L-functions.

TCC course, Oct–Dec 2008.

Kevin Buzzard

Chapter 0: Introduction.

http://tcc.maths.ox.ac.uk/syllabi/L-Functions.shtm

(or just google TCC Oxford) for books. Syllabus is mildly inaccurate (my fault): Tate didn't give a "new proof" of the functional equation of the Riemann zeta function—he conceptually explained an older one.

Basic definitions.

If $r > 0$ is real and s is complex, define $r^s := \exp(s.\log(r))$. Note that $|r^s| = r^{\mathrm{Re}(s)}$.

The Riemann zeta function is a holomorphic function of a variable s, whose definition for $\mathrm{Re}(s) > 1$ is

$$\zeta(s) := \sum_{n \geq 1} n^{-s}.$$

$$\zeta(s) := \sum_{n \geq 1} n^{-s}.$$

It's easily checked to converge to a holomorphic function (the convergence is absolute and locally uniform). The first big fact is that it has a meromorphic continuation to $s \in \mathbf{C}$ with a simple pole at $s = 1$ and no other poles. We'll see a proof of this in Lecture 2.

To explain the functional equation (relating $\zeta(s)$ to $\zeta(1-s)$ I'll need the Γ function

$$\Gamma(z) := \int_0^\infty t^{z-1} e^{-t} \mathrm{d}t$$

which converges (absolutely and locally uniformly) for $\mathrm{Re}(z) > 0$ and hence defines a holomorphic function there; we'll see that this also has a meromorphic continuation to $z \in \mathbf{C}$ but I want to state the functional equation before we get onto proofs.

The theorem (due to Riemann) is

$$\zeta(s) = 2^s \pi^{s-1} \sin(\pi s/2) \Gamma(1-s) \zeta(1-s).$$

This is an equality of meromorphic functions; one has to be a bit careful. For example if s is a positive even integer then the simple zero of $\sin(\pi s/2)$ cancels the simple pole of $\Gamma(1-s)$ on the RHS (when we get off the introduction and onto the details we'll see that Γ has some simple poles).

Here's a nicer (more symmetric) way of writing the functional equation: this is crucial. If we set

$$\xi(s) := \pi^{-s/2} \Gamma\left(\frac{s}{2}\right) \zeta(s)$$

$$\xi(s) := \pi^{-s/2}\Gamma\left(\frac{s}{2}\right)\zeta(s)$$

then the functional equation can be rewritten

$$\xi(s) = \xi(1-s)$$

We'll prove this before we do anything else because it's kind of important to us. And then we'll look at the proof and spend the rest of the course trying to generalise it to a conceptual proof of meromorphic continuation of a huge class of functions.

Remarks.

1) ξ is truly a "product of local factors"; one can check that $\zeta(s) = \prod_p (1-p^{-s})^{-1}$ (where the product is over all prime numbers; this is because every positive integer is uniquely the product of primes). The factor $(1-p^{-s})^{-1}$ is "the local factor at p". The stuff that was in ξ but not in ζ is the "local factor at ∞". We'll make this rigorous later. It was one of Tate's many insights that this could be formalised and massively generalised.

2) Why do we care about $\zeta(s)$, either for $\mathrm{Re}(s) > 1$, or for all $s \in \mathbf{C}$? It's a well-observed phenomenon that the zeta function (and its generalisations) encode arithmetic information, especially where ζ doesn't converge. Indeed, the general idea is that given an arithmetic object, it could have a zeta function, which will converge for $\mathrm{Re}(s)$ sufficiently large, and then it might be a tough theorem (or, more likely, a profound open conjecture) that this zeta function has a meromorphic continuation to the complex numbers, and then "special values" of this function (i.e. its values at certain carefully-chosen points) might tell you information about the original arithmetic object.

Examples of this phenomenon: $\zeta(2) = \pi^2/6$ and $\zeta(4) = \pi^4/90$ and

$$\zeta(12) = 691\pi^{12}/638512875$$

(denominator is $3^6 5^3 7^2 11.13$) and $\zeta(-11) = 691/32760$ (denominator is $2^3 3^2.5.7.13$; numerator is prime). These numbers are related to Bernoulli numbers—for example $B_{12} = -691/2730$. All this was known to Euler (1700s), in some sense. Bernoulli numbers tell us information about unramified extensions of cyclotomic fields: so in some sense the zeta function really is telling us that the class number of $\mathbf{Q}(\mu_{691})$ is a multiple of 691.

The Riemann zeta function has a simple pole at 1 (with residue 1). Hence there are infinitely many primes! (think about the representation of $\zeta(s)$ as a product). Dirichlet's theorem (about 150 years ago) pushed this idea a lot further: mild generalisations of the zeta function plus their behaviour at $s = 1$ give his famous theorem that there are "infinitely many primes in an AP".

The Riemann Hypothesis is that all the zeros of the Riemann zeta function, other than those at $s = -2, -4, -6, \ldots$, lie on the line $\mathrm{Re}(s) = 1/2$. This is a deep open problem which, were it to be true, would have lots of applications (it and its generalisations to other zeta functions give you all sorts of results about the error term in the prime number theorem, or the smallest quadratic nonresidue mod p, and so on).

Generalisations of the Riemann zeta function: I've already mentioned Dirichlet's "L-functions": Also, a number field K has a zeta function $\zeta_K(s)$, with $\zeta_{\mathbf{Q}}(s)$ being the classical Riemann zeta function. The function $\zeta_K(s)$ has a simple pole at $s = 1$ and the residue is

$$\lim_{s \to 1} (s-1)\zeta_K(s) = \frac{2^{r_1} \cdot (2\pi)^{r_2} \cdot h_K \cdot R_K}{w_K \cdot \sqrt{\mid D_K \mid}}$$

$$\lim_{s \to 1} (s-1)\zeta_K(s) = \frac{2^{r_1} \cdot (2\pi)^{r_2} \cdot h_K \cdot R_K}{w_K \cdot \sqrt{\mid D_K \mid}}$$

where, as usual, r_1 is the number of real embeddings $K \to \mathbf{R}$, r_2 is half the number of non-real embeddings $K \to \mathbf{C}$, h_K is the size of the class group of K, R_K is the regulator (this is to do with the logarithms of the fundamental units), w_K is the number of roots of unity in K and D_K is the discriminant of K.

Using the functional equation (this zeta function also has a functional equation) we can recast this statement as a statement about ζ_K near $s = 0$, and it turns out to say that $\zeta_K(s)$ has a zero of order $r_1 + r_2 - 1$ at $s = 0$ (the rank of the class group) and the power series expansion near $s = 0$ looks like

$$(-h_K.R_K/w_K)s^{r_1+r_2-1} + \ldots.$$

$$\zeta_K(s) = (-h_K.R_K/w_K)s^{r_1+r_2-1} + \ldots.$$

So in some sense the reason $\zeta(0) = -1/2$ is because the rational integers are a PID and the only units are the two roots of unity (and hence the regulator is 1).

Zeta functions hold profound arithmetic secrets. More general zeta functions are also called L-functions. Putting Dirichlet's ideas together with the generalisations to number fields gives the "correct" analogue and proof of Dirichlet's theorem for the integers of a number field.

Other things with zeta functions: automorphic forms, elliptic curves, algebraic varieties,. . . . Special values of L-functions and analogues of the class number formula above give profound conjectures. For example the Birch–Swinnerton-Dyer conjecture is just the analogue of the above theorem about $\zeta_K(s)$ near $s = 0$, but for the L-function of an elliptic curve.

Hecke proved Tate's theorem first; but Tate's proof was amenable to vast generalisations and has run and run.

Chapter 1: Meromorphic continuation and functional equation of the Riemann ζ function.

1.1 The Γ function.

Definition:

$$\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} \mathrm{d}t.$$

Converges for $\mathrm{Re}(z) > 0$ [integrand blowing up at zero if $\mathrm{Re}(z) < 1$ but not too badly: integral converges] to a holomorphic function.

Integrate by parts: for $\mathrm{Re}(z) > 0$ we have

$$\begin{aligned}\Gamma(z+1) &= \int_0^\infty t^z e^{-t} \mathrm{d}t \\ &= [-t^z e^{-t}]_0^\infty + \int_0^\infty z t^{z-1} e^{-t} \mathrm{d}t \\ &= z\Gamma(z).\end{aligned}$$

Hence for $\mathrm{Re}(z) > 0$ and $n \in \mathbf{Z}_{\geq 1}$ we have

$$\begin{aligned}\Gamma(z+n) &= (z+n-1)\Gamma(z+n-1) \\ &= (z+n-1)(z+n-2)\dots(z+1)z\Gamma(z)\end{aligned}$$

and hence $\Gamma(z) = \Gamma(z+n)/[(z+n-1)(z+n-2)\dots(z+1)(z)]$, and the right hand side is meromorphic for $\mathrm{Re}(z) > -n$, with (at worst) simple poles at $z = 0, -1, -2, \dots, 1-n$. So we can now regard Γ as a meromorphic function on the entire complex plane, satisfying $z\Gamma(z) = \Gamma(z+1)$.

Easy: $\Gamma(1) = 1$ (just compute the integral), and now it's an easy exercise from $z\Gamma(z) = \Gamma(z+1)$ to check that

- $\Gamma(n+1) = n!$ for $n \in \mathbf{Z}_{\geq 0}$

- $\Gamma(z)$ has a simple pole at $z = 0, -1, -2, \ldots$ and no other poles. (exercise: compute the residue at these poles).

To check the two versions of the functional equation that I gave in the first lecture are the same, one has to check

$$2^s \pi^{-1/2} \sin(\pi s/2) \Gamma(1-s) \Gamma(s/2) = \Gamma((1-s)/2).$$

but I won't use this because we'll never use the asymmetric functional equation. [It follows easily if you can prove

- Euler's reflection formula

$$\Gamma(1-z)\Gamma(z) = \pi / \sin(\pi z)$$

- and Legendre's duplication formula

$$\Gamma(z)\, \Gamma\left(z + \frac{1}{2}\right) = 2^{1-2z}\, \sqrt{\pi}\, \Gamma(2z).$$

]

1.2: Poisson summation.

Define

$$\theta(t) = \sum_{n \in \mathbf{Z}} e^{-\pi n^2 t^2}$$

(a function on the positive reals); it's easily checked to converge, and tends to one (very rapidly) as $t \to \infty$. Our goal here is to show the fundamental fact

$$\theta(1/t) = t\theta(t).$$

This is not at all obvious (to me)—for example $e^{16\pi} \sim 10^{20}$ so it's not surprising (looking at the definition) that

$$\theta(4) = 1.00000000000000000002958\ldots$$

but it is surprising (to me) that

$$\theta(1/4) = 4.00000000000000000000118322\ldots$$

(or equivalently, why, if $r = e^{-\pi/16} = 0.821724958\ldots$ then $r + r^4 + r^9 + r^{16} + r^{25} + r^{36} + r^{49} = 1.49999\ldots$.

The fundamental fact

$$\theta(1/t) = t\theta(t)$$

follows from the Poisson Summation formula, which follows from the general theory of Fourier series. Here's how a proof goes.

First let me remind you of a crucial integral:

$$\int_{-\infty}^{\infty} e^{-\pi x^2} \mathrm{d}x = 1$$

because if I denotes the integral then $I^2 = \int_{\mathbf{R}^2} e^{-\pi(x^2+y^2)} \mathrm{d}x\mathrm{d}y$ which is (recalling $\mathrm{d}x\mathrm{d}y = r\mathrm{d}r\mathrm{d}\theta$)

$$\int_{r\geq 0} \int_{0\leq\theta\leq 2\pi} e^{-\pi r^2} r\mathrm{d}r\mathrm{d}\theta$$

which is $2\pi[-e^{-\pi r^2}/2\pi]_0^\infty = 1$. As a consequence we deduce

$$\int_{-\infty}^{\infty} e^{-\pi y^2 + 2\pi i r y} \mathrm{d}y = e^{-\pi r^2} \qquad (*)$$

for $r > 0$ real (complete the square with $x = y - ir$ and use Cauchy's theorem). As another consequence we deduce $\Gamma(1/2) = \sqrt{\pi}$ (Exercise: comes straight from the definition after a simple substitution).

Now let's get back to θ. Fix $t > 0$ and define (for $x \in \mathbf{R}$) a function $f(x) = e^{-\pi t^2 x^2}$, and then define

$$F(x) = \sum_{n \in \mathbf{Z}} f(x+n)$$
$$= \sum_{n \in \mathbf{Z}} e^{-\pi t^2 (x+n)^2}.$$

Note that $F(0) = \theta(t)$. But note also that F is continuous and periodic with $F(x) = F(x+1)$ so by the theory of Fourier series [which we'll do in some generality later on, but let me just assume the classical theory now] we must have $F(x) = \sum_{m \in \mathbf{Z}} a_m e^{2\pi i m x}$ and we can compute

$$a_m = \int_0^1 F(x) e^{-2\pi i m x} \mathrm{d}x.$$

$$a_m = \int_0^1 F(x) e^{-2\pi i m x} \mathrm{d}x$$

and this comes out to be

$$\sum_{n \in \mathbf{Z}} \int_0^1 f(x+n) e^{2\pi i m x} \mathrm{d}x$$

$$= \sum_{n \in \mathbf{Z}} \int_0^1 f(x+n) e^{2\pi i m (x+n)} \mathrm{d}x$$

because changing x to $x+n$ changes things by $e^{2\pi i m n}$ which is 1. And now this is just

$$\int_{-\infty}^{\infty} f(x) e^{2\pi i m x} \mathrm{d}x$$

so we have proved that

$$a_m = \int_{-\infty}^{\infty} e^{-\pi t^2 x^2 + 2\pi i m x} \mathrm{d}x.$$

$$a_m = \int_{-\infty}^{\infty} e^{-\pi t^2 x^2 + 2\pi i m x} \mathrm{d}x.$$

Now setting $y = tx$ and $r = m/t$ we get

$$a_m = t^{-1} e^{-\pi m^2/t^2}$$

from $(*)$ above. So that's a_m and now (from the definitions)

$$\begin{aligned} \theta(t) &= F(0) \\ &= \sum_{m \in \mathbf{Z}} a_m \\ &= t^{-1}\theta(1/t) \end{aligned}$$

so we're done.

Corollary: $\theta(t) \sim 1/t$ for $t > 0$ small.

1.3: Meromorphic continuation of $\xi(s)$ and $\zeta(s)$.

Define, for $\mathrm{Re}(s) > 1$,

$$\xi(s) := \int_{t=0}^{\infty} (\theta(t) - 1) t^{s-1} \mathrm{d}t.$$

Note: this was not the definition of ξ we saw in the first lecture; but we'll prove it's the same. This function ξ is the Mellin Transform of $\theta(t) - 1$, and we'll now see that the relation between $\theta(t)$ and $\theta(1/t)$ translates into proof of meromorphic continuation and functional equation for $\xi(s)$.

If $\mathrm{Re}(s) > 1$ then this integral converges. Indeed the integral from 1 onwards is fine, because $\theta(t) - 1$ is decaying exponentially, and the integral from 0 to 1 is OK because $\theta(t)$ is like $1/t$ so we're just OK.

Now breaking up the integral at the point $t = 1$ we see

$$\xi(s) = \int_{t=1}^{\infty} (\theta(t)-1)t^{s-1}\mathrm{d}t + \int_{t=0}^{1} (\theta(t)-1)t^{s-1}\mathrm{d}t.$$

The first integral converges for all $s \in \mathbf{C}$, and using the fact that $\theta(t) = \theta(1/t)/t$ and subbing $u = 1/t$ we get that the second is

$$\int_{u=1}^{\infty} (u\theta(u) - 1)u^{-1-s}\mathrm{d}u$$

and we can break this up into two pieces as

$$\int_{u=1}^{\infty} \theta(u)u^{-s}\mathrm{d}u - \int_{u=1}^{\infty} u^{-1-s}\mathrm{d}s$$

(noting that both integrals converge in the region $\mathrm{Re}(s) > 1$). The second piece is just $-1/s$. Changing that θ back to $\theta - 1$ in the first piece by adding and subtracting $\int_1^{\infty} u^{-s}\mathrm{d}u = -1/(1-s)$, and putting everything together, gives us

$$\begin{aligned}\xi(s) = &\int_{t=1}^{\infty} (\theta(t) - 1)t^{s-1}\mathrm{d}t \\ &+ \int_{u=1}^{\infty} (\theta(u) - 1)u^{-s}\mathrm{d}u - 1/(1-s) - 1/s\end{aligned}$$

so

$$\xi(s) = \int_{t=1}^{\infty} (\theta(t) - 1)(t^{s-1} + t^{-s}) \mathrm{d}t \\ - 1/(1-s) - 1/s$$

Now that integral converges for all $s \in \mathbf{C}$ to a holomorphic function which is visibly invariant under $s \mapsto 1 - s$. We deduce that ξ has simple poles at $s = 0$ and $s = 1$ with residues -1 and $+1$ respectively, and no other poles, and satisfies $\xi(s) = \xi(1 - s)$.

So what's left is to check that $\xi(s)$ has got something to do with the zeta function! And we do this by now assuming $\mathrm{Re}(s) > 1$ again, and writing

$$\xi(s) = 2 \int_{t=0}^{\infty} \sum_{n \geq 1} e^{-\pi n^2 t^2} t^{s-1} \mathrm{d}t$$

and interchanging the sum and the integral, and observing that we can then do the inner integral: it's

$$\int_{t=0}^{\infty} e^{-\pi n^2 t^2} t^{s-1} \mathrm{d}t$$

$$\int_{t=0}^{\infty} e^{-\pi n^2 t^2} t^{s-1} \mathrm{d}t$$

and now setting $u = nt$ we get

$$n^{-s} \int_{u=0}^{\infty} e^{-\pi u^2} u^{s-1} \mathrm{d}u$$

and now setting $v = \pi u^2$ so $2\pi u \mathrm{d}u = \mathrm{d}v$ we get

$$n^{-s} \int_{v=0}^{\infty} e^{-v} (v/\pi)^{s/2-1} (2\pi)^{-1} \mathrm{d}v$$

which is

$$n^{-s} 2^{-1} \pi^{-s/2} \Gamma(s/2)$$

so $\xi(s) = \pi^{-s/2}\Gamma(s/2)\zeta(s)$ for $\mathrm{Re}(s) > 1$ and that was what we wanted!

Finally, let's think about poles. We saw $\xi(s)$ had simple poles at $s = 0$ and $s = 1$, with residues -1 and $+1$ respectively, and no other poles. So

$$\zeta(s) = \xi(s).\pi^{s/2}/\Gamma(s/2)$$

(an equation which now gives us the meromorphic continuation of the Riemann zeta function!) will have a simple pole at $s = 1$ with residue $\pi^{1/2}/\Gamma(1/2) = 1$, and will be holomorphic at $s = 0$ because Γ has a simple pole at $s = 0$. Furthermore, the only other poles of $\zeta(s)$ will come from zeros of the Γ function—but if the Γ function had a zero then $z\Gamma(z) = \Gamma(z + 1)$ implies it would have zeros with arbitrarily large real part and hence $\zeta(s)$ would have poles with arbitrarily large real part—but this is impossible because $\zeta(s)$ is holomorphic for $\mathrm{Re}(s) > 1$.

What Tate did was he managed to understand the above argument to such an extent that he could generalise it. Perhaps you can't see the wood from the trees at the minute, but somehow the ingredients are: clever definition of ξ, and two ways of evaluating it: one by "brute force" and one by viewing it as a Mellin transform of a theta function, breaking up the integral into two pieces, and using Poisson summation. This is the strategy that we shall generalise, once we have spent at least half of the course creating the necessary machinery.

Chapter 2: Local fields.

Let k be a field. A *norm* on k is $|.| : k \to \mathbf{R}$ with

(i) $|x| \geq 0$ with equality iff $x = 0$

(ii) $|xy| = |x||y|$

and some version of the triangle inequality, which varies from book to book. Let me use the following variant:

(iii) There's some constant $C \geq 1$ such that $|x| \leq 1$ implies $|1 + x| \leq C$.

We say that a pair $(k, |.|)$ consisting of a field k and a function $|.| : k \to \mathbf{R}$ satisfying the above axioms is a *normed field*.

(i) $|x| \geq 0$ with equality iff $x = 0$

(ii) $|xy| = |x||y|$

(iii) There's some constant $C \geq 1$ such that $|x| \leq 1$ implies $|1 + x| \leq C$.

There's now lots of things to do and no (for me) clear order in which to do them. We need comments about the axioms, basic properties deducible from the axioms, and elementary examples.

Let me first make a comment about the axioms: Why not the triangle inequality? Why not $|x + y| \leq |x| + |y|$ instead of (iii)? In fact (iii) is slightly weaker than the triangle inequality, as can be easily seen: if $|x + y| \leq |x| + |y|$ for all x and y then (iii) is satisfied with $C = 2$.

(i) $|x| \geq 0$ with equality iff $x = 0$

(ii) $|xy| = |x||y|$

(iii) There's some constant $C \geq 1$ such that $|x| \leq 1$ implies $|1 + x| \leq C$.

The "problem" with the triangle equality if you take a norm on a field which satisfies the triangle inequality, and then cube it, it might not satisfy the triangle inequality any more. On the other hand, one can easily check that $|.|$ is a norm in the sense above iff $|.|^r$ is for any $r > 0$ (easy exercise: replace C by C^r).

Definition We say that two norms $|.|$ and $|.|'$ on a field k are *equivalent* if there's some $r > 0$ such that $|x|^r = |x|'$ for all $x \in k$.

We only really care about norms up to equivalence.

(i) $|x| \geq 0$ with equality iff $x = 0$

(ii) $|xy| = |x||y|$

(iii) There's some constant $C \geq 1$ such that $|x| \leq 1$ implies $|1 + x| \leq C$.

Basic properties of a normed field: $|0| = 0$, and $|1| = |1|^2$ and hence $|1| = 1$ so $|-1|^2 = 1$ and hence $|-1| = 1$ and $|-a| = |a|$.

Examples: $k = \mathbf{R}$ (or any subfield, for example $\mathbf{Q}$), and $|x|$ is the usual norm: $|x| = x$ for $x \geq 0$ and $-x$ for $x < 0$.

Trivial example: the "trivial norm" on a field: $|0| = 0$ and $|x| = 1$ for all $x \neq 0$.

Back to the triangle inequality. I've already mentioned that if a function $|.|$ on a field k satisfies (i) and (ii) and $|x+y| \leq |x| + |y|$ for all x and y, then it clearly satisfies (iii) with $C = 2$.

Conversely,

Lemma. if $(k, |.|)$ is a normed field and if $|x| \leq 1$ implies $|1 + x| \leq 2$ (that is, if we can take $C = 2$ in the definition of the norm), then $|.|$ satisfies the triangle inequality.

I'll sketch a proof of this because the proof is slightly tricky and we use corollaries of this result quite a bit. Let me mention some corollaries first.

Corollary 1. Any norm is equivalent to a norm satisfying the triangle inequality.

Proof: $C^r \leq 2$ for some appropriate r.

Corollary 2. A norm defines a topology on k: if we say that a subset U of k is open iff for all $u \in U$ there's $\epsilon > 0$ such that $|u - v| < \epsilon$ implies $v \in U$, then the open sets satisfy the axioms for a topology.

Proof: equivalent norms define the same open sets, and if the norm satisfies the triangle inequality then $d(x, y) = |x - y|$ is a metric and the open sets for a metric form a topology.

OK, now onto the proof of the lemma.

Lemma. if $(k, |.|)$ is a normed field and if $|x| \leq 1$ implies $|1 + x| \leq 2$ (that is, if we can take $C = 2$ in the definition of the norm), then $|.|$ satisfies the triangle inequality.

Proof (sketch).

(a) The definition implies $|x+y| \leq 2 \max\{|x|, |y|\}$.

(b) Hence (induction) $|x_1 + x_2 + \ldots + x_{2^n}| \leq 2^n \max\{|x_1|, |x_2|, \ldots\}$

(c) Hence

$|x_1+x_2+\ldots+x_N| \leq 2N \max\{|x_1|, |x_2|, \ldots, |x_N|\}$

(choose n with $2^{n-1} < N \leq 2^n$ and use (b) with $x_{N+1} = \ldots = 0$).

(c)

$$|x_1+x_2+\ldots+x_N| \leq 2N \max\{|x_1|, |x_2|, \ldots, |x_N|\}$$

(d) Hence $|N| \leq 2N$ for all $N \in \mathbf{Z}_{\geq 0}$.

(e) Now use the binomial theorem, (c) and (d) to check that

$$|(x+y)^m| \leq 4(m+1)(|x|+|y|)^m$$

for all $m \in \mathbf{Z}_{\geq 1}$.

(f) Now let $m \to \infty$ and take mth roots to get the result.

There's a dichotomy: if we can take $C = 1$ in (iii) then $C = 1$ will also do for any equivalent norm. But if we need $C > 1$ in (iii) then by replacing $|.|$ with $|.|^N$ for some $N >> 0$ we can make C as large as we like (and equivalently as small as we like, subject to it being bigger than 1, by letting $N \to 0^+$).

Definition: a norm is *non-archimedean* if we can take $C = 1$ in (iii) above. A norm is *archimedean* if it's not non-archimedean. This definition is good on equivalence classes. Note that a norm is non-archimedean iff $|x + y| \leq \max\{|x|, |y|\}$ for all $x, y \in k$ (easy check). This is *much* stronger than the triangle inequality!

The usual norm on $\mathbf{R}$ is archimedean. The trivial norm is non-archimedean. Other examples: $k = \mathbf{C}$ and $|x + iy| = \sqrt{x^2 + y^2}$, or even $|x+iy| = x^2+y^2$: these are archimedean.

A less trivial example of a norm: $k = \mathbf{Q}$, choose a prime p, and define $|p| = p^{-1}$ (or indeed $|p| = r$ for any $0 < r < 1$) and $|q| = 1$ for any other prime q, and extend multiplicatively (and set $|0| = 0$). So we have

$$\left| p^n \cdot \frac{u}{v} \right| = p^{-n}$$

for $n \in \mathbf{Z}$ and u, v integers prime to p.

I claim that this is a non-archimedean norm (it's called the "p-adic norm" on $\mathbf{Q}$). This is such an important norm for us that I'll check the axioms.

Definition: $|0| = 0$ and

$$\left| p^n \cdot \frac{u}{v} \right| = p^{-n}$$

for $n \in \mathbf{Z}$ and u, v integers prime to p.

Check it's a norm: (i) and (ii) are obvious. So it suffices to check that for all x and y we have $|x + y| \leq \max\{|x|, |y|\}$; then (iii) will follow with $C = 1$.

Now $|x + y| \leq \max\{|x|, |y|\}$ is clear if any of x, y or $x + y$ is zero. In the general case we may assume $|x| \geq |y|$, so $x = p^n \frac{u}{v}$ and $y = p^m \frac{s}{t}$ with $n \leq m$, and we see that

$$\begin{aligned} x + y &= p^n \left(\frac{u}{v} + \frac{p^{m-n} s}{t} \right) \\ &= p^n \frac{ut + p^{m-n} sv}{vt} \\ &= p^{n'} \frac{u'}{v'} \end{aligned}$$

with $v' = vt$ and $n' \geq n$, so $|x + y| = p^{-n'} \leq p^{-n} = |x| = \max\{|x|, |y|\}$.

Easy exercise: check (by beefing up the proof) that in fact the p-adic norm on $\mathbf{Q}$ satisfies $|x+y| = \max\{|x|, |y|\}$ if $|x| \neq |y|$.

More fun: check that if $|.|$ is any non-arch norm on any field k then $|x| \neq |y|$ implies $|x+y| = \max\{|x|, |y|\}$. There's a one-line proof from the axioms which I sometimes struggle to find.

Natural generalisation of the p-adic norm: if K is any number field with integers R, and if P is a non-zero prime ideal of R, then there's a P-adic norm on K, defined by $|0| = 0$ and, for $0 \neq x \in K$, if we factor the fractional ideal (x) as $(x) = xR = P^e \cdot \prod_i P_i^{e_i}$ with the product finite and only involving prime ideals other than P, then we can define $|x| = r^e$ for any r with $0 < r < 1$; traditionally we take $r = 1/N(P)$ where $N(P)$ is the size of the finite field R/P. Again one checks that this is a norm, and indeed it's non-archimedean. These norms generalise the p-adic norm on $\mathbf{Q}$.

Note that if p factors into more than one prime in R, then there is more than one P-adic norm on K that induces a norm equivalent to the p-adic norm on $\mathbf{Q}$. For example, the $(2+i)$-adic norm on $\mathbf{Q}(i)$ is certainly not equivalent to the $(2-i)$-adic norm [because $|2+i| = 1/5$ for one of them and $|2+i| = 1$ for the other].

There are also natural generalisation of the usual archimedean norm on $\mathbf{Q}$ to a number field: if K is a number field then for any field homomorphism $K \to \mathbf{C}$ ($\mathbf{C}$ the complexes), the usual norm on $\mathbf{C}$ induces (by restriction) a norm on K. There is a subtlety here: if $\tau : K \to \mathbf{C}$ is a field homomorphism then $\overline{\tau}$, defined by $\overline{\tau}(x) = \overline{\tau(x)}$, is also a field homomorphism, and $\overline{\tau}$ may or may not be equal to τ, but τ and $\overline{\tau}$ induce the same norm on K, because $|\overline{z}| = |z|$ on $\mathbf{C}$. So in fact we're led to the following equivalence relation on field homomorphisms $\tau : K \to \mathbf{C}$ defined by: $\tau \sim \tau$ and $\tau \sim \overline{\tau}$, and nothing else. The equivalence classes are easily described: the maps $K \to \mathbf{R}$ each give one equivalence class (the standard notation is that there are r_1 of these) and the maps $K \to \mathbf{C}$ which don't land in $\mathbf{R}$ come in pairs $\{\tau, \overline{\tau}\}$ of equivalence classes: there are r_2 equivalence classes (and hence $2r_2$ embeddings). Let's stick with this notation throughout the course.

If $K = \mathbf{Q}(\alpha)$ and $P(X) \in \mathbf{Q}[X]$ is the minimal polynomial of α, then r_1 is the number of real roots of α, and r_2 is half the number of non-real roots. This shows that $r_1 + 2r_2$ is the degree of P. This leads us easily to a proof that $r_1 + 2r_2 = [K : \mathbf{Q}]$.

We won't logically need the following result so I won't prove it:

Theorem. If K is a number field, then any non-trivial norm on K is either equivalent to a P-adic norm for a unique P (this is iff it's non-archimedean), or equivalent to the valuation induced by an embedding $K \to \mathbf{C}$, for a unique equivalence class $\{\tau, \overline{\tau}\}$ of embeddings as above (this is iff it's archimedean).

The case $K = \mathbf{Q}$ is due to Ostrowski (an explicit elementary calculation), and the general case can be deduced from this case (after a little work).

Now here's a crucial property of norms. Given a normed field (and you can assume $C \leq 2$ if you like, because what we do here only depends on the equivalence class of the norm) there are obvious notions of a Cauchy sequence and a convergent sequence:

A sequence $(a_n)_{n\geq 1}$ is *Cauchy* if for all $\epsilon > 0$ there is $M > 0$ such that $m, n \geq M$ implies $|a_m - a_n| < \epsilon$.

A sequence $(a_n)_{n\geq 1}$ is *convergent* if there exists $b \in k$ such that for all $\epsilon > 0$ there's $M > 0$ with $n > M$ implies $|a_n - b| < \epsilon$.

Both notions only depend on the equivalence class of the norm. Every convergent sequence is Cauchy (easy). We say a normed field is *complete* if every Cauchy sequence is convergent.

Examples: $\mathbf{R}$ with the usual norm is complete. $\mathbf{Q}$ with the usual norm isn't.

In fact, if we're defining mathematics from the ground up we would build $\mathbf{R}$ by "completing" $\mathbf{Q}$. This is a process that works in much more generality!

Theorem. Given a normed field $(k, |.|)$ it has (up to unique isomorphism) a completion $(K, ||.||)$, by which I mean:

(i) a complete normed field $(K, ||.||)$, and

(ii) an inclusion $k \to K$ which preserves the norm (so if you're thinking of K as containing k then I'm just saying that for $x \in k$ we have $|x| = ||x||$),

such that

(iii) if we endow K with the topology induced by $||.||$, then the closure of k is K.

Note that (iii) is absolutely crucial: we want $\mathbf{R}$ to be the completion of $\mathbf{Q}$, whereas $\mathbf{Q} \subset \mathbf{C}$ satisfies (i) and (ii).

Proof. Let's do existence first. WLOG $|.|$ satisfies the triangle inequality. Let R denote the set of all Cauchy sequences in k; it's a ring with respect to pointwise addition and multiplication; the constant sequences give a map $k \to R$ of rings (so $1 \in R$ is the sequence $(1, 1, 1, \ldots)$).

One checks that if $(a_n)_{n \geq 1}$ is a Cauchy sequence in k then $(|a_n|)_{n \geq 1}$ is a Cauchy sequence of real numbers, so it's convergent. Say $\ell((a_n))$ is its limit.

Let I denote the ideal in R of sequences which tend to zero; this is easily checked to be an ideal. Let K denote the quotient R/I. It's easily checked that if $(a_n) - (b_n) \in I$ then $\ell((a_n)) = \ell((b_n))$. So ℓ induces a map $||.|| : K \to \mathbf{R}$.

The claim is that this works. Let's see why.

To check that K is a field one needs to check that I is maximal; this is because if (a_n) is Cauchy but doesn't tend to zero then $a_n \neq 0$ for all n sufficiently large and one can check that the sequence b_n defined by $b_n = 1/a_n$ (unless $a_n = 0$ in which case set $b_n = 59$) satisfies $(a_n)(b_n) - (1) \in I$; this is enough to prove that I is maximal.

It's clear that the obvious map $k \to K$ is a map of rings, and hence it's an injection because k is a field. The first two axioms for a norm are easily checked to be satisfied by $||.||$. We're assuming that $|.|$ satisfies the triangle inequality, and we deduce that $||.||$ does too. Hence (iii) is satisfied.

To check denseness of k in K we need to check that for any $(a_n) \in K$ and $\epsilon > 0$ we can find $a \in K$ with $||(a_n - a)|| < \epsilon$, but this is easy from the definition of Cauchy.

Finally, to check completeness of K we use the fact that Cauchy sequences in k converge in K, and that any Cauchy sequence in K can be approximated by a Cauchy sequence in k in a sufficiently sensible way to ensure that the limits coincide.

So we've done existence. For uniqueness we need to check that if $(K_1, ||.||_1)$ and $(K_2, ||.||_2)$ both work then there's a norm-preserving isomorphism of fields $K_1 = K_2$ which is the identity on k. The reason for this is that the map $k \to K_2$ can be extended to a map $K_1 \to K_2$ thus: write $\lambda \in K_1$ as the limit of a sequence in k; this sequence converges in K_2; send λ to this element. Now check that this gives a well-defined bijection—just follow your nose. $\square$

The only complete archimedean fields we care about in this course are the reals and the complexes (in fact Ostrowski proved that any field complete with respect to an archimedean norm was equivalent to $(\mathbf{R}, |.|)$ or $(\mathbf{C}, |.|)$, with $|.|$ denoting the usual norm, but we won't need this: see Chapter 3 of Cassels' "Local Fields").

So now we press on with the (arguably more subtle) theory of the structure of complete non-archimedean fields. It's easy to give examples of such things: for example let's define the p-*adic numbers* to be the completion of $\mathbf{Q}$ with respect to the p-adic norm—the usual notation for the p-adic numbers is $\mathbf{Q}_p$ (note: we haven't yet proved that $\mathbf{Q}$ with its p-adic norm isn't complete, or equivalently that $\mathbf{Q}_p \neq \mathbf{Q}$. But this will come out in the wash later).

It turns out that if k is a number field and P is a non-zero prime ideal of its integer ring, and if P contains the rational prime number p, then the completion of k with respect to the P-adic norm is naturally a finite extension of $\mathbf{Q}_p$ (I'll prove this later but it shouldn't surprise you because k is a finite extension of $\mathbf{Q}$), so in some sense the basic example of a complete non-archimedean field is the p-adic numbers, and the most general example we'll ever use in this course is a finite field extension of the p-adic numbers.

Before we start on the general structure theory, let me observe that the norm on $\mathbf{Q}_p$ or more generally k_P is "discrete", in the following sense: the P-adic norm on a number field k has the property that there's a real number $q > 1$ (the way I normalised it we have $q = N(P)$, the norm of P), such that *every* element of k had norm either equal to zero, or to an integer power of q. What does this imply about the norm on k_P?

