

EXTENSIONS OF VECTOR BUNDLES ON THE FARGUES-FONTAINE CURVE

CHRISTOPHER BIRKBECK¹, TONY FENG², DAVID HANSEN³,
SERIN HONG⁴, QIRUI LI⁷, ANTHONY WANG⁵ AND LYNNELLE YE⁶

¹*Department of Mathematics, University College London, Gower street,
WC1E 6BT (c.birkbeck@ucl.ac.uk)*

²*MIT Department of Mathematics, 182 Memorial Dr., Cambridge, MA 02142
(fengt@mit.edu)*

³*Max Planck Institute for Mathematics, Vivatsgasse 7, 53111 Bonn, Germany
(dhansen@mpim-bonn.mpg.de)*

⁴*Department of Mathematics, University of Michigan, 530 Church Street,
Ann Arbor, MI 48109, USA (serinh@umich.edu)*

⁵*Department of Mathematics, University of Chicago, 5734 S. University Avenue,
Chicago, IL 60637, USA (anthonyw@math.uchicago.edu)*

⁶*Department of Mathematics, Building 380, Stanford, California 94305
(lynnelle@stanford.edu)*

⁷*Department of Mathematics, Columbia University, 2990 Broadway, New York,
10 NY 10027, USA (qiruili@math.columbia.edu)*

(Received 4 April 2020; accepted 4 April 2020;
first published online 14 May 2020)

Abstract We completely classify the possible extensions between semistable vector bundles on the Fargues–Fontaine curve (over an algebraically closed perfectoid field), in terms of a simple condition on Harder–Narasimhan (HN) polygons. Our arguments rely on a careful study of various moduli spaces of bundle maps, which we define and analyze using Scholze’s language of diamonds. This analysis reduces our main results to a somewhat involved combinatorial problem, which we then solve via a reinterpretation in terms of the Euclidean geometry of HN polygons.

Keywords: Fargues–Fontaine; Diamonds; Vector bundles

2010 *Mathematics subject classification:* Primary 14D24

Secondary 14G45

DH is grateful to Christian Johansson for some useful conversations about the material in § 3.2, and Peter Scholze for providing early access to the manuscript [13] and for some helpful conversations about the results therein. The project group students (CB, TF, SH, QL, AW, and LY) thank DH and Kiran Kedlaya for suggesting the problem. TF gratefully acknowledges the support of an NSF Graduate Fellowship. LY gratefully acknowledges the support of the National Defense Science and Engineering Graduate Fellowship. We would also like to thank David Linus Hamann and the referee for their valuable feedback on the first version of this paper.

Contents

1 Introduction **488**

1.1 Statement of results 489

1.2 The basic strategy 491

2 Background on the Fargues–Fontaine curve **492**

2.1 The Fargues–Fontaine curve and vector bundles 492

2.2 Harder–Narasimhan filtrations and polygons 494

2.3 Geometric interpretation of degrees 497

3 Diamonds, vector bundles, and dimensions **499**

3.1 Recollections on diamonds 499

3.2 Dimension theory 502

3.3 Diamonds and moduli of bundle maps 506

3.4 Summary of the strategy 513

4 Step one: surjections of vector bundles **514**

4.1 The goal 514

4.2 Some lemmas 515

4.3 Proof of Theorem 4.1.1 516

5 Step two: extensions of semistable bundles **520**

5.1 The key inequality 520

5.2 Proof sketch of Theorem 5.1.1 521

5.3 Formal proof of Theorem 5.1.1 523

5.4 The extension theorem 528

6 Multistep filtrations **529**

References **531**

1. Introduction

Let E be a p -adic local field with residue field \mathbf{F}_q , and let F/\mathbf{F}_q be an algebraically closed complete nonarchimedean field of characteristic p . Given any such pair, Fargues and Fontaine [4] defined a remarkable scheme $X = X_{E,F}$, the so-called *Fargues–Fontaine curve*. Many constructions in classical p -adic Hodge theory can be reinterpreted ‘geometrically’ in terms of vector bundles on X .

One of the main results of [4] is a complete classification of vector bundles on X . In particular, Fargues and Fontaine proved that any bundle is determined up to isomorphism by its Harder–Narasimhan (HN) polygon, that the HN filtration of any bundle splits, and that there is a unique isomorphism class of stable bundles of any specified slope $\lambda \in \mathbf{Q}$.

In this paper, we study the question of which bundles can occur as *extensions* between two given vector bundles on X . This question turns out to have a rather combinatorial flavor, since vector bundles on the curve are determined by their HN polygons. Moreover, abstract slope theory imposes certain conditions on the HN polygon of any bundle \mathcal{E}

appearing as an extension of two specified bundles $\mathcal{F}_1, \mathcal{F}_2$. Our main result is that when both bundles \mathcal{F}_i are semistable, these necessary conditions are also sufficient; for a precise statement, see Theorem 1.1.2. We also prove a natural generalization characterizing bundles admitting multistep filtrations with specified semistable graded pieces; cf. Theorem 1.1.4.

The results in this paper have applications toward understanding the geometry of the stack \mathbf{Bun}_G of G -bundles on the Fargues-Fontaine curve, for G a connected reductive group over E . More precisely, this stack has a natural stratification into locally closed strata \mathbf{Bun}_G^b indexed by the Kottwitz set $B(G)$, and our results have implications for the problem of computing the closures of the individual strata. In the case $G = \mathrm{GL}_n$, the aforementioned strata are simply indexed by HN polygons, and the exact statement is given in Theorem 1.1.3. These applications are the subject of a companion paper by one of us [6].

1.1. Statement of results

Before stating our results, we briefly recall the classification of vector bundles on X . A more extensive discussion of this and related background will be given in §2.

Theorem 1.1.1 (Fargues-Fontaine, Kedlaya). [4, 9] *Vector bundles on X enjoy the following properties:*

- (1) *Every vector bundle \mathcal{E} admits a canonical HN filtration.*
- (2) *For every rational number λ , there is a unique stable bundle of slope λ on X , which is denoted by $\mathcal{O}(\lambda)$. Writing $\lambda = p/q$ in lowest terms, the bundle $\mathcal{O}(\lambda)$ has rank q and degree p .*
- (3) *Any semistable bundle of slope λ is a finite direct sum $\mathcal{O}(\lambda)^d$, and tensor products of semistable bundles are semistable.*
- (4) *For any $\lambda \in \mathbf{Q}$, we have*

$$H^0(\mathcal{O}(\lambda)) = 0 \quad \text{if and only if } \lambda < 0$$

and

$$H^1(\mathcal{O}(\lambda)) = 0 \quad \text{if and only if } \lambda \geq 0.$$

In particular, any vector bundle \mathcal{E} admits a splitting

$$\mathcal{E} \simeq \bigoplus_i \mathcal{O}(\lambda_i)$$

of its HN filtration.

This project began at the 2017 Arizona Winter School, as an investigation of the following question:

Given two vector bundles \mathcal{F}_1 and \mathcal{F}_2 on X , which polygons can occur as $\mathrm{HN}(\mathcal{E})$ for an extension

$$0 \rightarrow \mathcal{F}_1 \rightarrow \mathcal{E} \rightarrow \mathcal{F}_2 \rightarrow 0?$$




Figure 1. An illustration of the scenario $\text{HN}(\mathcal{E}) \leq \text{HN}(\mathcal{F}_1 \oplus \mathcal{F}_2)$ for semistable \mathcal{F}_i 's. Here $\text{HN}(\mathcal{E})$ (resp. $\text{HN}(\mathcal{F}_1 \oplus \mathcal{F}_2)$) is depicted as the dashed (resp. solid) polygonal segment.

Given such a short exact sequence, the total rank and degree of \mathcal{E} are clearly determined by \mathcal{F}_1 and \mathcal{F}_2 , and they determine the endpoints of $\text{HN}(\mathcal{E})$. By concavity, $\text{HN}(\mathcal{E})$ must lie above the straight line between its endpoints, which corresponds to the HN polygon of the unique semistable bundle of the correct degree and rank. Slope theory also imposes a less trivial constraint on $\text{HN}(\mathcal{E})$. To explain this, let us define a partial order on the set of HN polygons by writing $P \leq P'$ if P lies (nonstrictly) below P' and has the same right endpoint. It is then not difficult to show that $\text{HN}(\mathcal{E}) \leq \text{HN}(\mathcal{F}_1 \oplus \mathcal{F}_2)$; cf. Corollary 2.2.13.

Our first main result is that when \mathcal{F}_1 and \mathcal{F}_2 are semistable, every bundle \mathcal{E} satisfying these conditions is actually realized. Given a bundle \mathcal{E} , let $\mu(\mathcal{E})$ be the slope of \mathcal{E} .

Theorem 1.1.2. *Let \mathcal{F}_1 and \mathcal{F}_2 be semistable vector bundles on X such that $\mu(\mathcal{F}_1) < \mu(\mathcal{F}_2)$. Then any vector bundle \mathcal{E} such that $\text{HN}(\mathcal{E}) \leq \text{HN}(\mathcal{F}_1 \oplus \mathcal{F}_2)$ is realized as an extension*

$$0 \rightarrow \mathcal{F}_1 \rightarrow \mathcal{E} \rightarrow \mathcal{F}_2 \rightarrow 0.$$

For a more quantitative version of this result, see Theorem 5.4.2.

In § 1.2, we will give a detailed summary of the proof. For now let us simply remark that even though Theorem 1.1.2 is purely a statement about vector bundles on a Noetherian scheme, our proof uses *diamonds* in a crucial way. Furthermore, we believe that any natural proof of this result will make heavy use of diamonds; in our arguments, they arise as the correct framework for constructing moduli spaces of bundle maps with specified properties. We challenge the skeptical reader to produce a short exact sequence

$$0 \rightarrow \mathcal{O}(-\frac{1}{2})^2 \rightarrow \mathcal{O}(\frac{1}{3}) \oplus \mathcal{O}(\frac{6}{5}) \rightarrow \mathcal{O}(\frac{9}{4}) \rightarrow 0$$

by an argument that does not involve diamonds.

This theorem has a natural generalization to multistep filtrations, which appears as Theorem 1.1.4. Before giving the precise statement, we explain the motivation. As discussed above, the stack Bun_n of rank n vector bundles on the curve admits a stratification

$$\text{Bun}_n = \bigsqcup_P \text{Bun}_n^P,$$

where the stratum Bun_n^P parametrizes vector bundles with fixed HN polygon P . Immediately before the winter school, DH realized that Theorem 1.1.4, if true, could be used to determine the precise closure relations among these strata.

Theorem 1.1.3 (Hansen). *Let P be any HN polygon of width n . Then we have*

$$\overline{\text{Bun}}_n^P = \text{Bun}_n^{\geq P}$$

as substacks of Bun_n , where $\text{Bun}_n^{\geq P}$ denotes the substack parametrizing vector bundles whose HN polygon is $\geq P$.

This is in some sense a converse to the well-known fact that HN polygons jump up (in our convention) on closed subsets (cf. [11, Theorem 7.4.5]).

We now state our second main result, which is a generalization of Theorem 1.1.2.

Theorem 1.1.4. *Let $k \geq 2$ be arbitrary, and let $\mathcal{F}_1, \dots, \mathcal{F}_k$ be semistable vector bundles on X such that $\mu(\mathcal{F}_i) < \mu(\mathcal{F}_{i+1})$ for all $1 \leq i \leq k - 1$. Let \mathcal{E} be any vector bundle on X such that $\text{HN}(\mathcal{E}) \leq \text{HN}(\mathcal{F}_1 \oplus \dots \oplus \mathcal{F}_k)$. Then \mathcal{E} admits a filtration*

$$0 = \mathcal{E}_0 \subset \mathcal{E}_1 \subset \dots \subset \mathcal{E}_k = \mathcal{E}$$

such that $\mathcal{E}_i/\mathcal{E}_{i-1} \simeq \mathcal{F}_i$ for all $1 \leq i \leq k$.

Corollary 2.2.13 again implies that the condition $\text{HN}(\mathcal{E}) \leq \text{HN}(\mathcal{F}_1 \oplus \dots \oplus \mathcal{F}_k)$ is necessary for \mathcal{E} to admit a filtration with successive grades $\mathcal{F}_1, \dots, \mathcal{F}_k$, so this theorem is optimal.

Although Theorem 1.1.4 generalizes Theorem 1.1.2, in fact, the former follows quickly from the latter by induction on k , which we explain in §6. The rough idea is to triangulate the desired HN polygon from Theorem 1.1.4 into pieces, which are then handled by Theorem 1.1.2.

The bulk of the paper is thus concerned with establishing Theorem 1.1.2. We now outline the strategy of our proof.

1.2. The basic strategy

Theorem 1.1.2 asserts that under certain conditions, there exists an extension \mathcal{E} of \mathcal{F}_2 by \mathcal{F}_1 with a specified HN polygon. To prove this, we proceed in two steps (both under the hypotheses of Theorem 1.1.2):

Step (1): We show that \mathcal{E} admits some surjection $\mathcal{E} \twoheadrightarrow \mathcal{F}_2$.

Step (2): We show that if \mathcal{E} admits a surjection $\mathcal{E} \twoheadrightarrow \mathcal{F}_2$, then it admits such a surjection with kernel isomorphic to \mathcal{F}_1 .

Let us sketch the arguments for these steps.¹

Step (1). To carry out Step (1), we construct a moduli space $\text{Hom}(\mathcal{E}, \mathcal{F}_2)$ parametrizing bundle maps $\mathcal{E} \rightarrow \mathcal{F}_2$, as well as an open subspace $\text{Surj}(\mathcal{E}, \mathcal{F}_2)$ parametrizing surjective maps. These objects are easily defined as functors on the category of perfectoid spaces over F , but they typically are not representable by adic spaces. We prove that they are

¹Strictly speaking, the following arguments work only after an easy reduction of Theorem 1.1.2 to the case where $\mu(\mathcal{F}_2)$ is strictly larger than the maximal slope of \mathcal{E} ; see §3.4 for details.

diamonds in the sense of Scholze [13, 15]; in fact, we show that they are ‘locally spatial, equidimensional diamonds’. We emphasize that our use of diamonds here is genuinely necessary for the arguments. In particular, we crucially use the fact that a locally spatial diamond X has a naturally associated locally spectral topological space $|X|$, and that the Krull dimension of this space gives rise to a reasonable dimension theory for such X ’s; cf. §§ 3.1–3.2.

In any case, once we have proved that $\mathrm{Surj}(\mathcal{E}, \mathcal{F}_2)$ is a reasonable diamond, our task is to show that it is nonempty, which we deduce from the following claim.

Claim 1.2.1. The dimension of the space of maps $\mathcal{E} \rightarrow \mathcal{F}_2$ which factor through a proper subbundle of \mathcal{F}_2 is *strictly smaller* than the dimension of $\mathrm{Hom}(\mathcal{E}, \mathcal{F}_2)$.

Roughly speaking, we prove this by stratifying $\mathrm{Hom}(\mathcal{E}, \mathcal{F}_2)$ according to the isomorphy type of the image of the map $\mathcal{E} \rightarrow \mathcal{F}_2$ and then computing the dimensions of these individual strata. The computation of these dimensions is nontrivial, and relies on the theory developed in §§ 3.2–3.3; the final formula is stated in Lemma 3.3.10. Once we have this formula, Claim 1.2.1 is reduced to a certain finite collection of inequalities, whose verification is a mechanical computation in any given example, but which presents a difficult combinatorial problem in general. We eventually realized that it was natural to interpret the terms appearing in the inequality as sums of *areas* of certain polygons (related to the slopes in the HN polygons), and that the inequality could then be seen by a complicated transformation of the polygons after which one area would clearly dominate the other. The precise statement and its proof are given in § 4.

Step (2). The argument here is similar to that for Step (1). For any given bundle \mathcal{K} , we define a moduli space $\mathrm{Surj}(\mathcal{E}, \mathcal{F}_2)^{\mathcal{K}}$ parametrizing surjections $\mathcal{E} \rightarrow \mathcal{F}_2$ with kernel isomorphic to \mathcal{K} . Again, this is a locally spatial diamond. For $\mathcal{K} = \mathcal{F}_1$, this object is open in $\mathrm{Surj}(\mathcal{E}, \mathcal{F}_2)$ for essentially formal reasons (the openness of semistable loci). Again, the task is to show $\mathrm{Surj}(\mathcal{E}, \mathcal{F}_2)^{\mathcal{F}_1}$ is nonempty, which we achieve by showing that $\mathrm{Surj}(\mathcal{E}, \mathcal{F}_2)^{\mathcal{K}}$ has a strictly smaller dimension for any other isomorphism class of \mathcal{K} ’s. This again amounts to a certain collection of inequalities, which we can again interpret via comparison of areas, and which after some manipulation become geometrically evident. The precise arguments are given in § 5.

2. Background on the Fargues–Fontaine curve

2.1. The Fargues–Fontaine curve and vector bundles

We begin by recalling the definition of the Fargues–Fontaine curve $X := X_{E,F}$. This is already a slightly subtle issue, as there are several different incarnations of the curve: as a scheme (which was the original definition of Fargues and Fontaine), as an *adic space*, and as a *diamond*. In this paper, we will only need to use the classification of vector bundles on X as a black box, so we will not actually need any technical details of the construction of X itself. Therefore, we content ourselves with giving just a cursory introduction to the construction.

Definition 2.1.1. Let E be a finite extension of \mathbf{Q}_p , with uniformizer π , ring of integers E° , and residue field \mathbf{F}_q where $q = p^f$, and let F/\mathbf{F}_q be an algebraically closed perfectoid field, with ring of integers F° and pseudouniformizer ϖ .

Let $W_{E^\circ}(F^\circ) = W(F^\circ) \otimes_{W(\mathbf{F}_q)} E^\circ$ be the ramified Witt vectors of F° with coefficients in E° . Define

$$\mathcal{Y}_{E,F} = \text{Spa}(W_{E^\circ}(F^\circ)) \setminus \{ |p[\varpi]| = 0 \},$$

and let $\phi : \mathcal{Y}_{E,F} \rightarrow \mathcal{Y}_{E,F}$ be the Frobenius automorphism of $\mathcal{Y}_{E,F}$ induced by the natural q -Frobenius $\varphi_q = \varphi^f \otimes 1$ on $W_{E^\circ}(F^\circ)$. The (mixed-characteristic) *adic Fargues-Fontaine curve* $\mathcal{X}_{E,F}$ is

$$\mathcal{X}_{E,F} = \mathcal{Y}_{E,F} / \phi^{\mathbf{Z}}.$$

Proposition 2.1.2 (Kedlaya). *For any pair (E, F) as above, $\mathcal{X}_{E,F}$ is a Noetherian adic space over $\text{Spa } E$.*

Proof. This is one of the main results of [10]. □

Remark 2.1.3. There is also an equal-characteristic Fargues-Fontaine curve associated with a pair (E, F) as above, but where E is now taken to be a finite extension of $\mathbf{F}_p((t))$. In this paper, we will largely focus on the mixed-characteristic curve: our main results hold verbatim in the equicharacteristic setting, but the proofs are strictly easier.

