Fermat’s Last Theorem Blueprint

Contents

Contents i

1 Introduction 1
1.1 Which proof is being formalised? . . . 2
1.2 The structure of this blueprint . . . 2
1.3 Remarks . . . 2

2 First Reductions Of The Problem 4
2.1 Goal . . . 4
2.2 Overview . . . 4
2.3 Reduction To n >= 5 And Prime . . . 4
2.4 Frey Packages . . . 5
2.5 Galois Representations And Elliptic Curves . . . 6
2.6 The Frey Curve . . . 7
2.7 Reduction To Two Big Theorems . . . 7

3 Reducibility Of p-Torsion Of The Frey Curve 9
3.1 Overview . . . 9
3.2 Hardly Ramified Representations . . . 9

4 An overview of the proof 15
4.1 Potential modularity. . . . 15
4.2 A modularity lifting theorem . . . 16
4.3 Compatible families, and reduction at 3 . . . 17

5 An Example Of An Automorphic Form 18
5.1 Introduction . . . 18
5.2 A quaternion algebra . . . 19
5.3 Zhat . . . 19
5.4 More Advanced Remarks On Zhat Versus Q . . . 22
5.5 Qhat and tensor products. . . . 22
5.6 A crash course in tensor products . . . 22
5.7 Additive structure of Qhat. . . . 25
5.8 Multiplicative structure of the units of Qhat. . . . 26
5.9 The Hurwitz quaternions . . . 27

5.10 Profinite completion of the Hurwitz quaternions 30

6 Stating the modularity lifting theorems 32
6.1 Automorphic forms and analysis 32
6.2 Central simple algebras 33

7 Automorphic forms and the Langlands Conjectures 35
7.1 Definition of an automorphic form for GLn over Q 35
7.2 The Finite Adeles Of The Rationals 35
7.3 The Group GLn Of The Adeles 35
7.4 Smooth Functions 35
7.5 Slowly-Increasing Functions 36
7.6 Weights At Infinity 36
7.7 The Action Of The Universal Enveloping Algebra 36
7.8 Automorphic Forms 37
7.9 Hecke Operators 38

8 Miniproject: Frobenius Elements 39
8.1 Status 39
8.2 Introduction And Goal 39
8.3 Statement Of The Theorem 39
8.4 Examples 40
8.5 The Extension B/A 40
8.6 The Extension (B/Q)/(A/P) 41
8.7 The Extension L/K 42

9 Miniproject: Adeles 45
9.1 Status 45
9.2 Goals 45
9.3 Local compactness 46
9.4 Base change 47
9.5 Discreteness and compactness 59

10 Miniproject: Haar Characters 61
10.1 Goal 61
10.2 Initial definitions 61
10.3 Examples 63
10.4 Algebras 65
10.5 Left and right multiplication 66
10.6 Finite products 67
10.7 Some measure-theoretic preliminaries 68
10.8 Restricted products 69
10.9 Adeles 71

11 Miniproject: Fujisaki's Lemma 75
11.1 Initial definitions 75
11.2 Enter the adeles 76

11.3 The proof 76

12 Miniproject: Quaternion algebras 80
12.1 The goal 80
12.2 Initial definitions 80
12.3 Brief introduction to automorphic forms in this setting 81
12.4 Definition of spaces of automorphic forms 81
12.5 The main result 83

13 Miniproject: Hecke Operators 84
13.1 Status 84
13.2 Goal 84
13.3 Abstract theory 84
13.4 Restricted products 86
13.5 Some local theory 89
13.6 Adelic groups 91
13.7 Automorphic forms 92
13.8 Concrete Hecke operators 93
13.9 Analysis of the Hecke algebra 94

14 Appendix: A collection of results which are needed in the proof. 95
14.1 Results from class field theory 95
14.2 Structures on the points of an affine variety. 96
14.3 Algebraic groups. 98
14.4 Automorphic forms and representations 98
14.5 Galois representations 99
14.6 Algebraic geometry 99
14.7 Algebra 100

Dependency Graph 102

Blueprint Summary 103

Blueprint Bibliography 104

This repository is the Verso blueprint integration layer for the FLT project. It keeps the blueprint code at the repository root while treating `FLT/` as the upstream formalization checkout.

1 Introduction

Fermat's Last Theorem is the statement that if a, b, c, n are positive whole numbers with $n \geq 3$, then $a^n + b^n \neq c^n$. It is thus the claim that a family of Diophantine equations, $a^3 + b^3 = c^3$, $a^4 + b^4 = c^4$, and so on, has no positive integer solutions. Diophantus was a Greek mathematician who lived around `1800` years ago, and he would have been able to understand the statement of the theorem (he knew about positive integers, addition and multiplication).

Fermat's Last Theorem was explicitly raised by Fermat in `1637`, and was proved by Wiles, with the proof completed in joint work with Taylor, in `1994`. There are now several proofs, but all of them go broadly in the same direction, using elliptic curves and modular forms.

Lean is an interactive theorem prover; it checks mathematical arguments with super-human accuracy. Explaining a proof of Fermat's Last Theorem to Lean is in some sense like explaining the proof to Diophantus; for example, the proof starts by observing that before we go any further it is convenient to first invent or discover zero and negative numbers, and one can point explicitly at places in Lean's source code where these things happen. However, we will adopt a more efficient approach: we will assume all of the theorems both in core Lean and in its mathematics library `mathlib`, and proceed from there.

To give some idea of what this entails: `mathlib` at the time of writing contains most of an undergraduate mathematics degree and parts of several relevant Masters level courses (for example, the definitions and basic properties of elliptic curves and modular forms are in `mathlib`). Thus our task can be likened to teaching a graduate level course on Fermat's Last Theorem to a computer. The computer is quite a challenging audience member – it will insist on being given all technical details of all arguments, and it will not accept proof by intimidation or by appeal to higher authority. Most mathematicians know humans who also behave in this manner. Up until `2026`, the computer would only start filling in details of arguments by itself once the arguments are mathematically utterly obvious. This made formalization of mathematics a very time-consuming process. However, with the advent of AI computers are now getting better at filling in arguments themselves.

1.1 Which proof is being formalised?

At the time of writing, these notes do not contain anywhere near a proof of FLT, or even a sketch proof, although we are currently actively working on fixing this.

From 2024 to 2029 we will be beginning to build a proof of FLT, following a strategy basically due to Khare, and modified by Taylor, taking into account Buzzard's comments on what would be easy or hard to do in Lean. Our strategy uses refinements of the original Taylor–Wiles method by Diamond/Fujiwara, Khare–Wintenberger, Skinner–Wiles, Kisin, Taylor and others. One could call it a 21st century proof of the theorem. During this initial phase of the project, we shall also be assuming many nontrivial theorems without proof, as long as they were published by 31 December 1989. To get technical for just a second, we shall for example be assuming the existence of Galois representations attached to weight 2 Hilbert modular forms, Langlands' cyclic base change theorem for GL_2, Mazur's theorem bounding the torsion subgroup of an elliptic curve over the rationals, and several other nontrivial results which were known by the end of the 1980s.

The upshot of this is that, by 2029 at the end of this first phase, the project should contain a complete proof that FLT follows from results that were known to humanity in the 1980s. In particular, one naive way of understanding the goal is that it is a formalization of the papers of Wiles and Taylor–Wiles, assuming the results in the references of those papers. However, as noted above, we will actually be taking a slightly different path.

1.2 The structure of this blueprint

This blueprint is a nonlinear document, comprising many chapters. The chapters are not designed to be read in order. Each chapter is self-contained and has a well-defined goal, typically stated at the top of the chapter.

After this chapter, you should next read the first reductions chapter (p. 4), which explains how to reduce FLT to two highly nontrivial statements about the p-torsion in a certain elliptic curve (the Frey curve). One of these statements was proved in the 1970s by Mazur, and we shall not be concentrating on it until after the first phase is complete. The other is a theorem of Wiles, and this is what we will be concentrating on in the remainder of the blueprint.

1.3 Remarks

The actual blueprint currently also contains a lot of disorganised ideas. Currently these should be disregarded.

The overview chapter (p. 15) is an extremely sketchy overview of how the rest of the proof goes. This is currently being expanded and should be ignored right now.

All of the remaining chapters are experiments, and most of them are what I am currently calling mini-projects. A mini-project is a bottom-up project, typically at early graduate student level, with a concrete goal. The ultimate goal of many of these projects is to actually get some result into mathlib. We have had one success so far: the Frobenius mini-project is currently being PRed to mathlib by Thomas Browning. Several other mini-projects have been completed in this repo but have not yet been PRed to mathlib. Examples of completed miniproject which are not yet upstreamed include adeles mini-project and the quaternion algebra mini-project. These projects do not logically depend on each other for the most part, and one can pick and choose how one reads them.

There is also an appendix, which is again very sketchy, and comprises mostly of a big list of nontrivial theorems many of which we will be assuming without proof in the FLT project.

The next chapter to read, where the proof begins, is the first reductions chapter (p. 4).

2 First Reductions Of The Problem

2.1 Goal

The goal of this chapter is to reduce FLT to a deep theorem of Mazur and a deep theorem of Wiles about a Galois representation.

2.2 Overview

The proof of Fermat's Last Theorem is by contradiction. We assume that we have a counterexample $a^n + b^n = c^n$, and manipulate it until it satisfies the axioms of a Frey package, a basic concept which we will explain below. From the Frey package we build a Frey curve, an elliptic curve defined over $\mathbb{Q}$. We then look at a certain representation of a Galois group coming from this elliptic curve, and finally, using two very deep and independent theorems, one due to Mazur and the other due to Wiles, we show that this representation is both reducible and irreducible, the contradiction we seek.

2.3 Reduction To n >= 5 And Prime

Lemma 2.3.1

> If there is a counterexample to Fermat's Last Theorem, then there is a counterexample $a^p + b^p = c^p$ with p an odd prime.

Proof for Lemma 2.3.1

> This proof is already in mathlib, but the TeX blueprint spells it out for completeness.Say $a^n + b^n = c^n$ is a counterexample to Fermat's Last Theorem. Every positive integer is either a power of 2 or has an odd prime factor. If $n = kp$ has an odd prime factor p, then $(a^k)^p + (b^k)^p = (c^k)^p$, which is the counterexample we seek.It remains to deal with the case where n is a power of 2, so let us assume this. We have $3 \leq n$ by assumption, so $n = 4k$ is a multiple of 4, and thus $(a^k)^4 + (b^k)^4 = (c^k)^4$, giving a counterexample to Fermat's Last Theorem for $n = 4$. But Fermat's theorem for exponent 4, already in mathlib, says that this is impossible.

Euler proved Fermat's Last Theorem for $p = 3$.

Lemma 2.3.2

There are no solutions in positive integers to $a^3 + b^3 = c^3$.

Proof for Lemma 2.3.2

The proof in mathlib was formalized by a team from the "Lean For the Curious Mathematician" conference held in Luminy in March 2024 (its dependency graph can be visualised here[1]).

Corollary 2.3.3

If there is a counterexample to Fermat's Last Theorem, then there is a counterexample $a^p + b^p = c^p$ with p prime and $p \geq 5$.

Proof for Corollary 2.3.3

Follows from the previous two lemmas.

2.4 Frey Packages

For convenience we make the following definition.

Definition 2.4.1

A Frey package (a, b, c, p) is three nonzero pairwise-coprime integers a, b, and c, with $a \equiv 3 \pmod 4$ and $b \equiv 0 \pmod 2$, together with a prime $p \geq 5$, such that $a^p + b^p = c^p$.

Our next reduction is as follows.

Lemma 2.4.2

If Fermat's Last Theorem is false for p prime and $p \geq 5$, then there exists a Frey package.

Proof for Lemma 2.4.2

Suppose we have a counterexample $a^p + b^p = c^p$ for the given p; we now build a Frey package from this data.If the greatest common divisor of a, b, c is d, then $a^p + b^p = c^p$ implies $(a/d)^p + (b/d)^p = (c/d)^p$. Dividing through, we can thus assume that no prime divides all of a, b, c. Under this assumption we must have that a, b, c are pairwise coprime, as if some prime divides two of the integers a, b, c then by $a^p + b^p = c^p$ and unique factorization it must divide all three of them. In particular we may assume that not all of a, b, c

[1] https://pitmonticone.github.io/FLT3/blueprint/dep_graph_document.html

are even, and now reducing modulo 2 shows that precisely one of them must be even.Next we show that we can find a counterexample with b even. If a is the even one, then we can just switch a and b. If c is the even one, then we can replace c by -b and b by -c, using that p is odd.The last thing to ensure is that a is 3 mod 4. Because b is even, we know that a is odd, so it is either 1 or 3 mod 4. If a is 3 mod 4 then we are home; if however a is 1 mod 4, we replace a, b, c by their negatives, and this is the Frey package we seek.

2.5 Galois Representations And Elliptic Curves

To continue, we need some of the theory of elliptic curves over $\mathbb{Q}$. So let f(X) denote any monic cubic polynomial with rational coefficients and whose three complex roots are distinct, and let us consider the equation $E : Y^2 = f(X)$, which defines a curve in the (X,Y)-plane. This curve, or strictly speaking its projectivization, is a so-called elliptic curve over $\mathbb{Q}$.

If $E : Y^2 = f(X)$ is an elliptic curve over $\mathbb{Q}$, and if K is any characteristic zero field, and hence a $\mathbb{Q}$-algebra, then we write E(K) for the set of solutions to $y^2 = f(x)$ with $x, y \in K$, together with an additional "point at infinity" corresponding morally to $x = y = \infty$. It is an extraordinary fact, and not at all obvious, that E(K) naturally has the structure of an additive abelian group, with the point at infinity being the zero element, the identity. Fortunately this fact is already in mathlib. This additive group structure has the property that three distinct points P, Q, R in K^2 which are in E(K) will sum to zero if and only if they are collinear.

The group structure behaves well under change of field: if E is an elliptic curve over $\mathbb{Q}$ and if $K \to L$ is a homomorphism of characteristic zero fields, then the induced map $E(K) \to E(L)$ is a group homomorphism. Thus if $f : K \to L$ is an isomorphism of characteristic zero fields, the induced map $E(K) \to E(L)$ is an isomorphism of groups, with the inverse isomorphism being the map $E(L) \to E(K)$ induced by f^{-1}. This construction thus gives us an action of the multiplicative group Aut(K) of automorphisms of the field K on the additive abelian group $E(K)$, and hence also on the n-torsion of this group for any positive integer n. In particular, if $\overline{\mathbb{Q}}$ denotes an algebraic closure of the rationals, for example the algebraic numbers in $\mathbb{C}$, and if Gal$(\overline{\mathbb{Q}}/\mathbb{Q})$ denotes the group of field isomorphisms $\overline{\mathbb{Q}} \to \overline{\mathbb{Q}}$, then for any elliptic curve E over $\mathbb{Q}$ we have an action of Gal$(\overline{\mathbb{Q}}/\mathbb{Q})$ on the additive abelian group $E(\overline{\mathbb{Q}})$, and hence on its n-torsion subgroup $E(\overline{\mathbb{Q}})[n]$.

If furthermore $n = p$ is prime, then $E(\overline{\mathbb{Q}})[p]$ is naturally a vector space over $\mathbb{Z}/p\mathbb{Z}$, and therefore it inherits the structure of a mod p representation of Gal$(\overline{\mathbb{Q}}/\mathbb{Q})$. This is the mod p Galois representation attached to the elliptic curve E, and it is well-known to be 2-dimensional. We call it $\rho_{E,p}$.

In the next section we apply this theory to the elliptic curve coming from a counterexample to Fermat's Last Theorem.

2.6 The Frey Curve

Recall that a Frey package (a, b, c, p) is a prime $p \geq 5$ and nonzero pairwise-coprime integers a, b, and c satisfying $a^p + b^p = c^p$ and the congruences $a \equiv 3 \pmod 4$ and $b \equiv 0 \pmod 2$. We have shown above that if Fermat's Last Theorem is false, then a Frey package exists.

Definition 2.6.1

> Given a Frey package (a, b, c, p), the corresponding Frey curve is the elliptic curve over $\mathbb{Q}$ (considered by Frey and, before him, Hellegouarch) defined by the equation $Y^2 = X(X - a^p)(X + b^p)$.

Note that the roots of the cubic $X(X - a^p)(X + b^p)$ are distinct because a, b, and c are nonzero and $a^p + b^p = c^p$.

Given a Frey package (a, b, c, p) with corresponding Frey curve E, the mod p Galois representation $\rho_{E,p}$ attached to this package is the 2-dimensional representation of $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ on $E(\overline{\mathbb{Q}})[p]$ described above. Frey's observation is that this representation has some very surprising properties. We will make this remark more explicit in the next chapter. Here we shall show how these properties can be used to finish the job.

2.7 Reduction To Two Big Theorems

Recall that a representation of a group G on a vector space W is said to be irreducible if there are precisely two G-stable subspaces of W, namely 0 and W. The representation is said to be reducible otherwise.

Now say (a, b, c, p) is a Frey package. Consider the mod p representation of $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ coming from the p-torsion in the Frey curve $Y^2 = X(x-a^p)(X+b^p)$ associated to the package. Let's call this representation ρ, and we say that ρ is the mod p representation associated to the Frey package (a, b, c, p). Is ρ irreducible or not?

Theorem 2.7.1

> If ρ is the mod p Galois representation associated to a Frey package (a,b,c,p) then ρ is irreducible.

Proof for Theorem 2.7.1

> This follows from a profound and long result of Mazur (Mazur, 1977) from 1977, namely the fact that the torsion subgroup of an elliptic curve over $\mathbb{Q}$ can have size at most 16. In fact there is still a little more work which needs to be done to deduce the theorem from Mazur's result. A pre-1990 reference for the full proof of this claim is Proposition 6 in Section 4.1 of (Serre, 1987).

Note that in the first (pre-2029) phase of the FLT project, we will not be working on a formalization of this result, as it was known in the 1980s. We will however be thinking a lot about the next result, which says the exact opposite.

Theorem 2.7.2

If ρ is the mod `p` Galois representation associated to a Frey package `(a,b,c,p)` then ρ is reducible.

Proof for Theorem 2.7.2

This is the main content of Wiles' magnum opus. We omit the argument for now, although later on in this project we will have a lot to say about its proof.

Corollary 2.7.3

There is no Frey package.

Proof for Corollary 2.7.3

Follows immediately from the previous two theorems Theorem 2.7.1 and Theorem 2.7.2.

We deduce.

Corollary 2.7.4

Fermat's Last Theorem is true. In other words, there are no positive integers `a`, `b`, and `c` and natural numbers `n >= 3` such that $a^n + b^n = c^n$.

Proof for Corollary 2.7.4

Assume there is a there is a counterexample $a^n + b^n = c^n$. By Corollary Corollary 2.3.3, we may assume that there is also a counterexample $a^p + b^p = c^p$ with $p \geq 5$ and prime. Then Lemma 2.4.2 produces a Frey package, contradicting Corollary 2.7.3.

Because we are, for now at least, assuming Mazur's theorem, we now turn our attention to a proof of theorem Theorem 2.7.2. We start on this proof in the Frey-curve reducibility chapter (p. 9).

3 Reducibility Of p-Torsion Of The Frey Curve

3.1 Overview

In the first reductions chapter (p. 4) we reduced FLT, modulo a hard theorem from the 1970s, to theorem Theorem 2.7.2, the assertion that p-torsion in the Frey curve is reducible. In this chapter we deduce this assertion from three more complex claims about hardly ramified Galois representations. It is relatively straightforward to reduce one of these three claims to a result of Fontaine proved in the 1980s in his paper on the nonexistence of nontrivial abelian schemes over **Z**. The other two claims lie deeper, and their proofs use techniques initially developed by Wiles in the 1990s.

3.2 Hardly Ramified Representations

Let (a, b, c, p) be a Frey package, so in particular $p \geq 5$ is prime and $a^p + b^p = c^p$, let E be the corresponding Frey curve over $\mathbb{Q}$, and let $\rho : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{Aut}(E(\overline{\mathbb{Q}})[p])$ be the 2-dimensional Galois representation on the p-torsion of E. Recall that our goal is to prove that ρ is reducible.

What we need to leverage is the fact that ρ has very little ramification. To give a toy example before we start: if K is a number field, that is, a finite extension of $\mathbb{Q}$, and if the extension $K/\mathbb{Q}$ is unramified at all primes, then an old theorem of Minkowski tells us that $K = \mathbb{Q}$. We want to prove a theorem in a similar vein, namely that if a 2-dimensional mod p Galois representation is hardly ramified, then it is reducible. Below, we give a precise definition of what it means for a continuous 2-dimensional representation $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{GL}_2(R)$ to be hardly ramified. Before we do that, we need to say precisely which topological rings R we will allow. We say that a topological ring is pro-Artinian if it is a projective limit of Artin local rings each equipped with the discrete topology, and if it has the projective limit topology. We are only concerned with local pro-Artinian rings with finite residue field; such things can be checked to be the same thing as topological local rings with finite residue field whose underlying topological space is profinite, and such

that additive translates of open ideals form a basis for the topology. Let us call such rings coefficient rings for now.

We make some remarks to orient the reader.

- Any complete local Noetherian ring with finite residue field is a coefficient ring, if the ring is equipped with the $\mathfrak{m}$-adic topology where $\mathfrak{m}$ is the maximal ideal. In this case, all powers of $\mathfrak{m}$ are open.
- In particular finite fields, and integer rings of finite extensions of $\mathbb{Q}_p$, are coefficient rings.
- If R is a coefficient ring then R is isomorphic to the projective limit of the finite rings R/I as I runs over the open ideals of R.
- A non-Noetherian example of a coefficient ring is the projective limit over n of the rings $\mathbb{Z}/p\mathbb{Z}[\varepsilon_1, \ldots, \varepsilon_n]/(\forall i, j, \varepsilon_i\varepsilon_j = 0)$; these rings are convenient to include as coefficient rings for technical reasons; they make representability theorems easier.
- The category of coefficient rings is equivalent to the pro-category of the category of finite local rings.
- A coefficient ring is pseudocompact in the sense of Grothendieck. A pseudocompact local ring is however a more general concept as such a thing may have an infinite residue field and would thus not be profinite.
- If R is a coefficient ring with residue field of characteristic ℓ, then there is a unique continuous map $\mathbb{Z}_\ell \to R$. Indeed, it suffices to prove that there is a unique continuous map $\mathbb{Z}_\ell \to R/I$ for each open ideal I, but R/I is a finite local ring with residue field of characteristic ℓ. R/I is hence Artinian, so some power of the maximal ideal is zero by Nakayama. This means that $\ell^N = 0$ for some sufficiently large N, and hence R/I is a $\mathbb{Z}/\ell^N\mathbb{Z}$-algebra and thus admits admits a unique map from $\mathbb{Z}_\ell$.
- It will be more convenient to fix once and for all the integer $\mathcal{O}$ in a finite extension of $\mathbb{Q}_\ell$ and consider coefficient $\mathcal{O}$-algebras, namely coefficient rings R equipped with a continuous map $\mathcal{O} \to R$ which is a local homomorphism inducing an isomorphism on residue fields.

Because a coefficient ring R with residue field of characteristic ℓ is naturally a $\mathbf{Z}_\ell$-algebra, we can talk about the ℓ-adic cyclotomic character $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to R^\times$. We are now ready to define hardly ramified representations.

Definition 3.2.1

Let R be a coefficient ring with finite residue field of characteristic $\ell \geq 3$. Let V be a finite free R-module of rank 2, equipped with the product topology. A continuous representation $\rho : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{GL}_R(V)$ is said to be hardly ramified if it satisfies the following four conditions:

- $\det(\rho) : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to R^\times$ is the cyclotomic character.
- ρ is unramified outside 2ℓ.
- The restriction of ρ to $\mathrm{Gal}(\overline{\mathbb{Q}}_2/\mathbb{Q}_2)$ is reducible. More precisely, there is a short exact sequence $0 \to R \to V \to R \to 0$ which is stable under the action of $\mathrm{Gal}(\overline{\mathbb{Q}}_2/\mathbb{Q}_2)$, and the Galois action on the 1-dimensional quotient is an unramified representation of $\mathrm{Gal}(\overline{\mathbb{Q}}_2/\mathbb{Q}_2)$ whose square is trivial.
- The restriction of ρ to $\mathrm{Gal}(\overline{\mathbb{Q}}_\ell/\mathbb{Q}_\ell)$ is flat, by which we mean that for all open ideals I of R, the finite-image representation $\rho \bmod I : \mathrm{Gal}(\overline{\mathbb{Q}}_\ell/\mathbb{Q}_\ell) \to \mathrm{GL}_{R/I}(V/I)$ comes from a finite flat group scheme.

A well-known result, which basically goes back to Frey, is the following:

Theorem 3.2.2

The ℓ-torsion $\rho : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$ in the Frey curve associated to a Frey package (a, b, c, ℓ) is hardly ramified.

Proof for Theorem 3.2.2

This was well-known in the 1980s. A proof sketch is as follows. First note that $\ell \geq 5 > 3$ by definition of a Frey package. Let ρ denote the Galois representation on the ℓ-torsion of the Frey curve. The fact that ρ is 2-dimensional is Corollary III.6.4(b) of (Silverman, 2009), and the fact that its determinant is cyclotomic is Proposition III.8.3 of the same reference. These results hold for elliptic curves in general. The remaining claims are specific to the Frey curve and lie deeper. The fact that ρ is unramified outside 2ℓ is a consequence of `(4.1.12)` and `(4.1.13)` of (Serre, 1987). The fact that ρ at 2 has an unramified 1-dimensional quotient of order at most 2 follows from the fact that the Frey curve is semistable at 2 and the theory of the Tate curve. Finally, the claim that ρ is flat at ℓ is Proposition 5 and `(4.1.13)` of (Serre, 1987).

Note that irreducibility and absolute irreducibility for hardly ramified mod ℓ representations are the same, because our assumptions that $\ell \geq 3$ and that the determinant is cyclotomic imply that the image of complex conjugation has distinct eigenvalues defined over the ground field.

Theorem 3.2.3

If $\ell \geq 3$ is prime and $\rho : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{GL}_2(\mathbf{Z}/\ell\mathbf{Z})$ is hardly ramified, then ρ is reducible.

Note that this (deep) claim is a consequence of Serre's conjecture (Serre, 1987), now a theorem of Khare and Wintenberger (Khare and Wintenberger,

2009), and indeed we shall use methods introduced by Khare and Wintenberger to prove this special case of Serre's conjecture. Given this result, we can deduce theorem Theorem 2.7.2 (which we restate here) easily:

Theorem 3.2.4

> If $\overline{\rho}$ is the mod p Galois representation associated to a Frey package `(a,b,c,p)` then $\overline{\rho}$ is reducible.

Proof for Theorem 3.2.4

> Indeed, ρ is hardly ramified by theorem Theorem 3.2.2 and thus reducible by theorem Theorem 3.2.3.

Our job of reducing FLT to theorems of the 1980s is hence reduced to proving theorem Theorem 3.2.3.

Hardly Ramified Mod p Representations Are Reducible

In this section we will state three theorems, from which theorem Theorem 3.2.3 easily follows.

Firstly, we claim that an irreducible hardly ramified mod ℓ representation lifts to an ℓ-adic representation.

Theorem 3.2.1.1

> If $\ell \geq 3$ is prime and $\overline{\rho} : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{GL}_2(\mathbf{Z}/\ell\mathbf{Z})$ is hardly ramified and irreducible, then there exists a finite extension $K/\mathbb{Q}_\ell$ with integer ring $\mathcal{O}$ and maximal ideal $\mathfrak{m}$ and a hardly ramified representation $\rho : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{GL}_2(\mathcal{O})$ whose reduction modulo $\mathfrak{m}$ is isomorphic to $\overline{\rho}$.

Proof for Theorem 3.2.1.1

> Omitted for now. **TODO**

Next we claim that a hardly ramified ℓ-adic representation spreads out to a compatible family of hardly ramified q-adic representations for all odd primes q. Note that we have not made a definition of a hardly ramified 2-adic representation.

Theorem 3.2.1.2

> If $\ell \geq 3$ is prime, K is a finite extension of $\mathbb{Q}_\ell$ with integers $\mathcal{O}$ and if $\rho : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{GL}_2(\mathcal{O})$ is a hardly ramified representation whose reduction is irreducible, then there exists a number field M and, for each finite place μ of M of characteristic prime to 2, with completion M_μ having integer ring R_μ, a hardly ramified semisimple representation $\rho_\mu : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{GL}_2(R_\mu)$ (by which we mean the generic fibre is semisimple), with the following properties:

- There is some $\lambda \mid \ell$ of M such that $\rho_\lambda \cong \rho$, the isomorphism happening over some appropriate local field containing a copy of M_λ and a copy of K
- If μ_1 and μ_2 are two finite places of M with odd residue characteristics m_1 and m_2, and if $p \nmid 2m_1m_2$ is prime, then ρ_{μ_1} and ρ_{μ_2} are both unramified at p and the characteristic polynomials $\rho_{\mu_1}(\text{Frob}_p)$ and $\rho_{\mu_2}(\text{Frob}_p)$ lie in $M[X]$ and are equal

Proof for Theorem 3.2.1.2

Omitted for now. **TODO**

In particular, we can "move" from an irreducible hardly ramified mod ℓ representation to a hardly ramified 3-adic representation, and hence to a hardly ramified mod 3 representation.

