Contents

1 Notes on shtukas and Harris's conjecture[1]

1.1 Motivation

Fix a finite extension $E/\mathbf{Q}_p$ with uniformizer π and residue field $\mathbf{F}_q$. Set $\Gamma = \mathrm{Gal}(\overline{E}/E)$ and $\breve{E} = \widehat{E^{\mathrm{unr}}}$, and let $\sigma \in \mathrm{Aut}(\breve{E}/E)$ be the natural q-Frobenius.

Fix a reductive group $\mathbf{G}/E$, and set $G = \mathbf{G}(E)$. For simplicity I'll assume $\mathbf{G}$ is *quasisplit.* Fix, then, a maximal torus $\mathbf{T} \subset \mathbf{G}$ and a Borel subgroup $\mathbf{B} \supset \mathbf{T}$ both defined over E. The group Γ acts on $X_*(\mathbf{T})$, and we have a natural dominant cone $X_*(\mathbf{T})_{\mathrm{dom}}$.

Now consider triples $(\mathbf{G}, \mu, b)$ where $\mu \in X_*(\mathbf{T})_{\mathrm{dom}}$ is a $\mathbf{B}$-dominant cocharacter and $b \in \mathbf{G}(\breve{E})$ is an element whose Kottwitz class satisfies $[b] \in B(\mathbf{G}, \mu^{-1})$.[2] Then we have a moduli space of local shtukas with infinite level structure $\mathrm{Sht}_{\mathbf{G},\mu,b}$, defined as a certain functor

$$\mathrm{Sht}_{\mathbf{G},\mu,b} : \mathrm{Perf}_{/\breve{E}} \to \mathrm{Sets}.$$

For any $S \in \mathrm{Perf}_{/\breve{E}}$, the set $\mathrm{Sht}_{\mathbf{G},\mu,b}(S)$ consists of isomorphism classes of triples $(\mathcal{F}, u, \alpha)$ where

- $\mathcal{F}$ is a $\mathbf{G}$-bundle over $\mathcal{X}_{S^\flat}$,
- $u : \mathcal{F}|_{\mathcal{X}_{S^\flat} \smallsetminus S} \xrightarrow{\sim} \mathcal{E}_{b,S^\flat}|_{\mathcal{X}_{S^\flat} \smallsetminus S}$ is a $\mathbf{G}$-bundle isomorphism extending to a type-μ modification of the $\mathbf{G}$-bundle $\mathcal{E}_{b,S^\flat}/\mathcal{X}_{S^\flat}$, and
- $\alpha : \mathcal{E}_{\mathrm{triv},S^\flat} \xrightarrow{\sim} \mathcal{F}$ is a $\mathbf{G}$-bundle isomorphism.

Then $\mathrm{Sht}_{\mathbf{G},\mu,b}$ has natural commuting actions of $G = \mathbf{G}(E)$ and $\mathcal{J}_b : S \mapsto \mathrm{Aut}(\mathcal{E}_{b,S^\flat})$, and there are two period morphisms sitting in a diagram

$$\begin{tikzcd} \mathrm{Sht}_{\mathbf{G},\mu,b} \arrow[r, "\pi_{\mathrm{HT}}"] \arrow[d, "\pi_{\mathrm{GM}}"] & \mathrm{Gr}_{\mathbf{G},\mu^{-1}} \\ \mathrm{Gr}_{\mathbf{G},\mu} & \end{tikzcd}$$

where the map π_{GM} forgets α, and the map π_{HT} records $(u \circ \alpha^{-1})$ as giving a μ^{-1}-bounded modification of $\mathcal{E}_{\mathrm{triv},S^\flat}$ along S. Each object in this diagram is a diamond over $\breve{E}$. The targets of π_{GM} and π_{HT} have natural actions of $\mathcal{J}_b$ and G, respectively, and the period maps are equivariant for these commuting actions on Sht.

[1] Version of 5/24/2016; comments and corrections to hansen@math.columbia.edu

[2] By a recipe of Kottwitz, b gives rise to a canonical cocharacter $\nu_b \in X_*(\mathbf{G}_{\breve{E}}) \otimes \mathbf{Q}$ such that $\nu_{hb\sigma(h)^{-1}} = h\nu_b h^{-1}$; this cocharacter can be conjugated uniquely to a Galois-invariant and $\mathbf{B}$-dominant cocharacter factoring through $\mathbf{T}$, which we denote by $\nu_{[b]} \in X_*(\mathbf{T})_{\mathbf{Q},\mathrm{dom}}$. Let us say that b is *well-chosen* if $\nu_b = \nu_{[b]}$ and b is contained in the standard Levi $\mathbf{M}_{[b]} = \mathrm{Cent}(\nu_{[b]})$. Every class $[b] \in B(\mathbf{G})$ has a well-chosen representative.

The image of π_{GM} is always open and partially proper inside its target; this is the so-called admissible locus, denoted $\mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}}$, and defined intrinsically as follows: For any $S \in \mathrm{Perf}_{/\breve{E}}$, the set $\mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}}$ consists of the set of isomorphism classes of pairs $(\mathcal{F}, u)$ where

$$u : \mathcal{F}|_{\mathcal{X}_{S^\flat} \smallsetminus S} \overset{\sim}{\to} \mathcal{E}_{b,S^\flat}|_{\mathcal{X}_{S^\flat} \smallsetminus S}$$

is a $\mathbf{G}$-bundle isomorphism extending to a type-μ modification of the $\mathbf{G}$-bundle $\mathcal{E}_{b,S^\flat}/\mathcal{X}_{S^\flat}$ as before, with the further property that for any point $\mathrm{Spa}(K, \mathcal{O}_K) \to S$, the pullback of $\mathcal{F}$ to $\mathcal{X}_{\mathrm{Spa}(K^\flat, \mathcal{O}_K^\flat)}$ is semistable of degree zero. The image $\mathrm{Gr}_{\mathbf{G},\mu^{-1}}^b$ of π_{HT} is much stranger, and is open exactly when b is basic.

Example. Take $\mathbf{G} = \mathrm{GL}_n/E$ and $\mu = (1, 0, \ldots, 0)$, and let b_i be a representative of the σ-conjugacy class with slopes $(-1/(n-i), \ldots, -1/(n-i), 0, \ldots, 0)$ (where 0 appears i times) for $0 \leq i \leq n-1$. Then b_0 is basic and b_{n-1} is μ-ordinary, and the bundle $\mathcal{E}_b$ is generally $\simeq \mathcal{O}(\frac{1}{n-i}) \oplus \mathcal{O}^i$. In this case, $\mathrm{Gr}_{\mathbf{G},\mu^{-1}} \cong \mathbf{P}_E^{n-1}$, with the strata $\left(\mathbf{P}_E^{n-1}\right)^{(i)} = \mathrm{Gr}_{\mathbf{G},\mu^{-1}}^{b_i}$ described as follows:

- $\left(\mathbf{P}_E^{n-1}\right)^{(0)} \cong \Omega^{n-1}$ is the open stratum, and is isomorphic to Drinfeld space over E.
- $\left(\mathbf{P}_E^{n-1}\right)^{(n-1)} = \mathbf{P}^{n-1}(E)$ is the set of E-points of $\mathbf{P}^{n-1}$.
- The intermediate strata can be described as follows: Let

$$\mathbf{P}_{i,j} = \begin{pmatrix} \mathrm{GL}_i & \\ \mathbf{G}_a^{ij} & \mathrm{GL}_j \end{pmatrix}$$

  denote the standard maximal parabolic of type (i, j) in GL_{i+j}. Then $\mathbf{P}_{i,n-i}(E)$ acts naturally on Ω_E^{n-1-i} through its projection to $\mathrm{GL}_{n-i}(E)$, and $\left(\mathbf{P}_E^{n-1}\right)^{(i)} \cong \underline{\Omega_E^{n-1-i} \times_{\mathbf{P}_{i,n-i}(E)} \mathrm{GL}_n(E)}$. On C-points, a point $x \in \mathbf{P}^{n-1}(C)$ lies in the ith stratum iff the associated hyperplane $H_x \subset C^n$ satisfies $\dim_E(H_x \cap E^n) = i$; the projection $\left(\mathbf{P}_E^{n-1}\right)^{(i)} \to \mathbf{P}_{i,n-i}(E)\backslash \mathrm{GL}_n(E)$ simply records the flag $0 \subset H_x \cap E^n \subset E^n$. Note that $\mathbf{P}_{i,n-i}(E)\backslash\mathrm{GL}_n(E)$ is a "naive" p-adic manifold, although we may also give it the structure of a zero-dimensional adic space or diamond over E.

Theorem 1.1. *If $i > 0$, then $H_c^*(\mathrm{Sht}_{\mathbf{G},\mu,b_i}, \overline{\mathbf{Q}_\ell})$ is induced from $\mathbf{P}_{i,n-i}(E)$ as a representation of $\mathrm{GL}_n(E)$.*

Proof. Consider the $\mathrm{GL}_n(E)$-equivariant composite map

$$\tilde{\pi}_{\mathrm{HT}} : \mathrm{Sht}_{\mathbf{G},b_i,\mu} \overset{\pi_{\mathrm{HT}}}{\to} \left(\mathbf{P}_E^{n-1}\right)^{(i)} \overset{\mathrm{pr}}{\to} \mathbf{P}_{i,n-i}(E)\backslash\mathrm{GL}_n(E) = X_i.$$

Note that X_i contains a natural basepoint e, with stabilizer $\mathbf{P}_{i,n-i}(E)$. Then

$$\begin{aligned} H_c^*(\mathrm{Sht}_{\mathbf{G},b_i,\mu}, \overline{\mathbf{Q}_\ell}) &\cong H^0(X_i, R^*\tilde{\pi}_{\mathrm{HT}!}\overline{\mathbf{Q}_\ell}) \\ &\cong \mathrm{Ind}_{\mathbf{P}_{i,n-i}(E)}^{\mathrm{GL}_n(E)}\left(R^*\tilde{\pi}_{\mathrm{HT}!}\overline{\mathbf{Q}_\ell}|_e\right) \end{aligned}$$

where $R^*\tilde{\pi}_{\mathrm{HT}!}\overline{\mathbf{Q}_\ell}|_e$ denotes the stalk at e.

