frattini subgroup** of G, denoted $\Phi(G)$, is the intersection of all maximal subgroups of G.

Dually, we may consider the minimal subgroups of a group. These are the subgroups of prime order.

Definition 5.1.3 The **Frattini dual** of G, denoted $\Psi(G)$, is the subgroup generated by all the minimal subgroups of G.

Definition 5.1.4 An element $g \in G$ is a **non-generator** if it can be removed from any generating set without destroying the generating property.

The Frattini subgroup is of particular interest to us because of the following proposition:

Proposition 5.1.5 *$\Phi(G)$ is the set of non-generators of G.*

Proof. Suppose $g \in G$ is a non-generator, but $g \notin \Phi(G)$. Then there is some maximal subgroup $M < G$ such that $g \notin M$. Let $\{g_1, \ldots, g_n\}$ be a generating set of M. Then since M is maximal and $g \notin M$, $\{g, g_1, \ldots, g_n\}$ is a generating set of G, and further, g cannot be removed from this set without destroying the generating property. This contradicts our assumption that g is a non-generator, and shows that the set of non-generators is contained in $\Phi(G)$.

Suppose now that $g \in \Phi(G)$ but g is not a non-generator. Then there is some generating set, $S = \{g, g_1, \ldots, g_n\}$, such that $S \backslash g$ does not generate G, so $S \backslash g$ must be contained in some maximal subgroup $M < G$. But $g \in \Phi(G)$, the intersection of all maximal subgroups of G; in particular, $g \in M$, so S generates some subgroup of M, which is a proper subgroup of G. This contradicts our assumption that S generates G, and shows that $\Phi(G)$ is contained in the set of non-generators. □

Proposition 5.1.6 *Let G be finite, and let $S \subseteq G$. Then S irredundantly generates G iff the image of S in $G/\Phi(G)$ irredundantly generates $G/\Phi(G)$.*

Proof. Let $S = \{g_1, \ldots g_n\} \subseteq G$, and let $\bar{S} = \{g_1\Phi(G), \ldots g_n\Phi(G)\} \subseteq G/\Phi(G)$ be the image of S in $G/\Phi(G)$. It is easy to see that S generates $G \implies \bar{S}$ generates $G/\Phi(G)$. Contrapositively, if S does not generate G, then $\langle S\rangle \subseteq M \subsetneq G$, for some maximal subgroup, M. In particular, $\exists g \in G$ such that $g \notin M$, so $g \notin \Phi(G)$, and in the quotient $g\Phi(G) \notin \langle \bar{S}\rangle$, so $\bar{S}$ does not generate $G/\Phi(G)$. This shows that S generates $G \iff \bar{S}$ generates $G/\Phi(G)$.

But in fact, this is sufficient to show that S is irredundant generating $\iff \bar{S}$ is irredundant generating. Indeed, suppose S is irredundant generating but $\bar{S}$ is not irredundant (although still generating). Then some proper subset of $\bar{S}$ is generating, and the preimage of this proper subset is a proper subset of S which is generating, which contradicts our assumption that S is irredundant generating. Similarly, if $\bar{S}$ is irredundant generating, but S is not irredundant, then some proper subset of S is generating. The image of this proper subset of S is a proper subset of $\bar{S}$ that is generating, which contradicts the assumption that $\bar{S}$ is irredundant. □

Corollary 5.1.7 *For G finite nilpotent,*

$$\phi_k(G) = |\Phi(G)|^k \phi_k\big(G/\Phi(G)\big) \tag{5.1}$$

Proof. Let $S = \{g_1\Phi(G), \ldots, g_k\Phi(G)\}$ be an irredundant generating set in $G/\Phi(G)$. Then there are $|\Phi(G)|^k$ irredundant generating sets of G whose image in $G/\Phi(G)$ is S. By the proposition, this accounts for all irredundant generating k-sets of G. □

It is also obvious that the proposition gives:

Corollary 5.1.8 *For G finite nilpotent,*

$$r(G) = r(G/\Phi(G)) \tag{5.2}$$
$$m(G) = m(G/\Phi(G)) \tag{5.3}$$

Given groups G, H, we may form the direct product $G \times H$ and ask how $L(G \times H)$ is related to $L(G)$ and $L(H)$. A comprehensive discussion of the properties of $L(G \times H)$, and a proof of Proposition 5.1.10 may be found in [7].

Definition 5.1.9 Two groups, G, H are **relatively prime** if they have no non-trivial common quotient. i.e. if $\exists N \trianglelefteq G$ and $M \trianglelefteq H$ such that $G/N \cong H/M$, then both quotients are trivial.

Proposition 5.1.10 *If G, H are relatively prime, then $L(G \times H) = L(G) \times L(H)$.*

For our purposes, it suffices to know the following special case:

Proposition 5.1.11 *Let G, H be such that $(|G|, |H|) = 1$. Then $L(G \times H) = L(G) \times L(H)$.*

This proposition is a direct consequence of the following lemma.