By descent, a vector bundle on $\mathcal{X}_{E,F}$ is the same as a ϕ -equivariant vector bundle on $\mathcal{Y}_{E,F}$, that is, a vector bundle $\tilde{\mathcal{E}}$ on $\mathcal{Y}_{E,F}$ together with an isomorphism $\phi^* \tilde{\mathcal{E}} \xrightarrow{\sim} \tilde{\mathcal{E}}$.

Definition 2.1.4. If $\lambda = r/s$ is a rational number written in lowest terms with s positive, we define a vector bundle $\mathcal{O}(\lambda)$ on $\mathcal{X}_{E,F}$ as the descent of the trivial rank s vector bundle on $\mathcal{Y}_{E,F}$ equipped with the following ϕ -equivariant structure. Let v_1, \dots, v_s be a trivializing basis of $\tilde{\mathcal{E}} := \mathcal{O}_{\mathcal{Y}_{E,F}}^{\oplus s}$, and by abuse of notation view it as a trivializing basis for $\phi^* \tilde{\mathcal{E}}$ as well. Define $\phi^* \tilde{\mathcal{E}} \xrightarrow{\sim} \tilde{\mathcal{E}}$ by

$$\begin{aligned} v_1 &\mapsto v_2 \\ v_2 &\mapsto v_3 \\ &\dots \\ v_{s-1} &\mapsto v_s \\ v_s &\mapsto \pi^{-r} v_1. \end{aligned}$$

As previously discussed, there is also a *scheme-theoretic* Fargues-Fontaine curve $X_{E,F}$.

Definition 2.1.5. Let E and F be as in Definition 2.1.1. We define the *scheme-theoretic Fargues-Fontaine curve* $X_{E,F}$ to be

$$X_{E,F} := \text{Proj} \left(\bigoplus_{n \geq 0} H^0(\mathcal{X}_{E,F}, \mathcal{O}(n)) \right).$$

Remark 2.1.6. The original definition [4, 5] of the scheme-theoretic Fargues–Fontaine curve was given in terms of certain period rings of p -adic Hodge theory:

$$X_{E,F} = \text{“Proj} \left(\bigoplus_{n \geq 0} B^{\varphi_q = \pi^n} \right) \text{”}$$

[5, p. 22]; however, the vector spaces $B^{\varphi_q = \pi^n}$ coincide with $H^0(\mathcal{X}_{E,F}, \mathcal{O}(n))$, and so the definitions agree.

Proposition 2.1.7. *The scheme $X_{E,F}$ is Noetherian, connected, and regular of (absolute) dimension one.*

It is known (‘GAGA for the Fargues–Fontaine curve’, [11, Theorem 6.3.12]) that there is a natural map

$$\mathcal{X}_{E,F} \rightarrow X_{E,F},$$

which induces by pullback an equivalence of categories of vector bundles. Therefore, we can and do speak interchangeably about vector bundles on $\mathcal{X}_{E,F}$ and $X_{E,F}$. For ease of notation, we henceforth set $X = X_{E,F}$.

2.2. Harder–Narasimhan filtrations and polygons

The following fundamental result of Fargues and Fontaine is the key to developing a slope theory for bundles on X .

Proposition 2.2.1 [4]. *The curve X is complete in the sense that if $f \in k(X)$ is any nonzero rational function on X , then the divisor of f has degree zero.*

This implies, in particular, that if \mathcal{L} is a line bundle on X and s is any nonzero meromorphic section of \mathcal{L} , then $\text{deg } \mathcal{L} \stackrel{\text{def}}{=} \text{deg div}(s)$ is well defined independently of the choice of s .

Definition 2.2.2. If \mathcal{E} is a vector bundle on X , we define the *degree* of \mathcal{E} by $\text{deg } \mathcal{E} := \text{deg } \wedge^{\text{rank } \mathcal{E}} \mathcal{E}$ and the *slope* of \mathcal{E} by $\mu(\mathcal{E}) := \frac{\text{deg } \mathcal{E}}{\text{rank } \mathcal{E}}$.

Example 2.2.3. As a vector bundle on X , $\mathcal{O}(r/s)$ has rank s and degree r , and hence slope r/s . One may check by hand from Definition 2.1.4 that

$$\mathcal{O}\left(\frac{r}{s}\right) \otimes \mathcal{O}\left(\frac{r'}{s'}\right) \simeq \mathcal{O}\left(\frac{r}{s} + \frac{r'}{s'}\right)^{\oplus \text{gcd}(ss', rs' + r's)}.$$

In particular, $\mathcal{O}(r/s) \otimes \mathcal{O}(r'/s')$ has rank ss' , degree $rs' + r's$, and slope $r/s + r'/s'$.

We recall the usual notions of (semi)stability.

Definition 2.2.4. We say that a vector bundle \mathcal{E} on X is *stable* if it has no proper, nonzero subbundles $\mathcal{F} \subset \mathcal{E}$ with $\mu(\mathcal{F}) \geq \mu(\mathcal{E})$. We say that \mathcal{E} is *semistable* if it has no proper, nonzero subbundles $\mathcal{F} \subset \mathcal{E}$ with $\mu(\mathcal{F}) > \mu(\mathcal{E})$.

Stable and semistable vector bundles on X turn out to admit a complete classification.

Proposition 2.2.5 [4, 8]. *For any $\lambda \in \mathbf{Q}$, the vector bundle $\mathcal{O}(\lambda)$ is stable, and any finite direct sum $\mathcal{O}(\lambda)^{\oplus n}$ is semistable. Moreover, every semistable vector bundle of slope λ is isomorphic to some finite direct sum $\mathcal{O}(\lambda)^{\oplus n}$.*

Combining this classification with Example 2.2.3, one immediately deduces that tensor products of semistable bundles are semistable.

Definition 2.2.6. A *Harder–Narasimhan (HN) filtration* of a vector bundle \mathcal{E} is a filtration of \mathcal{E} by subbundles

$$0 = \mathcal{E}_0 \subset \mathcal{E}_1 \subset \dots \subset \mathcal{E}_m = \mathcal{E}$$

such that each successive quotient $\mathcal{E}_i/\mathcal{E}_{i-1}$ is a semistable vector bundle that is of slope μ_i , and such that

$$\mu_1 > \mu_2 > \dots > \mu_m.$$

Keeping this notation, the *Harder–Narasimhan (HN) polygon* of \mathcal{E} is the upper convex hull of the points $(\text{rank } \mathcal{E}_i, \text{deg } \mathcal{E}_i)$.

Remark 2.2.7. Our convention for HN polygons is opposed to the usual convention for Newton polygons, which are usually lower convex hulls.

Theorem 2.2.8 (Fargues–Fontaine). *Every vector bundle on X admits a canonical and functorial HN filtration.*

Proof. This follows from [4, Théorème 6.5.2] together with the discussion in [4, § 5.5]. \square

Example 2.2.9. In Figure 2, we depict an example of the HN polygon associated with a vector bundle whose HN filtration has the form

$$0 \subset \mathcal{E}_1 \subset \mathcal{E}_2 \subset \mathcal{E}_3 \subset \mathcal{E}_4 = \mathcal{E}$$

with $\mu_i := \mu(\mathcal{E}_i/\mathcal{E}_{i-1})$.

Combining the classification of semistable bundles with an explicit calculation of the cohomology groups $H^i(X, \mathcal{O}(\lambda))$, one deduces that the HN filtration of any vector bundle on X is split. The precise statement is as follows.

Theorem 2.2.10 (Kedlaya, Fargues–Fontaine [4, 9]). (1) *For any given rational number λ , we have*

$$H^0(\mathcal{O}(\lambda)) = 0 \text{ if and only if } \lambda < 0$$

and

$$H^1(\mathcal{O}(\lambda)) = 0 \text{ if and only if } \lambda \geq 0.$$




Figure 2. The HN polygon of a vector bundle with a 4-step HN filtration.

(2) Any vector bundle \mathcal{E} on X admits a direct sum decomposition

$$\mathcal{E} \simeq \bigoplus_i \mathcal{O}(\lambda_i),$$

where the λ_i 's run over the HN slopes of \mathcal{E} counted with the appropriate multiplicities.

Corollary 2.2.11. A vector bundle \mathcal{E} on X is determined up to isomorphism by its HN polygon $\text{HN}(\mathcal{E})$.

Definition 2.2.12. We define a partial order on HN polygons as follows: if HN_1 and HN_2 are two polygons, then we say that $\text{HN}_1 \leq \text{HN}_2$ if HN_1 lies on or below HN_2 and the polygons have the same endpoints.

Under this partial ordering, the HN filtration of a bundle \mathcal{E} satisfies the following extremality property among all filtrations with semistable graded pieces.

Corollary 2.2.13. Let \mathcal{E} be a vector bundle on X , and let

$$0 = \mathcal{E}_0 \subset \mathcal{E}_1 \subset \dots \subset \mathcal{E}_m = \mathcal{E}$$

be any filtration of \mathcal{E} such that each graded piece $\mathcal{E}_i/\mathcal{E}_{i-1}$ is a semistable vector bundle. Then

$$\text{HN}(\mathcal{E}) \leq \text{HN}(\mathcal{E}_1/\mathcal{E}_0 \oplus \dots \oplus \mathcal{E}_m/\mathcal{E}_{m-1}).$$




Figure 3. This depicts the order relation between HN polygons. If HN_1 is the polygon shaded in blue and HN_2 is the polygon shaded in green, then $HN_1 \leq HN_2$.

Proof. This is [8, Corollary 3.4.18]. □

Definition 2.2.14. Let \mathcal{E} be a vector bundle on X with HN filtration

$$0 = \mathcal{E}_0 \subset \mathcal{E}_1 \subset \dots \subset \mathcal{E}_m = \mathcal{E}.$$

Then for any $\lambda \in \mathbf{Q}$, we define $\mathcal{E}^{\geq \lambda}$ (resp. $\mathcal{E}^{> \lambda}$) to be the subbundle of \mathcal{E} given by \mathcal{E}_i for the largest value of i such that $\mu(\mathcal{E}_i/\mathcal{E}_{i-1}) \geq \lambda$ (resp. such that $\mu(\mathcal{E}_i/\mathcal{E}_{i-1}) > \lambda$). We also define $\mathcal{E}^{< \lambda} = \mathcal{E}/\mathcal{E}^{\geq \lambda}$ and $\mathcal{E}^{\leq \lambda} = \mathcal{E}/\mathcal{E}^{> \lambda}$.

2.3. Geometric interpretation of degrees

In a number of our later arguments, we will need to understand quantities of the form $\text{deg}(\mathcal{E}_1 \otimes \mathcal{E}_2^{\vee})^{\geq 0}$ for some fairly arbitrary bundles \mathcal{E}_1 and \mathcal{E}_2 . Here we develop some preliminary language for doing this. The results in this section will not be necessary until § 4.

Definition 2.3.1. For two vectors $v, w \in \mathbb{R}^2$ with nonzero x -coordinates, we write $v < w$ (resp. $v \leq w$) if the slope of v is less than (resp. not greater than) the slope of w .

We can determine this relation between v and w using their (two-dimensional) cross product as follows.

Lemma 2.3.2. We have the following characterization of the order $<$:

- if the x -coordinates of v and w have the same sign, $v < w$ (resp. $v \leq w$) if and only if $v \times w > 0$ (resp. $v \times w \geq 0$);
- if the x -coordinates of v and w have different signs, $v < w$ (resp. $v \leq w$) if and only if $v \times w < 0$ (resp. $v \times w \leq 0$).

Proof. This is straightforward. □

Let \mathcal{V} be a vector bundle on X . By Theorem 2.2.10, \mathcal{V} decomposes into a direct sum

$$\mathcal{V} = \bigoplus_{i=1}^s \mathcal{O}(d_i/h_i)^{m_i},$$

where $d_1/h_1 > d_2/h_2 > \dots > d_s/h_s$ and d_i, h_i are relatively prime for each $i = 1, 2, \dots, s$. Then the HN polygon of \mathcal{V} is built out of the sequence of vectors

$$v_1 \succ v_2 \succ \dots \succ v_s$$

, where $v_i = (m_i h_i, m_i d_i)$ is the i th edge in $\text{HN}(\mathcal{V})$.

Definition 2.3.3. Keeping the notation of the preceding paragraph, the *HN vectors* of \mathcal{V} will be denoted by $\overrightarrow{\text{HN}}(\mathcal{V}) := (v_i)_{1 \leq i \leq s}$. Note that the isomorphism class of \mathcal{V} is uniquely determined by $\overrightarrow{\text{HN}}(\mathcal{V})$.

Lemma 2.3.4. Let \mathcal{V} and \mathcal{W} be any vector bundles on X with $\overrightarrow{\text{HN}}(\mathcal{V}) = (v_i)$ and $\overrightarrow{\text{HN}}(\mathcal{W}) = (w_j)$. Then we have

$$\text{deg}(\mathcal{V}^\vee \otimes \mathcal{W}) = \sum_{i,j} v_i \times w_j \tag{2.1}$$

and

$$\text{deg}(\mathcal{V}^\vee \otimes \mathcal{W})^{\geq 0} = \sum_{v_i \leq w_j} v_i \times w_j. \tag{2.2}$$

Proof. Let us first consider the case where \mathcal{V} and \mathcal{W} are both semistable. In this case, both $\overrightarrow{\text{HN}}(\mathcal{V})$ and $\overrightarrow{\text{HN}}(\mathcal{W})$ consist of a single element, namely $(\text{rank}(\mathcal{V}), \text{deg}(\mathcal{V}))$ and $(\text{rank}(\mathcal{W}), \text{deg}(\mathcal{W}))$, respectively. Then (2.1) follows directly from Example 2.2.3, and (2.2) follows from (2.1) by Part 1 of Lemma 2.3.2.

The general case now follows by observing that both sides of (2.1) and (2.2) are linear in both \mathcal{V} and \mathcal{W} , and using Theorem 1.1.1 to see that every vector bundle decomposes as a direct sum of semistable bundles. □

Proposition 2.3.5. Let \mathcal{V} be a vector bundle on X . Then $\text{deg}(\mathcal{V}^\vee \otimes \mathcal{V})^{\geq 0}$ is equal to twice the area of the region enclosed between $\text{HN}(\mathcal{V})$ and the line segment joining the two endpoints of $\overrightarrow{\text{HN}}(\mathcal{V})$. In particular, $\text{deg}(\mathcal{V}^\vee \otimes \mathcal{V})^{\geq 0} = 0$ if and only if \mathcal{V} is semistable.

Proof. Write $\overrightarrow{\text{HN}}(\mathcal{V}) = (v_i)_{1 \leq i \leq s}$, and let $O = P_0, P_1, P_2, \dots, P_s$ be the breakpoints of $\text{HN}(\mathcal{V})$, listed in order of increasing x -coordinates (see Figure 4). Note that $v_i = \overrightarrow{P_{i-1}P_i}$.

By definition, we have $v_1 \succ v_2 \succ \dots \succ v_s$. Hence Lemma 2.3.4 yields

$$\begin{aligned} \text{deg}(\mathcal{V}^\vee \otimes \mathcal{V})^{\geq 0} &= \sum_{i \geq j} v_i \times v_j = \sum_{j=1}^s (v_j + v_{j+1} + \dots + v_s) \times v_j \\ &= \sum_{j=1}^s (\overrightarrow{P_{j-1}P_s} \times \overrightarrow{P_{j-1}P_j}) \\ &= \sum_{j=1}^s 2 \cdot \text{Area}(\triangle P_{j-1}P_jP_s), \end{aligned}$$




Figure 4. Visual depiction of Proposition 2.3.5.

which is clearly equal to twice the area of the region enclosed between $\text{HN}(\mathcal{V})$ and the line segment joining O and P_s . □

Corollary 2.3.6. *Let \mathcal{V} and \mathcal{W} be any vector bundles on X such that $\text{HN}(\mathcal{V}) \leq \text{HN}(\mathcal{W})$. Then $\deg(\mathcal{V}^\vee \otimes \mathcal{V})^{\geq 0} \leq \deg(\mathcal{W}^\vee \otimes \mathcal{W})^{\geq 0}$.*

Proof. This is an immediate consequence of Proposition 2.3.5. □

3. Diamonds, vector bundles, and dimensions

3.1. Recollections on diamonds

In this section, we (very) briefly introduce the language of diamonds. Everything here can be found in [13, §§ 7–8, §§ 10–11, § 18, § 21].

Definition 3.1.1. (1) A map $Y \rightarrow X$ of affinoid perfectoid spaces is *affinoid pro-étale* if Y is isomorphic to the limit of some cofiltered system of affinoid perfectoid spaces Y_i étale over X .

(2) A map $f : Y \rightarrow X$ of perfectoid spaces is *pro-étale* if X and Y can be covered by open affinoid perfectoid subsets $U_i \subset X, V_{ij} \subset Y$ such that $f|_{V_{ij}} : V_{ij} \rightarrow X$ factors through the inclusion of U_i and such that the induced map $V_{ij} \rightarrow U_i$ is affinoid pro-étale.

(3) A map $f : Y \rightarrow X$ of perfectoid spaces is a *pro-étale cover* if f is pro-étale and for any qc open subset $U \subset X$, there is some qc open subset $V \subset Y$ with $f(V) = U$.

(4) The *big pro-étale site* Perf is the site of perfectoid spaces in characteristic p , with covers given by pro-étale covers.

Proposition 3.1.2. *The site Perf is subcanonical.*

We slightly abuse notation and write $X = h_X$ for any $X \in \text{Perf}$, i.e., we freely identify a characteristic p perfectoid space X with its Yoneda embedding into sheaves on Perf . By the previous proposition, there is little harm in this.

Definition 3.1.3. A *diamond* is a sheaf Y on Perf such that

$$Y \simeq \text{Coeq}(R \rightrightarrows X),$$

where $R \rightrightarrows X$ is a pro-étale equivalence relation in characteristic p perfectoid spaces.

If Y is a diamond with a presentation $Y \simeq \text{Coeq}(R \rightrightarrows X)$, then $|R| \rightrightarrows |X|$ defines an equivalence relation on $|X|$, and we define $|Y| = |X|/|R|$ as the associated quotient, topologized via the quotient topology. The topological space $|Y|$ is well defined independently of the chosen presentation, and is functorial in Y .

Definition 3.1.4. A subfunctor $Y' \subset Y$ of a diamond Y is an *open subdiamond* if for any perfectoid space T with a map $T \rightarrow Y$, $Y' \times_Y T$ is representable and $Y' \times_Y T \rightarrow T$ is an open immersion.

One then checks that if Y is any diamond, the association $Y' \rightsquigarrow |Y'|$ defines an inclusion-preserving bijection between open subdiamonds of Y and open subsets of $|Y|$.

