Jason Starr
Fall 2016

MAT 312 Burnside's Lemma and Other Supplementary Topics

1 Group Actions

Let $(G, *)$ be a group with identity element e and with group inverse operation $g \mapsto g^{-1}$. Let X be a set. A **left action** of G on X is a function

$$\lambda : G \times X \to X,$$

such that (i) for every $x \in X$, $\lambda(e, x)$ equals x, and (ii) for every $g, h \in G$ and for every $x \in X$, $\lambda(g, \lambda(h, x))$ equals $\lambda(g * h, x)$. Similarly, a **right action** of G on X is a function

$$\rho : X \times G \to X,$$

such that (i) for every $x \in X$, $\rho(x, e)$ equals x, and (ii) for every $g, h \in G$ and for every $x \in X$, $\rho(\rho(x, h), g)$ equals $\rho(x, h * g)$.

For a left action λ of G on X, for every $x \in X$, the λ-**orbit** of x is the following subset of X,

$$G \cdot x := \{y \in X \mid \exists g \in G, y = \lambda(g, x)\} = \{\lambda(g, x) |\ g \in G\}.$$

The λ-**stabilizer** of x is the subset of G,

$$\mathrm{Stab}_\lambda(x) := \{g \in G \mid \lambda(g, x) = x\},$$

and every $g \in \mathrm{Stab}_\lambda(x)$ is said to **stabilize** x. Similarly, for every $g \in G$, the λ-**fixed locus** of g is the subset of X,

$$X^g := \{x \in X \mid \lambda(g, x) = x\},$$

and every $x \in X^g$ is said to be **fixed** by g. The action λ is **transitive** if X equals an orbit $G \cdot x$ for some $x \in X$, and hence for every $x \in X$. The action λ is **faithful** if the only element $g \in G$ that fixes every element of X is the identity element e.

Similarly for a right action ρ of G on X, for every $x \in X$, the ρ-**orbit** of x is the following subset of X,

$$x \cdot G := \{y \in X \mid \exists g \in G, y = \rho(x, g)\} = \{\rho(x, g) |\ g \in G\}.$$

Burnside's Lemma

The ρ-**stabilizer** of x is the subset of G,

$$\mathrm{Stab}_\rho(x) := \{g \in G \mid \rho(x, g) = x\}.$$

Similarly, for every $g \in G$, the ρ-**fixed locus** of g is the subset of X,

$$X^g := \{x \in X \mid \rho(x, g) = x\}.$$

Lemma 1.1. *For every $x \in X$, $Stab_\lambda(x)$ is a subgroup of G. Similarly, $Stab_\rho(x)$ is a subgroup of G.*

Proof. By the first axiom of an action, $\lambda(e, x)$ equals x, so e is contained in $\mathrm{Stab}_\lambda(x)$. For every pair $g, h \in \mathrm{Stab}_\lambda(x)$, by the second axiom,

$$\lambda(g * h, x) = \lambda(g, \lambda(h, x)) = \lambda(g, x) = x.$$

Thus, $g * h$ is in $\mathrm{Stab}_\lambda(x)$. Finally, for every $g \in \mathrm{Stab}_\lambda(x)$,

$$\lambda(g^{-1}, x) = \lambda(g^{-1}, \lambda(g, x)) = \lambda(g^{-1} * g, x) = \lambda(e, x) = x.$$

Thus g^{-1} is in $\mathrm{Stab}_\lambda(x)$. Therefore $\mathrm{Stab}_\lambda(x)$ is a subgroup of G. A similar argument proves that $\mathrm{Stab}_\rho(x)$ is a subgroup of G. □

For a left action λ of G on X, for elements $x, y \in X$, we define x to be λ-**equivalent** to y, $x \overset{\lambda}{\sim} y$, if y is an element of the λ-orbit of x, $G \cdot x$, i.e., if there exists $g \in G$ with $y = \lambda(g, x)$. Similarly, for a right action ρ of G on X, for elements $x, y \in X$, we define x to be ρ-**equivalent** to y, $x \overset{\rho}{\sim} y$, if y is an element of the ρ-orbit of x, $x \cdot G$, i.e., if there exists $g \in G$ with $y = \rho(x, g)$.

Lemma 1.2. *For every $x, y, z \in X$, (i) $x \overset{\lambda}{\sim} x$, (ii) if both $x \overset{\lambda}{\sim} y$ and $y \overset{\lambda}{\sim} z$, then $x \overset{\lambda}{\sim} z$, and (iii) if $x \overset{\lambda}{\sim} y$, then $y \overset{\lambda}{\sim} x$. The same holds for ρ-equivalence.*

Proof. By the definition of a group action, $\lambda(e, x) = x$, thus $x \overset{\lambda}{\sim} x$. If there exist $g, h \in G$ with $y = \lambda(h, x)$ and $z = \lambda(g, y)$, then z equals $\lambda(g, \lambda(h, x))$. By the definition of a group action, this equals $\lambda(g * h, x)$. Thus $x \overset{\lambda}{\sim} z$. Finally, if there exists $h \in G$ with $y = \lambda(h, x)$, then

$$\lambda(h^{-1}, y) = \lambda(h^{-1}, \lambda(h, y)) = \lambda(h^{-1} * h, y) = \lambda(e, y) = y.$$

Thus, if $x \overset{\lambda}{\sim} y$, then also $y \overset{\lambda}{\sim} x$. □

A reformulation of the lemma is that the λ-orbits $G \cdot x$ and $G \cdot y$ intersect if and only if they are equal. Similarly, ρ-orbits that intersect are equal. The λ-**orbit space**, denoted $G\backslash X$, is defined to be the set of all λ-orbits, resp. the ρ-orbit space, denoted X/G, is defined to be the set of all ρ-orbits. The λ-**quotient function** is defined to be the set function,

$$q_\lambda : X \to G\backslash X, \quad q_\lambda(x) = G \cdot x.$$

Similarly, the ρ-**quotient function** is defined to be the set function,

$$q_\rho : X \to X/G, \quad q_\rho(x) = x \cdot G.$$

By construction, each of these maps is onto. One more reformulation is that $x \overset{\lambda}{\sim} y$ if and only if $q_\lambda(x) = q_\lambda(y)$, respectively, $x \overset{\rho}{\sim} y$ if and only if $q_\rho(x) = q_\rho(y)$.

A subset $Y \subset X$ is λ-**preserved** if for every $y \in Y$ and for every $z \in X$ with $y \overset{\lambda}{\sim} z$, then also $z \in Y$. Similarly, $Y \subset X$ is ρ-**preserved** if for every $y \in Y$ and for every $z \in X$ with $y \overset{\rho}{\sim} z$, then also $z \in Y$.

Lemma 1.3. *A subset $Y \subset X$ is λ-preserved if and only if Y contains every λ-orbit that intersects Y. Similarly, Y is λ-preserved if and only if there exists a subset $Z \subset G\backslash X$ whose preimage $q_\lambda^{pre}(Z)$ equals Y, in which case Z equals $q_\lambda(Y)$. The same holds for ρ-preserved subsets.*

Proof. First assume that Y is λ-preserved. Let $G\cdot x$ be an orbit that intersects Y, say $y \in Y \cap (G\cdot x)$. Since $x \overset{\lambda}{\sim} y$, $G \cdot x$ equals $G \cdot y$. For every $z \in G \cdot y$, i.e., $y \overset{\lambda}{\sim} z$, also $z \in Y$. Thus Y contains $G \cdot y$, i.e., Y contains $G \cdot x$. Define $Z = q_\lambda(Y) = \{q_\lambda(y) | y \in Y\}$. Since $q_\lambda(Y)$ is contained in Z (in fact, equals Z), Y is contained in the preimage $q_\lambda^{\mathrm{pre}}(Z)$ of Z. For every $y \in Y$, i.e., for every $q_\lambda(y) \in Z$, Y contains $G \cdot y$, i.e., Y contains $q_\lambda^{\mathrm{pre}}(\{q_\lambda(y)\})$. Since this holds for every $q_\lambda(y)$, Y contains $q_\lambda^{\mathrm{pre}}(Z)$. Thus Y equals $q_\lambda^{\mathrm{pre}}(Z)$.

Conversely, assume that Y contains every λ-orbit that it intersects. For every $y \in Y$, Y intersects $G \cdot y$. Thus Y contains $G \cdot y$. Therefore, for every $z \in X$ with $y \overset{\lambda}{\sim} z$, i.e., with $z \in G \cdot y$, then $z \in Y$. Thus Y is λ-preserved.

Finally, assume that Y equals $q_\lambda^{\mathrm{pre}}(Z)$ for a subset Z of $G\backslash X$. Since q_λ is onto, $q_\lambda(Y)$ equals Z. For every $y \in Y$, for every $z \in X$ with $y \overset{\lambda}{\sim} z$, then $q_\lambda(z) = q_\lambda(y)$. Since $q_\lambda(y)$ is in $q_\lambda(Y) = Z$, and since z is in $q_\lambda^{\mathrm{pre}}(Z)$, also $z \in Y$. Thus Y is λ-preserved. □

For every subset W of X, the λ-preserved subset of X **generated** by W, denoted $G \cdot W$, is defined to be the preimage $q_\lambda^{\mathrm{pre}}(V)$ of the image $V = q_\lambda(Z)$. By the previous lemma, this is the λ-preserved subset that contains W and that is minimal with respect to set inclusion. It equals $\{\lambda(g, w) | g \in G, w \in W\}$. There is a similar notion for ρ.

For every left action λ of a group $(G, *)$ on a set X, for every group homomorphism $f : (H, \bullet) \to (G, *)$, the f-**pullback** of λ to H, λ^f, is the composition of λ with the set map $f \times \mathrm{Id}_X : H \times X \to G \times X$,

$$\lambda^f : H \times X \to X, \quad \lambda^f(h, x) = \lambda(f(h), x)$$

Similarly, for every right action ρ of a group $(G, *)$ on a set X, for every subgroup H of G, the f-**pullback** of ρ to H, ρ^f, is the composition of ρ with the set map $\mathrm{Id}_X \times f : X \times H \to X \times G$,

$$\rho^f : X \times H \to X, \quad \rho^f(x, h) = \rho(x, f(h)).$$

Lemma 1.4. *The function λ^f is a left action of H on X. For every $x \in X$, $Stab_{\lambda^f}(x)$ is the preimage $f^{pre}(Stab_\lambda(x))$. For every $h \in H$, the λ^f-fixed locus of h equals the λ-fixed locus of $f(h)$. Similarly, the function ρ^f is a right action of H on X, $Stab_{\rho^f}(x) = f^{pre}(Stab_\rho(x))$, and the ρ^f-fixed locus of h equals the ρ-fixed locus of $f(h)$.*

Proof. The homomorphism f maps the identity ϵ of H to the identity e of G. Thus, by the first axiom of a group action for λ,

$$\lambda^f(\epsilon, x) = \lambda(f(\epsilon), x) = \lambda(e, x) = x.$$

Thus λ^f satisfies the first axiom of a left group action. Similarly, for every $h, k \in H$,

$$\lambda^f(h, \lambda^f(k, x)) = \lambda(f(h), \lambda(f(k), x)) = \lambda(f(h) * f(k), x) = \lambda(f(h \bullet k), x) = \lambda^f(h \bullet k, x).$$

Thus λ^f satisfies the second axiom of a left group action. Thus λ^f is a left group action.

By definition, $\lambda^f(h, x)$ equals $\lambda(f(h), x)$. Thus, $h \in \mathrm{Stab}_{\lambda^f}(x)$ if and only if $f(h) \in \mathrm{Stab}_\lambda(x)$. Therefore $\mathrm{Stab}_{\lambda^f}(x)$ equals $f^{\mathrm{pre}}(\mathrm{Stab}_\lambda(x))$. Similarly, the λ-fixed locus of $f(h)$ equals the λ^f-fixed locus of h. □

Similarly, for every left action λ of $(G, *)$ on a set X, for every λ-preserved subset Y of X, $\lambda(G \times Y)$ is contained in Y. Thus, there is a unique set map,

$$\lambda_{-,Y} : G \times Y \to Y,$$

such that for every $g \in G$ and for every $y \in Y$, $\lambda_{-,Y}(g, y)$ equals $\lambda(g, y)$. Simiarly, for every right action ρ of $(G, *)$ on X and for every ρ-preserved subset Y, there is a unique set map,

$$\rho_{Y,-} : Y \times G \to Y,$$

such that for every $g \in G$ and for every $y \in Y$, $\rho_{Y,-}(g, y)$ equals $\rho(g, y)$.

Lemma 1.5. *For every λ-preserved subset Y of X, $\lambda_{-,Y}$ is a left action of G on Y. Similarly, for every ρ-preserved subset Y of X, $\rho_{Y,-}$ is a right action of G on Y.*

Proof. Since for every $g \in G$ and for every $y \in Y$, $\lambda_{-,Y}(g, y)$ equals $\lambda(g, y)$, the axioms for λ imply the axioms for $\lambda_{-,Y}$, and similarly for ρ. □

Recall that for every $g \in G$, there is a set function,

$$c_g : G \to G, \quad c_g(h) = g * h * g^{-1}.$$

Lemma 1.6. *For every left action λ of G on X, for every $g, k \in G$ and for every $x \in X$, $\lambda(k, \lambda(g, x))$ equals $\lambda(c_k(g), \lambda(k, x)$. Also, $Stab_\lambda(\lambda(k, x))$ equals $c_k(Stab_\lambda(x))$, and $\lambda(k, X^g)$ equals $X^{c_k(g)}$. Similarly, for every right action ρ of G on X, $\rho(\rho(x, g), k) = \rho(\rho(x, k), c_{k^{-1}}(g))$, $Stab_\rho(\rho(x, k)) = c_{k^{-1}}(Stab_\rho(x))$, and $\rho(X^g, k) = X^{c_{k^{-1}}(g)}$.*

Proof. Indeed, by the associativity, inverse, and identity axioms for a group,

$$c_k(g) * k = (k * g * k^{-1}) * k = (k * g) * (k^{-1} * k) = (k * g) * e = k * g.$$

Thus,

$$\lambda(c_k(g), \lambda(k, x)) = \lambda(c_k(g) * k, x) = \lambda(k * g, x) = \lambda(k, \lambda(g, x)).$$

Thus, $\lambda(g, x)$ equals x if and only if $\lambda(k, \lambda(g, x))$ equals $\lambda(k, x)$. Also, $\lambda(k, \lambda(g, x))$ equals $\lambda(c_k(g), \lambda(k, x))$. Thus g stabilizes x if and only if $c_k(g)$ stabilizes $\lambda(k, x)$, i.e., $c_k(\mathrm{Stab}_\lambda(x))$ equals $\mathrm{Stab}_\lambda(\lambda(k, x))$. Also, x is fixed by g if and only if $\lambda(k, x)$ is fixed by $c_k(g)$, i.e., $\lambda(k, X^g)$ equals $X^{c_k(g)}$. □

2 Examples of Group Actions

For every group $(G, *)$, the **left regular action** of G on itself is the function,

$$\lambda_{G,G} : G \times G \to G, \quad (g, x) \mapsto g * x.$$

Note, $e * x$ equals x for every $x \in G$ by the identity axiom. Also $g * (h * x)$ equals $(g * h) * x$ by the associativity axiom. Thus, $\lambda_{G,G}$ is a left action of G on itself. Similarly, for every group $(G, *)$, the **right regular action** of G on itself is the function,

$$\rho_{G,G} : G \times G \to G, \quad (g, x) \mapsto g * x.$$

Note, $x * e$ equals x for every $x \in G$ by the identity axiom. Also $(x * h) * g$ equals $x * (h * g)$ by the associativity axiom. Thus, $\rho_{G,G}$ is a right action of G on itself. Moreover, the left and right actions are compatible in the following sense,

$$\lambda_{G,G}(g, \rho_{G,G}(x, h)) = g * (x * h) = (g * x) * h = \rho_{G,G}(\lambda_{G,G}(g, x), h).$$

For every $x \in G$, the $\lambda_{G,G}$-orbit of x is all of G. Each of these actions is transitive, i.e., every orbit equals all of G. Thus the orbit space is a singleton set, and the quotient function is a constant function to this singleton set. Also, each of these actions is faithful. In fact, for every $x \in G$, the stabilizer of x is the trivial subgroup $\{e\}$.

For every group homomorphism $f : (H, \bullet) \to (G, *)$, the f-pullback of the left regular action, resp., the right regular action, is a left action of H on G, resp. a right action of H on G,

$$\lambda_{f,G} : H \times G \to G, \quad \lambda_{f,G}(h, x) = f(h) * x,$$

$$\rho_{G,f} : G \times H \to G, \quad \rho_{G,f}(x, h) = x * f(h).$$

For every $x \in G$, the stabilizer subgroup of x in H equals the **kernel** of f, i.e., the preimage subgroup $f^{\mathrm{pre}}(\{e\}) = \{h \in H | f(h) = e\}$. In particular, since the stabilizer of $\lambda_{f,G}(k, x)$ equals the stabilizer of x, the kernel of f is mapped to itself under the conjugation by k, $c_k(\mathrm{Ker}(f)) = \mathrm{Ker}(f)$. This is another proof (really the same proof) that the kernel of f is a normal subgroup. In particular,

the action $\lambda_{f,G}$ is transitive, resp. faithful, if and only if the function f is onto, resp. one-to-one, and the same holds for the action $\rho_{G,f}$.

Now consider the special case that $H \subset G$ is a subgroup and $f : H \to G$ is the inclusion set function. In this case, we denote the pullback actions by $\lambda_{H,G}$, resp. $\rho_{G,H}$. Since the inclusion is one-to-one, the action is faithful. However, the action is transitive if and only if H equals G, so the action is typically not transitive. The $\rho_{G,H}$-orbits, xH, are the **left H-cosets**. Similarly, the $\lambda_{H,G}$-orbits, Hx, are the **right H-cosets**. Thus the $\rho_{G,H}$-orbit space and quotient function is the set of left H-cosets and the quotient function,

$$q : G \to G/H,\ q(x) = xH.$$

A similar result holds for the right H-cosets and the $\lambda_{H,G}$-orbits. Because the left and right actions of G on itself are compatible, there is still a well-defined left G-action on G/H,

$$\lambda_{G,G/H} : G \times (G/H) \to G/H,\quad \lambda_{G,G/H}(g, xH) = (g * x)H.$$

Similarly, there is a well-defined right G-action on $H\backslash G$,

$$\rho_{H\backslash G,G} : (H\backslash G) \times G \to H\ G,\quad \rho_{H\backslash G,G}(Hx, g) = H(x * g).$$

Each of these actions is transitive, and the stabilizer of the "neutral" coset H is simply the subgroup $H \subset G$.

For any group homomorphism $f : (L, \bullet) \to (G, *)$, denoting the image subgroup $f(L)$ by H, the $\rho_{G,f}$-orbits, $x \cdot L$, are precisely the left H-cosets. Similarly, the $\lambda_{f,G}$-orbits, $L \cdot x$ are precisely the right H-cosets. Thus the $\rho_{G,f}$-orbit space and quotient function is precisely the set of left H-cosets and the quotient function $q : G \to G/H$. A similar result holds for the right H-cosets and the $\lambda_{f,G}$-orbits.

For every set T, recall that the **symmetric group** of T is the set $S(T)$ of all invertible functions $\sigma : T \to T$. The group operation on this group is composition, i.e., for invertible functions, $\sigma : T \to T$ and $\tau : T \to T$, $\sigma \circ \tau : T \to T$ is $(\sigma \circ \tau)(t) = \sigma(\tau(t))$. The inverse of $\sigma \circ \tau$ is $(\sigma \circ \tau)^{-1} = \tau^{-1} \circ \sigma^{-1}$, and the identity element for the symmetric group is the identity function $\mathrm{Id}_T : T \to T$, $\mathrm{Id}_T(t) = t$.

The **standard left action** of $S(T)$ on T is the function

$$\lambda_T : S(T) \times T \to T,\quad \lambda_T(\sigma, t) = \sigma(t).$$

Lemma 2.1. *The standard left action of the symmetric group $S(T)$ on T is a left action.*

Proof. Observe that $\lambda_T(\mathrm{Id}_T, t) = \mathrm{Id}_T(t) = t$ for every $t \in T$. Also, for every $\sigma, \tau \in S(T)$ and for every $t \in T$,

$$\lambda_T(\sigma, \lambda_T(\tau, t)) = \lambda_T(\sigma, \tau(t)) = \sigma(\tau(t)) = (\sigma \circ \tau)(t) = \lambda_T(\sigma \circ \tau, t).$$

Thus, λ_T is a left action of $S(T)$ on T. □

For every $t \in T$, denote by T_t the subset $T \setminus \{t\}$. For every invertible function $\sigma : T \to T$ that stabilizes t, since σ is one-to-one and onto, $\sigma(T_t) = \sigma(T) \setminus \sigma(\{t\}) = T \setminus \{t\} = T_t$. Thus the restriction $\sigma|_{T_t}$ defines an element in $S(T_t)$. For $\tau, \sigma \in \mathrm{Stab}_{\lambda_T}(t)$, $(\tau \circ \sigma)|_{T_t}$ equals $\tau|_{T_t} \circ \sigma|_{T_t}$. Thus, restriction defines a homomorphism $\mathrm{Stab}_{\lambda_T}(t) \to S(T_t)$.

Lemma 2.2. *For every nonempty set T, the standard left action of $S(T)$ on T is transitive. For every $t \in T$, restriction is an isomorphism from $Stab_{\lambda_T}(t)$ to $S(T_t)$. For every $\sigma \in S(T)$, the λ_T-fixed locus of σ equals $T^\sigma = \{t \in T | \sigma(t) = t\}$.*

Proof. Let $t, t' \in T$ be distinct elements of T. Then the 2-cycle $\sigma = (t, t')$ is the invertible function permuting t and t' and fixing every element of T different from t and t'. Since $\sigma(t) = t'$, t' is in the λ_T-orbit of t. Since this holds for every t, t', there is a unique λ_T-orbit.

For every $\sigma \in \mathrm{Stab}_{\lambda_T}(t)$, σ is uniquely recovered from its restriction $\sigma|_{T_t}$. Thus, the restriction homomorphism is an isomorphism from $\mathrm{Stab}_{\lambda_T}(t)$ to $S(T_t)$. □

For every group $(G, *)$ and for every left action

$$\lambda : G \times T \to T,$$

for every $g \in G$, define $\widetilde{\lambda}_g : T \to T$ by $\widetilde{\lambda}_g(t) = \lambda(g, t)$.