Lemma 5.1.12 *Let G, H be such that $(|G|, |H|) = 1$. Let $(g_1, h_1), \ldots, (g_k, h_k)$ be elements of $G \times H$. Then*

$$\big\langle (g_1, h_1), \ldots, (g_k, h_k)\big\rangle = \langle g_1, \ldots, g_k\rangle \times \langle h_1, \ldots, h_k\rangle$$

Proof. It is obvious that

$$\big\langle (g_1,h_1),\dots,(g_k,h_k)\big\rangle \le \langle g_1,\dots,g_k\rangle \times \langle h_1,\dots,h_k\rangle \tag{5.4}$$

The opposite inclusion follows from the Chinese Remainder Theorem. Since $(|G|,|H|)=1$, we have that for all $i\in[k]$, $\big(o(g_i),o(h_i)\big)=1$. By the Chinese Remainder Theorem, for each $i\in[k]$, $\exists n\in\mathbb{N}$ such that

$$\begin{aligned} n &\equiv 1 \pmod{o(g_i)}\\ n &\equiv 0 \pmod{o(h_i)}. \end{aligned}$$

This gives

$$(g_i,h_i)^n = (g_i,\mathrm{id}) \tag{5.5}$$

for all $i\in[k]$, so $\big\langle (g_1,h_1),\dots,(g_k,h_k)\big\rangle$ contains

$$\big\langle (g_1,\ \mathrm{id}),\dots,(g_k,\ \mathrm{id})\big\rangle = \langle g_1,\dots,g_k\rangle \times \{\ \mathrm{id}\}. \tag{5.6}$$

By a similar argument applied to the h_j's, we have that

$$\{\mathrm{id}\}\times\langle h_1,\dots,h_k\rangle \le \big\langle (g_1,h_1),\dots,(g_k,h_k)\big\rangle. \tag{5.7}$$

Since $\langle g_1,\dots,g_k\rangle\times\{\mathrm{id}\}$ and $\{\mathrm{id}\}\times\langle h_1,\dots,h_k\rangle$ together generate $\langle g_1,\dots,g_k\rangle\times\langle h_1,\dots,h_k\rangle$, we get the opposite inclusion

$$\big\langle (g_1,h_1),\dots,(g_k,h_k)\big\rangle \ge \langle g_1,\dots,g_k\rangle \times \langle h_1,\dots,h_k\rangle \tag{5.8}$$

which yields the lemma. □

This Lemma shows that all subgroups of $G\times H$ are of the form $A\times B$, where $A\le G$ and $B\le H$, so the ground set of $L(G\times H)$ is equal to that of $L(G)\times L(H)$. It is easy to check that $\le,\wedge,\vee$ are preserved as well, hence Proposition 5.1.11 holds. We will use this proposition to prove a few properties about $G\times H$, when $(|G|,|H|)=1$.

Proposition 5.1.13 *If $L(G\times H)=L(G)\times L(H)$, then $\Phi(G\times H)=\Phi(G)\times\Phi(H)$.*

Proof. Maximal subgroups of G are coatoms of $L(G)$, so $\Phi(G)=\bigwedge CoA(L(G))$. If $x_1,\dots,x_k$ are coatoms of $L(G)$, and $y_1,\dots,y_l$ are coatoms of $L(H)$, then the coatoms of $L(G)\times L(H)$ are $(x_1,H),\dots,(x_k,H),(G,y_1),\dots,(G,y_l)$. Since $L(G\times H)=L(G)\times L(H)$,

$$\begin{aligned} \bigwedge CoA(L(G\times H)) &= (x_1,H)\wedge\cdots\wedge(x_k,H)\wedge(G,y_1)\wedge\cdots\wedge(G,y_l)\\ &= (\Phi(G),H)\wedge(G,\Phi(H))\\ &= (\Phi(G),\Phi(H)). \end{aligned}$$

□

The Frattini subgroup of a group is always normal, so we may form the quotient $G/\Phi(G)$. If $L(G\times H)=L(G)\times L(H)$, then since $\Phi(G\times H)=\Phi(G)\times\Phi(H)$, we have

Corollary 5.1.14

$$(G\times H)/\Phi(G\times H) = (G/\Phi(G))\times(H/\Phi(H)).$$

Proposition 5.1.15 *If $L(G\times H)=L(G)\times L(H)$, then $\Psi(G\times H)=\Psi(G)\times\Psi(H)$.*

Proof. Minimal subgroups of G are atoms of $L(G)$, so $\Psi(G)=\bigvee A(L(G))$. Replacing coatoms by atoms, and $\wedge$ by $\vee$ in the proof of Proposition 5.1.13 gives the result. □

Recall that $r(G)$ and $m(G)$ are the cardinalities of the smallest and largest irredundant generating sets of G, respectively. We want to know the relation of $r(G \times H)$ and $m(G \times H)$ to $r(G), r(H), m(G), m(H)$ when $L(G \times H) = L(G) \times L(H)$. First, we reformulate the notion of an irredundant generating set in $L(G)$.

