Algebra

by

Stephen S. Shatz[*] and Jean Gallier[**]

[*]Department of Mathematics
[**]Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104, USA

August 12, 2006

To Peter A. Cassileth and to the memories of Stanley M.K. Chung and Ralph C. Marcove, physicians and friends all. Being mortal and denied the gift of life, he gives and they gave the next best thing: Time

To Anne, Mia, Philippe and Sylvie

Contents

Preface

A book on "Abstract" or "Modern" Algebra is a commonplace thing in today's mathematical milieu. Even a book for *well-prepared, serious* beginning graduate students who intend to become research mathematicians is not so strange any longer. But, the genesis of this book, which *is* intended for serious, *well-prepared* graduate students, is somewhat strange.

To begin with, it is a reworking of notes for a year long graduate course I gave several years ago–not in itself a strange thing. But, I possess no such notes nor did I ever make any and I never lecture with a written *aide memoir* of any sort. Rather, my method is to work out fully during lecture (at the board) each proof and example. Students will thus see what are the "inner workings" of the subject. Of course, this is pedagogically to their advantage and, furthermore, it slows me down.

Then where did the notes (to be reworked) come from? They were provided by my friend and colleague Jean H. Gallier (of the Computer Science Department at Penn). Determined to augment his mathematical knowledge, he began several years ago to audit some of my graduate courses. "Audit" for him means faithfully attending lectures, doing all the problem assignments, participating in each bi-weekly problem session (where he takes his turn presenting problems), writing excellent notes from my oral presentation *and rendering these notes in LaTeX form.*[1] That this book will appear is, in large measure, his doing. While I have been responsible for its writing, he has on occasion introduced results and/or alternate proofs that have rendered some material more perspicacious from a student's point of view–these have improved the text. He is in every sense a joint author, save that errors are solely my responsibility. There is no way I can thank him adequately here in plain words and I won't try except to say, *Je te remercie vivement, mon ami Jean, pour tout ton travail.*

Others should be thanked as well–in particular the members of the class that attended the course from which the book is formed.[2] By their interest and attention to detail, they kept me on my toes. One particular member of that class deserves special mention: Mathew Cross.[3] Mathew started the index and set the original 115 problems in LaTeX. He lightened our burden by a considerable amount.

The content of the book follows rather closely the oral lectures–with just a few exceptions. These are: In Chapter 3, the section on Integral Dependence is now augmented by proofs of all results, the original lectures had statements only of some of these (due to exigencies of time) and Gallier insisted on a full treatment. In Chapter 4, the sections on Norms and Traces as well as Kummer Theory and Transcendental Extensions are likewise augmented by full proofs. In Chapter 5, there is now more to the section on (co)homological functors and there are full proofs in the last section on the Koszul Complex. Otherwise, the material is just (a smoothed out version of) what was presented. One will have to move fast to present it to students in one year, at least I did.

But the heart of the book is the Problem section. Here, I've attempted to simulate at the beginning graduate level some of the features of real mathematical work. There is a jumbling of the problems *vis a*

[1] One must realize he maintains a full research and teaching schedule, directs Ph.D. students, attends to administrative duties and has a family life in addition to this "auditing"!

[2] The members of the class were: A. Bak, D. Boyarchenko, S. Brooks, M. Campbell, S. Corry, M. Cross, C. Daenzer, C. Devena, J. Gallier, S. Guerra, C. Hoelscher, T. Jaeger, J. Long, S. Mason, T. Zhu.

[3] Mathew spells his name with but one "t"; there is no misprint.

vis subject matter just as in real research one never knows what kind of mathematics will be needed in the solution of a problem. There is no hint of the level of difficulty of a problem (save for the few problems where suggestions are offered), and anyway the notion of difficulty is ill-defined. And, the problems refer to each other, just as in real work one is constantly reminded of past efforts (successful or not). In effect, as suggested in the preface for students, one should begin with the problems and use the text as a means to fill in knowledge as required to do them (as well as to do other problems assigned by an instructor in this course or another course).

This brings me to the text material itself. There is no attempt to be encyclopedic. After all, the material is a faithful copy of what was actually covered in a year and any competent instructor can add material that has been omitted. I regret not covering the Wederburn-Artin Theory of DCC rings, the Brauer Group, and some basic material on group representations. What is covered, however, is to my mind central to the education of any prospective mathematician who aspires to contribute to what is now the mainstream of mathematical endeavor. Also, while there are over 150 problems filling some 55 pages of text (some of the problems are rather long being multi-part), other problems of an instructor's choosing can certainly be assigned. As to the attribution of the origins of these problems, I have assigned names when they are known to me. If no name is assigned, the problem comes from some source in my past (perhaps one of my own teachers or their teachers) and in no way do I claim it as my own. Good problems from all sources are the treasure hoard of practicing mathematicians in their role as passers on of our common heritage.

I refer to the special symbols (DX) and the "curves ahead" road sign (appearing at odd places in the text) in the student preface; no repeat of the explanations I offer there is necessary. If you as instructor are lucky enough to have a class as interested and tough to satisfy as I did, you are lucky indeed and need no further assurance that mathematics will be in good hands in the future. I intend this book to be of service to such individuals as they begin their long climb to mathematical independence and maturity.

Tolda Santa Cotogna
Summer, 2006

For the Student

It may be surprising but the most important part of the book you now hold before you is the very last section–the one labeled "Problems". To learn mathematics one must *do* mathematics. Indeed, the best way to read this book is to turn immediately to the problem section and begin to do the problems. Of course, you will soon reach some unknown terminology or not have enough knowledge to meet the technical demands of a problem and this is where you turn to the text to fill in gaps, see ideas explained and techniques demonstrated. Then you plunge once more back into the problems and repeat the whole process.

The book is designed for serious, well-prepared students who plan on becoming research mathematicians. It presumes you have had previous acquaintance with algebra; in particular you have met the concepts of group, ring, field, vector space, homomorphism, isomorphism, and the elementary theorems about these things. No book on mathematics can be simply read, rather you must recreate the text yourself line by line checking at each stage all details including those omitted. This is slow work and, as you know, mathematics has very high density on the page.

In the text, you will find two special symbols: (DX) and a sign such as one sees on the road warning of dangerous curves ahead. The symbol (DX) stands for "diagnostic exercise", it means some elementary details have been omitted and that supplying them should be easy. However, if supplying them is not easy, then you should go back a page or two as something fundamental has skipped you by. In this way, the sign (DX) is like a medical test: failing it is sure to tell you if something is wrong (no false positives), however, if you pass it (supply the details), something still might be wrong. Just read on and anything wrong will surface later. As for the dangerous curves sign, it precedes counter-examples to naively made conjectures, it warns when things could go wrong if hypotheses are omitted, and generally forces you to slow down in the reading and recreating.

If you use this book in a course or even for self study, I recommend that you tackle the problems in a small group (two to four persons, total). This is because no person has a monopoly on ideas, a good idea or half-idea can germ in any head, and the working out of a problem by a committed group is akin to the actual way much research mathematics is accomplished. In your group, you want constant give and take, and there must be time to think alone so that a real contribution to the group's effort can be made.

The problems are all jumbled up by area and there is no signal given as to a problem's difficulty (exceptions are the few cases where hints or suggestions are given). In real mathematical life, no signs are given that a question being attacked involves a certain small area of mathematical knowledge or is hard or easy; any such sign is gleaned by virtue of experience and that is what you are obtaining by *doing* mathematics in these problems. Moreover, hard and easy are in the eyes of the beholder; they are not universal characteristics of a problem. About all one can say is that if a large number of people find a problem difficult, we may classify it so. However, we shouldn't be surprised when an individual solves it and claims that, "it was not that hard". In any case, guard against confusing mathematical talent either with overall intelligence or with mathematical speed. Some quick people are in fact talented, many are just quick. Don't be discouraged if you find yourself slower than another, the things that really count in doing mathematics (assuming talent) are persistence and courage.

I can think of no better lines to close with than these which come from B. Pasternak's poem entitled "*Night*"[4]

"And maybe in an attic
And under ancient slates
A man sits wakeful working
He thinks and broods and waits."

"He looks upon the planet,
As if the heavenly spheres
Were part of his entrusted
Nocturnal private cares."

"Fight off your sleep: be wakeful,
Work on, keep up your pace,
Keep vigil like the pilot,
Like all the stars in space."

"Work on, work on, creator–
To sleep would be a crime–
Eternity's own hostage,
And prisoner of Time."

Tolda Santa Cotogna
Summer, 2006

[4] From the collection entitled "*When It Clears Up*", 1956. Translated by Lydia Pasternak Slater (the poet's sister).

Chapter 1

Group Theory

1.1 Introduction

Groups are probably the most useful of the structures of algebra; they appear throughout mathematics, physics[1] and chemistry. They almost always occur as "groups of transformations" and that is the way we will use them at first. This allows of tremendous freedom, constrained only by the imagination in finding objects on which to let groups act, or, what is the same, in finding homomorphisms from the group to the "automorphisms" of some object or structure. Then we will look into groups *qua* groups, and here there is a sharp distinction between the finite case and the infinite case. In the finite case, there is a subtle interplay (not yet fully understood) between the order of a group and its structure, whereas in the infinite case"geometric" arguments and applications are more the norm.

1.2 Group Actions and First Applications; The Three Sylow Theorems

We begin by reviewing the notion of group action.

Definition 1.1 Let G be a group and S be a set. We say that *G acts on S (on the left)* (or that *there is a (left) G-action on S*) iff there is a map

$$\begin{array}{ccc} G\prod S & \longrightarrow & S \\ (\sigma,\, s) & \mapsto & \sigma \cdot s \end{array}$$

called the *action*, satisfying the two rules:

(1) $(\forall s \in S)(1 \cdot s = s)$

(2) $(\forall \sigma, \tau \in G)(\forall s \in S)(\sigma \cdot (\tau \cdot s) = (\sigma\tau) \cdot s)$.

Remarks:

(1) For every $\sigma \in G$, the map $s \mapsto \sigma \cdot s$ is a bijection of S to itself. Its inverse is the map $s \mapsto \sigma^{-1} \cdot s$. We let $\mathrm{Aut}(S)$ denote the set of all set theoretic bijections of S.

[1] The word group even occurs in Einstein's first paper [12] on special relativity; it is the only place to my knowledge where that word appears in Einstein's corpus of scientific work.

(2) Write $\theta(\sigma)$ for the element of $\mathrm{Aut}(S)$ given by remark (1), *i.e.*,

$$\theta(\sigma)(s) = \sigma \cdot s.$$

Then, the map $\theta\colon G \to \mathrm{Aut}(S)$ is a homomorphism of groups (where $\mathrm{Aut}(S)$ is a group under composition).

(3) Conversely, a n.a.s.c. that G act on S is that there is a *homomorphism* $\theta\colon G \to \mathrm{Aut}(S)$. (The action gives θ by remarks (1) and (2), and given θ, define the corresponding action by $\sigma \cdot s = \theta(\sigma)(s)$. Check that this is an action (DX).)

Say G acts on S, and for any given s consider

$$\mathrm{St}(s) = \{\sigma \in G \mid \sigma \cdot s = s\},$$

the *stabilizer* of s. It is always a subgroup of G. The set

$$\{t \in S \mid (\exists \sigma \in G)(\sigma \cdot s = t)\}$$

is the *orbit* of s under the action, and it is denoted $\mathrm{O}_G(s)$.

(4) There is a one-to-one correspondence between the elements of the orbit of s and the left cosets of $\mathrm{St}(s)$ in G. Namely, if $H = \mathrm{St}(s)$, there are maps

$$\begin{aligned} \sigma H &\mapsto \sigma \cdot s \\ \sigma \cdot s &\mapsto \sigma H, \end{aligned}$$

for any left coset, σH. The first map is well-defined because if $\sigma H = \tau H$, then $\tau = \sigma h$ for some $h \in H$, and

$$\tau \cdot s = (\sigma h) \cdot s = \sigma \cdot (h \cdot s) = \sigma \cdot s$$

as $h \in \mathrm{St}(s)$. The reader should check that the second map is well-defined (DX).

If G is finite or $(G : \mathrm{St}(s))$ is finite (here, $(G : H)$ denotes the index of the subgroup H in G, *i.e.*, the number of (left) cosets of H in G), then $\mathrm{O}_G(s)$ is a finite set and when G is finite, $\#(\mathrm{O}_G(s))$ divides $\#(G)$.

(5) Say $t \in \mathrm{O}_G(s)$ and $H = \mathrm{St}(s)$. Write $t = \sigma \cdot s$. What is $\mathrm{St}(t)$?

We have $\tau \in \mathrm{St}(t)$ iff $\tau \cdot t = t$ iff $\tau \cdot (\sigma \cdot s) = \sigma \cdot s$ iff $(\sigma^{-1}\tau\sigma) \cdot s = s$ iff $\sigma^{-1}\tau\sigma \in H$ iff $\tau \in \sigma H \sigma^{-1}$. In conclusion, we see that $\mathrm{St}(\sigma \cdot s) = \sigma \mathrm{St}(s) \sigma^{-1}$, a conjugate subgroup of $\mathrm{St}(s)$.

(6) The reader can check that the relation $\sim$ on the set S defined by

$$s \sim t \quad \text{iff} \quad t = \sigma \cdot s \quad \text{for some } \sigma \in G$$

is an equivalence relation on S, and that the equivalence classes of this relation are exactly the distinct orbits $\mathrm{O}_G(s)$. Thus, given two orbits, $\mathrm{O}_G(s)$ and $\mathrm{O}_G(t)$, either $\mathrm{O}_G(s) \cap \mathrm{O}_G(t) = \emptyset$ or $\mathrm{O}_G(s) = \mathrm{O}_G(t)$. As a conclusion,

$$S = \bigcup_{\text{distinct orbits}}^{\cdot} \mathrm{O}_G(s).$$

The *orbit space*, $G \setminus S$, is the quotient set $S/\sim$, *i.e.*, the collection of orbits, each considered as a distinct entity.

Obviously, we can define the notion of right action using a map $S\prod G \longrightarrow G$. It is obvious how to modify conditions (1) and (2) in Definition 1.1.

We now give some examples of group actions.

Example 1.1

(1) *Trivial action.* Let G be any group and S be any set. The action is

$$\sigma \cdot s = s,$$

that is, it leaves every element of S fixed.

(2) Let G be a group and H be a subgroup of G. Consider G as a set, H as a group, and the action $H\prod G \longrightarrow G$ given by

$$(\tau, s) \mapsto \tau \cdot s = \tau s \in G.$$

This action is called *translation.* Observe that

$$\mathrm{St}(s) = \{\tau \in H \mid \tau s = s\} = \{1\},$$

and

$$\begin{aligned} \mathrm{O}_H(s) &= \{t \in G \mid (\exists \sigma \in H)(\sigma \cdot s = t)\} \\ &= \{t \in G \mid (\exists \sigma \in H)(\sigma s = t)\} \\ &= Hs = \text{a right coset of } s. \end{aligned}$$

(3) Let G be a group and H be a subgroup of G. Consider G as a set, H as a group, and the action $H\prod G \longrightarrow G$ given by

$$(\tau, s) \mapsto \tau \cdot s = \tau s \tau^{-1} \in G.$$

This action is called *conjugation.* Note that

$$\begin{aligned} \mathrm{St}(s) &= \{\tau \in H \mid \tau s \tau^{-1} = s\} \\ &= \{\tau \in H \mid \tau s = s\tau\}, \end{aligned}$$

the collection of τ's in H which commute with s. When $H = G$, we see that $\mathrm{St}(s)$ is the *centralizer* of s in G, denoted $Z_G(s)$. For an arbitrary subgroup H of G, we get $\mathrm{St}(s) = Z_G(s) \cap H$. We also have

$$\mathrm{O}_H(s) = \{t \in G \mid (\exists \sigma \in H)(\sigma s \sigma^{-1} = t)\},$$

the *H-conjugacy class* of s, denoted $\mathrm{Cl}_H(s)$. When $H = G$, we get the *conjugacy class* of s, denoted $\mathrm{Cl}(s)$.

(4) Suppose the set S has some structure. Two very important special cases are:

(a) The set S is a vector space over a field. Then, we require $\theta\colon G \to \mathrm{Aut}(S)$ to land in the *linear* automorphisms of S, *i.e.*, in the invertible linear maps. In this case, our action is called a *(linear) representation* of G.

(b) The set S is an abelian group under addition, $+$. Then, we require $\theta\colon G \to \mathrm{Aut}(S)$ to land in the group of group automorphisms of S. Our action makes S into a *G-module*. Observe that in addition to the axioms (1) and (2) of Definition 1.1, a G-module action also satisfies the axiom

$$\sigma \cdot (a + b) = (\sigma \cdot a) + (\sigma \cdot b), \quad \text{for all } \sigma \in G \text{ and all } a, b \in S.$$

Now, assume that G is finite. Observe that the converse of Lagrange's theorem is false; namely, if G has order n and h divides n, then there isn't necessarily a subgroup of order h. Indeed, the group, A_4, of even permutations on four elements, has order 12 and $6 \mid 12$, yet A_4 has *no* subgroup of order 6. In 1872, Sylow (pronounce "Zŏloff") discovered the Sylow existence theorem and the classification theorem, known now as Sylow theorems I & II.

Theorem 1.1 *(Sylow, I) If G is a finite group of order g and p is a given prime number, then whenever $p^\alpha \mid g$ (with $\alpha \geq 0$), there exists a subgroup, H, of G of exact order p^α.*

To prove Theorem 1.1, we need an easy counting lemma. If m is an integer, write $\mathrm{ord}_p(m)$ for the maximal exponent to which p divides m (*i.e.*, $\mathrm{ord}_p(m) = \beta$ for the largest β such that $p^\beta \mid m$). The following simple properties hold (DX):

(1) $\mathrm{ord}_p(mn) = \mathrm{ord}_p(m) + \mathrm{ord}_p(n)$.

(2) $\mathrm{ord}_p(m \pm n) \geq \min\{\mathrm{ord}_p(m), \mathrm{ord}_p(n)\}$,
with equality if $\mathrm{ord}_p(m) = \mathrm{ord}_p(n)$.

(3) By convention, $\mathrm{ord}_p(0) = \infty$.

Lemma 1.2 *(Counting lemma) Let p be a prime, α, m positive integers. Then,*

$$\mathrm{ord}_p\binom{p^\alpha m}{p^\alpha} = \mathrm{ord}_p(m).$$

Proof. We know that

$$\binom{p^\alpha m}{p^\alpha} = \frac{p^\alpha m(p^\alpha m - 1)\cdots(p^\alpha m - (p^\alpha - 1))}{p^\alpha(p^\alpha - 1)\cdots 2 \cdot 1}.$$

Observe that for $0 < i < p^\alpha$, we have (DX)

$$\mathrm{ord}_p(p^\alpha m - i) = \mathrm{ord}_p(p^\alpha - i).$$

Thus,

$$\binom{p^\alpha m}{p^\alpha} = mK, \quad \text{where } K \text{ is prime to } p.$$

Therefore,

$$\mathrm{ord}_p\binom{p^\alpha m}{p^\alpha} = \mathrm{ord}_p(m),$$

as contended. □

Proof of Sylow I. (Wielandt, 1959) If S is any subset of G, let

$$\sigma \cdot S = \{\sigma t \mid t \in S\},$$

and note that $\sigma \cdot S$ is a subset of the same cardinality of that of S. Let

$$\mathcal{S} = \{S \subseteq G \mid \#(S) = p^\alpha\}.$$

Note that in the above definition, S is *any* subset of G, and not necessarily a subgroup of G. Of course,

$$\#(\mathcal{S}) = \binom{p^\alpha m}{p^\alpha}.$$

The group G acts on $\mathcal{S}$ by translation, *i.e., via*, $S \mapsto \sigma \cdot S$.

Claim. There is some $S \in \mathcal{S}$ so that

$$\mathrm{ord}_p(\#(\mathrm{O}_G(S))) \leq \mathrm{ord}_p(m).$$

If not, then for all $S \in \mathcal{S}$, we have $\mathrm{ord}_p(\#(\mathrm{O}_G(S))) > \mathrm{ord}_p(m)$. But we know that $\mathcal{S}$ can be written as a disjoint union of G-orbits,

$$\mathcal{S} = \dot{\bigcup_{\text{distinct orbits}}} \mathrm{O}_G(S).$$

So,

$$\#(\mathcal{S}) = \sum_{\text{distinct orbits}} \#(\mathrm{O}_G(S)).$$

Consequently,

$$\mathrm{ord}_p(\#(\mathcal{S})) \geq \min\{\mathrm{ord}_p(\#(\mathrm{O}_G(S)))\} > \mathrm{ord}_p(m).$$

But

$$\mathrm{ord}_p(\#(\mathcal{S})) = \mathrm{ord}_p\binom{p^\alpha m}{p^\alpha},$$

contradicting Lemma 1.2. This proves the claim.

Now, pick some $S \in \mathcal{S}$ so that $\mathrm{ord}_p(\#(\mathrm{O}_G(S))) \leq \mathrm{ord}_p(m)$. Let H be the stabilizer of S. We know that

(a) $\#(\mathrm{O}_G(S)) = (G : \mathrm{St}(S)) = (G : H)$.

(b) $p^\alpha m = \#(G) = \#(H)\#(\mathrm{O}_G(S))$.

From (b), applying the ord function, we get

$$\alpha + \mathrm{ord}_p(m) = \mathrm{ord}_p(\#(H)) + \mathrm{ord}_p(\#(\mathrm{O}_G(S))) \leq \mathrm{ord}_p(\#(H)) + \mathrm{ord}_p(m).$$

So, $\alpha \leq \mathrm{ord}_p(\#(H))$ and then, p^α divides $\#(H)$, and thus, $\#(H) \geq p^\alpha$. Now, H takes S elementwise to itself by translation, and for every $s \in S$,

$$\mathrm{St}(s) = \{\sigma \in H \mid \sigma s = s\} = \{1\}.$$

Therefore, $\#(H) = \#(\mathrm{O}_H(s))$ for every $s \in S$, and yet every orbit is contained in S. Thus,

$$\#(\mathrm{O}_H(s)) \leq \#(S) = p^\alpha,$$

from which we deduce that $\#(H) \leq p^\alpha$. We conclude that $\#(H) = p^\alpha$, and H is the required subgroup. □

Corollary 1.3 *(Original Sylow I) If p^β is the maximal power of p to divide $\#(G)$ and p is a prime number, then G possesses a subgroup of order p^β.*

The subgroups of maximal p-power order arising in Corollary 1.3 are called the *p-Sylow subgroups of G* (there can be more than one).

Corollary 1.4 *(Cauchy, 1840) Say G is a finite group and $p \mid \#(G)$, where p is a prime number. Then, there is some σ of order p in G.*

Nomenclature: A *p-group* is a finite group whose order is a power of the prime number p.

Corollary 1.5 *Say G is a p-group, with $\#(G) = p^r$. Then G possesses a descending chain*

$$G = G_0 > G_1 \cdots > G_{r-1} > G_r = \{1\},$$

so that $(G_i : G_{i+1}) = p$ for all i with $0 \leq i \leq r-1$. Hence, $\#(G_i) = p^{r-i}$.

Proof. By Sylow I, a subgroup G_1 of order p^{r-1} exists. An induction finishes the proof. □

Remark: It is not clear that G_{i+1} is normal in G_i. In fact, this is true, but it takes more work (see Proposition 1.10).

To prove Sylow II, we need the local embedding lemma. In order to state this lemma, we need to recall the concept of a normalizer. If $\mathcal{S}$ denotes the collection of all *subsets* of G, then G acts on $\mathcal{S}$ by conjugation: $S \mapsto \sigma S \sigma^{-1}$. This action preserves cardinality. For every $S \in \mathcal{S}$, we have

$$\mathrm{St}(S) = \{\sigma \in G \mid \sigma S \sigma^{-1} = S\}.$$

The group $\mathrm{St}(S)$ is called the *normalizer* of S in G, and it is denoted $N_G(S)$. If S is a subgroup of G, then S is normal in $N_G(S)$ (denoted $S \lhd N_G(S)$), and $N_G(S)$ is the biggest subgroup in which S is normal (DX).

The “philosophy” behind the local embedding lemma is that if P is any subgroup of a group G, then $N_G(P)$ is a “local neighborhood” of P in which P perhaps behaves nicely. We recall the following proposition which is used for proving Lemma 1.7.

Proposition 1.6 *Given a group G, for any two subgroups S and P, if $S \subseteq N_G(P)$, then $PS = SP$ is the subgroup of $N_G(P)$ generated by $S \cup P$, the subgroup P is normal in SP and $(SP)/P \cong S/(S \cap P)$.*

Proof. Since $S \subseteq N_G(P)$, we have $\sigma P \sigma^{-1} = P$ for all $\sigma \in S$, and thus, it clear that $SP = PS$. We have $\sigma\tau\sigma^{-1} \in P$ for all $\sigma \in S$ and all $\tau \in P$, and thus, for all $a, c \in S$ and all $b, d \in P$, we have

$$\begin{aligned} (ab)(cd) &= (ac)(c^{-1}bc)d \\ b^{-1}a^{-1} &= a^{-1}(ab^{-1}a^{-1}). \end{aligned}$$

The above identities prove that SP is a group. Since S and P contain the identity, this group contains S and P, and clearly any subgroup containing S and P contains SP. Therefore, SP is indeed the subgroup of $N_G(P)$ generated by $S \cup P$ and it is clear that P is normal in SP. Now, look at the composition φ of the injection $S \longrightarrow SP$ with the quotient map $SP \longrightarrow (SP)/P$. It is surjective, and $\varphi(\sigma) = \sigma P$ for every $\sigma \in S$. Thus, $\sigma \in \mathrm{Ker}\ \varphi$ iff $\sigma \in S \cap P$, and so $\mathrm{Ker}\ \varphi = S \cap P$, and the first isomorphism theorem yields
$(SP)/P \cong S/(S \cap P)$. □

After this short digression, we return to the main stream of the lecture.