However, we can essentially completely classify the hardly ramified mod 3 Galois representations:

Theorem 3.2.1.3

If k is a finite field of characteristic 3 and $\overline{\rho} : \text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \text{GL}_2(k)$ is hardly ramified, then $\overline{\rho}$ is an extension of the cyclotomic character by the trivial representation.

Proof for Theorem 3.2.1.3

Omitted for now. **TODO**

And we can use this to essentially completely classify the hardly ramified 3-adic Galois representations:

Theorem 3.2.1.4

If $L/\mathbb{Q}_3$ is a finite extension with integer ring $\mathcal{O}_L$ and $\rho_3 : \text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \text{GL}_2(\mathcal{O}_L)$ is hardly ramified, then, viewed as a representation to $\text{GL}_2(L)$, one has $\rho_3^{ss} = 1 \oplus \chi_3$ where 1 is the trivial character and χ_3 is the 3-adic cyclotomic character.

Proof for Theorem 3.2.1.4

Omitted for now. **TODO**

Theorem Theorem 3.2.3 (if $\ell \geq 3$ is a prime and $\overline{\rho} : \text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \text{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$ is hardly ramified, then $\overline{\rho}$ is reducible) is an easy consequence of these theorems, as we now show.

Proof for Theorem 3.2.3

Assume for a contradiction that $\overline{\rho}$ is irreducible. By theorem `hardly_ramified_lifts`, $\overline{\rho}$ lifts to a hardly ramified ℓ-adic representation ρ. By theorem `hardly_ramified_spreads_out`, ρ is part of a compatible family of q-adic Galois representations. By theorem `hardly_ramified_3adic_reducible`, any 3-adic member ρ_3 of this family has semisimplification $1 \oplus \chi_3$ and in particular for $p \nmid 6$ we have that the characteristic polynomial of $\rho_3(\mathrm{Frob}_p)$ is $(X-p)(X-1)$. By compatibility of the family we deduce that for $p \nmid 6\ell$ the characteristic polynomial of $\rho(\mathrm{Frob}_p)$ is $(X-p)(X-1)$, and thus the characteristic polynomial of $\overline{\rho}(\mathrm{Frob}_p)$ is $(X-p)(X-1)$. By the Chebotarev density theorem, $\overline{\rho}$ and $1 \oplus \chi$ have the same characteristic polynomials everywhere (here χ is the mod ℓ cyclotomic character). Thus by the Brauer-Nesbitt theorem, $\overline{\rho}$ is reducible, the contradiction we seek.

What remains then, modulo several results which were known in the 1980s, is to prove the three theorems Theorem 3.2.1.1, Theorem 3.2.1.2, and Theorem 3.2.1.4. By far the easiest is theorem Theorem 3.2.1.4; this follows from old estimates of Fontaine, ultimately relying on bounds for root discriminants due to Odlyzko and Poitou, originally developed to prove that there was no non-trivial abelian scheme over **Z**. The other two theorems are deeper, and both use modern variants of Wiles' $R = T$ machinery.

We have not yet written any more LaTeX on how to proceed further; the rest of this blueprint should be considered as more unfocussed thoughts.

4 An overview of the proof

So far we have seen that, modulo Mazur's theorem (and various other things which will still take some work to formalise but which are much easier), Fermat's Last Theorem can be reduced to the statement that there is no prime $\ell \geq 5$ and hardly-ramified irreducible 2-dimensional Galois representation $\rho : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$.

In this chapter we give an overview of our strategy for proving this, and collect various results which we will need along the way. Note that we no longer need to assume that ρ comes from the ℓ-torsion in an elliptic curve.

4.1 Potential modularity.

We will only speak about modularity for 2-dimensional representations of the absolute Galois group of a totally real field F of even degree over $\mathbb{Q}$. What we will mean by "modular" is "associated to an automorphic representation of the units of the totally definite quaternion algebra over F ramified at no finite places and at all infinite places". Such a quaternion algebra only exists if F has even degree, hence the restriction in our definition of modularity. We can furthermore even demand that the infinity type of the form is trivial (the analogue in this setting of classical weight 2 forms), as these are the only forms we shall need for FLT.

Assume we have a hardly-ramified representation ρ as above. Let K be the number field corresponding to the kernel of ρ. Our first claim is that there is some totally real field F of even degree, Galois over $\mathbb{Q}$, unramified at ℓ, and disjoint from K, such that $\rho|G_F$ is modular. The proof of this is very long, and uses a host of machinery. For example:

- Moret–Bailly's result (Moret-Bailly, 1989) on points on curves with prescribed local behaviour;
- several nontrivial results in global class field theory;
- the Jacquet–Langlands correspondence;
- The assertion that irreducible 2-dimensional mod p representations induced from a character are modular (this follows from converse theorems);

- A modularity lifting theorem.

Almost everything here dates back to the 1980s or before. The exception is the modularity lifting theorem, which we now state explicitly.

4.2 A modularity lifting theorem

Suppose $\ell \geq 5$ is a prime, that F is a totally real field of even degree in which ℓ is unramified, and that S is a finite set of finite places of F not dividing ℓ. Write G_F for the absolute Galois group of F.

If $v \in S$ then let F_v denote the completion of F at v, fix an inclusion $\overline{F} \to \overline{F_v}$, let $\mathcal{O}_v$ denote the integers of F_v and $k(v)$ the residue field. Let $I_v \subset G_F$ denote the inertia subgroup at v. Local class field theory (or a more elementary approach) gives a map $I_v \to \mathcal{O}_{F_v}^\times$ and hence a map $I_v \to k(v)^\times$. Let J_v denote the kernel of this map.

Let R be a complete local Noetherian $\mathbb{Z}_\ell$-algebra with finite residue field of characteristic ℓ. We will be interested in representations $\rho : G_F \to \mathrm{GL}_2(R)$ with the following four properties.

- $\det(\rho)$ is the cyclotomic character;
- ρ is unramified outside $S \cup \{\ell\}$;
- If $v \in S$ then $\rho(g)$ has trace equal to 2 for all $g \in J_v$;
- If $v \mid \ell$ is a place of F then ρ is flat at v.

In the last bullet point, "flat" means "projective limit of representations arising from finite flat group schemes". Let us use the lousy temporary notation "S-good" to denote representations with these four properties.

Say k is a finite extension of $\mathbb{Z}/\ell\mathbb{Z}$ and $\overline{\rho} : G_F \to \mathrm{GL}_2(k)$ is continuous, absolutely irreducible when restricted to $F(\zeta_\ell)$, and S-good. One can check that the functor representing S-good lifts of $\overline{\rho}$ is representable.

Theorem 4.2.1

> If $\overline{\rho}$ is modular of level $\Gamma_1(S)$ and $\rho : G_F \to \mathrm{GL}_2(\mathcal{O})$ is an S-good lift of $\overline{\rho}$ to $\mathcal{O}$, the integers of a finite extension of $\mathbb{Q}_\ell$, then ρ is also modular of level $\Gamma_1(S)$.

Right now we are very far from even stating this theorem in Lean.

I am not entirely sure where to find a proof of this in the literature, although it has certainly been known to the experts for some time. Theorem 3.3 of (Taylor, 2006) comes close, although it assumes that ℓ is totally split in F rather than just unramified. Another near-reference is Theorem 5.2 of (Gee, 2006), although this assumes the slightly stronger assumption that the image of ρ contains $\mathrm{SL}_2(\mathbb{Z}/p\mathbb{Z})$ (however it is well-known to the experts that this can be weakened to give the result we need). One reference for the proof is Richard Taylor's 2018 Stanford course.

Proof for Theorem 4.2.1

Sketch.The proof is a two-stage procedure and has a nontrivial analytic input. First one uses the Skinner–Wiles trick to reduce to the "minimal case", and this needs cyclic base change for $\mathrm{GL}(2)$ and also a characterisation of the image of the base change construction; this seems to need a multiplicity one result, which (because of our definition of "modular") will need Jacquet–Langlands as well.In the minimal case, the argument is the usual Taylor–Wiles trick, using refinements due to Kisin and others.

Given this modularity lifting theorem, the strategy to show potential modularity of ρ is to use Moret–Bailly to find an appropriate totally real field F, an auxiliary prime p, and an auxiliary elliptic curve over F whose mod ℓ Galois representation is ρ and whose mod $\mathfrak{p}$ Galois representation is induced from a character. By converse theorems (for example) the mod $\mathfrak{p}$ Galois representation is associated to an automorphic representation of GL_2/F and hence by Jacquet–Langlands it is modular. Now we use the modularity lifting theorem to deduce the modularity of the curve over F and hence the modularity of the ℓ-torsion.

4.3 Compatible families, and reduction at 3

We now use Khare–Wintenberger to lift ρ to a potentially modular ℓ-adic Galois representation of conductor 2, and put it into an ℓ-adic family using the Brauer's theorem trick in (Barnet-Lamb et al., 2014). Finally we look at the 3-adic specialisation of this family. Reducing mod 3 we get a representation which is flat at 3 and tame at 2, so must be reducible because of the techniques introduced in Fontaine's paper on abelian varieties over $\mathbb{Z}$ (an irreducible representation would cut out a number field whose discriminant violates the Odlyzko bounds). One can now go on to deduce that the 3-adic representation must be reducible, which contradicts the irreducibility of ρ.

We apologise for the sketchiness of what is here, however at the time of writing it is so far from what we are even able to state in Lean that there seems to be little point right now in fleshing out the argument further. As this document grows, we will add a far more detailed discussion of what is going on here. Note in particular that stating the modularity lifting theorem in Lean is the first target.

5 An Example Of An Automorphic Form

5.1 Introduction

The key ingredient in Wiles' proof of Fermat's Last Theorem is a modularity lifting theorem, sometimes called an $R = T$ theorem. For Wiles, the R came from elliptic curves, the T came from classical modular forms, and the fact that they're equal is basically the Shimura–Taniyama–Weil conjecture, now known as the Breuil–Conrad–Diamond–Taylor modularity theorem: any elliptic curve over the rationals is modular.

At the heart of the proof we shall formalise is also an $R = T$ theorem, however the T which we shall use will be associated not to classical modular forms, but to spaces of more general automorphic forms called quaternionic modular forms. Those of you who know something about classical modular forms might well know that the groups $\mathrm{SL}_2(\mathbb{R})$ and $\mathrm{SL}_2(\mathbb{Z})$ are intimately involved; these are the norm 1 units in the matrix rings $M_2(\mathbb{R})$ and $M_2(\mathbb{Z})$. In the theory of quaternionic modular forms, the analogous groups are the norm 1 units in rings such as Hamilton's quaternions $\mathbb{R} \oplus \mathbb{R}i \oplus \mathbb{R}j \oplus \mathbb{R}k$, and subrings such as $\mathbb{Z} \oplus \mathbb{Z}i \oplus \mathbb{Z}j \oplus \mathbb{Z}k$.

One of the main goals of the FLT project at the time of writing this sentence, is formalising the statement of the modularity lifting theorem which we shall use. So we are going to need to develop the theory of quaternionic modular forms, which is rather different to the theory of classical modular forms (for example, in the cases we need, the definition is completely algebraic; there are no holomorphic functions in sight, and the analogue of the upper half plane in the quaternionic theory is a finite set of points).

We could just launch into the general theory over totally real fields, which will be the generality which we'll need. But when I was a PhD student, I learnt about these objects by playing with explicit examples. So, whilst not logically necessary for the proof, I thought it would be fun, and perhaps also instructional, to compute a concrete example of a space of quaternionic modular forms. The process of constructing the example might even inform what kind of machinery we should be developing in general. Let's begin by discussing the quaternion algebra we shall use.

5.2 A quaternion algebra

Let's define D to be the quaternion algebra $\mathbb{Q} \oplus \mathbb{Q}i \oplus \mathbb{Q}j \oplus \mathbb{Q}k$. As a vector space, D is 4-dimensional over $\mathbb{Q}$ with $[1, i, j, k]$ giving a basis. It has a (non-commutative) ring structure, with multiplication satisfying the usual quaternion algebra relations $i^2 = j^2 = k^2 = ijk = -1$. You can think of D as an analogue of `2 × 2` matrices with rational coefficients, hence its units $D^\times$ are an analogue of the group $\mathrm{GL}_2(\mathbb{Q})$.

We will also need an analogue of the group $\mathrm{GL}_2(\mathbb{Z})$, which will come from an integral structure on D. We choose the Hurwitz order, namely the subring $\mathcal{O} := \mathbb{Z} \oplus \mathbb{Z}i \oplus \mathbb{Z}j \oplus \mathbb{Z}\omega$, where $\omega = \frac{-1+(i+j+k)}{2}$, a cube root of unity, as $(i + j + k)^2 = -3$. The simplest way to understand $\mathcal{O}$ is that it's quaternions $a + bi + cj + dk$ where either a, b, c, d are all integers or are all in $\frac{1}{2} + \mathbb{Z}$.

Note that $\mathcal{O}$ is a maximal order and a Euclidean domain, which is why we prefer it over the more obvious sublattice $\mathbb{Z} \oplus \mathbb{Z}i \oplus \mathbb{Z}j \oplus \mathbb{Z}k$.

In this chapter, we are going to compute a complex vector space which could be called something like the "weight 2 level 2 modular forms for $D^\times$". The main result will be that this space is 1-dimensional.

Note that mathlib has modular forms, but it doesn't have enough complex analysis to deduce that the space of modular forms of a given weight and level is finite-dimensional. If all the `sorry`s in this chapter are completed before mathlib gets the necessary complex analysis, then the first nonzero space of modular forms to be proved finite-dimensional in Lean will be a space of quaternionic modular forms.

We will use a modern "adelic" definition of our modular forms, so the first thing we need to do is to talk about profinite completions.

5.3 Zhat

Definition 5.3.1

> The profinite completion $\widehat{\mathbb{Z}}$ of $\mathbb{Z}$ is the set of all compatible collections $c = (c_N)_N$ of elements of $\mathbb{Z}/N\mathbb{Z}$ indexed by $\mathbb{N}^+ := \{1, 2, 3, \ldots\}$. A collection is said to be compatible if for all positive integers $D \mid N$, we have $c_N \bmod D$ equals c_D.

Classically automorphic forms were defined as functions on symmetric spaces (like the upper half plane) which transformed well under the action of certain discrete groups (for example $\mathrm{SL}_2(\mathbb{Z})$). However such definitions became combinatorially problematic when generalised to number fields with nontrivial class group, because the classical theory needed a number p to define the Hecke operator T_p, and in the case where p was a non-principal prime ideal in a number field, there was no appropriate number. One fix is to take disjoint unions of symmetric spaces indexed by the ideal class group of the field in question, but it is easier to work adelically, which is morally what we shall do. However we are able to avoid introducing the adeles explicitly; we

can work instead with the conceptually simpler object $\widehat{\mathbb{Z}}$, the profinite completion of $\mathbb{Z}$. So what is $\widehat{\mathbb{Z}}$? We offer a low-level definition of this object.

Lemma 5.3.2

$\widehat{\mathbb{Z}}$ is a subring of $\prod_{N\geq 1}(Z/N\mathbb{Z})$ and in particular is a ring.

Proof for Lemma 5.3.2

Follow your nose.

Lemma 5.3.3

$0 \neq 1$ in $\widehat{\mathbb{Z}}$.

Proof for Lemma 5.3.3

Recall that you can evaluate an element of $\widehat{\mathbb{Z}}$ at a positive integer. Evaluating 0 at 2 gives 0, and evaluating 1 at 2 gives 1, and these are distinct elements of $\mathbb{Z}/2\mathbb{Z}$, so $0 \neq 1$ in $\widehat{\mathbb{Z}}$.

Lemma 5.3.4

The map from the naturals into $\widehat{\mathbb{Z}}$ sending n to n is injective.

Proof for Lemma 5.3.4

Generalise the above idea. Feel free to write up a LaTeX proof and PR it.

Note that it follows easily that that the map from the integers to $\widehat{\mathbb{Z}}$ is injective.

But $\widehat{\mathbb{Z}}$ is **much** larger than $\mathbb{Z}$; it has the same cardinality as the reals in fact. Let's write down an explicit example of an element of $\widehat{\mathbb{Z}}$ which isn't obviously in $\mathbb{Z}$.

Definition 5.3.5

The infinite sum $0!+1!+2!+3!+4!+5!+\cdots$ looks like it makes no sense at all; it is the sum of an infinite series of larger and larger positive numbers. However, the sum is finite modulo N for every positive integer N, because all the terms from $N!$ onwards are multiples of N and thus are zero in $\mathbb{Z}/N\mathbb{Z}$. Thus it makes sense to define e_N to be the value of the finite sum modulo N. Explicitly, $e_N = 0! + 1! + \cdots + (N-1)!$ modulo N.

Lemma 5.3.6

The collection $(e_N)_N$ is an element of $\widehat{\mathbb{Z}}$.

Proof for Lemma 5.3.6

This boils down to checking that $D! + (D+1)! + \cdots + (N-1)!$ is a multiple of D.

Lemma 5.3.7

The element $(e_N)_N$ of $\widehat{\mathbb{Z}}$ is not in $\mathbb{Z}$.

Proof for Lemma 5.3.7

First imagine that $e = n$ with $n \in \mathbb{Z}$ and $0 \leq n$. In this case, choose j such that $0! + 1! + 2! + \cdots + j! > n$ and check also that the sum is less than $(j+1)!$. Now set $N = (j+1)!$ and let's compare e_N and $n_N = n$. The trick is that e_N must be $0! + 1! + \cdots + j!$ mod N, because all the terms beyond this are multiples not just of $(j+1)$ but of $(j+1)! = N$. Thus mod N we have $0 \leq n < e_N < N$ so $n \neq e$.Now we deal with $n = -t < 0$; choose j large such that $(j+1)! - (0! + 1! + \cdots + j!) > t$ (possible because the sum is at most $2 \times j!$) and then set $N = (j+1)!$ and we have $0 < e_N < N - t < N$ so we cannot have $e_N = -t$ in $\mathbb{Z}/N\mathbb{Z}$, so again $e \neq n$.

Lemma 5.3.8

If $0 < N$ is an integer then multiplication by N is injective on $\widehat{\mathbb{Z}}$.

Proof for Lemma 5.3.8

Suppose that $(z_i)_i \in \widehat{\mathbb{Z}}$ and $Nz = 0$. This means that $Nz_i = 0 \in \mathbb{Z}/i\mathbb{Z}$ for all i. Let us fix an arbitrary positive integer j; we need to prove that $z_j = 0 \in \mathbb{Z}/j\mathbb{Z}$. Consider the element $z_{Nj} \in \mathbb{Z}/Nj\mathbb{Z}$. By assumption, we have $Nz_{Nj} = 0$, meaning that if we lift z_{Nj} to an integer, we have $Nj \mid Nz_{Nj}$, and thus $j \mid z_{Nj}$. Thus by the compatibility assumption on the z_i we have that $z_j \in \mathbb{Z}/j\mathbb{Z}$ is the mod j reduction of z_{Nj} and hence is zero.

Lemma 5.3.9

The multiples of N in $\widehat{\mathbb{Z}}$ are precisely the compatible collections $(z_i)_i \in \widehat{\mathbb{Z}}$ with $z_N = 0$.

Proof for Lemma 5.3.9

Clearly $z_N = 0$ is a necessary condition to be a multiple of N. To see it is sufficient, take a general $(z_i) \in \widehat{\mathbb{Z}}$ such that $z_N = 0$, and now define a new element $(y_j)_j$ of $\widehat{\mathbb{Z}}$ by $y_j = z_{Nj}/N$. Just to clarify what this means: $z_{Nj} \in \mathbb{Z}/Nj\mathbb{Z}$ reduces mod N to $z_N = 0$ by the compatibility assumption, so it is in the subgroup $N\mathbb{Z}/Nj\mathbb{Z}$ of $\mathbb{Z}/Nj\mathbb{Z}$, which is isomorphic (via "division by N") to the group $\mathbb{Z}/j\mathbb{Z}$; this is how we construct y_j. It is easily checked that the y_j are compatible and that $Ny = z$.

5.4 More Advanced Remarks On Zhat Versus Q

This section can be skipped on first reading.

People who have seen some more advanced algebra might recognise this construction of $\widehat{\mathbb{Z}}$ as being the profinite completion of the additive abelian group $\mathbb{Z}$, so it is a fundamental object of mathematics in some sense. But usually, when building mathematics, after $\mathbb{Z}$ we go to $\mathbb{Q}$, a multiplicative localisation of $\mathbb{Z}$, and only complete after that (to get $\mathbb{R}$). The process of "completing before localising" gives us a far more arithmetic completion of $\mathbb{Z}$.

Even though $\mathbb{Q}$ is a divisible abelian group and hence its profinite completion vanishes, we can still attempt to "locally profinitely complete it" by defining $\widehat{\mathbb{Q}} := \mathbb{Q} \otimes_{\mathbb{Z}} \widehat{\mathbb{Z}}$. This object is more commonly known as the finite adeles of $\mathbb{Q}$. More generally if F is any number field then $F \otimes_{\mathbb{Z}} \widehat{\mathbb{Z}}$ is the ring of finite adeles of F. To get to the full ring of adeles of a number field F you need to take the product with the ring of infinite adeles of F, which is $F \otimes_{\mathbb{Q}} \mathbb{R}$: some kind of universal archimedean completion of F. I don't know a reference which develops the theory of adeles in this way, so this is what we shall do here.

5.5 Qhat and tensor products.

The definition of $\widehat{\mathbb{Q}}$ is easy if you know about tensor products of additive abelian groups.

Definition 5.5.1

> The profinite completion $\widehat{\mathbb{Q}}$ of $\mathbb{Q}$ is the tensor product $\mathbb{Q} \otimes_{\mathbb{Z}} \widehat{\mathbb{Z}}$, or $\widehat{\mathbb{Q}} = \mathbb{Q} \otimes \widehat{\mathbb{Z}}$ for short.

5.6 A crash course in tensor products

We've defined $\widehat{\mathbb{Q}}$ to be $\mathbb{Q} \otimes \widehat{\mathbb{Z}}$. Whatever does this mean? Well just to orient yourself, if A and B are additive abelian groups, then $A \otimes B$ is also an abelian group. And if A and B are commutative rings (as they are in our case), then $A \otimes B$ is also a commutative ring.

Even if A and B are completely concrete commutative rings, their tensor product $A \otimes B$ might be incomprehensible. For example $\mathbb{C} \otimes \mathbb{C}$ is completely incomprehensible (note that we are tensoring over the integers). It is not like the product of groups or the disjoint union of two sets, where you have a completely explicit unambiguous formula for each element.

In this sense, the theory of tensor products is a bit like the theory of continuous functions. Humanity started off studying concrete polynomial equations such as $x^2 + 1$ and then moved on to concrete analytic functions such as $\log(x)$ and $\sin(x)$, but eventually the abstract concept of a continuous function from the reals to the reals was born. There is no "formula" for a general continuous function, and continuous functions such as e^{-1/x^2} or $|x|$ have no

power series. Even if there were a formula for a specific continuous function of interest, it is not clear in general how to make sense of the claim that it's the "best" formula. In other words, there is no "canonical form" for a general continuous function, and yet we prove things about them anyway. We shall adopt the same attitude for elements of $A \otimes B$.

The first thing to know about the tensor product $A \otimes B$ of two abelian groups A and B is a "constructor" for the type. In other words, how can we make elements $A \otimes B$? Well, it turns out that given elements $a \in A$ and $b \in B$, we can form the element $a \otimes_t b \in A \otimes B$.

Recall that the sum of all the factorials is an element $e \in \widehat{\mathbb{Z}}$, and $22/7$ is certainly a rational number, so we can make the element $\frac{22}{7} \otimes_t e \in \widehat{\mathbb{Q}}$.

This example is in the Lean code.

Elements of the form $a \otimes_t b \in A \otimes B$ are known as pure tensors. In the literature, pure tensors are often written $a \otimes b$, but we shall follow `mathlib`'s convention in reserving the $\otimes$ symbol for groups like $A \otimes B$, and adorning it with a t when using it on elements of the groups (or, as Lean calls them, terms, which explains the notation).

Addition of pure tensors obeys the "distributivity" rules $a \otimes_t b_1 + a \otimes_t b_2 = a \otimes_t (b_1 + b_2)$ and $a_1 \otimes_t b + a_2 \otimes_t b = (a_1 + a_2) \otimes_t b$, but there is no rule which simplifies a general sum $a \otimes_t b + c \otimes_t d$ into a pure tensor. Indeed, in general it is not the case that every element of a tensor product $A \otimes B$ is of the form $a \otimes_t b$; there can be tensors which aren't pure. However every element of $A \otimes B$ is a finite sum of pure tensors, with the result that one can attempt to define additive maps from $A \otimes B$ by saying what they do on pure tensors, and then extending linearly.

Another thing worth understanding is that just like how rational numbers can be written as quotients of integers in several ways (for example $1/2 = 2/4 = 3/6 = \cdots$), a general pure tensor in $A \otimes B$ can be represented as $a \otimes_t b$ in many ways. For example, in $\widehat{\mathbb{Q}}$ we have $1 \otimes_t 2 = 2 \otimes_t 1$. A general rule for equality of pure tensors is that if $a \in A$ and $b \in B$ and $z \in \mathbb{Z}$, then $za \otimes_t b = a \otimes_t zb$; integers can move over the tensor symbol. But equality is hard: in general there may not be an algorithm to decide whether two pure tensors $a \otimes_t b$ and $c \otimes_t d$ are equal in $A \otimes B$.

A summary of the situation: if A and B are abelian groups, then every element of $A \otimes B$ can be written in the form $\sum_{i=1}^{N} a_i \otimes_t b_i$. It's just that this representation is highly nonunique, and furthermore given explicit elements $a_1, a_2 \in A$ and $b_1, b_2 \in B$ it might be a hard problem to figure out if $a_1 \otimes_t b_1 = a_2 \otimes_t b_2$.

For example, it turns out that $(\mathbb{Z}/2\mathbb{Z}) \otimes (\mathbb{Z}/3\mathbb{Z}) = 0$ and so in this tensor product all the $a \otimes_t b$ are equal to each other and to $0 \otimes 0$.

Having said all of that, one nice property of $\widehat{\mathbb{Q}}$ is that every tensor is pure; let's prove this now.

Lemma 5.6.1

Every element of $\widehat{\mathbb{Q}} := \mathbb{Q} \otimes \widehat{\mathbb{Z}}$ can be written as $q \otimes_t z$ with $q \in \mathbb{Q}$

and $z \in \widehat{\mathbb{Z}}$. Furthermore one can even assume that $q = \frac{1}{N}$ for some positive integer N.

Proof for Lemma 5.6.1

A proof I would write on the board would look like the following. Take a general element of $\widehat{\mathbb{Q}}$; we know it can be expressed as a finite sum $\sum_i q_i \otimes_t z_i$ with $q_i \in \mathbb{Q}$ and $z_i \in \widehat{\mathbb{Z}}$. Now choose a large positive integer N, the lowest common multiple of all the denominators showing up in the q_i, and then rewrite $\sum_i q_i \otimes_t z_i$ as $\sum_i \frac{n_i}{N} \otimes z_i$ with $n_i \in \mathbb{Z}$. Now using the fundamental fact that $na \otimes_t b = a \otimes_t nb$ for $n \in \mathbb{Z}$, we can rewrite the sum as $\sum_i \frac{1}{N} \otimes_t n_i z_i$ which is equal to the pure tensor $\frac{1}{N} \otimes (\sum_i n_i z_i)$.In Lean I would prove this using `TensorProduct.induction_on`, which quickly reduces us to the claim that the sum of two pure tensors is pure, which we can prove using the above technique whilst avoiding the general theory of finite sums.

Be careful though: just because every element of $\widehat{\mathbb{Q}}$ can be written as $q \otimes z$, this representation may not be unique. For example $2 \otimes 1 = 1 \otimes 2$. However, writing $\frac{1}{N} \otimes_t z$ as z/N does tempt us into the following definition.

Definition 5.6.2

If $N \in \mathbb{N}^+$ and $z \in \widehat{\mathbb{Z}}$ then we say that N and z are coprime if $z_N \in (\mathbb{Z}/N\mathbb{Z})^\times$. We write z/N as notation for the element $\frac{1}{N} \otimes_t z$.

Lemma 5.6.3

Every element of $\widehat{\mathbb{Q}}$ can be uniquely written as z/N with $z \in \widehat{\mathbb{Z}}$, $N \in \mathbb{N}^+$, and with N and z coprime.

Proof for Lemma 5.6.3

Existence: by the previous lemma, an arbitrary element can be written as z/N; let D be the greatest common divisor of N and z_N (lifted to a natural). If $D = 1$ then the fraction is by definition in lowest terms. However if $1 < D \mid N$ then z_D is the reduction of z_N and is hence `0`. By lemma Lemma 5.3.9 we deduce that $z = Dy$ is a multiple of D, and hence $z/N = \frac{1}{N} \otimes_t Dy = \frac{1}{E} \otimes y$, where $E = N/D$. Now if a natural divided both y_E and E then this natural would divide both z_N/D and N/D, contradicting the fact that D is the greatest common divisor.Uniqueness: if $z/N = w/M$, we deduce $1 \otimes_t Mz = 1 \otimes_t Nw$, and by injectivity of $\widehat{\mathbb{Z}} \to \widehat{\mathbb{Q}}$ we deduce that $Mz = Nw = y$. In particular, if L is the lowest common multiple of M and N then y_L is a multiple of both M and N and is hence zero, so $y = Lx$ is a multiple of L by Lemma 5.3.9, and we deduce from torsionfreeness that $z = (L/M)x$ and $w = (L/N)x$.