Medidating on this example, one reaches the following conclusion:

For any $(\mathbf{G}, \mu, b)$ such that $\mathrm{Gr}_{\mathbf{G},\mu^{-1}}^b$ fibers G-equivariantly over the E-points of a nontrivial flag variety for $\mathbf{G}$, the cohomology of $\mathrm{Sht}_{\mathbf{G},\mu,b}$ is induced.

So when does this happen?

1.2 The canonical retraction

For simplicity we fix fix $E = \mathbf{Q}_p$, $\mathbf{G} = \mathrm{GL}_n$, $\mathbf{B}$ upper-triangular, and μ the $\mathbf{B}$-dominant cocharacter with weights $(k_1 \geq \cdots \geq k_n)$. Fix $b \in \mathrm{GL}_n(\breve{\mathbf{Q}}_p)$ with $[b] \in B(\mathbf{G}, \mu^{-1})$, and let $\mathcal{E}_{b,S^\flat}$ denote the associated rank n vector bundle on $\mathcal{X}_{S^\flat}$ for any $S \in \mathrm{Perf}_{/\breve{\mathbf{Q}}_p}$; since this bundle is totally functorial with respect to any map

$\mathcal{X}_{S^\flat} \to \mathcal{X}_{T^\flat}$ induced by a map $S \to T$, we sometimes denote it agnostically by $\mathcal{E}_b$. (Whereby we're really regarding it as an $\overline{\mathbf{F}_p}$-point of the stack $\mathrm{Bun}_n(\mathcal{X})$.) For simplicity we assume $k_n \geq 0$,[3] so S-points in $\mathrm{Gr}_{\mathrm{GL}_n,\mu}$ give rise to modifications $u : \mathcal{F}|_{\mathcal{X}_{S^\flat} \smallsetminus S} \xrightarrow{\sim} \mathcal{E}_{b,S^\flat}|_{\mathcal{X}_{S^\flat} \smallsetminus S}$ which are *effective,* i.e. modifications for which u extends to an injection $u : \mathcal{F} \hookrightarrow \mathcal{E}_{b,S^\flat}$ of finite locally free $\mathcal{O}_{\mathcal{X}}$-modules. By definition, such a modification lies in the admissible locus $\mathrm{Gr}^{\mathcal{E}_b-\mathrm{adm}}_{\mathrm{GL}_n,\mu}$ if and only if $\mathcal{F}$ is pointwise semistable of slope zero at all points of $|S|$; we'll sometimes just say $(\mathcal{F}, u)$ is admissible.

Let $0 = \mathcal{E}_b^0 \subsetneq \mathcal{E}_b^1 \subsetneq \mathcal{E}_b^2 \cdots \subsetneq \mathcal{E}_b^s = \mathcal{E}_b$ denote the slope filtration of $\mathcal{E}_b$, where s denotes the number of distinct slopes of $\mathcal{E}_b$. Each $\mathcal{E}_b^i/\mathcal{E}_b^{i-1}$ is semistable. The condition $[b] \in B(\mathbf{G}, \mu^{-1})$ unwinds in this case to the usual relation between Newton and Hodge polygons; examination of polygons then produces the inequality

$$\deg(\mathcal{E}_b^i) \leq \sum_{1 \leq j \leq \mathrm{rank}(\mathcal{E}_b^i)} k_j$$

for any $1 \leq i \leq s$, with equality for $i = s$. Let $\mathcal{I} \subseteq \{1, \ldots, s\}$ denote the ordered set of integers for which this inequality is an *equality*; since $s \in \mathcal{I}$ always, $|\mathcal{I}| \geq 1$.

Definition 1.2. The datum $(\mathbf{G}, \mu, b)$ is *Hodge-Newton(HN)-reducible* if $|\mathcal{I}| \geq 2$. We say $\mathcal{E}_b^i$ is *HN-reducing* if $i \in \mathcal{I} \smallsetminus \{s\}$.

Remark. When μ is minuscule, it's easy to check that $|\mathcal{I}| \leq 3$. For non-minuscule μ, however, $|\mathcal{I}|$ can be arbitrarily large.

Suppose $(\mathbf{G}, \mu, b)$ is HN-reducible. Let $\{d_1, \ldots, d_k\}$ denote the ordered set

$$\{\mathrm{rank}(\mathcal{E}_b^{i_1}), \mathrm{rank}(\mathcal{E}_b^{i_2}/\mathcal{E}_b^{i_1}), \ldots, \mathrm{rank}(\mathcal{E}_b/\mathcal{E}_b^{i_{k-1}})\},$$

with $\mathcal{I} = \{i_1 < \cdots < i_k = s\}$ as above. Consider the standard Levi

$$\mathbf{M} = \begin{pmatrix} \mathrm{GL}_{d_1} & & & \\ & \mathrm{GL}_{d_2} & & \\ & & \ddots & \\ & & & \mathrm{GL}_{d_k} \end{pmatrix} \subset \mathbf{G}.$$

We write $\mathbf{P} = \mathbf{M} \cdot \mathbf{U}$ for the standard parabolic containing $\mathbf{M}$. After possibly replacing b by a σ-conjugate, we assume that $b \in \mathbf{M}(\breve{\mathbf{Q}}_p)$ and that $\nu_{b^{-1}}$ factors through $\mathbf{M}$ and is $\mathbf{M}$-conjugate to $\nu_{[b^{-1}]}$.[4] Writing b_m for the projection of b into the mth block of $\mathbf{M}$, we then get a decomposition $\mathcal{E}_b \cong \oplus_{1 \leq i \leq k} \mathcal{E}_{b_i}$, or equivalently a canonical reduction of $\mathcal{E}_b$ to an $\mathbf{M}$-bundle, such that the induced $\mathbf{P}$-bundle structure is a coarsening of the slope filtration of $\mathcal{E}_b$.

One easily checks that $(\mathbf{M}, \mu, b)$ defines a local shtuka datum, which is naturally a product of local shtuka data $\prod_{m=1}^k (\mathrm{GL}_{d_m}, \mu_m, b_m)$ (where again μ_m denotes the projection of μ resp. b into the mth block of $\mathbf{M}$) and we get canonical compatible isomorphisms

$$\mathrm{Gr}_{\mathbf{M},\mu} \cong \mathrm{Gr}_{\mathrm{GL}_{d_1},\mu_1} \times_{\mathrm{Spd}\,\breve{\mathbf{Q}}_p} \cdots \times_{\mathrm{Spd}\,\breve{\mathbf{Q}}_p} \mathrm{Gr}_{\mathrm{GL}_{d_k},\mu_k}$$

and

$$\mathrm{Gr}^{\mathcal{E}_b-\mathrm{adm}}_{\mathbf{M},\mu} \cong \mathrm{Gr}^{\mathcal{E}_{b_1}-\mathrm{adm}}_{\mathrm{GL}_{d_1},\mu_1} \times_{\mathrm{Spd}\,\breve{\mathbf{Q}}_p} \cdots \times_{\mathrm{Spd}\,\breve{\mathbf{Q}}_p} \mathrm{Gr}^{\mathcal{E}_{b_k}-\mathrm{adm}}_{\mathrm{GL}_{d_k},\mu_k}.$$

There is a canonical inclusion of period domains

$$i : \mathrm{Gr}_{\mathbf{M},\mu} \hookrightarrow \mathrm{Gr}_{\mathbf{G},\mu}$$

[3] This can always be achieved by a suitable twisting of the datum b, μ.

[4] This holds, for example, if b^{-1} is well-chosen; the unfortunate inverse here is due to the negation of slopes when passing from the isocrystal $(\breve{\mathbf{Q}}_p^n, b\sigma)$ to the bundle $\mathcal{E}_b$.

induced by the decomposition $\mathcal{E}_b \cong \oplus_{1\leq i\leq k}\mathcal{E}_{b_i}$ and therefore sending $\mathrm{Gr}_{\mathbf{M},\mu}^{\mathcal{E}_b-\mathrm{adm}}$ into $\mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}}$. These inclusions are equivariant for a natural action of the group $J_b = \mathbf{J}_b(\mathbf{Q}_p) \subset \mathbf{M}(\breve{\mathbf{Q}}_p)$. In fact, i is equivariant for the much larger group $\mathcal{J}_b^{\mathbf{M}} = \mathrm{Aut}_{\mathrm{Bun}(\mathbf{M})}(\mathcal{E}_b) = \prod_{1\leq m\leq k}\mathrm{Aut}(\mathcal{E}_{b_m})$. Note that the full automorphism group $\mathcal{J}_b = \mathrm{Aut}(\mathcal{E}_b)$ acts on $\mathrm{Gr}_{\mathbf{G},\mu}$, and that $\mathcal{J}_b$ canonically decomposes as the semidirect product $\mathcal{J}_b^{\mathbf{M}} \ltimes \mathcal{J}_b^{\mathbf{U}}$, where $\mathcal{J}_b^{\mathbf{U}}$ is the subgroup of elements

$$f \in \mathrm{Aut}(\mathcal{E}_b) \subset \mathrm{Hom}(\mathcal{E}_b, \mathcal{E}_b)$$

such that $f-1$ carries each $\mathcal{E}_b^{i_m}$ into $\mathcal{E}_b^{i_{m-1}}$. (To be clear, $\mathcal{J}_b$ and its decorated variants are not groups, but group sheaves, i.e. functors $\mathrm{Perf}_{/\breve{\mathbf{Q}}_p} \to \mathrm{Groups}$. E.g. $\mathcal{J}_b(S) = \mathrm{Aut}(\mathcal{E}_{b,S^\flat})$. These are all *group diamonds.*)

The first main result is the following theorem.

Theorem 1.3. *The inclusion $i : \mathrm{Gr}_{\mathbf{M},\mu}^{\mathcal{E}_b-\mathrm{adm}} \hookrightarrow \mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}}$ admits a canonical $\mathcal{J}_b^{\mathbf{M}}$-equivariant retraction*

$$r : \mathrm{Gr}_{\mathbf{G},\mu}^{b-\mathrm{adm}} \to \mathrm{Gr}_{\mathbf{M},\mu}^{b-\mathrm{adm}}.$$

In other words, any admissible type-μ modification of $\mathcal{E}_{b,S^\flat}$ along S admits a canonical reduction to a collection of admissible type μ_i-modifications of the bundles $\mathcal{E}_{b_i,S^\flat}$ along S.