Lemma 1.7 *(Local embedding lemma) Suppose that P is a p-Sylow subgroup of G. Then for every $\sigma \in N_G(P)$, if σ has p-power order then $\sigma \in P$. In particular, if H is a p-subgroup of $N_G(P)$, then $H \subseteq P$ and P is unique in $N_G(P)$.*

Proof. Let S be any p-subgroup of $N_G(P)$. Look at the group, H, generated by S and P in $N_G(P)$, denoted $\mathrm{Gp}\{S, P\}$. Since P is normal in $N_G(P)$, from Proposition 1.6, we have $H = SP = PS$, and $H/P = (SP)/P \cong S/(S \cap P)$. Thus,

$$(H : P) = (S : S \cap P),$$

and $(S : S \cap P)$ is a p-power, since S is a p-group. On the other hand, $(S : S \cap P)$ is prime to p, as $(G : P) = (G : H)(H : P)$ and $(G : P)$ is prime to p by definition of P. So, we must have $(H : P) = (S : S \cap P) = 1$, which implies that $H = P$. Thus, $S = S \cap P$, and $S \subseteq P$. We finish the proof by letting S be the cyclic p-group generated by σ. □

Theorem 1.8 *(Sylow II) If G is a finite group, write $\mathcal{S}\mathrm{yl}_p(G)$ for the collection of all p-Sylow subgroups of G, and $\mathcal{P}$ for the collection of* **all** *the p-subgroups of G, where p is a prime number. Then, the following hold:*

(1) $\mathrm{syl}_p(G) = \#(\mathcal{S}\mathrm{yl}_p(G)) \equiv 1\ (mod\, p)$.

(2) For all $S \in \mathcal{P}(G)$ and all $P \in \mathcal{S}\mathrm{yl}_p(G)$, there is some $\sigma \in G$ so that $S \subseteq \sigma P \sigma^{-1}$. In particular, any two p-Sylow subgroups of G are conjugate in G.

(3) $\mathrm{syl}_p(G)$ *divides $\#(G)$; in fact,* $\mathrm{syl}_p(G)$ *divides the prime to p part of $\#(G)$.*

Proof. (1) The group G acts by conjugation on $\mathcal{S}\mathrm{yl}(G)$ (drop the subscript p in the course of this proof). So

$$\mathcal{S}\mathrm{yl}(G) = \dot{\bigcup_{\text{distinct orbits}}} \mathrm{O}_G(P).$$

Any $S \in \mathcal{P}(G)$ also acts by conjugation on $\mathcal{S}\mathrm{yl}(G)$, and so

$$\mathcal{S}\mathrm{yl}(G) = \dot{\bigcup_{\text{distinct orbits}}} \mathrm{O}_S(P).$$

What is $\mathrm{St}(P)$? We have

$$\mathrm{St}(P) = \{\sigma \in S \mid \sigma P \sigma^{-1} = P\} = S \cap N_G(P).$$

But S has p-power order, so $S \cap N_G(P)$ is a p-subgroup of $N_G(P)$. The embedding lemma implies that $S \cap N_G(P) \subseteq P$, from which we deduce that $S \cap N_G(P) = S \cap P$. So,

$$\#(\mathrm{O}_S(P)) = (S : S \cap P).$$

Now, take for S one of the p-Sylow subgroups, say P. Then, $\#(\mathrm{O}_P(Q)) = (P : P \cap Q)$. If $Q \neq P$, then $P \cap Q < P$, and so, $(P : P \cap Q)$ is a nontrivial p-power (*i.e*, not equal to 1). If $P = Q$, then $(P : P \cap Q) = 1$. Therefore, in the orbit decomposition

$$\mathcal{S}\mathrm{yl}(G) = \dot{\bigcup_{\substack{\text{distinct orbits} \\ Q \in \mathcal{S}\mathrm{yl}(G)}}} \mathrm{O}_P(Q),$$

one orbit has cardinality 1, the rest having nontrivial p-power cardinalities. We conclude that

$$\#(\mathcal{S}\mathrm{yl}(G)) = 1 + \sum p\text{-powers},$$

and $\mathrm{syl}_p(G) = \#(\mathcal{S}\mathrm{yl}_p(G)) \equiv 1 \pmod{p}$, as claimed.

(2) Let $S \in \mathcal{P}(G)$ and look at $\mathrm{O}_G(P)$ where $P \in \mathcal{S}\mathrm{yl}(G)$. The subgroup S acts by conjugation on $\mathrm{O}_G(P)$. So, we have

$$\mathrm{O}_G(P) = \dot{\bigcup_{\substack{\text{distinct orbits} \\ Q \in \mathrm{O}_G(P)}}} \mathrm{O}_S(Q). \qquad (*)$$

If $Q \in \mathrm{O}_G(P)$, then consider the stabilizer of Q *in* S,

$$\mathrm{St}(Q) = \{\sigma \in S \mid \sigma Q \sigma^{-1} = Q\} = S \cap N_G(Q).$$

As before, by the embedding lemma, $S \cap N_G(Q) = S \cap Q$. Then, $\#(\mathrm{O}_S(Q)) = (S : S \cap Q)$. Take $S = P$ itself. If $Q = P$, then $(P : P \cap P) = 1$ and $\#(\mathrm{O}_P(P)) = 1$. On the other hand, if $P \neq Q$, then $(P : P \cap Q)$ is a nontrivial p-power. Thus, as before, using $(*)$, we deduce that

$$\#(\mathrm{O}_G(P)) \equiv 1 \pmod{p}.$$

Assume that (2) is false. Then, there exist some S and some P such that $S \not\subseteq \sigma P \sigma^{-1}$ for *any* $\sigma \in G$. Let this S act on $\mathrm{O}_G(P)$, for this P. But we have

$$\#(\mathrm{O}_G(P)) = \sum_{\substack{\text{distinct orbits} \\ Q \in \mathrm{O}_G(P)}} \#(\mathrm{O}_S(Q)), \qquad (**)$$

and $\#(\mathrm{O}_S(Q)) = (S\colon S\cap Q)$ where Q is a conjugate of P, so that $S \not\subseteq Q$, and therefore $(S : S\cap Q)$ is a nontrivial p-power. Then, $(**)$ implies

$$\#(\mathrm{O}_G(P)) \equiv 0 \pmod p,$$

a contradiction. Thus, neither S nor P exist and (2) holds.

(3) By (2), $\mathcal{S}\mathrm{yl}(G) = \mathrm{O}_G(P)$, for some fixed P. But the size of an orbit divides the order of the group. The rest is clear. □

Theorem 1.9 *(Sylow III) If G is a finite group and P is a p-Sylow subgroup of G, then $N_G(N_G(P)) = N_G(P)$.*

Proof. Let $T = N_G(N_G(P))$ and $S = N_G(P)$, so that $T = N_G(S)$ and $S \lhd T$.

Claim. For every $\sigma \in T$, if σ has p-power order then $\sigma \in P$.

The order of T/S is $(T : S)$. But

$$(G : P) = (G : T)(T : S)(S : P)$$

and $(G : P)$ is prime to p by definition of P. So, $(T : S)$ is prime to p. Consider $\overline{\sigma}$, the image of σ in T/S. The element $\overline{\sigma}$ has p-power order, yet $\#(T/S)$ is prime to p. Thus, $\overline{\sigma} = 1$, and so, $\sigma \in S$. The local embedding lemma yields $\sigma \in P$. Therefore, if H is a p-subgroup of T, we have $H \subseteq P$. Thus, any p-Sylow subgroup, H, of T is contained in P; but since H has maximal p-size, $H = P$. This implies that T has a single p-Sylow subgroup, namely P. By Sylow II, the group P is normal in T and so $T \subseteq N_G(P) = S$. Yet, $S \subseteq T$, trivially, and $S = T$. □

Remark: A p-Sylow subgroup is unique iff it is normal in G.

Definition 1.2 A group, G, is *simple* if and only if it possesses no nontrivial normal subgroups ($\{1\}$ and G itself are the two trivial normal subgroups).

Example 1.2

(1) Assume that G is a group of order pq, with p and q prime and $p < q$. Look at the q-Sylow subgroups. Write $\mathrm{syl}(q)$ for the number of q-Sylow subgroups of G. We know that

$$\mathrm{syl}(q) \equiv 1 \pmod q \quad \text{and} \quad \mathrm{syl}(q) \mid p.$$

This implies that $\mathrm{syl}(q) = 1, p$. But $p < q$, so that $p \equiv p \pmod q$, and the only possibility is $\mathrm{syl}(q) = 1$. Therefore, the unique q-Sylow subgroup is normal, and G is *not* simple.

(2) Assume that G is a group of order pqr, with p, q, r prime and $p < q < r$. Look at the r-Sylow subgroups. We must have

$$\mathrm{syl}(r) \equiv 1 \pmod r \quad \text{and} \quad \mathrm{syl}(r) \mid pq.$$

This implies that $\mathrm{syl}(r) = 1, p, q, pq$. Since $p < r$ and $q < r$, as above, p and q are ruled out, and $\mathrm{syl}(r) = 1, pq$.

Suppose that $\mathrm{syl}(r) = pq$. We see immediately that $r < pq$. Now, each r-Sylow subgroup is isomorphic to $\mathbb{Z}/r\mathbb{Z}$ (cyclic of prime order), and any two distinct such subgroups intersect in the identity (since, otherwise, they would coincide). Hence, there are $pq(r-1)$ elements of order r. We shall now show that if $\mathrm{syl}(r) = pq$, then $\mathrm{syl}(q) = 1$. Assume that $\mathrm{syl}(r) = pq$ and look at the q-Sylow subgroups of G. We have

$$\mathrm{syl}(q) \equiv 1 \pmod q \quad \text{and} \quad \mathrm{syl}(q) \mid pr.$$

This implies that $\mathrm{syl}(q) = 1, p, r, pr$ and, as before, p is ruled out since $p < q$. So, $\mathrm{syl}(q) = 1, r, pr$. Suppose that $\mathrm{syl}(q) = r$ or $\mathrm{syl}(q) = pr$, and call it x. Reasoning as before but now on the q-Sylow subgroups, we see that there are $x(q-1)$ elements of order q. Now, $q-1 \geq p$ and $x \geq r$. Thus, there are at least rp elements of order q. But $r > q$, so there are more than pq elements of order q. Now, since there are $pq(r-1)$ elements of order r and more than pq elements of order q, there are more than

$$pq(r-1) + pq = pqr - pq + pq = pqr$$

elements in G, a contradiction. So, either the r-Sylow subgroup is normal in G (which is the case when $r > pq$) or the q-Sylow subgroup is normal in G. In either case, G is *not* simple.

Cases (1) and (2) have the following generalizations:

(a) Frobenius (1890's) showed that if $\#(G) = p_1p_2\cdots p_t$, a product of *distinct* primes, then G is *not* simple. The proof uses group representations and characters.

(b) Burnside (1901) proved the "p^aq^b-theorem": If $\#(G) = p^aq^b$, where p, q are distinct primes and $a, b \in \mathbb{N}$, then G is *not* simple. There are three known proofs, all hard, and all but one use group representations.

Obvious generalizations of (a) and (b) are *false*. The easiest case is $\#(G) = 2^2 \cdot 3 \cdot 5 = 60$. Indeed, the alternating group, A_5, is simple. After proving (b), Burnside conjectured (*circa* 1902) that every nonabelian group of odd order is *not* simple. This conjecture was proved in 1961 by W. Feit and J. Thompson. The proof is **very** hard, and very long (over 200 pages).

A piece of the proof of (a) and (b) is the following proposition:

Proposition 1.10 *If G is a finite group and p is the smallest prime number which divides the order of G, then any subgroup, H, of index p is automatically normal in G.*

Proof. Take H so that $(G : H) = p$. Consider the set $\mathcal{S} = \{H_1 = H, H_2, \ldots, H_p\}$ of cosets of H in G. The group G acts on $\mathcal{S}$ by translation,

$$\sigma \cdot H_j = \sigma H_j = H_l, \quad \text{for some } l, \text{ with } 1 \leq l \leq p.$$

This action is nontrivial, that is, we get a nontrivial homomorphism $\theta\colon G \to \mathfrak{S}_p$ (where $\mathfrak{S}_p \cong \mathrm{Aut}(\mathcal{S})$ is the group of permutations on p elements), and $\mathrm{Im}\ \theta \neq \{1\}$. We shall prove that $H = \mathrm{Ker}\ \theta$, which yields $H \lhd G$.

Observe that $\#(G) = \#(\mathrm{Ker}\ \theta) \cdot \#(\mathrm{Im}\ \theta)$. We must have

(1) $\#(\mathrm{Im}\ \theta) \mid p!$

(2) $\#(\mathrm{Im}\ \theta) \mid \#(G)$.

But $\#(G) = p^\alpha K$, where K contains primes greater than p. Therefore, $\#(\mathrm{Im}\ \theta) = p^aJ$, where $J = 1$ or J contains primes greater than p. If $J \neq 1$, then J contains some prime $q > p$, and since $p^\alpha J$ divides $p! = p(p-1)\cdots 2 \cdot 1$, the prime q must divide $p!$. Since q is prime, q must divide one of the terms in $p!$, which is impossible, since $q > p$. We conclude that $J = 1$. Now, $a \geq 1$ since $\mathrm{Im}\ \theta$ is nontrivial. If $a \geq 2$, since $p^{a-1} \mid (p-1)\cdots 2 \cdot 1$, the prime p should divide $p - j$, for some j with $1 \leq j \leq p-1$. However, this is impossible, and so, $a = 1$. Therefore, $\#(\mathrm{Im}\ \theta) = p$ and $(G : \mathrm{Ker}\ \theta) = p$. Note that $\sigma \in \mathrm{Ker}\ \theta$ iff σ acts trivially on $\mathcal{S}$ iff $\sigma\tau H = \tau H$ iff $\tau^{-1}H\tau = H$ iff $\tau^{-1}\sigma\tau \in H$ for all τ iff $\sigma \in \tau H\tau^{-1}$ for all $\tau \notin H$ iff

$$\sigma \in \bigcap_{\tau \in G} \tau H \tau^{-1}.$$

We deduce that

$$\mathrm{Ker}\ \theta = \bigcap_{\tau \in G} \tau H \tau^{-1} \subseteq H.$$

As $(G : \text{Ker } \theta) = p = (G : H)$ and $\text{Ker } \theta \subseteq H$, we get $H = \text{Ker } \theta$, and H is indeed normal in G. $\square$

Note that we can now improve Corollary 1.5 as follows: If G is a p-group with $\#(G) = p^r$, then there is a descending chain of subgroups

$$G = G_0 > G_1 > \cdots > G_r = \{1\},$$

where each G_{j+1} is normal in G_j and each quotient G_{j+1}/G_j is simple; so, $G_{j+1}/G_j = \mathbb{Z}/p\mathbb{Z}$, a cyclic group of order p.

Definition 1.3 A *composition series* for a group G is a chain of subgroups

$$G = G_0 > G_1 > \cdots > G_t = \{1\}$$

in which each subgroup G_{j+1} is maximal, normal in G_j. The factor groups G/G_1, $G_1/G_2, \ldots$, $G_{t-1}/G_t = G_{t-1}$ are called the *composition factors* of the given composition series and each one is a simple group.

Remark: Every finite group possesses a composition series (DX).

Not every group possesses maximal subgroups, even maximal normal subgroups (such groups must be infinite).

However, finitely generated groups do possess maximal subgroups, but because such groups can be infinite, the proof requires a form of transfinite induction known as Zorn's lemma. Since this lemma is an important tool, we briefly digress to state the lemma and illustrate how it is used.

Recall that a *partially ordered set* or *poset* is a pair, $(S, \leq)$, where S is a set and $\leq$ is a *partial order* on S, which means that $\leq$ is a binary relation on S satisfying the properties: For all $a, b, c \in S$, we have:

(1) $a \leq a$ (reflexivity)

(2) If $a \leq b$ and $b \leq c$, then $a \leq c$ (transitivity)

(3) If $a \leq b$ and $b \leq a$, then $a = b$. (antisymmetry)

Observe that given $a, b \in S$, it may happen that neither $a \leq b$ nor $b \leq a$. A *chain*, C, in S is a linearly ordered subset of S (which means that for all $a, b \in C$, either $a \leq b$ or $b \leq a$). The empty set is considered a chain. An element, $b \in S$, is an *upper bound* of C (resp. a *lower bound* of C) if $a \leq b$ for all $a \in C$ (resp. $b \leq a$ for all $a \in C$). Note that an upper bound of C (resp. a lower bound of C) *need not* belong to C. We say that $C \subseteq S$ is *bounded above* if it possesses some upper bound (in S) (resp. *bounded below* if it possesses some lower bound (in S)). The notion of least upper bound (resp. greatest lower bound) of a chain is clear as is the notion of least or greatest element of a chain. These need not exist. A set, S, which is a chain, is *well ordered* iff every nonempty subset of S has a least element.

Remark: Obviously, the notions of upper bound (resp. lower bound), maximal (resp. minimal) element, greatest (resp. smallest) element, all make sense for arbitrary subsets of a poset, and not just for chains. Some books define a well ordered set to be a poset so that every nonempty subset of S has a least element. Thus, it is not required that S be a chain, but it is required that *every* nonempty subset have a least element, not just chains. It follows that a well ordered set (under this new definition) is necessarily a chain. Indeed, for any two elements $a, b \in S$, the subset $\{a, b\}$ must have a smallest element, so, either $a \leq b$ or $b \leq a$.

Hausdorff maximal principle: Every nonempty poset possesses a maximal chain.

From set theory, it is known that Hausdorff's maximal principle is equivalent to the axiom of choice, which is also equivalent to Zermelo's well ordering principle (every nonempty subset can be well ordered).

We say that a poset is *inductive* iff every nonempty chain possesses a least upper bound.

Zorn's lemma: Each inductive poset possesses a maximal element.

Proof. By Hausdorff. □

Remark: Some books define a poset to be inductive iff every nonempty chain is bounded above. Zorn's lemma still holds under this slightly weaker assumption. In practice, this makes little difference, because when proving that a chain is bounded above, one usually shows that this chain has a least upper bound.

Here are two illustrations of the use of Zorn's lemma.

Theorem 1.11 *Every finitely generated group, G, possesses a maximal subgroup.*

Proof. Consider the set, $\mathcal{S}$, of all proper subgroups, H, of G. Partially order $\mathcal{S}$ by inclusion (ı.e., $H \leq K$ iff $H \subseteq K$). Let $\{H_\alpha\}$ be a chain in $\mathcal{S}$. If $H = \bigcup_\alpha H_\alpha$, we see that H is a group and that it is the least upper bound of $\{H_\alpha\}$. We must show that $H \neq G$. If $H = G$, then as G is finitely generated, $H = G = \mathrm{Gp}\{\sigma_1, \ldots, \sigma_t\}$, with $\sigma_i \in H$ for $i = 1, \ldots, t$. This means that, for each i, there is some α_i so that $\sigma_i \in H_{\alpha_i}$. Since $\{H_\alpha\}$ is a chain, there is some s so that $H_{\alpha_j} \subseteq H_{\alpha_s}$ for $j = 1, \ldots, t$. Thus, $\sigma_1, \ldots, \sigma_t \in H_{\alpha_s}$, and so, $H_{\alpha_s} = G$, contradicting the fact that $H_{\alpha_s} \neq G$. Therefore, $\mathcal{S}$ is inductive, and consequently, by Zorn's lemma, it possesses a maximal element. Such an element is a maximal subgroup of G. □

As a second illustration of Zorn's lemma, we prove that every vector space has a Hamel basis. Given a vector space, V, over a field, k, a *Hamel basis* of V is a family, $\{e_\alpha\}_{\alpha\in\Lambda}$, so that:

(1) For every $v \in V$, there exists a finite subset of Λ, say I, and some elements of k for these α's in I, say c_α, so that
$$v = \sum_{\alpha\in I} c_\alpha e_\alpha.$$

(2) The e_α's are linearly independent, *i.e.*, given any finite subset I of Λ, if $\sum_{\alpha\in I} c_\alpha e_\alpha = 0$, then $c_\alpha = 0$, for all $\alpha \in I$.

Theorem 1.12 *Every vector space, V, possesses a Hamel basis.*

Proof. Let $\mathcal{S}^*$ be the collection of all subspaces, W, of V which possess a Hamel basis, together with a choice of a basis. Write $(W, \{e_\alpha\})$ for any element of $\mathcal{S}^*$. The collection, $\mathcal{S}^*$, is nonempty, since finitely dimensional vector spaces have bases. Partially order $\mathcal{S}^*$ by $(W, \{e_\alpha\}) \leq (\widetilde{W}, \{f_\beta\})$ iff

(a) $W \subseteq \widetilde{W}$ and

(b) $\{e_\alpha\} \subseteq \{f_\beta\}$, which means that the basis $\{f_\beta\}$ extends the basis $\{e_\alpha\}$.

We claim that $\mathcal{S}^*$ is inductive.

Given a chain, $\{W^{(\lambda)}, \{e_\alpha^{(\lambda)}\}\}$, in $\mathcal{S}^*$, take
$$W = \bigcup_\lambda W^{(\lambda)} \quad \text{and} \quad \{e_\gamma\} = \bigcup_\lambda \{e_\alpha^{(\lambda)}\} \subseteq W.$$

The reader should check that $\{e_\gamma\}$ is a basis for W (DX); therefore, $(W, \{e_\gamma\})$ is the least upper bound of our chain. By Zorn's lemma, there exists a maximal element of $\mathcal{S}^*$, call it $(W_0, \{e_\gamma\})$. We need to show that $W_0 = V$. If not, there is some $v \in V$ with $v \notin W_0$. Consider the subspace
$$Z = W_0 \amalg kv = \{w + \xi v \mid w \in W_0,\, \xi \in k\}.$$

The subspace, Z, strictly contains W_0 and $\{e_\gamma\} \cup \{v\}$ is a Hamel basis for Z (DX). However, this contradicts the maximality of W_0. Therefore, $W_0 = V$. □

Corollary 1.13 *If W is a subspace of V and $\{e_\alpha\}$ is a Hamel basis for W, then there exists a Hamel basis of V extending $\{e_\alpha\}$.*

Application: The field, $\mathbb{R}$, is a vector space over $\mathbb{Q}$, and $1 \in \mathbb{Q}$ is a Hamel basis for $\mathbb{Q}$. We can extend this basis of $\mathbb{Q}$ to a Hamel basis for $\mathbb{R}$ (over $\mathbb{Q}$), call it $\{e_\alpha\}_{\alpha\in\Lambda}$, and say, $e_0 = 1$; then, $\mathbb{R}/\mathbb{Q}$ is a vector space (over $\mathbb{Q}$) spanned by the e_α other than e_0. So, we have

$$\mathbb{R}/\mathbb{Q} \cong \coprod_{\alpha\in\Lambda,\alpha\neq 0} \mathbb{Q}.$$

1.3 Elementary Theory of p-Groups

Recall that for a group G, the *center* of G, denoted $Z(G)$, is given by

$$Z(G) = \{\sigma \in G \mid (\forall \tau \in G)(\sigma\tau = \tau\sigma)\}.$$

We write $[\sigma, \tau]$ for the element $\sigma\tau\sigma^{-1}\tau^{-1}$, called the *commutator of* σ *and* τ. Observe that $[\tau, \sigma] = [\sigma, \tau]^{-1}$. Also,

$$Z(G) = \{\sigma \in G \mid (\forall \tau \in G)([\sigma, \tau] = 1)\}$$

and $Z(G)$ is the centralizer of G under conjugation.

Let G act on itself by conjugation. When do we have $\mathrm{O}_G(\sigma) = \{\sigma\}$? This happens when

$$(\forall \tau \in G)(\tau\sigma\tau^{-1} = \sigma) \quad \text{i.e.} \quad (\forall \tau \in G)(\tau\sigma\tau^{-1}\sigma^{-1} = [\tau, \sigma] = 1).$$

Thus, $\sigma \in Z(G)$ iff $\mathrm{O}_G(\sigma) = \{\sigma\}$.

Remark: Obviously,

$$Z(G) = \bigcap_{\sigma \in G} Z_G(\sigma).$$

Moreover, it is obvious that $\sigma \in Z_G(\sigma)$ for every $\sigma \in G$. Thus, for every $\sigma \notin Z(G)$, we have $Z(G) < Z_G(\sigma)$ (obviously, $Z_G(\sigma) = G$ if $\sigma \in Z(G)$.) Therefore, if G is nonabelian, then $Z(G) < Z_G(\sigma)$ for all $\sigma \in G$.

Proposition 1.14 *The center, $Z(G)$, of a p-group, G, is nontrivial.*

Proof. If we let G act on itself by conjugation, we know that G is the disjoint union of distinct orbits, and since $\mathrm{O}_G(\sigma)$ is the conjugacy class of σ and $\sigma \in Z(G)$ iff $\mathrm{O}_G(\sigma) = \{\sigma\}$, we get

$$G = Z(G) \,\dot\cup \dot{\bigcup_{\substack{\text{distinct orbits} \\ \tau \notin Z(G)}}} \mathrm{O}_G(\tau).$$

Consequently, using the fact that $\#(\mathrm{O}_G(\tau)) = (G : \mathrm{St}(\tau))$, we get

$$\#(G) = \#(Z(G)) + \sum_{\substack{\text{distinct orbits} \\ \tau \notin Z(G)}} (G : \mathrm{St}(\tau)). \tag{$*$}$$

But $\#(G) = p^r$, so that each term $(G : \mathrm{St}(\tau))$ for $\tau \notin Z(G)$ is a nontrivial p-power. So, in $(*)$, all terms must be divisible by p. Therefore, $p \mid \#(Z(G))$. □

Note that $Z(G)$ is normal in G. Thus, $G/Z(G)$ is a p-group of strictly smaller order, providing a basis for induction proofs.

We make the following provisional definition (due to E. Galois, 1832). A finite group, G, is *solvable* iff it possesses a composition series all of whose factors are abelian, or equivalently iff it possesses a composition series all of whose factors are cyclic of prime order.

