hexagon:2611.01170v1

Essential dimension of principal congruence covers of Shimura varieties

Mark Kisin*

Department of Mathematics, Harvard University
Cambridge, MA 02138, USA

Abstract

We prove that the principal congruence cover $S(p) \to S$ of a Shimura variety is geometrically p-incompressible for all sufficiently large good primes p. For varieties of pre-abelian type, $p > \dim S + 3$ suffices. The proof reformulates the obstruction to compression in [4, 6] in terms of the classifying stack $B\mathcal{F}$ of the finite flat group scheme governing the local cover. An invariant differential on the special fibre of $B\mathcal{F}$ allows us to apply the prismatic restriction method of [5].

Contents

1 Introduction 1

2 Prismatic restriction for smooth Artin stacks 4

3 Finite flat group schemes and specialization 6

4 Integral models of principal congruence covers 8

5 p-incompressibility of principal congruence covers 13

1 Introduction

Let (G, X) be a Shimura datum with reflex field E. At a prime p where G is unramified, let $\mathcal{G}/\mathbb{Z}_p$ be a reductive model and put

$$K_p = \mathcal{G}(\mathbb{Z}_p), \qquad K_p(p) = \ker\big(\mathcal{G}(\mathbb{Z}_p) \longrightarrow \mathcal{G}(\mathbb{F}_p)\big).$$

For a sufficiently small prime-to-p level K^p, set

$$S = \mathrm{Sh}_{K^pK_p}(G, X), \qquad S(p) = \mathrm{Sh}_{K^pK_p(p)}(G, X), \qquad d = \dim S.$$

*The author was supported by NSF grant DMS-2502326.

Our main result is the following:

Theorem A. *Suppose that G is unramified at $p > d+3$. Then*

$$\mathrm{ed}_p(S(p)/S) = d \tag{1.1}$$

if (G, X) is of pre-abelian type or if $p \gg 0$.

Here $\mathrm{ed}_p(S(p)/S)$ denotes the *p-essential dimension* of the cover $S(p) \to S$ base changed to $\mathbb{C}$. Recall that for $T \to Y$ a finite étale cover of a smooth $\mathbb{C}$-scheme, its essential dimension at a prime p, denoted $\mathrm{ed}_p(T/Y)$, is the least dimension of a variety from which this cover can be obtained by rational pullback, after a generically finite base change of degree prime to p. The cover is *p-incompressible* if $\mathrm{ed}_p(T/Y) = \dim Y$.

Definition 1.1. Let C be an algebraically closed field of characteristic zero and T a torsor under a finite group Γ over a field K/C. Its *essential dimension at p* is

$$\mathrm{ed}_{C,p}(T) = \min \mathrm{trdeg}_C\, L,$$

where the minimum runs over finite extensions K'/K of degree prime to p and subfields $C \subset L \subset K'$ over which $T_{K'}$ descends as a Γ-torsor. For a finite cover $T \to Y$ of an integral C-variety with generic fibre a Γ-torsor, put $\mathrm{ed}_p(T/Y) = \mathrm{ed}_{C,p}(T_{\kappa(Y)})$, where $\kappa(Y)$ is the function field of Y. For a disconnected base we take the maximum over its components.

Descending the underlying finite étale algebra gives the same minimum. If the algebra descends to L, its simply transitive Γ-actions are parametrized by a finite étale L-scheme. The given action over K' therefore descends to a finite extension of L in K'. Taking $K' = L = \kappa(Y)$ gives $\mathrm{ed}_p(T/Y) \leq \dim Y$.

Recall that a Shimura datum is of pre-abelian type if its adjoint datum is isomorphic to the adjoint datum of a Shimura datum of Hodge type [10, Defn. 7.7.1].

We deduce Theorem A by applying the following local result (Theorem 3.4 below) to a finite flat group scheme $\mathcal{F}$, whose geometric generic fiber is identified with $\mathcal{U}(\mathbb{F}_p)$, where $\mathcal{U}$ is a unipotent subgroup of $\mathcal{G}$ associated to the μ-ordinary slope filtration.

Theorem B. *Let k be an algebraically closed field of characteristic p, and let $\mathcal{F}/W(k)$ be a finite flat group scheme whose special fibre H is nilpotent. Suppose that $\mathcal{F}$ acts freely on a connected smooth affine $W(k)$-scheme $\mathcal{X}$ with nonempty special fibre. Put $d = \dim_k \mathrm{Lie}(H)$ and $\mathcal{Y} = \mathcal{X}/\mathcal{F}$, and let $\overline{K}$ be an algebraic closure of $W(k)[1/p]$. If $p > d+3$, then*

$$\mathrm{ed}_p(\mathcal{X}_{\overline{K}}/\mathcal{Y}_{\overline{K}}) \geq d.$$

In [4], the author, together with Farb and Wolfson, used deformation theory at ordinary points to prove incompressibility of congruence covers. The later paper [5] uses prismatic cohomology to detect étale cohomology classes that remain nonzero on every dense open, through the nonvanishing of differential forms. This was used to prove incompressibility for a different class of congruence covers.

Fakhruddin–Saini extended the first method beyond the ordinary case for certain unitary Shimura varieties [6]. They construct smooth integral models of level covers with free actions of finite flat group schemes, which may be noncommutative, and obtain incompressibility by controlling stabilizers after specialization. Their formulation in terms of finite flat group actions also motivates the construction below.

Because of the restriction $p > d+3$, Theorem A does not subsume all the results of [4, 5]. In particular, under their respective hypotheses, the ordinary-locus results of [4] and the tube-type cases of [5, Thm. 3.3.17] impose no dimension-dependent lower bound on p.

To explain the proofs, we begin by sketching a proof of Theorem B. The key observation is that the infinitesimal obstructions in [4, 6] admit a reformulation in terms of the classifying stack $B\mathcal{F}$ of the finite flat group scheme acting on the local cover. This supplies the differential form needed to apply the prismatic restriction method of [5].

Use the notation of Theorem B. Let $e : \operatorname{Spec} k \to H$ be the identity section and $L_{H/k}$ the cotangent complex. The conjugation representation on cohomology in degree -1

$$V_H = H^{-1}(e^* L_{H/k})$$

has dimension d. Nilpotence of H gives V_H a filtration with trivial successive quotients, so its determinant is trivial. For $d < p$, Totaro's description of the Hodge cohomology of BH [14, Thm. 3.1] identifies any generator of $\bigwedge^d V_H$ with a nonzero class

$$\omega_H \in H^0(BH, \bigwedge^d L_{BH/k})$$

(Lemma 3.2). Since $\mathcal{F}$ acts freely on $\mathcal{X}$, the quotient map $\mathcal{X} \to \mathcal{Y} = \mathcal{X}/\mathcal{F}$ is an $\mathcal{F}$-torsor. Its classifying morphism $f : \mathcal{Y} \to B\mathcal{F}$ is smooth because $\mathcal{X}$ is smooth over $W(k)$. Consequently, $f_k^* \omega_H$ is a nonzero differential d-form on each component of $\mathcal{Y}_k$.

The stack $B\mathcal{F}$ is proper and smooth over $W(k)$, and its Hodge cohomology groups are finitely generated $W(k)$-modules. Smooth descent and the comparison theorems of [8, 9] allow us to extend the prismatic restriction argument of [5, §2.2] to this stack. Let C be an algebraically closed complete extension of $W(k)[1/p]$ and set $Y = \mathcal{Y}_C$. Combining the differential above with this restriction argument shows, for $p > d+3$, that

$$H^d(\mathcal{F}(C), \mathbb{F}_p) \longrightarrow H^d_{\text{ét}}(V, \mathbb{F}_p)$$

is nonzero for every dense open $V \subset Y$. If $\operatorname{ed}_p(\mathcal{X}_{\overline{K}}/\mathcal{Y}_{\overline{K}}) < d$, the argument of [5, proof of Prop. 2.3.3] forces this map to vanish for some V. This proves Theorem B.

For a Shimura variety of Hodge type, the construction of $\mathcal{F}$ can be described using the μ-ordinary Serre–Tate theory of Shankar–Zhou [12]. Take $k = \overline{\mathbb{F}}_p$ and a μ-ordinary point x of the integral model over $W(k)$. Choose a Hodge embedding, and let D be the p-divisible group of the resulting abelian scheme, with its tensors, on an étale neighborhood $\mathscr{U}$ of x. Let $D_x^{\mathrm{can}}/W(k)$ be the canonical lift of D_x. The group scheme $\mathcal{F}$ consists of the automorphisms of $D_x^{\mathrm{can}}[p]$ that preserve the tensors and the slope filtration and induce the identity on its graded object. Consider the scheme $\mathcal{I}$ of tensor-preserving isomorphisms

$$D_x^{\mathrm{can}}[p] \times_{W(k)} \mathscr{U} \xrightarrow{\sim} D[p].$$

The group $\mathcal{F}$ acts freely on $\mathcal{I}$ by precomposition. We use the results of Madapusi–Youcis [10] to show that $\mathcal{F}$ is flat, and that after shrinking $\mathscr{U}$, the scheme $\mathcal{I}$ is affine and smooth, with nonempty special fibre (Proposition 4.19). In the unitary PEL case, $\mathcal{F}$ and $\mathcal{I}$ appear in the paper of Fakhruddin–Saini [6, Lem. 3.4 and Prop. 3.8, 4.2].