Recall that in the definition of the completion of a field, the norm of a Cauchy sequence was the limit of the norms of the elements, and hence (easy calculation) we see that $|.| : k_P \to \mathbf{R}$ is also taking values in the set $\{0\} \cup \{\ldots, q^{-2}, q^{-1}, 1, q, q^2, q^3, \ldots\}$.

We say that a norm on a field K is *discrete* if there's some $\epsilon > 0$ such that $a \in K$ and $1 - \epsilon < |a| < 1 + \epsilon$ implies $|a| = 1$. In fact, because $|K^\times| := \{|a| : a \in K^\times\}$ is a subgroup of $\mathbf{R}_{>0}$ it's easy to check that if a norm on a field K is discrete then either $|K^\times| = \{1\}$ (the trivial norm) or there's some $q \in \mathbf{R}_{>1}$ with $|K^\times| = \{q^n : n \in \mathbf{Z}\}$.

The usual norm on the reals or complexes is of course not discrete, but the P-adic norm on a number field k is, and we've just seen that even the completion k_P of k with respect to this norm is a discretely-normed field. Don't get confused though—there are blah non-archimedean norms that aren't discrete—for example if k_n is the field $\mathbf{Q}(p^{1/2^n})$ (so we "keep square rooting p") and k_∞ is the union of the fields k_n (note that k_n is naturally a subfield of k_{n+1}) then the p-adic norm on $\mathbf{Q}_p$ extends to a non-discrete, non-archimedean norm on k_∞. Note that k_∞ isn't a number field though, it's an infinite extension of $\mathbf{Q}$.

Ok so let's let K now denote an arbitrary field equipped with a non-archimedean norm $|.|$ (so $|x+y| \leq \max\{|x|,|y|\}$ and so at least the triangle inequality holds). Let's set

$$R = \{x \in K : |x| \leq 1\}$$

and

$$I = \{x \in R : |x| < 1\}.$$

It's easy to check that R is a ring: the surprising part is that if $x, y \in R$ then $x+y \in R$, and this is because $|x+y| \leq \max\{|x|,|y|\} \leq 1$. We say R is the *integers* of K. Note that in the archimedean case this part already fails: the closed unit disc is not a subring of the complex numbers. The fact that R is a ring in the non-archimedean case is, perhaps initially at least, a little psychologically disturbing: for example it implies that the integers are bounded within K.

Once we have re-adjusted, it's easy to check from the axioms that I is an ideal of R, and in fact I is the unique maximal ideal of R because if $r \in R$ and $r \notin I$ then $|r| = 1$ so $r \neq 0$ and $s := 1/r \in K$ has $|s| = 1/|r| = 1$ so $s \in R$, and we see that r is a unit (exercise: this is enough).

We say that the field R/I is the *residue field* of K.

Example: $K = \mathbf{Q}$ with the p-adic norm. Then (writing a general rational as a/b in lowest terms)

$$R = \{a/b : a, b \in \mathbf{Z},\ p \nmid b\}$$

and

$$I = \{a/b \in R : p \mid a\} = pR.$$

$$R = \{a/b : a, b \in \mathbf{Z},\ p \nmid b\}$$

and

$$I = \{a/b \in R : p \mid a\} = pR.$$

I claim that $R/I = \mathbf{Z}/p\mathbf{Z}$, and to check this all I have to do is to check that $\{0, 1, 2, \ldots, p-1\}$ meets every coset $r + I$ exactly once, which follows easily from the statement that given $a, b \in \mathbf{Z}$ with $p \nmid b$ there's a unique blah blah $t \in \{0, 1, 2, \ldots, p-1\}$ with $a \equiv bt \bmod p$. So the residue field of $\mathbf{Q}$ with its p-adic norm is $\mathbf{Z}/p\mathbf{Z}$.

My aim now is to basically prove a structure theorem for characteristic zero blah non-archimedean fields K which are complete with respect to a discrete valuation; this will easily give us enough to show that if K is the completion of a number field at a non-zero prime ideal then K and $K^\times$ are "locally compact abelian groups", which is the buzz-word we'll need to do the abstract Fourier analysis we'll need for Tate's thesis later on.

Structure of complete discrete non-arch fields.

The first (and main) goal of this lecture is to explain "what a complete discretely-valued non-arch field looks like"—we'll end up with some kind of "structure theorem", analogous to the theorem that every real has an essentially unique decimal expansion, but with 0.1 replaced by a small number in the field—for example the role of 0.1 is played by p in $\mathbf{Q}_p$. This structure theorem (plus the associated exercises on the example sheet) will hopefully greatly clarify what fields like the p-adic numbers (and their finite extensions) look like. Hopefully, by the end of the lecture, we'll begin to have a concrete feeling about how to compute with $\mathbf{Q}_p$, just as we have a concrete feeling about how to compute with real numbers.

Let K be a field complete with respect to a non-arch norm (we don't need discreteness for the next few slides). I'm going to do some "abstract analysis" now—things which will be familiar from basic analysis classes, but which will work just as well in K because they work for any complete field, not just the reals or complexes.

Definition. If $x_1, x_2, x_3, \ldots \in K$ then we say that $\sum_{n\geq 1} x_n$ *converges* if the partial sums tend to a limit ℓ: we write $\sum_{n\geq 1} x_n = \ell$. Because we're assuming K is complete, the sum converges iff the partial sums $s_m = \sum_{i=1}^{m} x_i$ are Cauchy, and the standard argument shows that if the sum converges then $x_n = s_n - s_{n-1}$ had better tend to zero (Cauchyness implies $s_n - s_{n-1}$ gets arbitrarily small).

[$\sum x_n$ converges implies $x_n \to 0$].

The weird thing is that, *in the non-arch world*, the converse is true. Let $x_1, x_2, x_3, \ldots$ be a sequence in a complete non-arch field K.

Lemma. If $x_n \to 0$ as $n \to \infty$ then $\sum_{n \geq 1} x_n$ converges! Furthermore, if B is real and $|x_n| \leq B$ for all n then $\sum x_n = s$ with $|s| \leq B$ too.

Proof. By an easy induction on n, using the definition of a non-archimedean norm, we see that if $|x_i| \leq B$ for $1 \leq i \leq n$ then $|\sum_{i=1}^{n} x_i| \leq B$ (note that this is a finite sum). It's easy (but crucial) to deduce from this that a sequence (a_n) is Cauchy if and only if $a_n - a_{n-1}$ tends to zero as $n \to \infty$. Now apply this with $a_n = \sum_{i=1}^{n} x_i$ to deduce that $x_n \to 0$ implies that the the a_n form a Cauchy sequence, and hence converge. One way of doing the second part is to prove that if $a_n \to \ell$ as $n \to \infty$ then $|a_n| \to |\ell|$ in $\mathbf{R}$—this is true in any normed field (hint: WLOG triangle inequality holds; now use it judiciously).

Before we go any further, let me explain why the residue field of a non-archimedean normed field is the same as the residue field of the completion. This is easy. Let k be a non-archimedean normed field with completion $\widehat{k}$. Let R, I be the integers and maximal ideal for k, and let $\widehat{R}$, $\widehat{I}$ denote the corresponding things for $\widehat{k}$. There's a natural map $k \to \widehat{k}$ sending R to $\widehat{R}$ and I to $\widehat{I}$, and hence sending $\kappa = R/I$ to $\widehat{\kappa} = \widehat{R}/\widehat{I}$. We'll see a bit later that $\widehat{k}$ can be "much bigger than k" (analogous to $\mathbf{R}$ being much bigger than $\mathbf{Q}$). But...

Lemma. The map $\kappa \to \widehat{\kappa}$ is an isomorphism of fields.

Proof. Injectivity is clear ($k \to \widehat{k}$ is norm-preserving, so $R \cap \widehat{I} = I$). To get surjectivity, for $\widehat{r} \in \widehat{R}$ simply choose $r \in k$ with $|r - \widehat{r}| < 1$ (this is possible by denseness) and observe that this implies $|r| \leq 1$ and hence $r \in R$. Moreover $r - \widehat{r} \in \widehat{I}$, so $\widehat{R} = \widehat{I} + R$ which shows that $\kappa \to \widehat{\kappa}$ is surjective.

Corollary. The residue field of $\mathbf{Q}_p$ is $\mathbf{Z}/p\mathbf{Z}$. For this is the residue field of $\mathbf{Q}$ with the p-adic norm.

Exercise: let k be a number field and let P denote a non-zero prime ideal of its integer ring A. Show that the residue field of k equipped with the P-adic norm is canonically isomorphic to A/P (hint: if R and I are the usual things then construct a natural surjective ring homomorphism $R \to A/P$ with kernel equal to I). Deduce that the residue field of k_P is A/P.

Clarification: A is the integers of k (so, for example, $\mathbf{Z}$ is the integers of $\mathbf{Q}$) but we also referred to $R = \{x \in k : |x| \leq 1\}$ as the "integers of k"—this notion of course depends on the choice of a norm on k. Sorry. If $k = \mathbf{Q}$ above then $A = \mathbf{Z}$ and if $P = (p)$ then $R = \{a/b : p \nmid b\}$ and I'm saying that there's a natural map $R \to \mathbf{Z}/p\mathbf{Z}$.

As a consequence, we see that if k_P is the completion of a number field at a non-zero prime ideal then the residue field of k_P is *finite*—such fields have a much more arithmetic flavour than general complete normed fields (for example you can do analysis in any complete normed field, but if the norm is discrete and the residue field is finite then you can do local class field theory (i.e., arithmetic) too).

Exercise: consider the ring $\mathbf{C}[[T]]$ of power series $\sum_{n\geq 0} a_n T^n$ with complex coefficients (and no convergence conditions—just abstract power series) (NB this exercise would work if you replaced $\mathbf{C}$ by any field at all). Let $k := \mathbf{C}((T))$ denote its field of fractions. Check that a general element of $\mathbf{C}((T))$ is $\sum_{n\geq M} a_n T^n$ with M a possibly negative integer. Define a norm on $\mathbf{C}((t))$ by $|0| = 0$ and, for $f = \sum_{n\geq M} a_n T^n$ with $a_M \neq 0$, set $|f| = e^{-M}$ (where e could really be replaced by any real number greater than 1). Check that this is a non-arch norm on k, that the integers R are $\mathbf{C}[[T]]$, that the maximal ideal I is $T\mathbf{C}[[T]]$ and that the residue field is $\mathbf{C}$ again. So "we can do analysis in k but not arithmetic".

Before we go on to prove the structure theorem, let's play about a bit with Cauchy sequences in $\mathbf{Q}$ with the p-adic norm, and see if any of them converge.

Example 1: Consider the sequence

$$3, 33, 333, 3333, \ldots$$

in $\mathbf{Q}$ with the 5-adic norm.

(a) It's Cauchy! Because if a_n is "n threes" then for $n \leq m$ we have $10^n \mid (a_m - a_n)$ so $|a_m - a_n| \leq 5^{-n}$.

(b) In fact it's even convergent! Because $3a_n + 1 = 10^{n+1}$ which tends to zero in the 5-adic norm, so $a_n \to -\frac{1}{3}$.

Example 2: Let's put the 3-adic norm on $\mathbf{Q}$. Set $a_1 = 1$ and $a_2 = 4$ and note that $3^2 \mid (a_2^2 - 7)$. Let's try and find an integer a_3 with $3^3 \mid (a_3^2 - 7)$. Let's try $a_3 = 4 + 9n$; then $a_3^2 = 16 + 72n$ mod 27 so $a_3^2 - 7 \equiv 0$ mod 27 iff $1 + 8n \equiv 0$ mod 3 so let's set $n = 1$ and $a_3 = 13$; this works.

Can we pull this trick off in general? Say $m \geq 1$ and $a_m \equiv 1$ mod 3 and

$$a_m^2 \equiv 7 \text{ mod } 3^m.$$

Can we find a_{m+1} with $a_{m+1}^2 \equiv 7 \text{ mod } 3^{m+1}$? Let's try setting $a_{m+1} = a_m + 3^m n$ for some n to be determined. Then we see that

$$\begin{aligned} a_{m+1}^2 &\equiv a_m^2 + 2n.3^m \pmod{3^{m+1}} \\ &\equiv 7 + t_m.3^m + 2n.3^m \pmod{3^{m+1}} \end{aligned}$$

and we can solve $2n + t_m \equiv 0$ mod 3 for n, so we can indeed find a_{m+1} whose square is 3-adically close to 7, and by letting m go to infinity we can get as close as we want.

Upshot: we have a sequence $a_1, a_2, a_3, \ldots$ of elements of $\mathbf{Z}$, with $a_{n+1} - a_n$ a multiple of 3^n, and a_n^2 tending to 7 in $\mathbf{Q}$ with the 3-adic norm.

Because $a_{m+1} - a_m$ is a multiple of 3^m we see that the a_m are a Cauchy sequence, and their limit ℓ in $\mathbf{Q}_3$ is visibly going to satisfy $\ell^2 = 7$. Hence $\mathbf{Q}$ is not complete with respect to the 3-adic norm!

Exercise: check that -7 is a square in $\mathbf{Q}_2$ and $1-p<0$ is a square in $\mathbf{Q}_p$ for any $p>2$. Hence $\mathbf{Q}$ is not complete with respect to the p-adic norm, for any p.

Remark: I've collected up these exercises and put them on an example sheet. See the course web page.

Now let's assume that F is a field with a non-trivial non-archimedean discrete norm. In this case we have seen that $|F^\times| := \{|a| : a \in F^\times\}$ is $\{q^n : n \in \mathbf{Z}\}$ for some $q > 1$; set $\rho = 1/q < 1$ and let's choose $\pi \in F$ with $|\pi| = \rho$. We call π a *uniformiser* in F. As an example, if $F = \mathbf{Q}$ or $\mathbf{Q}_p$ with the p-adic norm then we can set $\rho = 1/p$ and $\pi = p$, and more generally if F is a number field k or a completion k_P at a prime ideal then $\rho = 1/N(P)$ and, even though P may not be principal, we can find $x \in k$ an algebraic integer with $(x) = PJ$ and $P \nmid J$ (for example, by uniqueness of factorization we have $P^2 \neq P$ and any $x \in P$ with $x \notin P^2$ will do), and then $|x| = 1/N(P) = \rho$ so x is a uniformiser for both k and k_P with their P-adic norms.

Now let R be the integers of K, and let I denote the maximal ideal of R. If π is a uniformiser, then $y \in I$ implies $|y| < 1$ and hence $|y| \leq |\pi|$, so $y = z\pi$ with $|z| \leq 1$ and we have proved that $I = (\pi)$ is a principal ideal.

Now here we go with the structure theorem. Let K be complete with respect to a non-trivial non-arch norm. Let R be the integers, I the maximal ideal of R, let π be a uniformiser (so $I = (\pi)$) and let κ denote the residue field R/I. Let S denote a subset of R, containing 0, such that the reduction map $S \to R/I$ is a bijection (so S is a set of representatives for R/I).

Theorem.

(a) If a_0, a_1, a_2,... is an arbitrary infinite sequence of elements of S, then the infinite sum $\sum_{n\geq 0} a_n\pi^n$ converges in R, and furthermore for every element r of R it's possible to write $r = \sum_{n\geq 0} a_n\pi^n$ with the a_n as above, in a unique way.

(b) If $0 \neq r \in R$ then $r = \sum_{n\geq 0} a_n\pi^n$ with at least one $a_n \neq 0$ and in fact $|r| = |\pi|^m$, where $m \geq 0$ is the smallest non-negative integer such that $a_m \neq 0$.

(c) A general non-zero element α of K can be written uniquely as $\alpha = \sum_{n \geq M} a_n \pi^n$ with $a_M \neq 0$, $a_n \in S$ for all n, and we have $|\alpha| = |\pi|^M$.

Before we go on, let's observe the consequences for $\mathbf{Q}_p$. Let $\mathbf{Z}_p$ denote $\{x \in \mathbf{Q}_p : |x| \leq 1\}$.

Corollary. A general element of $\mathbf{Z}_p$ can be written uniquely as $\sum_{n \geq 0} a_n p^n$ with each $a_n \in \{0, 1, 2, \ldots, p-1\}$. A general non-zero element of $\mathbf{Q}_p$ can be written $\sum_{n \geq M} a_n p^n$ with $M \in \mathbf{Z}$, $0 \leq a_n \leq p-1$ and $a_M \neq 0$.

Note that we now see why π is called a blah uniformiser—it's playing some kind of analogue to the role of a local uniformiser in the theory of complex analytic functions of one variable, with the theorem giving a power series expansion near a point.

Corollary. $\mathbf{Q} \neq \mathbf{Q}_p$. Indeed we see that $\mathbf{Q}_p$ is uncountable.

Exercise: if $\alpha \in \mathbf{Q}^\times$ and we write $\alpha = \sum_{n \geq M} a_n p^n$ with $0 \leq a_n < p$ then check that the sequence a_n is ultimately periodic. Hence a number like $\sum_{n \geq 1} p^{n!}$ is an explicit example of an element in $\mathbf{Q}_p$ but not $\mathbf{Q}$.

Let's state the theorem again because it's a new lecture.

Let K be complete with respect to a non-trivial non-arch norm. Let R be the integers, I the maximal ideal of R, and let π be a uniformiser (so $|\pi| = \rho$ with $0 < \rho < 1$ and $|K^\times| = \rho^{\mathbf{Z}}$, and $I = (\pi)$). Let κ denote the residue field R/I. Let S denote a subset of R, containing 0, such that the reduction map $S \to R/I = \kappa$ is a bijection (so S is a set of representatives for κ).

Theorem.

(a) If a_0, a_1, a_2,... is an arbitrary infinite sequence of elements of S, then the infinite sum $\sum_{n\geq 0} a_n \pi^n$ converges in R, and furthermore for every element r of R it's possible to write $r = \sum_{n\geq 0} a_n \pi^n$ with the a_n as above, in a unique way.

(b) If $0 \neq r \in R$ then $r = \sum_{n\geq 0} a_n \pi^n$ with at least one $a_n \neq 0$ and in fact $|r| = |\pi|^m$, where $m \geq 0$ is the smallest non-negative integer such that $a_m \neq 0$.

(c) A general non-zero element α of K can be written uniquely as $\alpha = \sum_{n\geq M} a_n \pi^n$ with $a_M \neq 0$, $a_n \in S$ for all n, and we have $|\alpha| = |\pi|^M$.

Proof of theorem.

If $a_0, a_1, a_2, \ldots$ are arbitrary elements of R then $|a_n| \leq 1$ so $|a_n\pi^n| \leq \rho^n \to 0$, where $0 < \rho < 1$ is the real number which generates the norm group $|K^\times|$ as above. So the sequence $(a_n\pi^n)$ tends to zero, so the sum $\sum_{n\geq 0} a_n\pi^n$ converges, and furthermore $|a_n\pi^n| \leq 1$ for all $n \geq 0$ and hence the sum converges in R. That's done the first part of (a), because $S \subseteq R$ by definition.

Next note that again by definition $a \in S$ implies that either $a = 0$ or $|a| = 1$. So now if $r = \sum_{n\geq 0} a_n\pi^n$ with $a_n \in S$ and not all of the a_n equal to zero, and if $m \geq 0$ is the smallest non-negative integer with $a_m \neq 0$, then

$$\begin{aligned} r &= \sum_{n\geq 0} a_n\pi^n \\ &= a_m\pi^m + \sum_{n\geq m+1} a_n\pi^n \end{aligned}$$

and $|a_m\pi^m| = \rho^m$ whereas each term in blah $\sum_{n\geq m+1} a_n\pi^n$ has norm at most $\rho^{m+1} < \rho^m$,

so the sum converges to something with norm at most ρ^{m+1}, so $|a_m\pi^m| > |\sum_{n\geq m+1} a_n\pi^n|$ and we see $|r| = |a_m\pi^m| = \rho^m$. This does (b).

Now the uniqueness in (a) is easy: if $r = \sum_{n\geq 0} a_n\pi^n = \sum_{n\geq 0} b_n\pi^n$ with the a_i and b_i in S then $0 = \sum_{n\geq 0}(a_n - b_n)\pi^n$. But it's easily checked that for $a, b \in S$, either $a - b = 0$ or $|a - b| = 1$. So the argument above shows that if $a_n \neq b_n$ for some n then $|\sum(a_n - b_n)\pi^n| > 0$, contradicting $\sum_{n\geq 0} a_n\pi^n = \sum_{n\geq 0} b_n\pi^n$. Hence "$\pi$-adic expansions" are unique, if they exist.

To finish (a) we need a construction proof: given $r \in R$ we need to find $a_n \in S$ with $r = \sum_{n\geq 0} a_n\pi^n$. There's a natural way to do this. Given $r \in R$ we consider the image $\overline{r}$ of r in κ, the residue field. Choose $a_0 \in S$ whose reduction is $\overline{r}$. Now $r - a_0$ reduces to zero in κ, so $|r - a_0| < 1$ and hence $|r - a_0| \leq \rho$.

Hence $r_1 := (r - a_0)/\pi$ satisfies $|r_1| \leq 1$ and we can apply the same trick to find a_1 with $|r_1 - a_1| \leq \rho$. Hence $|\pi r_1 - \pi a_1| \leq \rho^2$ and we deduce

$$|r - a_0 - \pi a_1| \leq \rho^2.$$

Set $r_2 = (r - a_0 - \pi a_1)/pi^2$ and continue in this way. At the Nth step we find

$$|r - \sum_{i=0}^{N} a_i \pi^i| \leq \rho^{N+1}$$

so, by definition, $\sum_{n \geq 0} a_n \pi^n = r$.

That's done (a) and (b). For (c) we just observe that any $\alpha \in K$ with $\alpha \neq 0$ we have $|\alpha| = \rho^M$ for some integer M, and hence $\pi^{-M}\alpha \in R$ (with norm 1). So

$$\pi^{-M}\alpha = \sum_{n \geq 0} b_n \pi^n$$

with $b_n \in S$ and b_0 equal to a lift of the reduction of $\pi^{-M}\alpha$ in κ, so $b_0 \neq 0$. So

$$\alpha = \sum_{n \geq M} a_n \pi^n$$

with $a_n = b_{n-M}$.

We're done with our structure theorem; now go and do some exercises on the example sheet.

A corollary whose importance will become clear later is:

Corollary. If K is complete with respect to a non-trivial non-arch discrete norm, and K has integer ring R with maximal ideal I and residue field κ, then R (with the topology induced from the metric $d(x,y) = |x-y|$) is compact iff κ is finite.

Proof. Because R is a metric space, compactness is equivalent to sequential compactness, which I'll remind you means that given a sequence $(r_m)_{m\geq 1}$ with $r_m \in R$ we can always find a convergent subsequence, that is $m_0 < m_1 < m_2 < m_3 < \ldots$ such that $(r_{m_j})_{j\geq 0}$ converges. Let's firstly assume κ is finite and prove that R is sequentially compact.

By the structure theorem we can write

$$r_m = \sum_{n\geq 0} a_{m,n}\pi^n$$

with $a_{m,n} \in S$ (a set of coset representatives for κ).

Now κ is finite so S is finite, so we can apply the usual trick: König's Lemma (which according to Wikipedia is due to Kőnig). Explicitly, we know that $a_{m,0}$ assumes at least one value in S infinitely often; call it a_0, and let m_0 be any m such that $a_{m,0} = a_0$. Now, amongst the infinitely many $m > m_0$ with $a_{m,0} = a_0$, we know that $a_{m,1}$ takes on a value infinitely often—call it a_1. Let m_1 be one of the infinitely many m with $m > m_0$, $a_{m,0} = a_0$ and $a_{m,1} = a_1$. Continue in this way and we see easily that $\sum_{n\geq 0} a_n\pi^n$ is the limit of the infinite subsequence $r_{m_0}, r_{m_1}, r_{m_2}, \ldots$.

Conversely, if κ is infinite, then here's an infinite open cover of R with no finite subcover: for any $s \in S$ the open disc centre s and radius 1 is everything of the form $s + \alpha$ with $|\alpha| < 1$, so it's $s + I$. Because S is a set of coset representatives for κ we see that R is the disjoint union of the open sets $s + I$ for $s \in S$, and this is an infinite disjoint cover of R by open sets, which visibly has no finite subcover.

Corollary. If k is a number field equipped with a P-adic norm, and if R is the integers of k_P, then R is compact.

Indeed, the residue field of k_P is A/P, where A is the integers of k in the sense of algebraic number theory.

That corollary is very important for Tate's thesis, as we'll see later on. I want to finish local fields today, so, rather than developing the theory in some kind of logical way (for example Hensel's Lemma would be a natural thing to do next) I am just going to prove the other main thing we'll need, which is that if k_P is the completion of a number field at a prime ideal then k_P is naturally a finite extension of $\mathbf{Q}_p$, and I'll say a little about the structure of such extensions.

Let k be a number field, and P a non-zero prime ideal of its integer ring. We can think of k as a finite-dimensional vector space over $\mathbf{Q}$. Now let's say that P contains the rational prime p. The restriction of the P-adic norm $|.|_P$ on k, to $\mathbf{Q}$, is easily checked to satisfy $|\ell|_P = 1$ for ℓ a prime with $\ell \neq p$, and $|p|_P = p^{-m}$ for some positive integer m, so $|.|_P$ on k induces a norm equivalent to the p-adic norm on $\mathbf{Q}$ (in fact it's just the mth power of the p-adic norm, where m is easily checked to be ef, where the size of k_P is p^f and where $(p) = P^e.J$ with J and ideal coprime to P.

Now there are inclusions of fields $\mathbf{Q} \to k \to k_P$, and k_P is complete. Of course k_P might not be the completion of $\mathbf{Q}$, because there's no reason to expect that $\mathbf{Q}$ is dense in k_P [the archimedean analogue of what's going on is that $\mathbf{C}$ is an archimedean completion of $\mathbf{Q}(i)$ with $i^2 = -1$ but the resulting map $\mathbf{Q} \to \mathbf{C}$ doesn't have dense image].

$[\mathbf{Q} \to k \to k_P]$

But we can certainly take the closure of $\mathbf{Q}$ in k_P. A closed subspace of a complete metric space is complete, and it's easy to check that the closure of $\mathbf{Q}$ in k_P is a field (limit of sum is sum of limits, limit of product is product of limits, limit of reciprocals is reciprocal of limits when this makes sense), and hence a normed field (the norm is induced from k_P hence the axioms are satisfied). Hence this closure must be the completion of $\mathbf{Q}$ with respect to the mth power of the p-adic norm (because it's *a* completion, we showed that completions are unique up to unique isomorphism).

We deduce that k_P contains a copy of $\mathbf{Q}_p$ (although, as already mentioned, the norm on k_P restricts on $\mathbf{Q}_p$ to a norm which is in general a non-trivial power of the usual p-adic norm). Now $k/\mathbf{Q}$ was a finite extension,

so it won't surprise you to learn that $k_P/\mathbf{Q}_p$ will also be a finite extension. Perhaps what will surprise you is that the degree of $k_P/\mathbf{Q}_p$ might be smaller than that of $k/\mathbf{Q}$. In fact let me prove something stronger, which will clarify what's going on.

Let me start with some abstract algebra.

Let L be a field and let K be a subfield of L. Then L is naturally a vector space over K. The dimension of L as a K-vector space might be finite (for example $\mathbf{C}$ has dimension 2 over $\mathbf{R}$, $\mathbf{Q}(i)$ has dimension 2 over $\mathbf{Q}$, $\mathbf{Q}(2^{1/57})$ has dimension 57 over $\mathbf{Q}$) or infinite (for example $\mathbf{C}$ has infinite dimension as a $\mathbf{Q}$-vector space). Recall that *by definition* a number field is a field k that contains a copy of $\mathbf{Q}$ and such that the $\mathbf{Q}$-dimension of k is finite.

Let's go back to the general case $K \subseteq L$ and let's assume that the dimension of L as K-vector space is a finite number n. We say "L is finite over K", and "L has degree n over K" or even "L/K has degree n". Note that L/K isn't a quotient, it's just notation.

Now for $\lambda \in L$, multiplication by λ is a map $L \to L$ which is L-linear and hence K-linear, so we can regard it as a linear map on an n-dimensional vector space, and as such it has a trace and a determinant.

We define the *trace* of λ, written $\mathrm{Tr}(\lambda)$ or sometimes $\mathrm{Tr}_{L/K}(\lambda)$, to be the trace of this linear map, and we define the *norm* of λ to be the determinant of that linear map and write $N(\lambda)$ or $N_{L/K}(\lambda)$.

Note that the trace and the norm of an element of L is an element of K. Moreover $\mathrm{Tr}(\alpha + \beta) = \mathrm{Tr}(\alpha) + \mathrm{Tr}(\beta)$ and $N(\alpha\beta) = N(\alpha)N(\beta)$.

Example: Multiplication by $x + iy \in \mathbf{C}$ is, when you think of $\mathbf{C}$ as $\mathbf{R}^2$ with basis $1, i$, represented by the matrix $\begin{pmatrix} x & -y \\ y & x \end{pmatrix}$ and hence has trace $2x$ and determinant $x^2 + y^2$. So $\mathrm{Tr}_{\mathbf{C}/\mathbf{R}}(x + iy) = 2x$ and $N_{\mathbf{C}/\mathbf{R}}(x + iy) = x^2 + y^2$.

Example: if $K \subseteq L$ and L is finite over K of degree n, and $\alpha \in K$ then $\mathrm{Tr}_{L/K}(\alpha) = n\alpha$, and $N_{L/K}(\alpha) = \alpha^n$ (proof: the matrix representing multiplication by α is scalar).

Because the norm N is multiplicative, it should be no surprise that it can be used to extend norms (i.e. maps of the form $|.|$).

Lemma. Let $K \subseteq L$ with L finite over K, of degree n. Assume furthermore that K is equipped with a non-archimedean norm $|.|$ that makes K complete.

Then there is a *unique* norm $||.||$ on L which restricts to $|.|$ on K. It's non-archimedean, it makes L into a complete normed field, and it is given by the formula

$$||\lambda|| = |N_{L/K}(\lambda)|^{1/n}$$

Proof. Omitted. On example sheet. Elementary but a little long.

Note that the uniqueness statement needs K to be complete. For example $\mathbf{Q}(i)$ is finite over $\mathbf{Q}$ but if $A = \mathbf{Z}[i]$ then in A we have $(5) = (2+i)(2-i) = PQ$ and the P-adic norm and the Q-adic norm on $\mathbf{Q}(i)$ both extend the 5-adic norm on $\mathbf{Q}$. So in fact the lemma gives another proof that $\mathbf{Q}$ isn't complete with respect to the 5-adic norm (and it's not much trouble to deduce that it's not complete with respect to any p-adic norm this way).

Using this lemma let's deduce its analogue in the "incomplete" case (although I'm really only interested in the case of number fields). So now say L/K is a finite extension of fields of characteristic zero (or more generally, a finite separable extension of fields, if you know what that means). As we've just seen, it is now no longer true that a non-arch norm on K extends uniquely to a non-arch norm on L. Indeed, if L and K are number fields and if P is a prime ideal of the algebraic integers of K, and P factors in L as $Q_1^{e1}Q_2^{e2}\dots Q_r^{er}$ in L, then we have at least r norms on L extending the P-adic norm on K (namely the appropriate powers of the Q_i-adic norms for each i). But it turns out to be true that in the general case there are only finitely many norms on L that extend a given non-arch norm on K.

Let's fix a norm $|.|$ on K. We're asking how to extend it to L. The key construction is the following. Let $\hat{K}$ denote the completion of K with respect to $|.|$. Then L and $\hat{K}$ both contain copies of K, so we can form the tensor product $L\otimes_K \hat{K}$. I can write down what this is explicitly:

We know that L can be written as $K(\alpha)$, for some $\alpha \in L$ (that is, L is the smallest field containing K and α). Hence we can write $L = K[X]/(P(X))$ where $P(X)$ is the minimal polynomial of α, that is the monic polynomial of smallest positive degree with coefficients in K and having α as a root. For example $\mathbf{C} = \mathbf{R}(i) = \mathbf{R}[X]/(X^2+1)$. Now if you're not completely certain about the tensor product, you can simply *define* $L \otimes_K \widehat{K}$ to be the ring

$$\widehat{K}[X]/(P(X)).$$

Now, considered as a polynomial with coefficients in K, $P(X)$ was irreducible, and hence $(P(X))$ was a maximal ideal of $K[X]$ (and thus L was a field!). However, $P(X)$ might not be irreducible in $\widehat{K}[X]$. One thing is for sure though, and that's that $P(X)$ has no repeated roots (because if it did then it would have a factor in common with its derivative, contradicting the fact that it's irreducible). So, in $\widehat{K}[X]$, if $P(X)$ factors, it will factor as $Q_1(X)Q_2(X)\dots Q_r(X)$ with the $Q_i(X) \in \widehat{K}[X]$ irreducible and pairwise coprime.

So by the Chinese Remainder Theorem we see that

$$\begin{aligned} L \otimes_K \widehat{K} &= \widehat{K}[X]/(P(X)) \\ &= \widehat{K}[X]/(\prod_{i=1}^{r} Q_i(X)) \\ &= \oplus_{i=1}^{r} \widehat{K}[X]/(Q_i(X)) \\ &= \oplus_{i=1}^{r} \widehat{L}_i \end{aligned}$$

(this last line is a definition) where $\widehat{L}_i = \widehat{K}[X]/(Q_i(X))$ is a field with a name that is currently only suggestive of what is to come rather than being any kind of completion of L. Note that there's a completely canonical natural map $L \to L \otimes_K \widehat{K}$, sending α to X, and hence a map $L \to \oplus_{i=1}^{r} \widehat{L}_i$ so, by projection, maps $L \to \widehat{L}_i$ for each i.

Theorem. L/K finite as above, and $|.|$ a norm on K. Then there are only r extensions $||.||_i$ $(1 \leq i \leq r)$ of $|.|$ to L, and if L_i denotes L equipped with the ith extension then the completion of L_i is (after re-ordering if necessary) isomorphic to $\widehat{L}_i$.

Proof. Let $||.||$ be any norm of L extending $|.|$ on K. Let $\hat{L}$ denote the completion of L with respect to this norm. Then the closure of K in $\hat{L}$ is isomorphic to $\hat{K}$ (we saw this argument once today already). Now consider the subfield $\hat{K}(\alpha)$ of $\hat{L}$. Clearly $\hat{K}(\alpha)$ is a finite extension of $\hat{K}$. Moreover $\hat{K}(\alpha)$ inherits a norm from $\hat{L}$. So by the lemma-to-be-proved-on-the-example-sheet, $\hat{K}(\alpha)$ is complete! In particular it's a closed subspace of $\hat{L}$ that contains L and hence it *is* $\hat{L}$. Let $Q(X)$ denote the minimal polynomial of α over $\hat{K}$. Then $Q(X)$ divides $P(X)$ (because $P(\alpha) = 0$) and hence $Q(X)$ is one of the $Q_i(X)$ above and

$$\begin{aligned}\hat{L} &= \hat{K}[X]/(Q(X)) \\ &= \hat{K}[X]/(Q_i(X)) \\ &= \hat{L}_i\end{aligned}$$

for some i.

Conversely, each $\hat{L}_i$ is visibly a finite extension of $\hat{K}$ and hence inherits a unique norm extending that on $\hat{K}$, and the inclusion $L \to \hat{L}_i$ induces a norm on L.

All that remains is to show that distinct i's induce non-equivalent norms on L. But this is clear—if the norms corresponding to two distinct is were equivalent, then the completions would be isomorphic as L-algebras, but $Q_i(\alpha) = 0$ in $\hat{L}_i$ whereas $Q_j(\alpha) \neq 0$ in $\hat{L}_i$ if $i \neq j$.

I'll remind you that I stated earlier in the course, without proof, a theorem saying that the only non-arch norms on a number field k were the P-adic norms; it's not hard to use the above argument to reduce this statement to the case of $k = \mathbf{Q}$, which can be checked directly using a brute force argument due to Ostrowski.

As a final remark, we can now deduce that k_P is a finite extension of $\mathbf{Q}_p$ if $p \in P$. For we've just shown that k_P is a direct summand of the ring $k \otimes_{\mathbf{Q}} \mathbf{Q}_p$ and hence the dimension of $k_P/\mathbf{Q}_p$ is at most the dimension of $k/\mathbf{Q}$.

[Remark: people who want to see more of the theory of fields with norms have two excellent choices for books—Cassels' "Local fields", which does everything I did here but which is also completely stuffed with beautiful applications of the theory to number fields and Diophantine equations and lots of other things, and Serre's "Local Fields" which is more highbrow in nature and which goes much further than Cassels, going as far as proofs of the main theorems of local Class Field Theory.]