If some prime divided L/M then it would have to divide N which means that z is not in lowest terms; similarly if some prime divided L/N then w/M would not be in lowest terms. We deduce that $L = M = N$ and hence $z = w$ by torsionfreeness.

If A and B are additive abelian groups then $A \otimes B$ is also an additive abelian group. However if A and B are commutative rings, then $A \otimes B$ also inherits the structure of a commutative ring, with $0 = 0 \otimes_t 0$ and $1 = 1 \otimes_t 1$. Pure tensors multiply in the obvious way: the product of $a_1 \otimes_t b_1$ and $a_2 \otimes_t b_2$ is $a_1a_2 \otimes_t b_1b_2$. There are ring homomorphisms $A \to A \otimes B$ and $B \to A \otimes B$ sending a to $a \otimes_t 1$ and b to $1 \otimes_t b$. In general such maps are not injective, but in the case of $\widehat{\mathbb{Q}} = \mathbb{Q} \otimes \widehat{\mathbb{Z}}$ both maps from $\mathbb{Q}$ and $\widehat{\mathbb{Z}}$ are inclusions.

Lemma 5.6.4

The ring homomorphism $\mathbb{Q} \to \widehat{\mathbb{Q}}$ sending q to $q \otimes_t 1$ is injective.

Proof for Lemma 5.6.4

We have seen that the map from $\mathbb{Z}$ to $\widehat{\mathbb{Z}}$ is injective. Now $\mathbb{Q}$ is a flat $\mathbb{Z}$-module, because it's torsion-free, so tensoring up we deduce that the map from $\mathbb{Q} = \mathbb{Q} \otimes \mathbb{Z}$ to $\widehat{\mathbb{Q}} = \mathbb{Q} \otimes \widehat{\mathbb{Z}}$ is also injective. There is no doubt a more elementary proof of this fact.

Lemma 5.6.5

The ring homomorphism $\widehat{\mathbb{Z}} \to \widehat{\mathbb{Q}}$ sending z to $1 \otimes_t z$ is injective.

Proof for Lemma 5.6.5

The map from $\mathbb{Z}$ to $\mathbb{Q}$ is injective, and we have seen that $\widehat{\mathbb{Z}}$ is a torsion-free and thus flat $\mathbb{Z}$-module, so the map from $\widehat{\mathbb{Z}}$ to $\widehat{\mathbb{Q}}$ is also injective.

5.7 Additive structure of Qhat.

Here we forget the ring structure on everything, and analyse $\widehat{\mathbb{Q}}$ as an additive abelian group, and in particular how the subgroups $\mathbb{Z}$, $\mathbb{Q}$ and $\widehat{\mathbb{Z}}$ sit within it.

The two results we prove in this section are that $\mathbb{Q} \cap \widehat{\mathbb{Z}} = \mathbb{Z}$ and that $\mathbb{Q} + \widehat{\mathbb{Z}} = \widehat{\mathbb{Q}}$. Using lattice-theoretic notation we could write these results as $\mathbb{Q} \sqcap \widehat{\mathbb{Z}} = \mathbb{Z}$ and $\mathbb{Q} \sqcup \widehat{\mathbb{Z}} = \widehat{\mathbb{Q}}$.

Lemma 5.7.1

The intersection of $\mathbb{Q}$ and $\widehat{\mathbb{Z}}$ in $\widehat{\mathbb{Q}}$ is $\mathbb{Z}$.

Proof for Lemma 5.7.1

Clearly $\mathbb{Z} \subseteq \mathbb{Q} \cap \widehat{\mathbb{Z}}$. Now suppose that $x \in \mathbb{Q} \cap \widehat{\mathbb{Z}}$. Because x is rational we can write it as $\frac{A}{B} \otimes_t 1$ for some fraction A/B in lowest terms, and hence $x = A/B$ where now we regard $A \in \widehat{\mathbb{Z}}$ and note that A/B is still in lowest terms. However $x \in \widehat{\mathbb{Z}}$ implies that $x = x/1$ is in lowest terms, so we deduce that $B = 1$ and thus $x = A \in \mathbb{Z}$.

Lemma 5.7.2

The sum of $\mathbb{Q}$ and $\widehat{\mathbb{Z}}$ in $\widehat{\mathbb{Q}}$ is $\widehat{\mathbb{Q}}$. More precisely, every element of $\widehat{\mathbb{Q}}$ can be written as $q + z$ with $q \in \mathbb{Q}$ and $z \in \widehat{\mathbb{Z}}$, or more precisely as $q \otimes_t 1 + 1 \otimes_t z$.

Proof for Lemma 5.7.2

Write $x \in \widehat{\mathbb{Q}}$ as $x = z/N$ in lowest terms. Lift z_N to an integer t and observe that $(z - t)_N = 0$, hence $z - t = Ny$ for some $y \in \widehat{\mathbb{Z}}$. Now $x = t/N + y \in \mathbb{Q} + \widehat{\mathbb{Z}}$.

5.8 Multiplicative structure of the units of Qhat.

We now forget the additive structure on the commutative ring $\widehat{\mathbb{Q}}$ and consider the multiplicative structure of its group of units $\widehat{\mathbb{Q}}^\times$ (which I couldn't get into the section title). We have the obvious subgroups $\mathbb{Q}^\times$, $\mathbb{Z}^\times$ and $\widehat{\mathbb{Z}}^\times$.

Lemma 5.8.1

The intersection of $\mathbb{Q}^\times$ and $\widehat{\mathbb{Z}}^\times$ in $\widehat{\mathbb{Q}}^\times$ is $\mathbb{Z}^\times$.

Proof for Lemma 5.8.1

Clearly the intersection is contained within $\mathbb{Q} \cap \widehat{\mathbb{Z}} = \mathbb{Z}$. If $n \in \mathbb{Z}$ is in $\widehat{\mathbb{Z}}^\times$ then $n \neq 0$ and its inverse $1/n = \pm 1/|n|$ is in lowest terms but also in $\widehat{\mathbb{Z}}$, and hence $|n| = 1$ by uniqueness of lowest term representation.

Note that by the previous lemma, this representation will be unique up to sign.

Lemma 5.8.2

The product of $\mathbb{Q}^\times$ and $\widehat{\mathbb{Z}}^\times$ in $\widehat{\mathbb{Q}}^\times$ is all of $\widehat{\mathbb{Q}}^\times$. More precisely, every element of $\widehat{\mathbb{Q}}^\times$ can be written as qz with $q \in \mathbb{Q}^\times$ and $z \in \widehat{\mathbb{Z}}^\times$.

Proof for Lemma 5.8.2

We already know that a general element of $\widehat{\mathbb{Q}}^\times$ can be written as x/N with N positive, so this reduces us to proving that a general element $x \in \widehat{\mathbb{Z}}$ which is invertible in $\widehat{\mathbb{Q}}^\times$ can be written as qz with $q \in \mathbb{Q}^\times$ and $z \in \widehat{\mathbb{Z}}^\times$.We know $1/x$ can be written in lowest terms as y/M, and multiplying up we deduce that $xy = M$, and hence x divides a positive integer. If $i : \mathbb{Z} \to \widehat{\mathbb{Z}}$ denotes the inclusion, then we've just seen that the preimage of the principal ideal (x), namely, $J := i^{-1}(x\widehat{\mathbb{Z}})$ is nonzero, as it contains M. Let $g \in J$ be the smallest positive integer; it's well-known that $J = (g)$.I claim that it suffices to show that $x\widehat{\mathbb{Z}} = g\widehat{\mathbb{Z}}$. Because knowing $g = yx$ and $x = gz$ for some $y, z \in \widehat{\mathbb{Z}}$ tells us that $g(1 - yz) = 0$, and we know that multiplication by g is injective, hence $yz = 1$, so z is a unit and we have written $x = gz$ with $g \in \mathbb{Q}^\times$ and $z \in \widehat{\mathbb{Z}}^\times$.It remains to prove the claim. By definition $g \in J \subseteq x\widehat{\mathbb{Z}}$ so this is one inclusion. For the other, it suffices to prove that $x_g = 0$. However if $0 < x_g < g$ lifts x_g to the naturals then I claim that $x_g \in J$, for $x_g - x$ is a multiple of g and hence of x, and this contradicts minimality of g.

5.9 The Hurwitz quaternions

Definition 5.9.1

The Hurwitz quaternions are the set $\mathcal{O} := \mathbb{Z} \oplus \mathbb{Z}\omega \oplus \mathbb{Z}i \oplus \mathbb{Z}i\omega$ (as an abstract abelian group or as a subgroup of the usual quaternions). Here $\omega = \frac{-1+(i+j+k)}{2}$ and note that $(i + j + k)^2 = -3$. We have $\overline{\omega} = \omega^2 = -(\omega + 1)$. A general quaternion $a + bi + cj + dk$ is a Hurwitz quaternion if either $a, b, c, d \in \mathbb{Z}$ or $a, b, c, d \in \mathbb{Z} + \frac{1}{2}$.

Lemma 5.9.2

The Hurwitz quaternions form a ring.

Proof for Lemma 5.9.2

Follow your nose.

This ring is isomorphic to $\mathbb{Z}^4$ as an additive group, and $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{R} = \mathbb{R} \oplus \mathbb{R}i \oplus \mathbb{R}j \oplus \mathbb{R}\omega$ is the usual Hamilton quaternions.

Definition 5.9.3

There's a conjugation map (which we'll call "star") from the Hurwitz quaternions to themselves, sending integers to themselves and purely imaginary elements like $2\omega + 1$ to minus themselves. It satisfies $(x^*)^* = x$, $(xy)^* = y^*x^*$ and $(x+y)^* = x^* + y^*$. In particular, the Hurwitz quaternions are a "star ring" in the sense of mathlib.

Definition 5.9.4

The Hurwitz quaternions come equipped with an integer-valued norm, which is $a^2 + b^2 + c^2 + d^2$ on $a + bi + cj + dk$ but needs to be modified a bit to deal with ω.

Lemma 5.9.5

We have $N(x) = x\overline{x}$.

Proof for Lemma 5.9.5

Easy calculation.

Lemma 5.9.6

The norm of 0 is 0.

Proof for Lemma 5.9.6

A calculation.

Lemma 5.9.7

The norm of 1 is 1.

Proof for Lemma 5.9.7

A calculation.

Lemma 5.9.8

The norm of a product is the product of the norms.

Proof for Lemma 5.9.8

A calculation.

Lemma 5.9.9

The norm of an element is nonnegative.

Proof for Lemma 5.9.9

It's a sum of rational squares.

Lemma 5.9.10

The norm of an element is zero if and only if the element is zero.

Proof for Lemma 5.9.10

It's a sum of rational squares.

Lemma 5.9.11

Given a "usual" quaternion $a = x + yi + zj + wk$ with $x, y, z, w \in \mathbb{R}$, there exists a Hurwitz quaternion q such that $N(a - q) < 1$.

Proof for Lemma 5.9.11

If $[r]$ denotes the nearest integer to the real number r, then $|r - [r]| \leq \frac{1}{2}$. Hence if $q = [x] + [y]i + [z]j + [w]k$ then $N(a - q) = |x - [x]|^2 + \cdots \leq \frac{1}{4} + \frac{1}{4} + \frac{1}{4} + \frac{1}{4} \leq 1$, with strict inequality unless $|x - [x]| = |y - [y]| = |z - [z]| = |w - [w]| = \frac{1}{2}$, in which case $a \in \mathcal{O}$ because $a - \omega$ has integer coordinates.

Lemma 5.9.12

Given two Hurwitz quaternions a and b with b nonzero, there exists q and r such that $a = qb + r$ and $N(r) < N(b)$.

Proof for Lemma 5.9.12

Let q be the Hurwitz quaternion obtained by applying lemma Lemma 5.9.11 to $a/b := ab^{-1}$; then $N(a/b - q) < 1$ and now everything follows after multiplying up.

Corollary 5.9.13

All left ideals of $\mathcal{O}$ are principal.

Proof for Corollary 5.9.13

If the ideal is `0`, use `0`. Otherwise, choose a nonzero element of smallest norm.

All right ideals are principal too, because there's another version of Euclid saying $a = bq + r$.

5.10 Profinite completion of the Hurwitz quaternions

We define $\widehat{\mathcal{O}}$ to be $\mathcal{O} \otimes \widehat{\mathbb{Z}}$, so it's elements $a + bi + cj + d\omega$ with $a, b, c, d \in \widehat{\mathbb{Z}}$. The basic thing we need is this:

Theorem 5.10.1

If N is a positive natural then the obvious map $\mathcal{O} \to \widehat{\mathcal{O}}/N\widehat{\mathcal{O}}$ is surjective.

Proof for Theorem 5.10.1

This is just four copies of the surjection $\mathbb{Z} \to \widehat{\mathbb{Z}}/N\widehat{\mathbb{Z}}$. Note that this latter map is surjective because $\mathbb{Z} \to \mathbb{Z}/N\mathbb{Z}$ is surjective, hence given $z \in \widehat{\mathbb{Z}}$ you can subtract an integer w such that $(z - w)_N = 0$, so $z - w$ is a multiple of N.

We define $D := \mathbb{Q} \otimes \mathcal{O} = \mathbb{Q} \oplus \mathbb{Q}i \oplus \mathbb{Q}j \oplus \mathbb{Q}\omega = \mathbb{Q} \oplus \mathbb{Q}i \oplus \mathbb{Q}j \oplus \mathbb{Q}k$. Finally, we define $\widehat{D} := D \otimes \widehat{\mathbb{Z}}$. Just as with $\widehat{\mathbb{Q}}$ we have

Lemma 5.10.2

Every element of $\widehat{D}$ can be written as z/N with $z \in \widehat{\mathcal{O}}$ and $N \in \mathbb{N}^+$.

Proof for Lemma 5.10.2

Same as the proof for $\widehat{\mathbb{Q}}$.

It is not hard to check that $\widehat{D}$ contains $\widehat{\mathcal{O}}$ and D as subrings, and that as additive abelian groups we have $\widehat{\mathcal{O}} \cap D = \mathcal{O}$ and $\widehat{\mathcal{O}} + D = \widehat{D}$. This is because $\mathcal{O}$ is just four copies of $\mathbb{Z}$ and we've proved the analogous result for $\mathbb{Z}$.

However the multiplicative structure is more interesting, especially as D is not commutative. For a general quaternion algebra it is not true that $(\widehat{D})^\times = D^\times(\widehat{\mathcal{O}})^\times$, because there are "class group obstructions". The double coset space is some kind of non-commutative analogue of a class group. However for our particular choice of D and $\mathcal{O}$ the result is true.

Theorem 5.10.3

The group of units of $\widehat{D}$ is $D^\times\widehat{\mathcal{O}}^\times$. More precisely, every element of $\widehat{D}^\times$ can be written as a product δu with $\delta \in D^\times$ and $u \in \widehat{\mathcal{O}}^\times$.

Proof for Theorem 5.10.3

Given an element x of $\widehat{D}^\times$, we can use lemma Lemma 5.10.2 to write it as z/N with N a positive integer and $z \in \widehat{\mathcal{O}}$. Note that N is central and in $D^\times$. Similarly, we can write x^{-1} as y/M with M a positive integer and $y \in \widehat{\mathcal{O}}$. Then $1 = xx^{-1} = zy/NM$ and so $zy = NM = MN$, and $1 = x^{-1}x = yz/MN$ so $yz = MN$ too. In

particular y both left and right divides a positive integer.Now consider the left ideal $\widehat{\mathcal{O}}y$ generated by y. We've just seen that this ideal has nontrivial intersection with $\mathcal{O}$, because it contains $MN > 0$. Hence its intersection with $\mathcal{O}$ is a nonzero left ideal of $\mathcal{O}$, which is hence principal by corollary Corollary 5.9.13. Write it as $\mathcal{O}\alpha$ with $0 \neq \alpha \in \mathcal{O}$.It suffices to show that $\widehat{\mathcal{O}}\alpha = \widehat{\mathcal{O}}y$. For this would imply that $u\alpha = y$ and $vy = \alpha$ for some $u, v \in \widehat{\mathcal{O}}$, and thus $(vu-1)\alpha = 0$ and $(uv-1)y = 0$, and both α and y are left divisors of positive integers (the norm of α, and MN respectively), so now using the fact that $\widehat{\mathcal{O}}$ is $\mathbb{Z}$-torsion-free (is the tensor product of torsion-free abelian groups torsion-free? That would be a cheap way of doing it. Otherwise use $\mathcal{O} = \mathbb{Z}^4$) we deduce that u and v are units, and thus $x^{-1} = \frac{1}{M}u\alpha$ so $x = (M\alpha^{-1})v \in D^\times\widehat{\mathcal{O}}^\times$.What remains is this. We have $y \in \widehat{\mathcal{O}}$ which left and right divides some positive integer. We've defined $0 \neq \alpha \in \mathcal{O}$ such that $\mathcal{O}\alpha$ is the pullback of the abelian group $\widehat{\mathcal{O}}y$ along the map $\mathcal{O} \to \widehat{\mathcal{O}}$. We need to show that when we push this ideal $\mathcal{O}\alpha$ forwards to $\widehat{\mathcal{O}}$ we get $\widehat{\mathcal{O}}y$ again. The fact that $\widehat{\mathcal{O}}\alpha \subseteq \widehat{\mathcal{O}}y$ is easy, because $\alpha \in \widehat{\mathcal{O}}y$ by definition. So it remains to show that $y \in \widehat{\mathcal{O}}\alpha$.Let's define T to be a positive integer which is both a left and right multiple of both y and α (for example $T = MN\alpha\overline{\alpha}$ will do). Now note that we have an isomorphism $\mathcal{O}/T\mathcal{O} = \widehat{\mathcal{O}}/T\widehat{\mathcal{O}}$, so we can choose some $\beta \in \mathcal{O}$ such that $\beta - y \in T\widehat{\mathcal{O}}$ is a multiple of T. Next note that $\beta \in y + \widehat{\mathcal{O}}T \subset \widehat{\mathcal{O}}y$ is in $\widehat{\mathcal{O}}y \cap \mathcal{O} = \mathcal{O}\alpha$, meaning $\beta = \gamma\alpha$ for some $\gamma \in \mathcal{O}$. Hence $y \in \beta + \widehat{\mathcal{O}}T \subseteq \widehat{\mathcal{O}}\alpha$.

6 Stating the modularity lifting theorems

I think that a nice and accessible goal (which will maybe take a month or two) would be to state the modularity lifting theorems which we'll be formalising. There are in fact two; one (the "minimal case") is proved using an extension of the original Taylor–Wiles techniques, and the other is deduced from it using various more modern tricks which were developed later. This chapter (currently work in progress) will contain a detailed discussion of all the things involved in the statement of the theorem.

6.1 Automorphic forms and analysis

Modular forms were historically the first nontrivial examples of automorphic forms, but by the 1950s or so it was realised that they were special cases of a very general notion of an automorphic form, as were Dirichlet characters! Modular forms are holomorphic automorphic forms for the group $\mathrm{GL}_2/\mathbb{Q}$, and Dirichlet characters are automorphic forms for the group $\mathrm{GL}_1/\mathbb{Q}$. It's possible to make sense of the notion of an automorphic form for the group G/k. Here k is a "global field" – that is, a field which is either a finite extension of $\mathbb{Q}$ (a number field) or a finite extension of $(\mathbb{Z}/p\mathbb{Z})(T)$ (a function field), and G is a connected reductive group variety over k.

The reason that the definition of a modular form involves some analysis (they are holomorphic functions) is that if you quotient out the group $\mathrm{GL}_2(\mathbb{R})$ by its centre and the maximal compact subgroup $O_2(\mathbb{R})$, you get something which can be naturally identified with the upper half plane, a symmetric space with lots of interesting differential operators associated to it (for example a Casimir operator). However if you do the same thing with $\mathrm{GL}_1(\mathbb{R})$ then you get a one point set, which is why a Dirichlet character is just a combinatorial object; it's a group homomorphism $(\mathbb{Z}/N\mathbb{Z})^\times \to \mathbb{C}^\times$ where N is some positive integer. It turns out that there are many other connected reductive groups where the associated symmetric space is 0-dimensional, and in these cases the definition of an automorphic form is again combinatorial. An example would be the group variety associated to the units of a totally definite quaternion algebra over a totally real field. In this case, the analogue of $\mathrm{GL}_2(\mathbb{R})$ would be the

units $\mathbb{H}^\times$ in the Hamilton quaternions, a maximal compact subgroup would be the quaternions of norm 1 (homeomorphic to the 3-sphere S^3) and quotienting out $\mathbb{H}^\times$ by its centre $\mathbb{R}^\times$ and S^3 again just gives you 1 point.

Before we talk about quaternion algebras, let's talk about central simple algebras.

6.2 Central simple algebras

Convention: in this section, fields are commutative, but algebras over a field may not be.

Recall that a central simple algebra over a field K is a nonzero K-algebra D such that K is the centre of D and that D has no nontrivial two-sided ideals.

Another way of saying that D has no nontrivial two-sided ideals: every surjective ring homomorphism $D \twoheadrightarrow A$ to any ring A is either an isomorphism, or the zero map to the zero ring. Note that this latter condition has nothing to do with K.

Lemma 6.2.1

If $n \geq 1$ then the $n \times n$ matrices $M_n(K)$ are a central simple algebra over K.

Proof for Lemma 6.2.1

We prove more generally that matrices with coefficients in `K` and indexed by an arbitrary nonempty finite type are a central simple algebra over `K`.They are clearly an algebra over `K`, with `K` embedded via scalar matrices as usual (the injectivity of the map from `K` comes from nonemptiness of the finite index type). The centre clearly contains `K`; to show that it equals `K`, we argue as follows. Let `e(i,j)` be the matrix with a `1` in the `i`th row and `j`th column, and zeros everywhere else. An element $Z = (Z_{s,t})_{s,t}$ of the centre commutes with all matrices `e(i,j)` for $i \neq j$ and these equations immediately imply that $Z_{i,j} = 0$ if $i \neq j$ and that $Z_{i,i} = Z_{j,j}$.It suffices to prove that any nonzero two-sided ideal `I` is all of $M_n(K)$. So say $0 \neq M \in I$ and let us fix (i, j) such that $M_{i,j} \neq 0$. One easily checks that $M_{i,j}\mathrm{id} = \sum_k e(k, i) \times M \times e(j, k) \in I$, where $\mathrm{id} \in M_n(K)$ is the identity matrix. Therefore, $\mathrm{id} \in I$, so $I = M_n(K)$.The definition also requires that the ring be nonzero, but this follows from the index type being nonempty.

Lemma 6.2.2

If D is a central simple algebra over K and L/K is a field extension, then $L \otimes_K D$ is a central simple algebra over L.

Proof for Lemma 6.2.2

This is not too hard: it's lemma b of section 12.4 in Peirce's "Associative algebras".

Next: define trace and norm.

7 Automorphic forms and the Langlands Conjectures

7.1 Definition of an automorphic form for GLn over Q

The global Langlands reciprocity conjectures relate automorphic forms to Galois representations. The statements for a general connected reductive group involve the construction of the Langlands dual group, and we do not have quite enough Lie algebra theory to push this definition through in general. However if we restrict the special case of the group $\mathrm{GL}_n/\mathbb{Q}$, the dual group is just $\mathrm{GL}_n(\mathbb{C})$ and several other technical obstructions are also removed. In this section we will explain the definition of an automorphic form for the group $\mathrm{GL}_n/\mathbb{Q}$, following the exposition by Borel and Jacquet in Corvallis (volume, 1979).

7.2 The Finite Adeles Of The Rationals

Mathlib already has the definition of the finite adeles $\mathbb{A}_{\mathbb{Q}}^f$ of the rationals as a commutative $\mathbb{Q}$-algebra, and the proof that it's a topological ring.

7.3 The Group GLn Of The Adeles

The adeles $\mathbb{A}_{\mathbb{Q}}$ of $\mathbb{Q}$ are the product $\mathbb{A}_{\mathbb{Q}}^f \times \mathbb{R}$, with the product topology. They are a topological ring. Hence $\mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}) = \mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}^f) \times \mathrm{GL}_n(\mathbb{R})$ is a topological group, where we are being a bit liberal with our use of the equality symbol.

7.4 Smooth Functions

Definition 7.4.1

A function $f : \mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}^f) \times \mathrm{GL}_n(\mathbb{R}) \to \mathbb{C}$ is smooth if it has the following three properties:

1. f is continuous.

2. For all $x \in \mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}^f)$, the function $y \mapsto f(x, y)$ is smooth.
3. For all $y \in \mathrm{GL}_n(\mathbb{R})$, the function $x \mapsto f(x, y)$ is locally constant.

Current state of this definition: I've half-formalised it; I don't know how to say the the function is smooth on the infinite part, because I have never used the manifold library before and I have no idea what my model with corners is supposed to be.

7.5 Slowly-Increasing Functions

Automorphic representations satisfy a growth condition which we may as well factor out into a separate definition.

We define the following temporary "size" function $s : \mathrm{GL}_n(\mathbb{R}) \to \mathbb{R}$ by $s(M) = \mathrm{trace}(MM^T + M^{-1}M^{-T})$, where M^{-T} denotes inverse-transpose. Note that $s(M)$ is always positive, and is large if M has a very large or very small, in absolute value, eigenvalue.

Definition 7.5.1

We say that a function $f : \mathrm{GL}_n(\mathbb{R}) \to \mathbb{C}$ is slowly-increasing if there is some real constant C and positive integer n such that $|f(M)| \leq Cs(M)^n$ for all $M \in \mathrm{GL}_n(\mathbb{R})$.

Note: the book says n is positive, but $\{M | s(M) \leq 1\}$ is compact so I don't think it makes any difference.

7.6 Weights At Infinity

Definition 7.6.1

The weight of an automorphic form for $\mathrm{GL}_n/\mathbb{Q}$ can be thought of as a finite-dimensional continuous complex representation ρ of a maximal compact subgroup of $\mathrm{GL}_n(\mathbb{R})$, and it's convenient to choose one (they're all conjugate) so we choose $O_n(\mathbb{R})$.

The Lean definition is incomplete right now – I don't demand irreducibility (I wasn't sure whether I was doing this the right way; if I used category theory then I might have struggled to say that the representation was continuous).

7.7 The Action Of The Universal Enveloping Algebra

Definition 7.7.1

There is a natural action of the real Lie algebra of $\mathrm{GL}_n(\mathbb{R})$ on the complex vector space of smooth complex-valued functions on $\mathrm{GL}_n(\mathbb{R})$.

Definition 7.7.2

This extends to a natural complex Lie algebra action of the complexification of the real Lie algebra on the smooth complex-valued functions on $\mathrm{GL}_n(\mathbb{R})$. This depends on Definition 7.7.1.

Definition 7.7.3

By functoriality, we get an action of the universal enveloping algebra of this complexified Lie algebra on the smooth complex-valued functions. This depends on Definition 7.7.2.

Definition 7.7.4

Thus the centre Z_n of this universal enveloping algebra also acts on the smooth complex-valued functions. This depends on Definition 7.7.3.

The centre we just defined is a commutative ring which contains a copy of $\mathbb{C}$. Note that Harish-Chandra, or possibly this was known earlier, showed that it is a polynomial ring in n variables over the complexes. We shall not need this.

7.8 Automorphic Forms

From here on there is no more Lean right now, only LaTeX.
Definition 7.8.1

A smooth function $f : \mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}^f) \times \mathrm{GL}_n(\mathbb{R}) \to \mathbb{C}$ is an $O_n(\mathbb{R})$-automorphic form on $\mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}})$ if it satisfies the following five conditions. This depends on Definition 7.4.1, Definition 7.5.1, Definition 7.6.1, and Definition 7.7.4.

1. Periodicity: for all $g \in \mathrm{GL}_n(\mathbb{Q})$, we have $f(gx, gy) = f(x, y)$.
2. It has a finite level: there exists a compact open subgroup $U \subseteq \mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}^f)$ such that $f(xu, y) = f(x, y)$ for all $u \in U$, $x \in \mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}^f)$, and $y \in \mathrm{GL}_n(\mathbb{R})$.
3. It has weight ρ: there exists a continuous finite-dimensional irreducible complex representation ρ of $O_n(\mathbb{R})$ such that for every $(x, y) \in \mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}})$, the complex vector space spanned by the functions $k \mapsto f(x, yk)$ is finite-dimensional and isomorphic, as an $O_n(\mathbb{R})$-representation, to a direct sum $\rho^{\oplus m}$ of copies of ρ for some m.

4. It has an infinite level: there is an ideal $I \subseteq Z_n$ of the centre Z_n described in the previous section, with finite complex codimension, and I annihilates the function $y \mapsto f(x, y)$ for all $x \in \mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}^f)$. This is a very fancy way of saying that the function satisfies some natural differential equations. In the case of modular forms, these are the Cauchy-Riemann equations, which is why modular forms are holomorphic.
5. It satisfies the growth condition: for every $x \in \mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}^f)$, the function $y \mapsto f(x, y)$ on $\mathrm{GL}_n(\mathbb{R})$ is slowly-increasing.