Let $\mathcal{P} = \mathcal{M} \ltimes \mathcal{U}$ be the parabolic in $\mathcal{G}$ preserving the slope filtration, with Levi subgroup $\mathcal{M}$ acting on the graded object. A trivialization of the geometric generic fibre of $D_x^{\mathrm{can}}[p]$ compatible with the tensors and the slope filtration gives

$$\mathcal{F}(C) \simeq \mathcal{U}(\mathbb{F}_p) = \ker\big(\mathcal{P}(\mathbb{F}_p) \to \mathcal{M}(\mathbb{F}_p)\big).$$

The special fibre $\mathcal{F}_k$ is nilpotent and $\dim_k \mathrm{Lie}(\mathcal{F}_k) = \dim S$; in the ordinary case, $\mathcal{F} \simeq \mu_p^{\dim S}$. On geometric generic fibres, $\mathcal{I} \to \mathscr{U}$ is obtained from $S(p) \to S$ by pullback. Applying Theorem B to the connected component of $\mathcal{I}$ through the identity isomorphism over x then gives $\mathrm{ed}_p(S(p)/S) \geq \dim S$. The results of Gardner–Madapusi and Madapusi–Youcis [7, 10] extend this construction to arbitrary Shimura data at all sufficiently large good primes, and to pre-abelian data at every good prime $p > 2$. We use their formalism in Section 4 to give a uniform construction of $\mathcal{F}$ and $\mathcal{I}$.

AI use

The ideas in this manuscript were produced by GPT Astra when we asked it to prove p-incompressibility of congruence covers for arbitrary Shimura varieties. We did no more than point out the results of [4] and [5], and it came as a pleasant surprise to us when Astra combined the ideas in these two papers in the way presented here. The key new idea is that one can apply the prismatic techniques of [5] directly to the classifying stack of an appropriate group scheme. An initial version of the paper was written by Astra, but this was heavily edited by the author for readability and accuracy.

2 Prismatic restriction for smooth Artin stacks

2.1. Let k be a perfect field of characteristic p. We use the Breuil–Kisin prism $(A, (E))$, where

$$A = W(k)[[u]], \qquad E = u - p, \qquad \varphi(u) = u^p$$

and φ acts on $W(k)$ by Witt-vector Frobenius. Put $A_0 = A/p = k[[u]]$. For a smooth Artin stack $\mathcal{B}$ of finite presentation over $W(k)$, let $C_{\mathcal{B}}$ denote the prismatic cohomology $R\Gamma_\Delta(\widehat{\mathcal{B}}/A)$ of its p-adic completion [9, Defn. 2.2.1 and Lem. 2.2.5]. Put

$$\overline{C}_{\mathcal{B}} = C_{\mathcal{B}} \otimes^{\mathbf{L}}_A A_0, \qquad M^i_{\mathcal{B}} = H^i(\overline{C}_{\mathcal{B}}).$$

We use the same notation for smooth p-adic formal schemes. The functor φ^* denotes extension of scalars along Frobenius on A or A_0. For a k-scheme or stack Z, write $Z^{(1)}$ for its pullback along Frobenius on k.

Lemma 2.2. *Let M be a finite A_0-module and let $h \geq 0$. Suppose there are A_0-linear maps $\Phi : \varphi^* M \to M$ and $V : M \to \varphi^* M$ whose two composites are multiplication by u^h. If $p > h + 1$, then M is free over A_0.*

Proof. By [5, Lem. 2.2.1], M is torsion free, hence free over the discrete valuation ring A_0. □

Lemma 2.3. *Let $f : Y \to Z$ be a morphism over $W(k)$ between smooth Artin stacks of finite presentation, or between smooth p-adic formal schemes. For $i \geq 0$ there is a map $M^i_Z \to \varphi^* M^i_Z$, functorial in Z, whose two composites with the linearized Frobenius are multiplication by u^{i+1}. If $p > i + 2$, the image of $f^* : M^i_Z \to M^i_Y$ is free over A_0 whenever it is finitely generated. Taking $f = \mathrm{id}_Z$ gives freeness of any finite M^i_Z.*

Proof. Put $r = i + 1$. For a smooth p-adic formal $W(k)$-scheme U, apply $\tau_{\leq r} R\Gamma(U_{\text{ét}}, -)$ to the maps in [2, Cor. 15.5], using E^r as generator of $(E)^r$. This gives functorial maps

$$\tau_{\leq r}\varphi^* C_U \xrightarrow{\Phi} \tau_{\leq r} C_U \xrightarrow{V_r} \tau_{\leq r}\varphi^* C_U$$

whose two composites are multiplication by E^r.

When Z is an Artin stack, choose a smooth affine hypercover $U_\bullet \to Z$. Smooth descent gives $C_Z \simeq \operatorname{Tot} C_{U_\bullet}$ [9, Lem. 2.2.5]. Applying $\tau_{\le r}\operatorname{Tot}$ to the diagram for $\widehat{U}_\bullet$ gives the displayed diagram with C_U replaced by C_Z. Indeed, totalization preserves lower cohomological bounds and commutes with φ^* since $\varphi : A \to A$ is finite free.

Since $A_0 = A/p$ has Tor-amplitude $[-1, 0]$ over A, the map $\tau_{\le i+1}C \to C$ induces an isomorphism on $H^i(- \otimes^{\mathbf{L}}_A A_0)$ for any A-complex C. Thus reducing the diagram modulo p and taking H^i gives both identities, since $E^r \bmod p = u^r$. By functoriality and flatness of φ, the same identities hold on $\operatorname{im}(f^*)$. The freeness assertions follow from Lemma 2.2. □

2.4. Fix an algebraically closed complete extension C of $W(k)[1/p]$ with valuation ring $\mathcal{O}_C$ and residue field k_C. For a $W(k)$-scheme or stack Y, put $Y_C = Y \times_{W(k)} C$ and write $\widehat{Y}_C$ for the base change to C of the Raynaud generic fibre of its p-adic completion. For a C-scheme W, write W^{ad} for its analytification.

Definition 2.5. A smooth quasi-compact and quasi-separated Artin stack $\mathcal{B}$ over a noetherian ring R is *Hodge-proper* if, for every $j \ge 0$, $R\Gamma(\mathcal{B}, \bigwedge^j L_{\mathcal{B}/R})$ is bounded below with finitely generated cohomology over R [9, Defn. 1.2.1]. We use derived exterior powers and cohomological indexing.

Proposition 2.6. *Let $f : Y \to \mathcal{B}$ be a morphism over $W(k)$, where Y is a smooth scheme of finite presentation and $\mathcal{B}$ is a smooth proper Artin stack of finite presentation. Suppose that a class*

$$\omega \in H^0\left(\mathcal{B}_k^{(1)}, \bigwedge^d L_{\mathcal{B}_k^{(1)}/k}\right)$$

pulls back along $f_k^{(1)}$ to a nonzero d-form on an irreducible component of $Y_k^{(1)}$, with $d \ge 1$. If $p > d + 3$, then for every dense open $W \subset Y_C$, the map

$$H^d_{\text{ét}}(\mathcal{B}_C, \mathbb{F}_p) \longrightarrow H^d_{\text{ét}}(W, \mathbb{F}_p)$$

induced by $f_C|_W$ is nonzero.

Proof. Since $(f_k^{(1)})^*\omega$ remains nonzero after extension to k_C, we may base change to $W(k_C) \subset \mathcal{O}_C$ and assume $k = k_C$. As in [5, proof of Prop. 2.2.11], let D be the closure of $Y_C \setminus W$ in $Y_{\mathcal{O}_C}$. Since D_k is nowhere dense in Y_k [13, Tag 0B2J], we may choose an affine open $V \subset Y$ with $V_k \cap D_k = \varnothing$ and $(f_k^{(1)}|_{V_k^{(1)}})^*\omega \neq 0$. Then $\widehat{V}_C \subset W^{\mathrm{ad}}$.