Definition 5.1.16 Given a lattice, L, and a finite subset $S \subseteq L$, S is an **irredundant generating k-set** of L if it satisfies the following properties:

1. $|S| = k$
2. $\bigvee S = \hat{1}_L$
3. $\forall x \in S, \bigvee (S \setminus x) \lneq \bigvee S$

A subset that satisfies condition 2 is **generating**, while a subset that satisfies condition 3 is **irredundant**.

It is easy to see that $\{g_1, \ldots, g_k\}$ is an irredundant generating k-set of G iff $\{\langle g_1 \rangle, \ldots, \langle g_k \rangle\}$ is an irredundant generating k-set of $L(G)$. Then $r(G)$ and $m(G)$ are the cardinalities of the smallest and largest irredundant generating sets of $L(G)$ *consisting of cyclic subgroups of* G. This qualification is necessary as not all subgroups of G are cyclic.

We now prove some results concerning $r(G \times H)$ and $m(G \times H)$ in the special case where $L(G \times H) = L(G) \times L(H)$. A more general and comprehensive treatment can be found in [1].

Proposition 5.1.17 *If $L(G \times H) = L(G) \times L(H)$, then $r(G \times H) = \max\{r(G), r(H)\}$.*

Proof. Let $k = r(G), l = r(H)$, and let $\{\langle g_1 \rangle, \ldots, \langle g_k \rangle\}$ and $\{\langle h_1 \rangle, \ldots, \langle h_l \rangle\}$ be irredundant generating sets of $L(G)$ and $L(H)$, respectively. Without loss of generality, let us suppose that $k \geq l$. Then

$$\big\{\big(\langle g_1 \rangle, \langle h_1 \rangle\big), \ldots, \big(\langle g_l \rangle, \langle h_l \rangle\big), \big(\langle g_{l+1} \rangle, \langle \mathrm{id} \rangle\big), \ldots, \big(\langle g_k \rangle, \langle \mathrm{id} \rangle\big)\big\}$$

is an irredundundant generating k-set of $L(G \times H)$. This shows that $r(G \times H) \leq \max\{r(G), r(H)\}$.

To get the opposite inequality, suppose $\{(\langle g_1 \rangle, \langle h_1 \rangle), \ldots, (\langle g_n \rangle, \langle h_n \rangle)\}$ is an irredundant generating set of $L(G \times H)$. Then $\{\langle g_1 \rangle, \ldots, \langle g_n \rangle\}$ is a generating set of $L(G)$ so $n \geq r(G)$. Similarly, $n \geq r(H)$, so $n \geq \max\{r(G), r(H)\}$. This completes the proof. □

There is a similar result for $m(G \times H)$, which is true regardless of whether $L(G \times H) = L(G) \times L(H)$.

Proposition 5.1.18 (Collins [1]) *For G, H finite groups,*

$$m(G \times H) = m(G) + m(H)$$

Proof. Let $k = m(G), l = m(H)$, and let $\{g_1, \ldots, g_k\}$ and $\{h_1, \ldots, h_l\}$ be irredundant generating sets of G and H, respectively. Then

$$\{(g_1, \mathrm{id}), \ldots, (g_k, \mathrm{id}), (\mathrm{id}, h_1), \ldots, (\mathrm{id}, h_l)\}$$

is an irredundant generating sequence of $G \times H$, so $m(G \times H) \geq m(G) + m(H)$.

The opposite inequality is harder. We refer readers to [1] for a proof of the fact that if $N \trianglelefteq K$ for some group K, then

$$m(K) \leq m(K/N) + m(N).$$

The proposition follows by letting $N = G \times \{\mathrm{id}\}$ and $K = G \times H$. □

5.2 Generators in Nilpotent Groups

Definition 5.2.1 Given G and subgroups $H, K \leq G$, their **commutator**, $[H, K]$, is the subgroup generated by all $h^{-1}k^{-1}hk$ for all $h \in H, k \in K$.

Definition 5.2.2 G is a **nilpotent group** if there exists a finite series of subgroups

$$\{1\} = G_0 \leq G_1 \leq \cdots \leq G_n = G$$

such that $[G, G_{i+1}] \leq G_i$ for all i. Such a series is a finite central series.

Definition 5.2.3 The **commutator subgroup** of G is the subgroup $G' := [G, G]$.