Automorphic forms of a fixed weight ρ form a complex vector space, and if we also fix the finite level U and the infinite level I then we get a subspace which is finite-dimensional; this is a theorem of Harish-Chandra. There is also the concept of a cusp form, meaning an automorphic form for which furthermore some adelic integrals vanish.

7.9 Hecke Operators

Lemma 7.9.1

The group $\mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}^f)$ acts on the space of automorphic forms for $\mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}})$ by the formula $(g \cdot f)(x, y) = f(xg, y)$.

Proof for Lemma 7.9.1

This is obvious. Note that the conjugate of a compact open subgroup is still compact and open.

A formal development of the theory of Hecke operators looks like the following.

Let U be a fixed compact open subgroup of $\mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}^f)$, and let us also fix a weight ρ. Let $M_\rho(n)$ denote the complex vector space of automorphic forms for $\mathrm{GL}_n/\mathbb{Q}$ of weight ρ. The level U forms $M_\rho(n, U)$ are just the U-invariants of this space. If $g \in \mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}}^f)$, then the double coset space UgU can be written as a finite disjoint union of single cosets g_iU: the double coset space is certainly a disjoint union of left cosets, but the double coset space is compact and the left cosets are open.

Define the Hecke operator $T_g : M_\rho(n, U) \to M_\rho(n, U)$ by $T_g(f) = \sum g_i \cdot f$.

Lemma 7.9.2

This function is well-defined, i.e. it sends a U-invariant form to a U-invariant form which is independent of the choice of g_i.

Proof for Lemma 7.9.2

Easy group theory.

8 Miniproject: Frobenius Elements

8.1 Status

This miniproject has been a success: the main results are sorry-free and merged into mathlib. As a result there will be no more work on this miniproject in the FLT repo. Below is a fairly detailed sketch of the argument used.

8.2 Introduction And Goal

When this project started, I had thought that the existence of Frobenius elements was specific to the theory of local and global fields, and a slightly more general result held for Dedekind domains. Then Joel Riou pointed out on the Lean Zulip an extremely general result from from Bourbaki's Commutative Algebra (Chapter V, Section 2, number 2, theorem 2, part (ii)). This beautiful result is surely what we want to see in mathlib. Before we state Bourbaki's theorem, let us set the scene.

In particular, $\mathrm{Aut}(L/K)$ is finite as a corollary. What is so striking about this theorem to me is that the only finiteness hypothesis is on the group `G` which acts; there are no finiteness or Noetherian hypotheses on the rings at all.

8.3 Statement Of The Theorem

The set-up throughout this project: G is a finite group acting (via ring isomorphisms) on a commutative ring B, and A is the subring of G-invariants.

Say Q is a prime ideal of B, whose pullback to A is the prime ideal P. Note that G naturally acts on the ideals of B. Let the decomposition group D_Q of Q be the subgroup of G which stabilizes Q (just to be clear: $g \in D_Q$ means $\{g \cdot q : q \in Q\} = Q$, not $\forall q \in Q, g \cdot q = q$).

Let L be the field of fractions of the integral domain B/Q, and let K be the field of fractions of the subring A/P. Then L is naturally a K-algebra. In this generality L/K may not even be finite or Galois, but we can still talk about $\mathrm{Aut}(L/K)$.

In the next definition we write down a group homomorphism ϕ from D_Q to $\mathrm{Aut}(L/K)$.

Definition 8.3.1

Choose $g \in D_Q$. Then the action of g on B gives us an induced A/P-algebra automorphism of B/Q which extends to a K-algebra automorphism $\phi(g)$ of L. This construction $g \mapsto \phi(g)$ defines a group homomorphism from D_Q to $\mathrm{Aut}(L/K)$ (all the proofs implicit in the definition here are straightforward).

The theorem we want in this mini-project is

Theorem 8.3.2

The map $g \mapsto \phi_g$ from D_Q to $\mathrm{Aut}(L/K)$ defined above is surjective.

The goal of this mini-project is to get this theorem formalised and ideally into mathlib.

In particular, $\mathrm{Aut}(L/K)$ is finite as a corollary. What is so striking about this theorem to me is that the only finiteness hypothesis is on the group G which acts; there are no finiteness or Noetherian hypotheses on the rings at all.

8.4 Examples

As a trivial consequence we get Frobenius elements for finite Galois extensions in both the local and global field setting, as $\mathrm{Aut}(L/K)$ is just a Galois group of finite fields in these cases, so by surjectivity we can lift a Frobenius element.

Even though `G` is finite, it is possible in characteristic $p > 0$ for the extension L/K to be infinite (and mostly inseparable). The theorem implies that $\mathrm{Aut}(L/K)$ is always finite; what is actually happening is that L/K is algebraic and normal, and its maximal separable subextension is finite of degree at most `|G|`. However, we can prove surjectivity directly without reference to this maximal separable subextension.

8.5 The Extension B/A

The precise set-up we'll work in is the following. We fix G a finite group acting on B a commutative ring, and we have another commutative ring A such that B is an A-algebra and the image of A in B is precisely the G-invariant elements of B. We do not ever need the map $A \to B$ to be injective so we do not assume this.

We start with a construction which is fundamental to everything, and which explains why we need G to be finite.

Definition 8.5.1

If $b \in B$ then define the characteristic polynomial $F_b(X) \in B[X]$ of b to be $\prod_{g \in G}(X - g \cdot b)$.

Clearly F_b is a monic polynomial.

Lemma 8.5.2

F_b is monic.

Proof for Lemma 8.5.2

Obvious.

It is also clear that F_b has degree $|G|$ and has b as a root. Also F_b is G-invariant, because acting by some $\gamma \in G$ just permutes the order of the factors. Hence we can descend F_b to a monic polynomial $M_b(X)$ of degree $|G|$ in $A[X]$. We will also refer to M_b as the characteristic polynomial of b, even though it is not even well-defined if the map $A \to B$ is not injective.

Lemma 8.5.3

F_b is the lift of some monic polynomial M_b in $A[X]$.

Proof for Lemma 8.5.3

The coefficients of F_b are G-invariant, and thus lie in the image of A.

Theorem 8.5.4

B/A is integral.

Proof for Theorem 8.5.4

Use M_b.

If `Q` and `Q'` are two primes above `p` then there is some $g \in G$ such that $gQ = Q'$, and one can deduce from this that `Frob_Q` and `Frob_{Q'}` are conjugate. In particular if `G` is abelian then `Frob_Q` and `Frob_{Q'}` are equal, so we can call them both `Frob_p`.

8.6 The Extension (B/Q)/(A/P)

Note that P and Q are primes, so the quotients A/P and B/Q are integral domains.

The following technical lemma constructs an element of B with nice characteristic polynomial modulo Q.

Lemma 8.6.1

There exist elements $a, b \in B$, with $a \notin Q$ and a in the image of A, such that for all $g \in G$:

- If $g \cdot Q = Q$, then $g \cdot b \equiv a \pmod{Q}$.

- If $g \cdot Q \neq Q$, then $g \cdot b \equiv 0 \pmod{Q}$.

Proof for Lemma 8.6.1

The ideals $g \cdot Q \neq Q$ are not contained in Q. Since Q is a prime ideal, this implies that the intersection of all $g \cdot Q \neq Q$ is still not contained in Q. Then we can find an element $c \notin Q$ with $c \in g \cdot Q$ for all $g \cdot Q \neq Q$. Let F_c be the characteristic polynomial of c, and write $F_c(X) \equiv X^j R(X) \pmod{Q}$. Let a be the coefficient of X^j in $F_c(X)$, and choose $R(X)$ so that $R(0) = a$. Let $b = R(0) - R(c)$. Note that $F_c(c) = 0$ and $c \not\equiv 0 \pmod{Q}$, so $R(c) \equiv 0 \pmod{Q}$. Then $b \equiv a \pmod{Q}$, so $g \cdot b \equiv a \pmod{Q}$ whenever $g \cdot Q = Q$. But if $g \cdot Q \neq Q$, then $c \equiv 0 \pmod{g \cdot Q}$. Then $b \equiv 0 \pmod{g \cdot Q}$, so $g \cdot b \equiv 0 \pmod{Q}$ whenever $g \cdot Q \neq Q$.

A slight modification allows us to take an element of B fixed by D_Q as input.

Lemma 8.6.2

Let $b_0 \in B$. Suppose that the image of b_0 in the quotient B/Q is fixed by the stabilizer subgroup D_Q. Then there exist elements $a, b \in B$, with $a \notin Q$ and a in the image of A, such that for all $g \in G$:

- If $g \cdot Q = Q$, then $g \cdot b \equiv a b_0 \pmod{Q}$.
- If $g \cdot Q \neq Q$, then $g \cdot b \equiv 0 \pmod{Q}$.

Proof for Lemma 8.6.2

Multiply the b from Lemma 8.6.1 by b_0.

8.7 The Extension L/K

Let L^{D_Q} denote the fixed field of the action D_Q on L. Our strategy for proving surjectivity of $D_Q \to \mathrm{Aut}(L/K)$ will be to write this map as the composition $D_Q \to \mathrm{Aut}(L/L^{D_Q}) \to \mathrm{Aut}(L/K)$.

The surjectivity of the first map is a general fact of Galois theory.

Theorem 8.7.1

If a finite group H acts on a field F by field automorphisms, then the map $H \to \mathrm{Aut}(F/F^H)$ is surjective.

Proof for Theorem 8.7.1

This is a general fact of Galois theory and was already in mathlib.

For surjectivity of the second map, we need to know that every element of L^{D_Q} is fixed by $\mathrm{Aut}(L/K)$. We first show this for elements of B/Q fixed by D_Q.

Theorem 8.7.2

Let $b_0 \in B/Q$. Suppose that b_0 is fixed by the stabilizer subgroup D_Q. Then b_0 is fixed by $\mathrm{Aut}(L/K)$.

Proof for Theorem 8.7.2

Let $a, b \in B$ be elements from `fixed_of_fixed1_aux2`. Let $M_b \in A[X]$ be the characteristic polynomial of b. We can map M_b to $L[X]$ in two different ways: via $B[X]$ and via $K[X]$. Going via $B[X]$ tells us that the image of $M_b(X)$ in $L[X]$ is exactly $(X - ab_0)^{|D_Q|}X^{|G|-|D_Q|}$. But going via $K[X]$ tells us that this image lies in $K[X]$, so we must have $ab_0 \in K$. Then ab_0 is fixed by $\mathrm{Aut}(L/K)$, and a is nonzero in L (since $a \notin Q$), so b_0 is fixed by $\mathrm{Aut}(L/K)$.

Now we upgrade this to elements of L fixed by D_Q. The following lemma will allow us to lift the denominator from B/Q to A/P.

Lemma 8.7.3

If R/S is an algebraic extension of integral domains, then any fraction a/b with $a, b \in R$ can be written as c/d with $c \in R$ and $d \in S$.

Proof for Lemma 8.7.3

If $f \in S[X]$ satisfies $f(b) = 0$, then $f(0) \in S$ is a multiple of b. If $f(0) = bx \in S$, then $a/b = (ax)/(bx)$ as desired.

Theorem 8.7.4

Let $x \in L$. Suppose that x is fixed by the stabilizer subgroup D_Q. Then x is fixed by the automorphism group $\mathrm{Aut}(L/K)$.

Proof for Theorem 8.7.4

Since $(B/Q)/(A/Q)$ is algebraic by `Algebra.IsInvariant.isIntegral`, `IsAlgebraic.exists_smul_eq_mul` let's us write $x = b/a$ for $b \in B/Q$ and $a \in A/P$. Then b is fixed by the stabilizer subgroup D_Q, and it suffices to show that b is fixed by the automorphism group $\mathrm{Aut}(L/K)$. But this is exactly `fixed_of_fixed1`.

Combining this with Theorem 8.7.1 finishes the proof.

Proof for Theorem 8.3.2

The map $D_Q \to \mathrm{Aut}(L/L^{D_Q})$ is surjective by `FixedPoints.toAlgAut_surjective`. For surjectivity of $\mathrm{Aut}(L/L^{D_Q}) \to \mathrm{Aut}(L/K)$, let σ be a field automorphism of L fixing K pointwise. We must show that σ automatically fixes L^{D_Q} pointwise. But this is exactly `fixed_of_fixed2`. Thus, the composition $D_Q \to \mathrm{Aut}(L/L^{D_Q}) \to \mathrm{Aut}(L/K)$ is surjective.

9 Miniproject: Adeles

9.1 Status

This is an active miniproject.

9.2 Goals

There are several goals to this miniproject.

1. Define the adeles $\mathbb{A}_K$ of a number field K and give them the structure of a K-algebra (status: now in mathlib thanks to Salvatore Mercuri).
2. Prove that $\mathbb{A}_K$ is a locally compact topological ring (status: also proved by Mercuri[1] but not yet in mathlib).
3. Base change: show that if L/K is a finite extension of number fields then the natural map $L \otimes_K \mathbb{A}_K \to \mathbb{A}_L$ is an isomorphism, both algebraic and topological; (status: not formalized yet, but there is a plan – see the project dashboard).
4. Prove that $K \subseteq \mathbb{A}_K$ is a discrete subgroup and the quotient is compact (status: not formalized yet, but there is a plan – see the project dashboard).
5. Get this stuff into mathlib (status: (1) done, (2)–(4) not done).

We briefly go through the basic definitions. Let K be a number field. Let $\widehat{\mathbb{Z}} = \text{proj}\lim_{N\geq 1}(\mathbb{Z}/N\mathbb{Z})$ be the profinite completion of $\mathbb{Z}$, equipped with the projective limit topology.

A cheap definition of the finite adeles $\mathbb{A}_K^\infty$ of K is $K \otimes_{\mathbb{Z}} \widehat{\mathbb{Z}}$, equipped with the $\widehat{\mathbb{Z}}$-module topology.

A cheap definition of the infinite adeles K_∞ is $K \otimes_{\mathbb{Q}} \mathbb{R}$ with the $\mathbb{R}$-module topology. This is a finite-dimensional $\mathbb{R}$-vector space, so this is just the usual topology on $\mathbb{R}^n$.

[1] `https://github.com/smmercuri/adele-ring_locally-compact`

A cheap definition of the adeles of K is $\mathbb{A}_K^\infty \times K_\infty$ with the product topology. This is a commutative topological ring.

However in the literature (and in mathlib) we see different definitions. The finite adeles of K are usually defined in the books as the so-called restricted product $\prod'_{\mathfrak{p}} K_{\mathfrak{p}}$ over the completions $K_{\mathfrak{p}}$ of K at all maximal ideals $\mathfrak{p} \subseteq \mathcal{O}_K$ of the integers of K. Here the restricted product is the subset of $\prod_{\mathfrak{p}} K_{\mathfrak{p}}$ consisting of elements which are in the integers $\mathcal{O}_{K,\mathfrak{p}}$ of $K_{\mathfrak{p}}$ for all but finitely many $\mathfrak{p}$. This is the definition given in mathlib. Mathlib also has the proof that they are a topological ring; furthermore the construction of the finite adeles in mathlib works for any Dedekind domain. The adeles are an arithmetic object, but the finite adeles are an algebraic object.

Similarly the infinite adeles of a number field K are usually defined as $\prod_v K_v$, the product running over the archimedean completions of K, and this is the mathlib definition.

The adeles of a number field K are the product of the finite and infinite adeles, and mathlib knows that they are a K-algebra and a topological ring.

9.3 Local compactness

As mentioned above, Salvatore Mercuri was the first to give a complete formalization of the proof that the adele ring is locally compact as a topological space. His work proved the result using the ad hoc topology on the adeles which we initially had. Since then, adeles have been refactored to have the direct limit topology and mathlib has `RestrictedProduct.locallyCompactSpace_of_addGroup`, the result that a restricted product of topological additive groups K_v over compact open subgroups A_v is locally compact.

What we need then is this (note that this is not true for a general Dedekind domain):

Theorem 9.3.1

> If K is a number field and v is a nonzero prime ideal of the integers of K, then the integers of K_v is a compact open subgroup.

Proof for Theorem 9.3.1

> Openness should follow from the fact that the integers are $\{x : v(x) < v(1/\pi)\}$ where π is a uniformizer. Compactness needs finiteness of the residue field $\mathcal{O}_K/v$.

Once we have this, the above result from mathlib gives us

Theorem 9.3.2

> The adeles of a number field are locally compact.

Proof for Theorem 9.3.2

The adeles of a number field are a product of the finite adeles and the infinite adeles so it suffices to prove that the finite and infinite adeles are locally compact. The infinite adeles are just isomorphic to $\mathbb{R}^n$ as a topological space, so they're certainly locally compact. As for the finite adeles, the mathlib theorem `RestrictedProduct.locallyCompactSpace_of_addGrou` says that a restricted product of locally compact additive groups with respect to open compact subgroups is locally compact, so this reduces us the previous result.

9.4 Base change

The "theorem" we want is that if L/K is a finite extension of number fields, then $\mathbb{A}_L = L \otimes_K \mathbb{A}_K$. This isn't a theorem though, this is actually a definition (the map between the two objects) and a theorem about the definition (that it's an isomorphism).

In fact the full claim is that it is both a homeomorphism and an L-algebra isomorphism. Before we can prove the theorem, we need to make the definition.

Recall that the adeles $\mathbb{A}_K$ of a number field is a product $\mathbb{A}_K^\infty \times K_\infty$ of the finite adeles and the infinite adeles. So our "theorem" follows immediately from the "theorems" that $\mathbb{A}_L^\infty = L \otimes_K \mathbb{A}_K^\infty$ and $L_\infty = L \otimes_K K_\infty$ (both of these equalities mean an algebraic and topological isomorphism). We may thus treat the finite and infinite results separately.

Base Change For Nonarchimedean Completions

As pointed out above, the theory of finite adeles works fine for Dedekind domains. So for the time being let `A` be a Dedekind domain. Recall that the height one spectrum of `A` is the nonzero prime ideals of `A`.

Note that because we stick to the literature, rather than to common sense, fields are Dedekind domains in mathlib, and the height one spectrum of a field is empty.

Lemma 9.4.1.1

If $i : K \to L$ denotes the inclusion then for $k \in K$ we have $e \times w(i(k)) = v(k)$, where e is the ramification index of w/v (recall that valuations here are written additively, unlike in mathlib).

Proof for Lemma 9.4.1.1

Standard (and formalized).

Definition 9.4.1.2

There's a natural ring map $K_v \to L_w$ extending the map $K \to L$. It is defined by completing the inclusion $K \to L$ at the finite

places v and w (which can be done because the previous lemma shows that the map is uniformly continuous for the v-adic and w-adic topologies).

Lemma 9.4.1.3

If $i_v : K_v \to L_w$ denotes the map of the previous definition then for $x \in K_v$ we have $e \times w(i(k)) = v(k)$, where e is the ramification index of w/v.

Proof for Lemma 9.4.1.3

Follows by continuity from lemma Lemma 9.4.1.1.

Lemma 9.4.1.4

The map $i_v : K_v \to L_w$ sends the integer ring A_v into B_w.

Proof for Lemma 9.4.1.4

The integer ring is defined by $v \geq 0$ (or $v \leq 1$ in mathlib, which uses multiplicative valuations) so the result follows from Lemma 9.4.1.3.

Theorem 9.4.1.5

Giving L_w the K_v-algebra structure coming from the natural map $K_v \to L_w$, the w-adic topology on L_w is the K_v-module topology.

Proof for Theorem 9.4.1.5

Any basis for L as a K-vector space spans L_w as a K_v-module, so L_w is finite-dimensional over K_v and the module topology is the same as the product topology. So we need to establish that the product topology on $L_w = K_v^n$ is the w-adic topology. But the w-adic topology is induced by the w-adic norm, which makes L_w into a normed K_v-vector space, and (after picking a basis) the product norm on $L_w = K_v^n$ also makes L_w into a normed K_v-vector space. So the result follows from the standard fact (see for example the lemma on p52 of Cassels-Froelich, formalized as `ContinuousLinearEquiv.ofFinrankEq` in mathlib) that any two norms on a finite-dimensional vector space over a complete field are equivalent (and thus induce the same topology).

we can view L_w as an $L \otimes_K K_v$-algebra.

Now instead of fixing w upstairs, we fix v downstairs and consider all w lying over it at once. So say v is in the height one spectrum of A.

Lemma 9.4.1.6

There are only finitely many primes w of B lying above v.

Proof for Lemma 9.4.1.6

This is a standard fact about Dedekind domains. The key input is mathlib's theorem `primesOver_finite`.

We write $w|v$ to denote the fact that w is a prime of B above v of A.

Definition 9.4.1.7

The product of the maps $K_v \to L_w$ for $w|v$ is a natural ring map $K_v \to \prod_{w|v} L_w$ lying over $K \to L$.

Because $K_v \to \prod_{w|v} L_w$ lies over $K \to L$, there's an induced L-algebra map $L \otimes_K K_v \to \prod_{w|v} L_w$. We are now able to state one of the key results in this section. The proof is probably the hardest proof in this section to formalize.

Theorem 9.4.1.8

The induced L-algebra homomorphism $L \otimes_K K_v \to \prod_{w|v} L_w$ is an isomorphism of rings.

Proof for Theorem 9.4.1.8

My current proposal to formalize this is as follows. The map is surjective because the image is dense and closed; this has been formalized already. It is also a K_v-algebra homomorphism if we give L_w the obvious K_v-algebra structure. Thus we can conclude the result if we can prove that both spaces are finite-dimensional and have the same dimension. The K_v-dimension of $L_K K_v$ is equal to the `K`-dimension of `L`, which is $\sum_{w|v} e_w f_w$ using the standard notation that e_w is the ramification index of `w` and f_w the residue degree. So it suffices to prove that $[L_w : K_v] = e_w f_w$. We already have that e_w (defined globally) is equal to the local ramification index (defined as the factor by which the valuations differ on `K`). So what is left is to prove that (i) the residue field extension induced by L_w/K_v has degree is equal to the globally-defined f_w, (ii) an extension of local fields has degree ef. Now (i) sounds straightforward given what we have (the map from A to $\mathcal{O}_v$ has kernel v and dense image) and (ii) is true for any complete discretely-valued field; I am not suggesting we formalize the following proof, but at least it represents a rigorous justification: A field complete with respect to a discrete valuation is stable in the sense of the book by Bosch-Guntzer-Remmert (Prop 3.6.2.1), so every finite extension of such a field is cartesian (def 3.6.1.1) and thus $ef = n$ (Prop 3.6.2.4, (iii) implies (ii)). Note that if you weaken the hypotheses too much then there are counterexamples; it's possible to have $ef < n$ and BGR goes into details.

Theorem 9.4.1.9

For v fixed, the product topology on $\prod_{w|v} L_w$ is the K_v-module topology.

Proof for Theorem 9.4.1.9

This is a finite product of K_v-modules each of which has the K_v-module topology by Theorem 9.4.1.5, and the product topology is the module topology for a finite product of modules each of which has the module topology.

Theorem 9.4.1.10

If we give $L \otimes_K K_v$ the K_v-module topology, then the L-algebra isomorphism $L \otimes_K K_v \cong \prod_{w|v} L_w$ is also a homeomorphism.

Proof for Theorem 9.4.1.10

Indeed, it is a K_v-algebra isomorphism between two modules each of which have the module topology, and any module map is automorphically continuous for the module topologies.

Theorem 9.4.1.11

The isomorphism $L \otimes_K K_v \to \prod_{w|v} L_w$ induces an isomorphism $B \otimes_A A_v \to \prod_{w|v} B_w$ for all v in the height one spectrum of A.

Proof for Theorem 9.4.1.11

Certainly the image of the integral elements are integral. The argument in the other direction is more delicate. My original plan was to follow Cassels–Froehlich, Cassels' article Global fields, section 12 lemma, p61, which proves it for all but finitely many primes, but a PR by Matthew Jasper gives another approach which works for all primes. Jasper's argument is to show that the closure of A in K_v is A_v for a valuation on a Dedekind domain, and then that the closure of A in $\prod_{v\in S} K_v$ is $\prod_{v\in S} A_v$ for S a finite set of valuations (using the Chinese remainder theorem). Applying this to B we get that the closure of B in $\prod_{w|v} L_w$ is $\prod_{w|v} B_w$. He then shows that this closure is the image of $B \otimes_A \mathcal{O}_v$ by showing that this image is closed because it's open), giving surjectivity; injectivity follows from the statement that $L \otimes_K K_v = \prod_{w|v} L_w$.

A summary of what we have so far: for all finite places v of A we have shown that the natural map $L \otimes_K K_v \to \prod_w L_w$ is an isomorphism of L-algebras, and that if $L \otimes_K K_v$ has the K_v-module topology and each L_w has the valuation topology then this map is also a homeomorphism. Furthermore we have shown that there is an induced algebraic isomorphism $B \otimes_A A_v \cong \prod_w B_w$ on the subrings of the left and right hand sides.

Recall that the finite adeles $\mathbb{A}_{A,K}^\infty$ is defined in mathlib to be the restricted product of the K_v with respect to the A_v, equipped with a certain restricted product topology (which is not the subspace topology of the product topology, indeed $\prod_v A_v$ is open in this topology). We have seen in definition Definition 9.4.1.2 that there's a map $K_v \to L_w$ if $w|v$, extending $K \to L$, and we have seen in theorem Lemma 9.4.1.4 that this sends A_v to B_w. We conclude

Definition 9.4.1.12

There's a natural ring homomorphism $\mathbb{A}_{A,K}^\infty \to \mathbb{A}_{B,L}^\infty$ lying over $K \to L$.

Hence there is a natural L-algebra homomorphism $L \otimes_K \mathbb{A}_{A,K}^\infty \to \mathbb{A}_{B,L}^\infty$.

We start with the following observation. If M is a K-module then there's a canonical map $B \otimes_A M \to L \otimes_K M$ sending $b \otimes m$ to $b \otimes m$ by the universal property of the tensor product. Our first goal is to show that this map is an isomorphism, so let us establish some lemmas first.

Lemma 9.4.1.13

If $0 \neq b \in B$ then there exists $0 \neq a \in A$ such that b divides the image of a in B.

Proof for Lemma 9.4.1.13

Let $a = N_{L/K}(b)$, the norm. This is known to take nonzero elements of L to nonzero elements of K because the norm is the determinant of an invertible linear map, and integral elements to integral elements. Furthermore $a/b \in L$ is the product of the conjugates of b in some normal closure of L, and hence it is integral, so it lies in B.

Corollary 9.4.1.14

The A-bilinear map $B \times K \to L$ sending (b, k) to bk is surjective.

Proof for Corollary 9.4.1.14

Given $\lambda \in L$, write it as n/d with $0 \neq d \in B$. Choose $0 \neq a \in A$ and $b \in B$ with $db = a$, and then note that $\lambda = nb \times a^{-1}$.

Corollary 9.4.1.15

The natural map $B \otimes_A K \to L$ is a B-algebra isomorphism.

Proof for Corollary 9.4.1.15

We write down an inverse. Regard $B \otimes_A K$ as a B-algebra via the action on the left. Note that at this point it's not even clear that $B \otimes_A K$ is a field. We have the structure map $B \to B \otimes_A K$ sending b to $b \otimes 1$, which is B-linear. I claim that every nonzero element of B gets sent to an invertible element of $B \otimes_A K$. Indeed, if $b \neq 0$ and, using the previous lemma, we choose $0 \neq a \in A$ such that $bb' = a$, then $(b \otimes 1)(b' \otimes \frac{1}{a}) = 1$. Thus by the universal property of localization, the B-linear map $B \to B \otimes_A K$ extends to a ring homomorphism from the field of fractions of B to $B \otimes_A K$, which we claim is our desired inverse. Checking that both composites are the identity should be straightforward. Starting with $B \otimes_A K$ we only have to check on elements of the form $b \otimes k$; starting with L we only have to check on elements of B. Hopefully both are straightforward.

Our next goal in this section is the following two results. First the algebraic claim.

Theorem 9.4.1.16

The natural map $L \otimes_K \mathbb{A}^\infty_{A,K} \to \mathbb{A}^\infty_{B,L}$ is an isomorphism.

Corollary 9.4.1.17

If M is any K-module then the canonical map $B \otimes_A M \to L \otimes_K M$ is an isomorphism.

Proof for Corollary 9.4.1.17

We can factor this map as $B \otimes_A M \cong B \otimes_A (K \otimes_K M) \cong (B \otimes_A K) \cong_K M \to L \otimes_K M$, and we just showed that the latter map was an isomorphism.

The reason we care about this is the following.

Theorem 9.4.1.18

B is a finitely presented A-module.

Proof for Theorem 9.4.1.18

A is Noetherian as it is a Dedekind domain, so it suffices to prove that B is finitely generated as an A-module. But this is in mathlib already: a proof is around line `101` of `BaseChange.lean` in FLT at the time of writing.

Theorem 9.4.1.19

If R is a commutative ring, if M is a finitely presented R-module and if N_i are a collection of R-modules, then the canonical map $M \otimes_R \prod_i N_i \to \prod_i (M \otimes_R N_i)$ is an isomorphism.