Following [5, proof of Prop. 2.2.3], put

$$M = \operatorname{im}\big(M^d_{\mathcal{B}} \xrightarrow{(f|_V)^*} M^d_V\big).$$

As a smooth proper Artin stack over the regular noetherian ring $W(k)$, $\mathcal{B}$ is Hodge-proper [8, Props. 2.2.12 and 2.2.7]. By [9, Cor. 2.4.4], $C_{\mathcal{B}}$ has bounded below, finitely generated cohomology over A. Hence the $M^i_{\mathcal{B}}$ are finite over A_0, as is M. Since $p > d + 3$, we may apply Lemma 2.3 in degrees d and $d + 1$ to conclude that M and $M^{d+1}_{\mathcal{B}}$ are free over A_0.

We first show that M has positive rank. Give k its A_0-algebra structure by $u \mapsto 0$ and Frobenius on k. For $Z = \mathcal{B}, V$, reducing the de Rham comparison for smooth Artin stacks modulo p gives a functorial isomorphism $\overline{C}_Z \otimes^{\mathbf{L}}_{A_0} k \simeq R\Gamma_{\mathrm{dR}}(Z_k/k)$ [9, Thm. 0.3.2 and Prop. 2.3.6]. Since $M^{d+1}_{\mathcal{B}}$ has no u-torsion, the edge map $M^d_{\mathcal{B}} \otimes_{A_0} k \to H^d_{\mathrm{dR}}(\mathcal{B}_k/k)$ is an isomorphism. By functoriality of the edge map for $\widehat{V} \to \widehat{\mathcal{B}}$, restriction factors as

$$H^d_{\mathrm{dR}}(\mathcal{B}_k/k) \xrightarrow{\sim} M^d_{\mathcal{B}} \otimes_{A_0} k \longrightarrow M \otimes_{A_0} k \longrightarrow H^d_{\mathrm{dR}}(V_k/k).$$

For $Z = \mathcal{B}, V$, the $W_2(k)$-lift $Z \bmod p^2$ splits the conjugate filtration on $H^d_{\mathrm{dR}}(Z_k/k)$, since $d < p$ [8, Thms. 1.3.21 and 1.3.23]. Under the Deligne–Illusie isomorphism

$$H^d_{\mathrm{dR}}(Z_k/k) \simeq \bigoplus_{i=0}^{d} H^{d-i}\left(Z_k^{(1)}, \bigwedge^i L_{Z_k^{(1)}/k}\right),$$

restriction on the $i = d$ summand is pullback by $(f_k|_{V_k})^{(1)}$, by functoriality for $(f|_V) \bmod p^2$. This pullback is nonzero on ω, so M has positive rank.

It remains to detect this rank in étale cohomology. Let $\mathcal{O}_C^\flat$ be the tilt of $\mathcal{O}_C$ and $C^\flat = \mathrm{Frac}(\mathcal{O}_C^\flat)$. Choose compatible p-power roots of p and write $p^\flat$ for the resulting element of $\mathcal{O}_C^\flat$. Under $A \to A_{\mathrm{inf}} = W(\mathcal{O}_C^\flat)$, $u \mapsto [p^\flat]$ and $E \mapsto [p^\flat] - p$, a generator of $\ker(\theta : A_{\mathrm{inf}} \to \mathcal{O}_C)$ [1, Rem. 3.11]. For $Z = \mathcal{B}, V$, prismatic base change [9, Prop. 2.2.13] gives

$$D_Z := \left(R\Gamma_\Delta(\widehat{Z}_{\mathcal{O}_C}/A_{\mathrm{inf}}) \otimes^{\mathbf{L}}_{\mathbb{Z}_p} \mathbb{F}_p\right)[1/p^\flat] \simeq \left(\overline{C}_Z \widehat{\otimes}^{\mathbf{L}}_{A_0} \mathcal{O}_C^\flat\right)[1/p^\flat].$$

Here $\widehat{\otimes}^{\mathbf{L}}$ denotes the derived u-completed tensor product. It is exact on derived u-complete A_0-modules [5, Lem. 2.1.8] and agrees with ordinary base change on finite A_0-modules, as one sees by taking a finite free resolution. Hence $H^j(D_{\mathcal{B}}) = M_{\mathcal{B}}^j \otimes_{A_0} C^\flat$, and the image of $H^d(D_{\mathcal{B}}) \to H^d(D_V)$ is $M \otimes_{A_0} C^\flat \neq 0$.

By [9, Prop. 4.3.1], $R\Gamma_{\text{ét}}(\widehat{Z}_C, \mathbb{F}_p)$ is the homotopy fibre of $1 - \varphi : D_Z \to D_Z$. This gives a natural map

$$H^d_{\text{ét}}(\widehat{Z}_C, \mathbb{F}_p) \otimes_{\mathbb{F}_p} C^\flat \longrightarrow H^d(D_Z).$$

For $Z = \mathcal{B}$ this is an isomorphism, as in the proof of [5, Lem. 2.1.6]. Indeed, each $H^j(D_{\mathcal{B}})$ is finite-dimensional with bijective semilinear Frobenius by Lemma 2.3. Since $C^\flat$ is algebraically closed, $H^j(D_{\mathcal{B}})$ has a basis of fixed vectors and $1 - \varphi$ is surjective. Surjectivity in degree $d - 1$ identifies $H^d_{\text{ét}}(\widehat{\mathcal{B}}_C, \mathbb{F}_p)$ with $H^d(D_{\mathcal{B}})^{\varphi=1}$.

By functoriality, the nonzero map $H^d(D_{\mathcal{B}}) \to H^d(D_V)$ therefore gives a nonzero restriction from $\widehat{\mathcal{B}}_C$ to $\widehat{V}_C$ in étale cohomology. Properness identifies $H^d_{\text{ét}}(\mathcal{B}_C, \mathbb{F}_p)$ with $H^d_{\text{ét}}(\widehat{\mathcal{B}}_C, \mathbb{F}_p)$ [9, Cor. 3.3.16 and Rem. 4.1.15]. Since $\widehat{V}_C \subset W^{\mathrm{ad}}$, this restriction factors through $H^d_{\text{ét}}(W, \mathbb{F}_p)$, so the map in the proposition is nonzero, as required. □

3 Finite flat group schemes and specialization

3.1. Throughout this section k is a perfect field of characteristic p. For a finite flat group scheme $\mathcal{F}$ over $W(k)$ with special fibre H, we need a differential form on BH to apply Proposition 2.6 to $B\mathcal{F}$. A group scheme is called nilpotent if its scheme-theoretic lower central series terminates.

Lemma 3.2. *Let H be a finite nilpotent group scheme over k, and put $d = \dim_k \mathrm{Lie}(H)$. If $d < p$, there is a class*

$$\omega_H \in H^0(BH, \bigwedge^d L_{BH/k})$$

whose pullback along any smooth morphism $f : Y \to BH$ from a nonempty smooth k-scheme is a nonzero d-form on Y.

Proof. For a finite k-group scheme J with identity $e_J : \operatorname{Spec} k \to J$, let $\ell_J = e_J^* L_{J/k}$ be its co-Lie complex, with the J-action induced by conjugation, and put $V_J = H^{-1}(\ell_J)$. The local ring of J at e_J is a zero-dimensional complete intersection [11, Exp. VII$_\mathrm{B}$, Cor. 5.5.1]. Hence ℓ_J has Tor-amplitude in $[-1, 0]$ and virtual rank zero. As $H^0(\ell_J) = \mathrm{Lie}(J)^\vee$, we have $\dim_k V_J = \dim_k \mathrm{Lie}(J)$.

In degree zero Totaro's formula [14, Thm. 3.1] reads

$$H^0(BJ, \bigwedge^r L_{BJ/k}) \simeq H^{-r}(J, \mathrm{Sym}^r \ell_J) \qquad (r \geq 0),$$

where $\mathrm{Sym}^r \ell_J$ is the derived symmetric power with its induced J-action, and $H^a(J, -)$ denotes degree a cohomology of derived J-invariants. For $r < p$, the symmetrizing idempotent on $\ell_J^{\otimes r}$ acts on $H^{-r}(\ell_J^{\otimes r}) = V_J^{\otimes r}$ by antisymmetrization, since V_J occurs in degree -1. Thus $H^{-r}(\mathrm{Sym}^r \ell_J) \simeq \bigwedge^r V_J$ as J-modules. Since $\mathrm{Sym}^r \ell_J$ has no cohomology in lower degrees, taking derived J-invariants gives

$$H^0(BJ, \bigwedge^r L_{BJ/k}) \simeq (\bigwedge^r V_J)^J.$$

For $J = H$ and $r = d$, the right hand side is $(\det V_H)^H$. We claim that the conjugation action of H on V_H has a filtration with trivial successive quotients. This is clear for commutative H. Otherwise put $N = Z(H)$ and $Q = H/N$, and let $i : N \hookrightarrow H$ be the inclusion. Since $H \to Q$ is flat and $N = H \times_{Q, e_Q} \mathrm{Spec}\, k$, base change gives $\mathbf{L}i^* L_{H/Q} \simeq L_{N/k}$. Pulling back along e_N gives $e_H^* L_{H/Q} \simeq e_N^* L_{N/k} = \ell_N$. Thus the transitivity triangle for $H \to Q \to \mathrm{Spec}\, k$ pulls back along e_H to $\ell_Q \to \ell_H \to \ell_N \to \ell_Q[1]$. The resulting long exact cohomology sequence of H-modules contains

$$0 \longrightarrow V_Q \longrightarrow V_H \longrightarrow V_N \longrightarrow \mathrm{Lie}(Q)^\vee.$$

Conjugation on V_N is trivial, and the action on V_Q factors through Q, whose nilpotency class is smaller. Induction gives the asserted filtration. Thus H acts trivially on $\det V_H$. Let ω_H correspond to a generator of this line.

