CHARACTERS OF FINITE ABELIAN GROUPS (SHORT VERSION)

KEITH CONRAD

1. Introduction

The theme we will study is an analogue on finite abelian groups of Fourier analysis on $\mathbf{R}$. A Fourier series on the real line is the following type of series in sines and cosines:

$$f(x) = \sum_{n \geq 0} a_n \cos(nx) + \sum_{n \geq 1} b_n \sin(nx).$$

This is 2π-periodic. Since $e^{inx} = \cos(nx) + i\sin(nx)$ and $e^{-inx} = \cos(nx) - i\sin(nx)$, a Fourier series can also be written in terms of complex exponentials:

$$f(x) = \sum_{n \in \mathbf{Z}} c_n e^{inx},$$

where $c_0 = a_0$, $c_n = \frac{1}{2}(a_n - b_n i)$ for $n > 0$, and $c_n = \frac{1}{2}(a_{|n|} + b_{|n|} i)$ for $n < 0$. The convenient algebraic property of e^{inx}, not shared by sines and cosines, is that it is a group homomorphism from $\mathbf{R}$ to the unit circle $S^1 = \{z \in \mathbf{C} : |z| = 1\}$:

$$e^{in(x+x')} = e^{inx} e^{inx'}.$$

We now replace $\mathbf{R}$ with a finite abelian group. Here is the analogue of the functions e^{inx}.

Definition 1.1. A *character* of a finite abelian group G is a homomorphism $\chi \colon G \to S^1$.

We will usually write abstract groups multiplicatively, so $\chi(g_1 g_2) = \chi(g_1)\chi(g_2)$ and $\chi(1) = 1$.

Example 1.2. The *trivial* character of G is the homomorphism $\mathbf{1}_G$ defined by $\mathbf{1}_G(g) = 1$ for all $g \in G$.

Example 1.3. Let G be cyclic of order 4 with generator γ. Since $\gamma^4 = 1$, a character χ of G has $\chi(\gamma)^4 = 1$, so χ takes only four possible values at γ, namely 1, -1, i, or $-i$. Once $\chi(\gamma)$ is known, the value of χ elsewhere is determined by multiplicativity: $\chi(\gamma^j) = \chi(\gamma)^j$. So we get four characters, whose values can be placed in a table. See Table 1.

	1	γ	γ^2	γ^3
$\mathbf{1}_G$	1	1	1	1
χ_1	1	-1	1	-1
χ_2	1	i	-1	$-i$
χ_3	1	$-i$	-1	i

Table 1.

When G has size n and $g \in G$, for all characters χ of G we have $\chi(g)^n = \chi(g^n) = \chi(1) = 1$, so the values of χ lie among the nth roots of unity in S^1. More precisely, the order of $\chi(g)$ divides the order of g (which divides $|G|$).

Characters on finite abelian groups were first studied in number theory, which is a source of many interesting finite abelian groups. For instance, Dirichlet used characters of the group $(\mathbf{Z}/(m))^\times$ to prove that when $(a, m) = 1$ there are infinitely many primes $p \equiv a \bmod m$. The quadratic reciprocity law of elementary number theory is concerned with a deep property of a particular character, the Legendre symbol. Fourier series on finite abelian groups have applications in engineering: signal processing (the fast Fourier transform [1, Chap. 9]) and error-correcting codes [1, Chap. 11].

To provide a context against which our development of characters on finite abelian groups can be compared, Section 2 discusses classical Fourier analysis on the real line. In Section 3 we discuss some properties of characters of finite abelian groups and introduce their dual groups. Section 4 uses characters of a finite abelian group to develop a finite analogue of Fourier series. In Section 5 we use characters to prove a structure theorem for finite abelian groups. In Section 6 we use characters to count solutions to a congruence mod p.

Our notation is completely standard, but we make two remarks about it. For a complex-valued function $f(x)$, the complex-conjugate function is usually denoted $\overline{f}(x)$ instead of $\overline{f(x)}$ to stress that conjugation creates a new function. (We sometimes use the overline notation also to mean the reduction $\overline{g}$ into a quotient group.) For $n \geq 1$, we write μ_n for the group of nth roots of unity in the unit circle S^1. It is a cyclic group of size n.

Exercises.

1. Make a character table for $\mathbf{Z}/(2) \times \mathbf{Z}/(2)$, with columns labeled by elements of the group and rows labeled by characters, as in Table 1.
2. Let G be a finite nonabelian simple group. (Examples include A_n for $n \geq 5$.) Show the only group homomorphism $\chi\colon G \to S^1$ is the trivial map.

2. Classical Fourier analysis

This section serves as motivation for our later treatment of finite abelian groups, where there will be no convergence issues (just finite sums!), so we take a soft approach and sidestep analytic technicalities that a serious treatment of Fourier analysis on $\mathbf{R}$ demands.

Fourier analysis for periodic functions on $\mathbf{R}$ is based on the functions e^{inx} for $n \in \mathbf{Z}$. Every "reasonably nice" function $f\colon \mathbf{R} \to \mathbf{C}$ of period 2π can be expanded into a series

$$f(x) = \sum_{n \in \mathbf{Z}} c_n e^{inx},$$

where the sum runs over $\mathbf{Z}$ and the nth Fourier coefficient c_n can be recovered as an integral:

$$c_n = \frac{1}{2\pi} \int_0^{2\pi} f(x) e^{-inx}\, \mathrm{d}x. \tag{2.1}$$

This formula for c_n can be explained by replacing $f(x)$ in (2.1) by its Fourier series and integrating termwise (for "reasonably nice" functions this termwise integration is analytically justifiable), using the formula

$$\frac{1}{2\pi} \int_0^{2\pi} e^{imx} e^{-inx}\, \mathrm{d}x = \begin{cases} 1, & \text{if } m = n, \\ 0, & \text{if } m \neq n. \end{cases}$$

Rather than working with functions $f\colon \mathbf{R} \to \mathbf{C}$ having period 2π, formulas look cleaner using functions $f\colon \mathbf{R} \to \mathbf{C}$ having period 1. The basic exponentials become $e^{2\pi inx}$ and the Fourier series and coefficients for f are

$$f(x) = \sum_{n\in\mathbf{Z}} c_n e^{2\pi inx}, \qquad c_n = \int_0^1 f(x)e^{-2\pi inx}\,\mathrm{d}x. \tag{2.2}$$

Note c_n in (2.2) is not the same as c_n in (2.1).

In addition to Fourier series there are Fourier integrals. The *Fourier transform* of a function f that decays rapidly at $\pm\infty$ is the function $\widehat{f}\colon \mathbf{R} \to \mathbf{C}$ defined by the integral formula

$$\widehat{f}(y) = \int_{\mathbf{R}} f(x)e^{-2\pi ixy}\,\mathrm{d}x.$$

The analogue of the expansion (2.2) of a periodic function into a Fourier series is the Fourier inversion formula, which expresses f in terms of its Fourier transform $\widehat{f}$:

$$f(x) = \int_{\mathbf{R}} \widehat{f}(y)e^{2\pi ixy}\,\mathrm{d}y.$$

Example 2.1. A Gaussian is a function of the form ae^{-bx^2}, where $b > 0$. For example, the Gaussian $(1/\sqrt{2\pi})e^{-(1/2)x^2}$ is important in probability theory. The Fourier transform of a Gaussian is another Gaussian:

$$\int_{\mathbf{R}} ae^{-bx^2}e^{-2\pi ixy}\,\mathrm{d}x = \sqrt{\frac{\pi}{b}}ae^{-\pi^2y^2/b}. \tag{2.3}$$

This formula shows that a highly peaked Gaussian (large b) has a Fourier transform that is a spread out Gaussian (small π^2/b) and *vice versa*. More generally, there is a sense in which a function and its Fourier transform can't both be highly localized; this is a mathematical incarnation of Heisenberg's uncertainty principle from physics.

There are several conventions for where 2π appears in the Fourier transform. Table 2 collects three different 2π-conventions. The first column of Table 2 is a definition and the second column is a theorem (Fourier inversion).