Proof for Theorem 9.4.1.19

If `M` is finite and free then Maddy Crim has already formalized this in FLT. For the general case present M as $R^a \to R^b \to M \to 0$ and use that tensor products and arbitrary products preserve surjections.

Definition 9.4.1.20

Let `V` and `W` be index sets, and let $f : W \to V$ be a map with finite fibres. Let X_v be sets, with subsets C_v, let Y_w be sets with subsets D_w, and say for all vV we're given a bijection $X_v \to \prod_{w|f(w)=v} Y_w$, identifying C_v with $\prod_{w:f(w)=v} D_w$. Then there's an induced bijection between the restricted products $\prod'_v(X_v, C_v)$ and $\prod'_w(Y_w, D_w)$.

Proof for Definition 9.4.1.20

Let `V` be the finite places of `K` and `W` the finite places of `L`, let X_v be $B \otimes_A K_v$, let C_v be $B \otimes_A A_v$, let Y_w be L_w, let D_w be B_w, and the result follows from the previous definition, given theorem Theorem 9.4.1.11.

Theorem 9.4.1.21

In the same setup as Definition 9.4.1.20 (V, W index sets, $f : W \to V$, $C_v X_v$ and $D_w Y_w$, bijections $b_v : X_v \to \prod_{w|f(w)=v} Y_w$ identifying C_v with $\prod_{w|f(w)=v} D_w$), if all the X_v and Y_w are furthermore topological spaces, all the C_v and D_w are open, and all the b_v are homeomorphisms, then the induced map $\prod'_v(X_v, C_v) \to \prod'_w(Y_w, D_w)$ is also a homeomorphism.

Proof for Theorem 9.4.1.21

I have only thought about the cofinite filter case, where this should follow easily from the definition of the topology.

Corollary 9.4.1.22

If `S` is a finite set of nonzero primes of `A` then the natural map $B \otimes ((\prod_{v \in S} K_v) \times (\prod_{v \notin S} A_v)) \to (\prod_{v \in S} (B \otimes_A K_v)) \times (\prod_{v \notin S} (B \otimes_A A_v))$ is an isomorphism.

Proof for Corollary 9.4.1.22

Follows from the previous two theorems.

Recall that $\mathbb{A}_K^\infty$ is the finite adeles of K, defined as the restricted product of the K_v with respect to the A_v, where v runs through the nonzero primes of A. Let R denote the restricted product of the $B \otimes_A K_v$ with respect to the $B \otimes_A A_v$.

Corollary 9.4.1.23

The natural map $B \otimes_A \mathbb{A}_K^\infty \to R$ is a B-algebra isomorphism.

Proof for Corollary 9.4.1.23

This follows from the previous corollary and the fact that tensor products commute with filtered colimits.

Corollary 9.4.1.24

The ring R introduced above, the restricted product of the $B \otimes_A K_v$ with respect to the $B \otimes_A A_v$, is isomorphic to $\mathbb{A}_L$.

Proof for Corollary 9.4.1.24

Let V be the finite places of K and W the finite places of L, let X_v be $B \otimes_A K_v$, let C_v be $B \otimes_A A_v$, let Y_w be L_w, let D_w be B_w, and the result follows from the previous definition, given theorem Theorem 9.4.1.11.

From this, we can deduce the theorem we claimed earlier:

Theorem 9.4.1.25

The natural map $B \otimes_A \mathbb{A}_K^\infty \to \mathbb{A}_L^\infty$ is a B-algebra isomorphism.

Proof for Theorem 9.4.1.25

This map factors through the auxiliary ring R so the result follows from the previous two constructions.

Because this map factors through the isomorphism $B \otimes_A \mathbb{A}_K^\infty \to L \otimes_K \mathbb{A}_K^\infty$ we can finally deduce that the natural map $L \otimes_K \mathbb{A}_K^\infty \to \mathbb{A}_L^\infty$ is an algebraic isomorphism.

Proof for Theorem 9.4.1.16

Follows immediately from theorem Theorem 9.4.1.25 and theorem Corollary 9.4.1.17.

Now $L \otimes_K \mathbb{A}_{A,K}^\infty$ is an $\mathbb{A}_{A,K}^\infty$-module and hence can be given the $\mathbb{A}_{A,K}^\infty$-module topology. We also claim

Theorem 9.4.1.26

The induced L-algebra morphism $L \otimes_K \mathbb{A}_{A,K}^{\infty} \to \mathbb{A}_{B,L}^{\infty}$ is a topological isomorphism.

Informally, the proofs are simple: componentwise we know that $L \otimes_K K_v$ is isomorphic both algebraically and topologically to $\prod_{w|v} L_w$, and that this isomorphism sends the open set $B \otimes_A A_v$ homeomorphically onto $\prod_{w|v} B_w$, so now it's "just a case of putting everything together". Formally, we really need to spell this out, as there is a lot going on. We do this in the next subsection.

Corollary 9.4.1.27

$\mathbb{A}_{\mathbb{L}}^{\infty}$ is homeomorphic to $\prod_v (B \otimes_A K_v, B \otimes_A A_v)$.

Proof for Corollary 9.4.1.27

Follows from the previous theorem with $X_v = B \otimes_A K_v$, $D_w = L_w$, etc.

Theorem 9.4.1.28

Let S be a set of infinite places of K. The image of K under the embedding $K \hookrightarrow (K_v)_{v \in S}$, $k \mapsto (k)_v$, is dense in the product topology.

Proof for Theorem 9.4.1.28

Let (K, v) denote K equipped with the topology induced by the infinite place v. It suffices to show that the image of K under the embedding $K \hookrightarrow \prod_{v|\infty}(K, v)$ is dense in the product topology. By a standard analytic argument, for each v it is possible to select a sequence $(x_n^{(v)})_n$ with the property that $x_n^{(v)} \to 1$ in v's topology, while $x_n^{(v)} \to 0$ in any other infinite place's topology. Let $(z_v)_v \in \prod_{v|\infty}(K, v)$. For each v, the sequence $y_n := \sum_{v|\infty} x_n^{(v)} z_v$ in K converges to z_v in v's topology, so the embedded image of y_n converges to $(z_v)_v$ in the product topology.

Theorem 9.4.1.29

For a fixed infinite place v of K, we have $\dim_{K_v} \prod_{w|v} L_w = \dim_{K_v} L \otimes_K K_v$.

Definition 9.4.1.30

Let v be an infinite place of K. There is a continuous K-algebra homomorphism $K_v \to \prod_{w|v} L_w$, whose restriction to K corresponds to the global embedding of K into $(L_w)_w$.

Definition 9.4.1.31

Let v be an infinite place of K. There is a natural L-algebra homomorphism $L \otimes_K K_v \to \prod_{w|v} L_w$, whose restriction to $1 \otimes_K K_v$ corresponds to the map in `NumberField.InfinitePlace.Completion.piExtension`.

A map in Definition 9.4.1.30 can be lifted to an L-algebra homomorphism defined on $L \otimes_K K_v$.

Theorem 9.4.1.32

For a fixed infinite place v of K, the local infinite-place base-change map $L \otimes_K K_v \to \prod_{w|v} L_w$ is surjective.

Proof for Theorem 9.4.1.32

Let $(x_i)_i$ be a K_v-basis of $\prod_{w|v} L_w$. By the density of L in $\prod_{w|v} L_w$ from Theorem 9.4.1.28, we can find $\alpha_i \in L$ arbitrarily close to x_i with respect to the sup norm when embedded globally in $\prod_{w|v} L_w$. In particular, it is possible to choose such α_i so that the matrix representing the vector $((\alpha_i)_{w|v})_i$ in the basis $(x_i)_i$ has nonzero determinant. Since $(\alpha_i)_{w|v}$ is the image of $1 \otimes \alpha_i$ under base change, we have that $(1 \otimes \alpha_i)_i$ also forms a basis of $L \otimes_K K_v$, and base change is surjective.

Theorem 9.4.1.33

For a fixed infinite place v of K, the local infinite-place base-change map $L \otimes_K K_v \to \prod_{w|v} L_w$ is injective.

Proof for Theorem 9.4.1.33

The L-algebra map $L \otimes_K K_v \to \prod_{w|v} L_w$ can equivalently be thought of as K_v-linear, which is injective, because it is surjective by theorem Theorem 9.4.1.32, and both sides have the same K_v-dimension by theorem Theorem 9.4.1.29.

Theorem 9.4.1.34

If $w \mid v$ is an infinite place of L lying above an infinite place v of K, then L_w has the K_v-module topology.

Proof for Theorem 9.4.1.34

Because L_w is a finite-dimensional normed K_v-vector space, there exists a K_v-linear homeomorphism $L_w \cong K_v^n$, from which L_w has the K_v-module topology.

Theorem 9.4.1.35

Let v be an infinite place of K. There is a natural L-algebra homeomorphism $L \otimes_K K_v \cong \prod_{w|v} L_w$, whose restriction to $1 \otimes_K K_v$ corresponds to the map in Definition 9.4.1.30.

Proof for Theorem 9.4.1.35

The map in Definition 9.4.1.31 is an L-algebra isomorphism by the surjectivity theorem Theorem 9.4.1.32 and the injectivity theorem Theorem 9.4.1.33. Every K_v-algebra isomorphism between two K_v-module topological spaces is a homeomorphism. Since the L-algebra isomorphism of Definition 9.4.1.31 can equivalently be viewed as a K_v-algebra isomorphism, it is also a homeomorphism.

Theorem 9.4.1.36

Let v be an infinite place of K. There is a natural K_v-linear homeomorphism $K_v^{[L:K]} \cong \prod_{w|v} L_w$.

Proof for Theorem 9.4.1.36

Compose the K_v-linear isomorphism $K_v^{[L:K]} \cong \prod_{w|v} L_w$ with theorem Theorem 9.4.1.35 to get the isomorphism in the statement. Since both sides have the K_v-module topology, then this is also a homeomorphism.

Theorem 9.4.1.37

There is a natural K_∞-linear homeomorphism $K_\infty^{[L:K]} \cong L_\infty$.

Proof for Theorem 9.4.1.37

Using the isomorphisms Theorem 9.4.1.36, we clearly have a bijection $K_\infty^{[L:K]} \cong \prod_v \prod_{w|v} L_w \cong \prod_w L_w$. The K_v-linearity of each component isomorphism extends to K_∞-linearity if the action of $\prod_v K_v$ on $\prod_w L_w$ is constant on the fibers of the restriction map on infinite places. In other words, if, for all $x \in K_\infty$ and $y \in L_\infty$, we have $(x \cdot y)_w = x_{v_w} \cdot y_w$, which is true by definition.

Theorem 9.4.1.38

L_∞ has the K_∞-module topology.

Proof for Theorem 9.4.1.38

Since L_∞ is homeomorphic to a finite product of K_∞ as a K_∞-vector space, it has the K_∞-module topology.

Theorem 9.4.1.39

There is a natural L-algebra isomorphism $L \otimes_K K_\infty \cong L_\infty$.

Proof for Theorem 9.4.1.39

This follows from the chain of isomorphisms $L \otimes_K K_\infty \cong \prod_v (L \otimes_K K_v) \cong \prod_v \prod_{w|v} L_w \cong L_\infty$. The first isomorphism is the standard L-algebra isomorphism $L \otimes_K \prod_v K_v \cong \prod_v (L \otimes_K K_v)$. The second isomorphism is given by the component L-algebra isomorphisms $L \otimes_K K_v \cong \prod_{w|v} L_w$ from Theorem 9.4.1.35.

Theorem 9.4.1.40

If `K -> L` is a ring homomorphism between two number fields, then there is a natural isomorphism, both topological and algebraic, $L \otimes_K K_\infty \cong L_\infty$.

Proof for Theorem 9.4.1.40

Since both sides of the L-algebra isomorphism in Theorem 9.4.1.39 have the K_∞-module topology, and since the isomorphism can equivalently be viewed as a K_∞-linear isomorphism, it is also a homeomorphism.

Base Change For Adeles

From the previous results we deduce immediately that if L/K is a finite extension of number fields then there's a natural (topological and algebraic) isomorphism $L \otimes_K \mathbb{A}_K \to \mathbb{A}_L$.

Theorem 9.4.2.1

If $K \to L$ is a ring homomorphism between two number fields then there is a natural isomorphism (both topological and algebraic) $L \otimes_K \mathbb{A}_K \cong \mathbb{A}_L$.

Proof for Theorem 9.4.2.1

Follows from the previous results.

Theorem 9.4.2.2

If $K \to L$ is a ring homomorphism between two number fields, then the topology on $\mathbb{A}_L$ is the $\mathbb{A}_K$-module topology, where the module structure comes from the natural map $\mathbb{A}_K \to \mathbb{A}_L$.

Proof for Theorem 9.4.2.2

Indeed $\mathbb{A}_L \cong L \otimes_K \mathbb{A}_K$ is a homeomorphism, and the right hand side has the $\mathbb{A}_K$-module topology.

9.5 Discreteness and compactness

We need that if K is a number field then $K \subseteq \mathbb{A}_K$ is discrete, and the quotient (with the quotient topology) is compact. Here is a proposed proof.

Theorem 9.5.1

There's an open subset of $\mathbb{A}_{\mathbb{Q}}$ whose intersection with $\mathbb{Q}$ is $\{0\}$.

Proof for Theorem 9.5.1

Use $\prod_p \mathbb{Z}_p \times (-1, 1)$. Any rational q in this set is a p-adic integer for all primes p and hence, writing it in lowest terms as $q = n/d$, satisfies $p \nmid d$, meaning that $d = \pm 1$ and thus $q \in \mathbb{Z}$. The fact that $q \in (-1, 1)$ implies $q = 0$.

Theorem 9.5.2

There's an open subset of $\mathbb{A}_K$ whose intersection with K is $\{0\}$.

Proof for Theorem 9.5.2

By a previous result, we have $\mathbb{A}_K = K \otimes_{\mathbb{Q}} \mathbb{A}_{\mathbb{Q}}$. Choose a basis of $K/\mathbb{Q}$; then K can be identified with $\mathbb{Q}^n \subseteq (\mathbb{A}_{\mathbb{Q}})^n$ and the result follows from the previous theorem.

Theorem 9.5.3

The additive subgroup K of $\mathbf{A}_K$ is discrete.

Proof for Theorem 9.5.3

If $x \in K$ and U is the open subset in the previous lemma, then it's easily checked that $K \cap U = \{0\}$ implies $K \cap (U + x) = \{x\}$, and $U + x$ is open.

Theorem 9.5.4

The quotient $\mathbb{A}_{\mathbb{Q}}/\mathbb{Q}$ is compact.

Proof for Theorem 9.5.4

The space $\prod_p \mathbb{Z}_p \times [0, 1] \subseteq \mathbb{A}_{\mathbb{Q}}$ is a product of compact spaces and is hence compact. I claim that it surjects onto $\mathbb{A}_{\mathbb{Q}}/\mathbb{Q}$. Indeed, if $a \in \mathbb{A}_{\mathbb{Q}}$ then for the finitely many prime numbers $p \in S$ such that $a_p \notin \mathbb{Z}_p$ we have $a_p \in \frac{r_p}{p^{n_p}} + \mathbb{Z}_p$ with $r_p/p^{n_p} \in \mathbb{Q}$, and if $q = \sum_{p \in S} \frac{r_p}{p^{n_p}} \in \mathbb{Q}$ then $a - q \in \prod_p \mathbb{Z}_p \times \mathbb{R}$. Now just subtract $\lfloor a_\infty - q \rfloor$ to move into $\prod_p \mathbb{Z}_p \times [0, 1)$, and we are done.

Theorem 9.5.5

The quotient $\mathbb{A}_K/K$ is compact.

Proof for Theorem 9.5.5

We proceed as in the discreteness proof above, by reducing to $\mathbb{Q}$. As before, choosing a $\mathbb{Q}$-basis of K gives us $\mathbb{A}_K/K \cong (\mathbb{A}_\mathbb{Q}/\mathbb{Q})^n$, so the result follows from the previous theorem.

10 Miniproject: Haar Characters

10.1 Goal

The goal of this miniproject is to develop the theory, that is, the basic API, of Haar characters. `Haar character` is a name I've made up to describe a certain character of the units of a locally compact topological ring. The main result we need here is that if `B` is a finite-dimensional algebra over a number field `K`, then $B^\times$ is in the kernel of the Haar character of $B \otimes_K \mathbb{A}_K$, where $\mathbb{A}_K$ is the ring of adeles of `K`. Most if not all of this should probably be in mathlib.

KMB would like to heartily thank Sebastien Gouezel for the help he gave during the preparation of this material.

10.2 Initial definitions

Scaling Haar Measure On A Group

Let A be a locally compact topological additive abelian group. There's then a regular additive Haar measure μ on A, unique up to a positive scalar factor. If $: (A, +)(A, +)$ is a homeomorphism and an additive automorphism of `A`, then we can push forward `μ` along `φ` to get a second measure `φ_* μ` on `A`, with the property that `(φ_* μ)(X) = μ(φ^{-1} X)` for any Borel subset `X` of `A`.

Definition 10.2.1.1

> If `A` is a locally compact topological additive abelian group, if μ is a regular Haar measure on `A`, and if $\phi : A \to A$ is an additive homeomorphism, then we let $d_A(\phi)$ denote the unique positive real number such that $\mu(X) = d_A(\phi)(\phi_*\mu)(X)$ for any Borel set `X`.

Now `φ_* μ` is a translation-invariant and regular measure, and hence also a Haar measure on `A`. It must thus differ from `μ` by a positive scalar factor, which we call $d_A()$.

There is a choice of normalization here between $d_A(\phi)$ and $d_A(\phi)^{-1}$, so let us be more precise.

To give an example, if `φ` is multiplication by 2 on the real numbers, if $X = [0, 1]$, and if `μ` is Lebesgue measure on the Borel subsets of $\mathbb{R}$, we have that $_*(X) = (^{-1}(X)) = ([0, 1/2]) = 1/2$, so $1 = d_A(\phi)/2$, meaning that $d_A(\phi) = 2$.

Similarly if φ is multiplication by -2 and $X = [0,1]$ then $^{-1}(X) = [-1/2, 0]$, which again has measure 1/2, so $d_A(\phi)$ is 2 again.

Strictly speaking our definition of $d_A()$ depends on the choice of regular Haar measure μ. Note that `mathlib` offers a fixed Borel regular Haar measure `MeasureTheory.Measure.haar` on any locally compact topological group, and the actual definition of d_A in the code uses this definition. Note also that the code defines everything for multiplicative groups and uses `@[to_additive]` to deduce the corresponding results for additive groups.

Lemma 10.2.1.2

$d_A(\phi)$ is independent of choice of regular Haar measure.

Proof for Lemma 10.2.1.2

If μ' is a second choice then μ' = λ μ for some positive real λ, and the λs on each side of $'(X) = d_A(\phi)(\phi_*')(X)$ cancel.

Lemma 10.2.1.3

If μ is any regular Haar measure on A then $d_A(\phi)(\phi_*\mu) = \mu$.

Proof for Lemma 10.2.1.3

This is a restatement of the previous result.

Corollary 10.2.1.4

If μ is any regular Haar measure on A then $d_A() = \phi^*$.

Proof for Corollary 10.2.1.4

This follows from the pushforward formula Lemma 10.2.1.3 applied to the regular Haar measure $\phi^*\mu$ and the fact that $\phi_*\phi^*\mu = \mu$.

Lemma 10.2.1.5

$d_A(id) = 1$.

Proof for Lemma 10.2.1.5

Clear.

Lemma 10.2.1.6

If X is a Borel set, then $(X) = d_A()(^1(X))$.

Proof for Lemma 10.2.1.6

This follows immediately from lemma `MeasureTheory.addEquivAddHaarChar_smul_map` and the definition of the pushforward of a measure.

Lemma 10.2.1.7

If $f : A\to$ is a Borel measurable function then $d_A()\int f(x)d(*)(x) = \int f(x)d(x)$.

Proof for Lemma 10.2.1.7

This also follows immediately from lemma.

We also have the following variant:

Lemma 10.2.1.8

If $f : A\to$ is a Borel measurable function then $d_A()\int f(x)d(x) = \int f(x)d(^*)(x)$.

Proof for Lemma 10.2.1.8

This is immediate from corollary Corollary 10.2.1.4.

Lemma 10.2.1.9

$d_A(\phi \circ \psi) = d_A(\phi)d_A(\psi)$.

Proof for Lemma 10.2.1.9

Here's one way: it suffices to prove that $d_A(\phi \circ \psi)(\phi \circ \psi)_*\mu = d_A(\phi)d_A(\psi)(\phi\circ\psi)_*\mu$ (because there exists a compact set with positive finite measure), and using lemma Lemma 10.2.1.3 and the fact that $(\phi \circ \psi)_*\mu = \phi_*(\psi_*\mu)$, one can simplify both sides to μ.

10.3 Examples

Lemma 10.3.1

If $R = \mathbb{R}$ then $\delta_R(u) = |u|$.

Proof for Lemma 10.3.1

Take μ to be Lebesgue measure and $X = [0, 1]$. We have $(u) = \mu(uX)$. If $u > 0$ then $u[0, 1] = [0, u]$, which has measure `u = |u|`, and if $u < 0$ then $u[0, 1] = [u, 0]$, which has measure `-u = |u|`.

Lemma 10.3.2

If $R = \mathbb{C}$ then $\delta_R(u) = |u|^2$.

Proof for Lemma 10.3.2

Multiplication by a positive real `r` sends a unit square to a square of area $r^2 = |r|^2$. Multiplication by $e^{i\theta}$ is a rotation and thus does not change area. The general case follows.

Scaling Haar Measure On A Ring

Now let R be a locally compact topological ring. The Haar character of R, or more precisely the left Haar character of R, is a group homomorphism $R^\times \to \mathbb{R}^\times$ defined in the following way. If $uR^\times$ then left multiplication by u, namely the map $\ell_u : (R,+) \to (R,+)$ defined by $\ell_u(r) = ur$, is a homeomorphism and an additive automorphism of (R,+), so the preceding theory applies to ℓ_u.

Definition 10.3.1.1

We define δ_R(u), or just δ(u) when the ring R is clear, to be $d_R(u)$.

Lemma 10.3.1.2

If $f : R \to \mathbb{R}$ is a Borel measurable function and $u \in R^\times$ then $\delta_R(u) \int f(ux)d\mu(x) = \int f(x)d\mu(x)$.

Proof for Lemma 10.3.1.2

A short calculation using lemma.

Lemma 10.3.1.3

If X is a Borel subset of R and $rR^\times$, then μ(rX) = δ_R(r) μ(X).

Proof for Lemma 10.3.1.3

Immediate from lemma.

Corollary 10.3.1.4

The function $_R : R^\times \to_{>0}$ is continuous.

Proof for Corollary 10.3.1.4

Fix a Haar measure μ on R and a continuous real-valued function f on R with compact support and such that $\int f(x)d\mu(x) \neq 0$. Then $r \mapsto \int f(rx)d\mu(x)$ is a continuous function from $R \to \mathbb{R}$, because a continuous function with compact support is uniformly continuous, and thus gives a continuous function $R^\times \to \mathbb{R}$. Thus the function $u \mapsto (\int f(ux)d\mu(x))/(\int f(x)d\mu(x))$ is a continuous function from $R^\times$ to $\mathbb{R}$ taking values in $\mathbb{R}_{>0}$. Hence δ_R^{-1} is continuous from lemma, and thus δ_R is too.

Lemma 10.3.1.5

If $R = \mathbb{Q}_p$ then $\delta_R(u) = |u|_p$, the usual p-adic norm.

Proof for Lemma 10.3.1.5

Normalise Haar measure so that $(\mathbb{Z}_p) = 1$. If u is a p-adic unit then $u\mathbb{Z}_p = \mathbb{Z}_p$ so multiplication by u didn't change Haar measure. If however $u = p$ then $u\mathbb{Z}_p$ has index p in $\mathbb{Z}_p$, and because $(i + p\mathbb{Z}_p) = (p\mathbb{Z}_p)$ we have that $(\mathbb{Z}_p) = p(p\mathbb{Z}_p)$ and thus $(p) = p^{-1}$. These elements generate $\mathbb{Q}_p^\times$, and two characters which agree on generators of a group must agree on the group.

If R is a finite extension of $\mathbb{Q}_p$ then $\delta_R(u)$ is the norm on R normalised in the following way: $\delta_R(\varpi) = q^{-1}$, where ϖ is a uniformiser and q is the size of the (finite) residue field.

10.4 Algebras

Say F is a locally compact topological ring, for example $\mathbb{R}$, $\mathbb{C}$, $\mathbb{Q}_p$, or the adeles of a number field, V is a finite free F-module, and $\varphi : V \to V$ is an invertible F-linear map. Then V with its module topology (which is the product topology if one picks a basis) is a locally compact topological abelian group, and φ is additive. One can check that linearity implies continuity, this is `IsModuleTopology.continuous_of_linearMap` in `mathlib`, so in fact φ is a homeomorphism and our theory applies. The following lemma gives a formula for the scale factor $d_V(\varphi)$.

Lemma 10.4.1

Assume that there's an F-basis for V such that ϕ is a product of elementary and diagonal matrices, note that this is automatic if F is a field and mathlib has this. Then $d_V(\phi) = \delta_F(\det(\phi))$, where $\det(\phi) \in F$ is the determinant of ϕ as an F-linear map.

Proof for Lemma 10.4.1

The proof is a generalization of `Real.map_matrix_volume_pi_eq_smul_volume_pi`, which crucially uses the induction principle `Matrix.diagonal_transvection_induction_of_det_ne` One checks it explicitly for diagonal matrices and for matrices which are the identity except that one off-diagonal entry is non-zero.Note: we assume that F is second countable, but it shouldn't be necessary.

Now say F is a locally compact topological field, and that R is a possibly noncommutative F-algebra. Recall that this means that F lies in the center of R, unless $R = 0$. Assume that R is finite-dimensional as an F-vector space. Then if we give R the F-module topology, which is just the product topology if we pick a basis, it is known that R becomes a topological ring. Now say $uR^\times$, and recall that $\ell_u : R \to R$ is left multiplication by u. Then ℓ_u is easily checked to be an F-linear homeomorphism.

Corollary 10.4.2

If $u \in R^\times$ then $\delta_R(u) = \delta_F(\det(\ell_u))$.

Proof for Corollary 10.4.2

Follows immediately from the preceding lemma.

10.5 Left and right multiplication

If R is a locally compact topological ring, and if multiplication on R is not commutative, then left and right multiplication by an element of R can scale Haar measure in different ways. For example if R is the upper-triangular 2×2 matrices with real entries, then left multiplication by $\begin{pmatrix} a & 0 \\ 0 & 1 \end{pmatrix}$ sends $\begin{pmatrix} x & y \\ 0 & z \end{pmatrix}$ to $\begin{pmatrix} ax & ay \\ 0 & z \end{pmatrix}$ and thus scales R's additive Haar measure by $|a|^2$, but right multiplication by $\begin{pmatrix} a & 0 \\ 0 & 1 \end{pmatrix}$ sends $\begin{pmatrix} x & y \\ 0 & z \end{pmatrix}$ to $\begin{pmatrix} ax & y \\ 0 & z \end{pmatrix}$ and thus scales R's additive Haar measure by a factor of $|a|$.

What's going on here is that if we regard left and right multiplication as $\mathbb{R}$-linear maps from R to R, then their associated matrices with respect to the obvious basis are `diag(a,a,1)` and `diag(a,1,1)`, which have different determinants.

However, if k is now any field and if B is a finite-dimensional central simple algebra over k, for example a quaternion algebra, the case we'll care about later, and if $u \in B^\times$ then $x \mapsto ux$ and $x \mapsto xu$ are both k-linear endomorphisms of B, and I claim that they have the same determinant.

Lemma 10.5.1

Say B is a finite-dimensional central simple algebra over a field k, and $uB^\times$. Let $\ell_u : B \to B$ be the k-linear map sending x to ux, and let $r_u : B \to B$ be the k-linear map sending x to xu. Then $det(\ell_u) = det(r_u)$.

Proof for Lemma 10.5.1

Determinants are unchanged by base extension, so without loss of generality k is algebraically closed. Then B is a matrix algebra, say $M_n(k)$. Now u can be thought of as a matrix with intrinsic determinant d, and B as a left B-module becomes a direct sum of n copies of the standard n-dimensional representation V of B. Thus $det(\ell_u) = d^n$. Similarly $det(r_u) = d^n$, so they are equal.

Corollary 10.5.2

If B is a central simple algebra over a locally compact field F, and if $u \in B^\times$, then $d_B(r_u) = \delta_B(u)$, recalling that the latter is defined to mean $d_B(\ell_u)$.

Proof for Corollary 10.5.2

If ℓ_u and r_u denote left and right multiplication by `u` on `B`, then we have seen in lemma `MeasureTheory.addEquivAddHaarChar_eq_ringHaarChar_det` that $d_B(r_u) = \delta_F(\det(r_u))$. Lemma `IsSimpleRing.mulLeft_det_eq_mulRight_det` tells us that this is $\delta_F(\det(\ell_u))$, and this is $\delta_B(u)$ again by corollary `MeasureTheory.addEquivAddHaarChar_eq_ringHaarChar_det`.

10.6 Finite products

Here are two facts which we will need about products.