Chapter 3: Haar measure and abstract Fourier theory.

3.1: Introduction.

If f is a continuous function $\mathbf{R} \to \mathbf{C}$ such that $\int_{-\infty}^{\infty} |f(x)| \mathrm{d}x$ converges, then f has a Fourier transform $\widehat{f} : \mathbf{R} \to \mathbf{C}$, defined by

$$\widehat{f}(y) = \int_{-\infty}^{\infty} f(x) e^{-iyx} \mathrm{d}x.$$

$$\widehat{f}(y) = \int_{-\infty}^{\infty} f(x) e^{-iyx} \mathrm{d}x.$$

In general $\widehat{f}$ bears little resemblance to f. Let's do an example to stress this: let's say $f(x) = 1/(1+x^2)$. Then

$$\widehat{f}(y) = \int_{-\infty}^{\infty} \frac{e^{-iyx}}{1+x^2} \mathrm{d}x.$$

We do this integral via closing up the contour and getting back from $+\infty$ to $-\infty$ via a big arc $|z| = R$. We have a choice of two arcs—upper half plane and lower half plane—and which one we choose turns out to depend on the sign of y.

Let's imagine closing up via the upper half plane. So y is always real, but now we're thinking of x as a complex number with big positive imaginary part. If we want the integral along the big arc to be small then we'd better make sure that the integrand is small. So closing up along the top will work if $y < 0$ (because then we're integrating something whose value is at most c/R^2 along an arc whose length is $O(R)$).

And so, for $y < 0$,

$$\begin{aligned}\widehat{f}(y) &= \int_{-\infty}^{\infty} \frac{e^{-iyx}}{1+x^2}\mathrm{d}x \\ &= \lim_{D} \int_{D} \frac{e^{-iyx}}{1+x^2}\mathrm{d}x\end{aligned}$$

where D is a contour that looks like a D lying on its back, and is getting bigger and bigger. Now this integral is just going to be $2\pi i$ times the sum of the residues at the poles of $e^{-iyx}/(1+x^2)$ for x in the upper half plane. The only pole is at $x = i$, the residue is $e^y/(2i)$ and we deduce

$$\widehat{f}(y) = \pi e^y$$

for $y < 0$.

A similar argument shows $\widehat{f}(y) = \pi e^{-y}$ if $y > 0$ (now using the lower half plane). Finally $\widehat{f}(0) = [\tan^{-1}(x)]_{-\infty}^{\infty} = \pi$ so we conclude

$$\widehat{f}(y) = \pi e^{-|y|}.$$

The purpose of this was just to show that $\widehat{f}$ is of an entirely different nature to f.

Summary: if $f(x) = 1/(1+x^2)$ then $\widehat{f}(u) = \pi e^{-|y|}$. So in this case $\widehat{f}$ is "rapidly decreasing" (this means $\widehat{f}(y).P(y)$ tends to zero as $|y| \to \infty$, for any polynomial $P \in \mathbf{C}[X]$) but not differentiable, whereas f was infinitely differentiable but decreasing not particularly quickly.

Two very elementary exercises about Fourier transform:

(1) If $g(x) = f(x+r)$ (r real) then $\widehat{g}(y) = \widehat{f}(y)e^{iry}$.

(2) If $g(x) = f(x)e^{i\lambda x}$ then $\widehat{g}(y) = \widehat{f}(y-\lambda)$.

[Proof: change of variables]

This also indicates that $\widehat{f}$ is very much "not like f": it's transforming in a different way.

But here's a nice thing: sometimes $\widehat{f}$ also has a Fourier transform (for example the $\widehat{f}$ we just saw is certainly continuous and integrable). So we can take the Fourier transform again! And (1) and (2) together imply that $\widehat{\widehat{f}}$ behaves in a similar way to f (for example if $g(x) = f(x+r)$ then $\widehat{\widehat{g}}(z) = \widehat{\widehat{f}}(z-r)$).

Now let's try our toy example $f(x) = 1/(1+x^2)$, so $\widehat{f}(y) = \pi e^{-|y|}$. Then

$$\widehat{\widehat{f}}(z) = \pi \int_{-\infty}^{\infty} e^{-|y|} e^{-izy} \mathrm{d}y$$

and this integral can be done easily because the integrand has an indefinite integral. Split the integral into $\int_{-\infty}^{0} + \int_{0}^{\infty}$; the first integral is

$$\begin{aligned} &\pi \int_{-\infty}^{0} e^{y-izy} \mathrm{d}y \\ &= \pi [e^{y(1-iz)}/(1-iz)]_{-\infty}^{0} \\ &= \pi/(1-iz) \end{aligned}$$

and the second one is $\pi/(1+iz)$ so the sum is $2\pi/(1+z^2)$ and $\widehat{\widehat{f}}$ is looking remarkably similar to f. In fact, for this f, we have $\widehat{\widehat{f}}(x) = 2\pi f(x) = 2\pi f(-x)$ because f was even.

But the general theorem is that if f is now an arbitrary function which is, say, infinitely differentiable and rapidly decreasing (much weaker conditions will do, but these will suffice for us), then $\widehat{f}$ is also infinitely differentiable and rapidly decreasing, so $\widehat{\widehat{f}}$ makes sense and

Theorem (Fourier Inversion Theorem)

$$\widehat{\widehat{f}}(x) = 2\pi f(-x).$$

Now not only will I freely confess that I have no idea (yet) how to prove the above statement, but also, more importantly, before I had read Tate's thesis, I would never have *believed* that there would or could be some "abstract" version of this theorem which would, say, work over the p-adic numbers (what would play the role of 2π, for example??

So now I know better. In fact the Fourier transform should be thought of as some sort of "duality" sending functions on one $\mathbf{R}$ to functions on "a dual $\mathbf{R}$", and the Fourier inversion theorem is some form of the statement that the dual of the dual is the function you started with (up to some fudge factors).

A good analogy is with finite abelian groups G. Say G is finite abelian, and let $\widehat{G}$ be the set of (1-dimensional) characters of G. Then $\widehat{G}$ is a group non-canonically isomorphic to G. Now for $f : G \to \mathbf{C}$, define $\widehat{f} : \widehat{G} \to \mathbf{C}$ by

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

Exercise: if $\widehat{\widehat{G}}$ is identified with G in the following STRANGE way: let $g \in G$ define a group homomorphism $\widehat{G} \to \mathbf{C}^\times$ by sending χ to $\chi(g^{-1})$ [NOT $\chi(g)$], then $\widehat{\widehat{f}}(g) = (1/|G|).f(g)$.

$[\hat{\hat{f}}(g) = (1/|G|).f(g).]$

Note the minus sign in the Fourier inversion theorem corresponds to the strange identification of G with its double dual, and the fudge factor 2π corresponds to the fudge factor $1/|G|$. The analogy in fact is more than an analogy—our goal in this chapter is to formulate and prove an "abstract" Fourier inversion theorem and both the above things will be special cases. We need to start by coming up with an integration theory that works in much more generality than blah Riemann/Lesbesgue integration. Before we do that, I need to introduce the objects we'll be integrating on: locally compact Hausdorff topological groups.

3.2: Locally compact Hausdorff topological groups.

So on the real numbers we have the Riemann Integral. I'm going to explain in this lecture and the next a far more general integration theory that will work on an arbitrary locally compact Hausdorff topological group. So I have to start by explaining what a locally compact Hausdorff topological group is.

A *topological group* is a group G equipped with a topology on G such that $m : G \times G \to G$ and $i : G \to G$ defined by $m(x, y) = xy$ and $i(x) = x^{-1}$, are continuous (where $G \times G$ is equipped with the product topology). Examples: any group, with the discrete topology. The real numbers with its usual topology. The non-zero real numbers with its usual topology. If K is any normed field then K with the topology coming from the norm.

[One might ask whether continuity of m implies continuity of i. It doesn't: for example if G is the integers with the order topology (so the open sets are the empty set, the whole thing, and all sets of the form $\{n, n+1, n+2, n+3, \ldots\}$ then multiplication is continuous but inverse isn't).]

Here's a slightly more subtle example: if K is a normed field, then $K^\times$, with the topology induced from K, is a topological group. The reason one has to be careful here is that one has to check that inverse is continuous—but it is (exercise), because the topology is coming from a metric. I'll perhaps make the cryptic remark that if R is an arbitrary topological ring (so $+$ and $-$ and $*$ are continuous) then its unit group, with the induced topology, is not always a topological group, because inverse really might not be continuous in this generality; this can however be fixed by embedding $R^\times$ into R^2 via $u \mapsto (u, u^{-1})$, and giving it the subspace topology—then $R^\times$ really is a topological group.

Pedantic exercise: if K is a normed field, then check that the two topologies I've just put on $K^\times$ (the subspace topology coming from K, and the one coming from K^2) coincide.

Back to examples: If K is a normed field and G is an algebraic group over K (for example GL_n or Sp_n or something) then $G(K)$ is a topological group (so for example $\mathrm{GL}_n(\mathbf{R})$ and $\mathrm{GL}_n(\mathbf{Q}_p)$ are topological groups, or $E(\mathbf{Q}_p)$ for $E/\mathbf{Q}_p$ an elliptic curve, and so on). I won't prove these things because we don't need them, but they're not hard.

If $g \in G$ then the map $G \to G \times G$ sending $h \in G$ to (g, h) is continuous (think about the definition of the product topology, or the universal property) and hence if we fix $g \in G$ then "left multiplication by g", the function $G \to G$ sending h to gh, is continuous, and similarly right multiplication by g is continuous—and even homeomorphisms, because left/right multiplication by g^{-1} is a continuous inverse. In particular G is blah "homogeneous"—if $a, b \in G$ then left multiplication by ba^{-1} is a homeomorphism $G \to G$ sending a to b. In words, G "looks the same at a and at b"; the group of homeomorphisms of G acts transitively on G, if you prefer.

A remark on non-Hausdorff groups: It turns out that if G is a topological group which isn't Hausdorff, and if e is the identity element, then $\{e\}$ isn't a closed set, and its closure H is a normal subgroup of G such that G/H is naturally a Hausdorff topological group. Using this argument, questions about topological groups can frequently be reduced to the Hausdorff case, and we'll only be concerned with Hausdorff groups in practice anyway—for example, all the examples we saw above were Hausdorff; moreover topologies will usually come from metrics and hence will automatically be Hausdorff.

So let's get on. Let G be a hausdorff topological group. We want to integrate a class of continuous functions $G \to \mathbf{C}$. Which ones? Well, probably not all of them—for example if $G = \mathbf{R}$ then $f(x) = 1$ for all x won't be integrable. So let's restrict, at least for the time being, to continuous functions which vanish outside a compact set—this is a good finiteness condition. Unfortunately, in this generality, there might not be any such things! For example if $G = \mathbf{Q}$ with its subspace topology coming from $\mathbf{R}$, a continuous function $G \to \mathbf{C}$ which vanishes outside a compact set must be identically zero (exercise). This is unsurprising—who would do integration on $\mathbf{Q}$?? Here's a nice condition which will at least ensure the existence of lots of functions which vanish outside a compact set:

Definition. If X is a topological space and $x \in X$, we say that $U \subseteq X$ is an *open neighbourhood* of $x \in X$ if U is open and $x \in U$. We say $S \subseteq X$ is a *neighbourhood* of X if x is in the interior of S. We say that X is *locally compact* if every $x \in X$ has a compact neighbourhood.

For the rest of this chapter, we are only interested in locally compact Hausdorff topological groups. Let's call them LCHTGs.

Note that to check that a topological group is locally compact, it suffices to find a compact neighbourhood of the identity (by homogeneity).

Good examples of LCHTGS: if K is either the real numbers, or the complex numbers, or k_P for k a number field, then K (considered as a group under addition) is a LCHTG; for the reals and the complexes a compact neighbourhood of the identity is the closed unit ball, and for k_P the ring of integers will work, once we remember that the residue field is finite in this setting. Moreover, I claim that $K^\times$ (with the subspace topology induced from that of K) is also a LCHTG–if $K = k_P$ then this follows because $1+\pi R$ is a compact neighbourhood of 1, and in the archimedean case consider the closed ball centre 1 radius 1/2.

Now if X is a Hausdorff topological space then for $f : X \to \mathbf{C}$ continuous, it's easily checked that the following are equivalent:

(1) f vanishes outside a compact set (that is, there's some compact $K \subseteq X$ such that $f(x) = 0$ if $x \notin K$), and

(2) the *support* of f is compact

[recall that the support of a function $f : X \to \mathbf{C}$ is just the closure of the set $\{x \in X : f(x) \neq 0\}$].

So for a LCHTG G, let's define $\mathcal{K}(G)$ to be the continuous functions $G \to \mathbf{R}$ with compact support. The point is that this space is very rich (and in particular non-zero!). We'll see this in a second, using Urysohn's Lemma. First let me remark that we'd really like to integrate more functions than those in $\mathcal{K}(G)$, but $\mathcal{K}(G)$ is a good start and turns out to be where the work is; we'll expand our horizons later by taking limits. First I'll show $\mathcal{K}(G)$ is big. We need a lemma in order to prove Urysohn's Lemma.

First let me remind you that if X is compact Hausdorff and $C \subseteq U \subseteq X$ with C compact and U open, then we can find V open and D compact with $C \subseteq V \subseteq D \subseteq U$. For C and $K := X \backslash U$ are compact and disjoint, and by Hausdorffness it's easy (and fun!) to find $C \subseteq V$ and $K \subseteq W$ with V and W open and disjoint.

Definition, for convenience: $S \subseteq T \subseteq X$ are subsets of a topological space, we say that T is a *neighbourhood of* S if T is a neighbourhood of s for all $s \in S$.

Lemma. If X is locally compact and Hausdorff and $C \subseteq X$ is compact, then every neighbourhood of C contains a compact neighbourhood of C.

Proof. It suffices to check that if $C \subseteq U \subseteq X$ with C compact and U open, then there exists V open and D compact with $C \subseteq V \subseteq D \subseteq U$. Here's the idea; each $c \in C$ has a compact neighbourhood N_c, with interior M_c. Now C is covered by finitely many of the M_c, and the union X' of the corresponding N_c is compact and a neighbourhood of C. Replacing X with X' and U with $U' = X' \cap U$ reduces us to the case where X is compact, which we just did.

Lemma (Urysohn's Lemma): if X is a locally compact Hausdorff topological space, and $C \subseteq U \subseteq X$ with C compact and U open, then there's a continuous function $f : X \to \mathbf{R}$ such that

(1) $f(x) = 1$ for $x \in C$

(2) $f(x) = 0$ for $x \notin U$

(3) $0 \leq f(x) \leq 1$ for all $x \in X$

(4) Supp(f) is contained in U and is compact.

Proof.

The proof is sort of "constructive" (but does involve infinitely many choices). We apply the previous lemma infinitely often, basically. Set $C(1) = C$ and define $C(0)$ to be any compact neighbourhood of C in U.

[$C = C(1)$; $C(0)$ a compact neighbourhood of $C(1)$]

Now by the previous lemma (applied to C and the interior of $C(0)$) we get a compact $C(1/2)$ such that $C(0)$ is a neighbourhood of $C(1/2)$ and $C(1/2)$ is a neighbourhood of $C(1)$. Applying this trick again we get $C(1/4)$ between $C(1/2)$ and $C(0)$, and $C(3/4)$ between $C(1)$ and $C(1/2)$. Continuing in this way, we construct compacts $C(i/2^n)$ for all $n \geq 1$ and positive odd i with $0 < i < 2^n$, such that $C(\alpha)$ is in the interior of $C(\beta)$ for $\alpha < \beta$.

Now here's the magic: for $0 \leq r \leq 1$ an arbitrary real, define $C(r) = \cap_{s \leq r, s=i/2^n} C(s)$. Now we have a decreasing sequence of compact sets; let's finish the job by defining $C(r) = \emptyset$ for $r > 1$ and $C(r) = X$ for $r < 0$.

We define $f : X \to \mathbf{R}$ by letting $f(x)$ be the supremum of the α with $x \in C(\alpha)$. This sup visibly exists and is at most 1, the support of f is within $C(0)$, and indeed all properties we require of f are easy to check, the one with the most work being continuity, which goes like this: $f(x) < \beta$ if and only if $x \notin \cap_{\alpha<\beta} C(\alpha)$ and the intersection is closed so the complement is open, and $f(x) > \gamma$ if and only if $x \in \cup_{\alpha>\gamma} \mathrm{Int}(C(\alpha))$ (with Int meaning interior), which is also open, and so the x with $\gamma < f(x) < \beta$ are an open set, and that's enough.

Corollary If G is a LCHTG then $\mathcal{K}(G)$ separates points and in particular is non-zero.

3.3: Haar Measure/integral on a LCHTG.

As people probably realise, I'm preparing these lectures on the fly, but I might actually have to number lemmas in this section because the results are elementary but sometimes slightly tricky. I'll sometimes be sketchy with the easier proofs however, for the following reason: the only groups for which we'll actually need Haar integrals/measures are: (i) the reals and complexes (where the Haar integral is just the Lesbesgue integral), (ii) the p-adics and finite extensions thereof (where you can define the measure by hand), and (iii) the adeles (we'll get to these) (where you can define Haar measure as just a product of things in (i) and (ii)). My main motivation for going through this stuff really is that it's the kind of thing I wish I had been taught as a graduate student.

Let G be a LCHTG. We've seen that $\mathcal{K}(G)$ is non-zero, and moreover if we define $\mathcal{K}_+(G)$ to be the $f : G \to \mathbf{R}$ in $\mathcal{K}(G)$ such that $f(x) \geq 0$ for all x and $f(x) > 0$ for some x (that is, f isn't identically zero), then we've also see that $\mathcal{K}_+(G)$ isn't zero either. The reason we're sticking with the reals rather than the complexes is that we're after some kind of "integral" $\int_G$ which will take an element of $\mathcal{K}(G)$ to a real number and is guaranteed to take an element of $\mathcal{K}_+(G)$ to a non-negative real number; if we worked with the complexes then we couldn't enforce this sort of "positivity" condition so easily.

It's easy to define what a Haar integral is—the hard part is existence and uniqueness. If $f : G \to \mathbf{C}$ and $x \in G$ then define $f^x : G \to \mathbf{C}$ by

$$f^x(g) = f(gx^{-1}).$$

So it's just f composed with right multiplication by x^{-1} (don't read anything significant into the x^{-1}; it's easier to TEX f^x than ${}^x f$). Note that $f \in \mathcal{K}(G)$ implies $f^x \in \mathcal{K}(G)$ and similarly for $\mathcal{K}_+(G)$.

Definition. A *Haar Integral* on G is a non-zero $\mathbf{R}$-linear map

$$\mu : \mathcal{K}(G) \to \mathbf{R}$$

such that

(1) $\mu(f) \geq 0$ for $f \in \mathcal{K}_+(G)$

(2) $\mu(f) = \mu(f^x)$

Idea: think $\mu(f) = \int_G f$.

Remark: as I've already mentioned, ideally one would like to integrate more functions than just those with compact support, but these will come later on without too much trouble. The hard work is all in

Theorem. If G is a LCHTG then a Haar integral exists on G, and furthermore if μ_1 and μ_2 are two Haar integrals, then there's some $c > 0$ such that $c\mu_1(f) = \mu_2(f)$ for all $f \in \mathcal{K}(G)$.

As I mentioned before, we will principally be interested in the case $G = K$ or $K^\times$ for K a completion of a number field; in the archimedean case we have K isomorphic to $\mathbf{R}$ or $\mathbf{C}$ and we can use Riemann integration or Lesbesgue integration to produce a Haar measure (don't forget that it's "dx/x" in the $K^\times$ case to make it invariant under multiplication), and in the non-arch case one can check that $\mathcal{K}(G)$ is generated by step functions so (by linearity and translation-invariance) we only need to define the measure of the characteristic function of $\pi^n R$, which can be q^{-n}, and that does existence in all the cases we need. So in some sense this lecture and the next are not really logically necessary. So I might race through the elementary-undergraduate-exercise parts of some proofs a bit, although I will stress all of the key ideas.

In fact the proof contains several ideas. If you think about how the Riemann integral works, we first integrate "rectangular" functions (like the characteristic function of an interval) and then bound more general functions above and below by rectangular functions. The problem at this level is that such a "rectangular" function might not be continuous. Moreover, G is only a group, not a field, so we can't *yet* say things like "this open set is twice as big as this one". We fix this by, instead of using "rectangular" functions, setting up an "approximate" theory using an arbitrary element of $\mathcal{K}_+(G)$—indeed our first goal is, for $F \in \mathcal{K}_+(G)$, to define an "approximate integral" μ_F.

Notation: for $f, g : G \to \mathbf{R}$ we say $f \leq g$ if $f(x) \leq g(x)$ for all $x \in G$, and we say $f < g$ if $f \leq g$ and $f \neq g$. So, for example, $\mathcal{K}_+(G)$ is the $f \in \mathcal{K}(G)$ with $f > 0$.

Lemma 1. Say $f, F \in \mathcal{K}_+(G)$. Then there exist real numbers $\alpha_1, \ldots, \alpha_n \geq 0$ and $x_1, x_2, \ldots x_n \in G$ such that

$$\sum_{1 \leq i \leq n} \alpha_i F^{x_i} \geq f.$$

Hence if μ is a Haar integral, we have $\mu(f) \leq (\sum_i \alpha_i)\mu(F)$.

[think: what does this lemma "mean"?]

Proof. We know $F(t) = r > 0$ for some $t \in G$, so by continuity there's some open neighbourhood U of $e \in G$ such that $F(ut) > r/2$ for all $u \in U$. Now the support of f is covered by translates Uh of U and hence by finitely many translates; this gives the x_i (if Uh is in the cover then one of the x_i will be $t^{-1}h$). Finally note that $f \in \mathcal{K}(G)$ so f is bounded, say by M; now we can just let all the α_i be $M/(r/2)$. The final statement follows immediately from positivity, linearity and translation-invariance.

□

$[\mu(f) \leq (\sum_i \alpha_i)\mu(F)]$

As a consequence, which will guide us later, we deduce that for any $F \in \mathcal{K}_+(G)$ and for any Haar integral μ we have $\mu(F) > 0$ [or else taking F with $\mu(F) = 0$ shows that μ is identically zero on $\mathcal{K}_+(G)$; but for $f \in \mathcal{K}(G)$ we have $2f = (|f| + f) - (|f| - f)$ and both bracketed terms on the right are in $\mathcal{K}_+(G) \cup \{0\}$, so μ is identically zero, contradiction].

If we pretend that F is one of those "rectangle functions" then this motivates the following definition: for $f, F \in \mathcal{K}_+(G)$ we set $(f : F)$ to be the inf of the $\sum_i \alpha_i$ over all the possibilities for α_i in the lemma, that is, the inf over all the possible ways of choosing α_i and x_i with $f \leq \sum_i \alpha_i F^{x_i}$.

Exercise: prove $(f : h) \leq (f : g)(g : h)$ by observing that if $f \leq \sum_i \alpha_i g^{x_i}$ and $g \leq \sum \beta_j h^{y_j}$ then $f \leq \sum_{i,j} \alpha_i \beta_j h^{y_j x_i}$.

Now $(f : F)$ looks like a good candidate for the integral of f, at least if "the support of F is small", but in fact as well as "rounding errors" caused by F not being fine enough, there's a "normalisation" issue: if we replace F by $2F$, say, we see $(f : 2F) = 2(f : F)$. So if we want to define the integral of f as some kind of limit of the $(f : F)$ as, say, the "support of F tends to zero", we need to scale things.

So here's a crucial remark that shows that scaling is possible. Let f, F be as above (both in $\mathcal{K}_+(G)$). We defined $(f : F)$ to be the inf of the $\sum_i \alpha_i$ such that $\sum_i \alpha_i F^{x_i} \geq f$. Now any function in $\mathcal{K}_+(G)$ has a positive supremum, which it attains. Furthermore $\sum_i \alpha_i F^{x_i} \geq f$ implies that $(\sum_i \alpha_i) \sup(F) \geq \sup(f)$ which gives us a lower bound $\sum_i \alpha_i \geq (\sup(f)/\sup(F))$. Hence blah blah blah blah $(f : F) \geq \sup(f)/\sup(F) > 0$ and we've shown that $(f : F) > 0$ for all $f, F \in \mathcal{K}_+(G)$.

So here's the next good idea:

FIX ONCE AND FOR ALL A FUNCTION $\eta \in \mathcal{K}_+(G)$ (it doesn't matter what it is). We know that if a Haar integral exists, it will take η to something positive. We want to define "approximate Haar integrals" and tease the existence of a Haar integral from these approximate ones. For the approximate Haar integrals to be "compatible" we will simply force each of them to integrate η to 1.

Definition. For $f, F \in \mathcal{K}_+(G)$ define

$$\mu_F(f) := (f : F)/(\eta : F).$$

The idea: μ_F might not be a Haar integral but it's a good first approximation. We'll see that in fact as the support of F gets smaller the μ_F become better and better approximations to a Haar integral.

Exercise: Use the previous exercise to check that $\mu_F(f) \leq (f : \eta)$.

Now this definition of μ_F is great because it's impervious to linear changes of F. In fact it almost does the job already, at least for functions in $\mathcal{K}_+(G)$: μ_F is positive on $\mathcal{K}_+(G)$, it's translation-invariant, and satisfies $\mu_F(\alpha f) = \alpha\mu_F(f)$ for $\alpha > 0$. It's normalised in the sense that $\mu_F(\eta) = 1$. Unfortunately it's not additive; it's trivial to check that $\mu_F(f_1 + f_2) \leq \mu_F(f_1) + \mu_F(f_2)$ (easy exercise) but there's no reason why equality should hold (and it won't, in general).

$[\mu_F(f) := (f : F)/(\eta : F)$ and $\mu_F(f_1 + f_2) \leq \mu_F(f_1) + \mu_F(f_2)]$

We need more than subadditivity, we need an "approximate additivity", which is given by

Lemma 2. Let G be a LCHTG. Say $f_1, f_2 \in \mathcal{K}_+(G)$. Say $\epsilon > 0$. Then there's a *symmetric* open neighbourhood V (that is, $V = \{v^{-1} : v \in V\}$) of the identity in G (depending on f_1 and f_2 and ϵ) such that, for *any* $F \in \mathcal{K}_+(G)$ with support in V, we have

$$\mu_F(f_1 + f_2) \geq \mu_F(f_1) + \mu_F(f_2) - \epsilon.$$

Proof. Let C be the union of the supports of f_1 and f_2. Choose (Urysohn) $q \in \mathcal{K}_+(G)$ with $q(x) = 1$ for $x \in C$. Choose some tiny $\delta > 0$ (we'll say how tiny later—we could take a vote on this issue if the audience is sufficiently offended by this idea though); I'll tell you now that $\delta < 1$ though.

Set $p = f_1 + f_2 + \delta q$, so $p(x) \geq \delta$ for $x \in C$. The key construction is to define (for $i = 1, 2$) the functions

$$h_i(x) = f_i(x)/p(x) \text{ (if } x \in C)$$
$$= 0 \text{ (if } x \notin C)$$

[before you ask—there were overflow vbox issues]

[$p = f_1 + f_2 + \delta q$

$$h_i(x) = f_i(x)/p(x) \text{ (if } x \in C\text{)}$$
$$= 0 \text{ (if } x \notin C\text{)}$$

]

We need δ to ensure that the h_i are continuous! The sum of h_1 and h_2 is approximately the characteristic function of C. One checks easily that $h_i \in \mathcal{K}_+(G)$ (the support of h_i is closed in C) and $0 \leq h_1 + h_2 \leq 1$. Now continuous with compact support implies uniformly continuous, by the usual argument, so we can let W be a sufficiently small open neighbourhood of the identity such that $|h_i(x) - h_i(y)| < \delta/2$ whenever xy^{-1} or $x^{-1}y \in W$. By replacing W with $W \cap \{w^{-1} : w \in W\}$ we may even ensure that $W = W^{-1}$. Note that V will be our W when we've decided upon a δ.

So now choose any $F \in \mathcal{K}_+(G)$ with support in W. By Lemma 1 we can find α_j and x_j with $p \leq \sum_j \alpha_j F^{x_j}$.

$[p \leq \sum_j \alpha_j F^{x_j}]$

We've just bounded p above by translates of F, and now we can bound the f_i above by translates of F too. First note that $F^{x_j}(t) = 0$ if $t \notin Wx_j$, and for $t \in Wx_j$ we have $|h_i(x_j) - h_i(t)| < \delta/2$ for $1 \leq i \leq 2$. So in either case we have

$$h_i(t)F^{x_j}(t) \leq (h_i(x_j) + \delta/2)F^{x_j}(t).$$

Hence

$$\begin{aligned} f_i = ph_i &\leq \sum_j \alpha_j h_i F^{x_j} \\ &\leq \sum_j \alpha_j (h_i(x_j) + \delta/2)F^{x_j}. \end{aligned}$$

Hence, by definition,

$$(f_i : F) \leq \sum_j \alpha_j (h_i(x_j) + \delta/2).$$

And, because $h_1(x_j) + h_2(x_j) \leq 1$, we deduce

$$(f_1 : F) + (f_2 : F) \leq \sum_j \alpha_j (1 + \delta).$$

$$(f_1 : F) + (f_2 : F) \leq \sum_j \alpha_j (1 + \delta).$$

Now the last thing we chose were the α_j and x_j with $p \leq \sum_j \alpha_j F^{x_j}$, so by letting these vary we deduce from the previous equation

$$(f_1 : F) + (f_2 : F) \leq (1 + \delta)(p : F)$$

and hence (dividing by $(\eta : F)$)

$$\mu_F(f_1) + \mu_F(f_2) \leq (1 + \delta)\mu_F(p)$$

for any $F \in \mathcal{K}_+(G)$ with support in W. You can now presumably see that we're on the right track, because p is approximately $f_1 + f_2$.

In fact $p = f_1 + f_2 + \delta q$, and we deduce (from subadditivity of μ_F) that

$$\begin{aligned}\mu_F(f_1) + \mu_F(f_2) &\leq (1 + \delta)(\mu_F(f_1 + f_2) + \delta\mu_F(q)) \\ &\leq \mu_F(f_1 + f_2) + \delta.R,\end{aligned}$$

with $R = \mu_F(f_1 + f_2) + 2\mu_F(q)$.

$[\mu_F(f_1) + \mu_F(f_2) \leq \mu_F(f_1 + f_2) + \delta.R$, with $R = \mu_F(f_1 + f_2) + 2\mu_F(q).]$

Now unfortunately R depends on F which depends on W which depends on δ, but fortunately you all did the exercise earlier which showed $\mu_F(f) \leq (f : \eta)$, and hence $R \leq (f_1 + f_2 : \eta) + 2(q : \eta)$, which were all chosen before δ. So now choose δ such that $\delta((f_1+f_2 : \eta)+2(q : \eta)) < \epsilon$, let V denote the corresponding open set W, and we're home.

□

Corollary 3. Given $f_1, f_2, \ldots, f_n \in \mathcal{K}_+(G)$ and $\epsilon > 0$ there exists a symmetric open neighbourhood V of the identity in G such that whenever $F \in \mathcal{K}_+(G)$ has support in V, we have

$$\mu_F(\sum_i f_i) \geq \sum_i \mu_F(f_i) - \epsilon.$$

Proof. Induction. □

In the last lecture we proved some lemmas and in this lecture we need one or two more, but we also need to actually prove existence and uniqueness of the Haar integral. There are several ways to do this; I'll use a method that I found in P. J. Higgins' book "An introduction to topological groups" because in my view the uniqueness proof is the least painful out of all the references I've seen (it's still pretty painful though :-( ). We're going to deduce existence and uniqueness of Haar integrals from some Zorn's Lemma argument applied within the real vector space $\mathcal{K}(G)$ of continuous functions with compact support. Here's the abstract linear algebra we'll need to pull this off.

Let V denote a real vector space. A non-empty subset E of V is called *convex* if $v, w \in E$ and $0 \leq \lambda \leq 1$ implies $\lambda v + (1 - \lambda)w \in E$. A subset C of V is called a *cone* if $c \in C$ and $\lambda > 0$ implies $\lambda c \in C$. One checks that for E non-empty, E is a convex cone iff

(i) $v \in E$ and $\lambda > 0$ implies $\lambda v \in E$

(ii) $v, w \in E$ implies $v + w \in E$.

(examples: $(r, 0) \in \mathbf{R}^2$ with $r > 0$ or $r \geq 0$). Finally, we say that a convex cone E is *open* if $w \in E$ and $f \in V$ implies that there's some $\delta > 0$ such that $w + \alpha f \in E$ for all α with $|\alpha| < \delta$. So the $(r, 0)$ examples above aren't open, but (r, s) with $r, s > 0$ would be. Note that the complement of a cone is a cone, but the complement of a convex cone might not be convex.

A subspace H of V is called a *hyperplane* if V/H is 1-dimensional.

[convex cone: $\lambda E = E$ for $\lambda > 0$, and $E + E \subseteq E$]

Note that a Haar integral is a non-zero $\mathbf{R}$-linear map $\mathcal{K}(G) \to \mathbf{R}$ and is hence, up a non-zero constant, determined by its kernel (the functions whose integral is zero). The kernel will be a hyperplane in $\mathcal{K}(G)$ and our existence and uniqueness proofs of Haar integrals will be done via existence and uniqueness of hyperplanes with certain properties.

Proposition. ("Haar integral machine") Say V is a real vector space, E is an open convex cone in V, and W is a subspace of V that doesn't meet E.

(i) There's a hyperplane $H \supseteq W$ such that $H \cap E$ is empty.

(ii) If furthermore $V \backslash E$ (the complement of E) is convex, then H is unique.

Proof. This is elementary, unsurprisingly.

(i) By Zorn's Lemma one can choose a maximal subspace H containing W and missing E and the claim is that it's a hyperplane. Set $D = H + E$. It's easy to check that D is an open convex cone (H and E are convex cones, and E is open). By assumption $H \cap E$ is empty, so $H \cap D$ is also empty.

First I claim that V is the disjoint union of H, D and $-D$. Disjointness is trivial (if $D \cap -D$ was non-empty then use convexity to show $0 \in D$ which is false). The fact that V is the union of H, D and $-D$ follows from maximality: if $v \in V$ with $v \notin H$ then $H + \mathbf{R}v$ intersects E and hence $\mathbf{R}v$ meets $H + E = D$, but $0 \notin D$ so $\pm v \in D$ for some choice of sign.

Now I claim that H is a hyperplane. Note that $H \neq V$ because E is non-empty, so V/H has dimension at least 1. Say $v, w \in V$ generate a 2-dimensional subspace of V/H and let's get a contradiction. Well, $w \notin H$ so (after changing sign if necessary) we may assume $w \in -D$. Similarly (after changing sign of v if necessary) we may assume $v \in D$.

Now consider the line joining v to w; think about it as the image of $[0,1]$. Because D is an open convex cone, one checks easily that the intersection of D with this line is an open interval containing 0 but not 1. Similarly the intersection of the line with $-D$ is a open interval containing 1 but not 0. But D and $-D$ are disjoint, and two disjoint open intervals can't cover a line, so we have $\lambda v + (1-\lambda)w \in H$ for some $0 < \lambda < 1$ and there's a linear relation in V/H between v and w, the contradiction we seek.

(ii) Let E^* denote the complement of E in V; then E^* is assumed convex and is hence a convex cone. Now $0 \notin E$ so $E \cap (-E)$ is empty; let X be the complement of $E \cup (-E)$. Now $X = (E^*) \cap (-E^*)$ so it's a convex cone, and $-X = X$, so X is a vector subspace of V. The argument from (i) (with X replacing H and E replacing D) shows that X is a hyperplane; moreover any subspace of V disjoint from E will be contained in X, so we're home. $\square$

The application is of course the following. Let G be a LCHTG, set $V = \mathcal{K}(G)$, let W be the subspace of V spanned by all functions of the form $f - f^x$ for $f \in \mathcal{K}(G)$ (so everything in L should integrate to zero), set $E = \mathcal{K}_+(G) + W$ (think of E as "everything for which the axioms imply that the integral should be positive"). I claim that the hypotheses of the "Haar integral machine" are satisfied. Let's check these in a second, but let's first observe that if they are, then (i) gives us a Haar integral, and (ii) (if it applies, that is, if E^* is convex) gives us uniqueness up to a positive scalar. For (i) gives us a hyperplane H; let μ denote any $\mathbf{R}$-linear isomorphism $V/H \to \mathbf{R}$. Then clearly μ is linear and translation-invariant; furthermore if $f, g \in \mathcal{K}_+(G)$ then the line from f to g lies within $\mathcal{K}_+(G)$ so doesn't meet H, and hence $\mu(f)$ and $\mu(g)$ have the same sign, so either μ or $-\mu$ is a Haar measure. Conversely any kernel of a Haar integral will contain W and be disjoint from E, so if (ii) applies then there's only one possibility for the kernel.