Lemma 2.3. *For the group identity e, $\widetilde{\lambda}_e$ equals Id_T. For every pair of elements $g, h \in G$, $\widetilde{\lambda}_g \circ \widetilde{\lambda}_h = \widetilde{\lambda}_{g*h}$.*

Proof. By the first axiom for a left action $\widetilde{\lambda}_e(t) = \lambda(e, t) = t$, so $\widetilde{\lambda}_e$ is the identity function Id_T. By the second axiom for a left action, for every $t \in T$,

$$(\widetilde{\lambda}_g \circ \widetilde{\lambda}_h)(t) = \widetilde{\lambda}_g(\widetilde{\lambda}_h(t)) = \lambda(g, \lambda(h, t)) = \lambda(g * h, t) = \widetilde{\lambda}_{g*h}(t).$$

Thus, $\widetilde{\lambda}_g \circ \widetilde{\lambda}_h$ equals $\widetilde{\lambda}_{g*h}$. □

Lemma 2.4. *For every $g \in G$, $\widetilde{\lambda}_g$ is an element of $S(T)$. The function $\widetilde{\lambda} : G \to S(T)$ by $g \mapsto \widetilde{\lambda}_g$ is a group homomorphism. This is the unique group homomorphism such that the pullback of λ_T equals λ.*

Proof. Since $\widetilde{\lambda}_e$ equals Id_T, also,

$$\widetilde{\lambda}_g \circ \widetilde{\lambda}_{g^{-1}} = \mathrm{Id}_T = \widetilde{\lambda}_{g^{-1}} \circ \widetilde{\lambda}_g.$$

Thus, $\widetilde{\lambda}_g$ is an invertible function with inverse $\widetilde{\lambda}_{g^{-1}}$. Since $\widetilde{\lambda}_g \circ \widetilde{\lambda}_h$ equals $\widetilde{\lambda}_{g*h}$, the function $\widetilde{\lambda}$ is a group homomorphism. Finally, for the group homomorphism $f = \widetilde{\lambda}_g$,

$$\lambda_T^f(g, t) = \lambda_T(f(g), t) = \lambda_T(\widetilde{\lambda}_g, t) = \widetilde{\lambda}_g(t) = \lambda(g, t).$$

Thus λ_T^f equals λ. For any group homomorphism with $\lambda_T^f = \lambda$, for every $g \in G$, for every $t \in T$,

$$f(g)(t) = \lambda_T(f(g), t) = \lambda_T^f(g, t) = \lambda(g, t) = \widetilde{\lambda}_g(t).$$

Thus $f(g)$ equals $\widetilde{\lambda}_g$, so that f equals $\widetilde{\lambda}$. □

In this sense, left actions λ of G on T are equivalent to group homomorphisms $\widetilde{\lambda}$ from G to $S(T)$. In particular, the action λ is faithful if and only if the group homomorphism $\widetilde{\lambda}$ is one-to-one.

For every group $(G, *)$, for every set X, and for every left action of G on X,

$$\lambda : G \times X \to X,$$

there is an **associated right action** of G on X defined by

$$\lambda^\dagger : X \times G \to X, \quad \lambda^\dagger(x, g) = \lambda(g^{-1}, x).$$

Similarly, for every right action of G on X,

$$\rho : X \times G \to X,$$

there is an **associated left action** of G on X defined by

$$\rho^\dagger : G \times X \to X, \quad \rho^\dagger(g, x) = \rho(x, g^{-1}).$$

Lemma 2.5. *For every left action λ of $(G, *)$ on X, also $\lambda^\dagger$ is a right action of $(G, *)$ on X. Similarly, for every right action ρ of $(G, *)$ on X, $\rho^\dagger$ is a left action of $(G, *)$ on X. Finally, $(\lambda^\dagger)^\dagger$ equals λ and $(\rho^\dagger)^\dagger$ equals ρ.*

Proof. Note first that $e^{-1} = e$, so that $\lambda^\dagger(x, e) = \lambda(e, x) = x$. This is the first axiom for a right action. Note second that $(h * g)^{-1} = g^{-1} * h^{-1}$, so that

$$\lambda^\dagger(\lambda^\dagger(x, h), g) = \lambda(g^{-1}, \lambda(h^{-1}, x)) = \lambda(g^{-1} * h^{-1}, x) = \lambda((h * g)^{-1}, x) = \lambda^\dagger(x, h * g).$$

Thus, $\lambda^\dagger$ satisfies the second axiom for a right action. Thus $\lambda^\dagger$ is a right action of G on X. By a similar argument, $\rho^\dagger$ is a left action of G on X. Finally, since $(g^{-1})^{-1}$ equals g, note that

$$(\lambda^\dagger)^\dagger(g, x) = \lambda^\dagger(x, g^{-1}) = \lambda((g^{-1})^{-1}, x) = \lambda(g, x).$$

Thus $(\lambda^\dagger)^\dagger$ equals λ. Similarly, $(\rho^\dagger)^\dagger$ equals ρ. □

Therefore the operation of passing from a left action to the associated right action, and vice versa, are inverse operations. In this sense, right actions are equivalent to left actions. In particular, the orbits of $\lambda^\dagger$ equal the orbits of λ, each stabilizer subgroup $\mathrm{Stab}_\lambda(x)$ equals the stabilizer subgroup $\mathrm{Stab}_{\lambda^\dagger}(x)$, and the $\lambda^\dagger$-fixed locus of g equals the λ-fixed locus of g^{-1}.

For every group $(G, *)$, denote by $\mathrm{Aut}(G) \subset S(G)$ the subset of all invertible functions $\sigma : G \to G$ that happen to be group homomorphisms, i.e., for every $g, h \in G$, $\sigma(g * h)$ equals $\sigma(g) * \sigma(h)$. These are **automorphisms** of G.

Lemma 2.6. *The subset $Aut(G) \subset S(G)$ is a subgroup of the symmetric group $S(G)$.*

Proof. The identity function $\mathrm{Id}_G : G \to G$ is a group homomorphism, $\mathrm{Id}_G(g * h) = g * h = \mathrm{Id}_G(g) * \mathrm{Id}_G(h)$. Thus Id_G is an element of $\mathrm{Aut}(G)$. The composition of invertible function is an invertible function, and the composition of group homomorphisms is a group homomorphism. Thus the composition of automorphisms is an automorphism. Finally, for every automorphism $\sigma : G \to G$, for every $g, h \in G$, there exist elements $s = \sigma^{-1}(g)$ and $t = \sigma^{-1}(h)$ such that $g = \sigma(s)$ and $h = \sigma(t)$. Since σ is a group homomorphism, $\sigma(s * t)$ equals $\sigma(s) * \sigma(t) = g * h$. Since $\sigma(s * t)$ equals $g * h$, also $s * t$ equals $\sigma^{-1}(g * h)$. In other words, for every $g, h \in G$,

$$\sigma^{-1}(g * h) = s * t = \sigma^{-1}(g) * \sigma^{-1}(h).$$

Therefore the inverse function σ^{-1} is an automorphism. Thus, $\mathrm{Aut}(G)$ is a subgroup of $S(G)$. □

The induced left action,

$$\lambda_{\mathrm{Aut}(G),G} : \mathrm{Aut}(G) \times G \to G, (\sigma, g) \mapsto \sigma(g),$$

is the **automorphism action**. By definition, this action is faithful. Please note, the orbit of this action on e is precisely $\{e\}$. So, unless G is a trivial group, the automorphism action is not transitive. Similarly, the stabilizer subgroup of e equals all of $\mathrm{Aut}(G)$, whereas the stabilizer subgroups of some other elements of G might be proper subgroups of $\mathrm{Aut}(G)$. For every automorphism $\sigma : G \to G$, the fixed locus of σ is a subgroup of G.

For every group $(L, \bullet)$, for every group $(G, *)$, a left action of L on G,

$$\lambda : L \times G \to G,$$

is a **action by automorphisms** if for every $s \in L$, for every $g, h \in G$, $\lambda(s, g * h) = \lambda(s, g) * \lambda(s, h)$.

Lemma 2.7. *A left action λ of $(L, \bullet)$ on the group $(G, *)$ is an action by automorphisms if and only if the associated group homomorphism $\widetilde{\lambda} : (L, \bullet) \to (S(G), \circ)$ factors through the subgroup $Aut(G) \subset S(G)$.*

Proof. The image of $\widetilde{\lambda}(L)$ is contained in $\mathrm{Aut}(G) \subset S(G)$ if and only if for every $s \in L$, the induced invertible set function,

$$\widetilde{\lambda}_s : G \to G,$$

is a group homomorphism, i.e., $\widetilde{\lambda}_s(g * h)$ equals $\widetilde{\lambda}_s(g) * \widetilde{\lambda}_s(h)$. Expanding, this holds if and only if λ is an action by automorphisms. □

For every group $(G, *)$, the **conjugation action** of $(G, *)$ on $(G, *)$ by automorphisms is defined to be

$$c : G \times G \to G, \quad (g, x) \mapsto g * x * g^{-1}.$$

Denote $c_g(x) = c(g, x) = g * x * g^{-1}$.

Lemma 2.8. *The conjugation action is a left action of $(G, *)$ on $(G, *)$ by automorphisms.*

Proof. Please note that for every $x, y \in G$,

$$c_g(x)*c_g(y) = (g*x*g^{-1})*(g*y*g^{-1}) = (g*x)*(g*g^{-1})*(y*g^{-1}) = (g*x)*e*(y*g^{-1}) = g*(x*y)*g^{-1} = c_g(x*y).$$

Thus, c_g is a group homomorphism. Moreover, for every $g, h \in G$, and for every $x \in G$,

$$(c_g \circ c_h)(x) = c_g(c_h(x)) = c_g(h*x*h^{-1}) = g*(h*x*h^{-1})*g^{-1} = (g*h)*x*(h^{-1}*g^{-1}) = (g*h)*x*(g*h)^{-1} = c_{g*h}(x)$$

Thus, $c_g \circ c_h$ equals c_{g*h}. In particular, since c_e equals Id_G, also,

$$c_g \circ c_{g^{-1}} = \mathrm{Id}_G = c_{g^{-1}} \circ c_g.$$

Thus, c_g is an invertible function with inverse $c_{g^{-1}}$. Therefore, c_g is an automorphism of G. Therefore, there is a function

$$\widetilde{c} : G \to \mathrm{Aut}(G), \quad g \mapsto c_g.$$

Since $c_g \circ c_h$ equals c_{g*h}, the function $\widetilde{c}$ is a group homomorphism. The pullback of the automorphism action by the group homomorphism $\widetilde{c}$ is a left action of G on G by automorphisms of G, and by computation it is the conjugation action,

$$c : G \times G \to G, (g, x) \mapsto g * x * g^{-1}.$$

□

For every $x \in G$, the orbit of x under the conjugation action is called the **conjugacy class** of x in G. For every $x \in G$, the stabilizer subgroup $\mathrm{Stab}_c(x)$ is called the **centralizer** of g, $\mathrm{Stab}_c(x) = \{g \in G | g * x = x * g\}$. Similarly, for every $g \in G$, the c-fixed locus of g in G also equals the centralizer $\mathrm{Stab}_c(g)$.

For every field $(F, 0, 1, +, \cdot)$ and for every F-vector space $(V, 0, +, \cdot)$, the **general linear group** of $(V, 0, +, \cdot)$ is the subset $\mathbf{GL}(V) \subset S(V)$ of all F-linear transformations $T : V \to V$ that are invertible.

Lemma 2.9. *The general linear group $\boldsymbol{GL}(V)$ is a subgroup of $S(V)$.*

Proof. The identity function $\mathrm{Id}_V : V \to V$ is a linear transformation. The composition of two linear transformations is a linear transformation. Finally, the inverse function of an invertible linear transformation is a linear transformation. □

The restriction to $\mathbf{GL}(V)$ of the standard action λ_V is the **standard action** of $\mathbf{GL}(V)$ on V,

$$\lambda_{\mathbf{GL}(V)} : \mathbf{GL}(V) \times V \to V, \quad (T, v) \mapsto T(v).$$

A left action λ of a group $(G, *)$ on V is a **linear representation** of $(G, *)$ on $(V, 0, +, \cdot)$ if every $\widetilde{\lambda}_g : V \to V$ is a linear transformation. By definition of $\mathbf{GL}(V)$, λ is a linear representation if and only if $\widetilde{\lambda}$ factors through the subgroup $\mathbf{GL}(V) \subset S(V)$.

3 Lagrange's Theorem

Let $(H, *)$ be a group, let X be a set, and let

$$\rho : X \times H \to X$$

be a right action of H on X. Let $x \in X$ be an element, and denote by $x \cdot H$ the ρ-orbit of x. Denote by $H_x \subset H$ the stabilizer subgroup $\mathrm{Stab}_\rho(x)$. Denote by

$$q : H \to H_x\backslash H$$

the quotient map to the set of right H_x-cosets in H. This is an onto function by definition. Recall that there is a right H-action on the coset space $H_x\backslash H$,

$$\rho_{H_x\backslash H,H} : (H_x\backslash H) \times H \to H_X\backslash H, \rho(H_x g, h) = H_x(g * h).$$

Denote by ρ^x the set function

$$\rho^x : H \to H \cdot x, \quad \rho^x(h) = x \cdot h.$$

This is an onto function by definition.

Lemma 3.1. *For every $h, h' \in H$, $\rho^x(h)$ equals $\rho^x(h')$ if and only if $q(h)$ equals $q(h')$, i.e., if and only if $H_x h$ equals $H_x h'$. There is a unique bijection $\widetilde{r}^x : H_X\backslash H \to H \cdot x$ such that r^x equals $\widetilde{r}^x \circ q$. This is right H-equivariant: $\widetilde{r}^x(\rho_{H_x\backslash H,H}(H_x g, h)) = \rho(\widetilde{r}^x(H_x g), h)$ for every $H_x g \in H_x\backslash H$ and for every $h \in H$. The analogous result also holds for left actions.*

Proof. Since ρ is a right action, $\rho(x, h)$ equals $\rho(x, h')$ if and only if $\rho(x, h' * h^{-1})$ equals $\rho(x)$, i.e., if and only if $h' * h^{-1} \in H_x$. Thus $\rho(x, h)$ equals $\rho(x, h')$ if and only if $H_x * h'$ equals $H_x * h$. So $\rho^x(h)$ equals $\rho^x(h')$ if and only if $q(h)$ equals $q(h')$. Since both ρ^x and q are onto, there is a unique bijection $\widetilde{r}^x : H_x \to H$ such that r^x equals $\widetilde{r}^x \circ q$. Since $q(g * h) = \rho_{H_x\backslash H,H}(q(g), h)$ and $r^x(g * h) = \rho(r^x(g), h)$, it follows that $\widetilde{r}^x$ is right H-equivariant. □

This lemma says that every orbit $H \cdot x$ of ρ is equivalent, as a set with a right H-action, to one of the "model" right H-actions of H on a coset space $H_x\backslash H$. In particular, the cardinality of the orbit $x \cdot H$ equals the cardinality of $H_x\backslash H$. The most important special case is when X is a group $(G, *)$, H is a subgroup of G, and ρ is the right regular action of H on G. Then the ρ-orbits are the same as left H-cosets. The stabilizer group H_x is the trivial group. Thus the lemma proves that all left H-cosets are in bijection with H itself.

For every group $(H, *)$ and for every subgroup K, the **index** of K in H, denoted $[H : K]$, is the cardinality of the coset space H/K. This is infinite if the coset space H/K is infinite, and it is the number of elements of H/K if H/K is finite. Recall that because left and right actions are equivalent, the left coset space H/K is in bijection with the right coset space $K\backslash H$. By the lemma, the cardinality of $x \cdot H$ equals the index $[H : H_x]$.

Theorem 3.2 (Lagrange's Theorem)**.** *Let ρ be a right action of $(H, *)$ on a set X. If any ρ-orbit is infinite, or if the orbit space X/H is infinite, then X is infinite. If both X/H is finite, and if every orbit $x\dot{H} \in X/H$ is finite, then X is finite. In that case,*

$$\#X = \sum_{x\cdot H \in X/H} \#(x \cdot H) = \sum_{x\cdot H \in X/H} [H : H_x].$$

In particular, if the stabilizer of every $x \in X$ is the trivial subgroup, $H_x = \{e\}$, then

$$\#X = \#(X/H) \cdot \#H.$$

Proof. Every orbit is a subset of X. Thus, if any orbit is infinite, then X is infinite. The quotient function $q_\rho : X \to X/H$ is onto. Thus, if X/H is infinite, then X is finite. Therefore assume that both X/H is finite and every orbit is finite. A union of finitely many finite sets is finite. Thus, also X is finite. Since X is the disjoint union of the subsets $q_\rho^{\mathrm{pre}}(\{x \cdot H\}) = x \cdot H$ over the finitely many distinct elements $x \cdot H$ in X/H, the cardinality of X equals the sum of the cardinalities of these disjoint subsets,

$$\#X = \sum_{x\cdot H \in X/H} \#(x \cdot H).$$

By the lemma, $\#(x \cdot H)$ equals $[H : H_x]$. Thus,

$$\#X = \sum_{x\cdot H \in X/H} [H : H_x].$$

If every H_x is trivial, then every $[H : H_x]$ equals $\#H$. So, in that case,

$$\#X = \sum_{x\cdot H \in X/H} \#H = \#(X/H) \cdot \#H.$$

□

The most important special case is when X is a group $(G, *)$, H is a subgroup of G, and ρ is the right regular action of H on G. Then the stabilizer of every element is the trivial group. So, in this case, the lemma gives

$$\#G = [G : H]\#H.$$

This is the case proved in the textbook.

4 Burnside's Lemma and the Cauchy-Frobenius Theorem

Let $(G, *)$ be a group, and let H be a subgroup. Denote the quotient map to the set of left H-cosets as usual,

$$q_{G,G/H} : G \to G/H, \quad k \mapsto kH.$$

Burnside's Lemma

The **inertia subset** of $G \times G$ is defined to be

$$I_{G,H} = \{(g,k) \in G \times G | (g * k)H = kH\}.$$

Equivalently, (g,k) is contained in $I_{G,H}$ if and only if g is in the stabilizer subgroup $\mathrm{Stab}_{\lambda_{G,G/H}}(kH)$. Since the stabilizer of the neutral coset H equals the subgroup H, then the stabilizer of the coset kH is $c_k(\mathrm{Stab}_{\lambda_{G,G/H}}(H)) = c_k(H)$. Therefore, there is a bijection,

$$\Psi : H \times G \to I_{G,H}, \quad (h,k) \mapsto (c_k(h),k) = (k * h * k^{-1}, k).$$

In particular, if G is finite, then this gives an equality,

$$\#H \cdot \#G = \#I_{G,H}.$$

For every $g \in G$, denote by $I_{g,H}$ the unique subset of G such that $\{g\} \times I_{g,H}$ equals the intersection of $I_{G,H}$ with the subset $\{g\} \times G$. Thus, k is in $I_{g,H}$ if and only if kH is in the fixed locus $(G/H)^g$ of g. For the right regular action of H on G, $I_{g,H}$ is an H-preserved set. Since the stabilizers are trivial for the right regular action, Lagrange's theorem gives,

$$\#I_{g,H} = \#(G/H)^g \cdot \#H.$$

Of course $I_{G,H}$ is the disjoint union of the subsets $\{g\} \times I_{g,H}$ as g varies among the elements of G. In particular, if G is finite, this gives an equality,

$$\#I_{G,H} = \sum_{g \in G} \#I_{g,H} = \sum_{g \in G} \#(G/H)^g \cdot \#H = \#H \cdot \sum_{g \in G} \#(G/H)^g.$$

Combined with the previous identity, this gives the *basic Burnside identity*,

$$\#G = \sum_{g \in G} \#(G/H)^g.$$

Now consider the apparently more general case that Y is a nonempty set and λ is a left action of G on Y that is transitive, i.e., $Y = G \cdot y$ for some y in Y. Denote by H the stabilizer subgroup $\mathrm{Stab}_\lambda(y)$. Then there is a left G-equivariant bijection of Y with G/H. For every $g \in G$, this G-equivariant bijection identifies the fixed locus of g in Y with the fixed locus of g in G/H. Thus, if G is finite, the basic Burnside identity gives,

$$\#G = \sum_{g \in G} \#Y^g.$$

Finally, consider the general case that X is a nonempty set and λ is a left action of G on X. Then G is partitioned by the λ-orbits Y. In particular, for every $g \in G$, the g-fixed locus X^g is partitioned by the g-fixed locus Y^g of each λ-orbit Y. In particular, if G and X/G are finite, then also X is finite and we have the *Burnside identity*,

$$\#G \cdot \#(X/G) = \sum_{Y \in X/G} \#G = \sum_{Y \in X/G} \sum_{g \in G} \#Y^g = \sum_{g \in G} \sum_{Y \in X/G} \#Y^g = \sum_{g \in G} \#X^g.$$

This proves the following.

Burnside's Lemma

Lemma 4.1 (Burnside's Lemma)**.** *Let $(G, *)$ be a finite group, let X be a nonempty set, and let λ be a left action of G on X. If the orbit space X/G or any g-fixed locus X^g is infinite, then X is infinite. If the orbit space X/G and every g-fixed locus is finite, then*

$$\#G \cdot \#(X/G) = \sum_{g \in G} \#X^g.$$

This has many applications among counting arguments. The original application is to group theory, and it predates Burnside's formulation of the lemma. Let $(K, *)$ be a finite group, and let p be a prime integer that divides $n = \#K$. Let G denote a cyclic subgroup $\langle \sigma \rangle$ of order p in S_p generated by a p-cycle σ, say $\sigma = (1, 2, \ldots, p-1, p)$. There is a left action of the group G on the set K^p of ordered p-tuples of elements of K in the natural way,

$$\lambda(\sigma, (k_1, k_2, \ldots, k_{p-1}, k_p)) = (k_2, k_3, \ldots, k_p, k_1).$$

Let X denote the following subset of K^p,

$$X = \{(k_1, k_2, \ldots, k_{p-1}, k_p) \in K^p | k_1 * k_2 * \cdots * k_{p-1} * k_p = e\}.$$

Consider what happens to both sides of the equation under conjugation by k_1^{-1},

$$e = c_{k_1^{-1}}(e) = c_{k_1^{-1}}(k_1 * k_2 * \cdots * k_{p-1} * k_p) = k_1^{-1} * k_1 * k_2 * \cdots * \cdots * k_{p-1} * k_p * k_1 = k_2 * k_3 * \cdots * k_p * k_1.$$

Thus, the subset X is λ-preserved. So there is a restriction action of G on X.