1.3 The canonical retraction on points

We first explain the idea in the case when S is a point, where things are technically simpler. Fix any perfectoid field $K/\mathbf{Q}_p$ with corresponding adic space $S = \mathrm{Spa}(K, \mathcal{O}_K)$, so we have the adic Fargues-Fontaine curve $\mathcal{X} = \mathcal{X}_{K^\flat,\mathbf{Q}_p}$. This is a locally Noetherian quasicompact adic curve, and we have a natural closed immersion $i : S \to \mathcal{X}$. Let

$$0 \to \mathcal{F} \xrightarrow{u} \mathcal{E} \to \mathcal{Q} \to 0$$

be a short exact sequence of coherent sheaves on $\mathcal{X}$, where $\mathcal{F}$ and $\mathcal{E}$ are both rank n vector bundles and $\mathcal{Q}$ is supported at the distinguished point $x(\infty) = i(S) \in |\mathcal{X}|$. Then the stalk $Q = \mathcal{Q}_{x(\infty)}$ is a finite torsion module over the DVR $\widehat{\mathcal{O}}_{\mathcal{X},x(\infty)} \cong \mathbb{B}_{\mathrm{dR}}^+(K)$, and we say $(\mathcal{F}, u)$ is a type-μ modification of $\mathcal{E}$ along $x(\infty)$ if there is an isomorphism $Q \simeq \oplus_{1\leq i\leq n}\mathbb{B}_{\mathrm{dR}}^+(K)/\xi^{k_i}$ (here ξ denotes any uniformier of $\mathbb{B}_{\mathrm{dR}}^+(K)$). (If $\mathcal{E} = \mathcal{E}_b$, the $(K, \mathcal{O}_K)$-points of $\mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}}$ are exactly the isomorphism classes of such pairs $(\mathcal{F}, u)$ with the further property that $\mathcal{F}$ is semistable of slope zero.)

Theorem 1.4. *With notation as in the preceding paragraph, let $\mathcal{E}^+ \subseteq \mathcal{E}$ be any saturated subbundle, and set $\mathcal{F}^+ = \mathcal{F} \cap \mathcal{E}^+$. Then if $\mathcal{F}$ is semistable of slope zero, we have the inequality*

$$\deg(\mathcal{E}^+) \leq \sum_{1\leq i\leq \mathrm{rank}(\mathcal{E}^+)} k_i,$$

and if equality holds then $\mathcal{F}^+$ is also semistable of slope zero.

Proof. Let Q^+ denote the image of the stalk $\mathcal{E}^+_{x(\infty)}$ in Q. It's easy to see the equality

$$\deg(\mathcal{E}^+) = \deg(\mathcal{F}^+) + \ell(Q^+),$$

where ℓ denotes length as a $\mathbb{B}_{\mathrm{dR}}^+(K)$-module. Since $\mathcal{F}$ is semistable of slope zero and $\mathcal{F}^+ \subseteq \mathcal{F}$ is saturated, $\mathcal{F}^+$ must have degree ≤ 0, so dropping $\deg(\mathcal{F}^+)$ from this equality gives $\deg(\mathcal{E}^+) \leq \ell(Q^+)$. If r denotes the rank of $\mathcal{E}^+$, clearly $\mathcal{E}^+_{x(\infty)}$ and then also Q^+ are generated by r elements, so Lemma 1.5 implies the inequality $\ell(Q^+) \leq \sum_{1\leq i\leq r} k_i$. Combining these inequalities, the first part of the theorem follows. Putting together the first equality with this second inequality, we also get

$$\deg(\mathcal{E}^+) - \sum_{1\leq i\leq r} k_i \leq \deg(\mathcal{F}^+),$$

so if the left-hand side is zero then $\mathcal{F}^+$ has degree zero. But then $\mathcal{F}^+$ must be semistable of degree zero, since otherwise it would have a positive-degree subbundle as a step in its slope filtration, contradicting the semistability of $\mathcal{F}$. ∎

Lemma 1.5. *Let R be a DVR with uniformizer π, and let M be a finite torsion R-module, so $M \simeq \oplus_{1\leq i\leq n} R/\pi^{k_i}$ for some uniquely determined sequence $\mu(M) = (k_1 \geq \cdots \geq k_n)$ with $k_n > 0$. Let $N \subseteq M$ be an R-submodule generated by j elements. Then $\ell(N) \leq k_1 + \cdots + k_j$, and if equality holds then N is a direct summand.*

We will somewhat abusively refer to elements of the ordered sequence $\mu(M)$ as "elementary divisors" of M.

Proof. For the first claim, it clearly suffices to show that $\ell(M/N) \geq \sum_{j<i\leq n} k_i$. For this we use Fitting ideals. Recall that for any finite torsion module Q over R with elementary divisors k_i, we have an equality $\mathrm{Fitt}_j(Q) = (\pi^{\sum_{j<i} k_i})$; in particular, $\mathrm{Fitt}(Q) = \mathrm{Fitt}_0(Q) = (\pi^{\ell(Q)})$, and $\mathrm{Fitt}_m(Q) = R$ if Q is generated by $\leq m$ elements. Returning to the situation at hand, we have an inclusion

$$\mathrm{Fitt}_j(N)\mathrm{Fitt}(M/N) \subseteq \mathrm{Fitt}_j(M) = (\pi^{\sum_{j<i\leq n} k_i})$$

(this is a special case of Proposition XIII.10.7 in Lang's Algebra). But $\mathrm{Fitt}_j(N) = R$ since N is generated by j elements, so we get

$$(\pi^{\ell(M/N)}) = \mathrm{Fitt}(M/N) \subseteq \mathrm{Fitt}_j(M) = (\pi^{\sum_{j<i\leq n} k_i}),$$

and this immediately implies the desired inequality.

For the second claim, we argue by induction on j; the case $j = 1$ is easy. For the induction step, choose a projection $f : M \to R/\pi^{k_1}$ onto a maximal-length cyclic direct summand, so $\ker f \simeq \oplus_{2\leq i\leq n} R/\pi^{k_i}$. Let $n_1, \ldots, n_j$ be a set of elements generating N. After rearranging the n_i's, we may assume that $f(N) = f(C)$ where $C = Rn_1 \subseteq N$, i.e. that $f(N)$ is generated by $f(n_1)$. After then possibly replacing n_i by $n_i - r_i n_1$ for all $2 \leq i \leq j$, we may assume that $\ker f$ contains the submodule N' generated by $n_2, \ldots, n_j$. Note that we have inequalities $\ell(N') \leq k_2 + \cdots + k_j$ and $\ell(C) \leq k_1$, the latter because π^{k_1} kills M and the former by applying the first half of the lemma to $N' \subseteq \ker f$. By assumption, we have $\ell(N) = k_1 + \cdots + k_j$ so now the chain of inequalities

$$\ell(N) = \ell(N' + C) \leq \ell(N') + \ell(C) \leq k_1 + \cdots + k_j = \ell(N)$$

forces equalities $\ell(N') = k_2 + \cdots + k_j$ and $\ell(C) = k_1$. Since N' and C are generated by $j-1$ elements and 1 element, respectively, they are both direct summands of M by the induction hypothesis. Finally, the above chain also forces the equality $\ell(N' + C) = \ell(N') + \ell(C)$, which implies that $N' \cap C = 0$ inside M, so $N \cong N' \oplus C \subseteq M$ is a direct summand of M. $\square$

Returning to the setting of Theorem 1.3, take $\mathcal{E} = \mathcal{E}_{b,S^\flat}$, and choose $(\mathcal{F}, u)$ corresponding to a $(K, \mathcal{O}_K)$-point of $\mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}}$. Then for any i such that $\mathcal{E}^i = \mathcal{E}^i_{b,S^\flat}$ is HN-reducing, the equality $\deg(\mathcal{E}^i) = \sum_{1\leq j\leq \mathrm{rank}(\mathcal{E}^i)} k_j$ holds by assumption, so by Theorem 1.4 we get that the bundle $\mathcal{F}^i = \mathcal{F} \cap \mathcal{E}^i$ is semistable of slope zero, and then another application of Lemma 1.5 shows that the module

$$Q^i = \mathrm{im}(\mathcal{E}^i \to Q) = \mathcal{E}^i/\mathcal{F}^i \subseteq Q$$

is a direct summand of Q such that $Q^i \simeq \oplus_{1\leq j\leq \mathrm{rank}(\mathcal{E}^i)} \mathbb{B}^+_{\mathrm{dR}}(K)/\xi^{k_j}$. Doing this for all the HN-reducing i's, we get a canonical flag $0 \subsetneq \mathcal{F}^{i_1} \subsetneq \cdots \subsetneq \mathcal{F}^{i_k} = \mathcal{F}$ with each step semistable of slope zero. Thus the quotients are all ss of slope zero too, and we can form the short exact sequences

$$0 \to \mathcal{F}^{i_m}/\mathcal{F}^{i_{m-1}} \xrightarrow{u_m} \mathcal{E}^{i_m}/\mathcal{E}^{i_{m-1}} \to Q^{i_m}/Q^{i_{m-1}} \to 0.$$

But then easy induction on m shows that

$$Q^{i_m}/Q^{i_{m-1}} \simeq \oplus_{\mathrm{rank}(\mathcal{E}^{i_{m-1}})<j\leq \mathrm{rank}(\mathcal{E}^{i_m})} \mathbb{B}^+_{\mathrm{dR}}(K)/\xi^{k_j},$$

so we conclude that each pair $(\mathcal{F}^{i_m}/\mathcal{F}^{i_{m-1}}, u_m)$ is canonically an admissible type-μ_m modification of $\mathcal{E}^{i_m}/\mathcal{E}^{i_{m-1}} \cong \mathcal{E}_{b_m}$. Thus we get a canonical $(K, \mathcal{O}_K)$-point of $\mathrm{Gr}_{\mathbf{M},\mu}^{\mathcal{E}_b-\mathrm{adm}}$, as desired.

1.4 Spreading out the canonical retraction

1.4.1 Finite projective virtuality

First we do some module theory over $\mathbb{B}_{\mathrm{dR}}^+$. The relevance of the material here will become clear in the next subsection.

Fix a perfectoid Tate ring $A/\mathbf{Q}_p$, and let $\xi \in \mathbb{B}_{\mathrm{dR}}^+(A)$ denote a choice of generator for the kernel of $\theta : \mathbb{B}_{\mathrm{dR}}^+(A) \twoheadrightarrow A$. This is a non-zero-divisor, unique up to a unit.