[W spanned by $f - f^x$; $E = \mathcal{K}_+(G) + W$]

So what is left to do? For existence of a Haar integral, we just need to check the hypotheses of the proposition (that is, that $W \cap E$ is empty and that E is an open convex cone; we've done the work to prove these easily though). For uniqueness up to positive scalar we need to check that the complement of E is convex (we need another lemma to do this). Let's do existence first because it actually helps with uniqueness.

Existence of Haar integral.

To check $W \cap E$ is empty we just have to check $W \cap \mathcal{K}_+(G)$ is empty. This isn't a surprising result, because everything in W should integrate to zero, and nothing in $\mathcal{K}_+(G)$ should. But let's give the proof. Now W is generated by things of the form $f - f^x$; furthermore using the $2f = (|f| + f) - (|f| - f)$ trick we can check that W is generated by things of the form $f - f^x$ for $f \in \mathcal{K}_+(G)$. So if $W \cap \mathcal{K}_+(G)$ is nonempty then we can find $f, f_i \in \mathcal{K}_+(G)$ and $x_i \in G$ with

$$f = \sum_i (f_i - f_i^{x_i}).$$

We rewrite as

$$f + \sum_i f_i^{x_i} = \sum_i f_i$$

and for any $\epsilon > 0$ we use Corollary 3 to find an open neighbourhood V of the identity such that for any $F \in \mathcal{K}_+(G)$ with support in V, we have

$$\mu_F(f + \sum_i f_i^{x_i}) \geq \mu_F(f) + \sum_i \mu_F(f_i^{x_i}) - \epsilon.$$

Now we easily get a contradiction, for choosing F as above (Urysohn), we see

$$\begin{aligned}\sum_i \mu_F(f_i) &\geq \mu_F(\sum_i f_i) \\ &= \mu_F(f + \sum_i f_i^{x_i}) \\ &\geq \mu_F(f) + \sum_i \mu_F(f_i^{x_i}) - \epsilon \\ &\geq (\eta : f)^{-1} + \sum_i \mu_F(f_i) - \epsilon\end{aligned}$$

so if we had chosen ϵ with $0 < \epsilon < (\eta : f)^{-1}$ then we get a contradiction.

All that's left for existence is the check that $\mathcal{K}_+(G) + W$ is an open convex cone. It's clearly a convex cone; the issue is openness. If $f = p+q$ is an arbitrary element of $\mathcal{K}_+(G)+W$, and $k \in \mathcal{K}(G)$ is arbitrary, we need to show $f \pm \lambda k \in \mathcal{K}_+(G) + W$ for $0 < \lambda$ small. Here's how. By Lemma 1 we can bound $|k|$ above by $\sum_i \alpha_i p^{x_i}$ and WLOG not all of the α_i are zero. So

$$\begin{aligned} f \pm \lambda k &\geq p + q - \lambda \sum_i \alpha_i p^{x_i} \\ &= p + q - \lambda \sum_i \alpha_i p - \lambda \sum_i \alpha_i (p^{x_i} - p) \\ &= p(1 - \lambda \sum_i \alpha_i) + q' \end{aligned}$$

with $q' \in W$, so $f \pm \lambda k \in \mathcal{K}_+(G)+W$ if $0 < \lambda < (\sum_i \alpha_i)^{-1}$, and we have proved the existence of the Haar integral.

For uniqueness we need to show that (with the above notation) E^* is convex. The reason we don't yet have enough is that we have "only approximated a function from above"—we now really need to approximate a function in $\mathcal{K}(G)$ uniformly across G. To do this we need the a standard application of the "bump functions" that Urysohn's lemma gives us.

Lemma 4. If G is a LCHTG and $f \in \mathcal{K}_+(G)$ and W is any neighbourhood of the identity in G, then one can find $x_1, x_2, \ldots, x_n$ all in the support of f and $f_1, f_2, \ldots, f_n \in \mathcal{K}_+(G)$ with the support of f_i in Wx_i, and $\sum_i f_i = f$.

Proof. This is easy. First choose a compact neighbourhood N of the identity in W, with interior U. Now the support of f is compact so it's covered by finitely many Ux_i, $1 \leq i \leq n$, with x_i all in the support of f. By Urysohn, there exists $h_i \in \mathcal{K}_+(G)$ which is identically 1 on Nx_i and whose support is contained within Wx_i. Now set $h = \sum_i h_i$ and $f_i(x) = f(x)h_i(x)/h(x)$ for x in the support of f, and $f_i(x) = 0$ otherwise. It's an easy check that this works.

Say $F : G \to \mathbf{C}$ is *symmetric* if $F(x) = F(x^{-1})$ for all x. The following lemma is the last piece of the puzzle.

Lemma 5. (Uniform approximation) If $f \in \mathcal{K}_+(G)$ and $\epsilon > 0$, then there exists some neighbourhood V of the identity in G such that for *every* symmetric $F \in \mathcal{K}_+(G)$ with support contained in V there are real numbers $\alpha_1, \alpha_2, \ldots, \alpha_n \geq 0$ and $x_1, x_2, \ldots, x_n \in G$ such that

$$|f(x) - \sum_i \alpha_i F^{x_i}(x)| < \epsilon$$

for all $x \in G$.

Proof. By uniform continuity we can choose a neighbourhood V of the identity such that $|f(x) - f(y)| < \epsilon/2$ whenever $y \in Vx$ and this V is going to work. Say $F \in \mathcal{K}_+(G)$ is symmetric with support in V. Then of course the support of F^x is within Vx, and one deduces easily that $|f(x) - f(y)|F^x(y) \leq \frac{\epsilon}{2}F^x(y)$ for all $x, y \in G$, so by definition we have (as functions of y)

$$|f(x)F^x - f.F^x| \leq \frac{\epsilon}{2}F^x.$$

$$|f(x)F^x - f.F^x| \leq \frac{\epsilon}{2}F^x \qquad (1)$$

(as functions of y). Now for any $\delta > 0$, by uniform continuity of F, we can find some neighbourhood W of the identity such that

$$|F(y) - F(z)| < \delta$$

for all $y \in Wz$, and hence, for any $x \in G$, $|F^x(y) - F^x(z)| < \delta$ for all $y \in Wz$.

Now by (bump function) Lemma 4, applied to f and W, we write $f = \sum_i f_i$ with $f_i \in \mathcal{K}_+(G)$ and the support of f_i in Wx_i. The same trick as above gives us

$$f_i(y)|F^x(y) - F^x(x_i)| \leq \delta f_i(y)$$

for all $x, y \in G$ (check separately for $y \in Wx_i$ and $y \notin Wx_i$).

A labour-saving observation now is that F is symmetric so $F^x(x_i) = F^{x_i}(x)$, and summing the last equation over i we get

$$|f(y)F^x(y) - \sum_i f_i(y)F^{x_i}(x)| \leq \delta f(y).$$

This latter equation is true for all $x, y \in G$, and hence (by definition)

$$|f.F^x - \sum_i F^{x_i}(x)f_i| \leq \delta f. \qquad (2)$$

Recall now equation 1:

$$|f(x)F^x - f.F^x| \leq \frac{\epsilon}{2}F^x \qquad (1)$$

and we get

$$|f(x)F^x - \sum_i F^{x_i}(x)f_i| \leq \epsilon/2F^x + \delta f \qquad (3)$$

(an inequality of functions of y) for all x. What we did here was used uniform continuity of f and uniform continuity of F^x to get two good approximations for $f.F^x$, and we reaped the consequences.

[$|f(x)F^x - \sum_i F^{x_i}(x)f_i| \leq \epsilon/2F^x + \delta f$ (3)]

A painless way to finish now is to assume the existence of a Haar integral! We have already proved this so it's OK. Apply a Haar integral μ to this last equation (observing that if $\phi \in \mathcal{K}(G)$ then $\phi \leq |\phi|$ and $-\phi \leq |\phi|$, so $|\mu(\phi)| \leq \mu(|\phi|)$) and deduce

$$|f(x)\mu(F) - \sum_i F^{x_i}(x)\mu(f_i)| \leq \epsilon/2\mu(F) + \delta\mu(f).$$

It's an easy check that $\mu(f) \leq (f : F)\mu(F)$ [look at the definition of $(f : F)$ and apply μ] and we deduce

$$|f(x)\mu(F) - \sum_i F^{x_i}(x)\mu(f_i)| \leq (\epsilon/2 + \delta(f : F))\mu(F).$$

Now divide by $\mu(F)$, set $\alpha_i = \mu(f_i)/\mu(F)$, let δ be $\epsilon/(3(f : F))$ and we get

$$|f(x) - \sum_i \alpha_i F^{x_i}(x)| < \epsilon$$

and we have won. □

Corollary 6. Set $E = \mathcal{K}_+(G) + W$ as before. Then, for any $f \in \mathcal{K}(G)$, there exists some $h \in \mathcal{K}_+(G)$ such that for every $\epsilon > 0$, either $f + \epsilon h \in E$ or $f - \epsilon h \in -E$.

Proof. Let C be the support of f; let D be a compact neighbourhood of C. By Urysohn there's $h \in \mathcal{K}_+(G)$ with $h(x) > 2$ for $x \in D$. This h will work. For we can write $f = f_1 - f_2$ with $f_i \in \mathcal{K}_+(G)$, and by the previous Lemma (uniform approximation) both f_1 and f_2 can be uniformly approximated by scalings of translates of any symmetric function with support in some V, which is WLOG symmetric and satisfies $VC \subseteq D$.

The trick now is if $F_0 \in \mathcal{K}_+(G)$ has support in V then $F(x) = F_0(x) + F_0(x^{-1})$ is symmetric with support in V, and we can uniformly approximate f_1 and f_2 using F, and hence we can find $\alpha_1, \ldots, \alpha_n \in \mathbf{R}$ with $|f(x) - \sum_i \alpha_i F^{x_i}(x)| < 2\epsilon$ for all $x \in G$.

We have rigged it so that f and F^{x_i} all have support in D, so we can conclude that

$$|f - \sum_i \alpha_i F^{x_i}| < \epsilon h.$$

Now if $\alpha = \sum_i \alpha_i$ then $k := \alpha F - \sum_i \alpha_i F^{x_i} \in W$ and

$$f - \epsilon h < \alpha F - k < f + \epsilon h.$$

But this implies what we want: if $\alpha \geq 0$ then we've shown $f + \epsilon h > -k \in W$ so $f + \epsilon h \in \mathcal{K}_+(G) + W$, and if $\alpha \leq 0$ then we've shown $f - \epsilon h < -k$ and hence $f - \epsilon h \in -E$. $\square$

Uniqueness of Haar integrals.

As usual W is generated by $f - f^x$, $E = \mathcal{K}_+(G) + W$ and all we need to do is to prove that the complement of E in $V = \mathcal{K}(G)$ is convex. The complement is certainly a cone, so we need to show it's a convex cone, so we need to check that if $f_1, f_2 \in E^*$ (the complement of E) and $f_1 + f_2 \in E$ then we have a contradiction. This is now easy. Write $f = f_1 - f_2$ and apply the previous corollary to deduce that there's some $h \in \mathcal{K}_+(G)$ such that for all $\epsilon > 0$ either $f + \epsilon h \in E$ or $f - \epsilon h \in -E$. But E is open and $f_1 + f_2 \in E$, so there's some $\epsilon > 0$ such that $f_1 + f_2 - \epsilon h \in E$. If $f + \epsilon h \in E$ then $2f_1 \in E$, and if $f - \epsilon h \in -E$ then $2f_2 \in E$, and either one is a contradiction.

We now consider consequences of the existence and uniqueness of Haar integrals, and extend our range of definition somewhat. We established the existence of a Haar integral which, by definition, was impervious to right translations. These are sometimes called "right Haar integrals". We could instead demand that the integral of $f \in \mathcal{K}(G)$ was equal to the integral of the function $g \mapsto f(xg)$ [invariance under left translation]. Such a gadget would then be called a "left Haar integral". But these exist and are unique too:

Theorem. Left Haar integrals exist and are unique up to a positive constant.

Proof. If $f \in \mathcal{K}(G)$ then the function $\tilde{f} : x \mapsto f(x^{-1})$ is also in $\mathcal{K}(G)$ and if we define $\mu_L(f) = \mu(\tilde{f})$ then μ_L is a left Haar integral iff μ is a right Haar integral.

□

Note that the left Haar integral might not be (a positive constant times) the right Haar integral! The moral reason for this is that it's not hard to find a LCHTG G with a subgroup H and $g \in G$ with gHg^{-1} a proper subset of H. If there were a left and right invariant Haar integral on G then a good approximation to the characteristic function of H would have the same measure as a good approximation to the characteristic function of gHg^{-1} which can't happen because gHg^{-1} is "strictly smaller than H".

Exercise: Let G be the matrices $\left(\begin{smallmatrix} a & b \\ 0 & 1 \end{smallmatrix}\right)$ in $\mathrm{GL}_2(\mathbf{R})$. Conjugating by $g := \left(\begin{smallmatrix} 2 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ sends $\left(\begin{smallmatrix} a & b \\ 0 & 1 \end{smallmatrix}\right)$ to $\left(\begin{smallmatrix} a & 2b \\ 0 & 1 \end{smallmatrix}\right)$ so if f is any continuous function on $\mathbf{R}$ with compact support and which is increasing on $(-\infty, 0)$, has $f(0) > 0$, and is decreasing on $(0, \infty)$, then the function F on G defined by $F(\left(\begin{smallmatrix} e^t & b \\ 0 & 1 \end{smallmatrix}\right)) = f(t)f(b)$ (and $F = 0$ if $a < 0$) is continuous with compact support and satisfies $x \mapsto F(x) - F(gxg^{-1}) \in \mathcal{K}_+(G)$, which is enough to show that no bi-invariant measure can exist ($\mu(F) > \mu(gFg^{-1})$ for any right Haar measure).

If μ is a (right) Haar integral on G then let's write

$$\int_G f(x)\mathrm{d}\mu(x)$$

for $\mu(f)$; it's a more suggestive notation. Then we have the following result:

Fubini's Theorem. If G and H are LCHTGs with Haar integrals μ and ν, and $f \in \mathcal{K}(G\times H)$ then

$$\int_G \left(\int_H f(x,y)\mathrm{d}\nu(y)\right)\mathrm{d}\mu(x)$$

and

$$\int_H \left(\int_G f(x,y)\mathrm{d}\mu(x)\right)\mathrm{d}\nu(y)$$

exist, are equal, and are both right Haar integrals on $G \times H$.

Proof. Existence is easy. First, f has compact support so has support within a compact set of the form $C \times D$ (projection of a compact set is compact), so certainly the inner integral $\int_H f(x,y)\mathrm{d}\nu(y)$ exists and, as a function of x, has compact support.

We need to check $\int_H f(x,y)\mathrm{d}\nu(y)$ is continuous as a function of x but this isn't hard (big hint: if f is supported within $C \times D$ and we choose $k \in \mathcal{K}_+(H)$ which equals 1 on D then by uniform continuity we have that x close to x' implies $|f(x,y) - f(x',y)| \leq \epsilon k(y)$ so changing x to x' changes the integral by at most $\epsilon\nu(y)$ which can be made arbitrarily small). It follows easily that both integrals are Haar integrals. To check that they're the same, it suffices to check that they agree on *one* positive function! So check it for $f(x,y) = p(x)q(y)$ with $p \in \mathcal{K}_+(G)$ and $q \in \mathcal{K}_+(H)$, where both integrals are just $\mu(p)\nu(q) \neq 0$.

□

Finally let me talk about extending our range of integrable functions. This is rather formal, really.

Let G be a LCHTG and let U denote all the functions $f : G \to \mathbf{R} \cup \{+\infty\}$ which are pointwise limits of increasing sequences $f_1 \leq f_2 \leq f_3 \leq \ldots$ with $f_n \in \mathcal{K}(G)$. If $f \in U$ then one can check that $\mu(f) := \lim_n \mu(f_n) \in \mathbf{R} \cup \{+\infty\}$ is well-defined and independent of the choice of f_n. Set $-U = \{-f : f \in U\}$, define μ on $-U$ by $\mu(-f) = -\mu(f) \in \mathbf{R} \cup \{-\infty\}$.

Definition. A function $f : G \to \mathbf{R} \cup \{\pm\infty\}$ is *summable* if there exists $g \in -U$ and $h \in U$ with $g \leq f \leq h$ and, crucially,

$$\sup_{g \leq f, g \in -U} \mu(g) = \inf_{h \geq f, h \in U} \mu(h).$$

The common value is defined to be $\mu(f) \in \mathbf{R}$ (note: it can't be infinite).

Note that a summable function certainly doesn't have to be continuous.

Exercise: if $G = \mathbf{R}$ then check that the characteristic function of $[0, 1]$ is summable, and has integral equal to 1 (if the Haar measure is normalised in the usual way). Similarly check that the characteristic function of a point is summable and has integral equal to zero.

If $\mathcal{L}^1(G)$ denotes all the summable functions, then there's a natural "norm" on $\mathcal{L}^1(G)$, defined by $||f|| = \mu(|f|)$ (one can check $|f| \in \mathcal{L}^1(G)$). Unfortunately there are plenty of functions in $\mathcal{L}^1(G)$ with $||f|| = 0$ (for example the characteristic function of a point, if $G = \mathbf{R}$). Say a function f is *null* if $||f|| = 0$.

Definition. $L^1(G)$ is defined to be $\mathcal{L}^1(G)$ modulo the null functions.

One can check that $L^1(G)$ is in fact a real Banach space. In fact more generally, if $1 \leq p < \infty$ one can define $\mathcal{L}^p(G)$ to be the functions $f : G \to \mathbf{R} \cup \{\pm\infty\}$ such that $|f|^p$ is summable, one can define a "norm" on $\mathcal{L}^p(G)$ by $||f||_p = \mu(|f|^p)^{1/p}$ and then let $L^p(G)$ be the quotient of $\mathcal{L}^p(G)$ by the subspace of f with $||f||_p = 0$. It turns out that these are all Banach spaces (this needs a little proof), which are absolutely fundamental to the further development of the theory. Note also that $L^2(G)$ is a real Hilbert space, because one can make sense of

$$\langle f, g \rangle = \int_G f(x)g(x)\mathrm{d}\mu(x)$$

for $f, g \in L^2(G)$. One can also tensor all these spaces with the complexes to get complex Banach spaces and a complex Hilbert space in the usual way. All these spaces are independent of the explicit choice of Haar measure, but the inner product on $L^2(G)$ does depend on the choice (it affects things by a scaling factor).

Convolution (definition below) defines a product on $L^1(G)$; this is not hard to check. To make this work we have to fix a choice of Haar measure μ. Now if $f, g \in L^1(G)$ then we define $f * g \in L^1(G)$ by

$$(f * g)(z) = \int_G f(zy^{-1})g(y)\mathrm{d}\mu(y).$$

One checks that this is defined on $\mathcal{L}^1(G)$, descends to $L^1(G)$, and is associative and norm-non-increasing ($||f * g||_1 \leq ||f||_1||g||_1$).

One can define a measure on G associated to the integral μ; one says that a subset $A \subseteq G$ is *measurable* if its characteristic function χ_A is summable, and one defines $\mu(A) = \mu(\chi_A)$.

3.4: Overview of Pontrjagin duality and Fourier inversion.

I've decided/realised that one simply needs to assume too much measure theory/spectral theory to give a reasonable presentation of this stuff :-(and, given that I do actually want to spend some lectures talking about Tate's thesis, I've decided that it's impossible to give full proofs here (it would probably take 6 or so lectures to go through the details) and hence I may as well just give an overview of results. The original paper by Cartan and Godement ("Théorie de la dualité et analyse harmonique dans les groupes abéliens localement compacts") is a good reference, and it seems to me that the 40-page Chapter 3 of Ramakrishnan–Valenza is, to a large extent, an English translation of this paper (and chapter 2 of Ramakrishnan-Valenza is 30 pages of spectral theory and so on which one needs as prerequisites).

If you want to see a complete presentation of this stuff then, these 70 pages are perhaps one place to look. It is possible to read this stuff, but it would be helpful if you knew e.g. what a Radon measure was and knew some of the basic spectral theory of Banach algebras, and a fair bit of functional analysis too (the Banach–Alaoglu theorem, the Krein–Milman theorem and so on). I don't know if there is a simpler way to get to the results in the cases that we're interested in. I do know a low-level proof of the Fourier Inversion theorem for p-adic fields but we also need this result in an "adelic setting".

I will prove some basic results, and then give precise statements of deeper theorems.

Let G be a topological group (later on it will be locally compact and Hausdorff, of course, but we don't need that yet). The big new assumption now that we *do* need, is that G must be *abelian*. The non-abelian story is much more subtle (it occupied much of Harish-Chandra's mathematical life and there are still plenty of questions left unanswered). Even the abelian case needs some work (c.f. those 70 pages I just mentioned).

So let G be an abelian topological group. Define $\widehat{G}$, the *dual* of G, to be the group of continuous group homomorphisms

$$\chi : G \to S^1$$

with $S^1 = \{z \in \mathbf{C} : |z| = 1\}$ the circle group (remark that if G isn't abelian then $\widehat{G}$ should probably contain some higher-dimensional unitary representations so the non-abelian theory diverges at this point).

[G an abelian topological group; $\widehat{G} = \{\chi : G \to S^1\}$]

One checks easily that $\widehat{G}$ is a group (the product of two continuous group homomorphisms is continuous, and if χ is continuous then so is $g \mapsto \chi(g)^{-1}$). Our first job is to make $\widehat{G}$ into a topological group. There are subtleties here. In the analogous linear theory (topological vector spaces) there is more than one way to topologise the dual space of a topological vector space (look up "Weak topology" on Wikipedia, for example).

Here's how we're going to topologise $\widehat{G}$. At the minute all we need to assume is that G is an abelian topological group. If K is a compact subset of G and V is a neighbourhood of the identity in S^1 then define

$$W(K,V) := \{\chi \in \widehat{G} : \chi(K) \subseteq V\}.$$

Note that 1, the identity character (the one sending all $g \in G$ to $1 \in S^1$) is in all $W(K,V)$. We define a topology on $\widehat{G}$ by letting the $W(K,V)$ be a base of neighbourhoods of 1. Explicitly, a subset U of $\widehat{G}$ is defined to be open iff for all $\psi \in U$ there is K and V (possibly depending on ψ) such that $W(K,V)\psi \subseteq U$.

Lemma. Let G be an abelian topological group. Then the construction above does define a topology on $\widehat{G}$, and moreover $\widehat{G}$ becomes a topological group with respect to this topology.

$[W(K,V) := \{\chi \in \widehat{G} : \chi(K) \subseteq V\}.]$

To check that we've defined a topology on $\widehat{G}$ we first need to check firstly that the empty set and the entire space are open (which just boils down to checking that at least one set of the form $W(K,V)$ exists; for example $K = \{e\}$ for $e \in G$ the identity will do). Next we need to check that an arbitrary union of open sets is open, which is obvious. Finally we need to check that the intersection of two open sets is open, which boils down to checking that $W(K_1,V_1) \cap W(K_2,V_2)$ contains some $W(K,V)$; but this is true because one can set $K = K_1 \cup K_2$ and $V = V_1 \cap V_2$.

To check that multiplication on $\widehat{G}$ is continuous with respect to this topology, we need to ensure that if $\phi\psi = \rho$ then, for all $W(K,V)$ there is $W(K_1,V_1)$ and $W(K_2,V_2)$ with

$$W(K_1,V_1)\phi W(K_2,V_2)\psi \subseteq W(K,V)\rho.$$

This immediately simplifies to

$$W(K_1,V_1)W(K_2,V_2) \subseteq W(K,V)$$

because G is abelian.

Given K and V, we need K_i, V_i with

$$W(K_1, V_1)W(K_2, V_2) \subseteq W(K, V).$$

If we set $K_1 = K_2 = K$ then all that's left is to check that for all neighbourhoods V of the identity in S^1 there exists V_1 and V_2 with $V_1 V_2 \subseteq V$, which is clear because S^1 is itself a topological group! $\square$

Of course we could have done this more explicitly: if V contains $e^{i\theta}$ with $-\epsilon < \theta < \epsilon$ then we could let $V_1 = V_2 = \{e^{i\theta} : |\theta| < \epsilon/2\}$. While we're at it,

Notation. If $0 \leq r < \infty$ then define $N(r) := \{e^{i\theta} : |\theta| < r\}$.

Exercise. Let $G = \mathbf{R}$ be the real numbers, under addition. Prove that the continuous maps $G \to S^1$ are precisely those of the form $r \mapsto e^{irt}$ with $t \in \mathbf{R}$ (hint: consider a neighbourhood $N(\epsilon)$ in S^1 and its pre-image in G; this gives a map $(-\delta, \delta) \to (-\epsilon, \epsilon)$ which is "additive" whenever this makes sense; draw some conclusions). Check that that this identification of $\widehat{\mathbf{R}}$ with $\mathbf{R}$ induces an isomorphism of topological groups $\widehat{\mathbf{R}} = \mathbf{R}$.

Proposition. Say G is an abelian topological group and $\widehat{G}$ is its dual, topologised as above.

(i) If G is discrete (that is, if all subsets are open) then $\widehat{G}$ is compact.

(ii) If G is compact then $\widehat{G}$ is discrete.

Proof.

(i) Consider $\widehat{G}$ as a subset of $\mathrm{Hom}(G, S^1)$, the (arbitrary set-theoretic) maps from G to S^1. This latter space is just $\prod_{g \in G} S^1$; give it the product topology (reminder: a basis for the product topology, when considering an infinite product, are the subsets which are products of open sets such that all but finitely many of the prodands are the entire space).

First I claim that $\widehat{G}$ is a closed subset of $\mathrm{Hom}(G, S^1)$; this is because its complement is clearly open, as if $\chi(ab) \neq \chi(a)\chi(b)$ then one can choose neighbourhoods V_x of $\chi(x)$ for $x \in \{a, b, ab\}$ with $V_{ab} \cap V_a V_b = \emptyset$.

Next I claim that the subspace topology on $\widehat{G}$ is the compact-open topology. The compact subsets of G are just the finite subsets, and with this in mind it's easy to check that a set is open in one topology iff it's open in the other (both topologies give $\widehat{G}$ the structure of a topological group and this reduces the question to one about neighbourhoods of the identity, which just follows from the definitions).

Finally I claim that this does it, and this is because Tychanov's theorem says that a product of compact spaces is compact, so $\prod_g S^1$ is compact, and a closed subspace of a compact space is compact.

(ii) If G is compact then I claim that subset $\{1\}$ containing only the trivial character is an open subset of $\widehat{G}$, and because $\widehat{G}$ is a topological group this will suffice to prove that $\widehat{G}$ is discrete. To check that it's compact we need to show that $W(K, V) = \{1\}$ for some K and V; take $K = G$ and $V = N(\epsilon)$ for any $\epsilon < \pi/3$. If $\chi(G) \subseteq N(1)$ then $\chi(G)$ is a subgroup of $N(\epsilon)$, but the only subgroup of S^1 contained in $N(\epsilon)$ is $\{1\}$; hence $W(G, N(\epsilon)) = \{1\}$ and we're done. $\square$

I will now start stating things without proof.

Theorem. If G is locally compact Hausdorff, then so is $\widehat{G}$.

I have seen a low-level proof of this, and a more abstract one. The low-level proof (which is long, but completely elementary—it could be a long example sheet question, with hints) goes as follows.

[**Theorem.** If G is locally compact Hausdorff, then so is $\widehat{G}$.]

One checks firstly that if K is a compact neighbourhood of the identity in G then blah $W(K, \overline{N(\pi/6)})$ is a compact neighbourhood of the identity in $\widehat{G}$ (a dull check), and secondly that this suffices (which boils down to checking that as K shrinks, $W(K, \overline{N(\pi/6)})$ gives a basis of neighbourhoods of 1).

The higher-level proof (which one needs later on in the theory anyway) proceeds by first introducing the *Fourier transform* of $f \in L^1(G)$. Notational note: from now on, $L^p(G)$ will *always* denote the complex L^p-functions, rather than the real-valued ones, so it's the thing I originally called $L^p(G)$, tensored over $\mathbf{R}$ with $\mathbf{C}$.

Fix a Haar measure on G. If $f \in L^1(G)$ then define $\widehat{f} : \widehat{G} \to \mathbf{C}$ by

$$\widehat{f}(\chi) = \int_G f(y)\overline{\chi(y)}\mathrm{d}y.$$

Note that $f \in L^1(G)$, and $|f(y)\chi(y)| = |f(y)|$, and it's easy to check that the integrand is also in $L^1(G)$ (do it!). In particular the integral makes sense. We call $\widehat{f}$ the *Fourier transform* of f.

Example: if $G = \mathbf{R}$ and $f \in L^1(G)$, and if we identify $\widehat{G}$ with $\mathbf{R}$ by associating the real number r with the character $x \mapsto e^{ixr}$, and if we use the usual Lesbesgue measure as our Haar measure, then we see that

$$\widehat{f}(r) = \int_{\mathbf{R}} f(x)e^{-irx}\mathrm{d}x$$

which is the definition of Fourier transform that I learned as an undergraduate.

$[\widehat{f}(\chi) = \int_G f(y)\overline{\chi(y)}\mathrm{d}y.]$

If however one identifies $\widehat{G}$ with $\mathbf{R}$ by identifying r with the character $x \mapsto e^{2\pi i x r}$ then one gets the definition of the Fourier transform which is used at the top of the Wikipedia page about Fourier transforms. Finally, if one sticks to $x \mapsto e^{ixr}$ but uses the Haar measure which is $\sqrt{2\pi}$ times Lesbesgue measure, then one gets a third way of normalising things, which according to Wikipedia is another popular choice. Which choice you prefer depends on why you're taking Fourier transforms, but the point of this discussion is that all three choices are covered by our definition.

Back to G locally compact and abelian. For every $f \in L^1(G)$ we get a function $\widehat{f} : \widehat{G} \to \mathbf{C}$ (pedantic remark that even though f "isn't a function" because two functions which differ on a null set are the same element of L^1, $\widehat{f}$ really is a function).

Define the *transform topology* on $\widehat{G}$ to be the weakest topology that makes every $\widehat{f}$ continuous. A computation which is basically elementary (if you know that the continuous functions with compact support are dense in $L^1(G)$ and that $L^1(G)$ is a Banach space, something I didn't prove) but long shows that the transform topology coincides with the compact-open topology (note that local compactness here is essential for this strategy even to make sense, as we used a Haar measure). So we get another proof of G-locally-compact-implies-$\widehat{G}$-locally-compact if we check that the transform topology is locally compact, which follows from Gelfand's theory of commutative Banach algebras, applied to $L^1(G)$.

This latter approach (via the Fourier transform) might seem heavy-handed, but in fact all of these techniques, and more, seem to be needed later on anyway.

The next step in the theory, at least in the development I've seen, is to consider an arbitrary Radon measure $\widehat{\mu}$ on $\widehat{G}$ with the property that $\widehat{\mu}(\widehat{G}) \leq 1$ (note: Haar measure may well not have this property! We are not demanding that μ is invariant under right translations). For such a measure we define its Fourier transform $T_{\widehat{\mu}}$ to be the function $G \to \mathbf{C}$ such that

$$T_{\widehat{\mu}}(y) = \int_{\widehat{G}} \chi(y) \mathrm{d}\widehat{\mu}(\chi)$$

and in some sense the crucial result seems to me to be an intrinsic characterisation of these functions on G; the functions $T_{\widehat{\mu}}$ that arise in this way are precisely the functions which are are essentially bounded by 1 and are "of positive type" (see below). The argument needs some graduate-level functional analysis and measure theory,

in the sense that it needs results which seem to be standard but which I didn't see in my undergraduate courses on functional analysis and measure theory).

Once one has all this, one can prove the first form of the Fourier inversion formula. Here G is an abelian LCHTG. First a definition. Say that $\phi : G \to \mathbf{C}$, continuous and bounded, is of *positive type* if for any $f \in \mathcal{K}(G)$ we have

$$\int_G \int_G \phi(s^{-1}t) f(s) \mathrm{d}s \overline{f(t)} \mathrm{d}t \geq 0.$$

Fourier inversion formula (first form).

There exists a Haar integral $\widehat{\mu}$ on $\widehat{G}$ with the following property: If $f \in L^1(G)$ with Fourier transform $\widehat{f} : \widehat{G} \to \mathbf{C}$, and if f is furthermore a $\mathbf{C}$-linear combination of functions of positive type, then

$$f(y) = \int_{\widehat{G}} \widehat{f}(\chi)\chi(y) \mathrm{d}\widehat{\mu}(\chi).$$

$$f(y) = \int_{\widehat{G}} \widehat{f}(\chi)\chi(y)\mathrm{d}\widehat{\mu}(\chi).$$

This is hard work. If one could interchange the integrals on the right hand side then it might perhaps be easier, but the problem is that $\int_G \chi(y)\overline{\chi}(t)\mathrm{d}y$ probably won't converge. I would almost certainly make a fool of myself were I to try and summarise the 16-page proof in Ramakrishnan-Valenza.

As a consequence of the Fourier inversion formula, and I know of no simple proof of this statement, we get

Theorem. If G is an abelian LCHTG and $z \in G$ is not the identity character, then there exists $\chi \in \widehat{G}$ with $\chi(z) \neq 1$.

Proof. If no such χ exists, then for every $f \in L^1(G)$ we would have $\widehat{f} = \widehat{f^z}$. Hence for every f for which the Fourier inversion formula applies, we would have $f = f^z$. But by Hausdorffness we can find a neighbourhood U of the identity with $U \cap Uz$ empty.

We next find a neighbourhood V with $V^2 \subseteq U$ and V symmetric; finally we observe that if ϕ is real-valued supported in V and $\phi(1) = 1$ then $f := \phi * \tilde{\phi}$ (with $\tilde{\phi}(g) = \phi(g^{-1})$) is of positive type but has support disjoint from that of f^z, a contradiction. □

Corollary. If G is an abelian LCHTG then the obvious map $G \to \hat{\hat{G}}$ is injective. □

It will of course turn out that if G is an abelian LCHTG then $G \to \hat{\hat{G}}$ is an isomorphism. But we have used analysis (rather than topology) to prove injectivity, and in particular we used that G was locally compact. If G is an arbitrary abelian topological group then one can still make sense of $\hat{\hat{G}}$ but I don't know, and very much doubt, if $G \to \hat{\hat{G}}$ is bijective (or even injective) in this generality; consider the case of double-duality of a vector space for some analogue of this—$V = V^{**}$ iff V is finite-dimensional.]

So from now on let's say G is an abelian LCHTC. I'll explain how the theory can be developed.

Next one checks that $G \to \hat{\hat{G}}$ has the property that the induced map from G to its image (with the subspace topology) is a homeomorphism onto a closed subspace (this argument is elementary).

Now one checks that for $f, g \in L^1(G)$ we have $\widehat{f * g} = \hat{f}\hat{g}$. This is just an unravelling of things (once one has realised that $*$ maps $L^1(G) \times L^1(G)$ to $L^1(G)$, which can be proved using Fubini: in fact $||f * g||_1 \leq ||f||_1 ||g||_1$.

If $f \in L^2(G)$ then set $\tilde{f}(x) = \overline{f(x^{-1})}$, so $\tilde{f} \in L^2(G)$, and define $h = f * \tilde{f}$ (so h is integrable and of positive type: this is some analogue of the fact that if A is a real matrix then A^tA is positive semidefinite). Now unravelling the definitions we see that if μ and $\widehat{\mu}$ are Haar measures normalised so that the first form of Fourier inversion holds, then

$$\begin{aligned}\int_G |f(x)|^2 \mathrm{d}\mu(x) &= h(1) \\ &= \int_{\widehat{G}} \widehat{h}(\chi) \mathrm{d}\widehat{\mu}(\chi) \\ &= \int_{\widehat{G}} |\widehat{f}(\chi)|^2 \mathrm{d}\widehat{\mu}(\chi)\end{aligned}$$

(the second $=$ is Fourier inversion, the other two are elementary). So the integrals of $|f|^2$ and $|\widehat{f}|^2$ coincide. This is the first form of the Plancherel theorem. But in fact, by a density argument one can now conclude

Plancherel Theorem.