For the identity element $(1) \in G$, the fixed locus is all of X. Note that X has cardinality n^{p-1}. Indeed, for every $(k_1, \ldots, k_{p-1}) \in K^{p-1}$, there is a unique choice of $k_p \in K$ such that $(k_1, \ldots, k_{p-1}, k_p) \in X$, namely

$$k_p = (k_1 * k_2 * \cdots * k_{p-1})^{-1}.$$

Thus, X is in bijection with K^{p-1} so that $\#X = \#(K^{p-1}) = n^{p-1}$.

On the other hand, for every $a \in \mathbb{Z}$ that is prime to p, there exists b such that $ab \cong 1 \pmod{p}$. Thus, $\sigma = (\sigma^a)^b$. Therefore, every element of X that is fixed by σ^a is fixed also by $\sigma = (\sigma^a)^b = \sigma^a \circ \cdots \circ \sigma^a$. Conversely, every element of X that is fixed by σ is fixed by $\sigma^a = \sigma \circ \cdots \circ \sigma$. Therefore, for the $p-1$ elements $\sigma^a \in G \setminus \{(1)\}$, the fixed locus of σ^a equals the fixed locus of σ. Therefore Burnside's Lemma gives,

$$p\#(X/G) = \#G \cdot \#(X/G) = \sum_{g \in G} \#X^g = \#X + (p-1)\#X^\sigma.$$

The first summand comes from the fixed locus of the identity element $(1) \in G$, and the $p-1$ remaining terms come from the elements $\sigma^a \in G \setminus \{(1)\}$. In particular, since $\#X = n^{p-1}$ is divisible by n^1, which in turn is divisible by p, this gives

$$(p-1)\#X^\sigma \cong 0 \pmod{p}.$$

Since $p-1$ is relatively prime to p, it follows that p divides $\#X^\sigma$.

The elements of K^p that are fixed by σ are precisely the elements of the form

$$(k_1, k_2, \ldots, k_{p-1}, k_p) = (k, k, \ldots, k, k),$$

for some element $k \in K$. The condition on such an element that it be contained in X, and hence be a σ-fixed element of K, is precisely that

$$k^p = e.$$

Thus, X^σ is bijective to the subset $\{k \in K | k^p = e\}$.

Theorem 4.2 (Cauchy-Frobenius Theorem)**.** *For every finite group $(K, *)$ whose order is divisible by a prime integer p, the subset $\{k \in K | k^p = e\}$ has cardinality divisible by p. In particular, e is not the only element in this subset, so there exists an element of K of order precisely p.*

The identity element e is in the subset, so the size of the subset is ≥ 1. Since it is divisible by p, it contains at least 2 elements. Thus there exists $k \in K$ such that $k^p = e$ and such that $k \neq e$. Since $k^p = e$, the order of k divides p. Since $k \neq e$, the order of k does not equal 1. Thus, the order of k equals precisely 1.

The Cauchy-Frobenius Theorem is a first step in the proof the Sylow Theorems, one of the key tools in the further study of finite groups.

5 Normal Subgroups and the Isomorphism Theorem

Among all subgroups of a given group $(G, *)$, the normal subgroups are characterized as those that occur as the kernel subgroup of a group homomorphism with domain equal to $(G, *)$. The relation between subgroups of $(G, *)$ and subgroups of the homomorphism image are called *isomorphism theorems*. The approach here is via left group actions. None of the material in this section is required later in the course; it is included for completeness.

For a group $(G, *)$, recall that a subgroup N is normal if for every $g \in G$, $c_g(N)$ equals N. For instance, if G is Abelian, then every subgroup of G is normal. As another example, in $S(3)$ the subgroup $A(3) = \{(1), (123), (132)\}$ is normal, yet the subgroups $\langle(12)\rangle$, $\langle(13)\rangle$, and $\langle(23)\rangle$ are not normal.

Lemma 5.1. *For a group $(G, *)$ and a subgroup H, for every $g \in G$, $c_g(H)$ equals H as subgroups of G if and only if both $c_g(H) \subseteq H$ and $c_{g^{-1}}(H) \subseteq H$. In particular, for every $g \in H$, $c_g(H)$ equals H.*

Proof. Assume first that g is an element of G such that $c_g(H) = H$. Applying the automorphism $c_{g^{-1}}$ to both sides, also $c_{g^{-1}}(c_g(H)) = c_{g^{-1}}(H)$. However, $c_{g^{-1}}$ and c_g are inverse isomorphisms.

Thus, $c_{g^{-1}}(c_g(H))$ equals H. Therefore, also $c_{g^{-1}}(H) = H$. In particular, both $c_g(H) \subseteq H$ and $c_{g^{-1}}(H) \subseteq H$.

Conversely, assume that both $c_g(H) \subseteq H$ and $c_{g^{-1}}(H) \subseteq H$. Applying c_g to the second equation, we have

$$c_g(c_{g^{-1}}(H)) \subseteq c_g(H).$$

As above, $c_g(c_{g^{-1}}(H))$ equals H. Thus $H \subseteq c_g(H)$. Since both $c_g(H) \subseteq H$ and $H \subseteq c_g(H)$, in fact, $c_g(H)$ equals H. Therefore, for every $g \in G$, $c_g(H)$ equals H if and only if both $c_g(H) \subseteq H$ and $c_{g^{-1}}(H) \subseteq H$.

Finally, for every $g \in H$ and for every $h \in H$, since H is a subgroup, $c_g(h) = g * h * g^{-1}$ is an element of H. Thus $c_g(H) \subseteq H$. Since H is a subgroup, also g^{-1} is an element of H. Thus, by the same argument, $c_{g^{-1}}(H) \subseteq H$. By the previous paragraph, $c_g(H)$ equals H. □

For a group $(G, *)$, for a subgroup H of G, an element g of G **normalizes** H if $c_g(H)$ equals H. The set of all elements $g \in G$ that normalize H is the **normalizer** of H in G, denoted $N_G(H)$. Thus, a subgroup N of G is normal if and only if $N_G(N)$ equals all of G, i.e., every element of G normalizes N.

Recall that for every set T, the **power set** of T is the set of all subsets of T, including both T and the empty set. In the standard axiomatic systems for set theory, for every set T, the power set of T is again a set (because of Russell's Paradox, this puts restrictions on the other axioms of set theory). For every group $(G, *)$ denote by Subgroup(G) the subset of the power set of G whose elements are all of the subgroups H of G. For instance, for every integer $n \geq 1$, for the cyclic group $G = \mathbb{Z}/n\mathbb{Z}$ with addition, Subgroup$(\mathbb{Z}/n\mathbb{Z})$ is the set $\{\langle [d]_n \rangle | d \geq 1, d \text{ divides } n\}$. Similarly, for $G = S(3)$,

$$\text{Subgroup}(S(3)) = \{\langle (1) \rangle, \langle (12) \rangle, \langle (13) \rangle, \langle (23) \rangle, A(3), S(3)\}.$$

The conjugation action of G on itself defines a set map,

$$c_{G,\text{Sub}(G)} : G \times \text{Subgroup}(G) \to \text{Subgroup}(G), \quad c_{G,\text{Sub}(G)}(g, H) = c_g(H).$$

Since $c_g : G \to G$ is a group automorphism, for every subgroup $H \subset G$, the restriction $c_g|_H : H \to G$ is a group homomorphism. Therefore the image $c_g(H)$ is a subgroup of G, so that $c_{G,\text{Sub}(G)}$ is a well-defined set map.

Lemma 5.2. *The set map $c_{G,Sub(G)}$ is a left action of G on Subgroup(G).*

Proof. For every subgroup H of X, $c_{G,\text{Sub}(G)}(e, H) = c_e(H) = \text{Id}_G(H) = H$. This is the first axiom for a left action. For every $g, k \in G$, $c_{G,\text{Sub}(G)}(g, c_{G,\text{Sub}(G)}(k, H)) = c_g(c_k(H)) = c_{g*k}(H) = c_{G,\text{Sub}(G)}(g * k, H)$. This is the second axiom for a left action. □

By definition, for every H in Subgroup(G), the stabilizer subgroup $\mathrm{Stab}_{c_{G,\mathrm{Sub}(G)}}(H)$ equals $N_G(H)$. For every subgroup K of G, denote by $c_{K,\mathrm{Sub}(G)}$ the restriction to $K \subset G$ of the left action of G on Subgroup(G),

$$c_{K,\mathrm{Sub}(G)} : K \times \mathrm{Subgroup}(G) \to \mathrm{Subgroup}(G), \quad c_{K,\mathrm{Sub}(G)}(k, H) = c_k(H).$$

The stabilizer subgroup $\mathrm{Stab}_{c_{K,\mathrm{Sub}(G)}}(H)$ of K equals $K \cap \mathrm{Stab}_{c_{G,\mathrm{Sub}(G)}}(H) = K \cap N_G(H)$. Denote this subgroup of K by $N_K(H)$.

Lemma 5.3. *The subgroup $N_K(H)$ of K is the unique maximal subgroup of K such that $N_K(H)H$ is a subgroup of G that contains H as a normal subgroup. In particular, for every subgroup K of G that contains H, $N_K(H)$ is the unique maximal subgroup of K that contains H as a normal subgroup. The coset space $G/N_G(H)$ is bijective to the $c_{G,Sub}$-orbit of H in Subgroup(G).*

Proof. Since $N_G(H)$ is the stabilizer of H, $G/N_G(H)$ is bijective to the orbit of H in Subgroup(G) by Lagrange's Theorem.

For every subgroup K of G, $K \cap N_G(H)$ is a subgroup of K that equals the set of all $k \in K$ such that $c_k(H) = H$, and this is equivalent to $c_{k^{-1}}(H) = H$ since c_k and $c_{k^{-1}}$ are inverse bijections. Thus, $N_K(H)H$ is a subgroup of G: $(k' * h') * (k * h) = (k' * k) * (c_{k^{-1}}(h') * h)$ and $(k * h)^{-1} = k^{-1} * c_k(h^{-1})$. This subgroup contains $H = \{e\}H$. For every $k*h$ in $N_K(H)H$, $c_{k*h}(H) = c_k(c_h(H)) = c_k(H) = H$. Thus, H is a normal subgroup of $N_K(H)H$.

For every subgroup L of K such that LH is a subgroup of G that contains H as a normal subgroup, then every element k of L satisfies $c_k(H) \subset H$ since $k \in LH$ and H is normal in LH. Since also k^{-1} is in L, also $c_{k^{-1}}(H) \subset H$. Therefore k is an element of $N_G(H)$, i.e., k is an element of $K \cap N_G(H)$. Thus L is a subgroup of $K \cap N_G(H) = N_K(H)$. Therefore $N_K(H)$ is the unique maximal subgroup of K such that $N_K(H)H$ is a subgroup of G that contains H as a normal subgroup.

In particular, for every subgroup K that contains H, since H is a subgroup of $N_G(H)$, also H is a subgroup of $K \cap N_G(H) = N_K(H)$. Thus $N_K(H)H$ equals $N_K(H)$. Therefore $N_K(H)$ is the unique maximal subgroup of K that contains H as a normal subgroup. □

For every group $(G, *)$ and for every subgroup H of G, for the quotient set function

$$q_H : G \to G/H, \quad q_H(g) = gH,$$

there is a unique left action of G on H such that q is left G-equivariant,

$$c_{G,G/H} : G \times (G/H) \to G/H, \quad c_{G,G/H}(g, kH) = c_{G,G/H}(g, q_H(k)) = q_H(g * k) = (g * k)H.$$

Theorem 5.4 (First Isomorphism Theorem)**.** *A subgroup N of G is normal if and only if $N_G(N)$ equals N. In this case, there is a unique group structure $\bullet$ on G/N such that q_N is a group homomorphism. Moreover, every left N-coset equals a right N-coset. Finally, the rule that associates to every subgroup L of G/N the preimage subgroup $q_N^{pre}(L)$ is an order-preserving bijection between the subgroups of G/N and the subgroups K of G that contain N. The preimage under q_N of a normal subgroup of G/N is a normal subgroup of G.*

Proof. This has mostly been discussed in the exercises and above. For every $g \in G$, for every $n \in N$, and for every $k \in G$, $g*n*k$ equals $g*k*c_{k^{-1}}(n)$. Since N is normal, $c_{k^{-1}}(n)$ is an element of N. Thus, $q_N(g*n*k)$ equals $q_N(g*k)$. Thus, there is a well-defined set function,

$$\bullet : (G/N) \times (G/N) \to G/N, \quad q_N(g) \bullet q_N(k) = q_N(g*k).$$

This is the unique set function such that $q_N(g*k)$ equals $q_N(g) \bullet q_N(k)$ for every $g, k \in G$. In particular, $q_N(g) \bullet q_N(e) = q_N(g*e) = q_N(g)$ and $q_N(e) * q_N(g) = q_N(e*g) = q_N(g)$. Thus $q_N(e) = N$ satisfies the identity axiom of a group. Also, for every $g, h, k \in G$,

$$(q_N(g) \bullet q_N(h)) \bullet q_N(k) = q_N(g*h) \bullet q_N(k) = q_N((g*h)*k) =$$

$$q_N(g*(h*k)) = q_N(g) \bullet (q_N(h*k)) = q_N(g) \bullet (q_N(h) \bullet q_N(k)).$$

Thus, the group operation satisfies the axiom of associativity. Finally, for every $g \in G$, $q_N(g) \bullet q_N(g^{-1}) = q_N(g*g^{-1}) = q_N(e)$ and $q_N(g^{-1}) \bullet q_N(g) = q_N(g^{-1}*g) = q_N(e)$ (technically that is redundant). Thus the group operation also satisfies the inverse axiom. Therefore $\bullet$ is a group operation on G/N.

Next, for every $g \in G$, since $c_g(N)$ equals N, $gN = (gNg^{-1})g = c_g(N)g = Ng$. Therefore every left N-coset equals a right N-coset.

For every subgroup L of G/N, since q_N is a group homomorphism, $q_N^{\mathrm{pre}}(L)$ is a subgroup of G. Since L contains $q(e) = q(N)$, $q_N^{\mathrm{pre}}(L)$ is a subgroup that contains N. Since preimage preserves set inclusion and intersection, this rule preserves inclusion of subgroups and intersection of subgroups. Similarly, for every subgroup K of G, $q_N(K)$ is a subgroup of G/N, and this rule also preserves inclusion of subgroups. Since q_N is onto, $q_N(q_N^{\mathrm{pre}}(L))$ equals L. Finally, for every subgroup K of G that contains N, for every $k \in K$, since K is a subgroup, K contains $kN = q_N^{\mathrm{pre}}(q_N(\{k\}))$. Thus, K equals $q_N^{\mathrm{pre}}(q_N(K))$. Therefore the operation of preimage group, $L \mapsto q_N^{\mathrm{pre}}(L)$, and of image group, $K \mapsto q_N(K)$, are inverse bijections between the set of all subgroups of G/N and the set of all subgroups of G that contains N. It was proved in lecture that for every group homomorphism q, the preimage under q of a normal subgroup is a normal subgroup of G. □

Let $(G, *)$ be a group, let H be a subgroup, and let $K \subset N_G(H)$ be a subgroup. Thus, KH is a subgroup of G that contains H as a normal subgroup. For the quotient set map

$$q_{KH,H} : KH \to (KH)/H, \quad q_{KH,H}(g) = gH,$$

there is a unique structure of group operation $\bullet$ on $(KH)/H$ such that $q_{KH,H}$ is a group homomorphism. This group homomorphism restricts on the subgroup $K \subset KH$ to a group homomorphism,

$$q_{K,H} : K \to (KH)/H, \quad q_{K,H}(g) = gH.$$

The kernel of $q_{K,H}$ equals $K \cap \mathrm{Ker}(q_{KH,H}) = K \cap H$, and this is a normal subgroup of K. For the quotient set map,

$$q_{K,K\cap H} : K \to K/(K \cap H), \quad q_{K,K\cap H}(g) = g(K \cap H),$$

there is a unique group operation $\bullet$ on $K/(K \cap H)$ such that $q_{K,K\cap H}$ is a group homomorphism. By Lagrange's Theorem, there is a left K-equivariant bijection

$$q_{K\cap H,KH} : K/(K \cap H) \to q_{K,H}(K), \quad q_{K\cap H,KH}(g(K \cap H)) = gH.$$

Theorem 5.5 (Second Isomorphism Theorem)**.** *The group homomorphism $q_{K,H}$ is onto, and the bijection $q_{K\cap H,KH}$ is a group isomorphism from $K/(K \cap H)$ to KH/H.*

Proof. By definition, every element of KH is of the form $k*h$ for some $k \in K$ and for some $h \in H$. Thus, $q_{KH,H}(k*h)$ equals $q_{KH,H}(k)$. This in turn equals $q_{K,H}(k)$. Therefore every element of KH/H is in the image of $q_{K,H}$, i.e., $q_{K,H}$ is onto. By Lagrange's Theorem, $q_{K\cap H,KH}$ is a bijection. To prove that $q_{K\cap H,KH}$ is a group isomorphism, it suffices to prove that $q_{K\cap H,KH}$ is a group homomorphism.

Since $q_{K,K\cap H}$ is onto, to prove that $q_{K\cap H,KH}$ is a group homomorphism, it suffices to prove for every $g, k \in K$ that $q_{K\cap H,KH}(q_{K,K\cap H}(g) \bullet q_{K,K\cap H}(k))$ equals $q_{K\cap H,KH}(q_{K,K\cap H}(g)) \bullet q_{K\cap H,KH}(q_{K,K\cap H}(k))$. Since $q_{K,K\cap H}$ is a group homomorphism, $q_{K,K\cap H}(g) \bullet q_{K,K\cap H}(k)$ equals $q_{K,K\cap H}(g*k)$. By definition of $q_{K\cap H,KH}$, $q_{K\cap H,KH}(q_{K,K\cap H}(j))$ equals $q_{KH,H}(j)$ for every $j \in K$. Thus, $q_{K\cap H,KH}(q_{K,K\cap H}(g*k))$ equals $q_{KH/H}(g*k)$. Since $q_{KH/H}$ is a group homomorphism, this equals $q_{KH,H}(g) \bullet q_{KH,H}(k)$. Reversing the steps, this equals $q_{K\cap H,KH}(q_{K,K\cap H}(g)) \bullet q_{K\cap H,KH}(q_{K,K\cap H}(k))$, as was to be shown. □

Let $(G, *)$ be a group, let X be a nonempty set, and let

$$\lambda : G \times X \to X,$$

be a left action of G on X. Let $N \subseteq G$ be a normal subgroup. Denote the quotient group homomorphism from G to G/N by

$$q_{G,N} : G \to G/N, \quad q_{G,N}(g) = gN.$$

An **induced left action** of G/N on X is a left group action

$$\lambda_{G/N,X} : (G/N) \times X \to X,$$

such that for every $g \in G$ and for every $x \in X$, $\lambda_{G/N,X}(q_{G,N}(g), x)$ equals $\lambda(g, x)$.

Proposition 5.6. *For a normal subgroup N of G, for every $x \in X$, if N is contained in $Stab_\lambda(x)$, then for every y in the λ-orbit of x, N is a subgroup of $Stab_\lambda(y)$. There is an induced left action of G/N on X if and only if N is contained in $Stab_\lambda(y)$ for every $y \in X$. In this case, the induced left action is unique.*

Proof. By definition of orbit, for every y in the λ-orbit of x, there exists $g \in G$ with $y = \lambda(g, x)$. Then $\mathrm{Stab}_\lambda(y)$ equals $c_g(\mathrm{Stab}_\lambda(x))$. Since N is contained in $\mathrm{Stab}_\lambda(x)$, also $c_g(N)$ is contained in $c_g(\mathrm{Stab}_\lambda(x))$. Since N is normal, $c_g(N)$ equals N. Therefore N is contained in $\mathrm{Stab}_\lambda(y)$.

Let N be a normal subgroup of G that is contained in $\mathrm{Stab}_\lambda(y)$ for every $y \in X$. Since N is contained in $\mathrm{Stab}_\lambda(y)$, for every $g \in G$ and for every $k \in N$, $\lambda(g * k, y) = \lambda(g, \lambda(k, y)) = \lambda(g, y)$. Thus, there is a unique well-defined set map

$$\lambda_{G/N,X} : (G/N) \times X \to X, \quad \lambda_{G/N,X}(gN, y) = \lambda(g, y),$$

such that for every $g \in G$ and for every $y \in X$, $\lambda_{G/N,X}(q_N(g), y)$ equals $\lambda(g, y)$, i.e., $\lambda_{G/N,X}$ is an induced left action of G/N on X. Note that this is indeed a left action: $\lambda_{G/N,X}(q_N(e), y) = \lambda(e, y) = y$ and $\lambda_{G/N,X}(q_N(g), \lambda_{G/N,X}(q_N(k), y)) = \lambda(g, \lambda(k, y)) = \lambda(g * k, y) = \lambda_{G/N,X}(q_N(g * k), y) = \lambda_{G/N,X}(q_N(g) \bullet q_N(k), y)$. Since λ is a left action of G on X, also $\lambda_{G/N,X}$ is a left action of G/N on X.