Definition 1.6. A $\mathbb{B}_{\mathrm{dR}}^+(A)$-module M is *finite projective virtually (fpv) over A* if M can be resolved by a short exact sequence

$$0 \to P_1 \to P_0 \to M \to 0$$

where P_0 and P_1 are finite projective $\mathbb{B}_{\mathrm{dR}}^+(A)$-modules, and some finite power of ξ kills M.

In other words, M is fpv over A if M is ξ-torsion and *1-fpd* as a $\mathbb{B}_{\mathrm{dR}}^+(A)$-module in the sense of [KL2]. Note the word "virtually": M is typically not an A-module, since $\mathbb{B}_{\mathrm{dR}}^+(A)$ is not an A-algebra. Note that any finite direct sum $\oplus_i \mathbb{B}_{\mathrm{dR}}^+(A)/\xi^{n_i}$ is fpv over A, and so is any finite projective A-module (regarded as a $\mathbb{B}_{\mathrm{dR}}^+(A)$-module via θ). Anyway, we regard these modules as a full subcategory of $\mathbb{B}_{\mathrm{dR}}^+(A)$-modules. By [KL2], Lemma 1.1.5, the property of being fpv over A is stable under formation of extensions. We record some further properties as a lemma.

Proposition 1.7. i. *Let*

$$0 \to M_1 \to M_2 \to M_3 \to 0$$

be an exact sequence of $\mathbb{B}_{\mathrm{dR}}^+(A)$-modules. If M_2 is finite projective and M_3 is fpv over A, then M_1 is finite projective. In fact, if M_2 is finite projective and M_3 is ξ-torsion, then M_1 is finite projective if and only if M_3 is fpv over A.

ii. *If M is fpv over A, then so are the submodules $\xi^n M$ and $M[\xi^n]$ for any n.*

iii. *If $0 \to M \to N \to L \to 0$ is an exact sequence of $\mathbb{B}_{\mathrm{dR}}^+(A)$-modules such that N and L are both fpv, then M is fpv.*

Proof. We just use the different portions of Lemma 1.1.5 of [KL2] a bunch of times. Part i. and part iii. are easy. For part ii., we note that $M/\xi^n M$ is 2-fpd, so then considering the sequence

$$0 \to \xi^n M \to M \to M/\xi^n M \to 0$$

Lemma 1.1.5(f) of [KL2] shows that $\xi^n M$ is 1-fpd, and hence fpv. But then looking at the sequence

$$0 \to M[\xi^n] \to M \to \xi^n M \to 0,$$

part iii. implies that $M[\xi^n]$ is 1-fpd. □

Proposition 1.8. *If M is a $\mathbb{B}_{\mathrm{dR}}^+(A)$-module which is fpv over A, and $N \subseteq M$ is a direct summand of M, then N is fpv over A.*

Proof. Let $e(M)$ be the smallest integer e such that ξ^e kills M. We prove the claim by induction on $e(M)$. When $e(M) = 1$, the result is clear: in this case, M is a finite projective A-module, and N is a direct summand thereof, so also finite projective over A. In general, we have a commutative diagram with exact rows

$$\begin{tikzcd}
0 \arrow[r] & N[\xi] \arrow[r] \arrow[d] & N \arrow[r] \arrow[d] & \xi N \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & M[\xi] \arrow[r] & M \arrow[r] & \xi M \arrow[r] & 0
\end{tikzcd}$$

where the vertical arrows identify the upper row as a direct summand of the lower row, in the evident sense. But $e(M[\xi]) = 1$ and $e(\xi M) = e(M) - 1$, so $N[\xi]$ and ξN are fpv over A by the induction hypothesis. Since the property of being fpv over A is stable under forming extensions, we get the result. □

Our next goal is a pointwise criterion for a $\mathbb{B}^+_{\mathrm{dR}}(A)$-module to be fpv over A. In order to do this, we introduce the following auxiliary notion.

Definition 1.9. A $\mathbb{B}^+_{\mathrm{dR}}(A)$-module M is *finite projective narrowly (fpn) over A* if M is finitely generated and the graded module

$$\mathrm{gr}_\xi(M) = \oplus_{i \geq 0}(\xi^i M/\xi^{i+1}M)$$

is a finite projective A-module.

It's easy to check that these conditions imply that M is ξ-adically separated and killed by a finite power of ξ. One also checks that if M is fpn over A, then M is fpv over A. The proof is again by induction on $e(M)$: the case of $e = 1$ is easy (since then M is just a finite projective A-module), and for the induction step one uses the sequence

$$0 \to \xi M \to M \to M/\xi M \to 0,$$

noting that $\mathrm{gr}_\xi(\xi M)$ is an A-module summand of $\mathrm{gr}_\xi(M)$.

Proposition 1.10. *Let N be a $\mathbb{B}^+_{\mathrm{dR}}(A)$-module which is finitely generated and ξ-torsion, and set $X = \mathrm{Spa}(A, A^\circ)$. If the elementary divisors of $N_x = N \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \mathbb{B}^+_{\mathrm{dR}}(K_x)$ are locally constant as a function of $x \in |X|$, then N is fpn over A, and hence fpv over A.*

Proof. Let $k_{1,x} \geq k_{2,x} \geq \dots$ be the elementary divisors of N_x as a $\mathbb{B}^+_{\mathrm{dR}}(K_x)$-module. These can be read off from $\mathrm{gr}(N_x)$ by the following recipe: $\dim_{K_x} \mathrm{gr}^i N_x =$the # of $k_{j,x}$'s with $k_{j,x} > i$. In particular, the hypotheses of the theorem guarantee that

$$\dim_{K_x} \mathrm{gr}_\xi(N_x)$$

is locally constant as a function of $x \in |X|$. On the other hand, there is a natural map

$$\mathrm{gr}_\xi(N) \otimes_A K_x \cong \mathrm{gr}_\xi(N) \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \mathbb{B}^+_{\mathrm{dR}}(K_x) \to \mathrm{gr}_\xi(N_x)$$

which one checks[5] is an isomorphism. Putting these two things together, we get that

$$\dim_{K_x} \mathrm{gr}_\xi(N) \otimes_A K_x$$

is locally constant as a function of x. Then since A is uniform, Proposition 2.8.4 of [KL1] implies that $\mathrm{gr}_\xi(N)$ is a finite projective A-module. This verifies that N is fpn over A, and thus fpv. □

The following thing is no longer used in the present draft, but I'll keep it for later reference.

Proposition 1.11. *Let M be a $\mathbb{B}^+_{\mathrm{dR}}(A)$ module which is fpv over A, and set $X = \mathrm{Spa}(A, A^\circ)$. Then the natural map*

$$M \to \prod_{x \in |X|} \left(M \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \mathbb{B}^+_{\mathrm{dR}}(K_x) \right)$$

is injective. In particular, if $m \in M$ is nonzero, its image in $M \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \mathbb{B}^+_{\mathrm{dR}}(K_x)$ is nonzero for some $x \in |X|$.

The proof which follows is the first one I worked out; shorter arguments are also possible.

Proof. Fix $\xi \in \mathbb{B}^+_{\mathrm{dR}}(A)$ generating $\ker \theta$; in what follows, we'll implicitly use the fact that for any continuous map $A \to B$ of perfectoid Tate rings over $\mathbf{Q}_p$, the image of ξ also generates $\ker \theta \subset \mathbb{B}^+_{\mathrm{dR}}(B)$.

To lighten notation, we define some natural functors on $\mathbb{B}^+_{\mathrm{dR}}(A)$-modules as follows:

- $M_x = M \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \mathbb{B}^+_{\mathrm{dR}}(K_x)$ for any $x \in |X|$.
- $T(M) = M \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \prod_{x \in |X|} \mathbb{B}^+_{\mathrm{dR}}(K_x)$.
- $\mathcal{T}(M) = \prod_{x \in |X|} M_x$.

Note the obvious natural transformations $\mathrm{id} \to T(-) \to \mathcal{T}(-) \rightleftarrows (-)_x$. It's not clear (and rather unlikely)

[5] Probably I should add a proof of this; one can use e.g. induction on $e(N)$.

that T and $\mathcal{T}$ have any good properties in general. However, we're going to prove that on the subcategory of $\mathbb{B}_{\mathrm{dR}}^{+}(A)$-modules M which are finite projective virtually over A, the functors T and $\mathcal{T}$ are naturally isomorphic, additive, exact and faithful. In particular, the natural map $M \to \mathcal{T}(M)$ is injective when M is fpv over A, which is exactly the claim we're trying to prove.

The key ingredient in showing all this is the injectivity of the map

$$\mathbb{B}_{\mathrm{dR}}^{+}(A) \to \prod_{x \in |X|} \mathbb{B}_{\mathrm{dR}}^{+}(K_x),$$

which we now check. Since both sides are ξ-adically separated, it suffices to check this modulo ξ^n for all n.

Set $\tilde{A}^{\circ} = \prod_{x \in |X|} K_x^{\circ}$ with the p-adic topology, and set $\tilde{A} = \tilde{A}^{\circ}[\frac{1}{p}] \subset \dot{A} = \prod_{x \in |X|} K_x$. Then $\tilde{A}$ is perfectoid, and since A is uniform, the map $A \to \tilde{A}$ is an injective homeomorphism onto its image by an old theorem of Berkovich. In particular, $A \to \dot{A}$ is injective. Now playing around, we deduce cascadingly that

- $A^{\circ\flat} \to \tilde{A}^{\circ\flat}$ is injective.
- Both arrows in the sequence of maps

$$W(A^{\circ\flat}) \to W(\tilde{A}^{\circ\flat}) \to \prod_{x \in X} W(K_x^{\circ\flat})$$

are injective (by looking at Teichmuller coordinates). Actually the second map is an isomorphism.

- All three arrows in the sequence

$$W(A^{\circ\flat})[\tfrac{1}{p}] \to W(\tilde{A}^{\circ\flat})[\tfrac{1}{p}] \to \left(\prod_{x \in X} W(K_x^{\circ\flat})\right)[\tfrac{1}{p}] \to \prod_{x \in |X|} W(K_x^{\circ\flat})[\tfrac{1}{p}]$$

are injective (easy).

