



Universiteit
Leiden

The Netherlands

Counting curves and their rational points

Spelier, P.

Citation

Spelier, P. (2024, June 12). *Counting curves and their rational points*.

Retrieved from <https://hdl.handle.net/1887/3762227>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3762227>

Note: To cite this publication please use the final published version (if applicable).

Chapter 7

Local heights on hyperelliptic curves and quadratic Chabauty

This chapter has already appeared as a preprint [BDRHS24]. This is joint work with Alex Betts, Juanita Duque–Rosero and Sachi Hashimoto.

Abstract. Local heights are arithmetic invariants used in the quadratic Chabauty method for determining the rational points on curves. We present an algorithm to compute these local heights for hyperelliptic curves at odd primes $\ell \neq p$. This algorithm significantly broadens the applicability of quadratic Chabauty to curves which were previously inaccessible due to the presence of non-trivial local heights. We provide numerous examples, including the first quadratic Chabauty computation for a curve with two non-trivial local heights.

Local heights are arithmetic invariants used in the quadratic Chabauty method for determining the rational points on curves. We present an algorithm to compute these local heights for hyperelliptic curves at odd primes $\ell \neq p$. This algorithm significantly broadens the applicability of quadratic Chabauty to curves which were previously inaccessible due to the presence of non-trivial local heights. We provide numerous examples, including the first quadratic Chabauty computation for a curve having two primes with non-trivial local heights.

7.1 Introduction

Smooth projective curves over the rational numbers $X/\mathbb{Q}$ provide a rich terrain for exploring diophantine questions. The quadratic Chabauty method has emerged as a powerful p -adic technique for computing the set of rational points $X(\mathbb{Q})$ on curves of genus $g > 1$ with Jacobian J . This method, introduced in [BD18, BD21], applies when the Mordell–Weil rank of $J(\mathbb{Q})$ equals g and the Picard number of J is at least 2. Combined with the Mordell–Weil sieve, quadratic Chabauty has had much success in determining $X(\mathbb{Q})$ [BDM⁺19, BDM⁺21].

The central object in quadratic Chabauty is a Nekovář p -adic height function h_Z , which depends on a choice of trace zero correspondence $Z \subset X \times X$ invariant under the Rosati involution. The global p -adic height function h_Z can be expressed as a sum of local height functions $h_Z = \sum_{\ell \text{ prime}} h_{Z,\ell}$. When $\ell \neq p$ is a prime of potentially good reduction for X , the local height function $h_{Z,\ell}$ vanishes. The local height at p is a locally analytic function, and while there exist established algorithms and implementations for computing $h_{Z,p}$ [BB12, BDM⁺19, GM23], the situation for the local height $h_{Z,\ell}$ at $\ell \neq p$ stands in stark contrast.

Until now, algorithmic calculations of $h_{Z,\ell}$ have been limited to the case where X is an elliptic curve or J factors as a product of two elliptic curves, in which case $h_{Z,\ell}$ is determined by arithmetic invariants of the regular model of the elliptic curve(s) [Sil88, CPS06, Bia20]. Beyond the case of elliptic and bielliptic curves, strategies for computing $h_{Z,\ell}$ have relied on constructing a regular model for X over $\mathbb{Z}_\ell$. For example, two recent local heights calculations hinged on both computing a regular semistable model for X over $\mathbb{Z}_\ell$ and the existence of an abundance of rational points on X [BDM⁺21].

The value of the local height function $h_{Z,\ell}$ at a point $z \in X(\mathbb{Q}_\ell)$ can be defined in terms of the local Néron–Tate height pairing as the height pairing of two divisors depending on the correspondence Z , the point z , and a chosen basepoint b . While algorithms exist for computing the local height *pairing* given a regular model for X over $\mathbb{Z}_\ell$ [Hol12, Mue14, vBHM20], computing the height *function* $h_{Z,\ell}$ in this way seems to be impractical, given that the defining equations for Z typically have very large degree, see Remark Remark 7.3.1. Consequently, beyond the case of (bi)elliptic curves, quadratic Chabauty has almost exclusively been applied in cases where all local height functions, except those at p , vanish.

This paper introduces an algorithm that marks the first practical method

for computing local heights away from p for a class of curves, outside of the (bi)elliptic case. Our algorithm computes the Nekovář local p -adic heights $h_{Z,\ell}$ for odd primes $\ell \neq p$ on hyperelliptic curves of genus $g > 1$. This algorithm significantly broadens the applicability of quadratic Chabauty to curves previously deemed inaccessible due to potentially having non-trivial local heights, offering a promising avenue for advancing our understanding of rational points on higher genus curves or those with larger conductors (see Theorem I and Corollary N).

We illustrate our algorithm by computing numerous examples of local heights on hyperelliptic curves with diverse reduction types. The computations for these examples use our Magma implementation of the algorithm, available at <https://github.com/sachihashimoto/local-heights>. We revisit several Atkin–Lehner quotients of $X_0(N)$ from [ACKP22], where they were not able to apply quadratic Chabauty due to the presence of potentially non-zero local heights away from p . We also study an Atkin–Lehner quotient of a Shimura curve $X_0(93, 1)/\langle w_{93} \rangle$; rational points on this type of Shimura curve quotient parametrise abelian surfaces with potential quaternionic multiplication. Our algorithm is practical even in high genus: we compute local heights on a genus 7 modular curve. As an application, we carry out the first quadratic Chabauty computation on a curve with more than one prime $\ell \neq p$ with non-trivial local heights, showing the following theorem.

Theorem I (cf. Section 7.7.6). *There are 10 rational points on the curve $X : y^2 = x^6 + 18/5x^4 + 6/5x^3 + 9/5x^2 + 6/5x + 1/5$.*

Our approach to computing local heights is based on a formula from [BD20] which describes $h_{Z,\ell}$ in terms of the action of Z_* on the homology of the reduction graph Γ of a semistable model of X and certain integers $\mathrm{tr}_v(Z)$ attached to the vertices of this graph. The problem, as noted in [BDM⁺21, §3.1], is that it is not a priori clear how to compute the action of Z_* on $H_1(\Gamma, \mathbb{Z})$, especially if X or Z has large genus, many components, or highly unstable reduction. Following a strategy suggested to us by Netan Dogra, we solve this problem using the Coleman–Iovita isomorphism for the curve X (Theorem 7.4.1, [CI10]), which relates the homology of Γ to the de Rham cohomology of X , where the action of Z_* is easier to compute. In order to use the Coleman–Iovita isomorphism, we verify that it commutes with the action of correspondences.

Theorem J (Push-pull compatibility of Coleman–Iovita, cf. Theorem 7.4.14). *Let $f : X \rightarrow X'$ be a finite morphism of smooth projective curves over $\mathbb{C}_\ell$. Then the pullback and pushforward maps on de Rham cohomology are compatible with the Coleman–Iovita isomorphism.*

Making this strategy viable requires several innovations. The first problem is to determine the reduction graph Γ of X over a finite extension of $\mathbb{Q}_\ell$ where X acquires semistable reduction. We do this using the theory of cluster pictures, which allows one to read off the graph Γ for a hyperelliptic curve in odd residue characteristic directly from the valuations of the differences of the roots of $f(x)$ [DDMM23].

The second problem is that, in order to explicitly compute the Coleman–Iovita isomorphism, we need to write down a suitable analytic covering of $X_{\mathbb{C}_\ell}$. We explain how to read off this semistable covering from the cluster picture.

Theorem K (cf. Theorem 7.5.14). *Let X/K be a hyperelliptic curve with split semistable reduction over a finite extension K of $\mathbb{Q}_\ell$, given by an equation $y^2 = f(x)$. Then there exists a semistable covering $\mathfrak{U} = (\tilde{U}_\mathfrak{s}^\pm)_\mathfrak{s}$ of X^{an} indexed by proper clusters $\mathfrak{s}$ in the cluster picture of f , where $\tilde{U}_\mathfrak{s}^\pm$ is as defined in Definition 7.5.11.*

It is interesting to note that this same semistable covering has already appeared in the literature without the explicit link to cluster pictures [Sto19, KK22]. In Section 7.5 we use our semistable covering to write down explicit formulas describing Coleman–Iovita isomorphism in this case.

The end result of this calculation is an ℓ -adic approximation to the action of Z_* on the homology of Γ . However, for applications to quadratic Chabauty, one needs an exact answer. The third innovation that we need is to prove bounds on the possible actions of Z_* and the traces $\text{tr}_v(Z)$ attached to vertices. This allows us to certify that the approximate values coming from our calculation are correct.

Theorem L (Boundedness of norms and traces, cf. Theorem 7.3.5). *Let $Z \subset X \times X$ be an effective correspondence, of degrees d_1 and d_2 over X , respectively. Then the operator norm of Z_* on $H_1(\Gamma, \mathbb{Z})$ and $|\text{tr}_v(Z)|$ have explicit bounds depending only on the degrees d_1 and d_2 .*

In summary, our method for determining the action of Z_* on $H_1(\Gamma, \mathbb{Z})$ proceeds in four steps. Here, K is a field over which X acquires split semistable reduction.


1. Use the cluster picture associated to X to write down a semistable covering of X^{an} and a basis of $H_1(\Gamma, \mathbb{Z})$, using Theorem K.
2. Compute the matrix M_{dR} representing the action of Z_* on $H_{\text{dR}}^1(X/K)$ in some basis of differentials of the second kind.
3. Compute a matrix T which is, to suitably high precision, an ℓ -adic ap-

proximation to the surjection $H_{\text{dR}}^1(X/K) \rightarrow K \otimes_{\mathbb{Z}} H_1(\Gamma, \mathbb{Z})$ coming from the Coleman–Iovita isomorphism.

4. By Theorem J, the matrix $M_{\Gamma} := T \cdot M_{\text{dR}} \cdot T^{-1}$ is then an ℓ -adic approximation of the action of Z_* on $H_1(\Gamma, \mathbb{Z})$, where T^{-1} is a right-inverse of T . Compute the action of Z_* on $H_1(\Gamma, \mathbb{Z})$ by rounding M_{Γ} to the unique integer matrix that satisfies the bounds from Theorem L.

One other consequence of our work is that, by linking together local heights and the combinatorics of the cluster picture, we obtain combinatorial constraints on local heights. This allows us to give various new criteria for when a hyperelliptic curve has all local heights equal to 0, for example:

Proposition M (cf. Example 7.5.20). *Suppose that $X/\mathbb{Q}_{\ell}$ is a genus 2 curve with the cluster picture shown below and that the leading coefficient of X is a unit in $\mathbb{Z}_{\ell}$. Then the local height of any $\mathbb{Q}_{\ell}$ -point on X is 0.*



Corollary N (cf. Section 7.7.2). *The Shimura curve quotient $X_0(93, 1)/\langle \omega_{93} \rangle$ has trivial local height at 31.*

In broad strokes, this overall strategy of computing local heights via the Coleman–Iovita isomorphism is viable for any smooth projective curve X . Our method relies on X being hyperelliptic in only one part. We specialise to hyperelliptic curves in order to use the machinery of cluster pictures to read off the reduction graph Γ and semistable covering of X , i.e. the data of the Berkovich skeleton of X . From this data, and the action of Z_* on $H_{\text{dR}}^1(X/\mathbb{Q})$, we produce the local heights on X . Thus, in order to generalise our algorithm to non-hyperelliptic curves, one would just need a method to determine the Berkovich skeleton. Moreover, our strategy seems promising for determining local heights without using explicit equations for X : if one has some *a priori* way of determining both the Berkovich skeleton and the action of Z_* on $H_{\text{dR}}^1(X/\mathbb{Q})$, say for X a modular curve, then we show that this determines the local heights.

The structure of the paper is as follows. In Section 7.2 we introduce the notions of semistable covering, semistable vertex set, and split semistable model, as well their equivalences. The definition of the local height function $h_{Z, \ell}$ and the local heights formula appear in Section 7.3. We also prove the boundedness of the operator norm and the traces $\text{tr}_v(Z)$ for correspondences $Z \subset X \times X$ in this section. Section 7.4 explains the Coleman–Iovita isomorphism and proves the compatibility with pushforward and pullback. We specialise to the case

of hyperelliptic curves in Section 7.5, and give an explicit description of the Coleman–Iovita isomorphism for these curves in this section. To do this, we introduce the machinery of cluster pictures, use this to construct a semistable covering for hyperelliptic curves, and explicitly describe the Berkovich skeleton of $X_{\mathbb{C}_\ell}^{\text{an}}$. We expand on the details of the explicit computation of the Coleman–Iovita isomorphism in Section 7.6. We show how to represent functions on an annulus, how to compute the action of Z_* on $H_{\text{dR}}^1(X/K)$, and finally how to turn an ℓ -adic matrix for the action of Z_* into an integer matrix, given the bounds from Theorem 7.3.5. Finally, Section 7.7 contains numerous worked examples.

Acknowledgements

We are very grateful to Netan Dogra for the inspiration for the ideas presented in this paper. We are also indebted to Jeroen Sijsling and John Voight for helping us with the implementation of the **Magma** code. We thank Jennifer Balakrishnan for providing numerous comments on an earlier draft of this paper. We are also grateful to Michael Stoll for sharing his code to demonstrate how to apply elliptic curve Chabauty on the example in Section 7.7.6. We would like to thank Matt Baker, Amnon Besser, Edgar Costa, Noam Elkies, Paul Helminck, David Holmes, Steffen Müller, Padmavathi Srinivasan, and Jan Vonk for helpful conversations, Laurent Moret-Bailly for a helpful response to a MathOverflow question, and Ole Ossen for inspiring one of the pictures. L.A.B. is supported by the Simons Collaboration on Arithmetic Geometry, Number Theory, and Computation under grant number 550031. J.D.R. was partially supported by Simons Foundation grant no. 550023. P.S. is supported by NWO grant VI.Vidi.193.006.

7.2 Semistable models, coverings, and vertex sets

Before we begin the paper proper, we first discuss some preliminaries regarding reduction graphs Γ of smooth projective curves X defined over ℓ -adic local fields. There are three equivalent ways that these graphs can be defined: in terms of a *split semistable model* of X ; in terms of a *semistable covering* of the rigid-analytification of X ; or in terms of a *semistable vertex set* inside the Berkovich analytification of X . All three perspectives have their advantages: split semistable models are the simplest conceptually and are the most widely known; semistable coverings are well-adapted to computations; and semistable

vertex sets are combinatorial in nature and well-adapted to proving theoretical results. In this paper, it will be vital to switch freely between these three different perspectives, and so we will begin by recalling the three definitions and their interrelationships, with a particular emphasis on how these notions behave with respect to morphisms of curves.

For the rest of this paper, we fix the following notation. We will fix a prime number ℓ , and will work over a field K which is (for simplicity) either equal to $\mathbb{C}_\ell$ or a finite extension of $\mathbb{Q}_\ell$. We denote the ring of integers, maximal ideal, and residue field of K by $\mathcal{O}_K$, $\mathfrak{m}_K$, and k , respectively. We always normalise the norm and valuation on K so that $|\ell| = \ell^{-1}$ and $v(\ell) = 1$. A curve X over K is always assumed to be smooth, projective, and geometrically integral.

Definition 7.2.1. A 1-dimensional separated scheme $\bar{\mathcal{X}}$ of finite type over k is called a *semistable curve* just when it is geometrically reduced and has at worst ordinary double points as singularities. It is called *strongly semistable* just when additionally every irreducible component is smooth; it is called *split semistable* just when additionally every component is geometrically irreducible, every singular point is k -rational, and the two tangent directions at every singular point are also k -rational.

A *model* of a smooth projective curve¹ X/K is a flat, proper, and finitely presented $\mathcal{O}_K$ -scheme $\mathcal{X}$ together with an isomorphism $\mathcal{X}_K \cong X$ of its generic fibre with X over K . (Any model $\mathcal{X}$ is automatically integral.) A model is called *semistable* (resp. *strongly semistable*, resp. *split semistable*) just when its special fibre is. Note that we do not require our semistable models to be regular, nor do we require them to be minimal.

Remark 7.2.2. The definitions of semistable models used across the literature vary slightly in exactly which properties they require of $\mathcal{X}$. For example, [Liu02, Definition 10.3.27] requires (for $\mathcal{O}_K$ Dedekind) that $\mathcal{X}$ be flat, projective, and normal; [BPR13, Remark 4.2(2)] requires (for $\mathcal{O}_K = \mathbb{C}_\ell$) that $\mathcal{X}$ be flat, proper, and integral; and [MC10, Definition 2.35] just requires that $\mathcal{X}$ be flat and proper. Our definition agrees with all three, as we will now show. Since $\mathcal{O}_K$ is an integral domain, it follows that any flat, finite type $\mathcal{O}_K$ -scheme is automatically finitely presented [RG71, Corollaire 3.47], and also any flat $\mathcal{O}_K$ -scheme with integral generic fibre is integral. For normality, we note that normality is étale local, so semistable curves are normal. And finally a semistable model $\mathcal{X}/\mathcal{O}_K$ is projective, as we can give a relatively very ample Cartier divisor by taking a sum of smooth points that meet every component

¹For us, a *curve* means a geometrically integral separated reduced K -scheme of dimension 1.

of the special fibre with sufficiently high multiplicity.

We will need to say a little about the local structure of a split semistable model $\mathcal{X}$ at a singular point $\bar{x}$. For this, if $a \in \mathfrak{m}_K \setminus \{0\}$, let

$$\mathfrak{S}(a) := \operatorname{Spec}(\mathcal{O}_K[s, t]/(st - a))$$

denote the standard algebraic annulus with parameter a . We write $\bar{0} \in \mathfrak{S}(a)(k)$ for the point defined by the $\mathcal{O}_K$ -algebra homomorphism sending s and t to 0 in k . According to (the proof of) [Sta18a, 0CBY], for any singular point $\bar{x} \in \bar{\mathcal{X}}(k)$ of the special fibre of a split semistable model $\mathcal{X}$, there exists some $a \in \mathfrak{m}_K \setminus \{0\}$, an $\mathcal{O}_K$ -scheme $\mathcal{U}$ with étale maps

$$\mathcal{X} \leftarrow \mathcal{U} \rightarrow \mathfrak{S}(a),$$

and a k -point $\bar{u} \in \bar{\mathcal{U}}(k)$ mapping to both $\bar{x} \in \bar{\mathcal{X}}(k)$ and $\bar{0} \in \bar{\mathfrak{S}}(a)(k)$. We call a diagram as above a *chart* at $\bar{x}$.

The maps appearing in a chart, being étale, induce isomorphisms on completed local rings. Thus, given two charts

$$(\mathcal{X}, \bar{x}) \leftarrow (\mathcal{U}, \bar{u}) \rightarrow (\mathfrak{S}(a), \bar{0}) \quad \text{and} \quad (\mathcal{X}, \bar{x}) \leftarrow (\mathcal{U}', \bar{u}') \rightarrow (\mathfrak{S}(a'), \bar{0}),$$

we find that the completed local rings of $\mathfrak{S}(a)$ and $\mathfrak{S}(a')$ at $\bar{0}$ are isomorphic. This implies that a and a' differ by multiplication by a unit in $\mathcal{O}_K$ [Kat00, Lemma 2.1]. ([Kat00, Lemma 2.1] is only stated when $\mathcal{O}_K$ is noetherian, but the conclusion we need is still valid when $K = \mathbb{C}_\ell$. The final part of the proof of [Kat00, Lemma 2.1] still works, with the caveat that rather than considering e.g. the completed local ring $\hat{\mathcal{O}}_{\mathcal{X}, \bar{x}}$, we instead need to consider the I -adic completion of $\mathcal{O}_{\mathcal{X}, \bar{x}}$, where I is a *finitely generated* ideal of $\mathcal{O}_{\mathcal{X}, \bar{x}}$ whose radical is the maximal ideal. For instance, in the localisation of $\mathcal{O}_K[s, t]/(st - a)$, one could take $I = (s, t, \varpi)$ where $\varpi \in \mathfrak{m}_K \setminus \{0\}$.) In any case, this implies that $v(a) = v(a')$, and so the positive rational number $v(a)$ is independent of the choice of chart at $\bar{x}$.

Definition 7.2.3. The value $v(a)$ is called the *thickness* of the singular point $\bar{x}$.

Remark 7.2.4. When K is a finite extension of $\mathbb{Q}_\ell$, then $\mathcal{X}$ is noetherian, and so a chart $\mathcal{X} \leftarrow \mathcal{U} \rightarrow \mathfrak{S}(a)$ induces an isomorphism

$$\hat{\mathcal{O}}_{\mathcal{X}, \bar{x}} \cong \mathcal{O}_K[[s, t]]/(st - a)$$

on completed local rings [Liu02, Proposition 4.3.26]. Hence our definition of the thickness of $\bar{x}$ agrees with the usual definition [Liu02, Definition 10.3.23]

up to a factor of the ramification degree of $K/\mathbb{Q}_\ell$ (coming from the fact that we normalise our valuation on K differently to [Liu02]). In particular, the thickness of $\bar{x}$ as we have defined it need not be an integer.

Much of the geometry of a split semistable model $\mathcal{X}$ can be captured by a combinatorial invariant known as the *reduction graph* $\Gamma = \Gamma_{\mathcal{X}}$. The vertices of Γ are the irreducible components of the special fibre $\bar{\mathcal{X}}$, and the unoriented edges of Γ are the singular points of $\bar{\mathcal{X}}$. Moreover, one can attach to Γ certain extra data making it into a combinatorial object known as a *metrised complex of k -curves*.

Definition 7.2.5 (Metrised complex of k -curves, [ABBR15, Definition 2.17]). In this paper, a *graph* always means a finite graph in the sense of Serre, i.e. a quadruple $\Gamma = (V(\Gamma), E(\Gamma), \partial_0, (-)^{-1})$ where $V(\Gamma)$ and $E(\Gamma)$ are finite sets and

$$\partial_0: E(\Gamma) \rightarrow V(\Gamma) \quad \text{and} \quad (-)^{-1}: E(\Gamma) \rightarrow E(\Gamma)$$

are functions, where $(-)^{-1}$ is an involution without fixed points. Elements of $V(\Gamma)$ and $E(\Gamma)$ are known as *vertices* and *oriented edges* of Γ , respectively. If $e \in E(\Gamma)$ is an oriented edge, then e^{-1} , $\partial_0(e)$, and $\partial_1(e) := \partial_0(e^{-1})$ are known as the *inverse*, *source*, and *target* of e , respectively. For a vertex v , we write $T_v(\Gamma)$ for the set of oriented edges with source v , and call $T_v(\Gamma)$ the set of *tangent directions* at (or out of) v . The set $E(\Gamma)^+$ of *unoriented edges* of Γ is the quotient of $E(\Gamma)$ by the equivalence relation $e \sim e^{-1}$.

A *metrised graph* is a graph Γ endowed with:

- for each unoriented edge $e \in E(\Gamma)^+$, a positive real number $l(e) \in \mathbb{R}_{>0}$ called the *length* of e .

A *metrised complex of k -curves* is a connected metrised graph Γ endowed with, additionally:

- for each vertex $v \in V(\Gamma)$, a smooth projective curve $\bar{\mathcal{X}}_v/k$, called the *vertex curve* at v ; and
- for each tangent direction $e \in T_v(\Gamma)$, a k -point $\bar{x}_e \in \bar{\mathcal{X}}_v(k)$.

We require that the points $\bar{x}_e$ for different tangent directions e are distinct.

Remark 7.2.6. The elements $\bar{x}_e$ appearing as part of the data of a metrised complex of k -curves are not relevant for our purposes in this paper, so we will mostly permit ourselves to omit mention of them in definitions and proofs.

To make the reduction graph Γ of a split semistable model into a metrised complex of k -curves, we adopt the following conventions.

- $V(\Gamma)$ is the set of irreducible components of the normalisation $\bar{\mathcal{X}}^\sim$ of the special fibre $\bar{\mathcal{X}}$. For a vertex $v \in V(\Gamma)$, we write $\bar{\mathcal{X}}_v$ for the component of $\bar{\mathcal{X}}^\sim$ to which it corresponds.
- $E(\Gamma)$ is the set of k -points of $\bar{\mathcal{X}}^\sim$ lying over singular points in $\bar{\mathcal{X}}$. For an oriented edge $e \in E(\Gamma)$, we write $\bar{x}_e \in \bar{\mathcal{X}}^\sim(k)$ for the point to which it corresponds.
- For an oriented edge $e \in E(\Gamma)$, $\partial_0(e)$ is the vertex such that $\bar{\mathcal{X}}_{\partial_0(e)} \ni \bar{x}_e$, and e^{-1} is the oriented edge different from e such that $\bar{x}_{e^{-1}}$ and $\bar{x}_e$ lie over the same singular point of $\bar{\mathcal{X}}$.
- The unoriented edges $e \in E(\Gamma)^+$ correspond to singular points $\bar{x}_e$ of $\bar{\mathcal{X}}$. The *length* of e is defined to be the thickness of $\bar{x}_e$ inside $\bar{\mathcal{X}}$ (Definition 7.2.3).
- The points $\bar{x}_e \in \bar{\mathcal{X}}_v(k)$ attached to tangent directions e at v are the elements above.

We easily verify that the construction of a metrised complex of k -curves to a split semistable model of a curve is compatible with base change, as per the following lemma.

Lemma 7.2.7. *Let K'/K be an extension of fields, each of which is either $\mathbb{C}_\ell$ or a finite extension of $\mathbb{Q}_\ell$, with residue field extension k'/k . Let $\mathcal{X}/\mathcal{O}_K$ be a split semistable model of a curve X/K , with reduction graph Γ (viewed as a metrised complex of k -curves). Then $\mathcal{X}_{\mathcal{O}_{K'}}$ is a split semistable model of $X_{K'}$, whose reduction graph Γ' is the metrised complex of k' -curves obtained from Γ by base-changing all the curves attached to vertices from k to k' , and leaving the underlying graph, metric and maps unchanged.*

Proof. This is clear. We remark that the fact that Γ and Γ' have the same metric is a consequence of our choice of normalisation of the valuations on K and K' , which ensures that the thickness of singular points of $\bar{\mathcal{X}}$ is unchanged upon base-change to $\mathcal{O}_{K'}$. $\square$

7.2.1 Semistable models, coverings and vertex sets

While semistable models of a curve X are well-behaved theoretically, they are rather difficult to perform explicit computations with – indeed, even writing down a semistable model of X is difficult computationally. Instead, a central part of our approach in this paper will be to replace semistable models with certain analytic data attached to X , equivalent to a choice of semistable model,

but much more amenable to direct computation. We will work with two kinds of analytic data. Firstly, in [Col89, MC10, CM88], Coleman defines the notion of a *semistable covering* of the rigid analytification X^{an} , and shows that (with some small caveats) semistable coverings correspond to semistable models. Secondly, in the case $K = \mathbb{C}_\ell$, Baker, Payne, and Rabinoff define the notion of a *semistable vertex set* inside the Berkovich analytification X^{an} of X , and again show that semistable vertex sets correspond to semistable models [BPR13]. We will presently recall the definitions of these semistable coverings and semistable vertex sets, but before we do so, we owe the reader a brief remark regarding our use of both rigid and Berkovich geometry.

For any complete valued extension K of $\mathbb{Q}_\ell$, the category of separated Berkovich strictly K -analytic spaces embeds as a full subcategory of the category of rigid K -analytic spaces [Ber90, Proposition 3.3.1], and most of the key notions of Berkovich and rigid geometry (finite maps, étale maps, ...) correspond under this embedding [Ber90, §3.3]. Every analytic space we consider will be a separated Berkovich strictly K -analytic space, so we will switch freely between the Berkovich and rigid perspectives.

Thus, if X/K is a smooth projective curve, we will feel free to use the notation X^{an} for both the rigid analytification and the Berkovich analytification of X . When we write $|X^{\text{an}}|$, we always mean the underlying topological space of X^{an} as a Berkovich space; the underlying set of X^{an} as a rigid space is the subset $|X^{\text{an}}|_{\text{rg}} \subset |X^{\text{an}}|$ of *rigid points*, i.e. points $x \in |X^{\text{an}}|$ whose completed residue field $\mathcal{H}(x)$ is a finite extension of K .

7.2.1.1 Semistable coverings

The first alternative perspective we will use is that split semistable models of a smooth projective curve X/K are equivalent to analytic coverings of X by rigid spaces of a certain kind. This perspective goes back to Coleman [Col89].

Definition 7.2.8 ([MC10, §2B]). Let W be a 1-dimensional smooth rigid space. We say that W is a *wide open* just when there exist affinoid subdomains $W_0 \subset W_1 \subset W$ such that:

- $W \setminus W_0$ is a disjoint union of finitely many open annuli;
- W_0 is relatively compact in W_1 (see [BGR84, §9.6.2]); and
- W_1 meets each component of $W \setminus W_0$ in a semi-open annulus.

The affinoid subdomains W_0 and W_1 are not part of the data of a wide open. We refer to W_0 as an *underlying affinoid* of W ; a pair (W, W_0) of a wide

open W and an underlying affinoid W_0 is called a *wide open pair*.

A *basic wide open* (resp. *strongly basic wide open*) is a pair $W = (W, W_0)$ consisting of a connected wide open W and an underlying affinoid W_0 , such that:

- the supremum seminorm of any element of $\mathcal{O}(W_0)$ is the norm of an element of K ;
- the canonical reduction of W_0 is an irreducible split² semistable curve (resp. a smooth curve); and
- the components of $W \setminus W_0$ are isomorphic over K to open annuli of inner radius 1.

(The first condition is equivalent to requiring that $\mathcal{O}(W_0)^\circ / \mathfrak{m}_K \mathcal{O}(W_0)^\circ$ is a reduced ring [BGR84, Theorem 6.4.3/1 & Proposition 6.4.3/4], where $\mathcal{O}(W_0)^\circ$ is the ring of power-bounded functions on W_0 . Then $\text{Spec}(\mathcal{O}(W_0)^\circ / \mathfrak{m}_K \mathcal{O}(W_0)^\circ)$ is the canonical reduction of W_0 .)

Basic wide opens arise as complements of discs in smooth projective curves. Let X/K be a smooth projective curve and let $\mathcal{X}/\mathcal{O}_K$ be a split semistable model whose special fibre is irreducible. Let $\bar{x}_1, \dots, \bar{x}_n$ be distinct smooth k -points of the special fibre of $\mathcal{X}$ for $n \geq 1$, with associated residue discs $D_i^\circ :=]\bar{x}_i[\subset X^{\text{an}}$. For $1 \leq i \leq n$, let $D_i^c \subset D_i^\circ$ be a closed subdisc. Then $W := X^{\text{an}} \setminus \bigcup_i D_i^c$ and $W_0 := X^{\text{an}} \setminus \bigcup_i D_i^\circ$ form a basic wide open (W, W_0) , and the canonical reduction of W_0 is canonically isomorphic to $\mathcal{X} \setminus \{\bar{x}_1, \dots, \bar{x}_n\}$. In particular, W is strongly basic if the model $\mathcal{X}$ was smooth. See [MC10, Proposition 2.21 & Corollary 2.23].

If $W = (W, W_0)$ is a basic wide open, then we refer to the components of $W \setminus W_0$ as the *bounding annuli* of W (they are sometimes called the *annulus ends* of W in the literature). Also, for any singular point $\bar{x}$ of the canonical reduction of W_0 , its residue class $]\bar{x}[$ is an open annulus [MC10, Proposition 2.10]. We refer to such annuli as *internal annuli* of W . We fix some terminology we will use when working with these annuli.

Definition 7.2.9. Let A be an open annulus over K , i.e. a rigid space isomorphic over K to the standard open annulus

$$A(r_1, r_2) = \{x : r_1 < |x| < r_2\}$$

for some $r_1 < r_2 \in \sqrt{|K^\times|}$. The quantity $\log_\ell(r_2/r_1) \in \mathbb{Q}$ is called the

²The definition in [MC10, Definition 2.35] does not explicitly use the word “split”, but their definition of “ordinary double point” includes the requirement that it should be a split node [MC10, Definition 2.9].

*width*³ of A , and is an isomorphism invariant (follows from [BPR13, Proposition 2.2(2)]).