We also recall the relation of diamonds with adic spaces.

Proposition 3.1.5. *There is a natural functor $X \mapsto X^\diamond$ from analytic adic spaces over $\text{Spa}(\mathbf{Z}_p, \mathbf{Z}_p)$ to diamonds, extending the functor $X \mapsto h_{X^\flat}$ on perfectoid spaces and inducing a functorial homeomorphism $|X^\diamond| \cong |X|$.*

If $\text{Spa}(A, A^+)$ is an affinoid adic space, we will often write $\text{Spd}(A, A^+) := \text{Spa}(A, A^+)^\diamond$ and $\text{Spd } A := \text{Spd}(A, A^\circ)$.

Next we recall a flexible variant of pro-étale morphisms in the setting of diamonds.

Definition 3.1.6. (1) A perfectoid space X is *strictly totally disconnected* if it is qcqs and every connected component of X is of the form $\text{Spa}(K, K^+)$ for some algebraically closed perfectoid field K .

(2) A map of diamonds $Y \rightarrow X$ is *quasi-pro-étale* if for any strictly totally disconnected perfectoid space T with a map $T \rightarrow X$, the sheaf $Y \times_X T$ is representable and the induced map of perfectoid spaces $Y \times_X T \rightarrow T$ is pro-étale.

We now recall some useful classes of diamonds and morphisms of diamonds.

Definition 3.1.7. (1) A diamond Y is *quasicompact (qc)* if it admits a presentation $Y \simeq \text{Coeq}(R \rightrightarrows X)$ with X quasicompact. A morphism of diamonds $Y \rightarrow X$ is *quasicompact* if $Y \times_X U$ is quasicompact for all quasicompact diamonds U with a map $U \rightarrow X$.

- (2) A diamond Y is *quasiseparated* (*qs*) if for any qc diamonds $U \rightarrow Y, V \rightarrow Y$, the fiber product $U \times_Y V$ is qc. A morphism of diamonds $Y \rightarrow X$ is *quasiseparated* if the relative diagonal $Y \rightarrow Y \times_X Y$ is quasicompact.
- (3) A diamond Y is *separated* (resp. *partially proper*) if it is quasiseparated and for all choices of a characteristic p perfectoid Tate ring R with a ring of integral elements R^+ , the restriction map

$$Y(R, R^+) \rightarrow Y(R, R^\circ)$$

is injective (resp. bijective). A morphism $Y \rightarrow X$ of diamonds is *separated* (resp. *partially proper*) if it is quasiseparated and for all choices of a characteristic p perfectoid Tate ring R with a ring of integral elements R^+ , the natural map

$$Y(R, R^+) \rightarrow Y(R, R^\circ) \times_{X(R, R^\circ)} X(R, R^+)$$

is injective (resp. bijective).

- (4) A diamond Y is *spatial* if it is quasicompact and quasiseparated, and if the subsets $|U| \subset |Y|$, for $U \subset Y$ running over arbitrary qc open subdiamonds of Y , give a neighborhood basis of $|Y|$.
- (5) A diamond is *locally spatial* if it admits a covering by spatial open subdiamonds.

We will mostly restrict our attention to locally spatial diamonds, which have several favorable properties.

Proposition 3.1.8. (1) *If Y is (locally) spatial, then $|Y|$ is a (locally) spectral topological space.*

- (2) *A locally spatial diamond Y is qc (resp. qs) if and only if $|Y|$ is qc (resp. qs).*
- (3) *If $U \rightarrow X \leftarrow V$ is a diagram of (locally) spatial diamonds, then $U \times_X V$ is (locally) spatial.*
- (4) *If $Y \rightarrow X$ is any map of locally spatial diamonds, the associated map of locally spectral spaces $|Y| \rightarrow |X|$ is spectral and generalizing.² Moreover, $|Y| \rightarrow |X|$ is a quotient map if $Y \rightarrow X$ is surjective.*
- (5) *If X is an analytic adic space, then X^\diamond is locally spatial, and X^\diamond is spatial exactly when X is qcqs.*

Finally, in some of the proofs in § 3.2, we will make use of Scholze’s results on ‘canonical compactifications’ of diamonds. Let us briefly explain what we need from this theory; the reader may wish to skip this discussion.

Fix a diamond S . In [13, § 18], Scholze defines a compactification functor $Y \mapsto \bar{Y}^S$ from separated S -diamonds to separated S -diamonds satisfying a number of useful compatibilities:

- (1) The map $\bar{Y}^S \rightarrow S$ is partially proper, and the diamond \bar{Y}^S comes with a functorial injection $Y \hookrightarrow \bar{Y}^S$, which is an isomorphism if Y is partially proper.

²Recall that a morphism of topological spaces $f : Y \rightarrow X$ is generalizing if for any point $y \in Y$ and any generalization x' of $f(y)$, there is some generalization y' of y such that $f(y') = x'$; cf. [7, Definition 1.3.3].

- (2) The functor $Y \mapsto \overline{Y}^S$ commutes with all limits.
- (3) If $Y \rightarrow Z$ is injective (resp. surjective, resp. quasi-pro-étale), then $\overline{Y}^S \rightarrow \overline{Z}^S$ is injective (resp. surjective, resp. quasi-pro-étale).
- (4) If $S = \mathrm{Spd}(A, A^+)$ and $Y = \mathrm{Spd}(B, B^+)$ are affinoid perfectoid, then $\overline{Y}^S = \mathrm{Spd}(B, B')$, where B' denotes the completed integral closure of $A^+ + B^\circ$ in B .
- (5) If W is a profinite set, then $\overline{Y \times W}^S \cong \overline{Y}^S \times W$.

We also need the following compatibility of compactifications with torsors.

Proposition 3.1.9. *Let G be a profinite group, and let $X \rightarrow Y$ be a \underline{G} -torsor in separated S -diamonds. Then $\overline{X}^S \rightarrow \overline{Y}^S$ is a \underline{G} -torsor.*

Proof. By property (3) above, the map $\overline{X}^S \rightarrow \overline{Y}^S$ is surjective and quasi-pro-étale. Functoriality of the compactification implies that the \underline{G} -action on X extends canonically to a \underline{G} -action on \overline{X}^S . It now suffices to check that

$$\underline{G} \times \overline{X}^S \rightarrow \overline{X}^S \times_{\overline{Y}^S} \overline{X}^S$$

is an isomorphism. By property (2), the target of this map is canonically isomorphic to

$$\overline{X \times_Y X}^S \cong \overline{\underline{G} \times \overline{X}^S},$$

which in turn is canonically isomorphic to $\underline{G} \times \overline{X}^S$ by property (5). This implies the desired result. □

3.2. Dimension theory

In this section, we work out some basic dimension theory for diamonds. Fix a complete nonarchimedean field K of residue characteristic p ; until further notice, all diamonds are assumed to live over $\mathrm{Spd}(K, K^\circ)$.

Definition 3.2.1. For X a locally spatial diamond, we say $x \in |X|$ is a *rank one point* if any one of the following equivalent conditions is satisfied:

- (1) x has no proper generalizations inside $|X|$.
- (2) There exist a perfectoid field L and a quasi-pro-étale map $\mathrm{Spd}(L, L^\circ) \rightarrow X$ with topological image x .
- (3) We have

$$x = \bigcap_{\substack{x \in U \subset |X| \\ U \text{ qc open}}} U$$

as subsets of $|X|$.

(The equivalence of these conditions follows easily from Proposition 3.1.8.) Note that any rank one point x has the structure of a diamond in its own right: more precisely, there is a quasicompact injection from a spatial diamond $x^\diamond \rightarrow X$ with $|x^\diamond| = x$, which is unique

up to unique isomorphism. We will sometimes confuse x and x^\diamond in what follows. We also note that for any map $Y \rightarrow X$ of locally spatial diamonds, the induced map $|Y| \rightarrow |X|$ carries rank one points onto rank one points.

Lemma 3.2.2. *Let X be a partially proper and locally spatial diamond over $S = \mathrm{Spd}(K, K^\circ)$. Choose a rank one point $x \in |X|$, and choose a perfectoid field L together with a quasi-pro-étale map $\mathrm{Spd}(L, L^\circ) \rightarrow X$ with topological image x . Then $\dim \{x\} = \mathrm{tr.deg}(l/k)$, where l and k denote the residue fields of L° and K° , respectively.*

Proof. Replacing L with \widehat{L} does not change the transcendence degree of l over k , so without loss of generality, we may assume that L is algebraically closed. By the proof of [13, Proposition 21.9], this gives rise to a presentation

$$x^\diamond \simeq \mathrm{Spd}(L, L^\circ)/\underline{G},$$

where G is a profinite group acting continuously and faithfully on L by K -linear automorphisms. Let L^{\min} be the completed integral closure of $K^\circ + L^\circ$ in L . Note that $\mathrm{Spd}(L, L^{\min})$ is simply the canonical compactification $\overline{\mathrm{Spd}(L, L^\circ)^{/S}}$. Note also that $|\mathrm{Spd}(L, L^{\min})|$ is homeomorphic to the Zariski–Riemann space of l/k' , where k' denotes the algebraic closure of k in l ; in particular, $|\mathrm{Spd}(L, L^{\min})|$ is spectral of dimension $\mathrm{tr.deg}(l/k') = \mathrm{tr.deg}(l/k)$.

Since $x^\diamond \rightarrow X$ is an injection of separated S -diamonds, with X partially proper and x^\diamond quasicompact, this extends to an injection $\overline{x^\diamond}^{/S} \rightarrow X$, and the injective map $|\overline{x^\diamond}^{/S}| \rightarrow |X|$ is a homeomorphism onto its (closed) image. We now claim that $\overline{\{x\}}$ and $|\overline{x^\diamond}^{/S}|$ are homeomorphic as subsets of $|X|$: this follows from observing that both subsets are closed and generalizing subsets of $|X|$, each containing x as their unique rank one point. Since $\mathrm{Spd}(L, L^\circ) \rightarrow x^\diamond$ is a \underline{G} -torsor, the map $\mathrm{Spd}(L, L^{\min}) \rightarrow \overline{x^\diamond}^{/S}$ is a \underline{G} -torsor by Proposition 3.1.9. Therefore we have natural homeomorphisms

$$\overline{\{x\}} \cong |\overline{x^\diamond}^{/S}| \cong |\mathrm{Spd}(L, L^{\min})|/G.$$

To complete the proof, it suffices to show the equality

$$\dim |\mathrm{Spd}(L, L^{\min})| = \dim |\mathrm{Spd}(L, L^{\min})|/G.$$

This is a special case of the next lemma. □

Lemma 3.2.3. *Let X be a spectral space, and let G be a profinite group acting continuously on X . Then X/G is spectral, the map $q : X \rightarrow X/G$ is spectral and generalizing, and $\dim X = \dim X/G$.*

Proof. Without loss of generality, G acts faithfully on X . One easily checks that the image R of the natural map $G \times X \xrightarrow{(xg, x)} X \times X$ is pro-constructible and defines an equivalence relation on X , such that the maps $s, t : R \rightarrow X$ are quasicompact, generalizing, and open. Spectrality of $X/R = X/G$ then follows from [13, Lemmas 2.9–2.10]; moreover, these same lemmas imply that $X \rightarrow X/G$ is spectral and generalizing.

For the equality of dimensions, note that since q is generalizing, we may lift any finite chain in X/G to a chain of the same length in X , and thus $\dim X \geq \dim X/G$. It remains to prove the opposite inequality. Since q sends chains in X to chains in X/G , it is enough to show that for any finite chain $C \subset X$, the length of C coincides with the length of $q(C)$. Suppose otherwise; then we can find two distinct points $x < y$ in X with $q(x) = q(y)$, or equivalently with $xG = yG$. Let³ $\lambda : |X| \rightarrow \mathbf{N} \cup \{\infty\}$ be the function sending a point $z \in X$ to the maximal length of any chain of generalizations of z . Clearly, $\lambda(x) \geq \lambda(y) + 1$. On the other hand, the G -action on X preserves the relation of one point generalizing another, so λ is constant on G -orbits. This is a contradiction. \square

Lemma 3.2.4. *Let $f : X \rightarrow Y$ be a map of partially proper and locally spatial diamonds over $\mathrm{Spd}(K, K^\circ)$. Choose any rank one point $x \in |X|$ with image $y = f(x)$, and let $X_y = X \times_Y y^\diamond$ denote the fiber of X over y (so, in particular, $x \in |X_y|$). Then*

$$\dim \overline{\{x\}}^X = \dim \overline{\{y\}}^Y + \dim \overline{\{x\}}^{X_y},$$

where $\overline{\{s\}}^S$ denotes the closure of $\{s\}$ in the topological space of S .

When X and Y are analytic adic spaces, this is exactly [1, Proposition 4.2.21]. The proof given below is essentially an adaptation of the argument in loc. cit., although we need to incorporate a rather elaborate pro-étale unscrewing.

Proof. Choose surjective quasi-pro-étale maps $\mathrm{Spd}(L, L^\circ) \rightarrow X$, $\mathrm{Spd}(M, M^\circ) \rightarrow Y$ with topological images x and y , respectively; without loss of generality, we may assume that L and M are algebraically closed, with residue fields l and m , respectively. Applying Lemma 3.2.2 twice, we get equalities $\dim \overline{\{x\}}^X = \mathrm{tr.deg}(l/k)$ and $\dim \overline{\{y\}}^Y = \mathrm{tr.deg}(m/k)$. By the additivity of transcendence degrees, it suffices to prove that there is *some* continuous K -algebra map $M \rightarrow L$ such that the induced extension $m \rightarrow l$ of residue fields satisfies $\mathrm{tr.deg}(l/m) = \dim \overline{\{x\}}^{X_y}$.

Let G and H be the profinite groups such that $\mathrm{Spd}(L, L^\circ) \rightarrow x^\diamond$ (resp. $\mathrm{Spd}(M, M^\circ) \rightarrow y^\diamond$) is a G -torsor (resp. an H -torsor) as before. Arguing as in the proof of Lemma 3.2.2, we have a canonical homeomorphism

$$\overline{\{x\}}^{X_y} \cong |\overline{x^\diamond/y^\diamond}|.$$

Base change along the map $\mathrm{Spd}(M, M^\circ) \rightarrow y^\diamond$ identifies $\overline{x^\diamond \times_{y^\diamond} \mathrm{Spd}(M, M^\circ)}/\mathrm{Spd}(M, M^\circ)$ with $\overline{x^\diamond/y^\diamond} \times_{y^\diamond} \mathrm{Spd}(M, M^\circ)$; in particular,

$$\overline{x^\diamond \times_{y^\diamond} \mathrm{Spd}(M, M^\circ)}/\mathrm{Spd}(M, M^\circ)$$

is naturally an H -torsor over $\overline{x^\diamond/y^\diamond}$, so

³Thanks to Christian Johansson for suggesting the following argument.

$$\begin{aligned} \dim \overline{\{x\}}^{X_y} &= \dim |\overline{x}^{\diamond}/y^{\diamond}| \\ &= \dim |\overline{x^{\diamond} \times_{y^{\diamond}} \mathrm{Spd}(M, M^{\circ})}^{\mathrm{Spd}(M, M^{\circ})}|, \end{aligned}$$

where the second equality follows from Lemma 3.2.3. Furthermore, the \underline{G} -torsor $\mathrm{Spd}(L, L^{\circ}) \rightarrow x^{\diamond}$ induces a \underline{G} -torsor structure on the map

$$\overline{\mathrm{Spd}(L, L^{\circ}) \times_{y^{\diamond}} \mathrm{Spd}(M, M^{\circ})}^{\mathrm{Spd}(M, M^{\circ})} \rightarrow \overline{x^{\diamond} \times_{y^{\diamond}} \mathrm{Spd}(M, M^{\circ})}^{\mathrm{Spd}(M, M^{\circ})},$$

so we are reduced to calculating the topological dimension of the source of this map. Since $\mathrm{Spd}(L, L^{\circ}) \times_{y^{\diamond}} \mathrm{Spd}(M, M^{\circ}) \rightarrow \mathrm{Spd}(L, L^{\circ})$ is an \underline{H} -torsor (and, in particular, a profinite étale map) with target a geometric point, there is an isomorphism

$$\mathrm{Spd}(L, L^{\circ}) \times_{y^{\diamond}} \mathrm{Spd}(M, M^{\circ}) \simeq \mathrm{Spd}(L, L^{\circ}) \times \underline{W}$$

for some profinite set W . In particular, any connected component V of

$$\mathrm{Spd}(L, L^{\circ}) \times_{y^{\diamond}} \mathrm{Spd}(M, M^{\circ})$$

is isomorphic to $\mathrm{Spd}(L, L^{\circ})$; choosing such a component induces a continuous map $M \rightarrow L$, and $\dim |\overline{V}^{\mathrm{Spd}(M, M^{\circ})}| = \mathrm{tr.deg}(l/m)$ by our previous discussions. We now conclude by observing that the compactification functor identifies any given connected component of $|\overline{\mathrm{Spd}(L, L^{\circ}) \times_{y^{\diamond}} \mathrm{Spd}(M, M^{\circ})}^{\mathrm{Spd}(M, M^{\circ})}|$ homeomorphically with $|\overline{V}^{\mathrm{Spd}(M, M^{\circ})}|$ for some uniquely determined connected component $V \subset \mathrm{Spd}(L, L^{\circ}) \times_{y^{\diamond}} \mathrm{Spd}(M, M^{\circ})$, so the result now follows by varying V . □

The next result conceptualizes the proof of [1, Proposition 4.2.23].

Lemma 3.2.5. *Let $f : X \rightarrow Y$ be a map of partially proper and locally spatial diamonds over $\mathrm{Spd}(K, K^{\circ})$. Suppose there is some integer $d \geq 0$ such that for every rank one point $y \in \mathrm{im}(|X| \rightarrow |Y|)$, the fiber $X_y = X \times_Y y^{\diamond}$ is of dimension d . Then*

$$\dim \mathrm{im}(|X| \rightarrow |Y|) = \dim |X| - d.$$

Remark 3.2.6. Let $f : X \rightarrow Y$ be as in the lemma. If f is quasicompact, then the subspace and quotient topologies on $\mathrm{im}(|X| \rightarrow |Y|)$ coincide, and $\mathrm{im}(|X| \rightarrow |Y|)$ is locally spectral and pro-constructible inside $|Y|$. For a general f , all of these outcomes can fail. It is true, however, that $\mathrm{im}(|X| \rightarrow |Y|)$ is always closed under generalization and specialization inside $|Y|$. In any case, we define $\dim \mathrm{im}(|X| \rightarrow |Y|)$ as the maximal length of any chain of specializations of any rank one point of $|Y|$ contained in the image of $|X|$.