- The map

$$W(A^{\circ\flat})[\tfrac{1}{p}]/\xi^n \to \prod_{x \in |X|} \left(W(K_x^{\circ\flat})[\tfrac{1}{p}]/\xi^n\right)$$

is injective for any n. We show this by induction on n. For $n = 1$ this becomes our original map $A \to \dot{A}$. For the induction step, look at the diagram

$$\begin{tikzcd}
0 \arrow[r] & \xi^n W(A^{\circ\flat})[\frac{1}{p}]/\xi^{n+1} \cong \xi^n A \arrow[r] \arrow[d] & W(A^{\circ\flat})[\frac{1}{p}]/\xi^{n+1} \arrow[r] \arrow[d] & W(A^{\circ\flat})[\frac{1}{p}]/\xi^n \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \prod_{x \in |X|} \left(\xi^n W(K_x^{\circ\flat})[\frac{1}{p}]/\xi^{n+1}\right) \cong \xi^n \dot{A} \arrow[r] & \prod_{x \in |X|} \left(W(K_x^{\circ\flat})[\frac{1}{p}]/\xi^{n+1}\right) \arrow[r] & \prod_{x \in |X|} \left(W(K_x^{\circ\flat})[\frac{1}{p}]/\xi^n\right) \arrow[r] & 0
\end{tikzcd}$$

(note that the rows are exact). Then the lefthand vertical arrow is injective, and the righthand vertical arrow is injective by our induction hypothesis, so now the snake lemma implies the claim.

- The map $\mathbb{B}_{\mathrm{dR}}^{+}(A) \to \prod_{x \in |X|} \mathbb{B}_{\mathrm{dR}}^{+}(K_x)$ is injective. Take the limit in the previous step.

Now suppose M is fpv, and fix a presentation

$$0 \to P_0 \to P_1 \to M \to 0$$

with P_i finite projective. Writing P_i as a summand of a finite free module, one easily checks (using the injectivity of $\mathbb{B}_{\mathrm{dR}}^{+}(A) \to \mathcal{T}(\mathbb{B}_{\mathrm{dR}}^{+}(A))$) that the natural map $P_i \to T(P_i)$ is injective and that the natural map $T(P_i) \to \mathcal{T}(P_i)$ is an isomorphism. Applying both functors to the presentation of M, and taking the previous sentence into account, we get a commutative diagram with exact rows

$$\begin{tikzcd}
T(P_0) \arrow[r] \arrow[d, "\wr"'] & T(P_1) \arrow[r] \arrow[d, "\wr"'] & T(M) \arrow[r] \arrow[d] & 0 \\
\mathcal{T}(P_0) \arrow[r] & \mathcal{T}(P_1) \arrow[r] & \mathcal{T}(M) \arrow[r] & 0
\end{tikzcd}$$

and then an easy diagram chase shows that $T(M) \to \mathcal{T}(M)$ is bijective. Next, since each $\mathbb{B}^+_{\mathrm{dR}}(K_x)$ is a DVR, the modules $P_{i,x}$ are finite free and the map $P_{0,x} \to P_{1,x}$ is injective: its kernel is torsion-free, hence free, of rank equal to the generic rank of M_x, which is zero. In particular, by the identification

$$0 = \ker(P_{0,x} \to P_{1,x}) = \mathrm{Tor}_1^{\mathbb{B}^+_{\mathrm{dR}}(A)}(M, \mathbb{B}^+_{\mathrm{dR}}(K_x)),$$

we get that $M \mapsto M_x$ is an exact functor on fpv modules. But then since $\mathcal{T}(M) = \prod M_x$, we get that $\mathcal{T}$ is an exact functor on fpv modules as well.

Finally, we prove the injectivity of the map $M \to \mathcal{T}(M)$. Again we argue by induction on $e(M)$. When $e(M) = 1$, the module M is a finite projective A-module, and the map $M \to \mathcal{T}(M)$ identifies with the map $M \to \prod_{x \in |X|}(M \otimes_A K_x)$; by writing M as a summand of a finite free A-module, injectivity of this map reduces to the aforementioned injectivity of $A \to \prod_{x\in|X|} K_x$. For the induction step, we use the commutative diagram

$$\begin{tikzcd}
0 \arrow[r] & M[\xi] \arrow[r] \arrow[d] & M \arrow[r] \arrow[d] & \xi M \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \mathcal{T}(M[\xi]) \arrow[r] & \mathcal{T}(M) \arrow[r] & \mathcal{T}(\xi M) \arrow[r] & 0
\end{tikzcd}$$

noting that both rows are exact (the lower row by the results of the previous paragraph). Since $e(M[\xi]) = 1$, the left vertical arrow is injective; since $e(\xi M) = e(M) - 1$, the right vertical arrow is injective by the induction hypothesis. But then the middle vertical arrow is injective by the snake lemma. □

1.4.2 Effective modifications in families

Fix an affinoid perfectoid space $S = \mathrm{Spa}(A, A^+)$ over $\mathbf{Q}_p$ with tilt $S^\flat = \mathrm{Spa}(R, R^+)$; we choose this notation for compatibility with [KL1]. Let $\mathcal{X} = \mathcal{X}_{S^\flat}$ denote the adic Fargues-Fontaine curve over $S^\flat$; there is a canonical Zariski-closed embedding $i : S \to \mathcal{X}_{S^\flat}$. Let $\mathcal{O}(1)$ be the canonical ample line bundle on $\mathcal{X}$, $P_R = \oplus_{i\geq 0} H^0(\mathcal{X}, \mathcal{O}(i))$. Then $X = X_{S^\flat} = \mathrm{Proj}(P_R)$ is the schematic FF curve associated with $S^\flat$. Set $Z = \mathrm{Spec}(A)$, so we have a canonical closed immersion $Z \to X$ such that the completion of X along Z is canonically identified with $\tilde{Z} := \mathrm{Spec}\, \mathbb{B}^+_{\mathrm{dR}}(A)$. Furthermore, the subscheme $X \smallsetminus Z$ is affine; we define $\mathbb{B}_e(A) = H^0(X \smallsetminus Z, \mathcal{O}_X)$ to be its coordinate ring.

Summing up, we have a canonical diagram of locally ringed spaces

$$\begin{tikzcd}
X \smallsetminus Z \coprod \tilde{Z} \arrow[r, "(f_e, f^+_{\mathrm{dR}})"] & X & \mathcal{X} \arrow[l, "f^{\mathrm{an}}"'] \\
 & Z \arrow[u] & S \arrow[l] \arrow[u]
\end{tikzcd}$$

over $\mathrm{Spec}\, \mathbf{Q}_p$, contravariantly functorial in morphisms $(A, A^+) \to (B, B^+)$ of perfectoid $\mathbf{Q}_p$-algebras.

Theorem 1.12 (Kedlaya-Liu)**.** *With the setup as above,*

i. *Pullback along the morphism f^{an} induces an equivalence of exact tensor categories from vector bundles on X to vector bundles on $\mathcal{X}$.*

ii. *Pulling back along the pair of morphisms (f_e, f^+_{dR}) and then passing to global sections induces an equivalence of exact tensor categories from vector bundles on X to B-pairs over A.*

In this context, a *B-pair over A* is a pair $M = (M_e, M^+_{\mathrm{dR}})$ where M_e is a finite projective $\mathbb{B}_e(A)$-module and M^+_{dR} is a finite projective $\mathbb{B}^+_{\mathrm{dR}}(A)$-lattice in the $\mathbb{B}_{\mathrm{dR}}(A)$-module $M_{\mathrm{dR}} = M_e \otimes_{\mathbb{B}_e(A)} \mathbb{B}_{\mathrm{dR}}(A)$. If $\mathcal{V}$ is a vector bundle on $\mathcal{X}$ (or X), we write $M(\mathcal{V}) = (M_e(\mathcal{V}), M^+_{\mathrm{dR}}(\mathcal{V}))$ for the associated B-pair; we denote the functor in the other direction by $M \mapsto \mathcal{V}(M)$.

We remark that by the functoriality of the above diagram, any point $x \in \mathrm{Spa}(A, A^+)$ gives rise to a morphism

$$s_x : \mathcal{X}_{\mathrm{Spa}(K_x, K_x^+)^\flat} \to \mathcal{X}_{S^\flat}.$$

If $\mathcal{E}$ is a vector bundle on $\mathcal{X}_{S^\flat}$, we abbreviate the pullback $s_x^*\mathcal{E}$ by $\mathcal{E}_x$. Note that $\mathcal{E}_x$ corresponds to the B-pair over K_x given by

$$(M_e(\mathcal{E}) \otimes_{\mathbb{B}_e(A)} \mathbb{B}_e(K_x), M_{\mathrm{dR}}^+(\mathcal{E}) \otimes_{\mathbb{B}_{\mathrm{dR}}^+(A)} \mathbb{B}_{\mathrm{dR}}^+(K_x)).$$

Definition 1.13. Notation and setup as above. An *effective modification along* S is a triple $(\mathcal{E}, \mathcal{F}, u)$ where $\mathcal{E}$ and $\mathcal{F}$ are vector bundles on $\mathcal{X}$, and $u : \mathcal{F} \hookrightarrow \mathcal{E}$ is an injective map of $\mathcal{O}_{\mathcal{X}}$-modules such that $\mathcal{E}/u(\mathcal{F})$ is killed by a finite power of the ideal sheaf cutting out S in $\mathcal{X}_{S^\flat}$.

When $\mathcal{E}$ is given, we also speak of $(\mathcal{F}, u)$ being an effective modification of $\mathcal{E}$ along S.