Since Y is nonempty and $|BH|$ is a point, f is a smooth surjection. A global d-form on BH is therefore a d-form on Y whose pullbacks under the two projections $Y \times_{BH} Y \rightrightarrows Y$ agree [14, §1], [8, Prop. 1.1.4]. In particular $f^* \omega_H \neq 0$. □

Lemma 3.3. *Let $\mathcal{F}$ be a finite flat group scheme over $W(k)$. Then $B\mathcal{F}$ is smooth, proper, and Hodge-proper over $W(k)$.*

Proof. The diagonal of $B\mathcal{F}$ is finite and $\mathrm{Spec}\, W(k) \to B\mathcal{F}$ is finite faithfully flat, so $B\mathcal{F}$ is proper. It is smooth by [9, Cor. A.1.19] applied to the action of $\mathcal{F}$ on $\mathrm{Spec}\, W(k)$. Hodge-properness now follows from [8, Props. 2.2.12 and 2.2.7]. □

Theorem 3.4. *Suppose that k is algebraically closed. Let $\mathcal{F}$ be a finite flat group scheme over $W(k)$ with nilpotent special fibre H, acting freely on a connected smooth affine $W(k)$-scheme $\mathcal{X}$ with nonempty special fibre. Put $\mathcal{Y} = \mathcal{X}/\mathcal{F}$ and $d = \dim_k \mathrm{Lie}(H)$. If $p > d + 3$, then for every algebraically closed complete extension C of $W(k)[1/p]$,*

$$\mathrm{ed}_p(\mathcal{X}_C/\mathcal{Y}_C) \geq d.$$

Proof. Assume $d > 0$, the case $d = 0$ being immediate. Since $\mathcal{F}$ is finite flat and acts freely on the affine scheme $\mathcal{X}$, the fppf quotient $\mathcal{Y}$ is a scheme and $\mathcal{X} \to \mathcal{Y}$ is an $\mathcal{F}$-torsor [13, Tag 07S7]. Let $f : \mathcal{Y} \to B\mathcal{F}$ classify this torsor. Its pullback along the faithfully flat cover $\mathrm{Spec}\, W(k) \to B\mathcal{F}$ is the smooth morphism $\mathcal{X} \to \mathrm{Spec}\, W(k)$. Hence f is smooth, and so is $\mathcal{Y}/W(k)$ by Lemma 3.3.

Apply Lemma 3.2 to $H^{(1)}$. The smooth map $f_k^{(1)}$ pulls $\omega_{H^{(1)}}$ back to a nonzero d-form on every component of $\mathcal{Y}_k^{(1)}$. By Lemma 3.3 and Proposition 2.6, the map

$$H^d_{\text{ét}}(B\mathcal{F}_C, \mathbb{F}_p) \longrightarrow H^d_{\text{ét}}(W, \mathbb{F}_p)$$

is nonzero for every dense open $W \subset \mathcal{Y}_C$.

As $\mathcal{X}$ is connected, so is $\mathcal{Y}$. Its smooth integral generic fibre is geometrically integral, since a k-point of $\mathcal{Y}_k$ lifts to $W(k)$ [13, Tags 0H74 and 0CDW]. Passing to the generic point of $\mathcal{Y}_C$ gives a nonzero map $H^d(\mathcal{F}(C), \mathbb{F}_p) \to H^d(\kappa(\mathcal{Y}_C), \mathbb{F}_p)$, since the source is finite dimensional. By the argument of [5, proof of Prop. 2.3.3], a torsor with this property has p-essential dimension at least d. □

Proof of Theorem B. Apply Theorem 3.4 with C the completion of $\overline{K}$. Extension of scalars from $\overline{K}$ to C preserves p-essential dimension by [4, 2.2.6]. □

4 Integral models of principal congruence covers

4.1. Finite flat group schemes in this section are not assumed to be commutative. Let k be a finite field of characteristic p and put $\mathcal{O} = W(k)$. Let $\mathcal{J}$ be a smooth affine group scheme over $\mathbb{Z}_p$ with connected fibres. We index weights so that $\mathbb{G}_m$ acts on the weight-i summand by $t \mapsto t^{-i}$. Fix a 1-*bounded* cocharacter $\lambda : \mathbb{G}_{m,\mathcal{O}} \to \mathcal{J}_{\mathcal{O}}$, that is, one whose adjoint weights on $\mathrm{Lie}(\mathcal{J}_{\mathcal{O}})$ are at most 1 [10, §1.7 and Setup 3.1.1].

4.2. For a p-complete $\mathcal{O}$-algebra A, write A^{syn} for the syntomification of $\operatorname{Spf} A$, a derived formal stack over $\operatorname{Spf} \mathbb{Z}_p$ [10, §2.1], and A^{syn}_n for its derived reduction modulo p^n ($n \geq 1$).

A $(\mathcal{J}, \lambda)$-*aperture* over A is a $\mathcal{J}$-torsor on A^{syn} of Hodge type λ. A *level-n $(\mathcal{J}, \lambda)$-aperture* is a $\mathcal{J}_{\mathbb{Z}/p^n\mathbb{Z}}$-torsor of the same Hodge type on A^{syn}_n. Hodge type λ means that at each geometric point of $\operatorname{Spf} A$ the associated graded of the Hodge filtration on the de Rham torsor admits a trivialization in which $\mathbb{G}_m$ acts by left multiplication through λ [10, Defn. 3.1.3].

Write $\mathrm{BT}_n^{\mathcal{J},\lambda}$ for the formal stack of level-n apertures and $\mathrm{BT}_\infty^{\mathcal{J},\lambda} = \varprojlim_n \mathrm{BT}_n^{\mathcal{J},\lambda}$ for the stack of apertures. The level-n truncation of an aperture Q is denoted $Q_n = Q|_{A^{\mathrm{syn}}_n}$. At any level a subscript B denotes base change along a map $A \to B$ of p-complete $\mathcal{O}$-algebras.

Lemma 4.3. *A level-n $(\mathcal{J}, \lambda)$-aperture Q over a p-complete $\mathcal{O}$-algebra A has an étale realization $Q_{\text{ét}}$, a $\mathcal{J}(\mathbb{Z}/p^n\mathbb{Z})$-torsor on $\operatorname{Spec} A[1/p]$, functorial in A and compatible with truncation and extension of structure group. For an aperture Q over A the torsors $(Q_n)_{\text{ét}}$ form a $\mathcal{J}(\mathbb{Z}_p)$-local system on $\operatorname{Spec} A[1/p]$.*

Proof. This is the functor $T_{\text{ét}}$ of [10, Constr. 3.8.1]. □

4.4. An *algebraic level-n $(\mathcal{J}, \lambda)$-aperture* over an $\mathcal{O}$-algebra A is a triple (Q, L, α). Here Q is a level-n $(\mathcal{J}, \lambda)$-aperture over the derived p-adic completion $\widehat{A}$, L is a $\mathcal{J}(\mathbb{Z}/p^n\mathbb{Z})$-torsor on $\operatorname{Spec} A[1/p]$, and

$$\alpha : Q_{\text{ét}} \xrightarrow{\sim} L|_{\operatorname{Spec} \widehat{A}[1/p]}.$$

The étale realization of $\mathcal{A} = (Q, L, \alpha)$ is $\mathcal{A}_{\text{ét}} = L$.

Lemma 4.5. *For finite n, the stack $\mathrm{BT}_n^{\mathcal{J},\lambda,\mathrm{alg}}$ of algebraic level-n apertures is a smooth Artin stack of relative dimension zero over $\mathcal{O}$ with affine diagonal of finite presentation. Its p-adic completion is $\mathrm{BT}_n^{\mathcal{J},\lambda}$ and its generic fibre is $B\big(\mathcal{J}(\mathbb{Z}/p^n\mathbb{Z})\big)$. These identifications commute with extension of structure group. For p-complete A, projection to Q induces an equivalence $\mathrm{BT}_n^{\mathcal{J},\lambda,\mathrm{alg}}(A) \simeq \mathrm{BT}_n^{\mathcal{J},\lambda}(A)$. Its inverse sends Q to its* algebraization $\widetilde{Q} = (Q, Q_{\text{ét}}, \mathrm{id})$.