For G an abelian LCHTG one can extend the Fourier transform uniquely to an isometric isomorphism

$$\widehat{\ } : L^2(G) \to L^2(\widehat{G}).$$

One has to be careful here: I am not asserting that if $f \in L^2(G)$ then the original definition of $\widehat{f}$ that I gave makes sense. All I'm saying is that the map, which we defined using an integral, extends to give a map on all of $L^2(G)$ in some way.

From this one gets, without too much trouble,

Pontrjagin duality.

If G is an abelian LCHTG then the obvious map $G \to \widehat{\widehat{G}}$ is a group-theoretic isomorphism and a topological homeomorphism.

And then finally this leads us to a cleaner version of the Fourier Inversion theorem:

Fourier Inversion Theorem (final form).

Fix Haar measures on G and $\widehat{G}$. Then there exists a positive real constant $c > 0$ such that if $f \in L^1(G)$ is continuous, and $\widehat{f} \in L^1(\widehat{G})$, and if we identify G with $\widehat{\widehat{G}}$, then

$$\widehat{\widehat{f}}(x) = cf(x^{-1})$$

for all $x \in G$. Furthermore, for any choice of Haar measure on G there's a unique Haar measure on $\widehat{G}$ which ensures $c = 1$.

We haven't given a complete proof of this. I do know complete proofs in certain explicit cases. I currently don't know whether the proofs I know suffice to cover the instances needed for Tate's thesis, but I'll probably find out within a few weeks.

Case studies.

1) $G = \mathbf{R}$. Here it's not so hard to give a proof. The trick is to introduce the following rapidly-decreasing functions: for $t > 0$ and $x \in \mathbf{R}$ fixed, consider $\phi(y) = e^{iyx - t^2 y^2}$. One explicitly computes the Fourier transform of this (good clean fun) and now, instead of integrating $\widehat{f}(y)e^{iyx}$ with respect to y, one integrates $\widehat{f}(y)\phi(y)$. The trick is that this is easily shown to be the integral of $f(r)\widehat{\phi}(r)$. Now one lets t tend to zero from above, and uses the dominated convergence theorem (and the fact that the Fourier transform is continuous, which also needs to be checked, and which also follows from the dominated convergence theorem).

2) $G = S^1$ with its usual topology, so $\widehat{G} = \mathbf{Z}$ with the discrete topology. In this case, the Fourier inversion theorem simply says that for a periodic function f on $\mathbf{R}$ (that is, a function on S^1), the Fourier series of f converges to f. This just boils down to the statement that if z is the inclusion $S^1 \to \mathbf{C}$ then the functions $z^n : n \in \mathbf{Z}$ is an orthonormal basis for the Hilbert space $L^2(S^1)$. Orthonormality is easy, and checking that the functions give a basis is just a standard application of the Stone-Weierstrass theorem (polynomials in z and $1/z$ separate points, and $\overline{z} = 1/z$ on S^1). Orthonormality also gives Plancherel's theorem (which is called Parseval's theorem in this context; Parseval was 1799 and thinking about Fourier series, Plancherel was 1910).

3) $G = \mathbf{Q}_p$ or k_P: we'll come back to these. In some sense they're much easier (in the sense that you don't have to remember what the dominated convergence theorem or the Stone-Weierstrass theorem are!). I'll treat these cases carefully in the next chapter.

Chapter 4. Local zeta functions.

The reference for this is chapter 2 of Tate's thesis. Throughout this section K will be a field which is either the reals, an algebraic closure of the reals (you can think "the complexes" but perhaps a more pedantic way of thinking about it is "the complexes except that there is no way of distinguishing between the two square roots of -1"), or a finite extension of $\mathbf{Q}_p$ for some p. In the global applications, K will be the completion of a number field k with respect to a norm. There is another case where everything in this section applies, and that is the "equicharacteristic case", that is, $K = \mathbf{F}_q((t))$, where $\mathbf{F}_q$ is a finite field with q elements and $\mathbf{F}_q((t))$ is the field obtained by adjoining $1/t$ to the integral domain $\mathbf{F}_q[[t]]$ of power series. Personal preferences mean that I will stick to the number field case, rather than the function field case, later on, but one might want to bear in mind that there are no obstructions to making this sort of thing work in the function field case.

In all cases ($K = \mathbf{R}$, $K \cong \mathbf{C}$, $K = k_P$ finite) we have got a canonical equivalence class of norms on K, and we have seen that K is a locally compact abelian group under addition (in the p-adic case the crucial observation was that the residue class field was compact). Now here's a completely wacky construction: we are going to single out a canonical norm in each equivalence class. Here's the idea. Choose a random Haar integral μ on K. For $\alpha \in K^\times$ consider the function $\mu_\alpha : \mathcal{K}_+(K) \to \mathbf{R}$ defined by

$$\mu_\alpha(f) = \mu(x \mapsto f(\alpha x)).$$

In words, μ_α is Haar measure, "stretched" by multiplication by α. One checks easily that μ_α is well-defined and is also a Haar measure, and hence $c\mu_\alpha = \mu$, where $c = c(\alpha)$ is a positive real number. One checks easily that $c(\alpha)$ does *not* depend on the choice of μ—it is truly intrinsic. The reason we didn't see this structure before is that we're not just thinking of K as an additive group, we're using its ring structure.

Let's write $|\alpha| := c(\alpha)$, and define $|0| = 0$.

$[c(\alpha)\mu_\alpha = \mu]$

This choice is a *canonical* choice of norm on K. One does need to check it's a norm—but this is easy by a brute force calculation, which I'll now do: in each case we see that we're reconstructing the norm I've already put on these fields—but now we see that the norm I put on them is "the natural norm".

1) If $K = \mathbf{R}$ then (think about a good approximation to the characteristic function of $[0, 1]$) $|\alpha|$ is just the usual absolute value of α.

2) If $K = \mathbf{C}$ then (think about the characteristic function of a square) we see $|x + iy| = x^2 + y^2$, so our canonical norm is the square of the usual norm (and hence doesn't satisfy the triangle inequality, which is the unique reason that I didn't make the triangle inequality an axiom earlier).

3) If $K = \mathbf{Q}_p$ then let's compute $|p|$. Well, if χ is the characteristic function of $\mathbf{Z}_p$, the integers of $\mathbf{Q}_p$, then χ really is continuous with compact support. Now $\mathbf{Z}_p$ is the disjoint union of $a + p\mathbf{Z}_p$ for $a = 0, 1, 2, \ldots, p-1$, so by finite additivity and translation-invariance of Haar measure we see that if ψ is the characteristic function of $p\mathbf{Z}_p$ then $p\mu(\psi) = \mu(\chi)$, so $\mu_p(\psi) = \mu(\chi) = p\mu(\psi)$ and hence that $|p| = p^{-1}$. So in fact the canonical norm on $\mathbf{Q}_p$ is just the usual p-adic norm, normalised the way I normalised it.

4) More generally (easy check) if $\pi \in k_P$ is a uniformiser and q is the size of the residue field A/P (A the global integers of the number field k), then we showed that the residue field of k_P has size q (residue fields don't change under completion) and one checks easily that the canonical norm sends π to $1/q$.

5) [optional extra] $K = \mathbf{F}_q((t))$. Then again we see that the index of $t\mathbf{F}_q[[t]]$ in $F_q[[t]]$ is q, so multiplication by t is making things q times smaller, so $|t| = q^{-1}$—again the norm of a uniformiser is the reciprocal of the size of the residue field.

4.1: the dual of $(K, +)$.

The next thing we'll do is to compute $\widehat{G}$, where $G = (K, +)$. As ever in this section, K is either the reals, the complexes, a completion of a number field at a prime ideal, or, if we're feeling adventurous, a the field of fractions of a power series ring over a finite field. Let's call these things "local fields" for simplicity, and let's always endow them with their canonical norms.

Theorem. If $G = (K, +)$ is a local field, considered as a group under addition, then $\widehat{G}$ is isomorphic to $(K, +)$ (not in a particularly canonical way, mind).

Remark. The case $K = \mathbf{R}$ was an exercise earlier.

Theorem. If $G = (K, +)$ is a local field, considered as a group under addition, then $\widehat{G}$ is isomorphic to $(K, +)$.

Remark. I am going to be lazy and give Tate's proof, which appears to me to assume (a consequence of) Pontrjagin duality, but which works for an arbitrary locally compact complete normed field (so, for example, it works for a power series field over a finite field). On the example sheet I'll give an explicit proof when $K/\mathbf{Q}_p$ is finite.

Proof. First let me *assume* that $\widehat{K} \neq 0$ (of course here K is considered as a group under addition). We'll check this later in a case-by-case way, although if you believe Pontrjagin duality then it's obvious because $\widehat{K} \neq 0$ implies $K = 0$. Anyway, let's fix once and for all a non-zero element χ of $\widehat{K}$, and later on I'll write one down explicitly just to prove one exists.

Now consider the map $i : K \to \widehat{K}$ (which depends on χ), defined by letting $i(\lambda)$ be the character $x \mapsto \chi(\lambda x)$ (so again we're crucially using both the additive and multiplicative structure of K). It's easily checked that $i(\lambda) \in \widehat{K}$ and that the induced map $i : K \to \widehat{K}$ is a group homomorphism. Injectivity is also easy: if $(i(\lambda))(x) = 1$ for all $x \in K$ then χ is trivial on λK which is impossible if $\lambda K = K$, because χ is non-trivial, so λ had better not have an inverse.

It's slightly more delicate to finish the job. We'll follow Tate and again assume Pontrjagin duality. A consequence of this duality is that one can show that the "annihilator" construction, sending a subgroup X of an abelian LCHTG G to the subgroup of $\widehat{G}$ consisting of characters which vanish on X, induces an order-reversing bijection between the closed subgroups of G and of $\widehat{G}$.

Apply this to the closure of $i(K) \subseteq \widehat{K}$ and we observe that the corresponding closed subgroup X of K must be be contained in the set $\{x \in K : (i(\lambda))(x) = 1 \forall \lambda \in K\}$ but this set is easily checked to be $\{1\}$. Hence the image of K is dense in $\widehat{K}$.

Next I claim that the map $i : K \to \widehat{K}$ is a homeomorphism onto its image. To check this we need to remember the definition of the topology on $\widehat{K}$: a general neighbourhood of the origin was given by $W(L, V)$ with $L \subseteq K$ compact and V a neighbourhood of the identity in S^1. So what we have to do is to first check that each $W(L, V)$ contains an $i(B(0, \epsilon))$ (the open ball centre zero radius ϵ in K), and conversely that each $i(B(0, \epsilon))$ contains $i(K) \cap W(L, V)$ for some L, V. Both of these are easy; I'll do the slightly harder of the two, which is the latter one. Given $\epsilon > 0$ we need to come up with with L and V such that if $\lambda \in K$ and $\chi(\lambda L) \subseteq V$ then $|\lambda| < \epsilon$, and we do this thus. Choose $k \in K$ with $\chi(k) \neq 1$. Let L be a huge closed disc centre 0 radius M (these are compact, as is easily checked), with the property that $|k| < \epsilon M$, and let V be any open neighbourhood of the identity in S^1 such that $\chi(k) \notin V$. Then for $\lambda \in K$, if $|\lambda| \geq \epsilon$ then $k \in \lambda L$ so $i(\lambda) \notin W(L, V)$ which is what we want.

Finally let's show that $i(K)$ is a closed subspace of $\hat{K}$; this will do us because we already know that it's dense. Because any compact set in K is bounded, it's easy to check that the identity in $\hat{K}$ has a countable basis of neighbourhoods (this isn't logically necessary, I don't think, but it's psychologically satisfying for what follows). For example if C_M denotes the closed disc centre zero radius M and V_M is $\{e^{i\theta} : |\theta| < 1/M\}$ then $N_M := W(C_M, V_M)$ will do. So now choose an arbitrary $\psi \in \hat{K}$. For each integer $M \geq 1$ choose $x_M \in K$ with $i(x_M) \in \psi N_M$ (we can do this by density of the image of i). It's easily checked (it's an argument similar to the one showing i was a homeo onto its image, but it seems to me to be not quite a formal consequence of what we already have) that the x_M form a Cauchy sequence, so $x_M \to x \in K$, and we have $i(x_M) \to \psi$ and $i(x_M) \to i(x)$, so by Hausdorffness we have $\psi = i(x)$ and we're home. □

Actually that's not really the end of the proof because I still need to exhibit a non-zero $\chi \in \hat{K}$. Let's do this on a case-by-case basis.

1) $K = \mathbf{R}$. Then define $\chi(x) = e^{-2\pi i x}$.

2) $K = \mathbf{Q}_p$. Then, by the structure theorem for elements of $\mathbf{Q}_p$ we can write any $k \in K$ as $k = \sum_{n=-N}^{\infty} a_n p^n$ with $a_n \in \{0, 1, 2, \ldots, p-1\}$. Set $q(k) = \sum_{n=-N}^{-1} a_n p^n \in \mathbf{Q}$ (so $q(k) = 0$ iff $k \in \mathbf{Z}_p$). It's an easy check that q is a group homomorphism $\mathbf{Q}_p/\mathbf{Z}_p \to \mathbf{Q}/\mathbf{Z}$ and hence that $\chi(k) = e^{2\pi i q(k)}$ will work.

3) K a finite extension of $K_0 := \mathbf{Q}_p$ or $\mathbf{R}$; then the trace map T_{K/K_0} is an additive map $K \to K_0$ and it's surjective (it's multiplication by $[K : K_0]$ on K_0), so we take this map and then just compose it with the relevant map coming from (1) or (2).

Note that if we write $\chi(y) = e^{2\pi i \Lambda(y)}$ with $\Lambda : K \to \mathbf{R}/\mathbf{Z}$ then our map i is just $(i(x))(y) = e^{2\pi i \Lambda(xy)}$.

4) $K = \mathbf{F}_q((t))$. Then "coefficient of t^{-1}" is a surjection $K \to \mathbf{F}_q$, and $\mathbf{F}_q$ is just a finite-dimensional vector space over $\mathbf{F}_p$, so choose a non-zero linear map $\mathbf{F}_q \to \mathbf{F}_p$ (if we want to fix one then we should use the trace map—but if you don't know about separable extensions it might not be immediately clear to you that this is non-zero), and finally $x \mapsto e^{2\pi i \tilde{x}/p}$ gets you from $\mathbf{F}_p$ to S^1, where $x \mapsto \tilde{x}$ is a lifting $\mathbf{Z}/p\mathbf{Z} \to \mathbf{Z}$.

It's convenient, but not essential, to fix a non-zero character of K once and for all; in cases (1)–(3) above I've written down precisely one character, so let's always use this one. Note that this is not a canonical choice however; it seems to me that K and $\widehat{K}$ are not canonically isomorphic.

While we're here, let's fix a choice of Haar integral on K. If $K = \mathbf{R}$ then let's choose the obvious one—the one such that the integral of the characteristic function of $[0, 1]$ is 1. If $K = \mathbf{C}$ then let's choose twice the obvious

one—the one such that the integral of the characteristic function of $[0,1] \times [0,1]$ is 2 (unsurprisingly, this 2 is related to the fact that our norm on $\mathbf{C}$ is the square of the usual norm).

If K is a finite extension of $\mathbf{Q}_p$ then we do something perhaps a bit more surprising. If K has degree n over $\mathbf{Q}_p$ then the integers R of K are isomorphic to $(\mathbf{Z}_p)^n$ as a $\mathbf{Z}_p$-module (this is elementary if you know that finitely-generated torsion-free modules over a PID are free) and one can define the *discriminant* in the usual way: choose a $\mathbf{Z}_p$-basis $\{e_1, e_2, \ldots, e_n\}$ for R, and define an $n \times n$ matrix A_{ij} whose (i,j)th entry is the trace of $e_i e_j$. The determinant of this matrix is in $\mathbf{Z}_p$ (easy) and generates an ideal called the *discriminant ideal* of $K/\mathbf{Q}_p$; it is a non-zero ideal (this is a standard fact from field theory, coming from separability) and is well-defined independent of all choices (easy). Say the discriminant ideal is $p^m\mathbf{Z}_p$. Let's define our Haar integral on K by letting the integral of the characteristic function of R be $p^{-m/2} \in \mathbf{R}$. [One can check that

if $k/\mathbf{Q}$ is a number field then the discriminant of $k_P/\mathbf{Q}_p$ will be $\mathbf{Z}_p$ if P is unramified in k, and in particular if you regard k as fixed then all but finitely many of its P-adic completions will have the property that their discriminant ideals will be $\mathbf{Z}_p$.]

Why are we labouring over these choices? Well, we have fixed a choice of Haar measure on K, and we have fixed an isomorphism $K \to \hat{K}$, so we get an induced Haar measure on $\hat{K}$, and we're now in a position to apply Fourier transforms twice. Recall that if f is continuous and $f \in L^1(K)$ with $\hat{f} \in L^1(\hat{K}) = L^1(K)$ then we know that $\hat{\hat{f}}(x) = cf(-x)$ where c is some constant depending only on our choices of Haar measure.

Proposition. With the choices we made, $c = 1$.

Non-proof. It suffices to check for just one function, and we'll have to compute loads of Fourier transforms quite soon, so I'll postpone this. We don't need this result at all, it's just psychologically satisfying.

4.2: The dual of $(K^\times, \times)$.

Let K be as usual. Actually not really interested in the Pontrjagin dual of $K^\times$, we're much more interested in the continuous group homomorphisms $K^\times \to \mathbf{C}^\times$, that is we are dropping the unitary assumption on our characters in this section. Let me call a continuous map $G \to \mathbf{C}^\times$ a *quasi-character* of G. The main results we need here are rather easy to prove. Let U denote $\{x \in K^\times : |x| = 1\}$—the *units* of K. Say that a quasi-character is *unramified* if it's trivial on U.

Lemma. The unramified quasi-characters $c : K^\times \to \mathbf{C}^\times$ are all of the form $\lambda \mapsto |\lambda|^s$ for s complex. If $K = \mathbf{R}$ or $\mathbf{C}$ then s is uniquely determined, but if K is p-adic or $\mathbf{F}_q((t))$ then s is only determined modulo $\frac{2\pi i}{\log(q)}\mathbf{Z}$ with q the size of the residue field. In either case, however, the unramified quasi-characters are naturally a 1-dimensional complex manifold.

Proof. Recall $|\lambda|^s$ means $e^{s.\log(|\lambda|)}$. Everything is immediate once we observe that $K^\times/U$ can be explicitly computed as the image $|K^\times|$ of the norm map in $\mathbf{R}_{>0}$. If $K = \mathbf{R}$ or $\mathbf{C}$ then $|.|$ gives an isomorphism $K^\times/U \to \mathbf{R}_{>0}$, and if K is non-arch then $|.| : K^\times/U \to q^{\mathbf{Z}}$ is an isomorphism. Assuming that you know that the continuous group homomorphisms $\mathbf{R}_{>0} \to \mathbf{C}^\times$ are all of the form $x \mapsto x^s$ then we're home. □

For simplicity we now fix a continuous splitting of the map $|.| : K^\times \to |K^\times|$, that is, we choose a subgroup V of $K^\times$ with the property that every element of $K^\times$ is uniquely of the form uv with $u \in U$ and $v \in V$. Again it's impossible, in general, to do this naturally.

If $K = \mathbf{R}$ or $\mathbf{C}$ then we can just let V be the positive reals. If K is non-arch then let's choose a uniformiser $\pi \in K$ and let V be $\pi^{\mathbf{Z}}$; this is easily checked to work. Now $K^\times = U \times V$ so $\mathrm{Hom}(K^\times, \mathbf{C}^\times) = \mathrm{Hom}(U, \mathbf{C}^\times) \times \mathrm{Hom}(V, \mathbf{C}^\times)$.

Notation: for $\alpha \in K^\times$, write $\alpha = \tilde{\alpha} v$ for $\tilde{\alpha} \in U$ and $v \in V$. This clearly depends on our choice of V but we'll only use this notation temporarily.

Corollary. The quasi-characters $c : K^\times \to \mathbf{C}^\times$ are all of the form $\alpha \mapsto \psi(\tilde{\alpha}).|\alpha|^s$ with $s \in \mathbf{C}$ and ψ a character of U.

Proof. This is just an explicit rephrasing of the statement $\mathrm{Hom}(K^\times, \mathbf{C}^\times) = \mathrm{Hom}(U, \mathbf{C}^\times) \times \mathrm{Hom}(V, \mathbf{C}^\times)$. □

Note that that corollary gives the group of all quasicharacters the structure of a 1-dimensional complex manifold (again given by the s variable; U is compact and we regard $\mathrm{Hom}(U, \mathbf{C}^\times) = \mathrm{Hom}(U, S^1) = \hat{U}$ as discrete). Pedants might like to check that this

complex structure is independent of the choice of V. In fact ψ is well-defined independent of the choice of V (it's just $c|U$) and s is well-defined up to the $2\pi i/\log(q)$ ambiguity mentioned earlier.

Explicitly, what is happening is that if $\hat{U}$ is the set of characters of U, then the quasi-characters of $K^\times$ are just one copy of either $\mathbf{C}$ or $\mathbf{C}/\frac{2\pi i}{\log(q)}\mathbf{Z}$ for each element of $\hat{U}$, the dictionary being that if ψ is any character of U and c is any quasi-character of $K^\times$ which restricts to ψ, then the connected component of c in the manifold of quasi-characters is just $c.|.|^s$ for $s \in \mathbf{C}$.

Let's do examples to see what these manifolds look like.

[$c : K^\times \to \mathbf{C}^\times$ are all of the form $\alpha \mapsto \psi(\tilde{\alpha}).|\alpha|^s$ with $s \in \mathbf{C}$ and ψ a character of U.]

The easiest example to think about is $K = \mathbf{R}$; then $K^\times = \pm\mathbf{R}_{>0}$, $U = \{\pm 1\}$, and a quasicharacter of $K^\times$ is just a sign (where we send -1) and a complex number (where we send the positive reals), and the complex structure is given by the complex number; the space is just two copies of the complexes. If $K = \mathbf{C}$ then $K^\times = S^1 \times \mathbf{R}_{>0}$, $U = S^1$, and the characters of S^1 are just $\mathbf{Z}$, so here the quasicharacters are countably infinitely many copies of the complex plane, indexed by the integers. If K is an algebraic closure of $\mathbf{R}$ then we get the same thing, but where $\mathbf{Z}$ is replaced by an infinite cyclic group.

If $K = \mathbf{Q}_p$ then we get a cylinder $\mathbf{C}/(\frac{2\pi i}{\log(p)}\mathbf{Z})$ for each primitive Dirichlet character $(\mathbf{Z}/p^n\mathbf{Z})^\times \to \mathbf{C}^\times$ and in the general non-arch case the picture is a generalisation of this.

Let me make two definitions before we go any further: using the corollary above one sees that if we have a quasi-character of $K^\times$ as in the corollary, then $|c| := |.| \circ c : K^\times \to \mathbf{R}_{>0}$ (note that the target $\mathbf{C}^\times$ always has the usual norm $|x + iy| = \sqrt{x^2 + y^2}$) must just be the map $\alpha \mapsto |\alpha|^\sigma$ with $\sigma = \mathrm{Re}(s)$ (note that this is well-defined even in the non-arch case). This real number σ is called the *exponent* of c, and "quasi-character of exponent at least σ" is going to be our analogue of "complex number with real part at least σ" later on. Also, let's say that two quasi-characters are *equivalent* if their ratio is unramified—this is just the same as demanding that they are in the same component of the quasi-character manifold.

We need now to fix a Haar integral on $K^\times$. If $f \in \mathcal{K}(K^\times)$ then the function $x \mapsto f(x)/|x|$ is a function on $K\backslash\{0\}$ and (by compactness of support) its extension to K (send 0 to 0) is continuous with compact support, so we can integrate it using our fixed choice of Haar measure on K; the resulting functional is easily checked to be a Haar integral μ_1 on $K^\times$. If K is archimedean then this will be our fixed choice of Haar integral on $K^\times$. If however K is p-adic then we're going to do something else.

Let's compute $\mu_1(\chi_U)$ (the characteristic function of the units) in the p-adic case. Because $|.|$ is trivial on U, we see $\mu_1(\chi_U) = \mu(\chi_U)$ and we normalised μ on K so that the integers R had measure $p^{-m/2}$, where $p^m\mathbf{Z}_p$ is the discriminant ideal of K. So, because $U = R\backslash\varpi R$, we see $\mu(\chi_U) = (1-1/q)\mu(\chi_R) = (1-1/q)p^{-m/2}$. It will be convenient to choose a Haar measure μ^* on $K^\times$ such that U has measure 1 for almost all P, as P runs through the primes of a number field k and $K = k_P$, so we define μ^* on $K^\times$ by $\mu^* = \frac{q}{q-1}\mu_1$ in the P-adic case. Then $\mu^*(\chi_U) = p^{-m/2}$ if K has discriminant p^m, and in particular if $K = k_P$ then $\mu^*(\chi_U) = 1$ for all but finitely many P (a number field has a discriminant and if P is coprime to this discriminant then the discriminant of k_P is just $\mathbf{Z}_p$).

4.3: Local analytic continuation and functional equations.

In some sense we've done nothing much in this chapter so far—apart from the check that K is isomorphic to $\widehat{K}$, all we've done is made explicit choices of things. Here's the first hint that something magic is happening though. Fix K as usual, and say $f : K \to \mathbf{C}$ is continuous, with $f \in L^1(K)$, and such that $\widehat{f}$ is also continuous and in $L^1(\widehat{K}) = L^1(K)$ (recall we have fixed an identification of K with $\widehat{K}$). Assume furthermore that $x \mapsto f(x)|x|^\sigma$ and $x \mapsto \widehat{f}(x)|x|^\sigma$ are both in $L^1(K^\times)$ for any $\sigma > 0$. One might ask whether any non-zero such functions exist, but we'll see plenty of examples later on, and in fact it's an easy exercise to come up with some examples. Let $Z = Z(K)$ be the set of such functions. In words these conditions imply that f is racing to zero more quickly than any polynomial in $|x|$ as $|x|$ gets big (that's the $L^1(K^\times)$ condition), and that f is bounded near zero (that's the $L^1(K)$ condition) and furthermore that $\widehat{f}$ has the same properties.

Easy exercise: why would asking $f \in L^1(K^\times)$ be asking a bit too much? [Hint: remember Haar measure on $K^\times$ isn't the same as that on K; consider what's happening near the origin].

Let Q denote the set of quasi-characters of $K^\times$. If $c \in Q$ then let's write $\mathrm{Re}(c)$ for the exponent of c. Given f as above, define a function $\zeta(f,-)$ on $\{c \in Q : \mathrm{Re}(c) > 0\}$ by

$$\zeta(f,c) = \int_{K^\times} f(t)c(t)\mathrm{d}\mu^*(t).$$

So we're using the multiplicative Haar measure on $K^\times$ defined above.

Lemma. The function $\zeta(f,-)$ (converges and) is holomorphic on the complex manifold $\{c \in Q : \mathrm{Re}(c) > 0\}$.

Proof. This rather fancy-sounding lemma is actually elementary to prove. Convergence is not an issue because our assumptions on f imply that the integrand defining this local zeta function is in $L^1(K^\times)$—we have $|f(x)c(x)| = |f(x)||x|^{\mathrm{Re}(c)}$ which is in $L^1(K^\times)$ by definition.

The complex structure near $c \in Q$ is given by $c.|.|^s$ for $s \in \mathbf{C}$ small, so all we have to do is to check that $\zeta(f, c.|.|^s)$ is holomorphic in s, for s small enough. So we have to differentiate $\zeta(f, c.|.|^s)$ with respect to s, and if you write out the definition of "differentiation", and remember all the boundedness assumptions we've made on f, you see that you can differentiate under the integral sign! This reduces us to checking that $\int_{K^\times} f(t)c(t)|t|^s \log(|t|)\mathrm{d}\mu^*(t)$ converges, but it does because for $|t|$ big, $|t|$ beats $\log(|t|)$ and the integral converges for all sufficiently large exponents, and for $|t|$ small, $|t|^{-\delta}$ beats $|\log(|t|)|$ (where $|t|$ is chosen so small that $\delta < \mathrm{Re}(s)$ works, which makes the integral converge by assumption). $\square$

But that's not the big local insight; the big insight is that $\zeta(f,-)$ has a meromorphic continuation to all of Q! The "global" zeta functions coming later will be products of local zeta functions for $\mathrm{Re}(c)$ sufficiently large. It is however important to note that this local analytic continuation result certainly does *not* give the meromorphic continuation of the global zeta functions that are coming later—an infinite product of meromorphic things might not be meromorphic (it might not even converge). Let me illustrate this by remarking that we'll shortly see that an example of this local meromorphic continuation statement is the statement that the function $\sum_{i\geq 0} p^{-is}$, which converges for $\mathrm{Re}(s) > 0$, can be rewritten as $1/(1-p^{-s})$, which is meromorphic for all $s \in \mathbf{C}$. This observation is clearly not enough to meromorphically continue blah $\prod_p (1-p^{-s})^{-1} = \zeta(s)$.

So let's prove this local meromorphic continuation, and even a local functional equation. We need an analogue of $s \mapsto 1-s$ for the functional equation; it's $c \mapsto \widehat{c}$, where $\widehat{c}$ is defined by $\widehat{c}(x) = |x|/c(x))$. Note that this hat has nothing to do with Fourier transforms, it's just an elementary definition. Note also that $\mathrm{Re}(\widehat{c}) = 1 - \mathrm{Re}(c)$. But also note that in general $\widehat{c}$ won't be equivalent to c—they could well lie on different components of Q.

Note that we haven't used any of our boundedness assumptions on $\widehat{f}$ so far, we've only used the L^1ness of f. We'll use L^1ness of $\widehat{f}$ now though.

Lemma. If $0 < \mathrm{Re}(c) < 1$ and $f, g \in Z$ then

$$\zeta(f,c)\zeta(\widehat{g},\widehat{c}) = \zeta(g,c)\zeta(\widehat{f},\widehat{c}).$$

[note that all integrals obviously converge.]

$[\zeta(f,c)\zeta(\widehat{g},\widehat{c}) = \zeta(g,c)\zeta(\widehat{f},\widehat{c}).]$

Perhaps I should remark that I have very little "true understanding" of this equation. I can prove it though, in fact it's dead easy to prove, it just follows from unravelling and Fubini. Let's see the proof.

First note that what we have to do, to prove the lemma, is to prove that if left hand side is $L(f,g)$, then $L(f,g) = L(g,f)$.

If we substitute in the definition of $\zeta(-,-)$ twice on the left hand side, we get

$$\int_{K^\times \times K^\times} f(\alpha)\widehat{g}(\beta)c(\alpha/\beta)|\beta|\mathrm{d}\mu^*(\alpha)\mathrm{d}\mu^*(\beta)$$

and by Fubini we can integrate in whatever order we want as long as $0 < \mathrm{Re}(c) < 1$ (I only proved Fubini for $\mathcal{K}(G \times H)$ but it extends to L^1). If I think about doing the integral over β first, for fixed α, then I can use invariance of $\mu^*(\beta)$ under multiplication to change β to $\gamma = \beta/\alpha$, and making this substitution shows that the integral equals

$$\int_{K^\times \times K^\times} f(\alpha)\widehat{g}(\alpha\gamma)c(\gamma^{-1})|\alpha\gamma|\mathrm{d}\mu^*(\alpha)\mathrm{d}\mu^*(\gamma).$$

$$\int_{K^\times \times K^\times} f(\alpha)\widehat{g}(\alpha\gamma)c(\gamma^{-1})|\alpha\gamma|\mathrm{d}\mu^*(\alpha)\mathrm{d}\mu^*(\gamma).$$

Now let's give a name to the "α" integral

$$H_{f,g}(\gamma) := \int_{K^\times} f(\alpha)\widehat{g}(\alpha\gamma)|\alpha|\mathrm{d}\mu^*\alpha;$$

then the integral we're trying to prove something about is

$$\int_{K^\times} H_{f,g}(\gamma)c(\gamma^{-1})|\gamma|\mathrm{d}\mu^*(\gamma).$$

Now *if* $H_{f,g}(\gamma) = H_{g,f}(\gamma)$ then visibly this latter integral (which has no other fs and gs in) will also be unchanged if we switch f and g, which is exactly what we wanted to prove. So we're now reduced to showing $H_{f,g} = H_{g,f}$. But recall that μ^* was, up to a constant κ which depended only on K, just $\mu(x)/|x|$, so

$$H_{f,g}(\gamma) = \kappa \int_K f(\alpha)\widehat{g}(\alpha\gamma)\mathrm{d}\mu(\alpha)$$

where note now the integral is over K, and now by definition of $\widehat{g}$ we see

$$H_{f,g}(\gamma) = \kappa \int_{K\times K} f(\alpha)g(\delta)e^{-2\pi i\Lambda(\alpha\gamma\delta)}\mathrm{d}\mu(\alpha)\mathrm{d}\mu(\delta)$$

recalling that our identification of K with $\widehat{K}$ sent x to the character $y \mapsto e^{2\pi i\Lambda(xy)}$. But now of course we're home, because switching f and g is just the same as changing notation $(\alpha, \delta) \to (\delta, \alpha)$. □

Reminder: we've just proved that if $0 < \mathrm{Re}(c) < 1$ and $f, g \in Z$ then

$$\zeta(f, c)\zeta(\widehat{g}, \widehat{c}) = \zeta(g, c)\zeta(\widehat{f}, \widehat{c}).$$

Corollary. If, for each component C of Q, we can find *one* explicit function $f = f_C \in Z$ such that $\zeta(\widehat{f}, \widehat{c})$ doesn't vanish identically on $\{c \in C : 0 < \mathrm{Re}(c) < 1\}$ and such that $\rho(c) := \zeta(f, c)/\zeta(\widehat{f}, \widehat{c})$ has a meromorphic continuation to all $c \in C$, then for *any* $g \in Z$, the function $\zeta(g, c)$ has meromorphic continuation to all of Q, and satisfies the functional equation

$$\zeta(g, c) = \rho(c)\zeta(\widehat{g}, \widehat{c})$$

for $c \in C$.

Proof. Trivial. The left hand side is holomorphic for $\mathrm{Re}(c) > 0$, our assumptions on ρ show that the right hand side is meromorphic for $\mathrm{Re}(c) < 1$, and the lemma we just proved shows that the two sides agree on the overlap.

4.4: Tidying up.

Here I'll give explicit examples of f_C as promised above, check that the corresponding $\rho(c)$ has meromorphic continuation, and also check the assertions about our choice of Haar measure being "self-dual", which will of course come out in the wash.

I have to start somewhere so let's start with $K = \mathbf{Q}_p$, the simplest case where we haven't done any explicit integrals yet. Recall that if $\mathbf{Z}_p$ is $\{k \in K : |k| \leq 1\}$, the integers of K, then $\mathbf{Z}_p$ is a ring, and the invertible elements $\mathbf{Z}_p^\times$ of this ring are easily seen to be just $\{k \in K : |k| = 1\} = U$.

What we have to do to perform the local meromorphic continuation is, for each character $\chi : \mathbf{Z}_p^\times \to S^1$, we need to find a function $f = f_\chi \in Z$ with $\zeta(\widehat{f}, \widehat{c})$ not identically zero on the region $0 < \mathrm{Re}(c) < 1$ of the component of Q corresponding to χ, and such that we can meromorphically continue $\rho(c) := \zeta(f, c)/\zeta(\widehat{f}, \widehat{c})$ to all of this component.

A reminder of normalisations: additive Haar measure μ on $\mathbf{Q}_p$ is normalised so that the characteristic function $\chi_{\mathbf{Z}_p}$ of $\mathbf{Z}_p$ has integral 1. Note that $\chi_{\mathbf{Z}_p} \in \mathcal{K}_+(\mathbf{Q}_p)$ and by additivity of Haar measure we have $\mu(\chi_{a+p^n\mathbf{Z}_p}) = p^{-n}$ for all $n \in \mathbf{Z}$. In particular locally constant functions with compact support are easy to integrate—but these things are easily checked to uniformly approximate anything in $K(\mathbf{Q}_p)$, so integration on $\mathbf{Q}_p$ is in fact easy. One last reminder: multiplicative Haar measure μ^* on $\mathbf{Q}_p^\times$ is $\frac{p}{p-1}$ times "$dz/|z|$", the usual trick to turn additive Haar measure into multiplicative Haar measure.

First let's do the component of Q corresponding to the trivial character of $\mathbf{Z}_p^\times$. Let f be the characteristic function of $\mathbf{Z}_p$. Now let's get it straight in our heads what we have to do.