If $\mathcal{O}(A)_1^\times$ denotes the group of invertible analytic functions f such that $|1 - f(x)| < 1$ for all points x of A , then the group

$$\mathcal{O}(A)^\times / K^\times \mathcal{O}(A)_1^\times$$

is infinite cyclic [BPR13, Proposition 2.2(1)]. An *orientation* of A is a choice of generator of this group (so there are two possible orientations of A). A *parameter* on an oriented annulus A is an element $t \in \mathcal{O}(A)^\times$ which maps to the chosen generator of $\mathcal{O}(A)^\times / K^\times \mathcal{O}(A)_1^\times$. The map $A \rightarrow \mathbb{G}_m^{\text{an}}$ induced by a parameter t is an isomorphism onto a standard annulus $A(r_1, r_2)$ [BPR13, Proposition 2.2(2)].

Remark 7.2.10. The bounding annuli A_i of a wide open pair (W, W_0) are oriented in a canonical way. Namely, if t is a parameter on A_i inducing an isomorphism $A \xrightarrow{\sim} A(r_1, r_2)$, then t maps $W_1 \cap A_i$ isomorphically onto a semi-open annulus, either $A(r_1, r']$ or $A[r', r_2)$ for some $r' \in (r_1, r_2)$. These two possibilities correspond to the two possible orientations of A_i ; we will always orient A_i with the orientation where $W_1 \cap A_i$ is identified with the outer annulus $A[r', r_2)$. This is the same convention as [Col89, Corollary 3.7a].

Example 7.2.11. Suppose that we are given closed discs $D^c, D_1^c, \dots, D_n^c$ inside $\mathbb{A}_K^{1, \text{an}}$, contained inside open discs $D^o, D_1^o, \dots, D_n^o$, respectively. Suppose that $D_i^o \subset D^c$ for all i and the D_i^o are pairwise disjoint. Then the domain $U = D^o \setminus \bigcup_{i=1}^n D_i^c$ is wide open, with underlying affinoid $U_0 = D^c \setminus \bigcup_{i=1}^n D_i^o$. The wide open pair (U, U_0) has $n + 1$ bounding annuli, namely $A = D^o \setminus D^c$ and $A_i = D_i^o \setminus D_i^c$ for $1 \leq i \leq n$. The orientation on each A_i is the standard one, i.e. a parameter is $t - \alpha_i$ where t is the standard coordinate on $\mathbb{A}_K^{1, \text{an}}$ and α_i is a centre of D_i^c . The orientation on A , however, is the *opposite* of the standard one, i.e. a parameter on A is $(t - \alpha)^{-1}$ where α is a centre of D^c .

One can study a curve X/K with split semistable reduction by taking an analytic cover by wide opens which intersect each other in a well-behaved manner. Such coverings are known as semistable coverings.

Definition 7.2.12 ([MC10, §2C]). Let X/K be a smooth projective curve. A (strongly) semistable covering of X^{an} is a finite admissible covering $\mathfrak{U} = (W_v)_{v \in V}$ of X^{an} by (strongly) basic wide opens $W_v = (W_v, W_{v,0})$ (as in, the W_v form an admissible covering) such that every pairwise intersection $W_v \cap W_{v'}$ is a union of bounding annuli of W_v , and every triplewise intersection $W_v \cap$

³Also called the *modulus* of A in [BPR13].

$\bar{\mathcal{X}} \setminus \bigcup_{v' \neq v} \bar{\mathcal{X}}_{v'}$ is the complement of the other irreducible components of the special fibre, then the residue class $W_{v,0} =]\bar{\mathcal{X}}_{v,0}[$ is an underlying affinoid, making $W_v = (W_v, W_{v,0})$ into a basic wide open. The bounding annuli (resp. internal annuli) of W_v are the residue classes of the points where $\bar{\mathcal{X}}_v$ intersects other components of $\bar{\mathcal{X}}$ (resp. intersects itself). The wide opens W_v for $\bar{\mathcal{X}}_v$ running over irreducible components of $\bar{\mathcal{X}}$, form a semistable covering of X^{an} .

Theorem 7.2.13 ([MC10, Theorem 2.36]). *Let K be either $\mathbb{C}_\ell$ or a finite extension of $\mathbb{Q}_\ell$, and let X/K be a smooth projective curve. Then the above construction sets up a bijective correspondence between split (strongly) semistable models of X whose special fibre has at least two irreducible components and (strongly) semistable coverings of X^{an} .*

One can read off the reduction graph of a split semistable model $\mathcal{X}$ from the corresponding semistable covering. For this, we attach a metrised complex of k -curves to any semistable covering as follows.

Definition 7.2.14 ([CI10, §3.5.1]). Suppose that $\mathfrak{U} = (W_v)_{v \in V}$ is a semistable covering of X^{an} indexed by a set V . We define a metrised complex of k -curves $\Gamma = \Gamma_{\mathfrak{U}}$ by:

- $V(\Gamma) = V$ is the indexing set V .
- $E(\Gamma)$ is the set of oriented open annuli $A \subset X^{\text{an}}$ which are either a bounding annulus or an internal annulus of some W_v . For $e \in E(\Gamma)$, we write A_e for the corresponding oriented open annulus.
- For an oriented edge $e \in E(\Gamma)$, we define $\partial_0(e)$ to be the unique vertex for which A_e is either a bounding annulus of W_v , equipped with the orientation described in Remark 7.2.10, or an internal annulus of W_v . We define e^{-1} to be the edge for which $A_{e^{-1}}$ is equal to A_e with the opposite orientation. (If A_e is a bounding annulus of W_v , then it is a component of an intersection $W_v \cap W_{v'}$ for $v' \neq v$, and then $A_{e^{-1}}$ is a bounding annulus of $W_{v'}$.)
- For an unoriented edge $e \in E(\Gamma)^+$, we define $l(e)$ to be the width of the annulus A_e .
- For a vertex $v \in V(\Gamma)$, the canonical reduction of the underlying affinoid $W_{v,0}$ is a reduced k -curve. We define $\bar{\mathcal{X}}_v$ to be the smooth compactification of its normalisation.

This construction recovers the reduction graph of a split semistable model, as we now prove carefully.

Lemma 7.2.15. *Let $\mathfrak{U}$ be a semistable covering of X^{an} , corresponding to the split semistable model $\mathcal{X}$ under Theorem 7.2.13. Then we have a canonical isomorphism*

$$\Gamma_{\mathfrak{U}} \cong \Gamma_{\mathcal{X}}$$

of metrised complexes of k -curves.

Proof. It is clear from the definition that there are canonical bijections

$$V(\Gamma_{\mathfrak{U}}) \cong V(\Gamma_{\mathcal{X}}) \quad \text{and} \quad E(\Gamma_{\mathfrak{U}})^+ \cong E(\Gamma_{\mathcal{X}})^+$$

between sets of vertices and unoriented edges, preserving incidence. There are three points which require some explanation:

1. How does one obtain a *canonical* bijection $E(\Gamma_{\mathfrak{U}}) \cong E(\Gamma_{\mathcal{X}})$ between oriented edges?
2. Why do the metrics on $\Gamma_{\mathfrak{U}}$ and $\Gamma_{\mathcal{X}}$ agree?
3. Why do the k -curves attached to vertices agree?

For the first point, let $\bar{x} \in \bar{\mathcal{X}}(k)$ be a singular point, with residue class $]\bar{x}[$. Let us write $\|\cdot\|$ for the sup norm on $\mathcal{O}(]\bar{x}[$, and $\mathcal{O}^\circ(]\bar{x}[) \subset \mathcal{O}(]\bar{x}[$ for the subring of elements of sup norm ≤ 1 . It follows from [MC10, Proposition 2.10] (or the argument below) that $]\bar{x}[$ is isomorphic to $A(|a|, 1)$ for some $a \in \mathfrak{m}_K \setminus \{0\}$, and so any parameter t on $]\bar{x}[$ can be rescaled so that $\|t\| = 1$. We let $s = at^{-1}$, so $\|s\| = 1$ also. According to [MC10, Lemma 2.8], the reduction of $\mathcal{O}^\circ(]\bar{x}[)$ is canonically isomorphic to $\hat{\mathcal{O}}_{\bar{\mathcal{X}}, \bar{x}}$. So, the reductions of s and t determine elements of $\hat{\mathcal{O}}_{\bar{\mathcal{X}}, \bar{x}}$ which generate the maximal ideal and whose product is zero (so $\hat{\mathcal{O}}_{\bar{\mathcal{X}}, \bar{x}} = k[[\bar{s}, \bar{t}]]/(\bar{s}\bar{t})$). The derivations $\frac{\partial}{\partial \bar{s}}$ and $\frac{\partial}{\partial \bar{t}}$ are then two tangent vectors to $\bar{\mathcal{X}}$ at $\bar{x}$ which span the two tangent directions. Changing t by an element of $K^\times \mathcal{O}(]\bar{x}[)_1^\times$ only changes the tangent vector $\frac{\partial}{\partial \bar{t}}$ by a scalar, and so we have described a canonical bijection between the two orientations on the annulus $]\bar{x}[$ and the two tangent directions to $\bar{\mathcal{X}}$ at $\bar{x}$ (t corresponds to $\frac{\partial}{\partial \bar{t}}$).

For the second point, choose a chart

$$\mathcal{X} \leftarrow \mathcal{U} \rightarrow \mathfrak{S}(a)$$

for some a . The choice of chart induces isomorphisms

$$]\bar{x}[\xleftarrow{\sim}]\bar{u}[\xrightarrow{\sim}]\bar{0}[= A(|a|, 1)$$

on the tubes of $\bar{x}$, $\bar{u}$, and $\bar{0}$ inside the formal completions of $\mathcal{X}$, $\mathcal{U}$, and $\mathfrak{S}(a)$, respectively. So the width of $]\bar{x}[$ is equal to $v(a)$, which is the thickness of $\bar{x}$, and so the metrics on $\Gamma_{\mathfrak{U}}$ and $\Gamma_{\mathcal{X}}$ agree.

For the third point, for any vertex v , the canonical reduction of $W_{v,0} = \text{red}^{-1}(\mathcal{X}_v^\circ)$ is $\mathcal{X}_v^\circ$, cf. the proof of [MC10, Proposition 2.36]. So the smooth compactification of the normalisation of $\bar{W}_{v,0}$ is $\mathcal{X}_v$. $\square$

In the proof, we used the following lemma, which is presumably well-known, but we could not find a reference in the literature.

Lemma 7.2.16. *Let $\mathcal{U} \rightarrow \mathcal{U}'$ be an étale morphism of reduced, flat, locally finitely presented $\mathcal{O}_K$ -schemes, and let $\bar{u} \in \bar{\mathcal{U}}(k)$ be a k -point on the special fibre of $\mathcal{U}$ mapping to a point $\bar{u}' \in \bar{\mathcal{U}}'(k)$. Then the induced map*

$$]\bar{u}[\rightarrow]\bar{u}'[$$

on residue classes is an isomorphism of rigid analytic spaces.

Proof. Shrinking $\mathcal{U}'$ and $\mathcal{U}$, we may assume that $\mathcal{U}' = \text{Spec}(A)$ and $\mathcal{U} = \text{Spec}(B)$ are affine, and that $\bar{u}$ is the unique point in the fibre above $\bar{u}'$ in $\bar{\mathcal{U}} \rightarrow \bar{\mathcal{U}}'$. Choose $f_1, \dots, f_n \in A$ whose reductions generate the ideal $\mathfrak{m}_{\bar{u}'}$ defining $\bar{u}'$. The residue class $]\bar{u}'[$ of $\bar{u}'$ is then the union of the affinoid rigid spaces $\text{Sp}(K \otimes_{\mathcal{O}_K} \hat{A}\langle \ell^{-1/m} f_1, \dots, \ell^{-1/m} f_n \rangle)$, where $\hat{A}\langle \ell^{-1/m} f_1, \dots, \ell^{-1/m} f_n \rangle$ is the ℓ -adic completion of the algebra

$$A[\ell^{-1/m} f_1, \dots, \ell^{-1/m} f_n] := A[z_1, \dots, z_n] / (\ell z_1 - f_1^m, \dots, \ell z_n - f_n^m).$$

Since the reductions of $f_1, \dots, f_m$ also generate the ideal in B defining $\bar{u}$, a similar description holds for the residue class of $\bar{u}$.

Now, for any positive integers m and r , the map

$$A[\ell^{-1/m} f_1, \dots, \ell^{-1/m} f_n] / \ell^r \rightarrow B[\ell^{-1/m} f_1, \dots, \ell^{-1/m} f_n] / \ell^r \quad (\star)$$

is étale (it is a base-change of the map $A \rightarrow B$), and its reduction modulo the ideal generated by $f_1, \dots, f_n$ and $\mathfrak{m}_K$ is an isomorphism (the identity on $k[z_1, \dots, z_n]$). Since this ideal consists of nilpotent elements, this implies that $(\star)$ itself is an isomorphism [Gro67, Théorème 18.1.2]. Taking the inverse limit over r shows that the map $\hat{A}\langle \ell^{-1/m} f_1, \dots, \ell^{-1/m} f_n \rangle \rightarrow \hat{B}\langle \ell^{-1/m} f_1, \dots, \ell^{-1/m} f_n \rangle$ is an isomorphism, and so $]\bar{u}[\rightarrow]\bar{u}'[$ is an isomorphism as claimed. $\square$

7.2.1.2 Semistable vertex sets

For this part, we specialise to the case that $K = \mathbb{C}_\ell$. For a smooth projective curve $X/\mathbb{C}_\ell$, write X^{an} for the Berkovich analytification of X , $|X^{\text{an}}|$ for its

underlying topological space, and $|X^{\text{an}}|_{\text{II}} \subset |X^{\text{an}}|$ for the set of type II points (see [Ber90, 1.4.4] for a discussion of types of points, or [Bak08, §2] for a more accessible discussion in the case $X = \mathbb{P}^1$).

A *semistable vertex set* in X^{an} is a finite subset $V \subset |X^{\text{an}}|_{\text{II}}$ such that $X^{\text{an}} \setminus V$ is a disjoint union of open balls and finitely many open annuli [BPR13, Definition 3.1]. These correspond to semistable models of X [BPR13, §4]. If $\mathcal{X}$ is a semistable model of X , then there is an associated *reduction map*

$$\text{red}: |X^{\text{an}}| \rightarrow |\bar{\mathcal{X}}|, \quad (7.2.1.2)$$

where $|\bar{\mathcal{X}}|$ denotes the underlying topological space of the special fibre. This extends the reduction map (7.2.1.1) on rigid points. The preimage of any generic point of $\bar{\mathcal{X}}$ under the reduction map is a single type II point of X^{an} , and the set of all these points as we range over generic points of $\bar{\mathcal{X}}$ is a semistable vertex set [BPR13, Corollary 4.7 & Remark 4.2(2)].

Theorem 7.2.17 ([BPR13, Theorem 4.11]). *Let $X/\mathbb{C}_\ell$ be a smooth projective curve. Then the above construction sets up a bijective correspondence between semistable models of X and semistable vertex sets in X^{an} .*

Again, one can also read off the reduction graph of a semistable model $\mathcal{X}$ from the corresponding semistable vertex set. For this, it is convenient for us to view graphs as metric spaces in the natural way (cf. [ABBR15, Definition 2.2]). That is, we can equivalently describe a metrised graph as a compact metric space⁴ Γ together with a distinguished finite set $V \subset \Gamma$ of *vertices* such that $\Gamma \setminus V$ is isometric to a finite disjoint union of open intervals. The set $T_v(\Gamma)$ of oriented edges with source v is then identified with the set *tangent directions* at v : of germs of isometric embeddings $[0, \epsilon) \rightarrow \Gamma$ taking 0 to v .

To produce a graph out of a semistable vertex set, recall that if $A(r_1, r_2)$ is the standard annulus of inner and outer radii r_1 and r_2 , then we define a map

$$\sigma: (0, \log_\ell(r_2/r_1)) \rightarrow |A(r_1, r_2)|$$

by sending s to the Gauss point of the open ball centred on 0 with radius $r_2\ell^{-s}$ [BPR13, §2.3]. The image of σ is called the *skeleton* $\text{sk}(A(r_1, r_2))$ of $A(r_1, r_2)$. More generally, if A is an open annulus, then we may choose a parameter t defining an isomorphism $t: A \xrightarrow{\sim} A(r_1, r_2)$ for some $r_1 < r_2$, and then define the skeleton of A to be $\text{sk}(A) := t^{-1}\text{sk}(A(r_1, r_2))$. The skeleton of A is independent of the choice of parameter t and is homeomorphic to the open interval $(0, \log_\ell(r_2/r_1))$. If A is oriented, then $\text{sk}(A)$ is also oriented.

⁴Our convention is that all metric spaces are length metric spaces, and the induced metric on a subspace means the induced length metric.

The *skeleton* Γ associated to a semistable vertex set $V \subset |X^{\text{an}}|_{\text{II}}$ is the union of V and the skeleta of the annuli in $|X^{\text{an}}| \setminus V$ [BPR13, Definition 3.3]. For any annulus A in $|X^{\text{an}}| \setminus V$, its skeleton $\text{sk}(A)$ is an open interval inside Γ whose closure is either a closed interval connecting two different elements in V , or a circle connecting an element in V to itself [BPR13, Lemma 3.2]. In this way, the skeleton Γ is a topological graph with vertex-set V .

In fact, the skeleton Γ is a metrised complex of k -curves in a natural way. If we write $\mathbb{H}_o(X) \subset |X^{\text{an}}|$ for the set of points of types II and III, then there is a canonical length metric on $\mathbb{H}_o(X)$ constructed by Baker–Payne–Rabinoff [BPR13, §5.3], and the restriction of this metric to Γ makes it into a metrised graph. This metric is characterised by the fact that the map σ tracing out the skeleton of an annulus is an isometry. Additionally, any element v of the semistable vertex set is a type II point of X^{an} , which means that the completed residue field $\mathcal{H}(v)$ is a complete valued extension of K whose residue field $\tilde{\mathcal{H}}(v)$ is a finitely generated extension of k of transcendence degree 1. We write $\bar{\mathcal{X}}_v$ for the unique smooth projective curve over k with function field $\tilde{\mathcal{H}}(v)$. Attaching to each point $v \in V$ the curve $\bar{\mathcal{X}}_v$ above makes the skeleton Γ into a metrised complex of curves [ABBR15, 3.22].

Remark 7.2.18. The perspective of semistable vertex sets makes it eminently clear that the reduction graph Γ associated to a semistable model of a curve $X/\mathbb{C}_\ell$ is independent of the choice of model Γ , up to certain simple operations. Indeed, let $V \subset |X^{\text{an}}|_{\text{II}}$ be a semistable vertex set, and let $v \in |X^{\text{an}}|_{\text{II}} \setminus V$ be another type II point. Then:

- if $v \in \text{sk}(A)$ lies in the skeleton of an open annulus A in $X^{\text{an}} \setminus V$, then $A \setminus \{v\}$ is a disjoint union of two open annuli and infinitely many open balls;
- if $v \in |B|$ lies in an open ball B in $X^{\text{an}} \setminus V$, then $B \setminus \{v\}$ is a disjoint union of one open annulus and infinitely many open balls.

In either case, $V' := V \cup \{v\}$ is another semistable vertex set in X^{an} . One can check (e.g. using [BPR13, Lemma 3.2(1)]) that the skeleton Γ' associated to V' is obtained from the skeleton Γ associated to V by either:

- (edge subdivision) adding a new vertex (namely v) at a point partway along an edge of Γ , without changing the underlying metric space; or
- (budding off a leaf) adding a new vertex v and a single edge connecting v to a vertex in V ,

respectively. In either case, the curve $\bar{\mathcal{X}}_v$ attached to the new vertex v is

isomorphic to $\mathbb{P}_k^1$.

Any two semistable vertex sets can be obtained from one another by doing and undoing operations of this form, and so the skeleton Γ associated to a semistable vertex set V is independent of V up to the above operations.

7.2.1.3 Semistable coverings versus semistable vertex sets

In the case $K = \mathbb{C}_\ell$, we know that semistable models of X correspond bijectively to semistable vertex sets in X^{an} , and semistable models whose special fibre has at least two irreducible components correspond bijectively to semistable coverings of X^{an} . In particular, semistable vertex sets of size ≥ 2 in X^{an} correspond bijectively to semistable coverings of X^{an} . Since this will be important later, we now describe this correspondence directly, without going via semistable models.

Proposition 7.2.19. *Let $X/\mathbb{C}_\ell$ be a smooth projective curve, and let $\mathcal{X}$ be a semistable model whose special fibre has at least two irreducible components. Let $V \subset |X^{\text{an}}|$ be the semistable vertex set corresponding to $\mathcal{X}$, and let $\mathfrak{U} = (W_v)_{v \in V}$ be the semistable covering of X^{an} corresponding to $\mathcal{X}$.*

Then for all $v \in V$:

- *the wide open W_v is the subspace of X^{an} given by the union of $\{v\}$ and all discs and annuli in $|X^{\text{an}}| \setminus V$ whose closure contains v ;*
- *the bounding annuli of W_v are exactly the oriented annuli in $|X^{\text{an}}| \setminus V$ corresponding to non-loop edges e in the skeleton with $\partial_0(e) = v$;*
- *the internal annuli of W_v are exactly the annuli in $|X^{\text{an}}| \setminus V$ corresponding to loop edges e with endpoints v ; and*
- *the function field of the canonical reduction of $W_{v,0}$ is $\tilde{\mathcal{H}}(v)$.*

Proof. For the first point, if $\bar{\mathcal{X}}_v$ is the irreducible component of the special fibre corresponding to a vertex v , then its inverse image under the reduction map is v (inverse image of the generic point $\bar{\eta}_v$) together with the open discs and annuli which are the inverse images of the closed points of $\bar{\mathcal{X}}_v$. Suppose first that $\bar{x} \in \bar{\mathcal{X}}_v(\bar{\mathbb{F}}_\ell)$ is a point which is smooth in $\bar{\mathcal{X}}$, so its residue class $]\bar{x}[$ is an open disc. The topological boundary of $]\bar{x}[$ inside $|X^{\text{an}}|$ is $\{v'\}$ for some $v' \in V$ [BPR13, Lemma 3.2(1)]. Since the reduction map is anti-continuous [Ber90, Corollary 2.4.2] and the subset $\{\bar{x}, \bar{\eta}_v\} \subset |\bar{\mathcal{X}}|$ is an intersection of open subsets, it follows that $]\bar{x}[\cup \{v\}$ is closed in $|X^{\text{an}}|$, and so we must have $v' = v$, i.e. v is the unique boundary point of $]\bar{x}[$.

If instead $\bar{x}$ is singular in $\bar{\mathcal{X}}$, then its residue class $]\bar{x}[$ is an annulus, whose topological boundary consists of one or two elements of V [BPR13, Lemma 3.2(2)]. So, a similar argument establishes that if $\bar{x}$ is a self-intersection point of $\mathcal{X}_v$, then the topological boundary of $]\bar{x}[$ is $\{v\}$, while if $\bar{x}$ is an intersection point of $\mathcal{X}_v$ and $\mathcal{X}_{v'}$, then the topological boundary of $]\bar{x}[$ is⁵ $\{v, v'\}$. Put together, this tells us that the inverse image of $\mathcal{X}_v$ under the reduction map is exactly the union of v and all of the open discs and annuli whose closures contain v . This gives the first three points.

For the final point, we first note that the underlying affinoid $W_{v,0}$ of W_v is the union of $\{v\}$ and all discs in $|X^{\text{an}}| \setminus V$ whose closure contains v . For any $f \in \mathcal{O}(W_{v,0})$, the maximum value of $|f|$ on $W_{v,0}$ is attained at v , e.g. because $|f|$ is a continuous function $|W_{v,0}| \rightarrow \mathbb{R}$, and on each open disc in $|W_{v,0}| \setminus \{v\}$, $|f|$ increases monotonically towards the boundary point v (this follows from the fact that the maximum value of $|f|$ on any *closed* disc is attained at its Gauss point). In other words, the supremum norm on $R := \mathcal{O}(W_{v,0})$ is just the multiplicative norm $|\cdot|_v$ attached to the Berkovich point v .

It then follows that the canonical reduction of $W_{v,0}$ is $\text{Spec}(R^\circ/R^{\circ\circ})$, where R° (resp. $R^{\circ\circ}$) denotes the subring of R (resp. ideal of R°) consisting of elements of $|\cdot|_v$ -norm ≤ 1 (resp. < 1). The residue field $\tilde{\mathcal{H}}(v)$, on the other hand, is $\text{Frac}(R)^\circ/\text{Frac}(R)^{\circ\circ}$, where the multiplicative norm $|\cdot|_v$ is extended uniquely to $\text{Frac}(R)$ [Ber90, Remark 1.2.2(i)]. Thus, we want to show that $\text{Frac}(R)^\circ/\text{Frac}(R)^{\circ\circ}$ is the fraction field of $R^\circ/R^{\circ\circ}$. There is certainly an $\mathbb{F}_\ell$ -algebra homomorphism

$$R^\circ/R^{\circ\circ} \rightarrow \text{Frac}(R)^\circ/\text{Frac}(R)^{\circ\circ}. \quad (*)$$

An element of the kernel of this homomorphism would be represented by some $f \in R^\circ$ for which $|\frac{f}{1}|_v < 1$, which implies that $f = 0$ in $R^\circ/R^{\circ\circ}$. So $(*)$ is injective, and thus induces a map

$$\text{Frac}(R^\circ/R^{\circ\circ}) \rightarrow \text{Frac}(R)^\circ/\text{Frac}(R)^{\circ\circ}. \quad (**)$$

⁵Actually, there is a small amount of justification missing here, which seems also to be omitted in [BPR13, §4.9]. Specifically, the argument given shows that the boundary is contained inside $\{v, v'\}$, but does not show that the two sets are equal. To show this, let $\mathcal{X}'$ be an admissible blowup of $\mathcal{X}$ centred at $\bar{x}$ (or rather, at a closed, finitely presented subscheme whose reduced subscheme is $\{\bar{x}\}$). The corresponding semistable vertex set V' consists of V and one new vertex v'' lying on the skeleton of $]\bar{x}[$. This splits the skeleton of $]\bar{x}[$ into two open intervals, which are the skeleta of two annuli in $X^{\text{an}} \setminus V'$ corresponding to intersections of the new component $\mathcal{X}'_{v''}$ with $\mathcal{X}'_v$ and $\mathcal{X}'_{v'}$. Accordingly, the boundaries of these two annuli are contained in $\{v, v''\}$ and $\{v', v''\}$. This forces the two limit points of the skeleton of $]\bar{x}[$ to be distinct, equal to v and v' .

Given any element $\frac{f}{g} \in \text{Frac}(R)^\circ$, we can rescale f and g by elements of $\mathbb{C}_\ell$ so that $|g|_v = 1$ and $|f|_v \leq 1$. So f and g determine an element $\frac{\bar{f}}{\bar{g}} \in \text{Frac}(R^\circ/R^{\circ\circ})$ mapping to the class of $\frac{f}{g}$ in $\text{Frac}(R)^\circ/\text{Frac}(R)^{\circ\circ}$. Thus $(**)$ is surjective, i.e. is an isomorphism of fields over $\overline{\mathbb{F}}_\ell$. This is what we wanted to prove. $\square$

As a consequence, the skeleton attached to the semistable vertex set V agrees with the graph attached to the semistable covering $\mathfrak{U}$ as a metrised complex of $\overline{\mathbb{F}}_\ell$ -curves (cf. [BPR13, §4.9] for the statement without metrics or vertex curves).

7.2.2 Harmonic morphisms of metrised complexes

The reduction graph Γ attached to a curve X/K via a choice of split semistable model is functorial with respect to finite morphisms $f: X \rightarrow X'$ of curves, at least after suitable choices of models. This functoriality turns out to be surprisingly subtle, so we devote some time to explaining this carefully. The first subtlety lies in the correct notion of morphisms of graphs.

Definition 7.2.20 ([ABBR15, Definitions 2.4 & 2.19]). Let Γ and Γ' be metrised graphs. A *finite morphism* $f: \Gamma \rightarrow \Gamma'$ is a pair of functions $V(\Gamma) \rightarrow V(\Gamma')$ and $E(\Gamma) \rightarrow E(\Gamma')$ compatible with edge-inversion and source maps, such that the quantity

$$d_e(f) := \frac{l(f(e))}{l(e)}$$

is a positive integer for all edges e . The quantity $d_e(f)$ is called the *degree* of f along e .

Equivalently, from the metric perspective, a finite morphism $f: \Gamma \rightarrow \Gamma'$ is a continuous map $f: \Gamma \rightarrow \Gamma'$ such that $f^{-1}(V(\Gamma')) = V(\Gamma)$, and such that every connected component e of $\Gamma \setminus V(\Gamma)$ (which is isometric to an open interval and maps homeomorphically onto a connected component $f(e)$ of $\Gamma' \setminus V(\Gamma')$) maps onto its image via a dilation of some scale factor $d_e(f) \in \mathbb{Z}_{>0}$. A finite morphism $f: \Gamma \rightarrow \Gamma'$ of metrised graphs is said to be *harmonic* of *degree* $d_v(f)$ at $v \in V(\Gamma)$ if for every $e' \in T_{f(v)}(\Gamma')$ we have

$$\sum_{\substack{e \in T_v(\Gamma) \\ f(e) = e'}} d_e(f) = d_v(f).$$

The map f is said to be harmonic if it is surjective and harmonic at every $v \in V(\Gamma)$. Then, for any $v' \in V(\Gamma')$ the sum $\deg(f) = \sum_{v: f(v)=v'} d_v(f)$ is

independent of v' and called the *degree* of f .

A *finite harmonic morphism* $f: \Gamma \rightarrow \Gamma'$ of metrised complexes of k -curves is a finite morphism between the underlying metrised graphs endowed with a choice of finite morphism

$$\bar{f}_v: \bar{\mathcal{X}}_v \rightarrow \bar{\mathcal{X}}'_{f(v)}$$

of k -curves for each $v \in V(\Gamma)$ satisfying the following three conditions.

1. For every vertex $v \in V(\Gamma)$ and every $e \in T_v(\Gamma)$, we have that $\bar{x}_{f(e)} = \bar{f}_v(\bar{x}_e)$, and the map $\bar{f}_v$ is ramified of degree $d_e(f)$ at $\bar{x}_e$.
2. Conversely, for every vertex $v \in V(\Gamma)$ and every $e' \in T_{f(v)}(\Gamma')$, every point in $\bar{f}_v^{-1}(\bar{x}_{e'})$ is $\bar{x}_e$ for some $e \in T_v(\Gamma)$ with $f(e) = e'$.
3. For every vertex $v \in V(\Gamma)$ we have $d_v(f) = \deg(\bar{f}_v)$. (This is automatic from the preceding conditions as soon as Γ has at least one edge.)

A finite morphism $f: X \rightarrow X'$ of smooth projective curves induces, after suitable choices, a finite harmonic morphism $f: \Gamma \rightarrow \Gamma'$ between their reduction graphs. Describing this in terms of split semistable models is rather subtle – for example, a finite morphism $\mathcal{X} \rightarrow \mathcal{X}'$ between models need not induce a finite map on reduction graphs – so we instead follow [ABBR15] and describe this in terms of semistable vertex sets. The key result is the following.

Theorem 7.2.21 ([ABBR15, Theorem A, Corollary 4.26 & §4.27]). *Take a finite morphism $f: X \rightarrow X'$ of smooth projective curves over $\mathbb{C}_\ell$. Let V'_0 be a finite set of type II points in X'^{an} . Then there exists a semistable vertex set $V' \subset |X'^{\text{an}}|_{\text{II}}$ for X'^{an} containing V'_0 such that the preimage $V := f^{-1}(V') \subset |X^{\text{an}}|_{\text{II}}$ is a semistable vertex set for X^{an} and such that the skeleton Γ associated to V is the preimage of the skeleton Γ' associated to V' .*

Moreover, for any such V' , the induced map $f: \Gamma \rightarrow \Gamma'$ on skeleta is a finite harmonic morphism of metrised complexes of k -curves, of degree equal to $\deg(f)$.

In the final part, for each vertex $v \in \Gamma$, the morphism $\bar{f}_v: \bar{\mathcal{X}}_v \rightarrow \bar{\mathcal{X}}'_{f(v)}$, part of the data of a morphism of metrised complexes of k -curves, is the one whose induced map on function fields is the pullback map $\tilde{\mathcal{H}}(f(v)) \rightarrow \tilde{\mathcal{H}}(v)$ induced by $f: X^{\text{an}} \rightarrow X'^{\text{an}}$, see [ABBR15, §4.20].

7.2.2.1 Morphisms of curves and semistable coverings

We will also need to translate Theorem 7.2.21 into the language of semistable coverings.