$\widehat{f}(y)$	$f(x)$
$\int_{\mathbf{R}} f(x)e^{-2\pi ixy}\,\mathrm{d}x$	$\int_{\mathbf{R}} \widehat{f}(y)e^{2\pi ixy}\,\mathrm{d}y$
$\int_{\mathbf{R}} f(x)e^{-ixy}\,\mathrm{d}x$	$\frac{1}{2\pi}\int_{\mathbf{R}} \widehat{f}(y)e^{ixy}\,\mathrm{d}y$
$\frac{1}{\sqrt{2\pi}}\int_{\mathbf{R}} f(x)e^{-ixy}\,\mathrm{d}x$	$\frac{1}{\sqrt{2\pi}}\int_{\mathbf{R}} \widehat{f}(y)e^{ixy}\,\mathrm{d}y$

TABLE 2.

A link between Fourier series and Fourier integrals is the *Poisson summation formula*: for a "nice" function $f\colon \mathbf{R} \to \mathbf{C}$ that decays rapidly enough at $\pm\infty$,

$$\sum_{n\in\mathbf{Z}} f(n) = \sum_{n\in\mathbf{Z}} \widehat{f}(n), \tag{2.4}$$

where $\widehat{f}(y) = \int_{\mathbf{R}} f(x)e^{-2\pi ixy}\,dx$. For example, when $f(x) = e^{-bx^2}$ (with $b > 0$), the Poisson summation formula says

$$\sum_{n\in\mathbf{Z}} e^{-bn^2} = \sum_{n\in\mathbf{Z}} \sqrt{\frac{\pi}{b}} e^{-\pi^2 n^2/b},$$

To prove the Poisson summation formula, we use Fourier series. Periodize $f(x)$ as

$$F(x) = \sum_{n\in\mathbf{Z}} f(x+n).$$

Since $F(x+1) = F(x)$, write F as a Fourier series: $F(x) = \sum_{n\in\mathbf{Z}} c_n e^{2\pi inx}$. Then

$$\begin{aligned} c_n &= \int_0^1 F(x)e^{-2\pi inx}\,dx \\ &= \int_0^1 \left(\sum_{m\in\mathbf{Z}} f(x+m)\right) e^{-2\pi inx}\,dx \\ &= \sum_{m\in\mathbf{Z}} \int_0^1 f(x+m)e^{-2\pi inx}\,dx \\ &= \sum_{m\in\mathbf{Z}} \int_m^{m+1} f(x)e^{-2\pi inx}\,dx \\ &= \int_{\mathbf{R}} f(x)e^{-2\pi inx}\,dx \\ &= \widehat{f}(n). \end{aligned}$$

Therefore the expansion of $F(x)$ into a Fourier series is equivalent to

$$\sum_{n\in\mathbf{Z}} f(x+n) = \sum_{n\in\mathbf{Z}} \widehat{f}(n)e^{2\pi inx}, \tag{2.5}$$

which becomes the Poisson summation formula (2.4) by setting $x = 0$.

Exercises.

1. Without dwelling on analytic subtleties, check from Fourier inversion that $\widehat{\widehat{f}}(x) = f(-x)$ (if the Fourier transform is defined suitably).
2. For a function $f\colon \mathbf{R} \to \mathbf{C}$ and $c \in \mathbf{R}$, let $g(x) = f(x+c)$. Define the Fourier transform of a function h by $\widehat{h}(y) = \int_{\mathbf{R}} h(x)e^{-2\pi ixy}\,dx$. If f has a Fourier transform, show g has Fourier transform $\widehat{g}(y) = e^{2\pi icy}\widehat{f}(y)$.
3. Assuming the Fourier inversion formula holds for a definition of the Fourier transform as in Table 2, check that for all α and β in $\mathbf{R}^\times$ that if we set

$$(\mathcal{F}f)(y) = \alpha \int_{\mathbf{R}} f(x)e^{-i\beta xy}\,dx$$

for all x then

$$f(x) = \frac{\beta}{2\pi\alpha} \int_{\mathbf{R}} (\mathcal{F}f)(y)e^{i\beta xy}\,dy.$$

(If $\beta = 2\pi\alpha^2$ then these two equations are symmetric in the roles of f and $\mathcal{F}f$ except for a sign in the exponential term.)

3. Finite Abelian Group Characters

We leave the real line and turn to the setting of finite abelian groups G. Our interest shifts from the functions e^{inx} to characters: homomorphisms from $G \to S^1$. The construction of characters of these groups begins with the case of cyclic groups.

Theorem 3.1. *Let G be a finite cyclic group of size n with a chosen generator γ. There are exactly n characters of G, each determined by sending γ to the different nth roots of unity in $\mathbf{C}$.*

Proof. We mimic Example 1.3, where G is cyclic of size 4. Since γ generates G, a character is determined by its value on γ and that value must be an nth root of unity (not necessarily of exact order n, *e.g.*, $\mathbf{1}_G(\gamma) = 1$), so there are at most n characters. We now write down n characters.

Let ζ be an nth root of unity in $\mathbf{C}$. Set $\chi(\gamma^j) = \zeta^j$ for $j \in \mathbf{Z}$. This formula is well-defined (if $\gamma^j = \gamma^k$ for two different integer exponents j and k, we have $j \equiv k \bmod n$ so $\zeta^j = \zeta^k$), and χ is a homomorphism. Of course χ depends on ζ. As ζ changes, we get different characters (their values at γ are changing), so in total we have n characters. □

To handle characters of non-cyclic groups, the following lemma is critical.

Lemma 3.2. *Let G be a finite abelian group, $H \subset G$ a proper subgroup, and $\chi\colon H \to S^1$ a character of H. For $g \in G - H$, there is an extension of χ to a character on $\langle H, g\rangle$.*

Proof. We want to extend χ to a character $\widetilde{\chi}$ of $\langle H, g\rangle$.

What is a possible value for $\widetilde{\chi}(g)$? Since $g \notin H$, $\widetilde{\chi}(g)$ is not initially defined. But some nonzero power of g is in H (*e.g.*, $g^{|G|} = 1 \in H$), and on these powers χ is defined. Pick $d \geq 1$ minimal with $g^d \in H$. That is, d is the order of g in G/H, so $d = [\langle H, g\rangle : H]$. If there is a character $\widetilde{\chi}$ on $\langle, H, g$ that extends χ on H then $\widetilde{\chi}(g)$ *must* be an d-th root of $\chi(g^d)$ since we must have $\widetilde{\chi}(g)^d = \widetilde{\chi}(g^d) = \chi(g^d)$. That is our clue: define $\widetilde{\chi}(g) \in S^1$ to be a solution to $z^d = \chi(g^d)$:

$$\widetilde{\chi}(g)^d = \chi(g^d). \tag{3.1}$$

Once we have chosen $\widetilde{\chi}(g)$ to satisfy (3.1), define $\widetilde{\chi}$ on $\langle H, g\rangle$ by

$$\widetilde{\chi}(hg^i) := \chi(h)\widetilde{\chi}(g)^i. \tag{3.2}$$

This formula covers all possible elements of $\langle H, g\rangle$, but is $\widetilde{\chi}$ well-defined? Perhaps H and $\langle g\rangle$ overlap nontrivially, so the expression of an element of $\langle H, g\rangle$ in the form hg^i is not unique. We have to show this doesn't lead to an inconsistency in the value of $\widetilde{\chi}$ in (3.2). Suppose $hg^i = h'g^{i'}$. Then $g^{i-i'} \in H$, so $i' \equiv i \bmod d$ since d is the order of g in G/H. Write $i' = i + dd'$, so $h = h'a^{i'-i} = h'g^{dd'}$. The terms h, h', and g^d are in H, so

$$\begin{aligned}
\chi(h')\widetilde{\chi}(g)^{i'} &= \chi(h')\widetilde{\chi}(g)^i\widetilde{\chi}(g)^{dd'} \\
&= \chi(h')\widetilde{\chi}(g)^i\chi(g^d)^{d'} \text{ since } \widetilde{\chi}(g)^d = \chi(g^d) \\
&= \chi(h'g^{dd'})\widetilde{\chi}(g)^i \\
&= \chi(h)\widetilde{\chi}(g)^i.
\end{aligned}$$

Therefore $\widetilde{\chi}\colon \langle H, g\rangle \to S^1$ is a well-defined function and it is easily checked to be a homomorphism. It restricts to χ on H. □

Theorem 3.3. *For a finite abelian group G and subgroup H, each character of H can be extended to a character of G.*

Proof. Let $\chi\colon H \to S^1$ be a character of H.