Proof. Algebraicity and the assertion about the diagonal are [10, Thm. 4.1.3]. The formal and generic identifications follow from the description by triples [10, Constr. 4.1.1 and Rem. 4.1.2]. □

4.6. Suppose now that $\mathcal{J} = \mathcal{G}$ is reductive, so that λ is minuscule. Choose a maximal torus $\mathcal{T}_0$ in a Borel subgroup of $\mathcal{G}$, both over $\mathbb{Z}_p$, and take λ to be the dominant representative of its conjugacy class in $\mathcal{T}_{0,\mathcal{O}}$. For a cocharacter η of $\mathcal{T}_0$ over a $\mathbb{Z}_p$-algebra R, write $\mathfrak{g}_i(\eta)$ for the weight-i summand of $\mathrm{Lie}(\mathcal{G}_R)$ under η. Let $\mathcal{P}_\eta \subset \mathcal{G}_R$ be the parabolic containing $\mathcal{T}_{0,R}$ with Lie algebra $\bigoplus_{i\geq 0}\mathfrak{g}_i(\eta)$, and $\mathcal{U}_\eta$ its unipotent radical. Put $d_\lambda = \mathrm{rk}_{\mathcal{O}}\,\mathfrak{g}_1(\lambda)$, the relative dimension of $\mathcal{G}_{\mathcal{O}}/\mathcal{P}_\lambda$.

Let φ be arithmetic Frobenius on $X_*(\mathcal{T}_0)$ and put $f = [k:\mathbb{F}_p]$. The cocharacter

$$\nu = \sum_{a=0}^{f-1} \varphi^a(\lambda)$$

is defined over $\mathbb{Z}_p$ [10, Setup 3.6.2]. For any $\mathbb{Z}_p$-algebra R, the R-points of the *slope parabolic* $\mathcal{P}_\nu \subset \mathcal{G}$ are the $g \in \mathcal{G}(R)$ for which $t \mapsto \nu(t)^{-1} g \nu(t)$ extends from $\mathbb{G}_{m,R}$ to a morphism $\mathbb{A}^1_R \to \mathcal{G}_R$. Evaluation at $t = 0$ gives the Levi projection $\mathcal{P}_\nu \to \mathcal{M}_\nu = Z_{\mathcal{G}}(\nu)$ with kernel $\mathcal{U}_\nu$. Since the $\varphi^a(\lambda)$ lie in one closed Weyl chamber, every root orthogonal to ν is orthogonal to λ. Thus λ is central in $\mathcal{M}_{\nu,\mathcal{O}}$. It is 1-bounded in $\mathcal{P}_\nu$ by restriction from $\mathcal{G}$.

4.7. At finite or infinite level, a $(\mathcal{G},\lambda)$-aperture over an algebraically closed extension of k is *λ-ordinary* if it is induced along $\mathcal{M}_\nu \hookrightarrow \mathcal{G}$ from an $(\mathcal{M}_\nu,\lambda)$-aperture of the same level. Write $\mathrm{BT}_n^{\mathcal{G},\lambda,\mathrm{ord}} \subset \mathrm{BT}_n^{\mathcal{G},\lambda}$ for the locus of apertures Q over A such that Q_κ is λ-ordinary for every geometric point $\mathrm{Spec}\,\kappa \to \mathrm{Spf}\,A$.

For a level-n aperture Q over A, $\mathrm{Aut}(Q)$ is the sheaf on $\mathrm{Spf}\,A$ assigning to a p-nilpotent A-algebra B the automorphisms of the torsor Q_B on B_n^{syn}. For an algebraic aperture over A we take the automorphism sheaf on $\mathrm{Spec}\,A$. Given $Q_M \in \mathrm{BT}_n^{\mathcal{M}_\nu,\lambda}(A)$, put $Q_P = Q_M \times^{\mathcal{M}_\nu} \mathcal{P}_\nu$. Let $F_{\nu,Q_M}(n)$ be the kernel of $\mathrm{Aut}(Q_P) \to \mathrm{Aut}(Q_M)$ induced by $\mathcal{P}_\nu \to \mathcal{M}_\nu$. Then

$$F_{\nu,Q_M}(n)(B) = \Gamma\big(B_n^{\mathrm{syn}},\, Q_{M,B} \times^{\mathcal{M}_\nu} \mathcal{U}_\nu\big),$$

where $\mathcal{M}_\nu$ acts on $\mathcal{U}_\nu$ by conjugation (see [10, proof of Prop. 3.6.8, after (3.6.8.1)]). As Q_M varies, these kernels define a group sheaf $F_\nu(n)$ over $\mathrm{BT}_n^{\mathcal{M}_\nu,\lambda}$.

Lemma 4.8. *For finite n the Levi stack $\mathrm{BT}_n^{\mathcal{M}_\nu,\lambda}$ is étale over $\mathrm{Spf}\,\mathcal{O}$ and is noncanonically isomorphic to $B\big(\mathcal{M}_\nu(\mathbb{Z}/p^n\mathbb{Z})\big)$ over this base. Extension of structure group from $\mathcal{P}_\nu$ to $\mathcal{G}$ identifies $\mathrm{BT}_n^{\mathcal{P}_\nu,\lambda}$ with the open substack $\mathrm{BT}_n^{\mathcal{G},\lambda,\mathrm{ord}}$. Via the projection $\mathcal{P}_\nu \to \mathcal{M}_\nu$, the stack $\mathrm{BT}_n^{\mathcal{P}_\nu,\lambda}$ is the classifying stack of $F_\nu(n)$ over $\mathrm{BT}_n^{\mathcal{M}_\nu,\lambda}$. The section induced by $\mathcal{M}_\nu \hookrightarrow \mathcal{P}_\nu$ classifies the trivial $F_\nu(n)$-torsor. For an algebraically closed extension κ of k, reduction gives an equivalence*

$$\mathrm{BT}_\infty^{\mathcal{M}_\nu,\lambda}(W(\kappa)) \xrightarrow{\sim} \mathrm{BT}_\infty^{\mathcal{M}_\nu,\lambda}(\kappa).$$

Proof. The finite-level assertions are [10, Props. 3.6.8 and 3.6.12]. Formal étaleness gives the reduction equivalence at each finite level, hence also in the inverse limit. Extending the resulting Levi lift to $\mathcal{G}$ gives the canonical lift of a λ-ordinary aperture [10, Defn. 3.6.18]. □

Lemma 4.9. *For finite n, $F_\nu(n)$ is finite flat over $\mathrm{BT}_n^{\mathcal{M}_\nu,\lambda}$. The ν-weight filtration of $\mathcal{U}_\nu$ induces a central filtration on $F_\nu(n)$ by finite flat subgroups. For $i > 0$ $\mathrm{gr}^i F_\nu(n)$ is a truncated Barsotti–Tate group of level n with geometrically connected special fibres, and*

$$h_i := \mathrm{ht}\,\mathrm{gr}^i F_\nu(n) = \mathrm{rk}_{\mathbb{Z}_p}\,\mathfrak{g}_i(\nu), \qquad d_i := \mathrm{rk}\,\mathrm{Lie}\big((\mathrm{gr}^i F_\nu(n))_k\big) = \mathrm{rk}_{\mathcal{O}}\big((\mathfrak{g}_i(\nu)\otimes_{\mathbb{Z}_p}\mathcal{O}) \cap \mathfrak{g}_1(\lambda)\big)$$

with the intersection taken in $\mathrm{Lie}(\mathcal{G}_{\mathcal{O}})$. In particular, $\sum_{i>0} h_i = \dim \mathcal{U}_{\nu,\mathbb{F}_p}$ and $\sum_{i>0} d_i = d_\lambda$.

Proof. The filtration and its Barsotti–Tate quotients are constructed in [10, Rem. 3.6.6, Lem. 3.6.7 and Prop. 3.6.8]. For $i \geq 1$, write $\mathcal{U}_{\geq i} \subset \mathcal{U}_\nu$ for the subgroup generated, after an unramified splitting extension, by the root groups of ν-weight at least i. For a p-complete $\mathcal{O}$-algebra A and $Q_M \in \mathrm{BT}_n^{\mathcal{M}_\nu,\lambda}(A)$, put $F = F_{\nu,Q_M}(n)$. The induced filtration on F is given by

$$F_{\geq i}(B) = \Gamma\big(B_n^{\mathrm{syn}},\, Q_{M,B} \times^{\mathcal{M}_\nu} \mathcal{U}_{\geq i}\big)$$

for every p-nilpotent A-algebra B. The commutator map $\mathcal{U}_{\geq i} \times \mathcal{U}_{\geq j} \to \mathcal{U}_{\geq i+j}$ is $\mathcal{M}_\nu$-equivariant, so taking sections gives $[F_{\geq i}, F_{\geq j}] \subset F_{\geq i+j}$. Since $F_{\geq 1} = F$, the filtration is central.