First we need to check $f \in Z$, which involves checking some boundedness conditions on f, computing $\widehat{f}$ and checking the boundedness conditions on this function too.

Next we need to compute $\zeta(f, c)$ and $\zeta(\widehat{f}, \widehat{c})$ for c in the component of $Q = \mathrm{Hom}(\mathbf{Q}_p^\times, \mathbf{C}^\times)$ corresponding to the trivial character of U (that is, for $c : \mathbf{Q}_p^\times \to \mathbf{C}^\times$ such that $c|\mathbf{Z}_p^\times$ is trivial).

Finally we need to check that $\zeta(\widehat{f}, \widehat{c})$ is not identically zero for such c, and that $\zeta(f, c)/\zeta(\widehat{f}, \widehat{c})$ has a meromorphic continuation to the entire component.

To do all of this we just need to unwind the definitions and then figure out the integrals.

To check $f \in Z$ we first need that f is continuous (easy), that f is integrable (it's even in $\mathcal{K}_+(\mathbf{Q}_p)$ so it's certainly integrable), and that $f(x)|x|^\sigma$ is in $L^1(\mathbf{Q}_p^\times)$ for $\sigma > 0$. This needs checking, not least because $f(x)$ is *not* in $\mathcal{K}_+(\mathbf{Q}_p^\times)$—$f$ is non-zero arbitrarily close to zero, and (think about the automorphism $x \mapsto 1/x$ of $\mathbf{Q}_p^\times$) this means f doesn't have compact support on $\mathbf{Q}_p^\times$. But f is visibly a pointwise increasing limit of functions with compact support (consider the characteristic functions f_n of $\mathbf{Z}_p \backslash p^n\mathbf{Z}_p$ for n large) so we had better compute the integrals of $f_n(x)|x|^\sigma$ on $\mathbf{Q}_p^\times$ and check they converge; if they do then we will have proved $f(x)|x|^\sigma$ is summable and hence in $L^1(\mathbf{Q}_p^\times)$.

Well

$$\begin{aligned}
&\int_{\mathbf{Q}_p^\times} f_n(x)|x|^\sigma \mathrm{d}\mu^*(x) \\
&= \int_{\mathbf{Z}_p \setminus p^n\mathbf{Z}_p} |x|^\sigma \mathrm{d}\mu^*(x) \\
&= \sum_{i=0}^{n-1} \int_{p^i\mathbf{Z}_p^*} |x|^\sigma \frac{p}{p-1}|x|^{-1}\mathrm{d}\mu(x) \\
&= \frac{p}{p-1}\sum_{i=0}^{n-1} \mu(p^i\mathbf{Z}_p^\times)p^{-i(\sigma-1)} \\
&= \frac{p}{p-1}\sum_{i=0}^{n-1} \frac{p-1}{p}p^{-i\sigma} \\
&= \sum_{i=0}^{n-1} p^{-i\sigma} \\
&\to \sum_{i=0}^{\infty} p^{-i\sigma} = \frac{1}{1-p^{-\sigma}}
\end{aligned}$$

as $n \to \infty$, if $\sigma > 0$. So the boundedness properties of f are satsified. Note: I am being lazy and writing $\mu(X)$ for $\mu(\chi_X)$, the characteristic function of X.

Now we need to compute $\widehat{f}$. Well, by definition

$$\widehat{f}(x) = \int_{\mathbf{Q}_p} f(y) e^{-2\pi i q(xy)} \mathrm{d}\mu(y)$$

where q is that function $\mathbf{Q}_p/\mathbf{Z}_p \to \mathbf{Q}/\mathbf{Z}$ defined by "take the fractional part". This is just

$$\int_{\mathbf{Z}_p} e^{-2\pi i q(xy)} \mathrm{d}\mu(y)$$

so let's do this integral. We always have $y \in \mathbf{Z}_p$ so if $x \in \mathbf{Z}_p$ then $q(xy) = 0$ and we just get $\mu(\mathbf{Z}_p) = 1$.

On the other hand, if $x \notin \mathbf{Z}_p$ and $e^{-2\pi i q(x)} = \zeta \neq 1$ and $|x| = p^n$, $n \geq 1$ (so ζ is a primitive p^nth root of unity), then $e^{-2\pi i q(xy)}$ only depends on y mod p^n and $\int_{\mathbf{Z}_p} e^{-2\pi i q(xy)} \mathrm{d}\mu(y) = \sum_{i=0}^{p^n-1} \zeta^i = 0$. Hence $\widehat{f} = f$.

But this is great—firstly we have proved that $c = 1$ in the Fourier Inversion theorem, as claimed earlier ($\hat{\hat{f}} = f(x) = f(-x)$), and secondly we have proved $f \in Z$, because $\hat{f} = f$ has all the same boundedness properties as f.

Next we need to compute $\zeta(f, c)$ and $\zeta(\hat{f}, \hat{c})$ on the region $0 < \mathrm{Re}(c) < 1$ of the component corresponding to the trivial character of U, the general element of which is $x \mapsto |x|^s$ for $s \in C = \mathbf{C}/\frac{2\pi i}{\log(p)}\mathbf{Z}$. We stick to $0 < \mathrm{Re}(s) < 1$. Now (being much lazier about the difference between L^1 and $\mathcal{K}(\mathbf{Q}_p^\times)$ this time)

$$\begin{aligned}
\zeta(f, c) &= \int_{\mathbf{Q}_p^\times} f(x)|x|^s \mathrm{d}\mu^*(x) \\
&= \frac{p}{p-1} \int_{\mathbf{Z}_p \setminus \{0\}} |x|^{s-1} \mathrm{d}\mu(x) \\
&= \frac{p}{p-1} \sum_{j=0}^{\infty} \int_{p^j \mathbf{Z}_p^\times} p^{-j(s-1)} \mathrm{d}\mu(x) \\
&= \sum_{j=0}^{\infty} p^{-js} = (1 - p^{-s})^{-1}
\end{aligned}$$

(which looks surprisingly familiar!). Note the last line is OK because Re$(s) > 0$ so $|p^{-s}| < 1$.

Similarly $\zeta(\widehat{f}, \widehat{c}) = \int_{\mathbf{Q}_p^\times} f(x)|x|^{1-s} \mathrm{d}\mu^*(x)$ which (recalling that we're assuming $0 <$ Re$(s) < 1$), by exactly the same calculation but with $1-s$ replacing s, comes out to be $(1-p^{-(1-s)})^{-1}$. Hence $\zeta(\widehat{f}, \widehat{c})$ is not identically zero on the component we're interested in, and $\rho(c) = \zeta(f, c)/\zeta(\widehat{f}, \widehat{c}) = \frac{1-p^{s-1}}{1-p^{-s}}$, which looks less familiar, but later on you'll see why you're not expected to recognise this function—these $\rho(c)$ will only show up explicitly at the "bad primes". The crucial observation that we need, however, is that $\rho(c)$ has a meromorphic continuation to all $s \in \mathbf{C}$, which is clear, and so our job is done.

Let me sketch the ramified case in this setting, because there $\rho(c)$ is quite different. Let χ denote a Dirichlet character of conductor p^n, $n \geq 1$, that is, a map $\chi : (\mathbf{Z}/p^n\mathbf{Z})^\times \to \mathbf{C}^\times$ (note that the image lands in S^1) that doesn't factor through $(\mathbf{Z}/p^{n-1}\mathbf{Z})^\times$. The natural map $\mathbf{Z}_p \to \mathbf{Z}/p^n\mathbf{Z}$ induces a map $U \to (\mathbf{Z}/p^n\mathbf{Z})^\times$ and hence a character of U, so we get a component of Q parameterised by $s \in \mathbf{C}/\frac{2\pi i}{\log(p)}\mathbf{Z}$ whose typical element sends $x = p^n u$ to $\chi(u)p^{-ns} = \chi(\tilde{x})|x|^s$ [we're setting $V = \{p^{\mathbf{Z}}\}$ with notation as above].

We need an $f = f_C$ for this component. Let's set $f(x) = 0$ if $|x| > p^n$, and $f(x) = e^{2\pi i q(x)}$ for $|x| \leq p^n$ (note $|x| \leq p^n$ implies $q(x) = a/p^n$ for some integer a). So again f is locally constant (indeed it's constant on cosets of $\mathbf{Z}_p$ in $p^{-n}\mathbf{Z}_p$) and takes values in the p^nth roots of unity. It's clear that $f \in L^1(\mathbf{Q}_p)$, and we can write $f = \chi_{\mathbf{Z}_p} + f'$ with $f' \in K(\mathbf{Q}_p^\times) \otimes \mathbf{C}$, and we already showed $\chi_{\mathbf{Z}_p}.|.|^\sigma \in L^1(\mathbf{Q}_p^\times)$ for $\sigma > 0$, which implies $f.|.|^\sigma \in L^1(\mathbf{Q}_p^\times)$ for $\sigma > 0$.

[$f(x) = e^{2\pi i q(x)}$ for $|x| \leq p^n$.] Next we need to compute the Fourier transform of f. We can do this by brute force or by a trick. Here's the brute force method:

$$\begin{aligned}\widehat{f}(x) &= \int_{\mathbf{Q}_p} f(y) e^{-2\pi i q(xy)} \mathrm{d}\mu(y) \\ &= \int_{p^{-n}\mathbf{Z}_p} e^{2\pi i q(y-xy)} \mathrm{d}\mu(y) \\ &= \int_{p^{-n}\mathbf{Z}_p} e^{2\pi i q(y(1-x))} \mathrm{d}\mu(y).\end{aligned}$$

The same "cancelling" phenomenon (adding roots of unity) hence shows $\widehat{f}(x) = 0$ if $|1 - x| > p^{-n}$ (we're adding roots of unity), but if $|1 - x| \leq p^{-n}$ then $|y(1 - x)| \leq 1$ for $y \in p^{-n}\mathbf{Z}_p$ and the integrand is just 1, showing that the integral is just $\mu(p^{-n}\mathbf{Z}_p) = p^n$ times the characteristic function of $1 + p^n\mathbf{Z}_p$. Note that this integral is bounded away from zero (as $n \geq 1$) so $\widehat{f} \in \mathcal{K}_+(\mathbf{Q}_p)$ and $\mathcal{K}_+(\mathbf{Q}_p^\times)$, and hence $f \in Z$.

[The trick way to do this last computation is to observe that the f here is just $\chi_{\mathbf{Z}_p}$ rescaled and multiplied by a character, and so we can compute the Fourier transform of our f from the Fourier transform of $\chi_{\mathbf{Z}_p}$ using basic properties of Fourier transforms.]

Next we need to compute $\zeta(f, c)$ and $\zeta(\widehat{f}, \widehat{c})$ for c of the form $p^n u \mapsto \chi(u)p^{-ns}$ for $0 < \mathrm{Re}(s) < 1$. These calculations are very similar to the ones we have already done (although perhaps slightly tougher, because in some cases it's trickier to check that certain sums of roots of unity are zero). I'll stick them on the example sheet and just tell you the answers: if I got it right then

$$\zeta(f, \chi(\tilde{x})|x|^s) = \frac{p^{ns+1-n}}{p-1} \sum_{j=1}^{p^n-1} \chi(j) e^{2\pi i j/p^n}$$

(the inner sum is called a Gauss sum) and $\zeta(\widehat{f}, \chi(\tilde{x})^{-1}|x|^{1-s})$ just turns out to be $p/(p-1)$, a constant!

So the ratio $\zeta(f,c)/\zeta(\widehat{f}/\widehat{c})$ is of the form $A.B^s$ with A a constant involving a Gauss sum, and B a positive real constant (a power of p in fact), which means that the ratio has meromorphic continuation to the component C and again we're done.

More precisely, the ratio is $p^{n(s-1)}\sum_{j=1}^{p^n-1}\chi(j)\zeta^j$ with $\zeta = e^{2\pi i/p^n}$. The following observation is now surely worth remarking on. The Dirichlet character χ we were just considering—we were doing local calculations with it, but we can also consider the global ζ function (or L-function, as it's more commonly known) attached to this character, which is (for $\mathrm{Re}(s) > 1$)

$$\sum_{m\geq 1}\chi(m)/m^s = \prod_{\ell}(1-\chi(\ell)\ell^s)^{-1},$$

the latter product being over all primes ℓ. This L-function has a meromorphic continuation to all of $\mathbf{C}$, which turns out to be holomorphic in this case, because we assumed the conductor was p^n for $n \geq 1$.

We have $\chi(-1) = \pm 1$ for some choice of sign. If $\chi(-1) = 1$ and if we multiply this L-function by the usual "fudge factor" $\pi^{-s/2}\Gamma(s/2)$, then we get a new function $\xi(\chi, s)$ satisfying

$$\xi(\chi, s) = \left(p^{-ns} \sum_{j=1}^{p^n - 1} \chi(j)\zeta^j \right) \xi(\overline{\chi}, 1 - s).$$

A similar sort of thing is true if $\chi(-1) = -1$ but then the fudge factors and the functional equation are slightly different. The moral is that this time the local integrals aren't showing up as components of the L-function, but the ratio $\rho(c)$ is showing up in the functional equation.

I am not going to plough through all the other cases. The computations are a little long but completely elementary and prime example sheet fodder. The crib is Tate's thesis, end of chapter 2. Here's the answers. If K is a finite extension of $\mathbf{Q}_p$ then the only extra subtlety is that we used the trace map to define $K \to \hat{K}$, and when doing the calculations one needs to compute $\{\alpha \in K : \mathrm{Tr}_{K/\mathbf{Q}_p}(\alpha v) \in \mathbf{Z}_p \forall v \in R\}$ where R is the integers of K. Clearly this set contains R, and is not all of K (because it doesn't contain p^{-n} for n large),so it's a fractional ideal of K, but what you may not know is that if we write it as $\pi^{-r}R$ then the norm to $\mathbf{Q}_p$ of π^r generates the discriminant ideal, which simplifies some constants a bit. The answers are "the same as in the $K = \mathbf{Q}_p$ case": for the unramified quasi-characters one lets f be the characteristic function of $\{\alpha \in K : \mathrm{Tr}_{K/\mathbf{Q}_p}(\alpha v) \in \mathbf{Z}_p \forall v \in R\}$ and one checks $\zeta(f, |.|^s) = p^{m(s-1/2)}/(1-q^{-s})$ and $\zeta(\hat{f}, |.|^{1-s}) = (1-q^{s-1})$, and in the ramified case one makes a sensible choice of f and ρ turns out to be of the form $A.B^s$ with A involving a Gauss sum.

If $K = \mathbf{R}$ then there are two components: on the component $x \mapsto |x|^s$ use $f(x) = e^{-\pi x^2}$, and on the component $x \mapsto \text{sgn}(x)|x|^s$, with $\text{sgn}(x)$ the sign of x, use $f(x) = xe^{-\pi x^2}$, and now use your 1337 Fourier Transform sk1llz to check that in the first case, when $c(x) = |x|^s$ we have $\zeta(f, c) = \pi^{-s/2}\Gamma(s/2)$ and $\zeta(\widehat{f}, \widehat{c}) = \pi^{-(1-s)/2}\Gamma((1-s)/2)$, so the ratio is meromorphic and furthermore *we have seen the ratio before!* The ratio shows up when writing $\zeta(1-s)/\zeta(s)$ as a product of simpler functions (i.e. the "fudge factors" in the functional equation). So now you're beginning to see some of the insights here—the "fudge factors" in the functional equation may have *local* explanations—for example the Γ factor is coming from the archimedean valuation on $\mathbf{Q}$. If you look at the functional equation for the zeta function for a number field, you will see several Γ factors, coming from the real and complex norms on the field, and one can check that they are the same factors that come up in these calculations.

The answer on the $\mathrm{sgn}(x)|x|^s$ component is similar, but one ends up with $\pi^{-\frac{s+1}{2}}\Gamma(\frac{s+1}{2})$, which is precisely the "fudge factor" that one has to use in the functional equation for the Dirichlet L-function when $\chi(-1) = -1$.

If $K = \mathbf{C}$ then the components are parametrized by the integers. Let's say the nth component is the quasicharacters whose restriction to $S^1 \subseteq \mathbf{C}$ is $z \mapsto z^n$. For $n \geq 0$ Tate chooses the function $f_n(x+iy) = (x-iy)^n e^{-2\pi(x^2+y^2)}$, and for $n \leq 0$ he chooses $f_n(x+iy) = (x+iy)^{-n}e^{-2\pi(x^2+y^2)}$. It turns out that $\hat{f_n} = c_n f_{-n}$ where c_n is an explicit root of unity (proof by basic integrals and induction on n) and the local zeta values are again just powers of π and Γ functions, for example if $n \geq 0$ then $\zeta(f_n, re^{i\theta} \mapsto r^s e^{in\theta}) = (2\pi)^{1-s+\frac{n}{2}}\Gamma(s+\frac{n}{2})$ and the other answers are similar. For an explicit list of the answers, look at the end of chapter 2 of Tate's thesis or the new example sheet.

Summary of what I just breezed through:

All local zeta functions have meromorphic continuations. The local zeta functions attached to our favourite functions (the fs we used) looked like $(1 - p^{-s})$ on the unramified non-arch components and involved the Γ function and πs in the real and complex cases. These local factors are precisely what one multiplies together to get the function $\xi(s)$ (the Riemann zeta function multiplied by the "fudge factors at infinity"). The local zeta functions on the ramified components in the non-arch case are messier, but the ratio $\zeta(f, c)/\zeta(\widehat{f}, \widehat{c})$ involves Gauss sums.

And let me stress once more that these local calculations do not even come close to analytically continuing the usual zeta function; we need more to do this.

Chapter 5. The adeles and ideles.

The Pontrjagin dual of $\mathbf{Z}$ (with the discrete topology) is $\mathbf{R}/\mathbf{Z}$. But the Pontrjagin dual of $\mathbf{Q}$ (with the discrete topology) turns out to be an absolutely huge uncountable compact topological group, rather surprisingly! The dual turns out to be related to some kind of infinite product of all the completions of $\mathbf{Q}$ at once, as we will see later on. But we have to be careful here: if I have infinitely many non-empty locally compact topological spaces X_i, their product turns out not to be locally compact in general (because the definition of the product topology has, as basic open sets, products of open sets U_i, but all but finitely many of the U_i have to be equal to X_i and this makes it hard to find a compact neighbourhood of such a product). So we have to be careful—the product over all p of $\mathbf{Q}_p$ isn't locally compact and hence we can't do Haar integration on it.

Here's a partial fix:

Lemma. If we have a collection X_i of locally compact Hausdorff topological groups, and *furthermore if all but finitely many of them are compact*, then the product of the X_i is a locally compact Hausdorff topological group.

Proof. Given a basic open neighbourhood $\prod_i U_i$ of a point (x_i) in the product, all but finitely many of the U_i are equal to X_i by definition, and are hence compact, so we leave them alone, and the rest of the U_i we can shrink to V_i, a compact neighbourhood of x_i, and the product of the V_i is a compact neighbourhood of (x_i) in $\prod_i U_i$. So the product (with its product topology) is locally compact, and the rest is easy (checking hausdorffness, and that multiplication and inverse are continuous).

The problem we now face is that the completions of $\mathbf{Q}$ with the p-adic and real norms are all locally compact, but none of them are compact. Here is the abstract construction that gets around this.

5.1: The restricted direct product.

Here's the set-up. We have a set I (typically infinite), a locally compact Haudsorff topological group G_i for all $i \in I$, and, for all but finitely many i, a given fixed subgroup H_i of G_i which is both *open* and *compact*. Say S_0 is the finite subset of I for which no H_i is given. Say S is any finite set containing S_0. Then we can form $G_S := \prod_{i \in S} G_i \times \prod_{i \notin S} H_i$; this is locally compact (with the product topology) and, as a set, sits naturally inside $\prod_i G_i$. But no one finite set S is better than any other, so we now take the union (within $\prod_i G_i$), as S gets bigger, of the G_S. Call this union G.

Then G is clearly a group (it's a directed union of groups; $G_S \cup G_T \subseteq S_{S \cup T}$). To make it a topological group we just have to give a basis for the topology near the identity, and we can do this by choosing any $S \supseteq S_0$ and saying that a basis of neighbourhoods of the identity in G is just a basis of neighbourhoods of the identity in the subgroup G_S. It's an elementary exercise to check that this independent of S (check that a basis of neighbourhoods is given by $\prod_i N_i$ with $1 \in N_i \subseteq G_i$ and $N_i = H_i$ for all but finitely many i) and that this construction makes G into a locally compact topological group.

If all the G_i are furthermore abelian then we have

$$0 \to \prod_{i \notin S_0} H_i \to G \to \left(\oplus_{i \notin S_0} G_i/H_i\right) \oplus \left(\oplus_{i \in S_0} G_i\right)$$

so G is a sort of mixture of a direct product with a direct sum.

Note that each G_i is naturally a subgroup of G. An element of G can be thought of as an element (g_i) of $\prod_i G_i$ with the property that $g_i \in H_i$ for all but finitely many i.

Notation:

$$G = \prod_i{}' G_i.$$

Not very good notation, because it doesn't say what the H_i are. Rotten luck.

From now on, assume that all the G_i are abelian.

The following things are all elementary to check and I will only hint at proofs.

[Reminder: $G \subseteq \prod_i G_i$ is the (g_i) with $g_i \in H_i$ for all but finitely many i]

1) If $c : G \to \mathbf{C}^\times$ is continuous, then $c_i := c|G_i$ is trivial on H_i for all but finitely many i, and hence one can make sense of the character $\prod_i c_i$ on G (because it's a finite product) and one can check that $\prod_i c_i = c$. [Proof: because c is continuous, if V is a small neighbourhood of 1 then $c^{-1}(V)$ is open in G and hence contains a subgroup of the form $\prod_{i \notin S} H_i$; but $c(\prod_{i \notin S} H_i)$ is now a subgroup of V and for V small enough the only subgroup is $\{1\}$].

2) If H_i is a compact open subgroup of G_i (note that open implies closed, because H_i is the complement of the open set $\cup_{g \notin H_i} gH_i$) and if we define H_i^* to be the annihiliator of H_i in $\widehat{G}_i$, then H_i^* is also compact and open. [Proof: the dual of H_i is $\widehat{G}_i/H_i^*$ so compactness of H_i implies discreteness of $\widehat{G}_i/H_i^*$ implies openness of H_i^* etc].

3) The Pontrjagin dual of $G = \prod_i' G_i$ (restricted product with respect to the H_i) is $\widehat{G} = \prod_i' \widehat{G}_i$ (restricted product with respect to the H_i^*). [Proof: we've seen that a unitary character c of G is a product of its components, and that conversely given a bunch of c_i all but finitely many of which are trivial on H_i we can multiply them together to get a c, and now one just unravels this.]

4) Say $S \supseteq S_0$, so $G_S = \prod_{i \in S} G_i \times \prod_{i \notin S} H_i$ makes sense and is a LCHTG. If we choose a Haar integral μ_i on each G_i ($i \in S$) and on each H_i ($i \notin S$) and we normalise the H_i ones such that $\mu(\mathbf{1}) = 1$ (where $\mathbf{1}$ is the constant function on H_i, which is in $\mathcal{K}(H_i)$), then there's a unique natural Haar measure $\prod_i \mu_i$ on G_S, with the property that if N_i is a subset of G_i for all i with the property that $N_i = H_i$ for all but finitely many i and that $\prod_i N_i \subseteq G_S$, and if χ_{N_i} is summable for all i, then $\mu(\chi_N) = \prod_i \mu_i(\chi_{N_i})$ [this is only a finite product of course].

5) Hence if we fix Haar measures μ_i on each G_i with the property that $\mu_i(\chi_{H_i}) = 1$ for all but finitely many i, we get a natural Haar integral μ on G, given by

$$\mu(f) = \lim_S \mu(f|_{G_S})$$

for any $f \in \mathcal{K}(G)$, the limit being taken over all $S \supseteq S_0$, and this limit will exist (because the support of f will be contained within one of the G_S, so in fact the sequence is ultimately constant).

6) (extension of 5 to summable functions). Say, for each i, we have a summable function f_i on G_i with the property that $f_i|H_i = 1$ for all but finitely many i. Define a function f on G by $f((g_i)) = \prod_i f_i(g_i)$ (a finite product!). Then the integral of $f|G_S$ is just $\prod_{i \notin S} \mu_i(f_i)$ and if the infinite product converges absolutely (for example if $f = \chi_{H_i}$ for all but finitely many i), we will have $\mu(f) = \prod_i \mu_i(f_i)$.

7) If for each i we have a continuous summable $f_i : G_i \to \mathbf{C}$ with the property that $\widehat{f_i} : \widehat{G_i} \to \mathbf{C}$ is also continuous and summable, and if furthermore $f_i = \chi_{H_i}$ for all but finitely many i, then $f = \prod_i f_i$ makes sense (and is a finite product wherever it is evaluated), it's continuous and summable, and $\widehat{f} : \widehat{G} \to \mathbf{C}$ is just $\prod_i \widehat{f_i}$, which is also continuous and summable.

8) Finally, if we fix Haar integrals on G_i and $\widehat{G_i}$ for all i with the property that the integrals are self-dual (that is $\widehat{\widehat{f}}(x) = f(-x)$, so the positive constant that may be involved is in fact 1 for each i) and if $\mu_i(H_i) = 1 = \widehat{\mu}(H_i^*)$ for all but finitely many i, then the product Haar integrals are also self-dual.

Recall last time: we had a collection G_i ($i \in I$) of locally compact Hausdorff topological groups, a finite set $S_0 \subseteq I$, and, for all $i \notin S_0$ (so, for all but finitely many i), we had a compact open subgroup H_i of G_i.

Given this data we can form $G := \prod_i' G_i$, the restricted product of the G_i with respect to the H_i. As a group it's the elements $(g_i) \in \prod_i G_i$ such that $g_i \in H_i$ for all but finitely many i (where this finite set is allowed to vary). The easiest way to think about the topology is to realise that $G_{S_0} := \prod_{i \notin S_0} H_i \times \prod_{i \in S_0} G_i$ is an open subgroup, with the usual product topology on it. It turns out that G is also locally compact and Hausdorff, its Haar measure can be thought of as "the product of the Haar measures on G_i" as long as these are normalised such that $\mu(H_i) = 1$ for all but finitely many H_i, and the Pontrjagin dual of G is just the restricted product of the $\widehat{G_i}$ with respect to H_i^*, the annihiliator of H_i in $\widehat{G_i}$.

In fact we only need two examples for Tate's thesis and in both cases the G_i (and hence G) will be abelian.

5.2: The adeles and ideles.

Let k be a number field, so a finite extension of $\mathbf{Q}$. [The theory works just as well for function fields—that is, finite extensions of $\mathbf{F}_p(T)$, but I'd like to emphasize the number field case, especially as I was too lazy to finish the proof of the meromorphic continuation of the local zeta functions in the function field case!]

Let I be the following set: there's an element of I for each non-zero prime ideal P of R, the algebraic integers in k, and there's also an element of I for each equivalence class of field homomorphisms $\tau : k \to \mathbf{C}$, with $\tau \sim \overline{\tau}$.

Recall from ages ago that each element of I gives us an equivalence class of norms on k; the prime ideals P give us P-adic norms, and the maps $k \to \mathbf{C}$ give us norms induced from the standard norm on $\mathbf{C}$.

The elements of I are called *places* of k, and a typical element of I is traditionally denoted v (for valuation, I guess, which is another word for norm). For each $v \in I$ let G_v denote the completion k_v of k with respect to the norm induced by v. Let S_0 denote the norms coming from $k \to \mathbf{C}$—these are called "the infinite places" [this set is empty in the function field case, and finite but non-empty in the number field case]. For $v \notin S_0$ (a "finite place") the completion $k_v = k_P$ of k has a ring of integers R_v; let this be H_v.

Define the *adeles of* k, written $\mathbf{A}_k$, to be the restricted product of the k_v with respect to the R_v.

Let's write this out explicitly in the case $k = \mathbf{Q}$: we have $\mathbf{A}_{\mathbf{Q}}$ is the subgroup (in fact it's easily checked to be a subring) of

$$\mathbf{Q}_2 \times \mathbf{Q}_3 \times \mathbf{Q}_5 \times \ldots \times \mathbf{R}$$

consisting of $(g_2, g_3, g_5, \ldots, g_\infty)$ with the property that $g_p \in \mathbf{Q}_p$ for all p, and $g_\infty \in \mathbf{R}$, and, crucially, that $g_p \in \mathbf{Z}_p$ for all but finitely many p.

Indeed in the general case one can easily check that the topological group $\mathbf{A}_k$ has a natural ring structure induced by componentwise multiplication (because H_i is a subring of G_i for all finite places i).

That's the first construction we will use. As you can see, the finite and the infinite places are behaving quite differently (the infinite places have no H_v) and it's common to write

$$\mathbf{A}_k = \mathbf{A}_k^f \times k_\infty$$

with $\mathbf{A}_k^f$ the "finite adeles", namely $\prod'_P k_P$, the restricted product over all the finite places, and k_∞ the "infinite adeles", namely the finite product $\prod_{[\tau]} k_\tau$, with $[\tau] = \{\tau, \overline{\tau}\}$ the equivalence class of τ, and where $k_\tau = \mathbf{R}$ if $\tau : k \to \mathbf{R}$ and $k_\tau \cong \mathbf{C}$ if $[\tau] = \{\tau, \overline{\tau}\}$ with $\tau : k \to \mathbf{C}$ with image not landing in $\mathbf{R}$.

An absolutely crucial observation is that the "diagonal map" $k \to \prod_v k_v$ sending λ to $(\lambda, \lambda, \lambda, \ldots)$ has image landing in $\mathbf{A}_k$; this is because any $\lambda \in k$ can be written $\lambda = a/b$ with $a, b \in R$, the integers of k, and $b \neq 0$, and the factorization of (b) into prime ideals only involves finitely many prime ideals of R, and if S is S_0 union this finite set then $\lambda \in H_v = R_v$ for all $v \notin S$.

That's the first construction we will use. The second is the *ideles* of k, which I'll denote $\mathbf{A}_k^\times$, and which is the restricted product of the $k_v^\times$ with respect to the $R_v^\times$. This is a topological group. As the name indicates,

Lemma. The ideles are the units in the ring of adeles.

Remark. Note that this is an algebraic statement; it says nothing about the toplogies of the adeles or ideles.

Proof. If $(g_v) \in \mathbf{A}_k$ has an inverse, then certainly all of the g_v are non-zero and the inverse is (g_v^{-1}). For both (g_v) and (g_v^{-1}) to be in $\mathbf{A}_k$ we need $g_v \in R_v$ for almost all v (n.b. "almost all" means "for all but finitely many") and $g_v^{-1} \in R_v$ for almost all v. This means $g_v \in R_v^\times$ for almost all v, which is precisely the assertion that (g_v) is an idele. $\square$

Historical note: ideles were invented/discovered before adeles. Ideles were introduced by Chevalley, and he actually called them "ideal elements", which he abbreviated "id.ele." which became "idele". It was later realised that they were the units of a ring, which Tate calls the "ring of valuation vectors" in his thesis.

It was Weil that introduced the terminology "adele", for "additive idele". If you look at Serre's CV (for example at the beginning of Vol. 1 of his collected works) you'll see that his mother's name was Adele, but Serre once told me that he had nothing to do with the introduction of the terminology, and merely found it ironic that his mother's name ended up being used in mathematics.

Pedantic/irrelevant remark (which we won't use later). The inclusion $\mathbf{A}_k^\times \to \mathbf{A}_k$ is continuous (because a basic open neighbourhood of the element $(1,1,1,1,\ldots)$ in $\mathbf{A}_k$ is $\prod_v N_v$ with $N_v = R_v$ for all but finitely many v, and hence its pullback to $\mathbf{A}_k^\times$ will contain $\prod_v M_v$ with $M_v = R_v^\times$ for all but finitely many v). However the inclusion is not a homeomorphism onto its image; the problem is that $\prod_{v<\infty} R_v^\times \times \prod_{v|\infty} K_v^\times$ is open in the ideles but not in the subspace topology (because any neighbourhood of 1 in the adeles will contain elements of the form $(1,1,1,1,\ldots,1,\pi,1,\ldots,1)$

(with π in the vth place and a uniformiser in k_v), for all but finitely many v. The way to fix this up turns out to be the trick I mentioned earlier: give $\mathbf{A}_k \times \mathbf{A}_k$ the product topology and embed $\mathbf{A}_k^\times$ into this product by sending u to $(u, 1/u)$; now the restricted product topology on $\mathbf{A}_k^\times$ is indeed the subspace topology.

An absolutely crucial function on the ideles of a number field is the *norm* function. For any completion k_v of a number field we have written down a canonical norm (the one where the norm of α is how much an additive Haar measure is "stretched" under multiplication by α). Let's call this norm $|.|_v$ now. Note that for v finite and $u_v \in R_v^\times$ we have $|u_v|_v = 1$. Hence there is a function

$$|.| : \mathbf{A}_k^\times \to \mathbf{R}_{>0}$$

defined by

$$|(g_v)| = \prod_v |g_v|_v$$

with the usual remark that, for any given v, this is a finite product. I'll refer to this function as "the global norm" but note that it's a continuous group homomorphism rather than a norm on a field in the sense we talked about earlier.

Unsurprisingly, given that a Haar integral on $\mathbf{A}_k$ can be thought of as a product of local Haar integrals, it turns out that this norm on $\mathbf{A}_k^\times$ is just the factor by which multipliaction by an idele is stretching the additive Haar integral on the adeles.

Our goal, of course, is to develop some machinery to work with the following sort of idea. Let me just stick to the case $k = \mathbf{Q}$. Let's define a function on the ideles $\mathbf{A}_\mathbf{Q}^\times$ thus: for p a prime number, define f_p on $\mathbf{Q}_p$ to be the characteristic function of $\mathbf{Z}_p$. Define f_∞ on $\mathbf{R}$ to be $e^{-\pi x^2}$. Define $f : \mathbf{A}_\mathbf{Q}^\times \to \mathbf{C}$ by $f((g_v)) = \prod_v (f_v(g_v))$ (a finite sum). Now consider the function

$$s \mapsto \int_{\mathbf{A}_\mathbf{Q}^\times} f(x)|x|^s \mathrm{d}\mu^*(x) \qquad (1)$$

where μ^* denotes the Haar measure on the ideles which is the product of the local Haar measures μ^* on $\mathbf{Q}_p^\times$ and $\mathbf{R}^\times$. This integral will not converge for a general $s \in \mathbf{C}$; the integrand isn't L^1. A sufficient condition for the integrand to be L^1 is that all the local integrands are L^1 and furthermore that the product of the local integrals is absolutely convergent. But we already worked these local integrals out, at least at the finite places: at the finite places we have

$$\int_{\mathbf{Q}_p^\times} f_p(x)|x|_p^s \mathrm{d}\mu^*(x)$$

and when we were meromorphically continuing local zeta functions we checked that this was L^1 for $\mathrm{Re}(s) > 0$ and that its value was $\sum_{j\geq 0} p^{-js} = (1-p^{-s})^{-1}$. At the infinite place, I skipped the calculation so let's do it now: we need to compute

$$\int_{\mathbf{R}^\times} e^{-\pi x^2}|x|^s(\mathrm{d}x/|x|)$$
$$= 2\int_0^\infty e^{-\pi x^2}x^{s-1}\mathrm{d}x$$

and setting $y = \pi x^2$ this is

$$\begin{aligned}&\pi^{-1}\int_0^\infty e^{-y}(y/\pi)^{\frac{s-2}{2}}\mathrm{d}y\\ &= \pi^{-s/2}\Gamma(s/2)\end{aligned}$$

by definition of the Γ function, if $\mathrm{Re}(s) > 0$ (and the integral doesn't converge absolutely at zero if $\mathrm{Re}(s) \leq 0$). Hence a necessary and sufficient condition for the adelic integral (1) to converge is that $\prod_p (1-p^{-s})^{-1}$ converges absolutely, and for $\mathrm{Re}(s) > 1$ this will be the case because the product is just $\sum_{n\geq 1} n^{-s} = \zeta(s)$. So for $\mathrm{Re}(s) > 1$ the adelic integral (1) will converge, and it will converge to

$$\xi(s) := \pi^{-s/2}\Gamma(s/2)\zeta(s).$$

I proved in the second lecture that $\xi(s)$ had a meromorphic continuation to $s \in \mathbf{C}$ and satisfied $\xi(s) = \xi(1-s)$. We now have an adelic interpretation of the statement.

If we can also give an adelic *proof* of $\xi(s) = \xi(1-s)$, by interpreting our original proof adelically, one might hope that the idea will generalise to all number fields. Indeed, our main theorem will be the meromorphic continuation of a wide class of integrals on idele groups, and we will recover a theorem of Hecke whose original proof was a real tour de force.