Conversely, if $N \subseteq G$ is a normal subgroup, and if $\lambda_{G/N,X}$ is an induced left action of G/N on X, then for every $y \in X$, for every $n \in N$, $\lambda(n, y) = \lambda_{G/N,X}(q_N(n), y) = \lambda_{G/N,X}(q_N(e), y) = \lambda(e, y) = y$. Thus, N is contained in $\mathrm{Stab}_\lambda(y)$ for every $y \in X$. □

Theorem 5.7 (Third Isomorphism Theorem)**.** *Let λ be a transitive left action of G on a set X. Let $N \subset G$ be a normal subgroup, and assume that there exists an induced left action $\lambda_{G/N,X}$ of G/N on X. Then the action $\lambda_{G/N,X}$ is transitive, and $Stab_{\lambda_{G/N,X}}(x)$ equals $q_{G,N}(Stab_\lambda(x))$ for every $x \in X$. There is a unique left G-equivariant bijection $q_{G/N,x} : (G/N)/q_{G,N}(Stab_\lambda(x)) \to X$ sending the coset $q_{G,N}(Stab_\lambda(x))$ to x.*

Proof. Since the action λ is transitive, for every $x, y \in X$, there exists $g \in G$ such that y equals $\lambda(g, x)$. By the defining property of $\lambda_{G/N,X}$, also $\lambda_{G/N,X}(gN, x)$ equals y. Thus $\lambda_{G/N,X}$ is transitive. By Lagrange's Theorem, there is a unique G-equivariant bijection

$$(G/N)/\mathrm{Stab}_{\lambda_{G/N,X}}(x) \to X$$

sending the coset $\mathrm{Stab}_{\lambda_{G/N,X}}(x)$ to x. Finally, by definition of $\lambda_{G/N,X}$, $\lambda_{G/N,X}(q_{G,N}(g), x)$ equals $\lambda(g, x)$. Thus, $q_{G,N}(g)$ is in the $\lambda_{G/N,X}$-stabilizer of x if and only if g is in the λ-stabilizer of x, i.e., $q_{G,N}(\mathrm{Stab}_\lambda(x))$ equals the $\lambda_{G/N,X}$-stabilizer of x. □

Let $(G, *)$ be a group, and let $H \subset G$ be a subgroup. Define X to be G/H with its natural left action $\lambda_{G,G/H}$ of G on G/H. This is a transitive action. The induced group homomorphism

$$\widetilde{\lambda}_{G,G/H} : G \to S(X),$$

has kernel $N = N_{G,H}$ equal to a normal subgroup of G that is contained in H. In particular, G/N is a quotient group of G, and H/N is a subgroup of G/N such that the coset space $(G/N)/(H/N)$ is bijective to G/H.

Lemma 5.8. *The kernel of $\widetilde{\lambda}_{G,G/H}$ equals the common intersection over all $g \in G$ of $c_g(H)$. If H has finite index m in G, then the index of N in G is a finite integer that divides $m!$, and the index of H/N in G/N equals m.*

Proof. For a left action λ of G on a set X, by the definition of the associated group homomorphism $\widetilde{\lambda} : G \to S(X)$, the kernel of $\widetilde{\lambda}$ equals the intersection over all $x \in X$ of $\mathrm{Stab}_\lambda(x)$. For the transitive action $\lambda_{G,G/H}$, for every $g \in G$, the stabilizer of $gH \in G/H$ equals $c_g(H)$. Thus N equals the intersection over all $g \in G$ of $c_g(H)$.

If H has finite index m in G, then X is a set with m elements by Lagrange's Theorem. Thus the symmetric group $S(X)$ is isomorphic to $S(m)$. This has $m!$ elements. Thus the image group $\widetilde{\lambda}_{G,H}(G) \subset S(X)$ is a subgroup of the finite group $S(X)$ that has order dividing $m!$, again by Lagrange's Theorem. By Lagrange's Theorem once more, the image of G in $S(X)$ is naturally bijective to G/N. Thus, the index of N in G divides $m!$, and equals the order of the image of G in $S(X)$. By the previous lemma, the index of H/N in G/N equals the index m of H in G. □

Let $(G, *)$ be a group, let $K \subset G$ be a subgroup, and let $H \subset K$ be a subgroup. Denote the quotient set maps to the coset spaces by

$$q_{G,G/K} : G \to G/K, \quad q_{G,G/K}(g) = gK,$$

$$q_{G,G/H} : G \to G/H, \quad q_{G,G/H}(g) = gH.$$

For every $g \in G$, gK contains gH. Thus, for every $g' \in gH$, $q_{G,G/K}(g')$ equals $q_{G,G/K}(g)$. Therefore, there exists a well-defined set function,

$$q_{G/H,G/K} : G/H \to G/K, \quad q_{G/H,G/K}(gH) = gK.$$

This is the unique set function such that $q_{G/H,G/K} \circ q_{G,G/H}$ equals $q_{G,G/K}$. For every $g \in G$, the preimage under $q_{G/H,G/K}$ over $\{gK\}$ is precisely $q_{G,G/H}(gK)$. For every $g \in G$, there is a bijection,

$$l_g : K/H \to q_{G,G/H}(gK), l_g(kH) = (g * k)H.$$

Proposition 5.9. *If K has infinite index in G, or if H has infinite index in K, then also H has infinite index in G. If both $[G : K]$ and $[K : H]$ are finite, then also $[G : H]$ is finite, and $[G : H]$ equals $[G : K][K : H]$.*

Proof. Since $q_{G/H,G/K}$ is onto, if G/K is infinite, then G/H is infinite. Similarly, since every fiber of $q_{G/H,G/K}$ is in bijection with K/H via l_g, if K/H is infinite then also G/H is infinite. On the other hand, if all of these are finite, then the fibers of the onto map $q_{G/H,G/K}$ form a partition of G/H. Thus,

$$\#(G/H) = \sum_{gK \in G/K} \#q^{\mathrm{pre}}_{G/H,G/K}(\{gK\}) = \sum_{gK \in G/K} \#(K/H) = \#(G/K)\#(K/H),$$

i.e., $[G : H]$ equals $[G : K][K : H]$. □

6 The Sylow Theorem

In lecture, the Sylow theorem was only mentioned in passing. The discussion here is only for completeness. The general Sylow theorem is used in the proof of the Structure Theorem for Finite Abelian Groups, but that theorem can be proved with much less than the general Sylow theorem.

For every prime integer p, a **p-group** is a group such that every element has finite order equal to a power of p (possibly $1 = p^0$ for the identity, and the power certainly depends on the element).

Lemma 6.1. *Every subgroup of a p-group is a p-group. For every group homomorphism $f : (P, *) \to (R, \bullet)$, P is a p-group if and only if both $f(P)$ and $Ker(f)$ are p-groups. A finite group P is a p-group if and only if the order of P is a power of p.*

Proof. Let $(P, *)$ be a p-group. For every subgroup H of P, since every element of P has finite order equal to a power of p, in particular, every element of H has finite order equal to a power of p. Thus H is a p-group.

Let $f : P \to R$ be a group homomorphism. First assume that P is a p-group. By the previous paragraph, the subgroup $\mathrm{Ker}(f)$ is a p-group. Next, for every $g \in P$, since there exists an integer $n \geq 0$ with $g^{p^n} = e$, then also $f(g)^{p^n} = f(e)$. Thus the order of $f(g)$ is finite and divides p^n. Therefore, the order of $f(g)$ is a power of p. So every $f(g) \in f(P)$ has finite order equal to a power of p, i.e., $f(P)$ is a p-group.

Next, assume that both $f(P)$ and $\mathrm{Ker}(f)$ are p-groups. For every $g \in P$, there exists an integer $n \geq 0$ such that $f(g)^{p^n} = f(e)$. Thus g^{p^n} is an element of $\mathrm{Ker}(f)$. Since $\mathrm{Ker}(f)$ is a p-group, there exists an integer $m \geq 0$ such that $(g^{p^n})^{p^m}$ equals e. Altogether, $g^{p^{n+m}}$ equals e. Therefore the order of g is finite and divides p^{n+m}. So for every $g \in G$, the order of g is finite and equals a power of p, i.e., G is a p-group.

If P is a finite group of order p^n, then by Lagrange's theorem, the order of every element of P is finite and divides p^n. Thus, the order of every element of P is finite and equal to a power of p, i.e., P is a p-group. Conversely, assume that P is a finite group whose order does not equal a power of p. By the Unique Factorization of Integers, there exists a prime integer $q \neq p$ such that q divides the order of P. By the Cauchy-Frobenius Theorem, there exists an element g of P of order equal to q. Thus P is not a p-group. Therefore, for a finite group P, P is a p-group if and only if the order of P equals a power of p. □

Corollary 6.2. *Let $(P, *)$ be a p-group. For every subgroup H of G of finite index, $[G : H]$ equals a power of p. Similarly, for every finite subgroup H of G, $\#H$ equals a power of p.*

Proof. By the previous lemma, every subgroup H of G is a p-group. If H is finite, then the previous lemma implies that $\#H$ equals a power of p.

Next let H be a subgroup of finite index. By the previous section, there exists a normal subgroup $N \subset P$ such that N is contained in H, and such that N has finite index in P. By the previous lemma, the quotient group P/N is a p-group, and it is finite. Thus P/N has order equal to a power

of p. In particular, the index of H/N in P/N divides the order of P/N by Lagrange's theorem. Thus the index of H/N in P/N also equals a power of p. By the third isomorphism theorem in the previous section, the index of H in P equals the index of H/N in P/N. Therefore the index of H in P equals a power of p. □

For a group $(G, *)$ a **p-subgroup** is a subgroup P of G such that $(P, *)$ is a p-group. A p-subgroup P of G is a **p-Sylow subgroup** of G if every p-subgroup Q of G that contains P equals P, i.e., P is a p-subgroup of G that is maximal with respect to set inclusion of p-subgroups.

Proposition 6.3. *Let $(G, *)$ be a group, and let $P \subset G$ be a p-subgroup. For every element $g \in G$ of p-power order such that $c_g(P) = P$, the subset $Q = \langle g \rangle P$ is a p-subgroup of G that contains P. If P is a p-Sylow subgroup, then g is an element of P.*

Proof. By the homework exercises, since $c_g(P)$ equals P, $\langle g \rangle P$ is a subgroup Q of G. By construction, P is a normal subgroup of Q. Thus the left Q-action on the P-coset space Q/P extends uniquely to a group structure on H/P such that the quotient function $q : Q \to Q/P$ is a group homomorphism. Since Q equals $\langle g \rangle$, the restriction of q to the subgroup $\langle g \rangle \subset Q$ is a group homomorphism $q_g : \langle g \rangle \to H/P$ that is surjective. In particular, the order of Q/P divides the order of $\langle g \rangle$. By hypothesis, the order of $\langle g \rangle$ is a power of p. Thus the same is true of the order of Q/P, i.e., the index $[Q : P]$ is a power of p (possibly $1 = p^0$). By Lagrange's Theorem, the order of Q equals $[Q : P] \# P$. Since both factors are powers of p, also the order of Q is a power of P. Thus Q is a p-subgroup of G that contains P. If P is a p-Sylow subgroup of G, then Q equals P. In that case, since g is an element of $Q = \langle g \rangle P$, also g is an element of P. □

For every element $g \in G$, conjugation by g is an automorphism of G with inverse automorphism $c_{g^{-1}}$. In particular, for every finite subgroup H, the order of $c_g(H)$ equals the order of H. Similarly, for every subgroup H of finite index in G, the index of $c_g(H)$ equals the index of H. In particular, for every p-subgroup P, also $c_g(P)$ is a p-subgroup of G.

Lemma 6.4. *For every $g \in G$, for every p-Sylow subgroup P of G, also $c_g(P)$ is a p-Sylow subgroup of G.*

Proof. Let Q be a p-subgroup of G that contains the p-subgroup $c_g(P)$. Then $c_{g^{-1}}(Q)$ is a p-subgroup of G that contains the p-subgroup P. If P is a p-Sylow subgroup, then $c_{g^{-1}}(Q)$ equals p, and then also $c_g(P)$ equals Q. □

Denote by $\mathrm{Syl}_p(G)$ the subset of $\mathrm{Subgroup}(G)$ whose elements are all of the p-Sylow subgroups of G. Please note, if the only element of G of p-power order is the identity element e (of order $1 = p^0$), then $\mathrm{Syl}_p(G)$ has only one element $\{e\}$. However, for every finite group $(G, *)$ whose order is divisible by p, then G contains elements of order p, hence $\mathrm{Syl}_p(G)$ is a nonempty finite set, every element of which is a p-Sylow subgroup of order $\geq p$.

For instance, for every integer $n \geq 1$, writing $n = p^r m$ with $\gcd(m,p) = 1$, then $\mathrm{Syl}_p(G)$ is the singleton subset containing the element $\langle [m]_n \rangle \cong \mathbb{Z}/p^r\mathbb{Z}$. Similarly, for the symmetric group $G = S(3)$ with order $6 = 2 \cdot 3$, the nontrivial p-Sylow subgroups are,

$$\mathrm{Syl}_2(S(3)) = \{\langle (12) \rangle, \langle (13) \rangle, \langle (23) \rangle\}, \quad \mathrm{Sym}_3(S(3)) = \{\langle (123) \rangle\} = \{\langle (132) \rangle\}.$$

Lemma 6.5. *For the conjugation action $c_{G,Sub}$, the subset $Syl_p(G)$ of Subgroup(G) is a $c_{G,Sub}$-preserved subset.*

Proof. By the previous lemma, $\mathrm{Syl}_p(G)$ is a $c_{G,\mathrm{Sub}}$-preserved subset of Subgroup(G). □

For every p-Sylow subgroup P of G, for every element g of $N_G(P)$ whose order is a power of p, g is an element of P by Proposition 6.3. Therefore, the normal subgroup P of $N_G(P)$ is the unique p-Sylow subgroup of $N_G(P)$, and every p-subgroup of $N_G(P)$ is contained in P.

Lemma 6.6. *Let P and Q be p-subgroups of G. Then $N_Q(P)P$ is a p-subgroup of G that contains P as a normal subgroup. If P is a p-Sylow subgroup, then $N_Q(P)$ equals $Q \cap P$.*

Proof. By the previous result, $K := N_Q(P)P$ is a subgroup of P that contains P as a normal subgroup. It only remains to prove that K is a p-group. Since P is a normal subgroup of K, there is a unique group structure on the coset space K/P such that the quotient map $q : K \to K/P$ is a group homomorphism, say

$$\bullet : (K/P) \times (K/P) \to K/P, \quad q(k) \bullet q(k') = q(k * k').$$

For every element $g \in N_Q(P) \subset Q$ and $h \in P$, $q(g * h)$ equals $q(g) \bullet q(h) = q(g) * q(e) = q(g)$.

Consider the quotient group homomorphism $q_{K,K/P} : K \to K/P$. By the Second Isomorphism Theorem, the restricted group homomorphism $q_{K,K/P}|_N : N_Q(P) \to K/P$ is onto and has kernel $N_Q(P) \cap P$. Since $N_Q(P)$ is a subgroup of Q, by Lemma 6.1, $N_Q(P)$ is a p-group. By that same lemma, the homomorphic image K/P of $q_{K,K/P}|_N$ is also a p-group. Since P is a p-group and K/P is a p-group, by the lemma once more, also K is a p-group.

Now assume that P is a p-Sylow subgroup. Since K is a p-group that contains P, K equals P. In particular, since $N_Q(P)$ is contained in K, $N_Q(P)$ is contained in P. Thus, $N_Q(P)$ is contained in $Q \cap P$. On the other hand, since P is a subgroup of $N_G(P)$, also $Q \cap P \subset Q \subset N_G(P) = N_Q(P)$. Therefore, $N_Q(P)$ equals $Q \cap P$. □

For the $c_{G,\mathrm{Sub}}$-preserved subset $\mathrm{Syl}_p(G)$ of Subgroup(G), denote by $c_{G,p}$ the restriction of $c_{G,\mathrm{Sub}}$,

$$c_{G,p} : G \times \mathrm{Syl}_p(G) \to \mathrm{Syl}_p(G), \quad c_{G,p}(g,P) = c_g(P).$$

This is a left action of G on $\mathrm{Syl}_p(G)$.

Let $X \subset \mathrm{Syl}_p(G)$ be a $c_{G,p}$-orbit, i.e., for every $P \in X$, X equals the set of all conjugates $c_g(P)$ of P. The left action of G on $\mathrm{Syl}_p(G)$ restricts to a left action of G on X, say

$$c_{G,X} : G \times X \to X, \quad c_{G,X}(g,P) = c_g(P).$$

For every subgroup K of G, denote by $c_{K,X}$ the restriction of $c_{G,X}$ to a left action of K on X,

$$c_{K,X} : G \times X \to X, \quad c_{K,X}(g, P) = c_g(P).$$

If G is a finite group, then $\mathrm{Syl}_p(G)$ is a subset of the power set of G, and thus also $\mathrm{Syl}_p(G)$ is finite. Therefore every orbit X is also finite.

Theorem 6.7 (The Sylow Theorem)**.** *Assume that there exists a $c_{G,p}$-orbit X of $Syl_p(G)$ that has finite cardinality n_X. Then for every P in X, $N_G(P)$ is a subgroup of G that has finite index $[G : N_G(P)] = n_X$. The positive integer n_X satisfies $n_X \equiv 1 \pmod p$. If also P has finite index in G, then $[N_G(P) : P]$ is prime to p so that $[G : P]$ is also prime to p. Finally, for every p-subgroup Q of G, there exists a p-Sylow subgroup P in X such that P contains Q. In particular, every p-Sylow subgroup of G is an element of X, i.e., $Syl_p(G)$ equals X.*

Proof. Let X be a $c_{G,p}$-orbit of $\mathrm{Sym}_p(G)$ that has finite cardinality $n_X \geq 1$. Let P be a p-Sylow subgroup that is an element of X. The stabilizer subgroup $\mathrm{Stab}_{c_{G,p}}(P)$ equals $N_G(P)$. Thus, by Lagrange's Theorem, $N_G(P)$ is a finite index subgroup with $[G : N_G(P)]$ equal to n_X.

For every p-subgroup Q of G, consider the $c_{Q,X}$-orbit of P. This is a subset of the $c_{G,X}$-orbit of P, i.e., the $c_{Q,X}$ orbit of P is a subset of X. Since X is finite, also the $c_{Q,X}$-orbit is finite. Thus the stabilizer subgroup $N_Q(P) = Q \cap N_G(P)$ is a finite index subset of Q. By Corollary 6.2, the index of $N_Q(P)$ in Q is a power of p. In particular, that power equals 1 if and only if Q is contained in $N_G(P)$. By Proposition 6.3, Q is contained in $N_G(P)$ if and only if Q is contained in P. In summary, the $c_{Q,X}$-orbit of P has cardinality a power of p, and that power equals 1 if and only if Q is contained in P.

Now consider the special case that Q equals P. In this case, since P is contained in P, the $c_{P,X}$-orbit of P equals $\{P\}$ and has size 1. For every P' in X with $P' \neq P$, then P is not contained in P', since this would contradict that P is a p-subgroup that is maximal for set inclusion. Thus, the $c_{P,X}$-orbit of P' has size equal to a power of p, in particular, it is divisible by p. Thus, for the action $c_{P,X}$ of P on X, there is precisely one orbit of size 1, and every other orbit has size equal to a multiple of p. Therefore, since the size n_X of X equals the sum of the sizes of distinct orbits by Lagrange's Theorem,

$$n_X \equiv 1 \pmod p.$$

Recall that P is the unique p-Sylow subgroup of $N_G(P)$, and P is normal in $N_G(P)$. For any p-subgroup R of $N_G(P)/P$, the preimage subgroup of $N_G(P)$ is a p-subgroup by Lemma 6.1, and it contains P. Thus, the preimage subgroup of $N_G(P)$ equals P. So the only p-subgroup of $N_G(P)/P$ is the trivial subgroup $\{e\}$. If $N_G(P)/P$ is finite, then by the Cauchy-Frobenius Theorem, p does not divide the order of $N_G(P)/P$, i.e., $[N_G(P) : P]$ is relatively prime to p. Since also $n_X = [G : N_G(P)]$ is relatively prime to p, by Proposition 5.9, also $[G : P]$ is a product of integers that are relatively prime to p. Since p is prime, $[G : P]$ is relatively prime to p.

Next let Q be any p-group, and consider the action $c_{Q,X}$ of Q on X. For every P in X such that Q is not contained in P, the $c_{Q,X}$-orbit of P is a power of p, hence is divisible by P. Thus, for the

$c_{Q,X}$-preserved subset X_Q of X that equals all P with Q not contained in P, the size of X_Q equals the sum over the distinct $c_{Q,X}$-orbits of X_Q of the size of that orbit, and that size is a multiple of p. Therefore the size of X_Q is a multiple of p. Since the size n_X of X is congruent to 1 modulo P, X_Q does not equal all of X. Therefore, there exists P in X that is not in X_Q, i.e., P is an element of X such that Q is contained in P. □

There are versions of Sylow's Theorem even when $\mathrm{Syl}_p(G)$ is infinite. However, there do exist infinite groups such that $\mathrm{Syl}_p(G)$ contains two distinct $c_{G,p}$-orbits.

A typical application of the Sylow Theorem is the following result.

Corollary 6.8. *Let p and q be distinct prime integers with $q < p$ and with $p \not\equiv 1 \pmod{q}$. Every group of order pq is isomorphic to $\mathbb{Z}/pq\mathbb{Z}$.*

Proof. Let $(G, *)$ be a group of order pq. Then Subgroup(G) is finite, so also Syl_p is finite. Therefore, by the Sylow Theorem, $n_p(G) = \#\mathrm{Syl}_p$ is congruent to 1 modulo p and divides $\#G = pq$. Thus $n_p(G)$ equals 1 or q. Since $q < p$, certainly $q \neq 1 \pmod{p}$. Thus, $n_p(G)$ equals 1. So there exists a unique p-Sylow subgroup P that is necessarily normal.

Similarly, $n_q(G)$ divides p, so $n_q(G)$ equals 1 or p. Since also $n_q(G)$ is congruent to 1 modulo q, by hypothesis, $n_q(G)$ does not equal p. Thus $n_q(G)$ equals 1. So there exists a unique q-Sylow subgroup Q that is necessarily normal.