Over an algebraically closed extension κ of k, the proof of [10, Lem. 3.6.7] shows that the i-th quotient is connected and that its Dieudonné module has underlying $W_n(\kappa)$-module $W_n(\kappa) \otimes_{\mathbb{Z}_p} \mathfrak{g}_i(\nu)$. The height of the quotient is the rank of this module, and its dimension is the multiplicity of the Hodge character $t \mapsto t^{-1}$ of λ on $\mathfrak{g}_i(\nu) \otimes_{\mathbb{Z}_p} \mathcal{O}$ [7, Thm. 11.3.3]. The first sum is the $\mathbb{Z}_p$-rank of $\mathrm{Lie}(\mathcal{U}_\nu)$. For the second, each root of λ-weight one has positive ν-weight since the conjugates of λ lie in one closed Weyl chamber. □

4.10. Let R be a k-algebra, and let F_R denote the absolute Frobenius of $\operatorname{Spec} R$. A *$\mathcal{G}$-zip of type λ* over R consists of a $\mathcal{G}_R$-torsor E, reductions $E^- \subset E$ and $E^+ \subset E$ to $\mathcal{P}_{-\lambda,R}$ and $\mathcal{P}_{\varphi(\lambda),R}$, respectively, and an isomorphism

$$\alpha : F_R^*(E^-/\mathcal{U}_{-\lambda,R}) \xrightarrow{\sim} E^+/\mathcal{U}_{\varphi(\lambda),R}$$

of torsors under the centralizer $Z_{\mathcal{G}_R}(\varphi(\lambda))$ [10, Defn. 3.5.2]. For a level-one $(\mathcal{G},\lambda)$-aperture Q over R, write Z_Q for its associated $\mathcal{G}$-zip, whose reductions E^- and E^+ are the Hodge and conjugate parabolic reductions [10, Prop. 3.5.5].

Lemma 4.11. *Let Q be a level-one $(\mathcal{G},\lambda)$-aperture over an algebraically closed extension κ of k. Then $L_Q = \ker(\mathrm{Aut}(Q) \to \mathrm{Aut}(Z_Q))$ is a finite commutative κ-group scheme killed by relative Frobenius, with $\dim_\kappa \mathrm{Lie}(L_Q) = d_\lambda$.*

Proof. Let $\mathscr{Z}$ be the stack of $\mathcal{G}$-zips of type λ over k. By [7, Thm. E and proof of Thm. 9.3.2], the morphism $\pi : (\mathrm{BT}_1^{\mathcal{G},\lambda})_k \to \mathscr{Z}$ is a gerbe banded by a finite flat commutative group scheme $\mathscr{L}$ over $\mathscr{Z}$, called the Lau group scheme, which is killed by relative Frobenius. For the point $z_Q : \operatorname{Spec}\kappa \to \mathscr{Z}$ represented by Z_Q, the banding identifies L_Q with $z_Q^*\mathscr{L}$.

Let E^- be the $\mathcal{P}_{-\lambda,\kappa}$-reduction in Z_Q, and put

$$N = \big(E^- \times^{\mathcal{P}_{-\lambda,\kappa}} \big(\mathrm{Lie}(\mathcal{G}_\kappa)/\mathrm{Lie}(\mathcal{P}_{-\lambda,\kappa})\big)\big)^\vee.$$

The F-zip associated to Z_Q by the cotangent representation $\mathrm{Lie}(\mathcal{G}_\kappa)^\vee[-1]$ has first Hodge-filtration step $N[-1]$. The construction of $\mathscr{L}$ in [7, Constr. 8.5.4 and Thm. 8.5.5], together with [7, Constr. 7.2.1 and Cor. 7.2.3], gives $\mathrm{rk}_\kappa L_Q = p^{\dim_\kappa N}$. Since L_Q is killed by relative Frobenius, $\mathrm{rk}_\kappa L_Q = p^{\dim_\kappa \mathrm{Lie}(L_Q)}$. Hence

$$\dim_\kappa \mathrm{Lie}(L_Q) = \dim_\kappa N = \dim(\mathcal{G}_\kappa/\mathcal{P}_{-\lambda,\kappa}) = d_\lambda.$$

□

Lemma 4.12. *Let Q_M be a level-one $(\mathcal{M}_\nu,\lambda)$-aperture over an algebraically closed extension κ of k, and put $H = F_{\nu,Q_M}(1)$. Then H is infinitesimal and nilpotent, and $\dim_\kappa \mathrm{Lie}(H) = d_\lambda$.*

Proof. By Lemma 4.9, H has a central filtration with finite connected successive quotients. Hence H is nilpotent and infinitesimal. Left exactness of Lie gives

$$\dim_\kappa \mathrm{Lie}(H) \le \sum_{i>0} d_i = d_\lambda.$$

For the reverse inequality, put $Q = Q_M \times^{\mathcal{M}_\nu} \mathcal{G}$. By Lemma 4.8, H is the kernel of $\mathrm{Aut}(Q) \to \mathrm{Aut}(Q_M)$. The target is étale and L_Q is connected, so $L_Q \subset H$. Lemma 4.11 gives $d_\lambda = \dim_\kappa \mathrm{Lie}(L_Q) \le \dim_\kappa \mathrm{Lie}(H)$. □

4.13. Let (G, X) and $\mathcal{G}$ be as in Section 1, and fix a place $v \mid p$ of the reflex field E. Since $E_v/\mathbb{Q}_p$ is unramified, we take $\mathcal{O} = \mathcal{O}_{E_v}$ and k to be its residue field.

A $\mathbb{Q}$-torus is *cuspidal* if it is isogenous to the product of a $\mathbb{Q}$-split torus and a torus anisotropic over $\mathbb{R}$. Let $Z(G)^{\mathrm{ac}}$ be the smallest $\mathbb{Q}$-subtorus $Z' \subset Z(G)^\circ$ such that $Z(G)^\circ/Z'$ is cuspidal, and put $G_c = G/Z(G)^{\mathrm{ac}}$ [10, Defns. 6.1.1 and 6.1.4]. The quotient of $\mathcal{G}$ by the schematic closure of $Z(G)^{\mathrm{ac}}_{\mathbb{Q}_p}$ is a reductive $\mathbb{Z}_p$-model $\mathcal{G}_c$ of $G_{c,\mathbb{Q}_p}$ [10, Setup 6.3.2].

Apply 4.6–4.12 to $\mathcal{G}_c$ with λ the dominant representative of the image of $-\mu_v$ in G_c, where μ_v is the Hodge cocharacter at v [10, Setup 6.4.1]. As $G^{\mathrm{ad}} = G_c^{\mathrm{ad}}$, the compact dual of X has dimension $d_\lambda = \dim S = d$. For this pair we call a λ-ordinary aperture μ-ordinary. To construct the free action required by Theorem 3.4, we compare at level one the aperture on the integral model of S in Hypothesis 4.15 with the canonical lift of its fibre at a μ-ordinary point.

Take K^p neat, and put $K = K^pK_p$, $K' = K^pK_p(p)$, $Z_K = Z(G)(\mathbb{Q}) \cap K$ and $\Delta = K/(K'Z_K)$.

Lemma 4.14. *The cover $S(p) \to S$ is a torsor under Δ. The reduction modulo p of the canonical $\mathcal{G}_c(\mathbb{Z}_p)$-local system on S is the $\mathcal{G}_c(\mathbb{F}_p)$-torsor*

$$T = S(p) \times^\Delta \mathcal{G}_c(\mathbb{F}_p),$$

where $\Delta \to \mathcal{G}_c(\mathbb{F}_p)$ is induced by projection to K_p followed by $\mathcal{G}(\mathbb{Z}_p) \to \mathcal{G}_c(\mathbb{F}_p)$. In particular $\mathrm{ed}_p(T/S) \le \mathrm{ed}_p(S(p)/S)$.