Since G is finite, it has a finite generating set $\{g_1, \ldots, g_k\}$ (*e.g.*, $\{g_i\}$ could be a listing of all the elements of G). Therefore we can build up a tower of subgroups from H to G by adjoining the elements g_i one at a time:

$$H \subset \langle H, g_1\rangle \subset \langle H, g_1, g_2\rangle \subset \cdots \subset \langle H, g_1, \ldots, g_k\rangle = G.$$

Each step along this tower has the form $H_i \subset \langle H_i, g_i\rangle$, where $H_0 = H$. By applying Lemma 3.2 at each step of the tower, χ can be extended as a character from H to H_1 to H_2, and so on up to $H_k = G$. □

Let's refine this to count the number of extensions of a character from H to G.

Theorem 3.4. *For a finite abelian group G and subgroup H, each character of H can be extended to a character of G in $[G : H]$ ways.*

Proof. We will induct on the index $[G : H]$. The result is clear when $[G : H] = 1$, *i.e.*, $H = G$, so suppose $[G : H] > 1$ and the theorem is proved for characters on subgroups of index smaller than $[G : H]$.

Pick $g \in G$ with $g \notin H$, so

$$H \subset \langle H, g\rangle \subset G.$$

To extend a character $\chi\colon H \to S^1$ to G, we at least need to be able to extend χ to a character $\widetilde{\chi}$ on $\langle H, g\rangle$. Let's count the number of ways to do that. Then we will use induction to count the number of extension of each character from $\langle H, g\rangle$ all the way up to G.

Let d be the smallest positive integer such that $g^d \in H$. An extension of χ on H to a character $\widetilde{\chi}$ on $\langle H, g\rangle$ is determined by $\widetilde{\chi}(g)$, and this value has to satisfy the condition $\widetilde{\chi}(g)^d = \chi(g^d)$. Each number in S^1 has d different d-th roots in S^1, so there are d potential values for $\widetilde{\chi}(g)$. The proof of Lemma 3.2 shows all of them really work.

The number of choices of $\widetilde{\chi}$ extending χ is the number of choices for $\widetilde{\chi}(g)$, which is $d = [\langle H, g\rangle : H]$. Since $[G : \langle H, g\rangle] < [G : H]$, by induction on the index there are $[G : \langle H, g\rangle]$ extensions of each $\widetilde{\chi}$ to a character of G, so the number of extensions of a character on H to a character on G is $[G : \langle H, g\rangle][\langle H, g\rangle : H] = [G : H]$. □

Theorem 3.5. *If $g \neq 1$ in a finite abelian group G then $\chi(g) \neq 1$ for some character χ of G. The number of characters of G is $|G|$.*

Proof. The cyclic group $\langle g\rangle$ is nontrivial, say of size n, so $n > 1$. In S^1 there is a cyclic subgroup of order n, namely the group μ_n of n-th roots of unity. There is an isomorphism $\langle g\rangle \cong \mu_n$, which can be viewed as a character of $\langle g\rangle$. By Theorem 3.3, this character of $\langle g\rangle$ extends to a character of G and does not send g to 1.

To show G has $|G|$ characters, apply Theorem 3.4 with H the trivial subgroup. □

We have used two important features of S^1 as the target group for characters: for each $d \geq 1$ the dth power map on S^1 is d-to-1 (proof of Theorem 3.4) and for each $n \geq 1$ there is a cyclic subgroup of order n in S^1 (proof of Theorem 3.5).

Corollary 3.6. *If G is a finite abelian group and $g_1 \neq g_2$ in G then there is a character of G that takes different values at g_1 and g_2.*

Proof. Apply Theorem 3.5 to $g = g_1 g_2^{-1}$. □

Corollary 3.6 shows the characters of G "separate" the elements of G: different elements of the group admit a character taking different values on them.

Corollary 3.7. *If G is a finite abelian group and $H \subset G$ is a subgroup and $g \in G$ with $g \notin H$ then there is a character of G that is trivial on H and not equal to 1 at g.*

Proof. We work in the group G/H, where $\overline{g} \neq \overline{1}$. By Theorem 3.5 there is a character of G/H that is not 1 at $\overline{g}$. Composing this character with the reduction map $G \to G/H$ yields a character of G that is trivial on H and not equal to 1 at g. □

It is easy to find functions on G that separate elements without using characters. For $g \in G$, define $\delta_g \colon G \to \{0,1\}$ by

$$\delta_g(x) = \begin{cases} 1, & \text{if } x = g, \\ 0, & \text{if } x \neq g. \end{cases} \tag{3.3}$$

These functions separate elements of the group, but characters do this too and have better algebraic properties: they are group homomorphisms.

Our definition of a character makes sense on nonabelian groups, but there will not be enough such characters for Theorem 3.5 to hold if G is finite and nonabelian: a homomorphism $\chi \colon G \to S^1$ must equal 1 on the commutator subgroup $[G,G]$, which is a nontrivial subgroup, so such homomorphisms can't distinguish elements in $[G,G]$ from each other. If $g \notin [G,G]$ then in the finite abelian group $G/[G,G]$ the coset of g is nontrivial so there is a character $G/[G,G] \to S^1$ that's nontrivial on $\overline{g}$. Composing this character with the reduction map $G \to G/[G,G]$ produces a homomorphism $G \to S^1$ that is nontrivial on g.

Definition 3.8. For a character χ on a finite abelian group G, the *conjugate character* is the function $\overline{\chi} \colon G \to S^1$ given by $\overline{\chi}(g) := \overline{\chi(g)}$.

Since a complex number z with $|z| = 1$ has $\overline{z} = 1/z$, $\overline{\chi}(g) = \chi(g)^{-1} = \chi(g^{-1})$.

Definition 3.9. The *dual group* of a finite abelian group G is the set of homomorphisms $G \to S^1$ with the group law of pointwise multiplication of functions: $(\chi\psi)(g) = \chi(g)\psi(g)$. The dual group of G is denoted $\widehat{G}$.

The trivial character of G is the identity in $\widehat{G}$ and the inverse of a character is its conjugate character. Note $\widehat{G}$ is abelian since multiplication in $\mathbf{C}^\times$ is commutative.

Theorem 3.5 says in part that

$$|G| = |\widehat{G}|. \tag{3.4}$$

In fact, the groups G and $\widehat{G}$ are *isomorphic*. First let's check this on cyclic groups.

Theorem 3.10. *If G is cyclic then $G \cong \widehat{G}$ as groups.*

Proof. We will show $\widehat{G}$ is cyclic. Then since G and $\widehat{G}$ have the same size they are isomorphic.

Let $n = |G|$ and γ be a generator of G. Set $\chi \colon G \to S^1$ by $\chi(\gamma^j) = e^{2\pi ij/n}$ for all j. For other characters $\psi \in \widehat{G}$, we have $\psi(\gamma) = e^{2\pi ik/n}$ for some integer k, so $\psi(\gamma) = \chi(\gamma)^k$. Then

$$\psi(\gamma^j) = \psi(\gamma)^j = \chi(\gamma)^{jk} = \chi(\gamma^j)^k,$$

which shows $\psi = \chi^k$. Therefore χ generates $\widehat{G}$. □

Lemma 3.11. *If A and B are finite abelian groups, there is an isomorphism $\widehat{A \times B} \cong \widehat{A} \times \widehat{B}$.*

Proof. Let χ be a character on $A \times B$. Identify the subgroups $A \times \{1\}$ and $\{1\} \times B$ of $A \times B$ with A and B in the obvious way. Let χ_A and χ_B be the restrictions of χ to A and

B respectively, *i.e.*, $\chi_A(a) = \chi(a,1)$ and $\chi_B(b) = \chi(1,b)$. Then χ_A and χ_B are characters of A and B and $\chi(a,b) = \chi((a,1)(1,b)) = \chi(a,1)\chi(1,b) = \chi_A(a)\chi_B(b)$. So we get a map

$$\widehat{A \times B} \to \widehat{A} \times \widehat{B} \tag{3.5}$$

by sending χ to (χ_A, χ_B). It is left to the reader to check (3.5) is a group homomorphism. Its kernel is trivial since if χ_A and χ_B are trivial characters then $\chi(a,b) = \chi_A(a)\chi_B(b) = 1$, so χ is trivial. Both sides of (3.5) have the same size by (3.4), so (3.5) is an isomorphism. □

Theorem 3.12. *If G is a finite abelian group then G is isomorphic to $\widehat{G}$.*

Proof. The case when G is cyclic was Theorem 3.10. Lemma 3.11 extends easily to several factors in a direct product:

$$(H_1 \times \cdots \times H_r)\widehat{\ } \cong \widehat{H}_1 \times \cdots \times \widehat{H}_r. \tag{3.6}$$