Lemma 10.6.1

If `(A,+)` and `(B,+)` are locally compact topological abelian groups, and if $\phi : A \to A$ and $\psi : B \to B$ are additive homeomorphisms, then $\phi \times \psi : A \times B \to A \times B$ is an additive homeomorphism, this is obvious, and $d_{A\times B}(\phi \times \psi) = d_A(\phi)d_B(\psi)$.

Proof for Lemma 10.6.1

We only need this result in the case where both `A` and `B` are second-countable, in which case `Prod.borelSpace` can be used to show that Haar measure on `A × B` is the product of Haar measures on `A` and `B`, and in this case the result follows easily. Without this assumption, the product of these measures may not even be a Borel measure and one has to be more careful. The proof in this case is explained by Gou"ezel. Here is the idea. Let ρ be a Haar measure on $A \times B$. Fix sets $X \subseteq A$ and $Y \subseteq B$ which are compact with nonempty interior. We can now pull back ρ to a measure ν on the Borel sigma-algebra of A defined as $\nu(s) = \rho(s \times Y)$, and this is easily checked to be a Haar measure on A. Then

$$\delta_{A\times B}(a,0)\nu(X) = \delta_{A\times B}(a,0)\rho(X\times Y) = \rho((a,0)(X\times Y)) = \rho(aX\times Y) = \nu(aX) = \delta_A(a)\nu(X)$$

. So $\delta_{A\times B}(a,0) = \delta_A(a)$. Similarly $\delta_{A\times B}(0,b) = \delta_B(b)$, and because $\delta_{A\times B}$ is a group homomorphism we're home.

Lemma 10.6.2

If A_i are a finite collection of locally compact topological abelian groups, with $\phi_i : A_i \to A_i$ additive homeomorphisms, then $d_{\prod_i A_i}(\prod_i \phi_i) = \prod_i d_{A_i}(\phi_i)$.

Proof for Lemma 10.6.2

Induction on the size of the finite set, using the previous lemma.

Lemma 10.6.3

If R and S are locally compact topological rings, then $\delta_{R\times S}(r,s) = \delta_R(r)\times\delta_S(s)$.

Proof for Lemma 10.6.3

Follows immediately from lemma Lemma 10.6.1.

Lemma 10.6.4

If R_i are a finite collection of locally compact topological rings, and $u_i \in R_i^\times$, then $\delta_{\prod_i R_i}((u_i)_i) = \prod_i \delta_{R_i}(u_i)$.

Proof for Lemma 10.6.4

Follows immediately from lemma Lemma 10.6.2.

10.7 Some measure-theoretic preliminaries

Lemma 10.7.1

Let A and B be locally compact topological groups, and let $f : A\to B$ be both a group homomorphism and an open embedding. The pullback along f of a Haar measure on B is a Haar measure on A.

Proof for Lemma 10.7.1

Translation-invariance is easy, compact sets are finite because continuous image of compact is compact, open sets are bounded because image of open is open.

Lemma 10.7.2

The pullback of a regular Borel measure along an open embedding is a regular Borel measure.

Proof for Lemma 10.7.2

Again this is because the image of compact is compact and the image of open is open, so all the properties of being a regular measure are easily checked.

Lemma 10.7.3

Say A is a compact topological additive group and $\quad : A\to A$ is an additive isomorphism. Then $d_A() = 1$.

Proof for Lemma 10.7.3

We have $d_A(\phi)\mu(A) = \mu(A)$, and $\mu(A)$ is positive and finite because `A` is open and compact, using Lemma 10.2.1.6.

Lemma 10.7.4

If $f : A \to B$ is a group homomorphism and open embedding between locally compact topological additive groups and if $: A \to A$ and $: B \to B$ are additive homeomorphisms such that the square commutes (i.e., `f ∘ α = β ∘ f`), then $d_A() = d_B()$.

Proof for Lemma 10.7.4

Choose a regular Haar measure `μ_B` on `B`. The pullback `μ_A := f^* μ_B` is also a regular Haar measure by the first two lemmas just cited. Now fix a continuous compactly-supported function `g` on `A` with $0 < \int g(a)d\mu(a) < \infty$. Then $d_A(\alpha) \int g(a)d\mu_A(a) = \int g(a)d(\alpha^*\mu_A)(a)$ by lemma. This equals $\int g(a)d(\alpha^* f^* \mu_B)(a)$ by definition, which is $\int g(a)d(f^*\beta^*\mu_B)(a)$ because pullback of pullback is pullback. This equals $d_B(\beta) \int g(a)d(f^*\mu_B)(a)$ by corollary Corollary 10.2.1.4, which is $d_B(\beta) \int g(a)d\mu_A(a)$ by definition, and so $d_A(\alpha) = d_B(\beta)$ as required.

10.8 Restricted products

Now say $A = \prod_i' A_i$ is the restricted product of a collection of types A_i with respect to the subsets C_i. Recall that this is the subset of $\prod_i A_i$ consisting of those elements whose `i`th coordinate lies in C_i for all but finitely many `i`. Say $B = \prod_i' B_i$ is the restricted product of types B_i over the same index set, with respect to subsets D_i. Say $_i : A_i \to B_i$ are functions with the property that $_i(C_i)D_i$ for all but finitely many `i`. It is easily checked that the `φ_i` induce a function $:= \prod_i' {}_i : A \to B$. It is also easily checked that if all the A_i and B_i are groups or rings or `R`-modules, the C_i and D_i are subgroups or subrings or submodules, and the `φ_i` are group or ring or module homomorphisms, then `φ` is a group or ring or module homomorphism. However topological facts lie a little deeper.

Lemma 10.8.1

If the A_i and B_i are topological spaces and the `φ_i` are continuous functions, then the restricted product $= \prod_i' {}_i$ is a continuous function.

Proof for Lemma 10.8.1

We use the universal property `RestrictedProduct.continuous_dom` of the topology in mathlib to reduce to the claim that for all finite `S`, the induced map $A_S := \prod_{i \in S} A_i \times \prod_{i \notin S} C_i \to B$ is continuous. Because the inclusion $A_S \to A_T$ is continuous for $S \subseteq T$,

we are reduced to checking this claim for S sufficiently large that it contains all of the i for which $\phi(C_i) \neq D_i$. For such S, this map factors as $A_S \to B_S \to B$, and $B_S \to B$ is continuous, so it suffices to prove that $A_S \to B_S$ is continuous, but this is just a product of continuous maps.

We now focus on the case that $B_i = A_i$ are locally compact groups, $D_i = C_i$ are compact open subgroups, and $_i : A_i \to A_i$ are group isomorphisms and homeomorphisms sending C_i onto C_i for all but finitely many i. Then the restricted product $A := \prod' A_i$ of the A_i with respect to the C_i is also a locally compact topological group, and the restricted product $ = \prod' _i$ of the φ_i is a group isomorphism and homeomorphism, so we can ask how $d_A()$ compares to the $d_{A_i}(_i)$. First note that $d_{A_i}(_i) = 1$ for all the i such that $_i(C_i) = C_i$, as $d_{A_i}(_i)$ can be computed as $(_i(C_i))/(C_i)$ and (C_i) is guaranteed to have positive finite measure as it is open and compact. Hence the product $\prod_i d_{A_i}(_i)$ is a finite product, in the sense that all but finitely many terms are 1. The following theorem shows that the value of this product is d(φ).

Theorem 10.8.2

With A, A_i, C_i, ϕ_i, ϕ defined as above, we have $\delta_A(\phi) = \prod_i \delta_{A_i}(\phi_i)$.

In the Lean file we make the additional assumption that the index set over which we're taking the product is countable, and that the A_i are second countable. This is because in this proof we make use of an infinite product $\prod_{i \notin S} C_i$ of topological spaces equipped with Haar measure, so each topological space gets equipped with the Borel sigma algebra. mathlib then gives the product a canonical product sigma algebra which, in the case where the index set is uncountable, may not be the Borel sigma algebra. However we also want to use Haar measure on the product, so we need the canonical sigma algebra on this product to be the Borel sigma algebra. Rather than taking the trouble to locally turn off this product sigma algebra construction in FLT, we simply restrict to a countable index set and add second countability assumptions, because then the product is a Borel sigma algebra and this is all we need in our applications to adeles.

Proof for Theorem 10.8.2

Assume $_i(C_i) = C_i$ for all $i \notin S$, a finite set, and work in the open subgroup $U := \prod_{i \in S} A_i \times \prod_{i \notin S} C_i$. Then φ induces an automorphism ϕ_S of this open subgroup U of A, and in particular lemma tells us that $\delta(\phi) = \delta_U(\phi_S)$. Next note that $\phi_S : U \to U$ can be written as a product of the automorphisms $\prod_{i \notin S} \phi_i|_{C_i}$ of $\prod_{i \notin S} C_i$ and $\prod_{i \in S} \phi_i$ of $\prod_{i \in S} A_i$, so by lemma we have $\delta(\phi) = \delta(\prod_{i \notin S} \phi_i|_{C_i}) \times \delta(\prod_{i \in S} \phi_i)$. Because $\prod_{i \notin S} C_i$ is a compact group we must have $\delta(\phi_{i \notin S} \phi_i|_{C_i}) = 1$ by lemma. Finally $\delta(\prod_{i \in S} \phi_i) = \prod_{i \in S} \delta(\phi_i)$ by lemma and we are home.

Corollary 10.8.3

If $u = (u_i)_i \in R^\times$, then $\delta_R(u) = \prod_i \delta_{R_i}(u_i)$.

Proof for Corollary 10.8.3

By definition of restricted product we have $u_i \in C_i$ for all but finitely many `i`. Note also that `u` has an inverse $v = (v_i)_i$ with $v_i \in C_i$ for all but finitely many `i`. The fact that $u_i v_i = v_i u_i = 1$ means that $u_i, v_i \in C_i^\times$ for all but finitely many `i`. Thus the previous theorem `MeasureTheory.addEquivAddHaarChar_restrictedProductCongrRight` applies.

10.9 Adeles

We finish this miniproject by proving some results about Haar characters for algebras over adele rings. So let `K` be a number field and let $\mathbb{A}_K$ be the adeles of `K`. We will prove some theorems about $\mathbb{A}_K$-algebras `R` which are finite and free as $\mathbb{A}_K$-modules. Such algebras can be given the $\mathbb{A}_K$-module topology, and this makes them into locally compact topological rings. In fact we shall only be concerned in applications with algebras of the form $B \otimes_K \mathbb{A}_K$ where `B` is a finite-dimensional `K`-algebra. So fix such a `B`, and write $B_\mathbb{A}$ for $B \otimes_K \mathbb{A}_K$. Let us first deal with a subtlety. Recall that if $K \subseteq L$ are number fields, then $\mathbb{A}_L$ is a module-finite $\mathbb{A}_K$-algebra and hence an $\mathbb{A}_K$-module, and theorem Theorem 9.4.2.2 tells us that $\mathbb{A}_L$ has the $\mathbb{A}_K$-module topology. Thus the next lemma applies.

Lemma 10.9.1

Say `R` and `S` are topological rings, and `S` is an `R`-algebra, finite as an `R`-module. Assume that the topology on `S` is the `R`-module topology. Now say `M` is an `S`-module, and give it the induced `R`-module structure. Then the `R`-module topology and `S`-module topology on `M` coincide.

Proof for Lemma 10.9.1

Let $i : R \to S$ denote the structure map. First observe that `S` has the `R`-module topology, so the `R`-action map $R \times S \to S$, explicitly defined by $(r, s) \mapsto i(r)s$, is continuous, and restricting to $s = 1$ we deduce that `i` is continuous.Now let M_R and M_S denote `M` with the `R`-module and `S`-module topologies respectively. It suffices to prove that the identity maps $M_R \to M_S$ and $M_S \to M_R$ are continuous. Equivalently, because the `A`-module topology on an `A`-module is the finest topology making it into a topological module, we need to prove that M_R is a topological `S`-module and that M_S is a topological `R`-module. We start with the latter claim.First observe that M_S is a topological `S`-module, so addition is continuous. Next note that the map $R \times M_S \to M_S$ factors through $S \times M_S$ and

is hence the composite of two continuous maps and thus continuous. Hence M_S is a topological R-module.It thus remains to check that M_R is a topological S-module, or equivalently that the map $S \times M_R \to M_R$ is continuous. But this map is R-bilinear, and by the result `Module.continuous_bilinear_of_finite` in `mathlib`, any R-bilinear map between modules with the R-module topology is automatically continuous if one of the source modules is finitely-generated. The result applies because S is assumed to be a finite R-module and the proof is complete.

Corollary 10.9.2

If K is a number field and V is a K-module, then the natural isomorphism $V \otimes_K \mathbb{A}_K = V \otimes_{\mathbb{Q}} \mathbb{A}_{\mathbb{Q}}$ induced by the natural isomorphism $\mathbb{A}_K = K \otimes_K \mathbb{A}_{\mathbb{Q}}$ is a homeomorphism if the left-hand side has the $\mathbb{A}_K$-module topology and the right-hand side has the $\mathbb{A}_{\mathbb{Q}}$-module topology.

Proof for Corollary 10.9.2

Lemma `IsModuleTopology.continuous_bilinear_of_finite_left` tells us that $V \otimes_K \mathbb{A}_K$ has the $\mathbb{A}_{\mathbb{Q}}$-module topology, and it is easily checked that the isomorphism is $\mathbb{A}_{\mathbb{Q}}$-linear and hence automatically continuous.Note that in Lean this is proved for a general extension L/K rather than only for $K/\mathbb{Q}$.

Theorem 10.9.3

Let B be a finite-dimensional central simple K-algebra. Say $uB_{\mathbb{A}}^{\times}$, and define ℓ_u and $r_u : B_{\mathbb{A}} \to B_{\mathbb{A}}$ by $\ell_u(x) = ux$ and $r_u(x) = xu$. Then $d_{B_{\mathbb{A}}}(\ell_u) = d_{B_{\mathbb{A}}}(r_u)$.

Proof for Theorem 10.9.3

We think of $B_{\mathbb{A}}$ as $B \otimes_K \mathbb{A}_K$. If $u = (u_v)$ as v runs through the places of K, then $d_{B_{\mathbb{A}}}(\ell_u) = \prod_v d_{B_v}(\ell_{u_v})$ by `MeasureTheory.addEquivAddHaarChar_restrictedProductCong` and the product is finite. By the central-simple-algebra formula established earlier this equals $\prod_v d_{B_v}(r_{u_v})$, and again by `MeasureTheory.addEquivAddHaarChar_restri` this is $d_{B_{\mathbb{A}}}(r_u)$.

The previous theorem only applies to inner forms of matrix algebras, but the below theorem, a generalization of the adelic product formula, is valid for any finite-dimensional K-algebra. Before we state it let's remind ourselves of the product formula for $\mathbb{Q}$, and restate it in the language of these Haar characters.

Lemma 10.9.4

If $x \in \mathbb{A}_{\mathbb{Q}}^{\times}$, then $\delta_{\mathbb{A}_{\mathbb{Q}}}(x) = \prod_v |x_v|_v$.

Proof for Lemma 10.9.4

By theorem `MeasureTheory.addEquivAddHaarChar_prodCongr` we have $\delta_{\mathbb{A}_{\mathbb{Q}}}(x) = \delta_{\mathbb{A}_{\mathbb{Q}}^\infty}(x^\infty) \times \delta_{\mathbb{R}}(x_\infty)$. By lemma `MeasureTheory.ringHaarChar_real` we have $\delta_{\mathbb{R}}(x_\infty) = |x|_\infty$, and by theorem `MeasureTheory.addEquivAddHaarChar_restrictedProduct` we have $\delta_{\mathbb{A}_{\mathbb{Q}}^\infty} = \prod_p \delta_{\mathbb{Q}_p}(x_p)$. By lemma `MeasureTheory.ringHaarChar_padic` we have $\delta_{\mathbb{Q}_p}(x_p) = |x_p|_p$, and putting everything together we get the result.

Now $\mathbb{A}_{\mathbb{Q}}$ is a nonzero $\mathbb{Q}$-algebra and hence we have an inclusion $\mathbb{Q}^\times \to \mathbb{A}_{\mathbb{Q}}^\times$. Here is our reinterpretation of the product formula.

Lemma 10.9.5

If $x \in \mathbb{Q}^\times \subseteq \mathbb{A}_{\mathbb{Q}}^\times$, then $\delta_{\mathbb{A}_{\mathbb{Q}}}(x) = 1$.

Proof for Lemma 10.9.5

By lemma `MeasureTheory.ringHaarChar_adeles_rat` we have $\delta_{\mathbb{A}_{\mathbb{Q}}}(x) = \prod_v |x|_v$. But the product formula says that this is `1`. A quick proof: if $x = \pm\prod_p p^{e_p}$, then $\prod_p |x|_p = \prod_p p^{-e_p}$ and $|x|_\infty = \prod_p p^{e_p}$, so they cancel.

Next we generalize this to finite-dimensional $\mathbb{Q}$-vector spaces.

So say `V` is an `N`-dimensional $\mathbb{Q}$-vector space, and define $V_{\mathbb{A}} := V \otimes_{\mathbb{Q}} \mathbb{A}_{\mathbb{Q}}$ with its $\mathbb{A}_{\mathbb{Q}}$-module topology. If we choose an isomorphism $V \cong \mathbb{Q}^N$, then $V_{\mathbb{A}} \cong \mathbb{A}_{\mathbb{Q}}^N$ as an additive topological abelian group. In particular, $V_{\mathbb{A}}$ is locally compact.

Fix a $\mathbb{Q}$-linear automorphism $\phi : V \to V$. By base extension, `φ` induces an $\mathbb{A}_{\mathbb{Q}}$-linear automorphism $\phi_{\mathbb{A}}$ of $V_{\mathbb{A}}$ which is also a homeomorphism of $V_{\mathbb{A}}$ if $V_{\mathbb{A}}$ is given the module topology as an $\mathbb{A}_{\mathbb{Q}}$-module. Our goal is

Theorem 10.9.6

In the above situation, where `V` is a finite-dimensional $\mathbb{Q}$-vector space, $\phi : V \cong V$ is $\mathbb{Q}$-linear, and $\phi_{\mathbb{A}}$ is the base extension to $V_{\mathbb{A}} := V \otimes_{\mathbb{Q}} \mathbb{A}_{\mathbb{Q}}$ as a continuous linear endomorphism of $V_{\mathbb{A}}$ with the $\mathbb{A}_{\mathbb{Q}}$-module topology, we have $d_{V_{\mathbb{A}}}(\phi_{\mathbb{A}}) = 1$.

Proof for Theorem 10.9.6

The original blueprint proof of this was that $\phi_{\mathbb{A}} : V_{\mathbb{A}} \to V_{\mathbb{A}}$ could be written as a restricted product of $\phi_v : V \otimes_{\mathbb{Q}} \mathbb{Q}_v \to V \otimes_{\mathbb{Q}} \mathbb{Q}_v$, and hence by Theorem 10.8.2 we have $d_{V_{\mathbb{A}}}(\phi_{\mathbb{A}}) = \prod_p d_{V_p}(\phi_p) \times d_{V_\infty}(\phi_\infty)$, and then applying Lemma 10.4.1 this is equal to $\prod_v \delta_{\mathbb{Q}_v}(\det(\phi_v)) \prod_v \delta_{\mathbb{Q}_v}(\det(\phi)) = 1$.This turned out to be a nightmare to formalize, because commuting the tensor product and the restricted product cannot be done naively, as the tensor product is over $\mathbb{Q}$ and the submodules in the

restricted product defining $\mathbb{A}_{\mathbb{Q}}$ are not $\mathbb{Q}$-modules. So one has to choose a $\mathbb{Z}$-lattice Λ in V and use the isomorphisms

$$V\otimes_{\mathbb{Q}}\mathbb{A} = \Lambda\otimes_{\mathbb{Z}}\mathbb{A} = \Lambda\otimes_{\mathbb{Z}}\mathbb{Q}_\infty\times\prod_p{}'[\Lambda\otimes_{\mathbb{Z}}\mathbb{Q}_p;\Lambda\otimes_{\mathbb{Z}}\mathbb{Z}_p] = V\otimes_{\mathbb{Q}}\mathbb{Q}_\infty\times\prod_p{}'[V\otimes_{\mathbb{Q}}\mathbb{Q}_p;im(\Lambda\otimes_{\mathbb{Z}}\mathbb{Z}_p)]$$

and check that all of these canonical maps are continuous. One of these claims boils down to yet another claim of the form "if you do something to the factors and then take the restricted product, then this is topologically the same as doing it to the restricted product", with the thing being $\Lambda\otimes_{\mathbb{Z}}$ in this case. The needed homeomorphism is recorded in Corollary 10.9.2.So here is the proof which we actually formalized. Say an automorphism of a finite free R-module is nice if it is a product of transvections and diagonal matrices with unit entries. Mathlib has the fact that if R is a field then all automorphisms are nice, and the base change of a nice morphism is nice. Hence $\phi_{\mathbb{A}}$ is nice, and we can simply prove Lemma 10.4.1 for nice endomorphisms over a commutative ring, which gives the result immediately by the product formula.

Corollary 10.9.7

If B is a finite-dimensional $\mathbb{Q}$-algebra, for example a number field, or a quaternion algebra over a number field, if $B_{\mathbb{A}}$ denotes the ring $B\otimes_{\mathbb{Q}}\mathbb{A}_{\mathbb{Q}}$, and if $b\in B^\times$, then $\delta_{B_{\mathbb{A}}}(b)=1$.

Proof for Corollary 10.9.7

Follows immediately from the previous theorem.

Corollary 10.9.8

If B is a finite-dimensional $\mathbb{Q}$-algebra and if $b\in B^\times$, then right multiplication by b does not change Haar measure on $B_{\mathbb{A}}$.

Proof for Corollary 10.9.8

Follows immediately from the previous theorem.

11 Miniproject: Fujisaki's Lemma

In fact there is even a noncommutative version of this statement. In John Voight's book (Voight, 2021) this is Main Theorem `27.6.14(a)` and Voight calls it Fujisaki's lemma. I know nothing of the history but I'm happy to adopt this name. In the quaternion algebra miniproject we will use this compactness result to prove finite-dimensionality of a space of quaternionic modular forms.

11.1 Initial definitions

Let `K` be a field. A central simple `K`-algebra is a `K`-algebra `B` not necessarily commutative with centre `K` such that `B` has exactly two two-sided ideals, namely `{0}` and `B`. We will be concerned only with central simple `K`-algebras which are finite-dimensional as `K`-vector spaces, and when `K` is clear we will just refer to them as central simple algebras. We remark that a 4-dimensional central simple algebra is called a quaternion algebra; we will have more to say about these later on.

Matrix algebras $M_n(K)$ are examples of finite-dimensional central simple `K`-algebras. If `K = ℂ` or more generally if `K` is algebraically closed, then matrix algebras are the only finite-dimensional examples up to isomorphism.

There are other examples over the reals: for example Hamilton's quaternions `ℍ := ℝ ⊕ ℝ i ⊕ ℝ j ⊕ ℝ k` with the usual rules $i^2 = j^2 = k^2 = -1$, $ij = -ji = k$ and so on are an example of a central simple $\mathbb{R}$-algebra, and matrix algebras over $\mathbb{H}$ are other central simple $\mathbb{R}$-algebras. For a general field `K` one can make an analogue of Hamilton's quaternions `K ⊕ Ki ⊕ Kj ⊕ Kk` with the same multiplication rules to describe the multiplication, and if the characteristic of `K` is not 2 then this is a quaternion algebra which may or may not be isomorphic to $M_2(K)$.

Some central simple algebras B are division algebras, meaning that they are division rings, or equivalently that every nonzero $b \in B$ has a two-sided inverse. For example Hamilton's quaternions are a division algebra over $\mathbb{R}$, because $(x + yi + zj + tk)(x - yi - zj - tk) = x^2 + y^2 + z^2 + t^2$, so the inverse of a nonzero $x + yi + zj + tk$ is $(x - yi - zj - tk)/(x^2 + y^2 + z^2 + t^2)$.

However 2×2 matrices over a field K, whilst being a central simple algebra over K, are never a division algebra, because a nonzero matrix with

determinant zero such as $\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ has no inverse.

11.2 Enter the adeles

Let K be a number field and let D/K be a finite-dimensional central simple K-algebra (later on D will be a division algebra, hence the name, but we do not need this yet). Then $D_{\mathbb{A}} := D \otimes_K \mathbb{A}_K$ is an $\mathbb{A}_K$-algebra which is free of finite rank, and if we give $D_{\mathbb{A}}$ the $\mathbb{A}_K$-module topology then it is a topological ring. Furthermore $D_{\mathbb{A}}$ is free of finite rank over the locally compact topological ring $\mathbb{A}_K$ and is thus also locally compact. So by the theory of Haar characters (see the Haar characters miniproject (p. 61)) there is a canonical character $\delta_{D_{\mathbb{A}}} : D_{\mathbb{A}}^\times \to \mathbb{R}_{>0}$ measuring how left multiplication by an element of $D_{\mathbb{A}}^\times$ changes the additive Haar measure on $D_{\mathbb{A}}$. Let $D_{\mathbb{A}}^{(1)}$ denote the kernel of $\delta_{D_{\mathbb{A}}}$, and give it the subspace topology coming from $D_{\mathbb{A}}^\times$.

One can furthermore check that if R is a finite K-algebra then the $\mathbb{A}_K$-module topologies and $\mathbb{A}_{\mathbb{Q}}$-module topologies on $R_{\mathbb{A}}$ coincide. Indeed, the topology on $\mathbb{A}_K$ is the $\mathbb{A}_{\mathbb{Q}}$-module topology, as $\mathbb{A}_K = \mathbb{A}_{\mathbb{Q}} \otimes_{\mathbb{Q}} K$ as topological $\mathbb{A}_{\mathbb{Q}}$-algebras, where the right hand side has the $\mathbb{A}_{\mathbb{Q}}$-module topology by definition.

11.3 The proof

Lemma 11.3.1

> There's a compact subset E of D_A with the property that for all $x \in D_A^{(1)}$, the obvious map $xE \to D\backslash D_A$ is not injective.

Proof for Lemma 11.3.1

> We know that if we pick a $\mathbb{Q}$-basis for D of size d then this identifies D with $\mathbb{Q}^d$, $D_{\mathbb{A}}$ with $\mathbb{A}_{\mathbb{Q}}^d$, and $D\backslash D_{\mathbb{A}}$ with $(\mathbb{Q}\backslash\mathbb{A}_{\mathbb{Q}})^d$. Now $\mathbb{Q}$ is discrete in $\mathbb{A}_{\mathbb{Q}}$ by theorem Theorem 9.5.3, and the quotient $\mathbb{Q}\backslash\mathbb{A}_{\mathbb{Q}}$ is compact by theorem Theorem 9.5.4. Hence D is discrete in $D_{\mathbb{A}}$ and the quotient $D\backslash D_{\mathbb{A}}$ is compact.Fix a Haar measure μ on $D_{\mathbb{A}}$ and push it forward to $D\backslash D_{\mathbb{A}}$; by compactness this quotient has finite and positive measure, say $m \in \mathbb{R}_{>0}$. Choose any compact $E \subseteq D_{\mathbb{A}}$ with measure $> m$ (for example, choose a $\mathbb{Z}$-lattice $L \cong \mathbb{Z}^d$ in $D \cong \mathbb{Q}^d$, define $E_f := \prod_p L_p \in D \otimes_{\mathbb{Q}} \mathbb{A}_{\mathbb{Q}}^\infty$, and define $E_\infty \subseteq D \otimes_{\mathbb{Q}} \mathbb{R} \cong \mathbb{R}^n$ to be a huge closed ball, large enough to ensure the measure of $E := E_f \times E_\infty$ is bigger than m). Then $\mu(xE) = \mu(E) > m$ so the map can't be injective.

Definition 11.3.2

We let E denote any compact set satisfying the hypothesis of the previous lemma.

Definition 11.3.3

Define $X := E - E := \{e - f : e, f \in E\} \subseteq D_{\mathbb{A}}$.

Definition 11.3.4

Define $Y := X.X := \{xy : x, y \in X\} \subseteq D_{\mathbb{A}}$.

Lemma 11.3.5

X is a compact subset of $D_{\mathbb{A}}$.

Proof for Lemma 11.3.5

It's the continuous image of the compact set $E \times E$.

Lemma 11.3.6

Y is a compact subset of $D_{\mathbb{A}}$.

Proof for Lemma 11.3.6

It's the continuous image of the compact set $X \times X$.

Lemma 11.3.7

If $\beta \in D_{\mathbb{A}}^{(1)}$ then $\beta X \cap D^\times \neq \emptyset$.

Proof for Lemma 11.3.7

Indeed by the previous lemma Lemma 11.3.1, the map $\beta E \to D\backslash D_{\mathbb{A}}$ is not injective, so there are distinct $\beta e_1, \beta e_2 \in \beta E$ with $e_i \in E$ and $\beta e_1 - \beta e_2 = b \in D$. Now $b \neq 0$ and D is a division algebra, so $b \in D^\times$. And $e_1 - e_2 \in X$ so $b = \beta(e_1 - e_2) \in \beta X$, so we are done.

Lemma 11.3.8

Similarly, if $\beta \in D_{\mathbb{A}}^{(1)}$ then $X\beta^{-1} \cap D^\times \neq \emptyset$.