Proof. By [10, Lem. 6.1.6 and Constr. 6.1.7], the Shimura tower is a pro-torsor under $K/\overline{Z_K}$, with finite quotient $S(p)$ of group Δ. Neatness gives $Z_K \subset Z(G)^{\mathrm{ac}}(\mathbb{Q})$, so the composite $K \to \mathcal{G}(\mathbb{Z}_p) \to \mathcal{G}_c(\mathbb{Z}_p)$ factors through $K/\overline{Z_K}$. Extension of structure group by this map gives the canonical local system. Modulo p the map kills K', giving the stated homomorphism from Δ. Thus T is the quotient of $S(p) \times \mathcal{G}_c(\mathbb{F}_p)$ by $(s, g) \cdot \delta = (s\delta, \bar\delta^{-1}g)$, where $\bar\delta$ is the image of δ. Any descent of the Δ-torsor induces one of T. □

Hypothesis 4.15. The variety S admits a smooth integral canonical model $\mathscr{S}/\mathcal{O}_{E,(v)}$ in the sense of [10, Defn. 6.5.2]. Write $\widehat{\mathscr{S}}$ for its p-adic completion. The aperture $\mathscr{Q}$ of the integral model is a $\mathcal{G}_c$-torsor of Hodge type λ on $(\widehat{\mathscr{S}})^{\mathrm{syn}}$. Its étale realization on the generic fibre of $\widehat{\mathscr{S}}$ is identified with the restriction of the canonical local system of Lemma 4.14. Glue $\mathscr{Q}_1$ to $T|_{S_{E_v}}$ by this identification modulo p, and let

$$q : \mathscr{S}_{\mathcal{O}_{E_v}} \longrightarrow \mathrm{BT}_1^{\mathcal{G}_c,\lambda,\mathrm{alg}}$$

classify the resulting algebraic aperture $\widetilde{\mathscr{Q}}_1$. The p-adic completion of q is smooth, and the special fibre of $\mathscr{S}$ contains a μ-ordinary point.

4.16. For a complete discretely valued extension F/E_v with perfect residue field, a point $y \in S(F)$ extends to an $\mathcal{O}_F$-point of $\mathscr{S}$ if and only if the canonical $\mathcal{G}_c(\mathbb{Z}_p)$-local system at y is crystalline [10, Defn. 1.1.4(2)]. The locus where $\mathscr{Q}$ is μ-ordinary is open and dense in the geometric special fibre of $\mathscr{S}$ [10, Thm. I(1)].

Proposition 4.17. *For (G, X) of pre-abelian type, Hypothesis 4.15 holds at every good prime $p > 2$ if K^p is sufficiently small. For an arbitrary fixed Shimura datum it holds at all sufficiently large good primes with sufficiently small prime-to-p level.*

Proof. The integral models exist by [10, Thm. C] in the pre-abelian case and by [10, Thm. D] for p sufficiently large. In both cases [10, Thm. H] gives smoothness and surjectivity of $\widehat{\mathscr{S}} \to \mathrm{BT}_1^{\mathcal{G}_c,\lambda}$. Its image therefore meets the μ-ordinary locus. □

4.18. Assume Hypothesis 4.15. Let k now denote an algebraically closed extension of the residue field at v, and put $K_0 = W(k)[1/p]$. Use the unique lift $\mathcal{O}_{E_v} \hookrightarrow W(k)$ of the residue field embedding. Fix a μ-ordinary point $x \in \mathscr{S}(k)$ and write $\mathscr{Q}_x = x^*\mathscr{Q}$ for the aperture over k.

Choose an $(\mathcal{M}_\nu, \lambda)$-aperture Q_M^0 over k and an isomorphism $Q_M^0 \times^{\mathcal{M}_\nu} \mathcal{G}_c \simeq \mathscr{Q}_x$. Lemma 4.8 gives a unique lift Q_M over $W(k)$ with specified reduction $Q_{M,k} \simeq Q_M^0$. Set

$$Q_P = Q_M \times^{\mathcal{M}_\nu} \mathcal{P}_\nu, \qquad Q^{\mathrm{can}} = Q_M \times^{\mathcal{M}_\nu} \mathcal{G}_c, \qquad \mathcal{F} = F_{\nu, Q_{M,1}}(1).$$

Thus Q^{can} is the canonical lift of $\mathscr{Q}_x$. By Lemma 4.9, $\mathcal{F}$ is finite flat over $\operatorname{Spf} W(k)$. Its coordinate Hopf algebra is finite free over $W(k)$ and defines a finite flat group scheme over $\operatorname{Spec} W(k)$, also denoted by $\mathcal{F}$.

Let $U \to \mathscr{S}_{W(k)}$ be an étale neighborhood of x with a chosen lift of x to $U(k)$. Pull back $\widetilde{\mathscr{Q}}_1$ from $\mathscr{S}_{\mathcal{O}_{E_v}}$ and $\widetilde{Q}_1^{\mathrm{can}}$ from $\operatorname{Spec} W(k)$ to U, and put

$$\mathcal{I} = \mathrm{Isom}_U\big(\widetilde{Q}_{1,U}^{\mathrm{can}}, \widetilde{\mathscr{Q}}_{1,U}\big).$$

Proposition 4.19. *One can choose U connected and affine so that $\mathcal{I}$ is affine and smooth over $W(k)$ and has a k-point above x. Precomposition gives a free right action of $\mathcal{F}$ on $\mathcal{I}$. For an algebraically closed extension C/K_0, a trivialization of $\widetilde{Q}_{1,\text{ét}}^{\mathrm{can}}$ over C identifies $\mathcal{I}_C \to U_C$ with the pullback of the canonical $\mathcal{G}_c(\mathbb{F}_p)$-torsor $T \to S$.*

Proof. Put $\mathcal{B} = (\mathrm{BT}_1^{\mathcal{G}_c,\lambda,\mathrm{alg}})_{W(k)}$. The map $q_{W(k)}|_U : U \to \mathcal{B}$ classifies $\widetilde{\mathscr{Q}}_{1,U}$. Its fibre over the $W(k)$-point of $\mathcal{B}$ defined by $\widetilde{Q}_1^{\mathrm{can}}$ is

$$\mathcal{I} = U \times_{\mathcal{B}} \operatorname{Spec} W(k).$$

The locally finitely presented morphism $q_{W(k)}$ has smooth p-adic completion by Hypothesis 4.15, so it is smooth near x by the formal criterion for smoothness. Choose U connected and affine in the smooth locus. Then $\mathcal{I}$ is smooth over $W(k)$ by base change and affine over U since $\mathcal{B}$ has affine diagonal. Truncating $Q_k^{\mathrm{can}} \simeq \mathscr{Q}_x$ at level one gives the required k-point.

Over $\operatorname{Spf} W(k)$, extension of structure group identifies $\mathrm{Aut}(Q_{P,1})$ with $\mathrm{Aut}(Q_1^{\mathrm{can}})$ by Lemma 4.8. The kernel $\mathcal{F}$ of $\mathrm{Aut}(Q_{P,1}) \to \mathrm{Aut}(Q_{M,1})$ is closed, since the target is finite étale. Put $A = \mathcal{O}(\mathcal{F})$. The universal element of $\mathcal{F}$ over $\operatorname{Spf} A$ acts on $Q_{P,1,A}$ and induces the identity on $Q_{M,1,A}$. Both A and $A \otimes_{W(k)} A$ are p-adically complete. By Lemma 4.5, algebraizing this automorphism over $\operatorname{Spec} A$ gives a map $\mathcal{F} \to \mathrm{Aut}(\widetilde{Q}_{P,1})$. Full faithfulness over $A \otimes_{W(k)} A$ shows that the map respects multiplication. The induced action on $\widetilde{Q}_{M,1}$ is trivial. Extension to $\mathcal{G}_c$ gives a homomorphism $\mathcal{F} \to \mathrm{Aut}(\widetilde{Q}_1^{\mathrm{can}})$. The map on coordinate rings is surjective modulo p by the closed immersion on formal completions, hence surjective by Nakayama's lemma since A is finite over $W(k)$. For any U-scheme V an isomorphism $\alpha \in \mathcal{I}(V)$ thus satisfies $\alpha \circ g = \alpha$ only for $g = 1$ in $\mathcal{F}(V)$. This proves freeness.

By Hypothesis 4.15 the étale realization of $\widetilde{\mathscr{Q}}_{1,U_C}$ is $T|_{U_C}$. Since $\mathcal{B}_C = B\big(\mathcal{G}_c(\mathbb{F}_p)\big)$, evaluation at $t \in \widetilde{Q}^{\mathrm{can}}_{1,\text{ét}}(C)$ gives

$$\mathcal{I}_C \xrightarrow{\sim} T|_{U_C}, \qquad \alpha \longmapsto \alpha_{\text{ét}}(t).$$

For $g \in \mathcal{F}(C)$ let $\bar{g} \in \mathcal{G}_c(\mathbb{F}_p)$ be determined by $g_{\text{ét}}(t) = t\bar{g}$. Evaluation carries precomposition by g to right multiplication by $\bar{g}$. □

Lemma 4.20. *A point of the Levi torsor $\widetilde{Q}_{M,1,\text{ét}}$ over an algebraically closed extension C/K_0 identifies $\mathcal{F}(C)$ with the subgroup*

$$\mathcal{U}_\nu(\mathbb{F}_p) = \ker\big(\mathcal{P}_\nu(\mathbb{F}_p) \longrightarrow \mathcal{M}_\nu(\mathbb{F}_p)\big)$$

of $\mathcal{G}_c(\mathbb{F}_p)$. In particular $\mathrm{rk}_{W(k)}\,\mathcal{F} = p^{\dim \mathcal{U}_{\nu,\mathbb{F}_p}}$.