When H_i is cyclic, $\widehat{H}_i \cong H_i$, so (3.6) tells us that that character group of $H_1 \times \cdots \times H_r$ is isomorphic to itself. Every finite abelian group is isomorphic to a direct product of cyclic groups, so the character group of a finite abelian group is isomorphic to itself. □

Although G and $\widehat{G}$ are isomorphic groups, there is not a *natural* isomorphism between them, even when G is cyclic. For instance, to prove $G \cong \widehat{G}$ when G is cyclic we had to *choose* a generator. If we change the generator, then the isomorphism changes.[1]

The double-dual group $\widehat{\widehat{G}}$ is the dual group of $\widehat{G}$. Since G and $\widehat{G}$ are isomorphic, G and $\widehat{\widehat{G}}$ are isomorphic. However, while there isn't a natural isomorphism from G to $\widehat{G}$, there *is* a natural isomorphism from G to $\widehat{\widehat{G}}$. The point is that there is a natural way to map G to its double-dual group: associate to each $g \in G$ the function "evaluate at g," which is the function $\widehat{G} \to S^1$ given by $\chi \mapsto \chi(g)$. Here g is fixed and χ varies. This is a character of $\widehat{G}$, since $(\chi_1\chi_2)(g) = \chi_1(g)\chi_2(g)$ by definition.

Theorem 3.13. *Let G be a finite abelian group. The homomorphism $G \to \widehat{\widehat{G}}$ associating to $g \in G$ the function "evaluate at g" is an isomorphism.*

Proof. Since a finite abelian group and its dual group have the same size, a group and its double-dual group have the same size, so it suffices to show this homomorphism is injective. If $g \in G$ is in the kernel then every element of $\widehat{G}$ is 1 at g, so $g = 1$ by Theorem 3.5. □

Theorem 3.13 is called *Pontryagin duality*. This label actually applies to a more general result about characters of locally compact abelian groups. Finite abelian groups are a special case, where difficult analytic techniques can be replaced by counting arguments. The isomorphism between G and its double-dual group given by Pontryagin duality lets us think about a finite abelian group G as a dual group (namely the dual group of $\widehat{G}$).

The isomorphism in Pontryagin duality is natural: it does not depend on *ad hoc* choices (unlike the isomorphism between a finite abelian group and its dual group).

Exercises.

[1] If G is trivial or of order 2, then it has a unique generator, so in that case we could say the isomorphism $G \cong \widehat{G}$ is canonical.

1. Let's find the characters of the additive group $(\mathbf{Z}/(m))^r$, an r-fold direct product.
   (a) For $k \in \mathbf{Z}/(m)$, let $\chi_k \colon \mathbf{Z}/(m) \to S^1$ by
   $$\chi_k(j) = e^{2\pi ijk/m},$$
   so $\chi_k(1) = e^{2\pi ik/m}$. Show $\chi_0, \chi_1, \ldots, \chi_{m-1}$ are all the characters of $\mathbf{Z}/(m)$ and $\chi_k\chi_l = \chi_{k+l}$.
   (b) Let $r \geq 1$. For r-tuples $\mathbf{a}$, $\mathbf{b}$ in $(\mathbf{Z}/(m))^r$, let
   $$\mathbf{a} \cdot \mathbf{b} = a_1b_1 + \cdots + a_rb_r \in \mathbf{Z}/(m)$$
   be the usual dot product. For $\mathbf{k} \in (\mathbf{Z}/(m))^r$, let $\chi_{\mathbf{k}}(\mathbf{j}) = e^{2\pi i(\mathbf{j}\cdot\mathbf{k})/m}$. Show the functions $\chi_{\mathbf{k}}$ are all the characters of $(\mathbf{Z}/(m))^r$ and $\chi_{\mathbf{k}}\chi_{\mathbf{l}} = \chi_{\mathbf{k}+\mathbf{l}}$.
2. Show the following are equivalent properties of a character χ: $\chi(g) = \pm 1$ for all g, $\overline{\chi}(g) = \chi(g)$ for all g, and $\chi^2 = \mathbf{1}_G$.
3. Describe the error in the following bogus proof of Theorem 3.4. Let $m = [G : H]$ and pick a set of coset representatives $g_1, \ldots, g_m$ for G/H. Given a character χ on H, define $\widetilde{\chi}$ on G by first picking the m $(= [G : H])$ values $\widetilde{\chi}(g_i)$ for $1 \leq i \leq m$ and then writing each $g \in G$ in the (unique) form g_ih and defining $\widetilde{\chi}(g) = \widetilde{\chi}(g_i)\chi(h)$. This defines $\widehat{\chi}$ on G, and since we had to make m choices there are m characters.
4. For finite nonabelian G, show the characters of G (that is, homomorphisms $G \to S^1$) separate elements modulo $[G, G]$: $\chi(g_1) = \chi(g_2)$ for all χ if and only if $g_1 = g_2$ in $G/[G, G]$.
5. This exercise will give an interpretation of characters as eigenvectors. For a finite abelian group G and $g \in G$, let $T_g \colon L(G) \to L(G)$ by $(T_gf)(x) = f(gx)$.
   (a) Show the T_g's are commuting linear transformations and each character of G is an eigenvector of each T_g.
   (b) If f is a simultaneous eigenvector of all the T_g's, show $f(1) \neq 0$ (if $f(1) = 0$ conclude f is identically zero, but the zero vector is not an eigenvector) and then after rescaling f so $f(1) = 1$ deduce that f is a character of G. Thus the characters of G are the simultaneous eigenvectors of the T_g's, suitably normalized.
   (c) Show the T_g's are each diagonalizable. Deduce from this and parts (a) and (b) that $\widehat{G}$ is a basis of $L(G)$, so $|\widehat{G}| = \dim L(G) = |G|$. (This gives a different proof that G and $\widehat{G}$ have the same size.)
6. For a subgroup H of a finite abelian group G, let
   $$H^\perp = \{\chi \in \widehat{G} : \chi = 1 \text{ on } H\}.$$
   These are the characters of G that are trivial on H. For example, $G^\perp = \{\mathbf{1}_G\}$ and $\{1\}^\perp = \widehat{G}$. Note $H^\perp \subset \widehat{G}$ and $H^\perp$ depends on H and G.
   Show $H^\perp$ is a subgroup of $\widehat{G}$, it is isomorphic to $\widehat{G/H}$, and $\widehat{G}/(H^\perp) \cong \widehat{H}$. In particular, $|H^\perp| = [G : H]$.
7. Let G be finite abelian and $H \subset G$ be a subgroup.
   (a) Viewing $H^{\perp\perp} = (H^\perp)^\perp$ in G using Pontryagin duality, show $H^{\perp\perp} = H$. (Hint: The inclusion in one direction is easy. Count sizes for the other inclusion.)
   (b) Show for each m dividing $|G|$ that
   $$|\{H \subset G : |H| = m\}| = |\{H \subset G : [G : H] = m\}|$$
   by associating H to $H^\perp$ and using a (fixed) isomorphism of G with $\widehat{G}$.

(c) For a finite abelian group G, part b says the number of subgroups of G with index 2 is equal to the number of elements of G with order 2. Use this idea to count the number of subgroups of $(\mathbf{Z}/(m))^\times$ with index 2. (The answer depends on the number of odd prime factors of m and the highest power of 2 dividing m.)

(d) Show, for a prime p, that the number of subspaces of $(\mathbf{Z}/(p))^n$ with dimension d equals the number of subspaces with dimension $n-d$.

8. For a finite abelian group G, let $G[n] = \{g \in G : g^n = 1\}$ and $G^n = \{g^n : g \in G\}$. Both are subgroups of G. Prove $G[n]^\perp = (\widehat{G})^n$ and $(G^n)^\perp = \widehat{G}[n]$ in $\widehat{G}$.

4. Finite Fourier series

Let G be a finite abelian group. Set

$$L(G) = \{f\colon G \to \mathbf{C}\},$$

the $\mathbf{C}$-valued functions on G. This is a $\mathbf{C}$-vector space of functions. Every $f \in L(G)$ can be expressed as a linear combination of the delta-functions δ_g from (3.3):

$$f = \sum_{g\in G} f(g)\delta_g. \tag{4.1}$$

Indeed, evaluate both sides at each $x \in G$ and we get the same value. The functions δ_g span $L(G)$ by (4.1) and they are linearly independent: if $\sum_g a_g\delta_g = 0$ then evaluating the sum at $x \in G$ shows $a_x = 0$. Thus the functions δ_g are a basis of $L(G)$, so $\dim L(G) = |G|$.