For every $g \in P$ and for every $h \in Q$, $c_h(g^{-1})$ is an element of P since P is a normal subgroup, and $c_g(h)$ is an element of Q since Q is a normal subgroup. Thus $g * h * g^{-1} * h^{-1} = c_g(h) * h^{-1}$ is an element of the subgroup Q. Similarly, $g * h * g^{-1} * h^{-1} = g * c_h(g^{-1})$ is an element of P. Thus $g * h * g^{-1} * h^{-1}$ is an element of $P \cap Q$. Since $P \cap Q$ is a subgroup of the p-group P, $P \cap Q$ is a p-group. Since $P \cap Q$ is a subgroup of the q-group Q, $P \cap Q$ is a q-group. Thus $P \cap Q$ is $\{e\}$ so that $g * h = h * g$. Therefore G is isomorphic to the direct product $P \times Q$. Since P has order p, P is isomorphic to $\mathbb{Z}/p\mathbb{Z}$. Similarly, Q is isomorphic to $\mathbb{Z}/q\mathbb{Z}$. Thus G is isomorphic to $\mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/q\mathbb{Z}$. By the Chinese Remainder Theorem, this is isomorphic to $\mathbb{Z}/pq\mathbb{Z}$. □

7 Finite Abelian Groups

Recall that for an indexing set I and a collection $(G_i, *_i)$ of groups indexed by $i \in I$ with identity elements $e_i \in G_i$, the **direct product** group is defined to be the product set

$$G = \prod_{i \in I} G_i = \{(g_i)_{i \in I} | g_i \in G_i\},$$

with the group operation

$$* : G \times G \to G, \quad (g_i)_{i \in I} * (h_i)_{i \in I} = (g_i *_i h_i)_{i \in I}.$$

Lemma 7.1. *The operation $*$ is a group structure on G. The group identity is $e = (e_i)_{i\in I}$. For every $g = (g_i)_{i\in I}$ in G, the inverse group element is $g^{-1} = (g_i^{-1})_{i\in I}$.*

Proof. First, $g * e = (g_i *_i e_i)_{i\in I} = (g_i)_{i\in I} = g$, and $e * g = (e_i *_i g_i)_{i\in I} = (g_i)_{i\in I} = g$. Thus e satisfies the identity axiom. For every element $h = (h_i)_{i\in I}$ of G, and for every element $k = (k_i)_{i\in I}$,

$$(g * h) * k = (g_i *_i h_i)_i * (k_i)_i = ((g_i *_i h_i) *_i k_i)_{i\in I} = (g_i *_i (h_i *_i k_i))_i = (g_i)_i * (h_i *_i k_i)_i = g * (h * k).$$

Thus the group operation satisfies the associative axiom. Finally, for every g, $g*g^{-1} = (g_i)_i*(g_i^{-1})_i = (g_i *_i g_i^{-1})_i = (e_i)_i = e$, and similarly $g^{-1} * g = e$. Thus the inverse axioms is also valid. Therefore $(G, *)$ is a group. □

For every $i \in I$, denote the projection from G to G_i by pr_i,

$$\mathrm{pr}_i : G \to G_i, \quad (g_j)_{j\in I} \mapsto g_i.$$

Proposition 7.2. *Every set function pr_i is a group homomorphism. For every group $(H, \bullet)$, and for every collection $(f_i)_{i\in I}$ of group homomorphisms $f_i : (H, \bullet) \to (G_i, *_i)$, there exists a unique group homomorphism $f : (H, \bullet) \to (G, *)$ such that for every $i \in I$, $pr_i \circ f$ equals f_i.*

Proof. By definition of $*$, for $g = (g_j)_{j\in I}$ and for $k = (k_j)_{j\in I}$, $g*k = (g_j *_j k_j)_{j\in I}$. Thus, $\mathrm{pr}_i(g*k) = g_i *_k k_i = \mathrm{pr}_i(g) *_i \mathrm{pr}_i(k)$. Therefore pr_i is a group homomorphism.

Let $(f_i)_{i\in I}$ be a collection of group homomorphisms. Define $f : H \to G$ by $f(h) = (f_j(h))_{j\in I}$. This is the unique set function such that for every $i \in I$, $\mathrm{pr}_i \circ f$ equals f_i. Moreover,

$$f(h \bullet h') = (f_j(h \bullet h'))_{j\in I} = (f_j(h) *_j f_j(h'))_{j\in I} = (f_j(h))_{j\in I} * (f_j(h'))_{j\in I} = f(h) * f(h').$$

Thus f is a group homomorphism. □

Let $(H, \bullet)$ be a group. Let I be a nonempty indexing set. For every $i \in I$, let $N_i \subset H$ be a normal subgroup. Denote the common intersection by N,

$$N = \cap_{i\in I} N_i.$$

Similarly, for every $i \in I$, denote by M_i the (normal) subgroup,

$$M_i = \cap j \in I, j \neq i N_j.$$

If I is a singleton set $\{i\}$, denote H by M_i. Denote the quotient group H/N_i by G_i, and denote the quotient homomorphism $H \to H/N_i$ by $q_i : H \to G_i$. By the previous lemma, there is a unique group homomorphism,

$$q : H \to G$$

such that for every $i \in I$, $\mathrm{pr}_i \circ q$ equals q_i.

Proposition 7.3. *The kernel of q equals N. In particular, q is one-to-one if and only if N equals $\{e\}$. If I is a finite set, and if M_iN_i equals H for every $i \in I$, then q is onto.*

Proof. By the definition of G, $g \in G$ equals e if and only if $\mathrm{pr}_i(g)$ equals e_i for every $i \in I$. Thus, for every $h \in H$, $q(h)$ equals e if and only if $q_i(h)$ equals e_i for every $i \in I$, i.e., if and only if $h \in N_i$ for every $i \in I$. Thus the kernel of q equals N.

Assume next that I is a finite set and that every M_iN_i equals H. Let $(h_iN_i)_{i\in I}$ be an element of G. Since M_iN_i equals H, for every $i \in I$, there exists $m_i \in M_i$ such that h_iN_i equals m_iN_i. Choose some ordering of I, say $I = \{i_1, \ldots, i_r\}$, and define $h = m_{i_1} * \ldots m_{i_r}$. For every i_j, $q_{i_j}(h)$ equals $q_{i_j}(m_{i_1}) * \cdots * q_{i_j}(m_{i_r})$. For every $i_k \neq i_j$, M_{i_k} is contained in N_{i_j}, so that $q_{i_j}(m_{i_k})$ equals e_{i_j}. Thus, $q_{i_j}(h)$ equals $q_{i_j}(m_{i_j}) = h_{i_j}N_{i_j}$ for every $i_j \in I$. Therefore $q(h)$ equals $(h_iN_i)_{i\in I}$. So q is onto. □

In particular, if N equals $\{e\}$ and if every M_iN_i equals H, then q is an isomorphism.

Now let $(A, *)$ be a finite Abelian group. If the order of A equals 1, then A equals $\{e\}$, and A is a trivial group. Thus, assume that the order of A is > 1. By the Unique Factorization of Integers, the set of primes dividing $\#A$ is a finite set, say a subset of a finite set of primes $I = \{p_1, \ldots, p_r\}$. Then $n = \#A$ satisfies $n = p_1^{e_1} \cdot \cdots \cdot p_r^{e_r}$ for integers $e_1, \ldots, e_r \geq 0$. For every $p_i \in I$, by the Sylow Theorem, there exists a p_i-Sylow subgroup $M_i \subset A$ that has order $p_i^{e_i}$. Since A is Abelian, M_i is a normal subgroup. Thus M_i is the unique p_i-Sylow subgroup of A.

For every $p_i \in I$, denote by $N_i \subset A$ the product over all $j \neq i$ of the normal subgroup M_j. Denote by G_i the quotient group A/N_i. Denote by $q_i : A \to G_i$ the quotient group homomorphism. Denote by G the product over all $p_i \in I$ of G_i, and denote by $q : A \to G$ the induced group homomorphism.

Since every element of M_j has order dividing $p_j^{e_j}$ by Lagrange's Theorem, and since the subgroups M_j commute with one another, every element of N_i has order dividing $\prod_{j\neq i} p_j^{e_j} = n/p_i^{e_i}$. In particular, no element of N_i has order equal to a positive power of p_i. Thus, by the Cauchy-Frobenius theorem, the order of N_i is prime to p_i. By Lagrange's Theorem, the order of N_i divides n. Thus the order of N_i divides $n/p_i^{e_i} = \prod_{j\neq i} p_j^{e_j}$. On the other hand, for every $j \neq i$, N_i contains M_j. So, again by Lagrange's Theorem, the order of N_i is divisible by each $p_j^{e_j}$ for $j \neq i$. Therefore the order of N_i equals $n/p_i^{e_i}$.

Since $M_i \cap N_i$ is a subgroup of M_i, it is a p_i-group. Since $n/p_i^{e_i}$ has order prime to p_i, $M_i \cap N_i$ is the trivial subgroup $\{e\}$. For every $j \neq i$, since M_j is a subgroup of N_i, also $M_i \cap M_j$ is a subgroup of $M_i \cap N_i$, so that $M_i \cap M_j$ is also the trivial subgroup. Since $[M_i : M_i \cap N_i]$ equals $\#M_i = p_i^{e_i}$, by the Second Isomorphism theorem, M_iN_i has order $[M_i : M_i \cap N_i]\#N_i = p_i^{e_i}(n/p_i^{e_i}) = n$. Thus, M_iN_i equals A.

Proposition 7.4. *For every finite Abelian group $(A, *)$ with order $n > 1$ divisible only by some of the primes in a finite subset $I = \{p_1, \ldots, p_r\}$, for the p_i-Sylow subgroups M_i and for $N_i = \prod_{j\neq i} M_j$, the restriction $q_i|_{M_i} : M_i \to (A/N_i)$ is a group isomorphism. The induced group homomorphism $q : A \to \prod_{p_i\in I} M_i$ is an isomorphism.*

Proof. The only thing that remains to prove is that $q_i|_{M_i} : M_i \to A/N_i$ is a group isomorphism. By the computation above, the kernel $M_i \cap N_i$ is a trivial group, so q_i is one-to-one. Also, by Lagrange's Theorem, $\#(A/N_i) = \#A/\#N_i = p_i^{e_i}$, and this equals $\#M_i$. Therefore, by the Pigeonhole Principle, also $q_i|_{M_i}$ is onto. □

By the proposition, every finite Abelian group is canonically isomorphic to the direct product of all of its nontrivial Sylow subgroups. Thus, to classify all finite Abelian groups, it suffices to classify the finite Abelian p-groups M. If M is nontrivial, then by the Cauchy-Frobenius Theorem there exists an element $g \in M$ of order p. A **maximal order element** of M is an element g whose order is maximal among the orders of all elements of M. Since M is a finite set, it has a maximal order element g. By the Cauchy-Frobenius Theorem, the p-power order p^n of g satisfies $n \geq 1$. Let C denote the cyclic subgroup $\langle g \rangle$. Denote the quotient group M/C by K, and denote by $q : M \to K$ the quotient group homomorphism.

Since C is a cyclic subgroup, every subgroup of C is of the form $\langle g^{p^s} \rangle$ for some integer $0 \leq s \leq n$. Let $B \subset M$ be another cyclic subgroup of order p^m. Then $B \cap C$ is a subgroup of C, and it is also a subgroup of B. Thus there exists a generator h of B such that h^{p^r} equals g^{p^s} for some integer $0 \leq r \leq m$. For $t = n - s$, $(h^{p^r})^{p^{t-1}} = (g^{p^s})^{p^{t-1}} = g^{p^{n-1}}$ does not equal e, yet $(h^{p^r})^{p^t} = g^{p^n}$ does equal e. Therefore the order of h divides p^{r+t}, yet the order does not divide p^{r+t-1}. Therefore the order of h equals p^{r+t}. Since g is a maximal order element $r + t \leq n = s + t$. Therefore also $r \leq s$. Now consider the element $h' = h * g^{-p^{s-r}}$. Then hC equals $h'C$, and $(h')^{p^r}$ equals e.

Proposition 7.5. *Every finite Abelian p-group is isomorphic to a direct product of finite cycle p-groups.*

Proof. This is proved by induction on the order of M. If the order of M equals 1, then M is a trivial group and so M is a cyclic group. Thus, by way of induction, assume that the order of M is > 1, and assume that the result has been proved for all finite Abelian p-groups of strictly smaller order.

As above, let g be a maximal order element, and let C be the cyclic subgroup generated by g. Denote by $q : M \to K$ the quotient group homomorphism with kernel C. By Lemma 6.1, K is also a p-group. Since M is finite and Abelian, the quotient K is also finite and Abelian. Since the order of g is p^n with $n \geq 1$, by Lagrange's Theorem, $\#K = \#M/p^n$ is strictly less than $\#M$. Thus, by the induction hypothesis, K is a direct product of cyclic p-subgroups. Precisely, let $k_2, \ldots, k_r \in K$ be elements such that K equals the direct product of the cyclic subgroups $\langle k_2 \rangle, \ldots, \langle k_r \rangle$.

For every $i = 2, \ldots, r$, denote the order of k_i by p^{r_i}. Let $B_i \subset M$ be a cyclic subgroup that is generated by an element that maps under q to k_i. By the previous argument, there exists a generator h_i of B_i such that $h_i^{p^{r_i}}$ equals g^{p^s}. By the argument, $r_i \leq s \leq n$, and there exists another element $g_i \in M$ such that $q(q_i) = q(h_i)$ and such that $g_i^{p^{r_i}}$ equals e. Thus, for every $i = 2, \ldots, r$, the restriction

$$q_i : \langle g_i \rangle \to \langle k_i \rangle$$

is an onto group homomorphism between groups of order p^{r_i}. By the Pigeonhole Principle, q_i is also one-to-one. Hence q_i is an isomorphism. Thus the product $K' \subset M$ of the subgroups $\langle g_2 \rangle, \ldots, \langle g_r \rangle$ is a subgroup of M such that

$$q_K : K' \to K$$

is onto. Thus K' has order $\geq \#K$. On the other hand, by the iterated Lagrange's Theorem, the order of K' is no greater than the product of the orders of $\langle g_2 \rangle, \ldots, \langle g_r \rangle$, and that also equals $\#K$. Thus q_K is an isomorphism. So K' is isomorphic to a direct product of cyclic groups.

Finally, set $g_1 = g$. Since q_K is onto, M equals $K'C$. By Lagrange's Theorem, $\#M = \#C\#K = \#C\#K'$. By the Second Isomorphism Theorem, $K' \cap C$ has order 1. Thus, by Proposition 7.3, M is isomorphic to the direct product $C \times K'$. Altogether, M is isomorphic to the direct product of the cyclic subgroups $\langle g_1 \rangle, \ldots, \langle g_r \rangle$. Thus, by induction, every finite, Abelian p-group is isomorphic to a direct product of copies of cyclic p-groups. □

As is clear from the proof, every nontrivial finite Abelian p-group is isomorphic to a product

$$\mathbb{Z}/p^{n_1}\mathbb{Z} \times \mathbb{Z}/p^{n_2}\mathbb{Z} \times \cdots \times \mathbb{Z}/p^{n_r}\mathbb{Z},$$

for an integer $r \geq 1$ and for integers $n_1, \ldots, n_r$ with $n_1 \geq n_2 \geq \cdots \geq n_r \geq 1$.

From now on we will denote the group operation on an Abelian A group by $+$, i.e., we write the group operation additively. For consistency, the group identity element will be denoted by 0. The group inverse operation applied to g will be denoted $-g$. Also, for every element $g \in A$, instead of denoting the n-fold operation $g * \cdots * g$ by g^n, we will denote $g + \cdots + g$ by $n \cdot g$. For every integer $n \geq 1$, define the **n-torsion subgroup** of A to be

$$A[n] = \{g \in G | n \cdot g = 0\}.$$

Lemma 7.6. *The subset $A[n]$ of A is a subgroup.*

Proof. Clearly 0 is an element of $A[n]$. For elements g and h of a group G that commute with each other, for the group product k of g and h, we have seen before that the order of k divides the least common multiple of the orders of g and h. Thus, if the order of g and the order of h both divide n, then also the order of k divides n. Thus $A[n]$ is preserved by products of elements in $A[n]$. Finally, for every group element g, the order of g equals the order of the group inverse of g. Thus, if g is in $A[n]$, then also the group inverse of g is in $A[n]$. Therefore $A[n]$ is a subgroup of A. □

Please be aware, this can fail badly for a non-Abelian group. For instance, the set of elements of $S(3)$ that have order dividing 2 equals $\{(1), (12), (13), (23)\}$, and this is not stable under products of pairs of elements in the subset.

For a direct product of Abelian groups, $A = A_1 \times \cdots \times A_r$, the inclusions $A_i[n] \subset A_i$ induce a direct product decomposition,

$$A[n] = A_1[n] \times \cdots \times A_r[n].$$

For an Abelian group A, and for integers $m, n \geq 1$ with $m|n$, then $A[m]$ is a subgroup of $A[n]$. Consider the quotient group $A[n]/A[m]$. For a direct product group, the direct product decompositions above determine a direct product decomposition,

$$A[n]/A[m] = (A_1[n]/A_1[m]) \times \cdots \times (A_r[n]/A_r[m]).$$

Now specialize to the case that A_i is a cyclic group of order p^{n_i}, $m = p^{e-1}$ and $n = p^e$. Then $A_i[p^e]$ is either $p^{n_i-e}A_i$ if $n_i \geq e$, or else $A_i[p^e]$ equals all of A_i if $e \geq n_i$. Thus, $A_i[p^e]/A_i[p^{e-1}]$ is either isomorphic to $\mathbb{Z}/p\mathbb{Z}$ if $e \leq n_i$, or $A_i[p^e]/A_i[p^{e-1}]$ is a trivial group if $e > n_i$.

Finally, consider a direct product group $A = A_1 \times \cdots \times A_r$ where each A_i is an Abelian group of order p^{n_i} with $r \geq 1$ and $n_1 \geq \cdots \geq n_r \geq 1$. For every integer $e > n_1$, $A[p^e]/A[p^{e-1}]$ is a trivial group. For every integer $e \leq n_1$, denote by i the unique integer such that $n_1 \geq \cdots \geq n_i \geq e > n_{i+1} \geq \cdots \geq n_r$. Then $A_j[p^e]/A_j[p^{e-1}]$ is isomorphic to $\mathbb{Z}/p\mathbb{Z}$ for $j \geq i$, and otherwise it is isomorphic to $\{0\}$. Thus, the group $A[p^e]/A[p^{e-1}]$ is isomorphic to a direct product of i copies of $\mathbb{Z}/p\mathbb{Z}$. In particular, by Lagrange's Theorem,

$$\#A[p^e]/\#A[p^{e-1}] = p^i.$$

Given the data of the sizes $\#A[p^e]/\#A[p^{e-1}]$, the sequence $(n_1, \ldots, n_r)$ can be uniquely recovered. Thus, we have proved the following.

Proposition 7.7. *For every finite Abelian p-group, the sequence $(n_1, \ldots, n_r)$ of integers with $n_1 \geq \cdots \geq n_r \geq 1$ satisfying $A \cong (\mathbb{Z}/p^{n_1}\mathbb{Z}) \times \cdots \times (\mathbb{Z}/p^{n_r}\mathbb{Z})$ is unique, and it can be effectively reconstructed from the function $e \mapsto \#A[p^e]$.*

For various reasons, it is sometimes inconvenient to decompose a finite Abelian group into a direct product of its Sylow subgroups. As indicated in lecture, using the Chinese Remainder Theorem, the structure theory above is equivalent to the following structure theorem.

Theorem 7.8 (Structure Theorem for Finite Abelian Groups)**.** *For every nontrivial finite Abelian group A, there exists a unique integer $r \geq 1$, and there exist a unique r-tuple of integers $(m_1, \ldots m_r)$ with $m_r > 1$ and with $m_2|m_1$, $m_3|m_2, \ldots, m_r|m_{r-1}$ such that A is isomorphic to the direct product of cyclic groups $(\mathbb{Z}/m_1\mathbb{Z}) \times \cdots \times (\mathbb{Z}/m_r\mathbb{Z})$. This sequence can be effectively computed from the function $m \mapsto \#A[m]$.*

The sequence of integers $(m_1, \ldots, m_r)$ is called the sequence of **elementary divisors** of the finite Abelian group A. It is important to note that, although the sequence of elementary divisors is unique, the isomorphism of A with the direct product of cyclic groups is almost never unique. The only canonical isomorphism in the structure theorem is the isomorphism of A with the direct product of the nontrivial Sylow subgroups of A.

Corollary 7.9. *A finite Abelian group A of order n is cyclic if and only if, for every positive integer divisor m of n, $\#A[m]$ is at most m.*

Proof. For every divisor m of n, $(\mathbb{Z}/n\mathbb{Z})[m]$ equals the cyclic subgroup generated by $[n/m]_n$, and this is isomorphic to $\mathbb{Z}/m\mathbb{Z}$. Thus, $\#(\mathbb{Z}/n\mathbb{Z})[m]$ equals m. So for every cyclic group A of order n, for every divisor m of n, $\#A[m]$ equals m. Thus it remains to prove the opposite implication: if every $\#A[m]$ is at most m, then A is cyclic.

This does follow from the structure theorem, but it is also straightforward to prove directly. The proof is by induction on n. If n equals 1, then A equals $\{0\}$, and so it is cyclic generated by 0. Thus, by way of induction, assume that $n > 1$, and assume that the result has been proved for all groups of order $< n$.

Consider first the case that n is a prime, $n = p$. In this case, by the Cauchy-Frobenius Theorem, there exists an element a of A of order p. Thus, the cyclic subgroup of A generated by a equals all of A. Therefore, if n is a prime, then A is cyclic.

Consider next the case that n is composite, divisible by more than one prime, say $n = \ell \cdot m$ with both $\ell, m > 1$ and $\gcd(\ell, m) = 1$. By Proposition 7.4, the subgroup $A[\ell]$, resp. $A[m]$, is the product of p-Sylow subgroups for all prime integers p dividing ℓ, resp. p dividing m. Thus A is the product $A[\ell] \times A[m]$. For every divisor d of ℓ, since d is relatively prime to m, $A[d]$ equals $(A[\ell])[d]$. Thus, the subgroup $A[\ell]$ satisfies the same hypothesis as A, $\#(A[\ell])[d] \leq d$. Since $\ell < n$, by the induction hypothesis, $A[\ell]$ is isomorphic to a cyclic group $\mathbb{Z}/\ell\mathbb{Z}$. By the same argument, also $A[m]$ is isomorphic to $\mathbb{Z}/m\mathbb{Z}$. By the results above, A is isomorphic to $A[\ell] \times A[m]$, which in turn is isomorphic to $(\mathbb{Z}/\ell\mathbb{Z}) \times (\mathbb{Z}/m\mathbb{Z})$. Since ℓ and m are relatively prime, by the Chinese Remainder Theorem, $(\mathbb{Z}/\ell\mathbb{Z}) \times (\mathbb{Z}/m\mathbb{Z})$ is isomorphic to $\mathbb{Z}/(\ell m)\mathbb{Z}$. Therefore A is cyclic.