Proof. Choose a rank one point $y \in \mathrm{im}(|X| \rightarrow |Y|)$, and choose some rank one point $x \in X_y$ lifting y such that $\dim \overline{\{x\}}^{X_y} = d$. Then

$$\dim \overline{\{y\}}^Y = \dim \overline{\{x\}}^X - \dim \overline{\{x\}}^{X_y} = \dim \overline{\{x\}}^X - d \leq \dim X - d$$

by Lemma 3.2.4. Taking the supremum over all y , we deduce the inequality

$$\dim \mathrm{im}(|X| \rightarrow |Y|) \leq \dim |X| - d.$$

For the opposite inequality, it clearly suffices to prove that

$$\dim X \leq \dim \text{im}(|X| \rightarrow |Y|) + d.$$

Let $|X|^{\text{rk} 1}$ denote the set of rank one points in $|X|$. Then

$$\begin{aligned} \dim X &= \sup_{x \in |X|^{\text{rk} 1}} \dim \overline{\{x\}}^X \\ &= \sup_{x \in |X|^{\text{rk} 1}} (\dim \overline{\{f(x)\}}^Y + \dim \overline{\{x\}}^{X_{f(x)}}) \\ &\leq \sup_{x \in |X|^{\text{rk} 1}} \dim \overline{\{f(x)\}}^Y + \sup_{x \in |X|^{\text{rk} 1}} \dim \overline{\{x\}}^{X_{f(x)}} \\ &= \dim \text{im}(|X| \rightarrow |Y|) + d. \end{aligned}$$

Here the first line is simply the definition of $\dim X$, the second line follows from Lemma 3.2.4, and the final line follows from the relevant definitions together with our assumption on the fibers of f . □

3.3. Diamonds and moduli of bundle maps

In this section, we define and study diamonds parametrizing maps between vector bundles on \mathcal{X} with various specified properties. We note in passing that every diamond considered here is a diamond in the sense of the original definition proposed by Scholze in 2014, i.e., they each admit a surjective representable pro-étale morphism from a perfectoid space.

Throughout this section, we fix a finite extension E/\mathbf{Q}_p with residue field \mathbf{F}_q and an algebraically closed perfectoid field F/\mathbf{F}_q , and we let $\mathcal{X} = \mathcal{X}_{E,F}$ denote the adic Fargues–Fontaine curve over $\text{Spa } F$. For any perfectoid space $S/\text{Spa } F$, we get a relative curve $\mathcal{X}_S = \mathcal{X}_{E,S}$ together with a natural map $\mathcal{X}_S \rightarrow \mathcal{X}$ (cf. [11, Chapters 7–8] for a thorough discussion of relative Fargues–Fontaine curves). If \mathcal{E} is any vector bundle on \mathcal{X} , we denote the vector bundle obtained via pullback along $\mathcal{X}_S \rightarrow \mathcal{X}$ by \mathcal{E}_S .

Definition 3.3.1. Fix vector bundles \mathcal{E}, \mathcal{F} on \mathcal{X} . Consider the following sheaves of sets on $\text{Perf}/\text{Spa } F$.

- (1) Let $\mathcal{H}^0(\mathcal{E}) : \text{Perf}/\text{Spa } F \rightarrow \text{Sets}$ be the functor sending $f : S \rightarrow \text{Spa } F$ to the set $H^0(\mathcal{X}_S, \mathcal{E}_S)$.
- (2) Let $\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})$ be the functor sending $f : S \rightarrow \text{Spa } F$ to the set of $\mathcal{O}_{\mathcal{X}_S}$ -module maps $m : \mathcal{E}_S \rightarrow \mathcal{F}_S$. Note that $\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F}) \cong \mathcal{H}^0(\mathcal{E}^\vee \otimes \mathcal{F})$.
- (3) Let $\text{Surj}(\mathcal{E}, \mathcal{F}) \subset \mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})$ be the subfunctor of $\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})$ whose S -points parametrize surjective $\mathcal{O}_{\mathcal{X}_S}$ -module maps.
- (4) Let $\text{Inj}(\mathcal{E}, \mathcal{F}) \subset \mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})$ be the subfunctor of $\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})$ whose S -points parametrize ‘fiberwise-injective’ $\mathcal{O}_{\mathcal{X}_S}$ -module maps. Precisely, this is the functor parametrizing $\mathcal{O}_{\mathcal{X}_S}$ -module maps $m : \mathcal{E}_S \rightarrow \mathcal{F}_S$ such that for every geometric point $\bar{x} = \text{Spa}(C, C^+) \rightarrow S$, the pullback of $m : \mathcal{E}_S \rightarrow \mathcal{F}_S$ along the induced map $\mathcal{X}_{\bar{x}} \rightarrow \mathcal{X}_S$ gives an injective $\mathcal{O}_{\mathcal{X}_{\bar{x}}}$ -module map.⁴

⁴The condition defining Inj is much stronger than the condition that $m : \mathcal{E}_S \rightarrow \mathcal{F}_S$ be injective; note

- (5) Let $\mathcal{H}^0(\mathcal{E})^\times := \text{Inj}(\mathcal{O}, \mathcal{E}) \subset \mathcal{H}^0(\mathcal{E})$ be the functor parametrizing sections of \mathcal{E} that are not identically zero on any fiber of $\mathcal{X}_S \dashrightarrow S$.
- (6) Let $\text{Aut}(\mathcal{E})$ be the functor sending $f : S \rightarrow \text{Spa } F$ to the group of $\mathcal{O}_{\mathcal{X}_S}$ -module automorphisms of \mathcal{E}_S .

Our first order of business is to show that these functors are all locally spatial and partially proper diamonds. We remark that the functors \mathcal{H}^0 and Hom are Banach–Colmez spaces [2, 12] and, in particular, can be given some meaningful geometric structure without appealing to the theory of diamonds. However, the other four functors defined above are not Banach–Colmez spaces (indeed, they are not even valued in \mathbf{Q}_p -vector spaces), and the theory of diamonds seems essential to our analysis of them.

Let

$$\tilde{\mathbf{D}}^n = \text{Spa}(F^\circ[[T_1^{1/p^\infty}, \dots, T_n^{1/p^\infty}]], F^\circ[[T_1^{1/p^\infty}, \dots, T_n^{1/p^\infty}]] \times_{\text{Spa}(F^\circ, F^\circ)} \text{Spa}(F, F^\circ)$$

denote the n -dimensional open perfectoid unit disk over F .

Proposition 3.3.2. (1) *The functor $\mathcal{H}^0(\mathcal{E})$ is a diamond: if \mathcal{E} has only positive slopes, then we can find an isomorphism*

$$\mathcal{H}^0(\mathcal{E}) \simeq \tilde{\mathbf{D}}^d / \underline{A},$$

where $d = \text{deg } \mathcal{E}$ and where $A \simeq \mathbf{Q}_p^m$ is an abelian locally profinite group acting freely on $\tilde{\mathbf{D}}^d$; in general, there is a natural isomorphism $\mathcal{H}^0(\mathcal{E}) \simeq \underline{\mathbf{Q}}_p^n \times \mathcal{H}^0(\mathcal{E}^{>0})$ for some $n \geq 0$.

- (2) *The diamond $\mathcal{H}^0(\mathcal{E})$ is partially proper and locally spatial, and equidimensional of dimension $\text{deg } \mathcal{E}^{\geq 0}$. Furthermore, any nonempty open subfunctor of $\mathcal{H}^0(\mathcal{E})$ has an F -point.*

We point out that the dimension of $\mathcal{H}^0(\mathcal{E})$ computed above coincides with the ‘principal dimension’ of this object in the language of Banach–Colmez spaces [2].

Remark 3.3.3. The partial properness of $\mathcal{H}^0(\mathcal{E})$ is a formal consequence of the fact that the category of vector bundles on a relative curve $\mathcal{X}_{\text{Spa}(R, R^+)}$ is canonically independent of the choice of R^+ ; cf. [11, Theorem 8.7.7]. In fact, every functor defined above is partially proper, for the same reason.

Proof. Writing $\mathcal{E} \simeq \mathcal{E}^{>0} \oplus \mathcal{O}^n \oplus \mathcal{E}'$, where \mathcal{E}' has only negative slopes, the isomorphism $\mathcal{H}^0(\mathcal{E}_1 \oplus \mathcal{E}_2) \cong \mathcal{H}^0(\mathcal{E}_1) \times_{\text{Spd } F} \mathcal{H}^0(\mathcal{E}_2)$ together with the identification $\mathcal{H}^0(\mathcal{O}) \simeq \underline{E}$ reduces us from the general case to the case of positive slopes. Writing $\mathcal{E} \simeq \oplus_i \mathcal{O}(\lambda_i)$ and observing that

$$\tilde{\mathbf{D}}^{d_1} / \underline{A}_1 \times_{\text{Spd } F} \tilde{\mathbf{D}}^{d_2} / \underline{A}_2 \cong \tilde{\mathbf{D}}^{d_1+d_2} / \underline{A}_1 \times \underline{A}_2,$$

that the association sending S to the set of injective ms is not even a presheaf. Note also that there is a natural transformation $\text{Surj}(\mathcal{F}^\vee, \mathcal{E}^\vee) \rightarrow \text{Inj}(\mathcal{E}, \mathcal{F})$; this turns out to be an open immersion, although it typically is not an isomorphism.

we reduce further to the case where $\mathcal{E} = \mathcal{O}(\lambda)$ for some $\lambda = d/h \in \mathbf{Q}_{>0}$. Let E'/E be the unramified extension of degree h , so we have a natural finite étale map $r : \mathcal{X}_{E',S} \rightarrow \mathcal{X}_S$ such that $r_*\mathcal{O}(d) \simeq \mathcal{O}(\lambda)$ functorially in S . Choose a short exact sequence

$$0 \rightarrow \mathcal{O}^{d-1} \rightarrow \mathcal{O}(1)^d \rightarrow \mathcal{O}(d) \rightarrow 0$$

of vector bundles on $\mathcal{X}_{E',F}$. Pushing forward along r_* and applying $\mathcal{H}^0(-)$ gives rise to a short exact sequence of abelian group sheaves on $\text{Perf}_{/\text{Spa } F}$, giving an identification

$$\mathcal{H}^0(\mathcal{O}(d/h)) \simeq \mathcal{H}^0(\mathcal{O}(1/h)^d)/\mathcal{H}^0(\mathcal{O}^{h(d-1)}).$$

It is easy to see that $\mathcal{H}^0(\mathcal{O}^{h(d-1)}) \simeq \underline{E}^{h(d-1)}$ as abelian group sheaves on $\text{Perf}_{/\text{Spa } F}$. Moreover, there is an isomorphism $\mathcal{H}^0(\mathcal{O}(1/h)) \simeq \tilde{\mathbf{D}}^1$: quite generally, for any given integers $0 < d \leq h$, one can exhibit an isomorphism

$$\mathcal{H}^0(\mathcal{O}(d/h)) \simeq \tilde{\mathbf{D}}^d$$

by identifying $\mathcal{H}^0(\mathcal{O}(d/h))$ with the universal cover of an isoclinic π -divisible \mathcal{O}_E -module of height h and dimension d (cf. [4, 16]). Putting things together, (1) follows.

For (2), assume for the moment that $\mathcal{H}^0(\mathcal{E})$ is locally spatial. Fix a presentation as in (1). After choosing an open profinite subgroup $A_0 \subset A$, the map

$$\tilde{\mathbf{D}}^d \rightarrow \tilde{\mathbf{D}}^d/A$$

factors as

$$\tilde{\mathbf{D}}^d \rightarrow \tilde{\mathbf{D}}^d/A_0 \rightarrow \tilde{\mathbf{D}}^d/A.$$

By [17, Proposition 4.3.2], the first arrow here is profinite étale surjective, and the second arrow is separated, étale and surjective by [13, Lemma 10.13]. The space $\tilde{\mathbf{D}}^d$ is equidimensional of dimension d , so $\tilde{\mathbf{D}}^d/A_0$ is equidimensional of dimension d by Lemma 3.2.3. We are now reduced to checking that if $f : X \rightarrow Y$ is a surjective étale map of locally spatial diamonds, then $\dim X = \dim Y$. Since f is generalizing, the inequality $\dim X \geq \dim Y$ is clear. For the opposite inequality, one easily reduces to the case where $X \rightarrow Y$ is a surjective finite étale map with Y connected, by [13, Lemma 11.31], which we then leave as an exercise for the interested reader.⁵

We still need to check that $\mathcal{H}^0(\mathcal{E})$ is locally spatial. Fix a presentation as in (1), and choose a quasicompact open subgroup $U_0 \subset \tilde{\mathbf{D}}^d$. Set $A_0 = U_0 \cap A$, so this is a profinite group sheaf acting freely on U_0 . Then $V = U_0/A_0$ is an open subdiamond of $\mathcal{H}^0(\mathcal{E})$, and it is spatial by the subsequent lemma. Since the p^{-n} -dilates of V cover all of $\mathcal{H}^0(\mathcal{E})$, the latter is locally spatial, as desired.

Finally, the statement on F -points follows from the explicit presentation in (1) together with the (easy) analogous statement for $\tilde{\mathbf{D}}^n$. □

⁵Sketch: By a standard argument, one can dominate X by a surjective finite étale map $X' \rightarrow X$ such that $X' \rightarrow Y$ is finite étale and Galois for some finite group G ; but then $\dim X \leq \dim X' = \dim Y$, where the latter equality follows from Proposition 3.2.3.

In the previous proof, we made use of the following lemma.

Lemma 3.3.4. *Let X be a spatial diamond with a free \underline{G} -action for some profinite group G . Then X/\underline{G} is a spatial diamond.*

Proof. This is immediate upon combining Lemma 10.13 and Proposition 11.24 from [13]. □

Proposition 3.3.5. *The functor $\mathcal{H}om(\mathcal{E}, \mathcal{F})$ is a locally spatial (and partially proper) diamond over $\mathrm{Spd} F$, equidimensional of dimension $\mathrm{deg}(\mathcal{E}^\vee \otimes \mathcal{F}) \geq 0$.*

Proof. This is immediate from the identification $\mathcal{H}om(\mathcal{E}, \mathcal{F}) \cong \mathcal{H}^0(\mathcal{E}^\vee \otimes \mathcal{F})$. □

Proposition 3.3.6. *The functors $\mathrm{Surj}(\mathcal{E}, \mathcal{F})$ and $\mathrm{Inj}(\mathcal{E}, \mathcal{F})$ are both open subfunctors of $\mathcal{H}om(\mathcal{E}, \mathcal{F})$. In particular, $\mathrm{Surj}(\mathcal{E}, \mathcal{F})$ and $\mathrm{Inj}(\mathcal{E}, \mathcal{F})$ are locally spatial (and partially proper) diamonds over $\mathrm{Spd} F$, each of which is either empty or equidimensional of dimension $\mathrm{deg}(\mathcal{E}^\vee \otimes \mathcal{F}) \geq 0$.*

Proof for Surj . Choose some $T \in \mathrm{Perf}$ together with a surjective quasi-pro-étale morphism $T \rightarrow \mathcal{H}om(\mathcal{E}, \mathcal{F})$. Over \mathcal{X}_T , we get a ‘universal’ $\mathcal{O}_{\mathcal{X}_T}$ -module map $m^{\mathrm{univ}} : \mathcal{E}_T \rightarrow \mathcal{F}_T$; let \mathcal{Q}_T be the cokernel of m^{univ} . By a standard argument, the support of \mathcal{Q}_T is Zariski-closed in \mathcal{X}_T , and we write $Z \subset |\mathcal{X}_T|$ for the associated closed subset.

Next, we observe that the map $|\mathcal{X}_T| \rightarrow |T|$ is closed. Indeed, this is a specializing quasicompact spectral map of locally spectral spaces, so the image of any closed subset is pro-constructible (by quasicompactness and spectrality) and stable under specialization, and hence closed by [14, Tag 0903]. In particular, the subset $V = \mathrm{im}(|\mathcal{X}_T| \rightarrow |T|)(Z) \subset |T|$ is closed. We also observe that a geometric point $x : \mathrm{Spd}(C, C^+) \rightarrow \mathcal{H}om(\mathcal{E}, \mathcal{F})$ defines a point of Surj (resp. $\mathcal{H}om \setminus \mathrm{Surj}$) if and only if the preimage of $|x|$ in $|T|$ is disjoint from V (resp. contained in V). In particular, the open subset $U = |T| \setminus V \subset |T|$ is the preimage of a subset $W \subset |\mathcal{H}om(\mathcal{E}, \mathcal{F})|$; since $|T| \twoheadrightarrow |\mathcal{H}om(\mathcal{E}, \mathcal{F})|$ is a quotient map, W is open. But now Surj can be identified with the open subdiamond of $\mathcal{H}om(\mathcal{E}, \mathcal{F})$ corresponding to the open subset W , so we win.

Proof for Inj . Set $r = \mathrm{rank}(\mathcal{E})$; by the formula

$$\mathrm{Inj}(\mathcal{E}, \mathcal{F}) \cong \mathcal{H}om(\mathcal{E}, \mathcal{F}) \times_{\mathcal{H}om(\wedge^r \mathcal{E}, \wedge^r \mathcal{F})} \mathrm{Inj}(\wedge^r \mathcal{E}, \wedge^r \mathcal{F}),$$

we reduce to the case where \mathcal{E} is a line bundle. After twisting, we reduce further to the case $\mathcal{E} = \mathcal{O}$; in other words, we need to prove that $\mathcal{H}^0(\mathcal{F})^\times$ is an open subfunctor of $\mathcal{H}^0(\mathcal{F})$. Fix an identification $\mathcal{F} = \bigoplus_{1 \leq i \leq n} \mathcal{O}(\lambda_i)$, and (for brevity) set $\mathcal{H}_i = \mathcal{H}^0(\mathcal{O}(\lambda_i))$ and $\mathcal{H}_i^\times = \mathcal{H}^0(\mathcal{O}(\lambda_i))^\times$. Under the identification

$$\mathcal{H}^0(\mathcal{F}) = \mathcal{H}_1 \times_{\mathrm{Spd} F} \cdots \times_{\mathrm{Spd} F} \mathcal{H}_n,$$

it is easy to see that the subfunctor $\mathcal{H}^0(\mathcal{F})^\times$ on the left-hand side is covered by the union of the subfunctors

$$U_i = \mathcal{H}_1 \times \cdots \times \mathcal{H}_{i-1} \times \mathcal{H}_i^\times \times \mathcal{H}_{i+1} \times \cdots \times \mathcal{H}_n, \quad 1 \leq i \leq n$$

on the right-hand side (here we have omitted the subscripted $\mathrm{Spd} F$'s for brevity). This, finally, reduces us to showing that $\mathcal{H}^0(\mathcal{O}(\lambda))^\times$ is an open subfunctor of $\mathcal{H}^0(\mathcal{O}(\lambda))$, for any fixed $\lambda \geq 0$. The case $\lambda = 0$ is easy and left to the reader. For $\lambda = d/h > 0$, writing $\mathcal{O}(\lambda)$ as the pushforward of $\mathcal{O}(d)/\mathcal{X}_{E',F}$ as in the proof of Proposition 3.3.2 reduces us to the case $h = 1$. In this case, the functors $\mathcal{H}(\mathcal{O}(d))$ and $\mathcal{H}(\mathcal{O}(d))^\times$ agree (by definition) with the functors denoted by $\mathbf{B}_{\mathrm{Spa} F}^{\varphi=\pi^d}$ and $(\mathbf{B}^{\varphi=\pi^d} \setminus \{0\})_{\mathrm{Spa} F}$ in [3, § 2.2], and our claim is exactly the content of [3, Lemme 2.10]. \square

Proposition 3.3.7. *The functor $\mathrm{Aut}(\mathcal{E})$ is an open and partially proper subdiamond of $\mathrm{Hom}(\mathcal{E}, \mathcal{E})$, equidimensional of dimension $\mathrm{deg}(\mathcal{E}^\vee \otimes \mathcal{E}) \geq 0$.*

Proof. This is immediate from the previous proposition together with the identification $\mathrm{Aut}(\mathcal{E}) \cong \mathrm{Surj}(\mathcal{E}, \mathcal{E})$, which holds by consideration of rank and degree. \square

We now explain how to reduce step one of the strategy outlined in Introduction to a combinatorial problem. The key result here is Theorem 3.3.11.