Theorem 1.14. *Let $\mathcal{E}$ be a vector bundle on $\mathcal{X}$. Then we have a natural identification between the set of isomorphism classes of effective modifications of $\mathcal{E}$ along S and the set of $\mathbb{B}_{\mathrm{dR}}^+(A)$-submodules $N \subseteq M_{\mathrm{dR}}^+(\mathcal{E})$ such that $M_{\mathrm{dR}}^+(\mathcal{E})/N$ is finite projective virtually over A.*

Proof. The functor in one direction sends $(\mathcal{F}, u)$ to $M_{\mathrm{dR}}^+(u) \circ M_{\mathrm{dR}}^+(\mathcal{F})$. For the functor in the other direction, set $M' = (M_e(\mathcal{E}), N)$. This is a B-pair, and by construction there is a natural injection of B-pairs $\iota : M' \to M(\mathcal{E})$ which is an isomorphism on the first factor. Set $\mathcal{F} = \mathcal{V}(M')$ and $u = \mathcal{V}(\iota)$; since $\mathcal{V}(-)$ is an equivalence, $u : \mathcal{F} \to \mathcal{E}$ is injective. The remaining verifications are an easy unwinding. □

Remark. If $\mathcal{E}$ is some vector bundle with associated B-pair (M_e, M_{dR}^+), and $N \subseteq M_{\mathrm{dR}}^+$ is any $\mathbb{B}_{\mathrm{dR}}^+(A)$ submodule such that $N[\frac{1}{\xi}] = M_{\mathrm{dR}}^+[\frac{1}{\xi}] = M_{\mathrm{dR}}$, then the following are equivalent:

1. N is a finite projective $\mathbb{B}_{\mathrm{dR}}^+(A)$-module.
2. M_{dR}^+/N is finite projective virtually over A.
3. The pair (M_e, N) is in the essential image of $M(-)$ (in which case $\mathcal{V}(M_e, N)$ is an effective modification of $\mathcal{E}$ along S).

Indeed, 1. and 2. are equivalent by Lemma 1.7.i, and 1. and 3. are equivalent by Theorem 1.12. This explains the condition in the previous theorem.

If $\mathcal{E}$ is a fixed vector bundle and $(\mathcal{F}, u)$ is an effective modification along S with associated $N \subseteq M_{\mathrm{dR}}^+(\mathcal{E})$, then for any point $x \in |S|$ we define the *type* of the modification at x, denoted $\mu_x(\mathcal{F}, u)$, as the ordered sequence of elementary divisors of the finite torsion $\mathbb{B}_{\mathrm{dR}}^+(K_x)$-module

$$(M_{\mathrm{dR}}^+(\mathcal{E})/N) \otimes_{\mathbb{B}_{\mathrm{dR}}^+(A)} \mathbb{B}_{\mathrm{dR}}^+(K_x).$$

The key result (which will easily imply Theorem 1.3) is now as follows. We have not stated the most general version.

Theorem 1.15. *Let $S = \mathrm{Spa}(A, A^+)$ be as above, and let $(\mathcal{E}, \mathcal{F}, u)$ be an effective modification along S of constant type $\mu = (k_1 \geq k_2 \geq \dots)$ such that $\mathcal{F}$ is pointwise semistable of slope zero. Let $\mathcal{E}^+ \subseteq \mathcal{E}$ be a subbundle with the property that for every point $x \in |S|$, $\mathcal{E}_x^+ \subseteq \mathcal{E}_x$ is saturated and we have an equality*

$$\deg(\mathcal{E}_x^+) = \sum_{1 \leq i \leq \mathrm{rank}(\mathcal{E}_x^+)} k_i.$$

Then the sheaf $\mathcal{F}^+ = \mathcal{F} \cap \mathcal{E}^+$ defines a sub-vector bundle of $\mathcal{F}$, and the bundle $\mathcal{F}^+$ is pointwise semistable of slope zero.

Proof. We argue at the level of B-pairs over A. Precisely, set $Q = M_{\mathrm{dR}}^+(\mathcal{E})/M_{\mathrm{dR}}^+(\mathcal{F})$; this is a $\mathbb{B}_{\mathrm{dR}}^+(A)$-module which is fpv over A by Theorem 1.14. Consider the $\mathbb{B}_{\mathrm{dR}}^+(A)$-submodule

$$Q^+ = \mathrm{im}(M_{\mathrm{dR}}^+(\mathcal{E}^+) \to Q)$$

of Q; this is finitely generated and ξ-torsion. We are going to prove that Q^+ is fpv over A. Granted this, Proposition 1.7 implies that

$$N = \ker(M^+_{\mathrm{dR}}(\mathcal{E}^+) \to Q^+) = M^+_{\mathrm{dR}}(\mathcal{E}^+) \cap M^+_{\mathrm{dR}}(\mathcal{F})$$

is a finite projective $\mathbb{B}^+_{\mathrm{dR}}(A)$-module. Then $(M_e(\mathcal{E}^+), N)$ defines a B-pair, and we obtain $\mathcal{F}^+$ as the associated vector bundle.

To show that Q^+ is fpv over A, we note that it sits in a short exact sequence

$$0 \to Q^+ \to Q \to Q_- \to 0,$$

where

$$\begin{aligned} Q_- &= M^+_{\mathrm{dR}}(\mathcal{E}) / \left(M^+_{\mathrm{dR}}(\mathcal{F}) + M^+_{\mathrm{dR}}(\mathcal{E}^+) \right) \\ &= \operatorname{coker}\left(M^+_{\mathrm{dR}}(\mathcal{F}) \oplus M^+_{\mathrm{dR}}(\mathcal{E}^+) \to M^+_{\mathrm{dR}}(\mathcal{E}) \right). \end{aligned}$$

Since Q is fpv, Proposition 1.7 shows it suffices to prove Q_- is fpv over A. We're going to do this by applying the pointwise criterion from Proposition 1.10.

Note that unlike Q^+ (at least a priori), Q and Q_- interact well with specializing to points $x \in |S|$. In particular, for any $x \in |S|$ we have a commutative diagram of $\mathbb{B}^+_{\mathrm{dR}}(K_x)$-modules

$$\begin{tikzcd} & 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] & \\ 0 \arrow[r] & M^+_{\mathrm{dR}}(\mathcal{E}^+_x) \cap M^+_{\mathrm{dR}}(\mathcal{F}_x) \arrow[r] \arrow[d] & M^+_{\mathrm{dR}}(\mathcal{E}^+_x) \arrow[r] \arrow[d] & T_x \arrow[r] \arrow[d] & 0 \\ 0 \arrow[r] & M^+_{\mathrm{dR}}(\mathcal{F}_x) \arrow[r] \arrow[d] & M^+_{\mathrm{dR}}(\mathcal{E}_x) \arrow[r] \arrow[d] & Q \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \mathbb{B}^+_{\mathrm{dR}}(K_x) \arrow[r] \arrow[d] & 0 \\ 0 \arrow[r] & S_x \arrow[r] \arrow[d] & M^+_{\mathrm{dR}}(\mathcal{E}_x)/M^+_{\mathrm{dR}}(\mathcal{E}^+_x) \arrow[r] \arrow[d] & Q_- \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \mathbb{B}^+_{\mathrm{dR}}(K_x) \arrow[r] \arrow[d] & 0 \\ & 0 & 0 & 0 & \end{tikzcd}$$

with exact rows and columns and with everything in the first two columns finitely generated and free. By hypothesis, the elementary divisors of $Q \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \mathbb{B}^+_{\mathrm{dR}}(K_x)$ are constant and simply given by the k_i's in the theorem. But now putting together the pointwise hypotheses in the theorem with Theorem 1.4 and Lemma 1.5, we see that T_x is a direct summand of $Q \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \mathbb{B}^+_{\mathrm{dR}}(K_x)$ with elementary divisors $k_1 \geq \cdots \geq k_{\mathrm{rank}(\mathcal{E}^+_x)}$ for any $x \in |S|$, so $Q_- \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \mathbb{B}^+_{\mathrm{dR}}(K_x)$ has elementary divisors

$$k_{\mathrm{rank}(\mathcal{E}^+_x)+1} \geq k_{\mathrm{rank}(\mathcal{E}^+_x)+2} \geq \ldots$$

for any $x \in |S|$. In particular, since $\mathrm{rank}(\mathcal{E}^+_x)$ is locally constant, the elementary divisors of $Q_- \otimes_{\mathbb{B}^+_{\mathrm{dR}}(A)} \mathbb{B}^+_{\mathrm{dR}}(K_x)$ are locally constant. Thus Proposition 1.10 applies, and so Q_- is fpv over A. This completes the proof. □

1.5 More about the retraction

Fix all data as in the leadup to Theorem 1.3. What can we say about the fibers of the retraction r?

Theorem 1.16. *The natural action map*

$$\mathrm{Gr}^{\mathcal{E}_b-\mathrm{adm}}_{\mathbf{M},\mu} \times_{\operatorname{Spd} \breve{\mathbf{Q}}_p} \mathcal{J}^{\mathbf{U}}_b \to \mathrm{Gr}^{\mathcal{E}_b-\mathrm{adm}}_{\mathbf{G},\mu}$$

is pro-étale-locally surjective.

Proof. We argue as follows. Let $f : S = \mathrm{Spa}(A, A^+) \to \mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}}$ be an S-point, with $(\mathcal{F}, u)$ the corresponding admissible type-μ modification of $\mathcal{E}_{b,S^\flat}$ along S. Let $0 \subsetneq \mathcal{F}^{i_1} \subsetneq \cdots \subsetneq \mathcal{F}^{i_k} = \mathcal{F}$ be the canonical $\mathbf{P}$-bundle structure on $\mathcal{F}$ (where $k = |\mathcal{I}|$ as before). Then applying the canonical retraction, i.e. looking at the point $r \circ f : \mathrm{Spa}(A, A^+) \to \mathrm{Gr}_{\mathbf{M},\mu}^{\mathcal{E}_b-\mathrm{adm}}$, we get a collection $(\mathcal{F}_m, u_m)_{1\leq m\leq k}$ of admissible type-μ_m modifications of the summands $\mathcal{E}_{b_m,S^\flat}$. The point $i \circ r \circ f$ then corresponds to viewing $\oplus_{1\leq m\leq k}(\mathcal{F}_m, u_m)$ as a type-μ modification of $\mathcal{E}_{b,S^\flat} \cong \oplus_{1\leq m\leq k}\mathcal{E}_{b_m,S^\flat}$. We're going to (pro-étale-locally on S) find an element $j \in \mathcal{J}_b^{\mathbf{U}}(S)$ which transports the point $i \circ r \circ f$ to the point f.