Let $f: X \rightarrow X'$ be a finite morphism of smooth projective curves over $\mathbb{C}_\ell$, and choose a semistable vertex set $V' \subset |X'^{\text{an}}|_{\text{II}}$ satisfying the conditions of Theorem 7.2.21. Enlarging V' if necessary, we may assume that $\#V' \geq 2$ and that V' is *strongly* semistable, meaning that the corresponding skeleton is loopless. This implies the same conditions for the semistable vertex set $V := f^{-1}V' \subset |X^{\text{an}}|_{\text{II}}$. So there are associated semistable coverings $\mathfrak{U} = (W_v)_{v \in V}$ and $\mathfrak{U}' = (W'_{v'})_{v' \in V'}$ of X^{an} and X'^{an} .

The relationship between semistable coverings and semistable vertex sets in Proposition 7.2.19 implies that the semistable coverings $\mathfrak{U}$ and $\mathfrak{U}'$ are compatible in the following sense.

Proposition 7.2.22. *In the setup of Theorem 7.2.21, if $\#V' \geq 2$ and V' is strongly semistable, then:*

- for any vertex $v' \in V'$, we have

$$f^{-1}W'_{v'} = \bigcup_{v \in f^{-1}(v')} W_v,$$

and for any $v \in f^{-1}(v')$, the restriction of f to a map $W_v \rightarrow W'_{v'}$ is a finite morphism of degree $d_v(f)$;

- for any oriented edge e' of the skeleton attached to V' , we have

$$f^{-1}A'_{e'} = \bigcup_{e \in f^{-1}(e')} A_e,$$

and for any $e \in f^{-1}(e')$, the restriction of f to a map $A_e \rightarrow A'_{e'}$ is a finite morphism of degree $d_e(f)$ which preserves orientations; and

- for any vertex $v' \in V'$, we have

$$f^{-1}W'_{v',0} = \bigcup_{v \in f^{-1}(v')} W_{v,0},$$

and for any $v \in f^{-1}(v')$, the restriction of f to a map $W_{v,0} \rightarrow W'_{v',0}$ is a finite morphism of degree $d_v(f)$ whose canonical reduction is the restriction of the map $\bar{f}_v: \bar{X}_v \rightarrow \bar{X}'_{f(v)}$.

Proof. The statements regarding the inverse images follow from the descriptions in Proposition 7.2.19. The fact that the maps on wide opens and bounding annuli are finite is immediate since f is finite. For the assertions regarding the degree, consider first the induced morphism $f|_{A_e}: A_e \rightarrow A'_{e'}$ on bounding

annuli. We know that the induced map on skeleta is orientation-preserving, and is dilation by the scale factor $d_e(f)$ by definition. Using the classification of morphisms between annuli in [BPR13, Proposition 2.2(2)], this implies that $f|_{A_e}$ has degree $d_e(f)$ as claimed. By harmonicity, this implies that the map $f|_{W_v} : W_v \rightarrow W'_{v'}$ has degree $d_v(f)$ (compute the total degree over any point in a bounding annulus).

For the final assertion, for any vertex $v \in V$, the morphism $f : X^{\text{an}} \rightarrow X'^{\text{an}}$ on Berkovich analytifications induces an extension $\mathcal{H}(f(v)) \hookrightarrow \mathcal{H}(v)$ on completed residue fields, and so an extension $\tilde{\mathcal{H}}(f(v)) \hookrightarrow \tilde{\mathcal{H}}(v)$ on their residue fields. The corresponding map $\mathcal{X}_v \rightarrow \mathcal{X}'_{f(v)}$ of smooth projective $\mathbb{F}_\ell$ -curves is, by definition, the morphism $\bar{f}_v$ [ABBR15, §4.20]. So it follows from Proposition 7.2.19 that this $\bar{f}_v$ restricts to the canonical reduction of $f|_{W_{v,0}} : W_{v,0} \rightarrow W'_{f(v),0}$. This completes the proof. $\square$

7.3 The local heights formula

If X is a smooth projective curve over a finite extension K of $\mathbb{Q}_\ell$ with a chosen basepoint $b \in X(K)$, then the *normalised local height* $\tilde{h}_{Z,\ell}$ with respect to a trace zero correspondence $Z \subset X \times X$ is defined as follows. Let $z \in X(K)$ be a K -rational point, and define a divisor $D_Z(b, z) \subset X$ by

$$D_Z(b, z) := i_\Delta^* Z - i_1^* Z - i_2^* Z, \quad (7.3.0.1)$$

where $i_\Delta : X \hookrightarrow X \times X$ is the inclusion of the diagonal, and $i_1, i_2 : X \hookrightarrow X \times X$ are the inclusions of the subvarieties $\{b\} \times X$ and $X \times \{z\}$ [BD18, Definition 6.2]. Let $\mathcal{X}/\mathcal{O}_K$ be a regular model, and let $\mathcal{D}_Z(b, z) \subset \mathcal{X}$ be a $\mathbb{Q}$ -linear divisor whose generic fibre is $D_Z(b, z)$ and whose intersection multiplicity with any vertical divisor is zero. Then the normalised local height is given by the intersection multiplicity

$$\tilde{h}_{Z,\ell}(z) := (z - b) \cdot \mathcal{D}_Z(b, z) \in \mathbb{Q},$$

where we extend $(z - b)$ to a divisor on $\mathcal{X}$ in the obvious way. This height has the property that $\chi(\varpi_K) \tilde{h}_{Z,\ell}(z)$ is the Coleman–Gross height pairing of $z - b$ and $D_Z(b, z)$, where $\chi : K^\times \rightarrow \mathbb{Q}_p$ is the character appearing in the definition of the Coleman–Gross height [CG89, Proposition 1.2].

Remark 7.3.1. There are algorithms for computing the intersection pairing on divisors given a regular model for X over $\mathbb{Z}_\ell$ [Hol12, Mue14, vBHM20]. These algorithms depend on equations for Z and X . Equations for Z can be

extremely complicated even for simple genus 2 curves [DRHS23, Appendix], and so the complexity of the computation skyrockets as the genus and the conductor of X grows. In principle, these intersection pairing methods have the potential to work, but from a practical standpoint, working with equations for the correspondence Z to construct an explicit divisor on a regular model is difficult. We did attempt to use these intersection theory algorithms to compute local heights on a genus 2 curve with RM, but were unable to surmount the difficulties that arose in practice with dealing with the complicated equations for the correspondence. As far as we know, there are no examples in the literature where these methods have been used to compute local heights in the setting of quadratic Chabauty for rational points.

Our approach to computing local heights in this paper is based on a formula due to the first author and Netan Dogra [BD20, Corollary 12.1.3], which avoids the need to compute a regular model over K , and instead gives a direct combinatorial interpretation of $h_{Z,\ell}(z)$ in terms of reduction graphs. In this section, we will recall the statement of this formula.

7.3.1 Homology of the reduction graph

Let Γ denote the reduction graph of a split semistable model of the curve X/K , thought of as a metric space in the usual way. The local heights formula will be expressed primarily in terms of the homology group $H_1(\Gamma, \mathbb{Z})$. By Lemma 7.2.7 and the discussion below Theorem 7.2.17, we could equivalently define this group to be the homology group of the skeleton associated to a semistable vertex set in $X_{\mathbb{C}_\ell}^{\text{an}}$ (after embedding K inside $\mathbb{C}_\ell$). This makes it clear how to interpret the group $H_1(\Gamma, \mathbb{Z})$ even when X does not have split semistable reduction over K , and also makes it clear that $H_1(\Gamma, \mathbb{Z})$ is independent of the choice of split semistable model / semistable vertex set up to canonical isomorphism (by Remark 7.2.18)⁶.

The homology group $H_1(\Gamma, \mathbb{Z})$ comes with a perfect $\mathbb{R}$ -valued pairing, depending only on the underlying metric space of the graph Γ , and combinatorial in nature. We recall the construction, following [CK21, §3]. Let Γ be a metrised graph, let $C_0(\Gamma, \mathbb{Z}) := \mathbb{Z} \cdot V(\Gamma)$ be the free $\mathbb{Z}$ -module generated by the vertices of Γ , and let $C_1(\Gamma, \mathbb{Z}) := \mathbb{Z} \cdot E(\Gamma) / \sim$ denote the free $\mathbb{Z}$ -module generated by the oriented edges of Γ , modulo the relations $e^{-1} \sim -e$ for all $e \in E(\Gamma)$. Then

⁶Another way to see this is that the inclusion $\Gamma \hookrightarrow |X_{\mathbb{C}_\ell}^{\text{an}}|$ is a strong deformation retract [Ber90, Proposition 4.1.6], so $H_1(\Gamma, \mathbb{Z})$ is canonically isomorphic to the homology group $H_1(|X_{\mathbb{C}_\ell}^{\text{an}}|, \mathbb{Z})$, which is obviously independent of any choices.

the homology $H_1(\Gamma, \mathbb{Z})$ is the kernel of the boundary map

$$\partial: C_1(\Gamma, \mathbb{Z}) \rightarrow C_0(\Gamma, \mathbb{Z})$$

sending an oriented edge e to $\partial(e) = \partial_0(e) - \partial_0(e^{-1})$.

There is a canonical positive-definite symmetric pairing $\langle \cdot, \cdot \rangle: \text{Sym}^2 C_1(\Gamma, \mathbb{Z}) \rightarrow \mathbb{R}$ given by

$$\langle e_1, e_2 \rangle := \begin{cases} l(e_1) & \text{if } e_2 = e_1, \\ -l(e_1) & \text{if } e_2 = e_1^{-1}, \\ 0 & \text{else.} \end{cases}$$

Since this pairing is positive-definite, its restriction to $H_1(\Gamma, \mathbb{Z}) \subset C_1(\Gamma, \mathbb{Z})$ also defines a positive-definite symmetric pairing on $H_1(\Gamma, \mathbb{Z})$, called the *intersection length pairing* (or *cycle pairing* in [CK21]). We remark that when Γ is the reduction graph associated to a split semistable model of X/K , the intersection length pairing on $H_1(\Gamma, \mathbb{Z})$ is also independent of the choice of model.

7.3.1.1 Functoriality

The assignment of the group $H_1(\Gamma, \mathbb{Z})$ is functorial with respect to finite morphisms of curves, both co- and contravariantly. The covariant functoriality is clear: given a finite morphism $f: X \rightarrow X'$ of curves, without loss of generality defined over $\mathbb{C}_\ell$, we have by Theorem 7.2.21 that f induces a finite harmonic morphism $\Gamma \rightarrow \Gamma'$ on suitably chosen skeleta, and hence a homomorphism

$$f_*: H_1(\Gamma, \mathbb{Z}) \rightarrow H_1(\Gamma', \mathbb{Z})$$

(which again is independent of the choice of skeleta).

We can also define a map

$$f^*: H_1(\Gamma', \mathbb{Z}) \rightarrow H_1(\Gamma, \mathbb{Z})$$

in the other direction, namely the adjoint of f_* with respect to the intersection length pairings on the homology groups. One thing which is not apparent from the definition is that f^* is defined over $\mathbb{Z}$, rather than over $\mathbb{R}$ or $\mathbb{Q}$. This is a consequence of the following calculation.

Proposition 7.3.2. *Let $f: \Gamma \rightarrow \Gamma'$ be a finite harmonic morphism of metrised graphs. Then the pullback map*

$$f^*: H_1(\Gamma', \mathbb{R}) \rightarrow H_1(\Gamma, \mathbb{R})$$

(adjoint to the pushforward) is given by

$$f^*\left(\sum_{e'} \lambda_{e'} \cdot e'\right) = \sum_e \lambda_{f(e')} d_e(f) \cdot e.$$

In the above, summations over e or e' implicitly mean sums over unoriented edges of Γ or Γ' (the summands are independent of the orientation of edges).

In particular, since each $d_e(f)$ is an integer, f^* restricts to a map

$$f^*: H_1(\Gamma', \mathbb{Z}) \rightarrow H_1(\Gamma, \mathbb{Z}).$$

Proof. For a homology class

$$\gamma' = \sum_{e'} \lambda_{e'} \cdot e' \in H_1(\Gamma', \mathbb{R}),$$

let us write

$$f'(\gamma') := \sum_e \lambda_{f(e)} d_e(f) \cdot e.$$

We claim that $f'(\gamma') \in H_1(\Gamma, \mathbb{R})$. For this, we compute

$$\begin{aligned} \partial(f'(\gamma')) &= \sum_e \lambda_{f(e)} d_e(f) \partial(e) \\ &= \sum_v \sum_{e \in T_v(\Gamma)} \lambda_{f(e)} d_e(f) \cdot v \\ &= \sum_v \sum_{e' \in T_{f(v)}(\Gamma')} \lambda_{e'} \sum_{e \in T_v(\Gamma), f(e)=e'} d_e(f) \cdot v \\ &= \sum_v \sum_{e' \in T_{f(v)}(\Gamma')} \lambda_{e'} \cdot \deg_v(f) \cdot v = 0 \end{aligned}$$

using harmonicity and that $\sum_{e' \in T_{v'}(\Gamma')} \lambda_{e'} = 0$ for all $v' \in V(\Gamma')$ since γ' is a homology class. Hence $f'(\gamma') \in H_1(\Gamma, \mathbb{R})$ as claimed.

Now for any $\gamma = \sum_e \mu_e \cdot e \in H_1(\Gamma, \mathbb{R})$ we have

$$\langle f'(\gamma'), \gamma \rangle = \sum_e \lambda_{f(e)} \mu_e \cdot l(f(e)) = \sum_{e'} \lambda_{e'} \left(\sum_{f(e)=e'} \mu_e \right) \cdot l(e') = \langle \gamma', f_*(\gamma) \rangle$$

using that $d_e(f) = \frac{l(f(e))}{l(e)}$ in the first line. Hence $f' = f^*$ as desired. $\square$

All told, we have shown that the assignment $X \mapsto H_1(\Gamma, \mathbb{Z})$ is a functor

$$\begin{aligned} &\{\text{curves over } K \text{ \& finite morphisms}\} \rightarrow \\ &\{\text{lattices with an inner product \& adjoint pairs of homomorphisms}\}. \end{aligned}$$

7.3.1.2 Correspondences and component traces

Combining the pushforward and pullback functoriality, we find that the homology group $H_1(\Gamma, \mathbb{Z})$ is functorial with respect to correspondences between curves. For our purposes, by a *correspondence* from X to Y , we mean a Weil divisor $Z \subset X \times Y$. The pushforward Z_* along a correspondence is defined as follows. If Z is a geometrically integral divisor and the projections $\pi_1: Z \rightarrow X$ and $\pi_2: Z \rightarrow Y$ are finite, then we define $Z_* := \tilde{\pi}_{2*} \circ \tilde{\pi}_1^*$ where $\tilde{\pi}_1$ and $\tilde{\pi}_2$ are the two projections from the normalisation $\tilde{Z}$ of Z . If Z is geometrically integral and one projection is not finite, then we set $Z_* = 0$. In general, we first extend the base field so that all integral components of Z are geometrically integral, and then extend by linearity. We will be especially interested in correspondences Z from a curve X to itself.

These correspondences also have so-called traces on the components, which we will need later for the local height formula.

Definition 7.3.3. Let $Z \subset X \times X$ be a correspondence defined over $\mathbb{C}_\ell$, and let Γ be a skeleton in X^{an} . Let $v \in V(\Gamma)$, and let $S \subset |Z^{\text{an}}|_\Gamma$ be the set of points w that map to (v, v) . If we write $\tilde{Z}_w$ for the $\overline{\mathbb{F}}_\ell$ -curve attached to a point $w \in S$, then $\tilde{Z}_{v,v} := \bigsqcup_{w \in S} \tilde{Z}_w$ is an effective divisor inside $\bar{\mathcal{X}}_v \times \bar{\mathcal{X}}_v$, with projections π_1, π_2 to $\bar{\mathcal{X}}_v$ of degree d_1 and d_2 in total. We define the *trace* at v to be $\text{tr}_v(Z) = d_1 + d_2 - \Delta \cdot \tilde{Z}_{v,v} \in \mathbb{Z}$, where Δ is the diagonal in $\bar{\mathcal{X}}_v \times \bar{\mathcal{X}}_v$.

The above definition makes it clear that the component traces $\text{tr}_v(Z)$ are integers. They can also be interpreted as traces of the action of $\tilde{Z}_{v,v,*}$ on cohomology.

Lemma 7.3.4. *In the above notation, $\text{tr}_v(Z)$ is equal to the trace of $\tilde{Z}_{v,v,*}$ acting on the rigid cohomology group $H_{\text{rig}}^1(\bar{\mathcal{X}}_v/\mathbb{C}_\ell)$. For $\ell' \neq \ell$ it is also the trace of $\tilde{Z}_{v,v,*}$ acting on the étale cohomology group $H_{\text{ét}}^1(\bar{\mathcal{X}}_{v,\bar{k}}, \mathbb{Q}_{\ell'})$.*

Proof. The Lefschetz trace formula [Ber74, Théorème VII.3.1.9]⁷ gives that

$$\Delta \cdot \bar{Z}_{v,v} = \sum_{i=0,1,2} (-1)^i \operatorname{tr}(\bar{Z}_{v,v,*} | H_{\text{rig}}^i(\bar{\mathcal{X}}_v/K)).$$

The trace of $\bar{Z}_{v,v,*}$ acting on $H_{\text{rig}}^i(\bar{\mathcal{X}}_v/\mathbb{C}_\ell)$ is d_2 for $i = 0$ and is d_1 for $i = 2$. Rearranging gives the desired identity for $\operatorname{tr}_v(Z)$. The same argument applies for étale cohomology, using the corresponding trace formula [Mil80, Theorem 25.1]. $\square$

To obtain certifiably correct outputs from our algorithms, it will be necessary to have some *a priori* control on the possible actions of a correspondence Z on $H_1(\Gamma, \mathbb{Z})$ and the traces $\operatorname{tr}_v(Z)$ attached to vertices. This is what we establish here.

Theorem 7.3.5. *Let $Z \subset X \times X$ be an effective correspondence, of degrees d_1 and d_2 over X , respectively. Then:*

- a) $Z_*: H_1(\Gamma, \mathbb{Z}) \rightarrow H_1(\Gamma, \mathbb{Z})$ has operator norm $\leq \sqrt{d_1 d_2}$ with respect to the intersection length pairing on $H_1(\Gamma, \mathbb{Z})$.
- b) For all $v \in V(\Gamma)$, we have

$$|\operatorname{tr}_v(Z)| \leq 2g(v) \cdot \max\{d_1, d_2\},$$

where $g(v)$ is the genus of v .

Now we prove Theorem 7.3.5, dealing with the two parts separately. For the first part, it suffices to deal separately with pushforward and pullback of cohomology classes.

Lemma 7.3.6. *Suppose that $f: \Gamma \rightarrow \Gamma'$ is a harmonic map of metrised graphs of degree d . Then the pushforward map $f_*: H_1(\Gamma, \mathbb{Z}) \rightarrow H_1(\Gamma', \mathbb{Z})$ has operator norm at most $\sqrt{d}$, i.e. we have*

$$\langle f_*(\gamma), f_*(\gamma) \rangle \leq d \cdot \langle \gamma, \gamma \rangle$$

for all $\gamma \in H_1(\Gamma, \mathbb{R})$.

⁷Strictly speaking, [Ber74, Théorème VII.3.1.9] only applies when $\bar{Z}_{v,v}$ is the graph of an endomorphism of $\bar{\mathcal{X}}_v$, but the argument is easily generalised to the case of correspondences. Also, [Ber74, Théorème VII.3.1.9] is a statement about crystalline cohomology rather than rigid cohomology, but these agree for smooth proper varieties [Ber86, Proposition 2].

Proof. If we write $\gamma = \sum_e \lambda_e \cdot e$ as usual, then we have

$$\begin{aligned}
 \langle f_*(\gamma), f_*(\gamma) \rangle &= \sum_{e'} \left(\sum_{f(e)=e'} \lambda_e \right)^2 \cdot l(e') \\
 &\leq \sum_{e'} \left(\sum_{f(e)=e'} \frac{1}{l(e)} \right) \left(\sum_{f(e)=e'} \lambda_e^2 l(e) \right) l(e') \\
 &= \sum_{e'} \left[\left(\sum_{f(e)=e'} d_e(f) \right) \left(\sum_{f(e)=e'} \lambda_e^2 l(e) \right) \right] \\
 &= d \sum_e \lambda_e^2 l(e) = d \langle \gamma, \gamma \rangle,
 \end{aligned}$$

using the Cauchy–Schwartz inequality in the second line. $\square$

Lemma 7.3.7. *Suppose that $f: \Gamma \rightarrow \Gamma'$ is a harmonic map of metrised graphs of degree d . Then we have*

$$\langle f^*(\gamma'_1), f^*(\gamma'_2) \rangle = d \cdot \langle \gamma'_1, \gamma'_2 \rangle$$

for all $\gamma'_1, \gamma'_2 \in H_1(\Gamma', \mathbb{R})$. In particular, f^* has operator norm exactly $\sqrt{d}$.

Proof. Let us write

$$\gamma'_1 = \sum_{e'} \lambda'_{1,e'} \cdot e' \quad \text{and} \quad \gamma'_2 = \sum_{e'} \lambda'_{2,e'} \cdot e'.$$

According to the pullback formula

$$f^*(\gamma'_i) = \sum_e \lambda_{if(e)} d_e(f) \cdot e$$

for $i = 1, 2$, and hence

$$\begin{aligned}
 \langle f^*(\gamma'_1), f^*(\gamma'_2) \rangle &= \sum_e \lambda_{1f(e)} \lambda_{2f(e)} d_e(f)^2 l(e) \\
 &= \sum_{e'} \left(\sum_{f(e)=e'} d_e(f) \right) \cdot \lambda_{1e'} \lambda_{2e'} \cdot l(e') \\
 &= d \sum_{e'} \lambda_{1e'} \lambda_{2e'} \cdot l(e') = d \langle \gamma'_1, \gamma'_2 \rangle.
 \end{aligned}$$

$\square$

Proof of Theorem 7.3.5 Item (a). We assume that all irreducible components of Z are geometrically irreducible, which can always be achieved after a finite extension of the base field. If Z itself is geometrically irreducible, then $Z_* = \pi_{2*}\pi_1^*$ is the composite of two maps which have operator norms at most $\sqrt{d_1}$ and $\sqrt{d_2}$, respectively, by the above lemmas and Theorem 7.2.21. So we are done in this case.

To prove the general case, we need to show that if we know the result for two effective correspondences Z_1 and Z_2 , then the result holds for $Z = Z_1 + Z_2$. For this, let (d_{11}, d_{12}) and (d_{21}, d_{22}) be the degrees of Z_1 and Z_2 over X , respectively, so that $d_1 = d_{11} + d_{21}$ and $d_2 = d_{12} + d_{22}$. For any $\gamma \in H_1(\Gamma, \mathbb{Z})$ we have

$$\begin{aligned} \langle Z_*(\gamma), Z_*(\gamma) \rangle &= \langle (Z_{1*} + Z_{2*})(\gamma), (Z_{1*} + Z_{2*})(\gamma) \rangle \\ &\leq \left(d_{11}d_{12} + 2\sqrt{d_{11}d_{12}d_{21}d_{22}} + d_{21}d_{22} \right) \cdot \langle \gamma, \gamma \rangle \\ &\leq (d_{11} + d_{21})(d_{12} + d_{22}) \cdot \langle \gamma, \gamma \rangle, \end{aligned}$$

using the Cauchy–Schwarz inequality in the second line and the AM–GM inequality in the last line. This proves the result we want for Z . $\square$

Now we turn to the second part of Theorem 7.3.5. The main calculation is the following.

Lemma 7.3.8. *Let $\bar{\mathcal{X}}$ be a smooth projective curve of genus g over k , and $\bar{\mathcal{Z}} \subset \bar{\mathcal{X}} \times \bar{\mathcal{X}}$ an effective correspondence, of degrees d_1 and d_2 over $\bar{\mathcal{X}}$. Then $d_1 + d_2 - \bar{\mathcal{Z}} \cdot \Delta$ has absolute value at most $2g \max\{d_1, d_2\}$.*

Proof. Note that by Lemma 7.3.4 we have $\text{tr}(\bar{\mathcal{Z}}_* | H_{\text{rig}}^1(\bar{\mathcal{X}}/K)) = d_1 + d_2 - \bar{\mathcal{Z}} \cdot \Delta$ for a correspondence $\bar{\mathcal{Z}}$ on $\bar{\mathcal{X}}$. The space $H_{\text{rig}}^1(\bar{\mathcal{X}}/K)$ is a $2g$ -dimensional vector space over K . We may assume that $g \geq 1$, else there is nothing to prove. If we write $\bar{\mathcal{Z}} = \bar{\mathcal{Z}}_0 + m\Delta$ where $m \geq 0$ and $\bar{\mathcal{Z}}_0$ is an effective divisor on $\bar{\mathcal{X}} \times \bar{\mathcal{X}}$ not containing Δ , then we have $\bar{\mathcal{Z}}_0 \cdot \Delta \geq 0$ by [Har77, Proposition V.1.4], and so $\bar{\mathcal{Z}} \cdot \Delta \geq m\Delta \cdot \Delta = 2m(1 - g) \geq (d_1 + d_2)(1 - g)$. Combined with the above we obtain

$$\text{tr}(\bar{\mathcal{Z}}_* | H_{\text{rig}}^1(X/K)) \leq g(d_1 + d_2).$$

Now applying the same logic to the n th iterate $\bar{\mathcal{Z}}^n$ of $\bar{\mathcal{Z}}$, we find that $\bar{\mathcal{Z}}_*^n$ has integer trace on $H_{\text{rig}}^1(\bar{\mathcal{X}}/K)$, which implies that all of the eigenvalues $\lambda_1, \dots, \lambda_{2g}$ of $\bar{\mathcal{Z}}_*$ on $H_{\text{rig}}^1(\bar{\mathcal{X}}/K)$ are algebraic integers. Moreover, we have

$$\text{tr}(\bar{\mathcal{Z}}_*^n | H_{\text{rig}}^1(\bar{\mathcal{X}}/K)) \leq g(d_1^n + d_2^n)$$

for all $n \geq 0$. So, if we view the eigenvalues $\lambda_1, \dots, \lambda_n$ as elements of $\mathbb{C}$ via some complex embedding $\mathbb{Q}(\lambda_1, \dots, \lambda_{2g}) \hookrightarrow \mathbb{C}$, then we have

$$\Re \left(\sum_{i=1}^n \lambda_i^n \right) \leq g(d_1^n + d_2^n)$$

for all $n \geq 0$. Since $(S^1)^{2g}$ is compact, we can choose $n \gg 0$ such that the argument of each λ_i^n is sufficiently close to 0 that $\Re(\lambda_i^n) \geq \frac{1}{2}|\lambda_i^n|$ for all i . So for these n we have

$$\sum_{i=1}^{2g} |\lambda_i^n| \leq 2g(d_1^n + d_2^n);$$

taking n sufficiently large with this property we obtain that $|\lambda_i| \leq \max\{d_1, d_2\}$ for all i . Hence

$$|\mathrm{tr}(\bar{\mathcal{Z}}_* | \mathrm{H}_{\mathrm{rig}}^1(\bar{\mathcal{X}}/K))| = \left| \sum_i \lambda_i \right| \leq 2g \max\{d_1, d_2\}.$$

□

Proof of Theorem 7.3.5 Item (b). Again, we may and do assume Z is geometrically irreducible. Letting $\bar{\mathcal{Z}}_{v,v} = \bigsqcup_{w \in S} \bar{\mathcal{Z}}_w$ be as in Definition 7.3.3, we obtain by Lemma 7.3.8 that

$$|\mathrm{tr}_v(Z)| \leq g(v) \cdot \max \left\{ \sum_{w \in S} d_w(\pi_1), \sum_{w \in S} d_w(\pi_2) \right\}.$$

But we have

$$\sum_{w \in S} d_w(\pi_i) \leq \deg(\pi_i) = d_i$$

for $i = 1, 2$ by Proposition 7.2.22, and so we are done. □

7.3.2 The local heights formula

We are now finally in a position to state the formula for local heights from [BD20]. To do so, recall that if Γ is a metrised graph (viewed as a metric space), then by a *piecewise polynomial function* we mean a continuous function $f: \Gamma \rightarrow \mathbb{R}$ which is given by a polynomial in the arc-length when restricted to any edge of Γ . By a *piecewise polynomial measure* we mean a formal sum

$$\mu = \sum_e g_e \cdot |\mathrm{d}s_e| + \sum_v \lambda_v \cdot \delta_v,$$

where the first sum is taken over unoriented edges of Γ (i.e. connected components of $\Gamma \setminus V(\Gamma)$), g_e is a function $e \rightarrow \mathbb{R}$ which is a polynomial in arc length, and $\lambda_v \in \mathbb{R}$. We think of $|ds_e|$ as the unit length measure on edge e and δ_v as a delta measure supported at the vertex v . Such a measure has a *total mass*

$$\sum_{e \in E(\Gamma)} \int_0^{\ell(e)} g_e(s_e) |ds_e| + \sum_{v \in V(\Gamma)} \lambda_v.$$

If f is a piecewise polynomial function and $\vec{v}$ is a tangent direction at a vertex $v \in V(\Gamma)$, then one can make sense of *derivative* $D_{\vec{v}}f(v)$ of f at v in direction $\vec{v}$ [BF06, Definition 3]. The *Laplacian* of f is the piecewise polynomial measure $\nabla^2(f)$ defined by

$$\nabla^2(f) := - \sum_e (f|_e)'' \cdot |ds_e| - \sum_v \left(\sum_{\vec{v} \in T_v(\Gamma)} D_{\vec{v}}f(v) \right) \cdot \delta_v,$$

where $(f|_e)''$ denotes the second derivative of f with respect to arc length along e [BF06, Definition 5]. The Laplacian defines an $\mathbb{R}$ -linear map from the space of piecewise polynomial functions to the space of piecewise polynomial measures; its kernel is the one-dimensional space of constant functions, and its image is the codimension-one space of measures of total mass 0.

Finding an explicit piecewise polynomial function f whose Laplacian is a given measure μ of total mass 0 is not difficult. By formally double-integrating g_e along each edge of Γ one finds a piecewise polynomial function f_0 such that $\mu - \nabla^2(f_0)$ is a sum of delta-measures supported at vertices of Γ . As in [BF06, §5], by finding a right inverse for the weighted Laplacian matrix one finds a piecewise affine function f_1 such that $\nabla^2(f_1) = \mu - \nabla^2(f_0)$, and then $f = f_0 + f_1$ is the desired Laplacian inverse of μ , well defined up to a constant function.

The local height formula of [BD20] gives a formula for the Laplacian of the normalised local height associated to a correspondence Z as an explicit piecewise polynomial measure on the reduction graph.

Theorem 7.3.9 ([BD20, Corollary 12.1.3]⁸). *Let K be a finite extension of $\mathbb{Q}_\ell$, X/K a smooth projective curve with base point $b \in X(K)$ and $Z \subset X \times X$ a correspondence representing a trace 0 endomorphism of $\mathfrak{Iac}(X)$ which is fixed*

⁸There are a couple of errors in the statement of [BD20, Corollary 12.1.3] which we correct here.

by the Rosati involution. Let Γ be the reduction graph of a split semistable model of X . Then the normalised local height function $\tilde{h}_{Z,\ell}: X(K) \rightarrow \mathbb{Q}$ factors through a piecewise polynomial function $\tilde{h}_{Z,\ell}: \Gamma \rightarrow \mathbb{R}$ whose Laplacian is given by

$$\nabla^2(\tilde{h}_{Z,\ell}) = 2 \sum_{e \in E(\Gamma)^+} \frac{1}{l(e)^2} \langle e, Z_*(\pi(e)) \rangle \cdot |ds_e| + \sum_{v \in V(\Gamma)} \text{tr}_v(Z) \cdot \delta_v,$$

where $\langle \cdot, \cdot \rangle$ denotes the intersection length pairing on $C_1(\Gamma, \mathbb{Z})$ and the map $\pi: C_1(\Gamma, \mathbb{R}) \rightarrow H_1(\Gamma, \mathbb{R})$ is the orthogonal projection. Furthermore, $\tilde{h}_{Z,\ell}$ vanishes at the reduction of the base point b , and these properties determine $\tilde{h}_{Z,\ell}$ completely.

Remark 7.3.10. Since all of the edge lengths in Γ are rational, it follows that all of the quantities appearing in the above expression are rational numbers. Hence the function $\tilde{h}_{Z,\ell}$ has the property that it takes points lying a rational distance along edges of Γ to rational numbers.