Finally, consider the case that n equals p^r for an integer $r \geq 2$. By the Cauchy-Frobenius Theorem, there exists an element a of A of order equal to p. Denote by N the cyclic subgroup of A generated by a. Since A is Abelian, the subgroup N is normal. Thus, there is a unique group structure on the coset space $Q = G/N$ such that the quotient map $q : G \to Q$ is a group homomorphism. By Lagrange's Theorem, $\#Q$ equals $\#G/\#N = p^r/p = p^{r-1}$.

For every integer $s \geq 0$, $q^{-1}(Q[p^s])$ is a subgroup of G that contains N. By Lagrange's Theorem, $\#q^{-1}(Q[p^s])$ equals $\#N \cdot \#Q[p^s]$. For every element $b \in q^{-1}(Q[p^s])$, $q(p^s b) = p^s q(b) = 0$, so that $p^s b$ is an element of N. Since N has order p, $p(p^s b)$ equals 0, i.e., $p^{s+1}b$ equals 0. Thus, $q^{-1}(Q[p^s])$ is contained in $A[p^{s+1}]$. By hypothesis, $\#A[p^{s+1}]$ is at most p^{s+1}. Thus, $\#q^{-1}(Q[p^s])$ is at most p^{s+1}. Combined with Lagrange's Theorem, $\#Q[p^s]$ is at most p^s. Thus, Q satisfies the same hypothesis as A, but $\#Q = p^{r-1}$ is less than $\#A = p^r$. Therefore, by the induction hypothesis, Q is cyclic of order p^{r-1}.

Let $b \in A$ be an element such that Q equals the cyclic subgroup generated by $q(b)$. In particular, $q(p^{r-1}b)$ equals 0, but $q(p^{r-2}b)$ does not equal 0. Thus, $p^{r-1}b$ is an element of N, but $p^{r-2}b$ is not in N. If $p^{r-1}b$ equals 0, then $p^{r-2}b$ is an element of order p in A, i.e., $p^{r-2}b \in A[p]$. Already N is a subgroup of $A[p]$ that has order p. Since $\#A[p]$ is at most p, N equals all of $A[p]$. Thus, there is no element $p^{r-2}b$ of A of order p that is not contained in N. Therefore, $p^{r-1}b$ is not zero. So the order of b is strictly greater than p^{r-1}, i.e., the order of b equals p^r. Therefore A is a cyclic group generated by b.

Thus, in all cases, A is a cyclic group. Therefore, by way of induction, for every finite Abelian group A, A is cyclic if and only if, for every positive integer divisor m of n, $\#A[m]$ is at most m. □

8 Finite Subgroups of Multiplicative Groups

Recall, a **commutative ring** is a datum $(R, +, 0, \cdot, 1)$ of a set R, a binary operation $+ : R \times R \to R$ called "addition", a specified element $0 \in R$, a binary operation $\cdot : R \times R \to R$ called "multiplication", and a specified element $1 \in R$ such that for every $r, s, t \in R$, all of the following hold.

(i) [Additive Associativity] $(r + s) + t$ equals $r + (s + t)$,

(ii) [Additive Identity] $r + 0 = r = 0 + r$,

(iii) [Additive Inverses] there exists an element $-r \in R$ such that $r + (-r) = 0 = (-r) + r$,

(iv) [Additive Commutativity] $r + s$ equals $s + r$,

(v) [Multiplicative Associativity] $(r \cdot s) \cdot t$ equals $r \cdot (s \cdot t)$,

(vi) [Distributivity] $r \cdot (s + t)$ equals $(r \cdot s) + (r \cdot t)$,

(vii) [Multiplicative Identity] $r \cdot 1 = r = 1 \cdot r$, and

(viii) [Multiplicative Commutativity] $(r \cdot s)$ equals $(s \cdot r)$.

In every ring $0 \cdot r = (0 + 0) \cdot r = (0 \cdot r) + (0 \cdot r)$. Thus, by additive cancellation, for every $r \in R$, $0 \cdot r$ equals 0. The main example of a commutative ring is $\mathbb{Z}$, the ring of integers with its standard addition and multiplication. Also, for every integer n, the set $\mathbb{Z}/n\mathbb{Z}$ with its standard addition and multiplication is an example of a commutative ring.

A **field** is a commutative ring $(F, +, 0, \cdot, 1)$ such that $0 \neq 1$ and such that for every $r \in F \setminus \{0\}$, there exists an element $r^{-1} \in F$ such that $r \cdot r^{-1} = 1 = r^{-1} \cdot r$. For every pair $r, s \in F \setminus \{0\}$, note that

$$(r \cdot s) \cdot (s^{-1} \cdot r^{-1}) = ((r \cdot s) \cdot s^{-1}) \cdot r^{-1} = (r \cdot (s \cdot s^{-1})) \cdot r^{-1} = (r \cdot 1) \cdot r^{-1} = r \cdot r^{-1} = 1.$$

Since $0 \cdot (s^{-1} \cdot r^{-1})$ equals 0 and $0 \neq 1$, $r \cdot s$ is not 0; moreover, $(r \cdot s)^{-1}$ equals $s^{-1} \cdot r^{-1}$. One example of a field is the ring $\mathbb{Q}$ of rational numbers. Also $\mathbb{R}$ is a field, as is $\mathbb{C}$. Finally, for every prime integer p (by definition, $p > 1$ and its only positive divisors are 1 and p), the ring $\mathbb{Z}/p\mathbb{Z}$ is a field, often denoted by $\mathbb{F}_p$.

For every field F, denote by F^* the subset $F \setminus \{0\}$. By the argument above, for every $r, s \in F^*$, also $r \cdot s$ is in F^*. Thus, multiplication restricts to a binary operation on F^*. By multiplicative

associativity, multiplicative identity, and the existence of inverses, F^* is a group with group operation given by multiplication and with group identity given by 1. By multiplicative commutativity, F^* is an Abelian group. Quite often F^* is infinite, e.g., if F equals $\mathbb{Q}$, $\mathbb{R}$, or $\mathbb{C}$. What are the finite subgroups of F^*?

Lemma 8.1. *For every integer $m \geq 1$, the subgroup $(F^*)[m]$ of F^* has order $\leq m$.*

Proof. Consider the polynomial $f_m(x) = x^m - 1$ in $F[x]$. The elements of $(F^*)[m]$ are precisely the zeroes of $f_m(x)$ in F. Each zero λ of $f_m(x)$ gives a linear polynomial $x - \lambda$ that divides $f_m(x)$. Of course $x - \lambda$ is an irreducible element of $F[x]$. By the Unique Factorization Theorem for $F[x]$, $f_m(x)$ has a unique factorization as a product of monic irreducible polynomials. Since the degree of f_m equals m, the number of monic irreducible factors is at most m. Thus, there are at most m distinct linear factors of $f_m(x)$. Therefore $(F^\times)[m]$ has at most m elements. □

Proposition 8.2. *The finite subgroups of F^* are precisely the cyclic subgroups generated by a root of unity, i.e., generated by an element of $(F^*)[m]$ for some integer $m \geq 1$.*

Proof. First of all, for every $m \geq 1$, by the previous lemma, $(F^*)[m]$ is a finite group of size $\leq m$. Thus, for every element of $(F^*)[m]$, the cyclic subgroup generated by that element is finite of size $\leq m$. Therefore, the cyclic subgroups generated by roots of unity are finite subgroups of F^*.

Next, let G be a finite subgroup of F^*. For every integer $m \geq 1$, $G[m]$ is a subgroup of $(F^*)[m]$. By the lemma, $(F^*)[m]$ has size $\leq m$. Thus, also $G[m]$ has size $\leq m$. By Corollary 7.9, G is a cyclic group of some finite order, say m. For every generator λ of G, since λ has order m, λ^m equals 1. Thus, λ is an element of $(F^*)[m]$, i.e., G is a cyclic group generated by a root of unity. □

Corollary 8.3. *For every field F that is finite as a set, the multiplicative group F^* is cyclic. In particular, for every prime integer $p > 1$, there exists an element $[a]_p \in G_p$ that is a cyclic generator of G_p of order $\phi(p) = p - 1$.*

Proof. Since F is a finite set, also the subset $F^* = F \setminus \{0\}$ is a finite set. Therefore F^* is a cyclic group. □

9 Field Extensions and Finite Fields

The following material was not covered in the book, nor in the lecture. If the semester continued for one more week, this is what we would have discussed next.

9.1 Algebras over a Field

For every field $(F, +, 0, \cdot, 1)$, a **F-algebra** is a pair of a commutative ring $(A, +, 0, \cdot, 1)$ and a ring homomorphism $u : F \to A$, i.e., $u(1)$ equals 1, and for every $f_1, f_2 \in F$, $u(f_1 + f_2)$ equals $u(f_1) + u(f_2)$ and $u(f_1 \cdot f_2)$ equals $u(f_1) \cdot u(f_2)$. If 1 equals 0 in A, then necessarily u is not injective: every element a of A equals 0 since $a = a \cdot 1 = a \cdot 0 = 0$.

Lemma 9.1. *For every F-algebra $u : (F, +, 1, \cdot, 0) \to (A, +, 1, \cdot, 0)$, if $1 \neq 0$ in A, then u is injective. Moreover, the binary operation $F \times A \to A$ by $(f, a) \mapsto u(f) \cdot a$ makes the Abelian group $(A, +, 0)$ into an F-vector space.*

Proof. For every $f \in F \setminus \{0\}$, there exists $f' \in F$ with $f \cdot f' = 1$. Thus, $u(f) \cdot u(f')$ equals $u(f \cdot f') = u(1) = 1$. Since $0 \cdot u(f')$ equals 0, and since $0 \neq 1$, $u(f)$ does not equal 0. Thus, u is injective.

Since $(A, +)$ is an Abelian group under addition, the additional axioms for an F-vector space are the identity axiom, scalar associativity, and distributivity, i.e., for every $f_1, f_2 \in F$ and for every $a_1, a_2 \in A$,

$$1{*}a_1 = a_1, \;\; f_1{*}(f_2{*}a_1) = (f_1{\cdot}f_2){*}a_1, \;\; (f_1{+}f_2){*}a_1 = (f_1{*}a_1){+}(f_2{*}a_1), \;\; f_1{*}(a_1{+}a_2) = (f_1{*}a_1){+}(f_1{*}a_2).$$

Each of these follows quickly from the corresponding axioms for the ring $(A, +, 0, \cdot, 1)$ and the fact that u is a ring homomorphism. □

For F-algebras, $(A, u : F \to A)$ and $(B, w : F \to B)$, a **F-algebra homomorphism** is a ring homomorphism $v : A \to B$ such that $v \circ u$ equals w. The identity ring homomorphism $\mathrm{Id}_A : A \to A$ is an F-algebra homomorphism. Moreover, every composition of F-algebra homomorphisms is again an F-algebra homomorphism.

A particularly important F-algebra is the algebra of polynomials in a specified variable x with coefficients in F, i.e., $F[x]$. A bit more generally, for every F-algebra A, $A[x]$ is defined to be the set of all formal linear combinations,

$$p(x) = a_0 + a_1 x + a_2 x^2 + \cdots + a_{d-1} x^{d-1} + a_d x^d,$$

where $d \geq 0$ is an arbitrary integer and where $a_0, \ldots, a_d$ are arbitrary elements of A. Equivalently, the polynomial $p(x)$ is an element $(a_0, a_1, \ldots, a_d) \in A^d$. For every integer $e \geq d$, $(a_0, \ldots, a_d) \in A^d$ is identified with $(a_0, \ldots, a_d, 0, \ldots, 0) \in A^e$, and in this way $A[x]$ is the union $\cup_{d \geq 0} A^d$. Addition is defined coefficient-by-coefficient, just as the usual addition in A^d. Multiplication is defined according to the usual rule,

$$(a_0 + \cdots + a_\ell x^\ell + \cdots + a_d x^d) \cdot (b_0 + \cdots + b_m x^m + \cdots + b_e x^e) = c_0 + \cdots + c_n x^n + \cdots + c_{d+e} x^{d+e},$$

$$c_n = \sum_{\ell + m = n} a_\ell b_m.$$

As checked in lecture, this does define a product making $A[x]$ into a commutative ring with 1. Moreover, there is a ring homomorphism $u_x : A \to A[x]$ sending every $a \in A$ to the constant polynomial with $a_0 = a$. In particular, $u_x \circ u : F \to A[x]$ makes $A[x]$ into an F-algebra. For every F-algebra homomorphism $v : A \to B$ as above, there is an associated F-algebra homomorphism $A[x] \to B[x]$, also denote by v,

$$v : A[x] \to B[x], \;\; a_d x^d + \cdots + a_\ell x^\ell + \cdots + a_0 x^0 \mapsto v(a_d) x^d + \cdots + v(a_\ell) x^\ell + \cdots + v(a_0) x^0.$$

The composition $v \circ u_x$ equals the composition $w_x \circ v$. Moreover, $v(x)$ equals x.

For every F-algebra B, and for every F-algebra homomorphism $\widetilde{v} : A[x] \to B$, the composition $\widetilde{v} \circ u_x : A \to B$ is an F-algebra homomorphism. Moreover, the element $\widetilde{v}(x)$ is a specified element of B.

Lemma 9.2. *For every pair of F-algebras, $u : (F, +, 0, \cdot, 1) \to (A, +, 0, \cdot, 1)$ and $w : (F, +, 0, \cdot, 1) \to (B, +, 0, \cdot, 1)$, for every F-algebra homomorphism $v : A \to B$, every element $b \in B$, there is a unique F-algebra homomorphism $\widetilde{v}_b : F[x] \to B$ such that $\widetilde{v}_b \circ u_x$ equals v and such that $\widetilde{v}_b(x)$ equals b. The image of $\widetilde{v}_b$, denoted $v(A)[b] \subset B$, is the smallest F-subalgebra of B that contains the image of v and contains b.*

Proof. For any F-algebra homomorphism whose composition with u_x equals v and sending x to b, because every polynomial is obtained by iterated addition and multiplication from x and from constants, the F-algebra homomorphism must be,

$$\widetilde{v}_b : A[x] \to B,\ a_d x^d + \cdots + a_\ell x^\ell + \cdots + a_0 x^0 \mapsto v(a_d)b^d + \cdots + v(a_\ell)b^\ell + \cdots + v(a_0)b^0.$$

It is straightforward to check that this is a ring homomorphism; the main step is checking that it is compatible with multiplication. Thus, $\widetilde{v}_b$ is the unique F-algebra homomorphism such that $\widetilde{v}_b \circ u_x$ equals v and such that $\widetilde{v}_b(x)$ equals b.

Since $\widetilde{v}_b$ is a homomorphism of F-algebras, the image $v(A)[b]$ is an F-subalgebra of B. The image contains $v(A) = \widetilde{v}_b(u_x(A))$ and it contains $b = \widetilde{v}_b(x)$. Every F-subalgebra of B that contains $v(A)$ and contains b contains every polynomial in b with coefficients in $v(A)$, since the F-subalgebra is stable for addition and multiplication. Thus, $v(A)[b]$ is the smallest F-subalgebra of B that contains $v(A)$ and b. □

The kernel of $\widetilde{v}_b$ is quite important.

Lemma 9.3. *For every F-algebra homomorphism $\widetilde{v}_b : A[x] \to B$, the kernel of $\widetilde{v}_b$ is an ideal in $A[x]$, i.e., an additive subgroup such that for every $p(x) \in Ker(\widetilde{v}_b)$ and for every $q(x) \in A[x]$, the product $q(x)p(x)$ is also in $Ker(\widetilde{v}_b)$. If $1 \neq 0$ in B, then the kernel is a proper ideal, i.e., $Ker(\widetilde{v}_b) \neq A[x]$. If, moreover, B, or even just the subspace $v(A)[b] \subset B$, has finite dimension as an F-vector space, then $Ker(\widetilde{v}_b)$ contains a nonconstant polynomial.*

Proof. As for every ring homomorphism, if $\widetilde{v}_b(p(x))$ equals 0, then for every $q(x) \in A[x]$, also

$$\widetilde{v}_b(q(x) \cdot p(x)) = \widetilde{v}_b(q(x)) \cdot \widetilde{v}_b(p(x)) = \widetilde{v}_b(q(x)) \cdot 0 = 0.$$

Thus, for every $p(x) \in \mathrm{Ker}(\widetilde{v}_b)$, for every $q(x) \in A[x]$, also $q(x) \cdot p(x) \in \mathrm{Ker}(\widetilde{v}_b)$.

If $1 \neq 0$ in B, then $\widetilde{v}_b(1) = 1 \neq 0$. Thus, 1 is not in $\mathrm{Ker}(\widetilde{v}_b)$.

Similarly, if B has finite dimension as an F-vector space, the infinite sequence of elements $(\widetilde{v}_b(x^0), \widetilde{v}_b(x^1), \ldots, \widetilde{v}_b($ cannot be F-linearly independent. Since $1 \neq 0$ in B, $\widetilde{v}_b(x^0) = 1 \neq 0$, so the least linear relation

involves terms of degree > 0. Thus, there exists a least integer $d > 0$, and there exist elements $c_0, \ldots, c_{d-1} \in F$ such that

$$-\widetilde{v}_b(x^d) = c_{d-1}\widetilde{v}_b(x^{d-1}) + \cdots + c_0\widetilde{v}_b(x^0).$$

Therefore the polynomial $p(x) = u(x^d + c_{d-1}x^{d-1} + \cdots + c_0) \in A[x]$ is a nonconstant element in $\mathrm{Ker}(\widetilde{v}_b)$. □

The most important special case is when the F-algebra A equals F. In this case, for every F-algebra $w : F \to B$, for every element $b \in B$, $\widetilde{v}_b : F[x] \to B$ is a homomorphism of F-algebras.

Lemma 9.4. *For every F-algebra $w : F \to B$, for every $b \in B$, the image of $\widetilde{v}_b : F[x] \to B$ is the smallest F-subalgebra of B that contains b. The kernel of $\widetilde{v}_b$ is nonzero if and only if the F-subspace $v(F)[b] \subset B$ has finite F-dimension d. In this case, $g_b(x) \in F[x]$ is a monic polynomial of degree d.*

Proof. Since $\widetilde{v}_b$ is an F-algebra homomorphism, the image is an F-subalgebra of B. The image contains b. Every F-subalgebra of B the contains b contains every polynomial in b, since the F-subalgebra is preserved by addition and multiplication. Therefore the image of $\widetilde{v}_b$ is the smallest F-subalgebra of B that contains b.

As proved above, the kernel of $\widetilde{v}_b$ is nonzero if and only if there exists an integer d such that $\widetilde{v}_b(x^d)$ is contained in the F-span of $\widetilde{v}_b(x^0), \ldots, \widetilde{v}_b(x^{d-1})$, i.e., there exists $c_0, \ldots, c_{d-1} \in F$ such that

$$-\widetilde{v}_b(x^d) = c_0\widetilde{v}_b(x^0) + \cdots + c_{d-1}\widetilde{v}_b(x^{d-1}).$$

The collection of such linear relations is in bijection with the monic polynomials

$$g(x) = x^d + c_{d-1}x^{d-1} + \cdots + c_1x^1 + c_0x^0$$

that are contained in the kernel of $\widetilde{v}_b$. In particular, choosing d to be minimal, the polynomial $g_d(x)$ is a generator of $\mathrm{Ker}(\widetilde{v}_d)$. In this case, $\mathcal{B} = (\widetilde{v}_b(x^0), \ldots, \widetilde{v}_b(x^{d-1}))$ is an ordered F-basis for $v(F)[b]$ since every $f(x) \in F[x]$ is congruent modulo $g_b(x)$ to a unique F-linear combination of $(x^0, \ldots, x^{d-1})$. □

In the case above, if $\mathrm{Ker}(\widetilde{v}_b)$ is nonzero, then the element b is called **integral** over F, or sometimes **algebraic** over F. In this case, the polynomial $g_b(x)$ is the **minimal polynomial** of b over F, sometimes denoted $m_{B/F,b}(x)$.

9.2 Field Extensions

Our main interest in F-algebras is the case when both the domain and the target are fields. For a field $(F, +, 0, \cdot, 1)$, a **field extension** is an F-algebra, $u : (F, +, 0, \cdot, 1) \to (E, +, 0, \cdot, 1)$, such that $(E, +, 0, \cdot, 1)$ is a field. In this case, since $1 \neq 0$ in F, u is injective. Psychologically, one often identifies F with the image of this injection, and thus one thinks of E as a "bigger field".

The field extension is a **finite** field extension if E has finite dimension as an F-vector space, otherwise the field extension is **infinite**. For a finite field extension, the finite dimension of E as an F-vector space is called the **degree** of the field extension, and it is denoted $\dim_F(E)$, $\deg(u)$, $\deg(E/F)$ or $[E:F]$.

PLEASE NOTE: Although we sometimes use the same notation $[E:F]$ for the degree of a field extension as we use for the index of a subgroup, it is almost never the case that the degree of the field extension equals the index of the additive subgroup $u(F)$ in E.

For field extensions $u : F \to E$ and $w : F \to K$, a **homomorphism of F-extensions** is an F-algebra homomorphism $v : E \to K$, i.e., a ring homomorphism such that $v \circ u$ equals w. For every F-extension $u : F \to E$, the identity map $\mathrm{Id}_E : E \to E$ is a homomorphism of F-extensions. Every composition of F-extensions is an F-extension.