We have shown that a p-group is solvable.

Remark: The above definition is provisional because it only works for finite group (c.f. Definition 1.7), but the concept of a solvable group can be defined for an arbitrary group.

Corollary 1.15 *Every p-group of order less than or equal to p^2 is abelian.*

Proof. Since $\#(G) = 1, p, p^2$ and G is obviously abelian in the first two cases, we may assume that $\#(G) = p^2$. We know that $Z(G)$ is non-trivial and we must prove that $Z(G) = G$. If $Z(G) < G$, then there is some $\sigma \in G$ so that $\sigma \notin Z(G)$. Clearly, $Z(G) \subseteq Z_G(\sigma)$ (where $Z_G(\sigma)$ denotes the centralizer of σ in G). But $\sigma \in Z_G(\sigma)$ implies that $(Z_G(\sigma) : Z(G)) \geq p$ and since $Z(G)$ is nontrivial, we must have $Z_G(\sigma) = G$. So, $\sigma \in Z(G)$, a contradiction. □

We now consider a nice property possessed by p-groups called *property (N)*. If G is any group, *G has property (N)* iff for every proper subgroup, H, of G, the group H is a proper subgroup of $N_G(H)$.

Remark: An abelian group has (N). Indeed, every subgroup of an abelian group is normal, and so, $N_G(H) = G$.

Proposition 1.16 *Every p-group has (N).*

Proof. We proceed by induction on $\#(G) = p^r$. Corollary 1.15 takes care of the base case of the induction. Next, let $\#(G) = p^{r+1}$ and assume that the induction hypothesis holds up to r. We know that $Z(G)$ is nontrivial, and so $\#(G/Z(G)) \leq p^r$. Thus, $G/Z(G)$ has (N). Pick H, any proper subgroup of G. Of course, $Z(G) \subseteq N_G(H)$, and we may assume that $Z(G) \subseteq H$ (since, otherwise, it is clear that $H < N_G(H)$). By the second homomorphism theorem, the question: $H < N_G(H)$? is reduced to the question: $\overline{H} < \overline{N_G(H)}$?, where the bar means pass to $G/Z(G)$. But in this case, as $Z(G) \subseteq H$, we see that (DX)

$$\overline{N_G(H)} = N_{\overline{G}}(\overline{H}),$$

and we just remarked that $\overline{G} = G/Z(G)$ has (N). Therefore, $N_{\overline{G}}(\overline{H}) > \overline{H}$, and so, $\overline{N_G(H)} > \overline{H}$, as desired. □

Groups that have property (N) tend to have good properties. Here are a few of them.

Proposition 1.17 *Say G is a finite group having (N), then each of its p-Sylow subgroups is unique and normal in G. Every maximal subgroup of G is also normal and has prime index.*

Proof. Look at P, a p-Sylow subgroup of G. Now, if $N_G(P) \neq G$, then by (N), we have $N_G(N_G(P)) > N_G(P)$, a contradiction to Sylow III. Thus, $N_G(P) = G$ and so, $P \lhd G$. Next, let H be a maximal subgroup. By (N), we have $N_G(H) > H$, yet H is maximal, so $N_G(H) = G$, and $H \lhd G$. It follows that G/H is a group with no nontrivial subgroup. But then, G/H is cyclic of prime order. □

Proposition 1.18 *Say G is a finite group and suppose that*

(a) $g = \#(G) = p_1^{a_1} \cdots p_t^{a_t}$ (where the p_i's are distinct primes)

(b) G has (N).

Write P_j for the p_j-Sylow subgroup of G. Then, the map

$$P_1 \coprod \cdots \coprod P_t \xrightarrow{\varphi} G$$

via $\varphi(\sigma_1, \ldots, \sigma_t) = \sigma_1 \cdots \sigma_t$ is an isomorphism of groups. Hence, G is isomorphic to a product of p-groups.

The proof depends on the following lemma:

Lemma 1.19 *Let G be a group and let H and K be normal subgroups of G. If $H \cap K = \{1\}$, then every element of H commutes with every element of K. Suppose that σ and τ are commuting elements in G, with orders r and s respectively. If r and s are relatively prime then the order of $\sigma\tau$ is rs.*

Proof. Look at $[\sigma, \tau]$, where $\sigma \in H$ and $\tau \in K$. We have

$$[\sigma, \tau] = \sigma\tau\sigma^{-1}\tau^{-1} = (\sigma\tau\sigma^{-1})\tau^{-1} = \sigma(\tau\sigma^{-1}\tau^{-1}).$$

Now, $\sigma\tau\sigma^{-1} \in K$, since $K \lhd G$. Thus, $(\sigma\tau\sigma^{-1})\tau^{-1} \in K$. Similarly, $\sigma(\tau\sigma^{-1}\tau^{-1}) \in H$. But $H \cap K = \{1\}$, and since we just proved that $[\sigma, \tau] \in H \cap K$, we have $[\sigma, \tau] = 1$. The second part of the lemma is left to the reader (DX). □

Proof of Proposition 1.18. By Proposition 1.17, each p_j-Sylow subgroup P_j is normal in G. First, we claim that the map $P_1 \prod \cdots \prod P_t \xrightarrow{\varphi} G$ is a group homomorphism. Now, because the orders of P_i and P_j are relatively prime if $i \neq j$, we have $P_i \cap P_j = \{1\}$. Since

$$\varphi((\sigma_1, \ldots, \sigma_t)(\tau_1, \ldots, \tau_t)) = \sigma_1\tau_1 \cdots \sigma_t\tau_t,$$

using Lemma 1.19, we can push each τ_j past $\sigma_{j+1} \cdots \sigma_t$, and we get

$$\varphi((\sigma_1, \ldots, \sigma_t)(\tau_1, \ldots, \tau_t)) = \sigma_1 \cdots \sigma_t\tau_1 \cdots \tau_t = \varphi(\sigma_1, \ldots, \sigma_t)\varphi(\tau_1, \ldots, \tau_t),$$

proving that φ is a homomorphism. The kernel of φ consists of those $\sigma = (\sigma_1, \ldots, \sigma_t)$ so that $\sigma_1 \cdots \sigma_t = 1$, or equivalently, $\sigma_t^{-1} = \sigma_1 \cdots \sigma_{t-1}$. Using Lemma 1.19 and an obvious induction, the order on the righthand side is $p_1^{l_1} \cdots p_{t-1}^{l_{t-1}}$ and the order on the left hand side in $p_t^{l_t}$, which implies that $l_1 = \cdots = l_t$, and thus, all $\sigma_j = 1$. Therefore, Ker $\varphi = \{1\}$ and φ is injective. One more application of Lemma 1.19 yields $\#(P_1 \prod \cdots \prod P_t) = g$. Since φ is injective, it is an isomorphism. □

Remark: The proof of Proposition 1.18 only uses the fact that every p-Sylow subgroup is normal in G.

Definition 1.4 Let G be any group, then the *Frattini subgroup* of G, denoted $\Phi(G)$, is the intersection of all the maximal proper subgroups of G. In case G has no maximal proper subgroup, we set $\Phi(G) = G$.

Remark: The additive abelian group $(\mathbb{Q}, +)$ has no maximal proper subgroup.

Definition 1.5 In a group, G, an element σ is a *non-generator* iff for every subset, A, if $G = \mathrm{Gp}\{A, \sigma\}$, then $G = \mathrm{Gp}\{A\}$ (where $\mathrm{Gp}\{A\}$ denotes the subgroup of G generated by A).

As an example, assume that G is a cyclic group of order p^r. Then, $\Phi(G)$ is the cyclic subgroup of order p^{r-1}.

Proposition 1.20 *The Frattini subgroup of G is a characteristic subgroup of G, i.e., for every automorphism, $\varphi \in \mathrm{Aut}(G)$, we have $\varphi(\Phi(G)) = \Phi(G)$. In particular, $\Phi(G)$ is normal in G. Furthermore, if G is finite, then*

$$\Phi(G) = \{\sigma \in G \mid \sigma \textit{ is a non-generator}\}.$$

Proof. Every automorphism permutes the collection of maximal subgroups of G. Therefore, $\Phi(G)$ is characteristic. Now assume G is finite, or, at least, that every proper subgroup is contained in a maximal subgroup.

Claim: If $\mathrm{Gp}\{A, \Phi(G)\} = G$, then $\mathrm{Gp}\{A\} = G$.

If not, $\mathrm{Gp}\{A\} \neq G$, and so, there exists a maximal subgroup, M, containing $\mathrm{Gp}\{A\}$. Now, $\Phi(G) \subseteq M$, therefore, $\mathrm{Gp}\{A, \Phi(G)\} \subseteq M \neq G$, a contradiction. This proves that $\Phi(G)$ is contained in the set of non-generators.

Conversely, assume that σ is a non-generator. Were $\sigma \notin \Phi(G)$, we would have a maximal subgroup, M, with $\sigma \notin M$. Take $M = A$ in the definition of a non-generator. Look at $\mathrm{Gp}\{M, \sigma\}$. Of course, $M \subseteq \mathrm{Gp}\{M, \sigma\}$ and $\sigma \in \mathrm{Gp}\{M, \sigma\}$, so $M < \mathrm{Gp}\{M, \sigma\}$. But M is maximal, and so, $\mathrm{Gp}\{M, \sigma\} = G$. By definition (since σ is a non-generator), $G = \mathrm{Gp}\{M\}$, and thus, $G = M$, a contradiction. □

Definition 1.6 A group G is an *elementary abelian p-group* iff

(1) It is abelian, and

(2) For every $\sigma \in G$, we have $\sigma^p = 1$.

Remark: Any elementary abelian p-group is, in a natural way, a vector space over $\mathbb{F}_p$. Conversely, for any vector space over the finite field $\mathbb{F}_p$, its additive group is an elementary abelian p-group. Under this correspondence, an endomorphism of G goes over to a linear map and an automorphism of G goes to an invertible linear map. The group G is finite iff the corresponding vector space is finite dimensional.

(Given G, write the group operation additively. Thus, we have

$$p \cdot \sigma = \underbrace{\sigma + \cdots + \sigma}_{p} = 0.$$

The finite field $\mathbb{F}_p$ acts on G as follows: If $\lambda \in \mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$, *i.e.*, $\lambda \equiv 0, 1, \ldots, p-1 \pmod{p}$, we set

$$\lambda \cdot \sigma = \underbrace{\sigma + \cdots + \sigma}_{\lambda \,(\mathrm{mod}\, p)\ \text{times}}.$$

The reader should check that scalar multiplication is indeed well defined and that the facts asserted in the previous remark are true (DX).)

Proposition 1.21 *For any p-group, G, the quotient group, $G/\Phi(G)$, is an elementary abelian p-group.*

Proof. Say H is a maximal subgroup of G. Since G has (N), the group, H, is normal in G and $(G : H) = p$. Therefore, G/H is cyclic of order p. Write $\overline{\sigma}$ for the image of σ in G/H. We know that $(\overline{\sigma})^p = 1$. So, $\overline{\sigma^p} = 1$, *i.e.*, $\sigma^p \in H$. But H is arbitrary, and so,

$$\sigma^p \in \bigcap_{H \text{ maximal}} H = \Phi(G).$$

Now, G/H is abelian since $G/H = \mathbb{Z}/p\mathbb{Z}$. This implies that $[G, G] \subseteq H$ (here $[G, G]$ is the subgroup of G generated by the commutators, called the *commutator group of* G; it is the smallest normal subgroup, K, of G such that G/K is abelian). Since H is arbitrary, we get

$$[G, G] \subseteq \bigcap_{H \text{ maximal}} H = \Phi(G).$$

This shows that $G/\Phi(G)$ is abelian. As $\sigma^p \in \Phi(G)$, we get $(\overline{\overline{\sigma}})^p = 1$ in $G/\Phi(G)$, where $\overline{\overline{\sigma}}$ is the image of σ in $G/\Phi(G)$. □

We now come to a famous theorem of Burnside.

Theorem 1.22 *(Burnside Basis Theorem) Say G is a p-group and let d be the minimal number of elements found among all minimal generating sets for G. The following properties hold:*

(1) Given any set of d elements in G, say $\sigma_1, \ldots, \sigma_d$, they generate G iff $\overline{\sigma_1}, \ldots, \overline{\sigma_d}$ are a basis of $G/\Phi(G)$.

(2) More generally, any set of t elements $\sigma_1, \ldots, \sigma_t$ in G generates G iff $\{\overline{\sigma_1}, \ldots, \overline{\sigma_t}\}$ spans $G/\Phi(G)$. Hence, any set of generators of G possesses a subset of exactly d elements which generates G. The number d is the dimension of $G/\Phi(G)$ over $\mathbb{F}_p$.

Proof. Everything follows from the statement: $\sigma_1, \ldots, \sigma_t$ generate G iff $\overline{\sigma_1}, \ldots, \overline{\sigma_t}$ generate $\overline{G} = G/\Phi(G)$ (DX).

The implication ($\Longrightarrow$) is trivial and always true. Conversely, if $\overline{\sigma_1}, \ldots, \overline{\sigma_t}$ generate $\overline{G}$, then

$$G = \mathrm{Gp}\{\sigma_1, \ldots, \sigma_t, \Phi(G)\}.$$

But then, as $\Phi(G)$ is the set of nongenerators, we have

$$G = \mathrm{Gp}\{\sigma_1, \ldots, \sigma_t, \Phi(G)\} = \mathrm{Gp}\{\sigma_1, \ldots, \sigma_t\},$$

as desired. $\square$

Let G be a group (possibly infinite). We set $\Delta^{(0)}(G) = G$, and $\Delta^{(1)}(G) = [G, G]$ and, more generally

$$\Delta^{(j+1)}(G) = [\Delta^{(j)}(G), \Delta^{(j)}(G)] = \Delta^{(1)}(\Delta^{(j)}(G)).$$

Observe that $\Delta^{(1)}(G) = [G, G]$ is the commutator group of G, and recall that for any normal subgroup, H, of G, we have $\Delta^{(1)}(G) \subseteq H$ iff G/H is abelian. Moreover, for a simple nonabelian group, $[G, G] = G$.

Proposition 1.23 *Suppose G is a group, then each $\Delta^{(j)}(G)$ is a characteristic subgroup of G and each group $\Delta^{(j)}(G)/\Delta^{(j+1)}(G)$ is abelian ($j \geq 0$). If G has property (N), then $\Delta^{(1)}(G) \subseteq \Phi(G) < G$ (provided maximal subgroups exist). If G is a p-group, then the chain*

$$G \supseteq \Delta^{(1)}(G) \supseteq \Delta^{(2)}(G) \supseteq \cdots \supseteq \Delta^{(t)}(G) \supseteq \cdots$$

is strictly descending and reaches $\{1\}$ after finitely many steps.

Proof. The group $\Delta^{(1)}(G)$ consists of products of the form

$$[\sigma_1, \tau_1] \cdots [\sigma_l, \tau_l], \quad l \geq 1.$$

If $\varphi \in \mathrm{Aut}(G)$, then

$$\varphi([\sigma_1, \tau_1] \cdots [\sigma_l, \tau_l]) = \varphi([\sigma_1, \tau_1]) \cdots \varphi([\sigma_l, \tau_l]),$$

and $\varphi([\sigma, \tau]) = [\varphi(\sigma), \varphi(\tau)]$, so $\Delta^{(1)}(G)$ is characteristic. We prove that $\Delta^{(j)}(G)$ is characteristic by induction on j. The base case $j = 1$ has just been established. Look at $\Delta^{(j+1)}(G)$. By the induction hypothesis, we have $\varphi(\Delta^{(j)}(G)) = \Delta^{(j)}(G)$. Therefore, φ is an automorphism of $\Delta^{(j)}(G)$. Yet, $\Delta^{(j+1)}(G) = \Delta^{(1)}(\Delta^{(j)}(G))$, and we proved that $\Delta^{(1)}(H)$ is characteristic for any group H (case $j = 1$). Now, $G/\Delta^{(1)}(G)$ is abelian for any group G, so $\Delta^{(j)}(G)/\Delta^{(j+1)}(G) = \Delta^{(j)}(G)/\Delta^{(1)}(\Delta^{(j)}(G))$ is abelian.

Say G has (N) and possesses maximal subgroups. If H is a maximal subgroup of G we know that $H \lhd G$ and H has prime index. So, G/H is abelian, and thus, $\Delta^{(1)}(G) \subseteq H$. Since H is arbitrary, we deduce that

$$\Delta^{(1)}(G) \subseteq \bigcap_{H \text{ maximal}} H = \Phi(G).$$

Now, assume that G is a p-group. Then, G has (N), and thus, $\Delta^{(1)}(G) \subseteq \Phi(G) < G$. But $\Delta^{(1)}(G)$ in turn is a p-group, so we can apply the argument to $\Delta^{(1)}(G)$ and we get $\Delta^{(2)}(G) < \Delta^{(1)}(G)$, etc. $\square$

Nomenclature.

(1) The group $\Delta^{(1)}(G)$ is called the *first derived group* of G (or *commutator group* of G).

(2) The group $\Delta^{(j)}(G)$ is the *j-th derived group* of G.

(3) The sequence
$$G = \Delta^{(0)}(G) \supseteq \Delta^{(1)}(G) \supseteq \Delta^{(2)}(G) \supseteq \cdots \supseteq \Delta^{(t)}(G) \supseteq \cdots$$
is the *derived series* of G.

(4) The smallest $t \geq 0$ for which $\Delta^{(t)}(G) = \{1\}$ is the *derived length* of G and if $\Delta^{(t)}(G)$ is never $\{1\}$ (e.g., in a nonabelian simple group) then the derived length is infinite. Write $\delta(G)$ for the derived length of G.

Look at the derived series of G:
$$G = \Delta^{(0)}(G) \supseteq \Delta^{(1)}(G) \supseteq \Delta^{(2)}(G) \supseteq \cdots \supseteq \Delta^{(t)}(G) \supseteq \cdots .$$
Each quotient $\Delta^{(j)}(G)/\Delta^{(j+1)}(G)$ is abelian. Suppose G is finite, then $\Delta^{(j)}(G)/\Delta^{(j+1)}(G)$ is finite abelian. Interpolate between $\Delta^{(j)}(G)$ and $\Delta^{(j+1)}(G)$ a sequence of subgroups, necessarily normal, each maximal in the previous one. If $\delta(G) < \infty$, we get a composition series all of whose factors are cyclic of prime order. This proves half of the

Proposition 1.24 *A necessary and sufficient condition that a finite group be solvable (in the sense of Galois) is that $\delta(G) < \infty$.*

Proof. We need only prove: If G is (Galois) solvable, then $\delta(G) < \infty$. Say
$$G = G_0 > G_1 > G_2 > \cdots > G_t = \{1\}$$
is a composition series with abelian factors. We have $G_1 < G$ and G/G_1 is abelian. Therefore, by a previous remark, $\Delta^{(1)}(G) \subseteq G_1$. Each quotient G_j/G_{j+1} is abelian, so $\Delta^{(1)}(G_j) \subseteq G_{j+1}$ for all j. Now, $\Delta^{(1)}(G) \subseteq G_1$ implies that $\Delta^{(1)}(\Delta^{(1)}(G)) \subseteq \Delta^{(1)}(G_1)$, and so,
$$\Delta^{(2)}(G) \subseteq \Delta^{(1)}(G_1) \subseteq G_2.$$
An easy induction yields $\Delta^{(r)}(G) \subseteq G_r$ (DX). Therefore, $\Delta^{(t)}(G) \subseteq \{1\}$, *i.e.*, $\delta(G) \leq t$. □

Observe that we actually proved more: The derived length, $\delta(G)$, of a solvable finite group is less than or equal to the length of any composition series for G.

Definition 1.7 An arbitrary group, G, is *solvable* iff $\delta(G) < \infty$.

Proposition 1.25 *Say G is a p-group of order at least p^2. Then, $(G : \Delta^{(1)}(G)) \geq p^2$.*

Proof. We may assume that G is nonabelian, else $\Delta^{(1)}(G) = \{1\}$ and so, $(G : \Delta^{(1)}(G)) = \#(G) \geq p^2$. As G is a p-group, if $(G : \Delta^{(1)}(G)) < p^2$, then $(G : \Delta^{(1)}(G)) = p$. We know that $\Delta^{(1)}(G) \subseteq \Phi(G)$. Therefore, $(G : \Phi(G)) = p$ and the Burnside dimension of G (*i.e.* $\dim_{\mathbb{F}_p} G/\Phi(G)$) is equal to 1. By the Burnside basis theorem, G is cyclic, so abelian, a contradiction. □

1.4 Group Extensions

Let G be a finite group and let

$$G = G_0 > G_1 > G_2 > \cdots > G_t = \{1\}$$

be a composition series. We have the groups $G_j/G_{j+1} = \overline{G}_j$, the composition factors of the composition series.

Problem: Given the (ordered) sequence $\overline{G_0}, \overline{G}_1, \overline{G}_2, \ldots, \overline{G}_{t-1}$, try to reconstruct G.

Say H and K are two groups, $\mathcal{G}$ is a "big" group and $H \lhd \mathcal{G}$ with $\mathcal{G}/H \xrightarrow{\sim} K$. Note, this is exactly the case at the small end of a composition series. We have
$G_{t-1} = \overline{G}_{t-1} = G_{t-1}/G_t$. We also have $G_{t-1} \lhd G_{t-2}$, and the quotient is $\overline{G}_{t-2}$, so we are in the above situation with $H = G_{t-1} = \overline{G}_{t-1}$, $K = \overline{G}_{t-2}$, $\mathcal{G} = G_{t-2}$, and $\mathcal{G}/H \xrightarrow{\sim} K$.

The above situation is a special case of an *exact sequence*. A diagram of groups and homomorphisms

$$0 \longrightarrow H \xrightarrow{\varphi} \mathcal{G} \xrightarrow{\psi} K \longrightarrow 0,$$

where the map $0 \longrightarrow H$ is the inclusion of $\{1\}$ into H and the map $K \longrightarrow 0$ is the surjection sending every element of K to 1 in the trivial group $\{1\}$, is called a *short exact sequence* iff the kernel of every homomorphism is equal to the image of the previous homomorphism on its left. This means that

(1) $\mathrm{Ker}\ \varphi = \{1\}$, so φ is injective, and we identify H with a subgroup of $\mathcal{G}$.

(2) $H = \mathrm{Im}\ \varphi = \mathrm{Ker}\ \psi$, so H is normal in $\mathcal{G}$.

(3) $\mathrm{Im}\ \psi = K$, so ψ is surjective. By the first homomorphism theorem, $\mathcal{G}/H \xrightarrow{\sim} K$.

(4) Properties (1), (2), (3) are equivalent to $0 \longrightarrow H \longrightarrow \mathcal{G} \longrightarrow K \longrightarrow 0$ is exact.

Going back to composition series, we have $G_{j+1} \lhd G_j$ and $\overline{G}_j = G_j/G_{j+1}$. So, a composition series is equivalent with a collection of short exact sequences

$$\begin{array}{l} 0 \longrightarrow G_{t-1} \longrightarrow G_{t-2} \longrightarrow \overline{G}_{t-2} \longrightarrow 0 \\ 0 \longrightarrow G_{t-2} \longrightarrow G_{t-3} \longrightarrow \overline{G}_{t-3} \longrightarrow 0 \\ \cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots \\ 0 \longrightarrow G_1 \longrightarrow G \longrightarrow \overline{G}_0 \longrightarrow 0. \end{array}$$

So our problem reduces to the problem of *group extensions*: Given H and K, groups, find (classify) all groups, $\mathcal{G}$, which can possibly fit into an exact sequence

$$0 \longrightarrow H \longrightarrow \mathcal{G} \longrightarrow K \longrightarrow 0.$$

The problem is very hard when H is nonabelian.

Definition 1.8 If A, G are groups, a group, $\mathcal{G}$, is an *extension of G by A* iff $\mathcal{G}$ fits into an exact sequence

$$(E) \qquad 0 \longrightarrow A \longrightarrow \mathcal{G} \longrightarrow G \longrightarrow 0.$$

Two such extensions (E), (E') are *equivalent* iff there exists a commutative diagram

$$\begin{tikzcd}
(E) & 0 \arrow[r] & A \arrow[r] \arrow[d, equal] & \mathcal{G} \arrow[r] \arrow[d, "\psi"] & G \arrow[r] \arrow[d, equal] & 0 \\
(E') & 0 \arrow[r] & A \arrow[r] & \mathcal{G}' \arrow[r] & G \arrow[r] & 0.
\end{tikzcd}$$

Remarks:

(1) The homomorphism, ψ, in the above diagram is an isomorphism of groups. So, the notion of equivalence is indeed an equivalence relation (DX).

(2) Equivalence of group extensions is stronger than isomorphism of $\mathcal{G}$ with $\mathcal{G}'$.

(3) The group $\mathcal{G}$ in (E) should be considered a "fibre space" whose base is G and whose "fibre" is A.

As we remarked before, the theory is good only when A is abelian. *From now on, we assume A is an abelian group.*

Proposition 1.26 *Say*

$$(E) \qquad 0 \longrightarrow A \longrightarrow \mathcal{G} \longrightarrow G \longrightarrow 0$$

is a group extension and A is abelian. Then, there exists a natural action of G on A; so, A is a G-module. Equivalent extensions give rise to the same action.