We have chosen Haar measures on k_v and isomorphisms $k_v = \widehat{k}_v$ in such a way that the Fourier inversion theorem on k_v is true on the nose (the fudge factor constant is 1). For each v our map $k_v \to \widehat{k}_v$ was of the form $x \mapsto (y \mapsto e^{2\pi i \Lambda_v(xy)})$ where Λ_v, which we called Λ at the time, was some explicitly given map $k_v \to \mathbf{R}/\mathbf{Z}$. For $\mathbf{Q}_p$ it was $\mathbf{Q}_p \to \mathbf{Q}_p/\mathbf{Z}_p \to \mathbf{Q}/\mathbf{Z} \to \mathbf{R}/\mathbf{Z}$, the middle map being called q. For $k_P/\mathbf{Q}_p$ finite it was the trace map $k_P \to \mathbf{Q}_p$ followed by the above map. For the reals it was $x \mapsto -x$ sending $\mathbf{R}$ to $\mathbf{R}/\mathbf{Z}$ and for the complexes it sent $(x+iy)$ to $-2x$.

Note that for all finite v, we see that R_v is in the kernel of Λ_v. So, by the usual trick, we get a map

$$\Lambda : \mathbf{A}_k \to \mathbf{R}/\mathbf{Z}$$

defined by

$$\Lambda((g_v)) = \sum_v \Lambda_v(g_v)$$

which is, as usual, a finite sum. I could now "cheat" and say that there was an induced map

$$\mathbf{A}_k \to \widehat{\mathbf{A}}_k$$

sending x to $y \mapsto e^{2\pi i \Lambda(xy)}$, which was a restricted product of isomorphisms, and is hence an isomorphism. But let me make a very pedantic remark: this last statement is true, but not completely formal: something needs to be checked. The problem is that $\widehat{\mathbf{A}}_k$ is the restricted product of the $\widehat{k}_v$ with respect to the R_v^*, the annihiliators of R_v. [Reminder: for G an abelian LCHTG and H a closed subgroup, $H^* \subseteq \widehat{G}$ is the characters of G which are trivial on H.]

Hence to make sure that we really do get a continuous map $\mathbf{A}_k \to \widehat{\mathbf{A}}_k$ this way, it would suffice to check that our fixed local isomorphisms $k_v = \widehat{k}_v$ sent R_v to R_v^* for all v. But they don't! By definition, R_v^* is the characters of k_v that vanish on R_v, whereas our local isomorphism sends $r \in R_v$ to the function $y \mapsto e^{2\pi i \Lambda_v(ry)}$, so maps R_v to the functions which vanish on $\{x \in k_v : \Lambda_v(xy) = 0 \forall y \in R_v\}$ and this is the "inverse different" of k_v, which is not always equal to R_v. However, an explicit calculation shows that if v is unramified in the extension $k/\mathbf{Q}$ then this inverse different is R_v again (this calculation would take me too far afield at this point, unfortunately), and hence R_v becomes identified with R_v^* for all but finitely many v, which is good enough to ensure that we get an isomorphism $\mathbf{A}_k \to \widehat{\mathbf{A}}_k$ this way.

Hence for $f \in L^1(\mathbf{A}_k)$ we can (using our fixed choice of normalisations of Haar integrals and our fixed map $\mathbf{A}_k \to \widehat{\mathbf{A}_k}$) consider its Fourier transform as a function on $\mathbf{A}_k$ again. Explicitly

$$\widehat{f}(x) = \int f(y) e^{-2\pi i \Lambda(xy)} \mathrm{d}\mu(y).$$

And because our local Fourier transforms satisfied Fourier inversion on the nose, we check (by using a non-zero test function which is a product of L^1 functions on the factors) that

$$\widehat{\widehat{f}}(x) = f(-x)$$

for $f \in L^1(\mathbf{A}_k)$.

Let me finish this chapter with some comments on the relationship between the adeles of a number field and the adeles of a finite extension of this field. I'll stick to the case of $k/\mathbf{Q}$ but what I say is true for general extensions.

We showed, when analysing extensions of norms to finite field extensions, that a given norm $|.|$ on the bottom extends in at least one, but at most finitely many ways to a norm on the top. We showed something more precise, in fact—we showed that if L/K was a finite extension of fields of characteristic zero (or more generally a finite separable extension), and $|.|$ was a norm on K, and $\hat{K}$ was its completion (note: this hat has nothing to do with Pontrjagin duality), then $L \otimes_K \hat{K}$ was a finite sum of fields, and these fields were precisely the completions of L at the norms on L which extend $|.|$.

Applying this to the extension $k/\mathbf{Q}$, we find that $k \otimes_{\mathbf{Q}} \mathbf{Q}_p$ will be isomorphic to the direct sum of all the completions of k at all the norms extending the p-adic norm on $\mathbf{Q}$, and one can re-interpret the classical result "$\Sigma_i e_i f_i = [k : \mathbf{Q}]$" (with $(p) = \prod_i P^{e_i}$) as simply saying that these extensions must just be the P-adic norms for $p \in P$.

[Alternatively one can prove this directly, as is done in Cassels' book, and then derive this formula $\sum_i e_i f_i = [k : \mathbf{Q}]$ from it; there is a little work to be done here though, which I won't do]. The upshot is that

$$k \otimes_{\mathbf{Q}} \mathbf{Q}_p = \oplus_{p \in P} k_P$$

and the analogous result at infinity is that

$$k \otimes_{\mathbf{Q}} \mathbf{R} = \oplus_{[\tau]} k_\tau.$$

Now the closure of R, the integers of k, in $\oplus_{p \in P} k_P$, is just its completion in each component, which is $\prod_P R_P$, and from this it follows that

$$\mathbf{A}_k = \mathbf{A}_{\mathbf{Q}} \otimes_{\mathbf{Z}} R = \mathbf{A}_{\mathbf{Q}} \otimes_{\mathbf{Q}} k.$$

More generally one checks that for L/k a finite extension of number fields, the same proof gives that $\mathbf{A}_L = \mathbf{A}_k \otimes_k L$.

One can also deduce from these decompositions that traces and norms "can be computed locally". For example

$$\mathrm{Tr}_{k/\mathbf{Q}}(\lambda) = \mathrm{Tr}_{k \otimes_{\mathbf{Q}} \mathbf{Q}_p / \mathbf{Q}_p}(\lambda) = \sum_{p \in P} \mathrm{Tr}_{k_P / \mathbf{Q}_p}(\lambda)$$

and similar results for norms, and similar results at the infinite places too.

Chapter 6: The main theorem.

As you have surely realised by now, our strategy is as follows. We're going to define "global zeta integrals" as integrals of $f(x).|x|^s$ on the ideles, for f carefully-chosen functions. We are going to use things we've proved in the course to meromorphically continue these functions to all $s \in \mathbf{C}$. In the local case these meromorphic continuation proofs were of the form "check it for one f and deduce it for all f by some trick involving Fubini's theorem". In the global setting the result is deeper and we will obtain our meromorphic continuation from some adelic version of Poisson summation.

Recall that the crucial fact in the proof of the meromorphic continuation of the Riemann zeta function was that $\theta(1/t) = t\theta(t)$ for some function θ, and the proof of that latter fact came from some concrete form of the Fourier inversion theorem, which was just the statement that the Fourier series of a periodic function $F(x)$ did in fact converge to $F(x)$.

Tate's insight, which has run and run, is that in this adelic setting, the correct analogue of the set $\mathbf{R}/\mathbf{Z}$ is the set $\mathbf{A}_k/k$. Let me run off a few things we know about the inclusion $\mathbf{Z} \to \mathbf{R}$. Firstly, $\mathbf{Z}$ is discrete, $\mathbf{R}$ is locally compact, $\hat{\mathbf{R}}$ (the Pontrjagin dual) is isomorphic to $\mathbf{R}$ again, and if we use the isomorphism $x \mapsto (y \mapsto e^{-2\pi ixy})$ to identify $\mathbf{R}$ with $\hat{\mathbf{R}}$ then we see that the annihilator $\mathbf{Z}^*$ of $\mathbf{Z}$ (that is, the elements $r \in \mathbf{R}$ such that $e^{-2\pi irn} = 1$ for all integers n) is just $\mathbf{Z}$ again.

Hence the Pontrjagin dual of the discrete group $\mathbf{Z}$ is the compact group $\mathbf{R}/\mathbf{Z}$, and the dual of the exact sequence

$$0 \to \mathbf{Z} \to \mathbf{R} \to \mathbf{R}/\mathbf{Z} \to 0$$

is itself. Finally the action of $\mathbf{Z}$ on $\mathbf{R}$ admits a natural "fundamental domain" (that is, a subset D of $\mathbf{R}$, namely $[0, 1)$, with the property that the induced map $D \to \mathbf{R}/\mathbf{Z}$ is a bijection), and the measure of D, with respect to the standard Haar measure on $\mathbf{R}$, is 1.

We're going to prove analogues of all of these things today, with $\mathbf{Z}$ replaced by a number field k, and $\mathbf{R}$ replaced by $\mathbf{A}_k$. For example we'll soon see that k embeds into $\mathbf{A}_k$ as a discrete subgroup. So what will be the analogue of our proof of the functional equation of the theta function?

When thinking about the θ function, we obtained our function $F(x)$ originally as $F(x) = \sum_{n \in \mathbf{Z}} f(x+n)$, with $f(x) = e^{-\pi t^2 x^2}$ a function on $\mathbf{R}$. The analogue of f in this setting will be a carefully-chosen function on $\mathbf{A}_k$ which is sufficiently "rapidly decreasing", and such that for all adeles x, the sum $\sum_{\lambda \in k} f(x + \lambda)$ converges absolutely. We then apply Fourier inversion to get some fact, and show that this fact is precisely what is needed to give us the meromorphic continuation and functional equation of the Riemann zeta function and a gazillion other functions too, all of which come out in the wash.

Historical interlude (non-examinable).

The theory of automorphic forms was really getting off the ground in the 1950s, when Tate's thesis was written, but the classical theory tended to revolve around considering functions on groups like $\mathrm{GL}_n(\mathbf{R})$ which were invariant, or transformed in some simple way, under the subgroup $\mathrm{GL}_n(\mathbf{Z})$. In the 1950s there was a move away from this setting to the adelic setting of functions on $\mathrm{GL}_n(\mathbf{A}_k)$ which were invariant under the discrete subgroup $\mathrm{GL}_n(k)$, and this insight enabled one to reformulate various notions such as Hecke operators in a purely local form. Indeed, Hecke operators could now be interpreted as operators in a purely local setting coming from the representation theory of $\mathrm{GL}_n(k_v)$, giving a huge new impetus to the representation theory of p-adic groups.

There were practical consequences too in that the theory of Hecke operators for Hilbert modular forms was very difficult to set up globally, if the integers of the base field were not a PID, because no natural analogue at P of the matrix $\left(\begin{smallmatrix} p & 0 \\ 0 & 1 \end{smallmatrix}\right) \in \mathrm{GL}_2(\mathbf{Q})$ existed if P was a non-principal prime. The adelic reformulation of the theory removes all of these problems because even though P is not a principal ideal, the element $(1, 1, 1, 1, \ldots, \pi\, 1, 1, \ldots)$ (with π a uniformiser at P) is still a perfectly good idele (indeed this was one of Chevalley's motivations for introducing these things).

6.1. The additive theory, and the adelic Poisson summation formula.

Let's prove that $\mathbf{Z}$ is to $\mathbf{R}$ as k is to $\mathbf{A}_k$. Here's the first big reason for believing this:

Proposition. The subspace topology on k coming from the embedding $k \to \mathbf{A}_k$ is the discrete topology (all sets are open). And the quotient $\mathbf{A}_k/k$ is a compact topological space.

We'll prove this soon; first we'll construct a fundamental domain for k in $\mathbf{A}_k$, analogous to $[0,1)$ in $\mathbf{R}$. Let's do this by trying to understand how k fits into $\mathbf{A}_k$ "at the finite places", and then thinking about the infinite places.

Consider the group that I called G_{S_0} when setting up the general theory of restricted products: this is just

$$\prod_{v<\infty} R_v \times \prod_{v|\infty} k_v.$$

The intersection of k (embedded diagonally) with this group is the elements of k which are integers at all finite places. If $0 \neq \lambda \in k$ and we write the fractional ideal (λ) as $\prod_i P_i^{e_i}$, and if one of the e_i is negative, then we have $\lambda \notin R_{P_i}$, by definition. Hence the intersection

$$k \cap \left(\prod_{v<\infty} R_v \times \prod_{v|\infty} k_v \right)$$

is just the elements of k which generate integral ideals, which is just another way of saying the (global) integers R of k.

Now let's think about what's going on in $\prod_{[\tau]} k_\tau$, the infinite adeles. Note that this space just looks like $\mathbf{R} \oplus \mathbf{R} \oplus \ldots \oplus \mathbf{R} \oplus \mathbf{C} \oplus \mathbf{C} \oplus \ldots \oplus \mathbf{C}$, where there are, say, r copies of $\mathbf{R}$, and s copies of $\mathbf{C}$, and we also note that the number of field homomorphisms $k \to \mathbf{C}$ is just $r + 2s$ (recalling that we only get one completion for each pair of complex conjugate maps $k \to \mathbf{C}$).

Now if $e_1, e_2, \ldots, e_n$ is a $\mathbf{Z}$-basis for the ring of integers R in k, then the definition of the discriminant of k is just (up to sign) the square of the determinant of the square matrix $(\sigma_i(e_j))_{i,j}$, where σ_i runs through the field maps $k \to \mathbf{C}$. Let us write $|d|$ for the absolute value of the discriminant of $k/\mathbf{Q}$. For later use it will be helpful to know

Lemma. The image of R in $k_\infty = \prod_{[\tau]} k_\tau$ (embedded diagonally) is a lattice, and, with respect to our choices of Haar integrals on the k_τ, the measure of a fundamental domain for this lattice is just $\sqrt{|d|}$, with d the discriminant of k.

Remark. A fundamental domain for a lattice $\Lambda \subseteq \mathbf{R}^n$ is just a connected set S with non-empty interior such that every element of $\mathbf{R}^n$ can uniquely be written $\lambda + s$ with $\lambda \in \Lambda$ and $s \in S$. One way of constructing such a thing is to write down a basis $e_1, e_2, \ldots, e_n$ for Λ and let S be $\{\sum_i \lambda_i e_i\}$ with $0 \leq \lambda_i < 1$ for all i—a "fundamental parallelogram" for Λ.

Proof of lemma. If k is totally real (that is, all $k \to \mathbf{C}$ land in $\mathbf{R}$) then the result is immediate: the volume of the fundamental domain of a lattice in $\mathbf{R}^n$ is just the absolute value of the matrix whose entries form a basis for the lattice. But if k has complex places then we have to be a little careful.

The problem is that if σ is a map $k \to \mathbf{C}$ whose image does not land in $\mathbf{R}$, and if $\sigma(e_j) = x + iy$, then in the usual discriminant calculation (which uses *all* embeddings, both σ and $\overline{\sigma}$) we will see a contribution from $x + iy$ and $x - iy$. But in the infinite adele computation we only see σ, taking values in something we can thinking of as $\mathbf{R}^2$, giving us coordinates of x and y. Now we have

$$\begin{pmatrix} x+iy \\ x-iy \end{pmatrix} = \begin{pmatrix} 1 & i \\ 1 & -i \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}$$

and the absolute value of the determinant of $\left(\begin{smallmatrix} 1 & i \\ 1 & -i \end{smallmatrix}\right)$ is 2.

So with respect to the *naive* measure on the infinite adeles, (which is dxdy at the complex places) the volume of a fundamental domain for the lattice R will be $\sqrt{|d|}.2^{-s}$, because we lose a factor of 2 at each complex place. However the normalisation of Haar measure that we chose for the complex infinite places was not the naive one—we inserted a factor of 2! Hence with our fixed choice of Haar measure the volume is again $\sqrt{|d|}$. □

Now it's convenient to make the following definition. Define D_∞ to be the following fundamental parallelogram for the lattice R in $k_\infty = \prod_{v|\infty} k_v$: choose a $\mathbf{Z}$-basis $(e_j)_{1 \leq j \leq n}$ for R and consider the e_j as elements of k_∞; they form a lattice (because the discriminant of a number field is non-zero!). Define D_∞ to be the "box" whose typical element is $\sum_{j=1}^{n} \lambda_j e_j$ with $0 \leq \lambda_j < 1$. Note that the closure $\overline{D_\infty}$ of D_∞ is obtained by letting the λ_j range through $[0, 1]$, and the interior D_∞^o is obtained by restricting to λ_j in $(0, 1)$. In particular D_∞ has compact closure and non-empty interior.

Now let's define $D \subseteq \mathbf{A}_k$ to be the product $D_f \times D_\infty$, with $D_f \subseteq \mathbf{A}_{k,f}$ simply being $\prod_v R_v$. Note that D_f is an open subgroup, and hence a closed subgroup, of $\mathbf{A}_{k,f}$, and hence the closure of D in $\mathbf{A}_k$ is simply $D_f \times \overline{D_\infty}$, which is compact, and the interior is $D_f \times D_\infty^o$. I now claim that D is a fundamental domain for the action of k on $\mathbf{A}_k$. More precisely,

Lemma. Any element of $\mathbf{A}_k$ can be written uniquely as $d + \lambda$ with $d \in D$ and $\lambda \in k$.

Proof. Given an adele (g_v) of k, it is in R_P at all but finitely many finite places P, by definition. Choose some $0 \neq b \in R$ whose prime factorization contains a sufficiently high power of each of the P for which g_P isn't integral to ensure $bg_P \in R_P$ for all P. Now for each prime ideal P dividing (b), say P^e exactly divides b and consider the equation $a \equiv bg_P$ modulo P^e. By the Chinese Remainder Theorem these equations can all be solved at once within R, and we set $\lambda_0 = a/b \in k$. Now $g_P - \lambda_0 \in R_P$ for all $P|(b)$ and hence for all P, because g_P and a/b are integral at all other finite places.

We have now rigged it so that $(g_v) - \lambda_0$ has finite part in D_f, but its infinite part might not be in D_∞; however this can be fixed by subtracting an appropriate $\lambda_1 \in R$ (because D_∞ is visibly a fundamental domain for R acting on k_∞).

We see that we have written $\mathbf{A}_k = D + k$ now, and all that is left is to prove that this decomposition is unique. But this is easy: if $d_1 + \lambda_1 = d_2 + \lambda_2$ then $t := d_1 - d_2 = \lambda_2 - \lambda_1 \in (D - D) \cap k$, and looking at the finite places we see $t \in k$ is in $D_f - D_f = D_f$ is integral at all finite places, so $t \in R$, and looking at the infinite places we see $t = 0$ because 0 is the only element of $R = \sum_i \mathbf{Z} e_i$ in $D_\infty - D_\infty = \{\sum_{i=1}^n \lambda_i e_i : -1 < \lambda_i < 1\}$. $\square$

We can now prove something promised earlier:

Proposition. $k \subseteq \mathbf{A}_k$ is discrete and the quotient is compact.

Proof. Discreteness follows because D has a non-empty interior. More precisely, if d is any adele in the (non-empty) interior D^o of D then $D^o - d$ is an open set in $\mathbf{A}_k$ containing 0, and conversely if $\lambda \in k$ is in $D^o - d$ then we have $d' - d = \lambda$ for $d, d' \in D$ and hence $d' = d + \lambda$, so $\lambda = 0$. Hence $D^o - d$ is an open set in $\mathbf{A}_k$ whose intersection with k is just $\{0\}$ and hence for any $\alpha \in k$, $D^o - d + \alpha$ is an open set in $\mathbf{A}_k$ whose intersection with k is $\{\alpha\}$.

Compactness follows because $\mathbf{A}_k/k$ is a continuous image of $\overline{D}$; the lemma implies that the map is surjective. □

Now let's prove $\mathbf{A}_k/k$-analogues of the other $\mathbf{R}/\mathbf{Z}$-results we mentioned earlier.

Proposition. The measure of (the characteristic function of) D (with respect to our fixed choice of Haar measuse on $\mathbf{A}_k$) is 1.

Proof. $D = D_f \times D_\infty$. We computed the measure of D_∞ as $\sqrt{|d|}$. The way we normalised our local Haar measures at the P-adic places was such that if R_P is the integers of k_P then $\mu(R_P) = p^{-m/2}$, where p^m was the (absolute value of the) discriminant of $k_P/\mathbf{Q}_p$. But the global discriminant of $k/\mathbf{Q}$ is just the product of the local discriminants, and hence the measure of D_f with respect to our choices is $|d|^{-1/2}$. Hence the measure of D is the product of the measures of D_f and D_∞, which is 1! $\square$

Proposition. Our fixed isomorphism $\mathbf{A}_k \to \widehat{\mathbf{A}}_k$ (defined by $x \mapsto (y \mapsto e^{2\pi i \Lambda(xy)})$) sends the closed subgroup k isomorphically onto the closed subgroup k^* of characters of $\widehat{\mathbf{A}}_k$ which are trivial on k.

Reminder. Our fixed map $\mathbf{R} \to \widehat{\mathbf{R}}$ sends x to $y \mapsto e^{-2\pi i x y}$, so sends $\mathbf{Z}$ to the characters $y \mapsto e^{-2\pi i n y}$ for $n \in \mathbf{Z}$. The $\mathbf{R}/\mathbf{Z}$ analogue of this proposition is the statement that the intersection of the kernels of all of these characters is precisely $\mathbf{Z}$ again.

Proof of proposition. We need to check that the set of characters $y \mapsto e^{2\pi i \Lambda(ry)}$, for $r \in k$, is precisely the set of characters that vanish on k. So we need to check

(i) If $\alpha \in k$ then $\Lambda(\alpha) = 0$

(ii) If $y \in \mathbf{A}_k$ and $\Lambda(\alpha y) = 0$ for all $\alpha \in k$ then $y \in k$.

Recall $\Lambda((g_v)) = \sum_v \Lambda_v(g_v)$, a finite sum, and the Λ_v are "trace down to $\mathbf{Q}_p$ or $\mathbf{R}$, and then use $q : \mathbf{Q}_p/\mathbf{Z}_p \to \mathbf{Q}/\mathbf{Z}$ or $x \mapsto -x : \mathbf{R} \to \mathbf{R}/\mathbf{Z}$".

(i) is true because Λ is a sum of local traces, and if $\alpha \in k$ then $\mathrm{Tr}_{k/\mathbf{Q}}(\alpha) \in \mathbf{Q}$, and this reduces (i) to the case $k = \mathbf{Q}$. It's clearly true that $\Lambda(n) = 0$ for $n \in \mathbf{Z}$ (because all the Λ_p are zero), so by additivity it suffices to check that $\Lambda(1/p^e)$ vanishes for all p prime and $e \geq 1$. Now finally I realise why Tate inserted the minus sign in his definition of his local Λ for the reals: we have $\Lambda_q(1/p^e) = 0$ for all $q \neq p$, we have $\Lambda_p(1/p^e) = 1/p^e$ and we have $\Lambda_\infty(1/p^e) = -1/p^e$, and the sum in $\mathbf{R}/\mathbf{Z}$ is zero. So (i) is proved.

For (ii) we use a trick. We have proved $\mathbf{A}_k/k$ is compact, so if k^* denotes the annihilator of k in $\widehat{\mathbf{A}}_k$, and if we identify $\widehat{\mathbf{A}}_k$ with $\mathbf{A}_k$ via our fixed isomorphism, then we know $k \subseteq k^*$ from (i). Now k^* is the annihilator of k and hence the Pontrjagin dual of $\mathbf{A}_k/k$ which we've seen is compact. Hence k^* is discrete, and a closed subgroup of $\mathbf{A}_k$. So k^*/k is discrete in $\mathbf{A}_k/k$, and closed too, so it's compact, so it's finite.

So for $\alpha \in k^*$ there's some positive integer n such that $n\alpha = \beta \in k$. But $\beta/n \in k$ and $\beta/n - \alpha$ is now a torsion element of $\mathbf{A}_k$, which contains no torsion other than zero. So $\alpha = \beta/n \in k$. $\square$

So now we really see that the inclusion $k \subset \mathbf{A}_k$ is very formally similar to the inclusion $\mathbf{Z} \subset \mathbf{R}$. In particular we see that the Pontrjagin dual of $\mathbf{A}_k/k$ is $k^* = k$ (with the discrete topology), and hence the Pontrjagin dual of k as $\mathbf{A}_k/k$ (analogous to the Pontrjagin dual of $\mathbf{Z}$ being $\mathbf{R}/\mathbf{Z}$). But in some sense the advantage of the adeles over the reals is that the adeles "can be broken into local factors", and the arithmetic of k is easier than the arithmetic of its integers.

Now we prove the analogue of the transformation property of the θ function. Instead of working with (the analogue of) $f(x) = e^{-\pi t^2 x^2}$ we set things up, for the time being at least, in more generality: we'll use a general function f for which we'll just assume everything converges.

First we observe that we have a natural Haar measure on the compact group $\mathbf{A}_k/k$: a function on $\mathbf{A}_k/k$ can be thought of as a "periodic" function on $\mathbf{A}_k$ (that is, one satisfying $f(x+\alpha)=f(x)$ for $\alpha \in k$) and, for a continuous function of this type, one checks easily that defining $\mu(f) = \int_D f(x)\mathrm{d}\mu(x)$ where μ is our fixed Haar measure on the adeles but *the integral is only over our fundamental domain* D, gives us a Haar measure on $\mathbf{A}_k/k$.

On the other side, if we endow k with the discrete topology then a natural Haar measure is just counting measure: a continuous function with compact support is just a function $f : k \to \mathbf{C}$ which vanishes away from a finite set, and we can define $\mu(f) = \sum_{\alpha \in k} f(\alpha)$. With these choices of Haar measure on k and $\mathbf{A}_k/k$, what is the constant in the Fourier inversion theorem? In other words, if we invert $F : k \to \mathbf{C}$ twice, we'll get $x \mapsto cF(-x)$. What is c?

Lemma. $c = 1$.

Proof. We just need to check this for one non-zero function. So let's let F be the characteristic function of $\{0\}$. Then $\widehat{F}$ is the function on $\widehat{k} = \mathbf{A}_k/k$ which sends a character $\chi : k \to S^1$ to $\sum_{\alpha \in k} F(\alpha)\overline{\chi(\alpha)} = \overline{\chi(0)} = 1$. Hence $\widehat{F}$ is the constant function on $\mathbf{A}_k/k$, sending everything to 1. Now we don't have to evaluate $\widehat{\widehat{F}}$ everywhere, we only need to evaluate it at 0, regarded as the trivial character of $\mathbf{A}_k/k$. By the choice of our Haar measure on $\mathbf{A}_k/k$, we see

$$\widehat{\widehat{F}}(0) = \int_D 1 \mathrm{d}\mu(x)$$

and we computed the integral of D as being 1, so $\widehat{\widehat{F}}(0) = F(0)$, and hence $c = 1$.

(new lecture) Summary of where we are:

1) k a number field. We have fixed an identification of $\mathbf{A}_k$ with its Pontrjagin dual $\widehat{\mathbf{A}}_k$. We checked that this isomorphism sends the closed (and discrete) subgroup k of $\mathbf{A}_k$ isomorphically onto its annihiliator k^* (recall that the annihilator of k is just by definition the elements of $\widehat{\mathbf{A}}_k$ which vanish on k). We deduce from this that the Pontrjagin dual of $\mathbf{A}_k/k$ is isomorphic to k (because it's canonically isomorphic to k^*). We checked that k was a discrete subgroup of $\mathbf{A}_k$ and that the quotient $\mathbf{A}_k/k$ was compact. I remarked (and indeed stressed) that this was very much analogous to $\mathbf{Z} \subset \mathbf{R}$ being discrete and $\mathbf{R}/\mathbf{Z}$ being compact, the identification of $\mathbf{R}$ with its dual sending $\mathbf{Z}$ to its own annihilator, and $\mathbf{Z}$ hence being the dual of $\mathbf{R}/\mathbf{Z}$.

2) We fixed a choice of Haar measure on $\mathbf{A}_k/k$, namely $\int_D$, where D is our fundamental domain for the action of k on $\mathbf{A}_k$ (analogous to $[0,1)$ in $\mathbf{R}$). This choice has the nice property that the integral of the constant function is 1 (because $\mu(D) = 1$). We fixed a choice of Haar measure on k with the discrete topology, namely the "counting measure". We computed enough of the Fourier transform of the Fourier transform of the characteristic function of $\{0\}$ on k to deduce that the Fourier inversion theorem holds in this situation with constant term equal to 1. I remarked that, much to my annoyance, I did not know the full proof of the Fourier inversion theorem in this generality (but that it was certainly true, because it was true for all locally compact abelian groups, it's just that the proof involves too much analysis for me.)

Now a reminder of something from long ago. The Fourier inversion theorem on $\mathbf{R}/\mathbf{Z}$, when unravelled, just tells us the classical fact that if F is a continuous function on $\mathbf{R}/\mathbf{Z}$, viewed as a periodic function on $\mathbf{R}$, and if $a_m = \int_D F(x)e^{-2\pi imx}\mathrm{d}x$ is its mth Fourier coefficient, where $m \in \mathbf{Z}$ and $D = [0,1)$, and if $\sum_m |a_m|$ converges, then $F(x) = \sum_{m\in\mathbf{Z}} a_m e^{2\pi imx}$. We applied this very early on to a function $F(x)$ of the form $F(x) = \sum_{n\in\mathbf{Z}} f(x+n)$ where f was a function which was rapidly decreasing, and we deduced

$$\sum_{n\in\mathbf{Z}} f(n) = F(0) = \sum_{m\in\mathbf{Z}} a_m.$$

And we computed a_m using this trick:

$$\begin{aligned} a_m &= \int_0^1 \sum_n f(x+n)e^{-2\pi imx}\mathrm{d}x \\ &= \int_0^1 \sum_n f(x+n)e^{-2\pi im(x+n)}\mathrm{d}x \\ &= \int_{\mathbf{R}} f(x)e^{-2\pi imx}\mathrm{d}x = \widehat{f}(m) \end{aligned}$$

and so $\sum_{n\in\mathbf{Z}} f(n) = \sum_{m\in\mathbf{Z}} \widehat{f}(m)$ as long as everything converges—this is the classical Poisson summation fomula.

Let's now do exactly the same thing, but on $\mathbf{A}_k/k$ instead of $\mathbf{R}/\mathbf{Z}$.

If $F \in L^1(\mathbf{A}_k/k)$ (that is, F is a function on the adeles and $F(x+\alpha) = F(x)$ for $\alpha \in k$, and furthermore if $\int_D F(x)\mathrm{d}\mu(x) < \infty$), then let's define $\widehat{F} : k \to \mathbf{C}$ by

$$\widehat{F}(\alpha) = \int_D F(x)e^{-2\pi i\Lambda(x\alpha)}\mathrm{d}\mu(x).$$

Lemma. With notation as above, if $\sum_{\alpha\in k} |\widehat{F}(\alpha)|$ converges, then

$$F(x) = \sum_{\alpha\in k} \widehat{F}(\alpha)e^{2\pi i\Lambda(\alpha x)}.$$

Proof. This is just the Fourier inversion theorem spelt out, together with the fact that $c = 1$, which we proved last time. □

Corollary. $F(0) = \sum_{\alpha\in k} \widehat{F}(\alpha)$. □

Remark. As I've mentioned already, I'm slightly "bothered" by the fact that I've not actually proved the Fourier inversion theorem. However the proof for $\mathbf{R}/\mathbf{Z}$ is not hard, and the proof for $\mathbf{Q}_p$ can be done by hand, and it looks to me like $\mathbf{A}_k/k$ is built up from things that look like this, and so I wonder whether one would be able to give a "hands-on" proof, avoiding all the functional analysis which I had to assume.

One last explicit definition: if $f \in L^1(\mathbf{A}_k)$ then, surprise surprise, define $\widehat{f} : \mathbf{A}_k \to \mathbf{C}$ by $\widehat{f}(y) = \int_{\mathbf{A}_k} f(x) e^{-2\pi i \Lambda(xy)} \mathrm{d}\mu(x)$, the usual Fourier transform, once we have identified $\mathbf{A}_k$ with its dual.

Theorem (Poisson summation, revisited.) If $f \in L^1(\mathbf{A}_k)$ is continuous, if $\sum_{\alpha \in k} f(x+\alpha)$ converges absolutely and uniformly for $x \in \mathbf{A}_k$, and if $\sum_{\alpha \in k} |\widehat{f}(\alpha)|$ also converges, then

$$\sum_{\beta \in k} f(\beta) = \sum_{\alpha \in k} \widehat{f}(\alpha).$$

Proof. (c.f. section 1.2.) Define $F : \mathbf{A}_k \to \mathbf{C}$ by $F(x) = \sum_{\beta \in k} f(x+\beta)$. Now by assumption the sum converges uniformly on $\mathbf{A}_k$, so F is continuous and periodic. Hence F, considered as a function on $\mathbf{A}_k/k$, is continuous with compact support and is hence L^1. Moreover, for $\alpha \in k$ we have (c.f. formula for a_m in 1.2)

$$\begin{aligned}
\widehat{F}(\alpha) &= \int_D F(x) e^{-2\pi i \Lambda(\alpha x)} \mathrm{d}\mu(x) \\
&= \int_D \sum_{\beta \in k} f(x+\beta) e^{-2\pi i \Lambda(\alpha x)} \mathrm{d}\mu(x) \\
&= \sum_{\beta \in k} \int_D f(x+\beta) e^{-2\pi i \Lambda(\alpha x)} \mathrm{d}\mu(x) \\
&= \sum_{\beta \in k} \int_D f(x+\beta) e^{-2\pi i \Lambda(\alpha (x+\beta))} \mathrm{d}\mu(x) \\
&= \int_{\mathbf{A}_k} f(x) e^{-2\pi i \Lambda(\alpha x)} \mathrm{d}\mu(x) \\
&= \widehat{f}(\alpha)
\end{aligned}$$

[where the interchange of sum and integral is OK because the sum converges uniformly on D, which has finite measure, and I've also used the fact (proved earlier) that $k \subset \ker(\Lambda)$, which I proved when showing $k = k^*$.] Hence

$$\begin{aligned}
\sum_{\beta \in k} f(\beta) &= F(0) \\
&= \sum_{\alpha \in k} \widehat{F}(\alpha) \\
&= \sum_{\alpha \in k} \widehat{f}(\alpha)
\end{aligned}$$

□

6.2 The multiplicative theory.

We just showed that $k \subseteq \mathbf{A}_k$ was discrete, with compact quotient. We'll now show that $k^\times \subseteq \mathbf{A}_k^\times$ is discrete, but perhaps one doesn't expect the quotient to be compact, because $\mathbf{R}^\times/\mathbf{Z}^\times \cong \mathbf{R}_{>0}$ isn't compact.

In fact here's a proof that $\mathbf{A}_k^\times/k^\times$ isn't compact. Recall that we have a norm function $|.| : \mathbf{A}_k^\times \to \mathbf{R}_{>0}$, defined as a product of local norms.

Lemma. If $\alpha \in k^\times$ then $|\alpha| = 1$.

Proof. Lazy proof: $\mathbf{A}_k = \mathbf{A}_\mathbf{Q} \otimes_\mathbf{Q} k$ and $|.|$ factors through the norm map $\mathbf{A}_k \to \mathbf{A}_\mathbf{Q}$ (if you believe that the P-adic norms are the only norms on k extending the p-adic norm on $\mathbf{Q}$, which is true and not hard and in Cassels, but I didn't prove it). This reduces us to the case $k = \mathbf{Q}$. In this case, by multiplicativity of the norm, we need only check the cases $\alpha = -1$ and $\alpha = p$ prime.

Now $\alpha = -1$ is a global unit so has local norm equal to 1 everywhere, and $\alpha = p$ also has global norm 1 because $|\alpha|_q = 1$ for all $q \neq p$, $|\alpha|_p = p^{-1}$ and $|\alpha|_\infty = p$, so the product of the local norms is 1. $\square$

Remark. It's not hard to give a direct computational proof for general k. Tate also notes that there's a "pure thought" proof which goes as follows: $|\alpha|$ is the factor by which additive Haar measure on the adeles is scaled, and because $\mu(D) = 1$ we will have $|\alpha| = \mu(\alpha D)$. But αD is a fundamental domain for $\alpha k = k$ and it's not hard to check now that $\mu(\alpha D)$ must then be $\mu(D)$ [consider $\alpha D = \cup_{\beta \in k}(\alpha D \cap (D + \beta))$ etc to see that fundamental domains have the same measure.]