Definition 3.3.8. For any vector bundles $\mathcal{E}, \mathcal{F}, \mathcal{Q}$, composition of bundle maps induces a natural map of diamonds

$$\mathrm{Surj}(\mathcal{E}, \mathcal{Q}) \times_{\mathrm{Spd} F} \mathrm{Inj}(\mathcal{Q}, \mathcal{F}) \rightarrow \mathrm{Hom}(\mathcal{E}, \mathcal{F}),$$

and we define

$$|\mathrm{Hom}(\mathcal{E}, \mathcal{F})_{\mathcal{Q}}| \subset |\mathrm{Hom}(\mathcal{E}, \mathcal{F})|$$

as the image of the induced map on topological spaces.

As the notation suggests, $|\mathrm{Hom}(\mathcal{E}, \mathcal{F})_{\mathcal{Q}}|$ is the underlying topological space of a subdiamond $\mathrm{Hom}(\mathcal{E}, \mathcal{F})_{\mathcal{Q}}$ of $\mathrm{Hom}(\mathcal{E}, \mathcal{F})$, which (more or less) parametrizes bundle maps $\mathcal{E} \rightarrow \mathcal{F}$ with image isomorphic to \mathcal{Q} at all geometric points. However, the diamond $\mathrm{Hom}(\mathcal{E}, \mathcal{F})_{\mathcal{Q}}$ is a little obscure (in particular, it is not entirely clear to us that it is locally spatial), so we avoid making any explicit use of this object. Fortunately, this does not cause any real complication.

Proposition 3.3.9. *For any $\mathcal{E}, \mathcal{F}, \mathcal{Q}$, $|\mathrm{Hom}(\mathcal{E}, \mathcal{F})_{\mathcal{Q}}|$ is stable under generalization and specialization inside $|\mathrm{Hom}(\mathcal{E}, \mathcal{F})|$.*

Proof. The subset in question is defined as the image of $|Y| \rightarrow |X|$, for some map of partially proper and locally spatial diamonds $Y \rightarrow X$. Quite generally, if $Y \rightarrow X$ is any map of locally spatial diamonds, the image of $|Y| \rightarrow |X|$ is stable under generalization by Proposition 3.1.8.iii. Orthogonally, one immediately checks from the valuative criterion that the image of $|Y| \rightarrow |X|$ is stable under specialization for $Y \rightarrow X$ any partially proper map of diamonds. Since any map between diamonds partially proper over a fixed base is automatically partially proper, the claim follows. \square

We now have the following crucial lemma.

Lemma 3.3.10. *For any $\mathcal{E}, \mathcal{F}, \mathcal{Q}$ as above, the subset $|\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})_{\mathcal{Q}}| \subset |\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})|$ is either empty, or of dimension*

$$\text{deg}(\mathcal{E}^\vee \otimes \mathcal{Q})^{\geq 0} + \text{deg}(\mathcal{Q}^\vee \otimes \mathcal{F})^{\geq 0} - \text{deg}(\mathcal{Q}^\vee \otimes \mathcal{Q})^{\geq 0}.$$

Proof. Recall that $|\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})_{\mathcal{Q}}|$ is defined as the image of the map on topological spaces associated with the map of locally spatial diamonds

$$f : \text{Surj}(\mathcal{E}, \mathcal{Q}) \times_{\text{Spd } F} \text{Inj}(\mathcal{Q}, \mathcal{F}) \xrightarrow{(s,i) \mapsto ios} \mathcal{H}\text{om}(\mathcal{E}, \mathcal{F}).$$

If either of the functors on the left is empty, there is nothing to prove. If both functors on the left are nonempty, then

$$\begin{aligned} \dim \text{Surj}(\mathcal{E}, \mathcal{Q}) \times_{\text{Spd } F} \text{Inj}(\mathcal{Q}, \mathcal{F}) &= \dim \text{Surj}(\mathcal{E}, \mathcal{Q}) + \dim \text{Inj}(\mathcal{Q}, \mathcal{F}) \\ &= \text{deg}(\mathcal{E}^\vee \otimes \mathcal{Q})^{\geq 0} + \text{deg}(\mathcal{Q}^\vee \otimes \mathcal{F})^{\geq 0} \end{aligned}$$

by Propositions 3.3.2 and 3.3.6. Moreover, the fiber of f over any rank one point y is an $\text{Aut}(\mathcal{Q}) \times_{\text{Spd } F} y$ -torsor, and hence can be identified with $\text{Aut}(\mathcal{Q}) \times_{\text{Spd } F} \bar{y}$ after taking a profinite étale covering of y by some geometric point $\bar{y} = \text{Spd } C$. Since

$$\dim \text{Aut}(\mathcal{Q}) = \text{deg}(\mathcal{Q}^\vee \otimes \mathcal{Q})^{\geq 0},$$

the result now follows from Lemma 3.2.5. □

Theorem 3.3.11. *Suppose \mathcal{E} and \mathcal{F} are vector bundles admitting a nonzero map $\mathcal{E} \rightarrow \mathcal{F}$ such that for any $\mathcal{Q} \subsetneq \mathcal{F}$ which also occurs as a quotient of \mathcal{E} , we have a strict inequality*

$$\text{deg}(\mathcal{E}^\vee \otimes \mathcal{Q})^{\geq 0} + \text{deg}(\mathcal{Q}^\vee \otimes \mathcal{F})^{\geq 0} < \text{deg}(\mathcal{E}^\vee \otimes \mathcal{F})^{\geq 0} + \text{deg}(\mathcal{Q}^\vee \otimes \mathcal{Q})^{\geq 0}.$$

Then there exists a surjective bundle map $\mathcal{E} \rightarrow \mathcal{F}$.

Proof. We need to show that $\text{Surj}(\mathcal{E}, \mathcal{F})(F)$ is nonempty. By Proposition 3.3.6, $\text{Surj}(\mathcal{E}, \mathcal{F})$ is an open subfunctor of $\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})$, and hence admits F -points if it is nonempty (by Proposition 3.3.2). Since $\text{Surj}(\mathcal{E}, \mathcal{F})$ is nonempty if and only if $|\text{Surj}(\mathcal{E}, \mathcal{F})|$ is nonempty, we can argue on topological spaces.

By the definitions above, we obtain a decomposition

$$X \stackrel{\text{def}}{=} |\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})| \setminus |\text{Surj}(\mathcal{E}, \mathcal{F})| = \coprod_{\mathcal{Q} \in S} |\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})_{\mathcal{Q}}|,$$

where S denotes the set of isomorphism classes of subbundles $\mathcal{Q} \subsetneq \mathcal{F}$, which also occur as quotients of \mathcal{E} . By Proposition 3.3.9, any chain in X is entirely contained in $|\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})_{\mathcal{Q}}|$ for some fixed \mathcal{Q} , so then

$$\dim X = \sup_{\mathcal{Q} \in S} \dim |\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})_{\mathcal{Q}}| \leq \sup_{\mathcal{Q} \in S} (\text{deg}(\mathcal{E}^\vee \otimes \mathcal{Q})^{\geq 0} + \text{deg}(\mathcal{Q}^\vee \otimes \mathcal{F})^{\geq 0} - \text{deg}(\mathcal{Q}^\vee \otimes \mathcal{Q})^{\geq 0})$$

by the previous lemma. By the assumptions of the theorem, we then deduce a strict inequality $\dim X < \text{deg}(\mathcal{E}^\vee \otimes \mathcal{F})^{\geq 0}$. If $|\text{Surj}(\mathcal{E}, \mathcal{F})|$ were empty, however, $X = |\mathcal{H}\text{om}(\mathcal{E}, \mathcal{F})|$ would have dimension $\text{deg}(\mathcal{E}^\vee \otimes \mathcal{F})^{\geq 0}$ by Proposition 3.3.5, and this is a contradiction. Thus $|\text{Surj}(\mathcal{E}, \mathcal{F})|$ is nonempty as desired. □

We now turn to step two of the strategy outlined in Introduction, which we will also reduce to a combinatorial problem. Here, the relevant diamonds are defined as follows.

Definition 3.3.12. For some fixed vector bundles $\mathcal{E}, \mathcal{F}, \mathcal{K}$ such that $\text{rank } \mathcal{K} + \text{rank } \mathcal{F} = \text{rank } \mathcal{E}$ and $\text{deg } \mathcal{K} + \text{deg } \mathcal{F} = \text{deg } \mathcal{E}$, define $\text{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}}$ as the subfunctor of $\text{Surj}(\mathcal{E}, \mathcal{F})$ whose S -points (for a given $S \rightarrow \text{Spd } F$) parametrize surjective bundle maps $q : \mathcal{E}_S \rightarrow \mathcal{F}_S$ with the property that $\ker q$ is isomorphic to \mathcal{K} after pullback along any geometric point $\bar{x} \rightarrow S$.

Equivalently, the map $S \rightarrow \text{Surj}(\mathcal{E}, \mathcal{F})$ corresponding to a surjection $q : \mathcal{E}_S \rightarrow \mathcal{F}_S$ defines a point of $\text{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}}$ exactly when the HN polygon of $\ker q$, regarded as a function on $|S|$, is constant and coincides with the HN polygon of \mathcal{K} .

Proposition 3.3.13. *The functor $\text{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}}$ is a locally spatial (and partially proper) diamond over $\text{Spd } F$.*

Proof. Choose a perfectoid space T together with a surjective and quasi-pro-étale map $f : T \rightarrow \text{Surj}(\mathcal{E}, \mathcal{F})$. Let \mathcal{V} be the bundle on \mathcal{X}_T defined by the kernel of the ‘universal’ quotient map $q^{\text{univ}} : \mathcal{E}_T \rightarrow \mathcal{F}_T$. Set $n = \text{rank}(\mathcal{E}) - \text{rank}(\mathcal{F})$; for a given HN polygon P of width n , let $|T|^{\geq P}$ (resp. $|T|^{\leq P}$) be the set of points x where $\text{HN}(\mathcal{V}_x) \geq P$ (resp. where $\text{HN}(\mathcal{V}_x) \leq P$). By [KL, Theorem 7.4.5], the locus $|T|^{\geq P}$ (resp. $|T|^{\leq P}$) is closed (resp. open) in $|T|$; moreover, these subsets are stable under generalization. Since $|f| : |T| \rightarrow |\text{Surj}(\mathcal{E}, \mathcal{F})|$ is a quotient map, one checks directly that $|T|^{\geq P}$ (resp. $|T|^{\leq P}$) is the preimage of a generalizing and closed (resp. open) subset $|\text{Surj}(\mathcal{E}, \mathcal{F})|^{\geq P}$ (resp. $|\text{Surj}(\mathcal{E}, \mathcal{F})|^{\leq P}$) of $|\text{Surj}(\mathcal{E}, \mathcal{F})|$. In particular,

$$|\text{Surj}(\mathcal{E}, \mathcal{F})|^{\geq P} \cap |\text{Surj}(\mathcal{E}, \mathcal{F})|^{\leq P}$$

is a locally closed generalizing subset of $|\text{Surj}(\mathcal{E}, \mathcal{F})|$, and hence corresponds to a locally spatial subdiamond

$$\text{Surj}(\mathcal{E}, \mathcal{F})^P \subset \text{Surj}(\mathcal{E}, \mathcal{F})$$

by the proof of [13, Proposition 11.20]. Taking $P = \text{HN}(\mathcal{K})$, the functor $\text{Surj}(\mathcal{E}, \mathcal{F})^P$ identifies with $\text{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}}$, and we conclude. □

Again, we have a clean dimension formula in certain cases.

Lemma 3.3.14. *Let \mathcal{E}, \mathcal{F} , and \mathcal{K} be as above, and suppose moreover that \mathcal{F} is semistable. Then $\text{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}}$ is either empty, or equidimensional of dimension*

$$\text{deg}(\mathcal{E} \otimes \mathcal{K}^\vee) \geq 0 - \text{deg}(\mathcal{K} \otimes \mathcal{K}^\vee) \geq 0.$$

Proof. Recall that $\text{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}}$ parametrizes surjective bundle maps $q : \mathcal{E} \rightarrow \mathcal{F}$ with kernel isomorphic to \mathcal{K} at all geometric points. Let

$$\text{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}, \heartsuit} \rightarrow \text{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}}$$

be the $\text{Aut}(\mathcal{K})$ -torsor parametrizing isomorphisms $\ker q \simeq \mathcal{K}$. Arguing as in the proof of Lemma 3.3.10, one easily checks that

$$\begin{aligned} \dim \operatorname{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}, \heartsuit} &= \dim \operatorname{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}} + \dim \operatorname{Aut}(\mathcal{K}) \\ &= \dim \operatorname{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}} + \deg(\mathcal{K}^\vee \otimes \mathcal{K}) \geq 0. \end{aligned}$$

Next, we observe that

$$\operatorname{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}, \heartsuit}$$

can also be described as the functor whose T -points parametrize isomorphism classes of short exact sequences

$$0 \rightarrow \mathcal{K}_T \rightarrow \mathcal{E}_T \rightarrow \mathcal{F}_T \rightarrow 0$$

on \mathcal{X}_T . Dualizing the set of such short exact sequences induces a canonical isomorphism

$$\operatorname{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}, \heartsuit} \cong \operatorname{Surj}(\mathcal{E}^\vee, \mathcal{K}^\vee)^{\mathcal{F}^\vee, \heartsuit}.$$

Returning to the matter at hand, we now assume that \mathcal{F} is semistable. Then the proof of Proposition 3.3.13 shows that $\operatorname{Surj}(\mathcal{E}^\vee, \mathcal{K}^\vee)^{\mathcal{F}^\vee} = \operatorname{Surj}(\mathcal{E}^\vee, \mathcal{K}^\vee)^{\leq \operatorname{HN}(\mathcal{F}^\vee)}$ is an open subfunctor of $\operatorname{Surj}(\mathcal{E}^\vee, \mathcal{K}^\vee)$. Moreover, the dimension formula above degenerates to the equality

$$\dim \operatorname{Surj}(\mathcal{E}^\vee, \mathcal{K}^\vee)^{\mathcal{F}^\vee, \heartsuit} = \dim \operatorname{Surj}(\mathcal{E}^\vee, \mathcal{K}^\vee)^{\mathcal{F}^\vee}.$$

Putting this together with Proposition 3.3.6, we deduce that

$$\dim \operatorname{Surj}(\mathcal{E}^\vee, \mathcal{K}^\vee)^{\mathcal{F}^\vee, \heartsuit} = \deg(\mathcal{E} \otimes \mathcal{K}^\vee) \geq 0$$

for semistable \mathcal{F} . But then

$$\begin{aligned} \dim \operatorname{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}} &= \dim \operatorname{Surj}(\mathcal{E}, \mathcal{F})^{\mathcal{K}, \heartsuit} - \deg(\mathcal{K}^\vee \otimes \mathcal{K}) \geq 0 \\ &= \dim \operatorname{Surj}(\mathcal{E}^\vee, \mathcal{K}^\vee)^{\mathcal{F}^\vee, \heartsuit} - \deg(\mathcal{K}^\vee \otimes \mathcal{K}) \geq 0 \\ &= \deg(\mathcal{E} \otimes \mathcal{K}^\vee) \geq 0 - \deg(\mathcal{K}^\vee \otimes \mathcal{K}) \geq 0, \end{aligned}$$

as desired. □

3.4. Summary of the strategy

The goal of §§ 4–5 is to establish Theorem 1.1.2, which we reproduce below for the reader’s convenience:

Theorem. Let \mathcal{F}_1 and \mathcal{F}_2 be semistable vector bundles on X with $\mu(\mathcal{F}_1) < \mu(\mathcal{F}_2)$. Then any vector bundle \mathcal{E} such that

$$\operatorname{HN}(\mathcal{E}) \leq \operatorname{HN}(\mathcal{F}_1 \oplus \mathcal{F}_2)$$

can be realized as an extension

$$0 \rightarrow \mathcal{F}_1 \rightarrow \mathcal{E} \rightarrow \mathcal{F}_2 \rightarrow 0.$$

Again, the strategy of our proof is as follows. We note that it suffices to show the theorem when the slope of \mathcal{F}_2 is strictly greater than the maximal slope of \mathcal{E} . Otherwise, we would have $\mathcal{F}_2 = \mathcal{O}(s)^m$ and $\mathcal{E} = \mathcal{E}' \oplus \mathcal{O}(s)^n$ for some \mathcal{E}' with maximal slope less than s and $n \leq m$. Solving the extension problem for

$$0 \rightarrow \mathcal{F}_1 \rightarrow \mathcal{E}' \rightarrow \mathcal{O}(s)^{m-n} \rightarrow 0$$

and then direct summing with $0 \rightarrow 0 \rightarrow \mathcal{O}(s)^n \rightarrow \mathcal{O}(s)^n \rightarrow 0$ gives the full theorem.

It then clearly suffices to show that the assumptions of Theorem 1.1.2 plus the additional assumption of strict inequality of slopes imply the following two assertions.

- (1) There exists a surjection $\mathcal{E} \rightarrow \mathcal{F}_2 \rightarrow 0$.
- (2) If \mathcal{E} admits a surjection $\mathcal{E} \rightarrow \mathcal{F}_2 \rightarrow 0$, then it admits such a surjection with kernel isomorphic to \mathcal{F}_1 .