Proof. Denote the surjective homomorphism $\mathcal{G} \longrightarrow G$ in (E) by bar (-). Pick $\xi \in G$ and any $a \in A$. There exists $x \in \mathcal{G}$ with $\overline{x} = \xi$. Consider xax^{-1}. Since $A \lhd \mathcal{G}$, we have $xax^{-1} \in A$. If $y \in \mathcal{G}$ and if $\overline{y} = \overline{x} = \xi$, then $x = y\alpha$ for some $\alpha \in A$. Then,

$$xax^{-1} = y\alpha a\alpha^{-1}y^{-1} = yay^{-1},$$

as A is abelian. Therefore, if we set

$$\xi \cdot a = xax^{-1},$$

this is a well-defined map. The reader should check that it is an action (DX). Assume we have an equivalence of extensions between (E) and (E'):

$$\begin{tikzcd}
(E) & 0 \arrow[r] & A \arrow[r] \arrow[d, equal] & \mathcal{G} \arrow[r] \arrow[d, "\psi"] & G \arrow[r] \arrow[d, equal] & 0 \\
(E') & 0 \arrow[r] & A \arrow[r] & \mathcal{G}' \arrow[r] & G \arrow[r] & 0.
\end{tikzcd}$$

Pick $\xi \in G$ and any $a \in A$. Denote the E-action by $\cdot$ and the E'-action by $\cdot\cdot$. Observe that

$$\xi \cdot a = \psi(\xi \cdot a) = \psi(xax^{-1}) = \psi(x)\psi(a)\psi(x)^{-1} = \psi(x)a\psi(x)^{-1},$$

since the left vertical arrow is the identity in the diagram, yet $\psi(x)$ lifts ξ in $\mathcal{G}'$, as the right vertical arrow is the identity in the diagram. However, by definition,

$$\xi \cdot\cdot\, a = \psi(x)a\psi(x)^{-1},$$

so, $\xi \cdot\cdot\, a = \xi \cdot a$ for all $a \in A$. $\square$

The *type* of (E) is the structure of A as G-module, i.e., the action of G on A. We get a first invariant of a *group extension*, its action (of G on A).

Fix the action of (E). Can we classify the extensions up to equivalence? Say we are given an extension

$$(E) \qquad 0 \longrightarrow A \longrightarrow \mathcal{G} \xrightarrow{\pi} G \longrightarrow 0.$$

There is always a set-theoretic section $s\colon G \to \mathcal{G}$, i.e., a set map, s, so that $\pi(s(\sigma)) = \sigma$ for all $\sigma \in G$. Write u_σ for the s-lift of σ, i.e., $s(\sigma) = u_\sigma$. So, $\pi(u_\sigma) = \overline{u_\sigma} = \sigma$. As s is **not** necessarily a group homomorphism, what is the obstruction? Consider

$$u_\sigma u_\tau (u_{\sigma\tau})^{-1} = f(\sigma, \tau). \qquad (*)$$

Note that $f(\sigma,\tau)=1$ iff $s\colon \sigma\mapsto u_\sigma$ is a group homomorphism. If we apply the homomorphism bar to $(*)$, we get $\overline{f(\sigma,\tau)}=1$, and so, $f(\sigma,\tau)\in A$. Observe that f is a function
$f\colon G\prod G\to A$. Given $x\in\mathcal{G}$, look at $\overline{x}$. We know that $\overline{x}=\sigma\in G$. If we apply bar to xu_σ^{-1}, we get 1, because $\overline{u_\sigma^{-1}}=\sigma^{-1}$ and $\overline{x}=\sigma$. So, we have $xu_\sigma^{-1}\in A$, which yields $x=au_\sigma$, for some $a\in A$.

Observe that:

(1) Each x determines *uniquely* a representation $x=au_\sigma$, with $a\in A$ and $\sigma\in G$.

(2) The map $A\prod G\longrightarrow\mathcal{G}$ (where $A\prod G$ is the product of A and G *as sets*) *via*

$$(a,\sigma)\mapsto au_\sigma$$

is a bijection of sets (an isomorphism in the category of sets).

(3) $\mathcal{G}$ (as a set) is just $A\prod G$ (product in the category of sets).[2]

Can we recover the group multiplication of $\mathcal{G}$? We have

$$\begin{aligned}(au_\sigma)(bu_\tau) &= a(u_\sigma b)u_\tau\\ &= a(u_\sigma bu_\sigma^{-1})u_\sigma u_\tau\\ &= a(\sigma\cdot b)u_\sigma u_\tau\\ &= a(\sigma\cdot b)f(\sigma,\tau)u_{\sigma\tau}\\ &= cu_{\sigma\tau},\end{aligned}$$

where $c=a(\sigma\cdot b)f(\sigma,\tau)$, and $c\in A$. Therefore, knowledge of the action and $f(\sigma,\tau)$ gives us knowledge of the group multiplication.

Thus, it is natural to try to go backwards and make $\mathcal{G}$ from the groups A and G, the action of G on A, and f. It is customary to use an additive notation for the group operation in A, since A is abelian. The underlying set of the group $\mathcal{G}$ is

$$A\prod G=\{\langle a,\sigma\rangle\mid a\in A,\,\sigma\in G\}.$$

Multiplication is given by

$$\langle a,\sigma\rangle\langle b,\tau\rangle=\langle a+\sigma\cdot b+f(\sigma,\tau),\sigma\tau\rangle. \tag{$\dagger$}$$

However, the multiplication defined by ($\dagger$) is supposed to make $\mathcal{G}$ into a group, and this imposes certain conditions on f. First, we deal with associativity. For this, we go back to the original $\mathcal{G}$ where we have the associative law:

$$(au_\sigma)((bu_\tau)(cu_\rho))=((au_\sigma)(bu_\tau))(cu_\rho).$$

Expanding the left hand side, we get

$$\begin{aligned}(au_\sigma)((bu_\tau)(cu_\rho)) &= (au_\sigma)(b(\tau\cdot c)f(\tau,\rho)u_{\tau\rho})\\ &= (a\sigma\cdot(b(\tau\cdot c)f(\tau,\rho))f(\sigma,\tau\rho)u_{\sigma(\tau\rho)}\\ &= a(\sigma\cdot b)(\sigma\tau\cdot c)(\sigma\cdot f(\tau,\rho))f(\sigma,\tau\rho)u_{\sigma(\tau\rho)}.\end{aligned}$$

Expanding the righthand side, we get

$$\begin{aligned}((au_\sigma)(bu_\tau))(cu_\rho) &= a(\sigma\cdot b)f(\sigma,\tau)u_{\sigma\tau})(cu_\rho)\\ &= a(\sigma\cdot b)f(\sigma,\tau)(\sigma\tau\cdot c)f(\sigma\tau,\rho)u_{(\sigma\tau)\rho}.\end{aligned}$$

[2] In (2) and (3) we give a foretaste of the language of categories to be introduced in Section 1.7.

Thus, the associative law becomes (writing RHS = LHS)

$$f(\sigma,\tau)(\sigma\tau\cdot c)f(\sigma\tau,\rho)=(\sigma\tau\cdot c)(\sigma\cdot f(\tau,\rho))f(\sigma,\tau\rho).$$

Now, all the above terms are in A, and since A is abelian, we can permute terms and perform cancellations, and we get

$$f(\sigma,\tau)f(\sigma\tau,\rho)=(\sigma\cdot f(\tau,\rho))f(\sigma,\tau\rho).\tag{††}$$

This identity is equivalent to the associativity law in $\mathcal{G}$.

Nomenclature: A function from $G\prod G$ to A is called a 2-*cochain on* G *with values in* A. Any 2-cochain satisfying (††) is called a 2-*cocycle with coefficients in* A.

Therefore, (†) is an associative multiplication in $A\prod G$ iff f is a 2-cocycle with values in A.

Does $A\prod G$ with multiplication (†) have an identity?

The original group, $\mathcal{G}$, has identity 1 and we have $1=u_1^{-1}u_1$, where $u_1\in A$, and so, $u_1^{-1}\in A$. For all $b\in A$ and all $\tau\in G$, we have

$$(u_1^{-1}u_1)(bu_\tau)=bu_\tau,$$

which yields

$$u_1^{-1}(1\cdot b)f(1,\tau)u_\tau=u_1^{-1}bf(1,\tau)u_\tau=bu_\tau.$$

Since A is abelian, we get

$$f(1,\tau)=u_1,$$

which shows that $f(1,\tau)$ is independent of τ. In particular, $u_1=f(1,1)$.

Question: Is (††) sufficient to imply that $f(1,\tau)=f(1,1)$ for all $\tau\in G$?

In (††), take $\sigma=1$. We get

$$f(1,\tau)f(\tau,\rho)=f(\tau,\rho)f(1,\tau\rho).$$

Again, since A is abelian, we deduce that $f(1,\tau)=f(1,\tau\rho)$. If we take $\tau=1$, we get $f(1,1)=f(1,\rho)$, for all ρ.

Therefore, (††) is sufficient and $A\prod G$ has an identity $\mathbf{1}=\langle f(1,1)^{-1},\,1\rangle$, or in additive notation (since A is abelian),

$$\mathbf{1}=\langle -f(1,1),\,1\rangle.\tag{$*$}$$

Finally, what about inverses? Once again, go back to our original $\mathcal{G}$.

We have $(au_\sigma)^{-1}=u_\sigma^{-1}a^{-1}$. Now,

$$\overline{u_\sigma^{-1}}=(\overline{u_\sigma})^{-1}=\sigma^{-1}=\overline{u_{\sigma^{-1}}}.$$

Therefore, there is some $\alpha\in A$ so that $u_\sigma^{-1}=\alpha u_{\sigma^{-1}}$. By multiplying on the right by u_σ, we get

$$1=\alpha u_{\sigma^{-1}}u_\sigma=\alpha f(\sigma^{-1},\sigma)u_{\sigma\sigma^{-1}}=\alpha f(\sigma^{-1},\sigma)u_1=\alpha f(\sigma^{-1},\sigma)f(1,1),$$

since $u_1=f(1,1)$. So, $\alpha=f(1,1)^{-1}f(\sigma^{-1},\sigma)^{-1}$. Consequently, we get

$$\begin{aligned}(au_\sigma)^{-1}&=u_\sigma^{-1}a^{-1}\\&=\alpha u_{\sigma^{-1}}a^{-1}\\&=\alpha(u_{\sigma^{-1}}a^{-1}u_{\sigma^{-1}}^{-1})u_{\sigma^{-1}}\\&=\alpha(\sigma^{-1}\cdot a^{-1})u_{\sigma^{-1}}\\&=f(1,1)^{-1}f(\sigma^{-1},\sigma)^{-1}(\sigma^{-1}\cdot a^{-1})u_{\sigma^{-1}}\\&=f(1,1)^{-1}f(\sigma^{-1},\sigma)^{-1}(\sigma^{-1}\cdot a)^{-1}u_{\sigma^{-1}}\\&=((\sigma^{-1}\cdot a)f(\sigma^{-1},\sigma)f(1,1))^{-1}u_{\sigma^{-1}}.\end{aligned}$$

Therefore, in $A\prod G$ (switching to additive notation since A is abelian), inverses are given by

$$\langle a,\sigma\rangle^{-1}=\langle -\sigma^{-1}\cdot a-f(\sigma^{-1},\sigma)-f(1,1),\,\sigma^{-1}\rangle. \tag{$**$}$$

We find that $A\prod G$ can be made into a group *via* (†), provided $f(\sigma,\tau)$ satisfies (††). The formulae ($*$) and ($**$) give the unit element and inverses, respectively. For temporary notation, let us write $(A\prod G;f)$ for this group. Also, since A is abelian, let us rewrite (††) in additive notation, since this will be more convenient later on:

$$\sigma\cdot f(\tau,\rho)+f(\sigma,\tau\rho)=f(\sigma\tau,\rho)+f(\sigma,\tau). \tag{††}$$

Go back to the original group, $\mathcal{G}$, and its set-theoretic section $s\colon G\to\mathcal{G}$ (with $s(\sigma)=u_\sigma$). We might have chosen another set-theoretic section, $t\colon G\to\mathcal{G}$, namely, $t(\sigma)=v_\sigma$. We get a 2-cocycle $g(\sigma,\tau)=v_\sigma v_\tau(v_{\sigma\tau})^{-1}$, i.e., $v_\sigma v_\tau=g(\sigma,\tau)v_{\sigma\tau}$.

What is the relation between f and g?

We know that $\overline{v_\sigma}=\sigma=\overline{u_\sigma}$, which implies that there is some $k(\sigma)\in A$ with $v_\sigma=k(\sigma)u_\sigma$. Then, we have

$$v_\sigma v_\tau=g(\sigma,\tau)v_{\sigma\tau}=g(\sigma,\tau)k(\sigma\tau)u_{\sigma\tau},$$

and also

$$v_\sigma v_\tau=k(\sigma)u_\sigma k(\tau)u_\tau=k(\sigma)(\sigma\cdot k(\tau))u_\sigma u_\tau=k(\sigma)(\sigma\cdot k(\tau))f(\sigma,\tau)u_{\sigma\tau}.$$

By equating these expressions, we get

$$g(\sigma,\tau)k(\sigma\tau)=k(\sigma)(\sigma\cdot k(\tau))f(\sigma,\tau).$$

But A is abelian, so we can write the above

$$g(\sigma,\tau)-f(\sigma,\tau)=\sigma\cdot k(\tau)-k(\sigma\tau)+k(\sigma). \tag{$*$}$$

Observe that $k\colon G\to A$ is a function of one variable on G. We call k a 1-*cochain on G with values in A*. For a 1-cochain, define a corresponding 2-cochain, called its *coboundary*, δk, by

$$(\delta k)(\sigma,\tau)=\sigma\cdot k(\tau)-k(\sigma\tau)+k(\sigma).$$

Remarks:

(1) Every coboundary of a 1-cochain is automatically a 2-cocycle (DX).

(2) Cocycles form a group under addition of functions denoted by $Z^2(G,A)$. The special 2-cocycles which are coboundaries (of 1-cochains) form a group (DX) denoted by $B^2(G,A)$. Item (1) says that $B^2(G,A)$ is a subgroup of $Z^2(G,A)$.

(3) The quotient group, $Z^2(G,A)/B^2(G,A)$, denoted $H^2(G,A)$, is the *second cohomology group of G with coefficients in A.*

(4) Equation ($*$) above says: If we change the choice of section from s to t, the corresponding cocycles, f and g, are *cohomologous*, i.e., $g-f=\delta k$, i.e., the image of f in $H^2(G,A)$ is the same as the image of g in $H^2(G,A)$. Thus, it is the cohomology class of f which is determined by (E).

Now, make $(A\prod G; f)$. Then, we can map A into $(A\prod G; f)$ *via*

$$a \mapsto \langle a - f(1,1), 1\rangle.$$

Claim. The set $\{\langle a - f(1,1), 1\rangle \mid a \in A\}$ is a subgroup of $(A\prod G; f)$. In fact, it is a normal subgroup and the quotient is G.

Proof. We have

$$\langle a - f(1,1), 1\rangle\langle b - f(1,1), 1\rangle = \langle a - f(1,1) + b - f(1,1) + f(1,1), 1\rangle = \langle a + b - f(1,1), 1\rangle,$$

and thus, the map $\lambda\colon a \mapsto \langle a - f(1,1), 1\rangle$ is a group homomorphism. We leave the rest as a (DX). □

Say $f - g = \delta k$, i.e., f and g are cohomologous, and make $(A\prod G; f)$ and $(A\prod G; g)$. Consider the map $\theta\colon (A\prod G; f) \to (A\prod G; g)$ given by

$$\theta\colon \langle a, \sigma\rangle \mapsto \langle a + k(\sigma), \sigma\rangle.$$

We claim that θ is a homomorphism. Since

$$\langle a, \sigma\rangle\langle b, \tau\rangle = \langle a + \sigma\cdot b + f(\sigma,\tau), \sigma\tau\rangle,$$

we have

$$\theta(\langle a, \sigma\rangle\langle b, \tau\rangle) = \langle a + \sigma\cdot b + f(\sigma,\tau) + k(\sigma\tau), \sigma\tau\rangle.$$

We also have

$$\begin{aligned}\theta(\langle a, \sigma\rangle)\theta(\langle b, \tau\rangle) &= \langle a + k(\sigma), \sigma\rangle\langle b + k(\tau), \tau\rangle)\\ &= \langle a + k(\sigma) + \sigma\cdot b + \sigma\cdot k(\tau) + g(\sigma,\tau), \sigma\tau\rangle).\end{aligned}$$

In order for θ to be a homomorphism, we need

$$k(\sigma) + \sigma\cdot k(\tau) + g(\sigma,\tau) = f(\sigma,\tau) + k(\sigma\tau),$$

that is, $f - g = \delta k$. Consequently, θ is a homomorphism, in fact, an isomorphism. Moreover, $(A\prod G; f)$ and $(A\prod G; g)$ fit into two extensions and we have the following diagram:

$$\begin{tikzcd}
(E)_f & 0 \arrow[r] & A \arrow[r] \arrow[d, equal] & (A\prod G; f) \arrow[r] \arrow[d, "\theta"] & G \arrow[r] \arrow[d, equal] & 0 \\
(E')_g & 0 \arrow[r] & A \arrow[r] & (A\prod G; g) \arrow[r] & G \arrow[r] & 0.
\end{tikzcd}$$

The rightmost rectangle commutes, but we need to check that the leftmost rectangle commutes. Going over horizontally and down from $(A\prod G; f)$, for any $a \in A$, we have

$$a \mapsto \langle a - f(1,1), 1\rangle \mapsto \langle a - f(1,1) + k(1), 1\rangle,$$

and going horizontally from the lower A, we have

$$a \mapsto \langle a - g(1,1), 1\rangle.$$

For the rectangle to commute, we need: $g(1,1) = f(1,1) - k(1)$. However,
$f(\sigma,\tau) = g(\sigma,\tau) + \delta k(\sigma,\tau)$ and $\delta k(\sigma,\tau) = \sigma\cdot k(\tau) - k(\sigma\tau) + k(\sigma)$. If we set $\sigma = \tau = 1$, we get

$$\delta k(1,1) = k(1) - k(1) + k(1) = k(1),$$

and it follows that $g(1,1) = f(1,1) - k(1)$, as desired.

Hence, cohomologous 2-cocycles give rise to *equivalent* group extensions (the action is fixed). Conversely, we now show that equivalent group extensions give rise to cohomologous 2-cocycles. Say

$$\begin{tikzcd}
(E) & 0 \arrow[r] & A \arrow[r] \arrow[d, equal] & \mathcal{G} \arrow[r] \arrow[d, "\psi"] & G \arrow[r] \arrow[d, equal] & 0 \\
(E') & 0 \arrow[r] & A \arrow[r] & \mathcal{G}' \arrow[r] & G \arrow[r] & 0.
\end{tikzcd}$$

is an equivalence of extensions (i.e., the diagram commutes). We know, up to the notion of being cohomologous, that we may adjust both cocycles f and g associated with (E) and (E') by choice of sections. In both cases, take $u_1 = 0$ (since we are using additive notation). Therefore, $f(1,1) = g(1,1) = 0$. From the commutativity of the diagram, ψ must be of the form

$$\psi\langle a, \sigma\rangle = \langle \varphi(a, \sigma), \sigma\rangle$$

for some function $\varphi\colon A \prod G \to A$. By the above choice, the maps $A \longrightarrow \mathcal{G}$ and $A \longrightarrow \mathcal{G}'$ are given by $a \mapsto \langle a, 1\rangle$ in both cases. Therefore,
$\psi(a, 1) = \langle \varphi(a, 1), 1) = (a, 1)$, and so,

$$\varphi(a, 1) = a, \quad \text{for all } a \in A.$$

Since ψ is a homomorphism, we have

$$\psi(\langle a, \sigma\rangle\langle b, \tau\rangle) = \psi(\langle a, \sigma\rangle)\psi(\langle b, \tau\rangle),$$

and this yields an identity relating f, g and φ. The left hand side of the above equation is equal to

$$\psi(\langle a + \sigma\cdot b + f(\sigma, \tau), \sigma\tau\rangle) = \langle \varphi(a + \sigma\cdot b + f(\sigma, \tau), \sigma\tau\rangle), \sigma\tau\rangle,$$

and the righthand side is equal to

$$\langle \varphi(a, \sigma), \sigma\rangle\langle \varphi(b, \tau), \tau\rangle = \langle \varphi(a, \sigma) + \sigma\cdot\varphi(b, \tau) + g(\sigma, \tau), \sigma\tau\rangle,$$

and by equating them, we get

$$\varphi(a + \sigma\cdot b + f(\sigma, \tau), \sigma\tau\rangle) = \varphi(a, \sigma) + \sigma\cdot\varphi(b, \tau) + g(\sigma, \tau). \tag{†††}$$

By taking $\tau = 1$ (using the fact that $\varphi(b, 1) = b$), we get

$$\varphi(a + \sigma\cdot b + f(\sigma, 1), \sigma\rangle) = \varphi(a, \sigma) + \sigma\cdot b + g(\sigma, 1). \tag{∗∗∗}$$

Now, (††) can be written as

$$\sigma\cdot f(\tau, \rho) - f(\sigma\tau, \rho) + f(\sigma, \tau\rho) - f(\sigma, \tau) = 0.$$

If we take $\rho = 1$, we get

$$\sigma\cdot f(\tau, 1) - f(\sigma\tau, 1) + f(\sigma, \tau) - f(\sigma, \tau) = 0.$$

which yields

$$\sigma\cdot f(\tau, 1) = f(\sigma\tau, 1).$$

If we take $\tau = 1$, we get $\sigma\cdot f(1, 1) = f(\sigma, 1)$, but $f(1, 1) = 0$, and so,

$$f(\sigma, 1) = 0.$$

Consequently, (∗∗∗) yields

$$\varphi(a + \sigma\cdot b, \sigma) = \varphi(a, \sigma) + \sigma\cdot b.$$

Writing $b = \sigma^{-1} \cdot c$, we get
$$\varphi(a + c, \sigma) = \varphi(a, \sigma) + c, \quad \text{for all } a, c \in A.$$

In particular, when $a = 0$, we get $\varphi(c, \sigma) = \varphi(0, \sigma) + c$. Let $\varphi(0, \sigma) = k(\sigma)$. Now, if we use $\varphi(a, \sigma) = \varphi(0, \sigma) + a$ in (†††), we get
$$a + \sigma \cdot b + f(\sigma, \tau) + k(\sigma\tau) = a + k(\sigma) + \sigma \cdot (b + k(\tau)) + g(\sigma, \tau),$$

which yields
$$f(\sigma, \tau) + k(\sigma\tau) = g(\sigma, \tau) + k(\sigma) + \sigma \cdot k(\tau),$$

that is, $f - g = \delta k$. Hence, we have proved almost all of the following fundamental theorem:

Theorem 1.27 *If G and A are groups and A is abelian, then each group extension*

$$(E) \qquad 0 \longrightarrow A \longrightarrow \mathcal{G} \xrightarrow{\pi} G \longrightarrow 0$$

makes A into a G-module; the G-module structure is the type of (E) and equivalent extensions have the same type. For a given type, the equivalence classes of extensions of G by A are in one-to-one correspondence with $H^2(G, A)$, the second cohomology group of G with coefficients in A. Hence, the distinct extensions of G by A (up to equivalence) are classified by the pairs $(\mathrm{type}(E), \chi(E))$, where $\chi(E)$ is the cohomology class in $H^2(G, A)$ corresponding to (E). In this correspondence, central extensions correspond to G-modules, A, with trivial action ((E) is central iff $A \subseteq Z(\mathcal{G})$). An extension of any type splits iff $\chi(E) = 0$ in $H^2(G, A)$. ((E) is split iff there is a group homomorphism $s\colon G \to \mathcal{G}$ so that $\pi \circ s = \mathrm{id}$).

Proof. We just have to prove the last two facts. Note that the type of extension is trivial iff
$$(\forall \sigma \in G)(\forall a \in A)(\sigma \cdot a = a)$$
iff
$$(\forall x \in \mathcal{G})(\forall a \in A)(x^{-1} a x = a)$$
iff
$$(\forall x \in \mathcal{G})(\forall a \in A)([x, a] = 1)$$
iff $A \subseteq Z(\mathcal{G})$.

Finally, the cohomology is trivial iff every cocycle is a coboundary iff every cocycle is cohomologous to 0 iff in (E) there is a map $\sigma \mapsto u_\sigma$ with $f(\sigma, \tau) = 0$. Such a map is a homomorphism. Thus, $\chi(E) = 0$ in $H^2(G, A)$ iff (E) has a splitting. □

Examples. (I) Find all extensions
$$0 \longrightarrow \mathbb{Z} \longrightarrow \mathcal{G} \longrightarrow \mathbb{Z}/2\mathbb{Z} \longrightarrow 0.$$

There are several cases to consider depending on the type and the cohomology class of the extension.

(a) Trivial type (the action of $\mathbb{Z}/2\mathbb{Z}$ on $\mathbb{Z}$ is trivial).

(a1) Split extension. We get $\mathcal{G} \xrightarrow{\sim} \mathbb{Z} \coprod (\mathbb{Z}/2\mathbb{Z})$.

(a2) Nonsplit extensions. In this case, we have to compute $H^2(\mathbb{Z}/2\mathbb{Z}, \mathbb{Z})$ (trivial action). We know from previous work that (up to cohomology) we can restrict ourselves to *normalized cochains*, $f(\sigma, \tau)$, i.e., cochains such that
$$f(\sigma, 1) = f(1, \sigma) = 0.$$

Elements in $\mathbb{Z}/2\mathbb{Z}$ are ±1. We need to know what $f(-1, -1)$ is. The reader should check that the cocycle condition, $\delta f = 0$, gives no condition on the integer $f(-1, -1)$, and thus, we have an isomorphism $Z^2(\mathbb{Z}/2\mathbb{Z}, \mathbb{Z}) \cong \mathbb{Z}$.

What about coboundaries: $f = \delta k$? Such k's are also normalized, and so, $k(1) = 0$. We have $k(-1) = b$, for any $b \in \mathbb{Z}$. Since

$$\delta k(\sigma, \tau) = \sigma \cdot k(\tau) - k(\sigma\tau) + k(\sigma),$$

using the fact that the action is trivial and that $k(1) = 0$, we get

$$\delta k(-1, -1) = (-1) \cdot k(-1) - k(1) + k(-1) = k(-1) + k(-1) = 2b.$$

So, we can adjust f, up to parity by coboundaries, and $H^2(\mathbb{Z}/2\mathbb{Z}, \mathbb{Z}) \cong \mathbb{Z}/2\mathbb{Z}$. Consequently, we have exactly one nonsplit, trivial-type extension

$$\mathcal{G} = \{(n, \pm 1) \mid n \in \mathbb{Z}\}.$$

The group operation is given by

$$\begin{aligned}(n, \pm 1)(m, 1) &= (n + m, \pm 1)\\ (n, 1)(m, \pm 1) &= (n + m, \pm 1)\\ (n, -1)(m, -1) &= (n + m + 1, 1),\end{aligned}$$

where in this last equation, we assumed without loss of generality that $f(-1, -1) = 1$.