Lemma 9.5. *For field extensions $u : (E, +, 0, \cdot, 1) \to (F, +, 0, \cdot, 1)$ and $v : (F, +, 0, \cdot, 1) \to (K, +, 0, \cdot, 1)$, the composition $v \circ u : (E, +, 0, \cdot, 1) \to (K, +, 0, \cdot, 1)$ is also a field extension. If either u is infinite or v is infinite, then $v \circ u$ is infinite. If both u and v are finite, then $v \circ u$ is finite and $\dim_E(K)$ equals $\dim_E(F) \cdot \dim_F(K)$.*

Proof. In fact, every composition of ring homomorphisms is a ring homomorphsm,

$$(v \circ u)(1) = v(u(1)) = v(1) = 1,$$

$$(v \circ u)(e_1 + e_2) = v(u(e_1 + e_2)) = v(u(e_1) + u(e_2)) = v(u(e_1)) + v(u(e_2)),$$
$$(v \circ u)(e_1 \cdot e_2) = v(u(e_1 \cdot e_2)) = v(u(e_1) \cdot u(e_2)) = v(u(e_1)) \cdot v(u(e_2)).$$

Thus $v \circ u$ is a field extension.

The ring homomorphism v is a homomorphism of E-vector spaces that is injective. Thus, if F has infinite dimension as an E-vector space, then so does K. Similarly, every generating set of K as an E-vector space is also a generating set of K as an F-vector space. Thus, if K has finite E-dimension, then it has finite F-dimension. Contrapositively, if K has infinite F-dimension, then K has infinite E-dimension. Thus, if either u or v is an infinite extension, then so is $v \circ u$.

Finally, assume that both v and u are finite field extensions. Then, as an F-vector space, K is isomorphic to F^e, where $e = \dim_F(K)$. Similarly, as an E-vector space, F is isomorphic to E^d, where $d = \dim_E(F)$. Thus, altogether, as an E-vector space, K is isomorphic to $(E^d)^f = E^{df}$. Therefore the degree of $v \circ u$ equals the product of the degree of v and the degree of u. □

For every homomorphism of F-extensions, $v : E \to K$, and for every $b \in K$, there is a unique F-algebra homomorphism $\widetilde{v}_b : E[x] \to K$ such that $\widetilde{v}_b \circ u_x$ equals v and such that $\widetilde{v}_b(x)$ equals b. The element b is **transcendental** over E if $\mathrm{Ker}(\widetilde{v}_b)$ is $\{0\}$, otherwise it is **algebraic** over E. The F-extension v is **algebraic** if every $b \in K$ is algebraic over E, otherwise it is **transcendental** (but not necessarily "purely transcendental").

Lemma 9.6. *The element b is algebraic over E if and only if the E-subalgebra $v(E)[b] \subset K$ has finite E-dimension d. In that case, $v(E)[b]$ is a subfield of K, and the monic generator $g_b(x)$ of $Ker(\widetilde{v}_b)$ is irreducible of degree d.*

Proof. As proved in lecture, in the Euclidean domain $E[x]$, every ideal is either $\{0\}$ or it is of the form $g(x)E[x]$ for some unique monic polynomial $g(x)$; the monic element of least degree in the ideal. If the kernel is $\{0\}$, then $\widetilde{v}_b$ is an injective E-linear transformation from an infinite-dimensional E-vector space $E[x]$ to K. Thus, in that case, Image$(\widetilde{v}_b)$ has infinite E-dimension.

Assume that $E[b] := \text{Image}(\widetilde{v}_b)$ has finite E-dimension equal to d. As proved above, $\text{Ker}(\widetilde{v}_b) = g_b(x)E[x]$ for a unique monic polynomial $g_b(x)$ of degree d. For every $c \in E[b]$, since $E[b]$ is an E-subalgebra of K, multiplication by c maps $E[b]$ to itself,

$$L_c : E[b] \to E[b], \quad L_c(k) = c \cdot k.$$

Since K is a field, for every nonzero $c, k \in K$, also $c \cdot k$ is nonzero. Thus L_c is an injective E-linear transformation from the d-dimensional E-vector space $E[b]$ to itself. By the Rank-Nullity Theorem, also L_c is surjective. Thus, there exists $k \in E[b]$ such that $c \cdot k$ equals 1, i.e., $E[b]$ is a field.

Finally, if $g_b(x)$ equals $h(x)k(x)$, then $\widetilde{v}_b(h(x)) \cdot \widetilde{v}_b(k(x)) = \widetilde{v}_b(g_b(x)) = 0$. As proved above, in every field K, every product of two nonzero elements is nonzero. Thus, either $h(x)$ is in $\text{Ker}(\widetilde{v}_b)$ or $k(x)$ is in $\text{Ker}(\widetilde{v}_b)$. Thus, either g_b divides $h(x)$ or g_b divides $k(x)$. Therefore $g_b(x)$ is irreducible. □

For applications in number theory (MAT 311), it is important to identify which elements in a given extension are algebraic. The following result clarifies this.

Proposition 9.7. *For field extensions $u : F \to E$ and $v : E \to K$, every element $b \in K$ that is F-algebraic is also E-algebraic. If u is finite, then every element $b \in K$ that is E-algebraic is also F-algebraic. For every pair $b, c \in K$ of F-algebraic elements, both $b + c \in K$ and $b \cdot c \in K$ are F-algebraic elements. For a nonzero F-algebraic element $b \in K$, also $1/b \in K$ is F-algebraic. The subset $Alg_{K/F} \subset K$ of all F-algebraic elements of K is an F-subextension of K.*

Proof. If b is F-algebraic, then for the minimal polynomial $m_{K/F,b}(x) \in F[x]$, also $u(m_{K/F,b}(x)) \in E[x]$ is in the kernel of $\widetilde{v}_b : E[x] \to K$. Thus, b is E-algebraic. Conversely, if b is E-algebraic, then $v(E)[b] \subset K$ is a finite E-subextension of K. If also $u : F \to E$ is finite, then the composition $F \to E \to v(E)[b]$ is a finite extension. Thus, since $v(u(E))[b]$ is an E-subspace of $v(E)[b]$, also $v(u(E))[b]$ has finite E-dimension, so that b is E-algebraic.

For every pair $b, c \in K$ of F-algebraic elements, the F-subextension $L = F[b] \subset K$ is a finite F-extension. Since c is finite over F, also c is finite over L. Thus, the L-subextension $L[c] \subset K$ is also finite. As the composition of finite extensions, $F \to L \to L[c]$ is a finite extension. Thus, for every $k \in L[c]$, since $F[k]$ is an F-subalgebra of $L[c]$, and since $L[c]$ has finite F-dimension, also $F[k]$ has finite F-dimension. Thus, k is algebraic over F. Applying this to $k = b + c \in L[c]$, resp. $k = b \cdot c \in L[c]$, resp. $k = 1/b \in L[c]$, each of these is F-algebraic. Thus, the subset $\text{Alg}_{K/F} \subset K$ is a subfield of K that contains the image of F. □

9.3 Characteristic of a Field

For every field $(F, +, 0, \cdot, 1)$, consider the multiplicative identity 1 as an element of the additive group $(F, +, 0)$; in particular, consider the "additive order" of 1 as an element of this group. The **characteristic** of F, denoted $\mathrm{char}(F)$, equals 0 if 1 has infinite additive order, and otherwise the characteristic of F equals the additive order of 1. When the characteristic of F is not 0, F is said to have **positive characteristic**.

Lemma 9.8. *If the characteristic of the field F is zero, then F is an infinite set. If the characteristic of the field F is positive, then it equals a prime integer.*

Proof. If a group contains an element of infinite order, then the group is infinite. Thus, if the additive order of 1 is infinite, then the field is an infinite set.

Since F is a field, 1 is not equal to 0. Thus, the characteristic does not equal 1. Assume that the characteristic is a positive integer > 1. In that case, for every integer r with $0 < r < \mathrm{char}(F)$, $r \cdot 1$ is nonzero. For every factorization of $\mathrm{char}(F)$ as a product of positive integers, say $\mathrm{char}(F) = m \cdot n$, using multiplicative associativity, distributitivity, and multiplicative identity,

$$(m \cdot 1) \cdot (n \cdot 1) = \mathrm{char}(F) \cdot (1 \cdot 1) = \mathrm{char}(F) \cdot 1 = 0.$$

In a field, the product of nonzero elements is nonzero. Thus, either $m \cdot 1$ equals 0 or $n \cdot 1$ equals 0. Thus, $\mathrm{char}(F)$ divides m or n. Therefore the characteristic is a prime integer. □

For this reason, fields of positive characteristic are also sometimes called fields of **prime characteristic**. Assume that F has positive characteristic equal to p. Associated to the element 1 in $(F, +, 0)$ of order p, there is a unique group homomorphism,

$$E_1 : (\mathbb{Z}/p\mathbb{Z}, +) \to (F, +),$$

such that $E_1([1]_p)$ equals 1.

Lemma 9.9. *The group homomorphism E_1 is a ring homomorphism, i.e., $E_1([1]_p)$ equals 1, and for every $[n]_p, [a]_p \in \mathbb{Z}/p\mathbb{Z}$, also $E_1([n]_p \cdot [a]_p)$ equals $E_1([n]_p) \cdot E_1([a]_p)$. Thus, E_1 is a field extension.*

Proof. By definition, $E_1([1]_p)$ equals 1. Group homomorphisms preserve exponentiation of group elements. Since these groups are written additively, this means that for every $[a]_p \in \mathbb{Z}/p\mathbb{Z}$, for every integer $n \in \mathbb{Z}$, $E_1(n \cdot [a]_p)$ equals $n \cdot E_1([a]_p)$. In particular,

$$E_1([n]_p) = E_1(n \cdot [1]_p) = n \cdot E_1([1]_p) = n \cdot 1.$$

But then also,

$$E_1([n]_p \cdot [a]_p) = E_1(n \cdot [a]_p) = n \cdot E_1([a]_p) = E_1([n]_p) \cdot E_1([a]_p).$$

Thus, E_1 is a ring homomorpism. □

By definition of p as the characteristic of F, E_1 is injective. The image of E_1 is a subfield of F that is called the **prime subfield**, but it is always identified with $\mathbb{Z}/p\mathbb{Z}$. When thinking of $\mathbb{Z}/p\mathbb{Z}$ as a field, it is more common to refer to this as $\mathbb{F}_p$ for "field with p elements", or sometimes $GF(p)$ for "Galois field with p elements".

For every field F of characteristic p, the **Frobenius homomorphism** is the set map

$$\mathrm{Frob}_{F,p} : F \to F, \quad f \mapsto f^p.$$

Lemma 9.10. *For every field F of characteristic p, the Frobenius homomorphism is a homomorphism of $\mathbb{F}_p$-extensions.*

Proof. By construction $\mathrm{Frob}_{F,p}(1)$ equals $1^p = 1$. Also, for every $f_1, f_2 \in F$, $\mathrm{Frob}_{F,p}(f_1 \cdot f_2)$ equals $(f_1 \cdot f_2)^p = f_1^p \cdot f_2^p$, and this in turn equals $\mathrm{Frob}_{F,p}(f_1) \cdot \mathrm{Frob}_{F,p}(f_2)$. The key issue is compatibility with addition. By the Binomial Theorem, for every $f_1, f_2 \in F$,

$$\mathrm{Frob}_{F,p}(f_1 + f_2) = (f_1 + f_2)^p = f_1^p + \left(\sum_{\ell=1}^{p-1} \binom{p}{\ell} f_1^{p-\ell} f_2^{\ell}\right) + f_2^p.$$

For every integer $\ell = 1, \ldots, p-1$, $(p-\ell)! \cdot \ell! \cdot \binom{p}{\ell}$ equals $p!$. Since p divides $p!$, yet p divides neither $\ell!$ nor $(p-\ell)!$, since p is a prime, p divides $\binom{p}{\ell}$. Thus, since F has characteristic p, $\binom{p}{\ell} f_1^{p-\ell} f_2^{\ell}$ equals 0. Therefore,

$$\mathrm{Frob}_{F,p}(f_1 + f_2) = f_1^p + f_2^p = \mathrm{Frob}_{F,p}(f_1) + \mathrm{Frob}_{F,p}(f_2).$$

Thus, $\mathrm{Frob}_{F,p}$ is a ring homomorphism.

For every integer $n \geq 1$, the claim is that $n^p \equiv n \pmod{p}$. This is proved by induction on n. For $n = 1$, since 1^p equals 1, the result holds. Thus, by way of induction, assume the result is proved for n and consider the case $n+1$. By the previous paragraph,

$$[n+1]_p^p = ([n]_p + [1]_p)^p = [n]_p^p + [1]_p^p = [n]_p + [1]_p = [n+1]_p.$$

Therefore, also $(n+1)^p \equiv n+1 \pmod{p}$. Thus, by induction on n, for every $n \geq 1$, $\mathrm{Frob}_{F,p}([n]_p) = [n]_p$. Therefore $\mathrm{Frob}_{F,p} : F \to F$ is a homomorphism of $\mathbb{F}_p$-extensions. □

Since composition of homomorphisms of $\mathbb{F}_p$-extensions is again a homomorphism of $\mathbb{F}_p$-extensions, for every integer $e \geq 1$, for $q = p^e$, the e-fold composition,

$$\mathrm{Frob}_{F,p^e} = \mathrm{Frob}_{F,p} \circ \mathrm{Frob}_{F,p^{e-1}} = \cdots = \mathrm{Frob}_{F,p} \circ \ldots \mathrm{Frob}_{F,p}$$

is also a homomorphism of $\mathbb{F}_p$-extensions,

$$\mathrm{Frob}_{F,q} : F \to F, \ f \mapsto f^q.$$

9.4 Finite Fields

A special case of a finite field extension is when the target field is finite as a set.

Lemma 9.11. *For every field extension $u : (E, +, 0, \cdot, 1) \to (F, +, 0, \cdot, 1)$, if F is finite, then so is E. In that case, the degree d of u is also finite. Finally, $\#F$ equals $(\#E)^d$. In particular, for $p = char(F)$, for e equal to the degree of $E_1 : \mathbb{F}_p \to F$, $\#F$ equals $q = p^e$. In that case, $Frob_{F,q} : F \to F$ equals the identity map, and each of q distinct elements of F is a root of the polynomial $g_q(x) = x^q - x \in \mathbb{F}_p[x]$.*

Proof. Since u is injective, E is isomorphic to the subfield $u(E)$ of F. Since F is finite, the subset $u(E)$ is also finite. Therefore E is finite.

Always, F is generated as an E-vector space by the subset of F equal to all of F. Since F is finite, this is a finite generating set. Since F has a finite generating set, F has finite dimension d as an E-vector space. As an E-vector space, F is isomorphic to E^d. Counting elements, $\#F$ equals $(\#E)^d$.

Finally, applying this to the prime field extension $E_1 : \mathbb{F}_p \to F$, $\#F$ equals p^e for $e = \deg(E_1)$.

By Corollary 8.3, F^* is a cyclic group of order $q - 1$. Thus, by Lagrange's Theorem, for every $\alpha \in F^*$, α^{q-1} equals 1. Thus, $\alpha^q = \alpha^{q-1} \cdot \alpha$ equals α. Similarly, for $\alpha = 0$, $\alpha^q = 0^q$ equals 0. Therefore, for every $\alpha \in F$, $\alpha^q = \alpha$. So $\mathrm{Frob}_{F,q}$ is Id_F, and every element of F is a root of $g_q(x) = x^q - x \in \mathbb{F}_p[x]$. □

9.5 Splitting Fields

The main source of finite field extensions arises from irreducible polynomials. Let F be a field. For every ideal $g(t)F[t]$ in $F[t]$, as discussed in lecture, the F-algebra $F[t]/g(t)F[t]$ is an F-vector space. If $g(t)$ equals 0, this is the infinite-dimensional F-vector space $F[t]$ with infinite basis $\{t^0, t^1, \dots, t^n, \dots\}$. If $g(t)$ is nonzero of degree d, then as discussed in lecture, $F[t]/g(t)F[t]$ is a finite-dimensional F-vector space with ordered basis $\mathcal{B} = ([t^0]_{g(t)}, [t^1]_{g(t)}, \dots, [t^{d-1}]_{g(t)})$. Thus, the dimension equals d, the degree of $g(t)$. Finally, as proved in lecture, if $g(t)$ is a nonzero, noninvertible element of $F[t]$ that is *irreducible*, then $K_g := F[t]/g(t)F[t]$ is a field. Then the F-algebra homomorphism of constant polynomials,

$$u_g : F \to K_g, \quad u_g(c) = [c]_g,$$

is a finite field extension. Moreover, there is a distinguished element $\gamma_g = [t]_g \in K_g$ that is a root of $u_g(g(x)) \in K_g[x]$; the proof in lecture used the Cayley-Hamilton Theorem, but there are more direct proofs.

Lemma 9.12. *For every field extension $w : (F, +, 0, \cdot, 1) \to (L, +, 0, \cdot, 1)$, for every element $\alpha \in L$ that is a root of $w(g(x)) \in L[x]$, there exists a unique homomorphism of F-field extensions, $\widehat{w}_\alpha : (K_g, +, 0, \cdot, 1) \to (L, +, 0, \cdot, 1)$ such that $\widehat{w}_\alpha(\gamma_g)$ equals α. Conversely, for every homomorphism $\widehat{w}$ of F-field extensions, $\alpha := \widehat{w}(\gamma_g)$ is an element such that $g(\alpha)$ equals 0, and thus $\widehat{w}$ equals $\widehat{w}_\alpha$.*

Proof. For every F-algebra, $w : F \to B$, composition with the projection homomorphism $F[t] \to F[t]/g(t)F[t]$ defines a one-to-one correspondence between F-algebra homomorphisms $\widehat{w} : F[t]/g(t)F[t] \to B$ and F-algebra homomorphisms $\widetilde{w} : F[t] \to B$ such that $\widetilde{w}(g(t))$ equals 0. By the universal property of $F[t]$, for every element $b \in B$, there is a unique F-algebra homomorphism $\widetilde{w}_b : F[t] \to B$ with $\widetilde{w}_b(t) = b$. By construction $\widetilde{w}_b(g(t))$ equals 0 if and only if b is a root of $w(g(t)) \in B[t]$. Thus, there is a one-to-one correspondence between F-algebra homomorphisms $\widehat{w} : F[t]/g(t)F[t] \to B$ and roots $b \in B$ of $w(g(t)) \in B[t]$. □

For every field extension $v : F \to M$, if $v(g(x)) \in M[x]$ factors as a product of linear factors, then the field extension M/F **splits** the polynomial $g(x)$. Because there exists a distinguished root γ_g of $g(x)$ in K_g, the polynomial $v_g(g(x)) \in K_g[x]$ factors as $(x - \gamma_g)h(x)$ for some $h(x)$ of smaller degree. This leads to an induction proof of the following.

Lemma 9.13. *For every field $(F, +, 0, \cdot, 1)$, for every nonzero, noninvertible polynomial $g(x) \in F[x]$, there exists a finite field extension $w : (F, +, 0, \cdot, 1) \to (M, +, 0, \cdot, 1)$ that splits $g(x)$. For every factored polynomial $g(x) = h(x)k(x)$ in $F[x]$, w splits $g(x)$ if and only if w splits both $h(x)$ and $k(x)$.*

Proof. First assume that $g(x)$ factors as $h(x)k(x)$ in $F[x]$. Let $w : F \to M$ is a field extension. For the irreducible factorizations of $w(h(x)) \in M[x]$ and $w(k(x)) \in M[x]$, the concatenation of these factorizations is an irreducible factorization of $w(g(x))$ in $M[x]$. By the Unique Factorization Theorem, it follows that $w(g(x))$ factors as a product of linear factors if and only if both $w(h(x))$ and $w(k(x))$ factor as products of linear factors. Thus, w splits $g(x)$ if and only if w splits both $h(x)$ and $k(x)$.

By Unique Factorization of Polynomials, $g(x)$ factors as a product of irreducible polynomials. The result is proved by induction on the greatest degree of an irreducible factor, and by induction on the number of irreducible factors of that degree. When the greatest degree of an irreducible factor equals 1, then $g(x)$ already factors as a product of linear polynomials. Thus, by way of induction, assume that the greatest degree $d = d(F, g)$ of an irreducible factor of $g(x)$ in $F[x]$ is > 1, assume that the number $m = m(F, g, d)$ of irreducible factors of degree d is ≥ 0, and assume that the result has already been proved for all pairs (E, h) of a field E and a polynomial h if either $d(E, h) < d(F, g)$ or if $d(E, h) = d(F, g)$ yet $m(E, h, d) < m(F, g, d)$.

Denote an irreducible factorization of $g(x)$ in $F[x]$ by

$$g(x) = g_1(x) \cdots g_m(x) \cdots g_r(x),$$

and assume that $g_1(x), \ldots, g_m(x)$ are the irreducible factors of degree d. For the F-algebra $K = K_{g_1}$, and for the distinguished root β of g_1 in K_{g_1}, the F-algebra homomorphism is a finite field extension, $u : F \to K$, and $\beta \in K$ is an element such that $g_1(\beta)$ equals 0.

In $K[x]$, a factorization of $u(g(x))$ is

$$u(g(x)) = u(g_1(x)) \cdots u(g_m(x)) \cdot u(g_r(x)),$$

and $\deg(u(g_i(x))) = \deg(g_i(x))$. Thus, factoring each $u(g_i(x))$ into irreducibles, it follows that the degree of the maximal irreducible dividing $u(g_i(x))$ is at most $\deg(g_i(x))$, and these are equal if and only if $u(g_i(x))$ is irreducible in $K[x]$. In particular, $d(K, u(g)) \leq d$ and $m(K, u(g), d) \leq m$, with equality if and only if every $u(g_1), \ldots, u(g_m)$ is irreducible in $K[x]$. However, by construction, $u(g_1(x))$ does factor in $K[x]$ as a product of $(x - \beta)$ and a polynomial of degree $d - 1$. Thus, the induction hypothesis holds for $(K, u(g))$. Therefore, by the induction assumption, there exists a finite field extension $v : K \to M$ such that $v(g(x))$ factors in $M[x]$ into a product of linear factors. Therefore, defining w to be $v \circ u$, then $w : F \to M$ is a finite field extension such that $w(g(x))$ factors into a product of linear factors. □

A **splitting field** of $g(x)$ is a finite field extension $v : F \to L$ that splits $g(x)$, and such that no proper F-subextension of L/F splits $g(x)$.

