The Skinner-Wiles level lowering trick

1. Background.

The trick in question is a generalization of an argument introduced in [SW] for $n = 2$. The background is as follows. Wiles' original paper on Fermat's Last Theorem divided the proof of the isomorphism $R \xrightarrow{\sim} \mathbb{T}$ into two parts. The first part was to prove it in the *minimal* case:

$$R_\emptyset \xrightarrow{\sim} \mathbb{T}_\emptyset.$$

The residual representation $\bar{\rho}$ was assumed ramified at a set S of primes of the base field F not dividing ℓ, together with the set S_ℓ of divisors of ℓ (for Wiles, $F = \mathbb{Q}$). The deformation ring $R_\emptyset$ classifies deformations

$$\rho_A : \Gamma_F \to GL(2, A)$$

that are unramified outside $S \cup S_\ell$, with the usual (crystalline or other) conditions at S_ℓ, and such that, for all $v \in S$, the natural map $\rho_A(I_v) \to \bar{\rho}(I_v)$ is an isomorphism. In other words, ρ_A is no more ramified than $\bar{\rho}$ at primes in S. This step was carried out in the Taylor-Wiles article. The second step deduces the general case – additional ramification allowed at $v \in S^*$, for some finite set S^* containing S – from the minimal case.

The structure of the proof is such that one needs to know, not only that $\bar{\rho}$ admits some lifting that comes from modular forms, but that it admits a modular (automorphic) lifting that is minimal in the above sense. In the case considered by Wiles, it was in fact known that the existence of some modular lifting implies the existence of a minimal modular lifting. The crucial step was due to Ribet, who showed that, if $\rho = \rho_{f,\ell}$ is the 2-dimensional Galois representation attached to a modular form f of weight 2 (say) and level N, and if q is a prime dividing N such that $\bar{\rho}$ is unramified at q, then f is congruent modulo ℓ to a form f' of level prime to q. Ribet's theorem was also the most subtle step in showing that the most precise form of Serre's conjecture on the modularity of 2-dimensional representations of $\Gamma_\mathbb{Q}$ over finite fields followed from the least precise form of the conjecture. This theorem was generalized to modular forms of other weights and levels but the basic argument was always the same.

Ribet's proof involved a very ingenious analysis of the geometry of Néron models of Shimura curves that could not be reproduced in other situations. It was long believed that this represented an insurmountable barrier to the generalization of Wiles' results to Galois representations of higher dimensions, though an unpublished manuscript of Harris-Taylor (the predecessor of [CHT]) showed how to generalize Taylor-Wiles, using the arguments of Diamond and Fujiwara. The article [SW] showed how to avoid Ribet's level lowering argument by substituting a base-change argument. Instead of finding an f' congruent to f of level prime to q, they find an f' congruent to f that is apparently more ramified at q than f, but whose ramification has finite (abelian) image. Thus after a finite totally real base change F'/F, $\rho_{f',\ell}$ is unramified at q. This leaves unsettled the interesting and important question whether $R \xrightarrow{\sim} \mathbb{T}$ over the original F, but to prove modularity of ρ, it suffices to establish $R \xrightarrow{\sim} \mathbb{T}$ over F'.

Once the minimal (Taylor-Wiles) isomorphism $R_\emptyset \xrightarrow{\sim} \mathbb{T}_\emptyset$ is established, the passage to the general case requires the techniques of level *raising*, also developed by Ribet. These arguments were completely original when Ribet invented them but appeared to be much less difficult than his level lowering theorem. All attempts to generalize them to higher dimension have failed thus far, however, because it has proved impossible to generalize a lemma of Ihara that Ribet used as his starting point. In [CHT] we outlined a strategy, developed by Taylor and his student Russ Mann, for proving the non-minimal case on the basis of an appropriate generalization of Ihara's Lemma. In [T], Taylor found a way to ignore level-raising entirely, using Kisin's technique of framed deformations and a geometric analysis of the level raising/lowering problem that is based on exactly the same principle as the Skinner-Wiles trick.

2. The Skinner-Wiles lemma.

Notation is as in the notes denoted **Introduction to the Proof**, §3; however, G_0 is the unitary group of a positive-definite hermitian form and there is no division algebra B. We consider the open compact subgroup $U = \prod_v U_v$ defined by (3.5.1-4). For $v \in R$,

$$U_v = I_v = \{u \in GL(n, \mathcal{O}_v) \mid u_{ij} \in \mathfrak{m}_{\mathcal{O}_v} \ \forall i > j\}.$$

In other words, the reduction of u modulo $\mathfrak{m}_{\mathcal{O}_v}$ is upper triangular. Consider as in [loc. cit] the subgroup $I(1)_v \subset I_v$:

$$I(1)_v = \{u \in I_v \mid u_{ii} \in 1 + \mathfrak{m}_{\mathcal{O}_v} \ \forall i\}.$$