(b) Nontrivial type. We need a nontrivial map $\mathbb{Z}/2\mathbb{Z} \longrightarrow \mathrm{Aut}(\mathbb{Z})$. Since $\mathbb{Z}$ is generated by 1 and -1, there is only one nontrivial action:

$$(-1) \cdot n = -n.$$

(Recall that $1 \cdot n = n$, always).

(b1) The split, nontrivial type extension. In this case

$$\mathcal{G} = \{(n, \sigma) \mid n \in \mathbb{Z},\ \sigma \in \mathbb{Z}/2\mathbb{Z}\},$$

with multiplication given by

$$(n, \sigma)(m, \tau) = (n + \sigma \cdot m, \sigma\tau).$$

Now, consider the map

$$(n, \sigma) \mapsto \begin{pmatrix} \sigma & n \\ 0 & 1 \end{pmatrix}.$$

Observe that matrix multiplication yields

$$\begin{pmatrix} \sigma & n \\ 0 & 1 \end{pmatrix}\begin{pmatrix} \tau & m \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} \sigma\tau & n + \sigma \cdot m \\ 0 & 1 \end{pmatrix}.$$

Therefore, $\mathcal{G}$ is isomorphic to the group of matrices

$$\begin{pmatrix} \sigma & n \\ 0 & 1 \end{pmatrix}$$

under matrix product. This is a nonabelian group, it is infinite and we claim that $\mathcal{G}$ is solvable with $\delta(\mathcal{G}) = 2$.

Indeed, we have $\mathcal{G}/\mathbb{Z} \cong \mathbb{Z}/2\mathbb{Z}$, an abelian group, and so $\Delta^{(1)}(\mathcal{G}) \subseteq \mathbb{Z}$. So, $\Delta^{(2)}(\mathcal{G}) \subseteq \Delta^{(1)}(\mathbb{Z}) = \{0\}$, and we conclude that $\delta(\mathcal{G}) = 2$.

(b2) Nonsplit, nontrivial type extension. We need to figure out what the cocycles are in order to compute $H^2(\mathbb{Z}/2\mathbb{Z}, \mathbb{Z})$. By the same reasoning as before, we need to know what is $f(-1, -1)$. We know that $\delta f(\sigma, \tau) = 0$. So, we have

$$\sigma \cdot f(\tau, \rho) - f(\sigma\tau, \rho) + f(\sigma, \tau\rho) - f(\sigma, \tau) = 0.$$

Let $\tau = \rho = -1$ in the above equation. We get

$$\sigma \cdot f(-1,-1) - f(-\sigma,-1) + f(\sigma,1) - f(\sigma,-1) = \sigma \cdot f(-1,-1) - f(-\sigma,-1) - f(\sigma,-1) = 0,$$

since $f(\sigma,1) = 0$. If we let $\sigma = -1$, since $f(1,-1) = 0$, we get

$$-f(-1,-1) - f(-1,-1) = 0,$$

and so, $2f(-1,-1) = 0$. Since $f(-1,-1) \in \mathbb{Z}$, we get $f(-1,-1) = 0$. Therefore, $f \equiv 0$ and the cohomology is trivial: $H^2(\mathbb{Z}/2\mathbb{Z}, \mathbb{Z}) = (0)$ (for nontrivial action).

As a conclusion, there exist three extension classes and three distinct groups, two of them abelian, the third solvable and faithfully representable by matrices.

(II) Let V be a finite dimensional vector space and consider V^+ as additive group. Let $G = \mathrm{GL}(V)$ and let the action of G on V be the natural one (i.e, for any $\varphi \in \mathrm{GL}(V)$ and any $v \in V$, $\varphi \cdot v = \varphi(v)$). We have the split extension

$$0 \to V \to \mathcal{G} \rightleftarrows \mathrm{GL}(V) \to 0.$$

The group, $\mathcal{G}$, in the above exact sequence is the *affine group* of V.

(III) Again, we restrict ourselves to split extensions. Let A be any abelian group and let $n \in \mathbb{N}$. The group

$$\underbrace{A \prod A \prod \cdots \prod A}_{n}$$

is acted on by the symmetric group, $\mathfrak{S}_n$, simply by permuting the factors. We have a split extension

$$0 \to \underbrace{A \prod A \prod \cdots \prod A}_{n} \to \mathcal{G} \rightleftarrows \mathfrak{S}_n \to 0.$$

The group, $\mathcal{G}$, is called the *wreath product* of A by $\mathfrak{S}_n$ and is denoted $A \wr \mathfrak{S}_n$. We denote the split extension of a given type of G by A by $A \mathbin{\times\!\!\!\times} G$ (note that this notation does not refer to the action).

Here are some useful facts on cohomology:

(1) If G is arbitrary and A is n-torsion, which means that $nA = 0$, then $H^2(G, A)$ is n-torsion.

(2) If G is a finite group, say $\#(G) = g$ and A is arbitrary, then $H^2(G, A)$ is g-torsion (this is not trivial to prove!).

(3) Suppose that A is n-torsion and G is finite, with $\#(G) = g$, and suppose that $(g, n) = 1$. Then, $H^2(G, A) = (0)$. (This is a clear consequence of (1) and (2).)

(4) Suppose that G is finite. We can define a homomorphism (of G-modules) $A \longrightarrow A$, called the *G-norm* and denoted $\mathcal{N}_G$ (we will usually drop the subscript G), defined by

$$\mathcal{N}_G(a) = \sum_{\sigma \in G} \sigma \cdot a.$$

Moreover, assume that G is a *finite cyclic group*. Then, for any A, there is an isomorphism

$$A^G/\mathcal{N}A \cong H^2(G, A),$$

where

$$A^G = \{a \in A \mid \sigma \cdot a = a, \quad \text{for all } \sigma \in G\}.$$

Here is an example of how to use the above facts.

(IV) Find all the groups of order pq (with p, q prime and $0 < p < q$).

We know that the q-Sylow subgroup is normal, namely, it is $\mathbb{Z}/q\mathbb{Z} = A \lhd \mathcal{G}$, and $G = \mathcal{G}/A = \mathbb{Z}/p\mathbb{Z}$. Therefore, whatever $\mathcal{G}$ is, it fits in the group extension

$$0 \longrightarrow \mathbb{Z}/q\mathbb{Z} \longrightarrow \mathcal{G} \longrightarrow \mathbb{Z}/p\mathbb{Z} \longrightarrow 0.$$

By (3), since $(p, q) = 1$, we have $H^2(G, A) = (0)$. So, we only have split extensions. What is $\mathrm{Aut}(\mathbb{Z}/q\mathbb{Z})$? Clearly, it is $\mathbb{Z}/(q-1)\mathbb{Z}$. So, we have to consider the homomorphisms

$$\mathbb{Z}/p\mathbb{Z} \longrightarrow \mathrm{Aut}(\mathbb{Z}/q\mathbb{Z}) = \mathbb{Z}/(q-1)\mathbb{Z}. \tag{$*$}$$

If $(*)$ is non-trivial, then $p \mid (q-1)$, i.e., $q \equiv 1 \pmod{p}$. So, if $q \not\equiv 1 \pmod{p}$, then we have trivial action and we find that

$$\mathcal{G} \cong (\mathbb{Z}/q\mathbb{Z}) \coprod (\mathbb{Z}/q\mathbb{Z}) \cong \mathbb{Z}/pq\mathbb{Z}.$$

If $q \equiv 1 \pmod{p}$, we also can have trivial action, and we get $\mathbb{Z}/pq\mathbb{Z}$, again. So, we now consider nontrivial actions. The unique cyclic group of order p in $\mathbb{Z}/(q-1)\mathbb{Z}$ is generated by $\lambda\frac{q-1}{p}$, where $\lambda = 1, 2, \ldots, p-1$. If we send $1 \in \mathbb{Z}/p\mathbb{Z}$ to $\lambda\frac{q-1}{p}$, the corresponding action is

$$n \mapsto n\lambda\frac{q-1}{p} \pmod{q}.$$

Thus, there are $p-1$ nontrivial (split) group extensions, (E_λ), with central groups

$$\mathcal{G}_\lambda = \{(n, \zeta^m) \mid 0 \leq m \leq p-1\}$$

(here the elements of $\mathbb{Z}/p\mathbb{Z}$ are $1, \zeta, \zeta^2, \ldots, \zeta^{p-1}$) and multiplication given by

$$(n, \zeta^m)(r, \zeta^s) = \left((n + rm\lambda\frac{q-1}{p}, \zeta^{m+s}\right).$$

Consider the map $\mathcal{G}_\lambda \longrightarrow \mathcal{G}_1$ given by

$$(n, \zeta^m) \mapsto (m, \zeta^{\lambda m}).$$

This is a group isomorphism. So, here we have all *inequivalent extensions*, (E_λ), with $p-1$ different actions, yet the groups $\mathcal{G}_\lambda$ are mutually isomorphic. Thus, $\mathcal{G}_1$ and $\mathbb{Z}/pq\mathbb{Z}$ are the two groups of order pq when $q \equiv 1 \pmod{p}$.

The second cohomology group, $H^2(G, A)$, has appeared naturally in the solution to the group extension problem. Consequently, it is natural at this stage to define cohomology groups in general.

The set up is: We have a group, G, and a G-action, $G \prod A \longrightarrow A$, where A is an abelian group. For every $n \in \mathbb{N}$, we define

$$C^n(G, A) = \{f \colon \underbrace{G \prod \cdots \prod G}_{n} \to A\},$$

where $\underbrace{G \prod \cdots \prod G}_{n}$ is the product of G with itself n times (in the category of sets). By convention, when $n = 0$, this set product is the one point set, $\{*\}$. The set $C^n(G, A)$ is an abelian group under addition of functions (*e.g*, $f + g$ is the function defined by $(f + g)(x) = f(x) + g(x)$ for all $x \in G$). The group

$C^n(G,A)$ is called the group of *n-cochains of G with coefficients in A*. We define the *coboundary map*, $\delta_n\colon C^n(G,A)\to C^{n+1}(G,A)$, for every $n\geq 0$, by the formula:

$$\begin{aligned}(\delta_n f)(\sigma_1,\ldots,\sigma_{n+1}) &= \sigma_1\cdot f(\sigma_2,\ldots,\sigma_{n+1})+\sum_{j=1}^{n}(-1)^j f(\sigma_1,\ldots,\sigma_{j-1},\sigma_j\sigma_{j+1},\sigma_{j+2},\ldots,\sigma_{n+1})\\ &\quad +(-1)^{n+1}f(\sigma_1,\ldots,\sigma_n),\end{aligned}$$

for all $f\in C^n(G,A)$ and all $\sigma_1,\ldots,\sigma_{n+1}\in G$.

(1) Check (DX): For all $n\geq 0$,

$$\delta_n(\delta_{n-1}f)\equiv 0.$$

(By convention, $\delta_{-1}=0$).

(2) Set $Z^n(G,A)=\mathrm{Ker}\ \delta_n$, a subgroup of $C^n(G,A)$, the group of *n-cocycles of G with coefficients in A*. We also let $B^n(G,A)=\mathrm{Im}\ \delta_{n-1}$, a subgroup of $C^n(G,A)$, the group of *n-coboundaries of G with coefficients in A*. Observe that since $\delta_{-1}=0$, we have $B^0(G,A)=(0)$. Furthermore, (1) implies that $B^n(G,A)\subseteq Z^n(G,A)$, for all $n\geq 0$.

(3) Set $H^n(G,A)=Z^n(G,A)/B^n(G,A)$; this is the *nth cohomology group of G with coefficients in A*.

Examples. (i) Case $n=0$: Then, $B^0=(0)$. The functions, f, in $C^0(G,A)$ are in one-to-one correspondence with the elements $f(*)$ of A, and so, $C^0(G,A)=A$. Note that for any $\sigma\in G$, if $f\in C^0(G,A)$ corresponds to the element a in A, we have

$$(\delta_0 f)(\sigma)=\sigma\cdot f(*)-f(*)=\sigma\cdot a-a.$$

Thus,

$$Z^0(G,A)=\{a\in A\mid \delta_0(a)=0\}=\{a\in A\mid \sigma\cdot a=a,\ \text{for all}\ \sigma\in G\}=A^G.$$

So, we also have $H^0(G,A)=A^G$.

(ii) Case $n=1$: Then, $C^1(G,A)$ is the set of all functions $f\colon G\to A$. For any $f\in C^1(G,A)$, we have

$$(\delta_1 f)(\sigma,\tau)=\sigma\cdot f(\tau)-f(\sigma\tau)+f(\sigma).$$

It follows that

$$Z^1(G,A)=\{f\in C^1(G,A)\mid \delta_1 f=0\}=\{f\in C^1(G,A)\mid f(\sigma\tau)=\sigma\cdot f(\tau)+f(\sigma)\}.$$

This is the set of *crossed* (or *twisted*) *homomorphisms* from G to A.

Remark: If A has trivial G-action, then $Z^1(G,A)=\mathrm{Hom}_{\mathcal{G}\mathrm{r}}(G,A)$.

We have $B^1(G,A)=\mathrm{Im}\ \delta_0=$ all functions, g, so that $g(\sigma)=(\delta_0(a))(\sigma)=\sigma\cdot a-a$, for some $a\in A$. Such objects are twisted homomorphisms, called *principal* (or *inner*) *twisted homomorphisms*.

Remark: If A has trivial G-action, then $B^1(G,A)=(0)$. So, $H^1(G,A)$ is the quotient of the twisted homomorphisms modulo the principal twisted homomorphisms if the action is nontrivial, and $H^1(G,A)=\mathrm{Hom}_{\mathcal{G}\mathrm{r}}(G,A)$ if the action is trivial.

(iii) Case $n=2$: We have already encountered this case in dealing with group extensions. We content ourselves with computing δ_2. Since $C^2(G,A)=\{f\colon G\prod G\to A\}$, we have

$$(\delta_2 f)(\sigma,\tau,\rho)=\sigma\cdot f(\tau,\rho)-f(\sigma\tau,\rho)+f(\sigma,\tau\rho)-f(\sigma,\tau).$$

We note that $Z^2(G, A)$ gives us back the group of "old" 2-cocycles, $B^2(G, A)$ gives us back the group of "old" 2-coboundaries, and $H^2(G, A)$ is in one-to-one correspondence with the equivalence classes of group extensions of a fixed type.

Remark: Given a group, G, Eilenberg and Mac Lane (1940's) constructed a topological space, $K(G, 1)$, unique up to homotopy type, with the following properties:

$$\pi_n(K(G,1)) = \begin{cases} G & \text{if } n = 1 \\ (0) & \text{if } n \neq 1. \end{cases}$$

Fact: If we compute the integral cohomology of $K(G, 1)$, denoted $H^n(K(G, 1), \mathbb{Z})$, we get

$$H^n(K(G,1), \mathbb{Z}) \cong H^n(G, \mathbb{Z}).$$

Here, the G-action on $\mathbb{Z}$ is trivial.

1.5 Solvable and Nilpotent Groups

Given a group, G, its derived series,

$$G = \Delta^{(0)}(G) \supseteq \Delta^{(1)}(G) \supseteq \Delta^{(2)}(G) \supseteq \cdots \supseteq \Delta^{(t)}(G) \supseteq \cdots ,$$

may decrease very quickly, and even though the solvable groups (those for which the derived series reaches $\{1\}$ after finitely many steps, i.e., those for which $\delta(G)$ is finite) are not as "wild" as groups for which $\delta(G) = \infty$, it desirable to delineate families of groups with an even "nicer" behavior. One way of doing so is to define descending (or ascending) chains that do not decrease (or increase) too quickly and then to investigate groups whose chains are finite. The collection of nilpotent groups is such a family of groups, and, moreover, nilpotent groups tend to show up as fundamental groups of spaces arising naturally in geometry. Every nilpotent group is solvable and solvability is inherited by subgroups and quotient groups, as shown in the following proposition:

Proposition 1.28 *If G is a group and G is solvable, then for every subgroup, H, of G, the group, H, is solvable. Moreover, if H is normal in G, then G/H is solvable. In fact, for both groups, $\delta(\text{either}) \leq \delta(G)$. Conversely, say G possesses a normal subgroup, H, so that both H and G/H are solvable. Then, G is solvable. In fact, $\delta(G) \leq \delta(H) + \delta(G/H)$.*

Proof. Let G be solvable. Then, $H \subseteq G$ implies $\Delta^{(1)}(H) \subseteq \Delta^{(1)}(G)$; therefore (by induction),

$$\Delta^{(j)}(H) \subseteq \Delta^{(j)}(G),$$

and we deduce that $\delta(H) \leq \delta(G)$. Consider $\overline{G} = G/H$ when $H \lhd G$. Then, $\overline{[x,y]} = [\overline{x}, \overline{y}]$ and this implies $\overline{\Delta^{(1)}(G)} = \Delta^{(1)}(\overline{G})$. Hence (by induction),

$$\overline{\Delta^{(j)}(G)} = \Delta^{(j)}(\overline{G}).$$

Therefore, $\delta(\overline{G}) \leq \delta(G)$.

Conversely, assume that H and G/H are solvable (with $H \lhd G$). We have $\overline{\Delta^{(j)}(G)} = \Delta^{(j)}(\overline{G})$ and if $j \geq \delta(\overline{G})$, then $\overline{\Delta^{(j)}(G)} = \{1\}$, which implies that $\Delta^{(j)}(G) \subseteq H$. So, $\Delta^{(k+j)}(G) \subseteq \Delta^{(k)}(H)$, and the latter is $\{1\}$ if $k = \delta(H)$. Therefore,

$$\Delta^{(\delta(\overline{G})+\delta(H))}(G) = \{1\},$$

and so, $\delta(G) \leq \delta(H) + \delta(G/H)$. □

Proposition 1.29 *Let (P) be some property of finite groups. Assume that (P) satisfies:*

(a) The trivial group has (P), every cyclic group of prime order has (P).

(b) Suppose G has (P), then $H \lhd G$ implies H and G/H have (P).

(c) If G has (P) (with $G \neq \{1\}$), then G is not simple unless G is cyclic of prime order.

Then, when G has (P), the group G is solvable.

Proof. We proceed by induction on $\#(G)$. The case $G = \{1\}$ is trivial, by (a) (nothing to check). Assume that the proposition holds for all G with $\#(G) \leq n$, and assume $\#(G) = n+1$. If $n+1$ is prime, then G is cyclic of prime order, which implies that it is solvable. Thus, we may assume that $n+1$ is not prime and that G has (P). By (c), the group G has some nontrivial normal subgroup, H. By (b), both H and G/H have (P), and the induction hypothesis implies that both H and G/H are solvable. Proposition 1.28 implies that G is solvable. □

Corollary 1.30 *(Burnside, Feit & Thompson) Every group G, of order $p^a q^b$ or odd order is solvable.*

Remark: Corollary 1.30 is not really proved. It depends on establishing (c) for the two properties: $p^a q^b$, odd order. As remarked just before Proposition 1.10, this is not easy.

Definition 1.9 Let G be any group. The *lower central series (LCS)* of G is the descending chain of subgroups

$$G = \Gamma_0 \supseteq \Gamma_1 \supseteq \cdots \supseteq \Gamma_d \supseteq \cdots ,$$

where $\Gamma_{j+1} = [G, \Gamma_j]$. The *upper central series (UCS)* of G is the ascending chain of subgroups

$$\{1\} = Z_0 \subseteq Z_1 \subseteq Z_2 \subseteq \cdots \subseteq Z_d \subseteq \cdots ,$$

where $Z_j =$ the inverse image in G of $Z(G/Z_{j-1})$.

Remarks:

(1) $\Gamma_1(G) = [G, \Gamma_0] = [G, G] = \Delta^{(1)}(G)$, and

$$\Gamma_2(G) = [G, \Gamma_1] = [G, \Delta^{(1)}(G)] \supseteq [\Delta^{(1)}(G), \Delta^{(1)}(G)] = \Delta^{(2)}(G),$$

and so, $\Gamma_2(G) \supseteq \Delta^{(2)}(G)$. The reader should check (DX) that $\Gamma_d(G) \supseteq \Delta^{(d)}(G)$, for all $d \geq 0$.

(2) $Z_1(G) =$ inverse image in G of $Z(G/Z_0) =$ inverse image of $Z(G)$, so $Z_1(G) = Z(G)$.

(3) If for some j, the equality $\Gamma_j(G) = \Gamma_{j+1}(G)$ holds, then $\Gamma_j(G) = \Gamma_d(G)$, for all $d \geq j$. The lower central series strictly descends until the first repetition.

(4) Similarly, if for some j, the equality $Z_j(G) = Z_{j+1}(G)$ holds, then $Z_j(G) = Z_d(G)$, for all $d \geq j$. The upper central series strictly ascends until the first repetition.

Proposition 1.31 *Suppose the lower central series of G reaches $\{1\}$ after r steps. Then, for every $j \leq r$, we have $\Gamma_{r-j} \subseteq Z_j$. Consequently, the upper central series reaches G after r steps. Conversely, suppose that the upper central series reaches G after r steps. Then, for every $j \leq r$, we have $\Gamma_j \subseteq Z_{r-j}$. Consequently, the lower central series reaches $\{1\}$ after r steps.*

Proof. By induction on j. For $j = 0$, we have $\Gamma_r = \Gamma_{r-0}$, and by hypothesis, $\Gamma_r = \{1\}$ and $Z_0 = \{1\}$, so the basis of the induction holds. Before we do the induction step, let us also consider the case $j = 1$. We need to show that $\Gamma_{r-1} \subseteq Z_1 = Z(G)$. But $\Gamma_r = \{1\}$, yet $\Gamma_r = [G, \Gamma_{r-1}]$. This means that for all $\sigma \in G$ and all $\tau \in \Gamma_{n-1}$, we have $[\sigma, \tau] \in \Gamma_r = \{1\}$. Thus, τ commutes with all $\sigma \in G$, and so, $\tau \in Z(G) = Z_1$. Let us now assume our statement, $\Gamma_{r-j} \subseteq Z_j$, for some j, and look at the case $j+1$. Now, $\Gamma_{r-j} = [G, \Gamma_{r-j-1}]$. By the induction hypothesis,

$$[G, \Gamma_{r-j-1}] \subseteq Z_j.$$

Consider the map $G \longrightarrow G/Z_j = \overline{G}$. Then,

$$[\overline{G}, \overline{\Gamma}_{r-j-1}] = \{1\} \quad \text{in } \overline{G}.$$

Therefore, Γ_{r-j-1} is contained in the inverse image of $Z(\overline{G}) = Z(G/Z_j) = Z_{j+1}$, concluding the induction step.

For the converse, again, use induction on j. When $j = 0$, we have $\Gamma_0 = G$ and $Z_r = Z_{r-0} = G$, by hypothesis, and the basis of the induction holds. Assume that $\Gamma_j \subseteq Z_{r-j}$ for some j, and consider the case $j+1$. We have

$$\Gamma_{j+1} = [G, \Gamma_j] \subseteq [G, Z_{r-j}],$$

by the induction hypothesis. Look at the map $G \longrightarrow G/Z_{r-j-1} = \overline{G}$. We have

$$\overline{\Gamma}_{j+1} \subseteq [\overline{G}, \overline{Z}_{r-j}].$$

But, by definition, $\overline{Z}_{r-j} = Z(\overline{G})$. Thus, $[\overline{G}, \overline{Z}_{r-j}] = \{1\}$ in $\overline{G}$. Therefore,
$\Gamma_{j+1} \subseteq \text{Ker } (G \longrightarrow \overline{G}) = Z_{r-j-1}$. □

Definition 1.10 A group, G, is *nilpotent* if and only if the lower central series reaches $\{1\}$ after finitely many steps. The smallest number of steps, say c, is the *nilpotence class* of G. We write $G \in \mathcal{N}\mathrm{ilp}(c)$. (We let $c = \infty$ if the LCS does not reach $\{1\}$ in finitely many steps.)

Remarks:

(1) $\mathcal{N}\mathrm{ilp}(0) =$ the class consisting only of the trivial group.
$\mathcal{N}\mathrm{ilp}(1) =$ the collection of abelian, nontrivial groups. If we let $\overline{\mathcal{N}\mathrm{ilp}}(c)$ denote the union of the collections $\mathcal{N}\mathrm{ilp}(k)$ for $k = 0, \ldots, c$, then it turns out that we have a strictly ascending chain

$$\mathcal{A}\mathrm{b} = \overline{\mathcal{N}\mathrm{ilp}}(1) < \overline{\mathcal{N}\mathrm{ilp}}(2) < \overline{\mathcal{N}\mathrm{ilp}}(3) < \cdots$$

of "worse and worse behaved" groups.

(2) We have $G \in \mathcal{N}\mathrm{ilp}(c)$ iff the UCS reaches G after c steps and c is minimal with this property.

(3) Each nilpotent group is automatically solvable, but the converse is false, even for finite groups, even for small finite groups. Indeed, we observed earlier that $\Delta^{(r)}(G) \subseteq \Gamma_r(G)$. Therefore, $\delta(G) \leq$ nilpotence class of G. For a counter-example, take $G = \mathfrak{S}_3$. This group has order 6, its center is trivial, and so $Z_1 = Z_0$ and G is *not* nilpotent. Yet, we have an exact sequence

$$0 \longrightarrow \mathbb{Z}/3\mathbb{Z} \longrightarrow \mathfrak{S}_3 \longrightarrow \mathbb{Z}/2\mathbb{Z} \longrightarrow 0,$$

and the extremes are solvable (even nilpotent, even abelian), so the middle is solvable.