The commutator subgroup gives a measure of how non-abelian a group is. More formally, we have the following standard result from group theory:

Proposition 5.2.4 *Let $N \trianglelefteq G$ be a normal subgroup of G. Then G/N is abelian $\iff$ $N \geq G'$.*

The following characterizes finite nilpotent groups. A proof, along with a more comprehensive characterization, can be found in [6].

Theorem 5.2.5 *Let G be finite. Then the following are equivalent:*

1. *G is nilpotent*
2. *$\Phi(G) \geq G'$*
3. *G is a direct product of p-groups*

Corollary 5.2.6 *Let G be finite nilpotent. Then $G/\Phi(G)$ is abelian. In fact,*

$$G/\Phi(G) \cong \mathbb{Z}_{p_1}^{m_1} \times \cdots \times \mathbb{Z}_{p_n}^{m_n} \tag{5.9}$$

where p_i are distinct primes.

Proof. By the second characterization, $\Phi(G) \geq G'$, so $G/\Phi(G)$ is abelian. By the third characterization, we can write $G = P_1 \times \cdots \times P_n$, where P_i are p_i-groups for distinct primes, p_i. From Proposition 5.1.13, $\Phi(G) = \Phi(P_1) \times \cdots \times \Phi(P_n)$, so we get:

$$G/\Phi(G) = P_1/\Phi(P_1) \times \cdots \times P_n/\Phi(P_n).$$

The Frattini quotient of a p-group is an elementary abelian group, yielding the result. □

Since $\mathbb{Z}_p^m$ is a vector space, $r(\mathbb{Z}_p^m) = m(\mathbb{Z}_p^m) = m$. If G is nilpotent, and $G/\Phi(G) \cong \mathbb{Z}_{p_1}^{m_1} \times \cdots \times \mathbb{Z}_{p_n}^{m_n}$, where p_i are distinct primes, then Propositions 5.1.17 and 5.1.18 together with Corollary 5.1.8 give

$$\begin{aligned} r(G) &= \max\{m_1, \ldots, m_n\} \\ m(G) &= m_1 + m_2 + \cdots + m_n. \end{aligned} \tag{5.10}$$

Corollary 5.1.7 reduces the problem of finding $\phi_k(G)$ to that of finding $\phi_k(G/\Phi(G))$. We thus turn our attention to direct products of elementary abelian groups.

5.3 Direct Products of Elementary Abelian Groups

For the rest of the chapter, fix n and the n-tuples $\mathbf{p} = (p_1, \ldots, p_n), \mathbf{m} = (m_1, \ldots, m_n)$ where p_i are distinct primes, and $m_i \in \mathbb{N}$. Let $G = G_{\mathbf{p}}^{\mathbf{m}} := \mathbb{Z}_{p_1}^{m_1} \times \cdots \times \mathbb{Z}_{p_n}^{m_n}$, where p_i are distinct primes. Elements of G will be written $g = (\mathbf{x}^1, \ldots, \mathbf{x}^n)$, where $\mathbf{x}^i \in \mathbb{Z}_{p_i}^{m_i}$ (here and elsewhere, the superscripts are merely indices, not exponents). We want to know when a set $\{g_1, \ldots, g_k\} \subseteq G$ is an irredundant generating set .

Let $\mathcal{C} = \{g_1, \ldots, g_k\}, g_i = (\mathbf{x}_i^1, \ldots, \mathbf{x}_i^n) \in G$. By Lemma 5.1.12, $\mathcal{C}$ is generating iff for all $j \in [n]$,

$$\langle \mathbf{x}_1^j, \ldots, \mathbf{x}_k^j \rangle = \mathbb{Z}_{p_j}^{m_j}. \tag{5.11}$$

Next we want to determine when $g_1, \ldots, g_k$ is irredundant. We reintroduce essentiality:

Definition 5.3.1 Let $\mathcal{C} = \{g_1, \ldots, g_k\}, g_i = (\mathbf{x}_i^1, \ldots, \mathbf{x}_i^n) \in G$. An entry $\mathbf{x}_i^j$ is **essential** in $\mathcal{C}$ if

$$\mathbf{x}_i^j \notin \left\langle \mathbf{x}_l^j, l \in [k], l \neq i \right\rangle.$$

Then the following characterization of irredundant sets of G should be obvious:

Proposition 5.3.2 *$\{g_1, \ldots, g_k\} \subseteq G$ is irredundant iff each g_i contains at least one essential entry.*

In fact, we can get a more useful characterization.

Definition 5.3.3 For $\mathcal{C} = \{g_1, \ldots, g_k\} \subseteq G$, let $X^j \subseteq \mathbb{Z}_{p_j}^{m_j}$ denote the set of essential entries that are from $\mathbb{Z}_{p_j}^{n_j}$, and let $B_i \subseteq \bigsqcup_{j=1}^n X^j$ denote the essential entries that are entries of g_i. Let $X(\mathcal{C}) := X^{\mathbf{m}} = \bigsqcup_{j=1}^n X^j$ be the set of essential entries, and $B(\mathcal{C}) := \{B_1, \ldots, B_k\}$ be the corresponding collection of essential subsets of g_i.