Proof for Lemma 11.3.8

Indeed, $\beta^{-1} \in D_{\mathbb{A}}^{(1)}$, and so left multiplication by β^{-1} does not change Haar measure on $D_{\mathbb{A}}$, so neither does right multiplication (by theorem `NumberField.AdeleRing.isCentralSimple_addHaarScalarFactor_left_mul_eq_righ`
So the same argument works: $E\beta^{-1} \to D\backslash D_{\mathbb{A}}$ is not injective so choose $e_1\beta^{-1} \neq e_2\beta^{-1}$ with difference $b \in D$ and then $(e_1 - e_2)\beta^{-1} \in D - \{0\} = D^\times$.

Definition 11.3.9

Let $T = Y \cap D^\times$.

Lemma 11.3.10

T is finite.

Proof for Lemma 11.3.10

It suffices to prove that $Y \cap D$ is finite. But $D \subseteq D_{\mathbb{A}}$ is a discrete additive subgroup, and hence closed. And $Y \subseteq D_{\mathbb{A}}$ is compact. So $D \cap Y$ is compact and discrete, so finite.

Definition 11.3.11

Define $C := (T^{-1}.X) \times X \subset D_{\mathbb{A}} \times D_{\mathbb{A}}$.

Lemma 11.3.12

C is compact.

Proof for Lemma 11.3.12

`X` is compact and `T` is finite.

Lemma 11.3.13

For every $\beta \in D_A^{(1)}$, there exist $b \in D^\times$ and $\nu \in D_A^{(1)}$ such that $\beta = b\nu$ and $(\nu, \nu^{-1}) \in C$.

Proof for Lemma 11.3.13

By lemma `NumberField.AdeleRing.DivisionAlgebra.Aux.X_meets_kernel`, $\beta X \cap D^\times \neq \emptyset$, and lemma `NumberField.AdeleRing.DivisionAlgebra.Aux.X_meets_kernel'`, $X\beta^{-1} \cap D^\times \neq \emptyset$, so we can write $\beta x_1 = b_1$ and $x_2\beta^{-1} = b_2$ with $b_i \in D^\times$ and $x_i \in X$. Note that $\beta \in D_{\mathbb{A}}^{(1)}$ and $b_i \in D^\times \subseteq D_{\mathbb{A}}^{(1)}$ by corollary Corollary 10.9.7, so $x_i \in D_{\mathbb{A}}^{(1)}$ as well. In particular $x_i \in D_{\mathbb{A}}^\times$ so x_1^{-1} makes sense.Multiplying the equations defining the x_i and b_i we deduce that $x_2x_1 = b_2b_1 \in Y \cap D^\times = T$ (recall that $Y = X.X$ and $T = Y \cap D^\times$ is finite); call this element t. Then $x_1^{-1} = t^{-1}x_2 \in T^{-1}.X$, and $x_1 \in X$, so if we set $\nu = x_1^{-1} \in D_{\mathbb{A}}^{(1)}$ and $b = b_1 \in D^\times$ then we have $\beta = b\nu$ and $(\nu, \nu^{-1}) \in C := (T^{-1}.X) \times X$. We are done!

Theorem 11.3.14

If D is a division algebra then the quotient $D^\times\backslash D_{\mathbb{A}}^{(1)}$ with its quotient topology coming from $D_{\mathbb{A}}^{(1)}$ is compact.

Proof for Theorem 11.3.14

Indeed, if M is the preimage of C under the inclusion $D_A^{(1)} \to D_A \times D_A$ sending ν to (ν, ν^{-1}), then M is a closed subspace of a compact space so it's compact (note that δ_{D_A} is continuous, by Corollary 10.3.1.4, so $D_A^{(1)}$ is a closed subset of $D_A^\times$ which is itself a closed subset of $D_A \times D_A$).Lemma 11.3.13 shows that M surjects onto $D^\times\backslash D_A^{(1)}$ which is thus also compact.

We note here some useful consequences.

Theorem 11.3.15

The quotient $D^\times\backslash(D \otimes_K \mathbf{A}_K^\infty)^\times$ is compact.

Proof for Theorem 11.3.15

There's a natural map α from $D^\times\backslash D_A^{(1)}$ to $D^\times\backslash(D \otimes_K \mathbf{A}_K^\infty)^\times$. We claim that it's surjective. Granted this claim, we are home, because if we put the quotient topology on $D^\times\backslash(D \otimes_K \mathbf{A}_K^\infty)^\times$ coming from $(D \otimes_K \mathbf{A}_K^\infty)^\times$ then it's readily verified that α is continuous, and the continuous image of a compact space is compact.As for surjectivity: say $x \in (D \otimes_K \mathbf{A}_K^\infty)^\times$. We need to extend x to an element $(x, y) \in (D \otimes_K \mathbf{A}_K^\infty)^\times \times (D \otimes_K K_\infty)^\times$ which is in the kernel of δ_{D_A}. Because $\delta_{D_A}(x, 1)$ is some positive real number, it will suffice to show that if r is any positive real number then we can find $y \in (D \otimes_K \mathbf{A}_K^\infty)^\times = (D \otimes_{\mathbf{Q}} \mathbf{R})^\times$ with $\delta_{D_A}(1, y) = r$, or equivalently, setting $D_{\mathbf{R}} = D \otimes_{\mathbf{Q}} \mathbf{R}$, that $\delta_{D_{\mathbf{R}}}(y) = r$. But $D \neq 0$ as it is a division algebra, and hence $\mathbf{Q} \subseteq D$, meaning $\mathbf{R} \subseteq D_{\mathbf{R}}$, and if $x \in \mathbf{R}^\times \subseteq D_{\mathbf{R}}^\times$ then $\delta(x) = |x|^d$ with $d = \dim_{\mathbf{Q}}(D)$, as multiplication by x is just scaling by a factor of x on $D_{\mathbf{R}} \cong \mathbf{R}^d$. In particular we can set $x = y^{1/d}$.

Theorem 11.3.16

If U is an open subgroup of $(D \otimes_K \mathbf{A}_K^\infty)^\times$, then the double coset space $D^\times\backslash(D \otimes_K \mathbf{A}_K^\infty)^\times/U$ is finite.

Proof for Theorem 11.3.16

The double cosets give a disjoint open cover of $(D \otimes_K \mathbb{A}_K^\infty)$ which descends to a disjoint open cover of the quotient space $D^\times\backslash(D \otimes_K \mathbb{A}_K^\infty)^\times$. However this space is compact by theorem Theorem 11.3.15.

12 Miniproject: Quaternion algebras

12.1 The goal

At the time of writing, `mathlib` still does not have a proof that the space of classical modular forms of a fixed weight, level and character is finite-dimensional. The main result in this miniproject is to prove that certain spaces of quaternionic modular forms are finite-dimensional. We need finiteness results like this in order to control the Hecke algebras which act on these spaces; these Hecke algebras are the `T`s which are isomorphic to the `R`s in the $R = T$ theorem which is the big first target for the FLT project.

12.2 Initial definitions

Our first goal in this miniproject is the definition of these spaces of quaternionic modular forms. We start with some preliminary work towards this.

Let K be a field. Recall that a quaternion algebra over K is a central simple K-algebra of K-dimension 4.

A fundamental fact about central simple algebras is that if D/K is a central simple K-algebra and L/K is an extension of fields, then $D \otimes_K L$ is a central simple L-algebra. In particular, if D is a quaternion algebra over K then $D \otimes_K L$ is a quaternion algebra over L. Some Imperial undergraduate students have established this fact in ongoing project work.

A **totally real field** is a number field F such that the image of every ring homomorphism $F \to \mathbb{C}$ is a subset of $\mathbb{R}$. We fix once and for all a totally real field F and a quaternion algebra D over F. We furthermore assume that D is **totally definite**, that is, that for all field embeddings $\tau : F \to \mathbb{R}$ we have $D \otimes_{F,\tau} \mathbb{R} \cong \mathbb{H}$. Because F has at least one real place, the totally definite hypothesis is enough to show that D is not a matrix algebra and thus must be a division algebra. Thus Fujisaki's theorem Theorem 11.3.15 from the Fujisaki miniproject applies, and we know that $D^\times \backslash (D \otimes_F \mathbb{A}_F)^{(1)}$ is compact.

The high-falutin' explanation of what is about to happen is that the units $D^\times$ of D can be regarded as a connected reductive algebraic group over F, and we are going to define certain spaces of automorphic forms for this algebraic

group. For a general connected reductive algebraic group, part of the definition of an automorphic form is that it is satisfies some differential equations (for example modular forms are automorphic forms for the algebraic group GL_2 over Q, and the definition of a modular form involves holomorphic functions, which are solutions to the Cauchy-Riemann equations). However under the assumption that F is totally real and D/F is totally definite, the "associated symmetric space is a 0-dimensional manifold", meaning in practice that the part of the definition of an automorphic form involving differential equations is vacuously satisfied in this setting.

As a consequence, the definitions we are about to give have a far more algebraic flavour. Crucially, the fact that we do not need to talk about differential equations at all means that one does not need to assume that our automorphic forms are $\mathbb{C}$-valued; our definition makes sense for forms valued in an arbitrary additive commutative group. In particular, it is possible to talk about mod p^n and p-adic automorphic forms in this setting without doing any complicated algebraic geometry.

12.3 Brief introduction to automorphic forms in this setting

Having made assumptions on D which makes the theory of automorphic forms over $D^\times$ far less technical, we will now make it a little more technical by using the modern adelic approach to the theory. Note that many results about the adeles of a number field are proved in the adele miniproject (p. 45). Our automorphic forms will be certain functions on the units of the ring $D_{\mathbb{A}^\infty} := D \otimes_F \mathbb{A}_F^\infty \cong D \otimes_{\mathbb{Q}} \mathbb{A}_{\mathbb{Q}}^\infty$.

To prove Fermat's Last Theorem it suffices to work in weight 2 and trivial character, and for simplicity we shall, at least for now, bake these assumptions into our definitions, even though they would be easy to remove. We remark again that there is no analogue of the holomorphicity condition that one sees in the classical theory, because the symmetric space is a finite set of points rather than the upper half plane. Also there is no analogue of the cuspidality condition because there are no cusps in this setting. Other than the number field F and the quaternion algebra D, the other variable we will see in our situation will be the level of the forms. The main result in this miniproject will be that the space of weight 2 automorphic forms of level U is finite-dimensional.

12.4 Definition of spaces of automorphic forms

Let us now give some precise definitions. Recall that by $\mathbb{A}_F^\infty$ we mean the finite adeles of the totally real number field F.

A level is a compact open subgroup U of $(D \otimes_F \mathbb{A}_F^\infty)^\times$. These are plentiful. The ring $D_f := D \otimes_F \mathbb{A}_F^\infty$ is a topological ring, and hence the units $D_f^\times$ of this ring are a topological group. This group is locally profinite, and hence has many compact open subgroups; we will see explicit examples later on.

We regard $\mathbb{A}_F^\infty$ as a subring of $D_{\mathbb{A}^\infty} := D \otimes_F \mathbb{A}_F^\infty$, which is possible because F is a subring of D. More precisely we embed $\mathbb{A}_F^\infty$ into $D \otimes_F \mathbb{A}_F^\infty$ via the map sending g to $1 \otimes g$. Because F is in the centre of D, we have that $\mathbb{A}_F^\infty$ is in the centre of $D_{\mathbb{A}^\infty}$ (in fact it is the centre, but we do not need this). As a consequence we can identify $(\mathbb{A}_F^\infty)^\times$ as a subgroup of $(D \otimes_F \mathbb{A}_F^\infty)^\times$. We may also regard D as a subring of $D \otimes_F \mathbb{A}_F^\infty$ via the map $d \mapsto d \otimes 1$, and hence we can think of $D^\times$ as a subgroup of $(D \otimes_F \mathbb{A}_F^\infty)^\times$.

Let R be an additive commutative group. Later on R will be a commutative ring but we will not need this for the definition.

Definition 12.4.1

The space of R-valued **automorphic forms** for $D^\times$ is the set of functions $f : D_{\mathbb{A}^\infty}^\times \to R$ satisfying the following axioms:

- $f(dg) = f(g)$ for all $d \in D^\times$ and $g \in D_{\mathbb{A}^\infty}^\times$.
- $f(gz) = f(g)$ for all $g \in D_{\mathbb{A}^\infty}^\times$.
- There exists a compact open subgroup $U \subseteq (D_{\mathbb{A}^f}^\times)$ such that $f(gu) = f(g)$ for all $g \in D_{\mathbb{A}^\infty}^\times$ and $u \in U$.

Let $S^D(R)$ denote the set of automorphic forms for $D^\times$. The space $S^D(R)$ is sometimes referred to as a space of "quaternionic modular forms" over R. Three basic observations about $S^D(R)$ are

Definition 12.4.2

Pointwise addition $(f_1 + f_2)(g) := f_1(g) + f_2(g)$ makes $S^D(R)$ into an additive abelian group. This depends on Definition 12.4.1.

Definition 12.4.3

If R is a commutative ring then pointwise scalar multiplication $(r \cdot f)(g) := r \cdot f(g)$ makes $S^D(R)$ into an R-module. This depends on Definition 12.4.1 and Definition 12.4.2.

Definition 12.4.4

The group $D_{\mathbb{A}^\infty}^\times$ acts on the additive abelian group $S^D(R)$ by $(g \cdot f)(x) = f(xg)$. This depends on Definition 12.4.1, Definition 12.4.2, and Definition 12.4.3.

If R is a commutative ring then the action of $D_{\mathbb{A}^\infty}^\times$ commutes with the R-action.

Now let U be a level, that is, a compact open subgroup of $D_{\mathbb{A}^\infty}^\times$.

Definition 12.4.5

The quaternionic modular forms of level U, with notation $S^D(U; R)$, are the U-invariants for the $D_{\mathbb{A}^\infty}^\times$-action on $S^D(R)$. This depends on Definition 12.4.1, Definition 12.4.2, Definition 12.4.4, and Definition 12.4.3.

The Hecke algebras involved in the main modularity lifting theorems needed in the FLT project will be endomorphisms of these spaces $S^D(U;R)$.

12.5 The main result

The point of this miniproject is the finite-dimensionality result below. This is an analogue of the result that classical modular forms of a fixed level, weight and character are finite-dimensional. In fact, by delicate results of Jacquet and Langlands this result, in the case $k=\mathbb{C}$, implies many cases of that classical claim, although of course the Jacquet–Langlands theorem is much harder to prove than the classical proof of finite-dimensionality.

Theorem 12.5.1

Let k be a field. Then the space $S^D(U;k)$ is a finite-dimensional k-vector space.

Proof for Theorem 12.5.1

The finite-dimensionality theorem is in fact an easy consequence of Fujisaki's lemma, proved in the Fujisaki miniproject (p. 75) chapter. Write $(D\otimes_F\mathbb{A}_F^\infty)^\times$ as a disjoint union of double cosets $\coprod_i D^\times g_i U$. This open cover descends to a disjoint open cover of $D^\times\backslash(D\otimes_F\mathbb{A}_F^\infty)^\times$, and this latter space is compact by theorem Theorem 11.3.15. Hence the cover is finite; write the double coset representatives as $g_1,g_2,\ldots,g_n$.We claim that the function $S^D(U;k)\to W^n$ sending f to $(f(g_1),f(g_2),\ldots,f(g_n))$ is injective and k-linear, which suffices by finite-dimensionality of W. k-linearity is easy, so let us talk about injectivity.Say f_1 and f_2 are two elements of $S^D(U;k)$ which agree on each g_i. It suffices to prove that $f_1(g)=f_2(g)$ for all $g\in(D\otimes_F\mathbb{A}_F^\infty)^\times$. So say $g\in(D\otimes_F\mathbb{A}_F^\infty)^\times$, and write $g=\delta g_i u$ for $\delta\in D^\times$ and $u\in U$. Then $f_1(g)=f_1(\delta g_i u)=f_1(g_i)$ by the definition of $S^D(U;k)$, and similarly $f_2(g)=f_2(g_i)$. Because $f_1(g_i)=f_2(g_i)$ by assumption, we deduce that $f_1(g)=f_2(g)$ as required.

13 Miniproject: Hecke Operators

13.1 Status

This is an active miniproject. The abstract theory is completely formalized; at the time of writing the concrete theory has no sorried definitions but it does have some sorried proofs.

13.2 Goal

The goal of this project is to get sorry-free definitions of Hecke operators acting on spaces of automorphic forms. These Hecke operators generate Hecke algebras, which are the rings called T in the modularity lifting theorems, or $R = T$ theorems, crucially introduced by Wiles in order to prove FLT.

The theory comes in two parts; the "abstract" theory, which is pure algebra, and the "concrete" theory where we apply the abstract constructions to produce endomorphisms of spaces of automorphic forms. The abstract theory is short (and completely formalized); the concrete theory still needs some work because to apply the theory to the adelic groups we care about we need to develop some more API around the theory of restricted products, and of compact open subgroups of matrix groups.

13.3 Abstract theory

The set-up: we have a commutative ring R, the coefficient ring, and all of our spaces which the operators act on will be R-modules.

We have a group G acting R-linearly on an R-module A.

We have subgroups U and V of G. We will be particularly interested in the R-modules A^U and A^V of invariant elements.

Given an element $g \in G$, then under a certain finiteness hypothesis we will be able to define an R-linear map T_g or $[UgV]$ from A^V to A^U. The finiteness hypothesis is that the double coset UgV can be written as a finite union of single cosets g_iV.

Definition 13.3.1

Assuming UgV is a finite union of cosets g_iV, we define $[UgV] : A^V \to A^U$ to be the map sending $a \in A^V$ to $\sum_i g_i a$.

Lemma 13.3.2

This function is well-defined (that is, independent of the choice of g_i), has image in A^U and is R-linear.

Proof for Lemma 13.3.2

Well-definedness is because if we change g_i to $g_i' := g_i v$ for some $v \in V$ then $g_i a = g_i' a$ because $a \in A^V$.The image lands in A^U because left multiplication by u fixes UgV and hence permutes the cosets g_iV.Finally R-linearity is because the G-action is R-linear.

The finiteness hypothesis that the decomposition $UgV = \coprod_i g_iV$ is into a finite union is necessary for the theory to work. If G is a topological group then here is a criterion which gives the finiteness hypothesis for free.

Lemma 13.3.3

If U and V are compact subgroups of a topological group G, if V is also open, and if $g \in G$, then the double coset space UgV is a finite union of left cosets g_iV.

Proof for Lemma 13.3.3

The subset UgV of G is a continuous image of the compact set $U \times V$ and is hence compact, and it is covered by the disjoint left cosets g_iV; this cover must thus be finite.

Lemma 13.3.4

Say $g, h \in G$ and suppose we have $UgU = \coprod_i g_iU$ and $UhU = \coprod_j h_j$ and that $g_ih_j = h_jg_i$ for all i, j. Then $[UgU][UhU] = [UhU][UgU]$, that is, the Hecke operators acting on A^U commute.

Proof for Lemma 13.3.4

We have $[UgU][UhU]a = \sum_i g_i(\sum_j h_j a) = \sum_{i,j} g_i h_j a$ and $[UhU][UgU]a = \sum_j h_j \sum_i g_i a = \sum_{j,i} h_j g_i a$ and these sums are equal because $g_ih_j = h_jg_i$.

13.4 Restricted products

In the concrete example of Hecke operators which we care about, the invariants A^G will be spaces of quaternionic automorphic forms by definition. We do not need to worry about the definition of A in this project at all. However we will need to do various computations with the specific groups G which we are interested in, and they are restricted products. So we now develop some theory for restricted products, starting by recalling the definition.

If I is an index set, if X_i are sets indexed by $i \in I$ and if Y_i are subsets, then the restricted product $\prod_i' X_i$ is defined to be the subset of the full product $\prod_i X_i$ consisting of those tuples (x_i) such that $x_i \in Y_i$ for all but finitely many i. We suppress the Y_i from the notation in this document, although in Lean we cannot do this and the restricted product looks something like $\prod^r i, [X\ i, Y\ i]$.

It is straightforward to check that if the X_i are groups or rings or R-modules, and the Y_i are subgroups or subrings or submodules, then the restricted product is a group, ring or R-module. Indeed the structure is inherited via the inclusion $\prod_i' X_i \subseteq \prod_i X_i$ and the fact that arbitrary products of groups, rings, and modules are groups, rings, and modules.

More subtle is the theory of topological space structures. If the X_i are topological spaces then we do not give $\prod_i' X_i$ the subspace topology coming from the product topology on $\prod_i X_i$; instead we give it the finest topology making all of the natural maps $\prod_{i\in S} X_i \times \prod_{i\notin S} Y_i \to \prod_i' X_i$ continuous, as S runs through all finite subsets of I. Here the product of X_is and Y_is has the product topology. For example if all of the Y_i are open then one can check that $\prod_i Y_i$ is an open subset of $\prod_i' X_i$, whereas it is not of the form $\left(\prod_i' X_i\right) \cap U$ for any open subset U of $\prod_i X_i$ in general; the map from $\prod_i' X_i$ to $\prod_i X_i$ is continuous but is not in general an embedding.

If you have seen automorphic forms before, then here is an obvious-sounding claim: because the adeles $\mathbb{A}_F$ of a number field are a restricted product of completions F_v with respect to the integer rings $\mathcal{O}_v$, then $GL_2(\mathbb{A}_F)$ is obviously topologically a restricted product of $GL_2(F_v)$ with respect to $GL_2(\mathcal{O}_v)$. We now spend some time justifying this claim, which is a little more intricate than it sounds.

Lemma 13.4.1

> If A_i is a family of topological spaces equipped with open subsets B_i, and C_i is a family of topological spaces equipped with open subsets D_i, and $A_i \times C_i$ is equipped with the open subset $B_i \times D_i$, then the natural bijection $\prod_i'(A_i \times C_i) = (\prod_i' A_i) \times (\prod_i' B_i)$ is a homeomorphism.

This may well not be true if B_i and D_i are not open, because filtered colimits and binary products do not appear in general to commute in the category of topological spaces. I don't know an explicit counterexample though.

Proof for Lemma 13.4.1

We need to check continuity in both directions. The easy way is continuity of the map from the restricted product to the map from the binary product; the lemma `RestrictedProduct.continuous_dom` in mathlib tells us that a map from a restricted product is continuous when its restriction to $\left(\prod_{i\in S}(A_i \times C_i)\right) \times \left(\prod_{i\notin S}(B_i \times D_i)\right)$ is continuous for all finite $S \subseteq I$; the universal property of the binary product tells us that the map into the binary product is continuous iff the maps into the factors are continuous, but the map into $\prod_i' X_i$ is a product of the natural maps from $\left(\prod_{i\in S}(A_i \times C_i)\right) \times \left(\prod_{i\notin S}(B_i \times D_i)\right)$ to $\left(\prod_{i\in S} A_i\right) \times \left(\prod_{i\notin S} B_i\right)$ and the inclusion, and both are known to be continuous (an arbitrary product of continuous maps is continuous, and the other claim is in the restricted product API in mathlib).The harder direction is the other way, because we are working against both universal properties. The trick is `RestrictedProduct.continuous_dom_prod` in mathlib (this is where we assume B_i and D_i are open), which tells us that a map out of a binary product of restricted products is continuous when its restriction to $\left(\left(\prod_{i\in S} A_i\right) \times \left(\prod_{i\notin S} B_i\right)\right) \times \left(\left(\prod_{i\in S} C_i\right) \times \left(\prod_{i\notin S} D_i\right)\right)$ is, for all finite S (note that the S in the mathlib lemma is actually our $I - S$). The map from this to the restricted product factors through $\left(\prod_{i\in S}(A_i \times C_i)\right) \times \left(\prod_{i\notin S}(B_i \times D_i)\right)$; the first map is a homeomorphism (use the fact that $\prod_i X_i \times Y_i$ is homeomorphic to $(\prod_i X_i) \times (\prod_i Y_i)$), and the second is continuous by definition of the topology on a restricted product.

Corollary 13.4.2

Restricted products with respect to open subspaces commute with finite products. In other words, if j runs through a finite set J and i runs through an arbitrary set I, and if X_{ji} are topological spaces equipped with open subspaces Y_{ji}, then the obvious bijection $\prod_i'(\prod_j X_{ji}) = \prod_j(\prod_i' X_{ji})$ is a homeomorphism.

Proof for Corollary 13.4.2

Induction on the size of the finite set, using Lemma 13.4.1 to get started.

Corollary 13.4.3

If X_i are topological spaces and the Y_i are open subspaces, then the obvious map $M_n(\prod_i' X_i) = \prod_i' M_n(X_i)$ is a homeomorphism.

Proof for Corollary 13.4.3

Immediate from the previous corollary Corollary 13.4.2.

Lemma 13.4.4

If M is a topological monoid and U is an open submonoid, then the units $U^\times$ of U are naturally an open subgroup of $M^\times$.

$M^\times$ does not get the subspace topology from `M`; it is embedded into $M \times M$ via $g \mapsto (g, g^{-1})$ and gets the subspace topology from the product. This makes it into a topological group.

Proof for Lemma 13.4.4

We have $U \times U$ is an open subset of $M \times M$, and if we view $M^\times$ embedded in $M \times M$ via $g \mapsto (g, g^{-1})$, then the intersection of this subgroup with $U \times U$ is open in $M^\times$ and consists of the elements of $M^\times$ which are in U and whose inverse is also in U, which is easily checked to be the copy of $U^\times$ we're talking about.

Lemma 13.4.5

If M is a Hausdorff topological monoid and U is a compact submonoid, then the units $U^\times$ of U are naturally a compact subgroup of $M^\times$.

Is Hausdorffness necessary?

Proof for Lemma 13.4.5

First I claim that $M^\times$ embedded in $M \times M$ via $g \mapsto (g, g^{-1})$ is a closed subset of $M \times M$. Indeed, if $p : M \times M \to M$ is $(a, b) \mapsto ab$ and $q : M \times M \to M$ is $(a, b) \mapsto ba$, then p and q are continuous, $M^\times \subseteq M \times M$ is the intersection $p^{-1}\{1\} \cap q^{-1}\{1\}$, and $\{1\}$ is closed because `M` is Hausdorff.We have $U \times U$ is a compact subset of $M \times M$, and so $U^\times = M^\times \cap U \times U$ is a closed subspace of a compact space and is thus compact.

Lemma 13.4.6

If U_i are topological monoids then the canonical group isomorphism $(\prod_i U_i)^\times = \prod_i (U_i^\times)$ is a homeomorphism.

Proof for Lemma 13.4.6

We prove that the maps in both directions are continuous. Let's start with the map from left to right.A map into a product is continuous when the maps to the factors are continuous. A map into the units of a monoid is continuous when the two projection maps to the monoid (the inclusion and the map $u \mapsto u^{-1}$) are continuous (because $M^\times$ has the topology induced from $M \times M$). This reduces us to checking that the maps $(\prod_i U_i)^\times \to U_j$ sending (u_i)

to u_j resp u_j^{-1} are continuous. But the former map is the continuous inclusion $(\prod_i U_i)^\times \to \prod_i U_i$ followed by the continuous projection to U_j, and the latter map is the continuous inclusion $(\prod_i U_i)^\times \to \prod_i U_i$ sending x to x^{-1} followed by the projection.To go the other way: because the units have the induced topology it suffices to check that the two maps $\prod_i (U_i^\times) \to \prod_i U_i$ sending (u_i) to (u_i) resp (u_i^{-1}) are continuous. A map to a product is continuous when the induced maps to the factors are. A projection from a project is continuous, and the identity and inverse are continuous maps $U_j^\times \to U_j$, and the maps we're concerned with are composites of these maps and are hence continuous.

Theorem 13.4.7

If M_i are a family of topological monoids equipped with open submonoids U_i, then the canonical map $(\prod_i' M_i)^\times \to \prod_i' (M_i^\times)$ is a homeomorphism.

Proof for Theorem 13.4.7

I don't know a clean way of showing that the map from left to right is continuous, so here is a direct proof that the map is a homeomorphism. It is certainly an abstract group isomorphism between topological groups. So to prove that it is a homeomorphism it suffices to prove that it is a homeomorphism near the identity, or equivalently that there are open neighbourhoods `X` and `Y` of the identity elements on each side such that the map induces a homeomorphism from `X` to `Y`. We choose the units of $\prod_i U_i$ and the product $\prod_i (U_i)$. Note that the former is open because of `Submonoid.units_isOpen`. The result now follows from the previous lemma `ContinuousMulEquiv.piUnits`.

We could work over a general nonarchimedean normed field but we still do not have them in mathlib, so we stick to the case of interest which is the completion of a number field K at a finite place v. Such a completion is a topological field K_v equipped with a discrete valuation, a ring of integers $\mathcal{O}_v$ having a principal maximal ideal (ϖ), and a finite residue field $k_v := \mathcal{O}_v/(\varpi)$.

There is no formal Lean code for the lemmas in this section; it would seem more sensible to prove them in the right generality, and we don't have a definition of nonarchimedean local field yet.

13.5 Some local theory

Lemma 13.5.1

$\mathcal{O}_v$ is an open subring of K_v.

Proof for Lemma 13.5.1

Openness is already in mathlib.

$\mathcal{O}_v$ is also a compact subring of K_v; we proved this in the adele miniproject.

Lemma 13.5.2

$M_2(\mathcal{O}_v)$ is an open subring of $M_2(K_v)$.