(4) Every p-group is nilpotent. This is because the center of a p-group is nontrivial, so the UCS is strictly ascending and our group is finite; so, this implies that our group is nilpotent.

Remark: The fundamental groups of many spaces arising in geometry tend to be nilpotent groups.

Proposition 1.32 *(Modified Sylow III) Say G is a finite group, P is a p-Sylow subgroup of G and H is some subgroup of G. If $H \supseteq N_G(P)$, then $N_G(H) = H$.*

Proof. (Frattini Argument). Pick $\sigma \in N_G(H)$. Then, $\sigma H \sigma^{-1} = H$ and $\sigma P \sigma^{-1} \subseteq \sigma H \sigma^{-1}$ (since $H \supseteq N_G(P)$). So, P and $\sigma P \sigma^{-1}$ are two p-Sylow subgroups of H, and by Sylow II, there is some $\tau \in H$ so that $\tau P \tau^{-1} = \sigma P \sigma^{-1}$. Thus, $\tau^{-1}\sigma P(\tau^{-1}\sigma)^{-1} = P$, and so, $\tau^{-1}\sigma \in N_G(P) \subseteq H$, by hypothesis. So, $\sigma \in \tau H = H$ (since $\tau \in H$). □

Theorem 1.33 *Let G be a finite group. Then, the following statements are equivalent:*

(1) G is nilpotent.

(2) G has property (N).

(3) Every maximal subgroup of G is normal.

(4) $\Delta^{(1)}(G) \subseteq \Phi(G)$.

(5) Every p-Sylow subgroup of G is normal in G.

(6) G is isomorphic to the product of its p-Sylow subgroups. (We write $G \cong \prod_p G_p$).

Proof. (1) ⇒ (2). Let H be a proper subgroup of G, we must prove that $N_G(H) > H$. Now, there is some c with $\Gamma_c = \{1\}$. Obviously, $\Gamma_c \subseteq H$, so pick a smallest d for which $\Gamma_d \subseteq H$, so that $\Gamma_{d-1} \not\subseteq H$.

Claim: $\Gamma_{d-1} \subseteq N_G(H)$.

If the claim holds, then $H < N_G(H)$, *i.e.*, G has property (N). Pick $\xi \in \Gamma_{d-1}$; so,

$$[H, \xi] \subseteq [H, \Gamma_{d-1}] \subseteq [G, \Gamma_{d-1}] = \Gamma_d.$$

Pick $h \in H$ and look at $[h^{-1}, \xi]$. The element $[h^{-1}, \xi]$ is in Γ_d, and so, in H (since $\Gamma_d \subseteq H$). Consequently, $h^{-1}\xi h\xi^{-1} \in H$, from which we deduce $\xi h\xi^{-1} \in H$, and since this is true for all $h \in H$, we have $\xi \in N_G(H)$, as desired.

(2) ⇒ (3). This has already been proved (c.f. Proposition 1.17).

(3) ⇒ (4). This has already been proved (c.f. Proposition 1.23).

(4) ⇒ (5). Let P be a p-Sylow subgroup of G. Look at $N_G(P)$. If $N_G(P) \neq G$, then $N_G(P)$ is contained is some maximal subgroup, M. By modified Sylow III, we get $N_G(M) = M$. Now, $\Delta^{(1)}(G) \subseteq \Phi(G) \subseteq M$, by hypothesis, and the second homomorphism theorem implies that M corresponds to a subgroup of $G/\Delta^{(1)}(G)$ and normal subgroups correspond to normal subgroups. Yet, $G/\Delta^{(1)}(G)$ is abelian, so all its subgroups are normal, which implies that M is normal, a contradiction.

(5) ⇒ (6). This has already been proved (c.f. Proposition 1.18).

(6) ⇒ (1). Since every p-group is nilpotent, the implication (6) ⇒ (1) follows from the following

Proposition 1.34 *Say $G_j \in \mathcal{N}\mathrm{ilp}(c_j)$, for $j = 1, \ldots, t$. Then,*

$$\prod_{j=1}^{t} G_j \in \mathcal{N}\mathrm{ilp}(\max_{1\leq j\leq t}\{c_j\}).$$

Proof. An obvious induction reduces us to the case $t = 2$. In this case, we use an induction on $\max\{c_1, c_2\}$. The cases $c_1 \leq 1$ and $c_2 \leq 1$ are trivial. Now, we have (DX)

$$Z(G_1 \prod G_2) \cong Z(G_1) \prod Z(G_2).$$

But then, $(G_1 \prod G_2)/Z(G_1 \prod G_2) \cong (G_1/Z(G_1)) \prod(G_2/Z(G_2))$; on the left hand side, the purported nilpotence class is down by 1 and on the righthand side, both are down by 1. We conclude by applying the induction hypothesis. □

This concludes the proof of Theorem 1.33. □

1.6 Ω-Groups and the Jordan-Hölder-Schreier Theorem

Let Ω be some set. If M is a group, we denote the monoid of group endomorphisms of M (under composition) by $\mathrm{End}_{\mathcal{G}\mathrm{r}}(M)$ and the group of (group) automorphisms of M by $\mathrm{Aut}_{\mathcal{G}\mathrm{r}}(M)$.

Definition 1.11 A group, M, is an *Ω-group* iff there exists a set map $\Omega \longrightarrow \mathrm{End}_{\mathcal{G}\mathrm{r}}(M)$. If Ω is itself a group, we demand that our map be a homomorphism (so, the image lies in $\mathrm{Aut}_{\mathcal{G}\mathrm{r}}(M)$). If Ω is a ring, we demand that M be an abelian group and that our map be a ring homomorphism taking $1 \in \Omega$ to the identity endomorphism of M.

Examples.

(1) When Ω is a group, we get an Ω-action on M (at first, as a set) and further, we obtain:

$$\begin{aligned} 1 \cdot m &= m, \\ \xi \cdot (\eta \cdot m) &= (\xi\eta) \cdot m, \\ \xi \cdot (mn) &= (\xi \cdot m)(\eta \cdot n). \end{aligned}$$

In particular, $(\xi \cdot m)^{-1} = \xi \cdot m^{-1}$.

(2) When Ω is a group and M is abelian, we just get an Ω-*module*.

(3) If Ω is a ring, then the nomenclature is Ω-*module* instead of Ω-group.

(4) When Ω is a field, then an Ω-module is a vector space over Ω.

(5) Being an $\mathbb{Z}$-module is equivalent to being an abelian group.

An Ω-*subgroup* of M (resp. Ω-*normal subgroup* of M) is just a subgroup (resp. a normal subgroup), N, of M stable under Ω, *i.e.*, for all $\xi \in \Omega$, for all $n \in N$, we have $\xi \cdot n \in N$.

Blanket Assertion (DX). The three isomorphism theorems of ordinary group theory are true for Ω-groups provided everywhere "subgroup" appears we substitute "Ω-subgroup", *mutatis–mutandis* for "normal subgroups."

Definition 1.12 A *normal flag* (*normal series, normal chain*) is a descending chain of Ω-subgroups of M:

$$M = M_0 \supseteq M_1 \supseteq M_2 \supseteq \cdots \supseteq M_r = \{1\}, \qquad (*)$$

each M_j being normal in the preceding M_{j-1}. A normal flag is *nonrepetitious* if for no j do we have $M_j = M_{j-1}$. Given a second normal flag:

$$M = M'_0 \supseteq M'_1 \supseteq M'_2 \supseteq \cdots \supseteq M'_s = \{1\}, \qquad (**)$$

the flag $(**)$ *refines* $(*)$ iff for every i the Ω-group M_i occurs as some M'_j. Two normal flags $(*)$ and $(**)$ are isomorphic iff the collection of their successive quotients, M_{i-1}/M_i and M'_{j-1}/M'_j may be rearranged so that, after rearrangement, they become pairwise isomorphic (in their new order). When this happens, the lengths r and s are equal.

Theorem 1.35 *(Schreier refinement theorem, 1928) For an Ω-group, any two normal flags possess isomorphic refinements. If both normal flags are nonrepetitious, so are their isomorphic refinements.*

The main corollary of the Schreier refinement theorem is:

Corollary 1.36 *(Jordan–Hölder theorem) Any two composition series for an Ω-group are isomorphic.*

Proof. A composition series has *no* refinements except itself—apply Schreier's theorem. □

Zassenhaus proved a lemma specifically designed to give the smoothest proof of Schreier's theorem—this is

Lemma 1.37 *(Zassenhaus' butterfly lemma) Say G is an Ω-group and A and C are subgroups. Suppose $B \lhd A$ and $D \lhd C$ are further Ω-subgroups. Then,*

$$(A \cap C)B/(A \cap D)B \cong (C \cap A)D/(C \cap B)D.$$

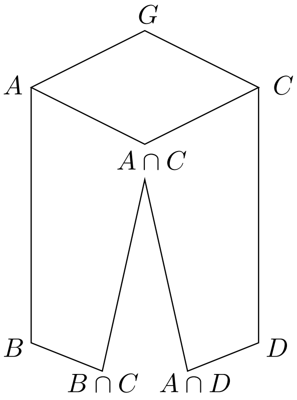

Figure 1.1: The butterfly lemma

Proof. Let $T = A \cap C = C \cap A$, $M = B \cap C$ and $N = A \cap D$. The conclusion of the lemma is

$$TB/NB \cong TD/MD.$$

First of all, there is right-left symmetry in the statement of the lemma and its conclusion ($A \leftrightarrow C$, $B \leftrightarrow D$; under these substitutions, $T \leftrightarrow T$ and $M \leftrightarrow N$). We must prove that $NB \lhd TB$. Pick $t \in G$ and look at $tNBt^{-1} = tNt^{-1}tBt^{-1}$. If $t \in A$, then $tBt^{-1} = B$, since $B \lhd A$. Thus, *if* $t \in A$ *then* $tNBt^{-1} = tNt^{-1}B$. If $t \in T \subseteq C$, then as $N = D \cap C \cap A = D \cap T$ and $D \lhd C$, we get

$$tNt^{-1} = tDt^{-1} \cap tTt^{-1} = tDt^{-1} \cap T = D \cap T = N.$$

Thus, *if* $t \in T$ *then* $tNBt^{-1} = NB$.

Say $\xi = tb \in TB$. Since $B \lhd A$ and $N \subseteq A$, we have $BN = NB$. Then, we find

$$\begin{aligned}
\xi NB\xi^{-1} &= tbNBb^{-1}t^{-1} \\
&= tbNBt^{-1} \\
&= tbBNt^{-1} \\
&= tBNt^{-1} \\
&= tNBt^{-1} \\
&= NB.
\end{aligned}$$

Therefore, $NB \lhd TB$. By symmetry, we get $MD \lhd TD$. Look at $TB/NB = TNB/BN$ (since $N \subseteq T$). By the third isomorphism theorem, we have

$$TB/NB \cong T/T \cap NB.$$

By symmetry,

$$TD/ND \cong T/T \cap MD.$$

If we prove that $T \cap NB = T \cap NM$ (and so, $T \cap MD = T \cap NM$, by symmetry), we will be done. Pick $\xi \in T \cap NB$. We can write $\xi = nb \in NB$, so $b = n^{-1}\xi \in NT = T$ (since $N \subseteq T$). Thus, $b \in B \cap T = B \cap C \cap A \subseteq M$, and so, $b \in M$. Consequently, $\xi = nb \in NM$ and since we also have $\xi \in T$, then $\xi \in T \cap NM$. This proves that $T \cap NB \subseteq T \cap NM$. The reverse inclusion is trivial, since $M \subseteq B$. Therefore, $T \cap NB = T \cap NM$, as claimed. □

Proof of Theorem 1.35. Let

$$M = M_0 \supseteq M_1 \supseteq M_2 \supseteq \cdots M_{i-1} \supseteq M_i \supseteq \cdots \supseteq M_r = \{1\}, \tag{$*$}$$

and

$$M = M'_0 \supseteq M'_1 \supseteq M'_2 \supseteq \cdots M'_{j-1} \supseteq M'_j \supseteq \cdots \supseteq M'_s = \{1\}, \tag{$**$}$$

be two normal nonrepetitious chains. Consider the groups

$$M_{i-1}^{(j)} = (M_{i-1} \cap M'_j)M_i.$$

As j varies, these groups start at M_{i-1} $(= M_{i-1}^{(0)})$ and end at M_i $(= M_{i-1}^{(s)})$ and we get a refinement of $(*)$ if we do this between any pair in $(*)$. Also consider the groups

$$M_{j-1}^{'(i)} = (M'_{j-1} \cap M_i)M'_j,$$

and let i vary. These groups interpolate between M'_{j-1} and M'_j, just as above. Look at the successive quotients

$$M_{i-1}^{(j-1)}/M_{i-1}^{(j)}; \quad M_{j-1}^{'(i-1)}/M_{j-1}^{'(i)}. \tag{$\dagger$}$$

If we let $A = M_{i-1}$, $B = M_i$ $(\lhd A)$, $C = M'_{j-1}$ and $D = M'_j$ $(\lhd C)$, we can write the first quotient group of $(\dagger)$ as

$$M_{i-1}^{(j-1)}/M_{i-1}^{(j)} = (M_{i-1} \cap M'_{j-1})M_i/(M_{i-1} \cap M'_j)M_i = (A \cap C)B/(A \cap D)B,$$

the left hand side of Zassenhaus' lemma. By symmetry, the second quotient group of $(\dagger)$ is the righthand side of Zassenhaus' lemma and we are done. □

1.7 Categories, Functors and Free Groups

Definition 1.13 A *category*, $\mathcal{C}$, is a pair: $\langle \mathcal{O}\mathrm{b}(\mathcal{C}), \mathcal{F}\mathrm{l}(\mathcal{C})\rangle$, in which $\mathcal{O}\mathrm{b}(\mathcal{C})$ and $\mathcal{F}\mathrm{l}(\mathcal{C})$ are classes, called the *objects of* $\mathcal{C}$ and the *morphisms* (or *arrows*) *of* $\mathcal{C}$, respectively. We require the following conditions:

(1) For all $A, B \in \mathcal{O}\mathrm{b}(\mathcal{C})$, there is a unique **set**, $\mathrm{Hom}_{\mathcal{C}}(A, B)$, called the collection of *morphisms from* A *to* B, and any two such are either disjoint or equal. Further

$$\mathcal{F}\mathrm{l}(\mathcal{C}) = \bigcup_{A,B} \mathrm{Hom}_{\mathcal{C}}(A, B).$$

For the morphisms, we also require:

(2) For every $u \in \mathrm{Hom}_{\mathcal{C}}(A, B)$ and $v \in \mathrm{Hom}_{\mathcal{C}}(B, C)$, there exists a unique morphism $w = v \circ u \in \mathrm{Hom}_{\mathcal{C}}(A, C)$, called the *composition* of v and u.

(3) For every $A \in \mathcal{O}\mathrm{b}(\mathcal{C})$, there is some arrow, $1_A \in \mathrm{Hom}_{\mathcal{C}}(A, A)$, so that for every $B \in \mathcal{O}\mathrm{b}(\mathcal{C})$ and $u \in \mathrm{Hom}_{\mathcal{C}}(A, B)$, we have

$$\begin{aligned} A \xrightarrow{1_A} A \xrightarrow{u} B &= A \xrightarrow{u} B \\ A \xrightarrow{u} B \xrightarrow{1_B} B &= A \xrightarrow{u} B. \end{aligned}$$

Note: This shows that 1_A is unique for each A (DX).

(4) We have the associativity law

$$u \circ (v \circ w) = (u \circ v) \circ w,$$

whenever the compositions all make sense.

Examples of Categories:

(1) $\mathcal{S}$ets, the category of sets; $\mathcal{O}\mathrm{b}(\mathcal{S}\mathrm{ets})$ = all sets, $\mathcal{F}\mathrm{l}(\mathcal{S}\mathrm{ets})$ = all maps of sets.

(2) $\mathcal{G}$r, the category of groups; $\mathcal{O}\mathrm{b}(\mathcal{G}\mathrm{r})$ = all groups, $\mathcal{F}\mathrm{l}(\mathcal{G}\mathrm{r})$ = all homomorphisms of groups. A special case is $\mathcal{A}$b, the category of abelian groups.

(3) Ω-$\mathcal{G}$r, the category of Ω-groups. Special cases are: The category of G-modules, $\mathcal{M}\mathrm{od}(G)$; the category of R-modules, $\mathcal{M}\mathrm{od}(R)$ (where R is a ring); and the category of vector spaces, Vect(k) (where k is a field). Also, $\mathcal{A}\mathrm{b} = \mathcal{M}\mathrm{od}(\mathbb{Z})$.

(4) TOP, the category of topological spaces; $\mathcal{O}\mathrm{b}(\mathrm{TOP})$ = all topological spaces, $\mathcal{F}\mathrm{l}(\mathrm{TOP})$ = all continuous maps.

(5) C^k-MAN, the category of C^k-manifolds; $\mathcal{O}\mathrm{b}(C^k\text{-MAN})$ = all (real) C^k-manifolds ($0 \leq k \leq \infty$ or ω), $\mathcal{F}\mathrm{l}(C^k\text{-MAN})$ = all C^k-maps of C^k-manifolds.

(6) HOL, the category of complex analytic manifolds; $\mathcal{O}\mathrm{b}(\mathrm{HOL})$ = all complex analytic manifolds, $\mathcal{F}\mathrm{l}(\mathrm{HOL})$ = all complex analytic maps of holomorphic manifolds.

(7) RNG, the category of all rings; $\mathcal{O}\mathrm{b}(\mathrm{RNG})$ = all rings (with unity), $\mathcal{F}\mathrm{l}(\mathrm{RNG})$ = all homomorphisms of rings. A special case is CR, the category of commutative rings.

A *subcategory*, $\mathcal{D}$, of $\mathcal{C}$ is a category, $\langle \mathcal{O}\mathrm{b}(\mathcal{D}), \mathcal{F}\mathrm{l}(\mathcal{D})\rangle$, so that

(a) $\mathcal{O}\mathrm{b}(\mathcal{D}) \subseteq \mathcal{O}\mathrm{b}(\mathcal{C})$.

(b) $\mathcal{F}\mathrm{l}(\mathcal{D}) \subseteq \mathcal{F}\mathrm{l}(\mathcal{C})$, in such a way that for all $A, B \in \mathcal{O}\mathrm{b}(\mathcal{D})$, we have $\mathrm{Hom}_{\mathcal{D}}(A, B) \subseteq \mathrm{Hom}_{\mathcal{C}}(A, B)$.

We say that $\mathcal{D}$ is a *full subcategory* of $\mathcal{C}$ iff for all $A, B \in \mathcal{O}\mathrm{b}(\mathcal{D})$, we have $\mathrm{Hom}_{\mathcal{D}}(A, B) = \mathrm{Hom}_{\mathcal{C}}(A, B)$.

Examples of Subcategories:

(1) The category, $\mathcal{A}$b, is a full subcategory of $\mathcal{G}$r; the category, CR, is a full subcategory of RNG.

(2) Recall that $u \in \mathrm{Hom}_{\mathcal{C}}(A, B)$ is an *isomorphism* (in $\mathcal{C}$) iff there is some $v \in \mathrm{Hom}_{\mathcal{C}}(B, A)$ so that

$$\begin{aligned} A \xrightarrow{u} B \xrightarrow{v} A &= A \xrightarrow{1_A} A \\ B \xrightarrow{v} A \xrightarrow{u} B &= B \xrightarrow{1_B} B. \end{aligned}$$

Take $\mathcal{D}$ so that $\mathcal{O}\mathrm{b}(\mathcal{D}) = \mathcal{O}\mathrm{b}(\mathcal{C})$, and morphisms, set $\mathrm{Hom}_{\mathcal{D}}(A, B) = \{u \in \mathrm{Hom}_{\mathcal{C}}(A, B) \mid u \text{ is an isomorphism}\}$ and $\mathcal{F}\mathrm{l}(\mathcal{D}) = \bigcup_{A,B} \mathrm{Hom}_{\mathcal{D}}(A, B)$. (Note that $\mathrm{Hom}_{\mathcal{D}}(A, B)$ may be empty.) The category, $\mathcal{D}$, is generally a nonfull subcategory of $\mathcal{C}$, for example when $\mathcal{C} = \mathcal{S}$ets.

Say $\mathcal{C}$ is a category, we can make a new category, $\mathcal{C}^D$, the *dual* or *opposite category*, as follows: $\mathcal{O}\mathrm{b}(\mathcal{C}^D) = \mathcal{O}\mathrm{b}(\mathcal{C})$ and reverse the arrows, *i.e.*, for all $A, B \in \mathcal{O}\mathrm{b}(\mathcal{C})$,

$$\mathrm{Hom}_{\mathcal{C}^D}(A, B) = \mathrm{Hom}_{\mathcal{C}}(B, A).$$

Definition 1.14 Let $\mathcal{C}$ and $\mathcal{C}'$ be categories. A *functor* (respectively, a *cofunctor*), F, from $\mathcal{C}$ to $\mathcal{C}'$ is a rule which associates to each object $A \in \mathcal{O}\mathrm{b}(\mathcal{C})$ an object $F(A) \in \mathcal{O}\mathrm{b}(\mathcal{C}')$ and to each arrow $u \in \mathrm{Hom}_{\mathcal{C}}(A, B)$ an arrow $F(u) \in \mathrm{Hom}_{\mathcal{C}'}(F(A), F(B))$ (resp. $F(u) \in \mathrm{Hom}_{\mathcal{C}'}(F(B), F(A))$) so that,

$$\begin{aligned} F(1_A) &= 1_{F(A)} \\ F(u \circ v) &= F(u) \circ F(v) \\ (\text{resp.}\quad F(u \circ v) &= F(v) \circ F(u), \quad \text{for cofunctors.}) \end{aligned}$$

Remark: Obviously, Definition 1.14 can be made more formal by defining a functor, F, from $\mathcal{C}$ to $\mathcal{C}'$ as a pair, $\langle F^{\mathrm{ob}}, F^{\mathrm{fl}} \rangle$, where $F^{\mathrm{ob}} \colon \mathcal{O}\mathrm{b}(\mathcal{C}) \to \mathcal{O}\mathrm{b}(\mathcal{C}')$ and $F^{\mathrm{fl}} \colon \mathcal{F}\mathrm{l}(\mathcal{C}) \to \mathcal{F}\mathrm{l}(\mathcal{C}')$, so that, for every $u \in \mathrm{Hom}_{\mathcal{C}}(A, B)$, we have $F^{\mathrm{fl}}(u) \in \mathrm{Hom}_{\mathcal{C}'}(F^{\mathrm{ob}}(A), F^{\mathrm{ob}}(B))$, and the conditions of Definition 1.14 hold (and similarly for cofunctors).

We use the notation $A \rightsquigarrow F(A)$ (or $u \rightsquigarrow F(u)$) to indicate that $F \colon \mathcal{C} \to \mathcal{C}'$ is a functor from $\mathcal{C}$ to $\mathcal{C}'$, and not just an ordinary function.

Examples of Functors

(1) For the categories in Examples (2)–(7), consider the rule:
$A \in \mathcal{O}\mathrm{b}(\mathcal{C}) \rightsquigarrow |A| =$ the underlying set of A, and
$u \in \mathcal{F}\mathrm{l}(\mathcal{C}) \rightsquigarrow |u| =$ the morphism, u, as a map of sets.
The functor, $|\ |$, is a functor from $\mathcal{C}$ to $\mathcal{S}$ets, called the *forgetful functor* or *stripping functor*.

(2) A cofunctor, $F \colon \mathcal{C} \to \mathcal{C}'$, is just a functor, $F \colon \mathcal{C}^{\mathcal{D}} \to \mathcal{C}'$ (equivalently, $F \colon \mathcal{C} \to \mathcal{C}'^{D}$).

(3) We have the functor, $\mathbb{G}_a \colon \mathrm{RNG} \to \mathcal{A}\mathrm{b}$, given by taking $\mathbb{G}_a(R) = R$ as an additive group, for every ring, R. The functor, $\mathbb{G}_a$, is called the *additive group functor*.

(4) For every integer, $n \geq 0$, we have the functor, $\mathrm{GL}_n\colon \mathrm{CR} \to \mathcal{G}\mathrm{r}$, where $\mathrm{GL}_n(A)$ is the group of invertible $n \times n$ matrices with entries in A. When $n = 1$, the group GL_1 is denoted $\mathbb{G}_m$. This is the *multiplicative group functor*, it takes CR to $\mathcal{A}\mathrm{b}$. The functor $\mathbb{G}_m$ can be promoted to a functor, $\mathrm{RNG} \longrightarrow \mathcal{G}\mathrm{r}$, taking the ring, A, to its group, A^*, of units.

(5) Let $(\mathrm{TOP}, *)$ be the category of topological spaces together with a base point. We have the subcategory $(\text{C-TOP}, *)$ consisting of connected and locally connected topological spaces with a base point. The morphisms of $(\text{C-TOP}, *)$ preserve base points. We have the functors (fundamental group)

$$\pi_1\colon (\text{C-TOP}, *) \to \mathcal{G}\mathrm{r},$$

and for $n > 1$ (nth homotopy group),

$$\pi_n\colon (\text{C-TOP}, *) \to \mathcal{A}\mathrm{b}.$$

(6) For every integer, $n \geq 0$, we have a functor (integral homology), $\mathrm{TOP} \longrightarrow \mathcal{A}\mathrm{b}$, given by $X \rightsquigarrow H_n(X, \mathbb{Z})$ and a cofunctor (integral cohomology), $\mathrm{TOP} \longrightarrow \mathcal{A}\mathrm{b}$, given by $X \rightsquigarrow H^n(X, \mathbb{Z})$.

(7) math.upenn.edu/ Given a group, G, for any integer, $n \geq 0$, we have a functor, $\mathcal{M}\mathrm{od}(G) \longrightarrow \mathcal{A}\mathrm{b}$, given by $A \rightsquigarrow H^n(G, A)$.