Proposition 5.3.4 *$\mathcal{C} = \{g_1, \ldots, g_k\} \subseteq G$ is irredundant $\iff$ $B(\mathcal{C})$ is a product k-partition of $X(\mathcal{C})$.*

Proof. If $\{g_1, \ldots, g_k\} \subseteq G$ is irredundant, then each g_i contains at least one essential entry, so $B_i \neq \varnothing$. All essential entries must belong to some g_i, so $\{B_1, \ldots, B_k\}$ is a k-partition of $\bigsqcup_{j=1}^n X^j$. Further, each group element, g_i, has as its entries only one element from each $\mathbb{Z}_{p_j}^{m_j}$, making this partition a product k-partition. Conversely, if $\{B_1, \ldots, B_k\}$ is a product k-partition of $\bigsqcup_{j=1}^n j$, then $B_i \neq \varnothing$ for all i, so each g_i contains an essential element, hence $\{g_1, \ldots, g_k\} \subseteq G$ is irredundant. □

This proposition is the crux of this thesis. It links the results of Chapter 3 and 4, and allows us to write the main theorem of this paper:

Theorem 5.3.5

$$\phi_k(G) = \sum_{\mathbf{0} \leq \mathbf{e} \leq \mathbf{m}} \left\{ \begin{matrix} \mathbf{e} \\ k \end{matrix} \right\} \prod_{j=1}^n \frac{1}{e_j!} F_{p_j}^{m_j}(k, e_j)$$

To prove the theorem, we will use the following lemma:

Lemma 5.3.6 *Let $E = (\mathbf{x}^1, \ldots, \mathbf{x}^e), \mathbf{x}^j \in \mathbb{Z}_p^m$ be an independent e-tuple. The number of k-tuples (of elements of $\mathbb{Z}_p^m$) whose essential index set is $[e]$, and whose initial tuple of e entries are exactly E is:*

$$\frac{1}{\binom{m}{e}_p (e)_p! \phi(p)^e p^{\binom{e}{2}}} F_p^m(k, e)$$

Proof. $F_p^m(k,e)$ is the number of k-tuples whose essential index set is $[e]$. For any e-tuple of independent elements, the number of these k-tuples whose initial segments are exactly that e-tuple is the total number of k-tuples divided by the number of ways to choose the initial e-tuple:

$$\frac{1}{\binom{m}{e}_p (e)_p! \phi(p)^e p^{\binom{e}{2}}} F_p^m(k,e)$$

□

Proof of theorem. By Proposition 5.3.4, we may construct irredundant generating k-sets in the following manner:

1. From each $\mathbb{Z}_{p_j}^{m_j}$, choose an independent e_j-set of elements, X^j, where $e_j \leq m_j$.
2. Choose a product k-partition, $\{B_1, \ldots, B_k\}$, of $X^{\mathbf{e}} := \bigsqcup_{j=1}^n X^j$ and write these B_i as rows of a matrix such that if an essential entry is from $\mathbb{Z}_{p_j}^{m_j}$, then it is in the j^{th} column of matrix. Leave all other entries empty.
3. For all $j \in [n]$, fill in the empty entries of the j^{th} column such that the resulting column generates $\mathbb{Z}_{p_j}^{m_j}$, but the only essential entries in the column are those that were entered in the previous step.
4. Let $g_1, \ldots g_k$ be the rows of the resulting matrix. Then $\{g_1, \ldots, g_k\}$ is an irredundant generating k-set of G.

The first two steps may be swapped, by first fixing some $\mathbf{e} = (e_1, \ldots, e_n)$, choosing a product k-partition of $X^{\mathbf{e}}$ (with elements x_i^j as placeholders), and then choosing independent elements from $\mathbb{Z}_{p_j}^{m_j}$ to be the elements of $X^{\mathbf{e}}$. This allows us to sum over $\mathbf{0} \leq \mathbf{e} \leq \mathbf{m}$, and for each $\mathbf{e}$, there are $\left\{ {\mathbf{e} \atop k} \right\}$ product k-partitions. By Proposition 4.1.15, for each of these partitions, there are

$$\prod_{j=1}^n \binom{m_j}{e_j}_{p_j} \frac{(e_j)_q!}{e_j!} \cdot \phi(p_j)^{e_j} \cdot {p_j}^{\binom{e_j}{2}}$$

ways to choose an independent e_j-set of elements. Note that these are sets, not tuples, since the product partition will take care of the ordering of these elements. Using the lemma, there are

$$\prod_{j=1}^n \frac{1}{\binom{m_j}{e_j}_{p_j} (e_j)_{p_j}! \phi(p_j)^{e_j} {p_j}^{\binom{e_j}{2}}} F_{p_j}^{m_j}(k, e_j)$$