Proof. The Levi point trivializes the étale realizations of $Q_{P,1}$ and Q_1^{can} by extension of structure group. Since the action of $\mathcal{F}$ on $\widetilde{Q}_{P,1}$ induces the identity on $\widetilde{Q}_{M,1}$, Proposition 4.19 gives an injection of $\mathcal{F}(C)$ into the displayed kernel. The generic fibre of $\mathcal{F}$ is étale, so Lemma 4.9 gives

$$|\mathcal{F}(C)| = \mathrm{rk}_{W(k)}\,\mathcal{F} = p^{\sum_{i>0} h_i} = p^{\dim \mathcal{U}_{\nu,\mathbb{F}_p}}.$$

Since the special fibre of $\mathcal{U}_\nu$ is split unipotent, its group of $\mathbb{F}_p$-points has the same order. Thus the injection is onto. □

4.21. Translating the chosen point of $\widetilde{Q}_{M,1,\text{ét}}(C)$ on the right by $m \in \mathcal{M}_\nu(\mathbb{F}_p)$ conjugates the isomorphism $\mathcal{F}(C) \simeq \mathcal{U}_\nu(\mathbb{F}_p)$ by m^{-1}. Thus over K_0 one has

$$\mathcal{F}_{K_0} \simeq \widetilde{Q}_{M,1,\text{ét}} \times^{\mathcal{M}_\nu(\mathbb{F}_p)} \mathcal{U}_\nu(\mathbb{F}_p).$$

Let $\pi : \mathcal{G} \to \mathcal{G}_c$ be the quotient morphism and put $\mathcal{P} = \pi^{-1}(\mathcal{P}_\nu) \subset \mathcal{G}$. Since $\ker \pi$ is central, π identifies the unipotent radical $\mathcal{U}$ of $\mathcal{P}$ with $\mathcal{U}_\nu$. Hence $\mathcal{F}(C)$ identifies with the subgroup $\mathcal{U}(\mathbb{F}_p) \subset \mathcal{G}(\mathbb{F}_p)$, up to conjugacy.

In the ordinary case $\varphi(\lambda) = \lambda$, so $\mathcal{P}_\nu = \mathcal{P}_\lambda$. Every root of $\mathcal{U}_\lambda$ has λ-weight one, so the filtration of Lemma 4.9 is concentrated in a single weight. Thus $\mathcal{F}$ is a level-one Barsotti–Tate group of height and dimension d. Its Cartier dual is étale and hence constant over the strictly henselian ring $W(k)$, so $\mathcal{F} \simeq \mu_p^d$. For the Siegel datum of genus g, $\mathcal{P}_\lambda$ is the Siegel parabolic. In a symplectic basis adapted to the slope filtration,

$$\mathcal{U}_\lambda(\mathbb{F}_p) = \left\{ \begin{pmatrix} I_g & B \\ 0 & I_g \end{pmatrix} \;\middle|\; B \in \mathrm{Mat}_g(\mathbb{F}_p),\ B = B^{\mathsf{t}} \right\}.$$

For $g = 1$ the parabolic is a Borel and $\mathcal{F}(C)$ identifies with its order-p unipotent subgroup.

5 p-incompressibility of principal congruence covers

Theorem 5.1. *Assume Hypothesis 4.15 at a good prime p. If $p > \dim S + 3$, then for every geometric connected component S^0 of S,*

$$\mathrm{ed}_p(S(p)|_{S^0}/S^0) \geq \dim S.$$

Proof. Fix $C = \mathbb{C}_p$, $k = \overline{\mathbb{F}}_p$, $K_0 = W(k)[1/p]$, with the embeddings determined by v. The cover and the geometric connected components of S are defined over $\overline{\mathbb{Q}}$, so we may compute the p-essential dimension on each component over C [4, 2.2.6].

Fix a component S^0 of S_C. It contains a CM point, which extends to $\mathscr{S}$ by [10, Lem. 7.1.9]. By smoothness, S^0 is therefore the geometric generic fibre of a connected component $\mathscr{S}^0$ of $\mathscr{S}_{W(k)}$ with nonempty special fibre. The μ-ordinary locus is dense in $\mathscr{S}^0_k$ by [10, Thm. I(1)]; choose $x \in \mathscr{S}^0(k)$ in this locus.

Apply Proposition 4.19 at x, and let $\mathcal{I}_0$ be the connected component of $\mathcal{I}$ containing the distinguished k-point above x. By Lemma 4.12, $\mathcal{F}_k$ is infinitesimal and nilpotent, with $\dim_k \mathrm{Lie}(\mathcal{F}_k) = d$; in particular, $\mathcal{F}$ preserves $\mathcal{I}_0$. Applying Theorem 3.4 to $\mathcal{I}_0$ and using Proposition 4.19 and Lemma 4.14 gives

$$d \leq \mathrm{ed}_p(T|_{S^0}/S^0) \leq \mathrm{ed}_p(S(p)|_{S^0}/S^0).$$

□

Corollary 5.2. *For a Shimura datum of pre-abelian type, the principal congruence cover is geometrically p-incompressible at each good prime $p > \dim S + 3$ with sufficiently small prime-to-p level.*

Proof. Since $p > 2$, Proposition 4.17 supplies the integral model required by Theorem 5.1. □

Proof of Theorem A. The pre-abelian case is Corollary 5.2. For a fixed Shimura datum, take p sufficiently large for Proposition 4.17 and for $p > \dim S + 3$, and apply Theorem 5.1. □

Remark 5.3. The further restriction on p for data outside pre-abelian type enters through the construction of the integral model by spreading out in [10, Thm. D and Rem. 1.1.8], as in the work of Bakker–Shankar–Tsimerman [3]. This construction gives no effective bound on the excluded primes.

References

[1] B. Bhatt, M. Morrow, P. Scholze, *Integral p-adic Hodge theory*, Publ. Math. Inst. Hautes Études Sci. 128 (2018), 219–397. arXiv:1602.03148v3.

[2] B. Bhatt, P. Scholze, *Prisms and prismatic cohomology*, Annals of Math. 196 (2022), 1135–1275. arXiv:1905.08229.

[3] B. Bakker, A. Shankar, J. Tsimerman, *Integral canonical models of exceptional Shimura varieties*. arXiv:2405.12392.

[4] B. Farb, M. Kisin, J. Wolfson, *The essential dimension of congruence covers*, Compositio Math. 157 (2021), 2407–2432. arXiv:1901.09013.

[5] B. Farb, M. Kisin, J. Wolfson, *Essential dimension via prismatic cohomology*, Duke Math. J. 173 (2024), 3059–3106. arXiv:2110.05534.

[6] N. Fakhruddin, R. Saini, *Finite groups scheme actions and incompressibility of Galois covers: beyond the ordinary case*, Documenta Math. 27 (2022), 151–182. doi:10.4171/DM/868.

[7] Z. Gardner, K. Madapusi, *An algebraicity conjecture of Drinfeld and the moduli of p-divisible groups*, version 4 (20 April 2026). arXiv:2412.10226.

[8] D. Kubrak, A. Prikhodko, *Hodge-to-de Rham degeneration for stacks*, International Mathematics Research Notices 2022, 12852–12939. The numbering used here is that of arXiv version 2. arXiv:1910.12665v2.

[9] D. Kubrak, A. Prikhodko, *p-adic Hodge theory for Artin stacks*, Memoirs of the American Mathematical Society 304, no. 1528 (2024). The numbering used here is that of the 2021 arXiv version. arXiv:2105.05319.

[10] K. Madapusi, A. Youcis, *On canonicity for integral models of Shimura varieties with hyperspecial level*, 2026 preprint, version 2 (27 April 2026). arXiv:2604.06442.

[11] *Schémas en groupes* (SGA 3), tome I, M. Demazure and A. Grothendieck (eds.), revised edition by P. Gille and P. Polo, Documents Mathématiques 7, Société Mathématique de France, 2011. Exposé $\mathrm{VII}_{\mathrm{B}}$.

[12] A. N. Shankar, R. Zhou, *Serre–Tate theory for Shimura varieties of Hodge type*, Math. Z. 297 (2021), 1249–1271. arXiv:1612.06456.

[13] The Stacks Project Authors, *The Stacks Project*. Tag 07S7.

[14] B. Totaro, *Hodge theory of classifying stacks*, Duke Math. J. 167 (2018), 1573–1621. doi:10.1215/00127094-2018-0003.