Remark 7.3.11. We have refined the statement from [BD20] so that our result is purely combinatorial, rather than a formula in terms of étale cohomology. Theorem 7.3.9 as stated above does not directly follow from [BD20]. The issue is the following. If we let $H_{\text{ét}}^1(X_{\overline{K}}, \mathbb{Q}_{\ell'})$ denote the étale cohomology group of $X_{\overline{K}}$, then this comes with a natural monodromy filtration, and the theorem of Picard–Lefschetz gives identifications of its graded pieces as

$$\text{gr}_i^M H_{\text{ét}}^1(X_{\overline{K}}, \mathbb{Q}_{\ell'}) \cong \begin{cases} H_1(\Gamma, \mathbb{Q}_{\ell'}) & \text{if } i = 2, \\ \bigoplus_{v \in V(\Gamma)} H_{\text{ét}}^1(\bar{\mathcal{X}}_{v,\bar{k}}, \mathbb{Q}_{\ell'}) & \text{if } i = 1, \\ H^1(\Gamma, \mathbb{Q}_{\ell'}) & \text{if } i = 0, \\ 0 & \text{else.} \end{cases}$$

By functoriality of étale cohomology, the correspondence Z induces a push-forward endomorphism Z_* of the group $H_{\text{ét}}^1(X_{\overline{K}}, \mathbb{Q}_{\ell'})$ preserving the monodromy filtration, and this in turn induces endomorphisms of $H_1(\Gamma, \mathbb{Q}_{\ell'})$ and each $H_{\text{ét}}^1(\bar{\mathcal{X}}_{v,\bar{k}}, \mathbb{Q}_{\ell'})$. The statement proved in [BD20] is valid for Z_* the endomorphisms induced this way via Picard–Lefschetz, rather than the endomorphisms we have defined graph-theoretically above. So to carefully deduce Theorem 7.3.9 in the form stated, we need to prove that the Picard–Lefschetz isomorphism is compatible with pushforwards and pullbacks; we will do this in a forthcoming appendix to this paper.

7.4 The Coleman–Iovita isomorphism

In order to apply the local height formula, we need to be able to compute the action of a correspondence $Z \subset X \times X$ on the homology $H_1(\Gamma, \mathbb{Z})$ of the dual graph Γ , as well as the traces $\mathrm{tr}_v(Z)$ of Z at vertices v . As outlined in the introduction, our strategy for doing this proceeds by first computing the action of Z_* on the de Rham cohomology of X , and then using this to deduce the action on $H_1(\Gamma, \mathbb{Z})$ via the Coleman–Iovita isomorphism. This isomorphism, whose statement we recall below, relates the de Rham cohomology of X to the homology of Γ and the rigid cohomology of the k -curves attached to vertices of Γ .

Theorem 7.4.1 (Coleman–Iovita isomorphism, [CI10, §3.5], [DR17, §3.1]). *Let K be $\mathbb{C}_\ell$ or a finite extension of $\mathbb{Q}_\ell$. Let X/K be a smooth projective curve, and let $\mathcal{X}/\mathcal{O}_K$ be a split strongly⁹ semistable model, with reduction graph Γ . Then the de Rham cohomology $H_{\mathrm{dR}}^1(X/K)$ carries a canonical monodromy filtration*

$$\begin{aligned} 0 &= M_{-1}H_{\mathrm{dR}}^1(X/K) \\ &\leq M_0H_{\mathrm{dR}}^1(X/K) \\ &\leq M_1H_{\mathrm{dR}}^1(X/K) \\ &\leq M_2H_{\mathrm{dR}}^1(X/K) = H_{\mathrm{dR}}^1(X/K) \end{aligned}$$

along with canonical identifications of the graded pieces

$$\mathrm{gr}_n^M H_{\mathrm{dR}}^1(X/K) \cong \begin{cases} H_1(\Gamma, K) & \text{if } n = 2, \\ \bigoplus_v H_{\mathrm{rig}}^1(\bar{\mathcal{X}}_v/K) & \text{if } n = 1, \\ H^1(\Gamma, K) & \text{if } n = 0. \end{cases}$$

Here, $H_1(\Gamma, K)$ and $H^1(\Gamma, K)$ denote the homology and cohomology of the graph Γ with coefficients in K , $H_{\mathrm{rig}}^1(\bar{\mathcal{X}}_v/K)$ denotes the rigid cohomology of the irreducible component $\bar{\mathcal{X}}_v$ of the special fibre $\bar{\mathcal{X}}$ attached to a vertex v of Γ , and the direct sum is taken over all vertices of Γ .

Since we will be using this isomorphism in computations, we will need to explicitly recall its construction, and prove some elementary (but fiddly) compatibility properties. For this, we first need an explicit description of the de Rham cohomology of wide opens.

⁹The same result is true without the word “strongly”, but the explicit description of the Coleman–Iovita isomorphism is a little more complicated.

7.4.1 Cohomology of wide opens

Let $(W_v, W_{v,0})$ be a wide open pair, with bounding annuli $A_1, \dots, A_n$, oriented as per Remark 7.2.10. Choose a parameter t_i on each oriented annulus A_i , giving an isomorphism $A_i \cong A(r_{i,1}, r_{i,2})$. Let X_v be the rigid space obtained by gluing together W_v and the open discs $D(r_{i,2})$ of radius $r_{i,2}$ for $1 \leq i \leq n$, along the isomorphisms $t_i: A_i \cong A(r_{i,1}, r_{i,2})$. Then X_v is the rigid analytification of a smooth projective curve over K [MC10, Theorem 2.18]. We call X_v a *compactification* of W_v . (The space W_v has many non-isomorphic compactifications, arising from different choices of the parameters t_i .)

For each bounding annulus A_i of W_v , its de Rham cohomology is isomorphic to K via the residue map

$$\mathrm{Res}_{A_i}: H_{\mathrm{dR}}^1(A_i/K) \xrightarrow{\sim} K,$$

see [MC10, Lemma 2.13]. This map depends only on the orientation on A_i , not on the particular parameter t_i .

Lemma 7.4.2. *Let $(W_v, W_{v,0})$ be a wide open pair as above. Then for all 1-forms $\omega \in \Omega^1(W_v)$, we have*

$$\sum_i \mathrm{Res}_{A_i}(\omega) = 0.$$

Proof. The proof of [MC10, Corollary 2.33] gives an exact sequence

$$0 \rightarrow H_{\mathrm{dR}}^1(X_v/K) \rightarrow H_{\mathrm{dR}}^1(W_v/K) \rightarrow \bigoplus_i K \xrightarrow{\Sigma} K \rightarrow 0, \quad (7.4.1.1)$$

in which the middle arrow is the direct sum of the residue maps Res_{A_i} . This implies the result. $\square$

The fact that the left-hand group $H_{\mathrm{dR}}^1(X_v/K)$ is non-canonical, i.e. depends on the parameters t_i , will cause us significant headaches when it comes to considering morphisms between curves in the next section. Thus, it is convenient for us to reinterpret this de Rham cohomology group as a rigid cohomology group in a manner which is independent of choices.

Recall that if $\bar{\mathcal{Y}}$ is a smooth variety over k , then the rigid cohomology of $\bar{\mathcal{Y}}$ can be defined by choosing a smooth proper frame

$$\bar{\mathcal{Y}} \hookrightarrow \bar{\mathcal{X}} \hookrightarrow \mathcal{P},$$

i.e. an open embedding of $\bar{\mathcal{Y}}$ in a proper k -variety $\bar{\mathcal{X}}$ and a closed embedding of $\bar{\mathcal{X}}$ in a formal $\mathcal{O}_K$ -scheme $\mathcal{P}$ which is smooth in a neighbourhood of $\bar{\mathcal{Y}}$ [LS07, Definitions 3.1.5, 3.3.5, 3.3.10]. The rigid cohomology¹⁰ $H_{\text{rig}}^\bullet(\bar{\mathcal{Y}}/K)$ of $\bar{\mathcal{Y}}$ is then defined by

$$H_{\text{rig}}^\bullet(\bar{\mathcal{Y}}/K) := H_{\text{dR}}^\bullet(\mathcal{P}_K, j^\dagger \mathcal{O}_{\bar{\mathcal{Y}}[\mathcal{P}]}) ,$$

i.e. it is the de Rham cohomology of the Raynaud generic fibre $\mathcal{P}_K$ with coefficients in the sheaf $j^\dagger \mathcal{O}_{\bar{\mathcal{Y}}[\mathcal{P}]}$ of overconvergent sections of the structure sheaf of the tube of $\bar{\mathcal{Y}}$ inside $\mathcal{P}$ [LS07, Definitions 8.2.5, Proposition 5.1.14]. As the notation suggests, the rigid cohomology of $\bar{\mathcal{Y}}$ is independent of the choice of smooth proper frame up to canonical isomorphism [LS07, Propositions 7.4.2, 8.2.1]. We will only be interested in the case that $\bar{\mathcal{Y}}$ is a smooth curve over k , $\bar{\mathcal{X}}$ is its smooth completion, and $\mathcal{P}$ is a deformation of $\bar{\mathcal{X}}$ to a smooth proper formal curve over $\mathcal{O}_K$.

Proposition 7.4.3. *Let $(W_v, W_{v,0})$ be a strongly basic wide open pair, let $\bar{\mathcal{Y}}_v$ be the smooth canonical reduction of $W_{v,0}$, with smooth completion $\bar{\mathcal{X}}_v$. Choose parameters on the bounding annuli of $(W_v, W_{v,0})$ giving rise to a compactification X_v of W_v .*

Then there are canonical isomorphisms

$$H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K) \cong H_{\text{dR}}^1(X_v/K) \quad \text{and} \quad H_{\text{rig}}^1(\bar{\mathcal{Y}}_v/K) \cong H_{\text{dR}}^1(W_v/K)$$

which fit into a commuting square

$$\begin{array}{ccc} H_{\text{dR}}^1(X_v/K) & \hookrightarrow & H_{\text{dR}}^1(W_v/K) \\ \parallel \wr & & \parallel \wr \\ H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K) & \longrightarrow & H_{\text{rig}}^1(\bar{\mathcal{Y}}_v/K) \end{array}$$

in which the horizontal maps are induced by the evident inclusions.

Proof of Proposition 7.4.3. According to [MC10, Theorem 2.27], the compactification X_v has a proper formal model $\mathcal{X}_v$ whose special fibre is $\bar{\mathcal{X}}_v$. Such a model is necessarily smooth (since $\bar{\mathcal{X}}_v$ is smooth), so

$$\bar{\mathcal{X}}_v = \bar{\mathcal{X}}_v \hookrightarrow \mathcal{X}_v \quad \text{and} \quad \bar{\mathcal{Y}}_v \hookrightarrow \bar{\mathcal{X}}_v \hookrightarrow \mathcal{X}_v$$

are smooth proper frames for $\bar{\mathcal{X}}_v$ and $\bar{\mathcal{Y}}_v$, respectively.

¹⁰The notation from [LS07] would be rather $H_{\text{rig}}^\bullet(\bar{\mathcal{Y}}/\mathcal{O}_K)$, but we prefer $H_{\text{rig}}^\bullet(\bar{\mathcal{Y}}/K)$ to emphasise that it is a K -vector space.

Calculating the rigid cohomology of $\bar{\mathcal{X}}_v$ with respect to this frame gives the desired identification

$$H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K) \cong H_{\text{dR}}^1(X_v/K),$$

cf. [LS07, Proposition 8.2.6(ii)].

Calculating the rigid cohomology of $\bar{\mathcal{Y}}_v$ with respect to its frame gives an identification

$$H_{\text{rig}}^1(\bar{\mathcal{Y}}_v/K) \cong \varinjlim_{W'} H_{\text{dR}}^1(W'/K),$$

where the colimit is taken over strict neighbourhoods W' of $W_{v,0}$ inside X_v [LS07, Proposition 5.1.12(ii)]. For $\lambda < 1$ in $\sqrt{|K^\times|}$, let $W^\lambda \subset X_v$ denote the complement of the closed discs $D[\lambda r_{1,i}^{-1}]$ inside the bounding discs $D(r_{1,i}^{-1})$. It follows from [LS07, Proposition 3.3.2] that the W^λ form a cofinal system of strict neighbourhoods of $W_{v,0}$ inside X_v .

Moreover, for λ sufficiently close to 1, we have $W^\lambda \subseteq W_v$, and W_v is obtained from W^λ by gluing the annuli $A(r_{2,i}^{-1}, r_{1,i}^{-1})$ along the subannuli $A(\lambda r_{1,i}^{-1}, r_{1,i}^{-1})$. Since each inclusion $A(\lambda r_{1,i}^{-1}, r_{1,i}^{-1}) \hookrightarrow A(r_{2,i}^{-1}, r_{1,i}^{-1})$ induces an isomorphism on de Rham cohomology, it follows by a Mayer–Vietoris argument that the inclusion $W^\lambda \hookrightarrow W_v$ also induces an isomorphism on de Rham cohomology. Hence the natural map

$$H_{\text{dR}}^1(W_v/K) \rightarrow \varinjlim_{\lambda} H_{\text{dR}}^1(W^\lambda/K) = \varinjlim_{W'} H_{\text{dR}}^1(W'/K)$$

coming from the fact that W_v is a strict neighbourhood of W_0 is an isomorphism, and so we have

$$H_{\text{rig}}^1(\bar{\mathcal{Y}}_v/K) \cong H_{\text{dR}}^1(W_v/K).$$

Finally, for the compatibility between these identifications, we use that the map $H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K) \rightarrow H_{\text{rig}}^1(\bar{\mathcal{Y}}_v/K)$ induced by the inclusion $\bar{\mathcal{Y}}_v \hookrightarrow \bar{\mathcal{X}}_v$ is the same as the base-change map in relative rigid cohomology associated to the commuting square

$$\begin{array}{ccc} \bar{\mathcal{Y}}_v & \hookrightarrow & \bar{\mathcal{X}}_v \\ \downarrow & & \downarrow \\ \text{Spec}(k) & \xlongequal{\quad} & \text{Spec}(k), \end{array}$$

cf. [LS07, Proposition 8.2.10]. A similar thing is true for the map induced on de Rham cohomology, so the desired compatibility follows from the fact that the base-change maps in rigid and de Rham cohomology are compatible [LS07, Proposition 8.2.11]. $\square$

Remark 7.4.4. When we say that the isomorphisms in Proposition 7.4.3 are canonical, we mean that the isomorphism $H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K) \cong H_{\text{dR}}^1(X_v/K)$ depends only on $(W_v, W_{v,0})$ and the chosen parameters t_i used to define X_v , while the isomorphism $H_{\text{rig}}^1(\bar{\mathcal{Y}}_v/K) \cong H_{\text{dR}}^1(W_v/K)$ depends only on $(W_v, W_{v,0})$ and *not* on the parameters t_i . In particular, the inclusion $H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K) \hookrightarrow H_{\text{dR}}^1(W_v/K)$ does not depend on the choice of parameters, so we can rewrite the exact sequence (7.4.1.1) as

$$0 \rightarrow H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K) \rightarrow H_{\text{dR}}^1(W_v/K) \rightarrow \bigoplus_i K \xrightarrow{\Sigma} K \rightarrow 0, \quad (7.4.1.2)$$

which depends only on the strongly basic wide open $(W_v, W_{v,0})$.

7.4.2 Description of the Coleman–Iovita isomorphism

Equipped with this description of the cohomology of wide opens, we are now in a position to explicitly describe the Coleman–Iovita isomorphism, following [DR17, §3]. This description will be purely rigid analytic. Let $\mathfrak{U} = (W_v)_v$ be the strongly semistable covering of X^{an} corresponding to a split strongly semistable model under Theorem 7.2.13. If $[\omega] \in H_{\text{dR}}^1(X/K)$ is a de Rham cohomology class, then for any oriented edge e of Γ , we can restrict $[\omega]$ to the annulus A_e to obtain a class in $H_{\text{dR}}^1(A_e/K)$, and then take its residue $\text{Res}_{A_e}(\omega)$. The residue of $[\omega]$ along the inverse annulus $A_{e^{-1}}$ is $\text{Res}_{A_{e^{-1}}}(\omega) = -\text{Res}_{A_e}(\omega)$, so these residues determine a well-defined 1-chain

$$\varphi_2(\omega) := \sum_e \text{Res}_{A_e}(\omega) \cdot e \in C_1(\Gamma, K).$$

By Lemma 7.4.2, this 1-chain is a 1-cycle, hence a homology class. The assignment $\omega \mapsto \sum_e \text{Res}_{A_e}(\omega) \cdot e$ thus defines a K -linear map

$$\varphi_2: H_{\text{dR}}^1(X/K) \rightarrow H_1(\Gamma, K).$$

If $[\omega] \in H_{\text{dR}}^1(X/K)$ lies in the kernel of φ_2 , then the restriction of $[\omega]$ to the wide open W_v has residue 0 over all bounding annuli, so lies in $H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K)$

using the exact sequence (7.4.1.2). This construction defines a K -linear map

$$\varphi_1: \ker(\varphi_2) \rightarrow \bigoplus_v H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K).$$

Finally, if $[\omega] \in H_{\text{dR}}^1(X/K)$ lies in the kernel of φ_2 and φ_1 , then $[\omega]$ is represented by a meromorphic 1-form ω of the second kind on X^{an} whose restriction to any wide open W_v is an exact form, so $d f_v$ for some meromorphic function f_v on W_v . If e is an oriented edge of Γ , then the difference $f_{\partial_0(e)}|_{A_e} - f_{\partial_1(e)}|_{A_e}$ is a constant function on A_e . The assignment

$$e \mapsto f_{\partial_0(e)}|_{A_e} - f_{\partial_1(e)}|_{A_e} \in K$$

is thus a 1-cocycle on Γ , and different choices of the f_v yield 1-cochains which differ by a 1-coboundary. Thus the cohomology class of this cocycle is well-defined, and we have defined a K -linear map

$$\varphi_0: \ker(\varphi_1) \rightarrow H^1(\Gamma, K).$$

Definition 7.4.5. The *monodromy filtration* on $H_{\text{dR}}^1(X/K)$ is defined by

$$M_n H_{\text{dR}}^1(X/K) := \begin{cases} H_{\text{dR}}^1(X/K) & \text{if } n \geq 2, \\ \ker(\varphi_2) & \text{if } n = 1, \\ \ker(\varphi_1) & \text{if } n = 0, \\ 0 & \text{if } n < 0. \end{cases}$$

It will be convenient in what follows to adopt the shorthand that if Γ is a metrised complex of k -curves, then

$$H_{\text{rig}}^1(\Gamma/K)$$

denotes the graded K -vector space given in degrees 0, 1 and 2 by

$$H^1(\Gamma, K), \quad \bigoplus_v H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K), \quad \text{and} \quad H_1(\Gamma, K),$$

respectively, and vanishing in all other degrees. When Γ is the reduction graph of a curve X/K with semistable reduction, the maps φ_2 , φ_1 and φ_0 thus induce a graded K -linear map

$$\varphi_\bullet: \text{gr}_\bullet^M H_{\text{dR}}^1(X/K) \rightarrow H_{\text{rig}}^1(\Gamma/K). \quad (7.4.2.1)$$

Lemma 7.4.6. *The map (7.4.2.1) is the Coleman–Iovita isomorphism of Theorem 7.4.1.*

Proof. Even though this is essentially contained in [DR17, §3.1], we will recap the proof, since our statement in Theorem 7.4.1 does not exactly match the corresponding statement in [DR17], and in any case, we will need to use parts of the proof again shortly. Let $\mathfrak{U} = (W_v)_v$ be a strongly semistable covering of X^{an} , then the de Rham cohomology of X^{an} can be computed as the Čech hypercohomology of the complex $\Omega_{X^{\text{an}}/K}^\bullet$ over the covering $\mathfrak{U}$, i.e. as the cohomology groups of the Čech complex

$$\check{C}^0(\mathfrak{U}, \Omega_{X^{\text{an}}/K}^\bullet) \rightarrow \check{C}^1(\mathfrak{U}, \Omega_{X^{\text{an}}/K}^\bullet) \rightarrow \check{C}^2(\mathfrak{U}, \Omega_{X^{\text{an}}/K}^\bullet).$$

Here, for clarity, we are using the alternating Čech complex of [Sta18a, 01FG], i.e.

$$\begin{aligned} \check{C}^0(\mathfrak{U}, \Omega_{X^{\text{an}}/K}^\bullet) &= \bigoplus_v \mathcal{O}(W_v) \\ \check{C}^1(\mathfrak{U}, \Omega_{X^{\text{an}}/K}^\bullet) &= \bigoplus_v \Omega^1(W_v) \oplus \left(\bigoplus_e \mathcal{O}(A_e) \right)^- \\ \check{C}^2(\mathfrak{U}, \Omega_{X^{\text{an}}/K}^\bullet) &= \left(\bigoplus_e \Omega^1(A_e) \right)^-, \end{aligned}$$

where the direct sums are taken over vertices v or oriented edges e of Γ , and the superscript $(\cdot)^-$ denotes the -1 eigenspace with respect to the involution given by switching the orientation on each edge. So, for example, elements of $\check{C}^1(\mathfrak{U}, \Omega_{X^{\text{an}}/K}^\bullet)$ are represented by tuples $((\omega_v)_v, (f_e)_e)$ where $\omega_v \in \Omega^1(W_v)$ for each vertex v and $f_e \in \mathcal{O}(A_e)$ for each oriented edge e , with $f_{e^{-1}} = -f_e$. The differentials on the Čech complex are given by

$$\begin{aligned} d((f_v)_v) &= ((df_v)_v, (f_{\partial_1(e)}|_{A_e} - f_{\partial_0(e)}|_{A_e})), \\ d((\omega_v)_v, (f_e)_e) &= ((\omega_{\partial_1(e)}|_{A_e} - \omega_{\partial_0(e)}|_{A_e} - df_e)_e). \end{aligned}$$

Now, the subspace $M_0H_{\text{dR}}^1(X/K)$ is the space of cohomology classes $[\omega]$ represented by Čech cocycles $((\omega_v)_v, (f_e)_e)$ where each ω_v is exact, or equivalently the space of cohomology classes which can be represented by Čech cocycles with $\omega_v = 0$ for all v . For all such cohomology classes, f_e is constant for all e , and the map φ_0 sends the cohomology class $[\omega]$ to the class of the map sending an edge e to f_e . It is clear from this description that φ_0 is injective.

Hence the map (7.4.2.1) is injective in all degrees. Since $\dim_K H_{\text{dR}}^1(X/K) = \dim_K H_{\text{rig}}^1(\Gamma/K)$ by [MC10, Proposition 2.34], it is an isomorphism. $\square$

Remark 7.4.7. It follows from Remark 7.2.18 (over $K = \mathbb{C}_\ell$, but this is sufficient for our purposes) that group $H_{\text{rig}}^1(\Gamma/K)$ is independent of the choice of strongly semistable covering. Moreover, one can check that the Coleman–Iovita isomorphism $\varphi_\bullet$ as we have defined it is also independent of the choice of strongly semistable covering.

7.4.3 Compatibility with Poincaré duality

In what follows, it will be important for us to know that the Coleman–Iovita isomorphism is compatible with all of the usual structures on cohomology: Poincaré duality and pullbacks and pushforwards along finite morphisms of curves. We begin by discussing compatibility with Poincaré duality. For this, if Γ is a metrised complex of k -curves, we equip $H_{\text{rig}}^1(\Gamma/K)$ with an antisymmetric pairing by defining

$$\langle \gamma_1 + \sum_v [\omega_{1,v}] + \gamma_1^*, \gamma_2 + \sum_v [\omega_{2,v}] + \gamma_2^* \rangle := \gamma_1^*(\gamma_2) + \sum_v \langle [\omega_{1,v}], [\omega_{2,v}] \rangle - \gamma_2^*(\gamma_1) \quad (7.4.3.1)$$

where γ_1 and γ_2 are homology classes on Γ , $[\omega_{1,v}]$ and $[\omega_{2,v}]$ are rigid cohomology classes on $\bar{\mathcal{X}}_v$, and γ_1^* and γ_2^* are cohomology classes on Γ . The pairing $\langle [\omega_{1,v}], [\omega_{2,v}] \rangle$ on the right-hand side denotes the Poincaré pairing on rigid cohomology.

Compatibility of the Coleman–Iovita isomorphism with Poincaré duality then amounts to the following.

Proposition 7.4.8. *The subspaces $M_0 H_{\text{dR}}^1(X/K)$ and $M_1 H_{\text{dR}}^1(X/K)$ are exact annihilators for the Poincaré pairing on $H_{\text{dR}}^1(X/K)$. Additionally, the Coleman–Iovita isomorphism*

$$\varphi_\bullet : \text{gr}_\bullet^M H_{\text{dR}}^1(X/K) \xrightarrow{\sim} H_{\text{rig}}^1(\Gamma/K)$$

is an isomorphism of K -vector spaces equipped with an antisymmetric pairing (the pairing on the left being induced by Poincaré duality, and the pairing on the right being (7.4.3.1)).

We want to give a description of the Poincaré pairing on $H_{\text{dR}}^1(X/K)$ in terms of Čech hypercohomology, as in the proof of Lemma 7.4.6. For this, we first need to describe the trace map $H_{\text{dR}}^2(X/K) \xrightarrow{\sim} K$. Any class $[\xi] \in H_{\text{dR}}^2(X/K)$

can be represented by a Čech 2-cocycle $(\omega_e)_e \in \check{C}^2(\mathfrak{U}, \Omega_{X^{\text{an}}/K}^\bullet)$, i.e. a tuple consisting of a 1-form $\omega_e \in \Omega^1(A_e)$ for each oriented edge e of Γ , subject to the constraint $\omega_{e^{-1}} = -\omega_e$.

Lemma 7.4.9. *The trace of the class $[\xi] \in H_{\text{dR}}^2(X/K)$ is given by*

$$\text{tr}([\xi]) = \sum_{e \in E(\Gamma)^+} \text{Res}_{A_e}(\omega_e).$$

Remark 7.4.10. The trace map $\text{tr}: H_{\text{dR}}^2(X/K) \rightarrow K$ on algebraic de Rham cohomology is determined, up to scaling by a global constant independent of X , by the requirement that for a finite morphism $f: X \rightarrow X'$ of curves, we have $\text{tr}(f^*[\xi']) = \deg(f) \cdot \text{tr}([\xi'])$ for all $[\xi'] \in H_{\text{dR}}^2(X'/K)$. However, finding the correct choice of normalisation (i.e. global constant) in the literature is surprisingly difficult. (For example, the definition in [Har75] involves composing several isomorphisms between one-dimensional vector spaces, and it is not clear, at least to the authors, how to normalise all of these isomorphisms, at least up to sign.)

The normalisation we will take for the trace map in this paper is as follows. Let $\mathfrak{U}$ be the usual affine covering of $\mathbb{P}_K^1$ by $\mathbb{P}_K^1 \setminus \{\infty\}$ and $\mathbb{P}_K^1 \setminus \{0\}$, and let $\xi \in \check{C}^2(\mathfrak{U}, \Omega_{X/K}^1)$ be the Čech 2-cocycle given by $\frac{dt}{t}$ on $\mathbb{G}_m$. Then we normalise the trace map so that

$$\text{tr}([\xi]) = 1.$$

We remark that this is the choice of normalisation for which, over the complex numbers, if $[\xi]$ is the class of a smooth 2-form η on $X(\mathbb{C})$, then $\text{tr}([\xi]) = \frac{1}{2\pi i} \int_{X(\mathbb{C})} \eta$.

Proof of Lemma 7.4.9. We first observe that the quantity $\sum_e \text{Res}_{A_e}(\omega_e)$ is independent of how we represent $[\xi]$ as a Čech 2-cocycle, and is also independent of the chosen strongly semistable covering by Remark 7.4.7. So the assignment $[\xi] \mapsto \sum_e \text{Res}_{A_e}(\omega_e)$ defines a K -linear map $\psi: H_{\text{dR}}^2(X/K) \rightarrow K$.

In order to prove that ψ is equal to the trace map, it suffices to consider only the case that $K = \mathbb{C}_\ell$. Given a finite morphism $f: X \rightarrow X'$, choose strongly semistable coverings $\mathfrak{U}$ and $\mathfrak{U}'$ of X^{an} and X'^{an} as in Proposition 7.2.22. If $[\xi'] \in H_{\text{dR}}^2(X'/K)$ is represented by a Čech 2-cocycle $(\omega_{e'})_{e'}$, then $f^*[\xi']$ is represented by the Čech 2-cocycle $(\omega_e)_e$ with $\omega_e = (f|_{A_e})^* \omega'_{f(e)}$. So

$$\text{Res}_{A_e}(\omega_e) = d_e(f) \cdot \text{Res}_{A'_{f(e)}}(\omega'_{f(e)}),$$

and so $\psi(f^*[\xi']) = \deg(f) \cdot \psi([\xi'])$. According to Remark 7.4.10, this implies that ψ is equal to the trace map, up to multiplication by some global constant.

To check that this global constant is 1, let $[\xi] \in H_{\text{dR}}^2(\mathbb{P}_K^1/K)$ be the cohomology class defined in Remark 7.4.10. With respect to the semistable covering of $\mathbb{P}_K^{1,\text{an}}$ consisting of the open ball of radius ℓ and the complement of the closed ball of radius 1 centred on 0, $[\xi]$ is represented by the Čech 2-cocycle whose value on the intersection $A(1, \ell)$ is $\frac{dt}{t}$. Thus $\psi([\xi]) = 1$ and we are done. $\square$

As a consequence, we can give a formula for the Poincaré pairing on $H_{\text{dR}}^1(X/K)$.

Lemma 7.4.11. *Let $[\omega_1], [\omega_2] \in H_{\text{dR}}^1(X/K)$ be represented by Čech 1-cocycles $\omega_1 = ((\omega_{1,v})_v, (f_{1,e})_e)$ and $\omega_2 = ((\omega_{2,v})_v, (f_{2,e})_e)$. Then the Poincaré pairing of $[\omega_1]$ and $[\omega_2]$ is given by*

$$\langle [\omega_1], [\omega_2] \rangle = \sum_{e \in E(\Gamma)^+} \text{Res}_{A_e}(f_{1,e}\omega_{2,\partial_1(e)} - f_{2,e}\omega_{1,\partial_0(e)}).$$

Remark 7.4.12. Lemma 7.4.11 contains as a special case the following. Suppose that W_v is a strongly basic wide open, and fix a compactification X_v . A strongly semistable covering of X_v is then given by W_v and the discs glued onto its bounding annuli. The corresponding reduction graph is a star, whose centre is the vertex v corresponding to W_v . If $\omega_1, \omega_2 \in \Omega^1(W_v)$ are differentials with residue 0 on each bounding annulus of W_v , then they determine de Rham cohomology classes on X_v via the exact sequence (7.4.1.1). These cohomology classes $[\omega_i]$ for $i = 1, 2$ are represented by Čech 1-cocycles $((\omega_{i,v'})_{v'}, (f_{i,e})_e)$ where $\omega_{i,v'} = \omega_i$ for $v' = v$ and $\omega_{i,v'} = 0$ otherwise. The edge functions f_e , for edges oriented away from v , satisfy $df_{i,e} = -\omega_i|_{A_e}$, i.e. $-f_{i,e} = \int \omega_i|_{A_e}$ is an antiderivative of $\omega_i|_{A_e}$. Thus Lemma 7.4.11 specialises to the formula

$$\langle [\omega_1], [\omega_2] \rangle = - \sum_e \text{Res}_{A_e} \left(\omega_1 \cdot \int \omega_2 \right) = \sum_e \text{Res}_{A_e} \left(\left(\int \omega_1 \right) \cdot \omega_2 \right)$$

where the sum is taken over the bounding annuli of W_v with their usual orientation, cf. Remark 7.2.10.

Remark 7.4.13. The case of Lemma 7.4.11 when $[\omega_1]$ and $[\omega_2]$ lie in the subspace $M_1 H_{\text{dR}}^1(X/K)$ is stated as [DR17, (109)], though is not proved. Our formula differs from that in [DR17] by a sign; we presume that this is due to differing conventions for how to orient bounding annuli of wide opens, which is not made explicit in [DR17].