The next theorem is the first step leading to an expression for each δ_g as a linear combination of characters of G, which will lead to a Fourier series expansion of f. It is the first time we *add* character values.

Theorem 4.1. *Let G be a finite abelian group. Then*

$$\sum_{g\in G}\chi(g) = \begin{cases} |G|, & \text{if } \chi = \mathbf{1}_G, \\ 0, & \text{if } \chi \neq \mathbf{1}_G, \end{cases} \qquad \sum_{\chi\in\widehat{G}}\chi(g) = \begin{cases} |G|, & \text{if } g = 1, \\ 0, & \text{if } g \neq 1. \end{cases}$$

Proof. Let $S = \sum_{g\in G}\chi(g)$. If χ is trivial on G then $S = |G|$. If χ is not trivial on G, say $\chi(g_0) \neq 1$. Then $\chi(g_0)S = \sum_{g\in G}\chi(gg_0) = \sum_{g\in G}\chi(g) = S$, so $S = 0$.

The second formula in the theorem can be viewed as an instance of the first formula via Pontryagin duality: the second sum is a sum of the character "evaluate at g" over the group $\widehat{G}$, and this character on $\widehat{G}$ is nontrivial when $g \neq 1$ by Pontryagin duality. □

Theorem 4.1 says the sum of a nontrivial character over a group vanishes and the sum of all characters of a group evaluated at a nontrivial element vanishes, so the sum of the elements in each row and column of a character table of G is zero except the row for the trivial character and the column for the identity element. Check this in Table 1.

Corollary 4.2. *For characters χ_1 and χ_2 in $\widehat{G}$ and g_1 and g_2 in G,*

$$\sum_{g\in G}\chi_1(g)\overline{\chi}_2(g) = \begin{cases} |G|, & \text{if } \chi_1 = \chi_2, \\ 0, & \text{if } \chi_1 \neq \chi_2, \end{cases} \qquad \sum_{\chi\in\widehat{G}}\chi(g_1)\overline{\chi}(g_2) = \begin{cases} |G|, & \text{if } g_1 = g_2, \\ 0, & \text{if } g_1 \neq g_2. \end{cases}$$

Proof. In the first equation of Theorem 4.1 let $\chi = \chi_1\overline{\chi}_2$. In the second equation of Theorem 4.1 let $g = g_1g_2^{-1}$. (Alternatively, after proving the first equation for all G we observe that the second equation is a special case of the first by Pontryagin duality.) □

The equations in Corollary 4.2 are called the *orthogonality relations*. They say that the character table of G has orthogonal rows and orthogonal columns when we define orthogonality of two n-tuples of complex numbers as vanishing of their Hermitian inner product in $\mathbf{C}^n$: $\langle (z_1, \ldots, z_n), (w_1, \ldots, w_n)\rangle := \sum_{k=1}^n z_k\overline{w_k}$.

By the second equation in Corollary 4.2 we can express the delta-functions in terms of characters:

$$\sum_{\chi\in\widehat{G}} \chi(g)\overline{\chi}(x) = |G|\delta_g(x) \Longrightarrow \delta_g(x) = \frac{1}{|G|}\sum_{\chi\in\widehat{G}} \overline{\chi}(g)\chi(x).$$

Substituting this formula for δ_g into (4.1) gives

$$\begin{aligned} f(x) &= \sum_{g\in G} f(g)\left(\frac{1}{|G|}\sum_{\chi\in\widehat{G}} \overline{\chi}(g)\chi(x)\right) \\ &= \sum_{\chi\in\widehat{G}}\sum_{g\in G} \frac{1}{|G|} f(g)\overline{\chi}(g)\chi(x) \\ &= \sum_{\chi\in\widehat{G}} c_\chi \chi(x), \end{aligned} \tag{4.2}$$

where

$$c_\chi = \frac{1}{|G|}\sum_{g\in G} f(g)\overline{\chi}(g). \tag{4.3}$$

The expansion (4.2) is the Fourier series for f.

Equation (4.3) is similar to the formula for the coefficient c_n of e^{inx} in (2.1): an integral over $[0, 2\pi]$ divided by 2π is replaced by a sum over G divided by $|G|$ and $f(x)e^{-inx}$ is replaced by $f(g)\overline{\chi}(g)$. The number e^{-inx} is the conjugate of e^{inx}, which is also the relation between $\overline{\chi}(g)$ and $\chi(g)$. Equation (4.2) shows $\widehat{G}$ is a spanning set for $L(G)$. Since $|\widehat{G}| = |G| = \dim L(G)$, $\widehat{G}$ is a basis for $L(G)$.

Definition 4.3. Let G be a finite abelian group. If $f \in L(G)$ then its *Fourier transform* is the function $\widehat{f} \in L(\widehat{G})$ given by

$$\widehat{f}(\chi) = \sum_{g\in G} f(g)\overline{\chi}(g).$$

By (4.2) and (4.3),

$$f(x) = \frac{1}{|G|}\sum_{\chi\in\widehat{G}} \widehat{f}(\chi)\chi(x). \tag{4.4}$$

Equation (4.4) is called the *Fourier inversion formula* since it tells us how to recover f from its Fourier transform.

Remark 4.4. Classically the Fourier transform of a function $\mathbf{R} \to \mathbf{C}$ is another function $\mathbf{R} \to \mathbf{C}$. The finite Fourier transform, however, is defined on the dual group instead of on the original group. We can also interpret the classical Fourier transform to be a function of characters. For $y \in \mathbf{R}$ let $\chi_y(x) = e^{ixy}$. Then $\chi_y \colon \mathbf{R} \to S^1$ is a character and $\widehat{f}(y)$ could be viewed as $\widehat{f}(\chi_y) = \int_{\mathbf{R}} f(x)\overline{\chi}_y(x)\,dx$, so $\widehat{f}$ is a function of characters rather than of numbers.

Example 4.5. Let $f = \delta_g$. Then $\widehat{f}(\chi) = \overline{\chi}(g) = \chi(g^{-1})$.

Let's look at Fourier transforms for functions on a cyclic group. By writing a cyclic group in the form $\mathbf{Z}/(m)$, we can make an isomorphism with the dual group explicit: every character of $\mathbf{Z}/(m)$ has the form $\chi_k\colon j \mapsto e^{2\pi ijk/m}$ for a unique $k \in \mathbf{Z}/(m)$ (Exercise 3.1). The Fourier transform of a function $f\colon \mathbf{Z}/(m) \to \mathbf{C}$ can be viewed as a function on $\mathbf{Z}/(m)$:

$$\widehat{f}(k) := \sum_{j \in \mathbf{Z}/(m)} f(j)\overline{\chi_k}(j) = \sum_{j \in \mathbf{Z}/(m)} f(j)e^{-2\pi ijk/m}. \tag{4.5}$$

This is like viewing the Fourier transform of a function on $\mathbf{R}$ as a function of $\mathbf{R}$.

Example 4.6. Let $f\colon \mathbf{Z}/(8) \to \mathbf{C}$ have the periodic values 5, 3, 1, and 1. Both f and its Fourier transform are in Table 3. This f has frequency 2 (its period repeats twice) and the Fourier transform vanishes except at 0, 2, 4, and 6, which are multiples of the frequency.

n	0	1	2	3	4	5	6	7
$f(n)$	5	3	1	1	5	3	1	1
$\widehat{f}(n)$	20	0	$8+4i$	0	4	0	$8-4i$	0

TABLE 3.

Example 4.7. Consider a function $f\colon \mathbf{Z}/(45) \to \mathbf{C}$ with the four successive repeating values 1, 8, 19, 17 starting with $f(0) = 1$. It is not a periodic function on $\mathbf{Z}/(45)$ since 4 does not divide 45, but the sequence 1, 8, 19, 17 repeats nearly 11 times. (The value of $f(44)$ is 1.) A calculation of $|\widehat{f}(n)|$, the *absolute value* of the Fourier transform of f, reveals sharp peaks at $n = 0, 11, 22, 23$, and 34. See a plot of $|\widehat{f}(n)|$ below.