Definition 1.15 Say F and F' are two functors $\mathcal{C} \longrightarrow \mathcal{C}'$. A *morphism*, θ, from F to F' is a collection $\{\theta_A \mid A \in \mathcal{O}\mathrm{b}(\mathcal{C})\}$, where:

(1) $\theta_A\colon F(A) \to F'(A)$ in $\mathcal{C}'$, so that (consistency)

(2) For every $v\colon A \to B$ in $\mathcal{C}$, the diagram

$$\begin{tikzcd} F(A) \arrow[r, "\theta_A"] \arrow[d, "F(v)"'] & F'(A) \arrow[d, "F'(v)"] \\ F(B) \arrow[r, "\theta_B"] & F'(B) \end{tikzcd}$$

commutes, for all $A, B \in \mathcal{O}\mathrm{b}(\mathcal{C})$.

A morphism of functors is also called a *natural transformation* of functors.

Examples of Morphisms of Functors:

(1) In the category $(\text{C-TOP}, *)$, we have the functors π_1 and $H_1(-, \mathbb{Z})$. The Hurewicz map

$$\pi_1(X) \xrightarrow{u_X} H_1(X, \mathbb{Z})$$

defines a morphism of functors.

(2) If G is a group and K is a subgroup of G, we have the obvious restriction functor res: $\mathcal{M}\mathrm{od}(G) \to \mathcal{M}\mathrm{od}(K)$, and it induces a morphism of functors $\mathrm{res}\colon H^n(G, -) \to H^n(K, -)$.

(3) The determinant, $\det\colon \mathrm{GL}_n \to \mathbb{G}_m$, is a morphism of functors (from CR to $\mathcal{A}\mathrm{b}$).

(4) Check (DX) that with the above notion of morphisms, the functors from $\mathcal{C}$ to $\mathcal{C}'$ form a category themselves. This category is denoted $\mathrm{Fun}(\mathcal{C}, \mathcal{C}')$.

Proposition 1.38 *Given a category, $\mathcal{C}$, each object, A, of $\mathcal{C}$ gives rise to both a functor, h_A, and a cofunctor, h_A^D, from $\mathcal{C}$ to* $\mathcal{S}$ets.

Proof. For any given $A \in \mathcal{O}b(\mathcal{C})$, let

$$\begin{aligned} h_A(B) &= \mathrm{Hom}_{\mathcal{C}}(A, B) \\ h_A^D(B) &= \mathrm{Hom}_{\mathcal{C}}(B, A). \end{aligned}$$

Moreover, for every $v \in \mathrm{Hom}_{\mathcal{C}}(B, C)$, define $h_A(v)\colon \mathrm{Hom}_{\mathcal{C}}(A, B) \to \mathrm{Hom}_{\mathcal{C}}(A, C)$ by composition, so that for every $u \in \mathrm{Hom}_{\mathcal{C}}(A, B)$,

$$h_A(v)(u) = v \circ u,$$

and, for every $v \in \mathrm{Hom}_{\mathcal{C}}(B, C)$, define $h_A^D(v)\colon \mathrm{Hom}_{\mathcal{C}}(C, A) \to \mathrm{Hom}_{\mathcal{C}}(B, A)$, again by composition, so that for every $u \in \mathrm{Hom}_{\mathcal{C}}(C, A)$,

$$h_A^D(v)(u) = u \circ v.$$

The reader should check that h_A and h_A^D are indeed functors (DX). □

The following proposition is half of the Yoneda embedding lemma:

Proposition 1.39 *Let A and $\widetilde{A}$ be two objects of $\mathcal{C}$ and suppose that the corresponding functors h_A and $h_{\widetilde{A}}$ are isomorphic, say by $\theta\colon h_A \to h_{\widetilde{A}}$. Then, A and $\widetilde{A}$ are isomorphic via a canonically determined isomorphism (dependent on θ).*

Proof. For every $B \in \mathcal{O}b(\mathcal{C})$, we have an isomorphism

$$\theta_B\colon \mathrm{Hom}_{\mathcal{C}}(A, B) \xrightarrow{\sim} \mathrm{Hom}_{\mathcal{C}}(\widetilde{A}, B),$$

and this is functorial. Let $B = A$, then $\theta_A\colon \mathrm{Hom}_{\mathcal{C}}(A, A) \xrightarrow{\sim} \mathrm{Hom}_{\mathcal{C}}(\widetilde{A}, A)$, and we set $\psi = \theta_A(1_A)$, a morphism in $\mathrm{Hom}_{\mathcal{C}}(\widetilde{A}, A)$. Now, if we let $B = \widetilde{A}$, we get $\theta_{\widetilde{A}}\colon \mathrm{Hom}_{\mathcal{C}}(A, \widetilde{A}) \xrightarrow{\sim} \mathrm{Hom}_{\mathcal{C}}(\widetilde{A}, \widetilde{A})$, and we set $\varphi = \theta_{\widetilde{A}}^{-1}(1_{\widetilde{A}})$, a morphism in $\mathrm{Hom}_{\mathcal{C}}(A, \widetilde{A})$. Pick any z in $\mathrm{Hom}_{\mathcal{C}}(A, B)$. We would like to understand what $\theta_B(z)$ is. We have the commutative diagram

$$\begin{tikzcd}
z \in \mathrm{Hom}_{\mathcal{C}}(A, B) \arrow[r, "\theta_B"] & \mathrm{Hom}_{\mathcal{C}}(\widetilde{A}, B) \\
1_A \in \mathrm{Hom}_{\mathcal{C}}(A, A) \arrow[r, "\theta_A"'] \arrow[u, "z\circ-"] & \mathrm{Hom}_{\mathcal{C}}(\widetilde{A}, A). \arrow[u, "z\circ-"']
\end{tikzcd}$$

Following the above commutative diagram clockwise, we get $\theta_B(z)$, and following it counterclockwise, we get $z \circ \psi$. We conclude that

$$\theta_B(z) = z \circ \psi.$$

Similarly, for any $\widetilde{z} \in \mathrm{Hom}_{\mathcal{C}}(\widetilde{A}, B)$, by considering the commutative diagram involving $\theta_{\widetilde{A}}^{-1}$ and θ_B^{-1}, we get

$$\theta_B^{-1}(\widetilde{z}) = \widetilde{z} \circ \varphi.$$

But then, we have

$$1_{\widetilde{A}} = \theta_{\widetilde{A}}(\varphi) = \varphi \circ \psi \quad \text{and} \quad 1_A = \theta_A^{-1}(\psi) = \psi \circ \varphi,$$

which shows that φ and ψ are inverse isomorphisms. Furthermore, φ (resp. ψ) determine θ, just as θ determines φ and ψ. □

Example. Recall that $\mathrm{Vect}(k)$ is the category of vector spaces over a field, k. There exists a cofunctor, $D\colon \mathrm{Vect}(k) \longrightarrow \mathrm{Vect}(k)$, given by: $V \rightsquigarrow V^D = \mathrm{Hom}_{\mathrm{Vect}(k)}(V, k) =$ the dual space of V; and for any linear map, $\theta\colon V \to W$, the map $\theta^D\colon W^D \to V^D$ is the adjoint of θ. By applying D again, we get a functor, $DD\colon \mathrm{Vect}(k) \longrightarrow \mathrm{Vect}(k)$. However, it is well-known that there exists a morphism of functors,
$\eta\colon \mathrm{id} \to DD$, where $\mathrm{id}(V) = V \xrightarrow{\eta_V} DD(V) = V^{DD}$, and this is functorial.

Two categories, $\mathcal{C}$ and $\mathcal{C}'$, are *equivalent* (resp. *isomorphic*) iff there exist functors $F\colon \mathcal{C} \to \mathcal{C}'$ and $F'\colon \mathcal{C}' \to \mathcal{C}$ so that $F' \circ F \cong 1_{\mathcal{C}}$ and $F \circ F' \cong 1_{\mathcal{C}'}$ (resp. $F' \circ F = 1_{\mathcal{C}}$ and $F \circ F' = 1_{\mathcal{C}'}$). Here $1_{\mathcal{C}}$ denotes the identity functor from $\mathcal{C}$ to itself.

Proposition 1.40 *(Yoneda's Embedding Lemma) The functor $A \rightsquigarrow h_A^D$ establishes an equivalence of the category, $\mathcal{C}$, with a full subcategory of $\mathrm{Fun}^D(\mathcal{C}, \mathcal{S}\mathrm{ets})$ (where $\mathrm{Fun}^D(\mathcal{C}, \mathcal{C}')$ denotes the category of cofunctors from $\mathcal{C}$ to $\mathcal{C}'$).*

Proof. We already know from Proposition 1.39 that if we have an isomorphism $\theta\colon h_A^D \to h_{\widetilde{A}}^D$, then θ determines uniquely two mutually inverse isomorphisms $\psi\colon A \to \widetilde{A}$ and $\varphi\colon \widetilde{A} \to A$. So, two objects A and $\widetilde{A}$ in $\mathcal{O}\mathrm{b}(\mathcal{C})$ give isomorphic cofunctors iff they themselves are isomorphic. Given any $v \in \mathrm{Hom}_{\mathcal{F}}(h_A^D, h_{\widetilde{A}}^D)$, where $\mathcal{F} = \mathrm{Fun}^D(\mathcal{C}, \mathcal{S}\mathrm{ets})$, we know (again) that there exists a morphism $\psi\colon A \to \widetilde{A}$, so that v is given by composing with ψ, *i.e.*, given a consistent family of morphisms, $v_B\colon h_A^D(B) \to h_{\widetilde{A}}^D(B)$, that is,
$v_B\colon \mathrm{Hom}_{\mathcal{C}}(B, A) \to \mathrm{Hom}_{\mathcal{C}}(B, \widetilde{A})$, we have $v_B(z) = \psi \circ z$, and our ψ is given by $\psi = v_A(1_A)$ (all this from the proof of Proposition 1.39). Hence, from v, we get a morphism $\psi\colon A \to \widetilde{A}$, thus

$$\mathrm{Hom}_{\mathcal{C}}(A, \widetilde{A}) \cong \mathrm{Hom}_{\mathcal{F}}(h_A^D, h_{\widetilde{A}}^D).$$

So, we indeed have an equivalence with a full subcategory of $\mathcal{F}$, namely the image consists of those cofunctors of the form h_A^D (easy details are left to the reader (DX)). $\square$

Remark: What does Yoneda's lemma say? It says that any object $A \in \mathcal{O}\mathrm{b}(\mathcal{C})$ is determined by its corresponding cofunctor h_A^D. The cofunctor, h_A^D, is a "collection of interconnected sets", $\mathrm{Hom}_{\mathcal{C}}(B, A)$ being the set associated with B.

Definition 1.16 Given a functor, F, from $\mathcal{C}$ to $\mathcal{S}\mathrm{ets}$ (resp. a cofunctor, G, from $\mathcal{C}^D$ to $\mathcal{S}\mathrm{ets}$), it is *representable* iff there exists a pair, (A, ξ), where $A \in \mathcal{O}\mathrm{b}(\mathcal{C})$ and $\xi \in F(A)$, so that F is *isomorphic* to h_A *via* the morphism of functors, $\widetilde{\xi}\colon h_A \to F$, given by the consistent family of morphisms $\widetilde{\xi}_B\colon \mathrm{Hom}_{\mathcal{C}}(A, B) \to F(B)$ defined *via*

$$\widetilde{\xi}_B(u) = F(u)(\xi),$$

(resp. G is *isomorphic* to h_A^D *via* the morphism of functors, $\widetilde{\xi}\colon h_A^D \to G$, given by $\widetilde{\xi}_B\colon \mathrm{Hom}_{\mathcal{C}}(B, A) \to G(B)$. Here, $\widetilde{\xi}_B$ is defined *via* $\widetilde{\xi}_B(u) = G(u)(\xi)$).

The notion of representable functor is a key concept of modern mathematics. The underlying idea is to "lift" as much as possible of the knowledge we have about the category of sets to other categories. More specifically, we are interested in those functors from a category $\mathcal{C}$ to $\mathcal{S}\mathrm{ets}$ that are of the form h_A for some object $A \in \mathcal{O}\mathrm{b}(\mathcal{C})$.

Remark: If (A, ξ) and (A', ξ') represent the same functor, then there exists *one and only one* isomorphism $A \xrightarrow{\sim} A'$ so that $\xi \in F(A)$ maps to $\xi' \in F(A')$. This is because we have the isomorphisms $\widetilde{\xi}\colon h_A \xrightarrow{\sim} F$ and $\widetilde{\xi'}\colon h_{A'} \xrightarrow{\sim} F$; and so, we have an isomorphism $\widetilde{\xi'}^{-1} \circ \widetilde{\xi}\colon h_A \xrightarrow{\sim} h_{A'}$. By Yoneda's lemma, $A \xrightarrow{\sim} A'$ *via* the isomorphism determined by $\widetilde{\xi}$ and $\widetilde{\xi'}$ and this maps ξ to ξ'. Uniqueness follows as everything is determined by ξ and ξ'.

Examples of Representable Functors:

(1) Let $\mathcal{C} = \mathcal{S}\mathrm{ets}$; consider the functor $F\colon \mathcal{S}\mathrm{ets}^D \to \mathcal{S}\mathrm{ets}$ given by: $F(S) =$ the collection of all subsets of S, and if $\theta\colon S \to T$ is a map of sets, the morphism $F(\theta)\colon F(T) \to F(S)$ is the map that sends every subset, V, of T to its inverse image, $\theta^{-1}(V)$, a subset of S. Is this a representable functor?

We need a set, Q, and an element, $\xi \in F(Q)$, i.e., some subset of Q, so that

$$h_Q^D(B) = \mathrm{Hom}_{\mathcal{S}\mathrm{ets}}(B, Q) \xrightarrow{\sim} F(B), \quad \textit{via } \widetilde{\xi}_B(u) = F(u)(\xi).$$

Now, we know that $F(u)\colon F(Q) \to F(B)$ is the map that sends a subset, S, of Q to its inverse image, $u^{-1}(S)$, a subset of B. So, $F(u)(\xi)$ is the inverse image of our chosen ξ.

Take $Q = \{0, 1\}$ and $\xi = \{1\} \subseteq Q$. Then, subsets of B are exactly of the form, $u^{-1}(1)$, for the various $u \in \mathrm{Hom}_{\mathcal{S}\mathrm{ets}}(B, Q)$, which are thus characteristic functions.

(2) Let $\mathcal{C} = \mathrm{RNG}$, and let $F\colon \mathrm{RNG} \to \mathcal{S}\mathrm{ets}$ be the stripping functor. Is it representable?

We need a ring, P, and an element, $\xi \in P$, so that for all rings, B,

$$\mathrm{Hom}_{\mathrm{RNG}}(P, B) \stackrel{\sim}{\longrightarrow} |B|,$$

via

$$u \in \mathrm{Hom}_{\mathrm{RNG}}(P, B) \mapsto u(\xi) \in |B|.$$

Take $P = \mathbb{Z}[T]$, the polynomial ring in one variable with integral coefficients, and $\xi = T$. Then, any ring homomorphism $u \in \mathrm{Hom}_{\mathrm{RNG}}(\mathbb{Z}[T], B)$ is uniquely determined by $u(T) = b \in |B|$, and **any** b can be used.

Definition 1.17 Let $F\colon \mathcal{C} \to \mathcal{C}'$ and $G\colon \mathcal{C}' \to \mathcal{C}$ be two functors. The functor F is the *left (resp. right) adjoint* of G iff for every $A \in \mathcal{O}\mathrm{b}(\mathcal{C})$ and $B \in \mathcal{O}\mathrm{b}(\mathcal{C}')$, we have functorial isomorphisms (in both A and B)

$$\begin{array}{rcl} \mathrm{Hom}_{\mathcal{C}'}(F(A), B) & \stackrel{\sim}{\longrightarrow} & \mathrm{Hom}_{\mathcal{C}}(A, G(B)). \\ (\text{resp.}\quad \mathrm{Hom}_{\mathcal{C}'}(B, F(A)) & \stackrel{\sim}{\longrightarrow} & \mathrm{Hom}_{\mathcal{C}}(G(B), A)). \end{array}$$

Observe that F is left-adjoint to G iff G is right-adjoint to F. Many so-called "universal constructions" arise from the existence of adjoint functors; this is a key concept in modern mathematics.

Remark: The concept of adjointness is related to the notion of representability of a functor, as shown by the following proposition whose simple proof is left to the reader:

Proposition 1.41 *A functor, $G\colon \mathcal{C}' \to \mathcal{C}$, has a left-adjoint if and only if, for every $A \in \mathcal{C}$, the functor $B \rightsquigarrow \mathrm{Hom}_{\mathcal{C}}(A, G(B))$ from $\mathcal{C}'$ to $\mathcal{S}\mathrm{ets}$ is representable. If $(F(A), \xi)$ represents this functor (so that $\widetilde{\xi}_B\colon \mathrm{Hom}_{\mathcal{C}'}(F(A), B) \cong \mathrm{Hom}_{\mathcal{C}}(A, G(B))$ is an isomorphism for every $B \in \mathcal{C}'$), then F is the object part of a left-adjoint of G for which the isomorphism $\widetilde{\xi}_B$ is functorial in B and yields the adjointness.*

A functor may have a right adjoint, but no left adjoint, and conversely (or no adjoint at all). For example, the functor, $G \rightsquigarrow G/[G, G] = G^{ab}$, from $\mathcal{G}\mathrm{r}$ to $\mathcal{A}\mathrm{b}$, is the left adjoint of the inclusion functor from $\mathcal{A}\mathrm{b}$ to $\mathcal{G}\mathrm{r}$. The inclusion views an abelian group just as a group. So, $G \rightsquigarrow G^{ab}$ has a right adjoint. However, we now prove that it has no left adjoint.

Suppose such a left adjoint, F, exists.

Claim 1: For any abelian group, H, the group $F(H)$ can never be simple unless $F(H) = \{1\}$, in which case, $H = \{1\}$.

The adjointness property states that for every group, G, we have a functorial isomorphism

$$\mathrm{Hom}_{\mathcal{G}\mathrm{r}}(F(H), G) \cong \mathrm{Hom}_{\mathcal{A}\mathrm{b}}(H, G^{ab}). \tag{$*$}$$

If we take $G = F(H)$ in $(*)$, we have

$$\mathrm{Hom}_{\mathcal{G}\mathrm{r}}(F(H), F(H)) \cong \mathrm{Hom}_{\mathcal{A}\mathrm{b}}(H, F(H)^{ab}).$$

If $F(H) \neq \{1\}$ and $F(H)$ is non-abelian simple, then, on the left hand side there are at least two maps (id and the constant map that sends all elements to 1), even though on the righthand side there is a single map, since $F(H)$ is non-abelian simple, a contradiction. If $F(H)$ is $\mathbb{Z}/p\mathbb{Z}$ for some prime p, take G in $(*)$ to be A_{3p}. Again, there are at least two maps in $\mathrm{Hom}_{\mathcal{G}\mathrm{r}}(F(H), A_{3p})$, namely: the constant map and an embedding. But A_{3p} is simple; so, the righthand side has only one element, again a contradiction. Now, if $F(H) = \{1\}$, take $G = H$ in $(*)$. In this case, the left hand side has a single map but the righthand side has at least two maps if $H \neq \{1\}$.

Claim 2: $F(H)$ has no maximal normal subgroups. If $M \lhd F(H)$ and M is maximal, then $F(H)/M$ is simple. Let $G = F(H)/M$ in $(*)$. If $F(H)/M$ is non-abelian, there are at least two maps on the left hand side, but only one on the righthand side, a contradiction. If $F(H)/M$ is abelian, say $\mathbb{Z}/p\mathbb{Z}$, again take $G = A_{3p}$ in $(*)$. There are two maps (at least) on the left hand side (stemming from the two maps $F(H) \longrightarrow F(H)/M$) and only one on the righthand side. So, if $F(H)$ exists, it is not finitely generated.

Take $H = G = \mathbb{Z}/2\mathbb{Z}$. Then, we have

$$\mathrm{Hom}_{\mathcal{G}\mathrm{r}}(F(\mathbb{Z}/2\mathbb{Z}), \mathbb{Z}/2\mathbb{Z}) \cong \mathrm{Hom}_{\mathcal{A}\mathrm{b}}(\mathbb{Z}/2\mathbb{Z}, \mathbb{Z}/2\mathbb{Z}).$$

Clearly, the righthand side has exactly two maps, and thus, so does the left hand side. But one of these maps is the constant map sending all elements to 1, so the other map must be surjective. If so, its kernel, K, is a subgroup of index 2, hence normal, and so, it must be maximal normal, a contradiction.

Therefore, the functor $G \rightsquigarrow G/[G,G] = G^{ab}$, from $\mathcal{G}\mathrm{r}$ to $\mathcal{A}\mathrm{b}$, has no left adjoint.

One often encounters situations (for example in topology, differential geometry and algebraic geometry) where the objects of interest are arrows "over" a given object (or the dual notion of arrows "co-over" a given object), for example, vector bundles, fibre bundles, algebras over a ring, etc. Such situations are captured by the abstract notion of "comma categories."

Definition 1.18 Let $\mathcal{C}$ be a category and fix some object, A, in $\mathcal{O}\mathrm{b}(\mathcal{C})$. We let $\mathcal{C}_A$, the *category over* A (or *comma category*), be the category whose objects are pairs (B, π_B), where B is some object in $\mathcal{O}\mathrm{b}(\mathcal{C})$ and π_B is a morphism in $\mathrm{Hom}_{\mathcal{C}}(B, A)$, and whose morphisms from (B, π_B) to (C, π_C) are the morphisms $u \in \mathrm{Hom}_{\mathcal{C}}(B, C)$ making the following diagram commute:

$$\begin{tikzcd}
B \arrow[rr, "u"] \arrow[dr, "\pi_B"'] & & C \arrow[dl, "\pi_C"] \\
& A &
\end{tikzcd}$$

Dually, we let $\mathcal{C}^A$, the *category co- over* A (also called *comma category*), be the category whose objects are pairs (B, i_B), where B is some object in $\mathcal{O}\mathrm{b}(\mathcal{C})$ and i_B is a morphism in $\mathrm{Hom}_{\mathcal{C}}(A, B)$, and whose morphisms from (B, i_B) to (C, i_C) are the morphisms $u \in \mathrm{Hom}_{\mathcal{C}}(B, C)$ making the following diagram commute:

$$\begin{tikzcd}
B \arrow[rr, "u"] & & C \\
& A \arrow[ul, "i_B"] \arrow[ur, "i_C"'] &
\end{tikzcd}$$

The notion of representable functor allows us to define products and coproducts in arbitrary categories.

Let $\mathcal{C}$ be any category. Say $\{A_\alpha\}_{\alpha \in \Lambda}$ is a set of objects in $\mathcal{O}\mathrm{b}(\mathcal{C})$.

(1) We get a cofunctor, F, from $\mathcal{C}^D$ to $\mathcal{S}$ets *via*

$$B \rightsquigarrow \prod_{\alpha} \mathrm{Hom}_{\mathcal{C}}(B, A_\alpha) = F(B),$$

where the above product is just the cartesian product of sets, and

(2) We get a functor, G, from $\mathcal{C}$ to $\mathcal{S}$ets *via*

$$B \rightsquigarrow \prod_\alpha \mathrm{Hom}_{\mathcal{C}}(A_\alpha, B) = G(B).$$

Are these (or either) representable?

First, consider (1). We need an object, $P \in \mathcal{O}\mathrm{b}(\mathcal{C})$ and some $\xi \in F(P)$, *i.e.* $\xi \in \prod_\alpha \mathrm{Hom}_{\mathcal{C}}(P, A_\alpha)$, which means that $\xi = \{pr_\alpha\}_\alpha$, where the pr_α are morphisms $pr_\alpha \colon P \to A_\alpha$.

Definition 1.19 When $(P, \{pr_\alpha\})$ exists, *i.e.* for every $B \in \mathcal{O}\mathrm{b}(\mathcal{C})$, there is a functorial isomorphism

$$\mathrm{Hom}_{\mathcal{C}}(B, P) \xrightarrow{\sim} \prod_\alpha \mathrm{Hom}_{\mathcal{C}}(B, A_\alpha),$$

via $u \mapsto (pr_\alpha \circ u)_\alpha$, the pair $(P, \{pr_\alpha\})$ is *the product* of the A_α's in $\mathcal{C}$. This product is denoted $\prod_\alpha A_\alpha$ (one usually drops the pr_α's). We have the (functorial) isomorphism

$$\mathrm{Hom}_{\mathcal{C}}(B, \prod_\alpha A_\alpha) \xrightarrow{\sim} \prod_\alpha \mathrm{Hom}_{\mathcal{C}}(B, A_\alpha). \qquad (*)$$

Remark: Definition 1.19 implies that for every family of morphisms, $\{f_\alpha \colon B \to A_\alpha\} \in \prod_\alpha \mathrm{Hom}_{\mathcal{C}}(B, A_\alpha)$, there is a *unique* morphism, $u \colon B \to \prod_\alpha A_\alpha$, so that

$$f_\alpha = pr_\alpha \circ u, \quad \text{for all } \alpha.$$

This is called the *universal mapping property* of products. In general, universal mapping properties are another name for representing a functor. The latter is a more general and supple notion and we will mainly stick to it.

Now, consider (2). We need an object, $Q \in \mathcal{O}\mathrm{b}(\mathcal{C})$, and some $\xi \in G(Q)$, *i.e.* $\xi \in \prod_\alpha \mathrm{Hom}_{\mathcal{C}}(A_\alpha, Q)$, which means that $\xi = \{i_\alpha\}_\alpha$, where the i_α are morphisms $i_\alpha \colon A_\alpha \to Q$.