Proof of Lemma 7.4.11. Fix a choice of orientation on the graph Γ (i.e. make a choice of one of each pair $\{e, e^{-1}\}$). The cup product of the cocycles ω_1 and ω_2 is the Čech 2-cocycle where

$$(\omega_1 \cup \omega_2)_e = f_{1,e} \omega_{2, \partial_1(e)}|_{A_e} - f_{2,e} \omega_{1, \partial_0(e)}|_{A_e}$$

for each edge e with the chosen orientation¹¹. Hence we have

$$\langle [\omega_1], [\omega_2] \rangle = \text{tr}(\omega_1 \cup \omega_2) = \sum_e \text{Res}_{A_e}(f_{1,e} \omega_{2, \partial_1(e)} - f_{2,e} \omega_{1, \partial_0(e)})$$

using Lemma 7.4.9, and we are done. $\square$

Now we are ready to prove the compatibility of the Coleman–Iovita isomorphism with Poincaré duality.

Proof of Proposition 7.4.8. If $[\omega_1]$ lies in $M_0 H_{\text{dR}}^1(X/K)$, then we may represent it by a Čech 1-cocycle $((\omega_{1,v})_v, (f_{1,e})_e)$ in which $\omega_{1,v} = 0$ for all v , and hence each $f_{1,e}$ is constant. It follows from Lemma 7.4.11 that

$$\begin{aligned} \langle [\omega_1], [\omega_2] \rangle &= \sum_e \text{Res}_{A_e}(f_{1,e} \omega_{2, \partial_1(e)}) \\ &= \sum_e f_{1,e} \text{Res}_{A_e}(\omega_{2, \partial_0(e)}) \\ &= \langle \varphi_0([\omega_1]), \varphi_2([\omega_2]) \rangle \end{aligned}$$

for all $[\omega_2] \in H_{\text{dR}}^1(X/K)$. In particular, if $[\omega_2] \in M_1 H_{\text{dR}}^1(X/K)$, then we have $\langle [\omega_1], [\omega_2] \rangle = 0$ and so the Poincaré pairing vanishes on $M_0 \otimes M_1$. Since M_0 and M_1 have complementary dimensions, they are exact annihilators.

The same formula establishes that the associated graded of the Poincaré pairing agrees with the pairing on $H_{\text{rig}}^1(\Gamma/K)$ in degrees 0 and 2. It remains to deal with the pairing in degree 1. Suppose now that $[\omega_1], [\omega_2] \in M_1 H_{\text{dR}}^1(X/K)$, so they are represented by Čech 1-cocycles with $\text{Res}_{A_e}(\omega_{i, \partial_0(e)}) = 0$ for all edges e and $i = 1, 2$. We may also assume for simplicity that $\omega_{1,v} = 0$ for all vertices v other than a single vertex v_1 . This implies that $f_{1,e}$ is constant for all edges e without v_1 as an endpoint; changing $[\omega_1]$ by an element

¹¹The need to choose an orientation here comes from the fact that this formula is the formula for the cup product of *non-alternating* Čech cochains, cf. [Sta18a, 01FP]. This formula induces the correct cup product on cohomology, independently of the choice of orientation.

in $M_0H_{\text{dR}}^1(X/K)$ (which does not affect the pairing $\langle [\omega_1], [\omega_2] \rangle$), we are free to assume that $f_e = 0$ for all such edges. We may assume the same thing for ω_2 at a vertex v_2 , which may or may not be equal to v_1 .

In the Poincaré pairing formula

$$\langle [\omega_1], [\omega_2] \rangle = \sum_e \text{Res}_{A_e}(f_{1,e}\omega_{2,\partial_1(e)} - f_{2,e}\omega_{1,\partial_0(e)}),$$

the summands on the right-hand side vanish for edges e which do not have v_1 and v_2 as their endpoints. If $v_1 \neq v_2$ are opposite endpoints of e , then the corresponding summand again vanishes (e.g. by orienting e from v_2 to v_1), and hence we see that $\langle [\omega_1], [\omega_2] \rangle = 0$ for $v_1 \neq v_2$. If instead $v_1 = v_2$ is the source of edge e then we have

$$\text{Res}_{A_e}(f_{1,e}\omega_{2,\partial_1(e)} - f_{2,e}\omega_{1,\partial_0(e)}) = -\text{Res}_{A_e}(f_{2,e}\omega_{1,v_1}) = \text{Res}_{A_e}\left(\omega_{1,v_1} \cdot \int \omega_{2,v_1}\right)$$

and so we find

$$\langle [\omega_1], [\omega_2] \rangle = \sum_{\partial_0(e)=v_1} \text{Res}_{A_e}\left(\omega_{1,v_1} \cdot \int \omega_{2,v_1}\right).$$

But by Remark 7.4.12, the right-hand side is equal to the Poincaré pairing of the classes $[\omega_{1,v_1}], [\omega_{2,v_1}] \in H_{\text{rig}}^1(\mathcal{X}_v/K)$ determined by ω_{1,v_1} and ω_{2,v_1} . This shows that $\langle [\omega_1], [\omega_2] \rangle = \langle \varphi_1([\omega_1]), \varphi_1([\omega_2]) \rangle$ for any $[\omega_1], [\omega_2] \in M_1H_{\text{dR}}^1(X/K)$ and so we are done. $\square$

7.4.4 Compatibility with pullbacks and pushforwards

The second property which we need to check – and the more important one – is that the Coleman–Iovita isomorphism is compatible with pullbacks and pushforwards along finite morphisms of curves, and hence that it is compatible with pushforwards along correspondences. If $f: \Gamma \rightarrow \Gamma'$ is a finite harmonic morphism of metrised complexes of k -curves (Definition 7.2.20), then we have already defined in Section 7.2.2 pullback and pushforward maps

$$f^*: H_1(\Gamma', \mathbb{Z}) \rightarrow H_1(\Gamma, \mathbb{Z}) \quad \text{and} \quad f_*: H_1(\Gamma, \mathbb{Z}) \rightarrow H_1(\Gamma', \mathbb{Z})$$

on homology, and similarly on cohomology. Moreover, since f is a morphism of metrised complexes of curves, it comes with a finite morphism $f_v: \mathcal{X}_v \rightarrow \mathcal{X}'_{f(v)}$

of curves over the residue field k for each vertex $v \in V$. Hence there are also pullback and pushforward maps

$$\begin{aligned} f^* &= \bigoplus_v f_v^* : \bigoplus_{v' \in V'} H_{\text{rig}}^1(\bar{\mathcal{X}}_{v'}/K) \rightarrow \bigoplus_{v \in V} H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K), \\ f_* &= \bigoplus_v f_{v,*} : \bigoplus_{v \in V} H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K) \rightarrow \bigoplus_{v' \in V'} H_{\text{rig}}^1(\bar{\mathcal{X}}_{v'}/K). \end{aligned}$$

All in all, these pullback and pushforward maps define graded K -linear maps

$$f^* : H_{\text{rig}}^1(\Gamma'/K) \rightarrow H_{\text{rig}}^1(\Gamma/K) \quad \text{and} \quad f_* : H_{\text{rig}}^1(\Gamma/K) \rightarrow H_{\text{rig}}^1(\Gamma'/K),$$

making $H_{\text{rig}}^1(\Gamma/K)$ both contra- and covariantly functorial in Γ with respect to finite harmonic morphisms. The maps f^* and f_* are adjoint with respect to the Poincaré pairing of (7.4.3.1).

We now carefully prove the following.

Theorem 7.4.14 (Push–pull compatibility of Coleman–Iovita). *Let $f : X \rightarrow X'$ be a finite morphism of smooth projective curves over $\mathbb{C}_\ell$, inducing a finite harmonic morphism $f : \Gamma \rightarrow \Gamma'$ between their reduction graphs as in Theorem 7.2.21. Then the pullback and pushforward maps*

$$f^* : H_{\text{dR}}^1(X'/K) \rightarrow H_{\text{dR}}^1(X/K) \quad \text{and} \quad f_* : H_{\text{dR}}^1(X/K) \rightarrow H_{\text{dR}}^1(X'/K)$$

on de Rham cohomology preserve the monodromy filtration (not necessarily strictly) and are compatible with the Coleman–Iovita isomorphism in the sense that both of the following squares commute.

$$\begin{array}{ccc} \text{gr}_{\bullet}^{\text{M}} H_{\text{dR}}^1(X'/K) & \xrightarrow{\sim} \varphi_{\bullet} & H_{\text{rig}}^1(\Gamma'/K) & \text{gr}_{\bullet}^{\text{M}} H_{\text{dR}}^1(X/K) & \xrightarrow{\sim} \varphi_{\bullet} & H_{\text{rig}}^1(\Gamma/K) \\ \downarrow f^* & & \downarrow f^* & \downarrow f_* & & \downarrow f_* \\ \text{gr}_{\bullet}^{\text{M}} H_{\text{dR}}^1(X/K) & \xrightarrow{\sim} \varphi_{\bullet} & H_{\text{rig}}^1(\Gamma/K) & \text{gr}_{\bullet}^{\text{M}} H_{\text{dR}}^1(X'/K) & \xrightarrow{\sim} \varphi_{\bullet} & H_{\text{rig}}^1(\Gamma'/K) \end{array} \quad (7.4.4.1)$$

For the proof, fix a strongly semistable vertex set $V' \subset |X'^{\text{an}}|_{\text{II}}$ of size ≥ 2 satisfying the conditions of Theorem 7.2.21, so that V' and $V = f^{-1}V'$ determine semistable coverings of X'^{an} and X^{an} . Per Proposition 7.2.22, for any bounding annulus $A'_{e'}$ in X'^{an} , the inverse image of $A'_{e'}$ in X^{an} is

$$f^{-1}A'_{e'} = \bigcup_{e \in f^{-1}(e')} A_e,$$

and the restriction of f to A_e is a finite morphism $A_e \rightarrow A'_{e'}$ of degree $d_e(f)$. According to [BPR13, Proposition 2.2], if t' and t are parameters on the annuli $A'_{e'}$ and A_e , then the restriction $f|_{A_e}$ is given by $t' = \alpha t^{d_e(f)}(1+g(t))$ where $g \in \mathcal{O}(A_e)$ is an analytic function with norm < 1 at all points of A_e . So

$$f^* \frac{dt'}{t'} \Big|_{A_e} = d \log \left(\alpha t^{d_e(f)} (1 + g(t)) \right) = d_e(f) \frac{dt}{t} + d \log(1 + g(t)).$$

Since $\log(1 + g(t))$ is convergent on the open annulus A_e , it follows that

$$\text{Res}_{A_e} \left(f^* \frac{dt'}{t'} \right) = d_e(f).$$

Since the class of $\frac{dt'}{t'}$ is a basis of $H_{\text{dR}}^1(A'_{e'}/K)$, it follows that

$$\text{Res}_{A_e}(f^* \omega') = d_e(f) \text{Res}_{A'_{e'}}(\omega')$$

for all $\omega' \in \Omega^1(A'_{e'})$. Fixing $[\omega'] \in H_{\text{dR}}^1(X'/K)$ and summing over edges of Γ gives

$$\begin{aligned} \varphi_2(f^* \omega') &= \sum_e \text{Res}_{A_e}(f^* \omega') \cdot e \\ &= \sum_{e'} \sum_{e \in f^{-1}(e')} d_e(f) \text{Res}_{A'_{e'}}(\omega') \cdot e = f^* \varphi_2(\omega') \end{aligned}$$

using Proposition 7.3.2. This tells us that $f^*(M_1 H_{\text{dR}}^1(X'/K)) \subseteq M_1 H_{\text{dR}}^1(X/K)$, and that the left-hand square in (7.4.4.1) commutes in degree 2.

Next, for any 1-form ω' on $A'_{e'}$, we have $(f|_{A_e})_*(f|_{A_e})^* \omega' = d_e(f) \omega'$ for all $e \in f^{-1}(e')$ and so

$$\text{Res}_{A'_{e'}}((f|_{A_e})_*(f|_{A_e})^* \omega') = d_e(f) \text{Res}_{A'_{e'}}(\omega') = \text{Res}_{A_e}((f|_{A_e})^* \omega').$$

This implies that $\text{Res}_{A'_{e'}}((f|_{A_e})_* \omega) = \text{Res}_{A_e}(\omega)$ for all 1-forms ω on A_e , and so we have

$$\begin{aligned} \varphi_2(f_* \omega) &= \sum_{e'} \text{Res}_{A'_{e'}}(f_* \omega) \cdot e' \\ &= \sum_{e'} \text{Res}_{A'_{e'}} \left(\sum_{e \in f^{-1}(e')} (f|_{A_e})_* \omega \right) \cdot e' \\ &= \sum_{e'} \sum_{e \in f^{-1}(e')} \text{Res}_{A_e}(\omega) \cdot e' = f_* \varphi_2(\omega). \end{aligned}$$

using Proposition 7.3.2. So $f_*(M_1 H_{\text{dR}}^1(X/K)) \subseteq M_1 H_{\text{dR}}^1(X'/K)$ and the right-hand square in (7.4.4.1) commutes in degree 2.

Next, for $[\omega'] \in M_1 H_{\text{dR}}^1(X'/K)$ represented by a Čech 1-cocycle $((\omega'_{v'})_{v'}, (f'_{e'})_{e'})$ we know that the inverse image of each basic wide open $W'_{v'}$ in X'^{an} is

$$f^{-1}W'_{v'} = \bigcup_{v \in f^{-1}(v')} W_v.$$

For each $v \in f^{-1}(v')$, the pullback $(f|_{W_v})^* \omega'_{v'}$ is a 1-form on W_v with residue zero on each bounding annulus. According to Proposition 7.2.22 (combined with [LS07, Proposition 8.2.10]), the class of $(f|_{W_v})^* \omega'_{v'}$ in rigid cohomology is the pullback of the class of $\omega'_{v'}$ along the map $f_v: \mathcal{X}_v \rightarrow \mathcal{X}'_{v'}$. Taking this over all v implies that

$$\varphi_1(f^*[\omega']) = f^* \varphi_1([\omega']),$$

which tells us that $f^*(M_0 H_{\text{dR}}^1(X'/K)) \subseteq M_0 H_{\text{dR}}^1(X/K)$, and that the left-hand square in (7.4.4.1) commutes in degree 1.

The remaining cases follow by Poincaré duality. For example, since pullback and pushforward are adjoint under the Poincaré pairing on $H_{\text{dR}}^1(X/K)$, it follows from the fact that f^* preserves M_1 that f_* preserves its annihilator M_0 . And then for any class $[\omega] \in M_1 H_{\text{dR}}^1(X/K)$ we have

$$\begin{aligned} \langle \varphi_1(f_*[\omega]), \varphi_1([\omega']) \rangle &= \langle f_*[\omega], [\omega'] \rangle = \langle [\omega], f^*[\omega'] \rangle = \langle \varphi_1([\omega]), \varphi_1(f^*[\omega']) \rangle \\ &= \langle \varphi_1([\omega]), f^*(\varphi_1([\omega'])) \rangle = \langle f_*(\varphi_1([\omega])), \varphi_1([\omega']) \rangle. \end{aligned}$$

Since the Poincaré pairing is perfect, this implies that $\varphi_1(f_*[\omega]) = f_*(\varphi_1([\omega]))$, and hence the right-hand square in (7.4.4.1) commutes in degree 1.

A similar argument shows that the squares in (7.4.4.1) commute in degree 0, completing the proof of Theorem 7.4.14. $\square$

7.4.4.1 Compatibility with correspondences

If $Z \subset X \times X'$ is a correspondence from X to X' , having reduction graphs Γ and Γ' , respectively, then we can define a pushforward

$$Z_*: H_{\text{rig}}^1(\Gamma/K) \rightarrow H_{\text{rig}}^1(\Gamma'/K)$$

in the usual way, namely take the composite $\tilde{\pi}_{2,*} \circ \tilde{\pi}_1^*$ where $\tilde{\pi}_1: \tilde{Z} \rightarrow X$ and $\tilde{\pi}_2: \tilde{Z} \rightarrow X'$ are the projections from the normalisation. It is a formal

consequence of Theorem 7.4.14 that the Coleman–Iovita isomorphism is also compatible with pushforward along correspondences.

In our computations, we will use this to compute the action of Z_* on $H_1(\Gamma, \mathbb{Z})$, as well as the vertex traces $\mathrm{tr}_v(Z)$. Once we know the action of Z_* on the space $H_{\mathrm{dR}}^1(X/K)$, we can use the Coleman–Iovita isomorphism to determine the action on $H_{\mathrm{rig}}^1(\Gamma/K)$. In degree 2, we read off the action of Z_* on $H_1(\Gamma, K)$, while in degree 1 we read off the action on $\bigoplus_v H_{\mathrm{rig}}^1(\mathcal{X}_v/K)$. Taking an appropriate block in the matrix representing this action and taking the trace, we obtain the vertex traces $\mathrm{tr}_v(Z)$ by Lemma 7.3.4.

7.5 Explicit Coleman–Iovita for hyperelliptic curves

We now specialialise to the case of hyperelliptic curves in odd residue characteristic, and begin to explain what the Coleman–Iovita isomorphism means explicitly. In this section we first explain background on cluster pictures. We then give a description of how to obtain a strongly semistable cover of a hyperelliptic curve over a finite extension $K/\mathbb{Q}_\ell$ in odd residue characteristic. We provide a definition of the bases of $H_1(\Gamma, \mathbb{Z})$, $H_{\mathrm{dR}}^1(X/K)$, and $H_{\mathrm{rig}}^1(\mathcal{X}_v/K)$ that we will use and give an explicit description of the Coleman–Iovita isomorphism with respect to these bases. We also explicitly describe the Berkovich skeleton of $X_{C_\ell}^{\mathrm{an}}$ using the semistable covering of X . We use this description to give a sufficient condition for when a component of the associated semistable model $\mathcal{X}$ has no $\mathbb{Q}_\ell$ -points reducing to it.

7.5.1 A semistable covering

Let $\ell > 2$ be a prime. Let $\pi : X \rightarrow \mathbb{P}^1$ be a nice hyperelliptic curve defined over a finite extension $K/\mathbb{Q}_\ell$. In the standard affine patch, suppose X is given by an equation $y^2 = f(x)$ and $\deg f \geq 3$, with f separable. Enlarging K if necessary, we may assume that K contains the biquadratic extension of the field generated by $\mathcal{R}$, the set of roots of f in $\bar{K}$. In this section, we use cluster theory [DDMM23, BBB⁺22] to write down a semistable covering of X^{an} without explicitly finding a semistable model of X/K .

Following [DDMM23] we define the following terms.

Definition 7.5.1 (Cluster vocabulary).

- A *cluster* $\mathfrak{s}$ is a non-empty subset of the roots of f cut out by a disc in

$\mathbb{A}^{1,\text{an}}$. That is, $\mathfrak{s} = D \cap \mathcal{R}$ for some open or closed disc $D \subset \mathbb{A}^{1,\text{an}}$.

- A cluster $\mathfrak{s}$ is *proper* if $\#\mathfrak{s} > 1$. It is *odd* if $\#\mathfrak{s}$ is odd and *even* if $\#\mathfrak{s}$ is even.
- The *top* cluster is $\mathcal{R}$ itself.
- A cluster $\mathfrak{s}'$ is a *child* of $\mathfrak{s}$ (and $\mathfrak{s}$ is the *parent* of $\mathfrak{s}'$), denoted by $\mathfrak{s}' < \mathfrak{s}$, if $\mathfrak{s}'$ is a maximal subcluster of $\mathfrak{s}$ (not equal to $\mathfrak{s}$ itself). We write $\mathfrak{s}'' \leq \mathfrak{s}$ to denote that $\mathfrak{s}'$ is a child or equal to $\mathfrak{s}$. For any non-top cluster we denote by $P(\mathfrak{s})$ the parent of $\mathfrak{s}$.
- An even cluster $\mathfrak{s}$ is called *übereven* if every child of $\mathfrak{s}$ is also even.
- The *depth* of a proper cluster $\mathfrak{s}$ is

$$d_{\mathfrak{s}} := \min_{\alpha_1, \alpha_2 \in \mathfrak{s}} v(\alpha_1 - \alpha_2).$$

If $\mathfrak{s}$ is not the top cluster, then the *relative depth* of $\mathfrak{s}$ is $\delta_{\mathfrak{s}} := d_{\mathfrak{s}} - d_{P(\mathfrak{s})}$.

- Let $\mathfrak{s} \wedge \mathfrak{s}'$ be the smallest cluster containing both $\mathfrak{s}$ and $\mathfrak{s}'$.
- Define the invariant

$$\nu_{\mathfrak{s}} := v(c) + \sum_{r \in \mathcal{R}} d_{\{r\} \wedge \mathfrak{s}}$$

for each proper cluster $\mathfrak{s}$, where c is the leading coefficient of f .


Remark 7.5.2. Our depths differ from those in [DDMM23] by a multiplicative constant, because they normalise their valuation to have valuation group $\mathbb{Z}$, while we normalise $v(\ell) = 1$.

Definition 7.5.3 (Open discs and annuli associated to clusters). For a cluster $\mathfrak{s}$, we define the following subsets of $\mathbb{P}_K^{1,\text{an}}$.

- Let $D_{\mathfrak{s}}^c$ be the smallest closed disc in $\mathbb{A}_K^{1,\text{an}}$ such that $\mathfrak{s} = D_{\mathfrak{s}}^c \cap \mathcal{R}$.
- For $\mathfrak{s}$ not the top cluster, let $D_{\mathfrak{s}}^o$ be the largest open disc in $\mathbb{A}_K^{1,\text{an}}$ such that $\mathfrak{s} = D_{\mathfrak{s}}^o \cap \mathcal{R}$. If $\mathfrak{s}$ is the top cluster we set $D_{\mathfrak{s}}^o = \mathbb{P}_K^{1,\text{an}}$.
- For $\mathfrak{s}$ not the top cluster, define the open annulus $A_{\mathfrak{s}} := D_{\mathfrak{s}}^o \setminus D_{\mathfrak{s}}^c$.
- Define

$$U_{\mathfrak{s}} := D_{\mathfrak{s}}^o \setminus \bigcup_{\substack{\mathfrak{s}' < \mathfrak{s} \\ \mathfrak{s}' \text{ proper}}} D_{\mathfrak{s}'}^c \subseteq \mathbb{P}_K^{1,\text{an}}.$$

Proposition 7.5.4. *The analytic spaces $U_{\mathfrak{s}}$ constitute a semistable covering of $\mathbb{P}_K^{1,\text{an}}$.*

Figure 7.3: A picture of the analytic space $U_{\mathfrak{s}}$.

Proof. Given $x \in \mathbb{A}_K^{1,\text{an}}(K)$, let D^c be the smallest closed disc containing x and at least two roots of f . Let $\mathfrak{s} := D^c \cap \mathcal{R}$, which is a proper cluster. Then $x \in D_{\mathfrak{s}}^o$. Indeed, since $D^c \cap D_{\mathfrak{s}}^o \neq \emptyset$, we have either $D^c \subset D_{\mathfrak{s}}^o$ and so $x \in D_{\mathfrak{s}}^o$, or $D_{\mathfrak{s}}^o \subset D^c$. In the latter case, there is a larger open disc $D_{\mathfrak{s}}^o \subsetneq D^o \subset D^c$ with $D^o \cap \mathcal{R} = \mathfrak{s}$ contradicting the maximality of $D_{\mathfrak{s}}^o$. By minimality of D^c , for any proper child $\mathfrak{s}' < \mathfrak{s}$, we have $x \notin D_{\mathfrak{s}'}^c$. So $x \in U_{\mathfrak{s}}$ and hence the $U_{\mathfrak{s}}$ cover $\mathbb{A}_K^{1,\text{an}}$.

To show the covering is semistable, we first note that for a cluster $U_{\mathfrak{s}}$, the pair

$$\left(U_{\mathfrak{s}}, D_{\mathfrak{s}}^c \setminus \bigcup_{\substack{\mathfrak{s}' < \mathfrak{s} \\ \mathfrak{s}' \text{ proper}}} D_{\mathfrak{s}'}^o \right)$$

forms a basic wide open. Secondly, the intersection $U_{\mathfrak{s}} \cap U_{\mathfrak{s}'}$ is by construction the annulus $A_{\mathfrak{s}'}$ if $\mathfrak{s}' < \mathfrak{s}$, and is empty if $\mathfrak{s} \not\leq \mathfrak{s}'$ and $\mathfrak{s}' \not\leq \mathfrak{s}$. Hence all double intersections are annuli or empty, and all triple intersections are empty. $\square$

This semistable covering also appears in [Sto19], where it is used to give uniform bounds for rational points on hyperelliptic curves of small Mordell–Weil rank and [KK22, Section 4] where it plays a role in defining a p -adic integral (where $p = \ell$) on bad reduction hyperelliptic curves. Here, we tie it to cluster pictures, as speculated about in [Kay21, Chapter 4].

We will repeatedly make use of the following lemma to describe equations for X restricted to annuli and wide opens of the semistable covering.

Lemma 7.5.5 ([KK22, Lemma 4.10]). *Let D^o be an open disc in $\mathbb{A}_K^{1,\text{an}}$ and let $D_1^c, \dots, D_n^c$ be pairwise disjoint closed subdiscs. Set $U := D^o \setminus \cup_i D_i^c$. Suppose that $h(x) \in K[x]$ is a polynomial such that*

- (a) $h(x)$ has no zeros in U , and
 (b) $h(x)$ has an even number of zeros in each D_i^c .

Let K' be the biquadratic extension of K . Then there is an invertible rigid analytic function $h(x)^{1/2} \in \mathcal{O}(U_{K'})^{\text{an}}$ whose square is $h(x)$. Moreover $h(x)^{1/2}$ is explicitly computable.

Proof. Write $h(x) = C \cdot \prod_{D_k^c} (\prod_{i=1}^{2s_k} (x - \alpha_i)) \cdot \prod_{j=1}^t (x - \beta_j)$ where $\alpha_i \in D_k^c$, $\beta_j \notin D^o$, and $C \in K$. Then

$$\left(\prod_{i=1}^{2s_k} (x - \alpha_i) \right)^{1/2} = (x - \alpha_1)^s \prod_{i=1}^{2s_k} \left(1 - \frac{\alpha_i - \alpha_1}{x - \alpha_1} \right)^{1/2}$$

converges outside of D_k^c for all $1 \leq k \leq n$. Furthermore,

$$\begin{aligned} \left(C \cdot \prod_{j=1}^t (x - \beta_j) \right)^{1/2} &= C^{1/2} \cdot \prod_{i=1}^t ((x - \beta_1) - (\beta_i - \beta_1))^{1/2} \\ &= C^{1/2} \cdot \prod_{i=1}^t \left((\beta_1 - \beta_i)^{1/2} \left(1 - \frac{x - \beta_1}{\beta_1 - \beta_i} \right)^{1/2} \right) \end{aligned}$$

converges on D^o . The product of the square roots is $h(x)^{1/2}$ and belongs to $\mathcal{O}(U_{K'})^{\text{an}}$. $\square$

Definition 7.5.6. We fix an ordering on $\mathcal{R}$. For every odd cluster $\mathfrak{s}$ (including singletons), choose the smallest element $\alpha_{\mathfrak{s}} \in \mathfrak{s}$. For a non-singleton cluster $\mathfrak{s}$, we define

$$g_{\mathfrak{s}}(x) := \prod_{\substack{\mathfrak{s}' < \mathfrak{s} \\ \mathfrak{s}' \text{ odd}}} (x - \alpha_{\mathfrak{s}'}) \quad \text{and} \quad h_{\mathfrak{s}}(x) := f(x)/g_{\mathfrak{s}}(x).$$

For a singleton cluster $\mathfrak{s}$ we define $g_{\mathfrak{s}}(x) := x - \alpha_{\mathfrak{s}}$ and $h_{\mathfrak{s}}(x) := f(x)/g_{\mathfrak{s}}(x)$.

Let $X_{\mathfrak{s}}$ be the projective hyperelliptic curve defined by the affine equation $u_{\mathfrak{s}}^2 = g_{\mathfrak{s}}(x)$.

Remark 7.5.7. If $\mathfrak{s}$ is übereven then $g_{\mathfrak{s}} = 1$ and so $X_{\mathfrak{s}} \simeq \mathbb{P}_K^{1,\text{an}} \sqcup \mathbb{P}_K^{1,\text{an}}$ is a disconnected hyperelliptic curve. This will not be a problem.

We can now describe equations for X restricted to subsets of the semistable covering and annuli.

Proposition 7.5.8. *For any cluster $\mathfrak{s}$, we have that $X|_{U_{\mathfrak{s}}}$ and $X_{\mathfrak{s}}|_{U_{\mathfrak{s}}}$ are isomorphic over $U_{\mathfrak{s}}$.*

Proof. By construction, the polynomial $h_{\mathfrak{s}}(x)$ has no roots in $U_{\mathfrak{s}}$ and an even number of roots in $D_{\mathfrak{s}}^c$, for each proper child $\mathfrak{s}' < \mathfrak{s}$. So it has a square root $h_{\mathfrak{s}}(x)^{1/2}$ by Lemma 7.5.5. The map

$$\begin{aligned} X|_{U_{\mathfrak{s}}} &\rightarrow X|_{U_{\mathfrak{s}}} \\ (x, u_{\mathfrak{s}}) &\mapsto (x, u_{\mathfrak{s}} \cdot h_{\mathfrak{s}}(x)^{1/2}) \end{aligned} \quad (7.5.1.1)$$

is the desired isomorphism. $\square$

Lemma 7.5.9. *If $\mathfrak{s}$ is a proper, non-top cluster, then $X|_{A_{\mathfrak{s}}}$ is isomorphic to*

1. *a disjoint union of two open annuli of width $\delta_{\mathfrak{s}}$, each mapping isomorphically to $A_{\mathfrak{s}}$ if $\mathfrak{s}$ is even;*
2. *an open annulus of width $\delta_{\mathfrak{s}}/2$ mapping to $A_{\mathfrak{s}}$ via the squaring map if $\mathfrak{s}$ is odd.*

Proof. If $\mathfrak{s}$ is an even cluster, then $g_{\mathfrak{s}}(x)^{1/2}h_{\mathfrak{s}}(x)^{1/2}$ is a square root of $f(x)$ on the annulus $A_{\mathfrak{s}}$, and so gives a splitting

$$\begin{aligned} A_{\mathfrak{s}} \times \{\pm 1\} &\simeq X|_{A_{\mathfrak{s}}} \\ (x, \varepsilon) &\mapsto (x, \varepsilon \cdot g_{\mathfrak{s}}(x)^{1/2}h_{\mathfrak{s}}(x)^{1/2}) \end{aligned} \quad (7.5.1.2)$$

of X over $A_{\mathfrak{s}}$. Otherwise, if $\mathfrak{s}$ is odd, then $\left(\frac{g_{\mathfrak{s}}(x)}{x - \alpha_{\mathfrak{s}}}\right)^{1/2}h_{\mathfrak{s}}(x)^{1/2}$ is a square root of $f(x)/(x - \alpha_{\mathfrak{s}})$ on $A_{\mathfrak{s}}$ and so, gives an isomorphism

$$(\text{curve given by } t_{\mathfrak{s}}^2 = (x - \alpha_{\mathfrak{s}}))|_{A_{\mathfrak{s}}} \simeq X|_{A_{\mathfrak{s}}}. \quad (7.5.1.3)$$

The curve $(t_{\mathfrak{s}}^2 = x - \alpha_{\mathfrak{s}})|_{A_{\mathfrak{s}}}$ is isomorphic to an annulus of width $\delta_{\mathfrak{s}}/2$ via projection onto the $t_{\mathfrak{s}}$ -coordinate. $\square$

Since the isomorphisms above depend on a choice of square root of $h_{\mathfrak{s}}$ and $g_{\mathfrak{s}}$ we fix these choices once and for all, as per the following definition.

Definition 7.5.10 (Choice of square roots). We fix the following data:

- for every cluster $\mathfrak{s}$, a square root $h_{\mathfrak{s}}(x)^{1/2}$ of $h_{\mathfrak{s}}$ on $U_{\mathfrak{s}}$.
- for every even cluster $\mathfrak{s}$ (not the top cluster), a square root $g_{\mathfrak{s}}(x)^{1/2}$ on the annulus $A_{\mathfrak{s}}$;

- for every odd cluster $\mathfrak{s}$ (not the top cluster), a square root $\left(\frac{g_{\mathfrak{s}}(x)}{(x-\alpha_{\mathfrak{s}})}\right)^{1/2}$ on the annulus $A_{\mathfrak{s}}$.