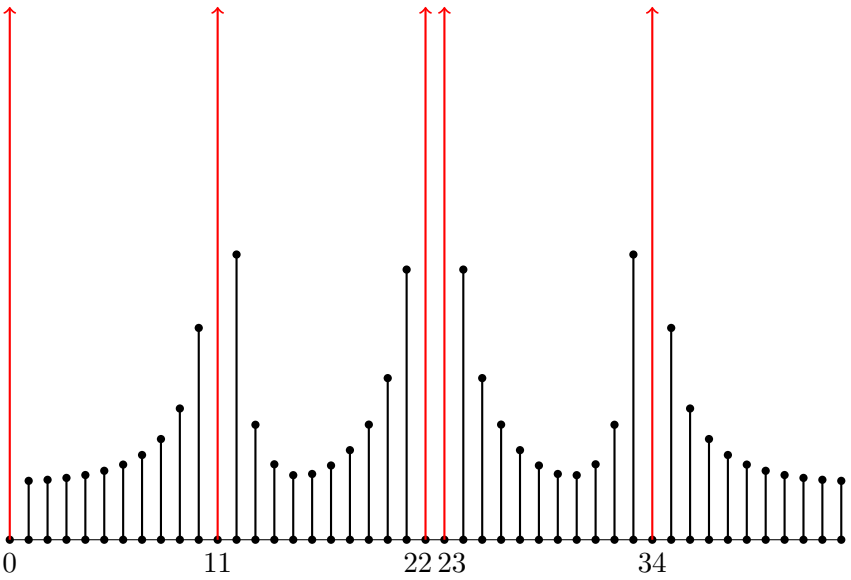

The red peaks are cut off because the lowest red bar would be around three times as tall as the highest black bar. Peaks in $|\widehat{f}(n)|$ occur approximately at multiples of the approximate frequency!

As Example 4.6 suggests, the Fourier transform of a periodic function on $\mathbf{Z}/(m)$ knows the frequency of the original function by the positions where the Fourier transform has nonzero values (Exercise 4.2). For *nearly* periodic functions on $\mathbf{Z}/(m)$, the approximate frequency is reflected in where the Fourier transform takes on its largest values. This idea is used in Shor's quantum algorithm for integer factorization [2], [3, Chap. 17].

Exercises.

1. Let $f\colon \mathbf{Z}/(8) \to \mathbf{C}$ take the four values a, b, c, and d twice in this order. Compute $\widehat{f}(n)$ explicitly and determine some values for a, b, c, and d such that $\widehat{f}(n)$ is nonzero for $n = 0, 2$, and 6, but $\widehat{f}(4) = 0$.
2. Let H be a subgroup of a finite abelian group G.
   (a) Suppose $f\colon G \to \mathbf{C}$ is constant on H-cosets (it is H-periodic). For $\chi \in \widehat{G}$ with $\chi \notin H^\perp$, show $\widehat{f}(\chi) = 0$. Thus the Fourier transform of an H-periodic function on G is supported on $H^\perp$.
   (b) If $f\colon \mathbf{Z}/(m) \to \mathbf{C}$ has period d where $d \mid m$, show $\widehat{f}\colon \mathbf{Z}/(m) \to \mathbf{C}$ is supported on the multiples of m/d. (See Example 4.6.)
3. Let $f\colon G \to \mathbf{C}$.
   a) Show $f(g) \in \mathbf{R}$ for all g if and only if $\overline{\widehat{f}(\chi)} = \widehat{f}(\overline{\chi})$ for all χ.
   b) Show $\widehat{f}(\chi) \in \mathbf{R}$ for all χ if and only if $\overline{f(g)} = f(g^{-1})$ for all g.
4. Let G be a finite abelian group and H be a subgroup. For a function $f\colon G \to \mathbf{C}$, Poisson summation on G says
$$\frac{1}{|H|}\sum_{h\in H} f(h) = \frac{1}{|G|}\sum_{\chi\in H^\perp} \widehat{f}(\chi),$$
where $H^\perp$ is as in Exercise 3.6. Prove this formula in two ways:
   a) Copy the classical proof sketched in Section 2 (start with the function $F(x) = \sum_{h\in H} f(xh)$, which is H-periodic so it defines a function on G/H) to obtain
$$\frac{1}{|H|}\sum_{h\in H} f(xh) = \frac{1}{|G|}\sum_{\chi\in H^\perp} \widehat{f}(\chi)\chi(x) \tag{4.6}$$
   for all $x \in G$ and then set $x = 1$.
   b) By linearity in f of both sides of the desired identity, verify Poisson summation directly on the delta-functions of G. (Corollary 3.7 and Example 4.5 will be useful.)

5. Structure of finite abelian groups

We will now put characters to work by using them to prove each finite abelian group is a direct product of cyclic groups. This result was used in the proof of Theorem 3.12, that $G \cong \widehat{G}$, but that work will not be used here, so no circular reasoning occurs.

The following theorem shows that every cyclic subgroup of maximal size in a finite abelian group can be split off as a direct factor. Characters get used in an essential way in the proof.

Theorem 5.1. *Let G be a finite abelian group and let $g \in G$ have maximal order in G. There is a subgroup $H \subset G$ such that $G \cong H \times \langle g\rangle$.*

Proof. Let n be the order of g. The subgroup $\langle g\rangle$ of G is cyclic of order n. In S^1 there is a cyclic subgroup of order n, namely μ_n. Since cyclic groups of the same order are

isomorphic, there is an isomorphism $\langle g\rangle \to \mu_n$, so g is mapped to a root of unity of order n. This isomorphism can be viewed as a character of $\langle g\rangle$. Extend this to a character of G (Theorem 3.3), so we have a character $\chi\colon G \to S^1$ such that $\chi(g)$ has order n. The image $\chi(G)$ contains μ_n, and it turns out to be no larger.

Claim: $\chi(G) = \mu_n$.

Since $\chi(G)$ is a finite subgroup of S^1, it is cyclic (all finite subgroups of S^1 are cyclic). Therefore $\chi(G) = \langle \chi(\gamma)\rangle$ for some $\gamma \in G$. Since $\chi(G)$ contains μ_n, $\chi(G) = \mu_{nn'}$ where $n' \geq 1$. Thus $\chi(\gamma)$ has order nn'. Let γ have order d in G, so $\gamma^d = 1$ in G and thus $\chi(\gamma)^d = 1$ in S^1. That implies $nn' \mid d$, so $nn' \leq d$. Since n is the maximal order of the elements in G, $d \leq n$. The relations $nn' \leq d$ and $d \leq n$ imply $n' = 1$, so $\chi(G) = \mu_n$. This proves the claim.

Set $H = \ker \chi$. Then $H \cap \langle g\rangle = \{1\}$ since χ is one-to-one on $\langle g\rangle$ by construction. For each $x \in G$, $\chi(x) \in \chi(G) = \mu_n = \chi(\langle g\rangle)$, so $\chi(x) = \chi(g^j)$ for some j. Therefore $h := xg^{-j}$ is in H and $x = hg^j$. This proves that the multiplication map $H \times \langle g\rangle \to G$ where $(h, g^j) \mapsto hg^j$ is surjective. It is a homomorphism and its kernel is trivial, so $G \cong H \times \langle g\rangle$. □

Theorem 5.2. *Every finite abelian group G is isomorphic to a product of cyclic groups:*

$$G \cong \mathbf{Z}/(n_1) \times \mathbf{Z}/(n_2) \times \cdots \times \mathbf{Z}/(n_k).$$

Proof. Induct on $|G|$. The result is clear when $|G| = 1$. When $|G| > 1$, let n be the maximal order of the elements of G, so $G \cong H \times \mathbf{Z}/(n)$ by Theorem 5.1. Since $|H| < |G|$, by induction H is isomorphic to a direct product of cyclic groups, so G is also isomorphic to a direct product of cyclic groups. □

Theorem 5.2 can be refined: G is a direct product of cyclic groups with the extra feature that $n_1 \mid n_2 \mid \cdots \mid n_k$. To prove this, use the fact that the order of each element of a finite abelian group G divides the maximal order of the elements of G.

Exercises.

1. What is the structure (as a direct product of cyclic groups) of the finite abelian groups whose nontrivial characters all have order 2?
2. Mimic the proof of Theorem 5.1 to decompose $(\mathbf{Z}/(20))^\times$ (of size 8) and $(\mathbf{Z}/(45))^\times$ (of size 24) into a direct product of cyclic groups.
3. Show by an explicit counterexample that the following is false: if two subgroups H and K of a finite abelian group G are isomorphic then there is an automorphism of G that restricts to an isomorphism from H to K.
4. For a finite abelian group G, show the maximum order of the elements of G and the number $|G|$ have the same prime factors. (Hint: If g has order n and there is an element h of prime order p where $p \nmid n$, what is the order of gh?)