Proposition 9.14. *For every finite field extension $v : F \to M$ that splits g, there exists a subfield $L \subset M$ that contains $v(F)$ such that the subextension $v : F \to L$ is a splitting field of $g(x)$. In particular, for every field F, for every nonzero, noninvertible $g(x)$ in $F[x]$, there exists a splitting field of $g(x)$.*

Proof. The dimension of M as an F-vector space is a finite integer d. Thus, for every F-subspace of M, the dimension of the subspace is also finite, and it is less than or equal to d with equality if and only if the subspace equals all of M. Thus, among all subfields L of M that contain $v(F)$ and such that L splits $g(x)$ (and $L = M$ is one such subfield), there exists one that has minimal dimension $d(F, g, v)$ as an F-vector space. For such a minimal subfield L, for every proper subfield K of L that contains $v(F)$, since the dimension of K as an F-vector space is strictly smaller than $d(F, g, v)$, then K does not split $g(x)$. Thus, $v : F \to L$ is a splitting field of $g(x)$.

By the previous lemma, there exists a finite field extension $F \to M$ that splits $g(x)$. Combined with the previous paragraph, there exists a splitting field of $g(x)$. □

If every irreducible factor of $g(x)$ is linear, then the identity $\mathrm{Id}_F : F \to F$ is already a splitting field of $g(x)$, and thus every splitting field $v : F \to L$ of $g(x)$ is also an isomorphism; since $v(F)$ is an F-subextension of L that splits $g(x)$, L must equals $v(F)$.

Theorem 9.15. *For every field extension $w : F \to K$ that splits $g(x)$, for every irreducible factor $h(x)$ of $g(x)$ in $F[x]$, for every root α of $h(x)$ in L, for every root β of $h(x)$ in K, there exists a homomorphism of F-extensions $\tau : L \to K$ such that β equals $\tau(\alpha)$.*

Proof. First consider the case that $h(x)$ is linear, say $c(x - \gamma)$ for some unique $c \in F^*$ and $\gamma \in F$. Then the unique root of $w(h(x))$ in K is $\beta = w(\gamma)$. Similarly, α equals $u(\gamma)$. Thus, for any homomorphism $\tau : L \to K$ of F-extensions, $\tau(\alpha) = \tau(u(\gamma)) = v(\gamma) = \beta$. Thus, when $h(x)$ has degree 1, every homomorphism τ of F-extensions satisfies $\tau(\alpha) = \beta$.

The result is proved by induction on the degree of L/F. When the degree equals 1, then $u : F \to L$ is an isomorphism, and so $g(x)$ already splits as a product of linear polynomials in F. Thus, by

the previous paragraph, the F-extension $\tau = w \circ u^{-1} : L \to K$ satisfies $\tau(\alpha) = \beta$. This proves the result when the degree of L/F equals 1.

Thus, by way of induction, assume that the degree d of L/F is greater than 1, and assume the result is true for all $(\widetilde{F}, \widetilde{g}(x), \widetilde{u} : \widetilde{F} \to \widetilde{L}, \widetilde{h}, \widetilde{\alpha}, \widetilde{\beta})$ with $\deg(\widetilde{L}/\widetilde{F})$ smaller than d. Since $d > 1$, $g(x)$ does not split as a product of linear polynomials in $F[x]$, i.e., there exists an irreducible factor of degree > 1. If $h(x)$ has degree > 1, define $\widehat{h}(x)$ to be $h(x)$, define $\widetilde{\alpha}$ to equal α, and define $\widetilde{\beta}$ to equal β. If $h(x)$ has degree 1, define $\widehat{h}(x)$ to be one of the irreducible factors of $g(x)$ of degree > 1, define $\widetilde{\alpha}$ to equal a root of $\widehat{h}(x)$ in L, and define $\widetilde{\beta}$ to equal a root of $\widehat{h}(x)$ in K.

Associated to the root $\widetilde{\alpha}$ of the irreducible polynomial $\widehat{h}(x) \in F[x]$, for the F-extension $K_{\widehat{h}} = F[t]/\widehat{h}(t)F[t]$ and its distinguished root $\widehat{\gamma} = [t]_{\widehat{h}(t)}$ of $\widehat{h}(x)$, there exists a unique homomorphism of F-extensions, $v_{\widetilde{\alpha}} : K_{\widehat{h}} \to L$ sending $\widehat{\gamma}$ to $\widetilde{\alpha}$. Similarly, associated to the root $\widetilde{\beta}$ of $\widehat{h}(x)$, there exists a unique homomorphism of F-extensions, $v_{\widetilde{\beta}} : K_{\widehat{h}} \to K$ sending $\widehat{\gamma}$ to $\widetilde{\beta}$. Since $\deg(L/F)$ equals $\deg(L/K_{\widehat{h}})\deg(K_{\widehat{h}}/F)$, and since $\deg(K_{\widehat{h}}/F) = \deg(\widehat{h}(x)) > 1$, $\deg(L/K_{\widehat{h}})$ is strictly less than d. Since L is a splitting field of $g(x)$ over F, also the extension $v_{\widehat{\alpha}}$ splits the polynomial $g(x) \in K_{\widehat{h}}[x]$. Since $v_{\widehat{\alpha}}$ is a homomorphism of F-extensions, the image of F in L is contained in the image of $K_{\widehat{h}}$ in L. Thus, every subfield of L that contains the image of $K_{\widehat{h}}$ also contains the image of F. Since $u : F \to L$ is a splitting field for $g(x) \in F[x]$, also $v_{\widetilde{\alpha}} : K_{\widehat{h}} \to L$ is a splitting field for $g(x) \in K_{\widehat{h}}[x]$.

Define $\widetilde{F}$ to be $K_{\widehat{h}}$. Define $\widetilde{L}$ to equal L. Define $\widetilde{K}$ to equal K. Define $\widetilde{v} : K_{\widehat{h}} \to L$ to be $v_{\widetilde{\alpha}}$. Define $\widetilde{w} : K_{\widehat{h}} \to K$ to be $v_{\widetilde{\beta}}$. Define $\widetilde{g}(x)$ to be the image of $g(x)$ in $K_{\widehat{k}}[x]$. Define $\widetilde{h}(x)$ to be $x - \widehat{\gamma}$. By construction, $\widetilde{\alpha}$ equals $\widetilde{v}(\widehat{\gamma})$, and $\widetilde{\beta}$ equals $\widetilde{w}(\widehat{\gamma})$. Thus, $\widetilde{\alpha}$ and $\widetilde{\beta}$ are roots of $\widetilde{h}(x)$ in the $\widetilde{F}$-extension $\widetilde{L}$, resp. $\widetilde{K}$. Thus, $(\widetilde{F}, \widetilde{g}(x), \widetilde{u} : \widetilde{F} \to \widetilde{L}, \widetilde{h}, \widetilde{\alpha}, \widetilde{\beta})$ is a datum as in the statement of the proposition, and $\deg(\widetilde{L}/\widetilde{F})$ is strictly less than $d = \deg(L/F)$. Thus, by the induction assumption, there exists a homomorphism of $\widetilde{F}$-extensions, $\widetilde{\tau} : L \to K$, such that $\widetilde{\beta}$ equals $\widetilde{\tau}(\widetilde{\alpha})$.

Define τ to equals $\widetilde{\tau}$. Since τ is a $\widetilde{F}$-extension, $\tau \circ v_{\widetilde{\alpha}}$ equals $v_{\widetilde{\beta}}$. Therefore,

$$\tau \circ v = \tau \circ (v_{\widetilde{\alpha}} \circ u_g) = (\tau \circ v_{\widetilde{\alpha}}) \circ u_g = v_{\widetilde{\beta}} \circ u_g = w.$$

Thus, τ is a homomorphism of F-extensions. If $h(x)$ has degree 1, then by the first paragraph, τ maps α to β. Finally, if $h(x)$ has degree > 1, then $\widehat{h}(x)$ equals $h(x)$, $\widetilde{\alpha}$ equals α, and $\widetilde{\beta}$ equals β. Since $\widetilde{\tau}$ maps $\widetilde{\alpha}$ to $\widetilde{\beta}$, τ maps α to β. Thus, the result is proved in all cases by induction on $\deg(L/F)$. □

Now let $g(x) \in F[x]$ be a nonzero, noninvertible polynomial, let $v : F \to L$ be a splitting field of $g(x)$, and let $v' : F \to L'$ be a splitting field of $g(x)$ (possibly v equals v').

Corollary 9.16. *For every irreducible factor $h(x)$ of $g(x)$ in $F[x]$, for every root α of $h(x)$ in L, for every root α' of $h(x)$ in L', there exists an isomorphism of F-extensions, $\tau : L \to L'$ such that $\tau(\alpha)$ equals α'.*

Proof. By the theorem, with $v : F \to L$ playing the role of the splitting field, there exists a homomorphism of F-extensions, $\tau : L \to L'$ such that $\tau(\alpha)$ equals α'. Since also $v' : F \to L'$ is a splitting field of $g(x)$, then applying the theorem once more, there exists a homomorphism of F-extensions $\tau' : L' \to L$ such that $\tau'(\alpha')$ equals α.

Consider the composite field extension $\sigma = \tau' \circ \tau$ from L to L. Since σ is a homomorphism of F-field extensions, it is an F-linear transformation. Since it is a ring homomorphism of fields, it is injective, so that the nullity of σ equals 0. Finally, since L is a finite-dimensional F-vector space, then by the Rank-Nullity Theorem, σ is also surjective. Thus, $\sigma : L \to L$ is an isomorphism of F-extensions. Therefore $\sigma^{-1} \circ \tau' : L' \to L$ is a left inverse of τ.

Permuting the roles of L and L', there is also a right inverse of τ. Thus, τ is invertible, i.e., τ is an isomorphism of F-extensions. □

9.6 Finite Fields are Splitting Fields

Because of the corollary, any two splitting fields of $g(x) \in F[x]$ are isomorphic as F-extensions. In particular, the degree of any of these splitting fields are equal, and this integer is called the **splitting degree** of $g(x)$. This leads to a characterization of finite fields.

Proposition 9.17. *For every finite field L with $q = p^e$ elements, $E_1 : \mathbb{F}_p \to L$ is a splitting field of $g_q(x) = x^q - x \in \mathbb{F}_p[x]$, and every element of L is a root of $g_q(x)$. In particular, any two finite fields with $q = p^e$ elements are isomorphic.*

Proof. We have already seen that every finite field L of characteristic p has $q = p^e$ elements for some integer $e \geq 1$, and every element α of L is a root of $g_q(x) = x^q - x$. Thus, for $u(g_q(x)) = x^q - x \in L[x]$, for every $\alpha \in L$, $x - \alpha$ is an irreducible factor for $x^q - x$ in $L[x]$. This is already q distinct linear factors of the degree q, monic polynomial $x^q - x$. Therefore, by the Unique Factorization of Polynomials, $g_q(x)$ factors in $L[x]$ as

$$g_q(x) = \prod_{\alpha \in L} (x - \alpha).$$

So L is a splitting field of $g_q(x)$.

Now, by the previous theorem, any two splitting fields of $g_q(x)$ are isomorphic. Thus, any two finite fields with $q = p^e$ elements are isomorphic as $\mathbb{F}_p$-extensions. □

Let F be a finite field with $q = p^e$ elements. Let $u : F \to L$ be a field extension. Recall that the function,

$$\mathrm{Frob}_{L,q} : L \to L, \ \alpha \mapsto \alpha^q,$$

is a homomorphism of $\mathbb{F}_q$-extensions.

Lemma 9.18. *The function $Frob_{L,q}$ is a homomorphism of F-extensions. In particular, it is an F-linear transformation. The kernel of the F-linear transformation $Frob_{L,q} - Id_L$ is precisely the image of u.*

Proof. By Lemma 9.11, $\mathrm{Frob}_{L,q}$ restricts as the identity on $u(F)$. Thus the homomorphism $\mathrm{Frob}_{L,q}$ of $\mathbb{F}_p$-field extensions is a homomorphism of F-field extensions.

For every $\alpha \in L$, if $\mathrm{Frob}_{L,q}(\alpha)$ equals α, then α is a root of $g_q(x) = x^q - x$. Since F is already the splitting field of $g_q(x)$, every root of $g_q(x)$ in L is contained in the image of u. Thus, the kernel of $\mathrm{Frob}_{L,q} - \mathrm{Id}_L$ is precisely the image of u. □

As a particular consequence, an element of $L[x]$ is in the image of $F[x]$ if and only if it is preserved by the induced ring homomorphism,

$$\mathrm{Frob}_{L,q} : L[x] \to L[x], \quad a_d x^d + \cdots + a_1 x^1 + a_0 x^0 \mapsto a_d^q x^d + \cdots + a_1^q x^1 + a_0^q x^0.$$

This leads to a characterization of the degree of an F-algebraic element $\alpha \in L^*$. Since $F[\alpha] \subset L$ is a finite F-extension of some degree d, by Corollary 8.3, $(F[\alpha])^*$ is a cyclic group of order $q^d - 1$. Thus, the multiplicative order m of $\alpha \in L^*$ is a finite integer dividing $q^d - 1$. In particular, $q^d \equiv 1 \pmod{m}$. Thus $[q]_m$ is an element of G_m that has some finite multiplicative order. By Lagrange's Theorem, this multiplicative order divides d.

Proposition 9.19. *For every F-algebraic element $\alpha \in L^*$ of multiplicative order m, the degree d of α as an F-algebraic element is the order of $[q]_m$ as an element of G_m. In particular, if L is a finite extension of F of degree e, then the degree d divides e.*

Proof. As remarked above, the multiplicative order r of $[q]_m \in G_m$ divides the degree d of α as an F-algebraic element. Conversely, form the sequence of r elements

$$(\alpha_0, \alpha_1, \ldots, \alpha_{r-1}) = (\alpha^{q^0}, \alpha^{q^1}, \ldots, \alpha^{q^{r-1}}),$$

of L^*. Observe that for every $i = 0, \ldots, r-2$, $\mathrm{Frob}_{L,q}(\alpha_i)$ equals α_{i+1}. Finally,

$$\mathrm{Frob}_{L,q}(\alpha_{r-1}) = (\alpha^{q^{r-1}})^q = \alpha^{q^r} = \alpha^{q^r - 1} \cdot \alpha = 1 \cdot \alpha = \alpha_0.$$

Thus, for the degree r, monic polynomial,

$$g_\alpha(x) = (x - \alpha_0) \cdot (x - \alpha_1) \cdots (x - \alpha_{r-2}) \cdot (x - \alpha_{r-1})$$

the image under $\mathrm{Frob}_{L,q}$ is,

$$\mathrm{Frob}_{L,q}(g_\alpha(x)) = (x - \alpha_1) \cdot (x - \alpha_2) \cdots (x - \alpha_{r-1}) \cdot (x - \alpha_0) = g_\alpha(x).$$

Since this polynomial is invariant under $\mathrm{Frob}_{L,q}$, it is the image under u of a unique, degree r, monic polynomial. For clarity, this polynomial is also denoted as $g_\alpha(x) \in F[x]$,

$$u(g_\alpha(x)) = (x - \alpha_0) \cdots (x - \alpha_{r-1}).$$

Since $\alpha = \alpha_0$ is a root of $u(g_\alpha(x))$, the minimal polynomial $m_{L/F,\alpha}(x)$ of $\alpha(x)$ divides $g_\alpha(x)$. Thus the degree d of $m_{L/F,\alpha}(x)$ is at most r. Since r divides d, d equals r.

If L/F has degree e, then L^* is a cyclic group of order $q^e - 1$. Thus, by Lagrange's Theorem, m divides $q^e - 1$, i.e., $[q]_m^e = [1]_m$. So the multiplicative order d of $[q]_m \in G_m$ divides e. □

Burnside's Lemma

Proposition 9.17 characterizes all finite fields with p^e elements as the splitting field of $g_{p^e}(x)$, assuming that there exists a finite field with p^e elements. It is natural to try to construct such a field as the splitting field of $g_{p^e}(x)$. In order to prove that this works, it is useful to recall some facts about the integer-coefficient polynomials

$$S_e(x) = x^0 + x^1 + x^2 + \cdots + x^m + \cdots + x^{e-1} + x^e = \sum_{m=0}^{e-1} x^m = \frac{x^e - 1}{x - 1}.$$

This is relevant since $g_{q^r}(x)$ equals

$$x^{q^r} - x = x(x^{q^r-1} - 1) = x(x-1)S_{q^r-1}(x).$$

Lemma 9.20. *For every integer $\ell \geq 1$, $S_\ell(0)$ equals 1, and $S_\ell(1)$ equals d. For every pair of integers $\ell, m \geq 1$, if ℓ divides m then $S_m(x) = S_{m/\ell}(x^\ell)S_\ell(x)$. More generally, for every pair of integers $m, n \geq 1$ with $\gcd(m,n) = \ell$, there exist integer-coefficient polynomials $B_{m,n}(x), B_{n,m}(x) \in \mathbb{Z}[x]$ such that*

$$B_{n,m}(x) \cdot S_n(x) + B_{m,n}(x) \cdot S_n(x) = S_\ell(x).$$

In particular, for every integer $a \geq 0$, the greatest common divisor of the integers $S_m(a)$ and $S_n(a)$ equals $S_\ell(a)$.

This is proved in the solutions to the Final Exam Review Sheet. In particular, consider the case that $\ell = q - 1$ and $m = q^r - 1$. Then m/ℓ equals $S_r(q) = 1 + q + \cdots + q^{r-1}$.

Lemma 9.21. *For a finite field F with $q = p^d$ elements that is the splitting field of $g_q(x) = x^q - x$, for every $\alpha \in F$, $S_{(q^e-1)/(q-1)}(\alpha^{q-1})$ equals 1. Thus, α is not a root of $S_{(q^e-1)/(q-1)}(x^{q-1})$. So F is not a splitting field of $g_{q^e}(x) = x(x-1)S_{q^e-1}(x)$.*

Proof. For $\alpha = 0$, α^{q-1} equals 0. For every integer $m \geq 0$, $S_m(0)$ equals 1. Next, for $\alpha \in F^*$, α^{q-1} equals 1. Also $(q^e - 1)/(q - 1)$ equals $S_r(q)$. Thus, $S_{(q^e-1)/(q-1)}(\alpha^{q-1})$ equals $S_{S_r(q)}(1) = S_r(q) \cdot 1$. Since $S_r(q) = 1 + q + \cdots + q^{r-1}$, $S_r(q)$ is congruent to 1 modulo p. Thus, $S_r(q) \cdot 1$ equals 1 in $\mathbb{F}_p$, hence also in F. Thus, for every $\alpha \in F$, $S_{(q^e-1)/(q-1)}(\alpha^{q-1})$ equals 1. Therefore α is not a root of $S_{(q^e-1)/(q-1)}(x^{q-1})$. Since $S_{q^e-1}(x) = S_{q-1}(x) \cdot S_{(q^e-1)/(q-1)}(x^{q-1})$, and since $S_{(q^e-1)/(q-1)}(x^{q-1})$ has no roots in F, $S_{q^e-1}(x)$ does not factor as a product of linear factors in $F[x]$. Since $g_{q^e}(x)$ equals $x(x-1)S_{q^e-1}(x)$, also $g_{q_e}(x)$ does not split in $F[x]$. □

Now we can prove that for every power of p, there does exist a finite field whose size equals that power of p.

Theorem 9.22. *For every finite field F with $q = p^m$ elements, for every integer $e \geq 1$, every splitting field $v : F \to L$ of $g_{q^e}(x) = x^{q^e} - x \in F[x]$ is a finite field with $q^e = p^{em}$ elements. Conversely, for every finite field L with q^e elements, there exists a ring homomorphism $v : F \to L$ that is a splitting field of $g_{q_e}(x)$.*

Proof. The last statement above is simply a reformulation of the previous proposition. The main issue is proving that every splitting field of $g_{q^e}(x) \in F[x]$ has degree e as an extension of F.

This is proved by induction on $e \geq 1$. For $e = 1$, then, by the proposition, the identity function $\mathrm{Id}_F : F \to F$ is already a splitting field of $g_q(x)$. Moreover, for every field L with q elements, also L is a splitting field of $g_q(x)$. Thus there exists an isomorphism $v : F \to L$.

Thus, by way of induction on e, assume that $e > 1$, and assume that the result is proved for all smaller values of e. Since $e > 1$, e is not invertible. Thus, there exists a prime $\ell > 1$ such that $e = d\ell$. Since $1 \leq d < e$, a splitting field of $g_{q^d}(x)$ over F is a finite field with q^d elements. Up to replacing F by this splitting field, replacing q by q^d, and replacing e by ℓ, assume that e is already a prime.

By Proposition 9.14, there exists a splitting field $v : F \to L$ of $g_{q^e}(x)$. Denote the degree $\deg(L/F)$ by d. By Lemma 9.21, d is greater than 1, i.e., F is not already a splitting of $g_{q^e}(x)$. By Corollary 8.3, L^* is a cyclic group of order $q^d - 1$. The subset of L^* of roots of $x^{q^e} - x$ is precisely the subgroup H of elements whose order m divides $q^e - 1$. In particular, this is a cyclic group. Let $\alpha \in L^*$ be a generator of this cyclic group, and let m denote the multiplicative order of α in L^*. Since α is a generator for H, every nonzero root of $x^{q^e} - x$ is a power of α, and hence the root is an element of the F-subextension $F[\alpha] \subset L$. Also 0 is in $F[\alpha]$. Thus, every root of $x^{q^e} - x$ in L is already contained in $F[\alpha]$, i.e., $x^{q^e} - x$ splits in $F[\alpha]$. Since L is a splitting field of $x^{q^e} - x$, L equals $F[\alpha]$. Thus, the degree d equals the degree of α over F. By Proposition 9.19, d equals the multiplicative order of $[q]_m$ in G_m. Since α^{q^e} equals α, also α^{q^e-1} equals 1. Thus, $[q]_m^e = [1]_m$, so d divides e. Since $d > 1$, and since e is prime, d equals e. Thus, the splitting field L of $g_{q^e}(x) \in F[x]$ is an extension of F of degree e, i.e., $\#L$ equals q^e. Thus, the theorem is proved by induction on e. □

As a consequence of the proof, please note that for every prime power $q = p^r$, for a finite field F of size $q = p^r$, for every integer $e \geq 1$, for a finite extension $F \to L$ of degree e, the smallest power of the F-algebra homomorphism $\mathrm{Frob}_{L,q} : L \to L$ that equals the identity function is e. This is essentially equivalent to the "Fundamental Theorem of Galois Theory" for the field extension $F \to L$.