ways to fill in the non-essential entries of the matrix. Multiplying these two expressions, the number of irredundant generating k-sets for a particular product k-partition is

$$\prod_{j=1}^n \frac{1}{e_j!} F_{p_j}^{m_j}(k, e_j).$$

This proves the theorem. □

The theorem also tells us the possible sizes for an irredundant generating set of G. Let k be the cardinality of an irredundant generating set of G. If $|\mathbf{e}| \leq |\mathbf{m}| < k$, then $\left\{ {\mathbf{e} \atop k} \right\} \leq \left\{ {|\mathbf{e}| \atop k} \right\} = 0$, so we must have $k \leq |\mathbf{m}|$. On the other hand, $F_{p_j}^{m_j}(k, e_j)$ is a multiple of $\left[\!\left[{k - e_j \atop k - m_j} \right]\!\right]_{p_j}$, which is zero if $k - m_j < 0$, so we must have $m_j \leq k$ for all j. Thus, we get (5.10) as a corollary of the theorem.

Chapter 6

Generalization to Direct Products of Lattices

This chapter generalizes the proof of the main theorem, to obtain a formula for counting the number of irredundant generating sets in direct products of lattices, provided we know certain facts about the lattices that make up the direct product. As a corollary, we will obtain the formula for $\phi_k(n)$. Throughout, we will assume our lattices are finite, and contain $\hat{0}$ and $\hat{1}$.

6.1 Irredundance and Essentiality in Lattices

We have already defined irredundance for $B_n(q)$. Here we define irredundance and essentiality for arbitrary lattices.

Definition 6.1.1 Let L be a finite lattice. A k-tuple $(x_1, \ldots, x_k), x_i \in L$ is **irredundant** if

$$\forall i, x_i \not\leq \left(\bigvee_{j \neq i} x_j \right).$$

Otherwise, it is called **redundant**.

Definition 6.1.2 Given a tuple $(x_1, \ldots, x_k), x_i \in L$, an entry x_i is **essential** if

$$x_i \not\leq \left(\bigvee_{j \neq i} x_j \right)$$

An index $i \in [k]$ for which $\mathbf{x}_i$ is essential is called an **essential index**. The **essential index set** of a tuple is the set of its essential indices.

As was the case with $B_n(q)$, a subset of L is irredundant iff all its members are essential. We want to count irredundant generating sets of L, but first we have to define what it means to generate L.

Definition 6.1.3 A subset $\{x_1, \ldots, x_k\} \subseteq L$ **generates** L if $\bigvee_{i=1}^{k} x_i = \hat{1}$.

An irredundant generating k-set of L is thus a subset of cardinality k that is irredundant and generates L, and $\phi_k(L)$ will denote the number of such k-sets. We will find a formula for $\phi_k(L)$ when $L = L_1 \times \cdots \times L_n$. We recall one last definition before presenting the theorem.

Definition 6.1.4 For $i \leq k$, let $F_L(k,i)$ denote the number of k-tuples of elements from L that generate L, and whose essential index sets are exactly $[i]$.

Theorem 6.1.5 *Let $L = L_1 \times \cdots \times L_n$. Let m_j denote the size of the largest irredundant subset of L_j, and let $\mathbf{m} = (m_1, \ldots, m_n)$. For ease of notation, let $F_j := F_{L_j}$. Then*

$$\phi_k(L) = \sum_{\mathbf{0} \leq \mathbf{e} \leq \mathbf{m}} \left\{ \begin{matrix} \mathbf{e} \\ k \end{matrix} \right\} \prod_{j=1}^{n} \frac{1}{e_j!} F_j(k, e_j)$$

The proof of the theorem follows the same construction as the proof of Theorem 5.3.5.

Definition 6.1.6 For $\mathcal{C} = \{\mathbf{x}_1, \ldots, \mathbf{x}_k\} \subseteq L, \mathbf{x}_i = (x_i^1, \ldots, x_i^n)$, $x_i^j \in L_j$, an entry x_i^j is essential if it is essential in the tuple $(x_1^j, x_2^j, \ldots, x_k^j)$. Let $X^j \subseteq L_j$ denote the set of essential entries that are from L_j, let $e_j = |X^j|$, and let $X(\mathcal{C}) = X^{\mathbf{e}} = \bigsqcup_{j=1}^{n} X^j$. Let $B_i \subseteq X(\mathcal{C})$ denote the essential entries that are entries of $\mathbf{x}_i$, and let $B(\mathcal{C}) = \{B_1, \ldots, B_k\}$.