The bulk of §4 is devoted to showing that the inequality

$$\deg(\mathcal{E}^\vee \otimes \mathcal{Q})^{\geq 0} + \deg(\mathcal{Q}^\vee \otimes \mathcal{F}_2)^{\geq 0} < \deg(\mathcal{E}^\vee \otimes \mathcal{F}_2)^{\geq 0} + \deg(\mathcal{Q}^\vee \otimes \mathcal{Q})^{\geq 0}$$

holds for any proper subbundle $\mathcal{Q} \subsetneq \mathcal{F}_2$, which also occurs as a quotient of \mathcal{E} . This collection of inequalities implies (1) by Theorem 3.3.11.

Assuming \mathcal{E} satisfies (1), Lemma 3.3.14 and Proposition 3.3.6 reduce (2) to the verification of the inequality

$$\deg(\mathcal{E} \otimes \mathcal{K}^\vee)^{\geq 0} - \deg(\mathcal{K} \otimes \mathcal{K}^\vee)^{\geq 0} < \deg(\mathcal{E}^\vee \otimes \mathcal{F}_2)^{\geq 0}$$

for a certain finite collection of \mathcal{K} 's (this implication will be fleshed out in §5.4). The proof of these inequalities forms the subject of §5; the precise statement is given in Theorem 5.1.1.

4. Step one: surjections of vector bundles

4.1. The goal

This section is devoted to establishing Step (1), as outlined in §3.4, wherein we show that the obvious necessary numerical conditions required for a bundle \mathcal{E} to surject onto a given semistable bundle \mathcal{F} are also sufficient. The precise statement is as follows.

Theorem 4.1.1. *Let \mathcal{F} be a semistable vector bundle on X . Let \mathcal{E} be a vector bundle on X such that*

- *rank $\mathcal{E} >$ rank \mathcal{F} , and*
- *the maximal slope of \mathcal{E} is less than or equal to $\mu(\mathcal{F})$.*

Then \mathcal{E} admits a surjection onto \mathcal{F} .

Note that, by an argument similar to the one in 3.4, it suffices to treat the case where the maximal slope of \mathcal{E} is strictly less than $\mu(\mathcal{F})$.

By this reduction and Theorem 3.3.11, we are now reduced to proving the inequality

$$\deg(\mathcal{E}^\vee \otimes \mathcal{Q})^{\geq 0} + \deg(\mathcal{Q}^\vee \otimes \mathcal{F})^{\geq 0} < \deg(\mathcal{E}^\vee \otimes \mathcal{F})^{\geq 0} + \deg(\mathcal{Q}^\vee \otimes \mathcal{Q})^{\geq 0} \tag{4.1}$$

for any \mathcal{Q} which is both a quotient of \mathcal{E} and a proper subbundle of \mathcal{F} , under the assumption that the maximal slope of \mathcal{E} is *strictly* less than $\mu(\mathcal{F})$. (The strict inequality is false without this latter strictness assumption, which is why we made the initial reduction to this case.)




Figure 5. Illustration of Corollary 4.2.2.

4.2. Some lemmas

We begin by translating the assumption that \mathcal{Q} is a quotient of \mathcal{E} into numerical constraints on the slopes of their HN polygons.

Lemma 4.2.1. *Suppose \mathcal{Q} is a quotient of \mathcal{E} . Then for every slope μ , the bundle $\mathcal{Q}^{\leq \mu}$ is a quotient of $\mathcal{E}^{\leq \mu}$. (Dually, if \mathcal{K} is a subbundle of \mathcal{E} , for every slope μ , the bundle $\mathcal{K}^{\geq \mu}$ is a subbundle of $\mathcal{E}^{\geq \mu}$.)*

Proof. The composite quotient map $\mathcal{E} \rightarrow \mathcal{Q} \rightarrow \mathcal{Q}^{\leq \mu}$ necessarily factors through $\mathcal{E}^{\leq \mu}$, since $\mathcal{E}^{> \mu}$ is a direct sum of stable bundles with slopes greater than μ . □

Since we will prove (4.1) by interpreting the terms on both sides as areas of polygons, it is convenient to have a ‘geometric’ reformulation of the preceding lemma.

Corollary 4.2.2. *Let \mathcal{Q} be a quotient of \mathcal{E} . Translate the HN polygons for \mathcal{Q} and \mathcal{E} in the plane so that both their right endpoints lie at the origin. Then for every integer i from 1 to $\text{rank } \mathcal{Q}$, the slope of the part of the HN polygon for \mathcal{Q} lying in the strip $[-i, -i + 1] \times \mathbf{R}$ is greater than or equal to the slope of the part of the HN polygon for \mathcal{E} lying in this strip. (See Figure 5.)*

In particular, the HN polygon for \mathcal{Q} always lies below the HN polygon for \mathcal{E} (when aligning right endpoints).

Proof. Suppose for the sake of contradiction that on some interval $[-i, -i + 1]$, the slope of $\text{HN}(\mathcal{Q})$ is less than that of $\text{HN}(\mathcal{E})$. Let μ be the slope of the HN polygon for \mathcal{Q} on this strip. Then we would have $\text{rank } \mathcal{E}^{\leq \mu} < i \leq \text{rank } \mathcal{Q}^{\leq \mu}$, contradicting Lemma 4.2.1. □

Remark 4.2.3. In this section, we will usually find it more convenient to align HN polygons with their *right endpoint* at the origin. (In future sections, it will usually be more convenient to align HN polygons with their *left endpoint* at the origin.)




Figure 6. Area interpretation of the left side of (4.2) in the case $s = 3, s' = 4$, with slope vectors $\vec{\mu}_1 \succeq \vec{\lambda}_1 > \vec{\mu}_2 > \vec{\mu}_3 \succeq \vec{\lambda}_2 > \vec{\mu}_4 \succeq \vec{\lambda}_3$.

4.3. Proof of Theorem 4.1.1

We are now ready to begin the proof of Theorem 4.1.1, the main result of this section. The idea is to interpret each term of the inequality as the area of a certain polygon, and to transform the polygons by area-preserving operations (specifically translations and shears) until they can be easily compared.

Proof of Theorem 4.1.1. Recall that we are trying to show

$$\text{deg}(\mathcal{E}^\vee \otimes \mathcal{Q})^{\geq 0} - \text{deg}(\mathcal{Q}^\vee \otimes \mathcal{Q})^{\geq 0} < \text{deg}(\mathcal{E}^\vee \otimes \mathcal{F})^{\geq 0} - \text{deg}(\mathcal{Q}^\vee \otimes \mathcal{F})^{\geq 0} \tag{4.2}$$

for any \mathcal{Q} which can occur as a quotient of \mathcal{E} and a proper subbundle of \mathcal{F} . As discussed before, we assume here that the maximum slope of \mathcal{E} is strictly less than the slope of \mathcal{F} , which is permissible by the initial reduction discussed in § 4.1.

Write $\vec{\text{HN}}(\mathcal{E}) = (\vec{\lambda}_j)_{1 \leq j \leq s}$ and $\vec{\text{HN}}(\mathcal{Q}) = (\vec{\mu}_i)_{1 \leq i \leq s'}$ (under the notation introduced in Definition 2.3.3). We will write λ_j for the slope of the vector $\vec{\lambda}_j$, and so on. By Lemma 2.3.4, we see that

$$\begin{aligned} \text{deg}(\mathcal{E}^\vee \otimes \mathcal{Q})^{\geq 0} - \text{deg}(\mathcal{Q}^\vee \otimes \mathcal{Q})^{\geq 0} &= \sum_{\vec{\lambda}_j \preceq \vec{\mu}_i} \vec{\lambda}_j \times \vec{\mu}_i - \sum_{\vec{\mu}_k \preceq \vec{\mu}_i} \vec{\mu}_k \times \vec{\mu}_i \\ &= \sum_i \left(\sum_{\vec{\lambda}_j \preceq \vec{\mu}_i} \vec{\lambda}_j - \sum_{\vec{\mu}_k \preceq \vec{\mu}_i} \vec{\mu}_k \right) \times \vec{\mu}_i \\ &= \sum_i \vec{a}_i \times \vec{\mu}_i, \end{aligned}$$

where

$$\vec{a}_i = \sum_{\vec{\lambda}_j \preceq \vec{\mu}_i} \vec{\lambda}_j - \sum_{\vec{\mu}_k \preceq \vec{\mu}_i} \vec{\mu}_k.$$

Then the left side of (4.2) is equal to two times the shaded area in Figure 6.

Note that for a fixed i , the left endpoint of $\sum_{\vec{\lambda}_j \preceq \vec{\mu}_i} \vec{\lambda}_j$ (whose x -coordinate is $-\text{rank } \mathcal{E}^{\leq \vec{\mu}_j}$) lies to the left of the left endpoint of $\sum_{\vec{\mu}_k \preceq \vec{\mu}_i} \vec{\mu}_k$ by Lemma 4.2.1.




Figure 7. Area interpretation of the right side of (4.2) in the case $s = 3, s' = 4$ with slope vectors $\bar{\mu}_1 \geq \bar{\lambda}_1 > \bar{\mu}_2 > \bar{\mu}_3 \geq \bar{\lambda}_2 > \bar{\mu}_4 \geq \bar{\lambda}_3$.

Now we analyze the right side of (4.2). Let $\text{HN}(\mathcal{F}) = (\bar{\gamma})$. Note that the slope γ of the semistable bundle \mathcal{F} must be at least μ_i for all i , and *strictly* greater than λ_j for all j . Therefore, by similar manipulations to those above, we have

$$\text{RHS of (4.2)} = \left(\sum_j \bar{\lambda}_j - \sum_i \bar{\mu}_i \right) \times \bar{\gamma}. \tag{4.3}$$

We note that the vector $\sum_j \bar{\lambda}_j - \sum_i \bar{\mu}_i$ has a positive x -coordinate because $\text{rank } \mathcal{E} > \text{rank } \mathcal{Q}$, and has slope $\leq \lambda_1 \leq \gamma$ because $\text{HN}(\mathcal{Q})$ lies below $\text{HN}(\mathcal{E})$ by Corollary 4.2.2.

If we align the *left* endpoints of the HN polygons for \mathcal{F} and \mathcal{E} , we then see that the right side of (4.2) is twice the area of the shaded triangle in Figure 7.

Therefore, we need to show that the shaded area in Figure 6 (corresponding to the left-hand side of (4.2)) is smaller than or equal to the shaded area in Figure 7 (corresponding to the right-hand side of (4.2)), with equality if and only if $\mathcal{Q} = \mathcal{F}$. The general idea of the proof is to apply shears to each of the triangles in turn so that their bases become horizontal, making it easier to compare the areas.

First, we apply a shear transformation

$$T = \begin{bmatrix} 1 & 0 \\ -\gamma & 1 \end{bmatrix}$$

to the whole figure so that the vector $\bar{\gamma}$ becomes horizontal. See Figure 8.




Figure 8. Effect of the shear T .

Remark 4.3.1. This shear is nearly equivalent to tensoring by $\mathcal{O}(-\mu(F))$. They are not quite the same because $\mathcal{O}(-\mu(F))$ does not have rank one. But for the purposes of inequality (4.2), tensoring by $\mathcal{O}(-\mu(F))$ would work just as well, as it multiplies both sides by the same factor.

Remark 4.3.2. We make some simple observations about the effect of such a shear:

- (1) The x -coordinates do not change, so the shear does not change the fact that the horizontal length of the polygon for \mathcal{Q} is at most the horizontal length of the polygon for \mathcal{F} , which is strictly less than the horizontal length of the polygon for \mathcal{E} . (This fact is simply a reformulation of the inequalities $\text{rank } \mathcal{E} > \text{rank } \mathcal{F} \geq \text{rank } \mathcal{Q}$.)
- (2) The slopes of all line segments decrease by γ , so the *difference* between slopes does not change. In particular, Corollary 4.2.2 still holds, and we may now assume that $\lambda_j < 0$ and $\mu_i \leq 0$ for all i and j .
- (3) The determinant of T is 1, so applying T preserves all areas. Thus, to prove the original inequality, we only need to prove the inequality for the transformed picture.

We now want to compute the area of each of the individual small triangles corresponding to terms in the sum for the left-hand side. Consider Figure 9.

Here A_i and A_{i-1} are the two endpoints of $T(\vec{\mu}_i)$, and the vertex B_i is the last vertex of the triangle whose area we are considering. By Corollary 4.2.2, B_i lies to the left of A_i . Extend the line $A_i A_{i-1}$ out leftward so that it meets the vertical line through B_i and call the point of intersection C_i .

We claim that C_i must lie below B_i . Indeed, the point A_i lies below the point (call it D_i) of $\text{HN}(\mathcal{E})$ on the vertical line through A_i . Moreover, by Corollary 4.2.2, the slope of the line from B_i to D_i has slope less than or equal to the slope of the line from C_i to A_i . These observations imply that C_i lies below B_i and the distance from C_i to B_i is at least as large as the distance from A_i to D_i .




Figure 9. Effect of the shear T_i .

Let h_i denote the distance from B_i to C_i , and x_i denote the x -coordinate of $\vec{\mu}_i$ (and $T(\vec{\mu}_i)$).

Claim 4.3.3. We have

$$\text{Area}(\Delta A_{i-1} A_i B_i) = \frac{1}{2} h_i x_i.$$

Proof of Claim 4.3.3. To see this, we apply a shear transformation

$$T_i = \begin{bmatrix} 1 & 0 \\ \gamma - \mu_i & 1 \end{bmatrix}.$$

Recall observations (1), (2), and (3) in Remark 4.3.2 that we made about the effect of such a shear.

Let $A'_{i-1}, A'_i, B'_i, C'_i$ be respectively the image of A_{i-1}, A_i, B_i, C_i under T_i (see Figure 9). By observations (1) and (2), we have the following facts:

- The line segment $A'_{i-1} A'_i$ is horizontal of length x_i .
- The line segment $B'_i C'_i$ is vertical of length h_i .

In particular, $\text{Area}(\Delta A'_{i-1} A'_i B'_i) = \frac{1}{2} h_i x_i$. Now the claim follows since $\text{Area}(\Delta A_{i-1} A_i B_i) = \text{Area}(\Delta A'_{i-1} A'_i B'_i)$ by observation (3).

We now let h denote the y -coordinate of the vector $\sum_j T(\vec{\lambda}_j) - \sum_i T(\vec{\mu}_i)$. Note that $h > 0$ since $\text{HN}(\mathcal{Q}) \leq \text{HN}(\mathcal{E})$, $\text{rank}(\mathcal{Q}) < \text{rank}(\mathcal{E})$, and the slopes of $T(\vec{\lambda}_i)$ are negative.

Claim 4.3.4. We have $h_i \leq h$ for all i , with equality achieved for all i only if \mathcal{Q} has only one slope μ_1 and $\mu_1 = \gamma$.

First let us see that Claim 4.3.4 finishes off the proof of (4.2). Since x_i is the x -coordinate of the vector $\vec{\mu}_i$, the sum $\sum_i x_i$ is equal to $\text{rank}(\mathcal{Q})$. On the other hand, the vector $T(\vec{\gamma})$ is a horizontal vector of length $\text{rank}(\mathcal{F})$. Thus by Claim 4.3.3, we find that

$$\begin{aligned} \text{LHS of (4.2)} &= 2 \sum_i \text{Area}(\Delta A_{i-1} A_i B_i) = \sum_i x_i h_i \\ &\leq h \sum_i x_i = h \cdot \text{rank}(\mathcal{Q}) \\ &\leq h \cdot \text{rank}(\mathcal{F}) = h \cdot \|T(\vec{\gamma})\| = \text{RHS of (4.2)}. \end{aligned}$$

Moreover, equality can only hold if $\text{rank}(\mathcal{Q}) = \text{rank}(\mathcal{F})$ and \mathcal{Q} is semistable of slope $\mu_1 = \gamma$, which implies that $\mathcal{Q} = \mathcal{F}$.

Thus, we have reduced to proving Claim 4.3.4.

Proof of Claim 4.3.4. To aid with proving this, we add in a helper horizontal line starting from the left endpoint of $T(\text{HN}(\mathcal{Q}))$ extending out to negative infinity to obtain an augmented polygon $T(\text{HN}(\mathcal{Q}))'$ (see the dashed line in Figure 8). The resulting figure is convex and retains the property described in Corollary 4.2.2.

Let E_i denote the point directly under B_i lying on this polygon (see Figure 9). By convexity, the length h_i is at most the length of the line from B_i to E_i , with equality if and only if the line segment $A_i C_i$ lies on the augmented polygon $T(\text{HN}(\mathcal{Q}))'$.

By the analogue of the property described in Corollary 4.2.2 for the augmented polygon $T(\text{HN}(\mathcal{Q}))'$, as we move left, the vertical distance from $T(\text{HN}(\mathcal{E}))$ to this augmented polygon is strictly increasing. But the distance at the far left when we reach the end of $T(\text{HN}(\mathcal{E}))$ is precisely h . This shows that $|B_i E_i| \leq h$. Equality occurs precisely when B_i is the left endpoint of $T(\text{HN}(\mathcal{E}))$ and $C_i = E_i$ is the point directly below the left endpoint of $T(\text{HN}(\mathcal{E}))$, which happens if and only if the line segment $A_i C_i$ is horizontal. But this occurs precisely when $\mu_j = \gamma$ for all $j \leq i$. Thus, if the equality is achieved for all i , \mathcal{Q} must have only one slope $\mu_1 = \gamma$.

We have finished proving Theorem 4.1.1. □

5. Step two: extensions of semistable bundles

5.1. The key inequality

The bulk of this section is devoted to proving the following result.

Theorem 5.1.1. *Let \mathcal{D} and \mathcal{F} be semistable vector bundles on X such that $\vec{\text{HN}}(\mathcal{D}) = (\vec{\beta})$, $\vec{\text{HN}}(\mathcal{F}) = (\vec{\gamma})$, and $\vec{\beta} \preceq \vec{\gamma}$. Suppose \mathcal{E} is a vector bundle on X such that $\vec{\text{HN}}(\mathcal{E}) \preceq \vec{\text{HN}}(\mathcal{D} \oplus \mathcal{F})$ with the same endpoints. Suppose also that the maximal slope of \mathcal{E} is strictly less than the slope of \mathcal{F} .*

Let \mathcal{K} be a vector bundle on X with the same rank and degree as \mathcal{D} satisfying the following two conditions:

- (i) *The maximal slope of \mathcal{K} is at most the maximal slope of \mathcal{E} .*
- (ii) *\mathcal{K} is not semistable.*




Figure 10. Comparison of the HN vectors for $\text{HN}(\mathcal{K} \oplus \mathcal{F})$ (the top solid lines) and $\text{HN}(\mathcal{E})$ (the bottom solid lines).

Then we have a strict inequality

$$\text{deg}(\mathcal{K}^\vee \otimes \mathcal{E})^{\geq 0} < \text{deg}(\mathcal{K}^\vee \otimes \mathcal{K})^{\geq 0} + \text{deg}(\mathcal{E}^\vee \otimes \mathcal{F})^{\geq 0}. \tag{5.1}$$

Setup and notation. Write

- $\overrightarrow{\text{HN}}(\mathcal{E}) = (\vec{\lambda}_j)_{1 \leq j \leq s}$,
- $\overrightarrow{\text{HN}}(\mathcal{K}) = (\vec{\mu}_i)_{1 \leq i \leq s'}$.