   This is false in general for nonabelian G, as shown in the table below where $g(n)$ (called Landau's function) is the maximal order of the elements of S_n. For $n \geq 3$ in the table, some prime factor of $n!$ does not divide $g(n)$.

n	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
$g(n)$	1	2	3	4	6	6	12	15	20	30	30	60	60	84	105

TABLE 4. Maximal order of elements of S_n

5. Let G be a finite abelian group and F be a field containing a full set of $|G|$th roots of unity. (So $x^{|G|} = 1$ has $|G|$ solutions in F.) Define characters of G to be group homomorphisms $\chi\colon G \to F^\times$ and write the set of all such characters as $\widehat{G}$.
   a) Construct a character table for $\mathbf{Z}/(4)$ and $(\mathbf{Z}/(2))^2$ when F is the field $\mathbf{Z}/(5)$.
   b) Prove every lemma, theorem, and corollary from Section 3 for the new meaning of $\widehat{G}$. There is no longer complex conjugation on character values, but the inverse of χ is still the function $g \mapsto \chi(g^{-1}) = \chi(g)^{-1}$. (Hint: For each d dividing $|G|$, $x^d = 1$ has d distinct solutions in $F^\times$, which form a cyclic group.)
   c) Prove Theorem 4.1 and Corollary 4.2 for F-valued characters of G.
   d) Set $L(G, F)$ to be the functions $G \to F$. This is an F-vector space in the same way that $L(G)$ is a complex vector space. For each function $f \in L(G, F)$, define its Fourier transform $\widehat{f} \in L(\widehat{G}, F)$ by $\widehat{f}(\chi) = \sum_{g\in G} f(g)\chi(g^{-1})$. Prove the Fourier inversion formula and Plancherel's theorem in this context. (Note: If the field F has characteristic p then $1/|G|$ in the Fourier inversion formula makes sense in F since p doesn't divide $|G|$ – why?)
   e) Check everything you have done goes through if the assumption that $x^{|G|} = 1$ has a full set of solutions in F is weakened to $x^m = 1$ having a full set of solutions in F, where m is the maximal order of the elements of G. For example, if $G = (\mathbf{Z}/(2))^d$ then $m = 2$ and we can use $F = \mathbf{Z}/(3)$.

6. Existence of solutions to a Mordell equation mod p

For $k \in \mathbf{Z}$, an equation of the form $y^2 = x^3 + k$ is called a *Mordell equation*. When $k \neq 0$, it is a hard theorem that such an equation has only finitely many integral solutions (x, y), which could include having no integral solutions.[2] For example, the integral solutions of $y^2 = x^3 - 4$ are $(2, \pm 2)$ and $(5, \pm 11)$, and the equation $y^2 = x^3 - 5$ has no integral solutions. Using characters, we will show the congruence $y^2 \equiv x^3 + k \bmod p$ modulo a prime p always has a solution.

For a character χ on $(\mathbf{Z}/(p))^\times$, extend χ to $\mathbf{Z}/(p)$ by setting $\chi(0) = 0$. Then $\chi(ab) = \chi(a)\chi(b)$ for all $a, b \in \mathbf{Z}/(p)$.

Lemma 6.1. *If p is prime and $d \mid (p-1)$, there is a character χ of $(\mathbf{Z}/(p))^\times$ with order d, and for each $a \in \mathbf{Z}/p\mathbf{Z}$,*

$$|\{x \in \mathbf{Z}/(p) : x^d \equiv a \bmod p\}| = 1 + \chi(a) + \chi(a)^2 + \cdots + \chi(a)^{d-1}.$$

Proof. The group $(\mathbf{Z}/(p))^\times$ is cyclic of order $p-1$,[3] so $|\{x \in \mathbf{Z}/(p) : x^d \equiv 1 \bmod p\}| = d$ since $d \mid (p-1)$. Thus

$$|\{x \in \mathbf{Z}/(p) : x^d \equiv a \bmod p\}| = \begin{cases} d, & \text{if } a \not\equiv 0 \bmod p \text{ and } a \text{ is a } d\text{th power mod } p, \\ 1, & \text{if } a \equiv 0 \bmod p, \\ 0, & \text{if } a \text{ is not a } d\text{th power mod } p. \end{cases} \tag{6.1}$$

The character group of $(\mathbf{Z}/(p))^\times$ is cyclic of order $p-1$ by Theorem 3.10, so $(\mathbf{Z}/(p))^\times$ has a character of order d. Call such a character χ. For each $a \in \mathbf{Z}/(p)$, we'll show that $1 + \chi(a) + \chi(a)^2 + \cdots + \chi(a)^{d-1}$ has the same values as in (6.1).

[2] When $k = 0$, the equation is $y^2 = x^3$ and has infinitely many integral solutions $(x, y) = (a^2, a^3)$ for $a \in \mathbf{Z}$.
[3] See https://kconrad.math.uconn.edu/blurbs/grouptheory/cyclicmodp.pdf for many proofs of this.

- If a is a nonzero dth power mod p, say $a \equiv b^d \bmod p$, then $\chi(a) = \chi(b^d) = \chi(b)^d = 1$ since χ^d is identically 1 on $(\mathbf{Z}/(p))^\times$, so $1 + \chi(a) + \chi(a)^2 + \cdots + \chi(a)^{d-1} = d$.
- If $a \equiv 0 \bmod p$ then $\chi(a) = 0$, so $1 + \chi(a) + \chi(a)^2 + \cdots + \chi(a)^{d-1} = 1$.
- Lastly, if a is in $(\mathbf{Z}/(p))^\times$ and is not a dth power, we'll show $\chi(a) \neq 1$, so by summing a finite geometric series,

$$1 + \chi(a) + \chi(a)^2 + \cdots + \chi(a)^{d-1} = \frac{\chi(a)^d - 1}{\chi(a) - 1} = \frac{1-1}{\chi(a)-1} = 0,$$

which would complete the proof.

Let g be a generator of $(\mathbf{Z}/(p))^\times$ and write $a = g^k$ for $k \in \mathbf{Z}$.

Step 1: $\chi(g)$ has order d. Since χ has order d, $\chi(g)^d = 1$, so $\chi(g)$ has order dividing d. Since $\chi((\mathbf{Z}/(p))^\times) = \langle \chi(g) \rangle$, if $\chi(g)$ has order less than d then χ as a character has order less than d, which χ doesn't. Thus $\chi(g)$ has order d.

Step 2: $\chi(a) \neq 1$. If $\chi(a) = 1$ then $\chi(g)^k = 1$, so $d \mid k$ by Step 1. Then $a = g^k$ is a dth power, which isn't the case, so $\chi(a) \neq 1$. □

Theorem 6.2. *For each prime p and $k \in \mathbf{Z}$, the congruence $y^2 \equiv x^3 + k \bmod p$ has at least two solutions (x, y) in $\mathbf{Z}/(p)$.*

Proof. We'll consider separately the cases $3 \nmid (p-1)$ and $3 \mid (p-1)$.

Case 1: $3 \nmid (p-1)$.

Since $(3, p-1) = 1$, cubing is a bijection $\mathbf{Z}/(p) \to \mathbf{Z}/(p)$ (on $(\mathbf{Z}/(p))^\times$ it is injective and thus surjective since the group is finite), so for each $y \in \mathbf{Z}/(p)$ there is a unique $x \in \mathbf{Z}/(p)$ such that $y^2 - k \equiv x^3 \bmod p$. Thus the number of solutions of $y^2 \equiv x^3 + k \bmod p$ is p, and $p \geq 2$.

Case 2: $3 \mid (p-1)$.

If $k \equiv 0 \bmod p$, then the congruence is $y^2 \equiv x^3 \bmod p$, which has the p solutions (a^2, a^3) for $a \in \mathbf{Z}/(p)$ (and in fact no further solutions mod p). So now we can assume $k \not\equiv 0 \bmod p$.

Since $3 \mid (p-1)$, there is a character χ on $(\mathbf{Z}/(p))^\times$ with order 3, and the inverse χ^2 of χ is the complex conjugate $\overline{\chi}$. Since p is odd, there is a quadratic character ψ on $(\mathbf{Z}/(p))^\times$ (it's the Legendre symbol). To count solutions to $y^2 \equiv x^3 + k \bmod p$ we will count solutions (a, b) to the simpler equation $b \equiv a + k \bmod p$ and then count how often a is a cube mod p and b is a square mod p.