Our choices of $g_{\mathfrak{s}}(x)^{1/2}$ and $h_{\mathfrak{s}}(x)^{1/2}$ give trivialisations

$$\begin{aligned} X|_{U_{\mathfrak{s}}} &\simeq U_{\mathfrak{s}} \times \{\pm 1\} && \text{for } \mathfrak{s} \text{ \u00fcbereven;} \\ X|_{A_{\mathfrak{s}}} &\simeq A_{\mathfrak{s}} \times \{\pm 1\} && \text{for } \mathfrak{s} \text{ even, not the top cluster.} \end{aligned}$$

Definition 7.5.11 (Notation for annuli and semistable covering). We write $\tilde{U}_{\mathfrak{s}}$ for $X|_{U_{\mathfrak{s}}}$ when $\mathfrak{s}$ is not \u00fcbereven. Otherwise, write $\tilde{U}_{\mathfrak{s}}^{\pm}$ for the two irreducible components of $X|_{U_{\mathfrak{s}}^{\pm}}$. We let $\tilde{A}_{\mathfrak{s}}$ denote $X|_{A_{\mathfrak{s}}}$, and $\tilde{A}_{\mathfrak{s}}^{\pm}$ the two different annuli if $\mathfrak{s}$ is even. If we write $\tilde{U}_{\mathfrak{s}}^{\pm}$ or $\tilde{A}_{\mathfrak{s}}^{\pm}$ for a non-\u00fcbereven or odd cluster, we mean $\tilde{U}_{\mathfrak{s}}$ or $\tilde{A}_{\mathfrak{s}}$.

We described what X looks like over each $U_{\mathfrak{s}}$ and $A_{\mathfrak{s}}$. We want to understand how to glue these together compatibly over pairwise intersections of $U_{\mathfrak{s}}$. For this, we ask the following compatibility condition on the trivialisations.

Definition 7.5.12 (\u00dcbereven compatibility criterion). We assume a compatibility condition between the various inverse square roots of $g_{\mathfrak{s}}$ and $h_{\mathfrak{s}}$, which takes place over the \u00fcbereven clusters. Namely, we suppose the following.

- If $\mathfrak{s}$ is \u00fcbereven and not the top cluster, then $g_{\mathfrak{s}}(x)^{1/2} = 1$. (Since $g_{\mathfrak{s}}(x) = 1$ here, this says we take the obvious choice of square root.)
- If $\mathfrak{s}$ is \u00fcbereven and $\mathfrak{s}'$ is a child of $\mathfrak{s}$, then

$$h_{\mathfrak{s}}(x)^{1/2}|_{A_{\mathfrak{s}'}} = g_{\mathfrak{s}'}(x)^{1/2}h_{\mathfrak{s}'}(x)^{1/2}|_{A_{\mathfrak{s}'}}.$$

(Both sides are square roots of f so we are again fixing a sign.)

It is always possible to make these conditions hold after possibly changing some of the $h_{\mathfrak{s}}(x)^{1/2}$ by a sign.

When $\mathfrak{s}$ is \u00fcbereven, for every $\mathfrak{s}' < \mathfrak{s}$, we know that $U_{\mathfrak{s}}^{\pm} \supset A_{\mathfrak{s}'}^{\pm}$, and $U_{\mathfrak{s}}^{\pm} \supset A_{\mathfrak{s}}^{\pm}$ if $\mathfrak{s}$ is not the top cluster. The compatibility conditions ensure that the trivialisations of $X|_{U_{\mathfrak{s}}}$ and $X|_{A_{\mathfrak{s}}}$ are compatible with one another (i.e., that $A_{\mathfrak{s}}^{\pm} \subset U_{\mathfrak{s}}^{\pm}$ instead of $A_{\mathfrak{s}}^{\mp} \subset U_{\mathfrak{s}}^{\pm}$).

The following lemma will be used in the proof of Theorem 7.5.15.

Lemma 7.5.13. *Let $\mathcal{R}'$ be $\mathcal{R}$ if $\deg(f)$ is even and $\mathcal{R} \cup \{\infty\}$ if $\deg(f)$ is odd. So $\mathcal{R}'$ is the set of ramification points of $X \rightarrow \mathbb{P}_K^1$. Let $D \subset \mathbb{P}_{\mathbb{C}_{\ell}}^{1,\text{an}}$ be an open disc (i.e. either an open disc in $\mathbb{A}_{\mathbb{C}_{\ell}}^{1,\text{an}}$ or the complement in $\mathbb{P}_{\mathbb{C}_{\ell}}^{1,\text{an}}$ of a*

closed disc in $\mathbb{A}_{\mathbb{C}_\ell}^{1,\text{an}}$) containing at most one ramification point. Then $X_{\mathbb{C}_\ell}|_D$ is isomorphic to

1. a disjoint union of two open discs, each mapping isomorphically to D if $\#D \cap \mathcal{R}' = 0$;
2. an open disc, mapping to D by the squaring map if $\#D \cap \mathcal{R}' = 1$.

Proof. The proof is analogous to the proof of Lemma 7.5.9. $\square$

The main theorem in this section is the description of a semistable covering of X^{an} .

Theorem 7.5.14. *Let $\pi : X^{\text{an}} \rightarrow \mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}}$ be the map induced by $\pi : X \rightarrow \mathbb{P}^1$. Then $\mathfrak{U} := \{\tilde{U}_\mathfrak{s}^\pm\}$ is a semistable covering of X^{an} (where $\tilde{U}_\mathfrak{s}^\pm = \tilde{U}_\mathfrak{s}$ when $\mathfrak{s}$ is not $\ddot{\text{u}}\text{bereven}$).*

Proof. Note that the canonical reduction of $\tilde{U}_\mathfrak{s}$ is a non-empty open inside the curve $y^2 = g_\mathfrak{s}(x)$. This automatically implies that $\tilde{U}_\mathfrak{s}$ is irreducible if $\mathfrak{s}$ is not $\ddot{\text{u}}\text{bereven}$, and splits as a disjoint union of two irreducibles if $\mathfrak{s}$ is $\ddot{\text{u}}\text{bereven}$. We see that for any $\mathfrak{s}$, the inverse image $\tilde{U}_{\mathfrak{s},0}^\pm$ of $D_\mathfrak{s}^c \setminus \bigcup_{\substack{\mathfrak{s}' < \mathfrak{s} \\ \mathfrak{s}' \text{ proper}}} D_{\mathfrak{s}'}^o \subset U_\mathfrak{s}$ inside $\tilde{U}_\mathfrak{s}^\pm$ makes $(\tilde{U}_\mathfrak{s}^\pm, \tilde{U}_{\mathfrak{s},0}^\pm)$ a basic wide open. Double intersections $\tilde{U}_\mathfrak{s} \cap \tilde{U}'_\mathfrak{s}$ are empty, or one of the annuli $\tilde{A}_\mathfrak{s}$ or $\tilde{A}'_{\mathfrak{s}'}$, and triple intersections are empty. $\square$

7.5.2 The minimal skeleton of a hyperelliptic curve

We can now write down the minimal skeleton of $X_{\mathbb{C}_\ell}^{\text{an}}$. The results here give an interpretation of the results in [DDMM23, Section 8] in the language of Berkovich spaces.

For a proper cluster $\mathfrak{s}$, let $\eta_\mathfrak{s} \in |\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}}|$ denote the Gauss point of the disc $D_\mathfrak{s}^c$. This is a type II point. If $\mathfrak{s}$ is not the top cluster, let $e_\mathfrak{s} \subset |\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}}|$ be the skeleton of the annulus $A_{\mathfrak{s},\mathbb{C}_\ell}$. This is an open interval of length equal to the width of $A_\mathfrak{s}$ (i.e. the relative depth $\delta_\mathfrak{s}$) embedded isometrically into $|\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}}|$. Moreover, the skeleton $e_\mathfrak{s}$ is canonically oriented where $\partial_0(e_\mathfrak{s}) = \eta_\mathfrak{s}$ and $\partial_1(e_\mathfrak{s}) = \eta_{P(\mathfrak{s})}$ are the “left and right hand limit points”.

Theorem 7.5.15. *Let $\pi : X \rightarrow \mathbb{P}^1$ be the projection, inducing the map $\pi : |X_{\mathbb{C}_\ell}^{\text{an}}| \rightarrow |\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}}|$ on the underlying topological spaces of Berkovich spaces. Then (a) for a proper cluster $\mathfrak{s}$, the preimage of $\eta_\mathfrak{s}$ in $|X_{\mathbb{C}_\ell}^{\text{an}}|$ consists of:*

- two type II points $\tilde{\eta}_\mathfrak{s}^+$ and $\tilde{\eta}_\mathfrak{s}^-$ of genus 0 if $\mathfrak{s}$ is $\ddot{\text{u}}\text{bereven}$;




Figure 7.4: Cluster, Berkovich skeleton, special fibre, and reduction map

- one type II point $\tilde{\eta}_{\mathfrak{s}}$ of genus $g(\mathfrak{s}) := \lceil \frac{\#\text{odd children of } \mathfrak{s}}{2} - 1 \rceil$ otherwise.
- (b) For a proper non-top cluster $\mathfrak{s}$, the preimage of $e_{\mathfrak{s}}$ in $|X_{\mathbb{C}_\ell}^{\text{an}}|$ consists of
- two embedded oriented intervals $\tilde{e}_{\mathfrak{s}}^+$ and $\tilde{e}_{\mathfrak{s}}^-$ of length $\delta_{\mathfrak{s}}$, each mapping oriented-isomorphically onto $e_{\mathfrak{s}}$ with $\partial_0(\tilde{e}_{\mathfrak{s}}^{\pm}) = \tilde{\eta}_{\mathfrak{s}}^{\pm}$ and $\partial_1(\tilde{e}_{\mathfrak{s}}^{\pm}) = \tilde{\eta}_{P(\mathfrak{s})}^{\pm}$ if $\mathfrak{s}$ is even. (We permit ourselves to write $\tilde{\eta}_{\mathfrak{s}}^{\pm} = \tilde{\eta}_{\mathfrak{s}}$ if $\mathfrak{s}$ is not übereven.)
 - one embedded oriented interval $\tilde{e}_{\mathfrak{s}}$ of length $\delta_{\mathfrak{s}}/2$ mapping onto $e_{\mathfrak{s}}$ by an oriented dilation of scale factor 2, with $\partial_0(\tilde{e}_{\mathfrak{s}}) = \tilde{\eta}_{\mathfrak{s}}$ and $\partial_1(\tilde{e}_{\mathfrak{s}}) = \tilde{\eta}_{P(\mathfrak{s})}$ if $\mathfrak{s}$ is odd.
- (c) The points $\tilde{\eta}_{\mathfrak{s}}^{\pm}$ for $\mathfrak{s}$ a proper cluster form a semistable vertex set V for $X_{\mathbb{C}_\ell}^{\text{an}}$, whose associated skeleton Γ is the union of the $\{\tilde{\eta}_{\mathfrak{s}}^{\pm}\}$ and the $\tilde{e}_{\mathfrak{s}}^{\pm}$. Moreover, the choices of $\tilde{\eta}_{\mathfrak{s}}^{\pm}$ for $\mathfrak{s}$ übereven and of $\tilde{e}_{\mathfrak{s}}^{\pm}$ for $\mathfrak{s}$ even are canonical given our choices of inverse square roots (i.e. our choices of inverse square roots allow us to distinguish between $\tilde{\eta}_{\mathfrak{s}}^+$ and $\tilde{\eta}_{\mathfrak{s}}^-$).

Remark 7.5.16. One can check that Γ is the minimal skeleton of $X_{\mathbb{C}_\ell}^{\text{an}}$. However, V is not necessarily a minimal vertex set for $X_{\mathbb{C}_\ell}^{\text{an}}$. A minimal vertex set is given by the points $\tilde{\eta}_{\mathfrak{s}}^{\pm}$ for proper non-top clusters $\mathfrak{s}$ except twins (clusters of size 2), cotwins (non-übereven clusters with a child of size 2 times the genus of X), as well as the top cluster if it has ≥ 3 children. See [DDMM23, Theorem 8.5].

Example 7.5.17. The left side of Figure 7.4 shows an example cluster picture where subscripts denote relative depths. We label the four proper clusters s_2, s_3, s_4 , and s_6 according to their size. Let us read off the minimal skeleton. We have five vertices, $\tilde{\eta}_2, \tilde{\eta}_3, \tilde{\eta}_4, \tilde{\eta}_6^+, \tilde{\eta}_6^-$ and five oriented edges $\tilde{e}_2^+, \tilde{e}_2^-, \tilde{e}_3, \tilde{e}_4^+, \tilde{e}_4^-$.

These are connected up as in the graph on the upper left of Figure 7.4. Moreover, $\tilde{\eta}_3$ has genus 1 and all other vertices have genus 0. The edge lengths are given by $l(\tilde{e}_2^\pm) = 1$, $l(\tilde{e}_3) = 1/2$, and $l(\tilde{e}_4^\pm) = 2$.

This Berkovich skeleton associated to X covers the Berkovich skeleton of a semistable model of $\mathbb{P}^1(\mathcal{R})$. This skeleton is depicted in the lower left. Moving right, we have an artistic depiction of the Berkovich spaces of X and $\mathbb{P}^{1,\text{an}}$ decomposed into wide opens and annuli, corresponding to the semistable covering. These come equipped with reduction maps to their special fibres, depicted to its right.

Proof of Theorem 7.5.15. (a) If $\mathfrak{s}$ is an übereven cluster, we know that $X|_{U_{\mathfrak{s}}} \simeq X_{\mathfrak{s}}|_{U_{\mathfrak{s}}} = U_{\mathfrak{s}} \times \{\pm 1\}$ via the isomorphism in (7.5.1.1). Since $\eta_{\mathfrak{s}} \in |U_{\mathfrak{s}, \mathbb{C}_\ell}|$, its preimage $\pi^{-1}(\eta_{\mathfrak{s}}) \subseteq X|_{U_{\mathfrak{s}}}$ thus consists of the two type II points $\tilde{\eta}_{\mathfrak{s}}^\pm = (\eta_{\mathfrak{s}}, \pm 1)$, both of genus 0.

If $\mathfrak{s}$ is an odd cluster, then the fibres of the map π are described in $(*)$ in the proof of [Ber90, Proposition 3.4.6]; this shows that the fibres are size 1 or 2. Therefore, the cover $|X_{\mathbb{C}_\ell}^{\text{an}}| \rightarrow |\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}}|$ is a ramified topological double covering. In particular, the ramification locus is closed. For any odd non-top cluster $\mathfrak{s}$, Lemma 7.5.9 shows that $e_{\mathfrak{s}}$ lies in the ramification locus (the squaring map on annuli is ramified along the skeleta). Then $\eta_{\mathfrak{s}}$ lies in the ramification locus, since it lies in the closure of $e_{\mathfrak{s}}$. If $\mathfrak{s}$ is the top cluster and odd, then f has odd degree, so $\infty \in \mathcal{R}'$. In this case, Lemma 7.5.13 implies that $X_{\mathbb{C}_\ell}^{\text{an}}|_{D_\infty}$ consists of a single disc mapping via the squaring map to $D_\infty = \mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}} \setminus D_{\mathfrak{s}}^c$. Hence the map $|X_{\mathbb{C}_\ell}^{\text{an}}| \rightarrow |\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}}|$ is again ramified over the open arc from $\eta_{\mathfrak{s}}$ to ∞ in $|\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}}|$, and so $\eta_{\mathfrak{s}}$ lies in its ramification locus. In either case, the point $\eta_{\mathfrak{s}}$ lies in the ramification locus, so has one preimage in $|X_{\mathbb{C}_\ell}^{\text{an}}|$, and therefore is of type II.

To finish part (a), we need to prove the assertion regarding the genus. For this we use the fact that, if $\bar{k}$ denotes the residue field of $\mathbb{C}_\ell$ (an algebraic closure of the residue field k of K), and if $\mathcal{X}_{\tilde{\eta}_{\mathfrak{s}}}$ denotes the smooth projective curve over $\bar{k}$ associated to $\tilde{\eta}_{\mathfrak{s}}$, then there is a bijection

$$\mathcal{X}_{\tilde{\eta}_{\mathfrak{s}}}(\bar{k}) \simeq T_{\tilde{\eta}_{\mathfrak{s}}} X_{\mathbb{C}_\ell}^{\text{an}} = \text{tangent directions at } \tilde{\eta}_{\mathfrak{s}} \in |X_{\mathbb{C}_\ell}^{\text{an}}|$$

[BPR13, Lemma 5.12(3)]. Moreover, this identification is natural with respect to finite morphisms of curves, so we have the induced map

$$\mathcal{X}_{\tilde{\eta}_{\mathfrak{s}}}(\bar{k}) \rightarrow P_{\eta_{\mathfrak{s}}}(\bar{k}) \simeq \mathbb{P}_{\bar{k}}^1(\bar{k}),$$

where $P_{\eta_{\mathfrak{s}}}(\bar{k})$ is the component of $\overline{\mathbb{P}_{\bar{k}}^1}$ given by $\overline{\text{red}(\eta)}$, see (7.2.1.2).

Our strategy now is to describe the tangent directions in $|\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}}|$ at $\eta_{\mathfrak{s}}$ and their lifts to tangent directions at $\tilde{\eta}_{\mathfrak{s}}$ in $|X_{\mathbb{C}_\ell}^{\text{an}}|$. There are three different classes of tangent directions. First, the “upward tangent direction” goes along $e_{\mathfrak{s}}$ for $\mathfrak{s}$ non-top; if $\mathfrak{s}$ is the top cluster then we mean the tangent direction on the arc from $\eta_{\mathfrak{s}}$ to ∞ . These lift to two tangent directions if $\mathfrak{s}$ is even, and one if $\mathfrak{s}$ is odd. Second, the “downward tangent direction” goes along $e_{\mathfrak{s}'}^{-1}$ for every $\mathfrak{s}' < \mathfrak{s}$ proper child. These lift to two tangent directions if $\mathfrak{s}'$ is an even proper child, and one if $\mathfrak{s}'$ is odd proper child. Third, there is a tangent direction for every disc component of $\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}} \setminus V$ whose closure contains $\eta_{\mathfrak{s}}$. In this case there are two lifts except if D contains a root of f (necessarily unique).

As a consequence, the number of ramification points of $\mathcal{X}_{\tilde{\eta}_{\mathfrak{s}}} \rightarrow \mathbb{P}_{\bar{k}}^1$ is equal to the number of odd children of $\mathfrak{s}$, plus 1 if $\mathfrak{s}$ itself is odd. Since $\mathcal{X}_{\tilde{\eta}_{\mathfrak{s}}} \rightarrow \mathbb{P}_{\bar{k}}^1$ has degree 2, this implies that $\mathcal{X}_{\tilde{\eta}_{\mathfrak{s}}}$ is a hyperelliptic curve over $\bar{k}$ of genus $g(\mathfrak{s})$.

(b) Let $\mathfrak{s}$ be a proper, non-top cluster. Since $e_{\mathfrak{s}}$ is the skeleton of $A_{\mathfrak{s}}$, it follows from Lemma 7.5.9 that for a proper, non-top even (resp. odd) cluster $\mathfrak{s}$, the set $\pi^{-1}(e_{\mathfrak{s}})$ is the union of two (resp. 1) embedded open intervals, each of length $\delta_{\mathfrak{s}}$ (resp. $\delta_{\mathfrak{s}}/2$).

It remains to compute the endpoints of $\pi^{-1}(e_{\mathfrak{s}})$. We know that $\partial_0(\tilde{e}_{\mathfrak{s}}^{\pm})$ must be a preimage of $\eta_{\mathfrak{s}} = \partial_0(e_{\mathfrak{s}})$. If $\mathfrak{s}$ is non-übereven, there is only one preimage. When $\mathfrak{s}$ is übereven, we want to show that $\partial_0(\tilde{e}_{\mathfrak{s}}^+) = \tilde{\eta}_{\mathfrak{s}}^+$ rather than $\tilde{\eta}_{\mathfrak{s}}^-$. This is a consequence of our compatibility assumptions. A similar argument deals with ∂_1 .

(c) Recall that a semistable vertex set V is a finite set of type II points such that $X_{\mathbb{C}_\ell}^{\text{an}} \setminus V$ is a disjoint union of open discs and finitely many open annuli. By definition, its skeleton is the union of V and the skeleton of the annuli in $X_{\mathbb{C}_\ell}^{\text{an}} \setminus V$.

Let $V_0 \subset |\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}}|$ be the set consisting of the points $\eta_{\mathfrak{s}}$, where $\mathfrak{s}$ ranges over all proper clusters. This is a semistable vertex set, since it is a finite set and $\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}} \setminus V_0$ is the disjoint union of the annuli $A_{\mathfrak{s}}$ (for $\mathfrak{s}$ proper, non-top) and open discs each of which contains at most 1 ramification point of f . Let Γ_0 be the skeleton of V_0 . Since $\pi : X \rightarrow \mathbb{P}^1$ is finite, it preserves types of points and so $\pi^{-1}(V_0) = V$ is a finite set of type II points. Moreover, $X_{\mathbb{C}_\ell}^{\text{an}} \setminus V$ is the disjoint union of $X_{\mathbb{C}_\ell}^{\text{an}}|_{A_{\mathfrak{s}}}$ (for $\mathfrak{s}$ proper, non-top) and $X_{\mathbb{C}_\ell}^{\text{an}}|_D$ for D an open disc component in $\mathbb{P}_{\mathbb{C}_\ell}^{1,\text{an}} \setminus V_0$. These are (disjoint unions of) open annuli and discs by Lemmas Lemma 7.5.9 and Lemma 7.5.13, so V is a semistable vertex set. Moreover, Lemma 7.5.9 shows that the preimage of the skeleton of $A_{\mathfrak{s}}$ is the skeleton of $X_{\mathbb{C}_\ell}^{\text{an}}|_{A_{\mathfrak{s}}}$, so $\pi^{-1}(\Gamma_0) = \Gamma$ is the skeleton associated to V , concluding

our proof. $\square$

Finally, we give a useful theorem to decide if a component of the dual graph contains the reduction of $\mathbb{Q}_\ell$ -points. First we set up some notation.

Definition 7.5.18 (Components of special fibres associated to clusters). Let $X/\mathbb{Q}_\ell$ be a hyperelliptic curve. Let $\mathfrak{s}$ be a cluster of X . Let $\mathcal{X}$ be the semistable model of $X_{\mathbb{C}_\ell}$ corresponding to the semistable covering from Theorem 7.5.14. We write $\bar{\mathcal{X}}_\mathfrak{s}^\pm$ for the the component(s) of the special fibre of $\mathcal{X}$ corresponding to $\mathfrak{s}$ and $g(\mathfrak{s})$ for the genus of these components.

Theorem 7.5.19. *Let $X/\mathbb{Q}_\ell$ be a hyperelliptic curve. Let $\mathfrak{s}$ be a proper cluster of X . Write $\bar{\mathcal{X}}_\mathfrak{s}^{\pm, \circ} = \bar{\mathcal{X}} \setminus \bigcup_{\mathfrak{s}' \neq \mathfrak{s}} \bar{\mathcal{X}}_{\mathfrak{s}'}^\pm$. If $\bar{\mathcal{X}}_\mathfrak{s}^{\pm, \circ}$ contains the reduction of a $\mathbb{Q}_\ell$ -point, then at least one of the following is true:*

1. $\mathfrak{s}$ has a singleton child which is $\mathbb{Q}_\ell$ -rational;
2. $\mathfrak{s}$ is the top cluster and X has a $\mathbb{Q}_\ell$ -rational point at infinity;
3. $\nu_\mathfrak{s} \in 2\mathbb{Z}$.

Furthermore, if $\bar{\mathcal{X}}_\mathfrak{s}^\pm \cap \bar{\mathcal{X}}_{\mathfrak{s}'}^\pm$ contains the reduction of a $\mathbb{Q}_\ell$ -point and $\mathfrak{s} = P(\mathfrak{s}')$ then $(\nu_\mathfrak{s}, \nu_{\mathfrak{s}'})$ contains an even integer.

Proof. Let $V_0 \subset |\mathbb{P}_{\mathbb{C}_\ell}^{1, \text{an}}|$ be the set of Gauss points $\eta_\mathfrak{s}$ attached to proper clusters $\mathfrak{s}$. Let Γ_0 be the skeleton of $\mathbb{P}_{\mathbb{C}_\ell}^{1, \text{an}}$ corresponding to V_0 , and let Γ_1 be the skeleton of the open curve $\mathbb{A}_{\mathbb{C}_\ell}^{1, \text{an}} \setminus \mathcal{R}$ where $\mathcal{R}$ is the set of roots of f . Though we have not explained the full definition of the skeleton of an open curve [BPR13, Definition 3.3], in this case, we can say explicitly that Γ_1 is the union of Γ_0 and a number of open rays isometric to $(0, \infty)$. We have one such ray for each root r of f , connecting $r \in |\mathbb{P}_{\mathbb{C}_\ell}^{1, \text{an}}|$ to $\eta_\mathfrak{s}$ where $\mathfrak{s}$ is the smallest proper cluster containing r , and one additional ray connecting the Gauss point attached to the top cluster to the point at ∞ . According to [BPR13, Definition 3.7], there is a canonical retraction $\tau_1 : |\mathbb{A}_{\mathbb{C}_\ell}^{1, \text{an}}| \setminus \mathcal{R} \rightarrow \Gamma_1$ [BPR13, Definition 3.7].

Now suppose that $(x, y) \in X(\mathbb{Q}_\ell)$ is $\mathbb{Q}_\ell$ -rational. We assume without loss of generality that (x, y) is not a point at infinity or a Weierstrass point. Note that (x, y) reduces onto $\bar{\mathcal{X}}_\mathfrak{s}^\pm$ if and only if $x \in U_\mathfrak{s}$, and (x, y) reduces onto $\bar{\mathcal{X}}_\mathfrak{s}^\pm \cap \bar{\mathcal{X}}_{P(\mathfrak{s})}^\pm$ if and only if $x \in A_\mathfrak{s}$. There are four cases to consider:

Firstly, suppose that $\tau_1(x)$ lies on the ray connecting a root r of f to $\eta_\mathfrak{s}$, where $\mathfrak{s}$ is the parent of $\{r\}$. This ray is the skeleton of a punctured open disc $D \setminus \{r\}$ where $D \ni x$. Since D contains a $\mathbb{Q}_\ell$ -rational point, it must be

setwise invariant under the action of the Galois group of $\mathbb{Q}_\ell$, and so too must be $\{r\} = D \cap \mathcal{R}$. So r is a $\mathbb{Q}_\ell$ -rational root of f which is a singleton child of $\mathfrak{s}$ and x reduces onto $\bar{\mathcal{X}}_{\mathfrak{s}}^{\pm, \circ}$.

Secondly, suppose that $\tau_1(x)$ lies on the ray connecting $\eta_{\mathfrak{s}}$ to ∞ for $\mathfrak{s}$ the top cluster. This ray is the skeleton of a punctured disc $D \setminus \{\infty\}$ where $D \subset \mathbb{P}_{\mathbb{C}_\ell}^{1, \text{an}}$ is an open disc around ∞ (complement of a closed disc in $\mathbb{A}_{\mathbb{C}_\ell}^{1, \text{an}}$). There are two possibilities, depending on the degree of f . Either the preimage of D inside $X_{\mathbb{C}_\ell}^{\text{an}}$ is an open disc mapping to D via a finite degree 2 map ramified over ∞ , or it is a disjoint union of two open discs mapping isomorphically to D . In either case, X has a $\mathbb{Q}_\ell$ -rational point at ∞ (e.g. in the second case, one of the two discs must contain (x, y) , and so be fixed by the natural Galois action). In this case, (x, y) reduces onto $\bar{\mathcal{X}}_{\mathfrak{s}}^{\pm, \circ}$ for $\mathfrak{s}$ the top cluster.

Thirdly, if $\tau_1(x) = \eta_{\mathfrak{s}}$ is the Gauss point of some cluster $\mathfrak{s}$, then we have

$$v(f(\eta_{\mathfrak{s}})) = v(f(x)) = 2v(f(y)) \in 2\mathbb{Z}$$

by the slope formula [BPR13, Theorem 5.15]. But it is easy to check, e.g. by factorising f , that $\nu_{\mathfrak{s}} = v(f(\eta_{\mathfrak{s}}))$, so $\nu_{\mathfrak{s}} \in 2\mathbb{Z}$. In this case, x lies in $U_{\mathfrak{s}}$ but not any of the annuli $A_{\mathfrak{s}'}$, so (x, y) reduces onto $\bar{\mathcal{X}}_{\mathfrak{s}}^{\pm, \circ}$.

Finally, if $\tau_1(x)$ lies in the skeleton of an annulus $A_{\mathfrak{s}}$, then arguing as above, we have that $v(f(\tau_1(x))) = v(f(x)) \in 2\mathbb{Z}$. But the skeleton of $A_{\mathfrak{s}}$ is an open interval in $|\mathbb{P}_{\mathbb{C}_\ell}^{1, \text{an}}|$ connecting the two Gauss points $\eta_{\mathfrak{s}}$ and $\eta_{P(\mathfrak{s})}$. Since the valuation of f is a linear function along this interval [BPR13, Theorem 5.15(2)], it follows that $v(f(\tau_1(x))) \in (\nu_{P(\mathfrak{s})}, \nu_{\mathfrak{s}})$ and we are done. $\square$

Example 7.5.20. Suppose that X has the cluster picture shown in Figure 7.5 at ℓ and that the leading coefficient of X is a unit in $\mathbb{Z}_\ell$. Then the local height of any $\mathbb{Q}_\ell$ -point on X is 0.




Figure 7.5: Cluster picture

7.5.3 Computing the Coleman–Iovita maps

In this section, we use the semistable covering $\mathcal{U}$ to describe the maps from Section 7.4.2 for hyperelliptic curves. Recall that φ_2 is needed to translate the action of the endomorphism Z_* on $H_{\text{dR}}^1(X/K)$ to the action of Z_* on $H_1(\Gamma, K)$.

In addition, we need φ_1 to compute the trace of Z_* on $H_{\text{rig}}^1(\bar{\mathcal{X}}_{\mathfrak{s}}/K)$ for higher genus clusters $\mathfrak{s}$.

The following lemma gives an explicit way of describing the map

$$\varphi_2 : H_{\text{dR}}^1(X/K) \rightarrow H_1(\Gamma, K)$$

for hyperelliptic curves by reducing the calculation to computing residues on $\mathbb{P}^1$.

Lemma 7.5.21. *Let X be a hyperelliptic curve, and $\mathfrak{s}$ be a non-top cluster with associated annulus $A_{\mathfrak{s}} \subseteq \mathbb{A}_K^{1,\text{an}}$. Let $\omega = \frac{P(x)dx}{2y} \in H_{\text{dR}}^1(X/K)$, where $P(x)$ is a polynomial. If $\mathfrak{s}$ is odd, then $\text{Res}_{\pi^{-1}(A_{\mathfrak{s}})}(\omega) = 0$. If $\mathfrak{s}$ is even, recall that $\pi^{-1}(A_{\mathfrak{s}}) = \tilde{A}_{\mathfrak{s}}^+ \sqcup \tilde{A}_{\mathfrak{s}}^-$. In this case,*

$$-\text{Res}_{\tilde{A}_{\mathfrak{s}}^-}(\omega) = \text{Res}_{\tilde{A}_{\mathfrak{s}}^+}(\omega) = \text{Res}_{A_{\mathfrak{s}}} \left(\frac{P(x)dx}{2g_{\mathfrak{s}}(x)^{1/2}h_{\mathfrak{s}}(x)^{1/2}} \right),$$

where t is any parameter on the annulus.

Proof. For any proper cluster $\mathfrak{s}'$ contained in $\mathfrak{s}$, applying Lemma 7.4.2 to $\tilde{U}_{\mathfrak{s}'}$ (or both components of $\tilde{U}_{\mathfrak{s}'}$, if $\mathfrak{s}'$ is *übereven*) shows that

$$\sum_{\mathfrak{s}'' < \mathfrak{s}'} \text{Res}_{\tilde{A}_{\mathfrak{s}''}^{\pm}}(\omega) - \text{Res}_{\tilde{A}_{\mathfrak{s}'}^+}(\omega) = 0$$

or

$$\sum_{\mathfrak{s}'' < \mathfrak{s}'} \text{Res}_{\tilde{A}_{\mathfrak{s}''}^{\pm}}(\omega) - \text{Res}_{\tilde{A}_{\mathfrak{s}'}^+}(\omega) - \text{Res}_{\tilde{A}_{\mathfrak{s}'}^-}(\omega) = 0,$$

according to whether $\mathfrak{s}'$ is odd or even. We use our conventions for the orientations of annuli as in Example 7.2.11. Summing this identity over all proper clusters $\mathfrak{s}'$ contained in $\mathfrak{s}$ shows that $\text{Res}_{\tilde{A}_{\mathfrak{s}}}(\omega) = 0$ if $\mathfrak{s}$ is odd, and $\text{Res}_{\tilde{A}_{\mathfrak{s}}^+}(\omega) + \text{Res}_{\tilde{A}_{\mathfrak{s}}^-}(\omega) = 0$ if $\mathfrak{s}$ is even.