By assumption $\text{HN}(\mathcal{E}) \leq \text{HN}(\mathcal{D} \oplus \mathcal{F})$, so

$$\vec{\lambda}_1 < \vec{\gamma} \quad \text{and} \quad \vec{\lambda}_s \geq \vec{\beta}.$$

In addition, we have $\vec{\mu}_1 \preceq \vec{\lambda}_1$ by assumption (i) of Theorem 5.1.1. On the other hand, assumption (ii) of Theorem 5.1.1 says $\vec{\mu}_{s'} < \vec{\beta}$, which implies $\vec{\mu}_{s'} < \vec{\lambda}_s$ as $\vec{\beta} \preceq \vec{\lambda}_s$.

We will first sketch the proof of Theorem 5.1.1 in geometric terms, because we feel that this conveys a better intuition for the formal argument, whose geometric meaning might be obscured by the algebraic manipulations. Then we will carefully present a thorough algebraic proof.

5.2. Proof sketch of Theorem 5.1.1

Let $O, P_1, \dots, P_{s'+1}$ be the vertices of $\overrightarrow{\text{HN}}(\mathcal{F} \oplus \mathcal{K})$, and $O, Q_1, \dots, Q_{s+1} = P_{s'+1}$ be the vertices of $\overrightarrow{\text{HN}}(\mathcal{E})$. (See Figure 10.)

First, we claim that the right-hand side of inequality (5.1) is twice the area of the convex polygon $OP_1 \dots P_{s'+1}$. Indeed, by Proposition 2.3.5, we see that $\text{deg}(\mathcal{K}^\vee \otimes \mathcal{K})^{\geq 0}$




Figure 11. Geometric interpretation of the left-hand side of (5.1).

is twice the area of the polygon $P_1 P_2 \dots P_{s'+1}$, while

$$\text{deg}(\mathcal{E}^\vee \otimes \mathcal{F}) \geq 0 = \sum_{i=1}^s \vec{\lambda}_i \times \vec{\gamma} = \overrightarrow{OP_{s'+1}} \times \overrightarrow{OP_1}$$

is twice the area of the triangle $\Delta OP_1 P_{s'+1}$. Summing then gives us the area of polygon $OP_1 \dots P_{s'+1}$. (See Figure 11.)

Now consider the left side of inequality (5.1). By Lemma 2.3.4, we have

$$\begin{aligned} \text{deg}(\mathcal{K}^\vee \otimes \mathcal{E}) \geq 0 &= \sum_{\vec{\mu}_i \preceq \vec{\lambda}_j} \vec{\mu}_i \times \vec{\lambda}_j \\ &= \sum_{i=1}^{s'} \vec{\mu}_i \times \left(\sum_{\vec{\lambda}_j \succeq \vec{\mu}_i} \vec{\lambda}_j \right) \\ &= \sum_{i=1}^{s'} \overrightarrow{P_i P_{i+1}} \times \overrightarrow{O Q_{j_i}}, \end{aligned}$$

where Q_{j_i} is some vertex of $\text{HN}(\mathcal{E})$. Now we claim that $\overrightarrow{P_i P_{i+1}} \times \overrightarrow{O Q_{j_i}}$ is at most twice the area of the triangle $OP_i P_{i+1}$, with equality if and only if Q_{j_i} lies on the line segment $P_i P_{i+1}$. Indeed, twice the area of $OP_i P_{i+1}$ is just $\overrightarrow{P_i P_{i+1}} \times \overrightarrow{O P_i}$, so

$$\overrightarrow{P_i P_{i+1}} \times \overrightarrow{O P_i} - \overrightarrow{P_i P_{i+1}} \times \overrightarrow{O Q_{j_i}} = \overrightarrow{P_i P_{i+1}} \times \overrightarrow{P_i Q_{j_i}}$$

is by convexity considerations just twice the area of $\Delta P_i P_{i+1} Q_{j_i}$, which is nonnegative. By convexity, it is 0 if and only if Q_{j_i} lies on the segment $P_i P_{i+1}$. Then we have

$$\begin{aligned} \deg(\mathcal{K}^\vee \otimes \mathcal{E})^{\geq 0} &\leq \sum_{i=1}^{s'} 2 \cdot \text{Area}(\Delta O P_i P_{i+1}) \\ &= 2 \cdot \text{Area}(O P_1 \dots P_{s'+1}) \\ &= \deg(\mathcal{K}^\vee \otimes \mathcal{K})^{\geq 0} + \deg(\mathcal{E}^\vee \otimes \mathcal{F})^{\geq 0}. \end{aligned}$$

To see equality cannot hold, observe that, as by assumption $\overrightarrow{\text{HN}}(\mathcal{E}) \leq \overrightarrow{\text{HN}}(\mathcal{F} \oplus \mathcal{D})$, all Q_{j_i} lie inside the triangle $\Delta O P_1 P_{s'+1}$. But since the largest slope of \mathcal{E} is strictly less than the slope of \mathcal{F} , we see that Q_{j_i} cannot be the point P_1 , so cannot lie on the segment $P_1 P_2$ as \mathcal{K} is not semistable. This gives the theorem. \square

5.3. Formal proof of Theorem 5.1.1

We now carefully go through a formal argument for Theorem 5.1.1. This is essentially just a thorough formalization of § 5.2, so the reader who is already convinced by § 5.2 can skip ahead to the next subsection.

Maintain the notation of § 5.1. By rearranging the vectors $\vec{\mu}_i$ and $\vec{\lambda}_j$ in order of decreasing slopes, we get a sequence of the form

$$\begin{aligned} \vec{\lambda}_1 > \dots > \vec{\lambda}_{j_1} \geq \vec{\mu}_1 > \dots > \vec{\mu}_{i_1} > \vec{\lambda}_{j_1+1} > \dots > \vec{\lambda}_{j_2} \geq \vec{\mu}_{i_1+1} > \dots > \vec{\mu}_{i_2} > \dots \\ > \vec{\lambda}_{j_{-1}+1} > \dots > \vec{\lambda}_{j_i=s} \geq \vec{\mu}_{i_{-1}+1} > \dots > \vec{\mu}_{i=s'}. \end{aligned}$$

For convenience, we set $i_0 = j_0 = 0$. Note that inequalities between $\vec{\lambda}_{j_k}$ and $\vec{\mu}_{i_{k-1}+1}$ are not strict. Let $\vec{\beta} = (r, d)$ and $\vec{\gamma} = (r', d')$.

5.3.1. A simple special case. Let us first consider the case where $t = 1$, which occurs precisely when $\vec{\mu}_i \leq \vec{\lambda}_j$ for all $1 \leq i \leq s'$ and $1 \leq j \leq s$. By Lemma 2.3.4, we compute

$$\begin{aligned} \deg(\mathcal{K}^\vee \otimes \mathcal{E})^{\geq 0} &= \sum_{\vec{\mu}_i \leq \vec{\lambda}_j} \vec{\mu}_i \times \vec{\lambda}_j = \sum_{i,j} \vec{\mu}_i \times \vec{\lambda}_j = \deg(\mathcal{K}^\vee \otimes \mathcal{E}) \\ &= \text{rank}(\mathcal{K}) \deg(\mathcal{E}) - \text{rank}(\mathcal{E}) \deg(\mathcal{K}) \\ &= r(d + d') - (r + r')d = rd' - r'd, \\ \deg(\mathcal{E}^\vee \otimes \mathcal{F})^{\geq 0} &= \sum_{\vec{\lambda}_j \leq \vec{\gamma}} \vec{\lambda}_j \times \vec{\gamma} = \sum_j \vec{\lambda}_j \times \vec{\gamma} = \deg(\mathcal{E}^\vee \otimes \mathcal{F}) \\ &= \text{rank}(\mathcal{E}) \deg(\mathcal{F}) - \text{rank}(\mathcal{F}) \deg(\mathcal{E}) \\ &= (r + r')d' - r'(d + d') = rd' - r'd. \end{aligned}$$

In particular, we have $\deg(\mathcal{K}^\vee \otimes \mathcal{E})^{\geq 0} = \deg(\mathcal{E}^\vee \otimes \mathcal{F})^{\geq 0}$. On the other hand, Proposition 2.3.5 yields $\deg(\mathcal{K}^\vee \otimes \mathcal{K})^{\geq 0} > 0$, as \mathcal{K} is not semistable. Hence we deduce the desired inequality, in this case.

5.3.2. Reduction to simpler slope relations. We now assume that $t \geq 2$. We now reduce to the case where the slope vectors satisfy a simpler ‘intertwining’ relation.

Define, for each index $k = 1, 2, \dots, t$,

$$\begin{aligned} \tilde{\lambda}_k &:= \vec{\lambda}_{j_{k-1}+1} + \dots + \vec{\lambda}_{j_k}, \\ \tilde{\mu}_k &:= \vec{\mu}_{i_{k-1}+1} + \dots + \vec{\mu}_{i_k}. \end{aligned}$$

Then we clearly have

$$\tilde{\lambda}_1 \geq \tilde{\mu}_1 > \tilde{\lambda}_2 \geq \tilde{\mu}_2 > \dots > \tilde{\lambda}_t \geq \tilde{\mu}_t. \tag{5.2}$$

Define two vector bundles $\tilde{\mathcal{K}}$ and $\tilde{\mathcal{E}}$ on \mathcal{X}_C by $\overrightarrow{\text{HN}}(\tilde{\mathcal{K}}) = (\tilde{\mu}_k)_{1 \leq k \leq t}$ and $\overrightarrow{\text{HN}}(\tilde{\mathcal{E}}) = (\tilde{\lambda}_k)_{1 \leq k \leq t}$. Then we have the following properties of $\tilde{\mathcal{K}}$ and $\tilde{\mathcal{E}}$:

- $\text{deg}(\tilde{\mathcal{K}}) = \text{deg}(\mathcal{K}) = d, \text{rank}(\tilde{\mathcal{K}}) = \text{rank}(\mathcal{K}) = r;$
- $\text{deg}(\tilde{\mathcal{E}}) = \text{deg}(\mathcal{E}) = d + d', \text{rank}(\tilde{\mathcal{E}}) = \text{rank}(\mathcal{E}) = r + r';$
- $\tilde{\mathcal{K}}$ is not semistable since $t \geq 2;$
- the maximal slope of $\tilde{\mathcal{K}}$ is less than or equal to the maximal slope of $\tilde{\mathcal{E}}$, as $\tilde{\mu}_1 \leq \tilde{\lambda}_1;$
- the maximal slope of $\tilde{\mathcal{E}}$ is strictly less than the slope of \mathcal{F} , as $\tilde{\lambda}_1 \leq \vec{\lambda}_1 < \vec{\gamma};$
- $\text{HN}(\tilde{\mathcal{E}}) \leq \text{HN}(\mathcal{E})$, and hence $\text{HN}(\tilde{\mathcal{E}}) \leq \text{HN}(\mathcal{D} \oplus \mathcal{F}).$

We claim that it suffices to prove Theorem 5.1.1 for $\tilde{\mathcal{K}}$ and $\tilde{\mathcal{E}}$.

Indeed, $\tilde{\mathcal{K}}$ and $\tilde{\mathcal{E}}$ satisfy the same assumptions made on \mathcal{K} and \mathcal{E} in Theorem 5.1.1. It is also straightforward to check the identities

$$\text{deg}(\mathcal{K}^\vee \otimes \mathcal{E}) \geq 0 = \text{deg}(\tilde{\mathcal{K}}^\vee \otimes \tilde{\mathcal{E}}) \geq 0 \quad \text{and} \quad \text{deg}(\mathcal{E}^\vee \otimes \mathcal{F}) \geq 0 = \text{deg}(\tilde{\mathcal{E}}^\vee \otimes \mathcal{F}) \geq 0$$

using Lemma 2.3.4 and the inequalities in (5.2). In addition, since $\text{HN}(\tilde{\mathcal{K}}) \leq \text{HN}(\mathcal{K})$, Corollary 2.3.6 yields

$$\text{deg}(\tilde{\mathcal{K}}^\vee \otimes \tilde{\mathcal{K}}) \geq 0 \leq \text{deg}(\mathcal{K}^\vee \otimes \mathcal{K}) \geq 0.$$

Hence for the purpose of establishing Theorem 5.1.1, it suffices to prove the inequality

$$\text{deg}(\tilde{\mathcal{K}}^\vee \otimes \tilde{\mathcal{E}}) \geq 0 < \text{deg}(\tilde{\mathcal{K}}^\vee \otimes \tilde{\mathcal{K}}) \geq 0 + \text{deg}(\tilde{\mathcal{E}}^\vee \otimes \mathcal{F}) \geq 0. \tag{5.3}$$

5.3.3. Some algebraic reductions. We now complete the proof of Theorem 5.1.1.

Using Lemma 2.3.4 and the inequalities in (5.2), we find

$$\begin{aligned} \text{deg}(\tilde{\mathcal{K}}^\vee \otimes \tilde{\mathcal{E}}) \geq 0 &= \sum_{i \geq j} \tilde{\mu}_i \times \tilde{\lambda}_j = \sum_{i=1}^t \left(\tilde{\mu}_i \times \sum_{j=1}^i \tilde{\lambda}_j \right) \\ \text{deg}(\tilde{\mathcal{K}}^\vee \otimes \tilde{\mathcal{K}}) \geq 0 &= \sum_{i \geq j} \tilde{\mu}_i \times \tilde{\mu}_j = \sum_{i=1}^t \left(\tilde{\mu}_i \times \sum_{j=1}^i \tilde{\mu}_j \right), \\ \text{deg}(\tilde{\mathcal{E}}^\vee \otimes \mathcal{F}) \geq 0 &= \sum_{j=1}^t \tilde{\lambda}_j \times \vec{\gamma} = \left(\sum_{j=1}^t \tilde{\lambda}_j \right) \times \vec{\gamma}. \end{aligned}$$

Note that $\vec{\gamma} + (\tilde{\mu}_1 + \dots + \tilde{\mu}_t) = \tilde{\lambda}_1 + \dots + \tilde{\lambda}_t$, so we may write

$$\deg(\tilde{\mathcal{E}}^\vee \otimes \mathcal{F})^{\geq 0} = \left(\vec{\gamma} + \sum_{i=1}^t \tilde{\mu}_i \right) \times \vec{\gamma} = \sum_{i=1}^t \tilde{\mu}_i \times \vec{\gamma}.$$

Substituting this above yields

$$\begin{aligned} \deg(\tilde{\mathcal{K}}^\vee \otimes \tilde{\mathcal{K}})^{\geq 0} + \deg(\tilde{\mathcal{E}}^\vee \otimes \mathcal{F})^{\geq 0} &= \sum_{i=1}^t \left[\tilde{\mu}_i \times \left(\vec{\gamma} + \sum_{j=1}^i \tilde{\mu}_j \right) \right] \\ &= \sum_{i=1}^t \left[\tilde{\mu}_i \times \left(\sum_{j=1}^t \tilde{\lambda}_j - \sum_{j=i+1}^t \tilde{\mu}_j \right) \right]. \end{aligned} \tag{5.4}$$

We thus have

$$\begin{aligned} \deg(\tilde{\mathcal{K}}^\vee \otimes \tilde{\mathcal{K}})^{\geq 0} + \deg(\tilde{\mathcal{E}}^\vee \otimes \mathcal{F})^{\geq 0} - \deg(\tilde{\mathcal{K}}^\vee \otimes \tilde{\mathcal{E}})^{\geq 0} &= \sum_{i=1}^t \left[\tilde{\mu}_i \times \left(\sum_{j=1}^t \tilde{\lambda}_j - \sum_{j=i+1}^t \tilde{\mu}_j \right) - \tilde{\mu}_i \times \sum_{j=1}^i \tilde{\lambda}_j \right] \\ &= \sum_{i=1}^t \left[\tilde{\mu}_i \times \left(\sum_{j=i+1}^t \tilde{\lambda}_j - \sum_{j=i+1}^t \tilde{\mu}_j \right) \right] \\ &= \sum_{i=1}^{t-1} \left[\tilde{\mu}_i \times \left(\sum_{j=i+1}^t \tilde{\lambda}_j - \sum_{j=i+1}^t \tilde{\mu}_j \right) \right]. \end{aligned}$$

Hence, to prove (5.3), it suffices to prove

$$0 < \tilde{\mu}_i \times \left(\sum_{j=i+1}^t \tilde{\lambda}_j - \sum_{j=i+1}^t \tilde{\mu}_j \right) \quad \text{for } i = 1, 2, \dots, t-1. \tag{5.5}$$

5.3.4. Completion of the proof. For the conclusion, it will be useful to refer to the geometry of the HN polygons. Let $O, P_1, P_2, \dots, P_{t+1} = P$ be the breakpoints of $\text{HN}(\tilde{\mathcal{K}} \oplus \mathcal{F})$, listed in order of increasing x -coordinates. Similarly, let $O = Q_1, Q_2, \dots, Q_{t+1} = P$ be the breakpoints of $\text{HN}(\tilde{\mathcal{E}})$, also listed in order of increasing x -coordinates. See Figure 12.

Then we have $\tilde{\mu}_i = \overrightarrow{P_i P_{i+1}}$ and $\tilde{\lambda}_j = \overrightarrow{Q_j Q_{j+1}}$, which yields

$$\begin{aligned} \sum_{j=i+1}^t \tilde{\lambda}_j - \sum_{j=i+1}^t \tilde{\mu}_j &= \sum_{j=i+1}^t \overrightarrow{Q_j Q_{j+1}} - \sum_{j=i+1}^t \overrightarrow{P_j P_{j+1}} \\ &= \overrightarrow{Q_{i+1} Q_{t+1}} - \overrightarrow{P_{i+1} P_{t+1}} \\ &= \overrightarrow{Q_{i+1} P_{i+1}}. \end{aligned}$$

Thus we may rewrite (5.5) as

$$0 < \overrightarrow{P_i P_{i+1}} \times \overrightarrow{Q_{i+1} P_{i+1}}. \tag{5.6}$$




Figure 12. Comparison of the HN vectors for $\text{HN}(\tilde{\mathcal{K}} \oplus \mathcal{F})$ (the top solid lines) and $\text{HN}(\tilde{\mathcal{E}})$ (the bottom solid lines).

We will prove inequality (5.6) by dividing into three cases depending on the sign of the x -coordinate of $\overrightarrow{Q_{i+1}P_{i+1}}$. Note that the x -coordinate of $\overrightarrow{P_iP_{i+1}}$ is always positive.