By Lemma 6.1, the number of ways a is a cube mod p is $1 + \chi(a) + \chi(a)^2 = 1 + \chi(a) + \overline{\chi}(a)$, and the number of ways b is a square mod p is $1 + \psi(b)$. Let $N_p(k)$ be the number of mod p solutions to $y^2 \equiv x^3 + k \bmod p$, so

$$N_p(k) = \sum_{(a,b)} (1 + \chi(a) + \overline{\chi}(a))(1 + \psi(b)),$$

where we sum over all (a, b) mod p for which $b \equiv a + k \bmod p$ (either a or b determines the other mod p). Expanding out the product, we get a sum of 6 terms over all the pairs (a, b) where $b \equiv a + k \bmod p$:

$$N_p(k) = \sum_{(a,b)} (1 + \chi(a) + \overline{\chi}(a) + \psi(b) + \chi(a)\psi(b) + \overline{\chi}(a)\psi(b)).$$

Split this up into 6 sums. The first sum is p since the number of possible (a, b) is p (both a and b determine each other mod p and each is free to take on any value). The second, third, and fourth sums are 0 since the sum of a nontrivial multiplicative character over $\mathbf{Z}/(p)$ is 0 (a term where $a = 0$ or $b = 0$ can be dropped since $\chi(0) = 0$ and $\psi(0) = 0$).

We're left with the sums of $\chi(a)\psi(b)$ and $\overline{\chi}(a)\psi(b)$, and at this point let's write b directly in terms of a (and k) so we can write the sums as running over all a mod p:

$$N_p(k) = p + \sum_a \chi(a)\psi(a+k) + \sum_a \overline{\chi}(a)\psi(a+k).$$

Since $k \not\equiv 0 \bmod p$, we can make the change of variables $a \mapsto ka$ in both sums and pull out the character values at k:

$$N_p(k) = p + \chi(k)\psi(k)\sum_a \chi(a)\psi(a+1) + \overline{\chi}(k)\psi(k)\sum_a \overline{\chi}(a)\psi(a+1).$$

Replace a mod p with $-a$ mod p in the sums:

$$N_p(k) = p + \chi(-k)\psi(k)\sum_a \chi(a)\psi(1-a) + \overline{\chi}(-k)\psi(k)\sum_a \overline{\chi}(a)\psi(1-a).$$

Set $S = \sum_a \chi(a)\psi(1-a)$, so $\overline{S} = \sum_a \overline{\chi}(a)\psi(1-a)$ since ψ-values are ± 1, and

$$N_p(k) = p + \chi(-k)\psi(k)S + \overline{\chi}(-k)\psi(k)\overline{S} = p + 2\mathrm{Re}(\chi(-k)\psi(k)S).$$

For each complex number z, $|\mathrm{Re}(z)| \leq \sqrt{|z|}$, so $|N_p(k) - p| \leq 2\sqrt{|S|}$. Since χ and ψ are nontrivial multiplicative characters mod p and $\chi\psi$ is nontrivial, $|S| = \sqrt{p}$.[4] Thus $|N_p(k) - p| \leq 2\sqrt{p}$, so $N_p(k) \geq p - 2\sqrt{p}$. The function $f(t) = t - 2\sqrt{t}$ is increasing for $t > 1$, the least prime $p \equiv 1 \bmod 3$ is 7, and $7 - 2\sqrt{7} \approx 1.7$, so $N_p(k) \geq 2$ when $3 \mid (p-1)$. □

Exercises.

1. For prime p and $n \in \mathbf{Z}^+$, set $d = (n, p-1)$. For $r \in \mathbf{Z}$, let $\varphi_r \colon (\mathbf{Z}/(p))^\times \to (\mathbf{Z}/(p))^\times$ by $\varphi_r(x) = x^r$.
   a) Show φ_n and φ_d have the same image and kernel. (Hint: d is a $\mathbf{Z}$-linear combination of n and $p-1$.)
   b) For nonzero a in $\mathbf{Z}/(p)$, use (a) to show the equations $x^n = a$ and $x^d = a$ have the same number of solutions in $\mathbf{Z}/(p)$.
   c) Find all solutions of $x^4 = 3$ and of $x^2 = 3$ in $\mathbf{Z}/(11)$. (There are two solutions in each case.)
   d) For nonzero a in $\mathbf{Z}/(p)$, show $x^n + y^n = a$ and $x^d + y^d = a$ have the same number of solutions in $\mathbf{Z}/(p)$.
2. Let p be prime, a be nonzero in $\mathbf{Z}/(p)$, and d be a positive factor of $p-1$. We want to estimate the number of solutions of $x^d + y^d = a$ in $\mathbf{Z}/(p)$.
   a) For a polynomial $f(x)$ with coefficients in $\mathbf{Z}/(p)$, let $N(f(x) = a)$ be the number of solutions of $f(x) = a$ in $\mathbf{Z}/(p)$. By Lemma 6.1, $(\mathbf{Z}/(p))^\times$ has a character χ of order d and we set $\chi(0) = 0$. Show

$$\begin{aligned} N(x^d + y^d = a) &= \sum_{\substack{b,c \in \mathbf{Z}/(p) \\ b+c=a}} N(x^d = b)N(y^d = c) \\ &= \sum_{b \in \mathbf{Z}/(p)} \left(1 + \sum_{i=1}^{d-1} \chi(b)^i\right)\left(1 + \sum_{j=1}^{d-1} \chi(a-b)^j\right). \end{aligned}$$

[4] See Corollary 2.4 in https://kconrad.math.uconn.edu/blurbs/gradnumthy/Gauss-Jacobi-sums.pdf.

b) Expand the product in (a) and rearrange terms to show

$$N(x^d + y^d = a) = p + \sum_{1 \le i,j \le d-1} \sum_{b \in \mathbf{Z}/(p)} \chi(b)^i \chi(a-b)^j$$
$$= p + \sum_{1 \le i,j \le d-1} \chi(a)^{i+j} \sum_{b \in \mathbf{Z}/(p)} \chi(b)^i \chi(1-b)^j.$$

c) For characters ψ and ψ' on $(\mathbf{Z}/(p))^\times$, set $J(\psi, \psi') = \sum_{b \in \mathbf{Z}/(p)} \psi(b)\psi'(1-b)$ (it is called a Jacobi sum), so by (b),

$$N(x^d + y^d = a) = p + \sum_{1 \le i,j \le d-1} \chi(a)^{i+j} J(\chi^i, \chi^j).$$

For a nontrivial character ψ on $(\mathbf{Z}/(p))^\times$, $J(\psi, \overline{\psi}) = -\psi(-1)$.[5] Use that to show

$$N(x^d + y^d = a) = p + 1 - N(x^d = -1) + \sum_{\substack{1 \le i,j \le d-1 \\ i+j \ne d}} \chi(a)^{i+j} \sum_{b \in \mathbf{Z}/(p)} \chi(b)^i \chi(1-b)^j.$$

d) When ψ, ψ', and $\psi\psi'$ are all nontrivial, $|J(\psi, \psi')| = \sqrt{p}$.[6] Use that and (c) to show

$$|N(x^d + y^d = a) - (p+1)| \le d + (d-1)(d-2)\sqrt{p}$$

e) Use part (d) and Exercise 6.1(d) to show for $n \in \mathbf{Z}^+$ and sufficiently large p (depending only on n) that each equation $x^n + y^n = a$ for $a \in (\mathbf{Z}/(p))^\times$ has a solution in $\mathbf{Z}/(p)$ where x and y are both nonzero.

References

[1] A. Terras, "Fourier Analysis on Finite Groups and Applications," Cambridge Univ. Press, Cambridge, 1999.

[2] P. Shor, *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*, http://arxiv.org/abs/quant-ph/9508027v2.

[3] W. Trappe and L. Washington, "Introduction to Cryptography with Coding Theory," Prentice-Hall, Upper Saddle River, NJ 2002.

[5] See Theorem 2.5 in https://kconrad.math.uconn.edu/blurbs/gradnumthy/Gauss-Jacobi-sums.pdf.

[6] See Corollary 2.4 in https://kconrad.math.uconn.edu/blurbs/gradnumthy/Gauss-Jacobi-sums.pdf.