Proposition 6.1.7 *$\mathcal{C} = \{\mathbf{x}_1, \ldots, \mathbf{x}_k\} \subseteq L$ is irredundant $\iff$ $B(\mathcal{C})$ is a product k-partition of $X(C)$.*

The proof is essentially the same as that of Proposition 5.3.4. With this proposition, we are able to construct irredundant generating k-sets of L in the following manner:

1. From each L_j, choose an irredundant e_j-set of elements, X^j, where $e_j \leq m_j$.

2. Choose a product k-partition, $\{B_1, \ldots, B_k\}$, of $X^{\mathbf{e}} := \bigsqcup_{j=1}^{n} X^j$ and write these B_i as rows of a matrix such that if an essential entry is from L_j, then it is in the j^{th} column of matrix. Leave all other entries empty.

3. For all $j \in [n]$, fill in the empty entries of the j^{th} column such that the resulting column is a tuple that generates L_j, but the only essential entries in the column are those that were entered in the previous step.

4. Let $\mathbf{x}_1, \ldots, \mathbf{x}_k$ be the rows of the resulting matrix. Then $\mathbf{x}_1, \ldots, \mathbf{x}_k$ is an irredundant generating k-set of L.

Once again, we may swap the first two steps, deciding instead to first choose a partition of an arbitrary set $X^{\mathbf{e}} = \bigsqcup_{j=1}^{n} X^j$, and then assigning elements from L_j to be the elements of X^j. $\left\{ \begin{matrix} \mathbf{e} \\ k \end{matrix} \right\}$ gives the number of such product k-partitions. $F_j(k, e_j)$ counts the number of k-tuples whose essential index sets are some fixed set with cardinality e_j. This set is specified by the product k-partition. We divide by $e_j!$ because the partition takes care of the ordering of essential elements. Once the essential elements are partitioned, each block is distinguished by the essential elements it contains. Summing over $\mathbf{e}$, which specifies the number of essential entries from each L_j, we get the theorem.

6.2 Minimal k-covers

We now apply this theorem to count minimal k-covers of $[n]$. This application is based on the following standard combinatorial fact which can be found in [8]:

Proposition 6.2.1

$$B_n \cong B_1 \times \cdots \times B_1 = (B_1)^n$$

B_1 is the poset consisting of only two elements, $\hat{0} \lneq \hat{1}$. The lattice isomorphism sends a subset $s \subseteq [n]$ to the element $(s_1, \ldots, s_n)$, where

$$s_j = \begin{cases} \hat{1} & \text{if } j \in s \\ \hat{0} & \text{otherwise} \end{cases} \tag{6.1}$$

If we write $0, 1$ instead of $\hat{0}, \hat{1}$, then this is just the characteristic function, where the 1's tell us which elements are in s. Under this isomorphism, a minimal k-cover of $[n]$ is precisely an irredundant generating k-set of $(B_1)^n$.

In B_1, there is only one irredundant subset, $\{\hat{1}\}$, so $m_j = 1$ for all $j \in [n]$. As we saw in Remark 3.2.7, if $\mathbf{e} \leq \mathbf{m}$ consists of only 0s and 1s, then $\left\{ {\mathbf{e} \atop k} \right\} = \left\{ {|\mathbf{e}| \atop k} \right\}$. Next, a k-tuple of elements from B_1 consists of only $\hat{0}$ or $\hat{1}$. It generates iff there is at least one entry that is $\hat{1}$, and it does so irredundantly iff there is exactly one entry with $\hat{1}$. Let $F(k,i) = F_{B_1}(k,i)$, then $F(k,1) = 1$, since all other entries must be $\hat{0}$, and $F(k,0) = 2^k - k - 1$, since of the 2^k ways to fill a k-tuple with $\hat{0}$ or $\hat{1}$, one of them does not generate (the one with all $\hat{0}$s), and k of them generate irredundantly (the ones with a $\hat{1}$ in one entry and $\hat{0}$s everywhere else). And of course, $\frac{1}{0!} = \frac{1}{1!} = 1$.

Putting all this together, we get

$$\begin{aligned} \phi_k(n) &= \sum_{\mathbf{0} \leq \mathbf{e} \leq \mathbf{m}} \left\{ {\mathbf{e} \atop k} \right\} \prod_{j=1}^{n} \frac{1}{e_j} F(k, e_j) \\ &= \sum_{\mathbf{0} \leq \mathbf{e} \leq \mathbf{m}} \left\{ {|\mathbf{e}| \atop k} \right\} \prod_{j=1}^{n} (2^k - k - 1)^{1 - e_j} \\ &= \sum_{\mathbf{0} \leq \mathbf{e} \leq \mathbf{m}} \left\{ {|\mathbf{e}| \atop k} \right\} (2^k - k - 1)^{n - |\mathbf{e}|} \\ &= \sum_{0 \leq i \leq n} \binom{n}{i} \left\{ {i \atop k} \right\} (2^k - k - 1)^{n-i} \end{aligned} \tag{6.2}$$

where the last equation follows from letting $|\mathbf{e}| = i$, and noting that there are $\binom{n}{i}$ tuples for each i. This proves Theorem 1.0.4.