Now, the fact that f and $i \circ r \circ f$ have the same retraction translates into the following fact: After choosing compatible isomorphisms $\iota_m : \mathcal{F}_1 \oplus \mathcal{F}_2 \oplus \cdots \oplus \mathcal{F}_m \simeq \mathcal{F}^{i_m}$ (which we can do pro-étale-locally on S), the compatible-in-m maps

$$\nu_m : u|_{\mathcal{F}^{i_m}} \circ \iota_m : \mathcal{F}_1 \oplus \mathcal{F}_2 \oplus \cdots \oplus \mathcal{F}_m \to \mathcal{E}^{i_m} \cong \oplus_{1\leq i\leq m}\mathcal{E}_{b_i}$$

and

$$\eta_m : u_1 \oplus \cdots \oplus u_m : \mathcal{F}_1 \oplus \mathcal{F}_2 \oplus \cdots \oplus \mathcal{F}_m \to \mathcal{E}^{i_m} \cong \mathcal{E}^{i_m} \cong \oplus_{1\leq i\leq m}\mathcal{E}_{b_i}$$

coincide after projection along $\mathcal{E}^{i_m} \twoheadrightarrow \mathcal{E}_{b_m}$. We are going to show that each $\nu_m \circ \eta_m^{-1}$, which is initially only a *meromorphic* endomorphism of $\mathcal{E}^{i_m}$, actually defines a global section of $(\mathcal{E}^{i_m})^\vee \otimes \mathcal{E}^{i_m}$ such that $\nu_m \circ \eta_m^{-1} - 1$ defines a section of the subbundle $(\mathcal{E}^{i_m})^\vee \otimes \mathcal{E}^{i_{m-1}}$. To do this, note that by an easy induction, each map $\eta_m - \nu_m : \mathcal{F}^{i_m} \to \mathcal{E}^{i_{m-1}}$ has zeros of order $\geq k_{d_1+\cdots+d_{m-1}}$ along $S \subset \mathcal{X}_{S^\flat}$. On the other hand, $u_m^{-1} : \mathcal{E}_{b_m} \to \mathcal{F}_m$ has poles of order $\leq k_{d_1+\cdots+d_{m-1}+1}$ along S.[6] Now, formally, we have the identity

$$\begin{aligned}
\nu_m \circ \eta_m^{-1} &= \nu_{m-1} \circ \eta_{m-1}^{-1} + \nu_m \circ u_m^{-1} \\
&= \nu_{m-1} \circ \eta_{m-1}^{-1} + (\nu_m - \eta_m + \eta_m) \circ u_m^{-1} \\
&= \nu_{m-1} \circ \eta_{m-1}^{-1} + (\nu_m - \eta_m) \circ u_m^{-1} + \eta_m \circ u_m^{-1}.
\end{aligned}$$

But $(\nu_m - \eta_m) \circ u_m^{-1} : \mathcal{E}_{b_m} \to \mathcal{E}^{i_{m-1}}$ is well-defined by our previous remarks on zeros and poles, and $\eta_m \circ u_m^{-1} : \mathcal{E}_{b_m} \to \mathcal{E}^{i_m}$ is just the canonical inclusion as a direct summand. Thus we get the desired properties of $\nu_m \circ \eta_m^{-1}$ by induction, noting that $\nu_1 \circ \eta_1^{-1} = \mathrm{id}$. But this analysis shows that the section

$$j = \nu_k \circ \eta_k^{-1} \in H^0(\mathcal{X}_{S^\flat}, \mathcal{E}_b^\vee \otimes \mathcal{E}_b)$$

defines an element of $\mathcal{J}_b^{\mathbf{U}}(S)$, and by construction it transports $\oplus_{1\leq m\leq k}(\mathcal{F}_m, u_m)$ to

$$(\mathcal{F}, u) \simeq (\oplus_{1\leq m\leq k}\mathcal{F}_m, u \circ \iota_k),$$

so we're done. □

1.6 Adding infinite level structure, and cohomological consequences

Setup as in Theorem 1.3. For brevity, we set $G, M, P = \mathbf{G}(\mathbf{Q}_p), \mathbf{M}(\mathbf{Q}_p)$,etc. We've already defined spaces $\mathrm{Sht}_{\mathbf{G},\mu,b}$ and $\mathrm{Sht}_{\mathbf{M},\mu,b}$. These sit in a diagram

$$\begin{tikzcd}
\mathrm{Sht}_{\mathbf{M},\mu,b} \arrow[r] \arrow[d, "\pi_{\mathbf{M}}"'] & \mathrm{Sht}_{\mathbf{G},\mu,b} \arrow[d, "\pi_{\mathbf{G}}"'] \\
\mathrm{Gr}_{\mathbf{M},\mu}^{\mathcal{E}_b-\mathrm{adm}} \arrow[r] & \mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}}
\end{tikzcd}$$

of diamonds over $\mathrm{Spd}\,\breve{\mathbf{Q}}_p$, where $\pi_{\mathbf{M}}$ (resp. $\pi_{\mathbf{G}}$) presents its source as a pro-étale M-torsor (resp. G-torsor) over its target. We've already proved that the lower horizontal arrow retracts in a canonical way. Now we'd like to study the upper arrow.

[6] It's not hard to make these statements about poles and zeros precise; the point is that the ideal sheaf cutting out $S \subset \mathcal{X}_{S^\flat}$ is locally principal and generated by a non-zero-divisor.

Following Mantovan, we do this by defining an intermediate space $\mathrm{Sht}_{\mathbf{P},\mu,b}$ of $\mathbf{P}$-shtukas. The precise definition is as follows: for S a perfectoid space over $\breve{\mathbf{Q}}_p$, the S-points of $\mathrm{Sht}_{\mathbf{P},\mu,b}$ consist of triples $(\mathcal{F}, u, \alpha)$ where $(\mathcal{F}, u)$ corresponds to an S-point of $\mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}}$ and $\alpha : \mathcal{O}^n_{\mathcal{X}_{S^\flat}} \xrightarrow{\sim} \mathcal{F}$ is a trivialization matching the flag

$$0 \subsetneq \mathcal{O}^{d_1}_{\mathcal{X}_{S^\flat}} \oplus 0 \subsetneq \mathcal{O}^{d_1+d_2}_{\mathcal{X}_{S^\flat}} \oplus 0 \subsetneq \cdots \subsetneq \mathcal{O}^n_{\mathcal{X}_{S^\flat}}$$

with the flag

$$0 \subsetneq \mathcal{F}^{i_1} \subsetneq \mathcal{F}^{i_2} \subsetneq \cdots \subsetneq \mathcal{F}.$$

In particular, we have inclusions of subfunctors

$$\mathrm{Sht}_{\mathbf{M},\mu,b} \overset{(1)}{\subset} \mathrm{Sht}_{\mathbf{P},\mu,b} \overset{(2)}{\subset} \mathrm{Sht}_{\mathbf{G},\mu,b},$$

and there is a natural action of P on $\mathrm{Sht}_{\mathbf{P},\mu,b}$ compatible with the M- and G-actions on $\mathrm{Sht}_{\mathbf{M},\mu,b}$ and $\mathrm{Sht}_{\mathbf{G},\mu,b}$. There is also a natural action of $\mathcal{J}_b = \mathcal{J}_b^{\mathbf{M}} \ltimes \mathcal{J}_b^{\mathbf{U}}$ on $\mathrm{Sht}_{\mathbf{P},\mu,b}$ making the inclusions (1) and (2) $\mathcal{J}_b^{\mathbf{M}}$-equivariant and $\mathcal{J}_b$-equivariant, respectively.

We are going to prove the following two things.

Theorem 1.17. *The inclusion* $\mathrm{Sht}_{\mathbf{P},\mu,b} \subset \mathrm{Sht}_{\mathbf{G},\mu,b}$ *induces a canonical equivariant identification*

$$\mathrm{Sht}_{\mathbf{G},\mu,b} \cong \mathrm{Sht}_{\mathbf{P},\mu,b} \times^{\underline{P}} \underline{G}.$$

In particular, $\mathrm{Sht}_{\mathbf{P},\mu,b} \to \mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}}$ *is a pro-étale P-torsor, and there is a canonical G-equivariant isomorphism*

$$H^*_{\mathrm{et}}\left(\mathrm{Sht}_{\mathbf{G},\mu,b} \times \mathrm{Spd}\, C, \mathbf{Q}_\ell\right) \cong \mathrm{Ind}_P^G\left(H^*_{\mathrm{et}}\left(\mathrm{Sht}_{\mathbf{P},\mu,b} \times \mathrm{Spd}\, C, \mathbf{Q}_\ell\right)\right)$$

preserving degrees and compatible with all additional structures; here $C/\breve{\mathbf{Q}}_p$ denotes any complete algebraically closed field. The same formula holds for compactly supported cohomology.

Indeed, we've essentially already shown this. On the other hand, we prove

Theorem 1.18. *The natural action map*

$$a : \mathrm{Sht}_{\mathbf{M},\mu,b} \times_{\mathrm{Spd}\, \breve{\mathbf{Q}}_p} \mathcal{J}_b^{\mathbf{U}} \to \mathrm{Sht}_{\mathbf{P},\mu,b}$$

is an isomorphism of diamonds.

In particular, the product $\mathrm{Sht}_{\mathbf{M},\mu,b} \times_{\mathrm{Spd}\, \breve{\mathbf{Q}}_p} \mathcal{J}_b^{\mathbf{U}}$ admits a canonical P-action; we caution the reader that although the action of $M \subset P$ is indeed the obvious one, given by its action on the first factor, the full P-action mixes both factors in a way which is a little tricky to describe directly.

Proof. We construct a two-sided inverse to a. Let $S \in \mathrm{Perf}_{/\mathrm{Spa}\, \breve{\mathbf{Q}}_p}$ and $(\mathcal{F}, u, \alpha) \in \mathrm{Sht}_{\mathbf{P},\mu,b}(S)$ be given. We need to construct a point

$$\prod_{1 \leq m \leq k} (\mathcal{G}_m, v_m, \beta_m) \in \mathrm{Sht}_{\mathbf{M},\mu,b}(S)$$

and an element $j \in \mathcal{J}_b^{\mathbf{U}}(S)$. The first is easier to find: applying the retraction on period domains to $(\mathcal{F}, u) \in \mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}}(S)$ gives a point

$$(\mathrm{gr}(\mathcal{F}), \mathrm{gr}(u)) = \prod_{1 \leq m \leq k} (\mathcal{F}_m, u_m) \in \mathrm{Gr}_{\mathbf{M},\mu}^{\mathcal{E}_b-\mathrm{adm}}(S).$$

Now by the definition of $\mathrm{Sht}_{\mathbf{P},\mu,b}$, it's easy to see check that "$\mathrm{gr}(\alpha)$" gives a well-defined sequence of trivializations $\alpha_m : \mathcal{O}^{d_m}_{\mathcal{X}_{S^\flat}} \xrightarrow{\sim} \mathcal{F}_m$, and this gives a point

$$(\mathrm{gr}(\mathcal{F}), \mathrm{gr}(u), \mathrm{gr}(\alpha)) = \prod_{1 \leq m \leq k} (\mathcal{F}_m, u_m, \alpha_m) \in \mathrm{Sht}_{\mathbf{M},\mu,b}(S)$$

as desired.