Proof for Lemma 13.5.2

Topologically $M_2(\mathcal{O}_v) \cong \mathcal{O}_v^4$ as a subset of K_v^4, so this follows from Lemma 13.5.1 because a product of compacts is compact and a product of opens is open.

Lemma 13.5.3

$M_2(\mathcal{O}_v)$ is a compact subring of $M_2(K_v)$.

Proof for Lemma 13.5.3

Topologically $M_2(\mathcal{O}_v) \cong \mathcal{O}_v^4$ as a subset of K_v^4, so this follows from Theorem 9.3.1 because a product of compacts is compact and a product of opens is open.

Lemma 13.5.4

$GL_2(\mathcal{O}_v)$ is a compact open subgroup of $GL_2(K_v)$.

Proof for Lemma 13.5.4

K_v is known to be Hausdorff, so $M_2(K_v)$ is Hausdorff and the result follows from Lemma 13.4.4 and Lemma 13.4.5.

Recall that there is a projection $\mathcal{O}_v \to k_v$ where k_v is the residue field of v, a finite field. This induces a ring homomorphism $M_2(\mathcal{O}_v) \to M_2(k_v)$ with kernel $M_2(\varpi\mathcal{O}_v)$, an ideal I of $M_2(\mathcal{O}_v)$ isomorphic to $(\varpi\mathcal{O}_v)^4$ and hence also compact and open.

Say Γ_v is a subgroup of $GL_2(k_v)$. Then Γ_v is finite. Consider it as a submonoid of the multiplicative monoid $M_2(k_v)$. Its preimage U_v in $M_2(\mathcal{O}_v)$ is easily checked to be a submonoid of $M_2(\mathcal{O}_v)$; furthermore it is a finite union of cosets of I and is hence compact and open as a submonoid of $M_2(\mathcal{O}_v)$ and hence of $M_2(K_v)$.

Lemma 13.5.5

U_v is a compact open subgroup of $GL_2(K_v)$.

Proof for Lemma 13.5.5

Γ_v is a group and hence its preimage U_v is a subgroup of the monoid $M_2(K_v)$. It is compact and open as we just saw. Hence its units (also U_v) are a subgroup of $GL_2(K_v)$, which is compact and open, again by Lemma 13.4.4 and Lemma 13.4.5.

Say now $\begin{pmatrix}1 & *\\ 0 & 1\end{pmatrix} \subseteq \Gamma_v \subseteq \begin{pmatrix}* & *\\ 0 & *\end{pmatrix}$ and let $U := U_v$ be its preimage in $GL_2(\mathcal{O}_v)$, considered as a compact open subgroup of $GL_2(K_v)$. Choose $0 \neq \alpha \in \mathcal{O}_v$ and define $g = \begin{pmatrix}\alpha & 0\\ 0 & 1\end{pmatrix} \in GL_2(K_v)$. Let's do an explicit double coset decomposition in preparation for a calculation with Hecke operators.

Lemma 13.5.6

The double coset space UgU is the disjoint union of g_tU as t ranges through $\mathcal{O}_v/\alpha\mathcal{O}_v$ and $g_t := \begin{pmatrix}\alpha & \tilde{t}\\ 0 & 1\end{pmatrix}$, where $\tilde{t}$ is any lift of t to $\mathcal{O}_v$.

Proof for Lemma 13.5.6

We first manipulate the statement into a statement about finite groups. We have $UgU = \coprod_t g_tU \iff UgUg^{-1} = \coprod_t g_tUg^{-1} = \coprod_t g_tg^{-1}(gUg^{-1})$. By the second isomorphism theorem this is true if $U = \coprod_t g_tg^{-1}(gUg^{-1} \cap U)$. So when is an element of U in gUg^{-1}? Equivalently, if $x \in U$, when is $g^{-1}xg \in U$? An explicit calculation of matrices shows us that this is true iff $g = \begin{pmatrix}a & b\\ c & d\end{pmatrix}$ with $\alpha \mid b$. Define U^α to be this subgroup of U. We have reduced the question to showing that the matrices $h_t := \begin{pmatrix}1 & \tilde{t}\\ 0 & 1\end{pmatrix}$ are a set of left coset representatives for the subgroup U^α of U.It thus suffices to show that if $u = \begin{pmatrix}a & b\\ c & d\end{pmatrix} \in U$ then $u \in h_tU^\alpha$ iff $b \in \mathcal{O}_v$ reduces mod α to $t \in \mathcal{O}_v/\alpha$. We do this by computing $h_t^{-1}u = \begin{pmatrix}a - \tilde{t}c & b - \tilde{t}d\\ c & d\end{pmatrix}$ and observing that its top right hand entry mod α is zero iff b mod α is t.

13.6 Adelic groups

We are finally ready to discuss the group G and the subgroups U which we will be using to define our Hecke operators. Let K be a number field, let D a quaternion algebra over K and let A_K^∞ be the finite adeles of K; recall that this is a commutative topological ring, defined to be the restricted product of the commutative topological fields K_v as v runs through the finite places of K, with respect to the compact open subrings $\mathcal{O}_v$.

The group G we are interested in for the rest of this miniproject is the group $(DotimesKAK\hat{}infty)\hat{}times$. We want to write down compact open subgroups of this group, but the first thing we need to do is to find a way of talking about elements of the group.

We will assume that there exists an $\mathbb{A}_K^\infty$-algebra isomorphism $D \otimes_K \mathbb{A}_K^\infty = M_2(\mathbb{A}_K^\infty)$ and we will fix such an isomorphism `r`, called a rigidification in the Lean code. We give both of these $\mathbb{A}_K^\infty$-algebras the $\mathbb{A}_K^\infty$-module topology, which is a fancy way of saying the product topology. They are both free of rank 4 as $\mathbb{A}_K^\infty$-modules, and the rigidification is then a homeomorphism because all $\mathbb{A}_K^\infty$-module maps between modules with the $\mathbb{A}_K^\infty$-module topology are continuous.

This means that our group G is isomorphic (both algebraically and topologically) to $GL2(AK\hat{}infty)$. Before we go any further, let say something about matrix rings over complete fields.

Theorem 13.6.1

> `G` is isomorphic and homeomorphic to the restricted product of $GL_2(K_v)$ with respect to the compact open subgroups $GL_2(\mathcal{O}_v)$.

Proof for Theorem 13.6.1

> This follows from lemma and lemma.

If `S` is a finite set of finite places of `K`, and for each $v \in S$ we choose a subgroup $_v$ of $GL_2(k_v)$ then we saw in the previous section how to create a compact open subgroup $\tilde{}_v$ of $GL_2(K_v)$. For $v \notin S$ define $\tilde{}_v = GL_2(\mathcal{O}_v)$. Then $\prod_v \tilde{}_v$ is a compact open subgroup of $\prod_v GL_2(\mathcal{O}_v)$. It is compact subgroups of this form which we shall be using.

13.7 Automorphic forms

We recall some of the definitions of spaces of automorphic forms, from the quaternion algebra project (p. 80).

We fix a totally real field F, a totally definite quaternion algebra D/F, and a coefficient additive abelian group R. Set $G = (D \otimes_F \mathbb{A}_F^\infty)^\times$ as in the previous section. Note that G naturally contains copies of $D^\times$ and $(\mathbb{A}_F^\infty)^\times$. Recall from definition Definition 12.4.1 that an `R`-valued weight 2 automorphic form is a function $f : G \to R$ satisfying the following axioms:

1. $f(dg) = f(g)$ for all $d \in D^\times \subseteq (D \otimes_F \mathbb{A}_F^\infty)^\times$.
2. There exists a compact open subgroup `U`, the level of `f`, such that $f(gu) = f(g)$ for all $g \in G$ and $u \in U$.
3. $f(gz) = f(g)$ for all $z \in (\mathbb{A}_F^\infty)^\times$.

It can be checked that the collection of all such forms is an additive abelian group, and if R is a ring then it is naturally an R-module. Let's call this group A for short. Then A has a left action of G, with $g \cdot f$ defined via $(g \cdot f)(x) := f(xg)$. Recall from definition Definition 12.4.5 that a weight 2 automorphic form of level U is simply an element of the fixed points A^U. In other words, the forms of level U are the forms satisfying the three axioms defining an automorphic form but with the compact open subgroup in the second axiom being U.

13.8 Concrete Hecke operators

Let F be a number field. For each finite place v we have the completion F_v of F at v, which is a normed field equipped with its integer ring $\mathcal{O}_v$, a local ring with finite residue field k_v.

For v a finite place of F, let Δ_v be a subgroup of $k_v^\times$ and consider the subgroup Γ_v of $GL_2(k_v)$ consisting of matrices $\begin{pmatrix} a & b \\ 0 & d \end{pmatrix}$ with $a, d \in k_v^\times$ and $a/d \in \Delta_v$.

It is easily checked that this is a subgroup, and that $\begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix} \subseteq \Gamma_v \subseteq \begin{pmatrix} * & * \\ 0 & * \end{pmatrix}$, so the standard coset-decomposition lemma applies. Let U_{Δ_v} be the preimage of this subgroup in $GL_2(\mathcal{O}_v)$. This is a compact open subgroup of $GL_2(\mathcal{O}_v)$, by the remarks above.

Let S be a finite set of finite places of F, and define $U_\Delta(S)$ to be the matrices in $\prod_v GL_2(\mathcal{O}_v)$ which are in U_{Δ_v} for all $v \in S$ (we put no condition at the places $v \notin S$). We can consider $U_\Delta(S)$ as a subgroup of $GL_2(\mathbb{A}_F^\infty)$; it is a product of compact subgroups and thus compact, and it is a product of opens all but finitely many of which are $GL_2(\mathcal{O}_v)$ and is thus open. Because the inclusion $\prod_v GL_2(\mathcal{O}_v) \to GL_2(\mathbb{A}_F^\infty)$ is an open embedding, we can regard $U_\Delta(S)$ as a compact open subgroup of $GL_2(\mathbb{A}_F^\infty)$.

If we fix r a rigidification, it induces an isomorphism $GL_2(\mathbb{A}_F^\infty) = (D \otimes_F \mathbb{A}_F^\infty)^\times$ and so we can identify $U_\Delta(S)$ with its image in $(D \otimes_F \mathbb{A}_F^\infty)$.

We introduce Hecke operators of two types.

First, for v any place not in S, we choose a uniformiser $\varpi_v \in F_v$, form the invertible 2×2 matrix $\begin{pmatrix} \varpi_v & 0 \\ 0 & 1 \end{pmatrix} \in GL_2(F_v)$, and extend this element to an element $g \in G$ by letting its component at all finite places $w \neq v$ be the identity. We define the Hecke operator $T_v : A^U \to A^U$ to be $[UgU]$, using the notation defined at the beginning of this section.

For the second kind of Hecke operator we choose $0 \neq \alpha \in \mathcal{O}_v$ and we consider the 2×2 matrix in $GL_2(\mathbb{A}_K^\infty)$ which is $\begin{pmatrix} \alpha & 0 \\ 0 & 1 \end{pmatrix}$ at v and 1 at all other local components. Via the rigidification r we obtain an element $g \in G$. We define the Hecke operator $U_{v,\alpha}$ to be $[UgU]$.

The Hecke algebra of interest to us will be generated by the Hecke operators T_v for $v \notin S$ and $U_{v,\alpha}$ for $v \in S$.

Theorem 13.8.1

Say R is a Noetherian ring. Then the subalgebra of the R-linear endomorphisms of A^U generated by the Hecke operators T_v for $v \notin S$ and $U_{v,\alpha}$ for $v \in S$ is a Noetherian commutative ring.

13.9 Analysis of the Hecke algebra

First we discuss commutativity of the Hecke operators. First, assume that $v \notin S$. Then $U = GL_2(\mathcal{O}_v) \times U'$ where U' is a subgroup of the restricted product of $GL_2(F_w)$ for $w \neq v$. We can use `RestrictedProduct.SubmonoidClass.isProductAt` to express this concept of being an internal direct product. If g is the element of G used to make T_v then g is also supported at w, so the double coset space UgU is just $(GL_2(\mathcal{O}_v) \begin{pmatrix} \varpi & 0 \\ 0 & 1 \end{pmatrix} GL_2(\mathcal{O}_v)) \times U'$ and in particular can be decomposed into single left U-cosets of the form $g_i U$ where g_i is also supported at v. This is `RestrictedProduct.mem_coset_and_mulSupport_subset_of_isProductAt`.

Similarly if $v \in S$, if $0 \neq \alpha \in \mathcal{O}_v$ and if $g_v = \begin{pmatrix} \alpha & 0 \\ 0 & 1 \end{pmatrix}$ and is 1 elsewhere, then the double coset space UgU can again be written as $\coprod_i g_i U$ with the g_i supported only at v.

We deduce immediately from lemma Lemma 13.3.4 that two Hecke operators associated to different finite places of F commute. What remains is to check that $U_{,v}$ and $U_{,v}$ commute. In fact we claim more, namely that $U_{,v}U_{,v} = U_{,v}$. This will suffice because $= $.

Lemma 13.9.1

If $v \in S$ and $\begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix} \subseteq \Gamma_v \subseteq \begin{pmatrix} * & * \\ 0 & * \end{pmatrix}$ then $U_{\alpha,v}U_{\beta,v} = U_{\alpha\beta,v}$.

Proof for Lemma 13.9.1

Follows easily from the explicit double coset decomposition proved above.

The reason that the Hecke algebra is Noetherian is that the main theorem of the Fujisaki miniproject immediately implies that A^G is a submodule of a finite free R-module and is hence Noetherian. Its endomorphism algebra is hence a Noetherian R-module, so the sub-R-algebra generated by the Hecke operators is also a Noetherian R-module and thus a Noetherian ring.

14 Appendix: A collection of results which are needed in the proof.

14.1 Results from class field theory

Theorem 14.1.1

The maximal unramified extension K^{un} in a given algebraic closure of K is Galois over K with Galois group "canonically" isomorphic to $\widehat{\mathbb{Z}}$ in two ways; one of these two isomorphisms identifies $1 \in \widehat{\mathbb{Z}}$ with an arithmetic Frobenius, the endomorphism inducing $x \mapsto x^q$ on the residue field of K^{un}, where q is the size of the residue field of K. The other identifies 1 with geometric Frobenius, defined to be the inverse of arithmetic Frobenius.

Definition 14.1.2

The topological group described above is called the Weil group of K.

Theorem 14.1.3

If K is a finite extension of $\mathbb{Q}_p$, then there are two canonical isomorphisms of topological abelian groups between $K^\times$ and the abelianization of the Weil group of K.

Proof for Theorem 14.1.3

This is the main theorem of local class field theory; see for example the relevant articles in the cited class-field-theory references or many other places.

Theorem 14.1.4

If M is finite then the cohomology groups $H^i(G_K, M)$ are finite.

Theorem 14.1.5

If M is torsion then $H^i(G_K, M) = 0$ if $i > 2$.

Theorem 14.1.6

$H^2(G_K, \mu_n)$ is canonically isomorphic to $\mathbb{Z}/n\mathbb{Z}$.

Theorem 14.1.7

If $\mu = \bigcup_{n\geq 1} \mu_n$ and $M' := \mathrm{Hom}(M, \mu)$ is the dual of M, then for $0 \leq i \leq 2$ the cup-product pairing $H^i(G_K, M) \times H^{2-i}(G_K, M') \to H^2(G_K, \mu) = \mathbb{Q}/\mathbb{Z}$ is perfect.

Theorem 14.1.8

If $h^i(M)$ denotes the order of $H^i(G_K, M)$, then $h^0(M) - h^1(M) + h^2(M) = 0$.

Theorem 14.1.9

If N is a finite extension of $\mathbb{Q}$ then there are two "canonical" isomorphisms of topological groups between the profinite abelian groups $\pi_0(\mathbb{A}_N^\times / N^\times)$ and Γ_N^{ab}; one sends local uniformisers to arithmetic Frobenii and the other to geometric Frobenii; each of the global isomorphisms is compatible with the local isomorphisms above.

Proof for Theorem 14.1.9

This is the main theorem of global class field theory; see for example Tate's article in the cited reference.

Theorem 14.1.10

Let S be a finite set of places of a number field K. For each $v \in S$ let L_v/K_v be a finite Galois extension. Then there is a finite solvable Galois extension L/K such that if w is a place of L dividing $v \in S$, then L_w/K_v is isomorphic to L_v/K_v as a K_v-algebra. Moreover, if K^{avoid}/K is any finite extension then L can be chosen linearly disjoint from K^{avoid}.

14.2 Structures on the points of an affine variety.

All rings and algebras in this section are commutative with a 1, and all morphisms send 1 to 1.

Let $X = \mathrm{Spec}(A)$ be an affine scheme of finite type over a field K. For example X could be an affine algebraic variety; in fact we shall only be interested in smooth affine varieties in the applications, but the initial definition and theorem are fine for all finite type schemes.

If R is any K-algebra then one can talk about the R-points X(R) of X, which in this case naturally bijects with the K-algebra homomorphisms from A to R.

Definition 14.2.1

If X is an affine scheme of finite type over K, and if R is a K-algebra which is also a topological ring, then we define a topology on the R-points $X(R)$ of X by embedding the K-algebra homomorphisms from A to R into the set-theoretic maps from A to R with its product topology, and giving it the subspace topology.

Theorem 14.2.2

If X is as above and $X \to \mathbb{A}^n_K$ is a closed immersion, then the induced map from $X(R)$ with its topology as above to R^n is an embedding of topological spaces (that is, a homeomorphism onto its image).

Proof for Theorem 14.2.2

See Conrad's notes[1].

We now specialise to the smooth case. I want to make the following conjectural definition:

Definition 14.2.3

Let K be a field equipped with an isomorphism to the reals, complexes, or a finite extension of the p-adic numbers. Let X be a smooth affine algebraic variety over K. Then the points X(K) naturally inherit the structure of a manifold over K.

Probably this is fine for a broader class of fields K.

Theorem 14.2.4

If X is as in the previous definition and $X \to \mathbb{A}^n_K$ is a closed immersion, then the induced map from $X(K)$ with its manifold structure to K^n is an embedding of manifolds.

Proof for Theorem 14.2.4

I'm assuming this is standard, if true.

Corollary 14.2.5

If G is an affine algebraic group of finite type over $K = \mathbb{R}$ or $K = \mathbb{C}$ then $G(K)$ is naturally a real or complex Lie group.

The corollary, for sure, is true! And it's all we need. I have not yet made any serious effort to find a reference for the definition or independence, although there seem to be some ideas here[2]. As a toy example, one can embed $\mathrm{GL}_n(\mathbb{R})$ into either $\mathbb{R}^{n^2+1}$ via $M \mapsto (M, \det(M)^{-1})$ or into $\mathbb{R}^{2n^2}$ via $M \mapsto (M, M^{-1})$ and the claim is that the two induced manifold structures are the same.

[1] https://math.stanford.edu/~conrad/papers/adelictop.pdf

[2] https://leanprover.zulipchat.com/#narrow/stream/116395-maths/topic/top.20space.20.2F.20manifold.20structure.20on.20points.20of.20alg.20varieties/near/431812525

14.3 Algebraic groups.

The concept of an affine algebraic group over a field K can be implemented in Lean as a commutative Hopf algebra over K, as a group object in the category of affine schemes over K, as a representable group functor on the category of affine schemes over K, or as a representable group functor on the category of schemes over K which is represented by an affine scheme. All of these are the same to mathematicians but different to Lean and some thought should go into which of these should be the actual definition, and which should be proved to be the same thing as the definition.

Definition 14.3.1

An affine algebraic group G of finite type over a field k is said to be connected if it is connected as a scheme, and reductive if $G_{\overline{k}}$ has no nontrivial smooth connected unipotent normal k-subgroup.

14.4 Automorphic forms and representations

Definition 14.4.1

A function $f : G(N_\infty) \to \mathbb{C}$ is slowly-increasing if there exist some $C > 0$ and $n \geq 1$ such that $|f(x)| \leq C||x||_\rho^n$.

Theorem 14.4.2

This is independent of the choice of ρ as above.

Definition 14.4.3

An automorphic form is a function $\phi : G(\mathbb{A}_N) \to \mathbb{C}$ satisfying the following conditions.

1. ϕ is locally constant on $G(\mathbb{A}_N^f)$ and C^∞ on $G(N_\infty)$. In other words, for every g_∞, $\phi(-, g_\infty)$ is locally constant, and for every g_f, $\phi(g_f, -)$ is smooth.
2. ϕ is left-invariant under $G(N)$.
3. ϕ is right-U_∞-finite, that is, the space spanned by $x \mapsto \phi(xu)$ as u varies over U_∞ is finite-dimensional.
4. ϕ is right K_f-finite, where K_f is one, or equivalently all, compact open subgroups of $G(\mathbb{A}_N^f)$.
5. ϕ is ‡-finite, where ‡ is the centre of the universal enveloping algebra of the Lie algebra of $G(N_\infty)$, acting via differential operators. Equivalently ϕ is annihilated by a finite index ideal of this centre, so morally ϕ satisfies many differential equations of a certain type.

6. For all g_f, the function $g_\infty \mapsto \phi(g_f g_\infty)$ is slowly-increasing in the sense above.

Definition 14.4.4

An automorphic form is cuspidal, or a cusp form, if it furthermore satisfies $\int_{U(N)\backslash U(\mathbb{A}_N)} \phi(ux)\, du = 0$, where P runs through all the proper parabolic subgroups of G defined over N and U is the unipotent radical of P, and the integral is with respect to the measure coming from Haar measure.

Definition 14.4.5

The group $G(\mathbb{A}_N)$ acts on itself on the right, and this induces a left action of its subgroup $G(\mathbb{A}_N^f) \times U_\infty$ on the spaces of automorphic forms and cusp forms. The Lie algebra $\mathfrak{g}$ of $G(N_\infty)$ also acts, via differential operators. Furthermore the actions of $\mathfrak{g}$ and U_∞ are compatible in the sense that the differential of the U_∞ action is the action of its Lie algebra considered as a subalgebra of $\mathfrak{g}$. We say that the spaces are $(G(\mathbb{A}_N^f) \times U_\infty, \mathfrak{g})$-modules.

For non-cuspidal representations, they do not decompose as a direct sum; there is a continuous spectrum which decomposes as a direct integral. We may not ever need these. As a result the definition of an automorphic representation has to be slightly modified in the non-cuspidal case.

Definition 14.4.6

An automorphic representation is an irreducible $(G(\mathbb{A}_N^f) \times U_\infty, \mathfrak{g})$-module isomorphic to an irreducible subquotient of the space of automorphic forms.

14.5 Galois representations

14.6 Algebraic geometry

Definition 14.6.1

We need the definition of the canonical model over F of the Shimura curve attached to an inner form of GL_2 with precisely one split infinite place, and likewise the Shimura surface associated to an inner form split at two infinite places and ramified elsewhere.

Theorem 14.6.2

Let K^{avoid}/K be a Galois extension of number fields. Suppose also that S is a finite set of places of K, and for each $v \in S$ let L_v/K_v be a finite Galois extension. Suppose also that T/K is a smooth, geometrically connected curve and that for each $v \in S$ we are given a nonempty, $\mathrm{Gal}(L_v/K_v)$-invariant open subset $\Omega_v \subseteq T(L_v)$. Then there is a finite Galois extension L/K and a point $P \in T(L)$ such that:

- L/K is Galois and linearly disjoint from K^{avoid} over K.
- If $v \in S$ and w is a prime of L above v, then L_w/K_v is isomorphic to L_v/K_v.
- $P \in \Omega_v \subseteq T(L_v) \cong T(L_w)$ via one such K_v-algebra morphism; this makes sense as Ω_v is $\mathrm{Gal}(L_v/K_v)$-invariant.

Note that we do not even have the definition of a curve over a field in Lean.

14.7 Algebra

Definition 14.7.1

A cuspidal automorphic representation is an irreducible $(G(\mathbb{A}_N^f) \times U_\infty, \mathfrak{g})$-module isomorphic to an irreducible summand of the space of cusp forms.

Theorem 14.7.2

The cusp forms decompose as a typically infinite direct sum of irreducible $(G(\mathbb{A}_N^f) \times U_\infty, \mathfrak{g})$-modules.

Theorem 14.7.3

An irreducible admissible $(G(\mathbb{A}_N^f) \times U_\infty, \mathfrak{g})$-module is a restricted tensor product of irreducible representations π_v of $G(N_v)$ as v runs through the finite places of N, tensored with a tensor product of irreducible $(\mathfrak{g}_v, U_{\infty,v})$-modules as v runs through the infinite places of N. The representations π_v are unramified for all but finitely many v.

Definition 14.7.4

Let N be a number field. A compatible family of d-dimensional Galois representations over N is a finite set of finite places S of N, a number field E, a monic degree d polynomial $F_{\mathfrak{p}}(X) \in E[X]$ for each finite place p of K not in S, and, for each prime number ℓ and field embedding $\phi : E \to \overline{\mathbf{Q}}_\ell$ (or essentially equivalently for each finite place of E), a continuous homomorphism $\rho : \mathrm{Gal}(\overline{K}/K) \to$

$\mathrm{GL}_2(\overline{\mathbf{Q}}_\ell)$ unramified outside S and the primes of K above ℓ, such that $\rho(\mathrm{Frob}_{\mathfrak{p}})$ has characteristic polynomial $P_\pi(X)$ if π lies above a prime number $p \neq \ell$ with $p \notin S$.

Theorem 14.7.5

Given an automorphic representation π for an inner form of GL_2 over a totally real field and with reflex field E, such that π is weight 2 discrete series at every infinite place, there exists a compatible family of 2-dimensional Galois representations associated to π, with S the places at which π is ramified and $F_{\mathfrak{p}}(X)$ the monic polynomial with roots the two Satake parameters for π at $\mathfrak{p}$.

Dependency Graph

The dependency graph is available in the HTML output.

Blueprint Summary

The Blueprint summary is available in the HTML output.

Blueprint Bibliography

«mazur-torsion» Barry Mazur (1977). "Modular curves and the Eisenstein ideal"[3]. *Publications Mathématiques de l'IHÉS.* ***47****(1), pp. 33–186.*

edix *Bas Edixhoven (1992). "The weight in Serre's conjectures on modular forms"*[4]. Inventiones Mathematicae. **109**(3), pp. 563–594.

kwII Chandrashekhar Khare and Jean-Pierre Wintenberger (2009). "Serre's modularity conjecture. II"[5]. *Inventiones Mathematicae.* ***178****(3), pp. 505–586.*

poitou_odlyzko_bounds *Georges Poitou (1977). "Sur les petits discriminants".* Séminaire Delange-Pisot-Poitou. **18**(1).

ddt Henri Darmon, Fred Diamond, and Richard Taylor (1995). "Fermat's last theorem". *Current Developments in Mathematics.* ***1995****(1), pp. 1–154.*

«serre-galcoh» *Jean-Pierre Serre (1964). "Cohomologie galoisienne".* Lecture Notes in Mathematics. **5**(1).

serrepropgal Jean-Pierre Serre (1972). "Propriétés galoisiennes des points d'ordre fini des courbes elliptiques"[6]. *Inventiones Mathematicae.* ***15****(4), pp. 259–331.*

serreconj *Jean-Pierre Serre (1987). "Sur les représentations modulaires de degré 2 de GQ"*[7]. Duke Mathematical Journal. **54**(1), pp. 179–230.

voightbook John Voight (2021). "Quaternion Algebras"[8]. *Open book draft.* ***1****(1).*

silverman1 *Joseph H. Silverman (2009). "The arithmetic of elliptic curves"*[9]. Graduate Texts in Mathematics. **106**(2).

[3] http://www.numdam.org/item?id=PMIHES_1977__47__33_0
[4] https://doi.org/10.1007/BF01232041
[5] https://doi.org/10.1007/s00222-009-0206-6
[6] https://doi.org/10.1007/BF01405086
[7] https://doi.org/10.1215/S0012-7094-87-05413-5
[8] https://quatalg.org/book/
[9] https://doi.org/10.1007/978-0-387-09494-6

«moret-bailly» Laurent Moret-Bailly (1989). "Groupes de Picard et problèmes de Skolem. I, II"[10]. *Annales scientifiques de l'École Normale Supérieure.* ***22****(2), pp. 161–179.*

«katz-mazur» *Nicholas M. Katz and Barry Mazur (1985). "Arithmetic moduli of elliptic curves".* Annals of Mathematics Studies. **108**(1).

corvallis1 Proceedings volume, 1979. "Automorphic forms, representations and L-functions. Part 1". In *Proceedings of Symposia in Pure Mathematics XXXIII.*

«taylor-mero-cont» *Richard Taylor (2006). "On the meromorphic continuation of degree two L-functions".* Documenta Mathematica. **11**(1), pp. 729–779.

blggt Thomas Barnet-Lamb, Toby Gee, David Geraghty, and Richard Taylor (2014). "Potential automorphy and change of weight"[11]. *Annals of Mathematics.* ***179****(2), pp. 501–609.*

«toby-modularity» *Toby Gee (2006). "Modularity lifting theorems for ordinary Galois representations".* Mathematische Annalen. **335**(4), pp. 681–715.

[10] http://www.numdam.org/item?id=ASENS_1989_4_22_2_161_0
[11] https://doi.org/10.4007/annals.2014.179.2.3