- (1) Suppose first that the x -coordinate of $\overrightarrow{Q_{i+1}P_{i+1}}$ is positive. In this case, (5.6) is equivalent by Lemma 2.3.2 to

$$\begin{aligned} \overrightarrow{P_iP_{i+1}} &< \overrightarrow{Q_{i+1}P_{i+1}} \\ \iff \overrightarrow{P_{i+1}P_i} &> \overrightarrow{P_{i+1}Q_{i+1}}, \end{aligned}$$

which clearly holds by the fact that $\text{HN}(\tilde{\mathcal{E}}) \leq \text{HN}(\tilde{\mathcal{K}} \oplus \mathcal{F})$ and $\tilde{\lambda}_1 < \tilde{\gamma}$.

- (2) Next consider the case where the x -coordinate of $\overrightarrow{Q_{i+1}P_{i+1}}$ is negative. In this case, (5.6) is equivalent to

$$\overrightarrow{P_iP_{i+1}} > \overrightarrow{Q_{i+1}P_{i+1}}.$$

In fact, one has a stronger inequality

$$\overrightarrow{P_{i+1}P_{i+2}} \geq \overrightarrow{Q_{i+1}P_{i+1}}$$

(this is stronger because $\overrightarrow{P_iP_{i+1}} > \overrightarrow{P_{i+1}P_{i+2}}$ by definition), which can be easily seen by the fact that $\text{HN}(\tilde{\mathcal{E}}) \leq \text{HN}(\tilde{\mathcal{K}} \oplus \mathcal{F})$.

- (3) It remains to consider the case where the x -coordinate of $\overrightarrow{Q_{i+1}P_{i+1}}$ is zero. In this case, (5.6) is equivalent to saying that the y -coordinate of $\overrightarrow{Q_{i+1}P_{i+1}}$ is positive. In fact, this is clearly nonnegative by the fact that $\text{HN}(\tilde{\mathcal{E}}) \leq \text{HN}(\tilde{\mathcal{K}} \oplus \mathcal{F})$, so it remains




Figure 13. Geometric interpretation of the left-hand side of (5.4).

to prove that this coordinate is never zero. Suppose for contradiction that the y -coordinate of $\overrightarrow{Q_{i+1}P_{i+1}}$ is zero, which means that $P_{i+1} = Q_{i+1}$. However, since $\text{HN}(\tilde{\mathcal{E}}) \leq \text{HN}(\tilde{\mathcal{D}} \oplus \mathcal{F})$, the only point on $\tilde{\mathcal{K}}$ that can also lie on $\tilde{\mathcal{E}}$ is P_1 . Hence we must have $P_1 = Q_1 = O$, which is clearly a contradiction. □

Remark 5.3.1. We explain a geometric perspective of (5.4), as an example of the translation from the geometric argument sketched in §5.2. As in the proof, let $O = P_0, P_1, P_2, \dots, P_{t+1}$ be the breakpoints of $\text{HN}(\tilde{\mathcal{K}} \oplus \mathcal{F})$, listed in order of increasing x -coordinates. Note that $\vec{\gamma} = \overrightarrow{OP_1}$, $\vec{\mu}_i = \overrightarrow{P_iP_{i+1}}$, and $\vec{\lambda}_1 + \dots + \vec{\lambda}_t = \overrightarrow{OP_{t+1}}$.

By Proposition 2.3.5, $\text{deg}(\tilde{\mathcal{K}}^\vee \otimes \tilde{\mathcal{K}}) \geq 0 = 2 \cdot \text{Area}(P_1P_2 \cdots P_{t+1})$. In addition, we have

$$\text{deg}(\tilde{\mathcal{E}}^\vee \otimes \mathcal{F}) \geq 0 = (\vec{\lambda}_1 + \dots + \vec{\lambda}_t) \times \vec{\gamma} = \overrightarrow{OP_{t+1}} \times \overrightarrow{OP_1} = 2 \cdot \text{Area}(\Delta OP_1P_{t+1}).$$

We may thus write (refer to Figure 13)

$$\begin{aligned} \text{deg}(\tilde{\mathcal{K}}^\vee \otimes \tilde{\mathcal{K}}) \geq 0 + \text{deg}(\tilde{\mathcal{E}}^\vee \otimes \mathcal{F}) \geq 0 &= 2 \cdot \text{Area}(P_1P_2 \cdots P_{t+1}) + 2 \cdot \text{Area}(\Delta OP_1P_{t+1}) \\ &= 2 \cdot \text{Area}(OP_1P_2 \cdots P_{t+1}) \\ &= 2 \sum_{i=1}^t \text{Area}(\Delta OP_iP_{i+1}) = 2 \sum_{i=1}^t \Delta_i, \end{aligned}$$

where we set $\Delta_i = \text{Area}(\Delta OP_iP_{i+1})$.




Figure 14. Geometric interpretation of the right-hand side of (5.4).

On the other hand, for each $i = 1, 2, \dots, t$, we have (refer to Figure 14):

$$\begin{aligned} 2\Delta_i &= \overrightarrow{P_i P_{i+1}} \times \overrightarrow{O P_{i+1}} = \overrightarrow{P_i P_{i+1}} \times (\overrightarrow{O P_{t+1}} - \overrightarrow{P_{i+1} P_{t+1}}) \\ &= \tilde{\mu}_i \times \left(\sum_{j=1}^t \tilde{\lambda}_j - \sum_{j=i+1}^t \tilde{\mu}_j \right). \end{aligned}$$

Hence we deduce (5.4).

5.4. The extension theorem

We can now complete the proof of the extension theorem.

Proof of Theorem 1.1.2. Let the notation and assumptions be as in the theorem. By the reduction in 3.4, it suffices to treat the case where the maximal slope of \mathcal{E} is strictly less than the slope of \mathcal{F}_2 . By Theorem 4.1.1, there exists a surjection $\mathcal{E} \rightarrow \mathcal{F}_2$; equivalently, $\text{Surj}(\mathcal{E}, \mathcal{F}_2)$ is nonempty. It now clearly suffices to show that $\text{Surj}(\mathcal{E}, \mathcal{F}_2)^{\mathcal{F}_1}$ is nonempty.

Let \mathcal{K} be a vector bundle on X , which is isomorphic to the kernel of a surjective map from \mathcal{E} to \mathcal{F}_2 . If \mathcal{K} is not semistable, it clearly satisfies the assumptions in Theorem 5.1.1, so we obtain an inequality

$$\text{deg}(\mathcal{K}^\vee \otimes \mathcal{E}) \geq 0 < \text{deg}(\mathcal{K}^\vee \otimes \mathcal{K}) \geq 0 + \text{deg}(\mathcal{E}^\vee \otimes \mathcal{F}_2) \geq 0.$$

Since $\text{Surj}(\mathcal{E}, \mathcal{F}_2)$ is not empty, we can insert this inequality into the dimension formulas for $\text{Surj}(\mathcal{E}, \mathcal{F}_2)^{\mathcal{K}}$ and $\text{Surj}(\mathcal{E}, \mathcal{F}_2)$ established in Proposition 3.3.6 and Lemma 3.3.14, respectively, yielding the strict inequality

$$\dim \text{Surj}(\mathcal{E}, \mathcal{F}_2)^{\mathcal{K}} < \dim \text{Surj}(\mathcal{E}, \mathcal{F}_2)$$

for any \mathcal{K} satisfying the assumptions of Theorem 5.1.1. Let S denote the set of isomorphism classes of such \mathcal{K} s. Now, arguing as the proof of Theorem 3.3.11, we deduce that the dimension of $X = |\text{Surj}(\mathcal{E}, \mathcal{F}_2)| \setminus |\text{Surj}(\mathcal{E}, \mathcal{F}_2)^{\mathcal{F}_1}|$ satisfies the inequality

$$\dim X \leq \sup_{\mathcal{K} \in S} \deg(\mathcal{K}^\vee \otimes \mathcal{E})^{\geq 0} - \deg(\mathcal{K}^\vee \otimes \mathcal{K})^{\geq 0} < \deg(\mathcal{E}^\vee \otimes \mathcal{F}_2)^{\geq 0},$$

which again yields a contradiction if $|\text{Surj}(\mathcal{E}, \mathcal{F}_2)^{\mathcal{F}_1}|$ is empty.

Remark 5.4.1. There is a more quantitative form of Theorem 1.1.2, which we would like to explain. To describe this result, fix bundles \mathcal{F}_i as in the statement of the theorem, and let $\mathcal{E}xt(\mathcal{F}_2, \mathcal{F}_1)$ be the sheafification of the presheaf sending $S \in \text{Perf}/\text{Spa } F$ to $H^1(\mathcal{X}_S, \mathcal{F}_{2,S}^\vee \otimes \mathcal{F}_{1,S})$. By arguments similar to those in §3.3, one checks that $\mathcal{E}xt(\mathcal{F}_2, \mathcal{F}_1)$ is a locally spatial and partially proper diamond over $\text{Spd } F$, which (by construction) parametrizes all isomorphism classes of extensions of \mathcal{F}_2 by \mathcal{F}_1 . For any given vector bundle \mathcal{E} , let $\mathcal{E}xt(\mathcal{F}_2, \mathcal{F}_1)^\mathcal{E}$ denote the locally closed subdiamond of $\mathcal{E}xt(\mathcal{F}_2, \mathcal{F}_1)$ parametrizing extensions which are isomorphic to \mathcal{E} at all geometric points. According to Theorem 1.1.2, $\mathcal{E}xt(\mathcal{F}_2, \mathcal{F}_1)^\mathcal{E}$ is nonempty if and only if $\text{HN}(\mathcal{E}) \leq \text{HN}(\mathcal{F}_1 \oplus \mathcal{F}_2)$. We then have the following result.

Theorem 5.4.2. *For any \mathcal{E} such that $\text{HN}(\mathcal{E}) \leq \text{HN}(\mathcal{F}_1 \oplus \mathcal{F}_2)$, the diamond $\mathcal{E}xt(\mathcal{F}_2, \mathcal{F}_1)^\mathcal{E}$ is equidimensional of dimension*

$$\deg(\mathcal{F}_2 \otimes \mathcal{F}_1^\vee) - \deg(\mathcal{E} \otimes \mathcal{E}^\vee)^{\geq 0}.$$

Using Proposition 2.3.5, one easily checks that the difference of degrees appearing here is twice the area of the region enclosed between $\text{HN}(\mathcal{E})$ and $\text{HN}(\mathcal{F}_1 \oplus \mathcal{F}_2)$.

Roughly, the idea of the proof is that after rigidifying the extension, we obtain an $\text{Aut}(\mathcal{E})$ -torsor

$$\mathcal{E}xt(\mathcal{F}_2, \mathcal{F}_1)^{\mathcal{E}, \heartsuit} \rightarrow \mathcal{E}xt(\mathcal{F}_2, \mathcal{F}_1)^\mathcal{E},$$

which can also be identified (more or less) with $\text{Surj}(\mathcal{E}, \mathcal{F}_2)^{\mathcal{F}_1, \heartsuit}$. One then argues using dimension theory as before.

6. Multistep filtrations

We now explain how to deduce Theorem 1.1.4 from Theorem 1.1.2; note that the latter is simply the special case $k = 2$ of the former.

As in the assumptions of Theorem 1.1.4, suppose vector bundles $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k$ are given, along with \mathcal{E} such that

$$\text{HN}(\mathcal{E}) \leq \text{HN}(\mathcal{F}_1 \oplus \dots \oplus \mathcal{F}_k).$$

We want to find a filtration on \mathcal{E} whose subquotients are \mathcal{F}_i . We induct on k , with the case $k = 1$ being trivial; we shall need to use the case $k = 2$, which is Theorem 1.1.2, in the inductive step.




Figure 15. Depiction of the HN polygons for \mathcal{E} and $\mathcal{F}_1 \oplus \dots \oplus \mathcal{F}_k$.




Figure 16. Construction of the first step of the filtration.

Proof of Theorem 1.1.4. Consider graphing the HN polygons of \mathcal{E} and $\mathcal{F}_1 \oplus \dots \oplus \mathcal{F}_k$. By assumption, we have

$$\text{HN}(\mathcal{F}_1 \oplus \dots \oplus \mathcal{F}_k) \geq \text{HN}(\mathcal{E}).$$

Since \mathcal{F}_k is semistable and has the largest slope, it lies above $\text{HN}(\mathcal{E})$. (See Figure 15.)

Now take the upper convex hull of $\text{HN}(\mathcal{F}_k)$ and $\text{HN}(\mathcal{E})$. This gives a polygon, which can be written as $\text{HN}(\mathcal{E}')$ for a bundle $\mathcal{E}' \simeq \mathcal{F}_k \oplus \mathcal{E}_{k-1}$ for some \mathcal{E}_{k-1} (shown in blue in Figure 16). We have not yet shown that \mathcal{E}_{k-1} is a subbundle of \mathcal{E} , but we will do so shortly, and then \mathcal{E}_{k-1} will indeed be the first step of the filtration claimed in Theorem 1.1.4.

Note that $\text{HN}(\mathcal{E}_{k-1})$ will consist of a subset $\text{HN}(\mathcal{H}_{k-1})$ of $\text{HN}(\mathcal{E})$ together with a line segment of maximal slope, which is $\text{HN}(\mathcal{G}_{k-1})$ for some \mathcal{G}_{k-1} , connecting \mathcal{F}_k to some vertex of \mathcal{E} . (See Figure 16.) Finally, $\text{HN}(\mathcal{E})$ is the union of $\text{HN}(\mathcal{H}_{k-1})$ and $\text{HN}(\mathcal{L}_{k-1})$ for some \mathcal{L}_{k-1} .

Since all HN filtrations split (Corollary 2.2.11), we know that

$$\mathcal{E}_{k-1} \simeq \mathcal{G}_{k-1} \oplus \mathcal{H}_{k-1} \tag{6.1}$$

and

$$\mathcal{E} \simeq \mathcal{L}_{k-1} \oplus \mathcal{H}_{k-1}. \tag{6.2}$$




Figure 17. The proof is completed by induction.

Now, by the $k = 2$ case, which is a reformulation of Theorem 1.1.4, we know that there exists an exact sequence

$$0 \rightarrow \mathcal{G}_{k-1} \rightarrow \mathcal{L}_{k-1} \rightarrow \mathcal{F}_k \rightarrow 0.$$

Pushing out this sequence with respect to $\mathcal{G}_{k-1} \rightarrow \mathcal{G}_{k-1} \oplus \mathcal{H}_{k-1}$ (i.e., taking the direct sum of the first two terms with $\mathcal{H}_{k-1} \xrightarrow{\text{Id}} \mathcal{H}_{k-1}$), we obtain an exact sequence (using (6.1) and (6.2))

$$0 \rightarrow \mathcal{E}_{k-1} \rightarrow \mathcal{E} \rightarrow \mathcal{F}_k \rightarrow 0.$$

Thus we have constructed a subbundle \mathcal{E}_{k-1} with $\mathcal{E}/\mathcal{E}_{k-1} = \mathcal{F}_k$.

By convexity of the HN polygons, we have $\text{HN}(\mathcal{E}_{k-1}) \leq \text{HN}(\mathcal{F}_1 \oplus \dots \oplus \mathcal{F}_{k-1})$ with the same endpoints. We can then conclude by induction (Figure 17). \square

Acknowledgments. The authors came together around these questions at the 2017 Arizona Winter School, under the umbrella of Kiran Kedlaya’s project group. The authors would like to heartily thank Kiran for giving them this opportunity. They would also like to thank the organizers of the Winter School for creating such a wonderful experience, and they gratefully acknowledge NSF support of the Winter School via grant DMS-1504537.

References

1. A. CARAIANI AND P. SCHOLZE, On the generic part of the cohomology of compact unitary Shimura varieties, *Ann. of Math. (2)*.
2. P. COLMEZ, Espaces de Banach de dimension finie, *J. Inst. Math. Jussieu* **1**(3) (2002), 331–439. MR 1956055.
3. L. FARGUES, Simple connexité des fibres d’une application d’Abel-Jacobi et corps de classe local, Preprint, <https://webusers.imj-prg.fr/~laurent.fargues/cdc.pdf>. To appear in *Annales scientifiques de l’École normale supérieure*.
4. L. FARGUES AND J.-M. FONTAINE, Courbes et fibrés vectoriels en théorie de Hodge p -adique, *Astérisque* (406) (2018), xiii+382. ISSN 0303-1179.
5. L. FARGUES AND J.-M. FONTAINE, Vector bundles on curves and p -adic Hodge theory, in *Automorphic forms and Galois representations. vol. 2*, London Mathematical Society

- Lecture Note Series, Volume 415, pp. 17–104 (Cambridge University Press, Cambridge, 2014). [MR 3444231](#).
6. D. HANSEN, Degenerating vector bundles in p -adic Hodge theory, Preprint, <http://www.math.columbia.edu/~hansen/degen.pdf>.
 7. R. HUBER, *étale cohomology of rigid analytic varieties and adic spaces*, Aspects of Mathematics, E30 (Friedr. Vieweg & Sohn, Braunschweig, 1996). [MR 1734903](#).
 8. K. KEDLAYA, Sheaves, stacks, and shtukas, Preprint, <http://swc.math.arizona.edu/aws/2017/2017KedlayaNotes.pdf>.
 9. K. S. KEDLAYA, Slope filtrations for relative Frobenius, *Astérisque* **I**(319) (2008), 259–301. Représentations p -adiques de groupes p -adiques. I. Représentations galoisiennes et (ϕ, Γ) -modules. [MR 2493220](#).
 10. K. S. KEDLAYA, Noetherian properties of Fargues-Fontaine curves, *Int. Math. Res. Not. IMRN* **2016**(8) (2016), 2544–2567. [MR 3519123](#).
 11. K. S. KEDLAYA AND R. LIU, Relative p -adic Hodge theory: foundations, *Astérisque* (371) (2015), 239. [MR 3379653](#).
 12. A.-C. LE BRAS, Espaces de Banach-Colmez et faisceaux cohérents sur la courbe de Fargues-Fontaine, *Duke Math. J.* **167**(18) (2018), 3455–3532.
 13. P. SCHOLZE, Étale cohomology of diamonds, Preprint, <http://www.math.uni-bonn.de/people/scholze/EtCohDiamonds.pdf>.
 14. The Stacks Project Authors, *Stacks Project*, <https://stacks.math.columbia.edu>, (2018).
 15. P. SCHOLZE AND J. WEINSTEIN, Lectures on p -adic geometry, *Annals of Mathematics Studies* (Princeton University Press). ISBN 9780691202099.
 16. P. SCHOLZE AND J. WEINSTEIN, Moduli of p -divisible groups, *Camb. J. Math.* **1**(2) (2013), 145–237. [MR 3272049](#).
 17. J. WEINSTEIN, $\text{Gal}(\mathbb{Q}_p/\mathbb{Q}_p)$ as a geometric fundamental group, *International Mathematics Research Notices* **2017**(10) (2016), 2964–2997. 1073-7928.