Finally, when $\mathfrak{s}$ is even, the isomorphism $A_{\mathfrak{s}} \simeq \tilde{A}_{\mathfrak{s}}^+$ is given by

$$x \mapsto (x, g_{\mathfrak{s}}(x)^{1/2}h_{\mathfrak{s}}(x)^{1/2})$$

by (7.5.1.2). So $\omega = \frac{P(x)dx}{2y}$ pulls back to the differential form $\frac{P(x)dx}{2g_{\mathfrak{s}}(x)^{1/2}h_{\mathfrak{s}}(x)^{1/2}}$ on $A_{\mathfrak{s}}$ and we are done. $\square$

In order to explicitly compute φ_1 , we use the semistable covering $\mathfrak{U}$. Let $\omega \in \ker(\varphi_2)$. Recall that $\omega|_{\tilde{U}_{\mathfrak{s}}}$ has residue 0 over all bounding annuli of $\tilde{U}_{\mathfrak{s}}$.

Therefore, $\omega \in H_{\text{dR}}^1(\bar{\mathcal{X}}_{\mathfrak{s}}/K)$. Our goal is to compute this restriction in terms of a basis $\eta_1, \dots, \eta_{2g(\mathfrak{s})}$ of $H_{\text{dR}}^1(\mathcal{X}_{\mathfrak{s}}/K)$. There exist $a_i \in K$ such that

$$\omega|_{\tilde{U}_{\mathfrak{s}}} = \sum_{i=1}^{2g(\mathfrak{s})} a_i \eta_i|_{\tilde{U}_{\mathfrak{s}}} + \text{exact form.} \quad (7.5.3.1)$$

To compute the coefficients a_i , we use global and local symbols.

Definition 7.5.22. Let X be a smooth projective curve over K , let $D_1^c, \dots, D_k^c$ be some pairwise disjoint closed discs in the analytification X^{an} , and let $U := X^{\text{an}} \setminus \bigcup_i D_i^c$ be the complement of these discs. Choose some open discs $D_i^o \supset D_i^c$, still pairwise disjoint, and let $A_i := D_i^o \setminus D_i^c \subset U$ be the corresponding annulus bounding the disc D_i^c . If ω, η are two differential forms on U , each with residue 0 on each annulus A_i , then the *global symbol* $\langle \omega, \eta \rangle$ is defined to be a sum of *local symbols* $\langle \omega, \eta \rangle_{A_i}$

$$\langle \omega, \eta \rangle := \sum_{A_i} \langle \omega, \eta \rangle_{A_i},$$

where

$$\langle \omega, \eta \rangle_{A_i} := -\text{Res}_{A_i} \left(\omega \cdot \int \eta \right) \in K.$$

Since η has residue 0 on each A_i , it has a formal antiderivative $\int \eta$ on A_i , which is unique up to an additive constant of integration. Because ω also has residue 0, the residue of $\omega \cdot \int \eta$ on A_i is independent of this constant of integration.

We now describe how to practically compute these local symbols in our use case, where $\omega \in \ker \varphi_2$, and η is a differential on $X_{\mathfrak{s}}$ for some cluster $\mathfrak{s}$. First, we remark that in any and all calculation, both ω and the η_i are chosen to be in the -1 eigenspace of the hyperelliptic involution ι of X .

Recall that the choice of a square root $h_{\mathfrak{s}}^{1/2}$ determines an isomorphism $X|_{U_{\mathfrak{s}}} \rightarrow X_{\mathfrak{s}}|_{U_{\mathfrak{s}}}$. We let $\tilde{\omega}$ denote the differential on $X_{\mathfrak{s}}$ that ω maps to under this isomorphism. The formula for local symbols on $X_{\mathfrak{s}}$ becomes

$$\langle \tilde{\omega}, \eta \rangle = - \sum_{\mathfrak{s}' \leq \mathfrak{s}} \text{Res}_{\tilde{A}_{\mathfrak{s}'}^{\pm}} \left(\tilde{\omega} \cdot \int \eta \right) \in K. \quad (7.5.3.2)$$

Here we are taking the orientation on $\tilde{A}_{\mathfrak{s}'}^{\pm}$ induced by their inclusion in the wide open $U_{\mathfrak{s}}$. To compute $\text{Res}_{\tilde{A}_{\mathfrak{s}'}^{\pm}}(\tilde{\omega} \cdot \int \eta)$, pick a parameter of $\tilde{A}_{\mathfrak{s}'}^{\pm}$. If $\mathfrak{s}$ is even, this parameter will be a parameter of $A_{\mathfrak{s}}^{\pm}$, and if $\mathfrak{s}$ is odd it will be a square

root of a parameter on $A_{\mathfrak{s}}$. We can then express everything in this parameter by using our choice of a local square root of $g_{\mathfrak{s}}$, the defining polynomial for $X_{\mathfrak{s}}$. Note that the resulting local symbol does not depend on the sign choice we made for $g_{\mathfrak{s}}^{1/2}$, since by assumption $\iota(\omega) = -\omega$ and $\iota(\eta) = -\eta$. A different sign choice would multiply the global symbol by $(-1)^2$. In particular, we find $\text{Res}_{\tilde{A}_{\mathfrak{s}}^+}(\tilde{\omega} \cdot \int \eta) = \text{Res}_{\tilde{A}_{\mathfrak{s}}^-}(\tilde{\omega} \cdot \int \eta)$. Write $\omega = a(x)dx/(2y)$ on the curve X . Write $\tilde{\omega} = a(x)dx/(2h_{\mathfrak{s}}^{1/2}u_{\mathfrak{s}})$ and $\eta = b(x)dx/(2u_{\mathfrak{s}})$ on the curve $u_{\mathfrak{s}}^2 = g_{\mathfrak{s}}(x)$. We compute the formula.

$$\langle \tilde{\omega}, \eta \rangle = \sum_{\mathfrak{s}' \leq \mathfrak{s}} \sigma_{\mathfrak{s}, \mathfrak{s}'} \begin{cases} \frac{1}{2} \text{Res}_{x=\alpha_{\mathfrak{s}'}} \left(a(x) h_{\mathfrak{s}}^{-\frac{1}{2}} g_{\mathfrak{s}}^{-\frac{1}{2}} dx \cdot \int b(x) g_{\mathfrak{s}}^{-\frac{1}{2}} dx \right) & \text{if } \mathfrak{s}' \text{ is even} \\ \text{Res}_{t_{\mathfrak{s}'}=0} \left(a(x) h_{\mathfrak{s}}^{-\frac{1}{2}}(x) g_{\mathfrak{s}, \mathfrak{s}'}^{-\frac{1}{2}}(x) dt_{\mathfrak{s}'} \cdot \int b(x) g_{\mathfrak{s}, \mathfrak{s}'}^{-\frac{1}{2}}(x) dt_{\mathfrak{s}'} \right) & \text{if } \mathfrak{s}' \text{ is odd} \end{cases} \in K \quad (7.5.3.3)$$

where $\sigma_{\mathfrak{s}, \mathfrak{s}'} = -1$ if $\mathfrak{s} = \mathfrak{s}'$ and $+1$ otherwise, and in the odd case we recall $t_{\mathfrak{s}'}$ is the coordinate on $X|_{A_{\mathfrak{s}'}}$ from (7.5.1.3), and $x = t_{\mathfrak{s}'}^2 + \alpha_{\mathfrak{s}'}$, and in the odd case $g_{\mathfrak{s}, \mathfrak{s}'}$ is shorthand for $g_{\mathfrak{s}}/(x - \alpha_{\mathfrak{s}'})$. The choice of square root of $g_{\mathfrak{s}, \mathfrak{s}'}$ does not matter, since a different choice multiplies the global symbol by $(-1)^2$. (However, we have already made a choice of square root of $g_{\mathfrak{s}, \mathfrak{s}'}$; this choice can be determined by asking that $h_{\mathfrak{s}}^{\frac{1}{2}} g_{\mathfrak{s}, \mathfrak{s}'}^{\frac{1}{2}}$ and $h_{\mathfrak{s}'}^{\frac{1}{2}} g_{\mathfrak{s}, \mathfrak{s}'}^{\frac{1}{2}}$ are the same square root of $f/(x - \alpha_{\mathfrak{s}'})$.)

Proposition 7.5.23. [Bes00, Proposition 4.10] cf. [BB12, Proposition 2.5] Suppose ω and η are restrictions of algebraic 1-forms of the second kind on X . Let $[\omega], [\eta] \in H_{\text{dR}}^1(X/K)$ be the associated de Rham cohomology classes.

$$\langle \omega, \eta \rangle = [\omega] \cup [\eta] \in H_{\text{dR}}^2(X/K) = K.$$

In particular, $\langle \omega, \eta \rangle$ depends only on the de Rham cohomology classes represented by ω and η .

Now we are ready to solve for the coefficients a_i from (7.5.3.1). For each $j = 1, \dots, 2g(\mathfrak{s})$, we compute

$$\langle \omega, \eta_j \rangle = \sum_{i=1}^{2g(\mathfrak{s})} a_i \langle \eta_i, \eta_j \rangle. \quad (7.5.3.4)$$

This gives a system of $2g(\mathfrak{s})$ linear equations in the $2g(\mathfrak{s})$ variables $a_1, \dots, a_{2g(\mathfrak{s})}$, and this system of equations has a unique solution since the cup product pairing on $H_{\text{dR}}^1(\mathcal{X}_{\mathfrak{s}}/K_{\ell})$ is perfect.

7.6 Computations

We give more details in the computation of the Coleman–Iovita morphisms, as well as how this results in the provably correct rational matrix of the action of Z_* on $H_1(\Gamma, \mathbb{Z})$, and integer traces $\mathrm{tr}_v(Z)$.

We first focus on the computations of the Coleman–Iovita morphisms from Section 7.5. As explained in Lemma 7.5.21, determining φ_2 reduces to computing residues over annuli $A_{\mathfrak{s}}$ on $\mathbb{P}^1$ for all even clusters $\mathfrak{s}$. To do this explicitly, we use Lemma 7.5.5 to construct an inverse square root of $f(t)$, where t is a choice of parameter in $A_{\mathfrak{s}}$. In order to have a provably correct result, it is important for our calculations to keep track of the p -adic precision of the approximations we make; this is why we do the following p -adic analysis.

In this section, we assume $(K, v : K \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\})$ is a field with a non-archimedean valuation. In applications, we will always take K to be a finite extension of $\mathbb{Q}_\ell$ together with its normalised valuation v .

Analytic functions on an annulus have a concrete characterisation. In order to carry out explicit computations, we need to work on the ring of analytic functions on *closed* annuli. That is, if $A[\ell^{-v_1}, \ell^{-v_2}]$ is the standard closed annulus centered at 0 with inner and outer radii ℓ^{-v_1} and ℓ^{-v_2} , we work in the ring of power series $\sum_{i \in \mathbb{Z}} a_i t^i \in \mathcal{O}(A)$ where the a_i satisfy

$$\lim_{i \rightarrow -\infty} v(a_i) + v_1 \cdot i = \infty \quad \text{and} \quad \lim_{i \rightarrow \infty} v(a_i) + v_2 \cdot i = \infty.$$

Definition 7.6.1. Let $A \subseteq \mathbb{P}_K^{1, \mathrm{an}}$ be a closed annulus centered at 0 with inner and outer radii ℓ^{-v_1} and ℓ^{-v_2} .

For each w in the closed interval $[v_2, v_1]$ define the valuation $v^w : \mathcal{O}(A) \rightarrow \mathbb{Q}$, given by $v^w(h) = \min v(a_i) + w \cdot i : i \in \mathbb{Z}$, where $h = \sum_{i \in \mathbb{Z}} a_i t^i \in \mathcal{O}(A)$. We let $\mathfrak{m} \subset \mathcal{O}(A)$ denote the ideal of elements h with $v^{v_1}(h) > 0$ and $v^{v_2}(h) > 0$.

Most of the time, we are interested in the units of the ring $\mathcal{O}(A)$, because many of the operations, like taking the inverse and taking a square root, are only applicable to units. To both recognise units, and easily work with them, we provide the following notion of decomposition for units.

By considering the Newton polygon, we find the following proposition characterising the units of $\mathcal{O}(A)$.

Proposition 7.6.2. Let $A \subseteq \mathbb{P}_K^{1, \mathrm{an}}$ be a closed annulus centered at 0 with inner and outer radii ℓ^{-v_1} and ℓ^{-v_2} . Let t be a parameter on A and let $h =$

$\sum_i a_i t^i \in \mathcal{O}(A)$. For $k = 1, 2$, we define the quantities

$$i_k := \operatorname{argmin}(v(a_i) + iv_k),$$

where $\operatorname{argmin}(v(a_i) + iv_k)$ is the minimal index that minimises the quantity $v(a_i) + iv_k$. Then, the following are equivalent.

(i) There is a factorisation

$$h(t) = ct^d(1 + g(t)),$$

where $d \in \mathbb{Z}$, $c \in K^\times$, and the valuation of $g(t)$ at the Gauss points of the closed balls $B(0, \ell^{-v_1})$ and $B(0, \ell^{-v_2})$ is positive;

(ii) the minimum for each i_1 and i_2 is achieved once, and $i_1 = i_2$;

(iii) h has no zeros on A ;

(iv) h is an unit in $\mathcal{O}(A)$.

Moreover, if these hold, then in the factorisation $d = i_1$, $c = a_d$, and $g(t) = h(t)t^{-d}/c$, and we call (c, d, g) a decomposition of h .

Proof. For the equivalence of (i), (iii), and (iv) see [BPR13, Proposition 2.2]. We will show the equivalence of (i) and (ii). This follows from the following observations. If we write

$$h = ct^d(1 + g(t))$$

for some $c \in K^\times$, $d \in \mathbb{Z}$, $g \in \mathcal{O}(A)$, then d minimises $(v(a_i) + iv_k)$ if and only if the valuation at the Gauss point of the closed ball $B(0, \ell^{-v_k})$ is non-negative, and d is the unique index where this is minimised if and only if the valuation is positive. $\square$

To do practical computations on an annulus A , we have to represent elements $h \in \mathcal{O}(A)$ up to finite precision in some specified way. For rings like $\mathbb{Q}_p$ and $\mathbb{R}$, this is typically done using ball arithmetic: one gives the center x (often taken to be in some countable subring, for example $\mathbb{Q}$), some bound on the error b , and this then represents any number y with $|y - x| \leq b$. These balls can be added and multiplied, and one can easily propagate the error bounds. The equivalent notion for $\mathcal{O}(A)$ has two error bounds, corresponding to the two different valuations v^{v_1}, v^{v_2} on $\mathcal{O}(A)$.

Definition 7.6.3. A *representation* of an element $h \in \mathcal{O}(A)$ is a triple of elements $(\tilde{h}, B_1, B_2)$ where $\tilde{h} \in K[t, t^{-1}]$, $B_i \in \mathbb{R} \cup \{\infty\}$ and $v^{v_i}(h - \tilde{h}) \geq B_i$ for $i \in \{1, 2\}$. We say $\tilde{h}$ *represents* h .

Remark 7.6.4. A representation $(\tilde{h}, B_1, B_2)$ represents infinitely many different f .

Remark 7.6.5. Let $h \in \mathcal{O}(A)$ be represented by $(\tilde{h}, B_1, B_2)$, where h has coefficients $(a_i)_{i \in \mathbb{Z}}$ and $\tilde{h}$ has coefficients $(\tilde{a}_i)_{i \in \mathbb{Z}}$. Then we see $v(a_i - \tilde{a}_i) \geq \max(-iv_1 + B_1, -iv_2 + B_2)$.

Lemma 7.6.6. *Let $(\tilde{h}, B_1, B_2)$ be a representation for $h \in \mathcal{O}(A)$. If $\tilde{h}$ has a decomposition $(c, d, \tilde{g})$ and $v^{v_i}(\tilde{h}) < B_i$ for $i = 1, 2$, then h has a decomposition (c, d, g) where $(\tilde{g}, B_1 - v(c) - v_1 \cdot d, B_2 - v(c) - v_2 \cdot d)$ represents g . If not, then $\tilde{h}$ represents an element which achieves the value 0 on A .*

Proof. If $\tilde{h}$ is not a unit (which happens if and only if it does not have a decomposition), then the lemma is clear. Otherwise, let $(c, d, \tilde{g})$ be the decomposition of $\tilde{h}$, and define $g = h/(ct^d) - 1$. Then $(\tilde{g}, B_1 - v(c) - v_1 \cdot d, B_2 - v(c) - v_2 \cdot d)$ represents g . If $v^{v_i}(\tilde{h}) < B_i, B_2$, then we see g must lie in $\mathfrak{m}$, and (c, d, g) represents h . Otherwise, there is a choice of g' represented by $(\tilde{g}, B_1 - v(c) - v_1 \cdot d, B_2 - v(c) - v_2 \cdot d)$ that achieves the value -1 , and then $h' = ct^d(1 + g')$ is represented by $(\tilde{h}, B_1, B_2)$ and achieves the value 0. $\square$

Lemma 7.6.7. *Assume $v(2) = 0$. Let $h \in \mathcal{O}(A)$ be represented by $(\tilde{h}, B_1, B_2)$. Assume $\tilde{h}$ has a decomposition $(c, d, \tilde{g})$ as in Lemma 7.6.6. If d is odd or c is not a square, then h has no (inverse) square root. If d is even and c is a square, then $h^{-\frac{1}{2}}$ is represented by*

$$\left((c^{-1/2} x^{d/2} \sum_{i=0}^k \binom{-\frac{1}{2}}{i} \tilde{g}^i, \min(v(c^{-1/2}) + dv_1 + (k+1)v^{v_1}(\tilde{g}), \frac{1}{2}B_1), \right. \\ \left. \min(v(c^{-1/2}) + dv_2 + (k+1)v^{v_2}(\tilde{g}), \frac{1}{2}B_2) \right)$$

for any $k \in \mathbb{Z}_{\geq 1}$.

Proof. This follows immediately from the power series expansion of $(1+x)^{-\frac{1}{2}}$, and the fact that $\binom{-\frac{1}{2}}{i}$ has valuation at least 0. $\square$

Let $\mathfrak{s}$ be an even non-top cluster. We use Proposition 7.6.2 and Lemma 7.6.7 to construct an inverse square root of $f(t)$ on the annulus $A_{\mathfrak{s}}$.

To compute the map φ_1 , we need to compute the local symbol explicitly. Let $\omega \in H_{\text{dR}}^1(X/K)$ be in the kernel of φ_2 . Fix $\mathfrak{s}$ a cluster such that $g(\mathfrak{s}) > 0$.

Recall that we have chosen a basis $\eta_1, \dots, \eta_{2g(\mathfrak{s})}$ for $H_{\text{dR}}^1(\mathcal{X}_{k,\mathfrak{s}}/K)$. We want to compute $\langle \omega, \eta_i \rangle$ and $\langle \eta_i, \eta_j \rangle$ for $1 \leq i, j \leq g(\mathfrak{s})$. The global symbol is a sum of local symbols as in (7.5.3.2). This motivates the following lemma on the precision analysis for integration of functions on an annulus.

Lemma 7.6.8. *Let $h \in \mathcal{O}(A)$ be a function with representation $(\tilde{h}, B_1, B_2)$, with the coefficient of a_{-1} being 0 for h and $\tilde{h}$. Let A' be an annulus with inner and outer radii $\ell^{-v_1+\varepsilon}$ and $\ell^{-v_2-\varepsilon}$. Let $\delta(\varepsilon)$ denote the minimal value of $\varepsilon \cdot i - v(i)$ for $i \in \mathbb{Z}_{\geq 0}$. Then the integral of h on A' is well-defined and represented on A' by $(\int \tilde{h}, B'_1, B'_2)$, where*

$$B'_1 = \min(B_1 + v_1 - \delta(\varepsilon), B_2 + v_2 - \delta(v_1 - \varepsilon - v_2))$$

and

$$B'_2 = \min(B_2 + v_2 - \delta(\varepsilon), B_1 + v_1 - \delta(v_1 - \varepsilon - v_2)).$$

Proof. We present a sample computation, showing that $B'_1 \geq B_1 + v_1 - \delta(\varepsilon)$. Assume for simplicity that $h = at^i$ with $i < -1$ and $\tilde{h} = bt^i$, with $v(a-b)+iv_i \geq B_i$. Then

$$\begin{aligned} v \left(\frac{1}{i+1} (a-b) \right) + (i+1)(v_1 - \varepsilon) &\geq B_1 + v_1 - v_p(i+1) - \varepsilon(i+1) \\ &\geq B_1 + v_1 - \delta(\varepsilon). \end{aligned}$$

The cases where $i \geq 0$, or we bound B_2 instead of B_1 proceed similarly. $\square$

In order to compute Z_* on $H_1(\Gamma, \mathbb{Z})$ and $\text{tr}_v(Z)$, we also need a matrix for the action of Z_* on $H_{\text{dR}}^1(X/K)$. We now describe a method to compute such a matrix with rational coefficients. This method is entirely new, and solves the problem of how to determine the matrix of an endomorphism on $H_{\text{dR}}^1(X/K)$ precisely. Previously methods for quadratic Chabauty use the Hecke operator $Z = T_p$ as the endomorphism and employ Eichler–Shimura to determine the action of T_p on de Rham cohomology to arbitrary p -adic precision [BDM⁺21, §3.5.2]. These methods are restricted to modular curves, and other proposed methods only compute the action of Z_* to some p -adic precision. In order to determine the local heights at $\ell \neq p$ as a rational multiple $r\chi(\varpi_K)$ for $r \in \mathbb{Q}$ (as in Section 7.3), we need to know the action of Z_* exactly. (Otherwise we can only determine r to some ℓ -adic precision.) To solve this problem, we produce a method to rigorously compute the integer matrix for the action of Z_* .

Proposition 7.6.9. *Let π_1 and π_2 be the projections from $X \times X \rightarrow X$. Let ω be a 1-form on X . Let $Z \subset X \times X$ be a correspondence. Let ω_0 be a fixed non-zero 1-form. Define $g := (\pi_1^*\omega/\pi_2^*\omega_0)$; this is a rational function on Z . Then*

$$Z_*\omega = \pi_{2,*}(\pi_1^*\omega) = (\pi_{2,*}g)\omega_0.$$

Here, the pushforward $\pi_{2,}g$ denotes the trace of g in the extension $K(Z)/K(X)$ of function fields.*

Proof. By the projection formula [Har75, Theorem 7.5]

$$\pi_{2,*}(\pi_1^*\omega) = \pi_{2,*}(g \cdot \pi_2^*\omega_0) = \pi_{2,*}(g) \cdot \omega_0.$$

□

Choose ω_i for $i = \{0, \dots, 2g-1\}$ be a basis for $H_{\text{dR}}^1(X/K)$ constructed by taking linear combinations of $x^i dx/2y$ for $i = 0, \dots, 2g$. We use Proposition 7.6.9 to compute $Z_*(\omega_i)$ for all $i = 1, \dots, 2g-1$. The last step is to rewrite these elements in terms of the basis we chose.

Note that the class of any differential form $g(x, y)dx \in H_{\text{dR}}^1(X/K)$, where $g(x, y) \in K(X)$, is equivalent in cohomology to the class of $2y(g(x, y) - g(x, -y))\frac{dx}{2y}$, where $2y(g(x, y) - g(x, -y))$ is a rational function only in x . So now we reduce to the case of forms $g(x)\frac{dx}{2y}$, where $g(x) \in K(x)$. Then, a linear algebra argument allows us to write $g(x)\frac{dx}{2y}$ in terms of the basis $\{\omega_i\}_{i=0}^{2g-1}$.

Proposition 7.6.10. *Let $s_1, \dots, s_N$ be clusters of f such that the set of edges corresponding to $\{A_{s_i}\}_{i=1}^N$ in the dual graph are the complement of a spanning tree.*

Let T be the matrix with

$$T_{ij} = \text{Res}_{A_{s_i}} \omega_j.$$

This has a right inverse T^{-1} . Let M be the matrix of Z_ acting on $H_{\text{dR}}^1(X/K)$.*

Then the matrix of Z_ on $H_1(\Gamma, K)$ is TMT^{-1} .*

Proof. We can identify $H_1(\Gamma, K)$ with $K^{\{s_1, \dots, s_N\}}$; the contraction of the complementary spanning tree in the dual graph is a bouquet of N circles. Note that T is the matrix of the K -linear morphism φ_2 sending $\omega_j \mapsto \sum_i \text{Res}_{s_j} A_{s_i} \omega_j$ under this identification. By Theorem 7.4.14, the pushforward Z_* is preserved under φ_2 . The matrix for Z_* acting on $H_1(\Gamma, K)$ is then given by conjugating M by T . □

Using Theorem 7.3.5, we can bound the operator norm of the matrix with entries in K obtained from Proposition 7.6.10. By computing the matrix T with enough ℓ -adic precision, this allows us to obtain a matrix with entries in $\mathbb{Q}$. In the rest of this section we determine precise bounds. We start by bounding the size of the integers appearing in the matrix.

Proposition 7.6.11. *Let $Z \subset X \times X$ be an effective correspondence of degrees d_1 and d_2 over X , respectively. Let M be the integer matrix representing Z_* on $H^1(\Gamma, K)$. Let C denote the matrix of the intersection length pairing on homology. Write $C = PP^T$ for some invertible matrix P . Let $\|\cdot\|_{\max}$ be the norm given by the maximum of the absolute values of the entries. Then we have the bound*

$$\|M\|_{\max} \leq \frac{2g^2}{\|P\|_{\max}\|P^{-1}\|_{\max}} \sqrt{d_1 d_2}.$$

Proof. By Theorem 7.3.5, the operator norm of Z_* with respect to this intersection pairing $\|Z\|_L$ is bounded by $\sqrt{d_1 d_2}$ where d_i is the degree of the projection $\pi_i : Z \rightarrow X$. Since C is a positive definite symmetric matrix, we can write $C = PP^T$ for some invertible matrix P . This allows us to give a bound on the entries of M . It can be checked by evaluating on standard basis vectors that $\|M\|_I^2 \geq \|M\|_{\max}^2$ where I is the identity matrix (and $\|M\|_{\max}$ is maximum over the absolute value of the entries). Therefore

$$\|Z\|_L^2 = \|PZP^{-1}\|_I \geq \|PZP^{-1}\|_{\max}^2.$$

We obtain the bound $\|M\|_{\max} \leq \frac{2g^2}{\|P\|_{\max}\|P^{-1}\|_{\max}} \sqrt{d_1 d_2}$. □

Corollary 7.6.12. *Let M, C, P be as above. Let $c := \frac{2g^2}{\|P\|_{\max}\|P^{-1}\|_{\max}} \sqrt{d_1 d_2}$. Let $k \geq 0$ be an integer such that $v_\ell(M'_{ij} - M_{ij}) \geq k$. Suppose that M' is an integer matrix such that $\|M'\|_{\max} \leq c$ and $c + \|M'\|_{\max} < \ell^k$. Then $M = M'$.*

Using the techniques for working with $\mathcal{O}(A)$ developed in this section, one can compute M up to arbitrary p -adic precision, and Corollary 7.6.12 allows us to determine when we have computed enough to uniquely determine M . An example of this is worked out in Section 7.7.2.

We end by discussing how to compute the trace of the endomorphism Z acting on the de Rham cohomology of a single vertex once we have determined M .

Proposition 7.6.13. *Let $\mathfrak{s}$ be a cluster. Let $\omega'_0, \dots, \omega'_d$ be a basis for $\ker(\varphi_2)$. Recall that we chose a basis $\eta_1, \dots, \eta_{g(\mathfrak{s})}$ of $H^1_{\text{dR}}(\bar{\mathcal{X}}_{\mathfrak{s}}, K)$. For each ω'_i , define $s_{ij} \in K$ such that $\omega'_i = \sum_{j=0}^{2g(\mathfrak{s})} s_{ij} \eta_j$ found using (7.5.3.4). Let $S_{\mathfrak{s}}$ be the matrix*



Figure 7.6: Cluster picture for example in Section 7.7.1

with entries $S_{ij} = s_{ij}$. Let S be the block matrix consisting of all S_s , representing the map $\varphi_1: H_{\text{dR}}^1(X/K) \rightarrow \bigoplus_v H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K)$. This has a right inverse S^{-1} . Let M be the matrix of Z_* acting on $H_{\text{dR}}^1(X/K)$. Then SMS^{-1} is the matrix representing the induced action of Z_* on $\bigoplus_v H_{\text{rig}}^1(\bar{\mathcal{X}}_v/K)$, and $\text{tr}(Z_s)$ is the trace of the diagonal block of SMS^{-1} corresponding to $H_{\text{dR}}^1(\bar{\mathcal{X}}_s, K) \cong H_{\text{rig}}^1(\bar{\mathcal{X}}_s/K)$.

Proof. Note that S is the matrix of the K -linear morphism φ_1 . By Theorem 7.4.14, the pushforward Z_* is preserved under φ_1 . The matrix for Z_* acting on $H_1(\Gamma, K)$ is then given by conjugating M by S . $\square$

7.7 Examples

In this section we work through some illustrative examples of our method.

7.7.1 Bielliptic curve

As a warm-up, let q be an odd prime number. We consider the genus 2 bielliptic curve with affine equation

$$X : y^2 = (x^2 - q^\alpha)((x - 1)^2 - q^\beta)((x + 1)^2 - q^\beta).$$

This curve has bad reduction at q . We will compute the local height contribution at q using the formula from Theorem 7.3.9 and via intersection theory. The bielliptic involution $\varphi : x \mapsto -x$ gives rise to a trace 0 endomorphism of the Jacobian of X ; we use the correspondence Z arising from this endomorphism.

First, we use the method described in Section 7.3.2. The cluster picture for this curve is presented in Figure 7.6. The reduction graph associated to the semistable cover is given in Figure 7.7 where $l(e_0) = \alpha$ and $l(e_1) = l(e_{-1}) = \beta$. The bielliptic involution defines the action $Z_*(e_0) = -e_0$, $Z_*(e_1) = -e_{-1}$, and $Z_*(e_{-1}) = -e_1$.