Definition 1.20 When $(Q, \{i_\alpha\})$ exists, *i.e.* for every $B \in \mathcal{O}\mathrm{b}(\mathcal{C})$, there is a functorial isomorphism

$$\mathrm{Hom}_{\mathcal{C}}(Q, B) \xrightarrow{\sim} \prod_\alpha \mathrm{Hom}_{\mathcal{C}}(A_\alpha, B),$$

via $u \mapsto (u \circ i_\alpha)_\alpha$, the pair $(Q, \{i_\alpha\})$ is *the coproduct* of the A_α's in $\mathcal{C}$. This coproduct is denoted $\coprod_\alpha A_\alpha$ (one usually drops the i_α's). We have the (functorial) isomorphism

$$\mathrm{Hom}_{\mathcal{C}}(\coprod_\alpha A_\alpha, B) \xrightarrow{\sim} \prod_\alpha \mathrm{Hom}_{\mathcal{C}}(A_\alpha, B). \qquad (**)$$

Of course, as above, there is a universal mapping property here, also.

Definition 1.21 The product in $\mathcal{C}_A$ is called the *fibred product over* A in $\mathcal{C}$. The coproduct in $\mathcal{C}^A$ is called the *fibred coproduct over* A in $\mathcal{C}$.

Remark: Given any family, $\{(A_\alpha, \pi_\alpha)\}_\alpha$, of objects in $\mathcal{C}_A$ (with $\pi_\alpha \colon A_\alpha \to A$), the fibred product of the A_α's over A in $\mathcal{C}$ is a pair, $(\prod_A A_\alpha, \xi)$, where $\prod_A A_\alpha$ is some object in $\mathcal{C}$ (together with a morphism, $\pi \colon \prod_A A_\alpha \to A$), and ξ consists of a family of morphisms, $pr_\alpha \colon \prod_A A_\alpha \to A_\alpha$, with

$$\pi_\alpha \circ pr_\alpha = \pi_\beta \circ pr_\beta \, (= \pi), \quad \text{for all } \alpha, \beta;$$

moreover, for any object, $B \in \mathcal{C}$, and any family of morphisms, $\{f_\alpha \colon B \to A_\alpha\}_\alpha$, with

$$\pi_\alpha \circ f_\alpha = \pi_\beta \circ f_\beta, \quad \text{for all } \alpha, \beta,$$

there is a *unique* morphism, $u \colon B \to \prod_A A_\alpha$, so that $f_\alpha = pr_\alpha \circ u$, for all α.

We leave it to the reader to unwind the definition of fibred coproducts over A in $\mathcal{C}$.

Examples of Products, Coproducts, Fibred Products and Fibred Coproducts:

(1) $\mathcal{C} = \mathcal{S}\text{ets}$. Given a family of sets, $\{A_\alpha\}_{\alpha \in \Lambda}$, does $\prod_\alpha A_\alpha$ or $\coprod_\alpha A_\alpha$ exist? If so, what are they?

For $\prod_\alpha A_\alpha$, we seek a set, P, and an element, ξ, in $F(P)$, where F is the cofunctor

$$T \rightsquigarrow F(T) = \prod_\alpha \mathrm{Hom}_{\mathcal{S}\text{ets}}(T, A_\alpha).$$

This means that $\xi \in F(P)$ is just a tuple of maps, $pr_\alpha \colon P \to A_\alpha$. Take P to be the ordinary cartesian product of the A_α's and $pr_\alpha \colon P \to A_\alpha$, the αth projection. Check that this works (DX).

For $\coprod_\alpha A_\alpha$, we seek a set, Q, and an element, ξ, in $G(Q)$, where G is our functor

$$T \rightsquigarrow G(T) = \prod_\alpha \mathrm{Hom}_{\mathcal{S}\text{ets}}(A_\alpha, T).$$

So, we need a family of maps $i_\alpha \colon A_\alpha \to Q$. Now, if Q is to work, then for every T, we need an isomorphism

$$\theta_T \colon \mathrm{Hom}_{\mathcal{S}\text{ets}}(Q, T) \xrightarrow{\sim} \prod_\alpha \mathrm{Hom}_{\mathcal{S}\text{ets}}(A_\alpha, T)$$

given by $\theta_T(\varphi) = (\varphi \circ i_\alpha)_\alpha$. Take $Q = \dot{\bigcup}_\alpha A_\alpha$ (the disjoint union of the A_α's). The rest of the construction is easy (DX).

(2) $\mathcal{C} = \mathcal{A}\text{b}$, more generally, $\mathcal{C} = \mathcal{M}\text{od}(R)$ (R a ring) or $\mathcal{C} = \mathcal{M}\text{od}(G)$ (G a group).

We begin with products. Given a family, $\{A_\alpha\}_{\alpha \in \Lambda}$, with each A_α in $\mathcal{M}\text{od}(R)$, we seek $P \in \mathcal{M}\text{od}(R)$ and maps $pr_\alpha \colon P \to A_\alpha$ in $\mathcal{M}\text{od}(R)$, so that for every $T \in \mathcal{M}\text{od}(R)$, there is an isomorphism

$$\theta_T \colon \mathrm{Hom}_R(T, P) \xrightarrow{\sim} \prod_\alpha \mathrm{Hom}_R(T, A_\alpha),$$

where $\theta_T(\varphi) = \{pr_\alpha \circ \varphi\}_\alpha$ (the notation $\mathrm{Hom}_R(A, B)$ is usually used, instead of the more accurate but more cumbersome notation $\mathrm{Hom}_{\mathcal{M}\text{od}(R)}(A, B)$). We see that P must be $\prod_\alpha A_\alpha$, the product in the category of sets, if this can be made an R-module. Now, $\prod_\alpha A_\alpha$ is an R-module *via* coordinatewise addition, with the R-action given by $r(\xi_\alpha) = (r\xi_\alpha)$. So, $\prod_\alpha A_\alpha$ is the product of the A_α's in $\mathcal{M}\text{od}(R)$.

Next, we consider coproducts. We seek $Q \in \mathcal{M}\text{od}(R)$ and maps $i_\alpha \colon A_\alpha \to Q$ in $\mathcal{M}\text{od}(R)$, so that for every $T \in \mathcal{M}\text{od}(R)$, there is an isomorphism

$$\theta_T \colon \mathrm{Hom}_R(Q, T) \xrightarrow{\sim} \prod_\alpha \mathrm{Hom}_R(A_\alpha, T),$$

where $\theta_T(\varphi) = \{\varphi \circ i_\alpha\}_\alpha$. The disjoint union $\dot{\bigcup}_\alpha A_\alpha$ may be a first approximation to Q, but it is not good enough. Instead, we let

$$Q = \Big\{\xi \in \prod_\alpha A_\alpha \mid pr_\alpha(\xi) = 0 \quad \text{for all but finitely many } \alpha\Big\}.$$

This is an R-submodule of $\prod_\alpha A_\alpha$. The isomorphism

$$\theta_T\colon \mathrm{Hom}_R(Q,T) \overset{\sim}{\longrightarrow} \prod_\alpha \mathrm{Hom}_R(A_\alpha,T)$$

can now be established. First, let $i_\alpha(u)=(\delta^u_\beta)_\beta$, where $\delta^u_\alpha=u$ and $\delta^u_\beta=0$ for all $\beta\neq\alpha$. Given a family, $(\varphi_\alpha)_\alpha$, of maps $\varphi_\alpha\colon A_\alpha\to T$, for any $\xi=(\xi_\alpha)_\alpha\in Q$, set $\varphi(\xi)=\sum_\alpha\varphi_\alpha(\xi_\alpha)\in T$. If $\varphi\in\mathrm{Hom}_R(Q,T)$ is given, define $\varphi_\alpha=\varphi\circ i_\alpha$. This shows that if we set $\coprod_\alpha A_\alpha$ to be our R-module Q and the i_α to be our maps, $i_\alpha\colon A_\alpha\to Q$, as above, we have proved the proposition:

Proposition 1.42 *The categories: $\mathcal{S}$ets, $\mathcal{A}$b, $\mathcal{M}$od(R), $\mathcal{M}$od(G) all possess arbitrary products and coproducts.*

How about fibred products and coproducts?

(3) Let us go back to $\mathcal{C}=\mathcal{S}$ets, and first consider fibred products over A. A first approximation to the product, P, in $\mathcal{S}\mathrm{ets}_A$, is $\prod_\alpha A_\alpha$. However, this is not good enough because there is no "structure map", $\pi\colon P\to A$, so that

$$\begin{tikzcd}
P \arrow[rr, "pr_\alpha"] \arrow[dr, "\pi"'] & & A_\alpha \arrow[dl, "\pi_\alpha"] \\
& A &
\end{tikzcd}$$

commutes for all α. We let

$$P_A=\Big\{\xi\in\prod_\alpha A_\alpha \mid \pi_\alpha(\xi_\alpha)=\pi_\beta(\xi_\beta),\quad \text{for all } \alpha,\beta\Big\}.$$

This is a set (possibly empty), and it lies over A; indeed, we can define $\pi\colon P_A\to A$ by $\pi(\xi)=\pi_\alpha(\xi_\alpha)$, for any chosen α, since this is well-defined by definition of P_A. We write $\prod\limits_A A_\alpha$ for P_A and, for every α, we define the map, $pr_\alpha\colon \prod\limits_A A_\beta\to A_\alpha$, as the restriction of $pr_\alpha\colon\prod A_\beta\to A_\alpha$ to $\prod\limits_A A_\alpha$. The reader should check that this yields products in $\mathcal{S}\mathrm{ets}_A$.

Coproducts are a bit harder. It is natural to try $\dot{\bigcup}_\alpha A_\alpha$ as a first approximation, but this is not good enough: this does not tell us what $i\colon A\to Q$ is. The difficulty is that $\dot{\bigcup}_\alpha A_\alpha$ is too big, and we need to identify some of its elements. To do so, we define an equivalence relation on $\dot{\bigcup}_\alpha A_\alpha$, in two steps. First, we define *immediate equivalence*. Given $\xi\in A_\alpha$ and $\eta\in A_\beta$, we say that ξ and η are *immediately equivalent*, denoted $\xi\approx\eta$, iff there is some $a\in A$, so that $\xi=i_\alpha(a)$ and $\eta=i_\beta(a)$. The relation $\approx$ is clearly reflexive and symmetric but it is not necessarily transitive. So, we define $\sim$ to be the equivalence relation generated by $\approx$. This means that $\xi\sim\eta$ iff there exist $x_0,\ldots,x_t\in\dot{\bigcup}_\alpha A_\alpha$, so that

$$\xi=x_0,\ x_0\approx x_1,\ x_1\approx x_2,\ \ldots, x_{t-1}\approx x_t,\ x_t=\eta.$$

(For example, if $\xi\approx x$ and $x\approx\eta$, then $\xi=i_\alpha(a)$, $x=i_\beta(a)$, $x=i_\beta(b)$ and $\eta=i_\gamma(b)$. Note that $i_\beta(a)=i_\beta(b)$.) We let $\coprod\limits_A A_\alpha=(\dot{\bigcup}_\alpha A_\alpha)\,/\sim$, and $i\colon A\to\coprod\limits_A A_\alpha$ is given by $i(a)=$ class of $i_\alpha(a)$, for any fixed α (this is well-defined, by definition of $\sim$). The verification that $\coprod\limits_A A_\alpha$ works is left as an exercise (DX). *Therefore, the category of sets has arbitrary fibred coproducts as well.*

(4) $\mathcal{C}=\mathcal{A}$b, $\mathcal{M}$od(R), $\mathcal{M}$od(G).

For fibred products, we use $\prod\limits_A A_\alpha$, as constructed for $\mathcal{S}$ets, but made into an R-module (resp. G-module), in the usual way.

For fibred coproducts, begin with $\coprod_\alpha A_\alpha$ (in $\mathcal{C}$), and define N to be the submodule generated by the elements $i_\alpha(a) - i_\beta(a)$ with $a \in A$ and α, β arbitrary. Take

$$\coprod_A A_\alpha = \left(\coprod_\alpha A_\alpha\right)/N.$$

Again, the reader should check that $\coprod\limits_A A_\alpha$ works (DX). *Therefore, $\mathcal{A}$b, $\mathcal{M}$od(R), $\mathcal{M}$od(G), all have arbitrary fibred products and coproducts.*

We now consider products and coproducts in the category of groups, $\mathcal{G}$r. There is no difficulty for products: Use $\prod_\alpha A_\alpha$, the usual cartesian product of the A_α's, as sets, and make $\prod_\alpha A_\alpha$ into a group under coordinatewise multiplication. The same idea works for fibred products. However, coproducts require a new idea.

Given the family of groups, $\{A_\alpha\}_{\alpha\in\Lambda}$, write $A^0_\alpha = A_\alpha - \{1\}$. Let

$$S = \dot{\bigcup_\alpha} A^0_\alpha,$$

and consider, S^n, the n-fold cartesian product of S. We can view S^n as the set of *words of length n* over the alphabet S; each word is an n-tuple, $(\sigma_{\alpha_1}, \ldots, \sigma_{\alpha_n})$, with $\sigma_\beta \in A_\beta$. We call such a word *admissible* iff $A_{\alpha_j} \neq A_{\alpha_{j+1}}$, for $j = 1, 2, \ldots, n-1$. Let S^{n*} denote the set of admissible words of length n, and let

$$Q = \left(\bigcup_{n\geq 1} S^{n*}\right) \cup \{\emptyset\}.$$

(The special word, $\emptyset$, is the "empty word".) Multiplication in Q is defined as follows:
Given $(\sigma) = (\sigma_{\alpha_1}, \ldots, \sigma_{\alpha_r})$ and $(\tau) = (\tau_{\beta_1}, \ldots, \tau_{\beta_s})$ in Q, set

$$(\sigma)(\tau) = (\sigma_{\alpha_1}, \ldots, \sigma_{\alpha_r}, \tau_{\beta_1}, \ldots, \tau_{\beta_s}),$$

the result of concatenating the r-tuple, (σ), with the s-tuple, (τ). In case one of (σ) or (τ) is $\emptyset$, the concatenation is just the non-empty word and $\emptyset\emptyset$ is $\emptyset$. The word $(\sigma)(\tau)$ is admissible of length $r+s$, except if $\alpha_r = \beta_1$, in which case we need to perform a reduction process to obtain an admissible word:

(1) Form $\sigma_{\alpha_r}\tau_{\beta_1}$ in $A_{\alpha_r} = A_{\beta_1}$. There are two cases:

(a) $\sigma_{\alpha_r}\tau_{\beta_1} \neq 1_{\alpha_r}$ $(= 1_{\beta_1})$; then

$$(\sigma_{\alpha_1}, \ldots, \sigma_{\alpha_{r-1}}, \sigma_r\tau_{\beta_1}, \tau_{\beta_2}, \ldots, \tau_{\beta_s})$$

is an admissible word of length $r+s-1$, and the reduction process ends with this word as output.

(b) $\sigma_{\alpha_r}\tau_{\beta_1} = 1_{\alpha_r}$ $(= 1_{\beta_1})$; then, omit σ_{α_r} and τ_{β_1}, form

$$(\sigma_{\alpha_1}, \ldots, \sigma_{\alpha_{r-1}}, \tau_{\beta_2}, \ldots, \tau_{\beta_s}),$$

a word of length $r+s-2$, and if necessary, go back to (1) above.

Since both step (a) and (b) decrease the length of the current word, the reduction process must end with some admissible word of length $l \leq r+s$, or the empty word.

The set Q with the above multiplication is indeed a group with identity element, $\emptyset$ (DX). (The map $i_\alpha \colon A_\alpha \to Q$ sends $\sigma \in A_\alpha$ to the length-one word (σ) if $\sigma \neq 1$ or to $\emptyset$ if $\sigma = 1$.) In summary, we get

Theorem 1.43 *The category of groups, $\mathcal{G}\mathrm{r}$, possesses arbitrary coproducts (old fashioned name: "free product of the A_α.")*

Definition 1.22 Given any set, S, define the *the free group on S* to be the group $\mathrm{Fr}(S) = \coprod_{S} \mathbb{Z}$.

We have just shown that coproducts exist in the category $\mathcal{G}\mathrm{r}$. What about coproducts in the category $\mathcal{G}\mathrm{r}^A$, where A is any group?

Given a family $\{(G_\alpha, i_\alpha)\}_{\alpha\in\Lambda}$ in $\mathcal{G}\mathrm{r}^A$, form $G = \coprod_\alpha G_\alpha$, in the category $\mathcal{G}\mathrm{r}$. In G, consider the collection of elements

$$\{i_\alpha(a)i_\beta^{-1}(a) \mid a \in A,\, i_\alpha\colon A \to G_\alpha,\ \alpha \,\text{and}\, \beta \in \Lambda\};$$

let N be the *normal* subgroup of G generated by the above elements. Then, $G/N \in \mathcal{O}\mathrm{b}(\mathcal{G}\mathrm{r}^A)$, because the map $i\colon A \to G/N$ given by $i(a) =$ image of $i_\alpha(a)$ in G/N (for any fixed α) is well-defined (since image of $i_\alpha(a) =$ image of $i_\beta(a)$ in G/N). Check that, (DX), $(G/N, i)$ is the fibred coproduct of the G_α's. (Old terminology: *amalgamated product of the G_α* over A.)

Examples of fibred coproducts: (1) Let U and V be two sets. Form the intersection $U \cap V$; we have inclusion maps $i_U : U \cap V \to U$ and $i_V : U \cap V \to V$. We know that $U \amalg V = U \uplus V$, the disjoint union of U and V, and then, the set-theoretic union of U and V is given by

$$U \cup V = U \coprod_{U\cap V} V.$$

(2) Consider the category $(\mathrm{TOP}, *)$ of ("nice", i.e., connected, locally connected) topological spaces with a base point. Given two spaces $(U, *)$ and $(V, *)$ in $(\mathrm{TOP}, *)$, consider $(U \cap V, *)$. Then, again,

$$(U \cup V, *) = (U, *) \coprod_{(U\cap V,*)} (V, *), \quad \text{in}\,(\mathrm{TOP}, *).$$

Van Kampen's theorem says that

$$\pi_1(U \cup V, *) = \pi_1(U, *) \coprod_{\pi_1(U\cap V,*)} \pi_1(V, *),$$

which may also be written as

$$\pi_1\left((U, *) \coprod_{(U\cap V,*)} (V, *)\right) = \pi_1(U, *) \coprod_{\pi_1(U\cap V,*)} \pi_1(V, *).$$

In other words, van Kampen's theorem says that π_1 commutes with fibred coproducts.

Go back to the free group, $\mathrm{Fr}(S)$. We have

$$\begin{aligned} \mathrm{Hom}_{\mathcal{G}\mathrm{r}}(\mathrm{Fr}(S), G) &= \mathrm{Hom}_{\mathcal{G}\mathrm{r}}(\coprod_S \mathbb{Z}, G) \\ &\cong \prod_S \mathrm{Hom}_{\mathcal{G}\mathrm{r}}(\mathbb{Z}, G) \\ &\cong \prod_S |G| = \mathrm{Hom}_{\mathcal{S}\mathrm{ets}}(S, |G|). \end{aligned}$$

Corollary 1.44 *The functor, $S \rightsquigarrow \mathrm{Fr}(S)$, from $\mathcal{S}\mathrm{ets}$ to $\mathcal{G}\mathrm{r}$ is the left adjoint to the stripping functor, $G \rightsquigarrow |G|$, from $\mathcal{G}\mathrm{r}$ to $\mathcal{S}\mathrm{ets}$.*

Corollary 1.45 *If $S \longrightarrow T$ is surjective, then $\mathrm{Fr}(S) \longrightarrow \mathrm{Fr}(T)$ is a surjection of groups. Also, $\mathrm{Fr}(S) \cong \mathrm{Fr}(T)$ iff $\#(S) = \#(T)$ (i.e., S and T have the same cardinality).*

Proof. If $u\colon S \to T$ is a surjection in $\mathcal{S}$ets, then there is a map $v\colon T \to S$ so that $u \circ v = 1_T$. Since Fr is a functor, we get homomorphisms $\mathrm{Fr}(u)\colon \mathrm{Fr}(S) \to \mathrm{Fr}(T)$ and $\mathrm{Fr}(v)\colon \mathrm{Fr}(T) \to \mathrm{Fr}(S)$; also, $\mathrm{Fr}(u) \circ \mathrm{Fr}(v) = 1_{\mathrm{Fr}(T)}$, which shows that $\mathrm{Fr}(u)$ is surjective.

If $\#(S) = \#(T)$, it is obvious that $\mathrm{Fr}(S) \cong \mathrm{Fr}(T)$. Conversely, assume that $\mathrm{Fr}(S) \cong \mathrm{Fr}(T)$. We know that

$$\mathrm{Hom}_{\mathcal{G}\mathrm{r}}(\mathrm{Fr}(S), G) \cong \mathrm{Hom}_{\mathcal{G}\mathrm{r}}(\mathrm{Fr}(T), G)$$

for all G. Take $G = \mathbb{Z}/2\mathbb{Z}$. Then, the left hand side is isomorphic to $\mathrm{Hom}_{\mathcal{S}\mathrm{ets}}(S, |\mathbb{Z}/2\mathbb{Z}|) = \mathcal{P}(S)$ (where $\mathcal{P}(S) =$ power set of S) and the righthand side is isomorphic to $\mathcal{P}(T)$. Therefore, $\#(\mathcal{P}(S)) = \#(\mathcal{P}(T))$; and so, $\#(S) = \#(T)$. □

Given a group, G, consider its underlying set, $|G|$, and then the group $\mathrm{Fr}(|G|)$. Since

$$\mathrm{Hom}_{\mathcal{G}\mathrm{r}}(\mathrm{Fr}(|G|), G) \cong \mathrm{Hom}_{\mathcal{S}\mathrm{ets}}(|G|, |G|),$$

the image of the identity map, $\mathrm{id}_G \in \mathrm{Hom}_{\mathcal{S}\mathrm{ets}}(|G|, |G|)$, yields a canonical surjection, $\mathrm{Fr}(|G|) \longrightarrow G$. If S is a subset of $|G|$, then, the inclusion map, $S \hookrightarrow |G|$, yields a morphism of groups, $\mathrm{Fr}(S) \longrightarrow G$.

Definition 1.23 A set, $S \subseteq |G|$, *generates* a group, G, iff the canonical map $\mathrm{Fr}(S) \longrightarrow G$ is surjective.

This definition agrees with our old use of *generation of a group* in previous sections. Say S generates G. Then, we have the exact sequence

$$0 \longrightarrow K \longrightarrow \mathrm{Fr}(S) \longrightarrow G \longrightarrow 0,$$

where K is the kernel of the surjective morphism, $\mathrm{Fr}(S) \longrightarrow G$ (so, K is normal in $\mathrm{Fr}(S)$). There is also a set, T, so that

$$\mathrm{Fr}(T) \longrightarrow K \longrightarrow 0 \quad \text{is exact.}$$

By splicing the two exact sequences, we get an exact sequence

$$\mathrm{Fr}(T) \longrightarrow \mathrm{Fr}(S) \longrightarrow G \longrightarrow 0,$$

called a *presentation of* G. Sometimes, a presentation is defined as a sequence

$$\mathrm{Fr}(T) \longrightarrow \mathrm{Fr}(S) \longrightarrow G \longrightarrow 0,$$

where the smallest *normal* subgroup containing $\mathrm{Im}\,(\mathrm{Fr}(T))$ is equal to the kernel of $\mathrm{Fr}(S) \longrightarrow G$. (Note that such a sequence is not necessarily exact at the group $\mathrm{Fr}(S)$.)

The following fundamental theorem about free groups was proved independently by J. Nielson and O. Schreier:

Theorem 1.46 *(Nielson-Schreier (1929)) Every subgroup of a free group is a free group.*

The original proof is quite messy. The theory of group actions on trees yields a more direct and more transparent proof.

We conclude this section on categories with one more interesting example of adjoint functors from homotopy theory.

Example: Consider the category, h-TOP, whose objects are the same as those of TOP, but whose morphisms, $\mathrm{Hom}_{\text{h-TOP}}(X, Y)$, are the homotopy classes of maps $X \longrightarrow Y$. Given any space, X, in h-TOP, we can form ΣX, the *suspension of* X: This is the space obtained by taking two new points, say 0 and 1, and forming the double cone obtained by joining 0 and 1 to every point of X, as illustrated in Figure 1.2.

We also have, ΩY, the *loop space on* Y, where ΩY consists of all continuous maps, $S^1 \longrightarrow Y$, from the unit circle to Y (say, mapping $(1,0)$ to the base point of Y). Then, we have the isomorphism

$$\mathrm{Hom}_{\text{h-TOP}}(\Sigma X, Y) \cong \mathrm{Hom}_{\text{h-TOP}}(X, \Omega Y)$$

i.e., suspension is left-adjoint to loops. For instance, given any $\theta \in \mathrm{Hom}_{\text{h-TOP}}(\Sigma X, Y)$, for any $p \in X$, send p to the image by θ of the loop $l(p)$ $(= (*, 0, p, 1, *)$ in $\Sigma X)$, in Y.

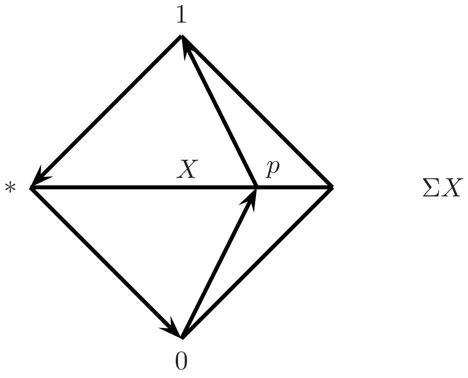

Figure 1.2: A suspension of X

1.8 Further Readings

> El que anda mucho y lee mucho,
> Vee mucho y sabe mucho.
> —**Miguel Cervantes**

Some group theory is covered in every algebra text. Among them, we mention Michael Artin [2], Lang [34], Hungerford, [27], Jacobson [29], Mac Lane and Birkhoff [37], Dummit and Foote [11], Van Der Waerden [47] and Bourbaki [4]. More specialized books include Rotman [43], Hall [22], Zassenhaus [52], Rose [42] and Gorenstein [19]. For group cohomology, see also Cartan and Eilenberg [9], Rotman [44], Mac Lane [36] and Serre [45]. Mac Lane [35] is a good reference for category theory.