In other words, the reduction of u modulo $\mathfrak{m}_{\mathcal{O}_v}$ is upper triangular unipotent. Let $U(1) = \prod_{v \notin R} U_v \times \prod_{v \in R} I(1)_v$, $D_R = \prod_{v \in R} I_v/I(1)_v = U/U(1) = \prod_{v \in R}(k(v)^\times)^n$. Recall that a character of $I_v/I(1)_v$ is denoted $\chi_v = (\chi_{1,v}, \ldots, \chi_{n,v})$ where each $\chi_{i,v}$ is a character of $k(v)^\times$. Let χ_v be as above, for $v \in R$, let A be either the coefficient ring $\mathcal{O}$, which we assume contains all the values of all possible χ_v, or its residue field k, and define

$$S_{\{\chi_v\}}(U, A) = \{f \in S(G_0, A) \mid f(gu) = \prod_{v \in R} \chi_v^{-1}(u_v) f(g)\}$$

for all $g \in G_0(\mathbf{A}_f)$ and $u = \prod u_v \in U$, as well as

$$S(U(1), A) = \{f \in S(G_0, A) \mid f(gu) = f(g)\},$$

which contains all the $S_{\{\chi_v\}}(U, A)$. This is the module on which our Hecke algebra $\mathbb{T} = \mathbb{T}^T(U(1))$ acts, as in [loc. cit.]. Recall that this Hecke algebra is generated by the (unramified, split) Hecke operators at primes outside T, and in particular the set of generators is independent of $\{\chi_v\}$. On the other hand, the action of D_R on $S(U(1), A)$ commutes with $\mathbb{T}$. Assuming U is sufficiently small – this is guaranteed by choosing an appropriate non-empty set S_1, as in [loc. cit., 3.5.3] – $S(U(1), A)$ is a free $A[D_R]$-module. The character $\{\chi_v\}$ defines a map $\mathcal{O}[D_R] \to \mathcal{O}^\times$ and there is a natural isomorphism

$$S(U(1), \mathcal{O}) \otimes_{\mathcal{O}[D_R], \{\chi_v\}} \mathcal{O} \xrightarrow{\sim} S_{\{\chi_v\}}(U, \mathcal{O}). \tag{2.1}$$

The same is true if $\mathcal{O}$ is replaced by k, in which case we write $\bar{\chi}_v$ instead of χ_v.

Lemma 2.2 [SW]. *Let* $\{\chi_v\}, \{\chi'_v\}$ *be two collections of characters of* D_R *as above. Let* $\mathfrak{m} \subset \mathbb{T}^T(U(1))$ *be a maximal ideal. Suppose* $\bar{\chi}_v = \bar{\chi}'_v$ *for all* v*. Then* $S_{\{\chi_v\}}(U, \mathcal{O})_{\mathfrak{m}} \neq 0$ *if and only if* $S_{\{\chi'_v\}}(U, \mathcal{O})_{\mathfrak{m}} \neq 0$.

Proof. This is certainly true if $\mathcal{O}$ is replaced by k in the statement of the lemma. But both $S_{\{\chi_v\}}(U, \mathcal{O})_{\mathfrak{m}}$ and $S_{\{\chi'_v\}}(U, \mathcal{O})_{\mathfrak{m}}$ are free $\mathcal{O}$-modules, so that

$$S_{\{\chi_v\}}(U, k)_{\mathfrak{m}} = S_{\{\chi_v\}}(U, \mathcal{O})_{\mathfrak{m}} \otimes k$$

and likewise for $S_{\{\chi'_v\}}(U, \mathcal{O})_{\mathfrak{m}}$. The Lemma is thus clear.

This is applied with $\chi'_v = 1$ for all v, χ_v *regular* in the sense that, for each v, $\chi_{i,v} \neq \chi_{j,v} \neq 1$ for all i, j. Recall that $\ell \mid |k(v)|^\times$ and $\ell > n$, so regular sets of χ_v exist. (If we don't want to assume $\ell > n$, we may replace the ground field F by an extension unramified at all $v \in R$ so that $k(v)^\times$ is divisible by a sufficiently large power of ℓ for all v.) An automorphic representation π of G_0 corresponding to an automorphic form in $S_{\{\chi_v\}}(U, \mathcal{O})$ defines an ℓ-adic Galois representation $\rho_{\pi,\ell}$ with the property that

$$\rho_{\pi,\ell}(I_v) \text{ is finite } \forall v \in R. \tag{2.3}$$

Thus an appropriate finite abelian base change of π eliminates all ramification at $v \in R$. If we now start with π' corresponding to $S_{\{1\}}(U, \mathcal{O})_{\mathfrak{m}}$ and assume π corresponds to $S_{\{\chi_v\}}(U, \mathcal{O})_{\mathfrak{m}}$, then $\bar{\rho}_{\pi',\ell} \xrightarrow{\sim} \bar{\rho}_{\pi,\ell}$. In this way, at the price of a possible base change F'/F, we can find an automorphic form $\pi_{F'}$ on some $G_{0,F'}$ such that

$$\bar{\rho}_{\pi_{F'},\ell} \xrightarrow{\sim} \bar{\rho}_{\pi,\ell} \mid_{\Gamma_{F'}}$$

but $\rho_{\pi_{F'},\ell}$ is unramified at all primes above R, no matter how ramified $\rho_{\pi,\ell}$ was at $v \in R$. This is the Skinner-Wiles trick.

References.

[SW] C.Skinner and A.Wiles, Base change and a problem of Serre, *Duke Math. J.*, **107** (2001), 1525.