To construct j, recall from the proof of Theorem 1.16 that after making any choices of compatible isomorphisms $\iota_m : \mathcal{F}_1 \oplus \mathcal{F}_2 \oplus \cdots \oplus \mathcal{F}_m \simeq \mathcal{F}^{i_m}$ $(1 \leq m \leq k)$, the two maps

$$\nu_k : u \circ \iota_k : \mathcal{F}_1 \oplus \mathcal{F}_2 \oplus \cdots \oplus \mathcal{F}_k \to \mathcal{E}_b \cong \oplus_{1 \leq i \leq k} \mathcal{E}_{b_i}$$

and

$$\eta_k : u_1 \oplus \cdots \oplus u_k : \mathcal{F}_1 \oplus \mathcal{F}_2 \oplus \cdots \oplus \mathcal{F}_k \to \mathcal{E}_b \cong \oplus_{1 \leq i \leq k} \mathcal{E}_{b_i}$$

have the property that $\nu_k \circ \eta_k^{-1}$ defines an element of $\mathcal{J}_b^{\mathbf{U}}$. Now we simply observe that at infinite level, there is a canonical choice for the ι_m's, as indicated by the diagram

$$\begin{tikzcd} \mathcal{F}_1 \oplus \mathcal{F}_2 \oplus \cdots \oplus \mathcal{F}_m & \mathcal{O}_{\mathcal{X}}^{d_1+\cdots+d_m} \arrow[l, "\sim"', "\alpha_1 \oplus \cdots \oplus \alpha_m"] \arrow[r, "\sim", "\alpha|(\mathcal{O}_{\mathcal{X}}^{d_1+\cdots+d_m} \oplus 0)"'] & \mathcal{F}^{i_m} \end{tikzcd} .$$

In other words, we take

$$\iota_k = \alpha \circ \left(\alpha_1^{-1} \oplus \cdots \oplus \alpha_k^{-1}\right) = \alpha \circ \mathrm{gr}(\alpha)^{-1},$$

and then

$$\nu_k \circ \eta_k^{-1} = u \circ \iota_k \circ \mathrm{gr}(u)^{-1} = u \circ \alpha \circ \mathrm{gr}(\alpha)^{-1} \circ \mathrm{gr}(u)^{-1} \in \mathcal{J}_b^{\mathbf{U}}(S)$$

is the element we seek. □

Corollary 1.19. *The retraction r lifts canonically to a retraction of the natural inclusion $\mathrm{Sht}_{\mathbf{M},\mu,b} \subset \mathrm{Sht}_{\mathbf{P},\mu,b}$ with fibers given by $\mathcal{J}_b^{\mathbf{U}}$-torsors. More precisely, the diagram*

$$\begin{tikzcd} & \mathrm{Gr}_{\mathbf{M},\mu}^{\mathcal{E}_b-\mathrm{adm}} \times_{\mathrm{Spd}\,\breve{\mathbf{Q}}_p} \mathcal{J}_b^{\mathbf{U}} \arrow[dr, "a"] & \\ \mathrm{Gr}_{\mathbf{M},\mu}^{\mathcal{E}_b-\mathrm{adm}} \arrow[ur] \arrow[rr, "i"'] & & \mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}} \arrow[ll, bend right, "r"'] \end{tikzcd}$$

has a canonical equivariant lifting to a diagram

$$\begin{tikzcd} & \mathrm{Sht}_{\mathbf{M},\mu,b} \times_{\mathrm{Spd}\,\breve{\mathbf{Q}}_p} \mathcal{J}_b^{\mathbf{U}} \arrow[dr, equal, "a"] & \\ \mathrm{Sht}_{\mathbf{M},\mu,b} \arrow[ur] \arrow[rr, "i"'] & & \mathrm{Sht}_{\mathbf{P},\mu,b} \arrow[ll, bend right, "r"'] \end{tikzcd}$$

with the map r given as a^{-1} followed by the projection $\mathrm{Sht}_{\mathbf{M},\mu,b} \times_{\mathrm{Spd}\,\breve{\mathbf{Q}}_p} \mathcal{J}_b^{\mathbf{U}} \to \mathrm{Sht}_{\mathbf{M},\mu,b}$.

The following diagram summarizes the situation in a manner which we hope is suggestive:

$$\begin{tikzcd} & \mathrm{Sht}_{\mathbf{P},\mu,b} \arrow[dl, bend right, "r"'] \arrow[dr, "\mathrm{Ind}_P^G(-)"] & \\ \mathrm{Sht}_{\mathbf{M},\mu,b} \arrow[ur, "i"'] \arrow[rr] \arrow[d, "\pi_{\mathbf{M}}"'] & & \mathrm{Sht}_{\mathbf{G},\mu,b} \arrow[d, "\pi_{\mathbf{G}}"] \\ \mathrm{Gr}_{\mathbf{M},\mu}^{\mathcal{E}_b-\mathrm{adm}} \arrow[rr, "i"'] & & \mathrm{Gr}_{\mathbf{G},\mu}^{\mathcal{E}_b-\mathrm{adm}} \arrow[ll, bend right, "r"'] \end{tikzcd}$$

With the previous two theorems in hand, it just remains to calculate the geometric étale cohomology of $\mathcal{J}_b^{\mathbf{U}}$, which turns out to be very simple:

Theorem 1.20. *For any complete algebraically closed field $C/\breve{\mathbf{Q}}_p$, we have*

$$R\Gamma_{\mathrm{et}}(\mathcal{J}_b^{\mathbf{U}} \times \operatorname{Spd} C, \mathbf{Q}_\ell) \simeq \mathbf{Q}_\ell[0]$$

and

$$R\Gamma_{\mathrm{et},c}(\mathcal{J}_b^{\mathbf{U}} \times \operatorname{Spd} C, \mathbf{Q}_\ell) \simeq \mathbf{Q}_\ell(-d)[-2d],$$

where

$$d = \dim \mathcal{J}_b^{\mathbf{U}} = \left\langle 2\rho_{\mathbf{U}}, \nu_{[b^{-1}]} \right\rangle = \dim \mathrm{Sht}_{\mathbf{G},\mu,b} - \dim \mathrm{Sht}_{\mathbf{M},\mu,b}.$$

Sketch. After unwinding, $\mathcal{J}_b^{\mathbf{U}}$ is a finite product of diamonds $\mathbb{B}_{\mathrm{crys}}^{+,\varphi=p^j}$ for varying j (viewed as diamonds over $\operatorname{Spd} C$). One computes their cohomology by induction on j. For $j=1$, use that $\mathbb{B}_{\mathrm{crys}}^{+,\varphi=p}$ is representable by a perfectoid space X; passing to the tilt, $X^\flat$ is just the perfection of the open unit disk over $C^\flat$, and the $\mathbf{Q}_\ell$-cohomology of the latter was calculated by Berkovich. For $j>1$, use the "fibration" sequence

$$0 \to \mathbb{B}_{\mathrm{crys}}^{+,\varphi=p^{j-1}} \xrightarrow{\cdot t} \mathbb{B}_{\mathrm{crys}}^{+,\varphi=p^j} \xrightarrow{\theta} \mathbf{A}^{1,\diamondsuit} \to 0$$

(which splits pro-étale-locally on $\mathbf{A}^{1,\diamondsuit}$) together with Berkovich's computation of the $\mathbf{Q}_\ell$-cohomology of $\mathbf{A}^1$. □

The main cohomological result is then as follows.

Theorem 1.21. *There are canonical G-equivariant isomorphisms*

$$H^*_{\mathrm{et}}\left(\mathrm{Sht}_{\mathbf{G},\mu,b} \times_{\operatorname{Spd} \breve{\mathbf{Q}}_p} \operatorname{Spd} C, \mathbf{Q}_\ell\right) \cong \mathrm{Ind}_P^G\left(H^*_{\mathrm{et}}\left(\mathrm{Sht}_{\mathbf{M},\mu,b} \times_{\operatorname{Spd} \breve{\mathbf{Q}}_p} \operatorname{Spd} C, \mathbf{Q}_\ell\right)\right)$$

and

$$H^*_{\mathrm{et},c}\left(\mathrm{Sht}_{\mathbf{G},\mu,b} \times_{\operatorname{Spd} \breve{\mathbf{Q}}_p} \operatorname{Spd} C, \mathbf{Q}_\ell\right) \cong \mathrm{Ind}_P^G\left(H^{*-2d}_{\mathrm{et},c}\left(\mathrm{Sht}_{\mathbf{M},\mu,b} \times_{\operatorname{Spd} \breve{\mathbf{Q}}_p} \operatorname{Spd} C, \mathbf{Q}_\ell\right)(-d)\right)$$

preserving degrees and compatible with all additional structures; here $d = \dim \mathrm{Sht}_{\mathbf{G},\mu,b} - \dim \mathrm{Sht}_{\mathbf{M},\mu,b}$, and $C/\breve{\mathbf{Q}}_p$ denotes any complete algebraically closed field.

We remark that there are no subtleties in defining the compactly supported (pro-)étale cohomology groups occurring here: all the diamonds in question are spatial and partially proper, so these cohomologies are just the derived functors of global sections with compact support. Strictly speaking, in the deduction of Theorem 1.21 from Theorem 1.20, we're appealing to some kind of Künneth formula/smooth base change theorem for $\mathbf{Q}_\ell$-cohomology of diamonds, but such a thing is a moral certainty in this situation: all diamonds in play are spatial and partially proper, and $\mathcal{J}_b^{\mathbf{U}}$ is smooth in a very strong sense.[7]

References

[KL1] K. Kedlaya and R. Liu, *Relative p-adic Hodge theory, I: Foundations*

[KL2] K. Kedlaya and R. Liu, *Relative p-adic Hodge theory, II: Imperfect period rings*

[7] It has a pro-étale covering by the perfection of a smooth rigid space over $C^\flat$.