We fix the midpoint of the edge e_0 as our base point. To apply the formula from Theorem 7.3.9, we first compute the orthogonal projections of the edges




Figure 7.7: Reduction graph for bielliptic example

in the first homology of the dual graph. For that, we note that the pairing on the span of the edges is given by

$$\langle e_i, e_j \rangle = \begin{cases} l(e_i) & \text{if } i = j; \\ 0 & \text{otherwise.} \end{cases}$$

Thus, the orthogonal projections are

$$\begin{aligned} \pi(e_0) &= \frac{\alpha}{2\alpha + \beta}(2e_0 - e_1 - e_{-1}); & \pi(e_1) &= \frac{1}{2\alpha + \beta}((\alpha + \beta)e_1 - \beta e_0 - \alpha e_{-1}); \\ \pi(e_{-1}) &= \frac{1}{2\alpha + \beta}((\alpha + \beta)e_{-1} - \beta e_0 - \alpha e_1). \end{aligned}$$

Then, we compute

$$\langle e_0, Z_*\pi(e_0) \rangle = -\frac{2\alpha^2}{2\alpha + \beta}; \quad \langle e_1, Z_*\pi(e_1) \rangle = \langle e_{-1}, Z_*\pi(e_{-1}) \rangle = \frac{\alpha\beta}{2\alpha + \beta}$$

to obtain the measure

$$\mu_Z = -\frac{4}{2\alpha + \beta}|ds_0| + \frac{2\alpha}{\beta(2\alpha + \beta)}|ds_1| + \frac{2\alpha}{\beta(2\alpha + \beta)}|ds_{-1}|,$$

where $|ds_i|$ is the arc-length measure along the edge e_i for $i = -1, 0, 1$. We conclude that $\tilde{h}_{Z,q}$ is given by

$$\begin{cases} -\frac{2}{2\alpha + \beta}s_{e_0}^2 + a_0s_{e_0} + b_0 & \text{on } e_0, \text{ with } 0 \leq s_{e_1} \leq \alpha; \\ \frac{\alpha}{\beta(2\alpha + \beta)}s_{e_1}^2 + a_1s_{e_1} + b_1 & \text{on } e_1, \text{ with } 0 \leq s_{e_1} \leq \beta; \\ \frac{\alpha}{\beta(2\alpha + \beta)}s_{e_{-1}}^2 + a_{-1}s_{e_{-1}} + b_{-1} & \text{on } e_{-1}, \text{ with } 0 \leq s_{e_{-1}} \leq \beta. \end{cases}$$

with constants a_i and b_i for $i = -1, 0, 1$. To find the exact values of the constants, we use that the polynomials must agree at endpoints, must vanish

at the chosen base point, and that $\Delta^2(\tilde{h}_{Z,q}) - \mu_Z = 0$ imposes a condition on the derivatives. This yields the local height function

$$\begin{cases} -\frac{2}{2\alpha+\beta} \left(\frac{\alpha}{2} - s_{e_0}\right)^2 & \text{on } e_0, \text{ with } 0 \leq s_{e_0} \leq \alpha; \\ \frac{\alpha}{\beta(2\alpha+\beta)} \left(\frac{\beta}{2} - s_{e_1}\right)^2 - \frac{\alpha}{4} & \text{on } e_1, \text{ with } 0 \leq s_{e_1} \leq \beta; \\ \frac{\alpha}{\beta(2\alpha+\beta)} \left(\frac{\beta}{2} - s_{e_{-1}}\right)^2 - \frac{\alpha}{4} & \text{on } e_{-1}, \text{ with } 0 \leq s_{e_{-1}} \leq \beta. \end{cases} \quad (7.7.1.1)$$

Now we compute local heights via intersection theory: the local height at q for a prime of bad reduction is defined in terms of intersection theory on the special fibre of a regular model $X_{\mathbb{Z}_q}$ as in the beginning of Section 7.3. Let $\mathcal{X}/\mathbb{Z}_q$ be the minimal regular model of Y . Then the following are true.

- The special fibre $\mathcal{X}_{\mathbb{F}_q}$ consists of two copies of $\mathbb{P}_{\mathbb{F}_q}^1$ connected by three chains of projective lines $\mathbb{P}_{\mathbb{F}_q}^1$ of lengths $\alpha-1$, $\beta-1$, and $\beta-1$, respectively.
- The action of φ on $\mathcal{X}_{\mathbb{F}_q}$ interchanges the two distinguished copies of $\mathbb{P}_{\mathbb{F}_q}^1$, reverses the chain of length $\alpha-1$, and reverses and interchanges the two chains of length $\beta-1$.

Let $Z = \Gamma_\varphi$ be the graph of φ . Let $P \in \mathcal{X}(\mathbb{Z}_q)$ be any point. Following (7.3.0.1), let

$$D_Z(b, P) = X^\varphi - \varphi(P) - \varphi(b)$$

be the divisor on the generic fibre, where X^φ is the set of fixed points under φ . We denote its closure in $\mathcal{X}$ also by

$$D_Z(b, P) = \mathcal{X}^\varphi - \varphi(P) - \varphi(b).$$


Our aim is then to find the vertical divisor V such that $\mathcal{D}_Z(b, P) := D_Z(b, P) + V$ has intersection number 0 with every vertical divisor. By normalising V so that the coefficient of the component containing $\bar{b}$ is zero, then the local height $\tilde{h}_{Z,q}(P)$ is the coefficient of the component of V containing $\bar{P}$.

Write the component of $\mathcal{X}_{\mathbb{F}_q}$ containing the point $\bar{P}$ as $[\bar{P}]$. Equivalently, without normalising, the height is given by

$$(\text{coefficient of } [\bar{P}] - \text{coefficient of } [\bar{b}]).$$

Write L for the matrix of the intersection pairing on components of the special fibre. We want a $\mathbb{Q}$ -linear combination V of components such that

$$L \cdot V = D(\bar{P}, \bar{b}) = [\varphi(\bar{P})] - [\bar{b}]$$

Figure 7.8: Special fibre for bielliptic example when $\alpha = 2$, $\beta = 1$

(using that $\mathcal{X}^\varphi$ and $\varphi(b)$ both reduce onto the same component as b). Therefore

$$\tilde{h}_{Z,q}(P) = ([\bar{P}] - [\bar{b}]) \cdot L^{-1}([\varphi(\bar{P})] - [\bar{b}]). \quad (7.7.1.2)$$

(One should be careful since L is not an invertible matrix. However $[\varphi(\bar{P})] - [\bar{b}]$ is in its image, and the local height is independent of the choice of preimage.)

Now L is also (by definition) the Laplacian matrix of the dual graph of $\mathcal{X}_{\mathbb{F}_q}$, and we recognise (7.7.1.2) as the formula for the height pairing on degree 0 divisors on a graph. So

$$\tilde{h}_{Z,q}(P) = \langle [\bar{P}] - [\bar{b}], [\varphi(\bar{P})] - [\bar{b}] \rangle. \quad (7.7.1.3)$$

Via the theory of electrical circuits on graphs, we can compute (7.7.1.3) explicitly, breaking it into cases depending on the component of the reduction of P . We obtain local heights agreeing with (7.7.1.1).

We work through one explicit example when $\alpha = 2$ and $\beta = 1$. By blowing up at the singular point $x = y = q = 0$, we compute a regular model $\mathcal{X}/\mathbb{Z}_q$, given by the equations

$$\begin{aligned} y^2 &= (x^2 - q^2)((x-1)^2 - q)((x+1)^2 - q), \\ Y^2 &= (X^2 - Q^2)((x-1)^2 - q)((x+1)^2 - q), \\ xY &= Xy, \quad yQ = Yq, \quad qX = Qx \end{aligned}$$

in $\mathbb{A}_{\mathbb{Z}_q}^2 \times \mathbb{P}_{\mathbb{Z}_q}^2$, which maps to the original curve by $(x, y; X, Y, Q) \mapsto (x, y)$. The special fibre $\mathcal{X}_{\mathbb{F}_q}$ is depicted in Figure 7.8. By counting intersections of components, we see that

$$L = \begin{bmatrix} -3 & 2 & 1 \\ 2 & -3 & 1 \\ 1 & 1 & -2 \end{bmatrix}$$

where the basis of L is the components in the order blue (left), pink (right), purple (top).

The graph $\Gamma_\varphi \subset (\mathbb{A}^2 \times \mathbb{P}^2)^2$ is given by the explicit equations $\{y = y', x = -x', X = -X', Y = Y', Q = Q'\}$. Then $\mathcal{X}^\varphi = \{(0, iq(1-q); 0, i(1-q),), (0, -iq(1-q); 0, -i(1-q), 1)\}$ are the fixed points under φ . This divisor reduces to two points, lying on the the purple (top) component. We choose a base point $b = (0, q; 0, 1, 1)$ reducing to the midpoint of the purple (top) component $Y = X^2 - Q^2$.

For any point P , we can now calculate the local height using intersection theory. For example, let P be a point reducing to the blue (left) component. Then $D_Z(b, P) = \mathcal{X}^\varphi - \varphi(P) - \varphi(b)$ has multidegree $(0, -1, +1)$ on the components. We now want to add to $D_Z(b, P)$ a vertical divisor V with no components containing b in order to make $D_Z(b, P)$ multidegree 0. Using linear algebra with L , we see V is $-3/5$ times the pink (right) component plus $-2/5$ times the blue (left) component, so the local height is $-2/5$, agreeing with (7.7.1.1).

Remark 7.7.1. This calculation of heights as an intersection multiplicity can be extended to work for any value of α and β and again gives the same values as in (7.7.1.1). As such, this gives an explicit check that the heights defined in terms of intersection multiplicities agree with those coming from Theorem 7.3.9.

7.7.2 A Shimura curve quotient

In [GY17] they determine equations for all geometrically hyperelliptic Shimura curves $X_0(D, N)$ and Atkin–Lehner operators on these curves. Since we know $X_0(D, N)(\mathbb{R}) = \emptyset$, for the study of rational points, it makes sense to consider Atkin–Lehner quotients of these curves. By applying Jaquet–Langlands and considering the sign of the L -function of the corresponding modular forms space, we see there are only two quotient curves whose Mordell–Weil rank are equal to their genus.

One is bielliptic, so its local heights at primes of bad reduction can be computed from the elliptic curve factors of its Jacobian. The other is the curve $X_0(93, 1)/\langle \omega_{93} \rangle$. The rational points on this curve correspond to abelian surfaces defined over $\mathbb{Q}$ with potential quaternionic multiplication by a quaternion algebra of discriminant 93. It turns out that this curve has an exceptional isomorphism to the modular curve quotient $X_0(93)^*$ whose rational points were previously determined in [BGX21] using elliptic Chabauty. We now show how to compute the local heights at primes of bad reduction for $X := X_0(93, 1)/\langle \omega_{93} \rangle$.




Figure 7.9: Reduction graph for Shimura curve quotient

The conductor of X is $3^2 \cdot 31^2$. At $\ell = 31$, the curve $X : y^2 = x^6 + 2x^4 + 6x^3 + 5x^2 - 6x + 1$ has the cluster picture in Figure 7.5 and so by Theorem 7.5.19 the local heights at 31 are trivial.

At $\ell = 3$ the cluster picture is the cluster in Figure 7.10 so more analysis is needed; the reduction graph Γ associated to the semistable covering contains three vertices and $\mathbb{Q}_\ell$ -points reduce to all three curves associated to the vertices. See Figure 7.9. Let $\mathfrak{t}_1$ be the twin cluster given by the two roots of f defined over $\mathbb{Q}_3$, and $\mathfrak{t}_2$ be the other twin cluster. These correspond to vertices v_1 and v_2 respectively. We label the edges v_1 to v_3 by $e_1^\pm$ and the edges v_2 to v_3 by $e_2^\pm$. We have $l(e_i^\pm) = 1$ for all i .




Figure 7.10: Cluster picture

Let Z be the endomorphism acting by $\sqrt{5}$ on the Jacobian. This acts by pushforward on the holomorphic differentials $\langle dx/(2y), xdx/(2y) \rangle$ by

$$M = \begin{bmatrix} -1 & 2 \\ 2 & 1 \end{bmatrix}.$$

By [BBB⁺22, Theorem 5.1], since $\nu_{\mathfrak{s}} = 0$ when $\mathfrak{s}$ is the top cluster and inertia fixes the twins setwise, X has a semistable model over $\mathbb{Q}_3$ and a split semistable model over $\mathbb{Q}_9$, the unramified quadratic extension of $\mathbb{Q}_3$.

For each $\mathfrak{t}_i$, we compute the map $A_{\mathfrak{t}_i}^+ \rightarrow X|_{A_{\mathfrak{t}_i}}$ given by $x \mapsto (x, g_{\mathfrak{t}_i}^{1/2} h_{\mathfrak{t}_i}^{1/2})$ from (7.5.1.2). We compute the matrix of $\varphi_2 : x^j dx/(2y) \mapsto \text{Res}_{A_{\mathfrak{t}_i}^+} x^j dx/(2y)$ for $i = 1, 2$ and $j = 0, 1$ to be

$$\begin{bmatrix} 20a + 20 + O(3^4) & 2a + 2 + O(3^4) \\ 2a + 2 + O(3^4) & 22a + 22 + O(3^4) \end{bmatrix}$$

where a has minimal polynomial $x^2 + 2x + 2$ over $\mathbb{Q}_3$. Using this change of basis, we conjugate M to obtain the action of Z_* on $H_1(\Gamma, \mathbb{Q}_9)$ by Proposition 7.6.10 using the complement e_1^-, e_2^- of a spanning tree

$$\begin{bmatrix} -1 + O(3^4) & 2 + O(3^4) \\ 2 + O(3^4) & 1 + O(3^4) \end{bmatrix}.$$

We now show this is an integer matrix using Proposition 7.6.11 and Corollary 7.6.12. Using [CMSV19] we explicitly construct a correspondence $Z \subset X \times X$ for M . We compute that Z has degrees $d_1 = 2$ and $d_2 = 10$ over X . We fix the (orthogonal) homology basis $\langle e_1^- - e_1^+, e_2^- - e_2^+ \rangle$. The intersection pairing on homology has the matrix $C = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$ in this basis. Writing $C = PP^T$, we see $\|P\|_{\max} = \sqrt{2}$ and $\|P^{-1}\|_{\max} = 1/\sqrt{2}$. Therefore the action of Z_* on homology is represented by an integer matrix whose entries have absolute value at most $16\sqrt{5}$. We have calculated that the entries are all integers inside this range, plus a multiple of 3^4 . Since $2 < 16\sqrt{5}$ and $3^4 > 16\sqrt{5} + 2$, we conclude that the action of Z_* on the homology of the graph is given by the integer matrix $\begin{bmatrix} -1 & 2 \\ 2 & 1 \end{bmatrix}$.

The projection of $e_1^\pm$ onto the basis of homology has coefficients $[\mp 1/2, 0]$, and the projection of $e_2^\pm$ onto the basis of homology has coefficients $[0, \mp 1/2]$. This computation yields the Laplacian $\nabla^2(\tilde{h}_{Z,\ell})$ of the piecewise polynomial height function as in Theorem 7.3.9

$$\mu_Z = 1 \cdot ds_{e_1^-} + 1 \cdot ds_{e_1^+} - 1 \cdot ds_{e_2^-} - 1 \cdot ds_{e_2^+}.$$

Let f_0 be the piecewise polynomial function $(-1)^i \frac{1}{2} s_{e_i^\pm} (s_{e_i^\pm} - 1)$, obtained from double integrating each part of μ_Z . Then f_0 is 0 at each vertex and $\mu_Z - \nabla^2(f_0) = 1 \cdot \delta_{v_1} - 1 \cdot \delta_{v_2} + 0 \cdot \delta_{v_3}$. The weighted Laplacian matrix L of the Γ is as follows. Using L we solve for the coefficient vector of a function whose Laplacian is the vector of the weights of the δ_{v_i} .

$$L := \begin{bmatrix} -2 & 0 & 2 \\ 0 & -2 & 2 \\ 2 & 2 & -4 \end{bmatrix}, \quad L \begin{bmatrix} -1 \\ 0 \\ -1/2 \end{bmatrix} = \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}.$$

Therefore the piecewise function $f_1 = \frac{1}{2} s_{e_1^\pm} - 1$ and $f_1 = -\frac{1}{2} s_{e_2^\pm}$ has the property that $\mu_Z = \nabla^2(f_0 + f_1)$. Finally, we adjust f_1' by constants to make the height function zero at a chosen base point. We fix a base point reducing to the component v_3 . This is equivalent to requiring that $f_1(0) = 0$ on each piecewise part, so $f_1 = (-1)^{i-1} \frac{1}{2} (s_{e_i^\pm} - 1)$. Our height functions is therefore

the piecewise quadratic function $\tilde{h}_{Z,\ell}$ given by $(-1)^i 1/2(s_{e_i^\pm}^2 - 2s_{e_i^\pm} + 1)$ on $e_i^\pm$ and $0 \leq s_{e_i^\pm} \leq 1$.

Using $\tilde{h}_{Z,\ell}$, we compute the heights at the known rational points of X . For each point, we first find what component of Γ it lies on. The points at infinity reduce to the unique vertex v_3 corresponding to the top cluster. For each finite point P , we compute the smallest cluster $\mathfrak{s}$ such that P belongs to $U_{\mathfrak{s}}$. For example, if $P = (1 : -3 : 1)$, picking a root $(-5 + O(3^3)) \in \mathfrak{t}_1$, then since $v_3(1 - (-5 + O(3^3))) > 0$ (where 0 is the absolute depth of the top cluster) we see that P lies in $U_{\mathfrak{t}_1}$. To distinguish if P reduces to the vertex v_1 or somewhere along an edge from $\mathfrak{t}_1$ to $P(\mathfrak{t}_1)$, we decide if P belongs to the bounding annulus $A_{\mathfrak{t}_1}$. In this case, since $v_3(1 - (-5 + O(3^3))) = 1$ is not strictly less than $\delta_{\mathfrak{t}_1} = 1$, the point P does not belong to the $A_{\mathfrak{t}_1}$.

Using the formula for $\tilde{h}_{Z,\ell}$ with $s_{e_1^-} = 1$ we see the height at P is $1/2$. Similar reasoning shows any finite $\mathbb{Q}_3$ -point $P = (x : y : z)$ where x/z is congruent to 1 modulo 3 reduces to v_1 and has height $1/2$. If x/z is congruent to 2 modulo 3, then P reduces to v_2 and has height $-1/2$, and if x/z is 0 modulo 3 then P reduces to v_3 and has height 0.

Remark 7.7.2. As the above example demonstrates, we are able to use very low ℓ -adic precision when determining the values of the local heights at primes of bad reduction, and the answers we obtain are rigorous.

7.7.3 Atkin–Lehner quotients of modular curves

In this subsection, we study the Atkin–Lehner quotients $X_0(N)^*$ of $X_0(N)$ by the full Atkin–Lehner group for $N = 147, 225, 330$. In [ACKP22] the rational points were computed using elliptic curve Chabauty instead of quadratic Chabauty, because they were unable to determine the local heights at primes of bad reduction. Here, we show all local heights are trivial.

7.7.3.1 $X_0(330)^*$

For $N = 330$, the primes of bad reduction are 3, 5 and 11. Note that $X = X_0(N)^*$ has endomorphism ring $\mathbb{Z}[\sqrt{2}]$, and hence we can take our trace 0 endomorphism Z of the Jacobian to be $\sqrt{2}$. Write $X : y^2 = x^6 + 8x^4 + 10x^3 + 20x^2 + 12x + 9$.

At $\ell = 5$ and $\ell = 11$ the cluster pictures are the cluster picture in Figure 7.10. By Theorem 7.5.19, since the leading coefficient of f is 1, the local heights at these primes are trivial on $\mathbb{Q}_\ell$ points.




Figure 7.11: Reduction graph for Atkin-Lehner quotient

For $\ell = 3$ however, the graph Γ associated to the semistable covering has two vertices v_1, v_2 corresponding to genus zero components, with a loop e_0 of length 1 at v_1 and two edges e_1, e_2 of length 1 connecting v_1 with v_2 , see Figure 7.11.

The curves attached to v_1 and v_2 both contain $\mathbb{F}_3$ -points. Since there are no higher genus vertices, we only need to compute the action of Z_* on the homology of the Γ . Like in Section 7.7.2, we can compute it 3-adically and then use Corollary 7.6.12 to obtain the following matrix of Z_* acting on the homology $H_1(\Gamma)$.

$$\begin{bmatrix} 0 & -2 \\ -1 & 0 \end{bmatrix} \quad (7.7.3.1)$$

Here the basis of the homology of Γ is $e_0, e_2 - e_1$. By the local heights formula Theorem 7.3.9 we then have that the Laplacian $\nabla^2(\tilde{h}_{Z,\ell})$ is 0, and hence $\tilde{h}_{Z,\ell} = 0$. We see that the local heights vanish.

7.7.3.2 $X_0(255)^*$

For $N = 255$, the primes of bad reduction are 5 and 17. Again, $X = X_0(N)^*$ has endomorphism ring $\mathbb{Z}[\sqrt{2}]$, and we can take our trace 0 endomorphism Z of the Jacobian to be $\sqrt{2}$. Write $X : y^2 = x^6 - 4x^5 - 12x^4 + 2x^3 + 8x^2 - 4x + 1$. At $\ell = 17$, the cluster picture is in Figure 7.10 and the leading coefficient of f is 1, so the heights vanish by Theorem 7.5.19. For $\ell = 5$, the semistable covering of the curve $X_0(N)^*$ has the same associated graph Γ as in the previous example, and the matrix of the action on $H_1(\Gamma)$ is also (7.7.3.1).

7.7.3.3 $X_0(147)^*$

For $N = 147$, the primes of bad reduction are 3 and 7. At $q = 7$, the curve $X_0(N)^*$ has potential good reduction, so the local height functions are trivial. At $q = 3$, following the same procedures as above we compute the action of Z_* on the associated homology graph to the semistable covering of $X_0(N)^*$

and find it is negative the matrix in (7.7.3.1) and therefore the heights are also trivial.

7.7.4 Higher genus example

We now provide an example of local height computations for the genus 7 modular curve with affine model $X : y^2 = x^{16} - 4x^8 + 16$. This curve has Cummins and Pauli label 48N7. This example shows that our method is practical even when the genus of X is large. We choose the endomorphism Z of the Jacobian to be associated to the diagonal matrix with entries $[4, 4, 4, -24, 4, 4, 4]$. The primes of bad reduction of X are 2 and 3; we compute local heights at $q = 3$. In this case, the cluster picture is as in Figure 7.12 and the reduction graph is as in Figure 7.13. By Theorem 7.5.19, after choosing the basepoint w_{-1} ,




Figure 7.12: Cluster picture for genus 7 curve




Figure 7.13: Reduction graph for genus 7 curve

the heights at all rational points will be trivial, but we can still determine the (non-trivial) height functions. We find that the action of Z_* on the homology of the dual graph is given by the matrix

$$\begin{bmatrix} 4 & 0 & 7 & 7 & 0 & 7 & 7 \\ 0 & 4 & 7 & 7 & 0 & 7 & 7 \\ 0 & 0 & -3 & -7 & 0 & -7 & -7 \\ 0 & 0 & -7 & -3 & 0 & -7 & -7 \\ 0 & 0 & 7 & 7 & 4 & 7 & 7 \\ 0 & 0 & -7 & -7 & 0 & -3 & -7 \\ 0 & 0 & -7 & -7 & 0 & -7 & -3 \end{bmatrix}.$$

This yields 8 distinct piecewise polynomial functions (since the two edges from each vertex v_i to $w_{\pm 1}$ have the same height function). For example, the edges from v_1 to $w_{\pm 1}$ have the height function

$$248873/44800s^2 - 248873/44800s + 248873/179200 \text{ where } 0 \leq s \leq 1/2.$$

7.7.5 A family of curves with genus 1 vertices

Let S_0 be a set of 4 roots in $\overline{\mathbb{Q}}_q$ such that $1 + q\mathbb{Z}_q$ is the smallest disc containing them. Let $n \in \mathbb{Z}$ with $n \mid (q - 1)$, let $\zeta = \zeta_n$ be a primitive n th root, and consider $S_k = \zeta^k S_0$ for $0 \leq k < n$. Let $S^{(n)} = \bigcup_{n > k \geq 0} S_k$, and let $f^{(n)} = \prod_{r \in S^{(n)}} (x - r)$.

Let X_n be the curve given by $y^2 = f^{(n)}(x)$. By construction, this has cluster picture Figure 7.14.




Figure 7.14: Cluster picture for higher genus family

Hence the semistable covering has the associated reduction graph Γ in Figure 7.15, where vertices w_0 and w_1 correspond to components of genus 0, and $v_0, \dots, v_{n-1}$ correspond to components of genus 1. Denote the edges from v_i to w_0 as e_i^+ , and those from v_i to w_1 as e_i^- .

Also by construction, the family of curves has an automorphism $\zeta : X_n \rightarrow X_n$, cyclically permuting the n clusters. On the level of graphs, this cyclically permutes the n genus 1 vertices, while keeping the genus 0 vertices invariant. We will also use ζ to denote the corresponding endomorphism of the Jacobian. Although the local height only makes sense for a trace 0 endomorphism, the




Figure 7.15: Reduction graph for higher genus family

formula for the measure in Theorem 7.3.9 makes sense for any endomorphism. This measure will be of mass 0 if the endomorphism is of trace 0. We let μ_k denote the measure corresponding to ζ_k .

We need to compute $\langle e_i^+, \zeta^k(\pi(e_i^+)) \rangle$ where $\pi : C_1(\Gamma) \rightarrow H_1(\Gamma)$ is the orthogonal projection. A straightforward calculation shows that

$$\pi(e_i^+) = \frac{n-1}{2n}(e_i^+ - e_i^-) - \frac{1}{2n} \sum_{j \neq i} (e_j^+ - e_j^-).$$

and hence $\langle e_i^+, \zeta^k(\pi(e_i^+)) \rangle$ is $\frac{n-1}{2n}$ if $k = 0$ and $-\frac{1}{2n}$ otherwise. By symmetry, the same holds for edges e_i^- .

Since ζ acts by cyclic permutation on the associated graph, we see that $\text{tr}_{v_i}(\zeta^k)$ is 2 if $k = 0$ and 0 otherwise. In total, we find

$$\begin{aligned} \mu_0 &= \frac{n-1}{n} \sum_i (|ds_{e_i^+}| + |ds_{e_i^-}|) + \sum_i 2\delta_v \\ \mu_k &= -\frac{1}{n} \sum_i (|ds_{e_i^+}| + |ds_{e_i^-}|) \text{ when } k \neq 0. \end{aligned}$$

These have total mass $2(n-1)+2n = 4n-2$ and -2 respectively. Hence we see a linear combination $\sum_k a_k \zeta_k$ has trace 0 if and only if $(2n-1)a_0 - \sum_{k \neq 0} a_k = 0$. For example, $Z = \zeta_0 - (2n-1)\zeta_1$ has trace 0, and measure

$$\mu_Z = - \sum_i (|ds_{e_i^+}| + |ds_{e_i^-}|) + \sum_i 2\delta_{v_i}.$$

We find an inverse Laplacian using the method explained in Section 7.3.2. First we find a quadratic piecewise polynomial function f_0 such that $\mu_z - \nabla^2 f_0$ is a sum of δ measures. We see that the function that is $\frac{-1}{2}s(1-s)$ on every edge works. Then $\mu_z - \nabla^2 f_0 = -\frac{n}{2}(\delta_{w_0} + \delta_{w_1}) + \sum_i \delta_{v_i}$, and we need to find a piecewise linear function f_1 with this Laplacian. We see any function f_1 which has slopes $-\frac{1}{2}$ along $e_i^\pm$ suffices.

Taking the basepoint to lie in w_0 uniquely determines f_1 by $f_1(w_0) = 0$, and we find that the normalised height is the piecewise polynomial given by

$$\frac{1}{2}(s-1)^2 \text{ on } e_i^\pm \text{ with } 0 \leq s \leq 1.$$

7.7.6 Applications to quadratic Chabauty computations

Consider the quadratic twist by 5 of the genus 2 curve with LMFDB label 18225.c.164025.1, given by the affine model $X : y^2 = x^6 + 18/5x^4 + 6/5x^3 +$




Figure 7.16: Reduction graph for quadratic Chabauty example at $\ell = 3$

$9/5x^2 + 6/5x + 1/5$. The conductor is $3^6 \cdot 5^4$. The endomorphism ring of $X/\mathbb{Q}$ is $\mathbb{Q}(\sqrt{13})$. We choose the endomorphism $\sqrt{13}$ on the Jacobian. We compute the local height functions at 3 and 5 and determine the set of rational points $X(\mathbb{Q})$ using quadratic Chabauty.

Remark 7.7.3. The rational points $X(\mathbb{Q})$ can also be determined using elliptic Chabauty.

Remark 7.7.4. The implementation [BDM⁺] chooses the endomorphism Z to be $-4\sqrt{13}$ when $p = 53$. For this reason, in our implementation we have to scale the heights accordingly.

Applying Proposition 7.6.9, after computing a correspondence $Z \subset X \times X$ for $\sqrt{13}$, we compute the action of Z_* on $H_{\text{dR}}^1(X)$ to be

$$\begin{bmatrix} 1 & 2 & 4/5 & -3/10 \\ 6 & -1 & 9/5 & 7/10 \\ 0 & 0 & -1 & 1 \\ 0 & 0 & 12 & 1 \end{bmatrix}.$$


where the basis for $H_{\text{dR}}^1(X)$ is

$$\{dx/(2y), xdx/(2y), x^3dx/(2y), (9/10)x^2dx/(2y) - (-1/2)x^4dx/(2y)\}.$$

Let $b = (-1/3 : 1/27 : 1)$ be the basepoint.

At $\ell = 3$, the cluster picture is Figure 7.17 and the curve has unstable reduction; in fact, the curve has unstable reduction over every tame extension. The associated reduction graph Γ has two genus 1 vertices v_1 and v_2 , linked by a genus 0 vertex, by edges e_1 and e_2 respectively, as in Figure 7.16. We must compute the traces $\text{tr}_{v_i}(Z)$ for each genus 1 vertex v_i . Applying Proposition 7.6.13, we compute the block matrices

$$\begin{bmatrix} -1 + O(3^3) & 0 + O(3^3) \\ 0 + O(3^3) & -1 + O(3^3) \end{bmatrix}, \quad \begin{bmatrix} 1 + O(3^3) & 0 + O(3^3) \\ 0 + O(3^3) & 1 + O(3^3) \end{bmatrix}$$

Figure 7.17: Cluster picture for quadratic Chabauty example at $\ell = 3$

to be the action of Z_* on $H_{\text{dR}}^1(\mathcal{X}_{v_i})$, for $i = 1, 2$ respectively. By Theorem 7.3.5, the traces of these matrices are bounded by $2 \cdot 1 \cdot \max\{2, 7\}$, so they are exactly -2 and 2 . Using the formula from Theorem 7.3.9, the Laplacian of the local height function is given by the measure

$$\mu_Z = -2\delta_{v_1} + 2\delta_{v_2}.$$

We solve for the height function, and find it is the piecewise polynomial function

$$\begin{cases} -2s + 1/2 & \text{on } e_1 \text{ with } 0 \leq s \leq 1/4 \\ 2s - 5/6 & \text{on } e_2 \text{ with } 0 \leq s \leq 5/12 \end{cases}$$

Using similar calculations to Section 7.7.2, we see that if $P = (x : y : z)$ is a point with $z \not\equiv 0 \pmod{3}$, then if $x/z \equiv 1 \pmod{3}$, then P lies at distance $1/12$ along edge e_2 . If $z \equiv 0 \pmod{3}$ then P reduces to the genus 0 vertex. (By Theorem 7.5.19, we cannot have $\mathbb{Q}_\ell$ -points reducing to $\mathcal{X}_{v_1}$.) Therefore the normalised local height of $P = (x : y : z)$ is

$$\begin{cases} 0 & \text{if } z \equiv 0 \pmod{3}, \\ -2/3 & \text{if } x/z \equiv 1 \pmod{3}. \end{cases}$$

Now consider $\ell = 5$ the other prime of bad reduction. Here the curve has unstable reduction. The cluster picture is Figure 7.18.

Figure 7.18: Cluster picture for quadratic Chabauty example at $\ell = 5$

Let $e_1^\pm, e_2^\pm$ be the edges as labeled in Figure 7.19, with $l(e_i^\pm) = 1/2$. We compute the following piecewise polynomial height function.

$$\begin{cases} 3s^2 & \text{on } e_1^\pm \text{ with } 0 \leq s \leq 1/2 \\ -3s^2 + 3s - 3/4 & \text{on } e_2^\pm \text{ with } 0 \leq s \leq 1/2 \end{cases}$$




Figure 7.19: Reduction graph for quadratic Chabauty example at $\ell = 5$

If $P = (x : y : z)$ is a rational point with $z \not\equiv 0 \pmod{5}$ and $x/z \equiv 2 \pmod{5}$, then P reduces to v_3 . If $z \equiv 0 \pmod{5}$ then P reduces to v_1 . Otherwise P reduces to v_2 . Therefore the normalised local height of $P = (x : y : z)$ is

$$\begin{cases} 3/4 & \text{if } z \equiv 0 \pmod{5}, \\ -3/4 & \text{if } x/z \equiv 2 \pmod{5}, \\ 0 & \text{otherwise.} \end{cases}$$

The curve X has Mordell–Weil rank 2, and therefore we can apply the quadratic Chabauty method at $p = 53$ (as implemented in [BDM⁺]) to find a finite set of p -adic points containing the rational points $X(\mathbb{Q})$. Combining this with a Mordell–Weil sieve we obtain the following theorem.

Theorem 7.7.5. *The rational points of $X : y^2 = x^6 + 18/5x^4 + 6/5x^3 + 9/5x^2 + 6/5x + 1/5$ are the 10 points*

$$\{(-1/3 : -1/27 : 1), (-1/3 : 1/27 : 1), (-1/5 : -21/125 : 1), (-1/5 : 21/125 : 1), (1 : 3 : 1), (1 : -3 : 1), (1 : 1 : 0), (1 : -1 : 0), (-1/2 : -3/8 : 1), (-1/2 : 3/8 : 1)\}.$$