6.3 Cyclic Groups of Squarefree Order

Let $m = p_1 p_2 \ldots p_n$, where p_i are distinct primes. Then $G = \mathbb{Z}_m$ is a cyclic group of square-free order, and is isomorphic to

$$\mathbb{Z}_{p_1} \times \cdots \times \mathbb{Z}_{p_n}. \tag{6.3}$$

Each $\mathbb{Z}_{p_i}$, has $\{\mathrm{id}\}, \mathbb{Z}_{p_i}$ as its only subgroups, so $L(\mathbb{Z}_{p_i})$ is the poset with two elements, $\hat{0} \lneq \hat{1}$, which is simply the lattice B_1. $L(G)$ is thus isomorphic to B_n, so

$$\phi_k(L(G)) = \phi_k(n). \tag{6.4}$$

To get $\phi_k(G)$, we compute

$$F^1_{p_i}(k,0) = (k)_{p_i} - k \tag{6.5}$$

$$F^1_{p_i}(k,1) = \phi(p_i) \tag{6.6}$$

which, when plugged into Theorem 5.3.5 gives

$$\begin{aligned}\phi_k(G) &= \sum_{\mathbf{0}\le\mathbf{e}\le\mathbf{1}} \left\{ \begin{matrix} \mathbf{e} \\ k \end{matrix} \right\} \prod_{j=1}^{n} F^1_{p_j}(k, e_j) \\ &= \sum_{\mathbf{0}\le\mathbf{e}\le\mathbf{1}} \left\{ \begin{matrix} |\mathbf{e}| \\ k \end{matrix} \right\} \prod_{j=1}^{n} \Big(\big((k)_{p_j} - k\big)^{1-e_j} \phi(p_j)^{e_j} \Big). \end{aligned} \tag{6.7}$$

Note that $\phi(2) = 1$, and $(k)_2 - k = 2^k - k - 1$, so if we set $p_j = 2$ for all j, we recover the formula for $\phi_k(n)$.

The similarity between $\phi_k(n)$ and $\phi_k(G)$, for $G = \mathbb{Z}_{p_1} \times \cdots \times \mathbb{Z}_{p_n}, p_i$'s distinct primes , was the original inspiration for this thesis. In fact, using ideas similar to (and simpler than) those presented in this thesis, we may define a polynomial:

$$f_{n,k}(x_1, \ldots, x_n) = \sum_{\boldsymbol{\alpha}} \left(\frac{1}{k^{n_{\boldsymbol{\alpha}}}} \left\{ \begin{matrix} n_{\boldsymbol{\alpha}} \\ k \end{matrix} \right\} \prod_{\alpha_i \in \boldsymbol{\alpha}} \binom{k}{\alpha_i} \right) \mathbf{x}^{\boldsymbol{\alpha}} \tag{6.8}$$

where $n_{\boldsymbol{\alpha}}$ is the number of entries of the n-tuple, $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$, that are equal to 1, and $\mathbf{x}^{\boldsymbol{\alpha}} = x_1^{\alpha_1} x_2^{\alpha_2} \ldots x_n^{\alpha_n}$. We get the nice result:

$$\begin{aligned}\phi_k(n) &= f_{n,k}(1, 1, \ldots, 1) \\ &= f_{n,k}\big(\phi(2), \phi(2), \ldots, \phi(2)\big) \qquad (6.9) \\ \phi_k(G) &= f_{n,k}\big(\phi(p_1), \phi(p_2), \ldots, \phi(p_n)\big). \qquad (6.10)\end{aligned}$$

Bibliography

[1] D. Collins. Generating sequences of finite groups.

[2] G.A. Grätzer. *General lattice theory*. Birkhäuser, 2003.

[3] T. Hearne and C. Wagner. Minimal covers of finite sets. *Discrete Mathematics*, 5(3):247–251, 1973.

[4] V.G. Kac and P. Cheung. *Quantum calculus*. Springer Verlag, 2002.

[5] A.J. Macula. Lewis Carroll and the enumeration of minimal covers. *Mathematics Magazine*, 68(4):269–274, 1995.

[6] M. Quick. Nilpotent groups. `http://www-groups.mcs.st-andrews.ac.uk/~martyn/teaching/5824/5824nilpotent.pdf`.

[7] R. Schmidt. *Subgroup lattices of groups*, volume 14. Walter de Gruyter, 1994.

[8] R.P. Stanley. *Enumerative combinatorics*, volume 1. Cambridge Univ Pr, 2001.

[9] A. Tarski. An interpolation theorem for irredundant bases of closure structures. *Discrete Mathematics*, 12(2):185–192, 1975.