Now it's clear that $|.| : \mathbf{A}_k^\times \to \mathbf{R}_{>0}$ is surjective (it's even surjective when you restrict to one infinite place), so certainly one can't hope that $\mathbf{A}_k^\times / k^\times$ is compact (because it has $\mathbf{R}_{>0}$ as a homomorphic image).

Definition. Let J be the kernel of $|.|$, with the subspace topology coming from $\mathbf{A}_k^\times$. We have "dropped one factor of $\mathbf{R}_{>0}$" going from $\mathbf{A}_k^\times$ to J. But it's enough, because

Proposition. $k^\times$ is a discrete subgroup of J and $J/k^\times$ is compact.

Proof. We follow the same strategy for showing k is discrete in $\mathbf{A}_k$, but we'll need some standard facts about class groups and unit groups of number fields, which of course I'll assume. In fact the proposition is equivalent to the union of the following statements: the rank of the unit group of k is $r+s-1$ (with r the number of real and s the number of complex places), the regulator is non-zero (which comes out of the standard proof of the unit group rank statement), the number of roots of unity in k is finite, and the class number of k is finite.

Don't take the following proof too seriously: we don't really need the precise volumes that come out. Just believe that the proof is "the same as in the additive case, but messier."

So here's how the argument goes (c.f. the construction of D). Define $\tilde{E}_f = \prod_{v<\infty} R_v^\times \subset (\mathbf{A}_k^f)^\times$. (I'm putting tildes on because the $\tilde{E}$ I'm about to build won't quite be a fundamental domain). Then $k^\times \cap \tilde{E}_f$ is the elements of $k^\times$ that are units at all finite places, and hence when written a/b have $(a) = (b)$; this is just the units $R^\times$ of $R \subset k$. Our choices of Haar measure imply that $\mu^*(\tilde{E}_f) = |d|^{-1/2}$.

At the infinite places we take logs: the product of the maps $\log(|.|) : k_\tau^\times \to \mathbf{R}$ give us a map $R^\times \to \mathbf{R}^{r+s}$ whose image lands in the hyperplane consisting of vectors the sum of whose entries is zero. Now it's a standard result that the image of $R^\times$ is a lattice in this hyperplane, and the kernel is the roots of unity. Let $\tilde{L}_\infty$ be a fundamental domain for this lattice, and we let $\tilde{E}_\infty$ be the pre-image of $\tilde{L}_\infty$ in $\ker(|.|) : k_\infty^\times \to \mathbf{R}_{>0}$; then $\tilde{E} := \tilde{E}_f \times \tilde{E}_\infty$ has measure $|d|^{-1/2}.2^r(2\pi)^s \operatorname{Reg}_k$.

Explanation: the discriminant factor comes from the finite places, the 2^r and $(2\pi)^s$ coming from the units at the infinite places, which were killed by the logs, and Reg_k is, by definition, the volume of the fundamental domain of $\tilde{L}_\infty$, which is by definition the *regulator* of the number field and is known to be non-zero and finite. Moreover, $\tilde{E}$ is *almost* a fundamental domain for $k^\times \subseteq J$. The problems are firstly that we lost track of the roots of unity (so $\tilde{E}$ is too big by a factor of the number of roots of unity) and secondly that we cannot multiply any finite idele by some element of $k^\times$ to put us in $\tilde{E}_f$ (the "multiplicative" version of the CRT argument fails), so $\tilde{E}$ is too small by a factor of $k^\times \backslash (\mathbf{A}_k^f)^\times / \tilde{E}_f$, and $(\mathbf{A}_k^f)^\times / \tilde{E}_f = \oplus_{v<\infty} \mathbf{Z} v$ is the group of fractional ideals, so its quotient by $k^\times$ is the class group of k, which is known to be finite. One now checks that $\tilde{E}$ can be modified "a finite amount" to ensure that it becomes a fundamental domain E for $k^\times$ in J, with measure $|d|^{-1/2}.2^r(2\pi)^s \,\mathrm{Reg}_k\, h/w$ with h the class number and w the number of roots of unity.

As I say, don't take all those delicate numbers too seriously, but do note that E has compact closure and non-empty interior, and that $J = \cup_{\alpha \in k^\times} \alpha E$ a disjoint union, so $k^\times$ is discrete in J with compact quotient. $\square$

Here's a nice consequence of compactness. Note that just as in the local case we consider quasicharacters of multiplicative groups, rather than just characters.

Corollary. If $c : k^\times \backslash \mathbf{A}_k^\times \to \mathbf{R}_{>0}$ is a continuous group homomorphism, then $c = |.|^\sigma$ for some real number σ.

Proof. $c(k^\times \backslash J)$ is a compact subgroup of $\mathbf{R}_{>0}$ and is hence $\{1\}$. So c factors through $\mathbf{A}_k^\times / J$ which, via the norm map, is $\mathbf{R}_{>0}$, and now taking logs we're done, because the only continuous group homomorphisms $\mathbf{R} \to \mathbf{R}$ are $x \mapsto \sigma x$.

6.3: Statement and proof of the main theorem.

Definitions. If $c : k^\times \backslash \mathbf{A}_k^\times \to \mathbf{C}^\times$ is a continuous group homomorphism then we say it's a *quasi-character* of $k^\times \backslash \mathbf{A}_k^\times$. We've just seen that $|c| : k^\times \backslash \mathbf{A}_k^\times \to \mathbf{R}_{>0}$ is of the form $x \mapsto |x|^\sigma$; define $\mathrm{Re}(c) = \sigma$. We let the set of quasi-characters of $k^\times \backslash \mathbf{A}_k^\times$ be a Riemann surface as in the local case, by letting the component of $c : k^\times \backslash \mathbf{A}_k^\times \to \mathbf{C}^\times$ be $\{c.|.|^s : s \in \mathbf{C}\}$. Note that in this case the Riemann surface is just an infinite union of copies of the complex numbers, indexed by the group $\widehat{J}$ of characters of J. If c is a quasi-character of $k^\times \backslash \mathbf{A}_k^\times$ then let $\widehat{c}$ be the character $x \mapsto |x|/c(x)$; note that $\mathrm{Re}(\widehat{c}) = 1 - \mathrm{Re}(c)$.

Remark. I know very little about $\widehat{J}$.

Recall that in the local setting we had a set Z consisting of "functions for which everything converged", and defined $\zeta(f, c)$ for $f \in Z$ and c a quasi-character with positive real part, as some sort of integral. Here's the analogy of this construction in the global setting.

Let Z denote the set of functions $f : \mathbf{A}_k \to \mathbf{C}$ satisfying the following "boundedness" conditions:

Firstly, we demand f is continuous and in $L^1(\mathbf{A}_k)$, and also that $\widehat{f} : \mathbf{A}_k \to \mathbf{C}$ is continuous and in $L^1(\mathbf{A}_k)$.

Secondly (a condition that wasn't present in the local setting), we demand that for every $y \in \mathbf{A}_k^\times$, the sums $\sum_{\alpha \in k} f(y(x+\alpha))$ and $\sum_{\alpha \in k} \widehat{f}(y(x+\alpha))$ converge absolutely, and moreover the convergence is "locally uniform" in the sense that it's uniform for $(x, y) \in D \times C$ for D our additive fundamental domain and C an arbitrary compact subset of $\mathbf{A}_k^\times$.

Thirdly, we demand that $f(y).|y|^\sigma : \mathbf{A}_k^\times \to \mathbf{C}^\times$ and $\widehat{f}(y).|y|^\sigma$ are in $L^1(\mathbf{A}_k^\times)$ for all $\sigma > 1$ (note: this was $\sigma > 0$ in the local setting).

What are the reasons for these conditions? The first two mean that we can apply Poisson summation to f and indeed to the map $x \mapsto f(yx)$ for any $y \in \mathbf{A}_k^\times$. The local uniform convergence in the second condition is so that we can interchange a sum and an integral at a crucial moment. The third condition means that our global "multiplicative zeta integral" will converge for $\mathrm{Re}(s) > 1$.

Definition. If $f \in Z$ and $c : k^\times \backslash \mathbf{A}_k^\times \to \mathbf{C}^\times$ is a quasi-character with $\mathrm{Re}(c) > 1$, define

$$\zeta(f, c) = \int_{\mathbf{A}_k^\times} f(y) c(y) \mathrm{d}\mu^*(y)$$

(the Haar measure on $\mathbf{A}_k^\times$ being, of course, the product of our fixed Haar measures μ_v^* on $k_v^\times$).

The last condition in the definition of Z ensures the integral converges. Our main goal is:

Theorem. If $f \in Z$ then the function $\zeta(f,.)$ is holomorphic on the Riemann surface of quasi-characters c with $\mathrm{Re}(c) > 1$, and has a meromorphic continuation to all quasi-characters. Assume furthermore that $f(0) \neq 0$ and $\widehat{f}(0) \neq 0$. Then $\zeta(f,.)$ has simple poles at the quasi-characters $c(x) = 1$ and $c(x) = |x|$, and no other poles (and \$1,000,000 attached to its zeros). Finally it satisfies the (very elegant!) functional equation

$$\zeta(f,c) = \zeta(\widehat{f},\widehat{c}).$$

We'll now start the proof of this, which of course is going to be a not-too-tough application of everything we have. But what else do we need to do in this course? Well the only other thing to do is to check that the theorem has some content—that is, that Z contains some non-zero functions and that, as special cases of the theorem, we are proving the meromorphic continuation of Dirichlet L-functions, zeta functions of number fields, zeta functions of Grössencharacters,. . . .

[extra $\widehat{f}$ L^1 condition.] Before we prove the theorem let me make some definitions and prove some lemmas. We have $J \subseteq \mathbf{A}_k^\times$, the kernel of the norm function. Just as in the local case let's split this by finding $I \subset \mathbf{A}_k^\times$ isomorphic to $\mathbf{R}_{>0}$ such that $\mathbf{A}_k^\times = I \times J$. We do this by just choosing an infinite place $[\tau_0]$ of k and letting I be the copy of the positive reals in $k_{\tau_0}^\times$. We identify I with $\mathbf{R}_{>0}$ so that the norm map induces the identity $\mathbf{R}_{>0} \to \mathbf{R}_{>0}$, so if τ_0 happens to be a complex place then, because our complex norms aren't standard, what we're doing here is letting I be the positive reals in $\mathbf{C}^\times$ but letting the map $\mathbf{R}_{>0} \to I$ be $t \mapsto \sqrt{t}$.

For $f \in Z$ and $\mathrm{Re}(c) > 1$, we firstly break off this factor of I in the definition of the zeta integral: we write

$$\begin{aligned}\zeta(f,c) &= \int_{I\times J} f(y)c(y)\mathrm{d}\mu^*(y) \\ &= \int_{t=0}^{\infty}\int_{b\in J} f(tb)c(tb)\mathrm{d}\mu^*(b)\mathrm{d}t/t \\ &= \int_{t=0}^{\infty} \zeta_t(f,c)\mathrm{d}t/t\end{aligned}$$

where our measure on J is the one such that its product with $\mathrm{d}t/t$ on I gives us μ^* on $\mathbf{A}_k^\times$, and the last line is the definition of $\zeta_t(f,c) := \int_J f(tb)c(tb)\mathrm{d}\mu^*(b)$.

Let's think a little about

$$\zeta_t(f,c) = \int_J f(tb)c(tb)\mathrm{d}\mu^*(b).$$

We know that the integral defining $\zeta(f,c)$ converges, by assumption, for $\mathrm{Re}(c) > 1$, and hence the integrals defining $\zeta_t(f,c)$ will converge (at least for all t away from a set of measure zero). But these integrals are very docile: for $b \in J$ we have $|b| = 1$ by definition, so if $\mathrm{Re}(c) = \sigma$ then $|c(tb)| = |tb|^\sigma = t^\sigma$ is constant on J, and hence if the integral defining $\zeta_t(f,c)$ converges for one quasi-character c (which it almost always does) then it converges for all of them.

$[\zeta_t(f,c) = \int_J f(tb)c(tb)\mathrm{d}\mu^*(b).]$

The problem, of course, is not in the convergence of the individual $\zeta_t(f,c)$; it's that as t goes to zero then $f(tb)$ will be approaching $f(0)$ and if this is non-zero, which it typically will be, then the integral of this function over the non-compact J might be getting very big, so $\int_{t=0}^{1} \zeta_t(f,c)\mathrm{d}t/t$ will probably diverge if, say, $\sigma < 0$ (because then t^σ is also getting big). This is the problem we have to solve.

Note also that we've written

$$\zeta(f,c) = \int_{t=0}^{\infty} \zeta_t(f,c)\mathrm{d}t/t$$

and that this is one of the crucial tricks. If $f = \prod_v f_v$ with f_v on k_v then we could compute the global integral as a product of local integrals—but in applications this would just tell us that our global zeta function is a product of local zeta functions, which will not help with the meromorphic continuation. The insight is to compute the integral in this second way. Note that Iwasawa independently had this insight in 1952.

In case you've not realised, let me stress that $\zeta(f, c)$ isn't a generalisation of the zeta function, it's a generalisation of $\xi(s)$, that is, the zeta function multiplied by the fudge factor at infinity, and the t in the integral above is precisely the t that we had at the beginning of section 1.3 right at the beginning of the course. The strategy is now clear: we break the integral over t up into two parts, one of which will converge for all c, and the other of which we will manipulate and, by making the substitution $u = 1/t$ and applying Poisson summation, turn into a form which also converges.

Recall that the closure of the fundamental domain E for $k^\times$ in J is compact, so the integrals below are finite (as the integrands are continuous). Using $J = k^\times.E$ we get

$$\begin{aligned}\zeta_t(f,c) &= \sum_{\alpha\in k^\times}\int_{\alpha E} f(tb)c(tb)\mathrm{d}\mu^*(b)\\ &= \sum_{\alpha\in k^\times}\int_{E} f(t\alpha b)c(tb)\mathrm{d}\mu^*(b)\\ &= \int_E\left(\sum_{\alpha\in k^\times} f(t\alpha b)\right)c(tb)\mathrm{d}\mu^*(b)\end{aligned}$$

where the first equality is the definition, the second uses the fact that μ^* is a multiplicative Haar measure on J and that c is trivial on $k^\times$, and the third is an interchange of a sum and an integral which is justified by our rather strong uniform convergence assumptions on $f \in Z$ and the observation that the closure of E is a compact subset of $\mathbf{A}_k^\times$.

Exactly the same argument (changing f to $\widehat{f} \in Z$, c to $\widehat{c}$ and t to $1/t$) shows that blah $\zeta_{1/t}(\widehat{f},\widehat{c}) = \int_E \left(\sum_{\alpha\in k^\times}\widehat{f}(\alpha b/t)\right)\widehat{c}(b/t)\mathrm{d}\mu^*(b)$.

$$[\zeta_t(f,c) = \int_E \left(\sum_{\alpha\in k^\times} f(t\alpha b)\right) c(tb)\mathrm{d}\mu^*(b)]$$

Now that sum over $k^\times$ looks almost like a sum over k, but firstly the term $\alpha = 0$ is missing (so we'll have to add it in) and secondly we're not summing $f(\alpha)$ but $f(t\alpha b)$. So we'll have to work out what the Fourier transform of $x \mapsto f(txb)$ is. In other words, we need to see how the additive Fourier transform scales under multiplication. In the application of the lemma below we'll have $\rho = tb$.

Lemma. If $f : \mathbf{A}_k \to \mathbf{C}$ is continuous and in $L^1(\mathbf{A}_k)$, if $\rho \in \mathbf{A}_k^\times$ is fixed and if $g(x) := f(x\rho)$ then $\widehat{g}(y) = \frac{1}{|\rho|}\widehat{f}(y/\rho)$.

Proof. An elementary computation. We have

$$\widehat{g}(y) = \int_{\mathbf{A}_k} f(x\rho)e^{-2\pi i\Lambda(xy)}\mathrm{d}\mu(x)$$

and setting $x' = x\rho$ we have $\mathrm{d}\mu(x') = |\rho|\mathrm{d}\mu(x)$ and hence

$$\begin{aligned}\widehat{g}(y) &= \int_{\mathbf{A}_k} f(x')e^{-2\pi i\Lambda(x'y/\rho)}\mathrm{d}\mu(x')/|\rho| \\ &= \frac{1}{|\rho|}\widehat{f}(y/\rho)\end{aligned}$$

as required. □

So now let's add in the missing $\alpha = 0$ term to $\zeta_t(f,c)$, apply Poisson summation, and see what happens. Recall we just showed that $\zeta_t(f,c) = \int_E \left(\sum_{\alpha\in k^\times} f(t\alpha b)\right) c(tb)\mathrm{d}\mu^*(b)$ and that $\zeta_{1/t}(\widehat{f},\widehat{c}) = \int_E \left(\sum_{\alpha\in k^\times} \widehat{f}(\alpha b/t)\right) \widehat{c}(b/t)\mathrm{d}\mu^*(b)$.

Key Lemma. For an arbitrary $t > 0$ and c we have

$$\begin{aligned}&\zeta_t(f,c) + f(0)\int_E c(tb)\mathrm{d}\mu^*(b) \\ &= \zeta_{1/t}(\widehat{f},\widehat{c}) + \widehat{f}(0)\int_E \widehat{c}(b/t)\mathrm{d}\mu^*(b).\end{aligned}$$

Proof. As we've already remarked, the formulas we have just derived for $\zeta_t(f,c)$ and $\zeta_t(\widehat{f},\widehat{c})$ involve sums of $\alpha \in k^\times$.

The LHS of the lemma is hence what you get when you add the missing $\alpha = 0$ term: it's

$$\int_E \left(\sum_{\alpha \in k} f(t\alpha b) \right) c(tb) \mathrm{d}\mu^*(b) \qquad (1).$$

Similarly the RHS is

$$\int_E \left(\sum_{\alpha \in k} \widehat{f}(\alpha b/t) \right) \widehat{c}(b/t) \mathrm{d}\mu^*(b) \qquad (2).$$

So we need to show (1) = (2). The internal sum over k screams out for an application of Poisson summation, which, when applied to the function $x \mapsto f(txb)$ (we're allowed to apply Poisson summation because of our assumptions on f) gives

$$\sum_{\alpha \in k} f(t\alpha b) = \sum_{\alpha \in k} \widehat{(x \mapsto f(txb))}(\alpha) = \sum_{\alpha \in k} \frac{1}{|tb|} \widehat{f}(\alpha/tb).$$

Hence formula (1) is equal to

$$\int_E \left(\sum_{\alpha \in k} \widehat{f}(\alpha/tb) \right) c(tb)/|tb| \mathrm{d}\mu^*(b)$$

and now making the substitution $b \mapsto 1/b$, which doesn't change Haar measure, this becomes

$$\int_E \left(\sum_{\alpha \in k} \widehat{f}(\alpha b/t) \right) c(t/b)|b|/|t| \mathrm{d}\mu^*(b)$$

$$= \int_E \left(\sum_{\alpha \in k} \widehat{f}(\alpha b/t) \right) \widehat{c}(b/t) \mathrm{d}\mu^*(b)$$

which is (2)! This proves the lemma. □

We're finally ready to meromorphically continue our global zeta integrals. But before we do, let's try and figure out exactly what that fudge factor was that we had to add to $\zeta_t(f,c)$ to make that argument work in that last lemma: we added $f(0)$ times

$$\int_E c(tb)\mathrm{d}\mu^*(b).$$

What is this? Well if $c(x) = |x|^s$ is trivial on J then $c(tb) = t^s$ is constant for $b \in E$ (indeed, for $b \in J$), so the integral is just $t^s\mu^*(E)$ and we computed the measure of E earlier to be $2^r(2\pi)^s hR/(w\sqrt{|d|})$—it's some finite non-zero number, anyway. But if c is non-trivial on J then, because it's always trivial on $k^\times$, it descends to a non-trivial character on the compact group $J/k^\times$"="E and the integral will hence be zero (distinct characters are orthogonal). So in fact we have

Corollary. If c is non-trivial on J and $f \in Z$ and $t > 0$ then $\zeta_t(f,c) = \zeta_{1/t}(\widehat{f},\widehat{c})$.

We're finally ready to prove the main theorem! I'll re-state it.

Theorem. If $f \in Z$ then the function $\zeta(f,.)$ is holomorphic on the Riemann surface of quasi-characters c with $\mathrm{Re}(c) > 1$, and has a meromorphic continuation to all quasi-characters. Assume furthermore that $f(0) \neq 0$ and $\widehat{f}(0) \neq 0$. Then $\zeta(f,.)$ has simple poles at the quasi-characters $c(x) = 1$ and $c(x) = |x|$, and no other poles, (and \$1,000,000 attached to its zeros). Finally it satisfies the functional equation

$$\zeta(f,c) = \zeta(\widehat{f},\widehat{c}).$$

Proof. For $\mathrm{Re}(c) > 1$ the LHS zeta integral converges (by assumption on f) and is holomorphic in the c variable (differentiate under the integral). By definition, $\zeta(f,c) = \int_{t=0}^{\infty} \zeta_t(f,c)\mathrm{d}t/t$, which converges by assumption for $\mathrm{Re}(c) > 1$, and now we break the integral up into two parts:

$$\zeta(f,c) = \int_{t=1}^{\infty} \zeta_t(f,c)\mathrm{d}t/t + \int_{t=0}^{1} \zeta_t(f,c)\mathrm{d}t/t.$$

Now just as in the argument for the classical zeta function, I claim that the integral for $t \geq 1$ converges for all c, because the ideles tb showing up in the integral all have $|tb| = |t||b| = |t| \geq 1$ so if the integral converges for e.g. Re$(c) = 2$ (which it does, by assumption, as $2 > 1$) then it converges for any c with Re$(c) < 2$ (because the integrand is getting smaller).

That term isn't the problem. The problem term is the integral from 0 to 1, which typically only converges for Re$(c) > 1$. So let's use the previous lemma, which has some content (Poisson summation) and see what happens. The simplest case is if $c(x) \neq |x|^s$ for any s (that is, c is non-trivial on J). In this case those extra fudge factors in the previous lemma disappear, and we see

$$\begin{aligned}\int_{t=0}^{1} \zeta_t(f,c)\mathrm{d}t/t &= \int_{t=0}^{1} \zeta_{1/t}(\widehat{f},\widehat{c})\mathrm{d}t/t \\ &= \int_{u=1}^{\infty} \zeta_u(\widehat{f},\widehat{c})\mathrm{d}u/u\end{aligned}$$

and this last integral also converges for all quasi-characters $k^\times\backslash\mathbf{A}_k^\times \to \mathbf{C}^\times$ because $u \geq 1$ so convergence again gets better as Re(c) gets smaller. Moreover our new expression for $\zeta(f,c)$, namely

$$\zeta(f,c) = \int_{t=1}^{\infty} \zeta_t(f,c)\mathrm{d}t/t + \int_{u=1}^{\infty} \zeta_u(\widehat{f},\widehat{c})\mathrm{d}u/u$$

converges for all c and makes it clear that $\zeta(f,c) = \zeta(\widehat{f},\widehat{c})$ (and that it's holomorphic for all c not in the component $|.|^s$). The proof is complete in this case!

We're not quite finished though: we need to deal with the component $c(x) = |x|^s$, where the argument is slightly messier because we pick up factors of $f(0)\int_E c(tb)\mathrm{d}\mu^*(b) = f(0)t^s\mu^*(E)$ and $\widehat{f}(0)\int_E \widehat{c}(\frac{1}{t}b)\mathrm{d}\mu^*(b)$. In this case (writing $c(x) = |x|^s$ now), the extra factors we'll see in the calculation will be (for Re(s) > 1)

$$f(0)\mu^*(E)\int_{t=0}^{1} t^s\mathrm{d}t/t$$
$$= f(0)\mu^*(E)[t^s/s]_0^1 = f(0)\mu^*(E)/s$$

and

$$\int_{t=0}^{1}(\widehat{f}(0)\int_E |b/t|^{1-s}\mathrm{d}\mu^*(b))\mathrm{d}t/t$$
$$= \widehat{f}(0)\mu^*(E)\int_{t=0}^{1} t^{s-2}$$
$$= \widehat{f}(0)\mu^*(E)[t^{s-1}/(s-1)]_0^1 = \widehat{f}(0)\mu^*(E)/(s-1).$$

These functions (cst $/s$ and cst $/(s-1)$) clearly have a meromorphic continuation to $s \in \mathbf{C}$! So we have, for $c(x) = |x|^s$ with $\mathrm{Re}(s) > 1$,

$$\begin{aligned}\zeta(f,c) &= \int_{t=1}^{\infty}\zeta_t(f,c)\mathrm{d}t/t + \int_{t=0}^{1}\zeta_t(f,c)\mathrm{d}t/t \\ &= \int_{t=1}^{\infty}\zeta_t(f,c)\mathrm{d}t/t + \int_{u=1}^{\infty}\zeta_u(\widehat{f},\widehat{c})\mathrm{d}u/u \\ &+ \mu^*(E)(-f(0)/s + \widehat{f}(0)/(s-1))\end{aligned}$$

and now we really have proved the theorem because this latter expression makes sense as a meromorphic function for all $s \in \mathbf{C}$, the integrals are all holomorphic for all $s \in \mathbf{C}$, and the expression is invariant under $(f,c) \mapsto (\widehat{f},\widehat{c})$. □

We've even computed the residues of $\zeta(f,|.|^s)$ at $s = 0$ and $s = 1$; they've come out in the wash.

They are $-f(0)\mu^*(E)$ and $\widehat{f}(0)\mu^*(E)$ respectively. Recall that we computed $\mu^*(E) = 2^r(2\pi)^s \operatorname{Reg}_k h/w\sqrt{|d|}$.

Short chapter 7: Applications!

We have left open the logical possibility that $Z = \{0\}$, in which case our theory is empty. Let's check it isn't!

Example of a non-zero $f \in Z$: let's build $f : \mathbf{A}_k \to \mathbf{C}$ as a product of f_v. If v is finite let's just let f_v be the characteristic function of R_v. If v is infinite and real set $f_v(x) = e^{-\pi x^2}$ and if v is complex set $f_v(x+iy) = e^{-2\pi(x^2+y^2)}$. At the infinite places we've rigged it so $\widehat{f}_v = \widehat{f}$. At the finite places, $\widehat{f}_v$ is $p^{-m/2}$ times the characteristic function of the inverse different of f, where p^m generates the discriminant ideal of k_v, so $\widehat{f}_v = f_v$ at the unramified places but not at the ramified places.

We now have a problem in analysis: we need to check $f \in Z$. First let's check f and $\widehat{f}$ are in $L^1(\mathbf{A}_k)$. Well, locally they are integrable, and at all but finitely many places the local integral is 1, so the infinite product trivially converges and gives the global integral.

Next let's check the third condition; we need to check that $f(y).|y|^\sigma$ is in $L^1(\mathbf{A}_f^\times)$ for $\sigma > 1$, and similarly for $\widehat{f}$. Well the local factors are certainly in L^1—indeed, they are in L^1 for $\sigma > 0$, because we checked this when we were doing our local zeta integrals. But this isn't enough to check that the product is L^1: we need to check that the infinite product of the local integrals converges. We evaluated the local integrals at the finite places, when doing our local calculations, and they were $(1 - p^{-\sigma})^{-1}$ for $k = \mathbf{Q}$ (I did these in class) and more generally $p^{-m/2}(1 - q^{-\sigma})^{-1}$ if k is a finite extension of $\mathbf{Q}$ and we're doing the computation at a P-adic place,

with residue field of size q and discriminant ideal (p^m) (I mentioned these on the example sheet; the proof is no more difficult). So we need to check that $\prod_P (1 - N(P)^{-\sigma})^{-1}$ converges for $\sigma > 1$—and it does; this is precisely the statement that the zeta function of a number field converges for $\mathrm{Re}(s) > 1$, which is proved by reducing to $k = \mathbf{Q}$ and then using standard estimates. This argument applies to both f and $\widehat{f}$, which are the same away from a finite set of places.

Finally we have to check the second condition (the one that let us apply Poisson summation and interchange a sum and an integral). Let y be a fixed idele, let x be a fixed adele, and let's first consider

$$\sum_{\alpha \in k} f(y(x + \alpha)).$$

First I claim that this sum converges absolutely. Because look at the support of f: at the finite places it's supported only on "integral ideles" $\mathbf{A}_k^f \cap \prod_{v<\infty} R_v$, so,

whatever y and x are, $f(y(x+\alpha))$ will actually equal zero if, at any place, the denominator of αy beats the denominator of xy. So this sum, ostensibly over all of $k^\times$, is really only over a fractional ideal in k, and now convergence is trivial because at the infinite places (and there is at least one infinite place) f is exponentially decreasing, and there are only finitely many lattice points with norm at most a given constant.

Now why is the convergence locally uniform? It's for the same reason. If y and x vary in a compact then the fractional ideal above might move but for compactness reasons the lattice won't get arbitrarily small (it's not difficult to write down a formal proof) and it's hence easy to uniformly bound the sums involved.

So the main theorem applies! What does it say in this case?

Well, $\zeta(f,|.|^s)$ and $\zeta(\widehat{f},|.|^{1-s})$ are closely related to, but not quite, the zeta function of k. Indeed if we write S_∞ for the infinite places of k and S_f for the finite places which are ramified in $k/\mathbf{Q}$ then $\zeta(f,|.|^s)=\prod_v \zeta(f_v,|.|^s)$ (the right hand integrals are local zeta integrals), which expands to

$$\prod_{v\in S_\infty}\zeta(f_v,|.|^s)\prod_{v\in S_f}(Nv)^{-m_v/2}\prod_P(1-N(P)^{-s})^{-1}$$

and $\prod_{v|\infty}\zeta_v(f_v,|.|^s)$ is a load of gamma factors—exactly the fudge factors which you multiply $\zeta_k(s)=\prod_P(1-N(P)^{-s})^{-1}$ by to get (definition) $\xi_k(s)$. So $\zeta(f,|.|^s)=\xi_k(s)|d|^{-1/2}$ with d the discriminant of k. Now $\zeta(\widehat{f},|.|^{1-s})$ is almost the same, except that $\widehat{f}\neq f$ at the finite ramified places: the local integral of f_v at the finite place is easily checked to be $p^{-ms}/(1-q^{s-1})$, so for $\mathrm{Re}(1-s)>1$ we have $\zeta(\widehat{f},|.|^{1-s})=\xi_k(1-s)|d|^{-s}$ and we deduce

$$\xi_k(1-s)=|d|^{s-1/2}\xi_k(s).$$

Slightly better: if we set

$$Z_k(s) = \xi_k(s).|d|^{s/2} = \zeta_k(s). \prod_{v|\infty} \zeta(f_v, s).|d|^{s/2}$$

then we get

$$Z_k(1-s) = Z_k(s).$$

This is the functional equation for the "Dedekind zeta function", that is, the zeta function of a number field.

Moreover, we know that the pole at $s = 1$ of $\zeta(f, |.|^s)$ is simple with residue $\widehat{f}(0)\mu^*(E) = \widehat{f}(0)2^r(2\pi)^s \operatorname{Reg}_k .h/(w\sqrt{|d|})$, and $\widehat{f}(0) = |d|^{-1/2}$, so the pole at $s = 1$ of $\xi_k(s) = \zeta(f, |.|^s)|d|^{1/2}$ has residue $2^r(2\pi)^s \operatorname{Reg}_k .h/(w\sqrt{|d|})$. Moreover the local zeta factors at the real infinite places are $\pi^{-s/2}\Gamma(s/2)$ which equals 1 at $s = 1$, and at the complex infinite places are $(2\pi)^{1-s}\Gamma(s)$ which is again 1 at $s = 1$, so we deduce

$$\lim_{s \to 1}(s-1)\zeta_k(s) = 2^r(2\pi)^s \operatorname{Reg}_k .h/(w\sqrt{|d|})$$

which is called the analytic class number formula and which is used crucially in both analytic arguments about densities of primes and in algebraic arguments in Iwasawa theory.

Remark. Iwasawa noted that applying the theory to the function above, *without* assuming the classical facts about class groups and unit groups that we needed when analysing $J/k^\times$, in fact showed that $\int_{J/k^\times} 1 < \infty$, and hence that one could *deduce* the finiteness of the class group and finite-generation of the unit group of a number field via this calculation.

Let's do one more example, if we have time: Dirichlet L-functions.

Let $N \geq 1$ be an integer, and $\chi : (\mathbf{Z}/N\mathbf{Z})^\times \to \mathbf{C}^\times$ be a character. By CRT we can write $\chi = \prod_{p|N} \chi_p$ with $\chi_p : (\mathbf{Z}/p^e\mathbf{Z})^\times \to \mathbf{C}^\times$, where $p^e||N$. Our calculations for a fundamental domain of $k^\times$ in J, when applied to $k = \mathbf{Q}$, show that $\mathbf{A}_\mathbf{Q}^\times = \mathbf{Q}^\times \times \prod_p \mathbf{Z}_p^\times \times \mathbf{R}_{>0}$, with the first factor embedded diagonally. Hence χ naturally gives rise to a character of $\prod_p \mathbf{Z}_p^\times$ (use χ_p if $p|N$ and 1 otherwise) and hence to a character $c : \mathbf{Q}^\times \backslash \mathbf{A}_\mathbf{Q}^\times \to \mathbf{C}^\times$ (make c trivial on $\mathbf{R}_{>0}$). We write $c = \prod_v c_v$. If $p \nmid N$ then $c_p : \mathbf{Q}_p^\times \to \mathbf{C}^\times$ is trivial on $\mathbf{Z}_p^\times$ and $c_p(p) = \chi(p)^{-1}$.

Let's now choose f so that $\zeta(f, c|.|^s)$ is not identically zero and let's see what the resulting function of s is. If $p \nmid N$ then we just let f_p be the characteristic function of $\mathbf{Z}_p$. If $p|N$ then we let f_p be the function we used on the example sheet when computing ρ on the component corresponding to χ_p. Note that we don't care what f_p is!

At infinity we let $f_\infty(x) = e^{-\pi x^2}$ if $\chi(-1) = 1$ and $f_\infty(x) = xe^{-\pi x^2}$ if $\chi(-1) = -1$; these are the function we used in our local calculations in the $\mathbf{R}$ case.

We set $f = \prod_v f_v$. The same arguments as above show $f \in Z$. We have

$$\zeta(f, c.|.|^s) = \zeta(f_\infty, c_\infty.|.|^s) \prod_{p|N} \zeta(f_p, c_p.|.|^s).L(\chi^{-1}, s)$$

because if $p \nmid N$ then one easily computes $\zeta(f_p, c_p.|.|^s) = (1 - \chi(p)^{-1}p^{-s})^{-1}$.

Similarly

$$\zeta(\hat{f}, \widehat{c.|.|^s}) = \zeta(\hat{f}_\infty, \widehat{c_\infty.|.|^s}) \prod_{p|N} \zeta(\hat{f}_p, \widehat{c_p.|.|^s}) L(\chi, 1-s)$$

and the trick here is *not* to attempt to work out $\zeta(f_p, c_p, |.|^s)$ or $\zeta(\hat{f}_p, \widehat{c_p.|.|^s})$ but to remember that these local zeta integrals both converge for $0 < \mathrm{Re}(s) < 1$ and that we worked out their ratio $\rho(c_p.|.|^s)$ when doing the local calculations!

The ratio was just $p^{e(s-1)} \sum_{j=1}^{p^e-1} \chi(j)\zeta_{p^e}^j$, the Gauss sum. [Note in passing that in particular we never used the local meromorphic continuation results to prove the global ones, we merely use the local ones to see the explicit form of the functional equation.] If $\xi(\chi, s)$ denotes $L(\chi, s)$ times the factor at infinity, we deduce

$$\xi(\chi^{-1}, s)N^s W = \xi(\chi, 1-s)$$

where W is an explicit algebraic number that depends only on χ and N and is basically a sum of roots of unity.

Finally I'll remark that there are more general quasi-characters $k^\times \backslash \mathbf{A}_k^\times \to \mathbf{C}^\times$ than those above. The general such thing is usually called a Hecke character or a Grössencharacter. If ψ is such a gadget, then ψ is unramified at all but finitely many finite places, and defining f_v at these unramified places to be just the characteristic function of R_v, and f_v at the other places to be the f_v we used when analysing the local ψ_v, the equation $\zeta(f, \psi.|.|^s) = \zeta(\widehat{f}, \widehat{\psi.|.|^s})$ unravels to become the meromorphic continuation and functional equation for the L-function of the Grössencharacter that Hecke discovered in his original tour de force!

THE END